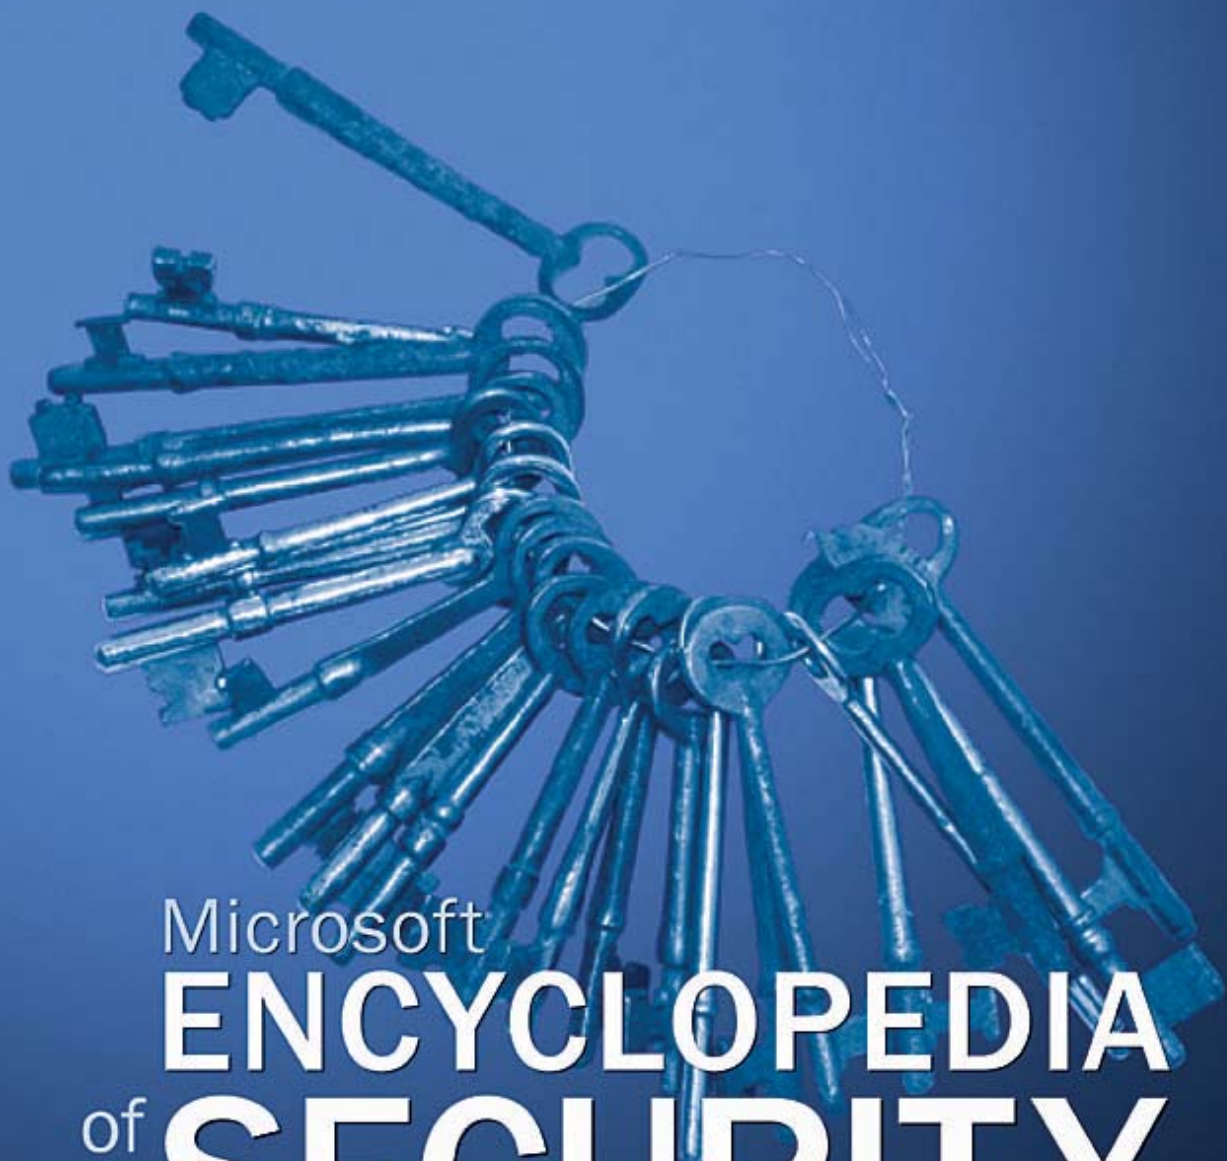


Microsoft



Microsoft
ENCYCLOPEDIA
of **SECURITY**

Mitch Tulloch

PUBLISHED BY

Microsoft Press
A Division of Microsoft Corporation
One Microsoft Way
Redmond, Washington 98052-6399

Copyright © 2003 by Mitch Tulloch

All rights reserved. No part of the contents of this book may be reproduced or transmitted in any form or by any means without the written permission of the publisher.

Library of Congress Cataloging-in-Publication Data
Tulloch, Mitch.

Microsoft Encyclopedia of Security / Mitch Tulloch.

p. cm.

ISBN 0-7356-1877-1

1. Computer security--Encyclopedias 2. Computer networks--Security measures--
Encyclopedias. I. Title.

QA76.9.A25T85 2003

005.'8'03--dc21

2003051323

Printed and bound in the United States of America.

1 2 3 4 5 6 7 8 9 QWT 8 7 6 5 4 3

Distributed in Canada by H.B. Fenn and Company Ltd.

A CIP catalogue record for this book is available from the British Library.

Microsoft Press books are available through booksellers and distributors worldwide. For further information about international editions, contact your local Microsoft Corporation office or contact Microsoft Press International directly at fax (425) 936-7329. Visit our Web site at www.microsoft.com/mspress. Send comments to mspinput@microsoft.com.

Active Directory, ActiveX, Authenticode, BackOffice, Hotmail, Microsoft, Microsoft Press, MS-DOS, MSDN, MSN, Windows, Windows NT, and Windows Server are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. Other product and company names mentioned herein may be the trademarks of their respective owners.

The example companies, organizations, products, domain names, e-mail addresses, logos, people, places, and events depicted herein are fictitious. No association with any real company, organization, product, domain name, e-mail address, logo, person, place, or event is intended or should be inferred.

Acquisitions Editor: Jeff Koch

Project Editor: Sandra Haynes

Dedicated to Neil Salkind, my agent and friend

Contents-

| | |
|---|--------------|
| <i>Acknowledgements</i> | <i>xxi</i> |
| <i>Introduction</i> | <i>xxiii</i> |
| <i>What Is Computer Security?</i> | <i>xxiii</i> |
| <i>Threats and Vulnerabilities</i> | <i>xxiii</i> |
| <i>Standards and Protocols</i> | <i>xxiii</i> |
| <i>Hacking and Cracking</i> | <i>xxiv</i> |
| <i>Tools and Procedures</i> | <i>xxv</i> |
| <i>Organizations and Certifications</i> | <i>xxv</i> |
| <i>Cryptography</i> | <i>xxv</i> |
| <i>Legal Issues</i> | <i>xxvi</i> |
| <i>Who This Work Is For</i> | <i>xxvi</i> |
| <i>How This Work Is Organized</i> | <i>xxvi</i> |
| <i>Disclaimer</i> | <i>xxvii</i> |
| <i>Comments and Questions</i> | <i>xxvii</i> |

Alphabetical Reference of Terms

Numbers

| | | | |
|--------------|---|---------------|---|
| 3DES | 1 | 802.11i | 3 |
| 802.1x | 1 | 2600 | 3 |

A

| | | | |
|-----------------------------------|----|--|----|
| A5 | 5 | ACK storm | 10 |
| AAA | 5 | ACL | 11 |
| acceptable use policy (AUP) | 5 | AclDiag | 11 |
| access | 6 | ACPA | 11 |
| access control | 6 | ACSA | 11 |
| access control entry (ACE) | 7 | ACSAC | 11 |
| access control list (ACL) | 7 | Active Directory | 11 |
| access list | 8 | adaptive proxy | 12 |
| access mask | 9 | Adaptive Security Algorithm (ASA) | 12 |
| access token | 9 | address-based authentication | 12 |
| account lockout | 9 | address munging | 13 |
| account lockout policy | 10 | address spoofing | 13 |
| account policy | 10 | Administrator | 13 |
| ACE | 10 | | |

Contents

| | | | |
|--|----|--|----|
| Admintool | 14 | arbitrary code execution attack | 23 |
| admnlock | 14 | Argus | 23 |
| ADMw0rm | 14 | ARP cache poisoning | 23 |
| Advanced Encryption | | ARP redirection | 24 |
| Standard (AES) | 14 | ARP spoofing | 24 |
| Advanced Security Audit Trail Analysis on UNIX | | Arpwatch | 24 |
| (ASAX) | 15 | AS | 25 |
| Advanced Transaction Look-up and Signaling | | ASA | 25 |
| (ATLAS) | 15 | ASAX | 25 |
| advisory | 16 | ASP.NET Forms authentication | 25 |
| adware | 16 | assets | 25 |
| AES | 16 | asymmetric key algorithm | 25 |
| AH | 17 | ATLAS | 26 |
| AKE | 17 | ATR string | 26 |
| alert | 17 | attack | 27 |
| alert flooding | 17 | attack map | 27 |
| Amap | 17 | auditing | 28 |
| amplification attack | 18 | audit log | 29 |
| Annual Computer Security Applications Conference | | Auditpol | 29 |
| (ACSAC) | 18 | audit policy | 29 |
| anomaly-based IDS | 18 | audit trail | 30 |
| anonymous access | 19 | Augmented Key Exchange (AKE) | 30 |
| anonymous proxy | 19 | AUP | 31 |
| anonymous Web browsing | 19 | AusCERT | 31 |
| Anticybersquatting Consumer Protection | | Australian Computer Emergency Response | |
| Act (ACPA) | 20 | Team (AusCERT) | 31 |
| AntiSniff | 20 | authentication | 31 |
| antivirus software | 20 | Authentication, Authorization, and | |
| application-level gateway | 20 | Accounting (AAA) | 32 |
| application-level proxy | 21 | Authentication Header (AH) | 32 |
| application protection system (APS) | 21 | authentication package | 33 |
| Application Security Tool (AppSec) | 21 | authentication server (AS) | 33 |
| Applications as Services (Srvany) | 21 | Authenticode | 33 |
| Applied Computer Security Associates (ACSA) .. | 22 | authorization | 34 |
| AppSec | 22 | authorization creep | 34 |
| APS | 22 | autologon | 34 |
| Apsend | 22 | autorooter | 34 |
| APSR | 23 | | |
| B | | | |
| backdoor | 35 | backup plan | 37 |
| Back Orifice | 35 | Badtrans.B | 37 |
| Back Orifice 2000 (BO2K) | 36 | bandwidth consumption attack | 38 |
| backup authority | 37 | banner grabbing | 38 |

| | | | |
|--|----|--|----|
| base content type | 38 | Blowfish | 44 |
| Basic authentication | 38 | BO2K | 45 |
| Basic Encoding Rules (BER) | 39 | boink attack | 45 |
| Bastille | 39 | bonk attack | 45 |
| bastion host | 39 | Brown Orifice | 45 |
| BBBOnLine | 40 | BRP | 46 |
| behavior-blocking software | 40 | brute-force attack | 46 |
| BER | 41 | bucket brigade attack | 46 |
| biometric identification | 41 | buffer overflow | 46 |
| BIOS cracking | 42 | buffer overrun | 46 |
| black hat | 43 | Bugtraq | 47 |
| Black Hat Briefings | 43 | bulk encryption key | 47 |
| blackholing | 43 | business continuity plan | 47 |
| BLOB | 44 | business resumption plan (BRP) | 47 |
| block cipher | 44 | | |
| C | | | |
| CA | 49 | certificate request | 54 |
| CA certificate | 49 | certificate revocation list (CRL) | 54 |
| CA hierarchy | 49 | certificate server | 54 |
| cache poisoning | 50 | certificate store | 55 |
| callback | 50 | certificate trust list (CTL) | 55 |
| Canadian Centre for Information Technology | | Certified Information Systems Security | |
| Security (CCITS) | 50 | Professional (CISSP) | 55 |
| canonicalization error | 50 | CFB | 56 |
| CAPI | 51 | chaining mode | 56 |
| CAPICOM | 51 | Challenge Handshake Authentication | |
| Carnivore | 51 | Protocol (CHAP) | 56 |
| CAS | 51 | challenge response authentication | 56 |
| CAST | 51 | CHAP | 57 |
| CBC | 51 | Chernobyl | 57 |
| CCA | 51 | chief security officer (CSO) | 57 |
| CCITS | 51 | chosen ciphertext attack | 58 |
| cDc | 52 | chosen plaintext attack | 58 |
| Center for Education and Research in Information | | chroot jail | 58 |
| Assurance and Security (CERIAS) | 52 | CIAC | 59 |
| Center for Internet Security (CIS) | 52 | cipher | 59 |
| CERIAS | 52 | cipher block chaining (CBC) | 59 |
| CERT/CC | 52 | cipher feedback (CFB) | 59 |
| CERT Coordination Center (CERT/CC) | 52 | cipher mode | 59 |
| certificate | 53 | ciphertext | 60 |
| certificate authority (CA) | 53 | ciphertext-only attack | 60 |
| certificate-based authentication | 53 | CIS | 60 |
| Certificate Information Systems Auditor (CISA) | 54 | CISA | 60 |

Contents

| | | | |
|---|----|---|----|
| CISSP | 60 | CRC | 69 |
| cleartext | 60 | credentials | 69 |
| clogging attack | 61 | CRL | 69 |
| code access permissions | 61 | cross-realm authentication | 69 |
| code access security (CAS) | 61 | cross-site scripting (CSS) | 69 |
| CodeRed | 61 | cryptanalysis | 70 |
| code signing | 62 | CryptoAPI (CAPI) | 70 |
| Common Criteria & Methodology for Information Technology Security Evaluation | 62 | cryptographic hash function | 71 |
| Common Cryptographic Architecture (CCA) | 63 | cryptographic service provider (CSP) | 71 |
| Common Vulnerabilities and Exposures (CVE) | 63 | cryptography | 71 |
| compromised system | 63 | cryptology | 72 |
| computer forensics | 64 | cryptosystem | 72 |
| Computer Incident Advisory Capability (CIAC) | 64 | CSD | 72 |
| Computer Security Division (CSD) | 65 | CSI | 72 |
| computer security incident response team (CSIRT) | 65 | CSIRT | 72 |
| Computer Security Institute (CSI) | 65 | CSO | 72 |
| confidentiality | 66 | CSP | 72 |
| confidentiality agreement | 66 | CSS | 72 |
| consensus baseline security settings | 66 | CTL | 73 |
| cookie poisoning | 67 | Cult of the Dead Cow (cDc) | 73 |
| covert channel | 67 | CVE | 73 |
| cracking | 68 | cybercrime | 73 |
| D | | cyclical redundancy check (CRC) | 73 |
| DAC | 75 | | |
| DACL | 75 | DESX | 80 |
| Data Encryption Algorithm (DEA) | 75 | DH | 81 |
| Data Encryption Standard (DES) | 75 | dictionary attack | 81 |
| data integrity | 75 | Diffie-Hellman (DH) | 81 |
| Data Protection API (DPAPI) | 76 | diffing | 81 |
| DCS-1000 | 76 | Digest authentication | 82 |
| DDoS | 77 | DigiCrime | 83 |
| DEA | 77 | digital certificate | 83 |
| decryption | 77 | digital fingerprinting | 83 |
| Defcon | 77 | digital forensics | 83 |
| defense in depth | 78 | Digital Millennium Copyright Act (DMCA) | 84 |
| demilitarized zone (DMZ) | 78 | Digital Rights Management (DRM) | 84 |
| denial of service (DoS) | 79 | digital signature | 85 |
| Department of Defense Information Technology Security Certification and Accreditation Process (DITSCAP) | 80 | Digital Signature Algorithm (DSA) | 86 |
| DES | 80 | Digital Signature Standard (DSS) | 86 |
| | | digital watermarking | 87 |
| | | disaster recovery plan (DRP) | 87 |
| | | discretionary access control (DAC) | 88 |

| | | | |
|---|-----|--|-----|
| discretionary access control list (DACL) | 88 | DPAPI | 92 |
| distributed denial of service (DDoS) | 89 | DRM | 92 |
| DITSCAP | 90 | DRP | 92 |
| DMCA | 90 | DSA | 92 |
| DMZ | 91 | Dsniff | 92 |
| DNS cache poisoning | 91 | DSS | 92 |
| DNS spoofing | 91 | dynamic packet filtering | 92 |
| DoS | 91 | dynamic proxy | 93 |
| dot bug vulnerability | 91 | | |
| E | | | |
| EAP | 95 | Encrypting File System (EFS) | 101 |
| EAP-TLS | 95 | encryption | 101 |
| EAP-TTLS | 95 | encryption algorithm | 101 |
| eavesdropping | 95 | end-to-end encryption | 102 |
| ECB | 96 | ENUM | 103 |
| ECC | 96 | enumeration | 103 |
| ECDSA | 96 | EoP | 104 |
| EFS | 96 | EPIC | 104 |
| egress filtering | 96 | E-SIGN Act | 104 |
| EICAR | 97 | ESP | 104 |
| EKE | 97 | /etc/passwd | 104 |
| Electronic Codebook (ECB) | 98 | Ethereal | 104 |
| Electronic Privacy Information Center (EPIC) | 98 | European Institute of Computer Anti-Virus Research (EICAR) | 105 |
| Electronic Signatures in Global and National Commerce (E-SIGN) Act | 98 | event logs | 105 |
| elevation of privileges (EoP) | 99 | exploit | 105 |
| El Gamal | 99 | exposure | 106 |
| elliptic curve cryptography (ECC) | 99 | Extensible Authentication Protocol (EAP) | 106 |
| Elliptic Curve Digital Signature Algorithm (ECDSA) | 100 | Extensible Authentication Protocol–Transport Layer Security (EAP-TLS) | 107 |
| Encapsulating Security Payload (ESP) | 100 | Extensible Authentication Protocol–Tunneled Transport Layer Security (EAP-TTLS) | 107 |
| Encrypted Key Exchange (EKE) | 100 | | |
| F | | | |
| Fair Information Practices (FIP) | 109 | Federal Information Technology Security Assessment Framework (FITSAF) | 111 |
| false negative | 109 | file integrity checker | 111 |
| false positive | 110 | File Signature Verification (FSV) | 112 |
| fast packet keying | 110 | file slack | 112 |
| FedCIRC | 110 | file system traversal attack | 113 |
| Federal Computer Incident Response Center (FedCIRC) | 110 | filter | 113 |
| Federal Information Processing Standard (FIPS) | 111 | Finger | 114 |

Contents

| | | | |
|--|-----|--|-----|
| fingerprinting | 114 | Forum of Incident Response and Security Teams (FIRST) | 117 |
| FIP | 115 | Fping | 118 |
| FIPS | 115 | Fpipe | 118 |
| firewall | 115 | Fport | 118 |
| FIRST | 116 | fragmentation | 119 |
| FITSAF | 117 | FSV | 119 |
| footprinting | 117 | FTP bounce attack | 119 |
| FORTEZZA | 117 | | |
| G | | | |
| GetAdmin | 121 | Goner | 122 |
| GIAC | 121 | Good Times | 123 |
| Global Information Assurance Certification (GIAC) | 121 | gray hat | 123 |
| GnuPG | 122 | Group Policy | 124 |
| GNU Privacy Guard (GnuPG) | 122 | guest account | 124 |
| H | | | |
| hacker | 125 | hijacking | 131 |
| Hackers On Planet Earth (HOPE) | 125 | HMAC | 131 |
| hacking | 126 | hoax | 131 |
| hacktivism | 126 | Honeynet Project | 132 |
| hardening | 127 | honeypot | 132 |
| hardware security module (HSM) | 127 | HOPE | 133 |
| hash | 128 | host-based IDS | 133 |
| hashing algorithm | 128 | host-based intrusion detection system (HIDS) | 133 |
| hash-based message authentication code (HMAC) | 129 | host-based security | 134 |
| headless server | 129 | hotfix | 134 |
| hex editor | 130 | Hping | 135 |
| hex encoding URL attack | 130 | HSM | 135 |
| HFNetChk | 130 | .htaccess | 135 |
| hidden file | 130 | HTTPS | 135 |
| HIDS | 131 | hybrid attack | 135 |
| hierarchy of trust | 131 | | |
| I | | | |
| IA | 137 | ICMP Traceback (itrace) | 138 |
| IASE | 137 | ICMP tunneling | 138 |
| IATF | 137 | IDEA | 139 |
| ICMP attacks | 137 | identity theft | 139 |
| ICMP enumeration | 137 | idle host scan | 140 |
| ICMP fingerprinting | 138 | IDS | 141 |
| ICMP flood | 138 | IIS Lockdown Tool | 141 |
| ICMP sweep | 138 | IKE | 141 |

| | | | |
|---|-----|---|-----|
| IKEv2 | 141 | International Information Systems Security | |
| ILOVEYOU | 141 | Certification Consortium (ISC) ² | 148 |
| impersonation | 141 | Internet Key Exchange (IKE) | 148 |
| incident | 141 | Internet Key Exchange version 2 (IKEv2) | 149 |
| incident response | 142 | Internet Protocol Security (IPSec) | 149 |
| incident response team | 142 | Internet Security and Acceleration (ISA) Server | 150 |
| infection | 142 | intrusion | 150 |
| information assurance (IA) | 143 | intrusion detection system (IDS) | 150 |
| Information Assurance Support | | intrusion prevention system (IPS) | 152 |
| Environment (IASE) | 143 | IP address–based authentication | 152 |
| Information Assurance Technical | | IP address restriction | 152 |
| Framework (IATF) | 143 | IP address spoofing | 153 |
| information leakage | 144 | IP fragmentation attack | 153 |
| Information Systems Audit and Control | | Iplog | 154 |
| Association (ISACA) | 144 | IPS | 154 |
| Information Systems Security Association (ISSA) | 145 | IPSec | 154 |
| Information Technology Security Evaluation | | IPSec filter | 154 |
| Criteria (ITSEC) | 145 | IPSec policy | 154 |
| infosec | 145 | IP spoofing | 155 |
| InfraGard | 145 | ISACA | 155 |
| ingress filtering | 146 | ISA Server | 155 |
| initialization vector | 146 | (ISC) ² | 155 |
| input validation attack | 146 | island-hopping | 155 |
| insider attack | 147 | ISO 17799 | 155 |
| integrity | 147 | ISSA | 156 |
| International Data Encryption Algorithm (IDEA) | 147 | itrace | 156 |
| | | ITSEC | 156 |
| J | | | |
| JFK | 157 | Jolt2 | 157 |
| Jill | 157 | Juggernaut | 157 |
| John the Ripper | 157 | Just Fast Keying (JFK) | 158 |
| K | | | |
| KDC | 159 | key management | 164 |
| Kensington security slot | 159 | key pair | 164 |
| Kerberos | 159 | key recovery | 164 |
| Kerberos policy | 160 | key ring | 165 |
| key | 160 | key rollover | 165 |
| key distribution center (KDC) | 161 | key search attack | 165 |
| keyed hash | 161 | keyspace | 165 |
| keyed-hash message authentication code | 162 | keystroke logger | 166 |
| key escrow | 162 | klaxon | 166 |
| key exchange | 163 | Klez | 166 |

Contents

| | | | |
|--|-----|---|-----|
| Knark | 167 | KryptoKnight | 167 |
| known plaintext attack | 167 | | |
| L | | | |
| L0phtCrack | 169 | locking down | 176 |
| L2TP | 169 | log analysis software | 176 |
| LaGrande Technology (LT) | 169 | log cleaning | 176 |
| LAND attack | 169 | log file monitor (LFM) | 177 |
| LAN Manager authentication | 170 | logic bomb | 177 |
| LANMAN authentication | 170 | Loginlog | 177 |
| Layer 2 Tunneling Protocol (L2TP) | 171 | logon | 177 |
| LEAP | 171 | logon identifier | 178 |
| least privilege | 171 | logon session | 178 |
| LFM | 172 | logon SID | 178 |
| Liberty Alliance Project | 172 | Loki | 178 |
| Lightweight Extensible Authentication Protocol (LEAP) | 172 | LoveLetter | 179 |
| Linsniff | 174 | LRA | 179 |
| listening port | 174 | LSA | 179 |
| LM authentication | 174 | Lsadbump2 | 179 |
| local attack | 175 | LSA Secrets | 179 |
| local exploit | 175 | Lsof | 180 |
| locally unique identifier (LUID) | 175 | LT | 180 |
| local registration authority (LRA) | 175 | LUCIFER | 180 |
| Local Security Authority (LSA) | 176 | LUID | 180 |
| local security policy | 176 | Luring attack | 180 |
| M | | | |
| MAC | 181 | MBSA | 188 |
| MAC duplication | 181 | MCSA: Security | 188 |
| MAC flooding | 181 | MCSE: Security | 188 |
| MAC spoofing | 181 | MD2 | 188 |
| macro virus | 182 | MD4 | 188 |
| Mafia Boy | 182 | MD5 | 188 |
| mail bombing | 182 | meet-in-the-middle attack | 188 |
| mail relaying | 183 | Melissa | 189 |
| malformed packet attack | 183 | message | 189 |
| malformed URL attack | 184 | message authentication code (MAC) | 190 |
| malicious code | 184 | message digest (MD) | 191 |
| malware | 184 | message digest 2 (MD2) | 191 |
| managed security service provider (MSSP) | 185 | message digest 4 (MD4) | 191 |
| mandatory access control (MAC) | 185 | message digest 5 (MD5) | 192 |
| man-in-the-middle (MITM) attack | 186 | message integrity code (MIC) | 192 |
| master key | 187 | MIC | 192 |

| | | | |
|--|-----|---|-----|
| Microsoft Baseline Security Analyzer (MBSA) .. | 192 | Microsoft Security Update | 197 |
| Microsoft Certified Systems Administrator (MCSA): Security | 193 | Microsoft Strategic Technology Protection Program (STPP) | 197 |
| Microsoft Certified Systems Engineer (MCSE): Security | 193 | Microsoft TechNet Security | 197 |
| Microsoft Challenge Handshake Authentication Protocol (MS-CHAP) | 194 | MITM | 198 |
| Microsoft Personal Security Analyzer (MPSA) .. | 195 | Morris worm | 198 |
| Microsoft Security & Privacy | 195 | MPSA | 198 |
| Microsoft Security Notification Service | 195 | MS-CHAP | 199 |
| Microsoft Security Response Center (MSRC) ... | 196 | MSRC | 199 |
| Microsoft Security Toolkit | 196 | MSSP | 199 |
| | | Mstream | 199 |
| N | | mutual authentication | 199 |
| NAT | 201 | | |
| National Computer Security Center (NCSC) | 201 | network mapper | 211 |
| National Fraud Information Center (NFIC) | 201 | network monitor | 211 |
| National Information Assurance Certification and Accreditation Process (NIACAP) | 202 | network monitoring | 211 |
| National Information Assurance Partnership (NIAP) | 203 | Network Security Hotfix Checker | 212 |
| National INFOSEC Education & Training Program (NIETP) | 203 | Newtear | 212 |
| National Infrastructure Protection Center (NIPC) | 203 | Next-Generation Secure Computing Base for Windows | 212 |
| National Institute of Standards and Technology (NIST) | 204 | NFIC | 213 |
| National Security Agency (NSA) | 204 | Ngrep | 213 |
| National Strategy to Secure Cyberspace | 205 | NIACAP | 213 |
| National Telecommunications and Information Administration (NTIA) | 206 | NIAP | 214 |
| Nbtscan | 206 | NIDS | 214 |
| Nbtstat | 206 | NIETP | 214 |
| NCSC | 206 | Nimda | 214 |
| Nessus | 206 | NIPC | 215 |
| Netbus | 207 | NIST | 215 |
| Netcat | 207 | Nmap | 215 |
| .NET Passport | 208 | nonce | 215 |
| Netstat | 208 | nonrepudiation | 216 |
| network address translation (NAT) | 209 | notice | 216 |
| network-based intrusion detection system (NIDS) | 210 | Npasswd | 216 |
| network-based security | 210 | NSA | 217 |
| network logon | 210 | Nslookup | 217 |
| | | NTBugtraq | 217 |
| | | NTFS | 217 |
| | | NTIA | 218 |
| | | NtLm | 218 |
| | | Ntrights | 219 |
| | | null session attack | 219 |

O

OAKLEY 221
 obscurity 221
 OCSP 221
 OCTAVE 222
 OFB 222
 one-time pad (OTP) 222
 one-time password (OTP) 222
 one-way authentication 222
 one-way encryption algorithm 223
 one-way function 223
 Onion Routing 223
 Online Certificate Status Protocol (OCSP) 224
 Online Personal Privacy Protection Act 224
 onward transfer 225

P

P3P 231
 packet filtering 231
 packet modification 233
 packet replay 233
 packet sniffer 233
 padding 233
 Palladium 233
 PAM 234
 PAP 234
 parking lot attack 234
 Passfilt.dll 234
 passive attack 234
 passphrase 234
 Passport 235
 Passprop 235
 password 235
 Password Authentication Protocol (PAP) 236
 password-based encryption (PBE) 236
 password cracking 237
 password grinding 238
 password hash 238
 password policy 238
 password recovery 238
 password shadowing 239
 patch 239
 PBE 239
 PCBC 239

OpenHack 225
 open mail relay 225
 OpenPGP 226
 OpenSSH 226
 OpenSSL 226
 open system 226
 Orange Book 227
 opt in 227
 opt out 227
 OTP 227
 Outlook E-mail Security Update 227
 out-of-band management 228
 output feedback mode (OFB) 228
 overt channel 229

PCT 239
 PEAP 240
 Peekabooty Project 240
 PEM 241
 penetration testing 241
 perfect forward secrecy (PFS) 241
 perimeter network 242
 permissions 242
 personal data 242
 personal identification device (PID) 242
 personal identification number (PIN) 243
 personal information 243
 personally identifiable information (PII) 243
 PFS 243
 PGP 243
 phishing 243
 Phrack 244
 phreaking 244
 physical security 244
 PIC 245
 PID 245
 PII 245
 pilfering 245
 PIN 245
 ping 246
 ping flood 246
 ping of death 246

| | | | |
|--|-----|--|-----|
| ping sweep | 246 | privacy statement | 255 |
| PKCS | 247 | Private Communication Technology (PCT) | 255 |
| PKCS #7 | 247 | private key | 255 |
| PKI | 247 | private key encryption | 255 |
| PKINIT | 247 | privilege escalation | 255 |
| PKIX | 247 | privileges | 255 |
| plaintext | 248 | PRNG | 256 |
| plaintext cipher block chaining (PCBC) | 248 | process table attack | 256 |
| Platform for Privacy Preferences (P3P) | 248 | promiscuous mode | 256 |
| playback | 248 | Protected Extensible Authentication Protocol (PEAP) | 256 |
| pluggable authentication module (PAM) | 248 | protocol analyzer | 256 |
| Point-to-Point Tunneling Protocol (PPTP) | 249 | pseudorandom number generator (PRNG) | 257 |
| port flooding | 249 | public key | 257 |
| port forwarding | 249 | public key cryptography | 258 |
| port numbers | 250 | public key cryptography standards (PKCS) | 259 |
| port redirection | 251 | public key encryption | 259 |
| port scanning | 251 | Public Key Infrastructure (PKI) | 259 |
| PPTP | 252 | Public-Key Infrastructure (X.509) (PKIX) | 260 |
| Pre-IKE Credential (PIC) | 252 | Publius Project | 260 |
| Pretty Good Privacy (PGP) | 253 | Pulist | 261 |
| principal | 253 | Pwdump | 261 |
| privacy | 254 | PWL file | 261 |
| Privacy Enhanced Mail (PEM) | 254 | | |
| privacy policy | 254 | | |
| Q | | | |
| Qchain | 263 | Queso | 263 |
| QFE | 263 | Quick Fix Engineering (QFE) | 264 |
| Qfecheck | 263 | | |
| R | | | |
| RA | 265 | recovery agent | 267 |
| race condition | 265 | Recovery Console | 268 |
| RADIUS | 265 | reflection attack | 268 |
| RAM slack | 265 | Regdmp | 269 |
| RAT | 266 | registration authority (RA) | 269 |
| RBAC | 266 | regression testing | 270 |
| RC2 | 266 | remote administration tool (RAT) | 270 |
| RC4 | 266 | Remote Authentication Dial-In User Service (RADIUS) | 270 |
| RC5 | 266 | replay attack | 271 |
| RC6 | 266 | repudiation | 271 |
| realm | 267 | resource exhaustion attack | 271 |
| recognizable plaintext attack | 267 | restrictive shell | 272 |
| reconnaissance | 267 | | |

Contents

| | | | |
|-------------------------------------|-----|---|-----|
| reverse Telnet | 272 | role-based security | 281 |
| reversible encryption | 272 | rollup | 282 |
| Rexec | 273 | root | 282 |
| .rhosts | 273 | root CA | 282 |
| rights | 273 | root certificate | 282 |
| Rijndael | 277 | rootkit | 284 |
| Rinetd | 278 | root rollover | 285 |
| RIP spoofing | 278 | route verification | 285 |
| risk assessment | 278 | Rpcdump | 285 |
| Rivest-Shamir-Adleman (RSA) | 279 | RSA | 286 |
| Rlogin | 280 | Rsh | 286 |
| Rnmap | 280 | rule | 287 |
| role-based access control (RBAC) | 281 | Runas | 287 |
| role-based authorization | 281 | Rwho | 288 |
| S | | | |
| SACL | 289 | Secure Hash Algorithm-1 (SHA-1) | 296 |
| sacrificial lamb | 289 | Secure Hash Standard (SHS) | 297 |
| Sadmin | 289 | Secure Hypertext Transfer Protocol (S-HTTP) | 297 |
| Safe Harbor Agreement | 290 | Secure/Multipurpose Internet Mail Extensions (S/MIME) | 297 |
| Safe Harbor Principles | 290 | Secure Shell (SSH) | 298 |
| SAINT | 291 | Secure Sockets Layer (SSL) | 298 |
| salt | 291 | Secure Windows Initiative (SWI) | 299 |
| SAM | 291 | Security+ | 300 |
| SAML | 291 | Security Accounts Manager (SAM) | 300 |
| Sam Spade | 291 | Security Administrator's Integrated Network Tool (SAINT) | 301 |
| sandbox | 291 | Security Assertion Markup Language (SAML) | 301 |
| Sandwich Test | 292 | Security Auditor's Research Assistant (SARA) | 301 |
| sanitized name | 292 | Security Configuration and Analysis | 302 |
| SANS Institute | 292 | security context | 302 |
| SARA | 293 | security descriptor | 302 |
| SAS | 293 | security identifier (SID) | 303 |
| SATAN | 293 | security log | 306 |
| scanning | 293 | security policy | 306 |
| screened subnet | 293 | security principal | 307 |
| screening router | 294 | security rollup package | 307 |
| script kiddie | 294 | security support provider interface (SSPI) | 307 |
| Sechole | 294 | security template | 308 |
| secondary data uses | 294 | security zone | 308 |
| secondary logon | 295 | SendIP | 309 |
| secret key | 295 | sensitive data | 310 |
| secret key encryption | 295 | SERPENT | 310 |
| secure attention sequence (SAS) | 295 | | |
| Secure Electronic Transaction (SET) | 296 | | |

| | | | |
|---|-----|---|-----|
| server certificate | 310 | Software Update Services (SUS) | 323 |
| server-gated cryptography (SGC) | 311 | source routing | 324 |
| service account | 311 | SP | 324 |
| service pack (SP) | 311 | spam | 324 |
| Service Release (SR) | 312 | SPAP | 325 |
| session hijacking | 312 | Spar | 325 |
| session key | 312 | special identities | 326 |
| SET | 313 | spoofing | 326 |
| SGC | 313 | spyware | 326 |
| SHA-1 | 313 | SR | 327 |
| SHA-2 | 313 | SSCP | 327 |
| shadow password file | 313 | SSH | 327 |
| shared secret | 313 | SSL | 327 |
| share-level security | 313 | SSL accelerator | 327 |
| Shiva PAP (SPAP) | 314 | SSPI | 327 |
| ShowAcls | 314 | Stacheldraht | 327 |
| ShowPriv | 314 | stealth scanning | 328 |
| SHS | 314 | stream cipher | 328 |
| S-HTTP | 314 | STPP | 329 |
| Sid2user | 314 | strong encryption | 329 |
| SIIA | 315 | Su | 329 |
| single sign-on (SSO) | 315 | subordinate CA | 329 |
| Sircam | 315 | SubSeven | 330 |
| site certificate | 316 | Sudo | 330 |
| Six/Four | 316 | SUID root | 330 |
| S/Key | 316 | superuser | 331 |
| Skipjack | 317 | SUS | 331 |
| Slammer | 318 | Swatch | 331 |
| Slashdot Effect | 318 | SWI | 331 |
| smart card | 318 | symmetric key | 331 |
| SMBRelay | 319 | symmetric key algorithm | 331 |
| SMB signing | 319 | symmetric key encryption | 332 |
| S/MIME | 319 | SYN flooding | 332 |
| Smurf attack | 319 | SYN scan | 333 |
| sniffing | 321 | Syskey | 333 |
| Snort | 321 | Syslog | 334 |
| social engineering | 321 | system access control list (SACL) | 334 |
| SOCKS | 322 | System Administrator Tool for Analyzing Networks (SATAN) | 335 |
| Software & Information Industry Association (SIIA) | 323 | System Security Certified Practitioner (SSCP) | 335 |
| software piracy | 323 | | |

T

TACACS 337
 TACACS+ 337
 TCPA 337
 Tcpdump 337
 Tcp_scan 337
 TCP session hijacking 338
 TCP SYN flooding 338
 TCP three-way handshake 338
 Tcp_wrapper 339
 TCT 339
 Teardrop attack 340
 Temporal Key Integrity Protocol (TKIP) 340
 Terminal Access Controller Access Control System
 (TACACS) 340
 TFN 341
 The Coroner’s Toolkit (TCT) 341
 threat 341
 ticket 341
 TKIP 342
 Tlist 342
 TLS 342
 Traceroute 342
 Tracert 343
 Transport Layer Security (TLS) 343

U

UDP scanning 353
 UDP tunneling 354
 URLScan 354

V

victim host 357
 virtual private network (VPN) 357
 virus 358
 virus protection software 359

W

wardialing 363
 wardriving 363
 Wassenaar Arrangement 364
 weak key 364
 Web anonymizer 364

trapdoor 343
 Trash2 344
 Tribal Flood Network (TFN) 344
 Tribal Flood Network 2000 (TFN2K) 344
 Trin00 345
 Trinoo 346
 Trinux 346
 Triple-A 346
 Triple DES 346
 Tripwire 346
 Trojan 347
 Trojan horse 348
 trust 348
 Trustbridge 348
 TRUSTe 349
 Trusted Computer System Evaluation
 Criteria (TCSEC) 349
 Trusted Computing Platform Alliance (TCPA) . . . 350
 Trustworthy Computing 350
 TSEnum 351
 tunneling 351
 two-factor authentication 352
 Twofish 352

User2sid 354
 UserDump 355
 user-level security 355

VLAD 360
 VPN 361
 vulnerability 361
 vulnerability scanner 361

Web bug 364
 web of trust 365
 Web permissions 365
 Web Services Security (WS-Security) 366
 WEP 366

| | | | |
|--|-----|--|-----|
| WFP | 366 | WinTrinoo | 372 |
| Whisker | 366 | Winux | 372 |
| white hat | 367 | Wired Equivalent Privacy (WEP) | 372 |
| Whois | 367 | workaround | 373 |
| Wi-Fi Protected Access | 368 | world-writable | 373 |
| Windows File Protection (WFP) | 368 | worm | 374 |
| Windows NT Challenge/Response | 369 | WPA | 374 |
| Windows Product Activation (WPA) | 369 | WRM | 374 |
| Windows Rights Management (WRM) | 370 | WS-Security | 374 |
| Windows Update | 370 | WWWhack | 375 |
| Winnuke | 372 | | |
| X | | | |
| X.509 | 377 | XMLENC | 380 |
| XACML | 378 | XML Encryption (XMLENC) | 380 |
| Xauth | 379 | XML Key Management Specification (XKMS) .. | 380 |
| Xhost | 379 | XML Signatures (XMLDSIG) | 381 |
| XKMS | 379 | Xscan | 381 |
| XMAS scan | 379 | Xterm | 382 |
| XMLDSIG | 380 | | |
| Y | | | |
| Ypgrab | 383 | | |
| Z | | | |
| Zap | 385 | | |
| zombie | 385 | | |
| Zombie Zapper | 385 | | |
| zone | 386 | | |
| <i>Appendix I: Applying Key Principles of Security</i> | | | 387 |
| <i>Appendix II: Understanding Your Enemy</i> | | | 395 |
| <i>Appendix III: Threats and Risk Assessment</i> | | | 405 |
| <i>Index</i> | | | 415 |

Acknowledgments

Thanks to my wife and business partner, Ingrid Tulloch, who was coauthor of our previous project with Microsoft Press, the *Microsoft Encyclopedia of Networking, Second Edition*. Thanks, Schatz, for contributing endless long hours of research and numerous helpful suggestions to assist me in writing this work. What a wife!

Thanks to the terrific team at Microsoft Press (mspress.microsoft.com) and nSight (www.nshintworks.com), including Jeff Koch, Sandra Haynes, Valerie Woolley, Susan McClung, Thomas Keegan, and Christina Palaia for their hard work and excellent suggestions. What a team!

Thanks to Neil Salkind of Studio B (www.studiob.com), my friend and literary agent who represented me in this project. What an agent!

Thanks to MTS Communications Inc. (www.mts.ca), for providing our company with Internet and Web hosting services. What a company!

And thanks finally to Ken and Bonnie Lewis, our best friends, and their four terrific kids, Karina, Alana, Sheri, and Vanessa, for encouraging us and praying for us as we worked on this project. What a family! And what friends!

Mitch Tulloch

www.mtit.com

Introduction

Welcome to the *Microsoft Encyclopedia of Security*, a general survey of computer security concepts, technologies, and tools. This work is intended to be a comprehensive, accurate, and up-to-date resource for students and practitioners, for policy and decision makers, for system and network administrators, and anyone else who works with computer, network, and information systems security.

What Is Computer Security?

Before we outline the scope of this work, let's begin with a simple question that has a surprisingly broad answer: What is computer security? We'll consider this question from seven different perspectives.

Threats and Vulnerabilities

Perhaps the most visible aspect of computer security today is the constant media attention surrounding vulnerabilities in software and the proliferation of viruses and other threats on the Internet. So one way of answering our question is that computer security is the science (and art) of dealing with threats and vulnerabilities.

Vulnerabilities generally arise from coding errors or bugs in software systems. This is not always the result of poor quality control of code development but instead is due to the ingenuity of hackers (good and bad) who explore and tinker with products looking for ways to circumvent security controls or simply see "what if" when unusual conditions or data arise. Some of the common vulnerabilities affecting software systems include

- Buffer overflows
- Input validation errors
- Uniform Resource Locator (URL) parsing errors

- Flawed password schemes
- Faulty implementation of Request for Comments (RFC) specifications
- Poorly configured default permissions
- Flawed security models
- Poor exception handling

Also garnering media attention these days are the various threats to computer security that are reported almost daily. Common threats that can affect the security of information system include

- Viruses, worms, Trojans, spyware, and other forms of malware
- Denial of service (DoS) and distributed denial of service (DDoS) attacks
- Other activity by black hat hackers, crackers, and script kiddies

Standards and Protocols

Another aspect of computer security is the various industry standards and protocols designed to provide confidentiality, integrity, and availability for data in information systems. Such standards may be

- Industry-wide efforts developed by such independent standards organizations as the Internet Engineering Task Force (IETF), the World Wide Web Consortium (W3C), and the Organization for the Advancement of Structured Information Systems (OASIS)
- Specifications developed by consortiums of vendors, such as the Wi-Fi Alliance, Liberty Alliance Project, and Trusted Computing Platform Alliance (TCPA)

- Standards developed by such government agencies and organizations as the National Institute of Standards and Technology (NIST), National Computer Security Center (NCSC), and National Security Agency (NSA)

Standards outlining specifications for commonly used security protocols are especially important because these protocols provide authentication, encryption, and other features that help keep computer networks secure. Some common examples of security protocols and mechanism include

- Network authentication protocols such as Kerberos and NT LAN Manager (NTLM)
- Protocols for secure exchange of data over the Internet such as Secure Sockets Layer (SSL) and Transport Layer Security (TLS)
- Protocols for wireless security such as Wired Equivalent Privacy (WEP), Temporal Key Integrity Protocol (TKIP), Wi-Fi Protected Access, and 802.11i
- Remote access protocols such as Password Authentication Protocol (PAP), Challenge Handshake Authentication Protocol (CHAP), and Microsoft Challenge Handshake Authentication Protocol (MS-CHAP)
- Protocols for secure virtual private networking such as Point-to-Point Tunneling Protocol (PPTP) and Layer 2 Tunneling Protocol (L2TP) combined with Internet Protocol Security (IPSec)
- Protocols for Authentication, Authorization, and Accounting (AAA) such as Remote Authentication Dial-In User Service (RADIUS), Terminal Access Controller Access Control System (TACACS), and TACACS+ Protocols for secure Extensible Markup Language (XML) Web services including Web Services Security (WS-Security), XML Encryption (XMLENC), and XML Signatures (XMLDSIG)
- International standards like ISO 17799 outlining best practices for information security

Hacking and Cracking

Another aspect of computer security involves the activities and exploits of individuals who seek to defeat it. These include hackers, crackers, phreakers, script kiddies, and the authors of viruses, worms, and Trojans. The term “hacker” is perhaps the most controversial one for security professionals, as it originally had no negative connotation and described individuals who were technically savvy and insatiably curious about everything having to do with computers. Today “hacker” is usually used pejoratively by the media, and to correct this influence the idea of “hats” was put forward, classifying hackers into white hats (good guys), gray hats (not so sure), and black hats (bad guys).

When we examine computer security from the perspective of hacking and cracking, we can talk about several issues, including

- General procedures used for breaking into computer networks, including footprinting, stack fingerprinting, enumeration, port scanning, address spoofing, session hijacking, elevation of privileges, root exploits, back channels, and log doctoring
- Common types of tools used to compromise systems, including sniffers, password crackers, rootkits, wardialers, vulnerability scanners, backdoors, remote administration tools (RATs), and malicious code
- Security tools that can be used for malicious purposes, ranging from sophisticated utilities such as Nmap, Fping, Snort, Netcat, and System Administrator Tool for Analyzing Networks (SATAN) to simple network troubleshooting tools such as Ping, Traceroute, Netstat, Finger, Nslookup, and Whois
- Popular exploits such as Smurf, Jolt, Bonk, Boink, Teardrop, Winnuke, Land, Fraggle, Trin00, and Stacheldraht, which can affect systems that are not properly patched with the latest fixes from vendors or exploit weaknesses in the fundamental design of Transmission Control Protocol/Internet Protocol (TCP/IP)

- Popular hacking and cracking Web sites, organizations, and media, such as *2600* magazine, Phrack, Attrition.org, Cult of the Dead Cow (cDc), and numerous others

Tools and Procedures

Yet another aspect of computer security is the tools and procedures used by businesses to protect the security of their systems, networks, and data. Security tools may either be commercial or free, proprietary or open source, and can be developed by legitimate security companies or borrowed from the black hat community.

At the simplest level are security technologies used to protect entry and control access to networks including

- Authentication mechanisms such as passwords, tickets, tokens, smart cards, and biometric systems
- Access control mechanisms such as discretionary access control (DAC) and mandatory access control (MAC)
- Permissions, rights, and other privileges that control system processes and tasks
- Auditing and logging mechanisms for recording security-related events and occurrences

Then there are tools and procedures used to protect networks from threats on the Internet, such as

- Firewalls and packet filtering routers
- Intrusion detection systems (IDSs) and honeypots
- Virus protection software and file system integrity checkers
- Vulnerability scanners and security auditing systems

Another important issue is the practices, procedures, and policies used to ensure network security, including

- Hardening systems and bastion hosts
- Penetration testing and security auditing
- Security policies and privacy policies

Organizations and Certifications

We've already mentioned Web sites frequented by black hats, but what about sites and organizations for legitimate security professionals? Numerous security advisory and support organizations exist that every security professional should be aware of, including

- SANS Institute
- Center for Internet Security (CIS)
- CERT Coordination Center (CERT/CC)
- Forum of Incident Response and Security Teams (FIRST)
- Microsoft Security Response Center (MSRC)

Certifications for security professionals are a way of ensuring one's skills are up to date and stand out from the crowd. Some of the popular certifications available include

- Certified Information Systems Security Professional (CISSP)
- System Security Certified Practitioner (SSCP)
- Certificate Information Systems Auditor (CISA)
- Global Information Assurance Certification (GIAC)

Cryptography

Ensuring the privacy and confidentiality of data stored on and transmitted between information systems is another important aspect of computer security, and this is built on the foundation of cryptography, the branch of mathematics concerned with procedures for encrypting and decrypting information. Every security professional should be familiar with the basics of this field, including knowledge of

- Public key cryptography with its elements of certificates, signatures, certificate authorities (CAs), and public key infrastructure (PKI)

- Secret key cryptography based on block ciphers, stream ciphers, one-time passwords (OTPs), session keys, and other constructs
- Encryption algorithms such as Blowfish, Rijndael, Twofish, MD5, RC4, Skipjack, Diffie-Hellman, and RSA
- Encryption standards such as Advanced Encryption Standard (AES), Data Encryption Standard (DES), and Digital Signature Standard (DSS)
- Encryption schemes for secure messaging such as Secure Multipurpose Internet Mail Extensions (S/MIME) and Pretty Good Privacy (PGP)
- Methods for cracking keys and passwords including brute-force and dictionary attacks

Legal Issues

Finally, there are the legal issues surrounding computer and information systems security. These include

- Software piracy and the technologies, laws, and initiatives designed to prevent it
- Privacy laws relating to what companies can do with personally identifiable information (PII) collected from individuals
- International agreements such as the Wassenaar Arrangement, which covers export control of dual-use technology such as encryption
- Technologies and initiatives for making computer systems more trustworthy such as Microsoft Corporation's Trustworthy Computing Initiative, the Trusted Computing Platform Alliance (TCPA), and Microsoft Corporation's Next-Generation Secure Computing Base for Windows, formerly called Palladium

Who This Work Is For

I think you can already see that the scope of this book is broad and wide, as an encyclopedia should be. This breadth of coverage is needed because computer security affects many different fields and requires that

security professionals have broad knowledge and skills concerning computer networking, operating systems, the Internet, code development, cryptography, incident response, forensics, and local, federal, and international law. What a business to be in! But what exciting times! Never before have professionals with computer security expertise been in so much demand to protect companies against a seemingly exponential rise in threats, attacks, and exploits against their systems and data.

The computer security field is growing in leaps and bounds, and this book is an attempt to provide a snapshot of everything and anything that has to do with the field. Future editions of this book are likely to include even more information as new exploits, tools, standards, and technologies are developed by both security professionals and black hat hackers. This present edition, however, is likely to be an invaluable reference work for the following kinds of individuals:

- Computer security professionals and practitioners in business, industry, government, and the military
- System and network administrators who work mainly with Windows, UNIX, and Linux platforms
- Students considering and/or pursuing academic degrees in computer science or industry certifications in information security
- Corporate policy makers, decision makers, and executives involved in MIS (Management Information Services), IS (Information Services), and IT (information technology)

How This Work Is Organized

The topics in this book are listed in alphabetical order and range from a few sentences to several paragraphs in length. Most articles include a definition and brief overview of a subject, while longer articles may include a description of how a technology is implemented, issues concerning its use, commercial and free products and services available for it in the marketplace, and brief notes or tips. Figures and diagrams have been included to explain some concepts, and URLs for finding further information on the subject have been provided. Most

articles also finish with cross-references to related topics found elsewhere in this book.

Disclaimer

The information contained in this work has been obtained from sources believed to be reliable. Although both the author and Microsoft Press have made every effort to be fair and accurate, neither the author nor the publisher assume any liability or responsibility for any inaccuracy or omissions contained within this book, or for any loss or damage resulting from application of the information presented therein. In other words, the information provided in this book is presented on an “as is” basis. Mention of organizations, vendors, products, and services in this work are not to be viewed as endorsements by either the author or by Microsoft.

Comments and Questions

If you have comments, questions, or suggestions regarding this encyclopedia, please direct them to Microsoft Press at MSPInput@microsoft.com or at the following postal address:

Microsoft Press
Attn: *Microsoft Encyclopedia of Security* Editor
One Microsoft Way
Redmond, WA 98052-6399
USA

Please note that product support is not offered through the above addresses.

You can also contact the author of this work directly through his Web site (www.mtit.com).

Numbers and Symbols

3DES

Also called Triple DES or EDE (encrypt, decrypt, encrypt), a secret key encryption algorithm based on repeated application of the Data Encryption Standard (DES).

Overview

3DES works by applying the DES algorithm three times in succession to 64-bit blocks of plaintext. It does this by using two independent 56-bit keys (K1 and K2) applied in the following manner:

- 1 Encryption with K1
- 2 Decryption with K2
- 3 Encryption with K1

Since this three-stage encryption process uses two different 56-bit keys, it has an effective key length of $2 \times 56 = 112$ bits, which is 2^{56} times stronger than DES. This means if you could crack a DES message in one hour, it would take 8 trillion years to crack 3DES using the same method! To decrypt a block of 3DES ciphertext you use the following procedure:

- 1 Decryption with K1
- 2 Encryption with K2
- 3 Decryption with K1

By setting $K1 = K2$ in the preceding encryption algorithm, 3DES defaults to DES in operation. This issue of backward compatibility with DES is one of the reasons that EDE is used instead of the equally plausible EEE (encrypt, encrypt, encrypt) for 3DES.

3DES is defined by ANSI standard X5.92 and complies with Federal Information Processing Standards (FIPS) 140-1 Level 1.

Implementation

3DES is commonly used to implement Internet Protocol Security (IPSec) encryption in firewalls and routers for building secure virtual private networks (VPNs). Due to its licensing requirements, 3DES is generally not included in enterprise software and must be obtained as an add-in, such as the *Microsoft Windows 2000 High Encryption Pack*. Support for 3DES is included in Microsoft Windows XP Professional for the Encrypting File System (EFS).

3DES is also on the way to replacing DES as a new standard for the electronic payment industries and is used to secure electronic transactions between banking and credit agencies and point-of-sale (POS) devices and automatic teller machines (ATMs). Both MasterCard and Visa, for example, are implementing end-to-end 3DES solutions for electronic funds transfers and payments.

Issues

The main drawback with 3DES is that it is slow because of the iterated nature of its algorithm. In principle, you could make DES even more secure by performing more than three iterations, but in practice the performance penalty is too great.

Notes

Some nonstandard implementations of 3DES employ three keys instead of two, with the difference being that the third iteration performs encryption using K3 instead of K1. The result is that the effective key length in these implementations of 3DES is 168 bits.

See Also: *Data Encryption Standard (DES)*

802.1x

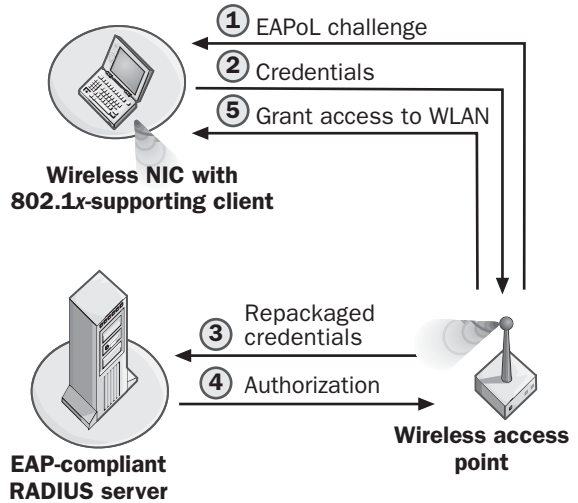
An IEEE standard for port-based network access control, particularly useful for securing 802.11 wireless local area networks (WLANs).

Overview

802.1x is a cornerstone of the Robust Security Network (RSN) initiative of the Institute of Electrical and Electronics Engineers (IEEE) and the emerging 802.11i standard. The 802.1x standard works by providing port-based access control to both wired and wireless networks. It is built on two standard network security protocols:

- **Extensible Authentication Protocol (EAP):** An extension to Point-to-Point Protocol (PPP) that is defined in RFC 2284 and allows an arbitrary authentication method to be negotiated during PPP session initialization
- **Remote Authentication Dial-In User Service (RADIUS):** A client/server security protocol that provides Authentication, Authorization, and Accounting (AAA) and is defined in RFCs 2138 and 2139

The 802.1x standard defines three types of entities: supplicant, authenticator, and authentication server. In a typical scenario, the supplicant is a remote user's laptop that has an 802.1x-compliant wireless network interface card (NIC) installed, while the authenticator is an 802.1x-compliant access point and the authenticator an EAP-compatible RADIUS server. When an authenticator detects a new supplicant that needs to be authenticated, it sends the supplicant a challenge message encapsulated using the EAP-over-LAN (EAPoL) security protocol defined by 802.1x. The supplicant then sends its credentials to the authenticator, which repackages them as a RADIUS message and forwards this to the authentication server. The authentication server then compares the submitted credentials against its authentication database or forwards them to another authentication server. Once the client has been authenticated, the authentication server informs the authenticator, which then allows the supplicant to access the network. The authentication server can also distribute a session key to the supplicant through the authenticator, and the supplicant and authenticator can then use this key for encrypted communications between them.



802.1x. How 802.1x authentication works.

When used in a switched Ethernet environment, the authenticator is typically a switch or router that enables a specific physical port to allow the client access to the network. In this scenario, 802.1x is referred to as providing port-based access control since it provides network access only through ports for which the client has been authenticated.

Implementation

There are several ways of deploying secure WLANs using 802.1x. The simplest scenario involves employing one or more RADIUS servers using a central authentication database (typically Lightweight Directory Access Protocol [LDAP]– or SQL–based) and managing wireless clients at a single site. In a distributed environment the authentication database can be replicated across multiple sites, and the RADIUS servers and access points for each site can be autonomous or managed centrally.

A number of vendors have started to incorporate 802.1x support into their switches, RADIUS servers, access points, and network adapters, including Cisco, Hewlett-Packard, Microsoft, Enterasys, Funk Software, Wind River, and several others. Interoperability issues between offerings from different vendors depend largely on how 802.1x authentication is being implemented. For example, Cisco has created an authentication method called

Lightweight Extensible Authentication Protocol (LEAP, or Lightweight EAP) that represents an interim step toward full 802.1x functionality. Other common authentication methods supported by EAP and used in 802.1x include MD5, Transport Layer Security (TLS), and Tunneler TLS (TTLS).

Issues

Because of its built-in security and support for AAA, 802.1x holds promise for simplifying how Internet service providers (ISPs) provision wireless Internet access in public spaces. However, researchers at the University of Maryland recently discovered that the present 802.1x standard is vulnerable to certain kinds of session hijacking or man-in-the-middle attacks. The cause of the problem is that 802.1x was designed mainly to secure the infrastructure (the WLAN access points and the wired network behind it) and not the clients themselves. A workaround is to supplement your WLANs with encrypted virtual private network (VPN) security, but in the meantime enterprises should use caution in deploying 802.1x as a panacea for their network security problems.

Notes

Microsoft also includes built-in support for 802.1x in its Microsoft Windows XP Professional and Windows Server 2003 operating systems.

See Also: 802.11i, *Extensible Authentication Protocol (EAP)*, *Remote Authentication Dial-In User Service (RADIUS)*, *wireless security*

802.11i

An emerging standard specifying security enhancements for the 802.11 wireless networking.

Overview

The development of 802.11i was motivated by serious flaws discovered in the earlier 802.11 security protocol called Wired Equivalent Privacy (WEP). The result was the Robust Security Network (RSN) initiative developed by the Institute of Electrical and Electronics Engineers (IEEE), of which the emerging 802.11i standard is the cornerstone. The 802.11i standard provides enhancements to the security of existing wireless local area network (WLAN) standards, including 802.11a,

802.11b, and 802.11g. These security enhancements include new authentication procedures, strengthened encryption schemes, and dynamic key allocation, all with the goal of ensuring WLANs are as secure as wired LANs. The 802.11i standard will include support for 802.1x port-based access control, Temporal Key Integrity Protocol (TKIP), Advanced Encryption Standard CBC-MAC Protocol (AES-CCMP) encryption, secure fast handoff, and secure deauthentication and disassociation.

The 802.11i standard is expected to be finalized by the IEEE in 2003. As an interim measure until the final 802.11i standard becomes available, the Wi-Fi Alliance has released an upgrade for WEP called Wi-Fi Protected Access (WPA), which is forward-compatible with 802.11i and can be implemented easily in existing wireless-networking equipment through firmware upgrades.

See Also: 802.1x, *Advanced Encryption Standard (AES)*, *Temporal Key Integrity Protocol (TKIP)*, *Wired Equivalent Privacy (WEP)*, *wireless security*

2600

A magazine devoted to hacking, cracking, and freedom of information.

Overview

Also called the *Hacker Quarterly*, *2600* is a nonprofit magazine edited by Eric Corley, who uses the pen name Emmanuel Goldstein after a character who leads an underground movement in *1984*, a novel by George Orwell. Since 1984, this magazine has been the best-known public voice in the underground hacking community and is available from bookstores and magazine stands everywhere. The magazine is widely read by security professionals and is often a valuable source of information about popular exploits and the tools and methods used to accomplish them. The name *2600* comes from the frequency of a whistle that used to be included in boxes of Captain Crunch cereal. It turned out this was also the frequency used by the old analog Plain Old Telephone System (POTS) for initiating operator-controlled calls, and in the early 1980s some hackers

discovered they could use the Captain Crunch whistle to make free long-distance calls, an activity called **phreaking** (phone hacking).

The 2600 team has also done other projects, including producing *Freedom Downtime*, a feature-length film about convicted cracker Kevin Mitnick, that aim to counter what hackers feel are unfair media portrayals of their subculture.

2600 also has a Web site (www.alt2600.com), and there is a series of newsgroups (alt.2600.*) used by the hacker community that contains a useful FAQ on security issues.

See Also: *hacker*

A5

A family of algorithms that is used to encrypt Global System for Mobile Communications (GSM) cellular communications.

Overview

A5 is a stream cipher that comes in two flavors: a “strong” form (A5/1) that is proprietary and a “weak form” (A5/2) that is in the public domain. In 1999, however, a crack for A5/1 was developed by Adi Shamir (the S in the Rivest-Shamir-Adleman or RSA algorithm) that can be run in real time using only a standard PC. This cryptographic feat meant that the privacy of cellular phone conversations of over 200 million users of GSM systems in Europe and Asia was endangered. As a result, a joint working party between the GSM Association Security Group and the 3rd Generation Partnership Project (3GPP) developed a newer and more secure algorithm called A5/3, which is based on the Kasumi algorithm and which is intended to ensure the privacy of GSM communications.

See Also: *cracking, cryptography, RSA algorithm, stream cipher*

AAA

Stands for Authentication, Authorization, and Accounting, a security framework for controlling access to network resources.

See: *Authentication, Authorization, and Accounting (AAA)*

acceptable use policy (AUP)

A policy that defines appropriate use of computing resources for a company or organization.

Overview

Developing an acceptable use policy for your network and communicating it clearly to employees are essential

for any good security policy. An acceptable use policy should generally have three goals:

- To communicate clearly which types of activities are not acceptable and why
- To provide legal notice concerning these unacceptable activities so that violators can be punished accordingly
- To protect the company from legal action for alleged violations of privacy

Examples of proscribed actions might include the following:

- Using another employee’s user account with or without that person’s permission
- Reading, copying, or altering files belonging to another employee without that person’s permission
- Using the company’s computing resources for personal gain
- Sending unsolicited commercial e-mail (UCE), more commonly known as spam, from your machine to others inside or outside the company
- Engaging in such practices as mail bombing that interfere with a user’s e-mail, regardless of whether or not the user is an employee of the company
- Downloading pornography from the Internet and storing it on your computer
- Releasing confidential information concerning the company or its network to outside parties
- Downloading and installing software on your computer without the knowledge or permission from the Helpdesk

Acceptable use policies should always be

- Clearly and concisely written

- Posted visibly in common areas such as the lunchroom
- Handed to new employees during their orientation period

Implementation

A typical outline for an acceptable use policy might look like this:

- 1 Introduction
- 2 Who must abide by this policy
- 3 What is acceptable conduct
- 4 What is unacceptable conduct
- 5 Consequences of violating this policy
- 6 Summary

See Also: security policy

access

Has a variety of meanings relating to privacy and the right to use resources.

Overview

In a general sense, the concept of access is related to privacy and has to do with an individual's ability to view, modify, and contest the accuracy of personal information collected about the individual. In this respect, access reflects the Fair Information Practices defined by the Privacy Act of 1974, legislation that protects personal information collected by the U.S. government.

In computer networking, **access** refers to the ability of an entity (typically a user or process) to connect to a resource (a Web site, database, shared folder, or some other network resource). Access can be managed several ways; for example, access to network resources is typically controlled by permissions implemented using access control lists (ACLs) that allow or deny various users and groups different levels of access to resources. Access to a network itself, such as a remote intranet, is often controlled by firewalls that use access lists allowing or denying access based on source IP address, port number, or Domain Name System (DNS) domain name. Finally, access to a local network is usually

controlled through authentication using a logon process that requires a user to submit credentials (user name and password) before gaining access to resources on the network.

See Also: access control list (ACL), access list, Fair Information Practices (FIP), firewall, permissions, personally identifiable information (PII)

access control

Any mechanism for controlling which resources a user can access or tasks users can perform.

Overview

Once a user has been authenticated and logged on to a system or network, access control takes over to enforce what the user is able to do. The most common way of implementing access control is using access control lists (ACLs) that specify a list of security protections applied to an object such as a file, folder, or process. Access control can also be managed using policy managers such as Cisco Secure Policy Manager (CSPM) on Cisco firewalls, virtual private network (VPN) gateways and intrusion detection systems (IDSs), or Group Policy on Microsoft Windows platforms.

There are two basic approaches to implementing access control:

- **Discretionary access control (DAC):** This method allows users to specify who can access resources they own and what level of access others have to these resources. DAC is used on Microsoft Windows platforms and most implementations of UNIX or Linux.
- **Mandatory access control (MAC):** In this method, the administrator controls access, usually by specifying a set of rules. MAC is more secure but less flexible than DAC, and most versions of UNIX and Linux support MAC in addition to DAC.

Implementation

Some examples of ways to configure DAC on different platforms include the following:

- Using the Permissions page of a file or folder's properties sheet on Microsoft Windows platforms to configure Windows NTFS permissions on the file or folder

- Using .htaccess files to control access to directories on an Apache server running on UNIX or Linux
- Configuring access lists on a Cisco router or access server

See Also: access, access control list (ACL), discretionary access control (DAC), .htaccess, mandatory access control (MAC), permissions

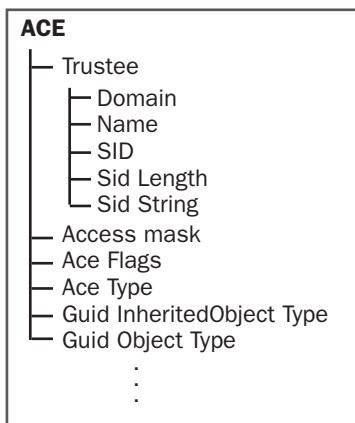
access control entry (ACE)

An entry in an access control list (ACL).

Overview

An ACE is a data structure that contains two things:

- A security identifier (SID) identifying the security principal whose access to a resource is being controlled by the entry.
- A set of access rights defining which operations the principal can perform on the resource. Examples of such operations might be read, open, create, execute, and so on. These operations can be either allowed or denied for the security principal.



Access control entry (ACE). Structure of a Win32 ACE.

See Also: access control list (ACL), security identifier (SID)

access control list (ACL)

A data structure associated with an object that specifies which users are authorized to access the object and what level of access they can have.

Overview

Access control lists (ACLs) are used on Microsoft Windows platforms to control access to securable objects such as files, processes, services, shares, printers, or anything else that has a security descriptor assigned to it. ACLs are composed of a series of access control entries (ACEs) that specify which operations each security principal (user or group) can perform on the object.

There are two types of ACLs on Microsoft Windows platforms:

- **Discretionary ACL (DACL):** These are ACLs that contain ACEs allowing or denying access to objects.
- **System ACL (SACL):** These can do the same thing as DACLs but can also generate auditing information using the security audit ACE.

Since an ACL must specify the actions that each user can perform on the object to which the ACL is attached, ACLs can rapidly grow in size as the number of users increases. To overcome this scaling problem, users can be assigned to groups, and these groups can then be assigned different privileges using ACLs. Special groups such as Everyone or World (depending on the platform) can be used to grant or deny privileges to all users using a single ACE.

Implementation

When a user account is created on a Microsoft Windows platform, it is assigned a security identifier (SID) that uniquely identifies the account to the operating system. When the user logs on using this account, an access token is created that combines the SID, the groups to which the account belongs, and a list of privileges for the account. This token is then copied to all processes and threads owned by the account. When the user tries to access an object secured using an ACL, the token is compared with each ACE in the ACL until a match is found and access is either allowed or denied.

On UNIX platforms, access to file system objects has traditionally been controlled using the user/group/other mechanism that is implemented with the change mode (Chmod) command. This is a rather coarse-grained approach to access control, however, since it is designed only to let you grant access to yourself, to groups to which you belong, or to everyone. If, however, you want to grant access to certain users only, this can be done only by creating a new group for these users, an approach that can cause the number of groups on a UNIX network to proliferate excessively. As a result, modern UNIX platforms such as HP-UX 11i also support implementing ACLs for more granular control of access to file system resources. These ACLs are implemented as text files that can be viewed and modified using the get ACL (Getacl) and set ACL (Setacl) commands or third-party utilities such as CalcMgr.

Other operating system platforms on which ACLs can be implemented include Novell NetWare, OpenVMS, and Solaris.

Notes

The term **access control list** has another meaning on Cisco devices and has to do with allowing or prohibiting traffic from passing through a network device such as a router or firewall. In this context, access control lists are generally referred to as **access lists** instead.

See Also: *access control entry (ACE), discretionary access control list (DACL), security descriptor, system access control list (SACL)*

access list

A list used for controlling traffic on Cisco devices.

Overview

Access lists are the Cisco equivalent of access control lists (ACLs) on Microsoft Windows platforms, except that while ACLs are generally used to control access to network objects (files and other resources), access lists control the flow of packets through a router or firewall. Access lists do this by examining various criteria such as the source address, destination address, or port number within a packet's header and then either forwarding the packet or blocking it from being passed through the device.

Access lists provide a number of important functions including these:

- **Security:** Access lists can be configured to block traffic from source addresses of malicious systems or networks.
- **Traffic flow:** Access lists can be used to filter certain types of traffic to prevent portions of a network from being overwhelmed with unnecessary traffic or to allow certain hosts access to specific portions of a network.

Implementation

Access lists on Cisco routers can be created and configured for each router interface. For Internet Protocol (IP), two separate access lists can be applied to a given interface, one for inbound traffic and the other for outbound, to provide greater control of traffic flow through the router. Each access list applied to an interface is defined by a unique name or number and can contain multiple access list statements. The order in which access list statements are added is important since these statements are processed in sequence. Also, you cannot reorder the statements within an access list; if you make a mistake and enter a statement out of order, you have to create the list over again. Note that at the end of each list of statements is an implied "deny all traffic" statement so that when a packet doesn't match any of the explicit statements, it is prohibited from passing through the configured interface.

Typically, when updating access lists on a Cisco router, you will create your lists on a Trivial File Transfer Protocol (TFTP) server and then download them to your router. The advantage of this approach is that you can create your access list statements using a text editor, which lets you reorder them as necessary before uploading them to the router.

Notes

Another name for an access list is **filter**, a term that is commonly used in reference to routers.

See Also: *access, access control list (ACL)*

access mask

A value specifying which rights are allowed or denied in an access control entry (ACE).

Overview

On Microsoft Windows platforms, access rights specified by ACEs are arranged in a specific order determined by a 32-byte access mask. The format specified by an access mask is as follows:

- Low-order bytes 0 through 15 are for object-specific access rights (varies with types of objects).
- Bytes 16 through 22 specify standard access rights (applies to most object types).
- Byte 23 specifies right to access system ACL (SACL).
- Bytes 24 through 27 are reserved.
- Bytes 28 through 31 specify generic access rights.

See Also: *access control entry (ACE)*

access token

A data structure containing the security information for a logon session.

Overview

When a user logs on to a Microsoft Windows–based network, the system creates an access token that determines which system tasks the user is able to perform and the securable objects the user is able to access. The access token contains information that identifies the user, the groups to which the user belongs, and the user’s level of privileges. The system attaches a copy of this token to every process executed on behalf of the user and uses the token to identify the user when threads interact with securable objects or attempt to perform system tasks requiring privileges.

Implementation

Access tokens include the following information:

- Security identifier (SID) for the user account
- SIDs for groups to which the user belongs
- Logon SID identifying the current logon session
- List of privileges held by the user account or groups to which the user belongs
- Owner SID
- SID for the primary group
- Default discretionary ACL (DACL) used by the operating system when the user creates a securable object without specifying a SID
- Source of the token
- Whether the token is a primary or impersonation type
- Optional list of restricting SIDs
- Current impersonation levels
- Other statistics

There are two types of access tokens:

- **Primary token:** A token created by the executive and assigned to a process to represent the default security information for that process. Primary tokens are used when process threads interact directly with securable objects.
- **Impersonation token:** A token that captures the security information of a client process to enable a server to “impersonate” a client process in security operations. Impersonation lets threads interact with securable objects using the client’s security context.

See Also: *access control*

account lockout

The condition in which a user account is disabled automatically for security reasons.

Overview

Account lockout protects user accounts by disabling an account temporarily when a specified number of failed logon attempts occur within a predetermined interval of time. The assumption behind this practice is that numerous incorrect logons within a short period of time may indicate an unauthorized person attempting to access the network. Another possibility, of course, is that the

user has simply forgotten his or her password, and this is often the case when companies require users to employ long, complex passwords. When a user's account becomes locked out, the user can either wait for the lockout condition to be reset automatically after a predetermined interval or contact an administrator or support person to reset the account manually.

Implementation

Most operating systems implement some form of account lockout. On Microsoft Windows platforms, account lockout is implemented using a policy-based method known as **account lockout policy**.

See Also: *account lockout policy, password*

account lockout policy

A policy that controls how account lockout is implemented for a system or network.

Overview

Account lockout policies are used on Microsoft Windows platforms to protect user accounts from attempts at unauthorized access. These policies are controlled by Active Directory service and define how the following settings are configured:

- **Account lockout duration:** This defines how long an account remains unavailable to the user once it is locked out. Possible values range from 0 to 99,999 minutes, with a value of 0 indicating the account remains locked out until manually reset by an administrator.
- **Account lockout threshold:** This specifies the number of failed logon attempts that must occur in order for the account to be locked out. Possible values range from 0 to 999 logon attempts.
- **Reset value:** This specifies the time interval after which the failed logon counter is reset to zero if the account is not locked out. For example, if this value is configured as 5 minutes, the counter keeping track of failed logon attempts will be reset to zero 5 minutes after the last failed logon, provided the account lockout threshold has not yet been exceeded. The purpose of this value is to provide the user who has forgotten his or her password with breather time to

try to remember the password before having the account locked out.

See Also: *account lockout, account policy*

account policy

A policy that controls the security of user accounts.

Overview

Account policies are used on Microsoft Windows platforms to protect user accounts in an Active Directory service scenario. Windows platforms basically support three types of account policies:

- **Account lockout policy:** This defines which actions will be taken after a specified number of failed logon attempts occur within a predetermined window of time.
- **Kerberos policy:** This specifies certain Kerberos parameters, including maximum ticket lifetime and clock synchronization tolerances between clients and servers.
- **Password policy:** This defines password restrictions such as minimum password length, password complexity requirements, and so on.

See Also: *account lockout policy, Kerberos policy, password policy*

ACE

Stands for **access control entry**, an entry in an access control list (ACL).

See: *access control entry (ACE)*

ACK storm

Generation of large numbers of Transmission Control Protocol (TCP) acknowledgment (ACK) packets, usually because of an attempted session hijacking.

Overview

ACK storms usually result when an intruder tries to hijack a TCP session by injecting spoofed packets into the session. What usually happens is that an intruder sends a forged packet to host B during a TCP session between hosts A and B. If the forged packet has the

correct TCP sequence number, host B responds by sending an acknowledgment (ACK) to host A, thinking that it was host A that sent the packet. Host A notices that host B has acknowledged a nonexistent packet (as far as it is concerned) and responds by returning the acknowledgment to host B along with what it thinks is the correct sequence number. Host B decides that host A has sent it a packet out of sequence and immediately responds with an acknowledgment to this effect, which causes host A to respond, which causes host B to respond, and so on. This flood of ACKs continues until the network becomes overloaded so that packets are dropped and the session times out.

If your packet sniffer or intrusion detection system (IDS) detects an ACK storm under way it is likely that your network is under attack. An intruder may be attempting to hijack a TCP session, usually something dangerous such as a telnet session, which can allow the intruder to execute arbitrary code on your hosts. If you don't have a sniffer or IDS running but your users begin to complain that the network has slowed down, an ACK storm is one possibility you should investigate immediately.

The potential for ACK storms is inherent within the operation of the TCP protocol and is one reason why you generally should never allow telnet sessions between remote users and your network. A better solution than telnet is to use Secure Shell (SSH), which can provide secure communications using 3DES or International Data Encryption Algorithm (IDEA) encryption.

See Also: 3DES, intrusion detection system (IDS), Secure Shell (SSH), sniffing

ACL

Stands for access control list, a list of security protections that applies to an object.

See: access control list (ACL)

Acldiag

A *Microsoft Windows 2000 Server Resource Kit* command-line tool for troubleshooting permissions problems.

Overview

Acldiag can be used to diagnose permissions problems with objects in Active Directory service. It does this by writing the information in the object's access control list (ACL) to a text file that can then be examined. When you use this tool, the only ACL entries that are written are those to which your currently logged on user account has rights.

For More Information

You can obtain the *Microsoft Windows 2000 Server Resource Kit* from Microsoft Press.

ACPA

Stands for Anticybersquatting Consumer Protection Act, a U.S. federal law that gives trademark owners legal remedies against domain name cybersquatters.

See: Anticybersquatting Consumer Protection Act (ACPA)

ACSA

Stands for Applied Computer Security Associates, a nonprofit association of computer security professionals whose goal is improving the understanding, theory, and practice of computer security.

See: Applied Computer Security Associates (ACSA)

ACSAC

Stands for Annual Computer Security Applications Conference, an annual conference on computer security organized and sponsored by Applied Computer Security Associates (ACSA).

See: Annual Computer Security Applications Conference (ACSAC)

Active Directory

The directory service for Microsoft Windows platforms.

Overview

While Active Directory service provides enterprise-level directory services for Windows operating systems, it is also important from the standpoint of network security

because Active Directory provides secure storage for credentials of users and computers. Active Directory is also responsible for authenticating users when they log on to the network and for authenticating computers when the network is started. Active Directory is implemented using domain controllers, special servers that contain copies of the directory database and make possible the single sign-on (SSO) feature that allows users to access the network from computers residing in any domain in the forest. Active Directory supports a variety of authentication methods including Kerberos, NTLM, and certificate-based Public Key Infrastructure (PKI).

For More Information

For more general information about Active Directory, see the *Microsoft Encyclopedia of Networking, Second Edition*, or the *Microsoft Windows 2000 Server Resource Kit*, both available from Microsoft Press.

See Also: authentication, Kerberos, NTLM, Public Key Infrastructure (PKI), single sign-on (SSO)

adaptive proxy

Also called **dynamic proxy**, an enhanced form of application-level gateway.

Overview

Application-level gateways are firewalls that look deep into packets to filter them according to Open Systems Interconnection (OSI) application-layer protocol information. For example, an application-layer gateway might be configured to accept all Hypertext Transfer Protocol (HTTP) GET requests except for those having certain values in their HTTP headers, such as those using cookies. The problem with such application-level gateways is that examining the application-layer information in every packet requires a great deal of processing power, which tends to make such firewalls relatively slow.

One solution is the adaptive proxy approach, which involves having the firewall examine application-layer information for only the initial packets of a Transmission Control Protocol (TCP) session. Once the session is determined to be legitimate, the firewall then stops looking inside the remaining packets and simply forwards them through the network layer. The advantage of the application proxy approach is improved speed

over traditional application-layer gateways. The disadvantage is a decrease in security since an intruder that hijacked a legitimate TCP session would have its packets passed through the firewall unhindered.

See Also: application-level gateway, firewall

Adaptive Security Algorithm (ASA)

A Cisco algorithm for managing stateful connections for PIX Firewalls.

Overview

The Adaptive Security Algorithm (ASA) uses security levels to describe whether a given firewall interface is inside (trusted) or outside (untrusted) relative to other interfaces. ASA security levels range from 0 (lowest) to 100 (highest), with 100 being the default for inside interfaces and 0 being the default for outside interfaces. Security levels 1 through 99 are typically used for interfaces connected to the demilitarized zone (DMZ).

In a typical configuration, inside interfaces are configured with higher security levels than outside ones. Packets entering the firewall through an interface with a higher security level can exit freely through one with a lower security level, while packets passing in the reverse direction are controlled by access lists or through a conduit.

See Also: access list, demilitarized zone (DMZ), firewall

address-based authentication

An authentication method that employs a network address as the credentials.

Overview

Address-based authentication was one of the first authentication methods employed on computer networks and was commonly used on UNIX and VMS platforms. It worked on the principle of One User, One Machine and used the machine's network address to identify the user to other machines on the network.

Address-based authentication is not often used on broadcast-based networks such as Ethernet or Token Ring, because it is fairly simple to impersonate a net-

work address. Combined with a password, however, address-based authentication is just as secure (and perhaps more convenient) than more common authentication methods that employ a user name/password combination.

See Also: *authentication, spoofing*

address munging

Any method of disguising an e-mail address to make it hard for Web crawlers to find.

Overview

Spam, or junk e-mail, is an ever-growing problem for most e-mail users. One way to avoid being added to junk mailing lists is to avoid posting your e-mail address to USENET newsgroups or on public World Wide Web sites. The reason is that companies that compile junk e-mail lists often employ Web crawlers, software that scans the Internet for e-mail addresses and automatically adds them to the list.

Of course, there are times when you would like to advertise your address publicly so others can send you mail, such as when you are trying to sell something on the Internet. One solution in this instance is to modify your e-mail address so that it becomes an invalid address as far as Web crawlers are concerned, but one that can still be recognized by human recipients as containing a valid e-mail address. This solution is commonly called address munging, and some munged examples of the e-mail address `mtulloch@microsoft.com` could be

- `mtulloch@nospam.microsoft.com`
- `mtulloch@remove-me.microsoft.com`
- `mtulloch@microsoft.com.nospam`
- `mtulloch@WGQ84FH7microsoft.com`
- `mtulloch AT Microsoft DOT com`

Note that if you munge your e-mail address in a newsgroup posting, you may need to post a notice such as, “Please do not reply to this message. If you want to

send me e-mail, remove ‘nospam’ from my address first” or something similar.

Notes

The term **munge** probably originated at MIT in the 1980s and means “mash until no good.” An alternative (and more trendy) recursive reading would be “munge until no good.”

See Also: *spam*

address spoofing

Usually simply called **spoofing**, the process of falsifying the source of Media Access Control (MAC) or Internet Protocol (IP) addresses of packets being sent on an Ethernet network.

See: *spoofing*

Administrator

The most powerful account on a Microsoft Windows-based network.

Overview

The Administrator account is a local user account that is created when Microsoft Windows Server 2003 (or Windows XP, Windows 2000, or Windows NT) is installed on a system. The Administrator account has basically all possible rights a user account can have and is a powerful account similar to root on UNIX platforms. As a result, an important aspect of security on Windows platforms is to protect the Administrator account from misuse. This can be accomplished in several ways:

- Give the account a strong password that is difficult to guess or crack.
- Rename the account from Administrator to something else that is hard to guess.
- Never use the account for performing ordinary user tasks such as checking e-mail or browsing the World Wide Web; use a second Domain Users account for this purpose.
- Avoid logging on using the account to perform routine network tasks; instead use secondary logon (the

Runas command) to execute programs using Administrator credentials while logged on to your console using your ordinary user account.

- Never use the Administrator account as credentials for mapping a virtual directory alias to a remote network share in Internet Information Services (IIS).

See Also: *rights, root*

Admintool

A tool on the Solaris platform used for configuring password policies for users.

Overview

Admintool is used to specify password expiration times and warnings, minimum password length, and whether passwords must be changed at first login. It is also used to specify the search order for authentication credentials and whether an account can be used for interactive login or su or both.

See Also: *authentication, password, su*

admnlock

A command-line tool in the *Microsoft Windows 2000 Server Resource Kit* that can be used to protect the Administrator account from abuse.

Overview

Admnlock can be used to lock out the Administrator account from being used for network logons. The account can still be used for interactive logons on the local machine, however, since this account is essential for administering the machine. Admnlock can be used only on machines running Windows 2000 Service Pack 2 or later.

See Also: *Administrator*

ADMw0rm

A worm developed by the hacker group ADM that exploits a buffer overflow in BIND.

Overview

The ADMw0rm exploits a buffer overflow in how BIND servers running on Linux platforms handle

inverse queries. The worm typically creates a “w0rm” user account with null password, creates the suid root shell /tmp/.w0rm, deletes /etc/hosts.deny, and replaces all index.html pages with the message “The ADM Inet w-rm is here!” The worm has been around since 1998, and the standard countermeasure is to ensure that you have upgraded your Linux name servers to the latest version of BIND.

The source code for ADM is available from *ftp://adm.freelbsd.net/ADM/* and consists of a number of scripts and programs.

See Also: *worm*

Advanced Encryption Standard (AES)

An encryption algorithm that has replaced the earlier Data Encryption Standard (DES) as the official U.S. government encryption standard.

Overview

When a 56-bit DES key was successfully cracked in 1997 using the idle processing time of thousands of ordinary computers connected to the Internet, it became apparent that a replacement was urgently needed for DES to ensure the confidentiality and integrity of electronic transmissions. A process was initiated by the National Institute of Standards and Technology (NIST) to find a suitable replacement for DES, and in 2001 a cryptographic algorithm called Rijndael (named after its Belgian developers Vincent Rijmen and Joan Daemen) was chosen to form the basis of the new Advanced Encryption Standard (AES). The U.S. government officially approved adoption of AES in May 2002 in Federal Information Processing Standard (FIPS) Publication 197.

Implementation

AES supports several different key lengths including 128, 192, and 256 bits, providing approximately 10^{38} , 10^{57} , and 10^{77} possible keys, respectively (DES provides only 10^{16} possible keys). The maximum key length of 256 bits is so secure that a brute-force cracking program capable of cracking DES in 1 second would take 150 trillion years to crack AES.

AES is implemented as a block cipher that encrypts 128, 192, or 256 bits of data at a time, depending on the key length used. The mathematical structure of AES is simple enough to make it feasible to implement AES on small-footprint devices such as cell phones and personal digital assistants (PDAs) that have limited processing power.

Issues

Although AES has become the official U.S. government standard for encryption, it will likely coexist with DES for the next few years because of the cost and effort of changing over to the new system.

Notes

A more secure cousin of DES is 3DES (Triple DES), which has a key length of 168 bits, which makes it much more secure than DES. While AES is a much faster algorithm than 3DES and requires less processing power to implement, the widespread use of 3DES makes it likely that 3DES will remain an approved U.S. government encryption standard (FIPS 46-3) for some time to come.

See Also: 3DES, Data Encryption Standard (DES), encryption, Federal Information Processing Standard (FIPS), National Institute of Standards and Technology (NIST)

Advanced Security Audit Trail Analysis on UNIX (ASAX)

A sequential file analysis tool for UNIX and Linux platforms that simplifies the analysis of audit information.

Overview

Advanced Security Audit Trail Analysis on UNIX (ASAX) is designed as a universal tool for audit trail analysis and includes a role-based language called Rule-Based Sequence Evaluation Language (RUSSEL) that can be used to create complex queries against audit information. To use RUSSEL, audit logs and other audit trail information must first be translated into a universal format called Normalized Audit Data Format (NADF). RUSSEL can then be used to create explicit rules that allow a normalized audit trail to be processed sequentially in a single pass.

For More Information

Version 1 of ASAX can be downloaded from <ftp://ftp.cerias.purdue.edu/pub/tools/unix/sysutils/asax/>.

See Also: auditing

Advanced Transaction Look-up and Signaling (ATLAS)

A system being developed by Verisign to replace BIND and to bridge between the network infrastructures of telephony and the Internet.

Overview

Advanced Transaction Look-up and Signaling (ATLAS) is designed to support the convergence of Internet and telephony technologies by providing a general platform for any communications system that relies on database lookups. ATLAS works by bridging together popular signaling and name resolution protocols such as the following:

- **Domain Name System (DNS):** The naming system used by the Internet
- **Session Initiation Protocol (SIP):** A signaling protocol developed by the Internet Engineering Task Force (IETF) for Internet conferencing, telephony, events notification, and instant messaging
- **Signaling System Seven (SS7):** The signaling protocol used to initiate calls in the Public Switched Telephone Network (PSTN)

The goal of ATLAS is to provide local number portability that enables users to communicate seamlessly between the Internet and telephone system using the full spectrum of communications devices currently available. Verisign plans to replace BIND with ATLAS on the 13 root name servers it manages, which together support the entire DNS naming scheme that makes the Internet work. Industry analysts expect that this move will help secure DNS and make it less prone to the type of denial-of-service (DoS) attacks that can slow or bring down portions of the Internet. Other advantages of ATLAS include faster propagation of changes to the DNS database (on the order of seconds instead of hours

today), greater scalability (up to 100 billion queries per day), and better support for Internet telephony.

For More Information

Find out more about ATLAS at Verisign's Web site, www.verisign.com.

See Also: denial of service (DoS)

advisory

A public warning concerning a security vulnerability in a software product.

Overview

Advisories (or security advisories) are issued to warn users about vulnerabilities that have been discovered in operating systems and applications. Advisories may be issued by different sources, including governmental agencies, public or private security watchdog organizations, or the vendor that produced the software. Advisories are typically posted to Web sites and mailing lists to ensure the widest possible distribution, and responsible administrators will subscribe to such lists or periodically visit such sites to ensure their networks and systems are secure and hardened against possible attack.

For More Information

RedHat maintains a Security Alerts and Advisories page on its site at www.redhat.com/solutions/security/news; this page includes a security mailing list you can subscribe to for the latest updates.

Cisco security advisories are published by its Product Security Incident Response Team (PSIRT) and are available at www.cisco.com/warp/public/707/advisory.html.

Microsoft maintains a list of security bulletins for all of its Windows platforms and products at www.microsoft.com/technet/security/current.asp and has a subscription-based Microsoft Security Notification Service at www.microsoft.com/technet/security/bulletin/notify.asp.

There are also numerous vendor-neutral organizations that maintain lists of current security advisories and fixes for various platforms. One of the more popular ones is the CERT Coordination Center (CERT/CC)

located at the Software Engineering Institute of Carnegie Mellon University (www.cert.org).

See Also: CERT Coordination Center (CERT/CC)

adware

Any software that installs itself on your system without your knowledge and displays advertisements when the user browses the Internet.

Overview

Adware is a type of "stealth" software and is usually installed on your system when you download and install shareware or free software from the Internet. There are dozens of different adware programs around, and some of them monitor your Web browsing habits and send this information to marketing companies so they can target advertising to you. Some commercial software applications also include adware components that may or may not be mentioned in the End-User License Agreement (EULA) for the product.

Most antivirus programs are not designed to detect adware since the intention of adware is not to harm a system maliciously but to support the cost of developing free software through targeted advertising. A few adware programs such as VX2 and WNAD.EXE, however, have been classified by some antivirus vendors as Trojans and are detected and removed by their products.

If you are concerned about your privacy and the possible presence of adware on your system, third-party utilities exist such as Ad-Aware from Lavasoft that can be used to detect and remove adware and other "spyware." There are also Web sites such as SpyChecker.com and Tom-Cat.com that maintain searchable lists of software that might contain adware or other forms of spyware.

See Also: malware, spyware, Trojan, virus protection software

AES

Stands for Advanced Encryption Standard, an encryption algorithm that has replaced the earlier Data Encryption Standard (DES) as the official U.S. government encryption standard.

See: Advanced Encryption Standard (AES)

AH

Stands for Authentication Header, a security protocol that provides authentication services for Internet Protocol Security (IPSec).

See: Authentication Header (AH)

AKE

Stands for Augmented Key Exchange, a key exchange protocol for public key cryptography systems.

See: Augmented Key Exchange (AKE)

alert

A message sent or event triggered in response to an intrusion, hardware failure, software problem, or some other condition.

Overview

Alerts are a way of quickly informing administrators that firewalls have been breached, networks are under attack, disk drives are full, and all sorts of other problems. Alerts can take different forms with different platforms and products including the following:

- Pop-up windows on Administrator console screens
- E-mail messages sent to Administrator console mailboxes
- Pager alerts or recorded cell phone messages
- Audible alarms, flashing lights, or other methods for gaining one's attention

alert flooding

An attack that tries to overwhelm an intrusion detection system (IDS) by deliberately causing it to generate too many alerts.

Overview

When an IDS detects a possible attack on your network, it typically generates an alert to notify administrators of the situation. This allows them to investigate the problem, determine whether a real attack is under way or whether a false positive has been generated, and take

corrective steps to block attacking systems or ignore similar alerts in the future.

One way in which an intruder may attempt to render an IDS ineffective is to send large numbers of packets that are deliberately designed to cause the IDS to generate alerts. The resulting flood of alerts can overwhelm busy administrators and hide less obvious attempts to probe and intrude upon the network. If too many alerts are generated, the attack can mimic the effects of a denial-of-service (DoS) attack and paralyze the defender's demilitarized zone (DMZ).

The simplest way to generate an alert flood is to package a large portion of the signature database of the IDS into packets and send them to the IDS. If the IDS is configured to generate alerts based on the first match in the database, the attacker usually tries to trigger matches for relatively benign signatures only, thus hiding attempts to breach the IDS using more serious attacks. This kind of attack is most effective against a known signature-based IDS and is less effective against anomaly-based IDSs.

See Also: attack, denial of service (DoS), intrusion detection system (IDS)

Amap

A network-scanning tool for identifying services and applications running on a network.

Overview

Amap is a tool developed by The Hacker's Choice (THC), a hacking community based in Germany. Amap works by sending handshake information for standard application-layer protocols to all TCP ports and is thus able to locate services running on nonstandard ports. For example, Lightweight Directory Access Protocol (LDAP) normally uses port 389 (or 639 if Secure Sockets Layer [SSL] is used), but some administrators might try changing this to a different port above 1023 to hide their network's LDAP services from intruders. However, an intruder using amap can simply scan all 65,535 possible port numbers, sending LDAP handshaking information and looking for the response that indicates which port number is assigned to LDAP.

For More Information

You can find THC online at www.thehackerschoice.com.

See Also: scanner

amplification attack

Any type of attack that magnifies the effect of a single attacking host.

Overview

Amplification attacks work by having one packet generate multiple responses. The resulting effect is that a single attacking host appears as multiple hosts, with the goal of intensifying the effect of the attack to bring down entire networks. Distributed denial-of-service (DDoS) attacks are classic examples of amplification attacks in which intermediary compromised hosts are used to multiply the malicious intent of a single intruder. The Smurf attack is another type of amplification attack and relies on the fact that a single spoofed Internet Control Message Protocol (ICMP) echo request will cause multiple hosts on a network to generate ICMP echo replies, the amplification factor here being the number of accessible hosts on the compromised network.

See Also: distributed denial of service (DDoS), Smurf attack

Annual Computer Security Applications Conference (ACSAC)

An annual conference on computer security organized and sponsored by Applied Computer Security Associates (ACSA).

Overview

Since 1985, the Annual Computer Security Applications Conference (ACSAC) has helped advance the principles and practices of computer security. Conference attendees work primarily in technical fields and include engineers, researchers, and practitioners in the field of computer security. Attendance at ACSAC averages around 250 people and is heavily weighted toward industry and government.

For More Information

For information on upcoming conference schedules and registration, see www.acsac.org.

See Also: Applied Computer Security Associates (ACSA)

anomaly-based IDS

An intrusion detection system (IDS) that uses a baseline instead of signatures to detect intrusions.

Overview

While signature-based (or rule-based) IDSs are more common, they are limited to recognizing known attacks and require their signature database to be updated regularly. An anomaly-based IDS takes a different approach and begins by capturing network traffic to form a profile or baseline of acceptable network events. Once this database has been created, an anomaly-based IDS then compares current traffic to baseline traffic and uses pattern-recognition algorithms to identify possible intrusion events by detecting traffic anomalies. To make the process more efficient, anomaly-based IDSs usually begin by filtering out known “safe” traffic such as Simple Mail Transfer Protocol (SMTP) mail or Domain Name System (DNS) lookups to reduce the amount of data they need to inspect.

Anomaly-based IDSs tend to be good at detecting the initial stage of an attack when an intruder is probing the network using port scans and sweeps. They can also detect when a new network service appears on any host on the network, indicating a possible breach of that host’s security.

The downside of anomaly-based IDSs is that they tend to be more difficult to configure than signature-based IDSs, because it is sometimes difficult to distinguish what constitutes “normal” traffic from “abnormal” and, as a result, they tend to generate more false alerts than signature-based ones. As a result, anomaly-based IDSs usually require a larger degree of human intervention in order to determine the status of “questionable” traffic and reconfigure the IDS to accept or reject such traffic in the future. Finally, anomaly-based IDSs usually need to be deployed in a distributed fashion across a network, close to the servers they are protecting, in order to

reduce the amount of noise they need to filter out. Signature-based IDSs, on the other hand, can often be deployed at network choke points such as firewalls and demilitarized zones (DMZs) provided they are powerful enough to process traffic at wire speed.

Marketplace

Some examples of vendors offering anomaly-based IDS software include Cisco Systems, Enterasys Networks, Intrusion.com, IntruVert, ISS, Lancope, NFR, OneSecure, Recourse Technologies, and Vsecure.

See Also: intrusion detection system (IDS), signature-based IDS

anonymous access

A form of authentication on Internet Information Services (IIS) that allows anonymous users access to a Web or FTP site.

Overview

Anonymous access is designed for public Web sites running on IIS machines connected to the Internet and is configured by default for newly created Web and File Transfer Protocol (FTP) sites. Anonymous access makes use of a special user account `IUSR_servername` for its credentials, where *servername* is the name of the IIS machine. Anonymous access works through firewalls and is compatible with other browsers besides Microsoft Internet Explorer.

See Also: authentication, Basic authentication, Digest authentication

anonymous proxy

A Web site that can be used for anonymous Web browsing.

anonymous Web browsing

Any method for browsing the World Wide Web anonymously.

Overview

Ordinary Web browsing is not an anonymous activity since Hypertext Transfer Protocol (HTTP) requires that your Internet Protocol (IP) address be known by the Web server so that it can return a response to your

request. Once the server has obtained your address, it can then track your browsing patterns and online transactions (cookies are also often used for this purpose). If you are concerned about your privacy and desire to protect your online identity when you browse the Web, there are several approaches you can take:

- Browse the Web from a public Internet terminal, for example, at an Internet cafe or public library. Be aware, however, that public computers may contain Trojans that can capture any credit card numbers or other sensitive information you submit to a Web site.
- Use a proxy server, either one on the perimeter of your company network or one residing at your Internet service provider (ISP) if it provides this service. Be sure you know your company's (or the ISP's) privacy policy before you try this, however, because your browsing history may be recorded in the server's log files and these may be open to inspection by government agencies on demand.
- Use an anonymous proxy service such as Anonymizer.com that performs Web caching and proxying but does not maintain log files. Companies that offer anonymous Web browsing typically offer a free version that is ad supported and a paid or subscription version free of ads. Some of these companies also provide other services such as anonymous e-mail messaging, pop-up ad blocking, and even support for anonymous Secure Sockets Layer (SSL) transactions. Some may also require that you download special client software to your machines to make the anonymous browsing experience transparent to users.

While the positive side of anonymous proxies is that you can use them to protect your privacy, there is a negative side: malicious hackers can sometimes use them to protect their identity when they launch attacks on networks (for example, using malformed URLs, or Uniform Resource Locators) since the proxy hides the true source IP address of the user performing the attack. The use of anonymous proxies can also make it difficult for companies to determine when employees violate acceptable use policies for Web browsing, and companies often block such sites to prevent employees from

surreptitiously downloading pornography and mp3 music files from the Internet.

See Also: *malformed URL attack, privacy*

Anticybersquatting Consumer Protection Act (ACPA)

A U.S. federal law that gives trademark owners legal remedies against domain name cybersquatters.

Overview

The Anticybersquatting Consumer Protection Act (ACPA) is a federal law that became effective in 1999 and was intended to help deal with the problem of domain name cybersquatting, which occurs when a company obtains a domain name “in bad faith,” that is, confusingly similar to a registered trademark for some other company. The company that feels its trademark is infringed or diluted has legal remedy to sue under this act to force the squatter to forfeit or transfer ownership of the contested domain name to the plaintiff and to claim statutory damages up to \$100,000 per domain name. A notable case under this act was **Electronics Boutique Holdings Corp. v. Zuccarini**.

An alternative route trademark owners can pursue to remedy such situations that does not require the expense and time of initiating a lawsuit is to file a complaint with the Internet Corporation for Assigned Names and Numbers (ICANN) under its Uniform Dispute Resolution Policy (UDRP).

AntiSniff

A tool developed by L0pht Heavy Industries (now @Stake) that can detect the presence of packet sniffers on a network.

Overview

Packet sniffers are usually used for troubleshooting network problems but can also be used maliciously to capture network traffic in order to obtain passwords and other sensitive information. AntiSniff was developed in 1999 by L0pht, then a group of hackers, as a network security tool that could detect the presence of sniffers

on a network. To do this, AntiSniff employs a number of techniques including flooding the network with traffic and looking for latency problems that might indicate a host’s network interface card (NIC) is running in promiscuous mode, which is a good indication that a sniffer could be installed on the host.

Versions of AntiSniff have been developed for Microsoft Windows NT, Solaris, OpenBSD, and Linux. @Stake, the security consulting company that L0pht evolved into, has discontinued sales and support for AntiSniff, but the tool is still widely used in the security community.

See Also: *sniffing*

antivirus software

Another name for **virus protection software**, applications for detecting computer viruses and preventing systems from becoming infected.

See: *virus protection software*

application-level gateway

Also called an **application-level proxy**, a type of firewall that establishes proxied connections for specific types of applications.

Overview

Application-level gateways operate similarly to circuit-level gateways in that they operate at the Open System Interconnection (OSI) session layer to monitor Transmission Control Protocol (TCP) handshaking to decide whether session requests should be allowed or denied. Application-level gateways must be specifically configured to support each application-layer protocol such as Hypertext Transfer Protocol (HTTP), File Transfer Protocol (FTP), or Simple Mail Transfer Protocol (SMTP) and to look deep inside packets to find this information. As a result, application-level gateways tend to have greater processing requirements than other types of firewalls and can become bottlenecks under heavy loading.

Application-level gateways are not transparent from the user perspective, as users’ client machines must

be specifically configured to use them. Because application-level gateways prevent direct connections from being made between internal and external hosts, they are particularly good at stopping certain types of network attacks such as protocol violations and buffer overflows.

See Also: *firewall*

application-level proxy

Another name for an **application-level gateway**, a type of firewall that establishes connections using a proxy.

See: *application-level gateway*

application protection system (APS)

Software that identifies hostile Hypertext Transfer Protocol (HTTP) traffic.

Overview

An application protection system (APS) is designed to complement an intrusion detection system (IDS) by examining HTTP traffic to look for suspicious patterns. It generally differs from an IDS in several ways:

- While an IDS examines traffic at the packet level, an APS looks at streams of traffic as a whole, particularly HTTP request/response sessions.
- While an IDS usually alerts administrators to the presence of suspicious traffic, an APS generally blocks such traffic from reaching the Web servers.
- While an IDS is generally located on the demilitarized zone (DMZ), an APS is generally placed immediately in front of the load balancer for the Web server farm.

APS software is usually managed using policies that define which types of HTTP traffic might be considered malicious or harmful to your Web servers.

Marketplace

A number of vendors offer APS software, including Kavado, Protegrity, Sanctum, and Stratum8. The advantage of deploying such systems to protect your Web server farms is that an APS can often detect new types

of attacks and exploits even before they are recognized by security watch organizations and patches are developed.

See Also: *demilitarized zone (DMZ), intrusion detection system (IDS)*

Application Security Tool (AppSec)

A *Microsoft Windows 2000 Server Resource Kit* utility that can be used to limit which applications a user can run.

Overview

Application Security Tool (AppSec) is a graphical user interface (GUI)-based utility that has two security levels:

- **Admin:** Can run any executable file on the machine
- **Non-Admin:** Can run only executables from the approved list

AppSec should generally be used in conjunction with Group Policy restrictions, which can restrict users from accessing such objects as the Start menu and desktop icons. AppSec takes application restriction a step further than Group Policy, however, because it restricts users from running applications even from the command line. This is particularly useful in a Terminal Services environment when you want to limit which applications users can run.

To use AppSec you simply specify the absolute path to those executables (*.exe files) that you want to allow logged-on users to run. The main limitations of AppSec are these:

- It can only be applied to computers, not to users.
- It works only with Win32 applications and not with older Win16 or MS-DOS applications.

For More Information

You can obtain the *Microsoft Windows 2000 Server Resource Kit* from Microsoft Press.

Applications as Services (Srvany)

A *Microsoft Windows 2000 Server Resource Kit* utility that can be used to enable applications to run as services.

Overview

The Applications as Services (Srvany) utility allows applications to be configured to run as background services on the machine. This has an important advantage as far as security is concerned, namely, that you can run applications within the context of a specified user account instead of the credentials of the logged-on user, which gives administrators greater control over the security context in which applications are run. This is particularly useful in a Terminal Services environment to add greater security to the terminal server. Other advantages of running applications as services include the following:

- The ability to run applications while no users are logged on to the system (works only with Win32 applications, not Win16 or MS-DOS ones)
- The ability of an application to continue running after the user logs off (works only with Win32 applications, not Win16 or MS-DOS ones)
- Elimination of the necessity for restarting applications manually after the machine reboots

Notes

Using Srvany requires that you edit the registry to specify the applications that you want to run as services.

Warning: Using Registry Editor incorrectly can cause serious problems that may require you to reinstall your operating system. Microsoft cannot guarantee that problems resulting from the incorrect use of Registry Editor can be solved. Use Registry Editor at your own risk.

For More Information

You can obtain the *Microsoft Windows 2000 Server Resource Kit* from Microsoft Press.

Applied Computer Security Associates (ACSA)

A nonprofit association of computer security professionals whose goal is improving the understanding, theory, and practice of computer security.

Overview

Applied Computer Security Associates (ACSA) was founded in 1985 as Aerospace Computer Security Associates and was renamed in 1996. The initial reason for creating ACSA was to provide ongoing support and funding for its Annual Computer Security Applications Conference (ACSAC), which was at first called the Aerospace Computer Security Conference.

ACSA also supports a number of activities and initiatives whose objectives are to advance the field of computer security. These include a virtual library of security resources, visiting lecture programs for universities, and various committees and workshops on security issues.

For More Information

Find out more about ACSA at www.acsac.org/acsa.

See Also: *Annual Computer Security Applications Conference (ACSAC)*

AppSec

Stands for Application Security Tool, a *Microsoft Windows 2000 Server Resource Kit* utility that can be used to limit which applications a user can run.

See: *Application Security Tool (AppSec)*

APS

Stands for application protection system, software that identifies hostile Hypertext Transfer Protocol (HTTP) traffic.

See: *application protection system (APS)*

Apsend

A free Linux utility for testing firewalls.

Overview

Apsend is a packet sender that lets you test firewalls and other network defenses by simulating SYN floods, User Datagram Protocol (UDP) floods, ping floods, and other forms of denial of service (DoS) attacks. Apsend is a powerful utility that can be used to build Ethernet frames for any type of protocol and that is configured

by default to support Internet Protocol (IP), Transmission Control Protocol (TCP), UDP, and Internet Control Message Protocol (ICMP).

Apsend is open source software written in Perl for the Linux platform and is available under the General Public License (GPL) from Tucows (www.tucows.com) and other shareware sites.

See Also: *denial of service (DoS), firewall, SYN flooding*

APSR

A network-testing tool that can send and receive arbitrary packets.

Overview

APSR is a project developed by the authors of *apsend*, an open source utility for testing firewalls by sending arbitrary Transmission Control Protocol/Internet Protocol (TCP/IP) packets. APSR is essentially an enhanced rewrite in C code of the original *apsend* Perl utility. APSR is currently under development and is commercially available for testing from www.aa-security.de. The authors intend to release a separate free General Public License (GPL) variant once development reaches version 1.

See Also: *Apsend*

arbitrary code execution attack

Any type of attack that enables an intruder to run arbitrary code on the target machine.

Overview

Arbitrary code execution attacks usually exploit application vulnerabilities such as buffer overflows or unchecked variables. Such vulnerabilities arise due to poor coding practices during application development, and writing secure code is essential in order to prevent such attacks from succeeding. Once an intruder has found a way to execute arbitrary code on a target machine, the machine is compromised and may need to be restored from backup because the footprints of the intrusion that follows may be difficult or impossible to follow. Necessary to the success of this attack, however,

is the fact that the intruder must also find a way to generate the code to execute on the target machine, either by copying files to the machine or by gaining control of the machine's file system and creating scripts using a text editor. Having a properly configured firewall and properly securing user accounts on your network can help prevent intruders from inserting such scripts on your machines.

See Also: *attack, buffer overflow, firewall*

Argus

An open source tool for monitoring network activity.

Overview

Argus is an Internet Protocol (IP) network scanner and auditing tool that monitors and records traffic pattern information and stores it in audit logs that can later be analyzed to troubleshoot network problems, verify whether network security policies are working, and be put to many other uses.

Argus 1.x was developed by the Software Engineering Institute (SEI) of Carnegie Mellon University and was first released into the public domain in 1996. Version 2 of Argus is owned by QoSient and is available under an open source licensing agreement.

Argus is named after a being from Greek mythology that had hundreds of eyes. Argus typically runs in the background as a daemon or service and is available for various UNIX platforms including Solaris, IRIX, FreeBSD, OpenBSD, NetBSD, and Linux.

For More Information

You can obtain the latest release of Argus from www.qosient.com.

See Also: *scanner*

ARP cache poisoning

Another name for ARP spoofing, the process of falsifying the source Media Access Control (MAC) addresses of packets being sent on an Ethernet network.

See: *ARP spoofing*

ARP redirection

Another name for ARP spoofing, the process of falsifying the source Media Access Control (MAC) addresses of packets being sent on an Ethernet network.

See: ARP spoofing

ARP spoofing

The process of falsifying the source Media Access Control (MAC) addresses of packets being sent on an Ethernet network.

Overview

Address Resolution Protocol (ARP) spoofing involves modifying the MAC address of packets to fool ARP into thinking they come from a different host than they actually do. ARP spoofing involves sending forged ARP replies to redirect network traffic to the attacking host. If the attacking host is only listening to traffic and not participating in it, legitimate hosts are usually unaware that the packets they are transmitting are being redirected to an attacker and are not reaching their intended destinations. ARP spoofing may be used either for initiating a man-in-the-middle type of attack or for denial of service (DoS) attacks on Ethernet networks.

There are several ways to combat ARP spoofing:

- Add all necessary persistent static entries to the ARP cache on each machine of your network. This is the only sure method of combating ARP spoofing, but it is really only manageable for relatively small networks.
- Use a sniffer to capture network traffic and examine the MAC addresses of Ethernet frames in detail. This is usually too much work to be a practical solution.
- Use a specialized tool such as Arpwatch, which monitors ARP traffic and maintains a global MAC-to-IP address table for all hosts on the segment.
- Use network-layer encryption such as IPSec or VPN over IP to secure all network transmissions.

- Avoid using any network tools that use unencrypted transmission of user credentials (for example, use Secure Shell [SSH] instead of telnet or File Transfer Protocol [FTP]).
- Configure port security on switches running Cisco IOS so that only one MAC address is allowed per port.

Notes

ARP spoofing is also called MAC address spoofing.

For More Information

For more information about ARP, see the *Microsoft Encyclopedia of Networking, Second Edition*, available from Microsoft Press.

See Also: Arpwatch, spoofing

Arpwatch

A command-line utility for UNIX/Linux platforms that monitors Address Resolution Protocol (ARP) tables for changes.

Overview

Arpwatch monitors your Ethernet network and maintains a database of MAC-to-IP address mappings for hosts on the network. Such changes that occur in this database may indicate several possibilities including the following:

- New hosts have been added to the network or existing ones removed.
- The host has obtained a new IP address using Dynamic Host Configuration Protocol (DHCP).
- The Media Access Control (MAC) address has been changed on the host using a vendor's network interface card (NIC) configuration utility.

When a change to the Arpwatch database occurs, an e-mail message is automatically sent to the local root user to notify concerning the change.

Arpwatch is freely available from numerous places on the Internet and is a useful tool to guard against ARP spoofing attacks.

For More Information

For more information about ARP, see the *Microsoft Encyclopedia of Networking, Second Edition*, available from Microsoft Press.

See Also: ARP spoofing

AS

Stands for authentication server, one of two types of servers in a Kerberos key distribution center (KDC).

See: authentication server (AS)

ASA

Stands for Adaptive Security Algorithm, a Cisco algorithm for managing stateful connections for PIX Firewalls.

See: Adaptive Security Algorithm (ASA)

ASAX

Stands for Advanced Security Audit Trail Analysis on UNIX, a sequential file analysis tool for UNIX and Linux platforms that simplifies the analysis of audit information.

See: Advanced Security Audit Trail Analysis on UNIX (ASAX)

ASP.NET Forms authentication

A secure forms-based Web site authentication method provided by ASP.NET on the Microsoft Windows Server 2003 platform.

Overview

ASP.NET Forms authentication uses Web pages with forms as the front end for authenticating users to Web sites. In a typical implementation, cookies would be used for storing authentication ticks and maintaining state management information, while a back-end SQL database would be used for storing user account information. ASP.NET supports the enabling of protection levels to ensure that sensitive information stored in cookies is encrypted or validated, making ASP.NET

Forms authentication more secure than traditional Active Server Pages (ASP) forms authentication using ActiveX Data Objects (ADO).

See Also: authentication

assets

What your company's network security plan is designed to protect.

Overview

A company's assets are its lifeblood, and the job of network security is to protect those assets. Examples of such assets include business plans, equipment, source code for commercial applications developed in-house, private cryptographic keys, credentials of employees, financial account information, and so on. The goal of a security plan is to employ a combination of hardware, software, policies, procedures, and personnel to ensure these assets are protected from intruders, competitors, and even disgruntled employees and contractors.

See Also: security policy

asymmetric key algorithm

A form of encryption in which two mathematically related keys are used.

Overview

Asymmetric key algorithms form the basis of public key cryptography and use two keys: a private key known only to the user and a public key available to everyone. The two most popular asymmetric algorithms are the following:

- **Diffie-Hellman:** Developed by Whitfield Diffie and Martin Hellman in 1976, this algorithm was the first published example of how public key cryptography could be performed. The Diffie-Hellman algorithm is relatively slow, however, and was intended not for encrypting data for transmission but mainly for securely transmitting a Data Encryption Standard (DES) session key to allow private key encryption to be used during the session. This approach is commonly known as a Public Key Distribution System (PKDS).

- **RSA:** Developed by Ron Rivest, Adi Shamir, and Leonard Adleman in 1977, RSA has similarities with Diffie-Hellman but is significantly faster and can be used to encrypt individual messages for secure transmission, an approach known as public key encryption (PKE). The RSA algorithm was originally proprietary but is now in the public domain.

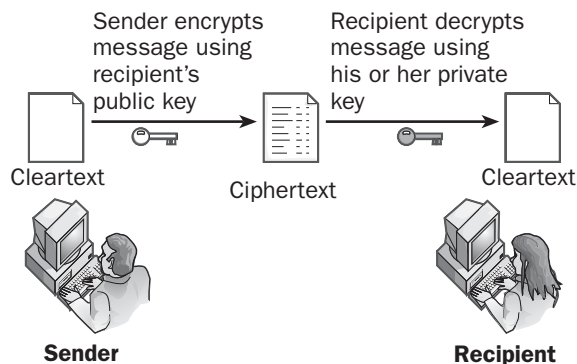
Besides the popular Diffie-Hellman and RSA, some other asymmetric algorithms in use today include these:

- **Elliptic Curve Cryptosystems (ECC) algorithms:** A family of algorithms based on elliptic curve theory that can provide a high degree of security even with a relatively small key size
- **El Gamal:** An algorithm based on calculating discrete logarithms

Key length in asymmetric algorithms is typically much larger than that used in symmetric algorithms such as DES. Common sizes for RSA keys include 1024 and 2048 bits in comparison with 56 bits for Data Encryption Standard (DES) and 128, 192, or 256 bits for Advanced Encryption Standard (AES).

Implementation

To send an encrypted message using asymmetric encryption, the sender first uses the recipient's public key to encrypt the message, transforming it from clear-text into ciphertext. The message is then sent to the recipient, who uses his or her own private key to reverse the process and decrypt the message.



Asymmetric key algorithm. How asymmetric key cryptography works.

While this process ensures confidentiality for the message, it does not guarantee the identity of the sender. For this purpose, the sender can attach a digital signature to the message, which verifies the identity of the sender to the recipient and ensures the integrity of the message.

Notes

Asymmetric key algorithms are also commonly referred to as **public key algorithms**.

See Also: *Data Encryption Standard (DES), Diffie-Hellman (DH), digital signature, RSA algorithm, symmetric key algorithm*

ATLAS

Stands for Advanced Transaction Look-up and Signaling, a system being developed by Verisign to replace BIND and to bridge between the network infrastructures of telephony and the Internet.

See: *Advanced Transaction Look-up and Signaling (ATLAS)*

ATR string

A string of bytes returned by a smart card when it is inserted into a smart card reader.

Overview

When a smart card is inserted into a reader, the reader generates a reset signal and the card responds by returning an Answer to Reset (ATR) string. This ATR string is used to identify the type of smart card, the status of the card, and information that optimizes the serial connection between the reader and the card. The format of ATR is described in the International Organization for Standardization (ISO) 7816-3 standard, which defines a maximum length for the string of 33 bytes.

In a typical implementation, before a smart card can be used, the setup utility must be run to initialize the card and assign it a friendly name, ATR string, and optional mask. Then when an application requests smart card authentication, it can connect to a card on a given reader, obtain the ATR string, and compare it to the ATR string of the requested card.

See Also: *smart card*

attack

Any method used to try to breach the security of a network or system.

Overview

Threats to a network's security can originate from a variety of sources including the following:

- External, structured threats from malicious individuals or organizations
- External, unstructured threats from inexperienced attackers such as script kiddies
- Internal threats from disgruntled employees or contractors

The overall approaches used by malicious individuals or organizations vary considerably, but can be broken down into several broad categories:

- **Access attacks:** The intruder tries to gain access to resources on your network by exploiting flaws in software such as buffer overflows and information leakage and by elevating the intruder's privileges to execute arbitrary code.
- **Denial of service (DoS) attacks:** The intruder tries to deny legitimate users access to resources on your network.
- **Reconnaissance attacks:** The intruder tries to map your network services in order to exploit vulnerabilities detected.

Another way of classifying attacks is according to their impact on the systems being attacked:

- **Active attacks:** These involve trying to modify data either during transmission or while stored on the target system. Examples include inserting backdoors and Trojans, deleting or modifying log files, disrupting services or communication, and so on.
- **Passive attacks:** The goal here is not to modify the target system but rather to capture data being transmitted by eavesdropping or by using a packet sniffer in order to obtain sensitive or confidential information such as passwords or credit card numbers. Passive attacks are also used for capturing

information that can help the attacker create a map of the target network's hosts and services, which usually forms the preamble of an active attack.

Some of the specific methods used by intruders for attacking networks include data modification, eavesdropping, impersonation, and packet replay attacks. Other common methods include exploiting coding vulnerabilities using buffer overflows, malformed Uniform Resource Locators (URLs), and other methods. Social engineering and Dumpster diving are different approaches that sometimes lead to immediate success in penetrating a network's defenses. Finally, phishing is a form of automated social engineering that sometimes bears fruit for the attacker.

See Also: *attack map, denial of service (DoS), eavesdropping, impersonation, packet replay, phishing, sniffing, social engineering, vulnerability*

attack map

A map of a network that an intruder plans to attack.

Overview

When an intruder wants to break into a company's network, the first stage of the attack is usually the reconnaissance stage, in which the intruder tries to create a map of the hosts and network services running on the network. The intruder might begin mapping a company's network by using a search engine to find the company's Web site and then use nslookup to find the Internet Protocol (IP) address of the server. Once the server's IP address is known, a whois query of the ARIN database can determine the range of IP addresses for the network and administrative contact information, which can indicate whether the server is being hosted on the company's demilitarized zone (DMZ) or at an Internet service provider (ISP) or hosting provider.

If the server is hosted by a service provider, the attacker could be out of luck. But if the company owns the IP address block to which the server belongs, then one host on the company's network has been identified. At this point, the intruder might scan the IP address block using a freely available tool such as Nmap to see whether any other hosts in the network are exposed

to the Internet (stealth mode is used for running Nmap to help avoid detection during the scanning process). Once exposed hosts are found and listening ports identified, the intruder has gained knowledge of which network services are running on these hosts and the attack map takes on shape.

The intruder might next try to determine which operating systems are running on the exposed hosts. For Web servers, this can be done by using telnet to send a Hypertext Transfer Protocol (HTTP) GET request to port 80, because the Web server's response to this request contains HTTP headers that typically contain such information. Having identified the operating system and network services running on hosts, the intruder has created a map that then allows it to test for common vulnerabilities that result from administrators failing to patch their systems appropriately.

See Also: *attack, Nmap, Nslookup, vulnerability*

auditing

A security principle that involves reviewing information concerning user and system activity to look for inappropriate actions.

Overview

Auditing is an essential part of any security program for any network, and most operating system platforms support some form of auditing. Auditing can be approached in two general ways:

- **Proactive auditing:** This involves regularly reviewing audit logs to look for signs of intrusion attempts or abnormal system behavior.
- **Reactive auditing:** This is basically a forensic activity that is performed after a system has been compromised.

By enabling auditing on a system, information is collected concerning specified events such as logons, resource access, and so on. This information is then stored in special log files called audit logs, which can then be reviewed to look for suspicious patterns of behavior or monitor resource usage activity for accounting purposes. Many applications such as

firewalls and intrusion detection systems (IDSs) also support various forms of auditing.

Implementation

Auditing is implemented in various ways on different operating system platforms and applications. In general, an auditing system comprises two components:

- A data collector that monitors the system or application and saves audit information in audit logs
- A data analyzer that allows administrators to display, query, and analyze audit logs to search for patterns and events

As far as analysis of auditing information is concerned, this may be performed in either of the following ways:

- Manually, by having administrators periodically examine audit logs using various tools
- Automatically, using statistical methods or rule-based expert systems, an approach generally used for IDSs

In addition, auditing systems can be implemented in either of the following ways:

- **Local auditing:** Each system is responsible for collecting and maintaining its own audit information.
- **Distributed auditing:** Collection of audit information is performed by different systems on a network and either stored centrally for processing or analyzed in a distributed fashion for load balancing of processing.

There is no single standard format for what information should be audited by a system or how it should be stored. Certain standards such as the Security Criteria for Distributed Systems developed by the Institute for Defense Analysis or the Trusted Computer Systems Evaluation Criteria from the National Computer Security Center are helpful in deciding what types of events a computer system should be able to audit, but different vendors usually implement such standards differently. And despite various attempts to standardize audit log formats, particularly on the UNIX platform, there is

currently no universal format that has achieved wide acceptance.

See Also: *audit log, audit policy, Security Criteria for Distributed Systems, Trusted Computer Systems Evaluation Criteria (TCSEC)*

audit log

A log file containing audit information for an application or system.

Overview

Audit logs provide a record of audit information for monitoring the security and accountability of applications and systems. There is no standard or universal format for audit logs, although there have been various attempts at developing such a format especially on UNIX platforms. Examples of such proposed standard audit file formats include

- Bishop's Standard Audit Trail Format
- Normalized Audit Data Format (NADF)

Implementation

On Microsoft Windows platforms, auditing is controlled by audit policies configured using Local Security Policy or Group Policy, depending on whether machines are running in a workgroup or Active Directory service scenario. Microsoft Windows platforms support auditing of file system objects, printers, Active Directory, and security events such as logons and privilege use.

Most UNIX systems maintain various types of auditing information including the following:

- **Logon logs:** These maintain a record of console logons, use of rsh, and sessions for telnet, File Transfer Protocol (FTP), and X. Logon logs are usually located either under the /etc or /var hierarchies.
- **System logs:** These maintain records of various system activities in a set of logs specified by /etc/syslog.conf.

Notes

When inspecting syslogs, look for suspicious events such as these:

- Attempts to access /etc/passwd, which may indicate intruders are trying to obtain copies of password hashes so they can run cracking utilities against them
- Failed attempts to use Su, which may indicate intruders are trying to gain root access to your system
- Missing log files, missing log entries, or unusual amounts of log activity at certain times, which may indicate the system has been compromised and log files modified

On Solaris 8 or greater, auditing can be configured using the auditconfig command, and audit logs are stored in the /var/audit directory.

See Also: *auditing, audit policy, Su*

Auditpol

A utility in the *Microsoft Windows NT 4.0 Server Resource Kit* for remotely managing auditing on servers running Windows NT.

Overview

Auditpol lets administrators enable, disable, or view auditing information on remote servers running Windows NT. Auditpol is sometimes used by intruders to disable auditing on servers running on networks they have penetrated. This is done to hide the intruder's footprints and make it harder to determine how the intrusion was accomplished or which actions were performed. Auditpol requires administrator credentials, however, so an important step in preventing this type of attack is to ensure your administrator accounts are secured and have strong passwords.

See Also: *auditing*

audit policy

A policy that specifies the level and type of auditing to be performed by a system or application.

Overview

Audit policies allow policy-based management of auditing of various system events such as logons, directory service access, privilege use, process tracking, and so on. Audit policies can be implemented on a variety of different platforms including Cisco, Microsoft Windows, and various UNIX platforms.

As an example, Secure IDS, an intrusion detection system (IDS) product from Cisco, lets you create audit policies for auditing network traffic through a router. Here, auditing is performed when IP packets arrive at a router interface and are compared with signatures configured for that interface. The Ip Audit command is used to create both a global audit policy and individual separate policies for inbound and outbound traffic at each interface.

On Microsoft Windows platforms, audit policies let administrators configure how auditing is performed. These policies are configured in a subnode of the Local Security Policy (or using Group Policy in an Active Directory service scenario) called Audit Policy. The types of events controlled by an audit policy include auditing the following:

- Account logon events such as successful or failed logons
- Account management events including creating and deleting accounts, changing or unlocking passwords, and so on
- Access to objects in Active Directory, the Windows file system, printers, or the Registry
- Policy change events such as when a password policy, audit policy, or some other policy is modified
- Privilege use when the user exercises system rights
- Process tracking including launching programs, indirect object access, and so on
- System events such as shutting down or rebooting the system

See Also: *Group Policy*

audit trail

A record of events generated by an application, system, or organization.

Overview

Audit trails are generally created to provide accountability for the actions of applications, systems, or individuals within an organization. They may be created automatically (for example, by enabling auditing on an application or operating system) or manually (for example, by going through records of memos and other paperwork having to do with an issue or individual's behavior).

Audit trails may also have other uses besides providing accountability, including the following:

- Allowing applications, systems, and business processes to be monitored to detect potential or impending problems, misuse of resources, or other purposes
- Allowing events to be forensically reconstructed after an intrusion, theft, or other criminal event to determine who was involved and what was done

For computer systems and applications, audit trails are generally automatically created when auditing is enabled, and the audit trail is stored in a format called an audit log.

See Also: *auditing, audit log*

Augmented Key Exchange (AKE)

A key exchange protocol for public key cryptography systems.

Overview

Augmented Key Exchange (AKE) is designed to provide mutual authentication and key agreement between users in a public key system. AKE was developed by Bellare and Merritt to address security shortcomings in their earlier Encrypted Key Exchange (EKE) protocol. AKE does this by requiring that verification servers must not store their passwords in plaintext. Otherwise, AKE works similarly to EKE; to understand such key exchange protocols, please refer to the article

Encrypted Key Exchange (EKE) elsewhere in this book.

See Also: *Encrypted Key Exchange (EKE), key exchange*

AUP

Stands for acceptable use policy, a policy that defines appropriate use of computing resources for a company or organization.

See: *acceptable use policy (AUP)*

AusCERT

Stands for Australian Computer Emergency Response Team, an independent nonprofit organization that monitors and evaluates global computer network threats and vulnerabilities.

See: *Australian Computer Emergency Response Team (AusCERT)*

Australian Computer Emergency Response Team (AusCERT)

An independent nonprofit organization that monitors and evaluates global computer network threats and vulnerabilities.

Overview

The Australian Computer Emergency Response Team (AusCERT) publishes security bulletins from various sources together with recommendations for prevention and mitigation of effects. AusCERT also provides training and consulting services, has an emergency response service for members, and hosts yearly Asia Pacific Information Technology Security Conferences. Some AusCERT services are available free to the general public, while others are offered to paid subscribers and the money is used to cover operating costs for the organization. AusCERT is also represented on the steering committee of the Forum for Incident Response and Security Teams (FIRST).

For more Information

Visit AusCERT online at www.auscert.org.au.

See Also: *advisory, threat, vulnerability*

authentication

The process of determining the identity of a user or other entity.

Overview

Authentication is a process that verifies that entities are in fact who they claim to be. Entities that may require authentication by computer systems include users, computers, and processes. On a typical computer network, user authentication is performed during the logon process when a user submits credentials usually consisting of a username and password. On Microsoft Windows-based networks that use Active Directory service, users may also be required to include their domain as part of their credentials.

Authentication is also employed in electronic messaging to determine the identity of the entity that signed a message (entity authentication) and to verify that the message has not been altered in transit (data authentication).

Implementation

Authentication can be implemented in many ways and forms including the following:

- **Address-based authentication:** A method that uses a host's network address as its identity for authentication purposes
- **Anonymous access:** A method used by Internet Information Services (IIS) to allow anonymous users access to public Web sites
- **ASP.NET Forms authentication:** A method for securely authenticating users to Web sites supported by Windows Server 2003
- **Basic authentication:** An RFC-compliant method for logging on to Web and FTP (File Transfer Protocol) sites
- **Biometric authentication:** Authentication that verifies identity using physical characteristics such as fingerprints or retinal scans

- **Certificate-based authentication:** A method that employs digital certificates and a Public Key Infrastructure (PKI) to authenticate users
- **Digest authentication:** A variant of Basic authentication that hashes passwords before transmitting them
- **Kerberos:** A secure authentication method defined in RFC 1510 and used by Microsoft Windows 2000 and Windows Server 2003
- **Smart cards:** An authentication method that employs cards with embedded chips containing encrypted information about the user
- **Windows NT Challenge/Response:** Also called NTLM (for NT LAN Manager); a secure authentication method used in Microsoft Windows NT and supported by later versions of the Windows operating system

These and several other authentication methods are discussed in more detail in separate articles in this book.

See Also: *address-based authentication, anonymous access, Basic authentication, biometric identification, certificate-based authentication, Digest authentication, Kerberos, smart card, Windows NT Challenge/Response*

Authentication, Authorization, and Accounting (AAA)

A security framework for controlling access to network resources.

Overview

Authentication, Authorization, and Accounting (AAA), or Triple-A, is a security framework that performs three functions:

- **Authentication:** Defining who can access a network
- **Authorization:** Determining what a user can access once authenticated
- **Accounting:** Keeping a record of what the user does once authenticated and authorized

AAA is currently not an Internet standard, but instead is classified by the Internet Engineering Task Force (IETF) as Experimental and is defined in RFC 2903, “Generic AAA Architecture,” and a series of Information RFCs including 2904 through 2906 and others.

Implementation

Numerous vendors have implemented AAA schemes, including Microsoft, Cisco, Hewlett-Packard (HP), and others. Cisco’s PIX Firewall product can forward authentication requests to an AAA server running Cisco Secure Access Control Server (CSACS) software, which then authenticates the user’s credentials, authorizes the user to access network resources, and tracks the user’s access to these resources.

HP’s Mobile AAA Server runs on HP-UX and can provide AAA requirements for mobile IP data services including 3G cellular systems. It includes a Lightweight Directory Access Protocol (LDAP) directory and session management tools.

The Internet Authentication Services (IAS) component of Microsoft Windows operating systems also provides AAA services for virtual private network (VPN) remote access through its implementation of the Remote Authentication Dial-In User Service (RADIUS) protocol.

For More Information

For more information about the RADIUS protocol, see the *Microsoft Encyclopedia of Networking, Second Edition*, available from Microsoft Press.

See Also: *authentication, authorization*

Authentication Header (AH)

A security protocol that provides authentication services for Internet Protocol Security (IPSec).

Overview

Authentication Header (AH) ensures that Internet Protocol (IP) packets have not been tampered with during IPSec sessions. It does this by acting like a digital signature for the packet, thereby ensuring data integrity. AH can be used either by itself or together with the Encapsulating Security Payload (ESP) protocol if data integrity is required. AH can optionally provide replay-detection services but does not provide data

encryption or decryption services. AH is described in RFC 2402.

Implementation

At the packet level, AH is implemented differently depending on how IPSec is configured to be used. Specifically, when IPSec is running in transport mode, the AH header follows the IP header and precedes the Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) header. When tunnel mode is used instead (not common), the AH header is placed between the new and original IP headers.

AH authentication is performed using a keyed message authentication code (MAC) or hash-based message authentication code (HMAC). The authentication algorithms usually used are either HMAC using MD5 or HMAC using SHA-1.

See Also: *Encapsulating Security Payload (ESP)*, *hash-based message authentication code (HMAC)*, *Internet Protocol Security (IPSec)*, *MD5*, *message authentication code (MAC)*, *Secure Hash Algorithm-1 (SHA-1)*

authentication package

Code that encapsulates the logic used for authenticating users.

Overview

In Microsoft Windows operating systems, authentication packages are implemented as dynamic link libraries (DLLs) and are used to implement features of security protocols. When the local security authority (LSA) receives a logon request, it authenticates the user by loading the appropriate authentication package based on information contained in the system Registry. The authentication package then determines whether the user should be allowed to log on to the system or network, establishes a new logon session for the user, and passes information to the LSA that enables it to generate a security token for the user.

The two authentication packages included by default with Microsoft Windows platforms are the following:

- MSV1_0 Authentication Package, included with Microsoft Windows platforms for Windows NT 3.51 and later
- Kerberos SSP/AP, included with Windows 2000 and Windows XP Professional

See Also: *authentication*

authentication server (AS)

One of two types of servers in a Kerberos key distribution center (KDC).

Overview

In a Kerberos implementation, the KDC employs two types of servers: the ticket-granting server (TGS) and the authentication server (AS). The AS performs the initial step of authenticating users to the TGS, which then performs the subsequent step of authenticating users to protected services. This two-stage approach precludes users from the need to reenter their password each time they want to access a service.

See Also: *Kerberos*, *key distribution center (KDC)*, *ticket-granting server (TGS)*

Authenticode

A feature of Microsoft Internet Explorer that enables users to know that software they download can be trusted.

Overview

Authenticode is a mechanism that allows digital certificates to be attached to software downloaded from the Internet, especially ActiveX controls, cabinet files, executable files, dynamic link libraries (DLLs), and catalog files. When a user tries to download a signed ActiveX control, a message appears indicating that the code originates with the developer and has not been altered by any third party. The user then decides whether to accept the message and download and run the control, or reject it.

For More Information

For more information about ActiveX and ActiveX controls, see the *Microsoft Encyclopedia of Networking, Second Edition*, available from Microsoft Press.

See Also: *digital certificate*

authorization

The process of granting rights to entities to allow them to access network resources.

Overview

In general, the process of authorization for accessing resources on a network can be approached two ways:

- **Role-based authorization:** Here, users are partitioned into logical roles in which members of a role share the same privileges. Network resources are then accessed using fixed identities, for example, the process identity for a Web application, and it is the responsibility of the application to correctly authorize users.
- **Resource-based authorization:** Here, resources are secured using access control lists (ACLs) that determine which users may access the resource and which actions they can perform, such as reading or modifying a file. In this scenario, network resources are accessed using impersonation.

Authorization and authentication go hand in hand because meaningful authorization to access network resources first requires that users be authenticated to access the network itself.

See Also: *authentication*

authorization creep

A term describing how users may possess unnecessarily high access privileges within a company or organization.

Overview

When users move from one department to another, they will sometimes maintain access credentials from their earlier position even if these are no longer needed. This scenario is sometimes called authentication creep since the term suggests a slow but invisible increase in the access privileges of authenticated users within a large organization. The solution to this problem is to ensure that when a user changes job roles or positions, the user's account credentials and access rights are reviewed by management and the ID department is immediately notified which former rights to disallow when new rights are granted to the user in the new position.

See Also: *access, authentication*

autologon

A logon method in which the user is automatically logged on to the system or network.

Overview

Autologon lets users log on without the need of specifying credentials each time they want to log on to their computers or networks. While autologon may seem like a convenience, it is generally not recommended except for computers used in kiosks. Servers that are physically secured in back rooms could also use it, but this is generally not recommended since autologon basically bypasses all security measures on a computer and allows anyone who can gain physical access to the machine access to resources commensurate with the user's level of privileges. In this respect, using autologon for Administrator accounts is clearly a bad idea.

On Microsoft Windows platforms, autologon can be enabled or disabled by editing certain registry settings specific to this feature.

See Also: *logon*

autorooter

An automated tool for discovering security vulnerabilities in networks.

Overview

The term **autorooter** is used mainly in the black hat community for any tool or collection of tools that can automatically scan large numbers of systems looking for vulnerabilities to exploit. Most autorooters work by first compiling a list of Internet Protocol (IP) addresses for live systems within a specified address range connected to the Internet (or on a compromised private IP network). Then the tool scans these systems to identify which operating systems are running and to identify any network services or applications running on them. Finally, the tool performs automated exploits against services and applications that have known vulnerabilities.

See Also: *black hat, scanning*

backdoor

Sometimes called back door, any hidden mechanism for accessing an application, system, or network.

Overview

Backdoors were originally mechanisms created by computer programmers to allow them special access to their programs, usually so they could fix the code when a bug caused a crash to occur. A famous example of this is when Ken Thompson admitted to the Association for Computing Machinery (ACM) in his 1983 Turing Award lecture that he had hard-coded a Trojan login program in the C compiler for early versions of UNIX, allowing him backdoor access to any UNIX system running on Bell Labs' internal network. This clever backdoor even protected itself against discovery and removal, for even if a user found the code in the compiler and removed it, the compiler still had to be recompiled (using itself!), and Thompson had inserted additional code that detected when the compiler was being used to recompile itself, and this additional code would then re-create the backdoor in the recompiled version!

Sometimes developers add backdoors to their programs for malicious (or at least suspicious) reasons. For example, a backdoor could be inserted into the code for an online shopping cart to enable the developer to surreptitiously obtain transaction information, including credit card numbers. Ostensibly, the reason for this may be to monitor the cart's operation to detect abuse, but users may rightly feel their privacy is being violated by such an action, especially if no mention of this is made in the privacy policy for the site.

The term **backdoor** was later co-opted by hackers to describe any mechanism by which an attacker could stealthily reaccess a compromised system or network without needing to repeat the exploit that provided the

attacker access in the first place. Typically, once a network has been compromised by exploiting some vulnerability of the application or system, attackers will proceed to cover their tracks by modifying or deleting log files and will then install a backdoor such as special software or a hidden account with administrator privileges. If the owner of the system or network discovers the intrusion and hardens the vulnerability to prevent further access but does not detect the presence of the installed backdoor software, the attacker has a stealthy way of reentering the system to cause further damage. Often the only way to be sure the backdoor has been removed is to wipe the system and reinstall from a backup known to be secure. A popular tool for installing backdoors on penetrated systems is Netcat, which can initiate or receive Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) connections on any port.

Another form of backdoor is key escrow, in which an agency such as the government is provided with keys that can decrypt encrypted messages. Public key encryption normally ensures the privacy of communications by ensuring that users' private keys are owned only by them. Key escrow thus provides a backdoor mechanism for reading users' private communications. The justification for doing this is concern for national security and prevention of criminal actions, especially terrorist attacks, but citizens often fear giving governments such capabilities, which may endanger their personal rights and freedoms.

See Also: *key escrow, Netcat, public key encryption, Trojan*

Back Orifice

A powerful Trojan program for the Microsoft Windows 95 and Windows 98 platforms.

Overview

Back Orifice was developed by a hacker group called Cult of the Dead Cow (cDc) and was released to the public at Defcon 6 in 1998. The program was identified by CERT (Computer Emergency Response Team) as a potentially serious vulnerability because it can give an attacker the level of privileges of the users who inadvertently install it on their systems. The original Back Orifice program targeted machines running Windows 95 and Windows 98, while a later version, Back Orifice 2000, also targeted Microsoft Windows NT systems. The tool's creators positioned it as a legitimate remote administration tool, but its potential for misuse caused most security advisories to classify it as malware.

Implementation

Back Orifice works as a client/server program, with the server portion running on the target computer and the client on the attacker's machine. In order for Back Orifice to work, users must first be tricked into installing the program on their computers, usually by downloading files from Web sites masquerading as friendly sites. Once the server portion is installed on the user machine, the attacker can use the client portion to monitor and control the user machine by logging keystrokes, running applications, viewing and modifying files, and so on. Communication between the client and server takes place using encrypted Transmission Control Protocol/Internet Protocol (TCP/IP) communications over port 31337, but this port is configurable.

An insidious aspect of Back Orifice is its ability to “piggyback” by attaching itself to a legitimate operating system process so that each time the system is booted the program automatically and stealthily starts. The original filename under which the program installed itself was “.exe” (a space followed by .exe), but even this can be configured, with the result that Back Orifice can be difficult to detect on a compromised system.

Notes

A program called BOSniffer appeared in 1998 and was reputed to be able to prevent Back Orifice from installing on systems and to find and remove backdoors installed by existing Back Orifice installs, but this is

actually a variant of Back Orifice itself and should not be installed.

For More Information

See Cult of the Dead Cow at www.cultdeadcow.com for more information.

See Also: Back Orifice 2000 (BO2K), malware, Trojan

Back Orifice 2000 (BO2K)

A well-known Trojan program for the Microsoft Windows NT and Windows 2000 platforms.

Overview

Back Orifice 2000 (BO2K) is a version of Back Orifice developed for the Windows NT and Windows 2000 platforms and can be used either legitimately as a remote administration tool or maliciously as a tool for monitoring and controlling compromised systems. By default, the background process for BO2K appears in Task Manager as UMGR32, but the program can also be configured to run stealthily as an invisible service not displayed in Task Manager.

BO2K was developed by Cult of the Dead Cow (cDc) and was released in 1999. With BO2K installed on a target system running Windows NT, an attacker can perform any action that a locally logged-on user can do, including view and modify files, run programs, perform encrypted file transfers, and so on. The architecture of BO2K is similar to the earlier Back Orifice but includes plug-in capability that extends the functionality of the basic tool.

Notes

There is also a Linux version of BO2K released under the General Public License (GPL).

For More Information

See Cult of the Dead Cow at www.cultdeadcow.com for more information.

See Also: Back Orifice, Trojan

backup authority

A trusted application running on a secure computer that provides secondary storage for session keys of clients.

Overview

Backup authorities are part of the Cryptographic API on Microsoft platforms and are used to store session keys as key binary large objects (binary BLOBs). These BLOBs are then encrypted using the public keys of the backup authority to secure them.

In order for an application to use a backup authority, it first encrypts the file, exports the session key used to encrypt the file into a simple key BLOB using the application's public key, and stores the key BLOB and file together. The session key is then exported a second time using the backup authority's public key to encrypt the key BLOB, and the key and its description are then sent to the backup authority. Then, should the key pair be lost, the keys can be recovered from the backup authority once the identity of the user has been established.

See Also: *key pair; session key*

backup plan

A plan for backing up important business information.

Overview

Backup plans are an essential part of any company's disaster recovery plan and specify when, which, and how data is backed up. Developing a backup plan involves determining answers to the following questions:

- Who is responsible for ensuring backups are done properly?
- What information should be backed up and how often?
- Which backup technologies, tools, media, and methods should be employed to ensure data can be recovered speedily after a disaster?
- Where can media be securely stored to ensure important business data cannot be irretrievably lost after a disaster?

- How can backups be properly tested to ensure the ability to recover from a disaster?

For More Information

For more information about backup technologies, tools, and methods, see the relevant articles in the *Microsoft Encyclopedia of Networking, Second Edition*, available from Microsoft Press.

See Also: *disaster recovery plan (DRP)*

Badtrans.B

A worm that targets Microsoft Windows-based messaging platforms.

Overview

Badtrans.B is an e-mail worm that is a variant of an earlier virus called Badtrans. The worm uses Microsoft's Messaging API (MAPI) in Microsoft Outlook to send copies of itself using different file names to everyone in the address book. It also creates the file Kdll.dll in the \system directory and uses this file to log keystroke activity on the user's machine. Infection usually occurs by opening infected e-mail attachments, and the best way to avoid infection is to block e-mail attachments used to spread viruses, including .exe, .bat, .vbs, .scr, .pif, and similar files.

Badtrans.B targets all 32-bit Windows platforms and, at its high point in 2002, it reached a threat level of Category 4 on the Symantec Security Response site. Protection against the worm involves standard messaging system security practices, including applying vendor patches such as the Microsoft Outlook security patch, blocking attachments with double extensions such as *.doc.exe, and so on.

Notes

The worm is identified as W32.Badtrans.B@mm by Symantec Security Response, where @mm identifies the worm as being of the mass mailer type.

For More Information

Search www.securityresponse.symantec.com for more information on Badtrans.B.

See Also: *worm*

bandwidth consumption attack

A type of denial of service (DoS) attack in which an attacker consumes all available bandwidth on the target network.

Overview

DoS attacks prevent legitimate users from accessing network resources. The most common way this is done is when an attacker attempts to utilize all available bandwidth on your network by “flooding your ports” with spurious Transmission Control Protocol (TCP) packets. For such attacks to be successful, attackers often employ large numbers of machines in the form of a distributed denial of service (DDoS) attack.

Bandwidth consumption attacks are different from SYN floods, which require low attack bandwidth and work by tying up a machine’s TCP connection resources. Instead, bandwidth consumption attacks use a flood of malicious packets to overwhelm the machine’s network connection resources and prevent legitimate packets from being received or transmitted by causing router interfaces to drop and discard them.

See Also: *denial of service (DoS), distributed denial of service (DDoS)*

banner grabbing

An attack designed to deduce the brand and/or version of an operating system or application.

Overview

Banner grabbing is used by attackers to profile target systems to allow them to select platform-specific methods for compromising them. For example, once a target system has been identified as running BSD/OS 4.3, the attacker can then consult a list of known vulnerabilities for this platform and attempt an exploit.

Common ports used by attackers for profiling target systems include FTP (File Transfer Protocol, port 21), SSH (Secure Shell, port 22), telnet (port 23), SMTP (Simple Mail Transfer Protocol, port 25), and HTTP (Hypertext Transfer Protocol, port 80). Fscan, a popular command-line port-scanning tool developed by

Foundstone Labs, is an example of a tool that can be used to perform banner grabbing.

See Also: *port scanning*

base content type

Type of data contained in a Public Key Cryptography Standards (PKCS) #7 message.

Overview

PKCS 7 is a standard that defines a general syntax for authentication and encryption. Base content types contain only data and cannot contain cryptographic enhancements such as hashes or signatures. The only base content type currently defined by PKCS 7 is the data content type, which contains simple strings of byte (octet) characters in unencrypted form.

See Also: *PKCS #7*

Basic authentication

A standard Hypertext Transfer Protocol (HTTP) authentication method.

Overview

Basic authentication is part of the HTTP 1 specification and can be used to authenticate users running Web browsers against Web sites running on Web servers. Basic authentication is supported by most Web browsers and Web servers, including Internet Information Services (IIS) on Microsoft Windows platforms.

Basic authentication passes a user’s credentials over network connections in unencrypted form (actually in Base64 encoding, but this is trivial to decode), making it vulnerable to sniffing attacks. The credentials received from the client are compared against either the local account database on the server or a network security controller (a domain controller in the Windows operating system case) in order to authenticate the user. To make Basic authentication secure, it can be combined with Secure Sockets Layer (SSL) to encrypt the user’s credentials.

See Also: *authentication*

Basic Encoding Rules (BER)

A set of rules used for encoding ASN.1-defined data into a bitstream.

Overview

Basic Encoding Rules (BER) is used to encode information formatted using ASN.1 into zeros and ones so it can be transmitted or stored. BER is thus the “transfer syntax” for ASN.1 and was designed by the Comité Consultatif International Télégraphique et Téléphonique (CCITT), the same group that created the ASN.1 specification. BER is described by the X.209 recommendation of the CCITT (now the International Telecommunications Union or ITU) and is also defined by the ISO 8825 standard. BER is a self-describing encoding scheme and is thus not especially bit-efficient for communications.

In Transmission Control Protocol/Internet Protocol (TCP/IP) networking, BER specifies the transfer syntax for sending Simple Network Management Protocol (SNMP) and Lightweight Directory Access Protocol (LDAP) messages. On Microsoft platforms, BER is used by the CryptoAPI (CAPI) application programming interface.

See Also: *CryptoAPI (CAPI)*

Bastille

A script used to harden the Linux operating system against attack.

Overview

Bastille is designed to “lock down” or secure Linux systems by implementing measures such as disabling unnecessary services, configuring permissions for maximum security, creating chroot jails, and so on. Bastille is implemented as a Perl script and can be run in two modes:

- **Interactively:** The user is prompted for action at each step of the hardening process, the advantage being that users are educated concerning how to harden their systems.
- **Noninteractively:** The script makes appropriate decisions about how to harden the system and does this automatically, the advantage being that a

standard secure operating system can be deployed quickly and easily on multiple systems.

Bastille also supports a revert feature that allows you to return your system to its prehardened state if problems arise running the script.

Marketplace

Bastille is currently available for several Linux distributions, including Debian, Mandrake, and Red Hat. It has also been ported to the HP-UX platform and is released under the General Public License (GPL).

For More Information

See www.bastille-linux.org for more information.

See Also: *chroot jail, hardening*

bastion host

A host that is fully exposed to attack on a public network.

Overview

Bastion hosts usually reside on the outside of a company’s demilitarized zone (DMZ) and are thus completely exposed to attack by malicious users on the Internet. In fact, bastion hosts generally must be exposed in order for them to perform their functions, and examples of such hosts include the following:

- Web servers
- File Transfer Protocol (FTP) servers
- Mail servers and Simple Mail Transfer Protocol (SMTP) forwarders
- Domain Name System (DNS) name servers
- Firewalls and gateways

Because these hosts are exposed, special care must be taken to harden them; that is, to make them **bastions**, a medieval word describing the highly fortified portion of a castle.

Implementation

Examples of hardening procedures necessary for bastion hosts include the following:

- Performing clean installation of operating system and server applications

- Applying service packs, patches, and hotfixes as soon as they become available
- Removing unnecessary server configuration tools and utilities
- Disabling unnecessary services and daemons
- Blocking unnecessary Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) ports
- Modifying access control lists (ACLs) on file system objects for maximum security
- Encrypting local password files and local user account databases
- Logging all system activity and regularly auditing system logs

See Also: demilitarized zone (DMZ), hardening

BBBOnLine

A reliability program developed by the Better Business Bureau (BBB) to help protect the privacy of consumers in online transactions.

Overview

With the growth of online shopping and e-commerce sites in recent years has come a heightened concern by consumers about their privacy in online transactions. To help alleviate such concerns and promote responsible business practices, the Better Business Bureau has developed its BBBOnLine program, which allows participating business to display a BBB seal of approval on their Web sites to demonstrate their commitment to ensuring consumer privacy. When visitors go to online shopping sites that display the BBBOnLine logo, they can click on the logo and be redirected to the BBB site, where they can view a reliability report concerning the past marketplace performance of the business. Based on this report, they can then make an informed decision about whether to shop there.

For More Information

Visit www.bbbonline.org for more information.

See Also: privacy

behavior-blocking software

Software that detects and prevents suspicious behavior from being executed on a system.

Overview

Behavior-blocking technology is designed to complement but not replace antivirus software and provide an additional layer of protection against worms, viruses, Trojans, and similar problems confronting today's networks. Traditional antivirus software uses a signature-based approach for recognizing and eradicating infections and attack vectors. The disadvantage of this approach is that signatures can usually only detect known attacks and may not be updated until hours or days after a new virus has been reported to the antivirus software vendor. In the meantime, a company's network may become infected and systems may be taken out, resulting in lost time and money.

Behavior-blocking software works differently and attempts to identify malware by the actions it tries to perform such as mass mailing to everyone in a user's address book or attempting to access the registry. As such, behavior-blocking technologies have an advantage over traditional antivirus software of providing real-time protection against new forms of attack that cannot be detected by a signature-based approach. The downside of behavior-blocking technologies, however, is that it is sometimes difficult to distinguish between legitimate and malicious behavior, with the result that false positives are often generated. This can even cause problems with legitimate applications whose actions may be interpreted by the behavior-blocking software as malicious, causing frustration for users and lost business time. Another disadvantage of this technology is that while antivirus software works automatically to provide protection, behavior-blocking software usually requires some form of user intervention to analyze blocked behaviors in order to distinguish between genuine attacks and false positives. Nevertheless, because behavior-blocking software is often the only way to defend against new types of threats, many companies are beginning to see it as an essential adjunct to more traditional security measures such as antivirus software and firewalls.

Implementation

There are two general approaches to implementing behavior blocking in a system. One approach is for the software to hook into the kernel and intercept system calls for file system access, registry access, Component Object Model (COM) object access, and so on. Interceptor modules trap these system calls and apply heuristics configured using policies to determine whether the call is legitimate and whether to allow or deny access to the resource being called.

Another approach to behavior blocking is to intercept incoming mobile code such as ActiveX objects, Java applets, and other executable code and scripts that can arrive by way of Web browsers or mail clients. This code is then “sandboxed” by restricting the level of access it has to system resources based on how the software’s policy settings are configured. Some behavior-blocking systems combine both of these approaches for greater flexibility.

Behavior-blocking software is generally installed on both servers and clients to provide maximum protection against infection by new agents. Real-world experience has demonstrated the usefulness of such software, which has been able to identify and stop the actions of dangerous worms such as Code Red and Nimda before antivirus vendors have been able to create signatures to recognize them.

Marketplace

While behavior-blocking technologies have been around for some time now, in the last couple of years interest in them has skyrocketed with the proliferation of Internet worms, viruses, and other threats impacting corporate networks through messaging systems and the Internet. A number of vendors have produced products for protecting systems and networks using behavior blocking; examples are eSafe Gateway from Aladdin Knowledge Systems (www.esafe.com), SurfinGate and SurfinShield from Finjan Software (www.finjan.com), SafeTNet from Pelican Security (www.pelicansecurity.com), and InterScan AppletTrap from Trend Micro (www.trendmicro.com). Other vendors of behavior-blocking software include Entercept, Granite Technology, Harris Corporation, Okena, and Sandbox Security.

Firewall companies such as CheckPoint have also begun to incorporate behavior-blocking technologies into their products following the lead of Tiny Software’s Personal Firewall 3 in this regard.

See Also: *firewall, intrusion detection system (IDS), virus, virus protection software, worm*

BER

Stands for Basic Encoding Rules, a set of rules used for encoding ASN.1-defined data into a bitstream.

See: *Basic Encoding Rules (BER)*

biometric identification

The process of using a person’s physical characteristics for identification purposes.

Overview

Biometrics is the science of identifying individuals using physical characteristics and behaviors. Examples of physical attributes that may be used for such purposes include using a person’s fingerprints, hand geometry, iris or retina, facial characteristics, voice pattern, or even body odor. A person’s DNA can also be used to uniquely identify him or her, but this is a more invasive process and requires a skin, tissue, or blood sample. Behaviors that can be used to identify people include computer keystroke dynamics, walking patterns, and how a person responds to a standard set of questions.

Biometrics has been around in nascent form since the 19th century, when police forces first used fingerprinting to identify possible criminals, and automated biometric technology was pioneered by defense agencies in the 1970s using voice-, iris-, retinal-, and fingerprint-scanning equipment to allow or deny individuals access to restricted sites. In the late 1990s, however, biometric hardware became a commodity that even small companies could afford, and biometric authentication is starting to become widespread in corporate networking environments, in the banking and financial industries, and in government.

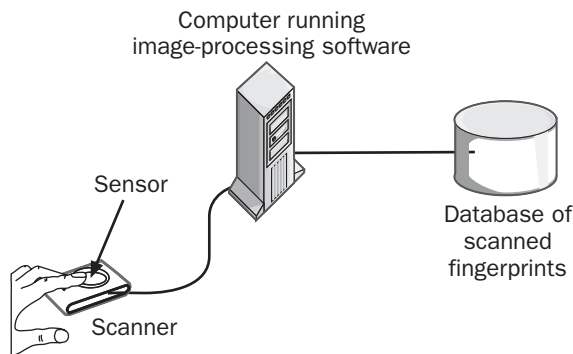
B

Implementation

The most popular biometric technologies at present are those used for fingerprint identification, iris scanning, and facial recognition. In general, any biometric system consists of three parts:

- A high-resolution scanning device that can be used to acquire an image of a person's physical characteristic and digitize it
- A storage system containing a database of digitized images of authorized individuals
- A computer system running image-processing software that can compare the acquired image with the database to recognize a match

In a typical biometric fingerprinting system, an individual places his or her index finger on a silicon sensor acting as a capacitor that is continually charged and discharged. The sense uses the ridges on the person's finger to generate an image of the fingerprint, which is then scanned at high resolution and converted into digital form. The scanned image is transferred into a computer using a universal serial bus (USB) or serial connection, where image-processing software compares it with a known database of digitized signatures. If a match is found, the system can be used to generate an authentication token that allows the individual access to the computer, network, or building controlled by the system.



Biometric identification. How a biometric fingerprinting system works.

Marketplace

A number of companies have established themselves as vendors in the emerging biometrics market. Examples in the fingerprint-imaging field include not only such industry heavyweights as Compaq and UNISYS but also smaller companies such as Digimarc, Identix, and Vitrix. Biometrika, Cognitec Systems, eTrue, and Visionics have developed facial recognition systems. Iridian is a leading vendor of iris scan technology, while voice recognition vendors include BioID, Nuance, and VeriVoice.

Issues

When automated biometric technologies first appeared in the 1970s, many people expressed concerns about their privacy being invaded by having digitized information about their physical characteristics stored in government databases. Others have argued that biometrics actually protects peoples' identities against the rising crime of identity theft. Biometric systems are not infallible, and while few people now argue with using biometrics for authentication purposes, civil rights advocates often argue that face recognition systems in public places such as airports are an invasion of privacy and that false positives may lead to harassment by airport authorities. In the post-9/11 world, however, the momentum for increasing use of biometric screening is likely to continue and grow.

See Also: authentication, identity theft

BIOS cracking

Compromising or resetting the password protecting a computer's basic input/output system (BIOS).

Overview

The BIOS contains the basic hardware configuration settings for a computer system, and best practices suggest that you should configure a password to protect your system's BIOS against unauthorized access. For example, you could configure the BIOS on desktop machines to disable the floppy disk drive and thus prevent users from installing shareware or becoming infected with viruses by sharing software on floppy disks. Another example would be to disable booting

from CD-ROM to prevent users from installing parallel operating systems to access files on a system's hard drive.

BIOS cracking refers to methods used to hack a computer's BIOS password or simply to reset it to null so that the BIOS can be accessed without a password. BIOS cracking usually requires interactive access to the local computer's console, so physically securing your systems is generally the best protection against this.

Notes

If you forget your BIOS password, some older BIOSes have backdoor passwords or reset procedures that may help. Contact your system manufacturer or search the Internet for more information on backdoor passwords. You may also be able to change a jumper on your motherboard or remove your complementary metal-oxide semiconductor (CMOS) battery to discharge your BIOS CMOS and clear the settings, after which you can flash your BIOS to restore the defaults. There are also tools such as Password Reminder from NewPowerSoft that can sometimes be used to display your BIOS password, depending on the type of BIOS.

See Also: password

black hat

Euphemism for a malicious hacker.

Overview

The term **black hat** can be used several ways. Malicious hackers who try to break into corporate networks to obtain sensitive information or simply to display their skills often wear this term as a badge of their participation in the underground hacking community. Alternatively, legitimate security experts (known as **white hats**) use the term as a pejorative to denounce the goals and intentions of malicious hackers. The origin of the term **black hat** is obscure but is probably linked to the practice of bad cowboys in the Old West who wore black hats to distinguish themselves from the good guys who wore white hats (or perhaps Hollywood's representation of such a practice!).

See Also: gray hat, hacker, white hat

Black Hat Briefings

Annual security conference held at various locations around the world.

Overview

Black Hat Briefings is a gathering of security experts, both legitimate and "underground," during which they spend two days discussing the latest security tools, problems, and incidents. The name of the conference is derived from the term **black hat** and suggests that the conference includes the participation of hackers with malicious intent. The philosophy behind this approach seems to be that to recognize activities of hackers you need to be one, and the tracks and sessions include coverage of both how to use hacking tools to penetrate a system or network and how to use tools to detect and prevent such intrusions. The conference is usually well attended by both "white hat" security experts and underground hackers, corporate security officers, media representatives, and law enforcement agencies. Immediately following the Las Vegas conference is another conference called Defcon, the largest hacker convention in the United States.

For More Information

Visit www.blackhat.com for more information.

See Also: black hat, Defcon, hacker, white hat

blackholing

Automated monitoring of entire networks for detecting threats such as worms or scans.

Overview

Blackholing is an extension of the honeypot concept, whereby a system emulates an entire network of systems, acting as a honeypot to try to attract and identify intrusions. While a honeypot is a system that mimics a legitimate network server, a blackholing monitor simulates general activity and traffic from the entire network by creating large numbers of virtual servers running services that you specify. Randomly targeted threats such as worms are generally sent to entire blocks of IP addresses, and the blackholing monitor responds to these threats for all unutilized addresses on the network,

mimicking legitimate hosts to gather data that can be used to identify trends of attacks.

An example of a tool that can be used to perform black-holing is Honeyd, released in 2002 as an open source package for UNIX platforms by Niels Provos of the University of Michigan.

For More Information

Find out more about Honeyd at www.citi.umich.edu/w/provos/honeyd/.

See Also: honeypot

BLOB

A generic sequence of bits used for storing data.

Overview

BLOB, which originally stood for binary large object (but is now simply known by the acronym), represents a generic data structure used for storing large amounts of data such as images, video, or attachments to e-mail messages. BLOBs typically contain one or more fixed-length header structures followed by data whose format depends on the context in which the data is being used.

BLOBs are used in Microsoft's CryptoAPI (CAPI) application programming interface in several places:

- **Attribute BLOB:** Contains an encoded representation of attribute information stored in a certificate request
- **Certificate BLOB:** Contains an encoded representation of data stored in a certificate
- **Certificate name BLOB:** Contains an encoded representation of name information (such as name of issuer or subject) stored in a certificate
- **Key BLOB:** Contains an encoded representation of an encrypted private key that is created by exporting the key

See Also: CryptoAPI (CAPI), digital certificate, key, signature

block cipher

Cipher algorithm that encrypts data in discrete chunks called blocks.

Overview

A block cipher is a cipher that encrypts or decrypts multiple bits of data simultaneously, usually 64 bits at a time (though Advanced Encryption Standard [AES] employs larger blocks of 128, 192, or 256 bits). As a result of this approach, block ciphers are generally faster than stream ciphers, which encrypt data as a continuous stream of bits. Each block of data is generally encrypted using the same encryption key, with the result that identical blocks of plaintext generate identical ciphertext (this can be avoided by using cipher block chaining, which makes identical blocks of plaintext generate different ciphertext by inserting additional information into each block).

Popular examples of block ciphers include Data Encryption Standard (DES), 3DES, International Data Encryption Algorithm (IDEA), AES, and Blowfish.

See Also: 3DES, Advanced Encryption Standard (AES), Blowfish, ciphertext, Data Encryption Standard (DES), encryption, International Data Encryption Algorithm (IDEA), plaintext

Blowfish

An unpatented, royalty-free encryption algorithm.

Overview

Blowfish was developed in 1993 by Bruce Schneier as a free alternative to existing encryption algorithms such as Data Encryption Standard (DES) and International Data Encryption Algorithm (IDEA). Blowfish is implemented as a standard 64-bit block cipher with a variable key length that can range from 32 to 448 bits. Blowfish runs much faster than DES or IDEA and provides strong encryption for applications and systems that use it. Blowfish is currently used in over 150 products, including the OpenBSD operating system and Linux kernel. Its source code is freely available from Counterpane Labs.

For More Information

Visit Counterpane Labs at www.counterpane.com for more information about Blowfish.

See Also: *Data Encryption Standard (DES), encryption algorithm, International Data Encryption Algorithm (IDEA), strong encryption*

B02K

Stands for Back Orifice 2000, a well-known Trojan horse program for the Microsoft Windows NT and Windows 2000 platforms.

See: *Back Orifice 2000 (BO2K)*

boink attack

A modified version of the bonk attack.

Overview

The original bonk attack allowed an attacker to crash machines running Windows 95 and Windows NT by sending corrupt User Datagram Protocol (UDP) packets to port 53 (the Domain Name System [DNS] port). The boink attack expands on this by allowing the attacker to attack multiple ports simultaneously. Both forms are variants of the more general “teardrop attack” and can be prevented by applying the most recent service packs and more generally by using a properly configured firewall.

Implementation

Boink (and bonk) work by manipulating the fragment offset field in Transmission Control Protocol/Internet Protocol (TCP/IP) packets in a way that causes the target system to think that incoming UDP packets are all part of a larger original packet that was fragmented. The target system tries to reconstruct the original packet, which turns out to be too large for the networking subsystem to handle, and the result is that the target machine hangs or crashes, resulting in services being denied to legitimate network traffic.

The boink attack thus falls under the classification of denial-of-service (DoS) attacks. No damage is generally done to the target system, and after a reboot, the system runs as usual or until the attack resumes. Boink attacks are usually detected by intrusion detection

systems (IDSs) by their signature, which involves unusual levels of fragmented packets.

Notes

The attack is named after the tool (Boink) used to perform it.

See Also: *bonk attack, denial of service (DoS), intrusion detection system (IDS), teardrop attack*

bonk attack

A variant of the teardrop attack.

Overview

The bonk attack originally targeted the Windows 95 and Windows NT platforms and could crash or deny services to these systems by sending malformed User Datagram Protocol (UDP) packets to port 53, the standard Domain Name System (DNS) port. A patch for this attack was developed for Winsock to prevent this attack, and in general by installing the latest service packs and hotfixes, variants of this attack can be prevented on all Windows platforms. The bonk attack is generally classified as an early example of a denial of service (DoS) attack. For an explanation of how bonk works, see the article **boink attack** earlier in this chapter.

Notes

The attack is named after the tool (Bonk) used to perform it.

See Also: *boink attack, denial of service (DoS), teardrop attack*

Brown Orifice

A backdoor that exploited a vulnerability in Netscape’s version of the Java Virtual Machine.

Overview

Brown Orifice, named after the famous Back Orifice remote administration tool, exploits a vulnerability in how Java is implemented in version 4.7 and earlier of Netscape Navigator. When a user simply visits a Web site on which the Brown Orifice applet is present, the applet runs on the client machine, turning it into a stealth Web server operating on port 8080 and allowing the attacker to gain full access to files stored on the

user's machine. Brown Orifice is not a security issue in the Java programming language, however, but in how the Java Virtual Machine was implemented in earlier versions of Navigator, and this vulnerability has been patched in later versions of the product. Brown Orifice was created by Dan Brumleve in 2000.

See Also: *Back Orifice, vulnerability*

BRP

Stands for business resumption plan, a detailed plan on how to resume normal business after a disaster.

See: *business resumption plan (BRP)*

brute-force attack

An attack based on systematically trying all possible keys for a secure system.

Overview

The brute-force approach originally referred to any computer program that relied on sheer processing power instead of intelligence. For example, solving a quadratic equation such as $x^2 + 7x - 44 = 0$, where x is an integer, using brute force simply involves writing a program that tries all possible integral values of x until an answer is found. The programmer's motto "when in doubt, use brute force" is attributed to Ken Thompson, a co-inventor of UNIX.

When this concept is applied to cryptography, the result is the brute-force attack, which tries to decode a cipher by guessing at every possible key until the correct one is found. The feasibility of such an approach obviously depends on the length of the key, the computational power available for the process, and the patience of the attacker.

Brute-force methods can also be used to try to crack passwords for secure systems, again by simply trying all possible strings of characters in succession. Such an approach is easily defeated by using a sufficiently long password string, and a dictionary attack is often more profitable for the attacker to pursue than simple brute force.

Another area in which brute force is often used is in trying to compromise the security of networks that use Simple Mail Transfer Protocol (SNMP). An attacker can launch a brute-force attack that tries to guess the SNMP community names in order to profile the devices and services running on the network. Again, the simplest way of defeating this approach is to use long and complex strings for community names.

A popular tool for using brute force to authenticate against Hypertext Transfer Protocol (HTTP), File Transfer Protocol (FTP), telnet, and other servers is Brutus, which can establish multiple connections with the server in order to speed the process. In general, the best defense against such tools is an intrusion detection system (IDS), which can detect the anomalous nature of network traffic when a brute-force attack is under way.

See Also: *cipher, dictionary attack, intrusion detection system (IDS), password cracking*

bucket brigade attack

More commonly called a **man-in-the-middle (MITM) attack**, an attack in which a third party intercepts an encrypted communication and masquerades as the other party to each party.

See: *man-in-the-middle (MITM) attack*

buffer overflow

Usually called buffer overrun, a condition resulting from adding more information to a buffer than it was designed to contain.

See: *buffer overrun*

buffer overrun

Also called buffer overflow, a condition resulting from adding more information to a buffer than it was designed to contain.

Overview

A **buffer** is a region of memory that is used as a temporary repository for holding information. A buffer overrun is a condition that may occur when too much data is

placed in the buffer, creating a vulnerability that an attacker can often exploit to harm the application or system. Buffer overruns are generally caused by poorly coded validation and error-handling routines and can be prevented by exercising proper coding procedures. They are particularly prevalent in programs coded in the C and C++ languages because of the way these languages handle memory violations.

When more data is placed in a buffer than it was originally designed to contain, the resulting buffer overrun can result in the application hanging or crashing. In some cases, if malicious code is included in the data that overruns the buffer into the stack, this code can execute on the system, causing damage or even providing attackers with elevated privileges.

In order to discover a potential buffer overrun in a product, an attacker needs deep knowledge of C/C++, assembly language, and the application programming interfaces (APIs) of the target system. Buffer overruns in popular software can have widespread impact; one of the earliest examples was in 1998 when a worm that exploited a buffer overflow in the UNIX finger service caused more than 6000 systems on the Internet to crash.

Software platforms and products from virtually all vendors have occasionally been found to have buffer overrun problems, mainly because of the rapid pace of software development in recent years, which has resulted in lower quality control for coding practices. The best way to handle this issue is to apply service packs and patches as soon as vendors release them.

See Also: *attack*

Bugtraq

A popular mailing list for announcing and discussing recently discovered security vulnerabilities.

Overview

Bugtraq is a security advisory mailing list that covers hacking and computer security and has been in existence since 1993. It currently has over 27,000 subscribers and is used for discussion of a wide variety of vulnerabilities, threats, and exploits and how to prevent and

recover from them. The list is moderated to keep the noise to a minimum and is managed by SecurityFocus, which has recently been acquired by Symantec.

For More Information

Visit www.symantec.com and www.securityfocus.com for more information.

See Also: *exploit, threat, vulnerability*

bulk encryption key

A session key used in encrypted messaging.

Overview

Session keys are one-time, temporary keys used in encrypted communications. Bulk encryption keys are session keys derived from the key used to encrypt the message. In secure messaging, the bulk encryption key is itself encrypted using the recipient's public key and is sent to the recipient together with the encrypted message. The encrypted bulk encryption key is sometimes referred to as a **lock box**. When the message and lock box are received, the lock box is decrypted to obtain the bulk encryption key, and the bulk encryption key is then used to decrypt and read the message. The reason for using this approach is that the session key is generated using symmetric key encryption, a method that is faster for encrypting and decrypting bulk data such as e-mail messages than the slower public key encryption method.

See Also: *key, public key cryptography, session key*

business continuity plan

Another name for business resumption plan (BRP), a detailed plan on how to resume normal business after a disaster.

See: *business resumption plan (BRP)*

business resumption plan (BRP)

A detailed plan on how to resume normal business after a disaster.

B**Overview**

Business resumption plans (BRPs) are an essential part of disaster recovery planning and are designed to facilitate the speedy, orderly, and systematic restoration of normal business activity after a disaster has occurred. Such plans should also specify how critical business functions can continue to operate normally during the recovery period. Developing a suitable business resumption plan for your business involves the following activities:

- Determining your critical business requirements
- Developing recovery strategies for different business elements
- Developing an emergency response team and problem escalation ladder

- Specifying those individuals who have the authority and responsibility for activating different portions of the plan
- Training staff in how to function in a recovery environment
- Testing the plan to ensure it works properly and keeping it current

Notes

Another name for such a plan is business continuity plan.

See Also: *disaster recovery plan (DRP)*

CA

Stands for certificate authority, a trusted entity that issues digital certificates.

See: *certificate authority (CA)*

CA certificate

A certificate that verifies the identity of a certificate authority (CA).

Overview

In order for certificates issued by CAs to be trusted, they must be signed by the CA itself. In order to sign certificates, the CA requires its own certificate, which is called a **CA certificate**. This CA certificate contains the public key of the CA and is used for signing certificates issued to users, applications, and systems requesting them.

If the CA is part of a chain or hierarchy of CAs, a CA certificate is usually signed by the CA immediately above it in the hierarchy. If the CA is at the top of the hierarchy (a root CA), the CA usually signs its own certificate, which is called a self-signed or root certificate.

See Also: *certificate authority (CA), digital certificate, root CA*

CA hierarchy

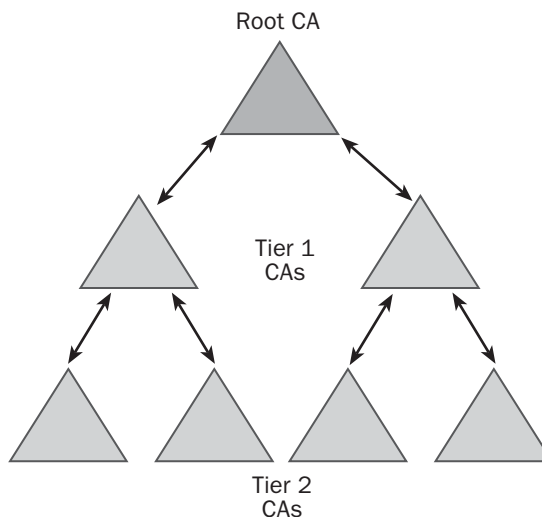
Also called a **hierarchy of trust**, a hierarchical collection of certificate authorities (CAs) bound together by trust relationships.

Overview

In a large organization such as a government agency or multinational enterprise, individual departments or locations may need to manage their own certificates by setting up their own CA. In order for certificates issued by one department to be accepted as valid by entities in

another department, trusts must be established between the CAs in different departments. The usual way of doing this is to establish a hierarchy of CAs, with each CA issuing and signing a certificate for the CAs immediately beneath it.

The top of a CA hierarchy is called the root CA and is universally trusted by all CAs in the agency or enterprise. The root CA signs its own certificate to guarantee its identity and issues signed certificates to lower-level or subordinate CAs beneath it to guarantee their own identities. The root CA itself may not issue certificates to other entities directly, and instead entities such as users, applications, or systems must contact the subordinate CA (otherwise known as an **issuing CA**) in their department to request a certificate.



CA hierarchy. Example of a CA hierarchy.

See Also: *certificate authority (CA), digital certificate, root CA*



cache poisoning

A method used for misdirecting certain types of network traffic.

Overview

Many forms of network services utilize caching to improve their performance. An example is the Domain Name System (DNS), in which DNS servers cache name resolution requests from clients in order to answer repeated requests more quickly. DNS cache poisoning can be prevented by patching DNS servers with the latest versions of their software, but because of the distributed nature of DNS and its use of recursive queries, cache poisoning can only be prevented by a cooperative effort of the entire Internet community, including agencies responsible for managing top-level domains.

Another example is the address resolution protocol (ARP), a Transmission Control Protocol/Internet Protocol (TCP/IP) protocol that resolves Internet Protocol (IP) addresses into Media Access Control (MAC)-layer addresses. ARP caches resolved address mappings to reduce unnecessary network traffic and speed communications between hosts on the network. Both of these protocols are subject to cache poisoning, in which spoofed packets are used to modify cached information so that future requests for such information result in misdirected traffic.

See Also: *ARP spoofing, DNS spoofing, spoofing*

callback

A security method used in remote access.

Overview

When a client tries to authenticate with a remote access server on which callback has been configured, the access server terminates the connection attempt and then calls the client back at a previously configured telephone number. This helps verify the identity of the client to the server, as only the client can respond from the configured number, although in reality this is relatively weak protection. Callback can also be used to reverse the charges for the client connection to help clients avoid paying long-distance calling charges.

Canadian Centre for Information Technology Security (CCITS)

An organization that provides education and research on computer security and high-tech criminal investigation.

Overview

Canadian Centre for Information Technology Security (CCITS) is a joint initiative of the University of British Columbia and the Justice Institute of British Columbia and offers a Certificate in Internet and Technology Security program that teaches best practices in information security. The program is designed for a wide range of professionals, including security managers, systems administrators, and law enforcement personnel, and provides comprehensive training in risk assessment, auditing, policies, procedures, and practices for securing information systems.

For More Information

Visit CCITS at www.ccits.org for more info.

canonicalization error

A coding error that can cause applications to be vulnerable to attack.

Overview

Canonicalization is the process by which different forms of a name are resolved into a single, standard name called the canonical name. A canonicalization error is a parsing error that allows an attacker to submit a malformed name (typically a malformed URL, or Uniform Resource Locator, submitted to a Web server) that causes incorrect permissions to be applied to the resource being accessed. File system resources typically inherit their permissions from the folder in which they reside, but when a canonicalization error occurs, the file may gain its permissions from a grandparent instead, that is, a folder higher up in its parentage chain. If the grandparent folder has less-restrictive permissions than the parent folder, the attacker has succeeded in gaining additional permissions, and it may be possi-

ble to utilize these extra permissions for launching some kind of attack.

See Also: *elevation of privileges (EoP)*

CAPI

Stands for Microsoft CryptoAPI, a set of application programming interfaces (APIs) for cryptography built into Microsoft Windows platforms.

See: *CryptoAPI (CAPI)*

CAPICOM

A Component Object Model (COM) interface for the Microsoft CryptoAPI (CAPI) programming interface.

Overview

CAPICOM is an ActiveX control that provides a COM interface to CryptoAPI (CAPI). CAPICOM exposes the cryptographic functions of CryptoAPI (CAPI) using COM so that developers can easily write applications that encrypt or decrypt data, digitally sign messages, generate and manage keys, and perform other cryptographic actions. Since CAPICOM is a COM interface, it can be accessed from a variety of programming environments including Active Server Pages (ASP) and ASP.NET, Visual Basic Scripting Edition (VBScript), JScript, C++, C#, and VB.NET. Because CAPICOM is implemented as an ActiveX control, it can easily be embedded in Web pages to cryptographically enable dynamic Web applications.

See Also: *CryptoAPI (CAPI)*

Carnivore

Now called **DCS-1000**, a surveillance technology used by the FBI for monitoring e-mail.

See: *DCS-1000*

CAS

Stands for code access security, a code security mechanism built into Microsoft Windows .NET Framework.

See: *code access security (CAS)*

CAST

A family of symmetric encryption algorithms.

Overview

CAST is a symmetric block cipher developed by cryptographer Carlisle Adams. CAST is similar to Data Encryption Standard (DES) in operation. Its original form, CAST-128, uses a 128-bit key with 16 successive rounds of application on 64-bit blocks of plaintext. An extension called CAST-256 uses a key twice the size of the original version.

CAST is available royalty-free for commercial or private use. CAST has been used in several products ranging from Pretty Good Privacy (PGP) to Microsoft CryptoAPI (CAPI).

The detailed operation of CAST is outlined in RFC 2144.

See Also: *block cipher, encryption algorithm*

CBC

Stands for cipher block chaining, a feedback mechanism commonly used in block ciphers.

See: *cipher block chaining (CBC)*

CCA

Stands for Common Cryptographic Architecture, a cryptographic architecture developed by IBM for its computing platforms.

See: *Common Cryptographic Architecture (CCA)*

CCITS

Stands for Canadian Centre for Information Technology Security, an organization that provides education and research on computer security and high-tech criminal investigation.

See: *Canadian Centre for Information Technology Security (CCITS)*

cDc

Stands for Cult of the Dead Cow, a notorious group of underground hackers.

See: *Cult of the Dead Cow (cDc)*

Center for Education and Research in Information Assurance and Security (CERIAS)

A center for research and education in information security at Purdue University.

Overview

The Center for Education and Research in Information Assurance and Security (CERIAS) is a well-known leader in research in computer, network, and information security and information assurance. CERIAS takes a multidisciplinary approach to research and education in these fields using the involvement of academia, government, and industry. The community of scholars involved in CERIAS works on solving fundamental problems in information security, participates actively in security organizations in government and industry, provides leadership in community-based education in information assurance and security, and assists organizations with their expertise.

The CERIAS affiliate program sponsors a variety of research projects at collaborating research centers and laboratories on subjects such as intrusion detection, denial of service (DoS) attacks, information privacy, network security, virtual computing, and many other topics. CERIAS also offers a graduate certificate for educators who want to develop information assurance programs at their colleges and universities.

See Also: *information assurance (IA)*

Center for Internet Security (CIS)

A nonprofit organization that helps organizations manage risk associated with information systems security.

Overview

The Center for Internet Security (CIS) is a cooperative of over 170 organizations from business, education, government, law enforcement, and professional associations that work together to provide tools and recommendations for measuring, monitoring, and improving information systems security. To meet these goals, CIS develops and publishes benchmarks that represent best practices in securing operating systems such as Windows 2000, Solaris, HP-UX, Linux, and IOS. These benchmarks provide detailed instructions for how to harden systems and include scoring tools for measuring systems against the benchmark and generating a variance report.

For More Information

Visit CIS at www.cisecurity.org for more information.

CERIAS

Stands for Center for Education and Research in Information Assurance and Security, a center for research and education in information security at Purdue University.

See: *Center for Education and Research in Information Assurance and Security (CERIAS)*

CERT/CC

Stands for CERT Coordination Center, a center of Internet security expertise operated by Carnegie Mellon University.

See: *CERT Coordination Center (CERT/CC)*

CERT Coordination Center (CERT/CC)

A center of Internet security expertise operated by Carnegie Mellon University.

Overview

The CERT Coordination Center (CERT/CC) is a federally funded research center that started in 1988 as a project of the Defense Advanced Research Projects Agency (DARPA). CERT/CC studies security

vulnerabilities in the Internet, publishes advisories and incident notes, recommends best practices for securing networks, and provides training and advice on how to develop computer security incident response teams. CERT/CC takes a technology-neutral approach but also provides specific recommendations for hardening specific operating system platforms. CERT/CC is widely recognized as a leader in information systems security and collaborates with business and industry to help make the Internet a safer place.

For More Information

Visit the CERT/CC Web site at www.cert.org for more information.

certificate

Properly called a **digital certificate**, encrypted information that guarantees that an encryption key belongs to a user.

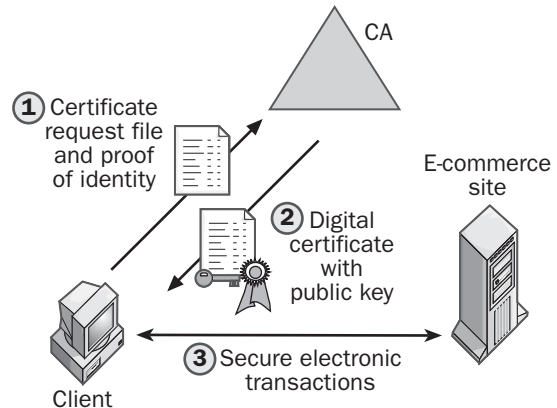
See: digital certificate

certificate authority (CA)

Also called **certification authority**, a trusted entity that issues digital certificates.

Overview

Certificate authorities (CAs) form the foundation of Public Key Infrastructure (PKI) systems and are responsible for issuing digital certificates in response to certificate requests, maintaining a certificate store of issued certificates, and maintaining and publishing a certificate revocation list (CRL) of expired, invalid, or compromised certificates. CAs can be stand-alone entities or part of a hierarchy or web of trust. At the top of a hierarchy sits the root CA, which issues certificates to other CAs to establish their identity (the root CA issues a certificate to itself to establish its own identity). Depending on how a PKI system is implemented, CAs may coexist with or cooperate with registration authorities (RAs) to validate the identity of users requesting certificates.



Certificate authority (CA). How a client obtains a digital certificate from a certificate authority.

Marketplace

CAs may include government agencies, commercial companies, or in-house authorities set up and managed by IT (information technology) departments of large organizations. Public certificate authorities widely recognized in the marketplace include Thawte, Verisign, and several others. Commercial software for enterprises to set up and manage their own internal certificate authorities is available from Microsoft, Sun, Netscape, RSA, and many other vendors.

See Also: certificate revocation list (CRL), digital certificate, Public Key Infrastructure (PKI), root CA

certificate-based authentication

Authentication of users by digital certificates.

Overview

Certificate-based authentication can be used to provide external users with secure access to resources on your network. The external user is first granted a certificate from a trusted certificate authority (CA). A user account is then created in the company directory and a mapping is established between the certificate and the account. When the external user wants to access

resources on the company network, the user presents the certificate to an authentication server that verifies it and grants access based on access control lists (ACLs) for the mapped account. One advantage of this approach is that a single certificate can be mapped to multiple accounts, allowing a department of one company, for example, to access resources in another company as part of a supply-chain relationship or business partnership.

Certificate-based authentication is supported by Active Directory directory service on the Microsoft Windows platform.

See Also: authentication, digital certificate

Certificate Information Systems Auditor (CISA)

A widely accepted certification in auditing, control, and security of information systems.

Overview

Certificate Information Systems Auditor (CISA) is a certification issued by the Information Systems Audit and Control Association (ISACA). The certification tests understanding of knowledge and practice in several areas, including disaster recovery and business continuity, protection of information assets, business process evaluation and risk management, and IS (information systems) audit processes. These areas form a foundation for good security practices for planning, implementing, and evaluating secure information systems. CISA is a recognized certification that has been around since 1978 and is held by more than 29,000 individuals worldwide.

For More Information

Visit the ISACA Web site at www.isaca.org for more information on CISA.

See Also: Information Systems Audit and Control Association (ISACA)

certificate request

A specially formatted message requesting a certificate from a certificate authority (CA).

Overview

In order for an entity such as a user or application to obtain a digital certificate, a request must be submitted to the appropriate CA. This request must be properly formatted and contain the information needed by the authority to grant the request. The entity then submits the request along with its public key to the CA, which then issues the requested certificate.

The standard format for certificate requests in Public Key Infrastructure (PKI) systems is the X.509 certificate request message format outlined in RFC 2511.

See Also: certificate authority (CA), digital certificate, public key

certificate revocation list (CRL)

A list of revoked certificates maintained by a certificate authority (CA).

Overview

Certificates are digitally signed statements issued by CAs to entities requesting them. These certificates can then be used to perform secure electronic transactions such as e-commerce or online banking. In order to prevent the abuse of such privileges, certificates that are lost, stolen, or expired must be readily identifiable to the parties involved, and for these purposes CAs maintain and publish a CRL of previously issued certificates that are no longer valid. By consulting such a list prior to completing a transaction, commercial parties are protected from liabilities arising from invalid certificates.

See Also: certificate authority (CA), digital certificate

certificate server

A server that issues a certificate for a certificate authority (CA)

Overview

Digital certificates are issued and managed by applications called **certificate servers**. These applications are designed to automatically process certificate requests, issue certificates, maintain a central store or database of issued certificates, and publish a certificate revocation

list (CRL) of expired, lost, or stolen certificates. Certificate servers form the basis of the operation of Public Key Infrastructure (PKI) systems upon which secure electronic transactions such as online banking and e-commerce depend.

Marketplace

Microsoft Windows Server 2003 includes a Certificate Services component that can be used to set up a PKI system for enterprise or commercial use. Certificate server applications from other vendors include Netscape Certificate Server, Sun ONE Certificate Server, Novell Certificate Server, PGP Certificate Server, and many others.

See Also: certificate authority (CA), certificate revocation list (CRL), certificate store, digital certificate, Public Key Infrastructure (PKI)

certificate store

A central database of certificates issued and maintained by a certificate authority (CA).

Overview

When a CA issues a certificate to an entity, the authority must maintain a copy of the certificate for reference purposes. These certificates are kept in a special database called a certificate store, which typically contains three things:

- Certificates issued to entities requesting them
- Certificate revocation lists (CRLs) of expired, lost, or stolen certificates
- Certificate trust lists (CTLs) of trusted certificate authorities and other trusted items

See Also: certificate authority (CA), certificate revocation list (CRL), certificate trust list (CTL), digital certificate

certificate trust list (CTL)

A group of items signed by a trusted certificate authority (CA).

Overview

Certificate trust lists (CTLs) can contain any information signed by a trusted entity, such as documents, lists of file names, or lists of hashes of certificates. By having these items signed by a trusted entity, their authenticity and ownership is validated and ensured. For example, CAs themselves maintain CTLs in their certificate stores to identify other CAs they themselves trust.

Another example would be Web servers that authenticate clients based on client certificates. Such servers can maintain their own CTLs containing information about which CAs are trusted by the server. If a client tries to authenticate using a certificate signed by an authority not present in the server's CTL, the server rejects the authentication attempt.

Web browsers also need to maintain their own CTLs that specify which CAs they trust. This is necessary when the browser needs to verify the identity of a Web server using the server's own certificate, for example, in secure online banking.

See Also: certificate authority (CA), digital certificate

Certified Information Systems Security Professional (CISSP)

A widely accepted certification for computer security professionals.

Overview

Certified Information Systems Security Professional (CISSP) is a certification issued by the International Information Systems Security Certification Consortium, or (ISC)². The certification has been widely recognized in the IT (information technology) community for more than a decade as a "gold standard" for security professionals. CISSP certification is difficult to achieve and is held by thousands of practitioners in more than 35 countries. The certification has an experience requirement, and candidates are required to pass a rigorous exam that tests mastery of a common body of knowledge covering 10 fields, including access control, systems development, cryptography, ethics, and security practices.

For More Information

Visit the (ISC)² Web site at www.isc2.org for more information.

See Also: *Global Information Assurance Certification (GIAC), International Information Systems Security Certification Consortium (ISC)²*

CFB

Stands for cipher feedback, a feedback mechanism used for block ciphers with low data rates.

See: *cipher feedback (CFB)*

chaining mode

A feedback mode of operation for block ciphers.

Overview

Feedback, which involves directing some of the output of a process into its input, is used extensively in cryptography to create a greater appearance of randomness in encrypted information. Some block ciphers can operate in chaining mode, in which part of the output of one application of the cipher is combined with the next block of plaintext to be processed. This has a distinct advantage over simple ciphers that process blocks of plaintext independently of one other, for such ciphers generate identical ciphertext when the plaintext is the same.

The most common type of chaining used in block ciphers is called cipher block chaining (CBC), which uses a simple mathematical algorithm that has minimal performance penalty on the operation of the cipher.

See Also: *block cipher, cipher block chaining (CBC), ciphertext, plaintext*

Challenge Handshake Authentication Protocol (CHAP)

A challenge response authentication scheme used in remote access.

Overview

Challenge Handshake Authentication Protocol (CHAP) is defined in RFC 1994 and is one of several authentication schemes used by Point-to-Point Protocol (PPP) and its derivatives. CHAP is based on challenge response mechanism and authenticates users without the need of transmitting their passwords over the connection in any form, either clear or encrypted. Instead, CHAP uses the industry standard Message Digest 5 (MD5) algorithm to hash user passwords and transmits this one-way hash instead.

To prevent replay attacks, CHAP continues to send challenges at random intervals during a client session. CHAP is supported by most access servers, including Cisco routers and the Routing and Remote Access Service (RRAS) on Microsoft Windows platforms.

See Also: *challenge response authentication, hashing algorithm, MD5*

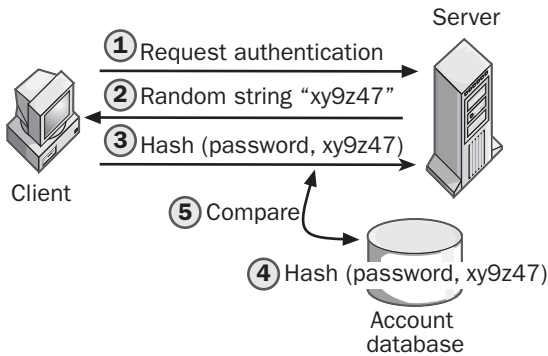
challenge response authentication

An authentication scheme in which passwords are not transmitted over the connection.

Overview

Challenge response authentication is a secure authentication scheme in which a client first contacts a server asking to be authenticated. The server responds by sending the client a randomly generated string of bytes called a **challenge**. The client hashes the challenge string with the user's password and sends the resulting response to the server. The server then performs the same hash using the challenge and the user's password, which it retrieves from its security accounts database. The server compares the response from the client with its own hash, and if the two are the same, the server authenticates the client and allows it access to the network.

Challenge response authentication forms the basis of LAN Manager (NTLM) authentication on the Microsoft Windows NT platform, which is still supported by the Microsoft Windows Server 2003 platform for backward compatibility but has largely been replaced by Kerberos authentication.



Challenge response authentication. How challenge response authentication works.

See Also: Kerberos

CHAP

Stands for Challenge Handshake Authentication Protocol, a challenge response authentication scheme used in remote access.

See: Challenge Handshake Authentication Protocol (CHAP)

Chernobyl

A notorious computer virus with a destructive payload.

Overview

The Chernobyl virus is a “space-filler virus” that fills up all available space on a computer’s hard drive. The virus also tries to overwrite flash basic input/output system (BIOS), which can render the system unbootable. Infection can result not only in data loss but also actual damage to BIOS chips and motherboards. Chernobyl was the first known virus that could physically damage a computer system.

Chernobyl first appeared in the wild in 1998 and wreaked havoc when its payload triggered on April 16, 1999, the 13th anniversary of the Chernobyl nuclear reactor incident in the former Soviet Union. The virus targeted systems running Microsoft Windows 95 and Microsoft Windows 98 platforms and affected

hundreds of thousands of systems, mainly in Asia and the Middle East.

The virus is sometimes called CIH for the initials of the developer, Chen Ing-hau, a computer engineering student in Taiwan. Chen wrote the program while on a tour of military duty, was arrested and released, and was later hired by a Taiwan technology firm.

New strains of CIH have continued to appear since the original virus was released. CIH can attach itself to other programs, including viruses, and infections have been seen in the Klez worm.

See Also: Klez, virus, worm

chief security officer (CSO)

Individual responsible for the security of a company’s network and communications systems.

Overview

With increasing concern over the security of information systems and resources, a new locus on the organizational sheet of large corporations has appeared: the chief security officer (CSO). Typical responsibilities for a CSO can include the following:

- Developing security policies and practices for authentication and access control and ensuring they are followed
- Procuring hardware and software necessary to ensure the security of network and communications systems and resources
- Training and educating users in security awareness and best practices
- Secure management of information assets for fixed and mobile users

Depending on the size of the company, the CSO may report to the chief information officer (CIO) or even the chief executive officer (CEO), and one or more of these roles may be combined in smaller companies. Many CSOs learn their skills working in the military or law enforcement environment in which security procedures are carefully outlined and rigorously followed. Other

names commonly used for this position include **corporate security officer**, **chief security architect**, **chief information security officer**, and **information security manager**.

See Also: *security policy*

chosen ciphertext attack

A cryptanalytic attack using chosen ciphertext to work with.

Overview

In a chosen ciphertext attack, the attacker decrypts chosen portions of ciphertext using the unknown key for the cryptosystem. By comparing the resulting plaintext with the chosen ciphertext using cryptanalytic methods, the attacker tries to determine the decryption key used by the system. This method can be effective against public key encryption systems for which one key is used to encrypt information and another to decrypt it.

There are two general approaches to performing a chosen ciphertext attack:

- **Batch method:** The attacker doesn't get to see any of the plaintext until after all chosen ciphertext has been decrypted.
- **Adaptive method:** The attacker gets to see plaintext generated from chosen ciphertext before choosing additional ciphertext to decrypt.

See Also: *cryptanalysis*

chosen plaintext attack

A cryptanalytic attack using chosen plaintext to work with.

Overview

In a chosen plaintext attack, the attacker encrypts chosen portions of plaintext using the unknown key for the cryptosystem. By comparing the resulting ciphertext with the chosen plaintext using cryptanalytic methods, the attacker tries to determine the encryption key used by the system. Since this method can determine only the encryption key, it is effective only against reversible

encryption systems that use the same key for encrypting and decrypting information.

There are two general approaches to performing a chosen plaintext attack:

- **Batch method:** The attacker doesn't get to see any of the ciphertext until after all chosen plaintext has been encrypted.
- **Adaptive method:** The attacker gets to see ciphertext generated from chosen plaintext before choosing additional plaintext to encrypt.

Notes

Public key systems do not use reversible encryption and are hence immune from this kind of attack.

See Also: *cryptanalysis*

chroot jail

A UNIX/Linux security measure for restricting file access.

Overview

A chroot jail is a security measure that prevents an application or daemon (service) from accessing files outside a specified directory tree. This limits the damage that can be done should the application or vulnerability be compromised by a malicious attacker.

For example, you could configure the Berkeley Internet Name Domain (BIND) daemon so that it runs "chrooted" to the directory `chroot/named`. The result is that BIND sees this directory as root ("`/`") and is thus unable to view or access anything outside the directory tree rooted at `/chroot/named`. Another common example is the File Transfer Protocol (FTP) daemon, where the FTP home directory appears as the machine's root directory to FTP users.

Notes

The UNIX `chroot` command is used to run commands using a specified root directory. Chroot can usually only be used by root, the UNIX superuser account.

See Also: *root*

CIAC

Stands for Computer Incident Advisory Capability, a branch of the U.S. Department of Energy that provides assistance when computer security incidents occur.

See: Computer Incident Advisory Capability (CIAC)

cipher

Another name for **encryption algorithm**, a mathematical procedure for converting plaintext into ciphertext.

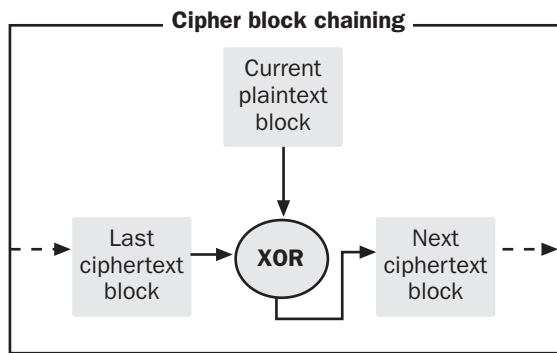
See: encryption algorithm

cipher block chaining (CBC)

A feedback mechanism commonly used in block ciphers.

Overview

Chaining refers to the process of combining previously generated ciphertext with new plaintext. Instead of encrypting each block of plaintext independently, a block of plaintext is first XORed with the most recently generated block of ciphertext, and then the block cipher is applied to the result. The first block of plaintext, having no antecedent block of ciphertext to use in this process, is instead XORed with a randomly generated “seed” called an **initialization vector**. The resulting ciphertext is more difficult to decrypt than if chaining were not employed since identical blocks of plaintext do not produce identical ciphertext.



Cipher block chaining (CBC). How cipher block chaining works.

Notes

CBC is the default cipher mode used by the base cryptographic provider of Microsoft CryptoAPI (CAPI).

See Also: block cipher, ciphertext, plaintext

cipher feedback (CFB)

A feedback mechanism used for block ciphers with low data rates.

Overview

Feedback is a mechanism used to prevent block ciphers from transforming identical blocks of plaintext into identical ciphertext. The most common feedback mechanism used is chaining mode, which combines whole blocks of plaintext and ciphertext together. Another approach sometimes used is cipher feedback, in which small increments of plaintext are transformed into cipher instead of processing entire blocks at a time.

For standard 64-bit block ciphers, the block typically is divided into eight sections of 8 bits, each using a shift register. Then for each encryption cycle, the shift register is first filled with the initialization vector, a random “seed” used to start the encryption process. The entire block is then encrypted, the leftmost 8 bits are combined with the first 8 bits of plaintext, and the result is 8 bits of ciphertext. The shift register then shifts 8 bits to the left, the 8 bits of cipher previously generated are moved to the rightmost 8 bits of the register, and the process repeats.

Cipher feedback is typically employed in situations in which the incoming data rate of plaintext is slow, for example, when data is originating from a keyboard.

See Also: cipher block chaining (CBC), cipher mode, ciphertext, plaintext

cipher mode

A mode of operation for a block cipher.

Overview

Block ciphers are encryption algorithms that encrypt plaintext in discrete chunks called blocks. This can be done in two ways:

- Each block of plaintext can be processed independently of others. This is the fastest method but suffers from the weakness that identical blocks of plaintext are transformed into identical blocks of ciphertext. This process is usually called Electronic Codebook (ECB), in recognition that this method bears resemblance to the codebooks used in wartime prior to the advent of electronic computers.
- Plaintext can be combined with ciphertext to further scramble the result, so that identical blocks of plaintext no longer produce identical ciphertext. One popular method for doing this is called cipher block chaining (CBC), which combines each block of plaintext with the previous block of ciphertext before encrypting the result. Another approach is cipher feedback (CFB), which combines smaller portions of plaintext with ciphertext before encrypting the result.

See Also: *block cipher, cipher block chaining (CBC), cipher feedback (CFB), ciphertext, Electronic Codebook (ECB), plaintext*

ciphertext

Information that has been encrypted.

Overview

Encryption is the process of transforming plaintext into ciphertext. Plaintext is information that is in human-readable form, for example, an e-mail message typed in a text editor. To prevent sensitive information from being read should it be intercepted by someone other than its intended recipient, the message can be encrypted using a mathematical procedure called an encryption algorithm. The result of applying this algorithm to the information is ciphertext, a string of bits that still contains the original information but which cannot be read by anyone unless it is first decrypted to convert it back into plaintext.

See Also: *encryption, encryption algorithm, plaintext*

ciphertext-only attack

A cryptanalytic attack using only ciphertext to work with.

Overview

In a ciphertext-only attack, the attacker has only a sample of ciphertext to work with. Nothing is known about the plaintext from which this sample has been generated, which makes it exceedingly difficult to crack the encryption system used. In general, ciphertext-only attacks can be successful only when a very large sample of ciphertext can be obtained in order to perform statistical analysis on it in conjunction with guessing certain properties of the original plaintext.

Another name for this attack is the **recognizable plaintext attack**.

See Also: *cryptanalysis*

CIS

Stands for Center for Internet Security, a nonprofit organization that helps organizations manage risk associated with information systems security.

See: *Center for Internet Security (CIS)*

CISA

Stands for Certificate Information Systems Auditor, a widely accepted certification in auditing, control, and security of information systems.

See: *Certificate Information Systems Auditor (CISA)*

CISSP

Stands for Certified Information Systems Security Professional, a widely accepted certification for computer security professionals.

See: *Certified Information Systems Security Professional (CISSP)*

cleartext

Another name for **plaintext**, information that is easily readable by human beings.

See: *plaintext*

clogging attack

A denial of service (DoS) attack against a public key cryptography system.

Overview

In a clogging attack, the attacker sends copies of public keys to a target user from spoofed source addresses of legitimate users. The target user quickly becomes overloaded with verifying these keys, and the result can be denial of service to legitimate users trying to communicate with the target. Certain encryption algorithms such as Diffie-Hellman (DH) are susceptible to clogging attacks. Oakley key exchange protocol is based on DH but has enhancements that prevent clogging. Simple Key-Management for Internet Protocols (SKIP), a protocol developed by Sun Microsystems for key management on IP networks, is also designed to be resistant to clogging attacks.

See Also: denial of service (DoS), Diffie-Hellman (DH), public key cryptography, spoofing

code access permissions

Permissions used in the Microsoft .NET Framework to protect resources accessed by code from unauthorized use.

Overview

Code access permissions are built into the common language runtime of .NET Framework and are used to enforce security restrictions on managed code and to implement code access security on the platform. Code access permissions make it easy to write secure code by providing built-in mechanisms for controlling access to protected operating system resources and operations. The Microsoft .NET Framework defines a number of built-in code access permissions that can be used to control access to directory services, Domain Name System (DNS), environment variables, event logs, file systems, message queues, performance counters, printers, the registry, services, and other resources. In addition, developers can also define their own custom permis-

sions when built-in permissions are insufficient for controlling access to a resource.

See Also: code access security (CAS), permissions

code access security (CAS)

A code security mechanism built into the Microsoft .NET Framework.

Overview

Code access security (CAS) is designed to protect computers from malicious code, for example, in software downloaded from the Internet. CAS is also designed to allow code from an untrusted origin to run safely and prevents trusted code from accidentally or intentionally compromising the security of a system. CAS defines different levels of trust that depend on where the code originates and employs mechanisms that enforce these trust levels. By employing CAS, the likelihood of running damaging code is reduced and the level of damage that can potentially result is lessened.

See Also: code access permissions

CodeRed

A worm that caused Web servers and routers to crash across the Internet.

Overview

The original CodeRed worm exploited a vulnerability in Internet Information Services (IIS) Web servers that resulted in Web sites being defaced with the message "Hacked by Chinese!" The worm first appeared on July 12, 2001, and caused little damage, but a variant appeared one week later and infected 360,000 machines worldwide in only 14 hours. In addition to Web site defacement and denial of service (DoS) due to degradation of Web server performance, this variant also caused routers, switches, and printers to crash when infected IIS machines tried to send them copies of the worm. Both the original version and its variant were memory-resident and could be removed by simply rebooting the server, though this did not prevent reinfection.

A totally new worm called CodeRedII, which exploited a buffer overflow vulnerability in IIS, appeared soon after on August 4, 2001. This new worm had almost no code in common with the original CodeRed worm, but since it included the string “CodeRedII” in its source code, it was named accordingly. CodeRedII was more dangerous since it also installed a backdoor into infected systems, allowing them to be used as platforms for launching distributed denial of service (DDoS) attacks against target networks.

Microsoft quickly issued patches to prevent both worms from infecting IIS machines.

See Also: *backdoor, buffer overrun, denial of service (DoS), distributed denial of service (DDoS), worm*

code signing

Signing code with digital certificates to validate its authenticity and integrity.

Overview

An important question for users who download software over the Internet is whether they can trust the software’s integrity and authenticity. In this context, **integrity** means the software has not been tampered with since it was developed, while **authenticity** guarantees that the software originates from where it says it does. One way of dealing with these issues is to sign code using digital certificates issued by a trusted authority, either an independent third-party certification authority (CA) or the software vendor itself. Microsoft Authenticode is a popular example of a code-signing mechanism to ensure for users that software they download from the Internet is authentic and has not been tampered with. Note that code signing does not necessarily signify that such software is safe for users to install, because software may be authentic and have integrity and yet be buggy.

Marketplace

Microsoft has developed criteria for software vendors to apply for and obtain a software-publishing certificate from a trusted authority, which they can then use to sign code they develop. For commercial certificates, the vendor must provide proof of identity, a pledge that the software does not contain known viruses or other code

that may harm a user’s computer, and a Dun & Bradstreet rating signifying the financial stability of the company. The criteria for individual certification are similar except that a Dun & Bradstreet number (DUNS Number) is not required.

Verisign provides services for issuing code-signing digital IDs for several software vendors, including Microsoft, Sun, Netscape, Macromedia, and others.

See Also: *Authenticode, certificate authority (CA)*

Common Criteria & Methodology for Information Technology Security Evaluation

Usually called Common Criteria, an international effort to standardize criteria for evaluating the security of information systems.

Overview

The Common Criteria is the outcome of a series of efforts by several nations that began in the early 1980s with the Trusted Computer Systems Evaluation Criteria (TCSEC), or *Orange Book*, developed by the U.S. Department of Defense. This effort combined with the European Information Technology Security Evaluation Criteria (ITSEC) in the early 1990s to create the Common Criteria Project, which issued version 1 of the Common Criteria in 1996. A revised version of these criteria evolved into the ISO 15408 standard in 1999, with which the current version 2.1 of these criteria closely aligns.

The Common Criteria provides a common language for defining the security requirements and describing the security capabilities of products. It also includes a series of evaluation assurance levels (EALs), an international program for accrediting laboratories for the testing and evaluation of security products.

For More Information

Visit www.commoncriteria.org for more information.

See Also: *Information Technology Security Evaluation Criteria (ITSEC), Trusted Computer Systems Evaluation Criteria (TCSEC)*

Common Cryptographic Architecture (CCA)

A cryptographic architecture developed by IBM for its computing platforms.

Overview

Common Cryptographic Architecture (CCA) defines a set of application programming interfaces (APIs) for providing cryptographic services to applications. These APIs include functions for confidentiality, data integrity, and message authentication. The architecture is based on the Data Encryption Standard (DES) and has found widespread use in the banking and financial industry in the IBM 4758, a tamper-resistant Peripheral Component Interconnect (PCI) card that plugs into PCs to provide cryptographic functions for secure communications. The IBM is encased in a hardened metal case and contains temperature, shock, and X-ray sensors to guard against tampering, and it is Federal Information Processing Standards (FIPS) 140-1 Level 4 certified.

Despite the hardened nature of this cryptographic device and the fact that it uses strong Triple DES (3DES) encryption, in 2002, an attack was devised by a team of researchers at Cambridge University's William Gates Computer Laboratory; using off-the-shelf hardware costing less than \$1,000, the team took less than a day to discover an encryption key used by CCA.

See Also: 3DES, cryptography, Data Encryption Standard (DES), encryption algorithm

Common Vulnerabilities and Exposures (CVE)

An emerging industry standard for naming vulnerabilities and other information security exposures.

Overview

Common Vulnerabilities and Exposures (CVE) is maintained by MITRE Corporation in collaboration with security experts, academic institutions, government agencies, and security tool vendors. CVE was developed to standardize the naming of security vulnerabilities so that information could be shared between different security databases and tools. CVE functions as

a kind of dictionary of all publicly known vulnerabilities and exposures for operating systems and applications. The National Institute of Standards and Technology (NIST) has recognized the importance of the CVE as an emerging industry standard.

For More Information

Visit MITRE at cve.mitre.org for more information.

See Also: vulnerability

compromised system

A computer system with unknown integrity because an attacker has gained illicit access.

Overview

The goal of a malicious individual attacking a computer system is to compromise the system. To **compromise** a system means to penetrate the security defenses of the system and gain access to some level of control over its processes and information. There are different levels at which a system can be compromised, ranging from relatively benign, such as Web site defacement, to extremely dangerous, such as gaining root access. Once a system has been compromised, the attacker is said to have achieved an **exploit**. This may then be the launching ground for further exploits, for example, as in distributed denial of service (DDoS) attacks in which compromised systems called **zombies** are used to launch attacks against other systems.

The CERT Coordination Center (CERT/CC) offers recommendations on procedures to follow in the event of a system being compromised. Recommended steps include these:

- Consultation with management, legal counsel, and law enforcement agencies
- Disconnecting the system from your network
- Imaging the system for analysis of the intrusion
- Searching for modifications in system, configuration, and data files
- Examining other systems on your network for evidence of compromise

- Reporting the incident to an incident response center
- Recovering your system using a clean install and hardening it against similar intrusions in the future

For More Information

Visit CERT/CC online at www.cert.org for more information.

See Also: *CERT Coordination Center (CERT/CC), exploit, intrusion, vulnerability*

computer forensics

Obtaining evidence of criminal activity from information systems.

Overview

Computer forensics involves the application of both computer technology and legal expertise to obtain from computer hardware and software evidence of intrusion, misuse, theft, or other criminal activities. Computer forensics is thus a branch of the more general subject of **forensics**, the application of science and technology to criminal investigation. With the rapid growth of the Internet and e-commerce, a corresponding growth in computer crime has occurred, and law enforcement agencies have had to apply high-tech approaches to tracking down and arresting those who commit such crimes.

Computer forensics is more than just recovering data erased from hard drives. It is a methodical process of extracting, identifying, documenting, and preserving digital information in forms that satisfy the needs of law enforcement officials, prosecutors, courts, insurance companies, and civil litigators. When performing a forensic investigation of computer media, certain requirements must be met; in particular, the integrity of the original media must not be affected. Best practice dictates that such examinations should never be performed on the original media, but on bit-image copies instead to lessen the danger of accidentally damaging evidence on the original media. Thorough documentation of data recovery procedures and careful storage of original media are also prerequisites for recovered evidence to stand up in court because of “chain of custody,” the legal requirement that evidence submitted in

court be accompanied by documentation of who had physical custody of that evidence and under what security conditions it was held by the parties holding it. Computer forensics experts must also be able to present evidence in court in a way that makes complex technology understandable to judges and jurors who may be laypersons in such technologies.

See Also: *cybercrime*

Computer Incident Advisory Capability (CIAC)

A branch of the U.S. Department of Energy that provides assistance when computer security incidents occur.

Overview

Computer Incident Advisory Capability (CIAC) was founded in 1989 shortly after an incident called the “Internet worm” brought down large portions of the Internet, an event that also prompted the formation of the CERT Coordination Center (CERT/CC) and other incident response bodies such as the Forum of Incident Response and Security Teams (FIRST). CIAC serves its constituents in the U.S. Department of Energy by providing technical assistance when requested in the event of a computer security incident. Such assistance can take the form of awareness training, threat evaluation, and the collection and analysis of data relating to vulnerabilities and exposures.

CIAC publishes bulletins regarding security vulnerabilities of different operating systems and articles on various aspects of securing network and computing resources. CIAC also maintains a famous database of Internet hoaxes and chain letters that has been maintained since 1995, and it provides advice on how to recognize hoaxes and distinguish them from genuine security threats.

For More Information

Visit CIAC online at www.ciac.org for more information.

See Also: *CERT Coordination Center (CERT/CC), Forum of Incident Response and Security Teams (FIRST)*

Computer Security Division (CSD)

A division of the National Institute of Standards and Technology (NIST) that focuses on information systems security.

Overview

Computer Security Division (CSD) is one of eight divisions within the Information Technology Laboratory at NIST, and its goal is to improve the security of information systems through several means:

- Researching vulnerabilities and devising techniques and procedures for overcoming them
- Developing standards and metrics for testing and validating security systems and products
- Providing guidance on how to plan and implement secure information systems
- Raising public awareness in the IT (information technology) community regarding risks, vulnerabilities, and dangers in the security area

The research and development areas CSD focuses on include the areas of cryptography, testing and evaluation, management, awareness training, and emerging technologies. CSD also maintains a Computer Security Resource Center (CSRC) from which it issues bulletins, reports news, and publishes information about upcoming workshops and events.

For More Information

Visit CSD online at csrc.nist.gov for more information.

See Also: *National Institute of Standards and Technology (NIST)*

computer security incident response team (CSIRT)

A term used by the CERT Coordination Center (CERT/CC) to describe a service organization that responds to computer security incidents.

Overview

CERT/CC is a leading center of Internet security expertise operated by Carnegie Mellon University, and one of

the services it provides is guidance on how to organize and run a computer security incident response team (CSIRT). These teams can range from larger organizations serving entire countries or regions such as AusCERT for the Asia-Pacific area or JPCERT/CC for Japan, to smaller organizations serving commercial enterprises or educational institutions, to corporate groups providing fee-based services on request. The job of CSIRTs, whether they are ad hoc or formalized, is to receive, review, and respond to reports of computer security incidents on behalf of their constituency. Such incidents may include threats, tampering, mischief, breaches, denial of service (DoS), or unauthorized use of computer hardware, software, services, and data.

Other common names for such teams include **incident response team**, **incident response center**, and **emergency response team**.

For More Information

Visit CERT/CC at www.cert.org for more information.

See Also: *CERT Coordination Center (CERT/CC)*, *Computer Incident Advisory Capability (CIAC)*, *Forum of Incident Response and Security Teams (FIRST)*

Computer Security Institute (CSI)

A membership organization dedicated to training information security professionals.

Overview

Computer Security Institute (CSI) is a San Francisco-based organization with thousands of members worldwide that provides information, programs, and training for information security practitioners in business, industry, and government. CSI sponsors conferences and exhibitions that include seminars on security awareness, intrusion management, data encryption, virtual private networking, and other topics. Membership benefits include the *ALERT* newsletter, a quarterly journal, and a *Buyer's Guide* of current products in the computer security field.

CSI also publishes an annual survey on computer crime and security, developed with the participation of the FBI's Computer Intrusion Squad. The purpose of this

survey is to help raise general awareness concerning computer crime and security issues in business and government.

For More Information

Visit CSI at www.gocsi.com for more information.

See Also: *cybercrime*

confidentiality

A security concept that implies safety from interception, viewing, or copying.

Overview

Confidentiality is an important element in the secure transmission of electronic information. On both wired and wireless networks, an attacker might try to eavesdrop to capture passwords or sensitive business information such as credit card numbers. To prevent eavesdropping from being effective, communications can be encrypted using Data Encryption Standard (DES) or Advanced Encryption Standard (AES) to ensure confidentiality. In addition, communications can be digitally signed to ensure their integrity—that is, to ensure that the information has not been tampered with during transit.

See Also: *authentication, encryption, integrity*

confidentiality agreement

An agreement between two parties to ensure the confidentiality of business information that they exchange.

Overview

Confidentiality agreements are common in many areas of business, including employer/employee contracts and supply chain agreements between business partners. A typical confidentiality agreement includes clauses regarding the following:

- The definition of the types of information considered confidential under the agreement

- Nondisclosure and nonuse obligations outlining the fact of nondisclosure of information and the parties to whom it should not be disclosed
- Exclusions outlining the parties with whom and conditions under which such information can be shared, typically requiring authorization in writing
- Ownership clause specifying who retains the rights for different types of confidential information
- Disclosure clause regarding the communication of third-party information
- Return clause outlining the handling of confidential information once the agreement has been terminated
- Term of the agreement and effective period for its application, usually called a nondisclosure agreement (NDA)

See Also: *confidentiality*

consensus baseline security settings

A set of guidelines for securing computers running Microsoft Windows 2000 Professional.

Overview

The consensus baseline security settings are an advanced set of recommendations developed by the Center for Internet Security (CIS), a nonprofit organization that helps organizations manage risk associated with information systems security, in conjunction with the President's Critical Infrastructure Protection Board, the National Security Agency (NSA), the General Services Administration, the National Institute of Standards and Technology (NIST), the Defense Information Systems Agency (DISA), and the SANS Institute.

The consensus baseline security settings are a level-2 baseline designed to provide system administrators with step-by-step procedures for ensuring that desktop computers running Windows 2000 Professional are

properly configured to protect them against attack. The consensus baseline security settings have been endorsed by a broad spectrum of industry and government agencies, including NIST, the General Services Administration (GSA), the NSA, and the SANS Institute. At present, these settings are viewed as recommendations by these agencies, and not as standards.

For More Information

You can obtain the consensus baseline security settings from CIS at www.cisecurity.org.

See Also: Center for Internet Security (CIS), National Institute of Standards and Technology (NIST), National Security Agency (NSA), SANS Institute

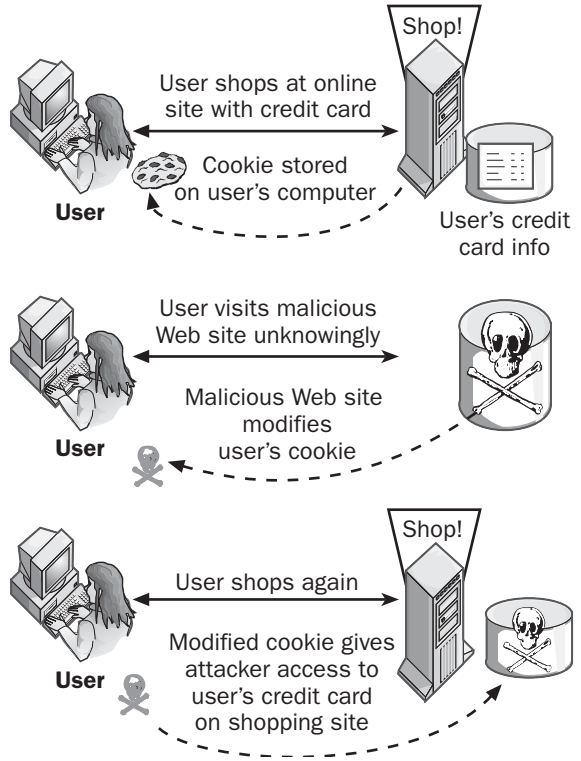
cookie poisoning

An attack involving modification of cookies on client computers.

Overview

Cookies are small files created on client computers when these systems browse certain Web sites. These cookies can contain information about the user’s shopping habits on e-commerce sites, personal information such as passwords or birth dates, or anything else the application running on the Web server chooses to implement. Cookie poisoning involves an attacker modifying a cookie on a client in order to impersonate the user, a form of identity theft. Using the modified cookie, the attacker can visit Web sites previously visited by the user and try to access personal information for the user stored on the site, such as the user’s credit card number.

The best protection against such an attack is for Web sites that use cookies to encrypt them so that attackers can’t read or edit information stored in them. Other names for this attack include **cookie hijacking** and **cookie snarfing**.



Cookie poisoning. How cookie poisoning works.

See Also: identity theft

covert channel

A communications channel that hides illicit information flow within a normal communications stream.

Overview

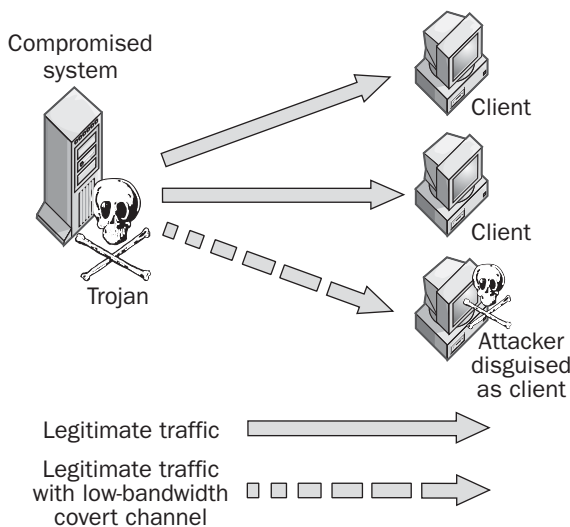
In computer networking, covert channels are methods for secretly sending information by hiding it in portions of packets not normally used for such purposes. Examples include hiding information in the identification field of Internet Protocol (IP) packets, the initial sequence number field of Transmission Control Protocol (TCP) packets, or the “bounce” or acknowledgment sequence number field in TCP packets. Covert transmission using



C

these methods can often pass through firewalls and stateless intrusion detection systems (IDSs) without being flagged or generating an alert, especially if included in traffic directed to ports normally open, such as TCP port 80 for Hypertext Transfer Protocol (HTTP) traffic. To the firewall or IDS, such packets appear to be innocuous HTTP packets, whereas in fact sensitive business information may be being transmitted by industrial espionage agents or disgruntled employees. Only by using special tools that can identify unusual network traffic patterns such as unsolicited SYN/ACK packets can such communications be detected.

Covert channels are exceedingly difficult to detect and protect against. They are often used by Trojans for clandestinely controlling a compromised system from a remote location.



Covert channel. How a covert channel works.

See Also: firewall, intrusion detection system (IDS)

cracking

Illegally modifying commercial software, circumventing authentication procedures, or deciphering encrypted communications.

Overview

Cracking software generally involves circumventing licensing and usage restrictions on commercial software by illegal methods. These methods can include modifying code directly through disassembling and bit editing, sharing stolen product keys, developing software to generate activation keys, and so on. Cracking is essentially a form of software piracy and is punishable under state and federal law.

The term **cracking** is also used to describe the act of breaking into a system or network by thwarting authentication procedures. A common example is password cracking, which involves guessing passwords to try to gain access to sensitive data such as credit card information stored in databases.

The word can also be used to describe attempts to guess session keys used for encrypting communications between two parties. With the proliferation of wireless networks, concern about privacy of wireless communications has become a significant issue for many businesses.

While hacking is an activity that has a long and venerable history in the computer world and is basically motivated by curiosity mixed with a fair degree of pride of accomplishment, cracking is essentially a criminal activity whose aim is theft or destruction of information or property.

Notes

A **cracker** is an individual who tries to crack software keys or network passwords, usually with malicious intent. Crackers are sometimes called **black hats** to distinguish them from **white hats**, or hackers with legitimate connection with the security community.

A **crack** can mean a stolen product key, guessed password, procedure for breaking into a network or application, or a tool to achieve such ends.

See Also: black hat, hacker, password cracking

CRC

Stands for cyclical redundancy check, a mathematical technique for ensuring the integrity of data.

See: *cyclical redundancy check (CRC)*

credentials

Information used to authenticate users on a system or network.

Overview

Credentials are pieces of information that users submit in order to gain access to resources on that network. The most common form of credential is the user account, which typically consists of three things:

- A user name or user account
- A password
- The domain or realm defining the network to which the user's account belongs

High-security environments may require more information to identify the user before granting access. An example is a digital certificate identifying the user and stored in a device such as a smart card or token. Biometrics can also be used for authenticating users, in which fingerprints, retinal scans, or some other physical characteristic scanning method is used to authenticate the user.

See Also: *biometric identification, digital certificate, password, smart card*

CRL

Stands for certificate revocation list, a list of revoked certificates maintained by a certificate authority (CA).

See: *certificate revocation list (CRL)*

cross-realm authentication

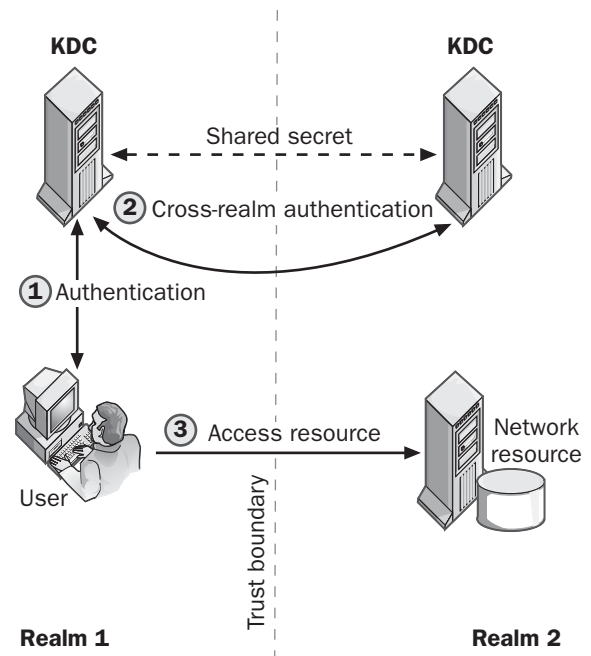
Authentication of a Kerberos principal in one realm by principals in another realm.

Overview

Kerberos is a security protocol for authenticating users and applications on a distributed network. The basic

unit of authentication in Kerberos is the **realm**, which is a network served by a single group of key distribution center (KDC) servers sharing a common authentication database. In general, a KDC can only authenticate users from its own realm. Using cross-realm authentication, however, KDCs in different realms establish trust through a shared secret called a cross-realm secret. This secret is used to prove the identity of a principal when it crosses the boundary between two realms.

Kerberos version 5 supports an enhanced form of cross-realm authentication called transitive cross-realm authentication. Using this method, a chain of realms can be established to allow principals to hop from one realm to another to be authenticated in the target realm.



Cross-realm authentication. How Kerberos cross-realm authentication works.

See Also: *Kerberos*

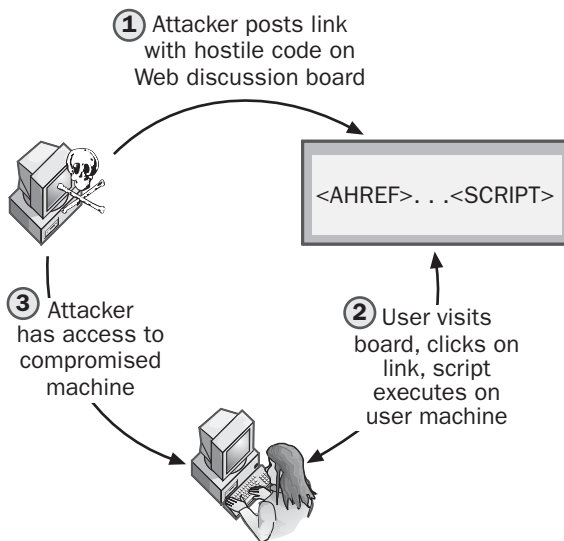
cross-site scripting (CSS)

A Web server vulnerability resulting from poor input validation.

Overview

Cross-site scripting (CSS) was first identified in early 2000 when the CERT Coordination Center (CERT/CC) issued an advisory warning against attacks against Web servers running Internet Information Services (IIS). In a typical cross-site scripting scenario, a malicious user typically posts a message to a Web discussion forum. This message contains a link to a Web site created by the attacker. When a user clicks on the link, hidden script in the link tag causes the response of the linked site to execute on the client with undesirable effect. By properly designing appropriate script code, the attacker can gain complete control over the client system and execute arbitrary code on it to damaging effect.

By disabling scripting in their Web browsers, users can prevent such attacks on their systems at the cost of reduced functionality. The alternative is for sites hosting public discussion boards to use proper input validation to ensure messages that are posted have no hidden script that can be used to launch such attacks.



Cross-site scripting. How cross-site scripting works.

See Also: vulnerability

cryptanalysis

The science of finding methods for breaking cryptosystems.

Overview

While cryptography is concerned with discovering new methods for encrypting information to ensure its privacy, cryptanalysis is concerned with the opposite question of how to crack encryption schemes. In theory, any encryption scheme can be cracked by using brute force to guess repeatedly what the decryption key might be. But with modern encryption algorithms, brute force is not good enough, since it could take the fastest computer in the world longer than the age of the universe to crack a single message encrypted with such algorithms. The job of the cryptanalyst is to devise more ingenious methods than brute force to exploit weaknesses in known algorithms or devise ways of gaining information about unknown ones in order to crack cryptosystems and decipher encrypted messages. Practitioners in the field of cryptanalysis are known as cryptanalysts.

Some examples of cryptanalytic attacks used to try to compromise cryptosystems include the chosen plaintext attack, chosen ciphertext attack, ciphertext-only attack, and known plaintext attack; see the related articles elsewhere in the book for more information about these types of attacks.

See Also: brute-force attack, chosen ciphertext attack, chosen plaintext attack, ciphertext-only attack, cryptography, encryption algorithm, known plaintext attack

CryptoAPI (CAPI)

A set of application programming interfaces (APIs) for cryptography built into Microsoft Windows-based platforms.

Overview

Microsoft CryptoAPI (CAPI) is a layer of cryptographic services that can be used by applications running on Windows-based systems. CryptoAPI provides five major functions:

- Base cryptographic functions, including context functions for connecting to a cryptographic service

provider (CPS), key generation functions for creating and storing cryptographic keys, and key exchange functions for transmitting and exchanging keys

- Certificate encoding and decoding functions used to encrypt, decrypt, and hash data
- Certificate store functions for storing and managing collections of digital certificates
- Simplified message functions for encrypting and decrypting messages and for signing and verifying digital signatures
- Low-level message functions for providing more granular control over encryption, decryption, signing, and verification of messages and their signatures

See Also: *cryptographic service provider (CSP), encryption*

cryptographic hash function

More commonly called **hash function**, a mathematical function that generates a fixed-size result from arbitrary amounts of data.

See: *hashing algorithm*

cryptographic service provider (CSP)

A provider of cryptographic functions to Microsoft CryptoAPI (CAPI).

Overview

CryptoAPI (CAPI) is the component of Microsoft Windows-based platforms that provides cryptographic services to applications. These services include the encryption and decryption of data, creation and verification of digital signatures, and generation and exchange of cryptographic keys. CryptoAPI acts as a wrapping layer around cryptographic service providers, which export functions called by CryptoAPI programming interfaces.

CryptoAPI (CAPI) comes with a basic set of cryptographic service providers that include a base provider that uses 512-bit RSA encryption, an enhanced provider that uses 1024-bit RSA, a strong provider, Digital Signature Standard (DSS) provider, and several others.

See Also: *CryptoAPI (CAPI), encryption, RSA*

cryptography

The science of discovering new methods for encrypting information.

Overview

Cryptography is a branch of mathematics and is concerned with discovering ways to ensure the privacy of communications between parties. Cryptography has a long history dating back to ancient times. One of the earliest examples was in 500 B.C., when Hebrew scribes used a reversed alphabet substitution cipher for writing down portions of the Book of Jeremiah. Cryptography has frequently been associated with military use, and Julius Caesar used ciphers involving shifting letters a fixed amount across the alphabet to obscure military communications during his campaign in Gaul. During the Second World War, the cracking of the Enigma machine's code by Polish mathematician Marian Rejewski was a turning point for the Allies in the war against Nazi Germany.

Modern cryptographic systems originated with a project called Lucifer that was developed by IBM in 1976. These modern systems or encryption algorithms are step-by-step procedures that use complex mathematics for transforming ordinary information called plaintext into ciphertext, which has the same information content but is no longer human readable. Cryptography is concerned with the theoretical basis of such systems and seeks not only to devise such systems but also to prove the degree to which they are difficult to crack. Modern cryptosystems are based on intrinsically difficult mathematical problems such as factoring large prime numbers and the complexity of elliptical functions. Practitioners in the field of cryptography are referred to as cryptographers.

Notes

The word **cryptography** comes from the Greek word **krypt**, meaning hidden or secret.

See Also: *ciphertext, cryptanalysis, encryption, encryption algorithm, plaintext*

cryptology

The science that combines cryptography and cryptanalysis.

See Also: *cryptanalysis, cryptography*

cryptosystem

A mathematical procedure for converting plaintext into ciphertext.

Overview

In general, modern cryptosystems can be broken down into two types of procedures:

- **Hashing functions:** These are one-way (nonreversible) procedures for generating ciphertext from plaintext and are used in challenge response authentication schemes and other areas.
- **Encryption algorithms:** These are reversible procedures that allow plaintext to be converted to ciphertext and then converted back again.

Encryption algorithms themselves can be classified as either of the following:

- **Symmetric key algorithm:** Uses a shared single key called a secret key to encrypt and decrypt data
- **Asymmetric key algorithm:** Uses two keys, a public key to encrypt data and a private key to decrypt it

See Also: *asymmetric key algorithm, challenge response authentication, ciphertext, encryption algorithm, hashing algorithm, plaintext, symmetric key algorithm*

CSD

Stands for Computer Security Division, a division of the National Institute of Standards and Technology (NIST) that focuses on information systems security.

See: *Computer Security Division (CSD)*

CSI

Stands for Computer Security Institute, a membership organization dedicated to training information security professionals.

See: *Computer Security Institute (CSI)*

CSIRT

Stands for computer security incident response team, a term used by the CERT Coordination Center (CERT/CC) to describe a service organization that responds to computer security incidents.

See: *computer security incident response team (CSIRT)*

CSO

Stands for chief security officer, the individual responsible in a company for the security of its network and communications systems.

See: *chief security officer (CSO)*

CSP

Stands for cryptographic service provider, a provider of cryptographic functions to Microsoft CryptoAPI (CAPI).

See: *cryptographic service provider (CSP)*

CSS

Stands for cross-site scripting, a Web server vulnerability resulting from poor input validation.

See: *cross-site scripting (CSS)*

CTL

Stands for certificate trust list, a group of items signed by a trusted certificate authority (CA).

See: *certificate trust list (CTL)*

Cult of the Dead Cow (cDc)

A notorious group of underground hackers.

Overview

Cult of the Dead Cow (cDc) was founded in 1985 and is one of the oldest communities of “black hat” hackers still active. Membership in the group consists of a few dozen influential people who have gained high profile in the hacking and security communities for their exploits and the techniques and tools they have developed. Sir Dystic, one member of cDc, created the infamous Back Orifice tool in 1998 that allows users to obtain remote control over machines running Microsoft Windows 95 and Windows 98 platforms. Another member named DilDog carried this further and in 1999 released at Defcon in Las Vegas Back Orifice 2000 (BO2K), a version of Back Orifice for Microsoft Windows NT.

For More Information

Visit cDc online at www.cultdeadcow.com for more information.

See Also: *Back Orifice, Back Orifice 2000 (BO2K), black hat, Defcon, hacker*

CVE

Stands for Common Vulnerabilities and Exposures, an emerging industry standard for naming vulnerabilities and other information security exposures.

See: *Common Vulnerabilities and Exposures (CVE)*

cybercrime

Criminal activities that take place in cyberspace (the Internet).

Overview

Cybercrime is a growing concern for both law enforcement officials and consumers as a result of the rapid expansion of the Internet into all forms of business and commerce. Like other forms of criminal activity, cybercrime can be directed toward persons, property, companies, or government authorities and can take many forms, including viruses, worms, Trojans, hoaxes, mail bombs, threats, harassment, stalking, fraud, theft, forgery, piracy, break-ins, child pornography, espionage, and terrorism. All aspects of the Internet are vulnerable to such activities, including the World Wide Web, e-mail, chat rooms, and newsgroups.

A survey in 2001 by the Computer Security Institute (CSI) in conjunction with the FBI revealed that the most common types of cybercrime experienced by companies were virus infection, insider abuse of network resources, and unauthorized access by insiders. Less common were system penetration, denial of service (DoS), theft of proprietary information, sabotage, fraud, and eavesdropping.

For More Information

Visit www.cybercrime.com, the Computer Crime and Intellectual Property Section (CCIPS) of the Criminal Division of the U.S. Department of Justice.

See Also: *computer forensics, Computer Security Institute (CSI), hoax, software piracy, Trojan, virus, worm*

cyclical redundancy check (CRC)

A mathematical technique for ensuring the integrity of data.

Overview

C

Data stored on a hard drive or transmitted over a network is subject to corruption from various sources such as noise or hardware errors. To ensure the integrity of such data during storage or transmission, a cyclical redundancy check (CRC) can be performed, which calculates a small numerical quantity called a checksum based on the totality of bits in the file or packet being transmitted.

In a network transmission using Ethernet, for example, a checksum is calculated for each frame and is appended to the frame. The recipient of the frame recalculates the checksum based on the binary value of the

frame received, and then compares this with the checksum appended to the frame. If the two values disagree, the frame has been modified in transit and must be re-sent.

CRC is designed to ensure the integrity of data only against random degradation caused by noise or other sources. It does not guarantee integrity against modification with malicious intent, since an attacker who modifies the contents of a frame could easily recalculate the checksum and replace the appended value to fool the recipient.

See Also: *integrity*

DAC

Stands for discretionary access control, a mechanism for controlling access by users to computing resources.

See: discretionary access control (DAC)

DACL

Stands for discretionary access control list, the most common type of access control list (ACL) used to control access to computer and network resources.

See: discretionary access control list (DACL)

Data Encryption Algorithm (DEA)

The name used by the American National Standards Institute (ANSI) for the Data Encryption Standard (DES).

See: Data Encryption Standard (DES)

Data Encryption Standard (DES)

An encryption standard used for many years by the U.S. federal government.

Overview

The Data Encryption Standard (DES) has been used since 1977 by federal agencies for protecting the confidentiality and integrity of sensitive information both during transmission and when in storage. DES is a secret key encryption algorithm defined by Federal Information Processing Standard FIPS 46-9. A stronger form of DES called 3DES or TDES (Triple DES) is also sometimes used by government agencies, but requires additional processing power because of the extra computation involved.

DES was cracked, however, in 1997, launching a search for a more secure replacement that would be faster than 3DES. The result of this process was the new Advanced Encryption Standard (AES), which is gradually being introduced in government agencies to phase out DES and 3DES.

Implementation

DES uses a 64-bit key, of which only 56 bits are used for encryption, while the remaining 8 bits are employed for error correction. The algorithm transforms 64 bits of plaintext into ciphertext blocks of the same size. Since DES is a symmetric key algorithm, both the sender and the receiver require the same key in order for secure communications to be implemented. To exchange a DES session key between two parties, an asymmetric key algorithm such as Diffie-Hellman (DH) or RSA can be employed.

DES can operate in several different modes, including cipher block chaining (CBC) and Electronic Codebook (ECB) mode. ECB uses DES directly to encrypt and decrypt information, while CBC chains blocks of ciphertext together.

Notes

The American Standards Institute (ANSI) refers to DES as the Data Encryption Algorithm (DEA).

See Also: 3DES, Advanced Encryption Standard (AES), asymmetric key algorithm, Diffie-Hellman (DH), RSA, symmetric key algorithm

data integrity

The validity of data that is transmitted or stored.

Overview

Maintaining data integrity is essential to the privacy, security, and reliability of critical business data.

There are many ways in which this integrity can be compromised:

- Corruption of data resulting from software bugs or the actions of malicious users
- Viruses infecting computer systems and Trojans masquerading as genuine applications
- Hardware failures caused by age, accident, or natural disasters
- Human error in entering, storing, or transmitting data over a network

To minimize these threats to data integrity, you should implement the following procedures:

- Back up important data regularly and store backups in a safe location.
- Use access control lists (ACLs) to control who is allowed to access data.
- Maintain and replace aging hardware to prevent unexpected failure.
- Include code in your applications for validating data input.
- Use digital signatures to ensure data has not been tampered with during storage or in transmission.

See Also: *backup plan, disaster recovery plan (DRP), Trojan, virus*

Data Protection API (DPAPI)

An application programming interface that is part of Microsoft CryptoAPI (CAPI) on Microsoft Windows platforms.

Overview

Data Protection API (DPAPI) implements Microsoft Windows Data Protection on Windows 2000, Windows XP, and Windows Server 2003 platforms. DPAPI is an operating system–level password-based data protection service that applications can use to encrypt and decrypt information. DPAPI uses the 3DES encryption algorithm

and strong keys generated from user passwords, typically the password of the currently logged-on user. Since multiple applications running under the same account might use the same password and have access to such encrypted data, DPAPI also allows an application to provide an additional “secret,” called secondary entropy, to ensure only that application can decrypt information it has previously encrypted. The process by which DPAPI generates a cryptographic key from a password is called Password-Based Key Derivation and is defined in the Public Key Cryptography Standards (PKCS) #5 standard.

Notes

DPAPI does not store encrypted information, and applications that use it must implement their own storage mechanisms for this purpose.

See Also: *3DES, password*

DCS-1000

Formerly known as Carnivore, a surveillance technology used by the FBI for monitoring e-mail.

Overview

Few actual details are known about DCS-1000 apart from the fact that it can be installed at an Internet service provider and configured to monitor various aspects of traffic in transit through the provider’s network. The Electronic Privacy Information Center (EPIC), concerned about the privacy of businesses and the public, has employed the Freedom of Information Act (FOIA) to force disclosure of some information concerning the platform, but the FBI has assured the public that it only uses the system to capture e-mail authorized for seizure by a court order, as opposed to unrestrictively capturing all online traffic.

For More Information

Further information can be found on the FBI Web site at www.fbi.gov/hq/lab/carnivore/carnivore.htm.

See Also: *privacy*

DDoS

Stands for distributed denial of service, a type of denial of service (DoS) attack that leverages the power of multiple intermediary hosts.

See: distributed denial of service (DDoS)

DEA

Stands for Data Encryption Algorithm, the name used by the American National Standards Institute (ANSI) for Data Encryption Standard (DES).

See: Data Encryption Standard (DES)

decryption

The process of converting ciphertext into plaintext.

Overview

Encryption and decryption are complementary aspects of cryptography. The first involves transforming plaintext (digital information containing human-readable content) into ciphertext (scrambled information that cannot be directly read by humans). Decryption is the reverse process, which recovers the meaning of an encrypted message by transforming it from ciphertext back into plaintext.

The approach used for decrypting messages depends on the method used to encrypt them. For example, in a symmetric (or secret) key algorithm, both the sender and the recipient use the same shared secret key to encrypt and decrypt the message. In asymmetric key algorithms such as those used by public key cryptography systems, two keys are used, one to encrypt the message and the other to decrypt it.

See Also: asymmetric key algorithm, cryptography, encryption, public key cryptography, symmetric key algorithm

Defcon

A popular hackers' convention held each fall in Las Vegas.

Overview

Defcon has been referred to by its organizers as the "annual computer underground party for hackers." In addition to papers and presentations on everything from how to hack a system to how to secure a system against attack by others, other topics discussed include phone phreaking, privacy issues, demonstration of new hacking and security tools, recently discovered vulnerabilities and how to exploit and correct them, advances in Trojan and remote-control technologies, and so on.

Defcon is generally well attended by hackers, security professionals, and representatives of government, law enforcement, and media agencies. Fun activities are usually included such as a capture-the-flag type of contest in which groups of hackers are pitted against each other to try to hack each other's networks while simultaneously defending their own networks against attack. Awards are often given; for instance, one was given at Defcon 9 to an individual who hacked the conference network itself in order to gain admission to the conference without a pass.

Defcon was founded by Jeff "Dark Tangent" Moss and had its 10th annual conference in August 2002, with attendance running around 5000 and some sessions being standing room only. Defcon has evolved somewhat from its early freewheeling days and has become more "respectable" as it began to attract IS managers concerned about their growing network security needs. Defcon immediately follows another conference called Black Hat Briefings, which brings legitimate and underground security experts together to discuss the latest network security issues and methodologies.

For More Information

Visit Defcon at www.defcon.org for information about upcoming conferences and archived information from previous ones.

See Also: *Black Hat Briefings, hacker, phreaking*

defense in depth

A layered approach to implementing network security.

Overview

The goal of defense in depth is to provide multiple barriers for attackers attempting to compromise the security of your network. These layers provide extra hurdles for the attacker to overcome, thus slowing down the attack and providing extra time for detecting, identifying, and countering the attack. For example, the first layer of defense against passive attacks such as eavesdropping might be implementing link- or network-layer encryption, followed by security-enabled applications as a backup defense. Defense against insider attacks can consist of layers such as physical security, authenticated access control, and regular analysis of audit logs.

From a more general perspective, the first line of defense for a network occurs at its perimeter where firewalls block unwanted traffic and intrusion detection systems (IDSs) monitor traffic passed through the firewall. Additional layers behind this can include host-based firewalls and IDSs, proper access control lists (ACLs) on server resources, strong password policies, and so on.

See Also: *access control list (ACL), firewall, intrusion detection system (IDS), password*

demilitarized zone (DMZ)

An isolated network segment at the point where a corporate network meets the Internet.

Overview

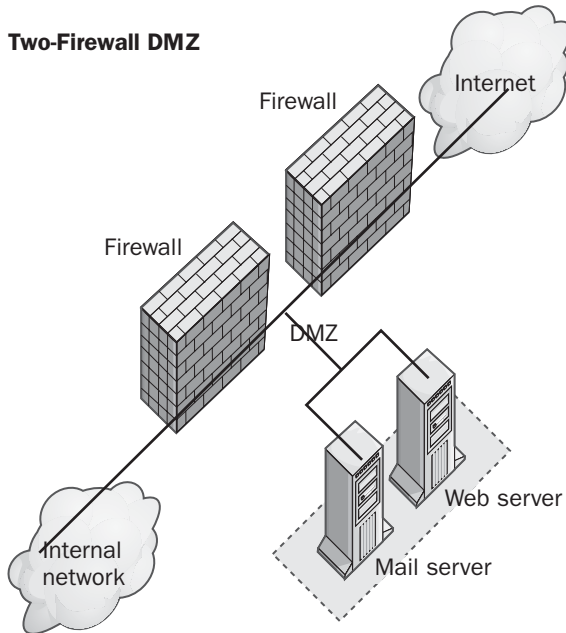
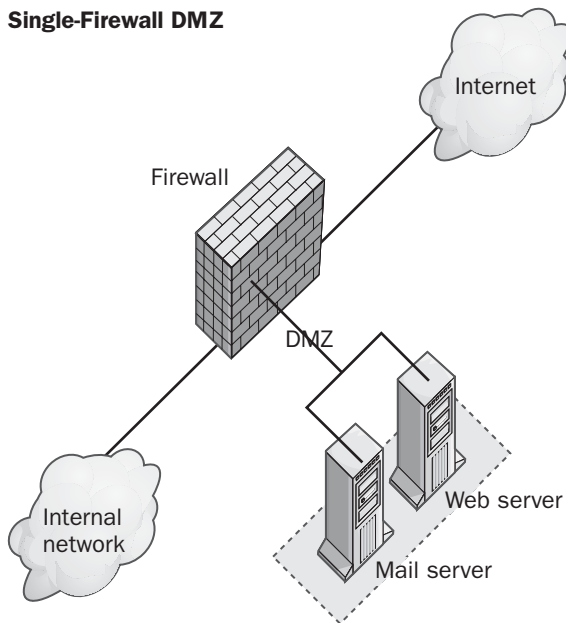
The demilitarized zone (DMZ) is a critical part of securing your network against attack. The term originated in the Korean War to refer to an area that both sides agreed to stay out of, which acted as a buffer zone to prevent hostilities from flaring up again.

In a networking scenario, the DMZ is used to segregate the private and public from each other while allowing essential network services such as Web site hosting, electronic messaging, and name resolution to function properly. To accomplish this, the DMZ is typically the location where hardened hosts such as Web, mail, and DNS servers are placed so they can handle traffic from both the internal and the external networks. This reduces the attack surface on both these hosts in particular and your network in general, for if these hosts were located outside the DMZ they would be more easily subject to attack, while if they were located inside the DMZ, compromising such a host could lead to penetration of your entire network.

Implementation

There are a variety of ways of implementing a DMZ, with two of the more popular being the following:

- **Dual-firewall DMZ:** Here, both the private and public networks terminate with firewalls, and the DMZ is the network segment connecting the two firewalls together. This approach is probably the most popular one in use today for implementing a DMZ.
- **Single-firewall DMZ:** This was the earliest approach to implementing a DMZ and consisted of a single firewall with three interfaces, one each for the private network, public Internet, and DMZ network segment.

Two-Firewall DMZ**Single-Firewall DMZ**

Demilitarized zone (DMZ). Single- and dual-firewall DMZ configurations.

Notes

The term **perimeter network** is more commonly used instead of DMZ in Microsoft networking environments.

See Also: [firewall](#)

denial of service (DoS)

A type of attack that tries to prevent legitimate users from accessing network services.

Overview

In a denial of service (DoS) attack, the attacker tries to prevent access to a system or network by several possible means, including the following:

- Flooding the network with so much traffic that traffic from legitimate clients is overwhelmed
- Flooding the network with so many requests for a network service that the host providing the service cannot receive similar requests from legitimate clients
- Disrupting communications between hosts and legitimate clients by various means, including alteration of system configuration information or even physical destruction of network servers and components

The earliest form of DoS attack was the SYN flood, which first appeared in 1996 and exploits a weakness in Transmission Control Protocol (TCP). Other attacks exploited vulnerabilities in operating systems and applications to bring down services or even crash servers. Numerous tools were developed and freely distributed on the Internet for conducting such attacks, including Bonk, LAND, Smurf, Snork, WinNuke, and Teardrop.

TCP attacks are still the most popular form of DoS attack. This is because other types of attack such as consuming all disk space on a system, locking out user accounts in a directory, or modifying routing tables in a router generally require networks to be penetrated first, which can be a difficult task when systems are properly hardened.

Defenses against DoS attacks include these:

- Disabling unneeded network services to limit the attack surface of your network
- Enabling disk quotas for all accounts including those used by network services
- Implementing filtering on routers and patch operating systems to reduce exposure to SYN flooding
- Baselining normal network usage to help identify such attacks in order to quickly defeat them
- Regularly backing up system configuration information and ensuring strong password policies

See Also: distributed denial of service (DDoS), SYN flooding

Department of Defense Information Technology Security Certification and Accreditation Process (DITSCAP)

A standardized approach for certifying the security of IT (information technology) systems.

Overview

The Department of Defense Information Technology Security Certification and Accreditation Process (DITSCAP) was developed to help guide U.S. Department of Defense (DoD) agencies by providing guidance for the accreditation process of IT systems. DITSCAP is a four-stage process involving

- Defining and documenting mission, function, requirements, and capabilities
- Recommending changes and summarizing them as a system security authorization agreement (SSAA), which summarizes specifications for the system being developed
- Validating the SSAA using vulnerability and penetration testing, resulting in full, interim, or withheld accreditation

- Postaccreditation monitoring and maintenance to ensure continued security

The goal of DITSCAP is to introduce integrated security into the life cycle of IT systems to minimize risks in shared infrastructures. DITSCAP was developed as a joint effort by the DoD, the Defense Information Systems Agency (DISA), and the National Security Agency (NSA). A related standard called National Information Assurance Certification and Accreditation Process (NIACAP) is employed for similar purposes between U.S. government agencies and contractors and consultants.

See Also: National Information Assurance Certification and Accreditation Process (NIACAP)

DES

Stands for Data Encryption Standard, an encryption standard used for many years by the U.S. federal government.

See: Data Encryption Standard (DES)

DESX

An enhanced version of the Data Encryption Standard (DES).

Overview

DESX, which stands for “DES XORed,” is a variant of DES developed by Ron Rivest in the 1980s. DESX performs similarly to DES but has greater resistance to exhaustive key search attacks. This is accomplished by XORing the input plaintext file with 64 bits of additional key material prior to encrypting the text using DES, a process sometimes called **whitening**, which is now implemented in other encryption schemes. Once DES has been applied to the whitened text, the result is again XORed with the same amount of additional key material.

See Also: Data Encryption Standard (DES)

DH

Stands for Diffie-Hellman, an algorithm used in public key cryptography schemes.

See: *Diffie-Hellman (DH)*

dictionary attack

A technique for cracking passwords.

Overview

The simplest but least efficient method for cracking passwords is the **brute-force attack**, which systematically tries all possible values in an attempt to guess the password. The dictionary attack is an improvement on this; it uses a dictionary (database) of common passwords derived from shared experiences of password crackers. Dictionary attacks can be performed online or offline, and readily available tools exist on the Internet for automating such attacks. A combination of a dictionary attack and a brute-force attack is called a hybrid attack.

In addition to cracking passwords, dictionary attacks have been used in other scenarios such as guessing community names on a network that uses Simple Network Management Protocol (SNMP). Once these names are guessed, the attacker can use SNMP to profile services on the targeted network.

See Also: *brute-force attack, hybrid attack*

Diffie-Hellman (DH)

An algorithm used in public key cryptography schemes.

Overview

Diffie-Hellman (DH) was the first algorithm developed for public key cryptography. It is used for key exchange by a variety of security protocols, including Internet Protocol Security (IPSec), Secure Sockets Layer (SSL), and Secure Shell (SSH), as well as many popular public key infrastructure (PKI) systems.

DH was developed by Whitfield Diffie and Martin Hellman in 1976 and was the first protocol developed for enabling users to exchange a secret over an insecure medium without an existing shared secret between

them. DH is not an encryption algorithm but a protocol for exchanging secret keys to be used for sending encrypted transmissions between users using Data Encryption Standard (DES), Blowfish, or some other symmetric encryption scheme.

Issues

DH in its simplest form is susceptible to man-in-the-middle attacks, though this can be mitigated by necessitating the use of digital signatures by all parties. The Station-to-Station (STS) protocol is an authenticated version of DH developed in 1992 that uses keys certified by certificate authorities (CAs) to prevent such attacks.

See Also: *public key cryptography*

diffing

A technique used by hackers that compares different versions of files to look for differences.

Overview

The word **diffing** derives from the diff utility on UNIX systems that performs bitwise comparison between two files. A variety of diffing tools exist that work at the file, database, and disk levels. These tools are sometimes used by hackers to compare a new version of a file with an earlier version for various reasons, including the following:

- Discovering where an application stores password information by entering a password, taking a bit-image snapshot of the application, changing the password, taking another snapshot, and diffing the two file images. This operation can show exactly where within the compiled code the password information is stored, and this may be of use in cracking other users' passwords.
- Determining what effects a patch has when applied to an application. When vendors create patches, they may not fully disclose the vulnerabilities corrected, and by diffing the application before and after the patch and examining the result, a hacker may learn more about the original vulnerabilities. Using this information, the hacker can then proceed

to attack unpatched versions of the application on other systems.

D

Examples of tools used for diffing include the Windows `fc` and UNIX `diff` commands. Once a file has been diffed to locate the section of code that has changed, the hacker can then use a hex editor such as Hackman to make bitwise modifications to the file if desired.

See Also: `hex editor`

Digest authentication

A Hypertext Transmission Protocol (HTTP) authentication scheme based on challenge–response authentication.

Overview

Digest authentication is a method used by Web servers to authenticate users trying to access sites. Digest authentication was proposed in RFC 2617 as a more secure method than Basic authentication, which passes user credentials across the connection in cleartext. Instead, Digest authentication encrypts user credentials as an MD5 hash to prevent credential theft by malicious users eavesdropping on the network.

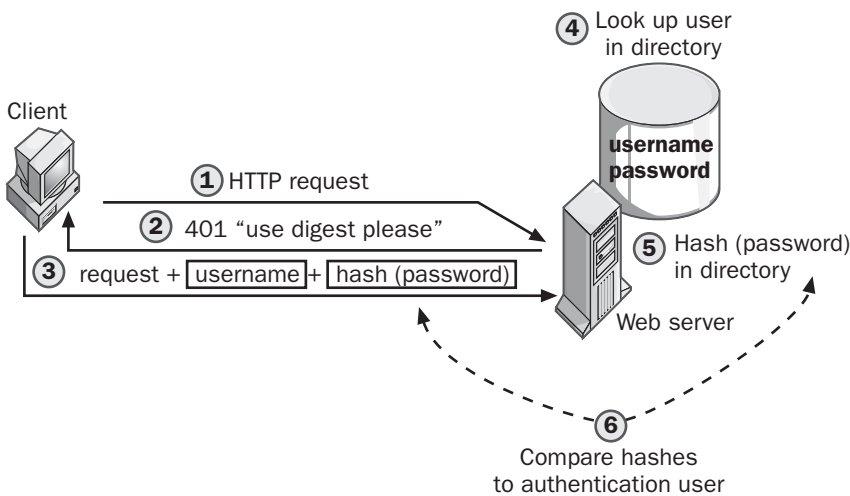
Digest authentication is supported by Internet Information Services (IIS) on Microsoft Windows server

platforms, the open source Apache Web server, the Jigsaw Web server developed by the World Wide Web Consortium (W3C), and many other platforms. Digest authentication can also be incorporated directly into Microsoft .NET–managed code, bypassing the version included in IIS on Microsoft Windows platforms.

Implementation

When a client browser tries to access a Web site on which Digest authentication is configured, the client begins by making an unauthenticated HTTP request to the server. The server responds with an HTTP 401 Unauthorized status code, sending a token called a nonce to the client and telling the client in the HTTP response header that it must use Digest authentication to access the site. The client then opens a dialog box to obtain the user’s name and password, hashes the password together with the nonce, and sends the username and hash to the server requesting authentication.

The server then generates the same hash using the copy of the user’s password stored in its security accounts database and compares this hash with the one received from the client. If the two hashes match, the client is allowed to download the requested resource from the server.



Digest authentication. How Digest authentication works.

Issues

Digest authentication is susceptible to replay attacks, but this can be minimized by time-limiting nonce values or using different values for each connection. While Digest authentication is more secure than Basic authentication, it is not as secure as Kerberos authentication or authentication based on client certificates. Another issue with the security of Digest authentication is that it requires passwords to be retrievable as cleartext.

See Also: authentication, Basic authentication, challenge response authentication, MD5, replay attack

DigiCrime

A Web site that humorously draws attention to information security issues.

Overview

DigiCrime (www.digicrime.com) is the brainchild of mathematician and computer scientist Kevin McCurley, and since 1996 this site has entertained the security community and informed the general public about potential issues in computer and online security. The site humorously promotes itself as offering “a full range of criminal services and products to our customers.” These “services” include identity theft, money laundering, airline ticket rerouting, telephone wiretapping, spamming, and more. The idea behind these “services” is to educate and inform the general public of potential dangers in blindly trusting online transactions and to challenge the security community and software vendors to take these dangers more seriously. The site includes a community of real individuals with tongue-in-cheek titles like Director of Disinformation, Chief of Insecurity, Illegal Counsel, and Chief Arms Trafficker, many of whom are security professionals or cryptography experts and who help contribute to the site.

digital certificate

Encrypted information that guarantees that an encryption key belongs to a user.

Overview

Sometimes simply called **certificates**, digital certificates are specially formatted digital information that is

used in secure messaging systems that employ public key cryptography. Certificates are used to verify the identity of the message sender to the recipient by generating a digital signature that can be used to sign the message. They are also used for providing the recipient of an encrypted message with a copy of the sender’s public key.

Digital certificates are issued by a certificate authority (CA) that is trusted by both the sender and recipient. The most common format used for certificates is the X.509 standard, which contains the user’s name and public key, a serial number, expiration date, the name and digital signature of the CA that issued the certificate, and other information. When a recipient receives an encrypted message with a certificate attached, the recipient uses the CA’s public key to decrypt the certificate and verify the sender’s identity.

See Also: digital signature, public key cryptography, X.509

digital fingerprinting

Another name for digital watermarking, a Digital Rights Management (DRM) antipiracy and copy-protection technology.

See: digital watermarking

digital forensics

The science of applying digital technologies to legal questions arising from criminal investigations.

Overview

Traditional forensic methods used in criminal investigations include looking for footprints, fingerprints, hair, fiber, and other physical evidence of an intruder’s presence. In computer crime, the evidence left behind is of a digital nature and can include data on hard drives, logs of Web server visits or router activity, and so on. Digital forensics is the science of mining computer hardware and software to find evidence that can be used in a court of law to identify and prosecute cybercriminals.

Many companies have deployed an intrusion detection system (IDS) on their network to monitor and detect

possible breaches of network security. When a breach has occurred, these companies may not have the necessary expertise to determine the extent of the breach or how the exploit was performed. In serious cases in which significant business loss has resulted, companies must establish an evidence trail to identify and prosecute the individuals responsible. In such cases, companies may enlist the services of digital forensic experts who can send in an incident response team to collect evidence, perform a “postmortem” by piecing together the evidence trail, help recover deleted files and other lost data, and perform “triage” to help restore compromised systems as quickly as possible.

Marketplace

Examples of companies offering digital forensics services include @stake, Computer Forensics, DigitalMedix, ESS Data Recovery, Guidance Software, Vigilinx, and others. Computer Sciences Corporation and Veridian share a significant portion of the digital forensics market for the U.S. federal government.

See Also: intrusion detection system (IDS)

Digital Millennium Copyright Act (DMCA)

Legislation that extends U.S. copyright law to cover digital content.

Overview

The Digital Millennium Copyright Act (DMCA) was enacted in 1998 as a vehicle for compliance toward treaties with the World Intellectual Property Organization (WIPO), a United Nations agency based in Geneva, Switzerland. The provisions of the DMCA include the following:

- Outlawing the circumvention of antipiracy measures such as Digital Rights Management (DRM) technologies built into commercial software. The law also outlaws the manufacture, sale, or distribution of devices or software to illegally crack or copy such software. Exceptions are allowed for those who conduct research and development of encryption and antipiracy technologies and for libraries

and other nonprofit organizations in certain circumstances.

- Requiring Internet service providers to remove any information on users’ Web sites that may constitute copyright infringement. Liability for simple transmission of such information by third parties is limited for these service providers, however, and for educational institutions hosting student Web sites.
- Requiring Web sites broadcasting copyrighted digital audio or video to pay licensing fees to companies producing such content.
- Upholding generally accepted “fair use” exemptions mandated by previous copyright legislation.

Issues

The DMCA has been widely praised by the entertainment and software industry but generally criticized by academics, librarians, and civil libertarians as part of larger issues surrounding the purposes and means of implementing DRM technologies in the consumer marketplace. A notable application of the DMCA was the arrest in 2001 of Russian programmer Dmitry Sklyarov, who was apprehended after a Defcon conference at which he presented a paper on how to circumvent copyright protection technology built into Adobe eBooks software.

See Also: Digital Rights Management (DRM)

Digital Rights Management (DRM)

Any technology used to protect the interests of copyright holders of commercial digital information products and services.

Overview

The last decade has seen the advent of consumer digital information products and services such as CD audio, DVD video, CD- and DVD-ROM software, and digital television. The potential for making illegal copies of digital products using standard computer hardware and software or through online file-sharing services has been viewed by the entertainment and software industries as

potentially reducing their revenues by opening a floodgate of copyright circumvention and software piracy. This danger is enhanced by the nature of digitized information, which allows such copies to contain exactly the same information as the original.

In response to this issue, companies such as Microsoft and others have developed various Digital Rights Management (DRM) technologies to protect commercial digital products and services. These technologies may control access to such products and services by preventing the sharing or copying of digital content, limiting the number of times content can be viewed or used, and tying the use or viewing of content to specific individuals, operating systems, or hardware.

Implementation

There are two general methods for implementing DRM:

- Encrypting the information so that only authorized users or devices can use it. An example is Microsoft Windows Media DRM, an end-to-end DRM system that provides content providers and retailers with the tools to encrypt Microsoft Windows Media files for broadcast or distribution.
- Including a “digital watermark” to secretly identify the product or service as copyrighted and to signal to the hardware displaying the content that the material is copy protected. A Federal Communications Commission (FCC) proposal to incorporate a “broadcast flag” into digital television signals is one example of this approach.

Various industry groups are working toward DRM standards, including the Internet Engineering Task Force (IETF), the MPEG Group, the OpenEBook Forum, and several others. Microsoft Corporation’s next-generation secure computing base, part of its Trustworthy Computing initiative, includes the incorporation of DRM technologies into the Microsoft Windows operating system platforms.

Issues

Critics of the encryption approach to DRM suggest that such technologies weaken the privacy of consumers by requiring them to provide personal information before

content can be viewed or used. Such collected information may then be used to profile consumer purchase patterns for marketing purposes and price discrimination, to limit access to certain kinds of material to certain classes of consumers, or to push users toward a pay-per-view licensing model to enhance the revenue stream for content providers.

For More Information

For information about Microsoft Windows Media DRM, see www.microsoft.com/windows/windowsmedia/drm.aspx.

See Also: digital watermarking, next-generation secure computing base

digital signature

Digital information used for purposes of identification of electronic messages or documents.

Overview

Digital signatures are a way of authenticating the identity of creators or producers of digital information. A digital signature is like a handwritten signature and can have the same legal authority in certain situations, such as buying and selling online or signing legal contracts. Digital signatures can also be used to ensure that the information signed has not been tampered with during transmission or repudiated after being received.

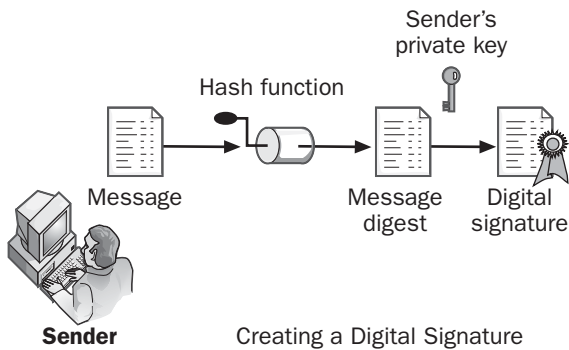
Digital signatures are dependent on public key cryptography algorithms for their operation. There are three public key algorithms that are approved Federal Information Processing Standards (FIPS) for purposes of generating and validating digital signatures:

- Digital Signature Algorithm (DSA)
- Elliptic Curve DSA (ECDSA)
- RSA algorithm

Implementation

To create a digital signature, the document or message to be transmitted is first mathematically hashed to produce a message digest. The hash is then encrypted using the sender’s private key to form the digital signature, which is appended to or embedded within the message.

Once the encrypted message is received, it is decrypted using the sender's public key. The recipient can then hash the original message and compare it with the hash included in the signature to verify the sender's identity. Nonrepudiation is guaranteed by the fact that the sender's public key has itself been digitally signed by the certificate authority (CA) that issued it.



Digital signature. Creating a digital signature.

Notes

Digital signatures are not the same as digital certificates. Digital certificates are like a driver's license you can use to identify yourself that is issued by a trusted third party, in the case of digital certificates, one called a certificate authority (CA). Included in your digital certificate are your private and public keys, which can be used to send encrypted messages and enable recipients to decrypt them. Your private key is then used to create your digital signature, so a digital certificate is a prerequisite for digitally signing documents.

See Also: *certificate authority (CA), digital certificate, Digital Signature Algorithm (DSA), Digital Signature Standard (DSS), Elliptic Curve Digital Signature Algorithm (ECDSA), hashing algorithm, public key cryptography, RSA*

Digital Signature Algorithm (DSA)

A public key cryptography algorithm used to generate digital signatures.

Overview

The Digital Signature Algorithm (DSA) is a public key algorithm used for creating digital signatures to verify the identity of individuals in electronic transactions. Signatures created using DSA can be used in place of handwritten signatures in scenarios such as legal contracts, electronic funds transfers, software distribution, and other uses. Although DSA is a public key algorithm, it is used mainly for digitally signing documents and not for encrypting them.

DSA is patented by the National Institute of Standards and Technology (NIST) and forms the basis of the Digital Signature Standard (DSS).

See Also: *digital signature, Digital Signature Standard (DSS), Federal Information Processing Standard (FIPS), National Institute of Standards and Technology (NIST), public key cryptography*

Digital Signature Standard (DSS)

A U.S. federal government standard defining how digital signatures are generated.

Overview

The Digital Signature Standard (DSS) is a Federal Information Processing Standard (FIPS) 186-2 issued in 1994. The goal of the standard is to promote electronic commerce by providing a way for documents and messages to be electronically signed using digital signatures. DSS employs two cryptographic algorithms for this purpose:

- **DSA:** A public key algorithm patented by the National Institute of Standards and Technology (NIST)
- **SHA-1:** A hashing algorithm standardized by NIST as FIPS 180

DSS is widely used in federal government and defense agencies for transmission of unclassified information.

See Also: *digital signature, public key cryptography, Secure Hash Algorithm (SHA-1)*

digital watermarking

A Digital Rights Management (DRM) antipiracy and copy-protection technology.

Overview

Digital watermarking enables digital content producers to insert hidden information in digital products and data streams to prevent them from being illegally used or copied. Such watermarks can be embedded into any form of commercially sold digital content, including audio CDs, DVD movies, software on CD- or DVD-ROMs, streaming audio and video, digital television, and so on. Watermarks can include information for copyright protection and authentication information to control who can use content and how such content can be used.

Implementation

There are two basic types of digital watermarks: visible and invisible. Visible watermarks resemble those formerly used to identify vendors of high-quality bond paper and are generally used to discourage copying of digital content. Visible watermarks do not prevent such copying from occurring, but instead may deter such copying by potentially providing legal evidence of copyright infringement through illegal copying of digital media. Invisible watermarks, on the other hand, can be used both for legal evidence and to implement invisible copy-protection schemes for media players designed to read them.

Most watermarking techniques involve manipulating digital content in the spatial or frequency domain using a mathematical procedure called fast Fourier transforms (FFT). Images of text can also be watermarked by subtly altering line and character spacing according to fixed rules.

Marketplace

A leading provider of digital-watermarking technologies and products is Digimarc (www.digimarc.com).

Notes

Another name used to refer to this procedure is **digital fingerprinting**.

See Also: *Digital Rights Management (DRM)*

disaster recovery plan (DRP)

A plan that helps a company recover data and restore services after a disaster.

Overview

Digital information is the lifeblood of today's companies, and loss of data means loss of business services and loss of revenue. Disasters that can destroy data can take many forms:

- Natural disasters such as floods and earthquakes
- Manufactured disasters such as terrorist attacks and criminal network intrusions
- Disasters caused by hardware failures or buggy software
- Accidental disasters from human error

Guarding against such disasters is important, but it's prudent to expect the worst and plan accordingly. Essential to the success of any company's IT (information technology) operations is a disaster recovery plan (DRP) to enable it to recover quickly after a disaster and restore services to customers. This can range from a simple plan to create a backup of the server every night in a small company, to the kind of technological redundancies and procedures that enabled Wall Street to recover from 9/11 after only a week. Clearly, a DRP is not a bandage you apply after things go wrong but a fundamental business practice a company should consider from day one of implementing its IT systems.

Implementation

Creating a good DRP begins with risk assessment and planning. Risk assessment determines the likelihood and scale of potential disasters, which aids in planning which technologies to implement and how much to budget. Planning involves determining which systems and data need to be backed up, how often they should be backed up, and where backed up data should be securely stored.

Selecting an appropriate backup technology and developing an appropriate backup plan for using such technology is important to avoid excessive costs and ensure reliable recovery after a disaster. Backup technologies

D

can include tape backup systems, recordable CDs and DVDs, backup to remote storage area networks (SANs) over secure virtual private network (VPN) connections, and backup to service provider networks. Outsourcing of backup needs is another option a company may consider if its IT department is small and can't manage such needs. The addition of hot-standby systems can greatly simplify the recovery process if financially feasible.

If your company uses IT services from service providers, it is essential to have service level agreements (SLAs) from these providers to help guarantee business continuity after a disaster. Establishing suitable information security policies and procedures is also essential to making a DRP work.

Once your DRP is up and running, it needs to be regularly tested and monitored to be sure it works. Verification of backups ensures information truly is being backed up, and periodic restores on test machines ensure that the DRP will work should it ever need to be implemented. If such monitoring and testing find weaknesses or problems in your plan, you need to modify the plan accordingly.

Having an external audit of your DRP by a company with expertise in this area can also be valuable. ISO 17799 is a recognized standard in IT security best practices, and auditing on this basis can be advantageous on a legal liability basis if your company provides information services to others.

Another essential component of a DRP is a business resumption plan (BRP), sometimes called a business continuity plan (BCP). This is a detailed step-by-step plan on how to quickly resume normal business after a disaster occurs.

Fundamentally, however, your DRP will never be fully tested until a significant disaster occurs.

See Also: backup plan, business resumption plan (BRP)

discretionary access control (DAC)

A mechanism for controlling access by users to computing resources.

Overview

Discretionary access control (DAC) is one of two basic approaches to implementing access control on computer systems, the other being mandatory access control (MAC). DAC specifies who can access a resource and which level of access each user or group of users has to the resource. DAC is generally implemented through the use of an access control list (ACL), a data structure that contains a series of access control entries (ACEs). Each ACE includes the identity of a user or group and a list of which operations that user or group can perform on the resource being secured.

Most computing platforms, including Microsoft Windows, Linux, and different flavors of UNIX, implement some form of DAC mechanism for controlling access to file system and other types of resources.

See Also: access control, access control entry (ACE), access control list (ACL), mandatory access control (MAC)

discretionary access control list (DACL)

The most common type of access control list (ACL) used to control access to computer and network resources.

Overview

Discretionary access control lists (DACLs) are one of two forms of ACLs, the other being system access control lists (SACLs). DACLs are the most general of these two types and are assigned to file system and other computing resources to specify who can access them and which level of access that user or group can have. In fact, when ACL is referred to in discussion, it can usually be assumed to refer to DACL unless system auditing is included. Using DACLs, an operating system can implement discretionary access control (DAC) for enforcing what users can or cannot do with system resources.

See Also: access control, access control list (ACL), discretionary access control (DAC), system access control list (SACL)

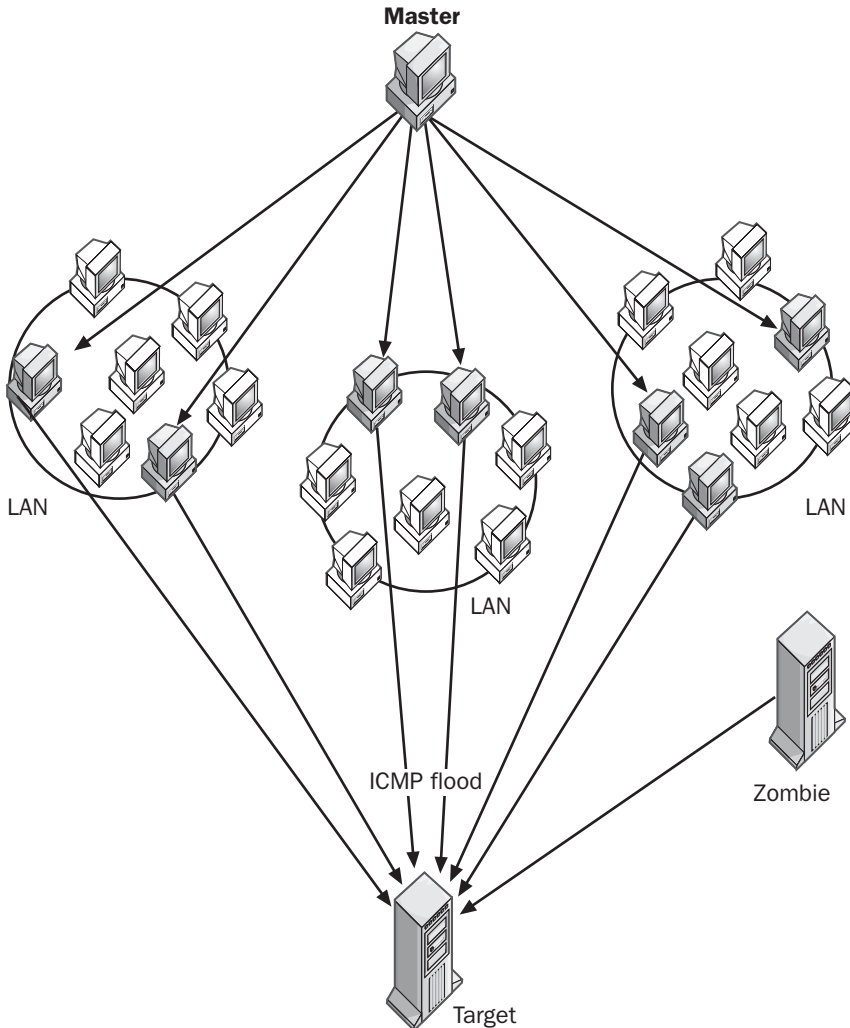
distributed denial of service (DDoS)

A type of denial of service (DoS) attack that leverages the power of multiple intermediary hosts.

Overview

Classic DoS attacks are one-to-one attacks in which a more powerful host generates traffic that swamps

the network connection of the target host, thus preventing legitimate clients from accessing network services on the target. The distributed denial of service (DDoS) attack takes this one step further by amplifying the attack manyfold, with the result that server farms or entire network segments can be rendered useless to clients.

D


Distributed denial of service. How a DDoS attack works.

D

DDoS attacks first appeared in 1999, just three years after DoS attacks using SYN flooding brought Web servers across the Internet to their knees. In early February 2000, a major attack took place on the Internet, bringing down popular Web sites such as Amazon, CNN, eBay, and Yahoo! for several hours. A more recent attack of some significance occurred in October 2002 when 9 of the 13 root DNS servers were crippled by a massive and coordinated DDoS attack called a ping flood. At the peak of the attack, some of these servers received more than 150,000 Internet Control Message Protocol (ICMP) requests per second. Fortunately, because of caching by top-level Domain Name System (DNS) servers and because the attack lasted only a half hour, traffic on the Internet was not severely disrupted by the attack.

Implementation

The theory and practice behind performing DDoS attacks is simple:

- 1 Run automated tools to find vulnerable hosts on other networks connected to the Internet. Once a vulnerable host is found, such tools can compromise the host and install a DDoS Trojan, turning the host into a zombie that can be controlled remotely by a master station that the attacker uses to launch the attack. Popular tools for launching such DDoS attacks include TFN, TFN2K, Trinoo, and Stacheldraht, all of which are readily available on the Internet.
- 2 Once enough hosts have been compromised, the attacker uses the master station to signal the zombies to commence the attack against the target host or network. This attack is usually some form of SYN flood or other simple DoS attack scheme, but the fact that hundreds or even thousands of zombie hosts are used in the attack creates a massive amount of network traffic that can quickly consume all Transmission Control Protocol (TCP) resources on the target and may even swamp the target's network connection to the Internet.

Almost all computer platforms are susceptible to being hijacked as zombies to conduct such an attack, including

Solaris, Linux, Microsoft Windows, and flavors of UNIX. The best way to defend against such attacks involves modifying router configurations at Internet service providers (ISPs), specifically:

- Filtering all RFC 1918 private Internet Protocol (IP) addresses using router access control lists
- Applying RFC 2267 ingress and egress filtering on all edge routers so that the client's side of the connection rejects incoming packets that have addresses originating within their own network, while the ISP's side accepts only packets that have addresses originating from the client's network
- Rate-limiting all ICMP and SYN packets on all router interfaces

For these practices to be most effective, the cooperation of the whole Internet community is required.

For More Information

A good resource on DDoS is the staff page of Dave Dittrich, senior security engineer at the University of Washington; see staff.washington.edu/dittrich/misc/ddos/.

See Also: denial of service (DoS), SYN flooding, zombie

DITSCAP

Stands for Department of Defense Information Technology Security Certification and Accreditation Process, a standardized approach for certifying the security of IT (information technology) systems.

See: Department of Defense Information Technology Security Certification and Accreditation Process (DITSCAP)

DMCA

Stands for Digital Millennium Copyright Act, legislation that extends U.S. copyright law to cover digital content.

See: Digital Millennium Copyright Act (DMCA)

DMZ

Stands for demilitarized zone, an isolated network segment at the point where a corporate network meets the Internet.

See: demilitarized zone (DMZ)

DNS cache poisoning

Another name for Domain Name System (DNS) spoofing, a method used for attacking DNS servers.

See: DNS spoofing

DNS spoofing

A method used for attacking Domain Name System (DNS) servers.

Overview

DNS spoofing provides DNS servers with false information to impersonate DNS servers. DNS spoofing can enable malicious users to deny access to authentic DNS servers, redirect users to different Web sites, or collect and read e-mail addressed to or sent from a given domain.

There are two basic approaches to DNS spoofing:

- By modifying a name server to provide false authoritative records in response to a recursive query, a malicious user can redirect all requests to a certain domain to an illicit DNS server. The result is that a user trying to access a popular site may be directed to a different site that looks the same but that has been set up to capture any personal information the user submits. A notorious example of this occurred in 1997 when Eugene Kashpureff used DNS spoofing to redirect users trying to access the InterNIC domain name registry to his own AlterNIC name registry.
- Another approach is to sniff a network connection over which DNS traffic regularly travels and spoof User Datagram Protocol (UDP) packets used in DNS queries. The attacker predicts the next query ID number and inserts this into a spoofed packet, thus hijacking the DNS query and redirecting the user to an illicit look-alike Web site.

The general approach to prevent such attacks includes patching DNS servers with the latest fixes, restricting zone transfers and dynamic updates, and turning off recursion if necessary. However, the real solution to the problem of DNS spoofing involves developing cryptographically authenticated DNS and deploying it across the Internet.

DNS spoofing can also be considered a form of denial of service (DoS) attack since it prevents users from accessing genuine DNS servers.

See Also: denial of service (DoS), spoofing

DoS

Stands for denial of service, a type of attack that tries to prevent legitimate users from accessing network services.

See: denial of service (DoS)

dot bug vulnerability

A type of coding vulnerability.

Overview

The dot bug vulnerability first appeared in 1997 when someone discovered that by appending two extra periods to the end of a Uniform Resource Locator (URL) requesting an Active Server Page (ASP) file from a Microsoft Internet Information Server (IIS) 3 Web server, you could view the ASP code instead of executing it. For example, browsing the URL `http://www.northwindtraders.com/somepage.asp` would cause the page to execute normally, while browsing `http://www.northwindtraders.com/somepage.asp..` would display the ASP code instead. Other similar exploits soon followed that had similar effect, including adding `2%e` in place of the period in `somepage.asp` and appending `::$DATA` to the end of the URL. A similar dot bug vulnerability that allowed scripts residing in cookies to be run and read information in other cookies was discovered in Microsoft Internet Explorer in February 2002.

Similar vulnerabilities have been found in other platforms and products. For instance, a dot bug vulnerability just like one found in ASP was later discovered in PHP,

another scripting platform for creating dynamic Web sites. A vulnerability was also discovered in the Hypertext Transfer Protocol (HTTP) server on the IBM AS/400 platform, whereby appending a forward slash (/) to the end of a URL would display the source code of the page.

Improved coding practices have generally resulted in fewer such bugs in the last few years.

See Also: *vulnerability*

DPAPI

Stands for Data Protection API, an application programming interface (API) that is part of CryptoAPI on Microsoft Windows platforms.

See: *Data Protection API (DPAPI)*

DRM

Stands for Digital Rights Management, any technology used to protect the interests of copyright holders of commercial digital information products and services.

See: *Digital Rights Management (DRM)*

DRP

Stands for disaster recovery plan, a plan that helps a company recover data and restore services after a disaster.

See: *disaster recovery plan (DRP)*

DSA

Stands for Digital Signature Algorithm, a public key cryptography algorithm used to generate digital signatures.

See: *Digital Signature Algorithm (DSA)*

Dsniff

A popular set of tools for network auditing and penetration testing.

Overview

Dsniff is a collection of tools used on UNIX/Linux platforms developed by Dug Song of the Center for Information Technology Integration at the University of Michigan. These tools are popular with network security professionals and hackers alike and in version 2.3 of Dsniff consist of the following:

- **Passive network monitoring tools:** Dsniff, Filesnarf, Mailsnarf, Msgsnarf, Urlnarf, and Webspay
- **Traffic interception tools:** Arpspoof, Dnsspoof, and Macof
- **Man-in-the-middle (MITM) attack tools:** shSmitm (for Secure Shell, SSH) and webmitm (for Hypertext Transfer Protocol Secure, HTTPS).

For More Information

See monkey.org/~dugsong/dsniff for more information.

See Also: *sniffer*

DSS

Stands for Digital Signature Standard, a U.S. federal government standard defining how digital signatures are generated.

See: *Digital Signature Standard (DSS)*

dynamic packet filtering

An advanced packet-filtering technology used by firewalls and some routers.

Overview

Packet filtering is used by routers and firewalls for filtering out undesired packets. Early routers employed static packet filtering, commonly called packet filtering, which allows routers to be manually configured to allow or block incoming or outgoing packets based on Internet Protocol (IP) address and port information found in packet headers. Dynamic packet filtering takes this a step further by opening ports only when required and closing them when no longer needed. Dynamic

packet filtering thus minimizes exposed ports and provides better security than static filtering.

Dynamic packet filtering is managed by creating policies that can rule for how long and when different ports should be opened or closed. All packets passing through the router or firewall are compared with these rules to determine whether to forward or drop them.

In addition to examining the packet header, some firewalls implementing dynamic packet filtering can inspect deeper layers of the TCP/IP protocol within each packet to create a state table containing information about each established connection. This allows them to filter packets not only by rules but also by state information concerning previous packets for that connection. This process is commonly called stateful inspection.

Marketplace

Microsoft Internet Security and Acceleration Server (ISA Server) supports policy-based dynamic packet filtering of IP traffic to enhance the security of your network. Most commercial firewalls also support some kind of dynamic packet filtering in their operation.

See Also: *firewall, packet filtering, stateful inspection*

dynamic proxy

Another name for adaptive proxy, an enhanced form of application-level gateway.

See: *adaptive proxy*

EAP

Stands for Extensible Authentication Protocol, a security extension for the Point-to-Point Protocol (PPP).

See: Extensible Authentication Protocol (EAP)

EAP-TLS

Stands for Extensible Authentication Protocol–Transport Layer Security, an encrypted authentication scheme based on Extensible Authentication Protocol (EAP).

See: Extensible Authentication Protocol–Transport Layer Security (EAP-TLS)

EAP-TTLS

Stands for Extensible Authentication Protocol–Tunneled Transport Layer Security, an encrypted authentication scheme based on Extensible Authentication Protocol (EAP) and easier to manage than Extensible Authentication Protocol–Transport Layer Security (EAP-TLS).

See: Extensible Authentication Protocol–Tunneled Transport Layer Security (EAP-TTLS)

eavesdropping

Secretly listening to traffic on a network.

Overview

Eavesdropping on telephone conversations (called wiretapping) requires specialized equipment and access to telephone-company switching facilities. Eavesdropping on Internet Protocol (IP) networks, however, is easy—just attach a “sniffer” to the network and capture all traffic

traveling on the network segment. The simplicity of network eavesdropping as opposed to wiretapping is caused by the inherent simplicity of the Transmission Control Protocol/Internet Protocol (TCP/IP) protocol itself, which has an open architecture that transmits data unencrypted over the network.

Why do people eavesdrop on networks? Innocent snooping may be one reason, but malicious hackers have more sinister reasons and want to capture passwords or credit card information to compromise systems or drain your bank account. Some forms of eavesdropping are of value, however, such as configuring firewalls to look deep inside incoming packets to determine whether their contents are safe.

How can you prevent attackers from eavesdropping on your network? Encryption is the surest method, and Internet Protocol Security (IPSec) is a popular protocol for encrypting IP communications. Some network managers have gone so far as to superimpose a second, virtual network on top of their physical one. In this scenario every network connection becomes a virtual private network (VPN) connection with end-to-end IPSec to secure it.

Besides encryption, another important step to prevent eavesdropping is to develop security policies and procedures and enforce them rigorously in the workplace. This will help prevent the kind of social engineering scenarios in which an attacker fakes its way into your company and installs a sniffer somewhere on your network. Antivirus software can also protect your network

against Trojans, which can be used to capture credentials and other sensitive information.

See Also: *Internet Protocol Security (IPSec), sniffing, social engineering, Trojan, virtual private network (VPN)*

ECB

Stands for Electronic Codebook, a mode of operation for block ciphers.

See: *Electronic Codebook (ECB)*

ECC

Stands for elliptic curve cryptography, cryptographic procedures based on elliptic curve mathematics.

See: *elliptic curve cryptography (ECC)*

ECDSA

Stands for Elliptic Curve Digital Signature Algorithm, an alternative to the Digital Signature Algorithm (DSA) based on elliptic curve cryptography (ECC).

See: *Elliptic Curve Digital Signature Algorithm (ECDSA)*

EFS

Stands for Encrypting File System, a Microsoft technology for protecting files stored on a hard drive.

See: *Encrypting File System (EFS)*

egress filtering

Filtering outgoing packets at a router or firewall.

Overview

Network managers are mostly concerned about protecting what comes into their networks, and traditionally firewalls and routers have been configured to filter incoming traffic from outside the corporate network,

a process called **ingress filtering**. Recently, however, the importance of egress filtering has become important for several reasons.

Malicious intruders who compromise networks may install software to sniff out passwords and other sensitive information and then transmit this information to the attacker. Egress filtering can help prevent such unauthorized information from leaking out of your network by blocking suspicious outbound traffic.

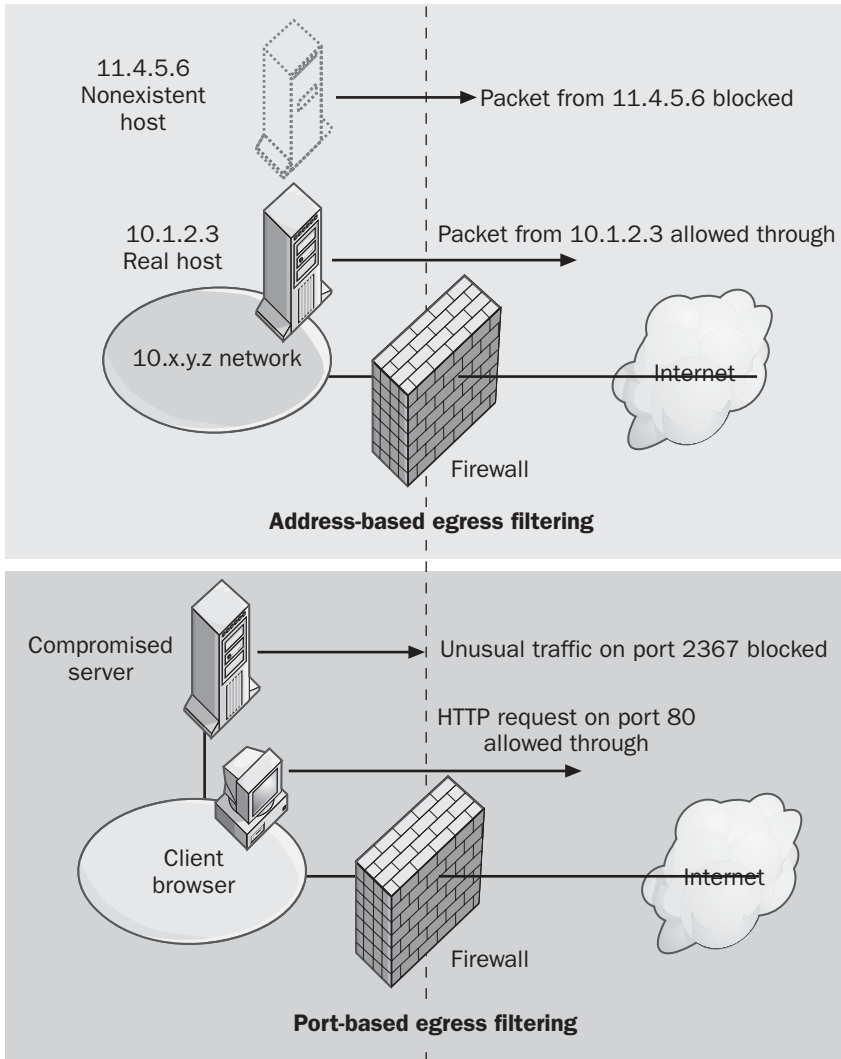
Attackers who penetrate your network can use it as a platform for launching distributed denial of service (DDoS) attacks on other networks. Using your own systems as remotely controlled “zombies,” a flood of outbound traffic can leave your network and wreak havoc with another company’s servers, whose intrusion detection system (IDS) points back to your network as the culprit.

To help keep the Internet a safer place for everyone, and to prevent your company from becoming the target of lawsuits as a result of being used to launch a DDoS attack, egress filtering can prevent any packets having invalid addresses or questionable port numbers from leaving your network.

Implementation

Egress filtering can be configured on routers and firewalls at two levels:

- IP address filtering can be configured to drop all outbound packets except those whose source address matches trusted hosts on your network.
- Port filtering can be configured to block all outbound packets except those well-known port numbers essential for Domain Name System (DNS), Hypertext Transfer Protocol (HTTP), Simple Mail Transfer Protocol (SMTP), and Post Office Protocol 3 (POP3) communications.



Egress filtering. Two types of egress filtering.

See Also: distributed denial of service (DDoS), firewall, intrusion detection system (IDS)

EICAR

Stands for European Institute of Computer Anti-Virus Research, an antivirus research organization.

See: European Institute of Computer Anti-Virus Research (EICAR)

EKE

Stands for Encrypted Key Exchange, a method of sharing a secret message between two parties that involves using a short password as the primary key.

See: Encrypted Key Exchange (EKE)

Electronic Codebook (ECB)

A mode of operation for block ciphers.

Overview

Block ciphers encrypt data in discrete chunks called blocks, usually 64 bits at a time. Electronic Codebook (ECB) is a mode of operation in which identical blocks of plaintext result in identical ciphertext. This is exactly how a traditional “code book” operates, from which the name of the mode derives.

The advantage of ECB is its speed, since the encryption algorithm can deal with each block of plaintext independently. The disadvantage is that eavesdropping of such encrypted communications can provide attackers with information that can be used to crack the ciphertext. This is especially the case if a man-in-the-middle attack can be mounted in which the attacker can submit arbitrary plaintext and examine the result.

Notes

The other main mode of operation for block ciphers is called cipher block chaining (CBC), in which each block of ciphertext is XORed with the previous block, resulting in different ciphertext for identical blocks of plaintext.

See Also: *block cipher, cipher block chaining (CBC)*

Electronic Privacy Information Center (EPIC)

A public interest organization focusing on civil liberties and privacy issues.

Overview

The Electronic Privacy Information Center (EPIC) was established in 1994 to focus public attention on issues relating to privacy in an information age. The organization has been involved in bringing attention to issues such as the sale of commercial data, privacy of medical records, establishment of national ID cards, development and use of the Clipper Chip, and many other issues. EPIC publishes a newsletter called *EPIC Alert*, which highlights the issue of civil liberties in the information age, and has an online bookstore devoted to online freedom. EPIC also publishes online guides to

privacy resources and tools, including how to obtain encryption software and how to browse the World Wide Web and send e-mail anonymously. EPIC is based in Washington, D.C., and is associated with Privacy International in the United Kingdom and the Internet Privacy Coalition.

For More Information

Visit www.epic.org for more information.

See Also: *privacy*

Electronic Signatures in Global and National Commerce (E-SIGN) Act

A U.S. law governing the use of digital signatures in business and commerce.

Overview

The Electronic Signatures in Global and National Commerce (E-SIGN) Act is a U.S. federal standard that gives electronic (digital) signatures the same legal force as traditional handwritten signatures. The purpose of the act is to promote e-business and e-commerce by speeding and facilitating the signing of contracts, purchase orders, credit card transactions, and other processes essential to the operation of business. E-SIGN is also designed to make the U.S. economy more efficient and competitive in an increasingly online global economy.

The main provisions of the E-SIGN Act took effect on October 1, 2000. The full implementation of the act, however, depends on the standardization of digital signature technologies. Currently, many different technologies are used, and the convergence of these technologies is essential before electronic signatures completely replace handwritten ones in the business, industry, and finance sectors. Public perception may also slow the adoption of electronic signatures in some sectors, with “hard” signatures written in ink on paper engendering more confidence than the invisible string of zeros or ones that constitutes a digital signature.

Marketplace

A number of vendors provide tools and services for digitally signing documents and transactions, including Entrust, Omtool, VeriSign, and others.

See Also: *digital signature*

elevation of privileges (EoP)

Method used by attackers to gain control of a system or network.

Overview

Elevation of privileges (EoP) refers to any approach whereby an attacker tries to fake or obtain credentials that provide broad access to system resources. Typically, this involves the attacker first gaining access to a low-privilege account such as the Guest account and then using a variety of methods to try to obtain a password for an Administrator account. These methods might include the following:

- Running a password-cracking program against a user account database
- Searching through registry keys for password information
- Reading e-mail and other documents in search of information about Administrator credentials
- Installing a Trojan to try and capture the credentials of a user with Administrator credentials
- Taking advantage of a bug in an application or network service to raise the privileges of a low-level account to Administrator level

Once the attacker has gained Administrator credentials, the attacker has access to virtually any resource or process and the system can be considered compromised. Common tools used by malicious hackers for cracking password files include Lsadbump2, LC3, Pwdump2, and John the Ripper.

See Also: *hacking, password, password cracking*

El Gamal

The encryption algorithm that forms the basis of the digital signature algorithm (DSA).

Overview

El Gamal is an asymmetric, or public key, algorithm similar to Diffie-Hellman (DH) and RSA. It is not widely used because it is slower than RSA and requires random seeding, but it was used by the National Institute of Standards and Technology (NIST) as the basis for its digital signature standard (DSS), the standardized implementation of the digital signature algorithm (DSA).

El Gamal can be used for both encrypting and signing digital messages, and was the first unpatented algorithm available for these purposes.

See Also: *Diffie-Hellman (DH), encryption algorithm, RSA*

elliptic curve cryptography (ECC)

Cryptographic procedures based on elliptic curve mathematics.

Overview

Elliptic curves are based on equations of the form $y^2 = Ax^3 + Bx^2 + Cx + D$. Under certain conditions, some of the points on an elliptic curve can have integral values that form a finite Abelian group and can be used as the basis for cryptographic transforms. Using elliptic curve cryptography (ECC), analogs can be created for traditional cryptographic algorithms such as Diffie-Hellman (DH), El Gamal, and RSA.

The advantage of the ECC approach, however, is that public key encryption schemes using it can have smaller keys and better performance than equivalent traditional cryptosystems. For example, an ECC system with a key size of 160 bits is approximately equivalent in security to a 1024-bit RSA system.

The elliptic curve digital signature algorithm (ECDSA) was one of the first commercial ECC systems and is

similar to the digital signature algorithm (DSA) used by the U.S. government for secure communication.

Marketplace

Commercial ECC implementations are available from a number of vendors, including Certicom and RSA Security. One place where ECC has found a niche is in smart card technology for which the smaller key size for digital signatures improves performance and reduces cost because of the limited processing power and memory of the card.

See Also: *cryptography, Diffie-Hellman (DH), Digital Signature Algorithm (DSA), El Gamal, Elliptic Curve Digital Signature Algorithm (ECDSA), RSA*

Elliptic Curve Digital Signature Algorithm (ECDSA)

An alternative to the Digital Signature Algorithm (DSA) based on elliptic curve cryptography (ECC).

Overview

DSA forms the basis of the Digital Signature Standard (DSS), a U.S. government standard for digital signatures published in 1994 as FIPS 186. The Elliptic Curve Digital Signature Algorithm (ECDSA) was proposed about the same time as DSA and has since been accepted and standardized at national and international levels as FIPS 186-2, ANSI X9.62, IEEE 1363-2000, and ISO 14888-3. The main advantage of ECDSA over other common encryption algorithms such as DSA and RSA is that shorter key lengths can be used, thus reducing the computation time for encryption and decryption of data without compromising security.

ECDSA has been proposed as an encryption algorithm for the emerging Extensible Markup Language (XML) digital signature standard XMLDSIG to provide message authentication and integrity for XML documents.

See Also: *Digital Signature Algorithm (DSA), elliptic curve cryptography (ECC)*

Encapsulating Security Payload (ESP)

An Internet Protocol Security (IPSec) security protocol that provides encryption.

Overview

Encapsulating Security Payload (ESP) is a security protocol defined by RFC 1827 that provides data integrity and confidentiality for IPSec. ESP can use 56-bit Data Encryption Standard (DES) or 168-bit Triple DES (3DES) to encrypt the information in an Internet Protocol (IP) packet. ESP does not provide authentication; however, this is provided by the other IPSec protocol Authentication Header (AH).

ESP can operate in one of two modes:

- **Tunnel mode:** Used primarily for secure communications between gateways and encrypts both the sender's IP address and the IP payload
- **Transport mode:** Used mainly for host-to-host virtual private networks (VPNs) and encrypts only the IP payload

See Also: *Internet Protocol Security (IPSec)*

Encrypted Key Exchange (EKE)

A method of sharing a secret message between two parties that involves using a short password as the primary key.

Overview

Most key exchange protocols use a long key to encrypt messages to guard against brute-force attacks. Encrypted Key Exchange (EKE) instead uses a short password coupled with a secret public key (ordinarily public keys are not secret). EKE uses this short password to encrypt the public key, which is used to keep the message secret. Because a short password can be susceptible to brute-force attacks, the algorithm produces well formed, but incorrect, public keys if the password is wrong. Thus, if an attacker tries to use a brute-force attack to crack the system, the attacker gets

a public key that looks correct but is not. Then, the attacker has to crack the public key, but then it often turns out to have been the wrong key anyway, so the attacker has to start over.

EKE is often used in conjunction with Diffie-Hellman (DH) to increase the amount of work involved in cracking a key exchange. The primary goal of EKE is to eliminate the weaknesses of brute-force attacks in the key exchange.

See Also: *brute-force attack, Diffie-Hellman (DH), key exchange, public key cryptography*

Encrypting File System (EFS)

A Microsoft technology for protecting files stored on a hard drive.

Overview

Encrypting File System (EFS) was first included with Microsoft Windows 2000 platform and provides a transparent way for users to store and read encrypted information on disk drives. EFS is built into version 3 or higher of the NTFS file system (NTFS) and is based on two industry-standard encryption algorithms: DESX and RSA.

Users can encrypt data on NTFS volumes several ways:

- By setting the encryption attribute for a file or folder using its properties sheet
- By creating, moving, or copying a file to a folder whose encryption attribute is set
- Using the cipher utility from the command line

To read an encrypted file, the user simply opens it using the appropriate application—EFS is built into the operating system kernel and automatically decrypts the file when needed.

See Also: *DESX, RSA*

encryption

Process of converting plaintext into ciphertext.

Overview

Encryption refers to any process that can convert readable data into secret code to prevent unauthorized users from reading the encrypted information, especially today in electronic transmission such as Web transactions, e-mail, and wireless networking. Unencrypted information is referred to as **plaintext**, while encrypted data is called **ciphertext**. An **encryption algorithm** is a mathematical procedure for converting plaintext into ciphertext. This is usually done using a numerical entity called a **key**, although other approaches to encryption exist that are not key based. In key-based encryption schemes, the strength of the scheme (degree of difficulty cracking encrypted messages) increases with the length (number of bits) of the key. Most encryption algorithms are also reversible to allow ciphertext to be converted back into plaintext, a process called **decryption**. An exception to this is hash functions, which are usually one-way encryption procedures.

There are two main approaches to encryption:

- **Symmetric encryption:** Also called secret key or private key encryption, a process in which both the sender and the recipient of an encrypted message use the same the shared secret (a secret key) to encrypt and decrypt the transmission or message
- **Asymmetric encryption:** Also called public key encryption, a process in which the sender and recipient use a pair of different but mathematically related keys, one for encrypting and the other for decrypting the transmission or message.

See Also: *asymmetric key algorithm, ciphertext, cryptography, encryption algorithm, hashing algorithm, key, plaintext, public key cryptography, symmetric key algorithm*

encryption algorithm

A mathematical procedure for converting plaintext into ciphertext.

Overview

Encryption algorithms form the basis of modern cryptographic systems and are mathematical procedures that

scramble information to make it unreadable to unauthorized users (and sometimes even to the owner of the information itself). Encryption has become the foundation of securing networks and communications systems in the information age.

Modern encryption algorithms trace their origin to the pioneering work of IBM in the 1960s. Out of this research came the first official U.S. government encryption standard, the Data Encryption Standard (DES), which employed the Data Encryption Algorithm (DEA). DES is a symmetric encryption algorithm in which a shared secret (a secret key) must first be securely delivered to all parties before these parties can engage in encrypted communications. The strength of an encryption algorithm generally increases with the length of the key, and while DES with its 56-bit key was considered secure for many years, it was eventually cracked and a replacement, described later, was devised. A variant of DES called 3DES or Triple DES applies DEA three times in succession, thus providing an effective key length of 112 bits, making it considerably more secure than DES (but unfortunately much slower as well).

A significant advance in encryption science occurred in 1976 when Whitfield Diffie and Martin Hellman published their paper “New Directions in Cryptography,” which outlined a scheme for public key cryptography. This advance was noteworthy because it provided a way of sidestepping the problem of providing parties with shared secrets in advance of performing encrypted communications. The first practical public key system was outlined the following year by Ron Rivest, Adi Shamir, and Leonard Adleman, who called the algorithm RSA after their last-name initials.

Other encryption algorithms include the International Data Encryption Algorithm (IDEA) developed by James Massey and Xuejia Lai in 1990, Pretty Good Privacy (PGP) developed by Phil Zimmerman in 1991, and Blowfish developed in 1993 by Bruce Schneier. The most recent and perhaps most important encryption algorithm is Rijndael, developed in 2001 by Belgian cryptographers Vincent Rijmen and Joan Daemen. Rijndael forms the basis of the new Advanced Encryption Standard (AES), which the U.S. government offi-

cially adopted in May 2002 as a replacement for the aging and no longer secure DES algorithm.

Marketplace

Many governments, including that of the United States, have export controls on strong encryption technologies. While there was a trend in the late 1990s to relax encryption export controls, international agreements such as the Wassenaar Arrangement have actually increased such restrictions. Before you implement strong encryption into a product targeted for foreign markets, be sure you are aware of current laws regarding export of encryption technologies. Note also that some encryption algorithms are patented and require additional permissions to use or implement in software or devices.

See Also: 3DES, Advanced Encryption Standard (AES), Blowfish, Data Encryption Algorithm (DEA), Data Encryption Standard (DES), International Data Encryption Algorithm (IDEA), Pretty Good Privacy (PGP), Rijndael, RSA

end-to-end encryption

Encrypted communications from sender to recipient.

Overview

End-to-end encryption protects transmissions from the time they originate at one host until the time they are received at another host. This protects the transmission from eavesdropping during the entire time of transit between the two hosts and at all intermediate transit points such as switches, routers, message queues, and disk-based storage. End-to-end encryption is the most secure way to transmit sensitive information across a network or communication system.

An example of how this can be implemented is Internet Protocol Security (IPSec), an extension of Internet Protocol (IP) that adds certificate-based encryption of data payload during transit between sending and receiving hosts. In a typical remote-access scenario, a connection is first negotiated and then IPSec is used to encrypt all data sent between the client and RAS server. IPSec can also be employed together with tunneling protocols to

create a secure virtual private network (VPN) between two hosts over the Internet.

See Also: *encryption, Internet Protocol Security (IPSec), virtual private network (VPN)*

ENUM

A proposed technology for mapping telephone numbers to the Domain Name System (DNS).

Overview

Storing contact information for business or personal use is complicated by the fact that individuals have so many different technologies by which they can be reached, such as “snail” mail, telephone, fax, and e-mail. ENUM is an attempt to bring convergence to such contact information by using an individual’s standard E.164 telephone number as that person’s primary contact information. By mapping these numbers to the DNS naming system of the Internet, you could send an e-mail message to someone by specifying the recipient’s telephone number instead of e-mail address.

ENUM works by using a special reverse DNS domain called e164.arpa that is used to store records for E.164 international telephone numbers. For example, the DNS name for someone whose telephone number is +44-6-2368572 would be 2.7.5.8.6.3.2.6.4.4.e164.arpa, constructed by reversing the digits and appending the e164.arpa domain name. A Naming Authority Pointer (NAPTR) record is then used to identify the services supported by this DNS name, such as telephone, e-mail, or fax. The NAPTR record effectively converts the E.164 telephone number into a Uniform Resource Identifier (URI).

Issues

ENUM promises to simplify Voice over IP (VoIP) communications by making it simpler to route calls over the Internet. However, the proposed global public database of ENUM contact information is seen by some industry watchers as a danger to privacy and a potential tool for spammers and mass marketers.

Notes

ENUM is described in RFC 2916.

For More Information

For an explanation of how DNS and VoIP work, see the *Microsoft Encyclopedia of Networking, Second Edition*, available from Microsoft Press.

See Also: *spam*

enumeration

Gathering information about a target system or network a hacker wants to compromise.

Overview

Enumeration is a collection of methods and procedures used by malicious hackers for gathering information that might be useful for launching an attack. Enumeration seeks to reveal poorly protected network resources that can be exploited for breaking into networks. Examples of such resources can include the following:

- Default user accounts that have no passwords
- Guest accounts that should normally be disabled
- Network services that are running but not needed

There are a variety of methods and approaches attackers use for enumerating systems and networks. One common method is to use port scanners to connect to standard Transmission Control Protocol (TCP) ports such as port 80 (Hypertext Transfer Protocol, HTTP) and send random data to the port to see what returns. If a Web server is listening on this port, it will usually respond with information identifying the vendor and version number. The attacker can then try compromising the server using known vulnerabilities of that version of the product, hoping that busy administrators have not had time to keep patches on the system up to date.

Some of the tools commonly used for enumeration include Netcat, Rcpdump, Dumpsec, Getmac, and many others.

See Also: *hacking, Netcat, port scanning*

EoP

Stands for elevation of privileges, any method used by attackers to gain control of a system or network.

See: elevation of privileges (EoP)

EPIC

Stands for Electronic Privacy Information Center, a public interest organization focusing on civil liberties and privacy issues.

See: Electronic Privacy Information Center (EPIC)

E-SIGN Act

Stands for Electronic Signatures in Global and National Commerce Act, a U.S. law governing the use of digital signatures in business and commerce.

See: Electronic Signatures in Global and National Commerce (E-SIGN) Act

ESP

Stands for Encapsulating Security Payload, an Internet Protocol Security (IPSec) protocol that provides encryption.

See: Encapsulating Security Payload (ESP)

/etc/passwd

A file used in most UNIX and Linux systems for storing user information.

Overview

The `/etc/passwd` file is a text file that typically contains the following information for each user on the system:

- The user's login name
- An encrypted version of the user's password
- A unique numerical ID (uid) for the user
- A numerical group ID (gid) for the user
- A comment field that can contain information such as the user's real name and address

- The location of the user's home directory
- The user's preferred shell

Implementation

As an example, the entry for user Denise Smith in `/etc/passwd` might be

```
dsmith:y29rf8er755:641:641:Northwind Traders:
home/dsmith:/bin/bash
```

The `etc/passwd` file is readable by all users, and even though passwords are stored in the file in encrypted form, this can constitute a security problem. One solution is to store only basic user information in `etc/passwd` and keep all passwords for users in a separate file called `etc/security/passwd`. Another solution is to implement shadow passwords, which store users' passwords in `/etc/shadow`, a file that can only be read by root. If shadow passwords are used, the preceding user's entry in `etc/passwd` usually looks like this:

```
dsmith:x:641:641:Northwind Traders:/home/
dsmith:/bin/bash
```

where `x` replaces the encrypted password and indicates that shadow passwords are being used.

See Also: password, shadow password

Ethereal

A free network protocol analyzer for UNIX and Microsoft Windows operating systems.

Overview

Ethereal is a free network “sniffer” created by Gerald Combs that allows you to capture and analyze traffic on a network. It works with a variety of data-link-layer protocols, including Ethernet, Token Ring, Fiber Distributed Data Interface (FDDI), Point-to-Point Protocol (PPP), and Classical IP over ATM. Display filters can highlight different types of packets in different colors, and captured data can be saved in plaintext or Post-Script format for further analysis and reporting.

Ethereal was released under the GNU General Public License and is freely available as open source software.

For More Information

You can download Ethereal from www.ethereal.com.

See Also: *sniffing*

European Institute of Computer Anti-Virus Research (EICAR)

An antivirus research organization.

Overview

The European Institute of Computer Anti-Virus Research (EICAR) is a consortium of industry, academic, technical, and legal experts united to prevent the proliferation of viruses, Trojans, and other malicious forms of code. EICAR also coordinates efforts to prevent computer crime including fraud, misuse of computers and networks, and other practices. EICAR also promotes a code of conduct that takes computer security issues seriously and prohibits the publishing of information that could be used to create viruses or other malicious code.

Since 1991, EICAR has hosted an annual conference in Europe designed to educate and inform IT (information technology) professionals about problems and solutions relating to computer viruses. EICAR also provides a downloadable antivirus test file that simulates a virus without doing damage and which can be used to test the effectiveness of antivirus products.

For More Information

Visit www.eicar.org for more information.

See Also: *virus, virus protection software*

event logs

Logs that record certain types of system information on Microsoft Windows platforms.

Overview

Microsoft Windows NT and later versions of the operating system support the logging of events, which are significant occurrences of operating system or application behavior. By default, systems running Microsoft Windows maintain three event logs:

- **System log:** Contains informational, warning, and critical events concerning the operation of operating system components including device drivers and network services
- **Application log:** Contains informational, warning, and critical events concerning the functioning of registered applications running on the system
- **Security log:** Contains Success and Failed audit events when auditing is enabled and configured on the system

Additional event logs may exist on systems that have network services such as Domain Name System (DNS) or Active Directory running. Microsoft Windows-based event logs are useful for several security-related reasons:

- Failure events in the Security log can indicate unsuccessful attempts by intruders to log on or access network resources.
- Success events in the Security log can be used to establish the identity of an intruder who has penetrated your system.
- Warning or Critical events in the System and Application logs can indicate components or applications infected by viruses or compromised by Trojans.

Notes

Manual analysis of event logs using Event Viewer can be tedious. *The Microsoft Windows 2000 Server Resource Kit* includes a tool called CyberSafe Log Analyst (CLA) that can be used to analyze event logs and generate reports of system activity.

exploit

Making use of vulnerabilities to compromise a network or system.

Overview

In hacker language, an exploit is an accomplishment that ends with successful intrusion into a network or system, acquiring an Administrator password and obtaining root access, installing a backdoor and erasing your tracks, executing arbitrary code to extract credit

card information from a database, defacing a public Web site, or just about anything an attacker would like to do. The term **exploit** is also used sometimes to refer to the tools and procedures by which the accomplishment is performed. By publishing information about such exploits, others are able to attempt and perform similar exploits.

Successful exploits usually depend on vulnerabilities in applications or operating systems. These vulnerabilities can include buffer overflows, unpatched systems, misconfigured network services, requiring strong passwords, and so on. Defending against exploits involves keeping systems and applications up to date with patches and service packs, installing intrusion detection systems (IDSs) to detect attacks when they occur, reviewing firewall logs, disabling unnecessary services, and other standard security measures.

See Also: *hacking, vulnerability*

exposure

Degree of protectedness in connection with a public network such as the Internet.

Overview

The Internet is a dangerous place nowadays, yet businesses have to connect to the Internet in order to communicate with suppliers, partners, and customers. A network that is directly connected to the Internet using a dedicated T1 or digital subscriber line (DSL) connection is highly exposed and at risk to attack. Using Network Address Translation (NAT) reduces exposure by hiding the company's IP address block from the outside world, and adding a properly configured firewall at the point where the network joins the Internet reduces exposure even further.

Exposure to threat can also be reduced in other ways. By disabling unnecessary network services and applying hotfixes or patches when vendors release them for your applications, your exposure is even further reduced. Employee training is necessary to prevent

users from opening dangerous attachments, infecting their systems with viruses, or giving out their passwords in response to social engineering attacks.

Yet, despite all these measures, some degree of residual exposure always remains, resulting from undiscovered coding errors, employee disregard of security procedures, new viruses for which signatures don't yet exist, old hardware still running that no one knows about, and many other sources. Good business sense views network exposure not as an absolute to be avoided but as a risk to be managed based on a cost/value equation.

See Also: *firewall, hotfix, patch, social engineering*

Extensible Authentication Protocol (EAP)

A security extension for the Point-to-Point Protocol (PPP).

Overview

The Extensible Authentication Protocol (EAP) extends PPP, an industry-standard wide area network (WAN) protocol, by providing support for additional authentication methods. Using EAP, a PPP session can authenticate using one-time passwords, token cards, smart cards, Kerberos, Public Key Infrastructure (PKI) certificates, and other methods. EAP provides an open architecture for incorporating virtually any authentication scheme to secure PPP sessions and is an important contributor to the rise in popularity of virtual private network (VPN) technologies. EAP is also used in wireless local area network (WLAN) technologies, where requests by clients for authentication are forwarded by access points to a Remote Authentication Dial-In User Service (RADIUS) server.

EAP is defined in RFC 2284.

See Also: *authentication, virtual private network (VPN)*

Extensible Authentication Protocol–Transport Layer Security (EAP-TLS)

An encrypted authentication scheme based on Extensible Authentication Protocol (EAP).

Overview

Extensible Authentication Protocol–Transport Layer Security (EAP-TLS) is a certificate-based authentication system for WAN (wide area network) and wireless local area network (WLAN) connections that combines the use of EAP for session negotiation and Transport Layer Security (TLS) for encrypted transmission of credentials.

EAP-TLS provides mutual authentication in which both the client and server require authentication from each other. This provides clients with confidence regarding the identity of the server they are trying to establish a connection with. A drawback of EAP-TLS is that certificates are required on both the server and the client sides, and managing these certificates adds extra overhead for network managers.

EAP-TLS is supported by Microsoft Windows 2000 as an authentication method for virtual private network (VPN) connections and is the standard 802.1x wireless security protocol used in Microsoft Windows XP.

See Also: *Extensible Authentication Protocol (EAP), Transport Layer Security (TLS)*

Extensible Authentication Protocol–Tunneled Transport Layer Security (EAP-TTLS)

An encrypted authentication scheme based on Extensible Authentication Protocol (EAP) that is easier to manage than Extensible Authentication Protocol–Transport Layer Security (EAP-TLS).

Overview

Extensible Authentication Protocol–Tunneled Transport Layer Security (EAP-TTLS) is used in wide area network (WAN) and wireless local area network (WLAN) networking to provide mutual, certificate-based authentication of both client and server. EAP-TTLS improves on EAP-TLS by requiring a certificate only on the server side and allowing clients to be authenticated using their credentials instead. Users' passwords are protected from eavesdropping by encrypting them using Transport Layer Security (TLS), an Internet Engineering Task Force (IETF) standardized version of Secure Sockets Layer (SSL) encryption developed by Netscape. Authenticating users using their credentials eliminates the need for client certificates, making EAP-TTLS considerably easier to manage than EAP-TLS in enterprise environments. EAP-TTLS is used in conjunction with 802.1x to provide strong security over wireless links and to simplify management of WLAN security.

EAP-TTLS was developed by Certicom and Funk Software and has been submitted to the IETF as a proposed Internet standard.

Marketplace

A number of WLAN vendors have endorsed EAP-TTLS, including Avaya, Enterasys, Intermec Technologies, and Proxim.

See Also: *Extensible Authentication Protocol (EAP), Extensible Authentication Protocol–Transport Layer Security (EAP-TLS), Transport Layer Security (TLS), tunneling*

Fair Information Practices (FIP)

Standards governing collection and use of personal data.

Overview

Protection and privacy of personal information is becoming increasingly important as e-commerce grows on the Internet. The concept of Fair Information Practices (FIP) can be traced back to the Privacy Act of 1974, U.S. legislation designed to protect personal information collected by government agencies. The Organization for Economic Cooperation and Development in Europe incorporated these practices into its Guidelines for the Protection of Personal Data and Transborder Data Flows in 1980, which evolved into the European Union Data Protection Directive in 1995.

FIP can be summarized in five basic principles:

- **Notice:** An agency collecting personal information from individuals must inform these individuals concerning its collection and use practices.
- **Choice:** Individuals must be able to determine how collected information should be used.
- **Access:** Individuals must be able to view, modify, and contest the accuracy of personal information collected about them.
- **Security:** Agencies collecting personal information must protect such information from unauthorized access.
- **Enforcement:** There should be legal mechanisms in place to enforce these practices to ensure their compliance.

Other important principles include these:

- **Data integrity:** Agencies collecting personal information must maintain the integrity of the data collected.

- **Onward transfer:** An agency collecting information from individuals must inform these individuals concerning its policies for passing such information on to other agencies.
- **Remedy:** Individuals must have avenues of remedy available should they determine that an agency holding personal information about them has misused this information or allowed it to be misused.

For More Information

The 1998 report “Privacy Online: A Report to Congress” by the Federal Trade Commission outlines the issues and practices surrounding FIP. You can download this report from www.ftc.gov/reports/privacy2000/privacy2000.pdf in PDF format.

See Also: *privacy*

false negative

Reporting of malicious events as benign by a security system.

Overview

False negatives occur when a firewall, intrusion detection system (IDS), or other network security device identifies a malicious event as benign. False negatives are therefore failures of these security systems to properly identify attempts to penetrate network defenses. They may be caused by misconfiguration of the security system or basic flaws in its design. Note that a malicious event resulting from a new form of exploit and ignored by a security system is not considered a flaw in the system, for no security system can completely defend against exploits that have not yet been conceived. (Heuristic methods try to anticipate new attacks but usually generate large numbers of false positives.)

False negatives can have catastrophic effects for the network the security device is protecting. Penetration of

the network's defenses can result in loss or theft of data and compromised systems being used for illicit purposes, such as launching a distributed denial of service (DDoS) attack against another network, with resulting legal liability for the compromised network. In general, it is better to tune a security system to eliminate false negatives rather than false positives, for while false positives require extra work for administrators to analyze, at least their network is protected against intrusion.

See Also: *false positive, firewall, intrusion detection system (IDS)*

false positive

Reporting of benign events as malicious by a security system.

Overview

False positives are certain types of events generated by firewalls, intrusion detection systems (IDSs), and other network security devices. False positives are generated when the system triggers because of traffic that appears dangerous but actually isn't. These may be triggered because the sensitivity of the security system is set too high or because of basic flaws in the design of the system.

False positives are undesirable since they increase the workload of administrators, who have to analyze them to distinguish them from genuine intrusion attempts. This drains resources and can increase the cost of maintaining the security system. By properly tuning a firewall or IDS, the proportion of false positives can usually be reduced to acceptable levels, and intelligent systems can also be programmed to learn how to distinguish false positives from genuine events.

False positives are less of a problem than false negatives, however, which indicate that a security system is not doing its job.

See Also: *false negative, firewall, intrusion detection system (IDS)*

fast packet keying

An enhancement for the RC4 algorithm used by Wired Equivalent Privacy (WEP).

Overview

WEP is a security protocol for protecting 802.11b wireless local area networks (WLANs) from eavesdropping. WEP suffers from weaknesses, however, that result from how RC4 encryption is implemented, and fast packet keying is one method of solving this problem. While standard WEP implementations use a unique RC4 secret key for each communication session, fast packet keying uses a unique key for each data packet that is transmitted, making it much harder to eavesdrop on wireless communications. Fast packet keying was proposed by two companies, RSA Security and Hifn, as a solution that can be implemented through firmware upgrades of existing 802.11b products.

See Also: *Wired Equivalent Privacy (WEP)*

FedCIRC

Stands for Federal Computer Incident Response Center, a U.S. government agency dealing with computer security issues for federal government departments.

See: *Federal Computer Incident Response Center (FedCIRC)*

Federal Computer Incident Response Center (FedCIRC)

A U.S. government agency dealing with computer security issues for federal government departments.

Overview

The Federal Computer Incident Response Center (FedCIRC) combines the efforts of the U.S. Department of Defense (DoD), law enforcement agencies, intelligence communities, and academia to coordinate the analysis of network intrusion incidents. FedCIRC acts as a trusted focal point for reporting incidents and receiving assistance in prevention and response. FedCIRC maintains a knowledge base and publishes advisories regarding Internet security problems and has a mailing list whose membership is restricted to users in the .gov and .mil domains.

FedCIRC operates under the auspices of the General Services Administration (GSA) and coordinates with

other computer incident response agencies, including the CERT Coordination Center (CERT/CC) and the National Information Protection Center (NIPC).

For More Information

Visit FedCIRC at www.fedcirc.gov for more information.

See Also: CERT Coordination Center (CERT/CC)

Federal Information Processing Standard (FIPS)

A series of standards developed by the National Institute of Standards and Technology (NIST) for federal computer systems.

Overview

Federal Information Processing Standard (FIPS) publications are intended for use by U.S. government agencies and their contract partners, and they cover various aspects of information technology. FIPS standards are developed mainly in response to needs for interoperability and security when no industry standards exist. FIPS publications are developed using an open process that provides interested parties a chance to comment on proposals. Examples of well-known FIPS publications that have significant impact on the technology industry include the following:

- FIPS 46-3 Data Encryption Standard (DES)
- FIPS 161-2 Electronic Data Interchange (EDI)
- FIPS 180-1 Secure Hash Standard (SHS)
- FIPS 186-2 Digital Signature Standard (DSS)
- FIPS 197 Advanced Encryption Standard (AES)
- FIPS 198 Keyed-Hash Message Authentication Code (HMAC)

For More Information

Visit NIST online at www.itl.nist.gov/fipspubs/ for more information on FIPS publications.

See Also: National Institute of Standards and Technology (NIST)

Federal Information Technology Security Assessment Framework (FITSAF)

A methodology for assessing the security of information systems.

Overview

The Federal Information Technology Security Assessment Framework (FITSAF) was proposed in 2000 by the Chief Information Officers (CIO) Council, a federal agency that helps other U.S. government agencies modernize their information services. FITSAF is designed to help other government agencies assess the readiness of their information security programs by ensuring that proper policies, programs, and procedures are in place and have been properly implemented, documented, and tested. The U.S. Internal Revenue Service (IRS) has adopted FITSAF as the basis for evaluating its own information security program, and other government agencies will likely follow.

For More Information

Visit the CIO Council at www.cio.gov for more information.

file integrity checker

Software that protects systems against having their files modified or replaced.

Overview

When attackers compromise a system, they often try to replace key system files with versions infected with Trojans to install backdoors for later entry. File integrity checkers help defeat such attacks by detecting when system files are modified or replaced. Typically, this is done by calculating a checksum of system files immediately after the operating system is installed. This can be done using 32-bit cyclic redundancy check (CRC) or more securely using cryptographic hash algorithms such as MD5 or Secure Hash Algorithm-1 (SHA-1). By comparing the initial value calculation with a later one, changes to the file can be detected, including modification of file attributes, permissions, modification time, size, and so on.

Marketplace

There are a number of free and commercially available file system checkers available from different vendors. Some common ones include AIDE, FileChecker, fsum, L5, integrity, SP1, Tripwire, and yafic. LANguard File Integrity Checker from Gfi is a popular freeware utility that can alert administrators of Microsoft Windows NT and Microsoft Windows 2000 systems when files have been added, deleted, or modified on a system.

Microsoft also included a file-checking feature called File Signature Verification (FSV) in its Microsoft Windows platforms starting with Windows 2000. Many host-based intrusion detection systems (IDSs) include file system checkers, as do some host-based firewalls and antivirus products.

See Also: File Signature Verification (FSV), hashing algorithm, intrusion detection system (IDS), MD5, Secure Hash Algorithm-1 (SHA-1)

File Signature Verification (FSV)

A file-checking feature of Microsoft Windows File Protection (WFP).

Overview

File Signature Verification (FSV) can be used for verifying that files on Microsoft Windows platforms have not been modified. Starting with Windows 2000, Microsoft signed key system files using digital signatures to guarantee and protect the integrity of such files. Using FSV, you can verify that signed files have not been modified and are in fact the original, unaltered files installed during setup. You can also use FSV to scan your system for any unsigned files that might be present. This protects your system both against attackers who try to modify files to install Trojans on compromised systems and against buggy applications that accidentally try to overwrite important system files.

See Also: digital signature, file integrity checker, Trojan, Windows File Protection (WFP)

file slack

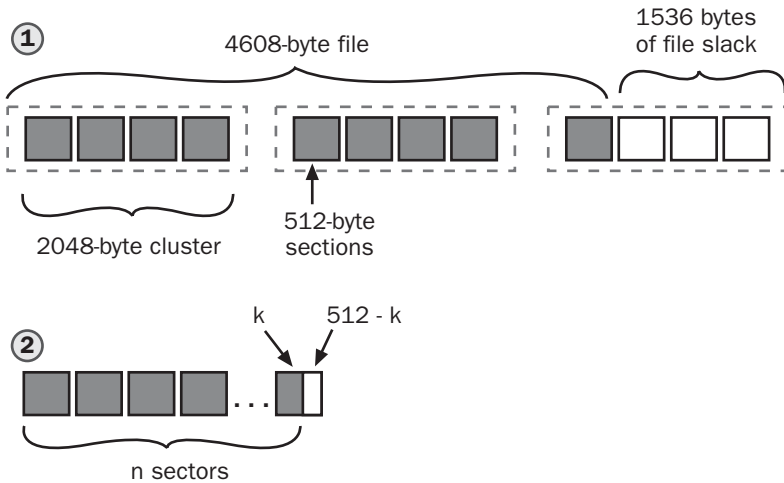
Unused space on hard disks that can hide important data.

Overview

An important task of computer forensics is to obtain evidence from examination of data stored on hard drives. What many do not realize is that simply deleting files from a hard drive does not necessarily prevent valuable information from being recovered. This is because deleting files simply deletes them from the file table rather than deleting the files themselves. To understand file slack and its related concept of RAM (random access memory) slack, you first have to understand how data is stored on disks.

Operating systems normally write digital information to disk drives in fixed-size blocks called sectors, which are typically 512 bytes in size. Actual files are written to larger blocks called clusters, which consist of one or more sectors and range from 512 bytes to 256 kilobytes (KB) and larger, depending on the size of the volume being formatted and the file system being used (NTFS or FAT). When a file is written to disk, an integral number of clusters is used, and the last cluster generally contains some unused space called **file slack**. This file slack typically contains data from an earlier file that used this cluster, and by examining this slack it may be possible to extract useful evidence in a forensic examination. On a large hard drive, there may even be gigabytes of file slack present containing bits and pieces of files previously deleted by the user, an amount of evidence that certainly is not trivial.

The last sector of the last cluster for a file can also contain another type of slack called RAM slack. This is because the operating system pads the data being written to the sector with data randomly taken from RAM in order to make it 512 bytes, the size of a hard disk sector. This RAM slack may be an even more important source of forensic evidence than file slack, for operating systems often maintain passwords and other valuable information in RAM and this information may find its way into RAM slack and thus be secretly persisted to disk.



File = $(512n + k)$ bytes

RAM slack = $(512 - k)$ bytes

File slack. Examples of file slack and RAM slack.

Marketplace

A variety of “disk-cleaning” programs are available that can overwrite the clusters and sectors of a hard drive to irretrievably erase all data from the disk. A popular commercial tool is Disk Wiper from Paragon Software Group. There are also many shareware disk-cleaning programs available such as DriveScrubber and System Shield.

Notes

There are other potential sources of forensic information that may be hidden on disks without users being aware of it. Examples include swap files such as the Microsoft Windows pagefile, spool files used in printing, temporary files created by applications and not deleted later, and deleted files in a recycle bin that has not been emptied.

See Also: computer forensics

file system traversal attack

A coding vulnerability allowing users to access files in parent directories.

Overview

File system traversal refers to the process of changing the current directory. For example, from the Microsoft Windows command prompt the command `cd ..\..` moves the current directory two levels up in the parent direction. A file system traversal attack involves inserting such syntax into command strings for Common Gateway Interface (CGI) or other Web applications, enabling attackers to access files in directories above the Web root directory. This may allow attackers to read files not intended for public use or even execute scripts or other applications and gain control over the server.

File system traversal attacks are a result of improper parsing in Web server code. They are easy to perform and require no special tools, making them a favorite for inexperienced hackers to perform. Vulnerabilities to such attacks were discovered in most Web server platforms and were later fixed. In its simplest form, this vulnerability is called the dot bug vulnerability.

See Also: dot bug vulnerability, vulnerability

filter

Any mechanism for removing unwanted data.

Overview

Many types of computer and network security systems implement filtering of some type. One example is the packet-filtering router, which forwards or blocks traffic based on Internet Protocol (IP) address and port information based on rules configured by administrators. Another common use of filters is in electronic messaging, where filters are used on mail servers and clients to block spam and other unwanted mail. Filters are also used in firewalls, intrusion detection systems (IDSs), Web caching servers, and in various other applications to protect systems and improve performance.

A typical filter (of any sort) usually consists of a series of rules. Each rule includes a condition and an action to be taken when the condition is fulfilled. For example, in an e-mail filter, the condition might be the presence of some objectionable word in the Subject line of an incoming message, and the action to be taken when the condition is fulfilled would be to discard such messages. Filters generally consist of a number of rules applied in order, but if there are too many rules, filters become hard to manage and their effect hard to predict.

See Also: *packet filtering, spam*

Finger

A command on UNIX platforms for obtaining information about users.

Overview

Finger is a command for obtaining information about users on a UNIX network. This information can include the user's full name, the user's default shell, and the last time the user logged in. To "finger" someone means to use the Finger command to display information about them. The syntax for fingering users is **finger user-id@domain**, where *domain* is the fully qualified domain name (FQDN) to which the user belongs.

For Finger to work, the underlying network must be running the Finger daemon (service), and since attackers might use this tool for footprinting a network they plan to attack, most organizations disable the Finger daemon on their UNIX networks nowadays, apart from a few academic institutions. There still exist a few "Finger gateways" on the Internet that can be used for fin-

gering systems, but these are rapidly losing their usefulness to attackers.

Notes

The Finger protocol is defined in RFC 1288.

See Also: *footprinting, whois lookup*

fingerprinting

Determining the identity of a remote system by analyzing packets it generates.

Overview

Fingerprinting is a technique used by attackers to determine product and version information about operating systems and applications running on remote systems. The technique is called **fingerprinting** because each platform or version number for a software product generally has its own specific ways of responding to different requests that uniquely identify it, similar to the way fingerprints are unique to each person. Once an attacker has "fingerprinted" a remote host and determined what operating system and version it runs, the attacker can consult a database of known vulnerabilities for that platform and launch an attack.

Fingerprinting can be either active or passive. In active fingerprinting, the attacker sends different kinds of packets to the target system and observes the result. In passive fingerprinting, the attacker analyzes normal traffic generated by the target system, for example, by intercepting e-mail messages and analyzing the headers. Some of the methods used for active fingerprinting of systems include the following:

- Sending valid requests to common ports (for example, Hypertext Transfer Protocol [HTTP] GET requests to port 80) and observing the result. Some Web servers respond to such requests by sending their product name and version number in the initial packets returned. This approach can also be used for other common protocols including File Transfer Protocol (FTP), Simple Mail Transfer Protocol (SMTP), and telnet.
- Sending invalid data to common ports and observing the results. The error messages returned by services are often more system- and version-specific

than normal responses to legitimate requests. One way of generating such data is to add special characters such as “~” or “*” to standard requests to try to exploit known vulnerabilities in certain applications and platforms. More complex methods involve creating invalid Internet Protocol (IP), Transmission Control Protocol (TCP), or Internet Control Message Protocol (ICMP) packets and analyzing how the target system responds to such packets.

- Use a port scanner such as Nmap to identify which ports are open on the target system and compare the results with a database of such information for different platforms and versions.

Notes

The term **fingerprinting** is sometimes called **stack fingerprinting**, referring to the TCP/IP protocol stack being probed by the attacker.

See Also: Nmap

FIP

Stands for Fair Information Practices, standards governing collection and use of personal data.

See: Fair Information Practices (FIP)

FIPS

Stands for Federal Information Processing Standard, a series of standards developed by the National Institute of Standards and Technology (NIST) for federal computer systems.

See: Federal Information Processing Standard (FIPS)

firewall

A device or application for protecting a network or host against hostile network traffic.

Overview

Firewalls monitor and control the flow of network traffic between two networks or between a host and its network. Firewalls are usually placed at the point at which a network connects to the Internet and act as a choke point for controlling what traffic can safely enter the

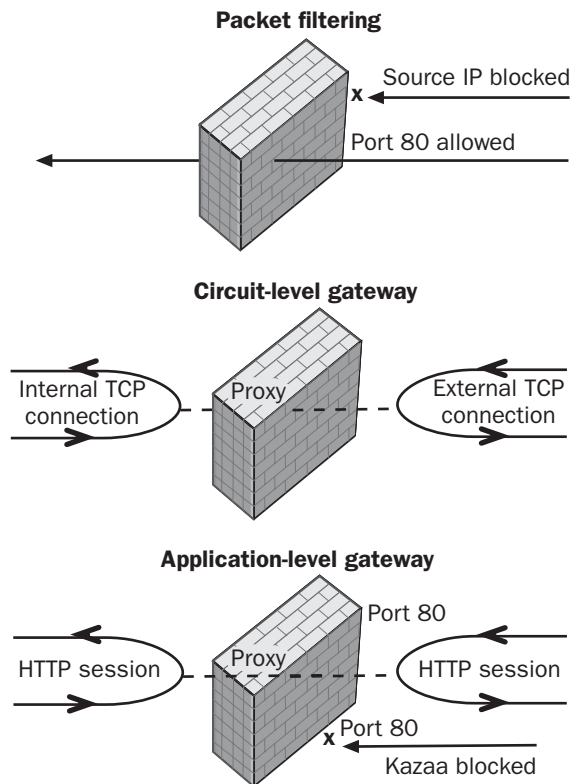
network. The firewall thus acts as a kind of gatekeeper for the network, controlling both what comes in and what goes out. For example, a firewall can be configured to allow Hypertext Transfer Protocol (HTTP) traffic on port 80 to pass freely between the network and the Internet while blocking all Telnet traffic on port 23.

There are several different types of firewalls from the perspective of how they operate, and these types overlap to a degree in commercial firewall products. The three basic types of firewalls are as follows:

- **Packet-filtering routers:** Also called screening routers, these firewalls are routers that can be configured with a series of rules to allow, reject, or drop inbound or outbound packets based on Internet Protocol (IP) address or port number. Packet-filtering routers operate at the network layer and combined with network address translation (NAT) can provide networks with a first level of defense. An advantage of packet-filtering routers is that they operate very fast and can process packets at line speed.
- **Circuit-level gateways:** These firewalls listen for Transmission Control Protocol (TCP) handshaking requests from external hosts and decide whether to accept or reject such requests based on port numbers. When the firewall accepts a handshaking request, a TCP session is established between the firewall and the remote host. The firewall then establishes a separate proxy session with the internal host that the remote host is trying to communicate with and then relays communication between the two sessions using an internal circuit connection it establishes within itself. Combining circuit-level screening with packet filtering provides a higher level of defense than packet filtering by itself.
- **Application-level gateway:** These are similar to circuit-level gateways but can also filter traffic based on the application layer protocol being used such as HTTP or File Transfer Protocol (FTP). While a circuit-level gateway might allow any protocol to establish a proxy TCP connection across port 80, an application-level gateway would only allow properly formatted HTTP traffic and block any other applications, such as a peer-to-peer (P2P)

file-sharing program, that might try to use that port. Application-level gateways can also log traffic, perform authentication, convert protocols, and perform other useful actions. They are more complex than other types of firewalls, however, since they require configuration for each application-layer protocol for which traffic will be allowed to pass. They are also very resource intensive and usually require special software or configuration by the user. A combination of application-level screening with packet filtering provides a high level of defense for networks.

Most firewalls incorporate a combination of all three methods described together with additional proprietary technologies developed by firewall vendors such as the stateful inspection technology developed by Check Point Software for its Firewall-1 product line.



Firewall. How the three basic types of firewalls work.

Marketplace

Firewalls come in all shapes and sizes today. Those acting as gatekeepers for large corporate networks are typically either commercial firewall applications such as Check Point's Firewall-1 installed on high-performance multihomed servers or advanced router applications such as Cisco's PIX Firewall. Host-based or desktop firewalls can provide companies with additional protection for their desktop computers against attack from inside their networks. An example of an integrated host-based firewall is the Internet Connection Firewall (ICF) component of Microsoft Windows XP and Microsoft Windows Server 2003. Commercial vendors of managed desktop firewalls include InfoExpress, Internet Security Systems, and Sybergen Networks.

Branch offices can secure their Internet connections using firewall appliances such as VelociRaptor from Symantec. Such appliances are simple to install and configure and easy to manage. For smaller Small Office/Home Office (SOHO) networks, firewalls are often built into digital subscriber line (DSL) or cable-modem routers that provide shared Internet access; common examples are products from Linksys, NetGear, and D-Link. Personal firewalls are applications that can be installed on individual computers to protect home or business users when connected to the Internet. Popular personal firewalls include BlackICE Defender from Internet Security Systems and ZoneAlarm Pro from Zone Labs.

See Also: demilitarized zone (DMZ), packet filtering

FIRST

Stands for Forum of Incident Response and Security Teams, an umbrella organization for computer security incident response centers around the world.

See: Forum of Incident Response and Security Teams (FIRST)

FITSAF

Stands for Federal Information Technology Security Assessment Framework, a methodology for assessing the security of information systems.

See: Federal Information Technology Security Assessment Framework (FITSAF)

footprinting

Method used by attackers to identify potential targets for attacking a network.

Overview

Footprinting is the first step performed in trying to hack into a network. Footprinting refers to the process of gathering as much information as possible about the network from publicly available sources. The goal is to create a map of the network to identify systems and applications that can be targeted for attack. Examples of ways an attacker might “footprint” a network include the following:

- Visiting the company’s Web site to look for publicly available information that might be useful
- Using search engines to try to find other useful information about the company such as anonymous File Transfer Protocol (FTP) sites and poorly secured intranet sites
- Using Whois at a domain registrar site to find out more about the company’s domain name and Internet Protocol (IP) address blocks
- Using Nslookup and other tools to try to perform zone transfers with Domain Name System (DNS) name servers
- Using Ping or Fping to test for the presence of hosts within the IP address block owned by the network
- Using Tracert to try to locate routers and map subnets for the target network
- Using Nmap to scan to identify operating system platforms and versions

Once an attacker has footprinted a network, the next step is usually enumeration of services running on the network to try to find vulnerable places to break in.

See Also: enumeration, Fping, hacking, Nmap, Nslookup, Ping, port scanning, Tracert, Whois lookup

FORTEZZA

Cryptographic technologies developed by the National Security Agency (NSA) for U.S. government use.

Overview

FORTEZZA was developed by the NSA as part of its Multi-Level Information Systems Security Initiative, and it defines a set of standard programming interfaces for cryptographic algorithms implemented in secure hardware devices. An example of a FORTEZZA device is the crypto card, a PCMCIA card that contains a cryptographic key in a special hardened case with built-in sensors to protect against tampering. Using this card, a government worker can securely log on to a restricted network such as a Defense Messaging System (DMS) to send and receive encrypted e-mail or communicate over an encrypted digital phone line.

FORTEZZA is supported by several Microsoft products including Microsoft Windows 2000, Microsoft Exchange 2000, and Microsoft Outlook 2000.

Notes

The name FORTEZZA is derived from an Italian word meaning “fortress” and indicates the secure nature of the technology.

See Also: cryptography

Forum of Incident Response and Security Teams (FIRST)

An umbrella organization for computer security incident response centers.

Overview

The Forum of Incident Response Security Teams (FIRST) is a global forum for coordinating the activities of incident response organizations from industry, government, defense, and academia. FIRST was founded in

1990 and has grown to include over 100 members from around the world including such well-known incident response centers as the following:

- **CERT/CC:** CERT Coordination Center
- **AusCERT:** Australian Computer Emergency Response Team
- **DFN-CERT:** German Computer Incident Response Team
- **JPCERT/CC:** Japan CERT Coordination Center
- **DOD-CERT:** U.S. Department of Defense CERT

The role of FIRST is to provide a trusted forum for computer security incident response teams to share information with each other. FIRST also hosts an annual conference on Computer Security Incident Handling and facilitates technical colloquiums on vulnerabilities, incidents, tools, and procedures. FIRST membership is open to any organization responsible for handling security incidents on condition of sponsorship from an existing FIRST member.

For More Information

Visit FIRST at www.first.org for more information.

See Also: *CERT Coordination Center (CERT/CC)*

Fping

A tool for testing network connectivity with hosts.

Overview

Fping is short for “fast ping” and is a command-line tool for pinging hosts to see if they are present and running on a Transmission Control Protocol/Internet Protocol (TCP/IP) network. Unlike Ping, which can only be used to ping one host at a time, Fping can be used to repeatedly ping a series of hosts specified in an associated text file. It can also be used to ping a series of hosts derived from a specified netmask or range of Internet Protocol (IP) addresses. Because of this enhanced functionality, Fping can be used in scripts for automatic pinging of networks in round-robin fashion, and the output can be parsed and analyzed to gain information about the target network.

Fping was developed by Roland Schemers of Stanford University and is freely available for use. Security professionals (“white hats”) can use this tool for monitoring their networks, while malicious hackers (“black hats”) can use it for footprinting networks targeted for attack.

For More Information

Visit www.fping.com to download the program and instructions.

See Also: *footprinting, hacking, Ping*

Fpipe

A tool for port redirection.

Overview

Fpipe is a free tool developed by Foundstone that can be used to create custom streams of Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) packets using source ports you specify. Using Fpipe, you could bypass the security of a firewall by sending any traffic you like through any open port on the firewall. Fpipe works by indirection and waits for a client to connect to its listening port, after which it establishes a TCP or UDP data stream with the target host inside the firewall. Fpipe can be installed either on the client host itself or can reside on a third-party host outside the network.

For More Information

Visit Foundstone online at www.foundstone.com and download Fpipe and other free security tools.

See Also: *firewall*

Fport

A tool for displaying which services are listening on a network.

Overview

Fport is a free tool developed by Foundstone that detects which Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) ports are listening on a target system or network. While other tools such as the Microsoft Windows Netstat command can be used for similar purposes, Fport also maps listening ports to their respective applications and can be used to determine

which network services are available on the target system. By displaying the open ports and listening applications, Fport can be used to help detect Trojans or backdoors installed on compromised systems. For example, if you run Fport and find such high-numbered ports as 31337 or 65000 open, it may indicate the presence of a Trojan listening on that port. Security professionals can also use this tool to verify that systems have been properly locked down by disabling unnecessary services and for baselining open port numbers on newly installed systems.

For More Information

Visit Foundstone online at www.foundstone.com to download Fport and other free security tools.

See Also: *Netstat, Trojan*

fragmentation

Breaking Internet Protocol (IP) packets into smaller parts.

Overview

Fragmentation is necessary to enable IP packets to traverse boundaries between media supporting different maximum packet sizes. Fragmentation can also be used, however, as a method for hiding an attack on a network by avoiding detection by intrusion detection systems (IDSs). An example of a tool that can be used for such purposes is Fragrouter, which takes a stream of IP packets and fragments it using various schemes for evading detection. Fragrouter also has usefulness to security professionals as a tool for testing IDSs to see whether they are able to detect and block such fragmented traffic. Fragrouter is freely available for Linux, FreeBSD, Solaris, and other UNIX platforms.

Fragmented data streams can also be used for denial of service (DoS) attacks on a variety of platforms, including Microsoft Windows NT and Cisco IOS. Keeping your operating system up to date with the latest security patches usually prevents your networks from being vulnerable to such attacks.

For More Information

Visit www.securityfocus.com to download Fragrouter.

See Also: *intrusion detection system (IDS)*

FSV

Stands for File Signature Verification, a file-checking feature of Microsoft Windows File Protection (WFP).

See: *File Signature Verification (FSV)*

FTP bounce attack

An attack that exploits a design flaw in File Transfer Protocol (FTP).

Overview

When an FTP client establishes a connection with an FTP server, the client usually sends a PORT command to tell the server which port number the client will use for its data channel connection to port 20 on the server. Since the FTP standard allows for any Internet Protocol (IP) address and port number to be used as a destination by the PORT command, an attacker could maliciously establish a data connection with an FTP server to circumvent a firewall or port scan a network without being detected.

To circumvent this problem, vendors generally modify their FTP server programs with nonstandard mechanisms for preventing bounce attacks from occurring. One common approach is to modify the PORT command so that it can send only the IP address of the client that has previously established the control channel connection to port 21 on the server. This prevents connections with arbitrary machines from being forced on the server by an attacking client.

For More Information

For a list of well-known port numbers and what they're used for, see the *Microsoft Encyclopedia of Networking, Second Edition*, available from Microsoft Press.

See Also: *port scanning*

GetAdmin

A well-known cracking tool for Microsoft Windows NT.

Overview

GetAdmin is an elevation-of-privileges tool used by attackers to obtain administrator access to Windows NT 4 systems. GetAdmin works by exploiting a flaw in the operating system kernel that failed to check the output address of a system call. In order to use GetAdmin, the attacker must first gain local access to the machine and log on using an ordinary user account. Once logged on, the attacker can run GetAdmin to grant administrator privileges to its account.

Although Microsoft Corporation issued a patch for GetAdmin in Service Pack 4 for Windows NT, the existence of this exploit and its widespread popularity highlight two important aspects of network security that every systems administrator should pay attention to:

- Systems that aren't physically secure are potentially vulnerable to serious exploits of this nature. If a system is locked away in a back room somewhere, an attacker cannot log on locally and the system is therefore immune to the GetAdmin exploit. Physical security is thus an essential part of any network security policy.
- Unpatched systems are more vulnerable to security breaches than those whose patches are up to date. Administrators who failed to apply Service Pack 4 or later to their Windows NT systems were leaving them open to exploits of this nature, and proper security measures include prompt application of patches issued by vendors.

GetAdmin was developed by Russian programmer Konstantin Sobolev.

See Also: *cracking, elevation of privileges (EoP), exploit*

GIAC

Stands for Global Information Assurance Certification, a certification program for computer security professionals developed by the SANS Institute.

See: *Global Information Assurance Certification (GIAC)*

Global Information Assurance Certification (GIAC)

A certification program for computer security professionals developed by the SANS Institute.

Overview

Global Information Assurance Certification (GIAC) is an independent certification program designed to validate the knowledge and experience of practitioners in different areas of system and network security. GIAC certifications cover a wide range of topics, including intrusion detection and analysis, firewalls, incident and response handling, auditing, and forensics. Certifications are also offered for Microsoft Windows and UNIX security administrators to independently validate expertise on these platforms. The GIAC Security Engineer (GSE) is a group of certifications for individuals demonstrating mastery in a wide range of security areas.

GIAC certifications have two components: a certification exam and a written assignment demonstrating practical experience with security issues, tools, and procedures. SANS requires that GIAC-certified individuals recertify every few years to ensure competency in the latest security standards and practices. GIAC has been widely recognized in the security community since its inception in 1999 as a valuable tool for ensuring that security professionals meet minimum standards of technical competency.

For more Information

Visit GIAC online at www.giac.org for more information.

See Also: *Certified Information Systems Security Professional (CISSP), SANS Institute*

GnuPG

Stands for GNU Privacy Guard, an open source tool implementing the OpenPGP encryption standard.

See: *GNU Privacy Guard (GnuPG)*

GNU Privacy Guard (GnuPG)

An open source tool implementing the OpenPGP encryption standard.

Overview

GNU Privacy Guard (GnuPG) is a command-line tool for encrypted communications and secure data storage based on the OpenPGP standard described in RFC 2440. Using GnuPG, you can encrypt and decrypt messages, digitally sign documents, and create and manage keys for public key encryption.

The GnuPG project is partially funded by the German government and is freely available under the General Public License (GPL) as open source software. GnuPG does not incorporate any patented encryption technologies and is therefore usually not subject to encryption export standards, though this may vary from country to country. In addition to its built-in support for Advanced Encryption Standard (AES), Blowfish, Data Encryption Standard (DES), Twofish, and other common encryption standards, GnuPG is also extensible so that it can easily support future encryption technologies.

GnuPG is available for a variety of platforms including UNIX/Linux, Microsoft Windows, and MacOS.

For More Information

Visit GnuPG online at www.gnupg.org for more information.

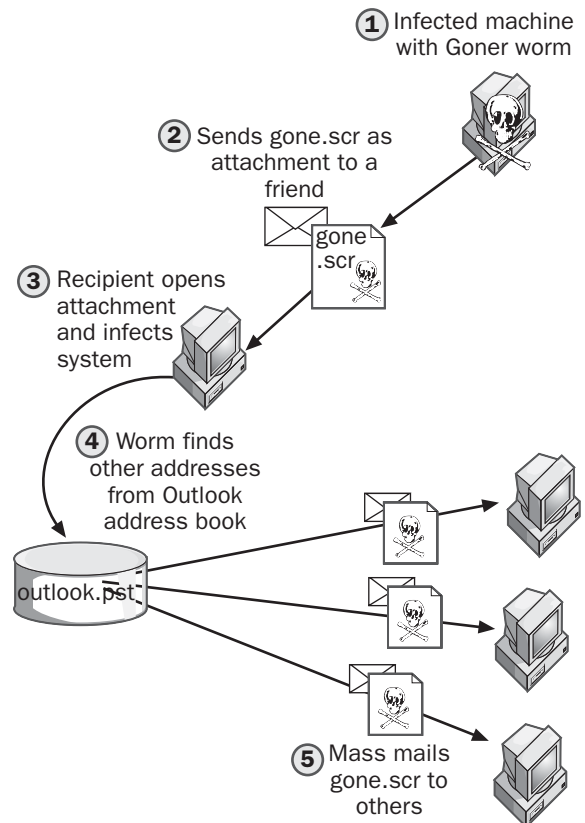
See Also: *OpenPGP, Pretty Good Privacy (PGP), public key cryptography*

Goner

A mass-mailer worm that spread across the Internet in December 2001.

Overview

Goner is a famous mass-mailer worm that spread rapidly in the form of e-mail messages with an attachment named `Gone.scr` that posed as a screen saver. Goner can infect systems two ways: through Microsoft Outlook and through ICQ instant messaging clients. Once a Goner e-mail message has infected a system by tricking a user into opening the attachment, the worm tries to disable antivirus and firewall applications and then spreads itself to others using the Outlook address book and ICQ contacts list. The worm does not damage user data or key operating system files, but by disabling security programs it can reduce the security of a system and expose it to further attack.



Goner. How a mass-mailer worm such as Goner works.

Other names for this worm include W32/Goner-A, W32/Goner@MM, and Pentagone. The best protection against such mass-mailer worms is for administrators to restrict the kinds of attachments that users can send or receive and to educate users concerning the dangers of opening attachments from unknown sources.

See Also: *worm*

Good Times

A famous example of a hoax that poses as an e-mail virus.

Overview

Good Times is perhaps the most famous example of an e-mail virus hoax. The hoax first appeared in 1994 as an e-mail message warning against a dangerous new virus called Good Times spreading across the Internet. The warning said that users could cause their systems to become infected simply by reading an e-mail message containing the virus and that severe harm would result, including erasing all files on users' hard drives. The warning concluded with the famous instruction that recipients should "Forward this to all your friends, it may help them a lot."

Driven by fear and a lack of understanding of how viruses can (and can't) infect systems, users began forwarding this message to their friends. The credibility of the hoax increased when news media began reporting the alleged danger and a number of government agencies and large companies started investing time and energy into researching the issue. Eventually, it came to be known that no such virus existed and worries generally subsided, but not until large resources had been committed to investigating the problem by IT (information technology) departments.

Although Good Times does not qualify as a computer virus in the traditional sense, it may be viewed as a virus in the sense that it had a significant effect on the time and productivity of those who work with computers, and its pattern of "infection" mirrored that of a normal virus. Remarkably, after almost 10 years, this hoax is still circulating, and unsuspecting users continue to

waste the time of help desk personnel and system administrators dealing with its nonexistent threat!

See Also: *hoax, virus*

gray hat

Euphemism for a hacker motivated by curiosity rather than malicious intent.

Overview

The term **gray hat** is used to describe hackers whose activities reside somewhere between those of black hats, who try to damage or steal information from systems, and white hats, who are IT (information technology) professionals who specialize in security. Gray hat is used in different senses depending on the literature you read, and its meaning can include the following:

- People who hack into company networks or software applications looking for vulnerabilities, inform the company of a vulnerability uncovered, but then also post the vulnerability to a public security forum. Some companies object to this action, claiming that by posting such information publicly before a company has had a chance to address the issue, such individuals are providing information that can be used by malicious hackers to try to attack its networks and products.
- People who break into networks for fun and leave harmless messages such as Web site defacements and messages on computer desktops. While such activities might seem to be pranks to some, defacing public Web sites is viewed by many companies as criminal damage to branding information and corporate reputation, and even messages on computer desktops require the time and energy of help desk personnel to deal with them, which costs companies money.
- People who are curious about how software applications work and like to poke around company networks looking for loopholes. In this sense, today's gray hats are the descendents of and closely resemble

the hackers of the classic age of computing in the 1970s and 1980s as described in the book *Hackers: Heroes of the Computer Revolution* by Steven Levy.

- People who enjoy tinkering with software and network security, but want to distance their motives from malicious black hats while maintaining independence from corporate white hat security professionals.

Whatever your view of gray hat hackers, increasing numbers of companies are expressing doubt that “ethical hacking” exists as a legitimate activity, and more and more such incidents result in legal action taken against the hacker who publicly reveals information that might compromise companies’ security or affect their business.

Notes

The origin of the term **gray hat** has been attributed by some to L0pht, a well-known hacking group.

See Also: *black hat, hacker, white hat*

Group Policy

A powerful tool for managing security in Microsoft Windows platforms.

Overview

In addition to its other uses, Group Policy can be used to lock down the security configuration of systems running Microsoft Windows 2000, Microsoft Windows XP, and Microsoft Windows Server 2003. Group Policy is integrated with Active Directory directory service to simplify the configuration and management of systems across large networks, and it includes configuration options for authentication methods, system auditing, event logging, password settings, registry access, Internet Protocol Security (IPSec) encryption, and many

other aspects of system and network security. Group Policy can be used in conjunction with security templates to easily create and deploy custom configurations for locked-down machines. Group Policy is even easier to manage on the Windows Server 2003 platform because the operating system introduced the Group Policy Management Console (GPMC), which simplifies the task of creating, implementing, and testing Group Policy Objects (GPOs).

For More Information

Visit www.microsoft.com/windowsserver2003/gpmc/ for more information about the Group Policy Management Console (GPMC) for Windows Server 2003.

See Also: *security template*

guest account

A local account used to allow limited access to system resources.

Overview

Popular computing platforms such as Microsoft Windows and UNIX include a guest account that can be used to provide anonymous users with access to system resources. Guest accounts should generally be disabled since they can often provide access to resources even for users who have not been authenticated. Best practices also indicate that the guest account should be renamed to make it more difficult for intruders to find once a system has been penetrated. Guest accounts should also have strong passwords that are difficult to crack, and in UNIX systems they should have a restricted shell.

See Also: *Administrator*

hacker

Someone who engages in the activity of hacking computer programs, systems, or networks.

Overview

Hackers can be motivated by a wide range of reasons, ranging from simple curiosity to criminal intent. As a result, the term **hacker** has different connotations in popular use, including the following:

- Programming geniuses who think up innovative ways of coding solutions to problems
- Whiz kids who download free tools from the Internet and use them to crack into systems and then boast of their exploits in chat rooms
- Social activists who deface or deny access to government or corporate Web sites as a form of protest
- Criminals who steal, contaminate, or destroy data for mischief or profit

Various alternative names have been invented for those who perform such activities including crackers, script kiddies, and others. A more recent way of classifying hackers is according to the “hat” they wear, that is, by their perceived or announced motives as follows:

- **Black hat:** A hacker with malicious intent to damage the data or services of an agency or business
- **White hat:** A legitimate security expert with knowledge and experience of black hat methods
- **Gray hat:** A hacker whose motivations and activities reside somewhere in between those of black and white hats

Notes

Interestingly, the term **hacker** has a similar but different use in other contexts. For example, in the context of

golfing, a “hacker” is someone who likes to golf but doesn’t really know how to do it properly and who has learned it on his or her own instead of taking lessons. For example, a golf hacker might not know the proper way of holding a putter, might cheat when a ball goes into the rough, and so on. The comparison with computer hackers is fairly obvious: hackers are generally not “professional” programmers but simply individuals who like to program, though computer hackers are generally more gifted than golf hackers!

See Also: *black hat, exploit, gray hat, hacking, script kiddie, white hat*

Hackers On Planet Earth (HOPE)

A popular series of conferences for black hat hackers.

Overview

Hackers On Planet Earth (HOPE) conferences are organized every few years by *2600* magazine and are held in New York City. The first conference took place in 1994 on the 10th anniversary of *2600* magazine, and recent conferences have included H2K (HOPE 2000) and H2K2 (HOPE 2002). Topics covered by speakers include hacking methodologies, security technologies, and civil liberties issues such as privacy and freedom of information.

For More Information

Visit www.h2k2.net for presentations from the H2K2 conferences and www.2600.com/hopes.html for archives of previous HOPE conferences.

See Also: *2600, black hat, hacking*

hacking

A term with a variety of meanings ranging from programming to network intrusion.

Overview

In a technical sense, **hacking** originally meant devising elegant solutions to difficult technical problems. The term entered popular use with the activities of the MIT Tech Model Railway Club in the 1950s, when members would “hack” switches and relays to improve performance or make them do things they weren’t originally designed to do. When interest in model trains gave way to programming minicomputers such as the PDP-1, hacking came to represent elegant solutions to difficult programming problems.

In modern usage, **hacking** has come to connote maliciously motivated activity, including attempting to penetrate the defenses of computer systems and networks to steal or destroy data. While hacking can be motivated by other reasons—simple curiosity, desire to show off, acts of social protest—the criminal activities of “hackers” are often those that gain the most attention in the media these days. Those more innocently motivated often protest against this semantic shift in meaning and prefer to identify the activity of malicious hacking under other terms such as **cracking** or **phreaking**. Law enforcement agencies tend to favor the words **cyber-crime** and **cyberespionage** to describe criminal hacking activities.

Implementation

Effective hacking is generally a systematic activity that in many ways mirrors traditional counterintelligence and espionage practices. To compromise or break into a network or system, a hacker generally uses steps similar to the following:

- **Footprinting:** Gathering information about a target system using publicly available sources
- **Scanning:** Gathering information about network services on a target system
- **Enumeration:** Gathering information about user accounts and applications on a target system
- **Gaining access:** Compromising the security of a target system by cracking passwords, exploiting buffer overflows, or some other technique
- **Elevating privileges:** Gaining increased control of a target system by elevating the rights of a cracked account
- **Installing backdoors:** Creating hidden mechanisms to allow attackers to reenter a compromised system at will
- **Covering tracks:** Erasing log files and other evidence of intrusion

Once a system has been compromised, the attacker has access to sensitive data and can leverage the system as a platform for attacking other systems, for example, as zombies in a distributed denial of service (DDoS) attack.

For More Information

For a fascinating account of the early days of computer hacking, see the book *Hackers: Heroes of the Computer Revolution* by Steven Levy (Penguin USA, 2001).

See Also: 2600, backdoor, cracking, distributed denial of service (DDoS), elevation of privileges (EoP), enumeration, footprinting, hacker, hacktivism, password cracking, Phrack, phreaking, social engineering, Trojan, zombie

hacktivism

Hacking for ideological reasons such as social or political protest.

Overview

Hacktivism is the online expression of activism and may be motivated for reasons similar to those for which individuals participate in protest marches, sit-ins, and similar activities. Hacktivism can take many forms that can disrupt business or government operations to various degrees. Web site defacement is one form of hacktivism similar to painting graffiti on signs and buildings. E-mail bombs are another form and involve sending large numbers of messages with large attachments in an attempt to overpower the ability of a mail server to

cope. The result is a denial of service (DoS) condition in which the server cannot handle legitimate e-mail traffic. Other forms of hacktivism can include electronic petitions, worms and viruses, computer break-ins, and so on. Most “hacktivists” can be classified as either trespass or blockage types and are subject to the legal penalties associated with such activities.

Opinions vary as to when such activities can be considered ethical. Some consider hacktivism as a legitimate form of electronic civil disobedience, while others label it criminal activity or cyberterrorism. As in most forms of civil disobedience, the perpetrators must be willing to resign themselves to criminal prosecution should their activities break state or federal laws.

For More Information

The Web site www.attrition.org maintains a gallery of famous examples of Web site defacements.

See Also: *hacking*

hardening

Configuring a host to make it more secure for a specific role.

Overview

Hardening refers to a combination of techniques to make special-purpose hosts secure against attack. Hosts that typically need to be hardened include Web servers, mail servers, Domain Name System (DNS) servers, firewalls, and other bastion hosts. The actual steps used to harden a server vary with the operating system platform used, but the general approach usually includes the following procedures:

- Removing or disabling any components that may have been installed by default but are unnecessary with respect to the server’s designated role
- Disabling unnecessary networking services to simplify the configuration of the server and provide only those services needed by clients
- Increasing access controls on critical system components such as system dynamic-link libraries (DLLs), configuration files, the registry, and other potential targets of attack

- Turning on password encryption and other cryptographic features that may not have been enabled by default during installation
- Configuring security policies to restrict access to critical system functions to the smallest possible user base
- Using file system checking and process tracking to record any unusual activity in the system logs

When hardening a server, it is generally best to perform a clean install of the operating system with the server disconnected from the production network. This will help ensure that no intrusion occurs during the hardening process and that the system is virus- and Trojan-free. Hardening is a two-edged sword since too much hardening can make the server difficult to administer, while too little leaves it vulnerable to attack.

While many security organizations and vendors have developed step-by-step procedures for hardening different platforms and products, hardening remains more of an art than a science, and servers being hardened should be carefully tested from both a security and a manageability perspective at each step along the way. Proper documentation of steps taken to harden a server must also be performed, and the configuration of hardened servers should be periodically reviewed and fine-tuned as necessary by applying patches and fixes supplied by vendors and security advisory services.

See Also: *bastion host*

hardware security module (HSM)

A hardware device used for protecting cryptographic keys.

Overview

A hardware security module (HSM) is a peripheral that attaches to a system and is used to generate and store keys used for encrypting information. The advantage of storing keys in this fashion is that it is more secure than storing them on a system’s hard drive, because if the system was compromised, the intruder would have access to the keys and could use them for impersonating

the identity of the user. HSMs are commonly used in Public Key Infrastructure (PKI) systems for storing critical keys such as root certificate authority (CA) keys.

The National Security Agency (NSA) and National Institute of Standards and Technology (NIST) have developed a set of criteria for evaluating the security of HSMs. These criteria are published as Federal Information Processing Standard (FIPS) 140-1, entitled “Security for Cryptographic Modules,” and rate such systems from level 1 (weakest) to level 4 (strongest) in terms of security.

Marketplace

Some well-known vendors of HSMs and their products include nShield from nCipher, Luna from Chrysalis-ITS, and Cryptoswift from Rainbow Technologies.

See Also: encryption, Federal Information Processing Standard (FIPS), key, Public Key Infrastructure (PKI), root certificate

hash

Another name for message digest, the result of applying a hashing algorithm to a message.

See: message digest

hashing algorithm

A mathematical procedure that generates a fixed-size result from arbitrary amounts of data.

Overview

Hashing algorithms are used in cryptography for creating messages digests, a kind of cryptographic checksum used to verify that an electronic message has not been modified in transit. Message digests are also used in digital signatures to verify the identity of the sender of a message.

Hashing algorithms are also used in some authentication schemes in which hashed values of users’ passwords are stored on the server for greater security. An example of an authentication scheme that uses hashing is challenge-response authentication, which compares two hashed values to determine whether to authenticate the user.

Hashing algorithms can also be used like simple checksum schemes to ensure the integrity of stored information. If data is modified in any way, its hash value changes and the user knows that the data has been compromised.

Implementation

A hashing algorithm is generally an iterative mathematical procedure that “scrambles” information, converting plaintext into unreadable ciphertext called a **hash**. Hashing algorithms are generally one-way functions in the sense that it is impossible or impractical to convert the hashed value back into its original form. The result of applying a hashing algorithm to any amount of plaintext is a fixed-size block of ciphertext that bears no resemblance to the original plaintext. If a well-designed hashing algorithm is applied to two portions of plaintext that differ only slightly, the result is two blocks of ciphertext that are completely different.

There are numerous examples of hashing algorithms commonly used in cryptography and authentication systems. Some of the more popular ones include these:

- **MD2, MD4, and MD5:** A series of hashing algorithms developed by Ron Rivest that creates a 128-bit message digest.
- **SHA-1:** A hashing algorithm defined by the Federal Information Processing Standard (FIPS) 180-1 Secure Hash Standard (SHS) that creates a 160-bit message digest. More recent variants include SHA-256, -384, and -512, which are collectively known as SHA-2.

Other less common hashing algorithms include HAVAL, Panama, RIPEMD, Snefru, and TIGER.

Notes

Another name for hashing algorithm is **hash function**.

See Also: ciphertext, cryptography, MD2, MD4, MD5, message digest (MD), plaintext, Secure Hash Standard (SHS), Secure Hash Algorithm-1 (SHA-1), SHA-2

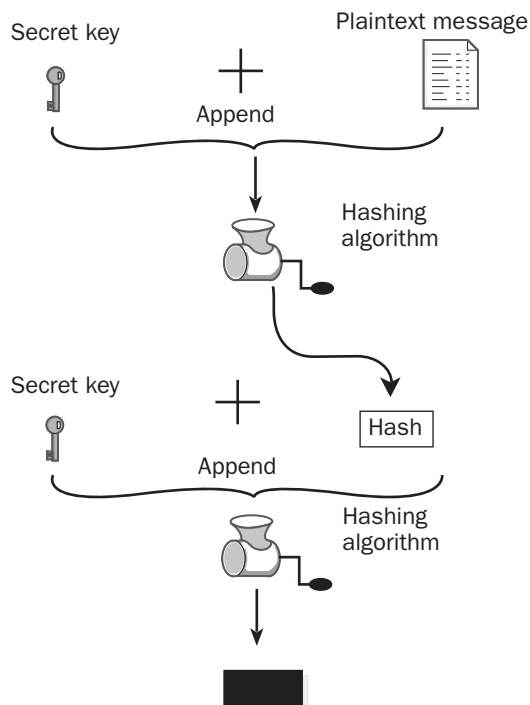
hash-based message authentication code (HMAC)

A message authentication code (MAC) algorithm that combines a hashing algorithm with a secret key.

Overview

Hash-based message authentication code (HMAC) is specified in RFC 2104 as a method for authenticating digital messages using a combination of standard hashing algorithms and symmetric key encryption. HMAC is a variant of MAC algorithms and can be used both to authenticate the source of a message and verify its integrity. In order to use HMAC, a shared secret key must be known by both the sender and the receiver.

HMAC is typically implemented using either the MD5 or SHA-1 hashing algorithms. The strength of an HMAC implementation depends on both the hashing algorithm and key length used. In typical usage, a combination of the plaintext message and shared secret are hashed and then the result is combined with the key and hashed again.



Hash-based message authentication code (HMAC). How HMAC works.

Notes

HMAC also is often called Keyed-hashing for Message Authentication, which is the form of the name used in RFC 2104, and Keyed-Hash Message Authentication Code, which is the form used in the Federal Information Processing Standard (FIPS) 198 and American National Standard Institute (ANSI) X9.71 standards.

See Also: *hashing algorithm, MD5, message authentication code (MAC), Secure Hash Algorithm-1 (SHA-1)*

headless server

A server without a keyboard, monitor, or mouse.

Overview

Headless servers are commonly used in service provider data centers where large numbers of servers are administered remotely from a central management station. Rack-mounted servers are typical examples of headless systems and require basic input/output systems (BIOS) and motherboards that support operation without a keyboard, mouse, or monitor (and often without a video card as well). This saves the energy and space used by having input/output peripherals attached to servers and the cost of keyboard, video, mouse (KVM) switches. Another advantage of operating servers in headless mode is increased physical security since malicious users are unable to interact with the server with no input device present.

In order to manage headless servers, out-of-band (OOB) connections are frequently used over serial ports. Alternative methods for managing headless servers include terminal services and Web interface applications over in-band network connections. Headless servers have been common for some time in UNIX environments. Microsoft Windows Server 2003 supports headless operation operating on Intel hardware platforms that support this feature.

See Also: *physical security*

hex editor

Tool used to modify binary files.

Overview

Hex editors go back to the early days of disk operating system (DOS) programming and allow binary files such as system dynamic-link libraries (DLLs) and executables to be displayed as hexadecimal characters (bytes) in similar fashion to how American Standard Code for Information Interchange (ASCII) files are displayed in text editors such as Microsoft Notepad. Using a hex editor to modify system files generally requires deep understanding of C/C++ programming and assembly language, but hackers can sometimes extract useful information from these files by manipulating them through programmatic means. For example, if an application has an embedded password or product key, this information can sometimes be extracted using a hex editor by changing the password or entering a different product key and then comparing the result byte by byte with the original file.

Some popular free and shareware hex editors include FRHED, HexIT, XVI32, HexEdit, and Hex Wizard. While the Notepad text editor on the Microsoft Windows platforms cannot be used to modify binary files because it adds nonprinting CR-LF characters when saving files, the legacy MS-DOS editor EDIT.COM did support a binary mode that could enable it to be used as a simple hex editor if required!

See Also: *hacking*

hex encoding URL attack

A form of file system traversal attack using hexadecimal characters.

Overview

In a file system traversal attack, the attacker uses “../” strings in Uniform Resource Locators (URLs) to access files outside the root directory of a Web site. A variant of this attack employs the string “%2e%2e%2f” instead, which represents the hexadecimal encoding of “../” in the URL. A hex-encoding URL attack takes this one step further by reencoding the string “%2e%2e%2f” itself in hexadecimal form—for exam-

ple, by representing “%2e” as “%25%32%65” and “%2f” as “%25%32%66” in order to circumvent the internal URL parsing routine of the Web server. If necessary, this procedure can be iterated further to circumvent Web server security or an intrusion detection system (IDS).

This vulnerability was exposed on the Microsoft Internet Information Services (IIS) version 4 and 5 platforms when the Nimda worm appeared in September 2001. Proper coding practices will prevent such attacks, but their very existence is testimony to the difficulty for developers to code their applications for every possible eventuality and the ingenuity of attackers who devise such schemes for compromising systems.

See Also: *file system traversal attack, intrusion detection system (IDS)*

HFNetChk

A Microsoft tool for keeping security patches up to date on a system.

Overview

HFNetChk is a command-line tool included in the Microsoft Baseline Security Analyzer (MBSA). Using HFNetChk, an administrator can determine which hotfixes are installed on a Microsoft Windows NT, Windows 2000, Windows XP, or Windows Server 2003 system. HFNetChk works by interacting with an online Extensible Markup Language (XML) file that Microsoft Corporation maintains on the Microsoft Download Center Web site. HFNetChk can be used to manage hotfixes on multiple systems by running it from a central administrator console and can also be used to manage hotfixes on Microsoft Exchange and Microsoft SQL Server.

See Also: *hotfix, Microsoft Baseline Security Analyzer (MBSA)*

hidden file

On UNIX/Linux platforms, hidden files are files whose names begin with a period (“.”), and they are therefore often called **dot files**. Hidden files are not displayed by default when browsing the file system from the command

line. Hidden files are generally important files relating to a user's environment and are a target for exploits by hackers. Examples of hidden files commonly found in home directories include `.login`, `.mailrc`, and `.forward`. Once an attacker has compromised a system, attackers may create hidden directories with unusual names such as "...." to hide utilities they may install such as backdoors and Trojans. The `find` command can be used to search for hidden files on a UNIX system.

See Also: *backdoor, Trojan*

HIDS

Stands for host-based intrusion detection system, an intrusion detection system (IDS) that monitors activity on a single host.

See: *host-based intrusion detection system (HIDS)*

hierarchy of trust

Another name for certificate authority (CA) hierarchy, a hierarchical collection of CAs bound together by trust relationships.

See: *CA hierarchy*

hijacking

In network security, theft of credentials, sessions, or identity.

Overview

Hijacking is a general term that applies to several types of malicious activity against computer systems and networks. In session hijacking, for example, software used to "sniff" network traffic can capture credentials and allow an attacker to impersonate a user at one end of a communication session. Domain-name hijacking occurs when an attacker convinces a domain name registration authority that the attacker is the legitimate owner of a domain name, with the result that traffic is redirected from the company's Web site to a site designed by the attacker. Examples of well-known domain names that have been hijacked in the past include *internet.com*, *nike.com*, *exodus.net*, and even

w3.org. Home page hijacking occurs when software constantly resets your browser home page to something other than what you desire, and often it occurs when software downloaded from the Internet includes adware or other "stealth" software.

See Also: *adware, session hijacking*

HMAC

Stands for hash-based message authentication code, a message authentication code (MAC) algorithm that combines a hashing algorithm with a secret key.

See: *hash-based message authentication code (HMAC)*

hoax

In the context of computer security, a phony virus threat.

Overview

Hoaxes are generally e-mail messages warning users about the dangers of supposed viruses. While many hoaxes may have been intended as pranks, their effect is often far from harmless as users overload help desk personnel with anxious requests for help in preventing infection. The result can be a considerable expenditure of time and energy by IT (information technology) personnel, with resulting costs incurred to the business until the hoax can be debunked and users reassured.

A famous example of a hoax was the Deeyenda virus hoax. This hoax warned users that by merely reading a certain e-mail message they would infect their system with a virus that would delete everything on their hard drive. This hoax took advantage of users who didn't understand that most viruses are propagated by attachments, not messages. The hoax included technical-sounding language that added plausibility to the message in the minds of readers. The hoax also purported to be a warning from the Federal Communications Commission (FCC), adding credibility to the message in the minds of ordinary users. The hoax also included the suggestion that recipients forward the message to all their friends, thus building a propagation mechanism into the hoax and resulting in a chain letter.

Other famous examples of hoaxes include AOL4FREE, Good Times, Happy New Year, Irina, PKZ300, and many others. Some hoaxes are almost a decade old by now and are still floating around the Internet, causing alarm to users and frustration to IT departments!

Hoaxes are fairly easy to recognize, however. If you receive an e-mail message that is an unsolicited warning concerning a threat or virus and the message suggests you delete certain files from your computer and/or requests that you forward it to friends, then it is probably a hoax and should be ignored.

For More Information

Visit the HoaxBusters page on the Department of Energy Computer Incident Advisory Capability (CIAC) site at hoaxbusters.ciac.org for more information on Internet hoaxes.

See Also: *Computer Incident Advisory Capability (CIAC), virus*

HoneyNet Project

A nonprofit initiative to learn the techniques used by hackers to break into computer networks.

Overview

The HoneyNet Project is a collaborative effort by a group of security professionals intent on learning the methods used by black hat hackers. The purpose of the project is to capture and record steps used by hackers to break into honeypots, which are decoy systems that mimic legitimate servers but that are actually intended as bait to lure hackers away from attacking real servers. The HoneyNet Project extends this idea of honeypots into entire networks called honeynets that mimic the operations of real networks but that are actually designed to lure hackers into attacking them and then recording how intrusions are performed. Information learned through the project is shared with such organizations as the SANS Institute and CERT Coordination Center (CERT/CC).

The HoneyNet Project is a volunteer effort started by Sun Microsystems engineer Lance Spitzer. It has helped raise awareness in the security community concerning black hat tools and methods. The success of the project has led to the formation of the HoneyNet Research Alliance involving participants in industry, government, business, academia, and the military.

For More Information

Visit the HoneyNet Project at www.project.honeynet.org for more information.

See Also: *black hat, hacking, honeypot, intrusion*

honeypot

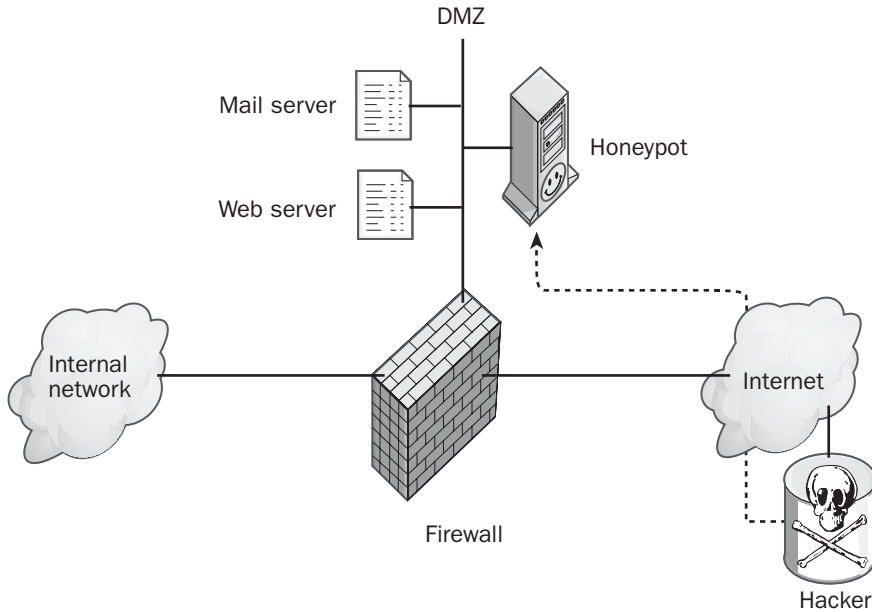
A host that is deliberately exposed to attack.

Overview

Honeypots are typically dummy hosts deployed on a demilitarized zone (DMZ) to attract an attacker away from legitimate hosts. The term **honeypot** suggests how such a host might draw attackers to it as bees are drawn to honey.

In a typical scenario, a dummy Web server might be deployed with phony information on it and left in a vulnerable state by leaving security patches unapplied. When an attacker enumerates the network segment and finds the vulnerable host, the attacker expends time and energy compromising a host while leaving more hardened legitimate hosts alone. A host-based intrusion detection system (HIDS) is typically installed on the dummy host to detect the attacker's activities and log them for further analysis.

In addition to deflecting attack, honeypots can also be used to collect information about how attackers try to compromise systems. Using this knowledge, sysadmins can better understand how to harden their systems against attack.



Honeypot. Using a honeypot to deflect attackers from real servers.

Marketplace

A number of vendors offer commercial honeypot systems you can deploy on your network as decoys to divert the interest of hackers from your real servers. Some popular examples are Cybercop Sting from Network Associates, ManTrap from Recourse Technologies, SPECTER from NETSEC, Back Officer Friendly from NFR Security, and Tripwire from Tripwire, Inc. There is also an open source honeypot called Honeyd that can be used to create multiple virtual honeypots on UNIX systems.

Issues

Before you deploy a honeypot on your network, consult your legal department about whether your setup constitutes enticement or entrapment to ensure you won't find yourself in front of a judge facing a lawsuit!

See Also: demilitarized zone (DMZ), hardening, Honeynet Project

HOPE

Stands for Hackers On Planet Earth, a popular series of conferences for black hat hackers.

See: Hackers On Planet Earth (HOPE)

host-based IDS

Stands for host-based intrusion detection system, an intrusion detection system (IDS) that monitors activity on a single host.

See: host-based intrusion detection system (HIDS)

host-based intrusion detection system (HIDS)

An intrusion detection system (IDS) that monitors activity on a single host.

Overview

A host-based intrusion detection system (HIDS) is generally agent software that resides on a host and monitors for suspicious activity. Such activity might include attempts to modify system files or the registry. HIDSs typically work by monitoring event logs and other logs on the system and by using file system notification to detect attempts to modify system dynamic-link libraries (DLLs). HIDSs may also monitor incoming network traffic, looking for suspicious events such as multiple failed authentication events. When suspicious activity is detected, the HIDSs notifies its management station, centralized software that controls HIDSs deployed on multiple hosts.

Implementation

Managing large numbers of HIDSs can be a complex task, so deployment is usually limited to critical servers exposed to hostile traffic, such as Web servers or mail servers located in the demilitarized zone (DMZ) of a firewall-protected network. HIDSs are generally used along with a network-based intrusion detection system (NIDS) to ensure the greatest level of security against intrusion. While NIDS are generally platform independent in operation, HIDSs are agent software designed for specific operating system platforms.

Marketplace

There are numerous HIDSs available in the market today. Some popular ones include Centrax from CyberSafe, RealSecure and Server Sensor from Internet Security Systems, Intruder Alert from Symantec, and Dragon from Enterasys Networks. Many firewall products such as SecureIIS from eEye also include basic HIDS features.

See Also: demilitarized zone (DMZ), intrusion detection system (IDS)

host-based security

Security implemented by configuring each host on a network.

Overview

Host-based security refers to the process of configuring the security of hosts, as opposed to network-based security, which refers to implementing security

measures that affect all hosts together. Implementing host-based security is similar to hardening systems and involves configuring file system permissions, configuring account policies, disabling unneeded network services, and implementing process accounting. Host-based security may also involve installing special security software on individual hosts including host-based firewall agents and host-based intrusion detection systems (HIDSs).

While host-based security is essential to consider in a multilayered approach to network security, the problem is that host-based security doesn't scale well since with increasing numbers of hosts there is correspondingly more work configuring them. System management tools such as Openview from Hewlett-Packard (HP) or Tivoli from IBM help in managing the security of distributed hosts but are themselves complex to deploy and operate. As a result, it is essential to complement the host-based security approach with network-based security by deploying high-performance firewalls at network choke points and using network-based intrusion detection systems (NIDSs) to monitor traffic on the network as a whole.

See Also: firewall, hardening, intrusion detection system (IDS), network-based intrusion detection system (NIDS), network-based security

hotfix

A security patch for a Microsoft product.

Overview

When vulnerabilities are discovered in Microsoft operating systems or applications, Microsoft engineers create a patch called a hotfix that can be downloaded and applied to affected systems to resolve the problem. These hotfixes can be distributed in several ways:

- By notifying users who have subscribed to the Microsoft Security Notification Service at www.microsoft.com/technet/security/bulletin/notify.asp
- By visiting the Microsoft Windows Update Web site at windowsupdate.microsoft.com and manually scanning your system for missing updates

- By using the Automatic Updates feature of Microsoft Windows XP, Microsoft Windows Server 2003, and *Service Pack 3 for Microsoft Windows 2000* to connect automatically to Windows Update and download updates on a scheduled basis
- By using the HFNetChk tool included in the Microsoft Baseline Security Analyzer (MBSA) to scan systems on your network for updates they may require
- By visiting the Security Bulletins page on Microsoft TechNet at www.microsoft.com/technet/security/current.asp and comparing the list of updates displayed with documentation of hotfixes you've installed on your servers
- By waiting and installing the latest service pack when it comes out, since hotfixes are consolidated for service packs

See Also: *HFNetChk, Microsoft Baseline Security Analyzer (MBSA), Microsoft Security Notification Service, service pack, Windows Update*

Hping

A security tool for testing and auditing Transmission Control Protocol/Internet Protocol (TCP/IP) networks.

Overview

Hping is an open source command-line tool that can be used to assemble and send custom TCP/IP packets including Transmission Control Protocol (TCP), User Datagram Protocol (UDP), Internet Control Message Protocol (ICMP), and raw Internet Protocol (IP) packets. Hping is similar in use and output to the Ping command but is much more powerful and can be used to perform port scanning, firewall testing, and fingerprinting remote systems. Hping can handle oversized and fragmented packets, supports covert channels, can be used for encapsulated file transfer, and has many other advanced features.

Hping was developed by Italian programmer Salvatore Sanfilippo and is freely available for UNIX/Linux platforms under the GNU Public License (GPL). The current version at time of writing is Hping2.

For More Information

Visit Hping online at www.hping.org for more information.

See Also: *fingerprinting, Ping, port scanning*

HSM

Stands for hardware security module, a hardware device used for protecting cryptographic keys.

See: *hardware security module (HSM)*

.htaccess

A configuration file for Apache Web servers.

Overview

The .htaccess file is a hidden or “dot” file on UNIX/Linux hosts running the Apache Web server. The .htaccess file is used to configure various aspects of the Web server, including the following security features:

- Password protecting directories
- Restricting access based on client Internet Protocol (IP) address
- Disabling directory listing
- Redirection to another site

The .htaccess file is an American Standard Code for Information Interchange (ASCII) file that can be edited using a text editor. Permissions should always be configured to prevent unauthorized users from viewing the contents of this file.

See Also: *hidden file*

HTTPS

Another name for Secure Sockets Layer (SSL) when used in conjunction with Hypertext Transfer Protocol (HTTP) communications.

See: *Secure Sockets Layer (SSL)*

hybrid attack

A combination of a brute-force attack and a dictionary attack.

Overview

Brute-force attacks are the most general method used to crack passwords, but they are often time-consuming and ineffective. Dictionary attacks try a database of common passwords to achieve the same purpose, but users who employ simple strategies such as appending numbers to their password strings can easily circumvent such attacks. The optimal approach to password cracking is the hybrid attack, which combines the features of brute-force and dictionary attacks. In a typical hybrid attack, the cracking program generates short strings of

characters and adds them to the beginnings and ends of dictionary words. For example, a password such as “daisy123” would likely succumb very quickly to a hybrid attack, which would try the word “daisy” with various short strings of characters appended.

Notes

LOphtcrack is a popular password cracking tool that can carry out hybrid attacks.

See Also: *brute-force attack, dictionary attack, LOphtCrack, password cracking*

IA

Stands for information assurance, methodologies for ensuring the security of information systems.

See: information assurance (IA)

IASE

Stands for Information Assurance Support Environment, a U.S. Department of Defense (DoD) clearing-house for information assurance (IA) information.

See: Information Assurance Support Environment (IASE)

IATF

Stands for Information Assurance Technical Framework, a framework for ensuring the security of information systems.

See: Information Assurance Technical Framework (IATF)

ICMP attacks

Attacks that exploit characteristics of Internet Control Message Protocol (ICMP).

Overview

ICMP is the portion of the Transport Control Protocol/Internet Protocol (TCP/IP) protocol suite responsible for sending error messages and providing methods for testing IP networks. Unfortunately, ICMP also has been exploited by malicious parties for performing various types of attacks against networks. Some of the common types of ICMP attacks include the following:

- **ICMP fingerprinting:** This technique uses ICMP to determine the operating system running on a host.

- **ICMP flood:** Also called a Smurf attack, this is a popular type of denial of service (DoS) attack.
- **ICMP sweep:** Also called a ping sweep, this is a technique for determining which hosts are active on a network.

Other approaches include these:

- Sending oversized ICMP messages to crash a target host
- Using ICMP route redirect messages to perform a man-in-the-middle (MIM) exploit
- Using ICMP router discovery messages to spoof a router and hijack traffic
- Using ICMP tunneling to set up a covert channel to leak information from a system

See Also: covert channel, hijacking, ICMP fingerprinting, ICMP tunneling, man-in-the-middle (MITM) attack, ping sweep, Smurf attack

ICMP enumeration

Using Internet Control Message Protocol (ICMP) messages to enumerate hosts on a network.

Overview

Enumeration is gathering information about a target system or network, such as which hosts are alive on a network. Firewalls are often configured to block ICMP echo messages (pings and traceroutes) but not other types of ICMP messages, such as time stamp or information request (also known as information reply) messages. By sending such ICMP messages to all possible Internet Protocol (IP) addresses on a remote network, an attacker can determine which hosts are alive (responding) on the remote network.

Notes

Popular tools used by attackers for performing ICMP enumeration include Icmpenum and Icmpquery.

See Also: enumeration, firewall

ICMP fingerprinting

Using Internet Control Message Protocol (ICMP) messages to fingerprint a host.

Overview

Fingerprinting is the process of determining the identity of a remote system by analyzing packets it generates. One way of fingerprinting Internet Protocol (IP) hosts is to send ICMP echo requests to the host and analyze the packets that are returned. Different operating systems sometimes implement ICMP differently in their Transmission Control Protocol/Internet Protocol (TCP/IP) stacks, resulting in slightly different bit patterns returned in response to ICMP echo requests. By comparing the response returned to a known database of ICMP echo response signatures, an attacker could determine which operating system is running on the remote host and use this information to better target the host for an exploit.

Another technique used in ICMP fingerprinting is to create specially crafted ICMP echo request packets that have nonstandard time stamps or other modifications. Different operating systems often respond to such nonstandard packets in unique ways.

See Also: fingerprinting, ICMP attacks

ICMP flood

Also called a Smurf attack, a denial of service (DoS) attack that uses Internet Control Message Protocol (ICMP) echo requests.

See: Smurf attack

ICMP sweep

Also called a ping sweep, a method of footprinting a network using Internet Control Message Protocol (ICMP) echo requests.

See: ping sweep

ICMP Traceback (itrace)

A proposed modification to Internet Control Message Protocol (ICMP) to enable Internet Protocol (IP) traffic to be traced back to its source.

Overview

ICMP Traceback (also known as itrace) is a proposed new ICMP message that is designed to help track down the source from which denial of service (DoS) attacks are originating. In order to reach a destination host on the Internet, an IP packet typically traverses a number of routers or hops. With ICMP Traceback enabled on them, routers would randomly emit one traceback message for every 20,000 packets they forward. This traceback message could be sent either to the source host from which the packet originated, or to the destination host to which the packet is targeted. By analyzing enough of these traceback messages, an administrator could determine the IP addresses of hosts from which a DoS attack is originating, and could then use this information to contact the owner of the network from which the attack was launched.

ICMP Traceback is currently an Internet draft standard. One limitation is that in order for ICMP Traceback to work, it must be implemented widely across the Internet, especially in backbone and edge routers of Internet service providers (ISPs) and enterprise networks. Because traceback messages constitute only 1 out of every 20,000 packets forwarded, they will have negligible effect on Internet traffic patterns.

See Also: denial of service (DoS)

ICMP tunneling

A method of using Internet Control Message Protocol (ICMP) to establish a covert channel.

Overview

Covert channels are communications channels that hide illicit information flow within a normal communications stream. One method of establishing a covert channel on an Internet Protocol (IP) network is to hide data in packets that normally don't carry payloads. An example is ICMP tunneling, which hides data in ICMP echo request/reply packets, the types of packets generated by Ping. If firewalls are configured to pass such traffic (and they often are since Ping is primarily a troubleshooting tool), then information can be leaked from the system without being detected by a firewall or intrusion detection system (IDS).

A common use of covert channels is communication with backdoors. Once an attacker has compromised a system and installed a backdoor, a covert channel allows the attacker to control the system or leak information from it using innocuous-looking ICMP echo packets. One tool that attackers can use for this purpose is Loki, a program first published in *Phrack* magazine. The best way of preventing ICMP tunneling is to block all ICMP traffic at the firewall.

See Also: *covert channel*, *Phrack*

IDEA

Stands for International Data Encryption Algorithm, a block cipher encryption algorithm developed by Xuejia Lai and James Massey.

See: *International Data Encryption Algorithm (IDEA)*

Identity theft

Impersonating someone's identity by stealing personal information.

Overview

While identity theft is not in essence a cybercrime, the increasing use of e-commerce has compounded the problem. Identity theft occurs when a criminal steals personal information concerning someone and then uses this information to obtain a driver's license, apply for credit cards, access bank accounts, and perform other actions that can harm the victim's financial security and reputation. With increasing numbers of

Web sites collecting personal information and storing it in databases accessible from the Internet, identity theft has become a major concern of law enforcement agencies and is costing financial and credit agencies billions of dollars each year.

To protect yourself against traditional methods of identity theft, you could do the following:

- Shred old financial statements and personal documents.
- Avoid giving out your Social Security number.
- Use a locked mailbox to protect your mail when you are not around.
- Check your credit card statements for accuracy.
- Review your credit information regularly for signs of misuse.

For individuals who use the World Wide Web for e-commerce, additional steps should be taken, such as the following:

- Ensure that Web sites on which you make purchases or perform financial transactions are using a secure server (shown by a padlock icon at the bottom of the browser window).
- Use different e-mail addresses for personal and business correspondence.
- Avoid giving personal information such as birth date or mother's maiden name when requested (such information often is used by financial institutions to identify you).
- Perform vanity searches (such as typing your name in Google and seeing what comes up) to detect misuse of your identity.

For More Information

Visit the Identity Theft Resource Center at www.idtheftcenter.org for more information. U.S. citizens can call the Federal Trade Commission (FTC) identity theft hotline at 877-IDTHEFT if they believe they may be victims of identity theft.

See Also: *cybercrime*

idle host scan

A stealth method of scanning a host on a network.

Overview

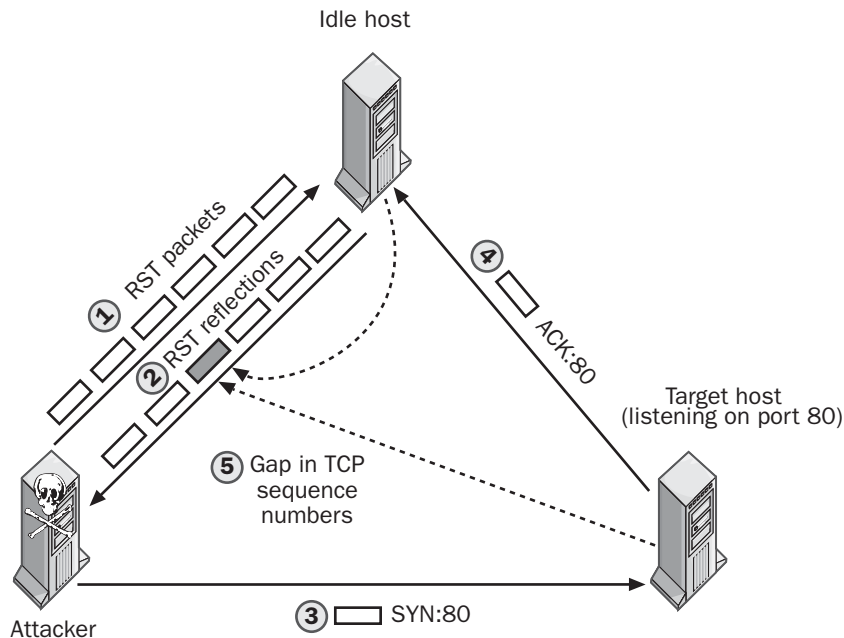
Port scanning is often used by attackers as a method of footprinting target systems or networks. Scanning determines which services are listening on the target, and this can provide attackers with useful information for attempting exploits using known vulnerabilities. An intrusion detection system (IDS) can detect when a host is being scanned, however, and ports can be closed to prevent the scan from being effective. Idle host scanning is a way of scanning ports on target hosts without an IDS detecting any unusual activity.

Implementation

Idle host scanning works by reflecting traffic off a third host and using gaps in Transmission Control Protocol (TCP) sequence numbers to determine which services are running on the target host. An attacker begins by sending a steady stream of TCP Reset (RST) packets to an unsuspecting third host on the network. This host should be relatively inactive (an idle host) so that it can generate a continuous sequence of RST packets in

response to the one received. The attacker then forges a synchronization (SYN) packet so that its source address is that of the idle host and its destination address is that of the target host. The forged packet is designed to open a session using a specific TCP or User Datagram Protocol (UDP) port in order to test whether that port is open on the target. When the target receives the packet, the acknowledgment (ACK) reply (if there is one) is sent to the idle host, causing an interruption in the TCP sequence numbers of the RST packets sent to the attacker. In other words, if the target is listening on the specified port, a “gap” is seen in the TCP sequence numbers of the reflected stream of packets returning to the attacker from the idle host, and the existence of this gap indicates the presence of an open port on the target system.

In order for idle host scanning to work, however, the Transmission Control Protocol/Internet Protocol (TCP/IP) stack on the idle host must generate TCP sequence numbers according to a predictable scheme that the attacker can decipher, and unfortunately this is the case with many TCP/IP stack implementations.



Idle host scan. How an idle host scan is performed.

Notes

Some tools often used by attackers for performing idle host scans are Hping and Idlescan.

See Also: port scanning

IDS

Stands for intrusion detection system, an application or device that identifies suspicious network activity.

See: intrusion detection system (IDS)

IIS Lockdown Tool

A downloadable tool for helping administrators secure Internet Information Services (IIS) versions 4 and 5.

Overview

The IIS Lockdown Tool facilitates securing IIS 4 and 5 by disabling unnecessary features to reduce the attack surface on Web servers. This wizard-based tool includes support for the following:

- Server roles, templates for configuring IIS in various scenarios
- URLscan integration for screening incoming Hypertext Transfer Protocol (HTTP) requests
- Selectively disabling HTTP, File Transfer Protocol (FTP), Simple Mail Transfer Protocol (SMTP), or Network News Transfer Protocol (NNTP) services

Notes

In IIS 6 on the Microsoft Windows Server 2003 platform, the IIS Lockdown Tool has been replaced by the Web Service Extensions node of Internet Services Manager.

For More Information

Visit www.microsoft.com/downloads/ to download the IIS Lockdown Tool.

See Also: URLscan

IKE

Stands for Internet Key Exchange, the key management protocol used by Internet Protocol Security (IPSec).

See: Internet Key Exchange (IKE)

IKEv2

Stands for Internet Key Exchange version 2, a proposed replacement for Internet Key Exchange (IKE).

See: Internet Key Exchange version 2 (IKEv2)

ILOVEYOU

Another name for the LoveLetter worm, a malicious VBScript program that spreads using the Microsoft Outlook address book.

See: LoveLetter

impersonation

Ability of a process to run using a different security context than the one that owns the process.

Overview

Impersonation is a feature of operating systems and applications that allows them to respond to client requests. Typically, a server impersonates a client to allow the client to access resources on the server. For example, Internet Information Services (IIS) uses impersonation to provide a secure context for responding to anonymous requests from clients.

An **impersonation token** is an access token that contains the security information of a client process and allows the server to impersonate the client to access resources.

See Also: authentication

incident

An adverse event affecting an information system.

Overview

Generally, an **incident** is any event that compromises the security of a system, a network, or data. An incident need not be real—even the threat of such an event is considered an incident in most cases. Examples can include malicious activities such as the following:

- Stealing hardware or software
- Using accounts or privileges without authorization

- Tampering with stored data
- Running malicious code that damages systems or data
- Disrupting service to legitimate users
- Misusing information for personal gain or industrial espionage
- Perpetrating hoaxes that cause stress and waste business resources

Incidents can also have accidental or natural origins, including these:

- Electrical power outages
- Hardware failures because of poor ventilation
- Civic disruption because of riot or vandalism
- Human error in entering data or configuring systems

See Also: *incident response, incident response team*

incident response

An action taken in response to an incident affecting information security.

Overview

Incident response is planned action in response to adverse events affecting systems, networks, and data. Response to an incident can range from recording the incident to alerting an incident response team to initiating legal action against malicious individuals. The best way to deal with incidents affecting information security is to follow a planned approach laid out in a carefully developed security policy. Such policies should outline what response is suitable for each type of incident, the individuals responsible for handling the situation, and appropriate escalation procedures to follow if necessary.

Incident response is a systematic activity designed to minimize the impact of information loss or theft, assist the company in recovering from the incident and

resuming normal business practice as quickly as possible, and help set in place procedures to prevent recurrence of such incidents in the future.

Notes

Incident response is generally limited to incidents whose origin is malicious in nature. Incidents caused by natural disaster or accident are more properly handled by disaster recovery teams.

See Also: *incident, incident response team*

incident response team

A team responsible for handling information security incidents when they occur.

Overview

Incident response teams can be either internally developed teams drawn from various departments or an external team brought in under contract. Incident response teams are trained to respond to computer security incidents in a careful, methodical manner that helps the affected company recover quickly from the incident and resume normal business activities as soon as possible. Incident response teams may also deal with legal issues regarding theft of information and may have legal counsel as part of their extended team.

The CERT Coordination Center (CERT/CC), a center of Internet security expertise operated by Carnegie Mellon University, provides training and advice on how to develop computer security incident response teams. CERT/CC refers to an incident response team as a Computer Security Incident Response Team (CSIRT) and offers a one-day course designed for managers tasked with implementing such a team for their companies.

See Also: *CERT Coordination Center (CERT/CC), incident, incident response*

infection

The act of a virus or worm establishing itself in a computer system.

Overview

Infection is the process by which computer viruses and worms cause damage to applications and data stored on computers. Infection can happen many ways:

- By using infected floppy disks borrowed from friends or taken home from work or school
- By downloading infected programs from the Internet, particularly shareware from untrusted sites
- By using pirated software that has previously been infected
- By opening infected attachments to e-mail messages

Once a worm becomes active on a computer, it can infect program or data files by copying and appending itself to files.

See Also: *virus, worm*

information assurance (IA)

Methodologies for ensuring the security of information systems.

Overview

Information assurance is the process of protecting and defending information systems and infrastructures against attack. **Assurance** means confidence that the security features of a product or system fulfill their stated aims, and information assurance provides policies and procedures for developing, testing, and implementing information products in a secure fashion.

Information assurance focuses on five elements of information security:

- Authentication
- Availability
- Confidentiality
- Integrity
- Nonrepudiation

Of these five elements, three of them (confidentiality, integrity, and availability) are often viewed as core ele-

ments of information security (infosec) and are generally referred to as the “CIA triad.” An increasingly popular approach for ensuring information assurance is the Common Criteria & Methodology for Information Technology Security Evaluation (usually called Common Criteria), an international effort to standardize criteria for evaluating the security of information systems outlined in the ISO 15408 standard.

See Also: *Common Criteria & Methodology for Information Technology Security Evaluation, Information Assurance Support Environment (IASE), Information Assurance Technical Framework (IATF), Information Technology Security Evaluation Criteria (ITSEC)*

Information Assurance Support Environment (IASE)

A U.S. Department of Defense (DoD) clearinghouse for information assurance (IA) information.

Overview

Information Assurance Support Environment (IASE) provides training, services, and guidance in procuring and implementing secure information systems. Most IASE services, including the Information Desk and Global Directory, are restricted to users in the .gov and .mil domains and require a digital certificate to access them. Publicly available IASE information includes a list of links related to IA and a list of free training products.

IASE is sponsored by the Defense Information Systems Agency (DISA).

For More Information

Visit iase.disa.mil for more information.

See Also: *information assurance (IA)*

Information Assurance Technical Framework (IATF)

A framework for ensuring the security of information systems.

Overview

The Information Assurance Technical Framework (IATF) was developed by the Information Assurance Technical Framework Forum (IATFF), an outreach activity of the U.S. National Security Agency (NSA), as a guide for developing and implementing secure information systems and protecting information infrastructures. The framework employs a layered defense-in-depth approach with four areas of focus:

- Defend the network and infrastructure
- Defend the enclave boundary
- Define the computing environment
- Support infrastructures

The IATF has been broadly adopted within U.S. government and defense industry, and its goal is to provide a framework for information assurance (IA) solutions in government, industry, and business.

For More Information

Visit www.iatf.net for more information.

See Also: *information assurance (IA)*

information leakage

Obtaining useful information through vulnerabilities in hardware or software.

Overview

Poorly designed hardware or software may “leak” information in unexpected ways, and attackers often can exploit such vulnerabilities to obtain information useful for furthering their exploits. Some of the many ways in which information leakage may occur include the following:

- Electromagnetic radiation from unshielded cabling can be intercepted using radio equipment and analyzed to determine the data being transmitted over the cabling.
- Ethernet drivers often respond to Internet Control Message Protocol (ICMP) echo requests by padding ICMP echo reply messages with kernel memory that can contain bits from traffic on other network segments.

- Blinking light-emitting diode (LED) indicator lights on communications equipment may sometimes be correlated with activities being performed by the device, allowing attackers with physical access to the equipment to gain useful information.
- Welcome messages generated by network communication tools may provide attackers with information about what authentication methods or hashing algorithms are being used.

See Also: *vulnerability*

Information Systems Audit and Control Association (ISACA)

A global organization concerned with information assurance (IA) and control.

Overview

The Information Systems Audit and Control Association (ISACA) is a recognized global leader in the field of IA and auditing and has over 26,000 members in more than 100 countries. Since 1969 the ISACA has provided the IT (information technology) community with training events and conferences, and it maintains for its members a global information repository of security information called K-NET.

The ISACA also administers the recognized standard in information systems auditing certification, the Certified Information Systems Auditor (CISA) designation, which is held by over 29,000 professionals worldwide. A new certification developed by ISACA is the Certified Information Security Manager (CISM) designation, which is geared toward experienced information security (infosec) professionals and covers design, implementation, and management of secure information systems at the conceptual level.

For More Information

Visit www.isaca.org for more information about the ISACA.

See Also: *Certified Information Systems Auditor (CISA)*

Information Systems Security Association (ISSA)

An independent organization of security professionals.

Overview

The Information Systems Security Association (ISSA) is a nonprofit organization that provides education, networking, and leadership opportunities for information security (infosec) professionals worldwide. Local chapters of the ISSA meet in different locations to provide opportunities for professional networking and exchange of information between peers. The ISSA also sponsors regional events and an annual conference, and it is a founding member of the International Information Systems Security Certification Consortium (ISC)². The ISSA also issues publications to enhance professional development of infosec practitioners.

For More Information

Visit www.issa.org for more information about the ISSA.

See Also: International Information Systems Security Certification Consortium (ISC)²

Information Technology Security Evaluation Criteria (ITSEC)

A set of criteria for information security.

Overview

Information Technology Security Evaluation Criteria (ITSEC) is a set of criteria developed by European countries for certifying the level of security an information product or system has. ITSEC evaluation involves demonstrating the compliance of the product being tested with a Security Target, a set of security requirements developed for the product by a commercial licensed evaluation facility (CLEF). The product or system being tested in this process is called a **target of evaluation** (TOE).

ITSEC certification is important for vendors both from marketing and procurement perspectives. By marketing their products as ITSEC-certified, vendors can demonstrate to potential clients their commitment to

information security. From the procurement perspective, many European government agencies require ITSEC-certified products and close their markets to uncertified vendors. Products and systems can be certified at various levels, ranging from E1 (the lowest) to E6 (highest), with assurance and functionality being separated into different levels. The higher the certification level, the greater rigor and attention to detail paid during the certification process.

Notes

The Common Criteria & Methodology for Information Technology Security Evaluation is an International Organization for Standardization (ISO) standard (ISO 15408) that is more widely recognized around the globe than ITSEC, which is primarily a European standard.

For More Information

Visit www.itsec.gov.uk for more information on ITSEC.

See Also: Common Criteria & Methodology for Information Technology Security Evaluation

infosec

Short for information security.

Overview

The term **infosec** is commonly used in several environments:

- Among professionally certified information security practitioners
- In the European context of information security
- In the military

The term is often capitalized as INFOSEC, especially in a military context. Related concepts are COMSEC, which stands for communications security, and RADSEC, which stands for radiation (electromagnetic) security.

See Also: information assurance (IA)

InfraGard

A cooperative effort for protecting critical information security (infosec) infrastructures in the United States.

Overview

InfraGard is an initiative based on an alliance between the following entities:

- The Federal Bureau of Investigation (FBI)
- The National Infrastructure Protection Center (NIPC)
- Local law enforcement agencies
- Business and industry
- Academia

The goal of InfraGard is to facilitate the timely sharing and analysis of information about intrusions, exploits, vulnerabilities, and threats to public and private information systems infrastructures. InfraGard functions using local chapters across the nation and is intended to help infrastructure companies and agencies guard against the following threats:

- Unstructured threats by insiders and recreational hackers
- Structured threats by terrorists, industrial spies, and organized crime (both physical and cyberthreats)
- National security threats

For More Information

Visit www.infragard.net for more information.

See Also: *cybercrime, hacking*

ingress filtering

Blocking incoming traffic whose source address is on the internal network.

Overview

Ingress filtering is a technique that can be used on firewalls and packet-filtering routers to help guard networks against denial of service (DoS) attacks that employ Internet Protocol (IP) address spoofing. Ingress filtering blocks any incoming packets that an attacker has forged to look like they originate from hosts residing on the internal network. Ingress filtering is a recommended practice for Internet services providers (ISPs)

to help protect their client networks from the increasing numbers of DoS attacks occurring on the Internet.

See Also: *denial of service (DoS), firewall, IP address spoofing, packet filtering*

initialization vector

A method of ensuring that initial blocks of ciphertext are always unique.

Overview

Block ciphers convert plaintext into ciphertext using mathematical algorithms. When two pieces of plaintext, such as e-mail messages, begin with identical information, such as message headers, it is important to ensure that the initial portions of ciphertext resulting from application of a block cipher are different. To accomplish this, a random series of bits called an initialization vector is appended to the beginning of the plaintext prior to application of the block cipher. The result is that the initial portion of ciphertext produced is always unique.

See Also: *block cipher, ciphertext, plaintext*

input validation attack

Any attack that exploits poor coding of the algorithms used to check the data that a user or program has entered.

Overview

Input validation is an essential part of good coding practice and involves checking input information to filter out undesirable input to ensure the program or its data cannot be harmed. However, good input validation is complex and difficult to implement, and ingenious programmers sometimes invent ways of circumventing the best input validation routines. The results of circumventing input validation can range from the ability to view files stored in parent directories to being able to run arbitrary code on the server.

There are several types of input validation attacks:

- Including special characters such as wildcards (*), script tags (<script>), directory transversal characters (.../), or escape characters to cause the program

to perform actions not intended to be performed by input information, such as running an executable or script to elevate the attacker's privileges

- Submitting input that is deliberately designed to generate errors, and then using these errors to profile the system; for example, to learn the names of tables on an SQL database server
- Submitting input that is deliberately designed to cause buffer overflows that crash the program, thus denying services to legitimate network users

See Also: *arbitrary code execution attack, dot bug vulnerability, elevation of privileges (EoP)*

insider attack

Compromise of network systems by company employees.

Overview

According to reports by Intergov and other organizations, the majority of information security incidents is perpetrated by insiders (some studies place this figure as high as 80 percent). Insider attacks are potentially more costly and more damaging than those perpetrated by outsiders. They are also harder to detect since they usually bypass firewalls and network intrusion detection systems (NIDSs), which often are not set up to look for such attacks. One way of detecting insider attacks is to set up a honeypot. For example, a decoy "payroll server" could be set up and monitored to trap an employee trying to access or manipulate payroll information in an unauthorized fashion.

While insider attacks may seem on the surface to be an information security (infosec) problem requiring a technical solution, they can sometimes be the result of poor management practices, such as favoritism in promotions, late payment of wages, unresponsiveness to suggestions for improving working conditions, and so on. In reality, however, such attacks are always criminal actions and the perpetrators are liable to be prosecuted.

See Also: *honeypot, intrusion*

integrity

Accuracy and completeness of a received message or retrieved file.

Overview

Integrity is an essential aspect of information security (infosec) and together with confidentiality and availability, it forms the "CIA triad." An information system that protects the integrity of data ensures that it has not been modified in transit (for messages) or during storage (for files). It is essential that intruders not be able to intercept and substitute legitimate data with false or forged data. Integrity of data can also be damaged accidentally by electrical discharges or natural disaster.

Physically securing storage systems and ensuring redundancy can protect the integrity of stored data. Protecting the integrity of transmitted information requires physical protection of transmission media, encryption of information so it can't be read or modified, and the use of checksums to detect when data has been modified.

See Also: *authentication, confidentiality, nonrepudiation*

International Data Encryption Algorithm (IDEA)

A block cipher encryption algorithm developed by Xuejia Lai and James Massey.

Overview

International Data Encryption Algorithm (IDEA) is a computationally fast block cipher that encrypts 64-bit blocks of plaintext into ciphertext blocks of the same size. IDEA uses a 128-bit key and performs encryption in eight rounds using 16-bit subkeys. IDEA was developed in 1991 and is patented by Ascom, a Swiss firm, but the company has been generous in granting permission for free noncommercial use. As a result, IDEA has found its way into popular encryption algorithms such as Pretty Good Privacy (PGP), a popular algorithm for encrypting e-mail. IDEA is considered to be a strong encryption algorithm and has resisted cryptanalytic attack to date.

See Also: *block cipher, encryption algorithm*

International Information Systems Security Certification Consortium (ISC)²

A nonprofit consortium for training and certifying information security (infosec) professionals.

Overview

Since 1989 the International Information Systems Security Certification Consortium (ISC)² has been the leading organization for certifying information security professionals. The (ISC)² administers the respected Certified Information Systems Security Professional (CISSP) and System Security Certified Practitioner (SSCP) standards, certifications that require years of field experience and passing rigorous exams to acquire. The (ISC)² also partners with other institutions, including the Information System Security Association (ISSA) and various academic and national information security organizations.

For More Information

Visit www.isc2.org for more information.

See Also: *Certified Information Systems Security Professional (CISSP), Information Systems Security Association (ISSA), System Security Certified Practitioner (SSCP)*

Internet Key Exchange (IKE)

The key management protocol used by Internet Protocol Security (IPSec).

Overview

Internet Key Exchange (IKE) defines methods for the endpoints of an intended IPSec session to mutually authenticate one another. IKE is a complex specification that involves several pieces:

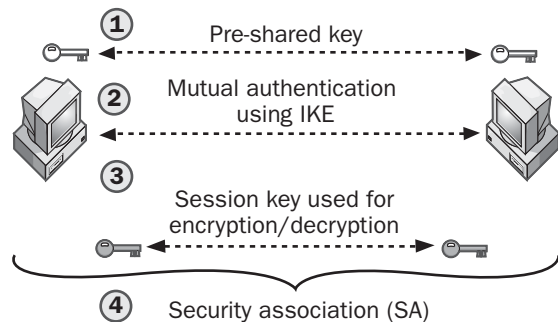
- Domain of Interpretation (DOI), defined in RFC 2407
- Internet Security Association and Key Management Protocol (ISAKMP), defined in RFC 2408
- IKE itself, defined in RFC 2409
- OAKLEY, defined in RFC 2412

All of these standards are interrelated, and as a result IKE is sometimes known as ISAKMP/IKE or ISAKMP/Oakley.

Implementation

IKE works in two phases:

- **Phase 1:** Mutual authentication of the two endpoints is performed using the preshared key, and two unique session keys are generated: an encryption key and an integrity key. The preshared key may be a shared secret key, a public encryption key, or a public signature-only key. The key exchange process can be performed two ways: aggressive mode or main mode.
- **Phase 2:** A security association (SA) is established between the endpoints using a key exchange process called quick mode, which negotiates the method used to encrypt information for secure communication between the endpoints.



Internet Key Exchange (IKE). How IKE uses a preshared key to generate a unique session key.

Issues

IKE suffers from several shortcomings that have plagued it since inception. These issues include the following:

- The high degree of complexity and even obscurity of portions of the Internet Engineering Task Force (IETF) standards defining IKE have resulted in interoperability problems with implementations from different vendors.

- The chatty nature of the negotiation methods used by IKE makes IPSec sessions vulnerable to denial of service (DoS) attacks.

As a result of these shortcomings, the IETF has been considering various replacements for IKE, including these:

- Internet Key Exchange version 2 (IKEv2)
- Just Fast Keying (JFK)
- Sigma

These replacements simplify IKE by reducing the number of features and restricting various options, resulting in key exchange methods that are more restrictive but simpler to implement. For example, the replacements will eliminate support for preshared keys and will support only digital signatures for authentication (IKE allows preshared keys and supports other authentication methods such as Remote Authentication Dial-In User Service protocol, or RADIUS, and electronic tokens). The result should be safer virtual private networks (VPNs) since there will be less opportunity for the kind of configuration errors that can happen because of IKE's complexity.

See Also: Internet Key Exchange version 2 (IKEv2), Internet Protocol Security (IPSec), Just Fast Keying (JFK)

Internet Key Exchange version 2 (IKEv2)

A proposed replacement for Internet Key Exchange (IKE).

Overview

Internet Key Exchange version 2 (IKEv2) is one of several proposed replacements for IKE, the key management protocol used by Internet Protocol Security (IPSec). IKEv2 preserves most of the key features of IKE but is easier to implement and less vulnerable to denial of service (DoS) attacks. While IKE supports eight different initial negotiation methods, IKEv2 supports only a single negotiation method. This reduced flexibility facilitates implementation of IKEv2; therefore, it is less likely to result in the

vendor interoperability problems that have affected IKE since its inception.

See Also: Internet Key Exchange (IKE), Internet Protocol Security (IPSec), Just Fast Keying (JFK)

Internet Protocol Security (IPSec)

Security extensions for Internet Protocol (IP).

Overview

Internet Protocol Security (IPSec) is a suite of network-layer protocols that extends IP by providing mechanisms for authentication, confidentiality, and integrity in IP communications. With the use of IPSec, a communication session between two hosts can be encrypted in a way that is transparent to applications running on the hosts. IPSec is widely used for implementing virtual private networks (VPNs) and in places where information security is a high priority.

Implementation

IPSec has two security protocols that can be implemented separately or together:

- **Authentication Header (AH):** Performs authentication of sender only. Authentication can be performed using Message Digest 5 (MD5), hash-based message authentication code (HMAC), or Secure Hash Algorithm-1 (SHA-1).
- **Encapsulating Security Protocol (ESP):** Performs both authentication of sender and encryption of data. Authentication can be performed using the algorithms described previously, while encryption can be performed using Digital Encryption Standard (DES), Triple DES (3DES), Blowfish, International Data Encryption Algorithm (IDEA), Cast, RC5, and other algorithms.

IPSec encryption can be implemented using two different modes:

- **Transport mode:** Only the payload (data portion) of a packet is encrypted, while the header remains unencrypted.
- **Tunnel mode:** Both the packet header and payload are encrypted.

To establish an IPSec security association (SA) between two hosts, the hosts must have previously shared a key (secret or public) or digital certificate. Key management in IPSec is performed using the Internet Key Exchange (IKE) protocol, which is sometimes referred to as ISAKMP/Oakley.

Notes

IPSec is defined in RFCs 2401 through 2412.

See Also: 3DES, Authentication Header (AH), Blowfish, Encapsulating Security Payload (ESP), Hashed-based message authentication code (HMAC), International Data Encryption Algorithm (IDEA), Internet Key Exchange (IKE), message digest (MD), message digest 5 (MD5), Secure Hash Algorithm-1 (SHA-1), virtual private network (VPN)

Internet Security and Acceleration (ISA) Server

Microsoft Corporation's firewall and secure application gateway product.

Overview

Internet Security and Acceleration (ISA) Server is Microsoft's International Computer Security Association (ICSA)-certified firewall product designed to protect enterprise networks from attack by intruders, worms, and other threats. ISA Server provides several layers of protection, including packet filtering, application-level filtering, stateful inspection, and an advanced proxy architecture. ISA Server also increases performance through Web caching to reduce network congestion and save bandwidth costs. ISA Server can restrict access by users and groups, type of application, content type, time of day, and destination sets. It also includes integrated logging, monitoring, alerting, and reporting features to help administrators block threats as they are detected.

For More Information

Visit www.microsoft.com/isaserver/ for more information about ISA Server.

See Also: *firewall*

intrusion

An attempt to compromise a system or network.

Overview

Intrusions are attempts by malicious individuals to discover and exploit vulnerabilities that may be used to compromise network security. Any suspicious network traffic that falls outside of normal or legitimate traffic patterns may be classified as an intrusion. The results of intrusion can take different forms, including the following:

- Destruction or theft of data
- Denial of service (DoS) to legitimate network users
- Hijacking of systems and communication sessions

To determine when an intrusion is taking place, an intrusion detection system (IDS) may be used.

See Also: *exploit, intrusion detection system (IDS), vulnerability*

intrusion detection system (IDS)

An application or device that identifies suspicious network activity.

Overview

An intrusion detection system (IDS) inspects inbound and outbound traffic on a host or network, analyzing it and looking for evidence of attempted intrusion. IDSs are of two basic types:

- **Network-based IDS (NIDS):** All traffic flowing through the network is analyzed for evidence of attempted intrusion. NIDS usually resides at a choke point on the perimeter of the network or on critical network segments where the servers reside. A limitation of an NIDS is that it is difficult to implement in switched networks, though some Ethernet switch vendors are starting to incorporate embedded IDS within switches and provide monitoring ports for connecting a NIDS to the switch's backplane.
- **Host-based IDS (HIDS):** The activity of an individual network host is monitored for evidence of attempted intrusion. HIDSs are usually placed on critical servers such as firewalls, mail servers, and Web servers exposed to the Internet.

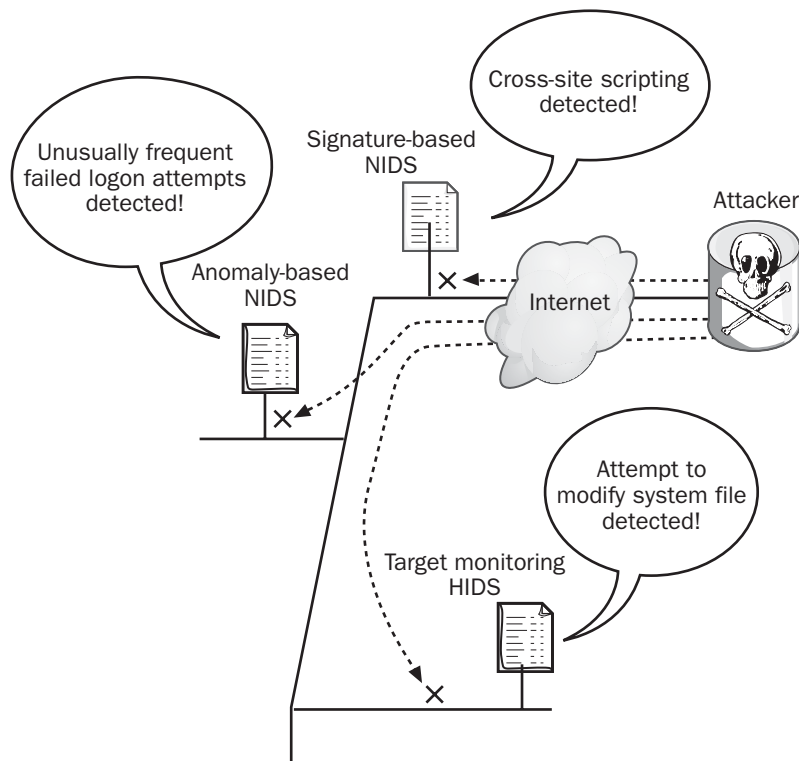
Most IDS products are passive systems whose job is merely to detect evidence of intrusion and alert administrators to possible attacks on their network. Recently, vendors have begun to develop reactive IDS products that can perform actions to protect against attacks when they are detected. Such actions might include closing certain ports or blocking certain Internet Protocol (IP) addresses. A reactive IDS product is sometimes called an intrusion prevention system (IPS). Some IDS products combine host- and network-based detection and are sometimes called hybrid systems.

Implementation

There are various approaches to how IDS products work. The most popular method is signature detection, which involves matching network traffic to a database of thousands of known intrusion signatures. Important to implementation of such systems is regular updating of the signature database.

Another approach is anomaly detection. This involves looking for unusual traffic patterns that may indicate an attack in progress and is generally accomplished using statistical techniques to compare current traffic with a baseline of normal traffic established previously. Anomaly detection has an advantage over signature detection in that it is able to detect new and undocumented forms of attack. The downside is that if the threshold for detection is set too high, large numbers of false positives are generated. The job of the administrator is thus complicated by having to sort out the real from the false events.

Another technique used in intrusion detection is monitoring file systems to look for attempts to replace or modify key system or log files. HIDS products generally incorporate this kind of approach, in addition to scanning system logs for unusual events.



Intrusion detection system (IDS). Ways in which intrusions can be detected.

Marketplace

Market leaders among NIDS products include Real-Secure from Internet Security Systems (ISS), Cisco Secure Intrusion Detection System from Cisco (which acquired NetRanger from WheelGroup), and eTrust Intrusion Detection from Computer Associates (which acquired SessionWall-3 from MEMCO).

Leaders in the HIDS market segment include Intruder Alert from Symantec, Computer Misuse Detection System from ODS Networks, and Kane Security Monitor from Security Dynamics.

Other IDS vendors include Cybersafe, Network Associates, Network Flight Recorder, Intellitactics, SecureWorks, and Security Wizards. The open source Snort is also popular as an intrusion detection tool.

Notes

The Internet Engineering Task Force (IETF) has started a working group for standardizing intrusion detection methodologies in order to promote interoperability between different IDS vendors.

See Also: *false positive, host-based intrusion detection system (HIDS), intrusion, intrusion prevention system (IPS), network-based intrusion detection system (NIDS), Snort*

intrusion prevention system (IPS)

An intrusion detection system (IDS) that can also react to intrusions by blocking them.

Overview

Traditional IDSs are passive systems that can detect intrusions but do nothing to block them. Instead, it is up to the administrators to review IDS logs and respond to alerts, closing ports on firewalls and taking other steps to prevent intruders from gaining a foothold.

A recent trend in firewall products is for vendors to include reactive intrusion detection technology that can automatically reconfigure the firewall when an intrusion is detected. Such products are sometimes called

intrusion prevention systems (IPSs) and indicate a blurring of the line separating firewall products from IDS platforms.

Marketplace

Examples of vendors who have implemented IPS features into their firewall products include Check Point Software, Cisco, and NetScreen. Vendors of Ethernet switches and load balancers are also beginning to incorporate IPS into their products as well.

See Also: *firewall, intrusion detection system (IDS)*

IP address-based authentication

Authenticating hosts based on their Internet Protocol (IP) addresses.

Overview

IP address-based authentication is an authentication method that uses the IP address of a remote host to determine whether that host should be able to access network services or other resources. IP address-based authentication is widely used on UNIX platforms where applications such as Rsh and Rlogin authenticate remote hosts based on information stored in .rhosts and other configuration files. IP address-based authentication is considered a weak authentication method since attackers may be able to circumvent such restrictions by spoofing the source addresses of IP packets.

See Also: *authentication, IP address spoofing, Rlogin, spoofing*

IP address restriction

Controlling access through Internet Protocol (IP) addresses.

Overview

IP address restriction is a method of controlling access to resources based on the IP address of the host trying to establish access. For example, Internet Information Services (IIS), which can restrict access to Web content for individual addresses or blocks of addresses defined by

network ID and subnet mask, uses IP address restriction. Access may then be either allowed or denied for each address or block of addresses. Another example is the Apache Web server, where access to Web content can be controlled using IP addresses by configuring the `.htaccess` file on UNIX platforms.

IP address restriction is considered a weak form of access control since attackers may be able to circumvent such restrictions by spoofing the source addresses of IP packets.

Notes

A related method of controlling access is domain name restriction, which restricts access based on the Domain Name System (DNS) domain to which the host trying to obtain access belongs.

See Also: *access control, .htaccess, IP address spoofing, Rlogin, spoofing*

IP address spoofing

The process of falsifying the source Internet Protocol (IP) address of IP packets.

Overview

IP address spoofing (or simply, IP spoofing) is a method used by intruders to impersonate trusted systems. By default, routers generally ignore source IP addresses when routing packets, and they use only destination IP addresses to ensure packets reach their intended destination. The result is that an attacker who forges IP packets containing source addresses of trusted systems may be able to circumvent router security and initiate denial of service (DoS) attacks, redirect traffic, or hijack sessions using man-in-the-middle (MITM) attacks.

IP spoofing is especially a hazard on UNIX platforms running such applications as Rsh or Rlogin that authenticate connections using source IP addresses stored in `.rhosts` files. IP address authentication is a weak form of authentication supported by many UNIX applications and should be replaced by password authentication to ensure security.

The standard approach for preventing IP spoofing attacks is to configure ingress filters on routers or firewalls in order to deny any inbound traffic whose source address is from a trusted host on your internal network. When an intrusion detection system (IDS) detects such traffic, there is a high probability that a spoofing attack is under way. Encryption of traffic between routers and external hosts is another effective way of protecting against spoofing attacks.

Notes

Tools used by attackers to launch spoofing attacks include Dsniff, Hunt, Ipspoof, and Spoofit.

See Also: *Dsniff, ingress filtering, .rhosts, spoofing*

IP fragmentation attack

An attack that uses fragmented Internet Protocol (IP) packets.

Overview

The IP standard supports fragmentation to allow IP packets to traverse different types of transmission media, for example, to travel between two local area networks (LANs) over a wide area network (WAN) connection. Fragmentation can also be used to attack IP hosts, however, and by deliberately crafting fragmented IP packets, it may be possible for attackers to circumvent firewall protection, hide traffic from intrusion detection systems (IDSs), or create denial of service (DoS) conditions to prevent legitimate users from accessing network services.

Early forms of fragmentation attacks were able to circumvent firewall restrictions because of the fact that firewall products didn't apply their rules until fragmented packets had been reassembled. As a result, firewall products were found to be vulnerable to DoS attack by continually sending them large numbers of forged initial fragments until the internal resources of the firewall were consumed. Tools used to initiate such attacks included Jolt2, Teardrop, and Nmap. Most firewall vendors have since modified their products to protect against such attacks. A tool called Fragrouter can

be used to test a firewall or IDS to see whether it is vulnerable to a whole series of different types of fragmentation attacks.

See Also: *denial of service (DoS), fragmentation, Jolt2, Nmap, Teardrop attack*

Iplog

An open source tool for logging Transmission Control Protocol/Internet Protocol (TCP/IP) traffic.

Overview

Iplog can be used for logging various types of TCP/IP traffic, including TCP, User Datagram Protocol (UDP), and Internet Message Control Protocol (ICMP) packets. The tool is useful for detecting various types of intrusions and attacks, including port scans, null scans, ping floods, and fragmentation attacks. You can also configure Iplog to run in promiscuous mode so that it monitors all network traffic on a segment and not just traffic on the local host.

Iplog is available for BSD, Linux, and Solaris platforms and is released under the General Public License (GPL).

For More Information

Visit www.sourceforge.net to download Iplog and other open source security tools.

See Also: *port scanning, promiscuous mode*

IPS

Stands for intrusion prevention system, an intrusion detection system (IDS) that can also react to intrusions by blocking them.

See: *intrusion prevention system (IPS)*

IPSec

Stands for Internet Protocol Security, security extensions for Internet Protocol (IP).

See: *Internet Protocol Security (IPSec)*

IPSec filter

A rule for filtering Internet Protocol (IP) traffic.

Overview

IPSec filters are rules that can be created in Internet Protocol Security (IPSec) policies to filter different types of IP traffic. Rules can either allow or deny traffic and can filter according to protocol type, source or destination address, or port number. Rules can apply to inbound traffic, outbound traffic, or both. You can create and manage IPSec filters using Group Policy or from the command line using the ipsecpol.exe utility.

See Also: *Internet Protocol Security (IPSec), IPSec policy*

IPSec policy

A policy for implementing Internet Protocol Security (IPSec).

Overview

IPSec policies specify authentication methods, encryption schemes, and filter actions for implementing secure network communication using IPSec. On Microsoft Windows Server 2003 and Windows 2000 platforms, IPSec policies are part of Group Policy and are stored in Active Directory directory service. An IPSec policy can contain one or more IPSec filters, providing more granular control over IP traffic than Transmission Control Protocol/Internet Protocol (TCP/IP) filtering in previous versions of Windows.

Windows Server 2003 and Windows 2000 include three default IPSec policies:

- **Client (Respond Only):** Used by workstations to respond to authorization requests from servers
- **Server (Request Security):** Used by servers in environments that contain systems that are not IPSec-aware to allow negotiation of authentication and encryption levels
- **Secure Server (Require Security):** Used by servers in environments that contain systems that are all IPSec-aware to deny all nonauthorized and unencrypted network traffic

See Also: *Internet Protocol Security (IPSec), IPSec filter*

IP spoofing

Short for Internet Protocol (IP) address spoofing, the process of falsifying the source IP address of IP packets.

See: IP address spoofing

ISACA

Stands for Information Systems Audit and Control Association, a global organization concerned with information assurance (IA) and control.

See: Information Systems Audit and Control Association (ISACA)

ISA Server

Stands for Internet Security and Acceleration Server, Microsoft's firewall and secure application gateway product.

See: Internet Security and Acceleration (ISA) Server

(ISC)²

Stands for International Information Systems Security Certification Consortium, a nonprofit consortium for training and certifying information security (infosec) professionals.

See: International Information Systems Security Certification Consortium (ISC)²

island-hopping

Using one compromised system or network to break into another.

Overview

One of the goals of an intruder who has compromised a system is to look for opportunities that could be exploited for compromising other targets. This practice is called **island-hopping** after the way the U.S. military captured one island after another in the Pacific during World War II. Common examples of island-hopping can include

- Cracking dial-up or remote access passwords to attack a branch office over a wide area network (WAN) connection
- Using a cracked local administrator password to obtain domain credentials that could be used to attack a remote trusted domain
- Attacking the network of an Internet service provider (ISP) from a compromised user's computer over a high-speed digital subscriber line (DSL) connection
- Compromising a router and then using spoofed routing protocol packets to attack other routers

See Also: hacking

ISO 17799

An international standard outlining best practices for information security.

Overview

ISO 17799 takes a generic approach to ensuring information security by outlining best practices for different aspects of information handling. The 10 areas of control outlined by this standard are as follows:

- Asset classification and control
- Business continuity planning
- Compliance
- Computer and operations management
- Personnel security
- Physical and environmental security
- Security organization
- Security policy
- System access control
- System development and maintenance

Compliance with these practices is the first step in achieving ISO 17799 certification, which is quickly becoming the internationally recognized security stan-

dard in industry and commerce. The goal of the standard is to facilitate electronic business by creating a trusted environment between certified partners.

For More Information

Visit www.iso-17799.com for more information on the ISO 17799 standard.

See Also: infosec

ISSA

Stands for Information Systems Security Association, an independent organization of security professionals.

See: Information Systems Security Association (ISSA)

itrace

Stands for ICMP Traceback, a proposed modification to Internet Control Message Protocol (ICMP) that would enable Internet Protocol (IP) traffic to be traced to its source.

See: ICMP Traceback (itrace)

ITSEC

Stands for Information Technology Security Evaluation Criteria, a set of criteria for information security.

See: Information Technology Security Evaluation Criteria (ITSEC)

JFK

Stands for Just Fast Keying, a proposed replacement for the Internet Key Exchange (IKE) protocol.

See: Just Fast Keying (JFK)

Jill

A tool for running arbitrary code on a machine running Microsoft Windows 2000.

Overview

Jill is an exploit that can open a remote shell on a machine running Windows 2000 and Internet Information Services (IIS) 5 and listening on port 80. Jill exploits a buffer overflow to start a shell in the security context of the LocalSystem account, allowing arbitrary code to run on the remote machine. Keeping Windows 2000 servers up to date with patches issued by Microsoft can prevent this from occurring.

Jill is written in C code and runs on UNIX/Linux platforms. Several related tools exist, including Jill-win32 (a Windows-based version of the exploit) and Iis5hack.

See Also: buffer overflow, exploit

John the Ripper

A popular password-cracking tool.

Overview

John the Ripper is a password-cracking tool available for the UNIX/Linux, OpenVMS, and Microsoft Windows platforms. This command-line tool is dictionary-based and can crack several popular encryption algorithms. It includes numerous rules for permuting dictionary entries to guess passwords that might be thought difficult to crack.

The intended use for this tool is to be able to detect weak UNIX passwords, but in practice its main use is for password cracking. The architecture of the tool is

extensible, allowing custom cracking modes to be defined using C code.

For More Information

Visit www.openwall.com/john/ for more information.

See Also: dictionary attack, password cracking

Jolt2

A denial of service (DoS) attack based on Internet Protocol (IP) packet fragmentation.

Overview

Jolt2 is an exploit that uses a stream of malformed fragments to drive the central processing unit (CPU) utilization of target hosts to 100 percent as they try to process the fragments. The result is that legitimate users are denied access to services on the target machine. Jolt2 appeared in 2000 and is the successor to the earlier Jolt exploit of 1997, which affected only Microsoft Windows 95 and Windows NT 4 systems. The new exploit, however, affected a much wider range of platforms, including Cisco routers, firewall products such as Checkpoint Software's Firewall-1 and Network Associate's Gauntlet, and all Microsoft Windows versions. Keeping these platforms up to date with patches released by the vendors can prevent Jolt2 exploits from happening.

See Also: fragmentation, IP fragmentation attack

Juggernaut

An open source packet-sniffing tool.

Overview

Juggernaut is a free sniffing tool that can be used to capture and hijack Transmission Control Protocol (TCP) sessions and kill connections. Juggernaut is open source software released under the General Public License (GPL) and runs on the Linux platform. In 1997, *Phrack* magazine originally released version 1 of Juggernaut.

For More Information

Visit www.phrack.org for more information.

See Also: *hijacking, packet sniffer*

Just Fast Keying (JFK)

A proposed replacement for the Internet Key Exchange (IKE) protocol.

Overview

Just Fast Keying (JFK) is one of several proposed replacements for IKE, the key management protocol used by Internet Protocol Security (IPSec). JFK is intended to overcome the deficiencies of IKE, which include its vulnerability to denial of service (DoS)

attacks, its complexity of operation, and its “chatty” nature (high number of rounds). To accomplish this, JFK has the following simplifications over IKE:

- JFK uses only one phase compared to two for IKE, making it much simpler to implement.
- JFK uses only two rounds with no option for additional rounds, which greatly reduces the chattiness of the protocol.

JFK has an architecture that resists memory and processor exhaustion attacks, making it less susceptible than IKE to DoS attacks.

See Also: *Internet Key Exchange (IKE), Internet Protocol Security (IPSec)*

KDC

Stands for key distribution center, which in Kerberos describes an entity that grants tickets to clients.

See: *key distribution center (KDC)*

Kensington security slot

A physical connector found on laptop computers that is used to link locks and cables developed by Kensington.

Overview

Surveys show that about 10 percent of business laptops are stolen each year, so laptop security is an essential part of protecting your company's assets. The Kensington security slot allows laptops to be physically secured using cables locked at one end to the machine and secured at the other end to some fixed structure, such as a desk or wall. Almost all laptops today include this security slot as a standard feature, and business users are well advised to make use of it whenever possible.

For More Information

Visit www.kensington.com/html/1356.html for the specifications of the Kensington security slot.

See Also: *physical security*

Kerberos

An authentication protocol developed by the Massachusetts Institute of Technology.

Overview

Kerberos was developed in the 1980s as a method for authenticating users on a large, distributed network. It uses secret key encryption with strong keys so that clients can both prove their identity to servers and also ensure the privacy and integrity of their communications with servers. The protocol is named after Kerberos, the three-headed dog in Greek mythology that

guarded the gates of Hades. The current version of the protocol is Kerberos version 5, outlined in RFC 1510 and described in the following section, and it has been implemented in many commercial platforms including Microsoft Windows 2000.

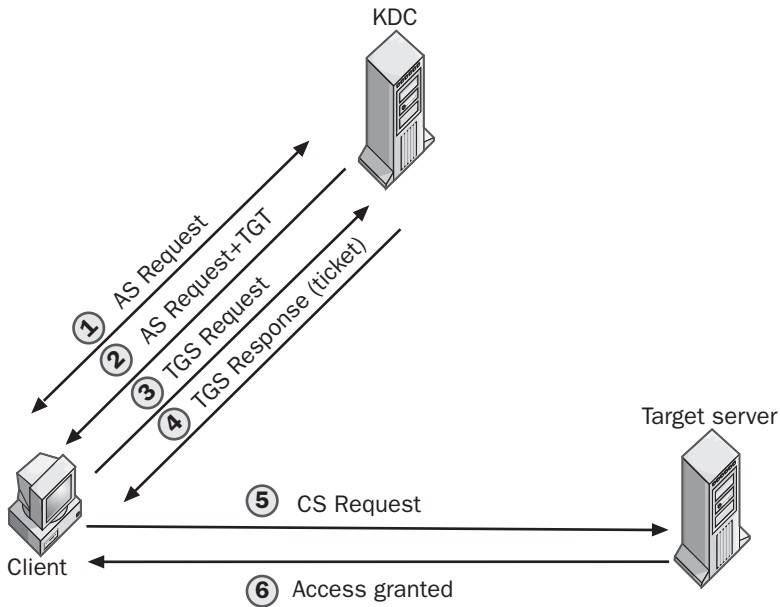
Implementation

Kerberos uses three subprotocols for its operation:

- **Authentication Service (AS) Exchange:** Used by the key distribution center (KDC) for providing clients with ticket-granting tickets (TGTs) and logon session keys
- **Ticket-Granting Service (TGS) Exchange:** Used by the KDC to distribute service session keys and their associated tickets
- **Client/Server (CS) Exchange:** Used by the client to present a ticket for admission to a service

A typical Kerberos authentication session between a client workstation and a network server looks like this:

- 1- The user's credentials are entered on the client, which submits a request to the KDC to access the TGS using the AS Exchange protocol. The request includes encrypted proof of the user's identity.
- 2- The KDC receives the request, looks up the master key of the user in Active Directory directory service, and decrypts the identify information contained in the request. If the user's identity is verified, the KDC responds by granting the user a TGT and a session key using the AS Exchange protocol.
- 3- The client then sends the KDC a TGS request containing the TGT granted earlier and requesting access to some service on a target server using the TGS Exchange protocol.



Kerberos. How Kerberos authentication works.

- The KDC receives the request, authenticates the user, and responds by granting the user a ticket and a session key for accessing the target server using the TGS Exchange protocol.
- The client then sends the target server a request containing the ticket granted earlier using the CS Exchange protocol. The server authenticates the ticket, replies with a session key, and the client can now access the server.

See Also: authentication, key distribution center (KDC)

Kerberos policy

Group Policy settings for Kerberos authentication in Microsoft Windows 2000.

Overview

Kerberos policy defines Kerberos settings for domain user accounts. These settings are stored in Active Directory directory service as part of domain security policy within Group Policy on Windows 2000. Kerberos pol-

icy includes settings covering maximum ticket lifetime, maximum lifetime for ticket renewal, maximum tolerance for computer clock synchronization, and enforcement of user logon restrictions.

See Also: Active Directory, Group Policy, Kerberos

key

A binary number used with an encryption algorithm.

Overview

An **encryption algorithm** is a mathematical procedure for converting plaintext into ciphertext. To eliminate the need to devise a new algorithm each time text must be encrypted, a numeric value called a key is used in conjunction with the algorithm. This way, the details of the algorithm can be made publicly known, while either the key can be kept secret or a new key can be generated each time encryption is required.

Keys come in several types:

- **Secret keys:** Also called symmetric keys, these are keys used with secret or symmetric encryption algorithms such as Data Encryption Standard (DES) and Advanced Encryption Standard (AES). To use such algorithms, both parties in a communication session must share a copy of the same key, which is sometimes called a shared secret.
- **Private and public keys:** These are keys used with public or asymmetric encryption algorithms such as the Rivest-Shamir-Adleman (RSA) algorithm and always come in pairs, with the private key known to only the owner and the public key available to everyone.
- **Session keys:** These are keys whose lifetime is restricted to a single communication session or part of a session. Session keys are generally secret keys that are created and exchanged using a public key algorithm in order for the two parties to encrypt a communication session.

Issues

Keys are fundamental to encrypted communication, but keys can be cracked if they aren't strong enough. The strength of an encryption key is related to its length; the longer the key, the harder it is to crack encryption performed using the key. In general, to provide the same level of security, keys for asymmetric or public key encryption systems must be larger (have more bits) than those used for symmetric or secret key encryption. For example, AES supports key lengths of 128, 192, and 256 bits; even the weakest keys of 128 bits currently are considered to be uncrackable. By comparison, 512-bit keys for the asymmetric RSA algorithm are considered crackable, so longer keys of 786 bits for individuals, 1024 bits for businesses, and 2048 bits for certificate authorities (CAs) are recommended.

See Also: encryption algorithm, key pair, private key, public key, secret key, session key

key distribution center (KDC)

In Kerberos, an entity that grants tickets to clients.

Overview

In a standard implementation of the Kerberos protocol, a key distribution center (KDC) hosts two services:

- **Authentication service (AS):** This service issues ticket-granting tickets (TGTs) to clients that must connect to the ticket-granting service (TGS) in their own domain or trusted domains.
- **Ticket-granting service (TGS):** This service issues tickets to clients that must access computers in their own domain or trusted domains.

In Microsoft Corporation's implementation of Kerberos, the KDC for a domain is located on domain controllers with the Kerberos account database stored in Active Directory directory service.

See Also: Kerberos, ticket

keyed hash

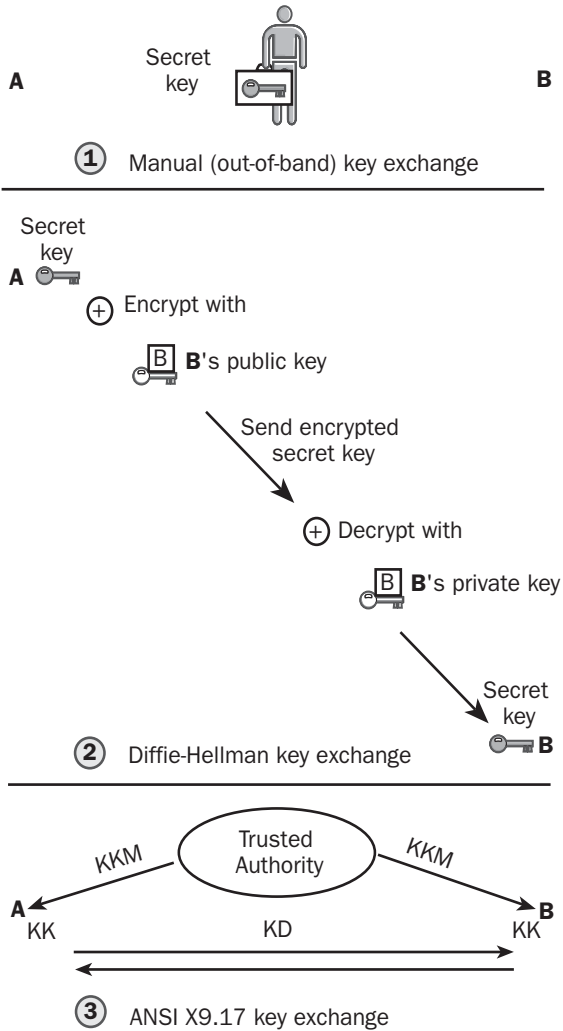
Combination of a hashing algorithm and a secret key.

Overview

A hashing algorithm is a mathematical procedure that generates from an arbitrary message a fixed-size result called a hash. To increase the security of the resultant hash, a secret key can be combined with the message prior to application of the hashing algorithm. The result is a keyed hash that can be calculated only by a user who knows the key.

Implementation

Keyed hashes are often used to generate a message authentication code (MAC) to ensure the integrity of messages being transmitted over insecure media. The sender appends a shared secret key to a message and hashes the result to produce a keyed hash. The sender then transmits the message together with the keyed hash to the recipient, who can verify the integrity of the message by creating a second keyed hash from the message using the same shared secret key and comparing this to the keyed hash sent with the message. If the two keyed hashes are the same, the recipient can be satisfied that the message was not tampered with in transit.



Keyed hash. How a keyed hash can be used to verify the integrity of a message sent over an insecure medium.

See Also: hashing algorithm, integrity, message authentication code (MAC), secret key

keyed-hash message authentication code

Another name for hash-based message authentication code (HMAC), a message authentication code (MAC)

algorithm that combines a hashing algorithm with a secret key.

See: hash-based message authentication code (HMAC)

key escrow

Providing a trusted third party with copies of cryptographic keys.

Overview

In order to prevent criminals and terrorists from communicating using encryption, governments may require that commercial cryptographic hardware and software implement key escrow, a method that provides law enforcement agencies with a “backdoor” to decrypt encrypted communications when necessary. The simplest form of key escrow is to require that all master keys for cryptographic systems, such as the private key of a certificate authority (CA) in a Public Key Infrastructure (PKI), have copies stored with trusted third parties that hold these keys “in escrow” for law enforcement agencies.

The idea of key escrow sometimes poses a concern for civil liberties advocates, who view it as an erosion of individual privacy. An example is the ongoing debate over the Clipper chip, a hardware-based encryption technology proposed by the U.S. government in 1994 and defined in the FIPS 185 Escrowed Encryption Standard (EES). Clipper is based on the classified Skipjack algorithm developed by the National Security Agency (NSA), and the original idea of the proposal was to make the inclusion of Clipper mandatory in computers, modems, telephones, television sets, and other communication devices. Since Clipper included built-in escrow technology, this would have provided government and law enforcement with unprecedented monitoring capability of all forms of electronic communications. Opposition from civil liberties organizations, computer manufacturers, and communications industries has delayed the implementation of Clipper, but in the wake of September 11, some lawmakers have renewed their efforts to mandate such technologies.

Implementation

Key escrow can be implemented various ways:

- By storing copies of entire keys in escrow (plain escrow) so that authorities have immediate access to them when required

- By storing only part of the key in escrow (partial-key escrow) so that authorities must expend computational effort to recover a key
- By splitting keys into two or more portions and distributing them to different escrow agents (shared escrow) so that authorities have to go through several legal steps in order to recover the portions and reassemble the key
- By using an authority's own public key to encrypt the session keys used to encrypt communications and then storing the encrypted session key in escrow (key encapsulation) so that authorities can decrypt each individual session as required but do not have the general ability to decrypt all communications by the user

See Also: certificate authority (CA), key, key recovery, Public Key Infrastructure (PKI)

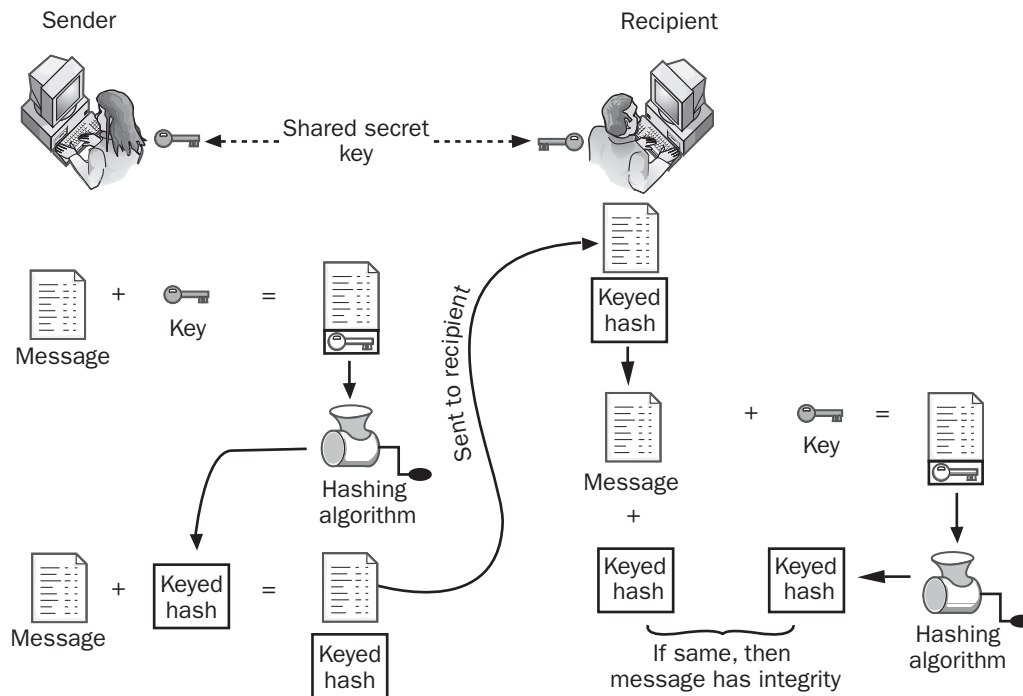
key exchange

Any method for sharing a secret key between two parties.

Overview

Symmetric (or secret) key encryption requires that the two parties involved share the same secret key. The main problem with this system is securely distributing the secret key, and there are various ways of doing this:

- **Out-of-band:** The secret key is distributed using a separate communication channel considered to be secure, for example, by hand delivery, registered mail, or some other method. This is the oldest method and can be quite secure but also expensive and time-consuming.



Key exchange. Three methods of exchanging secret keys between two parties.

- **Diffie-Hellman (DH):** Public key encryption is used to encrypt the secret key and transmit it to the second party. Once decrypted, the secret key can then be used as a session key for secure communication. This method is popular in many electronic-messaging systems.
- **ANSI X9.17:** This method is used in the financial industry and involves using a hierarchy of keys. At the top are master keys (KKMs), which are distributed manually and have long life spans. These KKMs are then used to encrypt key-encrypting keys (KKEs), which are distributed electronically and have shorter life spans. Once financial partners have copies of KKEs, they use them to exchange data keys (DKs), which are used for encrypting and decrypting messages for a single communication session.

See Also: *Diffie-Hellman (DH), key, secret key*

K

key management

An umbrella term describing various processes used for managing cryptographic keys.

Overview

Keys are essential to cryptography, and in order to prevent unauthorized entities from intercepting, decrypting, or hijacking encrypted communications, keys must be protected and managed appropriately. Some of the aspects involved in key management include the following:

- **Key generation:** Creating new keys when they are needed
- **Key storage:** Secure and safe storage of cryptographic keys
- **Key distribution:** Making public keys available to all who need them
- **Key exchange:** Methods for sharing a secret key so two parties can encrypt communication
- **Key revocation:** Mechanisms for revoking a key should it become lost or compromised
- **Key recovery:** Methods for recovering keys when they are lost or damaged

- **Key escrow:** Providing trusted authorities with access to keys for legal or supervisory requirements

See Also: *key, key escrow, key exchange, key recovery*

key pair

A mathematically related pair of cryptographic keys.

Overview

Key pairs are used in public key cryptography systems for which two keys are needed to encrypt or digitally sign messages. The two keys in a key pair are as follows:

- A private key possessed by the entity that owns it and known only to that entity
- A public key registered with a key distribution center (KDC) or certificate authority (CA) and available to anyone who requests it

These two mathematically related keys are generated at the same time. However, it is not computationally feasible to try to derive one of the keys from the other.

See Also: *key, private key, public key, public key cryptography, secret key, session key*

key recovery

Any method for re-creating a cryptographic key if it is lost, stolen, or damaged.

Overview

Key recovery is an essential part of key management for any cryptographic system, since if users lose their keys or have them stolen, their encrypted data would be inaccessible. Cryptographic storage systems therefore include key recovery agents that can be used to restore a lost or damaged key and decrypt otherwise indecipherable data. An example of a cryptographic storage system that employs recovery agents is Microsoft Corporation's Encrypting File System (EFS).

Notes

The term **key recovery** is sometimes used to describe key escrow, the process of providing a trusted third party with copies of cryptographic keys.

See Also: *Encrypting File System (EFS), key, key escrow, key management*

key ring

A data structure for storing public keys.

Overview

In some cryptographic schemes, users have key rings that contain the public keys of other users with whom they wish to communicate with encrypted messages. Key rings may also contain personal information of these other users and digital certificates used to sign documents. Different levels of trust may also be assigned to each key or certificate within the key ring. Users may also be allowed to share their key rings with other users to build a community database of trusted users.

An example of a cryptographic system in which key rings are used is Pretty Good Privacy (PGP), a popular encryption scheme used for sending encrypted e-mail.

See Also: *key, Pretty Good Privacy (PGP)*

key rollover

Changing keys during a cryptographic communication session.

Overview

Cryptographic keys generally have a useful lifetime before they become susceptible to cracking or misuse. That's why most cryptographic systems cause keys to expire after a period of time. Even with strong keys, it can be a good idea to change keys frequently to guard against attackers who might intercept an encrypted session and try to launch a man-in-the-middle (MITM) attack to hijack the session. For the highest level of security, keys can be changed repeatedly during a single communication session between two hosts, even to the extent of using a new key for each block of plaintext that must be encrypted. This process of changing session keys during an encrypted communication session is called key rollover.

Implementation

A simple way this can be done is to have one host select a random value for a new key, encrypt the value using the existing session key, and send it to the second host, who then decrypts the value and uses it as the new session key. An even more secure approach would be to

use Diffie-Hellman (DH) key exchange to send the new key to the host while also reauthenticating the host.

Another way of performing key rollover is to divide keys into several parts. For example, a key could be split into two portions, for which the larger portion is changed monthly while the smaller portion is changed more frequently such as every day or hour.

Another approach to key rollover is found in the Key Hopping technology developed by NextComm to enhance the security of 802.11a and 802.11b wireless networks. Key Hopping is implemented in hardware in integrated chips produced by NextComm for wireless networking vendors.

See Also: *key, key exchange*

key search attack

Attempting to guess a cryptographic key.

Overview

Exhaustive key search applies the brute-force method to cryptanalysis by trying all possible keys until one is found that can decrypt a given portion of ciphertext. Until a few years ago, popular encryption algorithms such as Data Encryption Standard (DES) were assumed to be immune to such attacks, which were viewed as computationally infeasible using current computing platforms. With the growth of the Internet, however, the potential for distributing the task of exhaustive key search to idle processing cycles on thousands of desktop PCs has become a reality. The result was that in 1998 a group led by Rocke Verser, Matt Curtin, and Justin Dolske succeeded in cracking a 56-bit DES key using the distributed processing power of users on the Internet. Even at the current fast rate of advances in computing power, however, it is unlikely that a 128-bit key such as the one used by Advanced Encryption Standard (AES) will be cracked in our current lifetime.

See Also: *Advanced Encryption Standard (AES), brute-force attack, ciphertext, cryptanalysis, Data Encryption Standard (DES), key*

keyspace

The scope of possibilities for a cryptographic key.

K

Overview

Keyspace is the name given to the collection of all possible values for a cryptographic key. The size of a key-space is related to the number of bits in the key. For example, a 56-bit Data Encryption Standard (DES) key has a keyspace of 2^{56} , which equals about 7×10^{16} possible values. The size of a keyspace is thus directly related to the difficulty in cracking a cryptographic system using a simple brute-force approach.

See Also: *Data Encryption Standard (DES), key*

keystroke logger

Hardware or software for capturing information entered on a keyboard.

Overview

Keystroke logging is a surveillance technique that records each key pressed on a keyboard. Keystroke logging can be implemented in two ways:

- Using a hardware device that is connected between the keyboard and the computer. Such devices are typically installed in high-security environments to keep track of what employees are doing, such as to prevent users from misusing company computers for personal use.
- Using software that can be installed either deliberately (for example, to monitor employees) or stealthily (for example, a Trojan installed by an intruder to steal information).

Marketplace

Commercially available hardware-based keystroke loggers include KeyKatcher from Allen Concepts and KeyGhost from KeyGhost Limited. Some commercial software-based keystroke loggers on the market include KeyLogger Stealth from Amecisco, KeyKey Monitor from KeyKey.com, and Spector Pro from SpectorSoft.

There are also programs available that can detect when keystroke-logging software has been installed on a computer. One example is SpyCop from the company of the same name.

Before businesses decide to implement keystroke logging technologies for the purpose of monitoring

employees' actions, however, they should consult their legal departments concerning the ethical and legal issues associated with such actions.

Notes

Many popular cracking tools such as Back Orifice and SubSeven include keystroke-logging tools. Media reports even indicate that the FBI has developed its own keystroke-logging software called Magic Lantern, which can be installed stealthily and run on remote systems similar to a Trojan.

See Also: *Trojan*

klaxon

A tool for detecting port-scanning attacks.

Overview

Klaxon is a tool developed by Doug Hughes of Auburn University that is useful for determining when your hosts are being port scanned with such tools as ISS or SATAN. Klaxon can detect and log port-scan connections on the host on which it runs. Klaxon runs on Linux and various UNIX platforms, including AIX and Solaris.

For More Information

Visit www.eng.auburn.edu/~doug/ to download Klaxon.

See Also: *port scanning*

Klez

A worm that targets Microsoft Windows messaging clients.

Overview

Klez is one of the most enduring worms ever to plague the Internet and was on the top 10 charts of antivirus vendors for almost the whole of the year 2002. Klez first appeared in November 2001 and targeted Microsoft Windows platforms by exploiting vulnerabilities in Microsoft Outlook and Outlook Express that allowed them to become infected simply when a user previewed or opened an e-mail message. When the worm infects a machine, it uses its own Simple Mail Transfer Protocol (SMTP) mailing engine to mass mail copies of itself to everyone in the user's address book.

The worm also includes a polymorphic virus called ElKern that can infect executable files.

Several variants of Klez have also appeared, including these:

- **Klex.D:** Also propagates itself to a user's ICQ database
- **Klez.E:** Essential version 2 of Klez, allowing the worm to infect files, spread across a network using mapped drives, kill virus protection software, and corrupt data
- **Klez.H:** No longer targets Outlook and has variations in its behavior that make it more difficult to identify and track

Notes

Installing the latest service packs and hotfixes prevents machines running on the Microsoft Windows platform from becoming infected with the worm and its variants.

See Also: virus, worm

Knark

A rootkit that targets the Linux platform.

Overview

A **rootkit** is a collection of tools installed by intruders on compromised systems to allow reentry without detection. Knark is a rootkit developed specifically to target Linux hosts, and what makes this rootkit unique is that it hides itself in the operating system kernel by using Loadable Kernel Modules (LKMs) for installation. This makes it more difficult to detect than traditional rootkits, which generally replace system files and can be detected by using file system verification tools.

In addition to providing a backdoor, Knark includes several other exploits that affect Berkeley Internet Name Domain (BIND), File Transfer Protocol (FTP), Line Printer Daemon (LPD), and other common network services.

See Also: rootkit

known plaintext attack

A cryptanalytic attack in which the cracker has some plaintext/ciphertext pairs to work with.

Overview

In a known plaintext attack, the cracker already knows the plaintext of one or more blocks of ciphertext. Using this information, it is generally fairly easy for the cracker to deduce the encryption key and decrypt additional blocks of ciphertext. An attacker might obtain the required plaintext/ciphertext pairs in these ways:

- As a result of older secret data being released into the public domain as plaintext
- By deducing that initial encryption blocks of a transmission represent standard document headers for Microsoft Word documents, Simple Mail Transfer Protocol (SMTP) e-mail, or some other common format.

See Also: cryptanalysis

KryptoKnight

A cryptographic authentication system developed by IBM.

Overview

KryptoKnight has an authentication architecture similar to the Kerberos protocol developed by Massachusetts Institute of Technology (MIT). While Kerberos uses key distribution centers (KDCs) distributed among different domains, KryptoKnight employs authentication servers managing different realms. An important difference, however, is that while Kerberos employs secret key encryption algorithms such as Data Encryption Standard (DES) for authentication and ticket encryption, KryptoKnight uses message digest (MD) functions instead for faster performance and easier compliance with cryptographic export controls. KryptoKnight can also employ random number challenges instead of necessarily relying on synchronized clocks between KDCs. Also, KryptoKnight does not support the advanced features of version 5 of Kerberos, including delegation, hierarchical realms, and renewable tickets.

Notes

KryptoKnight goes under the product name of Network Security Program (NetSP).

See Also: Data Encryption Standard (DES), Kerberos, message digest (MD)

L0phtCrack

A password-cracking tool from @stake (formerly L0pht Heavy Industries).

Overview

L0phtCrack is a popular tool for auditing account passwords and recovering lost passwords on Microsoft Windows platforms. Administrators can use L0phtCrack to audit their networks and detect weak passwords that could constitute security vulnerabilities. The tool can also be used as a password cracker, though a “Hide” feature allows administrators to configure it so that it does not divulge passwords it has cracked but rather simply displays auditing information such as password length.

L0phtCrack can crack or audit passwords obtained from several sources, including local computers, remote computers on the network, and by sniffing a network segment for NTLM authentication traffic. It works via a dictionary attack but can also be configured to perform a brute-force attack to recover passwords from machines running on the Microsoft Windows platform. The current version of L0phtCrack is 4 and is commonly referred to as LC4.

For More Information

Visit @stake at www.atstake.com for more information on L0phtCrack.

See Also: *brute-force attack, dictionary attack, John the Ripper, password cracking*

L2TP

Stands for Layer 2 Tunneling Protocol, a tunneling protocol used for virtual private networking.

See: *Layer 2 Tunneling Protocol (L2TP)*

LaGrande Technology (LT)

An emerging technology from Intel that integrates security features into processors and chipsets.

Overview

LaGrande Technology (LT) represents a hardware-based approach to enhancing computer security. The technology supports protected execution, memory, and storage to help ensure that programs and data are safeguarded as they enter or leave a system and are processed or stored. LT also protects input/output functions, such as keyboard input and video output, and helps protect systems against attack by Trojans, keystroke loggers, spyware, viruses, and other tools used to compromise systems. LT helps to implement the recommendations of the Trusted Computing Platform Alliance (TCPA) and is expected to be incorporated into Intel’s Prescott successor to the Pentium IV processor family and chipset.

See Also: *keystroke logger, spyware, Trojan, Trusted Computing Platform Alliance (TCPA), virus*

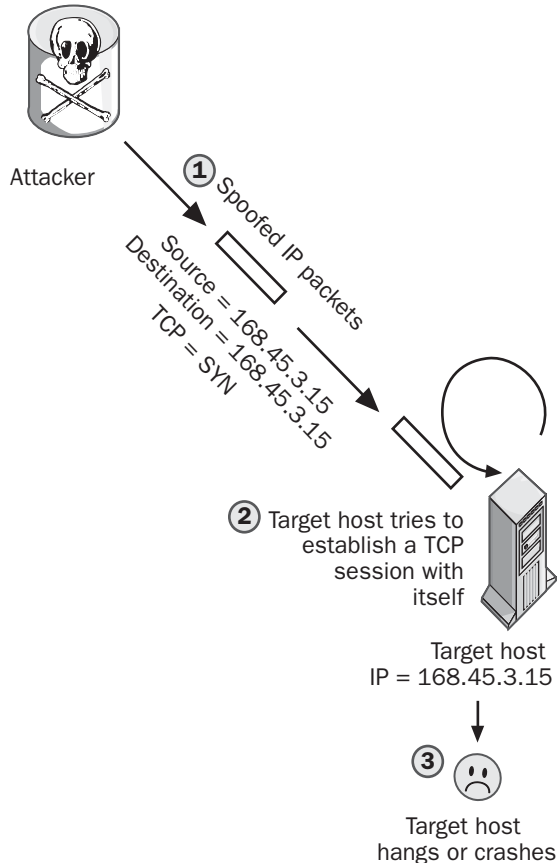
LAND attack

A well-known example of a denial of service (DoS) attack.

Overview

The LAND attack was developed in 1997 by Hugo Breton, a 16-year-old Montreal high school student who went by the moniker “meltman” or “m3lt.” The attack works by using spoofed Internet Protocol (IP) packets to trick the target host into trying to establish a Transmission Control Protocol (TCP) session with itself. IP packets are crafted with source and destination addresses set to the address of the target host, and the SYN flag is set in the packets to try to initiate a session on a designated port. The original exploit was discovered to crash machines running Microsoft Windows 95, but it was soon discovered that other platforms were also affected, including UNIX hosts, Cisco routers, and

network printers. The results of the attack varied with different platforms and ranged from temporary slow-downs to locking up, hanging, or crashing the machines. Patches were soon released to deal with the problem, but effects were widespread and brought attention to weaknesses in the Transmission Control Protocol/Internet Protocol (TCP/IP) protocol suite and the vulnerability of networks to cleverly crafted attacks.



LAND attack. How a LAND attack works.

Notes

The attack is named after land.c, the C code for the original exploit.

See Also: denial of service (DoS)

LAN Manager authentication

The authentication protocol used by legacy versions of the Microsoft Windows platform.

Overview

The LAN Manager authentication protocol was originally developed by IBM and used by Microsoft as the authentication method for Microsoft Windows 3.1, Windows for Workgroups 3.11, Windows 95, Windows 98, and Windows Millennium Edition (Windows Me). The protocol is supported by all versions of Windows but suffers from several features that make it vulnerable to compromise through eavesdropping, for the following reasons:

- User passwords are converted to uppercase before hashing, which makes hashed passwords more susceptible to cracking using dictionary attacks.
- Hashed passwords are padded with zeros and stored in 7-byte segments, which make them easier to crack than full-length passwords.

LAN Manager authentication was replaced by NTLM authentication in the Microsoft Windows NT platform, but LAN Manager password hashes were still stored together with NTLM hashes and both LAN Manager and NTLM responses were sent by default to clients requesting authentication. This was resolved in Windows NT 4 Service Pack 4, which provided the option of disabling LAN Manager authentication entirely.

In Windows 2000 the authentication protocols used can be configured using local security policy, and by default both LAN Manager and NTLM responses are sent to clients. In the Windows 2003 Server family, security has been tightened so that by default, only NTLM responses are sent to clients requesting authentication. LANMAN hashes are still stored in the registry, although Windows 2000 Server Pack 2 and later include a registry setting for disabling this.

See Also: authentication, Kerberos, NTLM

LANMAN authentication

Short for LAN Manager authentication, the authentication protocol used by legacy versions of the Microsoft Windows platform.

See: LAN Manager authentication

Layer 2 Tunneling Protocol (L2TP)

A tunneling protocol used for virtual private networking.

Overview

Layer 2 Tunneling Protocol (L2TP) is an industry standard tunneling protocol defined by RFC 2661. L2TP is based on two earlier tunneling protocols:

- Point-to-Point Tunneling Protocol (PPTP), developed by Microsoft
- Layer 2 Forwarding (L2F), developed by Cisco Systems

L2TP can be used to deploy virtual private networks (VPNs) over Internet Protocol (IP), Asynchronous Transfer Mode (ATM), frame relay, or X.25 networks. On IP networks, L2TP works by encapsulating Point-to-Point Protocol (PPP) frames into User Datagram Protocol (UDP) packets and provides encryption using Internet Protocol Security (IPSec).

L2TP has several advantages over PPTP:

- Encryption in L2TP begins prior to the PPP connection process instead of after authentication, as with PPTP.
- While PPTP uses RC4 (a relatively weak stream cipher) for encryption, L2TP supports Data Encryption Standard (DES) and Triple DES (3DES) encryption.
- The IPSec protocol used by L2TP provides additional security in the form of data integrity, confidentiality, and replay protection.

The main disadvantage of L2TP is that it can't be used in conjunction with Network Address Translation (NAT) as PPTP can.

See Also: 3DES, Data Encryption Standard (DES), Internet Protocol Security (IPSec), Point-to-Point Tunneling Protocol (PPTP), virtual private network (VPN)

LEAP

Stands for Lightweight Extensible Authentication Protocol, an authentication protocol developed by Cisco for wireless networks.

See: Lightweight Extensible Authentication Protocol (LEAP)

least privilege

A best practice regarding the rights and permissions that users and applications should have.

Overview

Least privilege means different things in different contexts, but the general idea is that entities (users, applications, or devices) should be assigned the minimum privileges (rights or permissions) they need to fulfill their purposes and no more. When this principle is applied to users, it means that user accounts should be granted just enough rights to do their jobs. For example, only administrators should have rights allowing them to back up servers, assign permissions, reset passwords, and perform other administrative tasks. Ordinary users should be able to run programs and access network resources they need, but should not be allowed to do the kinds of things administrators can do.

There are several reasons why the principle of least privilege is important:

- Users that are granted rights and permissions greater than they need may be tempted to use these privileges to access files or perform system tasks that they are not authorized to perform, which can result in data loss or damage (if they don't know what they're doing) or business loss (if they steal information or sabotage systems).
- If users have privileges greater than they need and a user's account is compromised by a malicious intruder, the intruder can use the elevated privileges of the user to cause damage, destroy data, or perform other harmful actions.

Another aspect of least privilege is that computing tasks should always be performed with the minimum credentials needed to perform them. For example, if administrators read their e-mail while logged on to their Administrator account, they violate the principle of least privilege since e-mail programs do not require administrative privileges to run. Best practice in this case would be for each administrator to have two separate accounts, an administrative-level account used to perform system tasks that require administrative privileges, and an ordinary user account with which to perform ordinary tasks, such as browsing the World Wide Web or checking

L

e-mail. On the Microsoft Windows platform, administrators can make use of the Runas command (the secondary logon) to perform administrative tasks while logged on with ordinary user credentials.

The least privilege principle can also be applied to programs, systems, or anything else in a computing environment. For example, a Web application should not run within the security context of the all-powerful LocalSystem account since, if the application were compromised, the attacker might be able to elevate its privileges and take over the system. This also applies to operating systems, which should run network services using the minimum credentials necessary.

See Also: *chroot jail, elevation of privileges (EoP), permissions, rights, secondary logon*

LFM

Stands for log file monitor, a tool that monitors log files looking for signs of intrusion.

See: *log file monitor (LFM)*

Liberty Alliance Project

An industry initiative to develop an open framework for managing network identity.

Overview

With the increase of e-commerce on the Internet, the growing problem of managing network identity has emerged. Most e-commerce sites manage their own database of client accounts, with the result that consumers must maintain multiple accounts and reenter them each time they want to access a different site. The Liberty Alliance Project is designed to address this issue by providing a federated single sign-on (SSO) solution that allows a consumer to enter credentials once to access a whole group of sites.

The goal of the Liberty Alliance is to simplify the process of managing online identity by developing an open standard for federated network identity while ensuring the privacy and security of all identity information. The standard will support a wide range of identity products and services and will be available to both commercial

and noncommercial organizations. The federated nature of the standard will allow consumers to decide which e-commerce sites should be linked to a given network identity, with the result that once the consumer has logged on to one site in a group, affiliate sites in the same group can be accessed without the need to reenter credentials. The result is that consumers benefit from the choice they receive in how they want to manage their identities and the convenience of being able to access multiple sites using an SSO approach. Companies and organizations that implement the standard will themselves benefit through new revenue and cost savings resulting from leveraging their relationships with customers and affiliates.

The Liberty Alliance has over 150 member companies and organizations, including American Express, AOL Time Warner, General Motors, Hewlett-Packard, MasterCard International, and Sun Microsystems. The current specification for the standard is version 1.2.

For More Information

Visit www.projectliberty.org for more information.

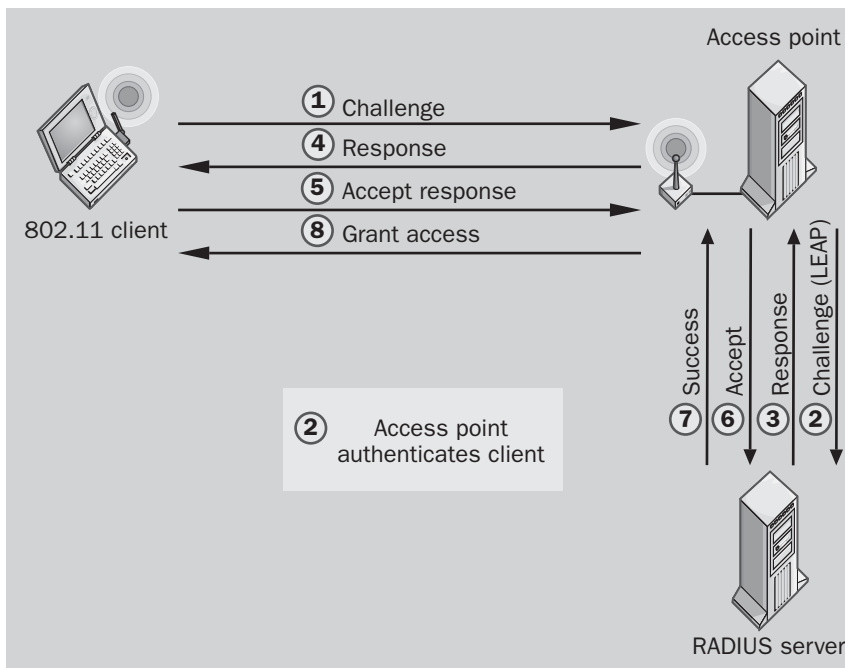
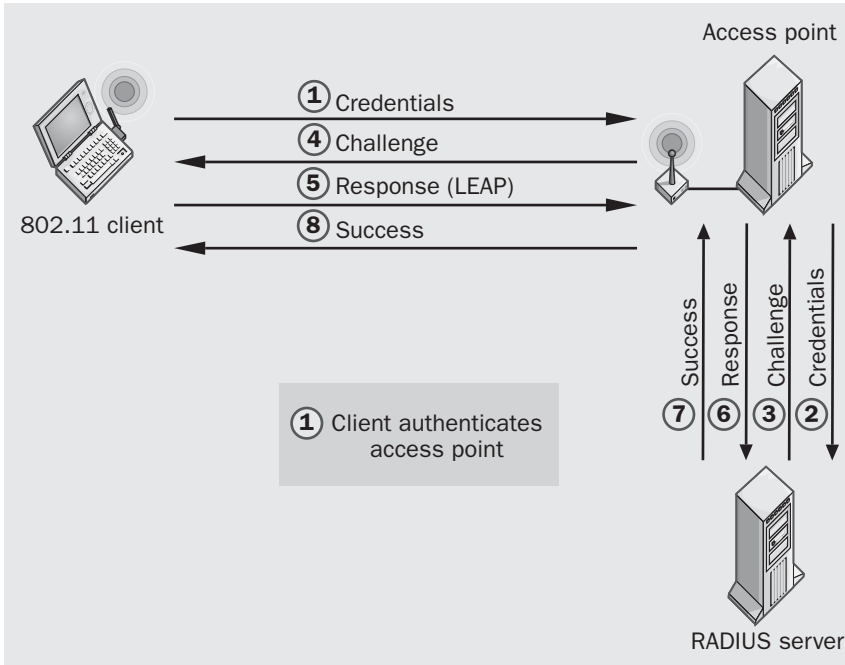
See Also: *.NET Passport, personally identifiable information (PII), single sign-on (SSO)*

Lightweight Extensible Authentication Protocol (LEAP)

An authentication protocol developed by Cisco for wireless networks.

Overview

Lightweight Extensible Authentication Protocol (LEAP) is a modified version of Extensible Authentication Protocol (EAP) developed by Cisco Systems for its Aironet line of wireless local area network (WLAN) products. LEAP is a prestandard implementation of 802.1x that provides an interim solution to security weaknesses inherent in Wired Equivalent Privacy (WEP), the original 802.11 security protocol. LEAP supports mutual authentication between WLAN adapters and access points, and it encrypts communications using dynamically generated WEP keys.



Lightweight Extensible Authentication Protocol (LEAP). How LEAP mutual authentication works.

Implementation

When a LEAP-enabled client tries to connect to a WLAN, it submits the user's credentials to the access point, which forwards them to an authentication server (an AS, typically a Remote Authentication Dial-In User Service, or RADIUS, server). The AS responds by sending a challenge string back to the access point, which forwards it to the client. The client combines the challenge string with the user password using the LEAP algorithm and sends the response string to the access point, which forwards it to the AS. The AS performs the same action on the challenge string and user password and compares the result with the response forwarded from the client. If the results match, the AS sends a success message to the access server, which forwards it to the client.

At this point the client has been authenticated, but since LEAP is a mutual authentication protocol, the client must now authenticate the access point. To do this, the client sends a challenge string to the access point, and a reverse LEAP authentication process takes place. Once the access point has been authenticated, the client sends a success message to the access server, which forwards it to the AS. The AS opens a port, and the client then can access the network.

See Also: 802.1x, Extensible Authentication Protocol (EAP), Wired Equivalent Privacy (WEP)

Linsniff

A password-sniffing tool for the Linux platform.

Overview

Linsniff is a tool for extracting Linux passwords from authentication traffic on Ethernet networks. The tool is similar in operation to Dsniff but doesn't support as many types of authentication protocols as Dsniff does. The C code for Linsniff can be downloaded from various sites on the Internet and compiled for use.

See Also: Dsniff, sniffer

listening port

A port on a server that is waiting for a client connection.

Overview

Listening (or open) ports are Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) ports that are used by running services to listen for a client attempting to establish a connection. Such ports are said to be running in a LISTENING state, and in general each network service running on a host has one or more listening ports. Listening ports are also open for attackers, however. Attackers can detect which ports are listening by performing a port scan on a host, and the results of such a scan can often fingerprint the system by providing information about which operating system is running, which version or service pack is applied, and which optional services or daemons are installed. In general, best practice suggests that administrators limit the number of listening ports on a host by disabling any unnecessary services running on the machine.

Another way to make it more difficult for attackers is to change the default port on which services listen. This approach has two disadvantages, however:

- It relies on “security through obscurity,” which is usually not considered a significant approach to enhancing the security of a platform.
- It makes it more difficult for clients to access network services and may require the added overhead of reconfiguring clients to utilize the new ports.

Notes

On machines running on the Microsoft Windows platform, you can use the Netstat command to determine which ports are currently listening for connections.

See Also: Netstat, port scanning

LM authentication

Short for LAN Manager authentication, the authentication protocol used by legacy versions of the Microsoft Windows platform.

See: LAN Manager authentication

local attack

An attack performed at the local console of a system.

Overview

A local attack is one that the attacker launches by interactively logging on to a computer. Local attacks are generally more dangerous than network attacks since network security measures such as firewalls are circumvented. To perform a local attack, the attacker requires two things:

- Physical access to the system
- A valid user account for logging on

Restricting physical access to systems is a fundamental principle of information security, and it prevents such attacks from being performed. Protecting user accounts with strong passwords is also critical since, once an attacker has gained local access to a system using an ordinary user account, the attacker may be able to elevate its privileges and gain control of the system.

See Also: *attack, physical security*

local exploit

Another name for local attack, an attack performed at the local console of a system.

See: *local attack*

locally unique identifier (LUID)

A value unique to a computer running on the Microsoft Windows platform.

Overview

A locally unique identifier (LUID) is a 64-bit value that is guaranteed to be unique on the computer on which it was generated. This uniqueness, however, is guaranteed only until the system restarts. LUIDs are not intended for direct manipulation and must be manipulated by applications using appropriate function calls.

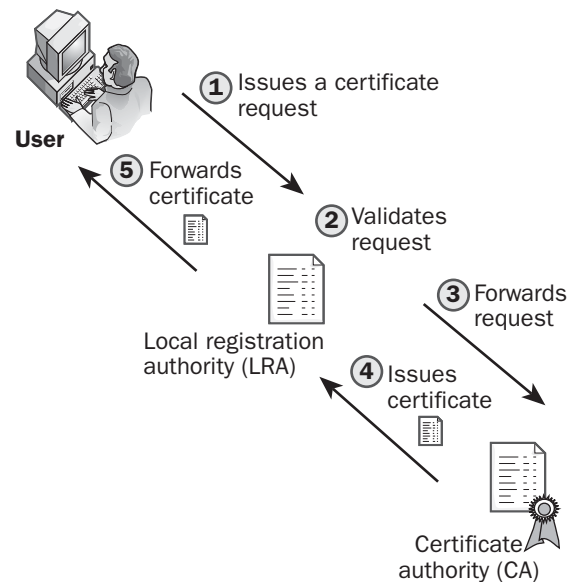
See Also: *security identifier (SID)*

local registration authority (LRA)

An intermediate registration authority (RA) in a Public Key Infrastructure (PKI).

Overview

A local registration authority (LRA) acts as an intermediary between users and a certificate authority (CA). In a typical scenario, users might submit their certificate requests to an LRA, which would validate the requests before forwarding them to the CA for issuing certificates. LRAs are optional components of PKI systems and can be employed to offload work from the CA by performing authentication, validation, and auditing tasks. LRAs may also be used for revoking certificates when they are lost or stolen. LRAs are most typically deployed in large, distributed PKI systems when users are at some distance from CAs.



Local registration authority (LRA). Processing certificate requests using an LRA.

See Also: *certificate authority (CA), Public Key Infrastructure (PKI), registration authority (RA)*

Local Security Authority (LSA)

A protected subsystem of computers running on the Microsoft Windows platform that performs authentication.

Overview

The Local Security Authority (LSA) authenticates users and interactively logs them on to the local system. The LSA also manages information concerning various aspects of local security stored in the local security policy of the system, including password restrictions and audit settings. The LSA is also responsible for generating access tokens that contain information about the user's group membership and level of security privileges on the system. The LSA runs as a user-mode process called lsass.exe within which various subcomponents run handling the different types of authentication supported by computers running versions of Microsoft Windows.

See Also: access token, authentication, logon

local security policy

A collection of settings relating to the security of computers running Microsoft Windows 2000.

Overview

Local security policy contains the following types of security information:

- Which domains are trusted for authentication of logon attempts
- Which user accounts are allowed to access the system and the way in which they can access it (interactively, through a network, or as a service)
- The various rights and privileges assigned to user accounts
- The audit policy for the machine
- Password and account lockout restrictions

Local security policy is managed by the Local Security Authority (LSA), a protected subsystem of computers running on the Microsoft Windows platform that performs authentication. Local security policy settings are stored in the registry as a set of LSA Policy Objects. In a domain environment, local security policy can be modified by using Group Policy.

See Also: Group Policy, Local Security Authority (LSA)

locking down

Another name for hardening, configuring a host to make it more secure for a specific role.

See: hardening

log analysis software

Software for generating reports from log files.

Overview

Log files are a key element of system security, and their analysis can detect when intrusion has occurred. Log files can also be analyzed for other reasons, such as for monitoring system or application performance to determine usage trends for upgrade planning or business expansion. Hundreds of different types of log analysis software are available in the market, but good log analysis software should at the minimum include the following features:

- Support for a variety of log types and log file formats
- Advanced filtering and query options
- Robust reporting capabilities, including summary and detailed reports
- Automation for real-time analysis and report generation
- A simple and easy-to-use interface

See Also: log file monitor (LFM)

log cleaning

Removal of evidence from log files after a successful intrusion.

Overview

Log cleaning is a step commonly performed by an attacker after compromising a system. It involves removing or modifying entries in system logs to erase all trace of the exploit to help hide the fact that the system has been compromised. Log cleaning may be performed manually if needed, but most rootkits include tools for automatic removal of log entries to cover an intruder's tracks. Some popular tools for manually cleaning log files are Clean and Zap2.

Sometimes intruders become overzealous and erase all entries in log files, and a careful system administrator who notices this may conclude that the system was compromised. To prevent log cleaning, store log files on a secured, separate system using remote logging. Frequent review of log files is another important way to detect when intruders have invaded systems. Finally, using a log file monitor (LFM) to scan logs automatically in real time and notify administrators of suspicious actions performed on them is a useful enhancement for system security.

See Also: *log file monitor (LFM), rootkit*

log file monitor (LFM)

A tool that monitors log files looking for signs of intrusion.

Overview

A host-based intrusion detection system (HIDS) often includes a log file monitor (LFM) to analyze log files on the fly, looking for evidence of attempted intrusion. These monitors may scan various types of logs, including system logs, security logs, or Web logs, depending on the operating system platform used and the applications running. LFMs can generally be configured to perform various actions when a suspicious log entry is detected, for example, sending an automatic e-mail message to an administrator for notification purposes.

Marketplace

Examples of popular LFMs for UNIX platforms include LogSentry from Psionic Software, LoFiMo from SourceForge, swatch (Simple WATCHer), and LogSurfer. Examples of LFMs for Microsoft Windows platforms include Monitor Magic from Advanced Toolware and SiteScope from Mercury Interactive.

See Also: *intrusion detection system (IDS)*

logic bomb

A program that triggers when certain conditions are met.

Overview

Logic bombs are programs deliberately written to produce certain results when certain conditions are met. For example, a program could trigger erasure of files on a hard drive on a certain day of the month or year.

Viruses and worms sometimes contain logic bombs; for example, the infamous “Friday the 13th” virus, which replicated itself each Friday and on the 13th of each month. Another famous example was the Michelangelo virus, which activated on the sixth of March and tried to wipe hard drives.

Logic bombs are often employed for insider attacks. Several famous instances of employees programming logic bombs in their companies’ servers have been documented in the media. They did it in such a way that the bombs had to be reset regularly to prevent them going off. Then, should the employee with malicious intent be fired and forced to leave the premises, there would be no one around to reset the bomb, which, at a predetermined time and date, would go off, wreaking havoc with company programs and data.

See Also: *backdoor, Trojan, virus, worm*

Loginlog

A UNIX tool for logging failed logons.

Overview

Loginlog is a command available on UNIX platforms that records failed logon attempts. Administrators can use this tool to detect attempts at breaking into a system since a high number of failed logons within a short period of time is a classic signature of intrusion. Loginlog records failed logons in the file `/var/adm/loginlog`, with each record in the file corresponding to one failed attempt and specifying the logon name, time, and tty specification. Loginlog records failed logons only if five or more attempts are made.

See Also: *intrusion*

logon

Authenticating credentials submitted by an entity seeking access to a system or network.

Overview

When users, applications, or devices wish to access resources on a system or network, they first need to log on to the system or network. **Logging on** is the process of submitting credentials, having them authenticated, and gaining access to the system or network that

L

performs the authentication. There are several types of logons supported by Microsoft Windows platforms:

- **Interactive logon:** Authentication of users on a computer by entering credentials on the local console
- **Network logon:** Proxy authentication for logon sessions on a remote computer
- **Service logon:** Authentication of Microsoft Windows services using LocalSystem or some other credentials as a security context in which to run
- **Batch logon:** Authentication of applications that run as batch jobs, such as COM servers

See Also: *authentication, secondary logon*

logon identifier

A locally unique identifier (LUID) that identifies a logon session.

Overview

A logon identifier is created when a user logs on to a computer running on the Microsoft Windows platform, and it remains valid until the user logs off. This logon identifier is unique while the computer is running; no other logon session that is started can have the same logon identifier. When the computer is rebooted, however, the set of possible logon identifiers is reset and can be reused. The logon identifier is part of the access token generated for the session and can be retrieved using the `GetTokenInformation` function for `TokenStatistics`.

See Also: *access token, logon, logon session*

logon session

A session that is started when a user logs on to a computer running on the Microsoft Windows platform.

Overview

The primary access token generated when a user logs on to a computer running on the Windows platform contains a logon identifier that uniquely identifies the logon session started on the computer. The access token also contains other information concerning the security context of the logon session, including the security identifier (SID) for the currently logged-on user and the logon SID.

There are four basic types of logon sessions that can be created: interactive, network, batch, and service.

See Also: *logon, logon identifier, logon SID*

logon SID

A security identifier (SID) that identifies a logon session.

Overview

A SID is a variable-length data structure that identifies a security principal (a user, group, or computer account) on a computer running on the Microsoft Windows platform. A logon SID is a SID created for a logon session and is valid for the duration of the session until the user logs off of the computer. The logon SID is unique to the computer, and no other logon session started on that machine can have the same SID. However, once the machine reboots, the slate of possible logon SIDs is reset and can be reused by new logon sessions. The logon SID for a logon session can be used in a discretionary access control list (DACL) to control access to resources during the session.

See Also: *logon, logon identifier, logon session*

Loki

A tool used to test or circumvent firewalls.

Overview

Loki is a tool that employs Internet Control Message Protocol (ICMP) tunneling to try to circumvent firewall protection for networks. ICMP tunneling is a method of using ICMP to establish a covert channel. Loki works by using a client (Loki) to encapsulate Internet Protocol (IP) packets from the attacker within the headers of ICMP packets and then transmit these packets to a server (Lokid) running on a system inside the firewall. Loki thus provides a type of backdoor through which systems can be remotely controlled across a firewall, though many firewall products now have been patched to resist such activity.

Notes

Loki was first published in *Phrack* magazine.

See Also: *firewall, ICMP tunneling*

LoveLetter

A malicious Visual Basic Script (VBScript) program that spreads using the Microsoft Outlook address book.

Overview

LoveLetter (also known as ILOVEYOU or The Love Bug) is a mass-mailer worm that appeared in May 2000. The worm is written in VBScript and is delivered as an attachment to e-mail messages. If the user's computer has the Microsoft Windows Scripting Host (WSH) enabled and opens the attachment, the script executes and sends copies of itself to everyone in the Outlook address book. The worm also performs other actions, including overwriting certain types of files (mostly multimedia files), copying itself to the system folder to ensure it reappears after a reboot, modifying the start page for Microsoft Internet Explorer so it points to a page that will download a Trojan, and using mIRC, if installed, to propagate itself over Internet Relay Chat (IRC).

What made the worm especially dangerous was the subject line "ILOVEYOU" and attachment name "LOVE-LETTER-FOR-YOU.TXT.vbs, which tempted recipients to open the attachment out of natural curiosity. Once released in the wild, the LoveLetter worm spread at a rapid rate across the Internet, costing businesses billions of dollars. To date, over 80 variants of the worm have been detected in the wild, making it one of the most popular and dangerous worms of all time.

Notes

The polymorphic worm called NewLove is similar to LoveLetter but more dangerous because it infects system and data files and can mutate itself to prevent detection by virus protection software. NewLove is not a variant of LoveLetter, though it employs some of the code base of LoveLetter.

See Also: *Melissa, worm*

LRA

Stands for local registration authority, an intermediate registration authority (RA) in a Public Key Infrastructure (PKI).

See: *local registration authority (LRA)*

LSA

Stands for Local Security Authority, a protected subsystem of computers running on the Microsoft Windows platform that performs authentication.

See: *Local Security Authority (LSA)*

Lsadump2

A cracking tool that displays the contents of LSA Secrets on computers running Microsoft Windows NT.

Overview

LSA Secrets is a portion of the Windows NT registry where the Local Security Authority (LSA) stores security information on behalf of applications. If attackers can gain local access to a machine running Windows NT using Administrator privileges, they can use the Lsadump2 tool to dump the contents of LSA Secrets and gain access to cached passwords for domain accounts, passwords for service accounts, and other important security information. Lsadump2 works by using a process called dynamic-link library (DLL) injection, which bypasses access controls and is therefore unsupported by Microsoft and could have unintended consequences on the machine it runs. Lsadump2 can also be used legitimately by administrators as a security auditing tool.

For More Information

Visit razor.bindview.com for more information.

See Also: *LSA Secrets, password, Pwdump2*

LSA Secrets

A portion of the Microsoft Windows NT registry where the Local Security Authority (LSA) stores security information on behalf of applications.

Overview

LSA Secrets contains cached passwords for domain accounts, passwords for service accounts, and other important security information critical to protect on systems running Windows NT. These passwords are stored in HKLM\SECURITY\Policy\Secrets, a secret portion of the Windows NT registry that is inaccessible even to members of the Administrators group on the local machine. Normally, the only security principal

that can access this information is the highly privileged LocalSystem account, but several local exploits exist that can provide attackers with access to information stored in LSA Secrets, including scheduling Regedt32.exe to run interactively using the At command, which starts Registry Editor using LocalSystem credentials, and using the cracking tool Lsdump2.

The existence of these simple exploits emphasizes the importance of ensuring the physical security of computers at all times. This includes not logging on to desktop machines using Domain Admin credentials since cached credentials could be displayed by crackers gaining physical access to the machine subsequently.

See Also: local attack, Lsdump2

Lsof

A tool for listing open files on a system.

Overview

Lsof (which stands for LiSt Open Files) is a tool that can be used for intrusion detection on UNIX platforms. The tool displays a list of all open files on a system, and by scanning this list a knowledgeable sysadmin can detect evidence of unauthorized access to a system. The tool also includes options for listing open files associated with a given process ID and displaying open ports that are listening for Transmission Control Protocol (TCP) connection attempts. By detecting an unusual file association with a process or an unexpected listening port, an administrator can determine whether an intruder has compromised a system.

For More Information

Visit <ftp://vic.cc.perdue.edu> to download Lsof.

See Also: intrusion, intrusion detection system (IDS)

LT

Stands for LaGrande Technology, an emerging technology from Intel that integrates security features into processors and chipsets.

See: LaGrande Technology (LT)

LUCIFER

An early block cipher developed by IBM.

Overview

LUCIFER was a block cipher that encrypted 128-bit blocks of plaintext using a 128-bit key. The cipher performed 16 rounds of iteration in which each round encrypted the left half of the block using a subkey, XORed the result with the right half, and then swapped the halves. The special significance of LUCIFER is that it later formed the basis of the Data Encryption Standard (DES) algorithm, an encryption standard used for many years by the U.S. federal government. LUCIFER is not a very secure encryption algorithm because of the regularity of its key schedule, despite the fact that it employs a larger key than DES.

See Also: block cipher, Data Encryption Standard (DES)

LUID

Stands for locally unique identifier, a value unique to a computer running on the Microsoft Windows platform.

See: locally unique identifier (LUID)

Luring attack

A type of attack that exploits trusted code to elevate privileges for untrusted code.

Overview

Luring attacks occur when malicious code causes trusted code to perform something that the privileges of the malicious code don't allow it to accomplish on its own. The way it usually works is that the malicious code somehow tricks the trusted code into calling a portion of the malicious code using the privileges of the trusted code. Early versions of the Java Language Specification were vulnerable to this type of attack; it can be difficult to develop code that is resistant to such attacks. The code access security feature (CAS) of the Microsoft Windows .NET Framework employs a "stack walk" method that protects user-developed code against such attacks.

See Also: code access security (CAS)

MAC

1. Stands for mandatory access control, a mechanism for controlling access by users to computing resources.
2. Stands for message authentication code, a keyed hashing algorithm used to guarantee the integrity of a message.

See: mandatory access control (MAC), message authentication code (MAC)

MAC duplication

A type of denial of service (DoS) attack against switched networks.

Overview

MAC duplication is an attack that involves forging packets that have the same source Media Access Control (MAC) address and then sending them to two different ports on the switch, making the switch think that the same host resides on two separate network segments. Some switches respond to this condition by hanging or crashing, which results in legitimate hosts being unable to send traffic to portions of the network. Other switches may respond by simply forwarding the traffic from both ports without any further consideration.

Another use of MAC duplication is to redirect traffic on a switched network. An attacker first compromises a host on a remote network by exploiting some vulnerability, and then changes the MAC address of the compromised host to that of another host being targeted. As a result, all traffic sent to the target host is also received by the compromised host.

Notes

For more information about Media Access Control and MAC addresses, see the *Microsoft Encyclopedia of Networking, Second Edition*, available from Microsoft Press.

See Also: denial of service (DoS), MAC flooding, MAC spoofing

MAC flooding

A type of denial of service (DoS) attack against switched networks.

Overview

MAC flooding is an attack that tries to flood the internal memory of Ethernet switches using large numbers of spoofed Media Access Control (MAC) addresses. Ethernet switches generally learn the MAC addresses of hosts on a local segment by listening to traffic on the port to which the segment connects. By spoofing large numbers of packets, each having different MAC addresses, the address table in the switch's memory can become full, which can prevent legitimate hosts from using the switch to send traffic. Some switches may even stop switching entirely when their tables fill up, and they begin forwarding traffic like shared hubs instead, which can sometimes allow attackers to capture traffic on parts of the network that were not previously visible to them. An example of a tool that can be used to launch such an attack is Macof, which can generate hundreds of thousands of spoofed frames per minute.

On Cisco Catalyst switches, port security can be used to mitigate the effects of such attacks. Port security specifies the maximum number of hosts that can be connected to a port on the switch. If this number is exceeded, the switch determines that a MAC flooding attack is underway and automatically shuts down the port. The administrator can then determine the source of the attack, resolve the problem, and turn the port back on.

See Also: denial of service (DoS), MAC spoofing

MAC spoofing

An attack that involves spoofing the Media Access Control (MAC) address of legitimate hosts.

Overview

MAC address spoofing involves forging the source MAC address of packets in an attempt to gain trusted access to a network. Such attacks are common on wireless local area networks (WLANs) on which malicious hosts try to masquerade as either a legitimate client or the network access point to bypass other access control mechanisms. It's relatively easy to spoof MAC addresses and there are a number of tools available for launching spoofing attacks, including AirJack and FakeAP.

MAC spoofing can also be used to launch a denial of service (DoS) attack on an Ethernet network by sending large numbers of forged Address Resolution Protocol (ARP) replies to a target host.

There can even be legitimate reasons for spoofing MAC addresses. For example, if your LAN is connected to a service provider using a router whose MAC address is authenticated by the provider, and you need to replace the router, you may be able to reconfigure the MAC address of the new router and avoid having to request that your provider reconfigure security on its end.

See Also: MAC flooding, spoofing

macro virus

A virus that exploits a macro programming language.

Overview

The first known macro virus was the Concept virus (also known as the Prank virus or Macro virus), which appeared in 1995. This virus exploited the macro language of Microsoft Word to automatically replicate itself into new Word documents. The Concept Virus was nondestructive but was a portent of more malicious viruses soon to come, the first of which was the WinWord virus. Soon macro viruses were appearing that could delete files, format drives, and perform other harmful actions, and within a couple of years viruses such as CAP and Wazzu were topping the charts of vendors of virus protection software. Macro viruses have also appeared for other Microsoft Office products including Microsoft Excel, Microsoft PowerPoint, and

Microsoft Access, and for products from other companies, including AmiPro from Lotus.

Implementation

Word macro viruses leverage the power of Visual Basic for Applications (VBA), the programming language built into Office. A macro virus generally infects a system by opening an infected Word document obtained from someone else. The virus then copies itself to Word template files so that other documents that are created or opened on the system automatically become infected as well. When an infected document is opened, the virus payload runs, performing whatever action the virus is programmed to perform.

The most certain way of guarding against macro viruses is to disable macros, but this results in a loss of the functionality that macros provide. Virus protection products easily detect existing macro viruses and can notify users when macros attempt to run.

See Also: Melissa, virus

Mafia Boy

A teenager who brought down many of the largest sites on the Internet.

Overview

Mafia Boy was the moniker of a 15-year-old cracker from Quebec, Canada, who in 2001 was charged with conducting denial of service (DoS) attacks that brought down Amazon.com, CNN.com, eBay, Yahoo!, and other popular Internet sites for more than six hours. The estimated losses due to the attack were placed at \$1.2 billion, including stock value (capitalization) losses, revenue losses, and recovery time. Since Mafia Boy was a juvenile, he was sentenced to a \$650 fine and two years of detention. The Mafia Boy incident brought cybercrime to the forefront of media attention and highlighted the vulnerability of the Internet to such DoS attacks.

See Also: cybercrime, denial of service (DoS)

mail bombing

A denial of service (DoS) attack on a user's mailbox.

Overview

Mail bombing is an activity that used to be fairly common in the early days of the Internet. Typically, someone posted something to USENET that upset someone else (usually for persistently violating “netiquette” or USENET customs), and the second person sometimes responded by sending a large amount of mail (typically messages with large attachments) to the poster. This resulted in the poster’s mailbox becoming full, which prevented the poster from receiving legitimate mail from others until the mailbox was cleared. In an age when most users were connected to the Internet using slow modem connections, mail bombing was particularly annoying since it could sometimes take hours to download the malicious messages and their attachments before the user could start receiving legitimate mail again.

With the arrival of high-speed Internet, mail bombing is not nearly as effective and its use has declined greatly. If you think you are a victim of mail bombing, contact your mail administrator, who can block mail from the attacking host.

Notes

There have even been examples of worms that were designed to perform mail-bombing attacks, such as the DoS.Storm worm that exploited the Web server folder transversal vulnerability in Internet Information Services (IIS) 4 and 5.

See Also: *denial of service (DoS), worm*

mail relaying

A method used by spammers for sending junk mail.

Overview

Mail relaying is a mechanism whereby a Simple Mail Transfer Protocol (SMTP) host or mail server is configured to forward messages regardless of their source or destination. Normally mail servers should forward only messages that are either

- From hosts belonging to the mail server’s local domain
- To hosts belonging to the mail server’s local domain

If a mail server forwards any other type of mail, it is performing mail relaying and is called an open mail

relay because it is open to forwarding all types of mail. Spammers utilize mail relays for two purposes:

- To offload the work of sending large amounts of mail
- To disguise the source of the mail

The negative impact of allowing mail relaying on mail servers includes theft of Internet bandwidth and central processing unit (CPU) cycles. Furthermore, if a mail server becomes recognized by others as an open mail relay, it may get “blacklisted” by being added to a public database of open mail relays, which could result in legitimate mail forwarded by your server being rejected by other SMTP hosts on the Internet.

Marketplace

Most mail server products such as Sendmail and Microsoft Exchange are configured to disable mail relaying by default. Examples of Web sites that maintain blacklists of mail relays include Open Relay Database (www.ordb.org), Mail Abuse Prevention Systems, LLC (mail-abuse.org), and Distributed Server Boycott List (dsbl.org).

See Also: *spam*

malformed packet attack

Any attack that utilizes nonstandard packets.

Overview

The protocols of the Transmission Control Protocol/Internet Protocol (TCP/IP) suite have specific limitations on their size, format, and the types of information each portion of the packet can contain. When packets violate these restrictions, they are said to be malformed. Such packets can arise either accidentally through hardware or software issues or can be created deliberately by individuals seeking to exploit vulnerabilities in services, operating systems, or devices such as routers. Some examples of attacks based on creating malformed packets include the following:

- **Chargen:** Malformed User Datagram Protocol (UDP) echo request packets result in bandwidth exhaustion.

- **LAND attack:** Internet Protocol (IP) packets whose source and destination IP addresses are identical cause the target host to crash or reboot.
- **Ping of Death:** Oversized Internet Control Message Protocol (ICMP) echo request packets cause the target host to hang or crash.
- **Teardrop:** Two fragments that cannot be reassembled cause the target host to crash or reboot.
- **WinNuke:** Out-of-band data sent to a certain port cause the target host to crash.

Most intrusion detection systems (IDSs) and firewalls are capable of detecting and preventing common types of malformed packet attacks.

See Also: *firewall, intrusion detection system (IDS), LAND attack, ping of death, Teardrop attack, Winnuke*

malformed URL attack

An attack that utilizes a nonstandard Uniform Resource Locator (URL).

Overview

A malformed URL attack is any exploit that attacks weaknesses in the URL-parsing algorithms of a Web server. An example of vulnerability to such an attack was found in the Internet Information Services (IIS) 5 platform and caused a memory allocation error that resulted in denial of service (DoS) to legitimate clients trying to connect to the Web server. By applying patches issued by vendors such coding weaknesses usually are resolved quickly.

See Also: *denial of service (DoS), dot bug vulnerability, input validation attack*

malicious code

Code that can cause harm to software or data.

Overview

While traditional malware usually includes viruses, worms, and Trojan horses, there are other kinds of mobile code that can cause harm to your systems and data when they run. Such code can arrive in your network through several routes, including e-mail attachments,

visits to Web sites, or wide area network (WAN) connections. Examples of code that can potentially perform malicious actions include ActiveX controls, Java applets, scripts on Web sites and in e-mail attachments, and Hypertext Markup Language (HTML) e-mail.

Signature-based virus protection software generally can't handle much of this code since attackers constantly are developing new exploits. As a result, various security vendors, including antivirus software vendors, have developed behavior-blocking products to handle the growing problem of malicious code.

Implementation

Behavior-blocking software monitors incoming mobile code in real time, identifies potentially harmful code by the actions it attempts to perform, and then blocks these actions from occurring. Behavior-blocking software generally works one of two ways:

- By confining all mobile code into a "sandbox" that restricts the ability of the code to access key operating system functions
- By intercepting kernel system calls and blocking actions attempted by mobile code

Marketplace

Examples of popular behavior-blocking software include eSafe Gateway from Aladdin, SurfinGate and SurfinShield from Finjan, InterScan AppletTrap from Trend Micro, and SafeTNet from Pelican Security.

See Also: *malware, virus protection software*

malware

Short for **malicious software**, a program developed for doing harm.

Overview

Malware is a term used to describe a wide range of software developed for malicious purposes that range from mischief to destruction of information. Examples of different types of malware include the following:

- **Viruses:** Programs that are spread manually by user action and infect other programs or data

- **Worms:** Programs that spread automatically by replicating themselves and infect other programs or data
- **Trojans:** Programs that masquerade as legitimate programs but perform malicious actions

In addition, the following could also be construed as forms of malware:

- **Spyware:** Programs that monitor a user's actions and secretly send information to third parties
- **Hoaxes:** Messages that masquerade as virus warnings and cause recipients alarm, forcing them to spend time and effort to resolve the nonexistent problem

See Also: hoax, spyware, Trojan, virus, worm

managed security service provider (MSSP)

A company that provides outsourced security services to businesses.

Overview

Outsourcing is a popular trend for companies seeking to reduce costs, and one of the hottest areas of outsourcing is security. With the growing proliferation of threats present on the Internet today, businesses pay a high price for not paying attention to the security of their networks, but the costs of internally managing network security can be high because of the special expertise required. Security professionals are in high demand and training internal staff may not be a cost-effective option, so many companies, both large and small, have chosen instead to outsource their security needs.

Managed security service providers (MSSPs) range from large organizations providing a broad range of security services to small companies targeting specific needs such as intrusion detection or incident response. Some Internet service providers (ISPs) are also beginning to offer managed security services to their clients as well. MSSP services involve provisioning security systems on the client site and then remotely monitoring and managing these systems. These security systems may be hardware or software and may include firewalls, intrusion detection systems (IDSs), virus protection

software, virtual private networks (VPNs), and content-filtering services.

Marketplace

Examples of popular MSSPs include Guardent, Counterpane, Foundstone, RipTech, TruSecure, and many others. As in any service provider environment, the landscape is constantly changing and businesses thinking of engaging the services of an MSSP should exercise due diligence in the selection process.

See Also: firewall, intrusion detection system (IDS), virtual private network (VPN), virus protection software

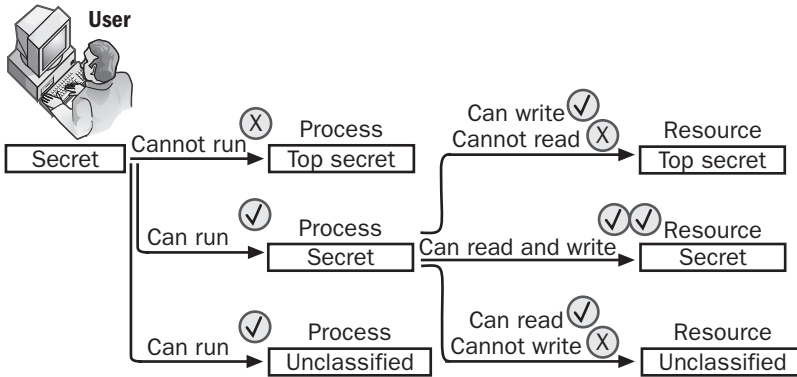
mandatory access control (MAC)

A mechanism for controlling access by users to computing resources.

Overview

Mandatory access control (MAC) is one of two basic approaches to implementing access control on computer systems; the other is discretionary access control (DAC). MAC systems control resources by confining them within security perimeters and enforcing policies that prevent resources from being moved from less secure to more secure environments. MAC systems work by assigning a security label to each user, process, or resource and then enforcing the following rules:

- A user is only allowed to run a process whose label is the same as or below that of the user's own label.
- A process is only allowed to read from a resource whose label is the same as or below that of the process's own label (no read-up allowed).
- A process is only allowed to write to a resource whose label is the same as or higher than that of the process's own label (no write-down allowed). Note that if a process writes to a resource whose label is higher than the process's own label, the process will subsequently be unable to read the information it has written.



Mandatory access control (MAC). Rules for mandatory access control (MAC).

The latter two points are known as the Bell-LaPadula Model, developed by D.E. Bell and L.J. LaPadula of the MITRE Corporation. Most MAC systems are based largely on work done in the 1970s by Bell and LaPadula.

MAC systems are more secure than those based on DAC, but are also more complex to manage. DAC systems give the creator (owner) of a resource the discretion to decide who is allowed to access the resource and what level of access that user can have, and the owner then may grant such access by configuring permissions on the resource. DAC thus assumes that users and processes are trustworthy.

By contrast, MAC takes the opposite approach and views users and processes as untrustworthy so that the creator of a resource does not have full control over its disposition. MAC systems thus give the site total control over who is allowed to access resources and what level of access they can have.

Marketplace

Few commercial operating systems support MAC because of the complexity of implementing and managing such systems. AIX 4.3.2 from IBM and Trusted Solaris 8 from Sun Microsystems are two commercial products that include support for MAC. A research

project called Security-Enhanced Linux (SELinux) undertaken by the Information Assurance Research Group of the National Security Agency (NSA) aims toward implementing MAC on the Linux platform.

See Also: access control, discretionary access control (DAC)

man-in-the-middle (MITM) attack

An attack in which the attacker impersonates both ends of a secure communication channel.

Overview

In a man-in-the-middle (MITM) attack, the attacker eavesdrops on a secure communication session to gain information that enables the attacker to impersonate both parties communicating. Some public key encryption systems are susceptible to MITM attacks, which require two things in order to be successful:

- The attacker must gain physical access to the communication channel to be able to capture the traffic when the two parties attempt to establish a secure communication channel.

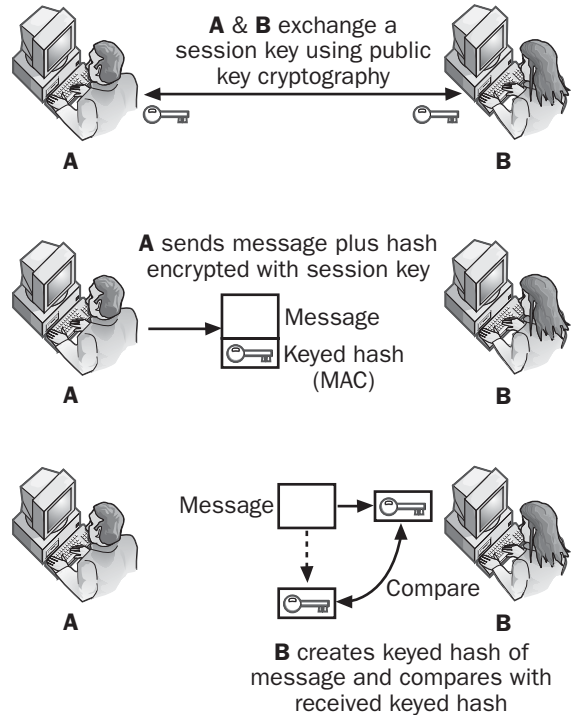
- The attacker must be able to intercept messages between the two parties and then relay them to maintain the session.

Some public key cryptography systems such as Diffie-Hellman (DH) are vulnerable to MITM attacks. The lack of authentication in DH means that recipients have no way of knowing whether a session key that has been exchanged with a sender actually belongs to the sender or to someone impersonating the sender.

Implementation

The way a MITM attack works is that at the beginning of a communication session, the sender requests the public key of the recipient so that the sender can use this key to encrypt a message to be sent to the recipient. The malicious party intercepts this request and returns its own public key instead of that of the recipient and masquerades as the recipient. The sender then exchanges a session key with the attacker, whom it thinks is the intended recipient, and all encrypted messages sent by the sender can now be read by the attacker. Meanwhile, the malicious party also masquerades as the sender and requests the public key of the recipient, exchanges a session key, and can also engage in encrypted communication with the recipient.

By incorporating digital signatures into public key systems, MITM attacks can be prevented. MITM attacks are also possible in other types of communication, including Transmission Control Protocol (TCP) sessions, but a more proper way of describing such attacks is TCP hijacking. The term **man-in-the-middle attack** is also applied sometimes to attacks against certain authentication protocols such as Challenge Handshake Authentication Protocol (CHAP), but this is not correct since the attacker impersonates only one side of the session—in a MITM attack the attacker impersonates both sides.



Man-in-the-middle (MITM) attack. How a man in the middle (MITM) attack works.

Notes

The MITM attack is sometimes called the bucket brigade attack, which derives from the ancient practice of putting out a fire by passing buckets of water from one person to the next between a source of water and a fire. The term **man in the middle** has a different source, deriving from a game in which two people try to throw a ball back and forth and a third person tries to intercept it.

See Also: Diffie-Hellman (DH), digital signature, public key, TCP session hijacking

master key

A key used for generating session keys.



Overview

Mutual authentication protocols generally employ a master key for generating session keys that can be used for encrypting data sent during a communication session. This master key is usually a shared secret key known to both parties and may be exchanged between the parties using a public key encryption system. An example of an authentication system that uses master keys is Kerberos, an authentication protocol developed by the Massachusetts Institute of Technology.

See Also: Kerberos, key, session key

MBSA

Stands for Microsoft Baseline Security Analyzer, an enhanced tool for identifying common security misconfigurations in Microsoft products.

See: Microsoft Baseline Security Analyzer (MBSA)

MCSA: Security

Stands for Microsoft Certified Systems Administrator: Security, a certification from Microsoft Corporation intended for systems administrators who focus on security in their job.

See: Microsoft Certified Systems Administrator (MCSA): Security

MCSE: Security

Stands for Microsoft Certified Systems Engineer: Security, a certification from Microsoft Corporation intended for systems engineers who focus on security in their job.

See Also: Microsoft Certified Systems Engineer (MCSE): Security

MD2

Stands for message digest 2, a hashing algorithm defined in RFC 1319.

See: message digest 2 (MD2)

MD4

Stands for message digest 4, a hashing algorithm defined in RFC 1320.

See: message digest 4 (MD4)

MD5

Stands for message digest 5, a hashing algorithm defined in RFC 1321.

See: message digest 5 (MD5)

meet-in-the-middle attack

A method for attacking secret key cryptographic systems.

Overview

If a portion of plaintext and its associated ciphertext can somehow be obtained, it may be possible to mount a meet-in-the-middle attack. Secret key systems that rely on an even number of keys are particularly susceptible to such attacks. The way such an attack could be mounted against an algorithm that successively encrypts plaintext using two different secret keys is as follows:

- 1 Create table 1, which contains all possible keys in column 1 and the result of encrypting the known portion of plaintext with each key in column 2.
- 2 Create table 2, which contains all possible keys in column 1 and the result of decrypting the known portion of ciphertext with each key in column 2.
- 3 Sort the two tables according to their second columns and then compare the second columns look-

ing for matches to find potential candidates for the two keys.

4 Test each match found to find the actual keys.

The amount of effort required to perform such an attack is only a few times more that of applying brute force to look for a single key. The possibility of mounting a meet-in-the-middle attack explains why Triple DES (3DES) employs three iterations instead of two.

See Also: 3DES, brute-force attack, ciphertext, plaintext

Melissa

A notorious macro virus that affects Microsoft Word.

Overview

Melissa first appeared in March 1999 as a Word macro virus that propagated rapidly across the Internet through attachments to e-mail messages whose subject line usually said “Important Message From <user>.” Opening the attachment caused the virus to infect the system and perform the following actions:

- Lowering the macro security settings in Word to allow macros to run without notifying the user.
- Sending copies of itself by e-mail to the first 50 entries of the user’s Microsoft Outlook Address Book.
- Infecting the Normal.dot template causing Word documents using this template to become infected with the virus.
- Modifying some infected documents by including additional text with the following phrase from the TV show *The Simpsons*: “Twenty-two points, plus triple-word-score, plus fifty points for using all my letters. Game’s over. I’m outta here.”

Additional effects of the virus included the following:

- Potential information leakage through sending Word documents on the infected user’s system to other users
- Heavy traffic through mail servers, often resulting in denial of service (DoS) conditions

See Also: macro virus, virus

message

Data that has been encoded for transmission or delivery between two or more parties.

Overview

Messages represent the content of communication systems, and in a business environment the security of such transmissions is paramount. The security of messages and messaging systems has a number of aspects, including the following:

- **Confidentiality:** The assurance that the content of a message is known only to its intended recipients. Confidentiality of messages is generally achieved through encryption.
- **Integrity:** The assurance that the content of a message has not been modified in transit. Integrity of messages can be achieved by creating a message digest (MD) or digital signature.
- **Availability:** The assurance that the content of a message can be accessed when required by those allowed to access it.
- **Authentication:** The assurance that the identity of the sender of a message can be proved to the recipient as correct.
- **Nonrepudiation:** The assurance that the identity of the sender of a message can be proved to a third party as correct.

Together these five aspects of message security comprise what is known as information assurance (IA).

See Also: authentication, digital signature, encryption, information assurance (IA), integrity, message digest (MD), nonrepudiation

message authentication code (MAC)

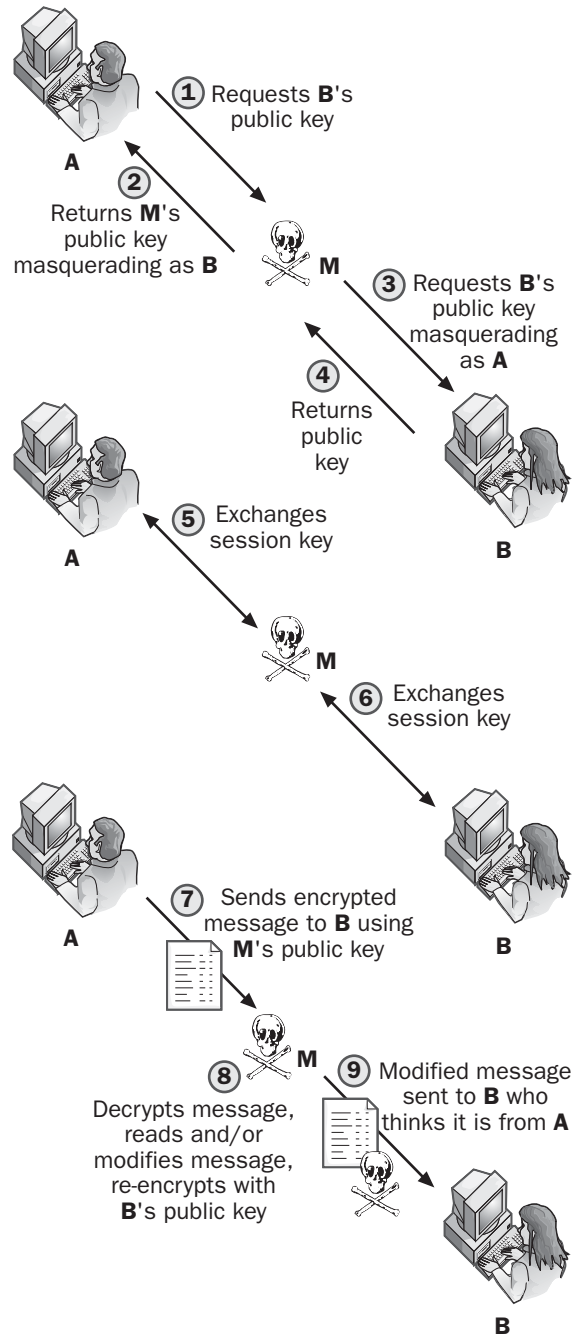
A keyed hashing algorithm used to guarantee the integrity of a message.

Overview

A message authentication code (MAC) encrypts a message digest with a session key to provide assurance that the content of a message has not been modified in transit. A hashing algorithm is first applied to the message to generate a **hash**, a short, fixed-length cryptographic string that uniquely represents the message. The sender then encrypts the hash using a session key, which is a shared secret key known to both sender and recipient. The resulting MAC is then attached to the message and sent. When the message is received, the recipient decrypts the MAC using the same session key to recover the hash. The recipient then hashes the original message and compares this to the received hash. If the two hashes match, the recipient knows that the message has integrity and has not been modified in transit.

Examples of message authentication codes include the following:

- Cipher block chaining message authentication code (CBC-MAC)
- Hash-based message authentication code (HMAC), an Internet standard defined in RFC 2104
- Nested message authentication code (NMAC)
- Universal-hashing message authentication code (UMAC)



Message authentication code (MAC). How a MAC is used to verify the integrity of a message.

Notes

Another method for guaranteeing the integrity of a message is to attach a digital signature to the message. The difference between MACs and digital signatures is that MACs use session (secret) keys, while signatures use public key encryption. The session keys used by MACs are themselves usually exchanged using public key encryption.

See Also: *digital signature, hash-based message authentication code (HMAC), hashing algorithm, integrity, message digest (MD)*

message digest (MD)

A cryptographic checksum used to verify that an electronic message has not been modified in transit.

Overview

Message digests (MDs) are used to verify the integrity of electronic messages to provide assurance that their content has not been modified in transit. MDs are based on hashing algorithms, mathematical procedures for generating a fixed-size result from arbitrary amounts of data. MDs perform a similar function to cyclical redundancy checks (CRCs) used in networking and telecommunication, but are cryptographically stronger and better able to protect information against accidental or intentional modification during transmission.

Examples of popular hashing algorithms used to create MDs include message digest 5 (MD5) and secure hash algorithm 1 (SHA-1). To guarantee message integrity, MDs are combined either with secret key cryptography to create a message authentication code (MAC) or with public key cryptography to create a digital signature.

See Also: *digital signature, hashing algorithm, integrity, message authentication code (MAC), message digest 2 (MD2), message digest 4 (MD4), message digest 5 (MD5), Secure Hash Algorithm-1 (SHA-1)*

message digest 2 (MD2)

A hashing algorithm defined in RFC 1319.

Overview

Message digest 2 (MD2) was developed by Ron Rivest in 1989 as one of the first algorithms for creating mes-

sage digests (MDs), cryptographic checksums used to verify that an electronic message has not been modified in transit. MD2 was optimized for 8-bit processors and has since been replaced by message digest 4 (MD4) and message digest 5 (MD5).

Implementation

To apply MD2 to a message of arbitrary length, first pad the message to make its number of bytes a multiple of 16. Next, append to the end of the message a 16-byte checksum that is mathematically derived from the message. Finally, iteratively process each 16 bytes of the message until a 16-byte (128-bit) message digest results.

Notes

Message digest 1 (MD1 or simply MD) was a proprietary algorithm that was never published.

See Also: *hashing algorithm, message digest (MD), message digest 4 (MD4), message digest 5 (MD5)*

message digest 4 (MD4)

A hashing algorithm defined in RFC 1320.

Overview

Message digest 4 (MD4) was developed by Ron Rivest in 1990 as the successor of his earlier message digest 2 (MD2) algorithm. MD4 was optimized for 32-bit processors but was later shown to be insecure and was replaced by message digest 5 (MD5).

Implementation

To apply MD4 to a message of arbitrary length, first pad the message by adding a single 1 bit followed by a string of 0 bits so that the result is a string that is 64 bits less than a multiple of 512. Append to this a 64-bit number equal to the number of bits in the original message modulo 2^{64} . The result is a string whose length is a multiple of 512 bits, which equals sixteen 32-bit words. This is then iteratively processed 512 bits at a time using a three-stage compression function until a 128-bit (four 32-bit word) message digest finally results.

Notes

There was in fact a message digest 3 (MD3), but it was superseded by MD4 and never published.

See Also: *hashing algorithm, message digest (MD), message digest 2 (MD2), message digest 5 (MD5)*

message digest 5 (MD5)

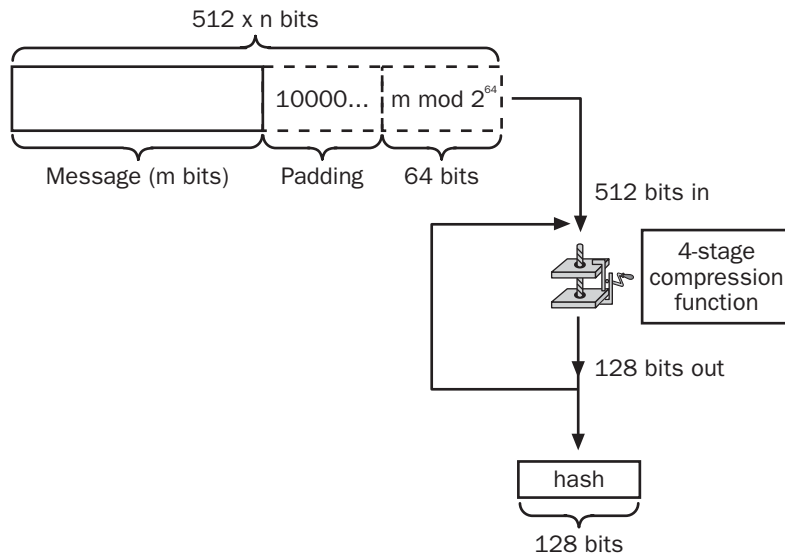
A hashing algorithm defined in RFC 1321.

Overview

Message digest 5 (MD5) was developed by Ron Rivest in 1991 as a modified version of his earlier message digest 4 (MD4) algorithm, which was found to be insecure because of collisions. MD5 is a popular algorithm optimized for 32-bit processors and widely used in cryptographic systems today.

Implementation

To apply MD5 to a message of arbitrary length, first pad the message by adding a single 1 bit followed by a string of 0 bits so that the result is a string that is 64 bits less than a multiple of 512. Append to this a 64-bit number equal to the number of bits in the original message modulo 2^{64} . The result is a string whose length is a multiple of 512 bits, which equals sixteen 32-bit words. This is then iteratively processed 512 bits at a time using a four-stage compression function until a 128-bit (four 32-bit word) message digest finally results. The main differences between MD4 and MD5 are the complexity and number of passes of the compression function.



Message digest 5 (MD5). How MD5 works.

See Also: hashing algorithm, message digest (MD), message digest 2 (MD2), message digest 4 (MD4)

message integrity code (MIC)

Another name for message authentication code (MAC), a keyed hashing algorithm used to guarantee the integrity of a message.

See: message authentication code (MAC)

MIC

Stands for message integrity code, another name for message authentication code (MAC), a keyed hashing algorithm used to guarantee the integrity of a message.

See: message authentication code (MAC)

Microsoft Baseline Security Analyzer (MBSA)

An enhanced tool for identifying common security misconfigurations in Microsoft products.

Overview

Microsoft Corporation developed Microsoft Baseline Security Analyzer (MBSA) as part of its Strategic Technology Protection Program (STPP) to help customers ensure the security of their systems. MBSA scans local and remote systems looking for common configuration problems, including missing service packs and security updates. Version 1.1 of MBSA can detect misconfigurations in the following products and applications:

- Microsoft Windows NT 4
- Microsoft Windows 2000
- Microsoft Windows XP
- IIS 4 and 5
- SQL Server 7 and 2000
- Internet Explorer 5.01 or later
- Microsoft Office 2000 and 2002

MBSA also incorporates the Hotfix Checker HfNetChk and support for Software Update Service (SUS) during security scans. MBSA generates Extensible Markup Language (XML) security reports for each system scanned and can display these reports in Hypertext Markup Language (HTML) format.

Notes

MBSA was developed for Microsoft by Shavlik Technologies LLC (www.shavlik.com). MBSA replaces the earlier Microsoft Personal Security Advisor (MPSA), also developed by Shavlik.

For More Information

Visit www.microsoft.com/technet/security/tools/Tools/MBSAhome.asp for more information.

See Also: HfNetChk, hotfix, service pack (SP), Software Update Services (SUS)

Microsoft Certified Systems Administrator (MCSA): Security

A certification from Microsoft Corporation intended for systems administrators who focus on security in their job.

Overview

Microsoft Certified Systems Administrator (MCSA): Security is a certification developed by Microsoft based on the Microsoft Certified Systems Administrator (MCSA) credential. The certification allows systems administrators to demonstrate deep, role-based skills in implementing security on the Microsoft Windows 2000 or Windows Server 2003 platform and also highlights a specialty appropriate to creating a secure computing environment.

To obtain MCSA: Security certification, individuals must demonstrate a security specialty and pass four core exams in the following areas:

- Client Operating System (1 exam)
- Networking System (2 exams)
- Security Implementation (1 exam)

The security specialty requirement can be met by passing Microsoft Certified Professional (MCP) Exam 70-227, *Installing, Configuring, and Administering Microsoft Internet Security and Acceleration Server 2000, Enterprise Edition*, or by obtaining CompTIA Security+ certification. Other options may also be available.

For More Information

Visit www.microsoft.com/traincert/mcp/mcsasecurity/ for more information.

See Also: Certified Information Systems Security Professional (CISSP), Global Information Assurance Certification (GIAC), Microsoft Certified Systems Engineer (MCSE): Security, Security+

Microsoft Certified Systems Engineer (MCSE): Security

A certification from Microsoft Corporation intended for systems engineers who focus on security in their job.

Overview

Microsoft Certified Systems Engineer (MCSE): Security is a certification developed by Microsoft based on the Microsoft Certified Systems Engineer (MCSE) credential. The certification allows systems engineers to

demonstrate deep, role-based skills in designing and implementing security on the Microsoft Windows 2000 or Windows Server 2003 platform and also highlights a specialty appropriate to creating a secure computing environment.

To obtain MCSE: Security certification, individuals must demonstrate a security specialty and pass core exams (six on Windows 2000, seven on Windows Server 2003) in the following areas:

- Client Operating System (1 exam)
- Networking System (3 exams on Windows 2000, 4 exams on Windows Server 2003)
- Security Design (1 exam)
- Security Implementation (1 exam)

The security specialty requirement can be met by passing Microsoft Certified Professional (MCP) Exam 70-227, *Installing, Configuring, and Administering Microsoft Internet Security and Acceleration Server 2000, Enterprise Edition*, or by obtaining CompTIA Security+ certification. Other options may also be available.

For More Information

Visit www.microsoft.com/traincert/mcp/mcsesecurity/ for more information

See Also: *Certified Information Systems Security Professional (CISSP), Global Information Assurance*

Certification (GIAC), Microsoft Certified Systems Administrator (MCSA): Security, Security+

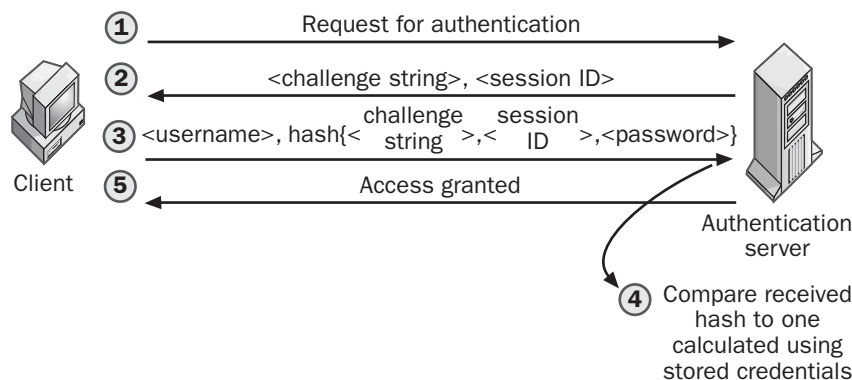
Microsoft Challenge Handshake Authentication Protocol (MS-CHAP)

An authentication protocol used with Point-to-Point Protocol (PPP).

Overview

Microsoft Challenge Handshake Authentication Protocol (MS-CHAP) was developed by Microsoft Corporation as an enhanced version of Challenge Handshake Authentication Protocol (CHAP), an industry standard wide area network (WAN) authentication protocol defined in RFC 1994. MS-CHAP authenticates remote access users using a handshaking process as follows:

- 1 The authentication server sends a challenge string and session identifier to the client attempting authentication.
- 2 The client responds with the user name and a non-reversible encryption of a string made up of the challenge string, session ID, and user's password.
- 3 The authentication server generates a similar encrypted response string using stored user credentials and compares the two strings.
- 4 If the two strings are the same, the client is authenticated and allowed to connect.



Microsoft Challenge Handshake Authentication Protocol (MS-CHAP). How MS-CHAP authenticates users.

MS-CHAP was originally included in the Microsoft Windows NT platform for authenticating remote access clients over PPP connections. An enhanced version of MS-CHAP called MS-CHAPv2 was later developed for Microsoft Windows 2000. MS-CHAPv1 is outlined in RFC 2433, and MS-CHAPv2 in RFC 2759.

See Also: authentication

Microsoft Personal Security Analyzer (MPSA)

A tool for identifying common security misconfigurations in Microsoft products, now supplanted by the Microsoft Baseline Security Analyzer (MBSA).

Overview

Microsoft Personal Security Analyzer (MPSA) was designed to scan desktop computers running Microsoft Windows NT 4 Workstation or Microsoft Windows 2000 Professional to look for common configuration problems and recommend how to correct them. Examples of such problems included weak passwords, missing patches, insecure macro settings in Microsoft Office, and weak security settings in Microsoft Internet Explorer and Microsoft Outlook Express.

MPSA was a standalone application implemented as an ActiveX control and could scan only the machine on which it was installed. MPSA has now been replaced by MBSA, which can perform both local and remote scans and incorporates additional features such as the Hotfix Checker HfNetChk.

See Also: HfNetChk, Microsoft Baseline Security Analyzer (MBSA)

Microsoft Security & Privacy

The portion of Microsoft Corporation's Web site devoted to the security of its products.

Overview

The Microsoft Security & Privacy site is the primary source of information about the security of Microsoft platforms and products for individuals. Information on the site is targeted toward several groups of users, including the following:

- **IT (information technology) professionals:** Security tools, checklists, best practices, and how to handle recently discovered vulnerabilities
- **Developers:** Samples of secure code, technical articles, and core documentation of Microsoft products
- **Businesses:** White papers, risk analysis, and strategies for protecting business assets
- **Home users:** Step-by-step tutorials and tips for protecting privacy and securing computer systems

The site also includes a Communities section, which provides access to security-related newsgroups, technical chats, and webcasts.

For More Information

Visit www.microsoft.com/security/ for more information.

See Also: Microsoft Security Notification Service

Microsoft Security Notification Service

A service that notifies customers about threats to Microsoft products and platforms.

Overview

With the growing proliferation of worms, viruses, and other threats on the Internet, it's critical for administrators to be aware of the latest dangers and how to protect their systems against them. Microsoft Security Notification Service provides customers with a free e-mail notification service that provides bulletins that include the following:

- The nature of the threat
- What Microsoft products it affects
- How you can protect your systems against it
- How Microsoft plans to address the problem

Microsoft Corporation recommends that all its customers subscribe to this service to be made aware in a timely fashion of the latest patches and security updates as they become available. All security bulletins are digitally signed so recipients can verify that they were in fact published by Microsoft.

For More Information

Visit www.microsoft.com/technet/security/bulletin/notify.asp for information on how to subscribe to this service.

To notify Microsoft of a suspected security vulnerability in one of its products, visit www.microsoft.com/technet/security/bulletin/alertus.asp and submit the required information.

To review past security bulletins, which are archived on Microsoft TechNet, visit www.microsoft.com/technet/security/current.asp.

See Also: *Microsoft Security & Privacy*

Microsoft Security Response Center (MSRC)

A team of security professionals at Microsoft responsible for responding to security threats involving Microsoft products.

Overview

Microsoft Security Response Center (MSRC) is a virtual team of hundreds of security professionals at Microsoft that acts as the hub of security-related activity for the company. The center handles all customer communication involving security-related issues and interacts with development teams who create patches to fix vulnerabilities. The team also works closely with development teams to plan the security of future Microsoft products.

When MSRC receives a report from a customer of a potential vulnerability in a Microsoft product, it begins a “triaging” process and assigns a tracking number to track every stage of the investigation. The appropriate development teams are informed and a team is created to try to reproduce the problem, determine its cause, and develop a solution. The customer is kept in the loop at each stage and is offered the opportunity to provide further feedback if required. Once a solution is developed for the issue, it is thoroughly tested prior to being released.

Solutions to security vulnerabilities can take several forms:

- A patch or hotfix that is issued immediately to resolve the problem if its severity warrants it
- A workaround that is provided as an interim solution while a patch is being developed

Patches are also later consolidated into service packs, and applying noncritical patches can often be delayed until the next service pack is released.

Only about 10 percent of all reports received by MSRC pass the initial triaging stage, and only about 10 percent of these are eventually determined as vulnerabilities that require patches. Vulnerabilities identified by MSRC are rated according to the following scale:

- **Critical:** Requires immediate patching to protect systems against a severe threat such as an Internet worm.
- **Important:** Systems should be patched immediately to prevent compromise of availability, confidentiality, or integrity of user’s data.
- **Moderate:** There are significant mitigating factors that make exploiting this vulnerability unlikely, but cautious administrators may want to immediately patch their systems just in case.
- **Low:** Exploiting this vulnerability is extremely unlikely, and administrators can decide whether to patch their systems immediately or wait for the next service pack.

For More Information

Visit www.microsoft.com/technet/security/bulletin/alertus.asp to report a suspected security vulnerability to the MSRC team.

See Also: *hotfix, Microsoft Security & Privacy, service pack (SP), workaround*

Microsoft Security Toolkit

A set of tools from Microsoft for helping customers protect their systems.

Overview

The Microsoft Security Toolkit is designed to help customers protect their systems against common security threats and vulnerabilities. The toolkit includes an assortment of tools to provide baseline security for servers connected to the Internet. It also includes guides for securing systems and patches for vulnerabilities identified by the Microsoft Security Response Center (MSRC) as being potentially dangerous to the security of servers. The toolkit applies specifically to the Microsoft Windows NT 4 platform and is available in CD-ROM format as a standalone product and as part of a Microsoft TechNet subscription.

For More Information

Visit www.microsoft.com/security/kitinfo.asp for more information.

See Also: Microsoft Security Response Center (MSRC)

Microsoft Security Update

A service that notifies home and small business users about threats to Microsoft products and platforms.

Overview

Microsoft Security Update is a simplified version of Microsoft Security Notification Service, which is targeted mainly toward enterprise customers. The Security Update is an e-mail alert service that notifies consumers whenever security updates become available for Microsoft products. The articles are written in nontechnical language and contain links to more information on Microsoft Corporation's Security & Privacy Web site.

For More Information

Visit www.microsoft.com/security/security_bulletins/decision.asp for information on how to subscribe to this service.

See Also: Microsoft Security & Privacy, Microsoft Security Notification Service

Microsoft Strategic Technology Protection Program (STPP)

An initiative launched by Microsoft Corporation in October 2001 to help protect its customers against threats from the Internet.

Overview

Microsoft Strategic Technology Protection Program (STPP) was launched largely in response to Code Red and Nimda, two Internet worms that wreaked havoc on Microsoft platforms and systems. STPP was developed to facilitate security-related product support for enterprise customers and initially included several new initiatives:

- A Security Tool Kit containing current service packs and critical security patches for Microsoft Windows NT 4, Microsoft Windows 2000, Microsoft Internet Information Services (IIS), and Microsoft Internet Explorer
- Free virus-related product support through a toll-free hotline 1-866-PC SAFETY within the United States and Canada
- Security rollups, packages that contain cumulative hotfixes and can be downloaded and applied using Microsoft Windows Update
- Microsoft Personal Security Analyzer (MPSA), a tool for identifying common security misconfigurations in Microsoft products

More recent initiatives of the program included these:

- Microsoft Baseline Security Analyzer (MBSA), an enhanced tool for identifying common security misconfigurations in Microsoft products
- Software Update Services (SUS), a tool for automatically deploying critical updates across the enterprise

For More Information

Visit www.microsoft.com/security/mstpp.asp for more information.

See Also: Microsoft Baseline Security Analyzer (MBSA), Microsoft Personal Security Analyzer (MPSA), security rollup package, Software Update Services (SUS)

Microsoft TechNet Security

The portion of the Microsoft TechNet Web site devoted to security issues.

Overview

The Microsoft TechNet Security site contains a vast amount of information useful to IT (information technology) professionals who need to ensure Microsoft platforms and products are deployed, configured, and administered securely. Some of the topics covered by the site include the following:

- Access to the latest security bulletins from Microsoft Corporation and information on how to subscribe to the Microsoft Security Notification Service
- Security resources, including assessment tools, checklists, best practices, how-to tutorials, case studies, security tips, service packs, rollup packages, and hotfixes
- Access to security-related newsgroups and information on how to contact the Microsoft Security Team
- Information on how to protect, detect, defend, recover, and manage security-related issues
- Links to security-related training, books, and third-party products and services

For More Information

Visit www.microsoft.com/technet/security/ for more information.

See Also: *Microsoft Security & Privacy*

MITM

Stands for man-in-the-middle attack, an attack in which the attacker impersonates both ends of a secure communication channel.

See: *man-in-the-middle (MITM) attack*

Morris worm

A notorious Internet worm that also acted as a virus.

Overview

The Morris worm was one of the first worms to cause damage to systems and achieve widespread media recognition. The worm was developed by Robert Morris, Jr., a

student at Cornell University, and though originally designed simply to spread and not cause harm, a coding error caused the worm to repeatedly replicate itself until it consumed available memory, filled free space on hard drives, and drove processor utilization to 100 percent. The result was denial of service (DoS) for legitimate users and systems that crashed and needed their hard drives to be cleaned before they could restart.

The Morris worm first appeared in November 1988 and spread rapidly across the Internet, infecting Sun servers and VAX minicomputers by exploiting vulnerabilities in the Sendmail, Fingerd, Rsh, and Exec daemons on UNIX platforms. The worm infected at least 6000 systems, which at the time represented about 10 percent of the Internet, and the resulting cleanup costs and business downtime was estimated at \$98 million.

One of the results of the Morris worm was the formation of the Computer Emergency Response Team (CERT), later the CERT Coordination Center (CERT/CC), at Carnegie Mellon University, to respond to such incidents in the future. Another result was Morris's conviction under the Computer Fraud and Abuse Act, a U.S. federal law that was first applied in the Morris case. After several appeals, Morris was eventually sentenced to three years probation, a \$10,050 fine, and 400 hours of community service, and he went on to become an assistant professor at Massachusetts Institute of Technology (MIT). Interestingly, Morris's father, Robert Morris, Sr., was a famous cryptographer at the National Computer Security Center (NCSC) of the National Security Agency (NSA).

See Also: *CERT Coordination Center (CERT/CC), worm*

MPSA

Stands for Microsoft Personal Security Analyzer, a tool for identifying common security misconfigurations in Microsoft products, now supplanted by the Microsoft Baseline Security Analyzer (MBSA).

See: *Microsoft Personal Security Analyzer (MPSA)*

MS-CHAP

Stands for Microsoft Challenge Handshake Authentication Protocol, an authentication protocol used with Point-to-Point Protocol (PPP).

See: Microsoft Challenge Handshake Authentication Protocol (MS-CHAP)

MSRC

Stands for Microsoft Security Response Center, a team of security professionals at Microsoft responsible for responding to security threats involving Microsoft products.

See: Microsoft Security Response Center (MSRC)

MSSP

Stands for managed security service provider, a company that provides outsourced security services to businesses.

See: managed security service provider (MSSP)

Mstream

A tool for launching distributed denial of service (DDoS) attacks.

Overview

Mstream is a DDoS tool that uses a handler/agent architecture similar to Trin00 and other common exploits. The signature of an Mstream attack is a flood of Transmission Control Protocol (TCP) packets that have their acknowledgment (ACK) flag set. These packets generally have random source Internet Protocol (IP) addresses and random source and destination TCP socket numbers. The target host responds with large numbers of TCP Reset (RST) packets sent to nonexistent hosts, resulting in bandwidth starvation and excessive central processing unit (CPU) utilization.

The Mstream attack is a modified version of the Stream exploit, an older denial of service (DoS) attack.

See Also: distributed denial of service (DDoS), Trin00

mutual authentication

Authentication of both ends of a communication session.

Overview

Traditional network authentication systems have centered around having the server authenticate the credentials of the client. They ignore authentication of the server by the client since it is assumed that the server is always a trusted entity. However, it is sometimes possible to spoof the identity of a server, especially in an Internet scenario in which information is sent over an insecure public communication system and is subject to eavesdropping, interception, and hijacking. Although simple consumer transactions such as users buying goods online may suffice with one-way authentication of clients by e-commerce servers, more costly business-to-business (B2B) and financial industry transactions need both ends of a communication channel to be authenticated before establishing a session and performing a transaction. **Mutual authentication** is the general term for any scheme by which both parties authenticate the other prior to sending sensitive information to each other.

One protocol that was developed for mutual authentication is Kerberos, a popular authentication protocol developed by the Massachusetts Institute of Technology (MIT) and used by Active Directory directory service in Microsoft Windows 2000 and Windows Server 2003. Other mutual authentication protocols include the following:

- Microsoft Challenge Handshake Authentication Protocol version 2 (MS-CHAPv2)
- Extensible Authentication Protocol/Transport Layer Security (EAP/TLS)
- Symmetric-Key Three-Pass Mutual Authentication Protocol defined in the ISO 9798 standard

See Also: authentication, Kerberos

NAT

Stands for network address translation, a mechanism for translating Internet Protocol (IP) addresses between two networks.

See: network address translation (NAT)

National Computer Security Center (NCSC)

An initiative of the National Security Agency (NSA) focused on information security (infosec).

Overview

The National Computer Security Center (NCSC) began in 1981 as the Department of Defense Computer Security Center and was a partnership between government, industry, and academia devoted to promoting research and development in information systems security.

Together with the Trusted Product Evaluation Program (TPEP), another NSA initiative, the NCSC operates a program for evaluating commercially developed computing equipment designed for high-security environments to ensure their capability for securely processing classified information. Together with other government agencies such as the National Institute of Standards and Technology (NIST), the NCSC also develops and publishes criteria and standards for developing trusted information systems.

For over two decades the NCSC promoted infosec awareness through an annual National Information Systems Security Conference, but this was discontinued in 2000. The NCSC also developed and published the legendary *Orange Book*, the Trusted Computer System Evaluation Criteria (TCSEC) used by the Department of Defense for designing secure information systems.

For More Information

Visit www.nsa.gov/isso/partners/ncsc.htm for more information.

See Also: infosec, National Security Agency (NSA), Trusted Computer System Evaluation Criteria (TCSEC)

National Fraud Information Center (NFIC)

An organization helping consumers and law enforcement agencies fight fraud.

Overview

The National Fraud Information Center (NFIC) was formed by the National Consumers League (NCL) in 1992 and provides a national toll-free hotline (1-800-876-7060 from 9 to 5 Monday through Friday) that consumers can call if they think they are victims of telemarketing or Internet fraud. NFIC also operates Internet Fraud Watch, which provides consumers with advice concerning various promotions and fraud schemes propagated on the Internet through Web sites and e-mail, including tips on how to recognize a fraud such as the following:

- A bogus credit card offer
- Amazingly cheap computer equipment and software
- Pyramid schemes and Nigerian money offers
- So-called advance fee loans
- Charity and scholarship scams
- Credit repair services and credit card loss protection schemes
- Business opportunities and work-at-home scams
- Fraudulent online auctions

For more information

Visit www.fraud.org for more information.

See Also: privacy

National Information Assurance Certification and Accreditation Process (NIACAP)

A U.S. national standards process for information assurance (IA) accreditation.

Overview

National Information Assurance Certification and Accreditation Process (NIACAP) outlines a standardized process for certification and accreditation (C&A) of the security of information systems for all departments of the executive branch of the U.S. government and any contractors and consultants that have dealings with it. It outlines a set of activities, tasks, and accompanying management structure to ensure information systems meet their documented security requirements and will continue to do so throughout their life cycle. NIACAP is defined as a National Security Telecommunications and Information System Security Instruction (NSTISSI) developed and issued by the Committee on National Security Systems (CNSS), which was formerly called the National Security Telecommunications and Information Systems Security Committee (NSTISSC).

Implementation

NIACAP specifies a four-step accreditation process:

- **Definition phase:** Defines the C&A level of effort resulting in agreement regarding methods to be used for implementing the security requirements of the information system under consideration. The agreement documents the purposes, architecture, target environment, security requirements, and access policies for the system.
- **Verification phase:** This involves testing the compliance of the system with the previously developed agreement and may involve refining the agreement or modifying the system as appropriate.

- **Validation phase:** The actual certification and accreditation phase in which the system has been integrated and measured as complying with security requirements within an acceptable range of variation.
- **Postaccreditation phase:** Activities needed to ensure the system will continue to operate securely within an acceptable range of variation, including ongoing monitoring, maintenance, reviewing of the agreement, and validation of compliance.

Some of the many areas covered by NIACAP include the following:

- Design analysis of hardware, software, and system architectures
- Life cycle management analysis
- Risk assessment and risk management
- Contingency planning
- Audit trails
- Access control
- Data integrity
- Penetration testing
- Personnel security
- Physical access control
- Threat assessment

Notes

The parallel process for IA certification in the U.S. defense industry is the Department of Defense Information Technology Security Certification and Accreditation Process (DITSCAP).

For More Information

Visit www.nstissc.gov for more information.

See Also: Department of Defense Information Technology Security Certification and Accreditation Process (DITSCAP), information assurance (IA)

National Information Assurance Partnership (NIAP)

A cooperative agency for promoting information security (infosec) in U.S. government agencies and private industry.

Overview

The National Information Assurance Partnership (NIAP) is a collaboration between the National Security Agency (NSA) and the National Institute of Standards and Technology (NIST) that was formed in 1997 in response to the President's Commission on Critical Infrastructure Protection. The mandate for NIAP is to promote information assurance (IA) by fostering the development of standards, practices, testing methods, tools, techniques, accreditation, and metrics. NIAP also tries to promote internal standards for IT (information technology) security and facilitate the growth of the IA industry within the United States. To foster information security (infosec) in the federal government, NIAP is developing guidelines for agencies to lock down their networks to protect them from cyberattack.

For More Information

Visit niap.nist.gov for more information.

See Also: *information assurance (IA), infosec, National Institute of Standards and Technology (NIST), National Security Agency (NSA)*

National INFOSEC Education & Training Program (NIETP)

An information security training program from the National Security Agency (NSA).

Overview

The National INFOSEC Education & Training Program (NIETP) is designed to help safeguard national security information systems by developing the information security (infosec) skills of the U.S. workforce. The NIETP works through education and training programs for schools and the workplace, primarily through nationwide leadership and advocacy of infosec training through community-based education. The goal of NIETP is to ensure that all government personnel are

trained and knowledgeable in how to safeguard national information systems, primarily by providing initiatives to multiply the number of trained and certified information security professionals working in government agencies. The NIETP directly supports the goals and aims of the Committee on National Security Systems (CNSS), which was formerly called the National Security Telecommunications and Information Systems Security Committee (NSTISSC).

For More Information

Visit www.nsa.gov/isso/programs/nietp/intro.htm for more information.

See Also: *infosec, National Security Agency (NSA)*

National Infrastructure Protection Center (NIPC)

A cooperative agency set up to help protect critical information system infrastructures.

Overview

The mandate of the National Infrastructure Protection Center (NIPC) is to act as a focal point for helping the U.S. government and law enforcement agencies detect, assess, warn, and respond to threats to the national information infrastructure. Such threats may involve unlawful acts that threaten information technologies, including both physical and cyberattacks. The NIPC is responsible for managing investigation of intrusion into federal information systems, supporting law enforcement in responding to threats and acts of cyberterrorism, and coordinating training of forensic cyberinvestigators for both government and industry.

The NIPC, a cooperative effort that includes federal, state, and local government agencies and the private sector, was propelled into existence through recommendations of the President's Commission on Critical Infrastructure Protection. The NIPC was formed in 1998 as a joint initiative of the Department of Justice and the Federal Bureau of Investigation (FBI) and operates out of FBI headquarters in Washington, D.C. In March 2003 the NIPC began transitioning to the newly formed Department of Homeland Security, from where it will operate in the future.

For More Information

Visit www.nipc.gov for more information.

See Also: *cybercrime, infosec*

National Institute of Standards and Technology (NIST)

A U.S. federal government agency that develops standards for government and private industry sectors.

Overview

The National Institute of Standards and Technology (NIST) is a nonregulatory agency that operates within the Technology Administration division of the U.S. Commerce Department. NIST's mandate is to develop and promote standards in all areas of technology in order to enhance business productivity and facilitate trade. NIST conducts its activities at several laboratories and through different cooperative programs in partnership with government and the private sector.

In the field of computer security, NIST has a Computer Security Division (CSD) that maintains a Computer Security Resource Center (CSRC) covering work in five areas:

- Cryptographic standards and applications
- Security assessment and validation
- Research into emerging security technologies
- Development of guidelines for secure management of resources
- Security awareness, training, and education (ATE) outreach programs

In addition to its many contributions to the field of information security (infosec), NIST issues a series of Federal Information Processing Standards (FIPS) publications that define standards in data processing, encryption, security, and related areas. An important initiative of NIST in the computer security area is NIST Special Publication 800-37, "Guidelines for the Security Certification and Accreditation of Federal Information Technology Systems," a set of guidelines released

in 2002 designed to help protect sensitive federal information systems.

For More Information

Visit www.nist.gov for more information.

See Also: *Computer Security Division (CSD), Federal Information Processing Standard (FIPS)*

National Security Agency (NSA)

A U.S. agency responsible for protecting national information systems and producing foreign intelligence information.

Overview

The National Security Agency (NSA) and its partner agency, the Central Security Service (CSS), the liaison between the NSA and the Armed Forces, is the primary U.S. government agency responsible for the development and use of cryptographic technologies. The NSA/CSS is one of 13 different federal agencies that comprise the U.S. Intelligence Community; some others are the Central Intelligence Agency (CIA), the Federal Bureau of Investigation (FBI), the Defense Intelligence Agency (DIA), and the Department of Energy (DOE).

The role of the NSA in national information systems security involves the protection of sensitive and classified information through encryption technologies. The NSA in fact is the leading employer of mathematicians in the United States and is a world leader in research and development related to cryptology. Mathematicians working at the NSA are generally involved in one of two activities:

- Devising new encryption algorithms, technologies, and devices
- Trying to crack the encryption technologies of other nations

From an operational perspective, the mission of the NSA/CSS is twofold:

- Developing information assurance (IA) solutions to protect critical national information systems, assets, and infrastructures

- Obtaining signal intelligence (SIGINT) by monitoring and decoding foreign communications for the protection of national interests

Some initiatives and programs of the NSA include the following:

- **National Computer Security Center (NCSC):** An initiative of the NSA focused on information security (infosec)
- **National Information Assurance Partnership (NIAP):** A cooperative agency for promoting information security (infosec) in U.S. government agencies and private industry, operated in conjunction with the National Institute of Standards and Technology (NIST)
- **National INFOSEC Education & Training Program (NIETP):** An information security training program from the NSA

The NSA was also involved in the development of the first large-scale computer in the 1950s and the first solid-state computer in the 1970s, which underlines its close involvement with high technology and information security (infosec). The NSA has also developed a series of Security Recommendation Guides outlining steps for secure configuration of different operating systems, including Microsoft Windows NT, Windows 2000, Windows XP, and Cisco IOS. These guides are being used by many government agencies as baselines for ensuring the security of their information systems.

Notes

The NSA operates a National Cryptologic Museum that is open to the general public.

For More Information

Visit www.nsa.gov for more information.

See Also: *cryptology, National Computer Security Center (NCSC), National Information Assurance Partnership (NIAP), National INFOSEC Education & Training Program (NIETP), National Institute of Standards and Technology (NIST)*

National Strategy to Secure Cyberspace

The information security (infosec) component of the U.S. National Strategy for Homeland Security.

Overview

Together with the National Strategy for the Physical Protection of Critical Infrastructures and Key Assets, the National Strategy to Secure Cyberspace is a framework for organizing and prioritizing efforts to enhance cyberspace security. The strategy outlines steps that can be followed by state and local governments, private companies and organizations, and even individual citizens to collectively improve national information security. The strategic objectives of the strategy are as follows:

- To prevent cyberattacks against critical information infrastructures
- To reduce the vulnerability of national information systems to such attacks
- To minimize the amount of damage such attacks can cause and the time it takes to recover

The aim of the strategy is to achieve these goals by establishing an architecture between the public and private sectors for responding to national-level cyberattacks, developing methods and tools for vulnerability assessment and the strategic and tactical analysis of such attacks, and improve the sharing of information concerning threats, vulnerabilities, and attacks. The strategy was developed by the President's Critical Infrastructure Protection Board.

For More Information

Visit www.whitehouse.gov/pcipb/ for more information.

See Also: *infosec, National Telecommunications and Information Administration (NTIA)*

National Telecommunications and Information Administration (NTIA)

A U.S. government agency that takes a leadership role in a variety of information technology issues including security.

Overview

The National Telecommunications and Information Administration (NTIA) has the job of developing executive policy to ensure greater innovation, competition, and consumer choice in telecommunication products and services. An essential part of this is the protection of the national telecommunications infrastructure, which is foundational for the operation of both government and the private sector. As a result, one responsibility of the NTIA has been to review and refine the National Strategy to Secure Cyberspace, the information security (infosec) component of the U.S. National Strategy for Homeland Security.

For More Information

Visit www.ntia.doc.gov for more information.

See Also: *National Strategy to Secure Cyberspace*

Nbtscan

A tool for automating NetBIOS scans of remote networks.

Overview

The simple Nbtstat command often is used by crackers to try to obtain useful information for compromising NetBIOS-enabled Internet Protocol (IP) networks, but it can be used to scan only one remote host at a time. A tool called Nbtscan simplifies the process by allowing attackers to scan ranges of IP addresses as they look for legacy servers running on the Microsoft Windows NT platform and other NetBIOS-enabled machines. Using this tool, an attacker can quickly gather information about user and domain names on Windows NT-based networks and then may employ password-cracking tools to compromise the security of the target network. By configuring a firewall to block User Datagram Protocol (UDP) ports 137 through 139, however, administrators can easily prevent such scans.

For More Information

Visit www.inetcat.org/software/nbtscan.html for more information.

See Also: *cracking, Nbtstat*

Nbtstat

A command-line tool for displaying NetBIOS over TCP/IP (NetBT) protocol statistics and other NetBIOS information.

Overview

Nbtstat can be used for troubleshooting NetBIOS-related issues on networks that include legacy computers running on the Microsoft Windows NT platform. Crackers can also use Nbtstat to gain useful information that might help them launch a password-cracking attack on a NetBIOS-enabled network. By using the Nbtstat command with -a or -A options, crackers may obtain the following information for targeted systems:

- User name of logged-on user
- Windows NT domain name
- Services running on the machine
- Media access control (MAC) address of network interface

Using the first two pieces of information, the attacker only has to find a third piece (the user's password) to compromise the security of the target system or network. By configuring a firewall to block ports 137 through 139, however, administrators can prevent such remote Nbtstat scans.

See Also: *cracking, Nbtscan*

NCSC

Stands for National Computer Security Center, an initiative of the National Security Agency (NSA) focused on information security (infosec).

See: *National Computer Security Center (NCSC)*

Nessus

An open source security scanning tool.

Overview

Nessus is a freely available tool that can be used to remotely audit a network to look for vulnerabilities that might be exploited by crackers. Nessus is different from many similar tools in that it doesn't rely on services utilizing well-known port numbers but instead scans for all listening ports and performs tests to recognize them. Nessus uses a client/server architecture in which the server part scans the network while the client collects the results of the scan. Nessus can be configured to operate in nondestructive mode to safely check for vulnerabilities, or it can be configured to exploit any vulnerabilities found to verify them (though this can sometimes bring down a server).

Nessus has a fully extensible architecture that uses plug-ins to provide functionality for the server part of the tool. Nessus can be scripted to automate scans and can scan any number of hosts simultaneously, depending on the processor speed and network connection available. The current version is Nessus 2 for UNIX/Linux systems, but a Nessus client for Microsoft Windows systems exists called NessusWX (the server part of Nessus is available only for UNIX/Linux platforms).

For More Information

Visit www.nessus.org for more information.

See Also: *scanning, vulnerability*

Netbus

A notorious Trojan and remote administration tool.

Overview

Netbus is a double-edged tool similar to Back Orifice in that it can be used either maliciously as a Trojan or as a legitimate tool for remote administration. Netbus, developed by Carl-Frederik Neikter, uses a client/server architecture in which the attacker uses the client component to remotely control the server component running on a target host. In fact, any Netbus client can control any Netbus server if it can access it over the Internet, so once your server is compromised it is wide

open to manipulation by any cracker. Netbus can perform virtually any system action, including controlling the mouse, logging keystrokes, capturing display screens, transferring files, and so on.

Notes

There is a program floating around called NetBuster that emulates the Netbus server to let you see who is trying to control your server, but it's generally not a good idea to use such programs as security tools unless you are absolutely sure that they don't themselves perform some malicious action such as install a backdoor on your server.

See Also: *Back Orifice, Trojan*

Netcat

A tool for port scanning and transferring information over network connections.

Overview

Netcat is a tool that can read or write data over network connections using Transmission Control Protocol (TCP), User Datagram Protocol (UDP), or any arbitrary port number. The tool can be used for various purposes, both good and bad, such as the following:

- Performing file transfers across a network
- Covertly transferring data to or from compromised systems
- Testing and debugging TCP or UDP communication problems
- Testing network services as a replacement for Telnet
- Scanning a target network for listening services

Netcat is an extremely flexible tool that can use any local source address and port to initiate a connection. It includes built-in port-scanning capabilities, can easily be scripted, and can operate in slow-send mode for covert channels. It can function both as a client to send data to a specific Internet Protocol (IP) address and port or as a server to listen for incoming connections on a

particular port. In fact, Netcat is so versatile that it is often referred to as the “Swiss army knife” for TCP/IP networks. Netcat is available from @Stake for both Microsoft Windows and UNIX/Linux platforms.

For More Information

Visit www.atstake.com/research/tools/network_utilities/ for more information.

See Also: *scanning*

.NET Passport

A system for managing online identity developed by Microsoft Corporation.

Overview

.NET Passport is an online service from Microsoft that allows users to be authenticated for multiple Web sites and services using a single set of credentials. This single sign-on (SSO) technology makes it easier for users to use e-commerce sites, perform transactions with business partners, and access resources across the Internet. Using a single name and password, a user can access any site that participates in the .NET Passport program without having to reenter credentials. To accomplish this, .NET Passport uses cookies on users' Web browsers to track their sign-on information and identifies users internally using a unique 64-bit number that is encrypted for greater security.

.NET Passport also allows users to store their personal information in a .NET Passport profile and to share this information with participating sites. Users have full control over information in their profile and can share it with sites of their own choosing. A user's .NET Passport profile may contain some or all of the following information, depending on the registering site:

- E-mail address (some sites may require an MSN or Hotmail account)
- First and last name
- State or territory
- Country or region

- ZIP or postal code
- Language preference
- Local time zone
- Gender
- Birth date
- Occupation

Also included in user profiles is .NET Passport wallet, a feature for securely storing users' telephone numbers, credit card numbers, and billing addresses so they won't have to reenter it every time they revisit a participating site.

For More Information

Visit www.passport.net for more information.

See Also: *authentication, Liberty Alliance Project, single sign-on (SSO), TrustBridge*

Netstat

A command-line tool for displaying Transmission Control Protocol/Internet Protocol (TCP/IP) statistics.

Overview

Netstat is a useful tool for troubleshooting TCP/IP network connections and can be used for displaying all current TCP connections and listening ports on a host. Both TCP and User Datagram Protocol (UDP) ports can be displayed in either numerical or label form; for example, **80** or **http** for port 80, the standard Hypertext Transfer Protocol (HTTP) server port. Netstat is also a useful security tool because, by displaying a list of listening ports on a server, it is possible to detect the presence of Trojans and other unauthorized server applications that listen for connections on unusual high-value ports.

Implementation

To display a list of listening ports in numerical form, type **netstat -an** at the command prompt. Typical output might look something like this:

```
C:\>netstat -an
```

Active Connections

| Proto | Local Address | Foreign Address | State |
|-------|-------------------|-------------------|-------------|
| TCP | 0.0.0.0:21 | 0.0.0.0:0 | LISTENING |
| TCP | 0.0.0.0:25 | 0.0.0.0:0 | LISTENING |
| TCP | 0.0.0.0:80 | 0.0.0.0:0 | LISTENING |
| TCP | 0.0.0.0:135 | 0.0.0.0:0 | LISTENING |
| TCP | 0.0.0.0:443 | 0.0.0.0:0 | LISTENING |
| TCP | 0.0.0.0:445 | 0.0.0.0:0 | LISTENING |
| TCP | 0.0.0.0:1025 | 0.0.0.0:0 | LISTENING |
| TCP | 0.0.0.0:1026 | 0.0.0.0:0 | LISTENING |
| TCP | 0.0.0.0:1028 | 0.0.0.0:0 | LISTENING |
| TCP | 0.0.0.0:3372 | 0.0.0.0:0 | LISTENING |
| TCP | 0.0.0.0:3744 | 0.0.0.0:0 | LISTENING |
| TCP | 0.0.0.0:54320 | 0.0.0.0:0 | LISTENING |
| TCP | 172.16.15.220:80 | 172.16.15.33:2187 | ESTABLISHED |
| TCP | 172.16.15.220:139 | 0.0.0.0:0 | LISTENING |
| UDP | 0.0.0.0:135 | *:* | |
| UDP | 0.0.0.0:445 | *:* | |
| UDP | 0.0.0.0:1027 | *:* | |
| UDP | 0.0.0.0:1029 | *:* | |
| UDP | 0.0.0.0:3456 | *:* | |
| UDP | 172.16.11.220:137 | *:* | |
| UDP | 172.16.11.220:138 | *:* | |
| UDP | 172.16.11.220:500 | *:* | |

Here, the presence of listening port number 54320 could indicate Back Orifice 2000 (BO2K) is installed on the server. Netstat isn't foolproof in detecting Trojans because many of them utilize low-value well-known port numbers for their operation, and the presence of such listening ports often would not look suspicious.

Notes

Netstat is available on both Microsoft Windows and UNIX/Linux platforms, and its implementation varies slightly between different platforms.

See Also: Trojan

network address translation (NAT)

A mechanism for translating Internet Protocol (IP) addresses between two networks.

Overview

Network address translation (NAT) is commonly used to securely connect corporate networks to the Internet. In a typical scenario, private (nonroutable) IP addresses

are used for hosts on the internal network, which is connected to the external Internet using a NAT-enabled router. NAT is defined in RFC 1631 and was originally created to address the problem of growing depletion of available public IP addresses, but NAT also enhances the security of business networks by hiding IP addresses of hosts from the outside world, making it more difficult for attackers to penetrate and compromise a network.

NAT is more than just security by obscurity since private IP addresses are not routable, but NAT by itself cannot protect against network intrusion from the Internet. What NAT does protect against is Internet users directly accessing Web and File Transfer Protocol (FTP) servers used internally on the private network. By configuring inbound mappings on your NAT router, however, you can map public IP addresses to expose internal hosts such as FTP servers to allow external users to access them. If you do this, however, you must ensure your exposed hosts are adequately protected by firewalls, locked down, up to date with patches, and monitored for possible intrusion.

Notes

NAT functionality is built into most routers, commercial firewall products, and network operating systems such as Microsoft Windows Server 2003 and various flavors of UNIX. NAT is usually called port address translation (PAT) when referring to Cisco routers and IP masquerading when referring to Linux platforms.

For More Information

For more information about NAT, see the *Microsoft Encyclopedia of Networking, Second Edition*, available from Microsoft Press.

See Also: *firewall*

network-based intrusion detection system (NIDS)

An intrusion detection system (IDS) that monitors activity on a network.

Overview

A network-based intrusion detection system (NIDS) scans traffic on a network, looking for anything that seems suspicious that might indicate an attack. A NIDS generally identifies malicious traffic by comparing traffic patterns to a database of known attack signatures. If a traffic event exceeds the threshold level, the event is logged and other actions can be taken such as alerting an administrator or closing a port to block the traffic. Configuring the threshold level is a delicate matter since setting the threshold too high can result in numerous false positives that waste administrators' time as they investigate spurious attacks, and configuring it too low can allow intruders to penetrate network defenses and compromise network security.

Marketplace

Some of the popular NIDS products on the market include RealSecure from Internet Security Systems, SecureNet from Intrusion.com, NID from NFR, and Cisco Secure IDS from Cisco Systems. The open source tool Snort is a free NIDS tool that is popular with the security community.

Issues

A NIDS works well when deployed either at a choke point on the perimeter of a network where all inbound traffic must enter or on a shared network segment such as a local area network (LAN) with hubs. A NIDS has difficulty with switched networks since switch ports are designed to isolate traffic from each other. This limitation can be overcome in several ways:

- By installing agents (also called taps or monitors) on remote network segments to capture and forward traffic to the NIDS
- By using switches that support port spanning to allow traffic to be copied to a special monitoring port where the NIDS is attached

See Also: *host-based intrusion detection system (HIDS), intrusion detection system (IDS)*

network-based security

The practice of hardening the elements of a network to protect other devices.

Overview

Firewalls, routers, and switches all play an integral part in the network. Most of these devices can be configured to play some part in network security. Routers and switches can remove unwanted traffic from the network before it even makes it to a firewall. Network-based security usually involves securing these devices so that the servers, clients, and other devices are protected by the network itself. Often, an organization will subscribe to this method, investing much time and money into securing network elements while the hosts go ignored. This can become problematic as more attacks use seemingly benign traffic to cause harm. A better solution is "defense in depth," which involves securing each layer to maximize protection.

See Also: *firewall*

network logon

Logging on to a computer using network credentials.

Overview

On stand-alone computers and in simple workgroup scenarios, individual machines manage their own sets of user accounts. Users who log on to a stand-alone machine using local accounts are performing **inter-active logons** since they are “interacting” with the machine through the console. In business environments, however, where security is an important consideration, user accounts are usually stored on a special server called a network authentication server (NAS). In a Microsoft Windows network that uses Active Directory directory service, such NAS servers are called domain controllers. To log on to such a network, users enter their credentials on the console of their local machine, which securely transmits these credentials over the network to the NAS server, which authenticates them to allow access to network resources. This type of authentication process is called a **network logon** since user credentials are sent over the network.

See Also: *authentication, logon*

network mapper

A tool for generating information that can be used to map or diagram the arrangement of hosts on a network.

Overview

Network mappers have several important functions, including inventorying resources on a network and monitoring these resources in case some go down. Attackers often use these tools as well in order to create a map of network hosts to help focus energy for an attack on those displaying vulnerabilities. The tools used by sysadmins and crackers tend to be different, however, with enterprise-level network mappers capable of displaying network information in graphical form using vector-based graphics and friendly icons, while mappers used by crackers tend to be cryptic command-line tools with primitive display capabilities. Both for attacking network and in defense of planned attack, network mappers are useful tools for security professionals and intruders alike.

Network mappers acquire the information they need to map a network in various ways. Such tools may query

router tables, scan predefined address ranges, sweep for listening ports, and use a variety of autodiscovery techniques to generate detailed maps of subnets, hosts, and services running on hosts.

Notes

A popular network mapping tool used by both security professionals and black-hat hackers is Nmap, which also stands for “network mapper.”

See Also: *Nmapscanning*

network monitor

Another name for protocol analyzer, a tool used to view network traffic at the packet level.

See: *protocol analyzer*

network monitoring

Collecting information about traffic patterns and health of a network.

Overview

There are many reasons for monitoring the behavior and operation of computer networks, including capacity planning for future upgrades, early detection of network problems, and network security. From a security perspective, network monitoring can take several forms:

- Reviewing firewall and proxy server logs for signs of intrusion
- Analyzing traffic patterns to detect indications of hosts being compromised
- Scanning listening hosts to identify evidence of backdoors or Trojans
- Reviewing audit logs to thwart attempted password cracking or unauthorized resource access
- Generating alerts when services or hosts become unavailable or their performance profiles change

Tools such as protocol analyzers and intrusion detection systems (IDSs) can provide administrators with detailed information about network traffic both in real time and statistically on the average, and such tools are

an essential part of every security professional's toolkit. Knowing how to use these tools generally requires a good understanding of network protocols, ports, and addressing schemes.

See Also: *backdoor, firewall, intrusion detection system (IDS), protocol analyzer, Trojan*

Network Security Hotfix Checker

Another name for HFNetChk, a Microsoft tool for keeping security patches up to date on a system.

See: *HFNetChk*

Newtear

A denial of service (DoS) exploit against machines running on the Microsoft Windows 95 and Windows NT 4 platforms.

Overview

Newtear exploits vulnerabilities in the Transmission Control Protocol/Internet Protocol (TCP/IP) stacks of systems running on Windows 95 and Windows NT 4 platforms that have not been patched. Newtear is a type of IP fragmentation attack that works by sending a pair of malformed fragments to the Domain Name System (DNS) port on the target host. When the host receives the fragments, it reassembles them into an invalid User Datagram Protocol (UDP) packet that hangs or crashes the machine, resulting in a DoS condition for legitimate users.

Newtear is a variant of an earlier exploit called the Teardrop attack. An exploit similar to Newtear is the boink attack, which also targets the DNS port using malformed UDP fragments.

See Also: *boink attack, denial of service (DoS), IP fragmentation attack, Teardrop attack*

Next-Generation Secure Computing Base for Windows

A set of features for upcoming versions of Microsoft Windows platforms that provides enhanced data security, personal privacy, and system integrity.

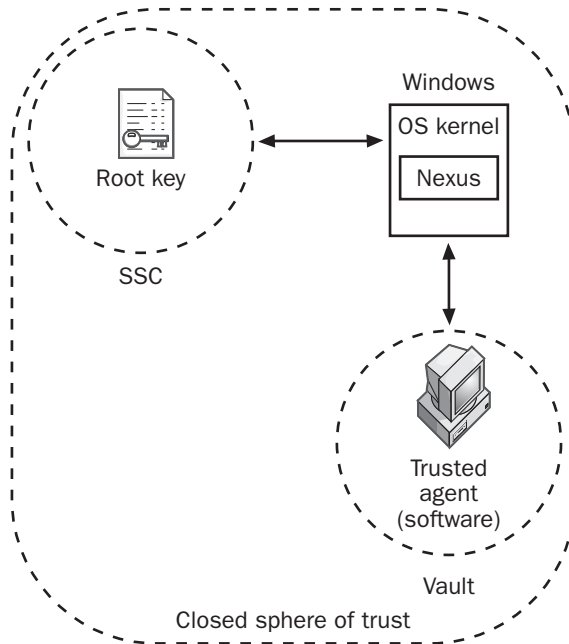
Overview

Formerly called Palladium, the Next-Generation Secure Computing Base for Windows is an initiative from Microsoft designed to provide a framework for developing trusted software that protects the privacy and security of user data. The Next-Generation Secure Computing Base will consist of architectural enhancements to both hardware (processors, chipsets, and peripherals) and the Microsoft Windows kernel that will provide a trusted execution subsystem within which security-enhanced applications can run. The Next-Generation Secure Computing Base is an entirely opt-in solution, and supporting systems will be shipped with these features turned off since the goal is to provide users with ultimate control of their systems, applications, and data.

Implementation

The Next-Generation Secure Computing Base combines public key encryption technologies, advanced hardware, and enhancements to the Microsoft Windows operating system to create closed spheres of trust that bind data and services to users and trusted applications. The Security Support Component (SSC) residing on system hardware provides a master root key that forms the basis of cryptographic storage and communications within the system. Security-enhanced applications called Trusted Agents interact with the SSC through the Nexus, an enhanced feature of the Windows kernel that manages trust between applications, the operating system, and hardware. Data and services are bound to users and applications within a closed sphere of trust that allows the Nexus to securely store data within a private storage area called a Vault.

Trusted code runs within physically isolated memory called Trusted Space that is inaccessible to the rest of the system, which helps protect systems and data against unauthorized programs such as Trojans and worms. An authentication mechanism called Sealed Storage is used to store secrets that cannot be read by untrusted programs, even if a duplicate operating system is installed or a hard drive is transferred to a different machine. Attestation is used to allow users to reveal selected characteristics of their system to external requestors.



Next-Generation Secure Computing Base for Windows.

How the Next-Generation Secure Computing Base for Windows works.

For More Information

Visit www.microsoft.com/presspass/features/2002/jul02/0724palladiumwp.asp for more information.

See Also: public key cryptography, Trojan, worm

NFIC

Stands for National Fraud Information Center, an organization helping consumers and law enforcement agencies fight fraud.

See: National Fraud Information Center (NFIC)

Ngrep

A tool for “grepping” (searching for) specific information in network packets.

Overview

Grep is a familiar utility on UNIX platforms that can use regular expressions to search files for lines that

match a specific pattern. Ngrep is a tool designed to work similarly except that it searches network traffic instead of files. Ngrep uses hexadecimal regular expressions to search the data payloads of packets for matching information, and it works with Transmission Control Protocol (TCP), User Datagram Protocol (UDP), and Internet Control Message Protocol (ICMP) packets on a variety of network interfaces, including the following:

- Local area network (LAN) interfaces such as Ethernet, Token Ring, and Fiber Distributed Data Interface (FDDI)
- Wide area network (WAN) interfaces such as Point-to-Point Protocol (PPP) and Serial Line Internet Protocol (SLIP)
- Null interfaces

Ngrep is available for both UNIX/Linux and Microsoft Windows platforms. Like most security tools, Ngrep can be used for good or ill. For example, network administrators can use it to troubleshoot various kinds of TCP/IP communication problems, while crackers could use it to sniff network traffic for passwords and other sensitive information (though other more powerful sniffers such as Dsniff are often preferred).

Notes

For some hacker humor, check out the song “Grepping in a UNIX Wonderland” at www.speelman.net/humor/xmas/winter7.html.

For More Information

Visit www.packetfactory.net for more information.

See Also: cracking, Dsniff, sniffer

NIACAP

Stands for National Information Assurance Certification and Accreditation Process, a standardized process for information assurance (IA) accreditation.

See: National Information Assurance Certification and Accreditation Process (NIACAP))

NIAP

Stands for National Information Assurance Partnership, a cooperative agency for promoting information security (infosec) in U.S. government agencies and private industry.

See: National Information Assurance Partnership (NIAP)

NIDS

Stands for network-based intrusion system, an intrusion detection system (IDS) that monitors activity on a network.

See: network-based intrusion detection system (NIDS)

NIETP

Stands for National INFOSEC Education & Training Program, an information security training program from the National Security Agency (NSA).

See: National INFOSEC Education & Training Program (NIETP)

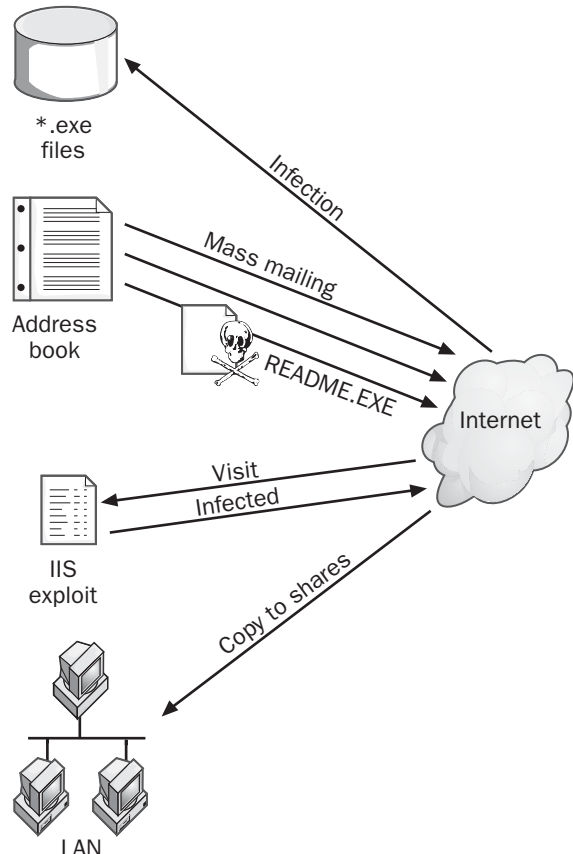
N

Nimda

A dangerous e-mail worm that appeared in September 2001.

Overview

Nimda (which is **Admin** spelled backward) is a severe mass-mailing worm that affects most versions of Microsoft Windows operating systems. The worm appeared a week after the World Trade Center terrorist attacks and rapidly spread across the Internet, causing massive damage. Nimda combines the features of Code Red and SirCam together with additional advanced features that allow it to infect Web servers running Internet Information Services (IIS). The worm is capable of giving attackers full access to infected systems, but by applying patches issued by Microsoft, administrators can prevent infection.



Nimda. The life cycle of the Nimda worm.

Implementation

Nimda's life cycle consists of four stages:

- **File infection:** Nimda infects systems by assimilating executable (*.exe) files on the target system.
- **Mass mailing:** The worm exploits a vulnerability in the implementation of the Messaging Application Programming Interface (MAPI) to locate e-mail addresses of recipients stored in the Windows address book. It then sends a message to each recipient that includes an attachment called Readme.exe. On some systems, this attachment could be opened automatically, causing infection. Nimda also can scan locally

stored Hypertext Markup Language (HTML) files for additional e-mail addresses to which to send itself.

- **IIS exploit:** The worm scans for IIS Web servers on the Internet and tries to exploit known vulnerabilities to compromise them. Compromised servers have their Web pages modified so that users visiting the pages automatically become infected with Nimda.
- **Local area network (LAN) propagation:** Nimda also scans the local network for file shares and tries to place a hidden file Riched20.dll into directories containing *.doc and *.eml files. When users try to open Microsoft Word documents or e-mails in these directories, Riched20.dll executes and infects their systems.

See Also: CodeRed, SirCam, worm

NIPC

Stands for National Infrastructure Protection Center, a cooperative agency set up to help protect critical information system infrastructures.

See: National Infrastructure Protection Center (NIPC)

NIST

Stands for National Institute of Standards and Technology, a U.S. government agency that develops standards for government and private industry sectors.

See: National Institute of Standards and Technology (NIST)

Nmap

An open source tool for network security auditing.

Overview

Nmap (which stands for **network mapper**) is a free tool that can be used to scan networks to obtain a large amount of information about their configuration and vulnerability. Like most security tools, Nmap can be used for good or ill purposes, and its use can range from auditing the security of enterprise networks to finding vulnerabilities to exploit for launching an attack. By

using Nmap, administrators or crackers can determine the following bits of information:

- The hosts available on a network
- The services running on each host (listening ports)
- The make and version of the operating system running on each host

Nmap can enumerate remote hosts by sending invalid Transmission Control Protocol (TCP) packets and comparing the results to a database of operating system signatures. Since each operating system platform and version tends to have its own unique implementation of Transmission Control Protocol/Internet Protocol (TCP/IP), this type of scan can usually provide attackers with a great deal of information about the targeted hosts. This enables the attacker to target efforts against known vulnerabilities to more easily compromise the target network.

Nmap is available for several UNIX platforms, including Linux, and a version for Microsoft Windows is also available.

For More Information

Visit www.insecure.org/nmap/ for more information.

See Also: enumeration, fingerprinting, scanning

nonce

A number that is used only once in an algorithm.

Overview

Nonces are employed in many encryption schemes and authentication protocols for inserting an unpredictable factor into their algorithms. Nonces may be created using several methods, including pseudorandom number (PRN) generators, time stamps by system clocks, or sequence numbers. In general, large random numbers make the best nonces because the chance of reusing them later is negligible. A typical method for generating such numbers is to use a complicated function of a time stamp and then encrypt it using a key known only to the party creating the nonce.

The purpose of a nonce is to prevent a replay attack from occurring. If an attacker eavesdrops on an authentication session that employs a nonce, the captured session cannot be replayed later since the nonce ensures that future sessions will be different.

See Also: authentication, encryption algorithm, pseudorandom number generator (PRNG)

nonrepudiation

The ability to prove identity to a third party.

Overview

Nonrepudiation is the ability to prove who performed an action such as sending a message, deleting a file, or rebooting a system. Systems that implement nonrepudiation can provide a legal basis for establishing evidence that can hold up in a court of law. In a transaction on such a system, both parties can legally prove the involvement of the other party should either of them choose to deny their involvement. Nonrepudiation thus involves both proof of sending the transaction and proof of its receipt, concepts which are known respectively as **nonrepudiation of origin** and **nonrepudiation of delivery**. The weaker concepts of proof of origin, proof of delivery, and proof of submission involve proving to the other party in the transaction who performed an action. Nonrepudiation is a stronger concept because it involves providing proof to a disinterested third party.

On computer systems, nonrepudiation is typically provided by audit logs that register significant actions such as deleting files by recording the identity (ID) of the user who deleted them and the day and time they were deleted. In an e-mail system, nonrepudiation is typically established by signing messages with digital signatures using keys and certificates issued by trusted third parties. Messaging systems can support a third variety of nonrepudiation called **nonrepudiation of submission**, which prevents message transfer agents (MTAs) or mail servers from denying that a message was submitted to them for delivery.

Notes

Nonrepudiation is one of the five core attributes of information assurance (IA), a set of methodologies for ensuring the security of information systems; the other

attributes are authentication, availability, confidentiality, and integrity.

See Also: auditing, digital signature, identity theft, information assurance (IA)

notice

Disclosure information concerning privacy practices.

Overview

Notice is a privacy principle that ensures reasonable disclosure is made of what a company does with identity information collected from consumers. Such identity information is typically called personally identifiable information (PII) and may consist of a person's name, address, e-mail address, credit card number, Social Security number, other kinds of ID numbers, Internet Protocol (IP) address, or any other unique identifier that is associated with the person or his or her computer system.

Notice is generally provided to consumers by a statement of privacy policy, which should be visibly placed on the public Web site for the business. Notice is a requirement of Fair Information Practices (FIP), a set of standards governing the collection and use of personal data that derive from the legislation such as the Privacy Act of 1974.

See Also: Fair Information Practices (FIP), identity theft, personally identifiable information (PII), privacy, privacy policy

Npasswd

A replacement for the standard Passwd utility on UNIX platforms.

Overview

Npasswd can be used on UNIX systems as a replacement for the weaker Passwd utility to enhance security. It accomplishes this by the following methods:

- Maintaining a password history for each user to prevent passwords from being reused too frequently
- Verifying the minimum size and complexity of passwords to prevent them from being easily guessed

- Preventing users from choosing easily guessed passwords such as telephone or Social Security numbers or passwords derived from the user's name, address, or other stored personal information
- Testing passwords against cracking dictionaries to ensure they cannot easily be cracked

Npasswd is available for most UNIX platforms and for Linux.

For More Information

Visit www.utexas.edu/cc/unix/software/npasswd/dist/ for more information.

See Also: *cracking, /etc/passwd, password*

NSA

Stands for National Security Agency, a U.S. government agency responsible for protecting national information systems and producing foreign intelligence information.

See: *National Security Agency (NSA)*

Nslookup

A command-line utility for querying Domain Name System (DNS) name servers.

Overview

Nslookup is a standard utility on most operating systems, including Microsoft Windows versions and different flavors of UNIX/Linux, and is generally used as a tool for troubleshooting resource record issues on DNS name servers. It can also be used by crackers, however, as a tool for footprinting, a method used by attackers to identify potential targets for attacking a network by gathering as much information as possible about the network from publicly available sources. Nslookup can also be used together with other tools for certain exploits against name servers that can compromise servers, hijack domain names, and redirect name query traffic.

Notes

An updated tool called Dig is available for UNIX and Microsoft Windows environments and has enhanced features compared to Nslookup.

See Also: *DNS spoofing, footprinting*

NTBugtraq

A mailing list for Microsoft Windows security issues.

Overview

NTBugtraq began as a vehicle for openly discussing bugs and vulnerabilities on the Windows NT platform, but later expanded to include Windows 2000, Windows XP, and related applications. The list is managed by Russ Cooper, who acts as editor and also as surgeon general for TruSecure Corporation. The list is modeled after an earlier security list called Bugtraq that was run by hacker Aleph One and that is now hosted on Security Focus (www.securityfocus.com).

For More Information

Visit www.ntbugtraq.com for more information.

See Also: *vulnerability*

NTFS

An enhanced file system used on Microsoft Windows NT and later versions of the operating system.

Overview

NTFS is the preferred file system for implementing secure data storage on Windows NT, Windows 2000, Windows XP, and Windows Server 2003. In addition to its features for data recoverability and fault tolerance, NTFS also includes advanced security features that enable users to control access to resources stored on disk systems. Files and directories are implemented in NTFS as securable objects, and access to files and directories can be restricted to specific users and groups using NTFS permissions, which include both standard and advanced permissions.

For More Information

For more information about NTFS and NTFS permissions, see the *Microsoft Encyclopedia of Networking, Second Edition*, available from Microsoft Press.

See Also: *access control, access control list (ACL), permissions*

NTIA

Stands for National Telecommunications and Information Administration, a U.S. government agency that takes a leadership role in a variety of information technology issues including security.

See: *National Telecommunications and Information Administration (NTIA)*

NTLM

The authentication protocol used by Microsoft Windows NT.

Overview

NTLM is a challenge-response authentication protocol based on the earlier LAN Manager (LM) authentication protocol originally developed by IBM and used by Microsoft as the authentication method for Windows 3.1, Windows for Workgroups 3.11, and Windows 95. NTLM is the default authentication protocol in Windows NT and is supported by Windows 2000, Windows XP, and Windows Server 2003 for backward-compatibility reasons. The Kerberos protocol replaces NTLM as the default authentication protocol for Windows 2000 and Windows Server 2003.

Implementation

NTLM can be used for both local (interactive) and network authentication including pass-through authentication. NTLM credentials consist of a domain name, user name, and a one-way hash of the user's password. NTLM employs an encrypted challenge-response method to authenticate users without requiring that their passwords be transmitted over the connection. For network authentication, NTLM performs the following

steps when a user on a client machine tries to access resources on a server:

- 1- User credentials are entered on the client machine.
- 2- The client calculates a hash of the user's password and discards the actual password.
- 3- The client sends the user's name to the server in cleartext.
- 4- The server generates a 16-byte random challenge string, or nonce, and sends it to the client.
- 5- The client encrypts the challenge string using the hash of the user's password and sends this response to the server.
- 6- The server sends the user's name, challenge string, and client response to a domain controller.
- 7- The domain controller retrieves the hash of the user's password from its Security Account Manager (SAM) database.
- 8- The domain controller encrypts the challenge string using the retrieved password hash.
- 9- The domain controller compares the encrypted challenge it calculated with the response string from the client.
- 10- If these are identical, the domain controller notifies the server that the client has been authenticated.

Notes

NTLM is also known as Windows NT Challenge/Response authentication. NTLM is not natively supported by Windows 95, Windows 98, or Windows Millennium Edition (Windows Me), though by installing a downloadable directory client NTLM is supported on these platforms. An enhanced version of NTLM called NTLMv2 has improved security and is supported by Windows NT 4 Service Pack 4 and later.

See Also: *authentication, hashing algorithm, Kerberos, nonce*

Ntrights

A tool for assigning rights to Microsoft Windows NT users or groups.

Overview

Ntrights was included in the *Microsoft Windows NT Server 4 Resource Kit* as a tool for assigning specific rights to users or groups from the command line. Crackers often use the tool as well for elevation of privileges (EoP) once a Windows NT-based system or network has been compromised. The tool can operate remotely to manage rights for users across a network connection and can be used to both grant and revoke user rights.

See Also: *elevation of privileges (EoP), rights*

null session attack

An exploit that uses unauthenticated NetBIOS connections to enumerate a target host.

Overview

Null sessions are unauthenticated NetBIOS sessions that are established with no user name or password. Null sessions were included in the Microsoft Windows

NT operating system by design to allow the enumeration computers, shares, and users on the network, but a vulnerability was later discovered that null sessions could be exploited using port 139 to allow access to the registry using the credentials of the Everyone built-in identity. This was fixed with Windows NT 4.0 Service Pack 3, but some security professionals feel that null sessions still constitute a security vulnerability in Windows platforms since they allow attackers to obtain useful target information using such NetBIOS enumeration tools as Dumpsec, Enum, Hunt, NBTenum, and Wininfo. Generally, administrators deal with this issue using one or more of the following methods:

- By blocking port 139 (and port 445 for Windows 2000 or later versions) on the firewall
- By using an intrusion detection system (IDS) that includes signatures for recognizing null sessions
- By disabling NetBIOS completely on machines running on Windows platforms

See Also: *enumeration*

OAKLEY

A key determination protocol used for encrypted communications.

Overview

OAKLEY defines a protocol by which two parties in an authenticated communication session can agree with each other regarding a shared secret key. OAKLEY is based on the Diffie-Hellman (DH) key exchange algorithm, a mechanism that allows two parties to agree on a shared value without the need of encryption. OAKLEY also supports Perfect Forward Secrecy, a condition that makes it impossible for an eavesdropper to decrypt a conversation even if the entire encrypted session can be captured.

OAKLEY is an Internet standard protocol that is defined in RFC 2412. OAKLEY is typically used together with Internet Security Association and Key Management Protocol (ISAKMP), a protocol for managing security associations, forming a combination called ISAKMP/Oakley that is now commonly called Internet Key Exchange (IKE).

See Also: *Diffie-Hellman (DH), Internet Key Exchange (IKE), perfect forward secrecy (PFS), secret key*

obscurity

A way of trying to enhance the security of a system by hiding aspects of its internal operation.

Overview

The principle of “security through obscurity” involves modifying aspects of the way a system works to hide the presence of resources or services. As an example, the standard Transmission Control Protocol (TCP) port on which Web servers listen for client connections is port 80, a well-known port number that every cracker knows. A simple port scan of a system that reveals it is

listening on port 80 is a clear indication that at least one of the roles of the system is as a Web server. By changing the default Hypertext Transfer Protocol (HTTP) port to some unexpected value like 13625, an administrator can use security through obscurity to hide the fact that the system is a Web server (clients, of course, must be informed of the port change so they will still be able to connect). In practice, however, this measure only minimally enhances the security of the system since the tools employed by crackers can easily discover the new listening port 13625 and perform simple tests to determine that it is listening for HTTP connections.

Most administrators therefore take the stance that security by obscurity is largely ineffective, creates unnecessary complexity by requiring client reconfiguration, and can actually have the negative effect of giving the impression of greater security when in fact this is not so. Taking this to its logical conclusion, some security professionals argue that only open systems whose architecture is fully available to the public can be completely secure and that proprietary systems in which vendors hide the implementation of operational features may actually be less secure in the long run. In practice this argument is mitigated by the increasing complexity of modern software, which makes it difficult to eliminate all vulnerabilities even from completely open systems and by the amazing ingenuity of hackers who, for whatever motive, try to break or compromise such software.

See Also: *open system*

OCSP

Stands for Online Certificate Status Protocol, a protocol for verifying whether digital certificates are valid or have expired.

See: *Online Certificate Status Protocol (OCSP)*

OCTAVE

A methodology for evaluating the security risks associated with information systems.

Overview

OCTAVE was developed by the CERT Coordination Center (CERT/CC), a center of Internet security expertise operated by Carnegie Mellon University. **OCTAVE** stands for Operationally Critical Threat, Asset, and Vulnerability Evaluation and is a methodology for self-directed security risk evaluation that organizations can perform on their information systems. OCTAVE measures an organization's information systems, policies, and procedures against industry-accepted best practices and helps an organization develop strategies for protecting information assets from common threats and vulnerabilities. OCTAVE is primarily designed for large organizations of several hundred employees or more, but a modified methodology called OCTAVE-S is being developed that can be used easily by small organizations.

For More Information

Visit www.cert.org/octave/ for more information.

See Also: *CERT Coordination Center (CERT/CC), information assurance (IA), threat, vulnerability*

OFB

Stands for output feedback mode, a type of stream cipher employing a one-time pad.

See: *output feedback mode (OFB)*

one-time pad (OTP)

A simple yet unbreakable symmetric cipher.

Overview

One-time pad (OTP) encryption is an uncrackable cipher that uses a randomly generated pad (series of bits) the same length as the message to be encrypted. The cipher encrypts the entire plaintext message simply by XORing it with the pad to create ciphertext. Since the cipher is symmetric, the recipient simply XORs the received ciphertext using the same pad to recover the original plaintext message. The scheme is completely uncrackable even using brute force to try all possible

pads since there is no way of knowing which of all possible plaintexts is the original message. The scheme relies on the fact that each message requires a unique, randomly generated pad and that no portion of this pad is ever reused for encrypting future messages.

Issues

The disadvantage of this scheme, of course, is that some other method is required to distribute the pad to both parties in advance, such as a public key encryption (PKE) scheme like RSA or Diffie-Hellman (DH). Another weakness is that if a pseudorandom number generator (PRNG) is used to create a sufficiently large pad, it may exhibit nonrandom characteristics that may enable an eavesdropper to crack the cipher. By generating true random numbers, for example, using the decay of a radioactive sample, this limitation can be overcome, but the result is added technical complexity.

See Also: *cipher*

one-time password (OTP)

An authentication scheme that requires a new password each time authentication is performed.

Overview

One-time passwords (OTPs) are a way of combating eavesdropping on open network connections. Since a new password is used each time the user authenticates with the network, it is impossible for an attacker to mount a replay attack to capture and replay authentication traffic in order to hijack a session.

Internet standard schemes for using OTPs are outlined in RFCs 1938 and 2289. These schemes are based on the S/KEY technology developed by Bellcore and defined in RFC 1760, and they generate OTPs using the message digest 4 (MD4), message digest 5 (MD5), or Secure Hash Algorithm-1 (SHA-1).

See Also: *eavesdropping, hijacking, message digest 4 (MD4), message digest 5 (MD5), password, replay attack, Secure Hash Algorithm-1 (SHA-1)*

one-way authentication

Authentication of only one end of a communication session.

Overview

Traditional authentication schemes have been one-way schemes in which one end (the client) is authenticated by the other end (the server) before a session can be established. One-way authentication schemes can be based on either passwords or a shared secret key. While such a scheme satisfies the server regarding the identity of the client, it leaves open the possibility of an attacker impersonating the server, which can result in a session that can leak information and compromise the security of the client.

To avoid such problems, mutual authentication (authentication of both ends of a communication session) can be used instead. An example of an authentication protocol that supports mutual authentication is Kerberos, an authentication protocol developed by the Massachusetts Institute of Technology.

See Also: authentication, Kerberos, mutual authentication

one-way encryption algorithm

Another name for a hashing algorithm, a mathematical procedure that generates a fixed-size result from arbitrary amounts of data.

See: hashing algorithm

one-way function

A mathematical function whose results are not easily reversed.

Overview

One-way functions are mathematical functions for which the inverse is extremely difficult (or preferably impossible) to compute. One-way functions are used in cryptography as the basis of hashing algorithms, mathematical procedures that generate fixed-size results from arbitrary amounts of data such that no two input values generate the same output (collisionless function). One-way functions are also used to construct pseudorandom number generators (PRNGs) used to create nonces, one-time pads (OTPs), and other encryption components. Public key encryption (PKE) systems also rely on one-way functions for their operation.

One-way functions can employ a variety of mathematical techniques, including modular arithmetic, logarithms, permutations, and iterative calculation. Although the various one-way functions used in cryptography are believed to be irreversible, none have been rigorously proved so, and it is possible that someday advanced mathematics may find a way to reverse some of them, which could lead to the immediate obsolescence of certain cryptosystems.

Notes

A **trapdoor function** is a one-way function that is reversible if a user knows a secret associated with the function.

See Also: hashing algorithm, nonce, one-time pad (OTP)

Onion Routing

An experimental system to prevent eavesdropping on the Internet.

Overview

Onion Routing was a research project conducted by the U.S. Navy Research Lab. Its purpose was to develop technologies for ensuring the privacy of communications sent over public networks such as the Internet by preventing eavesdropping and traffic analysis attacks. The Onion Routing project ran a prototype network using Hypertext Transfer Protocol (HTTP) traffic on a Sun Solaris network for several years to test proof of concept, and the initial stage of the project concluded in January 2000 with a second-generation system pending.

Implementation

Just as onions consist of layers, Onion Routing adds another layer to traditional Internet Protocol (IP) traffic to support both private and anonymous communications. Socket connections are moved beneath the application layer and are modified to be independent of the application being used. Normal IP applications communicate with anonymous sockets using proxies that anonymize the data stream by removing all identifying information. The proxies then establish anonymous connections from the source host through one or more onion routers to the destination host. To ensure both privacy and anonymity, each onion router along the

communication path adds a layer of encryption to the data being sent, which is decrypted by the next onion router along the path. As a result, the data appears different at each router along a path, which prevents traffic analysis from being used to track the origin and nature of the data.

For More Information

Visit www.onion-router.net for more information.

See Also: [eavesdropping](#), [privacy](#)

Online Certificate Status Protocol (OCSP)

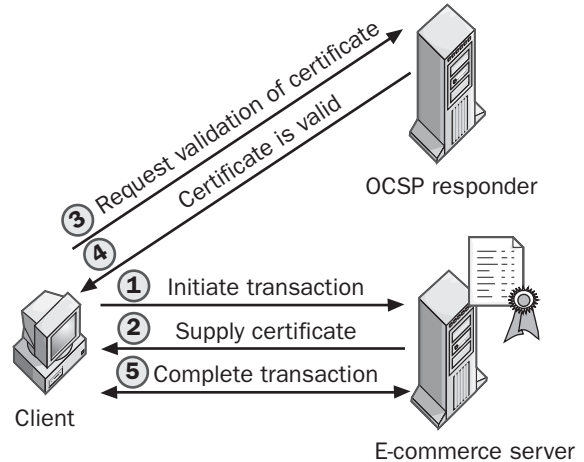
A protocol for verifying whether digital certificates are valid or have expired.

Overview

Online Certificate Status Protocol (OCSP) is a security protocol defined in RFC 2560 and used in Public Key Infrastructure (PKI) systems. OCSP can be used as either a replacement for or complementary to traditional certificate revocation lists (CRLs) to identify expired or revoked digital certificates. The advantage of OCSP over the CRL approach is that CRLs are difficult to manage and can become out of date if they are not updated frequently. By contrast, OCSP provides a real-time managed solution for providing information about the revocation status of digital certificates upon request.

Implementation

Digital certificates are typically used to verify the identity of e-commerce sites and other services. When a client initiates a transaction with a site that has a certificate, the site responds by sending its certificate to the client for validation. Using OCSP, the client can then request verification of the site's certificate by forwarding it to an OCSP responder. The responder replies with either an acknowledgment of the validity of the certificate or an error message indicating an expired or revoked certificate. The client can then decide whether to continue or abort the transaction based on the response received from the responder.



Online Certificate Status Protocol (OCSP). How OCSP works.

Marketplace

Commercial OCSP products such as Servant OCSP from SmartTrust are available in the marketplace. There are also open source and public domain implementations of OCSP available such as the `Ocsp` command that is part of OpenSSL.

See Also: [certificate revocation list \(CRL\)](#), [digital certificate](#), [Public Key Infrastructure \(PKI\)](#)

Online Personal Privacy Protection Act

Proposed U.S. legislation regulating the privacy of information collected from individuals on the Internet.

Overview

The Online Personal Privacy Protection Act proposes regulations for the collection, use, and disclosure of personally identifiable information (PII) collected by Internet service providers (ISPs), commercial Web site operators, and other online services. The act specifies the following notice and consent requirements:

- Clear and conspicuous notice must be given to individuals concerning what information will be

collected, how that information will be used, and what disclosure practices govern the communication of collected information to third parties.

- Individuals must be allowed to opt in or out regarding the collection of sensitive PII and must be given robust notice regarding opt-out options for non-sensitive PII.
- Consent granted or denied shall remain in effect until the associated individual chooses to change this, and such consent rules also apply to any successor entities that have taken over from the original service provider.

The act also outlines practices for notification of changes in privacy policies, lists exceptions for disclosure to law enforcement and national security agencies, ensures individuals have access to collected PII in order to make changes or corrections, and requires service providers to maintain reasonable procedures for protecting the confidentiality, integrity, and security of PII they collect.

See Also: *personally identifiable information (PII), privacy*

onward transfer

Transfer of personally identifiable information (PII) to another recipient.

Overview

Onward transfer is defined as the transfer of PII by the recipient of the original data to a second recipient. For example, the transfer of PII from a recipient in Canada to a recipient in the United States constitutes onward transfer of that data. Onward transfer is covered by Fair Information Practices (FIP), a set of standards governing collection and use of personal data.

See Also: *Fair Information Practices (FIP), personally identifiable information (PII)*

OpenHack

A series of online security challenges organized by *eWeek* magazine.

Overview

OpenHack was first organized in 1999 by *PC Week* magazine (now known as *eWeek*) as a test network set up on the Internet as a challenge for crackers to compromise. As recorded by firewall and intrusion detection logs, the network was quickly subjected to a variety of attacks, including port scans, spoofing attacks, and denial of service (DoS) attacks. Some of these attacks were opportunistic, while others were coordinated and involved diversionary tactics to mask nefarious activity. The network was soon hacked and this illustrated the importance of locking down critical servers by disabling unnecessary services and features to keep configuration simple and the attack surface small. Some attacks demonstrated sophisticated programming skills, and the challenge also highlighted the importance of keeping systems up to date with security patches from vendors.

OpenHack 2 in 2000 demonstrated further lessons concerning how to secure networks against attack. These lessons were well utilized in OpenHack 3 in 2001, which survived all challenges and remained uncompromised at the end of its 17-day existence, demonstrating how much the practice of network security had advanced since the inception of the challenge. OpenHack 4 in 2002, however, was cracked in only a few hours, indicating the need for continued vigilance and avoidance of complacency. The exploit that cracked the network involved a cross-site scripting vulnerability in an Oracle application, but because defense in depth had been implemented, the core services of the network remained secure.

For More Information

Visit www.eweek.com/openhack for more information.

See Also: *attack, hacking*

open mail relay

A mail server that supports mail relaying, a method used by spammers for sending junk mail.

See: *mail relaying*

OpenPGP

Open source implementation of the Pretty Good Privacy (PGP) encryption scheme.

Overview

OpenPGP is an Internet standard defined in RFCs 2440 and 3156. Like its antecedent PGP, OpenPGP leverages encryption technologies both for ensuring the privacy of electronic communication and for securely storing information on disks and other storage systems.

Implementation

OpenPGP employs symmetric encryption using Data Encryption Standard (DES) and Triple DES (3DES) for encrypting data. OpenPGP also includes support for digital signatures using the El Gamal algorithm and Digital Signature Standard (DSS). OpenPGP employs the same binary schemes as PGP for its message and certificate formats, uses the Secure Hash Algorithm-1 (SHA-1) for message hashing, and supports Multipurpose Internet Message Extensions (MIME) for encapsulation of both encrypted and signed data.

For More Information

Visit www.openpgp.org for more information.

See Also: 3DES, Data Encryption Standard (DES), Digital Signature Standard (DSS), Pretty Good Privacy (PGP), Secure Hash Algorithm-1 (SHA-1)

OpenSSH

A free version of the Secure Shell (SSH).

Overview

SSH is a set of protocols and tools that provides more secure replacements to Telnet, File Transfer Protocol (FTP), and other UNIX utilities. OpenSSH is a free version of SSH developed mainly by the OpenBSD Project. It is available for a variety of UNIX/Linux platforms and also for Mac OS X. To keep OpenSSH “open,” the project had to remove some features and support for patented encryption technologies such as International Data Encryption Algorithm (IDEA). OpenSSH nevertheless supports a wide variety of encryption schemes, including Triple DES (3DES), Advanced Encryption Standard (AES), Blowfish, CAST128, and Rivest-Shamir-Adleman (RSA).

Notes

There is a vulnerability in versions 2.2.9 through 3.3 of OpenSSH that could allow an attacker to execute arbitrary code using root privileges, but this has been fixed in later versions of the product.

For More Information

Visit www.openssh.org for more information.

See Also: Secure Shell (SSH)

OpenSSL

A free version of Secure Sockets Layer (SSL).

Overview

SSL is a protocol for establishing a secure communications channel and is widely used on the Internet for encrypting e-business and e-commerce traffic. OpenSSL is an open source toolkit primarily used for implementing SSL on the Apache Web server platform. OpenSSL also includes a library of cryptographic functions and supports the Internet standard Transport Layer Security (TLS) protocol defined in RFC 2246.

Notes

There is a known vulnerability in OpenSSL up to version 0.9.7beta2 that was exploited by the Slapper worm and could allow an attacker to execute arbitrary code using root privileges, but this has been fixed in later versions of the product.

For More Information

Visit www.openssl.org for more information.

See Also: Secure Sockets Layer (SSL), Transport Layer Security (TLS)

open system

A system whose specifications are fully available to anyone who wants to see them.

Overview

Open systems are hardware or software whose architecture or code is available to the public. They are generally developed and maintained by a group or community using a process that implements changes based on consensus in order to detect and correct flaws, enhance features, or modify functionality. The security

of open systems cannot be enhanced by “security through obscurity” because there is nothing obscure or hidden about an open system.

Open systems are the opposite of proprietary systems, which are developed in secrecy by vendors who guard the implementation details from public view. Another name for proprietary systems is commercial off-the-shelf (COTS) products.

Notes

The open source software movement is based on the open system approach combined with a licensing scheme called the General Public License (GPL) to govern how such systems are developed.

See Also: *obscurity*

Orange Book

Formally known as the Trusted Computer System Evaluation Criteria (TCSEC), a set of security classifications for computer systems developed by the U.S. Department of Defense.

See: *Trusted Computer System Evaluation Criteria (TCSEC)*

opt in

To explicitly consent to participate.

Overview

Opt in provides individuals with an element of choice in how their personally identifiable information (PII) is used by e-commerce sites, marketing programs, and other offerings. An example of opt in would be granting consent for the use of collected PII beyond the express purpose for which it was originally collected. Opt in is thus an essential aspect of privacy and is covered by Fair Information Practices (FIP), a set of standards governing collection and use of personal data.

See Also: *Fair Information Practices (FIP), opt out, personally identifiable information (PII), privacy*

opt out

To explicitly decline to participate.

Overview

Opt out provides individuals with an element of choice in how their personally identifiable information (PII) is used by e-commerce sites, marketing programs, and other offerings. An example of opt out would be denying consent for the use of collected PII beyond the express purpose for which it was originally collected. Opt out is thus an essential aspect of privacy and is covered by Fair Information Practices (FIP), a set of standards governing collection and use of personal data.

Notes

Many privacy advocates say that opt out is a poorer choice for sites to offer than opt in, arguing that people should not have to explicitly decline to participate in something, but rather should actively ask to participate.

See Also: *Fair Information Practices (FIP), opt in, personally identifiable information (PII), privacy*

OTP

1. Stands for one-time pad, a simple yet unbreakable symmetric cipher. **2.** Stands for one-time password, an authentication scheme that requires a new password each time authentication is performed.

See: *one-time pad (OTP), one-time password (OTP)*

Outlook E-mail Security Update

An update for Microsoft Outlook that helps protect against malicious e-mail messages.

Overview

In response to vulnerabilities discovered in how Outlook handles messages with attachments, Microsoft Corporation released a series of security updates that help protect users' systems against common threats prevalent on the Internet. The Outlook E-mail Security Update works by classifying messages with attachments into three categories based on the file extensions of the attachments and then taking appropriate action based on the level of threat for each category. The different levels defined are as follows:

- **Level 1:** Attachments that are potentially unsafe, including executables, scripts, or those that perform system-related functions. Examples in this category

include batch files (*.bat), compiled Hypertext Markup Language (HTML) Help files (*.chm), MS-DOS programs (*.com), Control Panel extensions (*.cpl), Microsoft Windows Installer packages (*.msi), screen savers (*.scr), VBScripts (*.vbs), and other attachments that can potentially cause harm to systems when they are opened. The security update handles level 1 attachments by blocking them (preventing users from opening or saving them on their machines).

- **Level 2:** Attachments that are not considered unsafe. By default all file extensions not defined by level 1 are considered level 2, and the security update handles level 2 attachments by prompting the user to save them to disk while preventing them from being opened within the Outlook program itself. Users have the additional option of adding specific file extensions to level 2 if desired.
- **Level 3:** All attachments that are not defined by level 1 or 2 are in this category, and the security update handles them by allowing the user to open them from within Outlook or save them to disk as desired. By default no file extensions are defined as level 3, but when a new application is installed that creates a new file type, the file type is classified as level 3.

See Also: *virus, worm*

out-of-band management

An alternate connection for remotely administering a system or device.

Overview

Remotely administering servers, routers, and other network hardware is usually performed using in-band management, in which commands are sent over the same network connection that the server uses for sending data. If the network goes down, however, it is important to still be able to remotely manage such devices, and this is the purpose of out-of-band management. A typical

method of remotely managing network hardware out-of-band is to use a modem connection to a serial port on the hardware. This serial connection can be used to send commands to the operating system in order to reboot or reconfigure the device as necessary.

For example, if an attacker compromises a router and corrupts its routing table, the router will no longer be able to forward network traffic properly. As a result, the administrator will be unable to manage the router using the normal network connection since this network connection is down. By using a modem-based out-of-band connection to the router's serial port, however, the administrator can log on to the router and repair the routing table to bring the network connection back up again.

Notes

Microsoft Windows Server 2003 includes a feature called Emergency Management Services (EMS) that allows servers running these operating systems to be remotely managed using an out-of-band serial connection even when the server is hung, low on resources, or has blue screened.

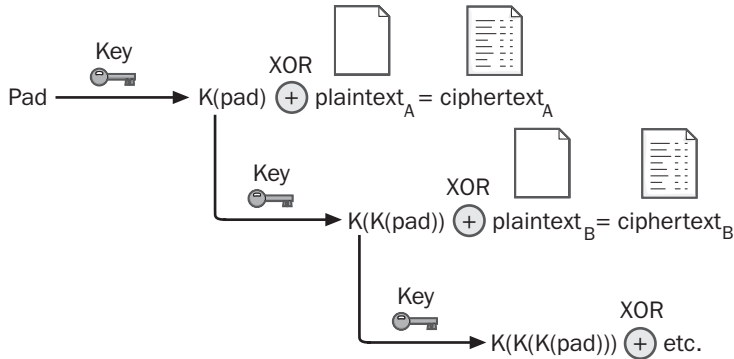
See Also: *exploit*

output feedback mode (OFB)

A type of block cipher employing a one-time pad (OTP).

Overview

Output feedback mode (OFB) is a stream cipher that encrypts plaintext using an OTP generated by the cipher. The pad is a random number of fixed length that is successively encrypted using a secret key known to both sender and recipient. The plaintext blocks are encrypted by XORing them with each successive encrypted pad, and the resulting ciphertext is transmitted to the recipient together with the original random number. The recipient then re-creates the series of pads by successively encrypting the received random number using the same secret key, and then XORs each block of ciphertext with the associated pad to recover the original plaintext blocks.



Output feedback mode (OFB). How the OFB cipher works.

OFB has several advantages as an encryption mechanism:

- The cipher is extremely fast since the series of pads can be created in advance before the plaintext is introduced.
- The cipher is resistant to noise because if a portion of ciphertext is damaged, only the corresponding plaintext is affected and not the entire message.
- Blocks of arbitrary size can be encrypted without having to pad undersized blocks of plaintext before applying the cipher.

The main disadvantage of the cipher is that it is susceptible to masquerading because if an attacker can obtain a portion of plaintext and its associated ciphertext, it is easy to forge an arbitrary message and send it to the recipient.

See Also: block cipher, cipher, one-time pad (OTP)

overt channel

The normal communication channel over which a system or network transfers information.

Overview

Overt channels are authorized channels for transmission of data. By contrast, a **covert channel** is a communications channel that hides illicit information flow within a normal communications stream, usually for purposes of information leakage or clandestine control of remote systems. In networking, an example of an overt channel would be a Transmission Control Protocol (TCP) session established between two authorized hosts, while a covert channel could hide information in the identification field of an Internet Protocol (IP) packet where such information is normally not found.

See Also: covert channel

P3P

Stands for Platform for Privacy Preferences, a mechanism for providing Internet users with privacy for their personally identifiable information (PII).

See: Platform for Privacy Preferences (P3P)

packet filtering

A mechanism that blocks packets based on a list of predetermined rules.

Overview

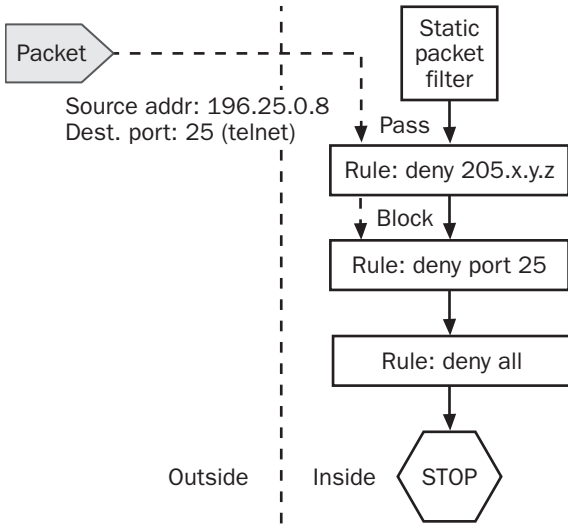
Packet filtering is a type of security technology used to control what kinds of packets are allowed to enter or leave a system or network. Packet filtering is typically used for blocking malicious traffic based on source address, port number, protocol type, and other criteria. Packet filtering is a standard feature of most operating systems, including UNIX/Linux and Microsoft Windows versions, and in most routers and firewall products. The rest of the discussion focuses on packet-filtering routers.

Implementation

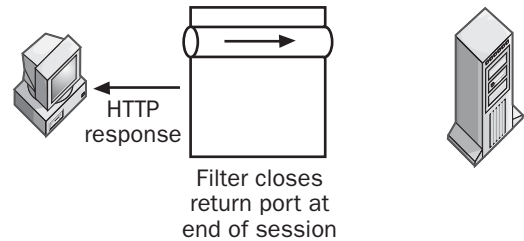
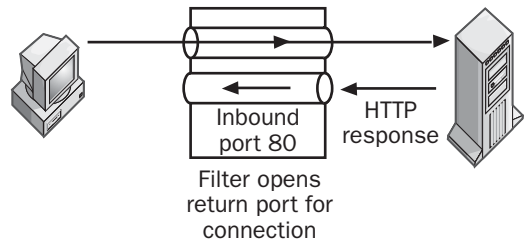
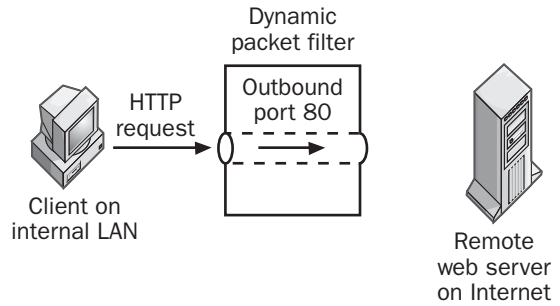
Packet filters come in two types: static and dynamic. Static packet filters determine whether to accept or block each packet based on information stored in the header of the packet, such as the source address or destination port of the packet. Static packet filters are typically found in operating systems and routers and use a series of rules for determining the fate of each packet. Administrators create these rules as an ordered list, and each packet that arrives at the filter is compared to each rule in succession until a match is found. If no match is found, the default rule, which is typically **deny all**, is applied. Rules can accept or deny packets based on header information describing the source, destination, and nature of the packet. Most packet filters allow two sets of rules to be created, one for inbound traffic and the other for outbound.

Dynamic packet filters operate similarly to static filters but also maintain session information that enables them to control the two-way flow of packets in a session between two hosts by dynamically opening and closing ports as required. Dynamic packet filters are commonly implemented in firewall products where they can be used to control the flow of traffic into and out of a network. For example, a dynamic packet filter could be configured so that the only inbound Hypertext Transfer Protocol (HTTP) traffic that can enter the network is traffic in response to requests from HTTP clients inside the network. To do this, outbound traffic through Transmission Control Protocol (TCP) port 80 is allowed, which enables HTTP requests from clients inside the network to reach the outside Internet. When an outbound HTTP request passes through the filter, the filter inspects the packets to obtain TCP session information for the request, and then opens port 80 for inbound traffic only in response to that request. When the HTTP response arrives, it passes through port 80 into the network, and then the filter closes port 80 for inbound traffic again. This way the only inbound HTTP traffic that is allowed to enter the network is traffic in response to HTTP requests from clients inside the network. All other inbound HTTP traffic is blocked.

This kind of approach is impossible with a static filter, which can be configured only to allow or block all inbound traffic to port 80 and not a portion of such traffic. Note, however, that dynamic packet filtering is not foolproof because an attacker could hijack a session and forge incoming traffic that would be allowed into the network. Note also that dynamic packet filtering is possible only with TCP packets and not with User Datagram Protocol (UDP) or Internet Control Message Protocol (ICMP) packets because UDP and ICMP are connectionless protocols that do not establish sessions for communication.



① Static packet filtering



② Dynamic packet filtering

Packet filtering. How static and dynamic packet filters work.

When configuring a packet filtering router, it's best to start by explicitly configuring a **deny all** rule even if this rule is implicitly used as the default rule. In other words, start with your packet filter in a completely locked down state and then gradually open it only to the degree that is necessary (an example of the least privi-

lege principle of network security at work). Filtering packets based on port number is typically used for protecting systems against known or potential vulnerabilities in network services, whereas filtering packets based on source address is used to protect networks against Internet Protocol (IP) spoofing attacks.

Notes

Another name for a dynamic packet filter is **stateful • packet filter** because it filters packets based on state (session) information contained in the packets.

See Also: *firewall, least privilege, port numbers, rule*

packet modification

Modifying information stored in network packets.

Overview

Packet modification is a technique used by attackers to gain control of target systems and networks. In a typical scenario, an attacker eavesdropping on a connection hijacks the session and then modifies information contained in session packets for malicious purposes. The best defense against such attacks is usually to encrypt all network traffic using Internet Protocol Security (IPSec) or some similar mechanism.

See Also: *eavesdropping, hijacking, Internet Protocol Security (IPSec)*

packet replay

Capturing and resending packets on a network.

Overview

Packet replay is a technique used by attackers to gain control of a communication session. The method typically involves capturing and recording traffic between two hosts, analyzing the packets and possibly modifying some of them, and then sending the captured packets back into the data stream to hijack sessions or perform other malicious actions. Packet replay is often employed to crack authentication sequences by enabling attackers to replay captured packets to become authenticated by the unsuspecting host or network.

There are various ways of protecting communication systems against packet replay attacks. Time-stamping packets and keeping track of sequence numbers of packets received can help prevent such attacks, while encrypting authentication sessions can provide additional defense against packet replay. Protocols such as Internet Protocol Security (IPSec) also include special

fields in packets that have a unique value for each security association or secure communication session and help prevent packet replay attacks from being performed.

See Also: *hijacking, Internet Protocol Security (IPSec), packet modification*

packet sniffer

Another name for protocol analyzer, a tool used to view network traffic at the packet level.

See: *protocol analyzer*

padding

A technique used in cryptography for simplifying the operation of encryption algorithms.

Overview

Padding is used in block ciphers and other encryption algorithms for adding bits to plaintext messages to make them an integral number of whole octets (bytes) or evenly divisible by some number. In message digest 5 (MD5), for example, the plaintext message must be an exact multiple of 512 bits (64 bytes) before the algorithm can be applied, and this is accomplished by appending to the end of the message a single “1” bit followed by a number of “0” bits until the message is 64 bits less than a multiple of 512. Then a 64-bit quantity that is a function of the number of bits in the original (unpadded) message is further appended to result in a padded message that is an exact multiple of 512 bits.

Padding is used in many other protocols for similar purposes. In Kerberos, for example, padding is included after the Timestamp field to ensure Kerberos messages are an exact multiple of 64 bits (8 octets) so they can be encrypted easily for secure transmission.

See Also: *block cipher, message digest 5 (MD5)*

Palladium

The former name for Next-Generation Secure Computing Base for Windows, a set of features for upcoming versions of Microsoft Windows operating system that

provides enhanced data security, personal privacy, and system integrity.

See: Next-Generation Secure Computing Base for Windows

PAM

Stands for pluggable authentication module, a UNIX model for extensible authentication architecture.

See: pluggable authentication module (PAM)

PAP

Stands for Password Authentication Protocol, a remote access authentication protocol supported by Point-to-Point Protocol (PPP).

See: Password Authentication Protocol (PAP)

parking lot attack

Another name for wardriving, a technique for finding poorly secured wireless networks.

See: wardriving

Passfilt.dll

Used for enhancing password security on systems running on the Microsoft Windows NT platform.

Overview

Passfilt.dll was included in *Windows NT 4.0 Service Pack 2* to increase the strength of passwords used to secure user accounts. Once Passfilt.dll is registered on the system, it modifies the password policy to ensure that all passwords are at least six characters in length and do not contain any portion of the name of the user. Passfilt.dll enforces password complexity by requiring that all passwords contain at least three of the following types of characters: lowercase letters, uppercase letters, numbers, and nonalphanumeric characters such as \$ or %. On later Windows platforms this functionality has been built into the operating system and is managed using Local Security Policy.

Notes

A Trojan horse named Passfilt.dll is also available on the Internet and allows passwords to be captured from systems running Windows NT and sent to an attacker.

See Also: password, Trojan horse

passive attack

Any form of attack that does not modify network traffic.

Overview

A passive attack is essentially a “listening attack” in which the attacker “listens” (captures) network traffic but doesn’t modify packets or insert new packets into the traffic stream. Passive attacks are “stealthy” in nature and are thus difficult to detect by administrators monitoring the security of their systems or networks. By contrast, an active attack is one that involves direct intrusion into network traffic through transmitting, modifying, or replaying packets.

See Also: attack, sniffing

passphrase

A phrase or sentence used in the same way a password is used.

Overview

Some applications such as the encryption tools employ passphrases instead of passwords. A passphrase can generally be any length and can contain spaces and might even express meaning like an ordinary sentence. The only real difference between a passphrase and password is length since typical passwords employed by users are usually around 6 to 12 characters in length to make them easy to remember.

In encryption tools such as Pretty Good Privacy (PGP), password phrases are used for generating unique session keys for encrypting messages. Such passphrases generally are long (50 to 100 characters) to ensure the resulting keys are strong enough to resist brute-force attempts to crack them. Passphrases for encrypting messages should be easy for the user to remember but hard for others to guess. For example, “To be or not to

be” would not be a good passphrase because a simple dictionary attack based on popular quotations would easily crack this phrase. By rearranging these words and adding numbers or special characters (for example, “be TO 468 NOT to # or”) the passphrase becomes more difficult to crack but can still be remembered with some effort. The most secure passphrases are, of course, strings of randomly generated characters, but human beings are generally not very good at remembering a string of 50 or 100 random characters!

See Also: *password*, *Pretty Good Privacy (PGP)*, *pseudorandom number generator (PRNG)*

Passport

Another name for .NET Passport, a system for managing online identity developed by Microsoft Corporation.

See: *.NET Passport*

Passprop

A tool for enhancing password security on Microsoft Windows NT.

Overview

Passprop is a command-line tool included in the *Windows NT 4.0 Server Resource Kit*. With this tool, administrators can enhance the security of Windows NT-based networks with the following methods:

- Enforcing password complexity to ensure passwords include a mix of upper- and lowercase letters, numbers, and symbols
- Enabling the default Administrator account to be locked out for interactive logons on all computers except domain controllers

See Also: *account lockout*, *Passfilt.dll*, *password*

password

A string of characters used to verify the identity of a user logging on to an application, system, or network.

Overview

Passwords are a fundamental element of the security of most systems and networks. They are also prime targets for intruders trying to break in and compromise such systems. A fundamental characteristic of password-protected systems is that the longer and more complex a password is, the more secure it is from being cracked. But this is also the fundamental weakness of such systems because the longer and more complex a password is, the harder it is for users to remember and the more likely it will be that users will expose their password for misuse in some fashion. For example, most users would have difficulty remembering a password like “t6Aq79J4rkM” and would therefore be likely to write it on a sticky note and stick it somewhere hidden like underneath the keyboard or in the bottom of a desk drawer. The problem is that one of the first things a **social engineer** (an attacker who gains physical entry to a company) might typically do is check under keyboards and in drawers for hidden passwords, and once a password is found this attacker can use it to gain access to sensitive information stored on the company network.

Passwords are therefore generally a trade-off between security and usability, and most users choose passwords that are 6 to 10 characters in length. Companies have several options to prevent passwords from being guessed or “cracked” by attackers:

- Providing users with a written security policy that governs the creation and use of employee passwords. For example, such a policy might require that all passwords have a minimum of eight characters and include some letters and some numbers. The policy might also prohibit users from employing parts of their names, addresses, or phone numbers in passwords.
- Enforcing a password policy using operating system features or add-on products. For example, in Microsoft Windows 2000, administrators can configure Password Policy, a part of Local Security Policy, to enforce such settings as minimum password length or complexity.

- Educating users about passwords by providing them with guidelines about how to create good passwords. A common suggestion for creating good passwords is to think of a phrase and then use the first letter of each word to form the password. For example, the phrase “I feel really bad for the way I treated you” might generate the password “ifRB4twity” if the user in fact felt “really bad” in this instance. Users should also be discouraged from using things like pets’ names or favorite movies as the basis for forming passwords and from thinking that by adding a simple numeric “123” to the end of a word a secure password results. Users should also be educated to guard against social-engineering attacks and to never give out their passwords except to known administrators or technical support people.
- Requiring that users change their passwords frequently either by creating a written policy or by enforcing the requirement in the operating system. Note that this can sometimes have the opposite effect, however, because users may find it more difficult to remember which password is their current one and may therefore be more tempted to write passwords down and keep them close at hand.

Because of the weakness of simple password-protected authentication, many businesses augment such systems with additional security measures, including smart cards and biometric identification systems. A high-security environment might employ three-factor authentication in which a user must enter a password, insert a smart card, and allow an iris scan in order to obtain access to the network. Such systems are far more difficult to crack than simple password-protected networks.

Notes

Many hardware devices such as routers come with default passwords that should be changed when the device is installed to protect against compromise.

See Also: authentication, biometric identification, one-time password (OTP), passphrase, password cracking, smart card

Password Authentication Protocol (PAP)

A remote access authentication protocol supported by Point-to-Point Protocol (PPP).

Overview

Password Authentication Protocol (PAP) is the simplest authentication protocol supported by PPP and transmits the user’s credentials (user name and password) over the connection in cleartext. As a result, PAP is also the least secure PPP authentication method and generally should not be used unless the client and access server cannot negotiate a more secure authentication protocol like Challenge Handshake Authentication Protocol (CHAP) or Microsoft Challenge Handshake Authentication Protocol (MS-CHAP). As outlined in RFC 1334, support for PAP is mandatory in PPP, but in practice it is usually needed only for connecting to older UNIX-based access servers that do not support other methods of authentication.

See Also: authentication, password

password-based encryption (PBE)

A method for generating a cryptographic key from a password.

Overview

Password-based encryption (PBE) algorithms are schemes that use passwords to generate secret keys for purposes of secrecy and data integrity. PBE algorithms are commonly used for secure storage of files or for protecting a user’s private key store on a system, but they also can be used for encrypting and signing electronic messages. Two public key cryptography standards (PKCSs) from RSA Security, PKCS #5 and #12, define PBE algorithms that can be used for generating secret keys from passwords.

Implementation

In a typical PBE scheme, the user’s password is appended with a **salt**, a pseudorandom number used to enlarge the space of possible passwords to reduce the susceptibility of the algorithm to brute-force key

search. The combination of password and salt is then hashed using a cryptographic hashing algorithm such as message digest 5 (MD5) or Secure Hash Algorithm-1 (SHA-1) to produce the secret key used for encryption. In some schemes the hashing function is iteratively applied a number of times to make it more difficult to crack the resulting key. Once the key has been generated, it can be employed with a standard symmetric key encryption algorithm such as Data Encryption Standard (DES) to encrypt the information to be protected.

See Also: *hashing algorithm, password, public key cryptography standards (PKCS)*

password cracking

Guessing the password for an application or system until the right one is found.

Overview

Since passwords form one of the foundations of security for most systems and networks, guessing or “cracking” passwords is high on the list of priorities for attackers trying to break into and compromise such systems. Cracking passwords can be approached two ways:

- **Online cracking:** This approach generally involves “sniffing” network traffic to capture authentication sessions and try to extract passwords from captured information. This is generally slow and difficult to accomplish, but there are tools available that are specifically designed for sniffing out passwords from network traffic.
- **Offline cracking:** This is the preferred method and involves compromising a system through some exploit to gain access to its password file or database, and then running a tool called a password cracker to try to guess valid passwords for user accounts. Offline cracking can be performed on the compromised machine or the password file can be “grabbed” and copied to a machine located outside the compromised network to be cracked at leisure. Even some worms such as DoubleTap and li0n can automatically grab passwords from infected systems.

Implementation

Password crackers guess passwords using two main techniques:

- **Dictionary attack:** This involves trying all words in a dictionary (a list of words typically used for passwords) to see if a valid match can be found. Sophisticated password crackers also use rules to generate complex combinations and variations of words in the dictionary; for example, by systematically varying between lower- and uppercase letters or appending simple numeric strings like “123” to the ends of words. Combining a dictionary and rule-based approach is often called a **hybrid attack**.
- **Brute-force attack:** When dictionaries fail, brute force is usually the only alternative. A brute-force attack simply involves trying all possible combinations of letters, numbers, and special characters to generate all possible passwords of every possible length until either the correct password is found or the program or attacker gives up.

The ease with which passwords can be cracked varies between different platforms and systems. Operating systems such as Microsoft Windows Server 2003 store passwords securely in encrypted form. To crack such passwords usually requires at the minimum physical access to the system using administrative credentials, and even then brute force is usually the only approach for extracting passwords. User applications such as office productivity tools can protect documents with passwords, and these are generally easier to crack than passwords for user accounts. Older platforms such as Windows 95 stored password information in *.pwl files that were weakly encrypted and easy to crack.

Marketplace

Two popular tools used by attackers for cracking passwords are L0phtCrack (whose current version is named LC4) and John the Ripper. While password crackers are frequently used for ill purposes, they also have valid uses in business environments. For example, an administrator might use a password cracker to audit the strength of user passwords to ensure guidelines outlined in the company security policy are being followed.

Companies like ElcomSoft (www.elcomsoft.com) and Password Crackers Inc. (www.pwcrack.com) also provide legitimate tools and services to companies that need to recover lost passwords in order to access password-protected documents or an administrator account or to disable screen savers.

Notes

Some devices such as routers and switches often have documented procedures for recovering passwords when passwords have been lost or forgotten. Refer to the vendor's Web site for more information.

See Also: *brute-force attack, dictionary attack, John the Ripper, L0phtCrack, password, Pwdump*

password grinding

Manually trying to guess passwords for an application, system, or network.

Overview

Password grinding is a primitive form of password cracking in which the attacker simply attempts to log on repeatedly to the target machine, trying different passwords until either the correct one is guessed or the system locks out the attacker. While this might seem like a fruitless activity, it is amazing how many users employ the word **password** as their passwords and how many administrators fail to change or disable the default passwords included with devices such as routers they install on their networks. Even considering the marked exaggeration of hacking abilities depicted in movies like *WarGames* and *Mission Impossible*, a knowledgeable cracker can occasionally succeed using this simple method and then leverage the obtained password to further compromise a target system or network.

See Also: *password, password cracking*

password hash

Stored passwords in encrypted form.

Overview

Password hashing is a security mechanism that uses encryption to protect passwords from unauthorized viewing. A hashing algorithm irreversibly converts passwords into unrecognizable form so that if an attacker can obtain a copy of the password file it will be

more difficult for the attacker to recover the original passwords. Password hashes are used in challenge-response authentication schemes such as NTLM for securely authenticating users without transmitting the password over the connection.

See Also: *hashing algorithm, password, password cracking*

password policy

A policy enforced by an operating system regarding attributes of passwords for user accounts.

Overview

Most operating systems today include support for password policies, a feature that allows administrators to configure what forms of passwords are acceptable for accounts and how these passwords are managed. On Microsoft Windows 2000, for example, Local Security Policy can be configured with the following password policy settings:

- Minimum allowed length for passwords
- Whether passwords can be simple (e.g., **password**) or complex (e.g., **paSS4321**)
- Whether a password history (list of old passwords) will be maintained or not, and the number of passwords maintained
- Minimum password age (time until password must be changed)
- Maximum password age (time until password expires unless it has been changed)
- Whether the password is stored internally using reversible or irreversible encryption

See Also: *password*

password recovery

Another name for password cracking, guessing the password for an application or system until the right one is found. Usually used in the context of legitimate activity.

See: *password cracking*

password shadowing

A technique used on UNIX platforms for hiding the location of passwords.

Overview

On UNIX systems user passwords, together with user names and other information concerning users, are stored in a world-readable file called `/etc/passwd`. One of the main goals of attackers trying to compromise such systems is “grabbing” the `passwd` file and then trying to crack the passwords it contains. Password shadowing separates the sensitive information (such as passwords) in this file from its public information (such as user names) and stores the sensitive information in a different file called a shadow file. Permissions on this shadow file are then configured as root-readable, which means only root (superuser) can access its contents, making it much more secure than the `passwd` file that anyone can access. The location of the shadow file varies with different platforms; for example, `/etc/shadow` on Linux and `/etc/security/passwd` on AIX.

See Also: `/etc/passwd`, `password`

patch

A fix for a flaw or bug in an application or operating system.

Overview

Patches are software fixes released by vendors to correct flaws in software products that can make them unreliable and result in loss or damage of data. Some flaws make products vulnerable to being compromised by attackers, in which case security patches are issued to correct the problems. The large numbers of security patches being released by vendors does not necessarily indicate the products were poorly designed. Instead, they often indicate vigilance on the part of the software vendor in response to the steadily increasing level of attacks being launched from the Internet. Nevertheless, patch management has become a major concern of many enterprises as the need to roll out new patches to large numbers of systems on a timely basis becomes important for the maintenance of secure networks.

One example of a patch management tool is Software Update Service (SUS), a tool from Microsoft Corporation

for deploying critical software updates across a network containing machines running Microsoft Windows 2000, Windows XP, or Windows Server 2003. Microsoft also provides services for notifying customers by e-mail when patches become available for newly discovered vulnerabilities in Microsoft products.

Because in our Internet-connected world network security affects everyone and not just the companies who own the networks, timely application of patches for known security vulnerabilities should be a priority for every business and organization. To help companies ensure their systems are up-to-date with security patches, the SANS Institute and the Federal Bureau of Investigation (FBI) work together to maintain a Top 20 List of the 20 most critical Internet security vulnerabilities. This list can be found at www.sans.org/top20/ and is updated periodically.

Notes

In Microsoft parlance, patches are known as hotfixes.

See Also: *hotfix*, *Microsoft Security Notification Service*, *Microsoft Security Response Center (MSRC)*, *Microsoft Security Update*, *Software Update Services (SUS)*, *vulnerability*

PBE

Stands for password-based encryption, a method for generating a cryptographic key from a password.

See: *password-based encryption (PBE)*

PCBC

Stands for plaintext cipher block chaining, a block cipher used in Kerberos authentication.

See: *plaintext cipher block chaining (PCBC)*

PCT

Stands for Private Communication Technology, a protocol for providing private communications over the Internet.

See: *Private Communication Technology (PCT)*

PEAP

Stands for Protected Extensible Authentication Protocol, an authentication protocol developed by Cisco for wireless networking.

See: Protected Extensible Authentication Protocol (PEAP)

Peekabooty Project

A project to develop software to bypass censorship restrictions on the Internet.

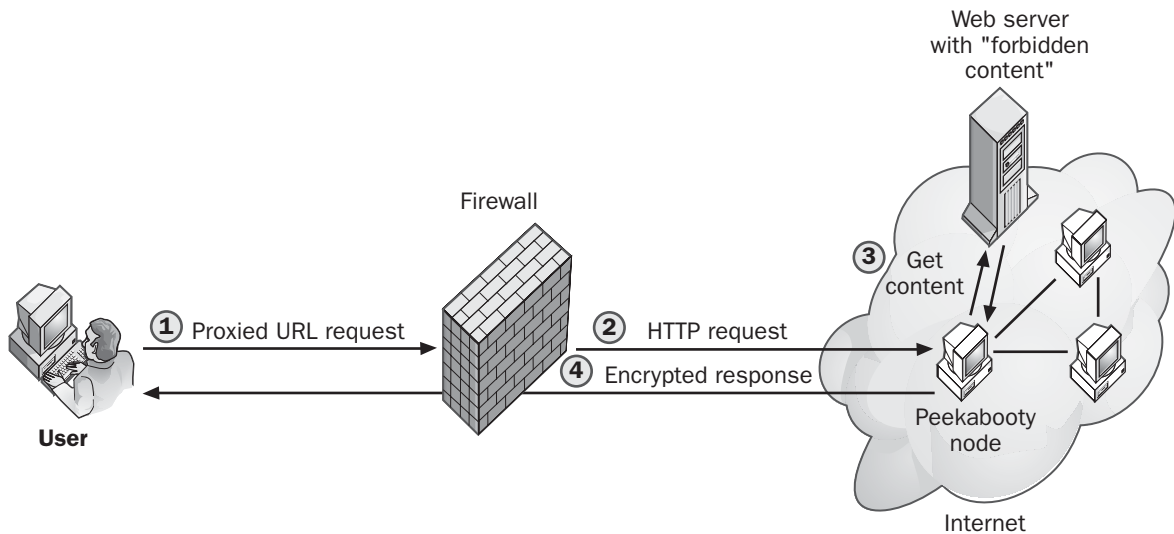
Overview

According to some estimates, almost two dozen countries censor some portions of the World Wide Web from their citizens. This is done by configuring firewalls at Internet service providers (ISPs) to prevent users from accessing certain Uniform Resource Locators (URLs) and by monitoring Internet traffic for content that is counter to laws or practices in these countries. The Peekabooty Project is a project for developing a peer-to-peer (P2P) network that can allow citizens in these countries to secretly access content that is otherwise censored by local authorities.

Peekabooty was originally developed by Hacktivism, a hacker group opposed to all forms of censorship on the Internet that was founded by a member of Cult of the Dead Cow (cDc). The project is currently run by Paul Baranowski, whose hacker pseudonym is Drunken Master. Peekabooty is an open source project released under the GNU Public License (GPL).

Implementation

Peekabooty consists of P2P software that resides on computers of willing users around the Internet. These computers act as proxies for relaying content to users behind firewalls in locations where such content is banned. To access a prohibited Web page, the user submits a URL to any Peekabooty node (a computer running Peekabooty P2P software) outside the firewall. The Peekabooty node retrieves the content and returns it to the user in encrypted form using a Secure Sockets Layer (SSL) connection. The firewall monitoring the connection is thus unable to differentiate the request from one made to a standard SSL-protected e-commerce site, and the result is the user has secretly accessed prohibited content.



Peekabooty Project. How Peekabooty works.

In order to use Peekabooby, users are not required to install any software on their client computers (installing such software might be interpreted by local authorities as an illegal act and may get users into trouble). All users must do to use Peekabooby is configure the proxy settings for their Web browser to forward requests for URLs to Peekabooby nodes on the Internet. The hope behind Peekabooby is that so many civil libertarians worldwide will eventually allow their computers to be used as Peekabooby nodes that countries censoring Internet content will be unable to block all possible nodes, enabling users in these countries to always find new nodes for accessing banned content. This distributed model is common in P2P computing and makes it difficult to control once it is deployed, which is the whole idea behind Peekabooby: to set up something that authorities can't control.

For More Information

Visit www.peekabooby.org for more information.

See Also: *firewall, privacy, Publius Project*

PEM

Stands for Privacy Enhanced Mail, a scheme for ensuring the privacy of e-mail sent over the Internet.

See: *Privacy Enhanced Mail (PEM)*

penetration testing

Testing the security of network defenses.

Overview

Configuring networks so they are secure is one thing; testing configurations to see whether they are secure is another. Penetration testing is an important part of network security and involves testing various aspects of network defense to see whether they really work. Penetration testing can uncover a variety of weaknesses in network defenses, including vulnerable services, procedural weaknesses, ineffective policies, and configuration problems. Penetration testing can test every aspect of a network including the internal local area network (LAN), servers, workstations, dial-in and leased-line wide area network (WAN) links, firewalls, operating systems, and applications.

There are two ways to perform penetration testing on a network:

- **Remote penetration testing:** Trying to uncover weaknesses in defense from outside the network. This can be done either with no prior knowledge of network configuration (no-information penetration testing) or in conjunction with network documentation provided by the company whose network is being tested. Although no-information penetration testing might seem preferable since it aligns more closely with how attackers usually work, in practice it can be less useful because penetration testing is usually a time-limited evaluation phase and attackers often have lots more time on their hands.
- **Internal penetration testing:** Analyzing the security of the network from within by examining system configurations and performing various tests. This approach can be more comprehensive than remote testing, but best practice is usually to combine both types of testing to ensure potential vulnerabilities are not overlooked.

Marketplace

A number of organizations provide penetration services for other companies, including En Garde Systems, KSAJ Inc., the NCC Group, and ProCinct Security. Companies should perform due diligence prior to hiring organizations that perform such tests because improperly conducted tests could actually result in damage or harm to systems or data. Companies with trained security personnel may be able to perform their own penetration tests using popular security tools such as Nmap and Nessus.

See Also: *Nessus, Nmap, vulnerability*

perfect forward secrecy (PFS)

A property of an encryption scheme that makes it difficult to compromise.

Overview

If an encryption scheme has perfect forward secrecy (PFS), attackers cannot compromise a communication session even if they could eavesdrop to obtain a transcript of an entire conversation and also break into each

party's system and steal their long-term secrets. Typical encryption schemes that have PFS are those that use session keys with the following characteristics:

- Uniquely generated for each session
- Not derived from long-term secrets stored by participants
- Forgotten completely after a session is over

In addition, such session keys are usually securely exchanged between the parties using a public key cryptography scheme such as Diffie-Hellman (DH). Kerberos does not have PFS because session keys included in tickets are encrypted with long-term secrets.

See Also: *eavesdropping, session key*

perimeter network

Another name for demilitarized zone (DMZ), an isolated network segment at the point where a corporate network meets the Internet.

See: *demilitarized zone (DMZ)*

permissions

Rules governing how objects such as files can be accessed.

Overview

Permissions are an essential component of the security of applications, systems, and networks. Permissions are used to control who has access to objects and what level of access they have. Types of objects that are typically secured using permissions include files, printers, and objects stored in directories.

Permissions are typically either allowed or denied, with permissions denied typically overriding permissions allowed. For example, if a user belonging to a group has permissions allowed over an object while the group has permissions denied, the user will typically be denied access to the object. Levels of permissions are often cumulative as well, so that a user who explicitly has full control permission over a file implicitly has the lesser read permission as well. Other rules for combining

permissions depend on the type of permissions being considered.

On Microsoft platforms, some of the common kinds of permissions include the following:

- NTFS permissions for files stored on NTFS volumes
- Shared folder permissions for folders that are shared for network access
- Printer permissions managing access to network printers
- Active Directory permissions for controlling access to objects in the Active Directory directory service
- Code-access, identity, and role-based permissions in the Microsoft.NET Framework

See Also: *NTFS, rights*

personal data

Another name for personally identifiable information (PII), information regarding the identity of a person.

See: *personally identifiable information (PII)*

personal identification device (PID)

A device used to establish a person's identity.

Overview

Personal identification devices (PIDs) typically are used to authenticate users so they can access systems or networks. PIDs are typically small devices that can easily be carried around; examples range from plastic cards with magnetic strips to handheld objects containing embedded memory chips and biometric fingerprint scanners. PIDs can contain anything from a person's name and company ID number to passport number, driver's license number, or whatever other personally identifiable information (PII) is required for use and operation. PIDs often are used in conjunction with passwords or personal identification numbers (PINs) so that if the PID is lost it can't be used by unauthorized parties.

Marketplace

A number of vendors offer PIDs of various kinds; some of the popular ones are Digipass Go from Mertek Systems, IDDisk from Immtec Inc., SKV from Secure Systems, and DigiPass Pro from Vasco.

See Also: password, personally identifiable information (PII)

personal identification number (PIN)

A unique identifier used together with a personal identification device (PID).

Overview

Personal identification numbers (PINs) are used to protect the security of PIDs by providing added proof that the person trying to use the PID is in fact the authorized owner. PINs are known only to the person who owns the PID and should never be divulged to anyone. PINs are typically numbers with four or more digits; the length of the number often is a characteristic of the type of PID being used. PINs help ensure that an individual's personally identifiable information (PII) stored on the PID remains private and does not fall into the hands of unauthorized parties.

See Also: personally identifiable information (PII), privacy

personal information

Another name for personally identifiable information (PII), information regarding the identity of a person.

See: personally identifiable information (PII)

personally identifiable information (PII)

Information regarding the identity of a person.

Overview

Personally identifiable information (PII) is a term used in government, finance, and advertising to refer to

personal information collected from individuals stored and for verifying their identity later. For example, an e-commerce Web site typically collects PII the first time a consumer purchases something from the site, and then stores this information in a database so the consumer won't have to reenter it every time he or she returns. PII can include such things as name, country, street address, e-mail address, credit card number, Social Security number, government ID number, Internet Protocol (IP) address, or any other unique identifier associated with the individual. Fair Information Practices (FIP), a set of standards governing collection and use of personal data that dates back to the U.S. Privacy Act of 1974, help protect the privacy of PII collected from individuals by industry and government.

See Also: Fair Information Practices (FIP), identity theft, privacy

PFS

Stands for perfect forward secrecy, a property of an encryption scheme that makes it difficult to compromise.

See: perfect forward secrecy (PFS)

PGP

Stands for Pretty Good Privacy, a popular e-mail encryption technology.

See: Pretty Good Privacy (PGP)

phishing

Conning someone into telling you his or her password or other sensitive information.

Overview

While password cracking is an entirely technical approach to trying to obtain a user's password, social-engineering approaches often are faster, easier, and have a higher rate of success. One organization performed a study and found that four out of five individuals working for a company would tell you their password if you asked them in the right way; for exam-

ple, by pretending to be a technical support person or network administrator. This clearly highlights the fact that network security is more than just a technical issue but a human one as well. Many e-mail scams are based on phishing for other useful information such as bank account numbers or credit card numbers.

See Also: *password, social engineering*

Phrack

One of the oldest online hacking magazines.

Overview

Phrack describes itself as “a Hacker magazine by the community, for the community” and says, “Those who know us know what we do; others do not have to.”

Phrack originated in 1985 as an ASCII-formatted “philes,” or articles, containing information about hacking, cracking, phreaking, and general anarchy distributed on bulletin boards and mailing lists. The magazine targets mainly the black hat community, but articles are usually of high technical quality and often contain information useful to legitimate network security professionals as well.

For More Information

Visit www.phrack.org for more information.

See Also: *2600, cracking, hacking, phreaking*

P

phreaking

Hacking and cracking telephone and telecommunications networks.

Overview

Phreaking became popular in the 1960s when early hackers, motivated largely by curiosity concerning anything technical, began to investigate what was then the largest network in the world, the telephone system. Soon “phreakers” learned how to map out the various switches and trunk lines of the Plain Old Telephone System (POTS) and learned how to fool the system into providing them with free long-distance calls using equipment as simple as a whistle included in a box of Cap’n Crunch cereal (the whistle generated an audible tone of 2600 Hz, which was the tone used for triggering

telephone switches and which later became the title of the earliest magazine for hackers and phreakers). The general idea was not to steal services from the telephone company but to display technical prowess to peers and discover undocumented secrets concerning the technology’s operation.

One popular phreaking activity in the 1970s was **boxing**, the construction of devices for fooling the telephone system into performing different actions. Some of the different types of boxes that were designed included these:

- **Blue box:** Generates the 2600-Hz tone described previously used for switching trunk lines, bumping the operator, and other activities
- **Black box:** Made the phone company think your phone was out of order so you wouldn’t be billed
- **Cheese box:** Made your phone behave like a public pay phone, often used by bookies for rerouting calls to hide their origin
- **Red box:** Simulated the sound of a coin dropping into a pay phone so free long-distance calls could be made from pay phones.

Phreaking declined in the 1980s when law enforcement agencies began to crack down on individuals manipulating the phone system to avoid paying for services, and with the disappearance of POTS and the rise of the modern digital phone system many of the early hacks performed by phreakers no longer worked. With the emergence of mobile cellular systems, however, phreaking has reemerged to some extent as an underground activity that requires considerable technical ingenuity to perform.

See Also: *hacking*

physical security

Securing computer systems by physically isolating and protecting them.

Overview

Physical security is a sometimes-neglected aspect of network security because it is generally viewed as “low tech” compared to other aspects of defending networks. Physical security can be as simple as placing key servers in a back room and locking the door to prevent unauthorized access. Although most companies invest wisely in protecting their networks from attack from without, insider attack by disgruntled employees or intruders who clandestinely have entered premises using social-engineering techniques may actually constitute a bigger threat to the security and integrity of business information systems. By simply entering an unlocked server room, an attacker may be able to boot a server from a CD-ROM and gain access to critical data or install backdoors for stealthy remote control of network resources.

Physical security may involve some or all of the following activities, depending on the degree of security required:

- Placing critical computing systems in locked rooms and limiting who has access to those rooms (and even hiding the location of those rooms)
- Employing electronic keycard locking systems for server rooms that log all entries to keep track of who comes in
- Monitoring server rooms using video cameras with remote recording facilities for protecting resulting tapes
- Disabling or removing hardware such as floppy disks and CD-ROMs so that physically compromised machines cannot be taken over by an attacker
- Ensuring that backup media containing sensitive company data are physically secured off premises in vaults or other locked containers

See Also: *headless server*

PIC

Stands for Pre-IKE Credential, a proposed replacement for the Internet Key Exchange (IKE) protocol.

See: *Pre-IKE Credential (PIC)*

PID

Stands for personal identification device, a device used to establish a person’s identity.

See: *personally identifiable information (PII)*

PII

Stands for personally identifiable information, information regarding the identity of a person.

See: *personally identifiable information (PII)*

pilfering

Grabbing as much information as possible after compromising a system or network.

Overview

Once an exploit has been performed, pilfering is an activity performed by attackers that compromises the security of a system. The goal usually is not theft of company data (unless that was the original intent of the attack) but obtaining password hashes and other information that can be leveraged later to enable the attacker to compromise other systems on the network. The next stage after pilfering usually involves installing a backdoor, a hidden mechanism to allow the attacker to reenter the system secretly later without having to reperform the original exploit. After a backdoor is installed, the attacker wipes logs to remove evidence of the attack and then moves on to attack other systems.

Notes

Other activities sometimes referred to as pilfering involve siphoning off bandwidth for remote access or Internet connections.

See Also: *backdoor, hacking, Pwdump*

PIN

Stands for personal identification number, a unique identifier used together with a personal identification device (PID).

See: *personally identifiable information (PII)*

ping

A utility that verifies the integrity of a network connection.

Overview

The Ping command is one of the first commands to use when troubleshooting communication problems on a Transmission Control Protocol/Internet Protocol (TCP/IP) network. To use Ping, you open a command line window and type **ping** followed by either the IP address or the fully qualified domain name (FQDN) of the host for which you want to test network connectivity. Internet Control Message Protocol (ICMP) echo packets are then transmitted to the host, and if connectivity is working, an equal number of echo replies are received. The replies show the packet size in bytes, response time in milliseconds, and Time to Live (TTL) of the echo reply. The TTL is decremented for each hop along the way and indicates the number of routers (hops) passed through along the network path.

Many firewalls, routers, and secure hosts discard ICMP echo packets and do not respond to Ping. This decreases the likelihood that crackers will find them. The ICMP echo packet has been used in many attacks, including ping floods and the ping of death. The best way to prevent these attacks is to make sure that systems are kept up to date.

See Also: *ICMP attacks, ping flood, ping of death*

ping flood

Flooding a system or network with Internet Control Message Protocol (ICMP) echo requests.

Overview

Ping flood was one of the earliest forms of denial of service (DoS) attacks and involves sending large numbers of ICMP echo request packets to a target host or network. If the bandwidth of the attacker is significantly greater than that of the target, the network can become saturated with responses to the requests and connections from legitimate users are denied. Ping floods have been used to attack Web sites, Internet Relay Chat (IRC) servers, firewalls, and other hosts connected to the network. Attackers typically use such tools as

Smurf that can send large numbers of requests simultaneously using spoofed source addresses to make them harder to trace. Filtering source addresses of incoming packets and blocking ICMP traffic entirely are the typical approaches used to deal with such attacks.

See Also: *denial of service (DoS), ping of death, Smurf attack*

ping of death

Crashing target systems by sending oversized ping requests.

Overview

The ping of death attack appeared in late 1996 and exploited a weakness in the design of Internet protocols. The Transmission Control Protocol/Internet Protocol (TCP/IP) standards limited the allowable size of packets to 65,535 bytes, but by creating a packet larger than this and fragmenting it into several portions, an attacker could crash or hang a TCP/IP host connected to the Internet. It worked in this way: when the fragments arrived at the target, the target tried to reassemble them, but because the result exceeded the buffer of the target's TCP/IP stack, an error condition resulted. Depending on whether the target was a computer system, router, network printer, or some other TCP/IP device, the target either would hang, crash, or reboot.

Internet Control Message Protocol (ICMP) was the first TCP/IP protocol used for exploiting this issue since early implementations of the ping utility allowed attackers to construct oversized ICMP echo requests (legitimate ping packets are only 64 bytes in length). Later other TCP/IP protocols such as User Datagram Protocol (UDP) were exploited for similar purposes. Patches have been developed to correct this flaw for products, and the ping of death is rarely seen nowadays but serves to highlight the fact that TCP/IP was not originally designed with security as its foremost concern.

See Also: *IP fragmentation attack, Jolt2, ping flood*

ping sweep

A method for footprinting a network using Internet Control Message Protocol (ICMP) echo requests.

Overview

Ping sweeping is a scanning technique that uses ping requests to try to determine which hosts are alive and listening on a network. A ping sweep generates a series of ping requests addressed to a targeted range of Internet Protocol (IP) addresses. The addresses that respond with ICMP reply messages are assumed to be present, while addresses that don't reply are likely unused. While the ping utility included with both Microsoft Windows and UNIX/Linux platforms is a legitimate network testing tool, ping sweeps generally utilize more sophisticated tools that can automatically scan an entire range of addresses and perform other kinds of tests to enumerate hosts that are detected.

A number of tools are around that can be used to perform ping sweeps on remote networks, including Fping, Gping, Nmap, Pinger, Ping Sweep, and Rhino9. To evade firewalls that block ICMP echo requests, many of these tools can use other types of packets, including ICMP time stamp and address mask requests as alternatives.

Notes

Ping sweeping is also known as ICMP sweeping.

See Also: enumeration, footprinting, Fping, Nmap, scanning

PKCS

Stands for public key cryptography standards, a series of specifications for Public Key Infrastructure (PKI) implementation.

See: public key cryptography standards (PKCS)

PKCS #7

A specification for cryptographic message syntax.

Overview

PKCS #7 is the most widely implemented of the PKCS de facto standards issued by RSA Security. PKCS #7 forms the basis of the Cryptographic Message Syntax (CMS) standard of RFC 2630 that outlines how to authenticate, digest, encrypt, and sign digital messages. Uses for PKCS #7 include certificate requests for Public Key Infrastructure (PKI) and digital signatures for

the Secure/Multipurpose Internet Mail Extensions (S/MIME) secure messaging standard. The current version of PKCS #7 is 1.6, but version 1.5 is used as the basis for S/MIME.

See Also: public key cryptography standards (PKCS), Public Key Infrastructure (PKI), Secure/Multipurpose Internet Mail Extensions (S/MIME)

PKI

Stands for Public Key Infrastructure, a set of technologies and policies for authenticating entities using public key cryptography.

See: Public Key Infrastructure (PKI)

PKINIT

An extension to Kerberos that adds public key cryptography.

Overview

PKINIT enhances the Kerberos specification RFC 1510bis by allowing Kerberos clients to have their initial authentication performed using public key cryptography. PKINIT is derived from "Public Key cryptography for INITIAL authentication" and is currently an Internet-Draft standard being considered by the Internet Engineering Task Force (IETF). The advantages of incorporating public key cryptography into Kerberos include easier key management and the ability to leverage emerging Public Key Infrastructure (PKI) systems. The PKINIT draft standard specifies how preauthentication data fields and error data fields in Kerberos messages can be used for carrying public key data.

See Also: Kerberos, public key cryptography

PKIX

Stands for Public-Key Infrastructure (X.509), a set of standards for implementing an X.509-based public key infrastructure (PKI).

See: Public-Key Infrastructure (X.509) (PKIX)

plaintext

Information that is unencrypted.

Overview

Encryption is the process of transforming plaintext into ciphertext. Plaintext is information that is in human-readable form; for example, an e-mail message typed in a text editor. To prevent sensitive information from being read if it is intercepted by someone other than its intended recipient, the message can be encrypted using a mathematical procedure called an encryption algorithm. The result of applying this algorithm to the information is ciphertext, a string of bits that still contains the original information but cannot be read by anyone unless it is first decrypted to convert it back into plaintext.

See Also: *ciphertext, encryption, encryption algorithm*

plaintext cipher block chaining (PCBC)

A block cipher used in Kerberos authentication.

Overview

Plaintext cipher block chaining (PCBC) is a modified form of cipher block chaining (CBC), a feedback mechanism commonly used in block ciphers. PCBC provides a mechanism for detecting when encrypted communications are compromised by ensuring that, if a portion of an encrypted message is changed, the content of the remaining part of the message is garbage (indecipherable). By including a standard block of data at the end of each Kerberos message, a recipient can test whether a message has been tampered with by seeing if this standard data decrypts properly.

See Also: *cipher block chaining (CBC), Kerberos*

Platform for Privacy Preferences (P3P)

A mechanism for providing Internet users with privacy of their personally identifiable information (PII).

Overview

Platform for Privacy Preferences (P3P) is a project of the World Wide Web Consortium (W3C) intended as an industry standard for protecting the privacy of users who submit personal information to Web sites they visit. P3P provides a mechanism for implementing privacy policies on Web sites and provides users with clear and unambiguous information about how sites will handle their personal information. P3P allows privacy policies for Web sites to be implemented in a standardized machine-readable format so that P3P-supporting Web browsers can automatically compare a site's policy to the user's privacy preferences configured in the browser.

The P3P 1.0 Recommendation was released in April 2002, and the W3C has published a number of guides and tools for how to implement P3P on both the client and server sides. Version 6 of Microsoft Internet Explorer Web browser supports many aspects of P3P as does Mozilla and other browsers.

For More Information

Visit www.w3.org/P3P/ for more information.

See Also: *privacy*

playback

Another name for packet replay, capturing and resending packets on a network.

See: *packet replay*

pluggable authentication module (PAM)

A UNIX programming model for extensible authentication architecture.

Overview

Pluggable authentication module (PAM) is a framework that allows new authentication services to be installed on UNIX systems to enhance their security. PAM consists of a library, configuration file, and pluggable modules, each of which implement a different authentication scheme. By stacking such modules, a UNIX system can

be configured to try multiple authentication methods when attempting to authenticate a user without the need for the user to reenter credentials for each module. Some of the popular PAM modules include those for implementing authentication using Kerberos, Rivest-Shamir-Adleman (RSA) public keys, smart cards, and S/Key.

See Also: authentication

Point-to-Point Tunneling Protocol (PPTP)

A method of encapsulating network traffic used for virtual private networking.

Overview

Point-to-Point Tunneling Protocol (PPTP) is a tunneling protocol used to allow remote clients secure access to private networks over the Internet, thus creating a virtual private network (VPN) that overlays a portion of the public Internet. PPTP works by encapsulating Point-to-Point Protocol (PPP) packets in Internet Protocol (IP) packets for sending them over the Internet. PPTP was developed by a consortium of Microsoft and other companies, and its details are outlined in RFC 2637.

See Also: tunneling, virtual private network (VPN)

port flooding

Sending large numbers of Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) packets to a particular port.

Overview

Port flooding is a technique used by attackers for performing various kinds of denial of service (DoS) attacks, including the following:

- Preventing legitimate network users from connecting to services on network services
- Filling up process tables on routers and switches, causing them to crash or reboot

- Causing Ethernet switches to go into an error state that makes them work like hubs so the attacker can sniff traffic on remote segments

See Also: denial of service (DoS), port numbers

port forwarding

A method used by Secure Shell (SSH) for secure communications over the Internet.

Overview

Port forwarding is used to create an encrypted communication session between an SSH server and SSH client. Port forwarding works by mapping a specific port on the server to one on the client to securely “pipe” traffic through a firewall and over the Internet.

Implementation

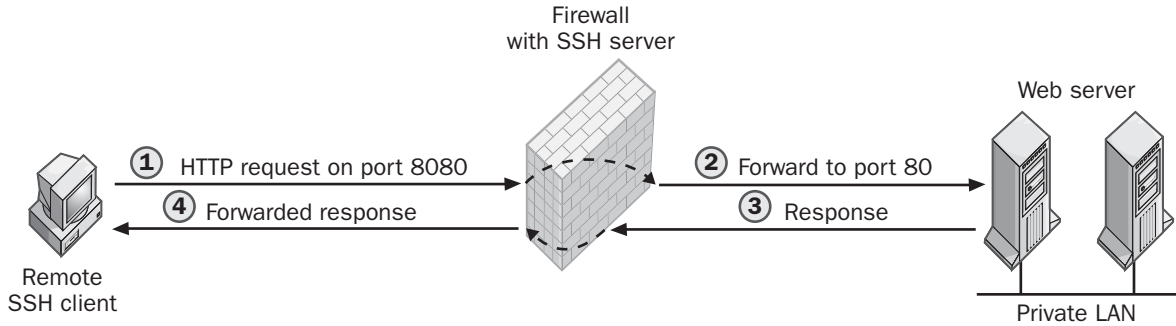
In a typical scenario, a company runs an SSH server on its firewall. A remote client running an SSH client wants to connect to a server on the company’s internal network. To make things interesting, assume the internal network is using Network Address Translation (NAT) so that its hosts have private Internet Protocol (IP) addresses assigned and that the firewall also blocks all incoming Hypertext Transfer Protocol (HTTP) traffic on port 80. As a result, the internal Web server is for intranet use and is not accessible to normal HTTP clients on the Internet. To connect to the internal server, the remote SSH client establishes a connection to a different port (for example, 8080) on the firewall, and the SSH server forwards traffic from the client to port 80 (the standard HTTP port) on the internal Web server. Here, SSH is used to implement local port forwarding by mapping one port (8080) to another (80) to circumvent the firewall and bypass NAT security to access a private IP address from the Internet. The opposite approach also works: remote port forwarding can be configured to allow internal clients to securely access remote servers using mapped ports to circumvent firewall restrictions.

Notes

Although port forwarding can be used as a legitimate means of circumventing firewall security, attackers also

can exploit it if they can compromise a firewall and install an SSH server on the firewall host. Some NAT devices can themselves be configured to support a form of port forwarding called inbound mappings to allow

external clients to access internal servers that have private IP addresses, but this is not a secure approach since encryption is not used as it is in SSH.



Port forwarding. How port forwarding works.

See Also: network address translation (NAT), port numbers, port redirection, Secure Shell (SSH)

port numbers

Identifiers for ports representing different network services.

Overview

A port is a kind of open door on a computer system or network device that is used to listen for connection attempts from clients trying to access services. For example, the Hypertext Transfer Protocol (HTTP) port is used on Web servers to listen for client machines trying to connect and download Web pages using HTTP, one of several application-layer Internet protocols. Each port representing a different network service has its own unique port number to identify it to clients; for example, port number 80 is used by Web servers listening for HTTP connections. In addition, each port number can be used two ways:

- For establishing communication sessions using Transmission Control Protocol (TCP), usually for the purpose of transferring data between hosts

- For connectionless communications using User Datagram Protocol (UDP), often for making enquiries concerning network services

Port numbers range from 0 to 65,535 ($2^{16} - 1$) and are organized into three categories:

- **System (well-known) port numbers:** These range from 0 to 1023 and represent standardized port numbers assigned by the Internet Assigned Numbers Authority (IANA); they are recognized across the industry.
- **User (registered) port numbers:** These range from 1024 to 49,151, and IANA does not control their assignment but does record how they are commonly used by applications from different vendors.
- **Dynamic and/or private port numbers:** These range from 49,152 to 65,535 and can be dynamically assigned by applications as needed.

Port numbers are important from the perspective of network security in several ways:

- Port numbers represent the doors through which intruders try to gain entry to systems and networks. Attackers typically do this by scanning a system or network to see which ports are listening (open), and

they use this information to identify services that might be vulnerable to attack. By locking down systems to remove nonessential services (thus closing their ports) and by configuring firewalls to allow traffic through only a limited number of ports, administrators can protect networks from many common forms of attack.

- Port numbers can be changed from standard values to nonstandard ones to try to hide network services from attackers, an approach called “security through obscurity.”
- Many common Trojans masquerade as legitimate network services by assuming control of well-known port numbers, while others have their own specific port numbers, which can be of help in identifying when a system has been infected with a Trojan.

For More Information

Visit IANA at www.iana.org/assignments/port-numbers for the latest list of well-known and assigned port numbers.

See Also: *hardening, obscurity, scanning, Trojan*

port redirection

A method used by attackers for circumventing firewall security.

Overview

Port redirection involves compromising an intermediary host, installing a port redirection program, and using the intermediary host to establish communications with another host (the target). The attacker then sends packets to the intermediary host, which redirects them to the target host so that they appear to have originated from the intermediary host instead of the attacker. If the intermediary host is a network firewall and the target host a server on the internal network, port redirection provides the attacker with access to resources on the target server.

Implementation

A popular tool used for port redirection is Fpipe from Foundstone. Although Fpipe can be used legitimately

for activities such as penetration testing, it (and similar tools like Datapipe or even Netcat) also can be used for port redirection by attackers trying to crack a network. Once the attacker has succeeded in compromising the firewall, the attacker installs Fpipe on the firewall host and configures Fpipe to listen on some inbound port that is often left open on firewalls such as User Datagram Protocol (UDP) port 53, which is used for Domain Name System (DNS) zone transfers. Fpipe is then configured to shovel (redirect) all incoming packets on UDP port 53 to port 23, allowing the attacker to connect to a Telnet server hidden behind the firewall. Because the firewall will typically block incoming connections from high-numbered dynamic ports of external clients, Fpipe also can be configured to spoof the source address of the incoming packets to UDP port 53 so the firewall thinks it is allowing an incoming DNS zone transfer to occur.

Notes

Secure Shell (SSH), a secure replacement for Telnet and other UNIX utilities, employs a similar technique called port forwarding.

See Also: *firewall, Fpipe, penetration testing, port forwarding, port numbers, Secure Shell (SSH)*

port scanning

A method for determining which ports are “listening” (open) on a target system or network.

Overview

Port scanning is a method used by attackers for gathering information about which services are running on a target system or network of systems. Network services commonly employ standard port numbers to identify themselves to clients wanting to connect to them. For example, port number 80 is the standard Transmission Control Protocol (TCP) port that Web servers listen on for incoming Hypertext Transfer Protocol (HTTP) requests from Web browsers. If a port scan determines that a target host has port 80 open, there’s a good chance that the host is a Web server, and the attacker’s next step is to enumerate the target by gathering information

about user accounts and applications on the system, searching for a known vulnerability to exploit.

Implementation

Some of the different approaches used for port scanning include these:

- **Vanilla scan:** The attacker tries to connect to all possible ports on the remote system from port number 0 through 65,535 by sending TCP SYN packets to each port of each address. This is also called a SYN scan because it uses TCP SYN packets.
- **Strobe scan:** The attacker tries to connect to a specific set of ports commonly open on Microsoft Windows or UNIX/Linux hosts (another form of SYN scan but faster than a vanilla scan).
- **UDP scan:** The attacker sends empty User Datagram Protocol (UDP) packets to different ports for a range of addresses and looks at the response. Some operating system platforms respond with Internet Control Message Protocol (ICMP) error packets when empty UDP packets are received by listening ports, while closed UDP ports typically respond with “port unreachable” packets. ICMP packets can also be used for similar purposes.
- **FTP bounce:** The attacker performs the scan through an intermediary File Transfer Protocol (FTP) server to disguise the location of the attacker’s machine.
- **Sweep:** The attacker scans a large range of Internet Protocol (IP) addresses looking for systems that have one specific port open (such as port 23 for Telnet servers).
- **FIN scan:** The attacker sends a TCP FIN packet to all (or some) ports for a range of addresses. The FIN packet indicates the sender wants to close a TCP session. If the port is closed already, the target usually replies with a TCP RST packet, but if the port is open (connected to some other host), the FIN packet is dropped. This is a stealthy form of scanning since it does not actually involve establishing connections with target hosts and such attempts often are not logged by the target system.

- **Passive scan:** The attacker captures all network traffic entering or leaving the remote network (perhaps by compromising the network’s firewall and installing a sniffing program) and analyzes the traffic to determine which ports are open on which hosts on the network (all the scans described previously are active scans, which are more common but less stealthy).

A simple tool such as Telnet can be used to scan a remote host one port at a time. For example, by starting a Telnet client session and trying to open port 80 on a target machine, an attacker can manually issue HTTP GET commands in the proper format and determine whether the target responds like a Web server would. More sophisticated tools such as Nmap and Nessus can be used to automate port scans against a range of IP addresses and can enumerate additional information that may help hackers perform exploits to compromise target systems. Other popular port-scanning tools include Netcat, Strobe, Pscan, and SATAN.

See Also: enumeration, hacking, Netcat, Nmap, port numbers, SATAN

PPTP

Stands for Point-to-Point Tunneling Protocol, a tunneling protocol used for virtual private networking.

See: Point-to-Point Tunneling Protocol (PPTP)

Pre-IKE Credential (PIC)

A proposed replacement for the Internet Key Exchange (IKE) protocol.

Overview

Pre-IKE Credential (PIC) is one of several proposed replacements for IKE, the key management protocol used by Internet Protocol Security (IPSec). PIC is intended to overcome some of the deficiencies of IKE, including its complexity of operation and its lack of support for legacy authentication methods widely used in the marketplace. PIC works by “bootstrapping” IKE authentication in which a user is first authenticated using a legacy method and then the authentication

server generates IKE-acceptable credentials. PIC is based on ISAKMP combined with Extensible Authentication Protocol (EAP) and requires no modifications to IKE itself.

See Also: *Internet Key Exchange (IKE), Internet Protocol Security (IPSec), Just Fast Keying (JFK)*

Pretty Good Privacy (PGP)

A popular e-mail encryption technology.

Overview

Pretty Good Privacy (PGP) is a scheme developed by Phil Zimmermann for ensuring the confidentiality and integrity of e-mail messaging and secure file storage. PGP was developed at a time when export of encryption technologies was strongly controlled by the U.S. government, and PGP was released as “guerrilla freeware” to place encryption technology in the hands of ordinary users. Several legal challenges to PGP resulted but were later dismissed. Since then PGP has evolved into several forms (some of which are incompatible) that have spread around the world, including these:

- **PGP Classic (PGP versions 2.6.2 and 6.5.8):** Uses Rivest-Shamir-Adleman (RSA) and International Data Encryption Algorithm (IDEA) encryption and is freely available from a Web site of the Massachusetts Institute of Technology (MIT) for noncommercial use by U.S. and Canadian citizens only
- **PGP 8.0.x:** The current commercial version available from PGP Corporation, which purchased the assets for PGP from Network Associates
- **OpenPGP:** A family of freely available versions derived from the work of an Internet Engineering Task Force (IETF) committee
- **GNU Privacy Guard (GnuPG):** An open source version of OpenPGP available under the GNU Public License (GPL)

Implementation

PGP uses public key cryptography to generate personal long-term keys for users. Unlike traditional Public Key Infrastructure (PKI) systems that employ a hierarchical chain of certificate authorities (CAs) for issuing and verifying digital certificates, PGP requires neither digi-

tal certificates nor CAs for managing them. Instead, each user simply decides which other users to trust and then obtains the public keys for those users by any means possible; for example, by mailing the key information or swapping keys on floppy disks at a conference. Such an anarchic scheme is called a web of trust, and although it gives users complete control over who they want to engage in cryptographic communications with, the model scales poorly compared with traditional PKI systems (commercial PGP does address this issue, however). Key revocation is also performed in a similar informal fashion.

To use PGP you first download and install the software on your computer and then use it to generate a private key, which is protected using a password and then hashed for secure storage. Keys for other users you communicate with are stored in key rings, which can be either local structures or shared databases on the Internet.

For More Information

Visit web.mit.edu/network/pgp.html and www.pgp.com for more information.

See Also: *encryption, OpenPGP, Secure/Multipurpose Internet Mail Extensions (S/MIME)*

principal

The identity of an individual in the Kerberos protocol.

Overview

Kerberos identifies entities (persons and processes) using principals, which have the general form `primary/instance@realm` where the components have the following meanings:

- **Primary:** The user’s name, ID number, or some other personal identifier
- **Instance:** An optional field that allows the entity to have multiple identities, each with their own level of access control to resources
- **Realm:** The Kerberos realm of the entity, which is usually omitted when authentication is being performed in the local realm but which must be included for authentication in remote realms

Notes

In the Microsoft Windows operating system, the term **principal** also means an account holder that is automatically assigned a security identifier (SID) to control access to resources.

See Also: *Kerberos, realm*

privacy

Preventing information from being viewed by unauthorized parties.

Overview

Privacy has several contexts in regard to information security (infosec). From the perspective of consumers in the online marketplace, privacy is the control users have over the collection, use, and distribution of their personally identifiable information (PII). More generally, in online communication and telecommunication, privacy is ensuring that a message can be read only by its intended recipients, something that is also called confidentiality. Privacy in a military context is often called secrecy.

Encryption is the primary means of ensuring the privacy of information, whether stored or transmitted. End-to-end privacy ensures that an encrypted message intercepted through eavesdropping at any point along the transmission path cannot be used to compromise the privacy of the participants in the communication session.

Privacy laws are an additional means of protecting privacy by enforcing penalties for the unauthorized use, procurement, and disclosure of encrypted information. A privacy policy is an organization's policy outlining the procedures it employs for protecting PII obtained from users.

See Also: *encryption, infosec, personally identifiable information (PII), privacy policy*

Privacy Enhanced Mail (PEM)

A scheme for ensuring the privacy of e-mail sent over the Internet.

Overview

Privacy Enhanced Mail (PEM) was one of the earliest schemes developed for encrypting e-mail communication, and it is defined in RFCs 1421 through 1424. PEM was developed when e-mail messages were text only. When Multipurpose Internet Mail Extensions (MIME) was developed to handle binary attachments, PEM became the basis of a new encryption scheme called Secure/Multipurpose Internet Mail Extensions (S/MIME). PEM was based on a Public Key Infrastructure (PKI) and used random session keys for encrypting messages.

See Also: *Secure/Multipurpose Internet Mail Extensions (S/MIME)*

privacy policy

A policy outlining the requirements an organization follows for complying with privacy regulations and directives.

Overview

A privacy policy details how an organization will collect, use, and share personally identifiable information (PII) obtained from users. A privacy policy typically includes the following:

- What type of information constitutes PII for the purposes of the policy
- The reason for collecting such information and what the organization does with it
- The conditions under which such information may be disclosed to other organizations, including law enforcement agencies
- The responsibilities for the organization to protect collected information
- The organizational structure surrounding the policy, including a privacy officer and mechanisms for enforcement and submitting concerns

See Also: *privacy, privacy statement*

privacy statement

A document summarizing the privacy policy of an organization that is published in a format and location that allows users to access it easily.

Overview

A privacy statement informs consumers of an organization's policy for the collection, use, and sharing of personally identifiable information (PII) collected from them. Privacy statements are usually prominently displayed on e-commerce sites to inform consumers of the privacy and security of their PII and how they can access and modify their PII. Privacy statements also may include information about specific technologies used to manage PII such as cookies and mass-mailing lists.

See Also: *privacy, privacy policy*

Private Communication Technology (PCT)

A security protocol developed by Microsoft for private communications over the Internet.

Overview

Private Communication Technology (PCT) was a variant of Secure Sockets Layer (SSL) version 2, a protocol developed by Netscape for secure communication of private information over insecure public networks such as the Internet. Like SSLv2, PCT operated over Transmission Control Protocol (TCP) for reliable network connections and could be used both for authenticating communication sessions and for encrypting data to ensure the privacy of such sessions. PCT also included support for digitally signing messages using a hash-based message authentication code (HMAC) to ensure the integrity of communications. PCT improved upon SSLv2 in several ways, however, including these:

- Shorter messages and a simpler round structure for faster authentication
- Greater flexibility in negotiating which encryption algorithm to use in a session
- Different keys used for signing messages rather than for encrypting them

Microsoft developed PCT in response to weaknesses in Netscape's implementation of SSLv2, but when SSLv3 was released it resolved its predecessor's weaknesses and PCT never became an Internet standard.

See Also: *Secure Sockets Layer (SSL)*

private key

A key known only to its owner in a public key cryptography system.

Overview

Each user in a public key cryptography system has two keys: a private key known only to the user and a public key available for anyone who wants to obtain it. Typically, private keys are used for the following purposes:

- Decrypting messages received from other users who encrypted them using the recipient's public key
- Digitally signing messages to prove to the recipient that the messages' integrity is intact

See Also: *key, public key, public key cryptography*

private key encryption

A term sometimes used to represent symmetric (or secret) key encryption, although this is actually a misnomer since private keys are part of public key cryptography systems.

See: *secret key encryption*

privilege escalation

Another name for elevation of privileges (EoP), a method used by attackers to gain control of a system or network.

See: *elevation of privileges (EoP)*

privileges

Another name for rights, rules governing what tasks a user can perform on a system.

See: *rights*

PRNG

Stands for pseudorandom number generator, software for generating a string of apparently random numbers or characters.

See: pseudorandom number generator (PRNG)

process table attack

A type of denial of service (DoS) attack against UNIX systems.

Overview

On some older implementations of UNIX and with certain network daemons, incoming Transmission Control Protocol (TCP) connections cause new processes to be instantiated on the host, usually by forking an existing process. If no limits are placed on the number of processes that can be created, a flood of such TCP connections can cause the process table to overflow on the target system. The result is that even root (superuser) is unable to execute any commands on the system, including commands to kill existing processes, and a hard reboot must be performed to reclaim control of the console. Most UNIX platforms are no longer vulnerable to such an attack because they monitor the number of processes to prevent the process table from becoming full like this.

See Also: denial of service (DoS)

P

promiscuous mode

A mode of operation of a network adapter in which it accepts all traffic whether directed to the local host or other hosts on the network.

Overview

In order for a network host to be used for scanning all traffic on a network segment, the network adapter on that host must be running in promiscuous mode. This causes the adapter to pick up any frames traveling on the network and pass them up the Transmission Control Protocol/Internet Protocol (TCP/IP) stack to the scanner for analysis. Generally, promiscuous mode is enabled (or disabled) using vendor-supplied software

for adapters, and once enabled (or disabled) it may not be possible to change it.

Marketplace

Several network security tools exist that can be used to detect adapters running in promiscuous mode on a network. Ifstatus from Cymru.com and proDETECT from SecuriTeam are two free products for “sniffing out sniffers” on a network.

See Also: protocol analyzer

Protected Extensible Authentication Protocol (PEAP)

An authentication protocol developed by Cisco for wireless networking.

Overview

Protected Extensible Authentication Protocol (PEAP) is an 802.1x wireless networking authentication protocol that uses passwords to authenticate clients and certificates to authenticate servers. On the server side PEAP uses Extensible Authentication Protocol—Transport Layer Security (EAP-TLS), while client authentication is designed to support legacy password-based authentication protocols and also one-time passwords (OTPs). PEAP is based on an Internet-Draft standard that Cisco, together with Microsoft and RSA Security, has submitted to the Internet Engineering Task Force (IETF). PEAP is supported by Cisco’s Aironet line of wireless network access point products.

See Also: 802.1x, Extensible Authentication Protocol (EAP), Extensible Authentication Protocol—Transport Layer Security (EAP-TLS)

protocol analyzer

A device or software for displaying information contained in packets traveling on a network.

Overview

Protocol analyzers are network troubleshooting tools that provide a detailed view of network traffic. They are

frequently used for troubleshooting mysterious network problems when other troubleshooting tools have failed. Protocol analyzers work like telephone taps because they allow the user to monitor “conversations” on a network, and like most security tools they can be used for good or ill purposes. Crackers often use “sniffers” (a common term for software-based protocol analyzers) for secretly monitoring networks in order to obtain useful information such as passwords.

Using protocol analyzers effectively, whether for cracking or troubleshooting purposes, generally requires deep understanding of how the various protocols of the Transmission Control Protocol/Internet Protocol (TCP/IP) suite work. For example, some analyzers provide network information in a raw state by dumping the hexadecimal representation of each byte in a packet. Some analyzers translate such information into American Standard Code for Information Interchange (ASCII) for easier reading of data payloads and identify different fields in header bytes using labels for each field. Many analyzers support various query options for filtering different types of traffic to find the type of information the user is looking for more easily.

Marketplace

There are protocol analyzers available on the market for every kind of networking technology around, including Ethernet, wireless, Asynchronous Transfer Mode (ATM), Fiber Channel, Small Computer System Interface (SCSI), and serial connections. A market leader in commercial products is Network Associates with its Sniffer line of analyzer tools. Other vendors of protocol analyzers include Frontline Test Equipment, LANSleuth, and Logix Communications. Some popular free protocol analyzers include Ethereal, EtherPeek, Snort, Tcpdump, and Windump. Microsoft Windows 2000 includes a limited-feature version of Network Monitor, a protocol analyzer whose full version is part of System Management Server (SMS), a distributed systems management platform from Microsoft.

Notes

Other names commonly used for protocol analyzers are **network analyzer**, **monitor**, and **packet sniffer**. The term **Sniffer** is actually a trademark of Network Associates, but **sniffer** and **sniffing** are as widely used

throughout the networking and security communities as **Kleenex** is for facial tissue.

See Also: *sniffing, Tcpdump*

pseudorandom number generator (PRNG)

Software for generating a string of apparently random numbers or characters.

Overview

Pseudorandom number generators (PRNGs) are fundamental to many cryptographic operations, including creating seeds for block ciphers and generating keys from passwords. The string of numbers generated by a PRNG is not actually random, however, since PRNGs operate using deterministic mathematical formulas. Based on some initial seed, a complicated mathematical process is typically used to generate a string of numbers that has the same statistical distribution as perfectly random numbers generated from a natural process such as radioactive decay. The hitch is that if the same seed is reused sometime later, the identical string of “random” digits is produced by the PRNG, which is why the word **pseudo** (meaning “false in appearance”) is used to describe the process.

In order for pseudorandom numbers to be truly random, a different random seed must be specified each time the PRNG is used. A typical way to generate such a seed on a computer system is to combine together information from several real-time sources including the internal clock, the location of the mouse pointer, the size of a currently open file on the hard drive, and so on. This information is then hashed using a one-way function to create a fixed-length seed that is entered into the PRNG and used to generate a unique and unpredictable string of (essentially) random numbers.

See Also: *cryptology, encryption algorithm*

public key

A key known to everyone in a public key cryptography system.

Overview

Each user in a public key cryptography system has two keys: a private key known only to the user and a public key available for anyone who wants to obtain it. Typically, public keys are used for the following purposes:

- Encrypting messages sent to other users who then decrypt them using their own private key
- Verifying digital signatures attached to messages to prove that the messages' integrity is intact

See Also: *key, public key, public key cryptography*

public key cryptography

An encryption scheme that allows private communications to take place without prior existence of a shared secret.

Overview

Traditional or secret key cryptography relies on the existence of a shared secret known by the parties involved. Secret key cryptography is highly private, but the weakness in this system is securely exchanging this secret between the two parties, a necessary step before encrypted communications can be performed between parties and a difficult task to complete. Public key cryptography solves this problem by providing a way to share a secret between two parties over an insecure public connection such as the Internet.

Several popular algorithms used for public key cryptography are these:

- **Diffie-Hellman (DH):** This was the first algorithm developed for public key cryptography.
- **Rivest-Shamir-Adleman (RSA) algorithm:** A proprietary algorithm whose patent expired in 2000, meaning that the algorithm is now in the public domain.
- **Digital Signature Algorithm (DSA):** A proprietary algorithm patented by the National Institute of Standards and Technology (NIST).

Other less-used public key algorithms include the following:

- El Gamal

- Shamir Three Pass
- Massey-Omura
- Efficient Probabilistic Public Key Encryption (EPOK)

Implementation

In public key cryptography, each user is issued a pair of keys: a public key available to anyone who requests it, and a private key known only to the user who owns it. If one of these two keys is used to encrypt a message, the other can be used to decrypt it. In order for user A to encrypt a message and send it to user B, user A could first obtain user B's public key (which is readily available somewhere) and encrypts the message using this key. User A then sends the encrypted message to user B, who decrypts it using its own private key (since user B's private key can undo whatever user B's public key has done).

Although public key systems could be used that way for encrypting communication between users, in practice they are not used this way. Instead, public key cryptography is generally used to securely exchange a session key (a secret key used only for a single communication session and then discarded) between the two users, and once both parties have this session key they can use it to encrypt and decrypt messages sent between them. This is done because public keys are much longer (1024 bits or more) than secret keys (56 to 256 bits, typically), and public key algorithms are more complex than secret key algorithms, so exchanging a session key like this makes communications faster and more efficient than using public key cryptography alone.

Public key cryptography also is used sometimes for another purpose: signing messages to confirm that they have not been tampered with in transit. Digital signatures are used to verify the integrity of electronic messages, and signing a message using public key cryptography is the opposite of encrypting the message. If user A wants to sign a message and send it to user B, user A uses his or her own private key to sign the message, attaches the signature to the message, and sends it to user B. When user B receives the signed message, he or she separates the signature from the message and verifies it using user A's public key, which can be

obtained from public sources. Although this system works, it can be computationally expensive, and in practice hash-based message authentication codes (HMACs), a form of keyed message digest (MD), are often used instead for signing digital messages.

See Also: Diffie-Hellman (DH), Digital Signature Algorithm (DSA), hash-based message authentication code (HMAC), private key, private key, Rivest-Shamir-Adleman (RSA), secret key encryption

public key cryptography standards (PKCS)

A series of specifications for Public Key Infrastructure (PKI) implementation.

Overview

Public key cryptography standards (PKCS) are a series of formal and de facto standards developed by RSA Security in conjunction with industry members and academia. The standards outline various protocols and specifications for implementing various aspects of PKI. Several standards have been incorporated into other industry standards, including Public-Key Infrastructure (X.509) (PKIX), Secure/Multipurpose Internet Mail Extensions (S/MIME), Secure Sockets Layer (SSL), and the X9 standards from the American National Standards Institute (ANSI).

The PKCS list includes the following:

- **PKCS #1:** RSA Cryptography Standard
- **PKCS #3:** Diffie-Hellman Key Agreement Standard
- **PKCS #5:** Password-Based Cryptography Standard
- **PKCS #6:** Extended-Certificate Syntax Standard
- **PKCS #7:** Cryptographic Message Syntax Standard
- **PKCS #8:** Private-Key Information Syntax Standard
- **PKCS #9:** Selected Attribute Types
- **PKCS #10:** Certification Request Syntax Standard
- **PKCS #11:** Cryptographic Token Interface Standard
- **PKCS #12:** Personal Information Exchange Syntax Standard

- **PKCS #13:** Elliptic Curve Cryptography Standard
- **PKCS #15:** Cryptographic Token Information Format Standard

The PKCS #2 and #4 omitted from the preceding list have been incorporated into PKCS #1.

See Also: public key cryptography, Public Key Infrastructure (PKI), Secure/Multipurpose Internet Mail Extensions (S/MIME)

public key encryption

Another name for public key cryptography, an encryption scheme that allows private communications to take place without prior existence of a shared secret.

See: public key cryptography

Public Key Infrastructure (PKI)

A set of technologies and policies for authenticating entities using public key cryptography.

Overview

Public Key Infrastructure (PKI) represents a set of standards, policies, software, and procedures for implementing strong authentication using public key cryptography in the marketplace. PKI enables digital signatures to be used for business-to-business (B2B) and business-to-consumer (B2C) transactions, including online banking, e-commerce, and in other places where verification of identity in electronic transactions is essential. PKI enables organizations to issue, manage, validate, and revoke digital certificates used for authenticating individuals, organizations, applications, and business processes. PKI also enables organizations to issue and manage public keys used for validating digital signatures.

PKI in itself does not provide encryption for electronic messaging, only verification of identity through certificate-based authentication. PKI also does not implement any form of access control for securing access to resources in an enterprise or network. As a result, PKI is usually combined with other technologies, including Diffie-Hellman (DH) key exchange for secure sharing of session keys and permissions-based access control

for secure management of shared resources. PKI is used in several protocols, including Secure/Multipurpose Internet Mail Extensions (S/MIME), Secure Sockets Layer (SSL), Pretty Good Privacy (PGP), and other secure communication systems.

Implementation

The main elements of a typical PKI implementation include the following:

- **Certificate authority (CA):** A trusted entity that issues digital certificates to entities (users, applications, or organizations) when enrollment is requested. The CA signs the certificates it issues using its own private key to guarantee certificate authenticity. An organization can have a single CA or, for greater scalability, several CAs arranged in a hierarchy of trust with a root CA at the top. CAs can also be arranged in mesh topology to form more complex certification paths, but this is less common.
- **Certificate store:** A central database or directory of certificates issued and maintained by a CA. Certificates expire after a period of time, and, if they are lost or compromised, they can be revoked either by adding them to a certificate revocation list (CRL) or by using the Online Certificate Status Protocol (OCSP).
- **Digital certificate:** Encrypted information that guarantees that an encryption key belongs to a user.
- **Registration authority (RA):** A trusted entity that validates requests for digital certificates and forwards these requests to a CA. In some cases the role of the RA is incorporated into that of the CA.

Marketplace

Major vendors of PKI technologies and services include VeriSign, Entrust, RSA Security, and Baltimore Technologies. Software vendors such as Microsoft Corporation and Sun Microsystems include components in their operating system platforms for building PKI systems. Major system integrators such as EDS and IBM

also provide products and services for implementing PKI systems in the enterprise.

See Also: CA hierarchy, certificate authority (CA), certificate store, digital certificate, Online Certificate Status Protocol (OCSP), public key cryptography, registration authority (RA), root CA

Public-Key Infrastructure (X.509) (PKIX)

A set of standards for implementing an X.509-based Public Key Infrastructure (PKI).

Overview

The X.509 standard from the International Telecommunications Union (ITU) defines the format used to identify entities in an X.500 directory. Although the X.500 directory standard was never widely implemented because of its complexity, the X.509 standard has gained widespread usage as a standard format for digital certificates used in Public Key Infrastructure (PKI) systems. Public-Key Infrastructure (X.509) (PKIX) represents the efforts of an Internet Engineering Task Force (IETF) working group to implement X.509 in PKI, and its work profiles which X.509 options should be supported in X.509-based PKI systems, including a series of extensions for X.509 version 3.

See Also: Public-Key Infrastructure (X.509) (PKIX)

Publius Project

A system for anonymously publishing information on the World Wide Web.

Overview

Publius was designed to circumvent mechanisms used for censoring access to certain forms of content on the Web. When a user publishes content using Publius, the content is published on Web servers in random-looking form designed to evade detection by firewalls configured to block certain types of content. By connecting to a Publius proxy server (or by downloading and installing a proxy client on their machines), users can access such random-looking content and display it in its original form.

The Publius Project is managed by the Center for Democracy & Technology (CDT), a civil liberties organization. The name of the system originates from the pen name used by Alexander Hamilton, John Jay, and James Madison, the authors of the Federalist Papers that were influential in helping to ratify the U.S. Constitution.

For More Information

Visit *publius.cdt.org* for more information.

See Also: *firewall, Peekabooby Project, privacy*

Pulist

A utility for displaying running processes on machines running Microsoft Windows NT or later versions of the operating system.

Overview

Pulist is a command-line tool included in the Windows NT and Microsoft Windows 2000 Resource Kits that displays the process name and process ID for running processes on local or remote computers. Pulist can also display the security context in which each process runs by displaying the user name associated with the process, provided the caller has sufficient system rights to obtain this information from the target machine. Pulist can be used to detect rogue processes such as Trojans and backdoors running on compromised machines, though some rootkits can hide such processes to make them invisible to Pulist.

See Also: *backdoor, rootkit, Tlist, Trojan*

Pwdump

A utility for dumping passwords on machines running Microsoft Windows NT or later versions of the operating system.

Overview

The original Pwdump is a command-line tool used for dumping the password database stored in the registry of machines running Windows NT. Pwdump extracts the user name, relative ID (RID), LANMAN password hash, and NTLM password hash for each user account

and dumps the information into a text file using a UNIX-style password file format. *Windows NT 4.0 Service Pack 2* enhanced the security of the SAM database using SYSKEY encryption, which prevented the Pwdump exploit from working. However, a new version called Pwdump2 soon developed that was able to overcome SYSKEY by using dynamic-link library (DLL) injection, an unsupported method for circumventing access control mechanisms on Microsoft Windows platforms that works by injecting additional code into the address space used by the Local Security Authority Subsystem (Lsass.exe). A third version, called Pwdump3, later developed and could extract password hashes remotely using Server Message Block (SMB) protocol. To utilize Pwdump2 or Pwdump3, however, the attacker must first gain Administrator credentials on the target machine.

See Also: *password cracking*

PWL file

A password file on Microsoft Windows 95, Windows 98, and Windows Millennium Edition (Windows Me).

Overview

PWL files, which are files having *.pwl as their file extension, were used for caching passwords on legacy Windows systems. They were found in the \Windows directory and contained passwords for dial-up networking, share-level security, and logon credentials for Windows NT and NetWare networks. Passwords stored in PWL files on machines running Windows 95 were weakly encrypted and easy to crack. Windows 95 OEM Service Release 2 (OEM2) strengthened the encryption algorithm, but an exploit called Pwltool was soon developed to crack such files using a dictionary attack. Because of the vulnerability of PWL files, it's best to disable password caching on legacy Windows platforms; this can be done by editing the DisablePwdCaching setting in the registry.

See Also: *password cracking*

Q

Qchain

A tool from Microsoft Corporation to simplify the task of installing multiple hotfixes.

Overview

A **hotfix** is a security patch for a Microsoft product, such as Microsoft Windows XP or Microsoft Exchange Server. Hotfixes are released whenever new vulnerabilities and bugs are discovered in products, and most hotfixes require the system to be rebooted after they are installed. To speed up the task of installing multiple hotfixes and to reduce system downtime, Microsoft created a command-line tool called Qchain. Using this tool, you can “chain together” multiple hotfixes, installing them sequentially, and then run Qchain afterward and reboot your system only once instead of multiple times.

Qchain is included in the Resource Kits for Windows NT 4 and Windows 2000 operating systems and is also available from the Microsoft Download Center (www.microsoft.com/downloads/). There are certain limitations in using Qchain, which are described in article #296861 in the Knowledge Base on the Microsoft Product Support Services (PSS) Web site (www.support.microsoft.com). The functionality of Qchain is also built into hotfixes for Windows XP, Windows Server 2003, and Windows 2000 after Service Pack 3.

See Also: *hotfix*

QFE

Stands for Quick Fix Engineering, a Microsoft team that produces critical updates, and by extension the names of the updates themselves.

See: *Quick Fix Engineering (QFE)*

Qfecheck

A tool from Microsoft for enumerating installed hotfixes on a system.

Overview

When hotfixes are released to patch vulnerabilities and bugs discovered in Microsoft products, these hotfixes must be applied to ensure the security of systems. Managing hotfixes can be lots of work, however, and to simplify this task Microsoft Corporation created a command-line tool called Qfecheck that inspects the registry on systems running Microsoft Windows XP and Windows 2000 to determine which hotfixes are already installed and list them by their article numbers from the Microsoft Knowledge Base.

Qfecheck replaces an earlier tool called Qfechkup, the Update Information Tool, which was first released with *Windows 95 Service Pack 1*. A more recent tool called HFNetChk, included as part of the Microsoft Baseline Security Analyzer (MBSA), can check installed hotfixes on both local and remote systems.

For More Information

Visit www.microsoft.com/downloads/ to download the Microsoft Baseline Security Analyzer (MBSA).

See Also: *HFNetChk, hotfix, Microsoft Baseline Security Analyzer (MBSA)*

Queso

A tool for fingerprinting a target system.

Overview

Queso can be used to determine the operating system running on a remote host. Queso works by sending Transmission Control Protocol (TCP) packets with

Q

special flags set to a specified port on the host. Queso then analyzes the response from the host by comparing the response to an internal table of possible responses for different operating system platforms and versions. By default, Queso sends packets to port 80, the standard port for Hypertext Transfer Protocol (HTTP) traffic, but this is configurable to any TCP port number.

Queso is not a port scanner but a tool for active stack fingerprinting. Nmap, a popular port scanner, has similar fingerprinting capability in addition to its many other features.

For More Information

Visit packetstormsecurity.nl/UNIX/scanners/ to obtain Queso and other tools.

See Also: *fingerprinting, Nmap*

Quick Fix Engineering (QFE)

A Microsoft team that produces critical updates for a product, and by extension the name of the updates themselves.

Overview

Quick Fix Engineering (QFE) updates (or QFE fixes or QFEs) are updates issued by Microsoft Corporation to patch critical vulnerabilities found in Microsoft products. The term **QFE** is essentially a synonym for the more commonly used term **hotfix**. Some QFEs, however, are available only through special channels for enterprise customers who need them. QFE teams exist for each family of Microsoft products such as Microsoft Exchange Server and Microsoft Systems Management Server. Some QFE teams are now called Sustained Engineering Teams instead.

See Also: *hotfix, service pack (SP)*

R

RA

Stands for registration authority, a trusted entity that acts as an intermediary between entities requesting digital certificates and a certificate authority (CA) issuing them.

See: registration authority (RA)

race condition

A condition in which an application or system tries to perform two operations at the same time.

Overview

A race condition is when a system on which operations must normally be performed sequentially tries to perform such operations simultaneously. For example, trying to both read and write the same file at the same time creates a race condition, as does simultaneous attempts to gain control of a channel for communicating over a network. Race conditions are generally bad for two reasons:

- They can result in errors being generated that can crash or hang the system.
- They can provide attackers with a brief window in which the system is vulnerable to compromise, allowing the attacker to execute a privileged operation that he or she normally could not perform.

An example of a race condition is a UNIX vulnerability discovered in 1996 that affected how UNIX used signals to handle asynchronous communication between processes, and an exploit was developed that used File Transfer Protocol (FTP) to compromise affected servers. More common exploits involving race conditions are denial of service (DoS) attacks such as the Unicast Service Race Condition vulnerability that was identified in 2000 with regard to Microsoft Windows Media

Services 4 on Microsoft Windows NT. Race conditions are difficult to guard against when coding software, but by employing proper schemes for prioritizing process interrupts they can generally be avoided and the security vulnerabilities associated with them eliminated.

See Also: denial of service (DoS)

RADIUS

Stands for Remote Authentication Dial-In User Service, a security protocol used for centralized authentication, authorization, and accounting of network access.

See: Remote Authentication Dial-In User Service (RADIUS)

RAM slack

Unused space on hard disks that can contain data copied from random access memory (RAM), meaning physical memory.

Overview

On most computing platforms, files are written to disk using a whole number of clusters regardless of whether the entire last cluster is required or not. The last sector of the last cluster used for writing a file is generally padded with data copied from memory and is called RAM slack. Forensic examination of hard drives can display the contents of this RAM slack, which may contain user names, passwords, or other sensitive information copied from memory. Disk-cleaning programs can eliminate RAM slack from hard drives and ensure the protection of sensitive information when computer systems are lost or discarded.

See Also: computer forensics, file slack

R

RAT

Stands for remote administration tool, a program used to covertly control a remote host.

See: remote administration tool (RAT)

RBAC

Stands for role-based access control, a security model used by Sun Microsystem's Solaris platform.

See: role-based access control (RBAC)

RC2

A block cipher developed by Ron Rivest.

Overview

RC2 is a block cipher with variable key length that can be used as a replacement for Data Encryption Standard (DES) for encryption and decryption of information. It is used particularly in software exported internationally since RC2 is not subject to the same restrictions as DES is for export control. RC2 employs a block size of 64 bits and adds a random "salt" to encryption keys to reduce the possibility of certain types of cryptanalytic attack. RC2 is several times faster than DES on comparable hardware and can be made more secure than DES by choosing a suitably long key.

Notes

The **RC** in RC2 (and in RC4, RC5, and RC6) derives either from "Rivest Cipher" or "Ron's Code," depending on the source.

See Also: block cipher, Data Encryption Standard (DES), RC4, RC5, RC6

RC4

A stream cipher developed by Ron Rivest.

Overview

RC4 is a stream cipher with variable key length that uses an algorithm based on random permutations. RC4 performs encryption by generating a one-time pad and XORing it bitwise with the stream of plaintext to create a corresponding stream of ciphertext. The RC4 algorithm was originally proprietary and belonged to RSA

Security but was "outed" in 1994, which has allowed it to be extensively analyzed by the cryptanalytic community. As a result, the cipher has been found to be secure as long as the initial portion of the one-time pad is discarded. RC4 can also be used as a pseudorandom number generator (PRNG) and is used for encrypting Hypertext Transfer Protocol (HTTP) traffic in Secure Sockets Layer (SSL) communications on the Internet.

See Also: one-time pad (OTP), RC2, RC5, RC6, Secure Sockets Layer (SSL), stream cipher

RC5

A fast block cipher developed by Ron Rivest.

Overview

RC5 is a block cipher that is parameterized to support the following:

- A variable key length from 0 to 2048 bits
- A variable block size from 32 to 128 bits
- A variable number of rounds from 0 to 255

RC5 uses key expansion to generate a key table that is then used for encrypting plaintext and decrypting the corresponding ciphertext. RC5 encryption employs a scheme of bitwise XOR, integer addition, and data-dependent rotation operations as part of its encryption process. The security of RC5 is somewhat weakened since a 64-bit RC5 key was cracked in 2002 using idle central-processing unit (CPU) cycles on 331,000 computers across the Internet, though 128-bit RC5 keys are currently viewed as impermeable to brute-force cracking.

See Also: block cipher, RC2, RC4, RC6

RC6

A fast block cipher developed by Ron Rivest with Ray Sidney and Yiqun Yin.

Overview

RC6 is a modified version of RC5 that includes the additional operations of integer multiplication and 32-bit working registers. RC6 was proposed as a candidate for the Advanced Encryption Standard (AES), but

Rijndael was the winning finalist for the competition. RC6 is patented and owned by RSA Security.

For More Information

Visit Rivest's MIT page at theory.lcs.mit.edu/~rivest/ for several publications on RC6.

See Also: *Advanced Encryption Standard (AES), block cipher, RC2, RC4, RC5, Rijndael*

realm

A network served by a single Kerberos database and group of key distribution centers (KDCs).

Overview

To ensure that Kerberos authentication can scale to arbitrarily large networks, realms can be created to partition the network into different portions, each under the authority of a group of KDCs sharing the same Kerberos database. In popular Kerberos systems such as the Microsoft Active Directory directory service, such realms are identified using Domain Name System (DNS) names, but some Kerberos systems use X.500 names instead. A user in one realm can be authenticated by a different realm using cross-realm authentication. In Kerberos V4 such authentication could take place directly only between two realms that trust each other, but in Kerberos V5 a chain of trust can be established in order to traverse a hierarchy of realms and authenticate a user.

See Also: *Kerberos, key distribution center (KDC)*

recognizable plaintext attack

Another name for a ciphertext-only attack, a cryptanalytic attack in which the attacker has only ciphertext to work with.

See: *ciphertext-only attack*

reconnaissance

A cracking term describing preliminary activities in preparation for compromising a target.

Overview

Reconnaissance is a term sometimes used to describe activities more commonly known as **footprinting**, the

process of gathering as much information as possible about the network from publicly available sources, and **enumeration**, different methods used for gathering information that reveals poorly protected network resources that can be exploited for breaking into networks. Types of activities usually included under the umbrella term **reconnaissance** include port scans, ping sweeps, Whois lookups, and other methods.

Common network reconnaissance tools used by attackers include Hping, MingSweeper, Nmap, and many others. One way to think of network reconnaissance is that someone is “rattling the door handles” of your network’s security perimeter, and intrusion detection systems (IDSs) are generally able to detect when a network is being “reconned” preliminary to an all-out attack. Legitimate network reconnaissance for the purposes of evaluating the security of your network defenses is commonly called penetration testing.

See Also: *enumeration, footprinting, intrusion detection system (IDS), Nmap, penetration testing, ping sweep, port scanning*

recovery agent

A designated user that can decrypt encrypted files when a private key is lost.

Overview

The Encrypting File System (EFS) in Microsoft Windows 2000 and later versions can be used to encrypt data files using private keys, each owned by a user. Should users lose their private keys through hardware failure, system compromise, or some other action, the encrypted data would be forever lost unless a recovery agent is available. By default, the local Administrator account is the default recovery agent in a workgroup scenario and the domain administrator in an Active Directory environment. Additional recovery agents can be delegated if desired, and the private key of the recovery agent should also be exported and stored in a safe location to ensure recovery of encrypted data if needed. By default, the recovery policy for EFS forces a recovery agent to be designated before encryption of data can be performed.

Notes

Note that the EFS recovery agent does not actually have access to the private keys of each user, which would be key escrow, the process of providing a trusted third party with copies of cryptographic keys. Instead, the recovery agent can access only the randomly generated keys used to encrypt individual files. The public/private EFS keys for users are not used for actual file encryption but simply for protection of these random file encryption keys.

See Also: *Encrypting File System (EFS), public key, private key*

Recovery Console

A tool for troubleshooting startup problems on Microsoft Windows 2000 or later versions of the operating system.

Overview

The Recovery Console provides limited command-line access to machines running Windows for troubleshooting startup problems caused by missing or corrupt system files and device drivers. Using the Recovery Console, an administrator can enable or disable services, format drives, and read, delete, copy, or move files on NTFS volumes. The Recovery Console can be started two ways:

- By installing it on your system and selecting the Recovery Console option from the boot loader menu at startup
- By booting your system from the Windows product CD and selecting the Recovery Console option

Since the Recovery Console provides an alternative method for logging on to a system, there are potential security issues associated with its use. For example, if a machine running Windows XP is booted from a Windows 2000 CD, the Windows 2000 Recovery Console can be started without a password to gain access to selected system folders on the target machine. While this may seem at first like a flaw in the design of the Recovery Console, it actually highlights the importance of physical security as a part of every network security

program, since a user who has physical access to a machine could also use a different CD to boot to some other non-Windows operating system and compromise the machine, or even install a parallel version of the operating system. One way of securing machines against this exploit is to configure the basic input/output system (BIOS) to prevent booting from a CD-ROM drive and then protect the BIOS settings by configuring a BIOS password.

See Also: *physical security*

reflection attack

A type of distributed denial of service (DDoS) attack similar to Smurf but involving spoofed source addresses.

Overview

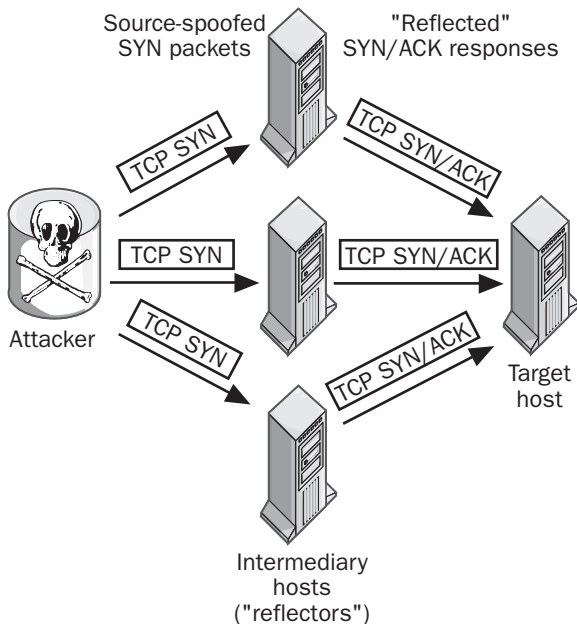
Denial of service (DoS) attacks leverage weaknesses in the Transmission Control Protocol (TCP) to allow attackers to prevent legitimate users from accessing network services on target machines. Distributed denial of service (DDoS) attacks take this one step further, amplifying DoS attacks by using multiple compromised intermediary systems called zombies. A reflection attack “reflects” spoofed TCP packets off of large numbers of intermediary hosts to the target to overwhelm it and render its services unavailable to legitimate users.

Implementation

To launch a reflection attack, the attacker uses a packet-spoofing tool to create or modify TCP packets so that they seem to originate from the target host. The particular type of TCP packet spoofed is a SYN packet, which is the initial packet in a TCP three-way handshake and represents a request to establish a TCP connection. The attacker then sends these spoofed packets to hundreds or thousands of publicly available high-bandwidth hosts on the Internet, such as powerful Border Gateway Protocol (BGP) routers, high-availability Web servers, and heavy-duty Domain Name System (DNS) servers. These streams of packets sent to intermediary hosts are kept at a sufficiently low rate that they avoid triggering any intrusion detection system (IDS) monitoring the host.

When the intermediary hosts receive these spoofed SYN packets, they interpret it as a request for a TCP connection, and they immediately respond with SYN/ACK packets that acknowledge the request and try to initiate a connection with the target host (not the attacker's machine) because of the spoofed source addresses of the packets. The target host is thus flooded with SYN/ACK packets because of the magnification factor of the attacker using hundreds or thousands of intermediary hosts to concentrate the effect.

The best solution to this attack and other DDoS attacks is generally configuring egress filtering on border routers located on Internet service provider (ISP) networks.



Reflection attack. How a reflection attack works.

Notes

The term **reflection attack** is also used to describe an attack against certain forms of mutual authentication protocols. To foil such attacks, the initiator of the authentication session should be the first one to have to prove its identity during the exchange.

For More Information

See the paper by Vern Paxson entitled "An Analysis of Using Reflectors for Distributed Denial-of-Service Attacks" in *Computer Communication Review* 31(3), July 2001.

See Also: *denial of service (DoS), distributed denial of service (DDoS), intrusion detection system (IDS), TCP three-way handshake*

Regdmp

A Microsoft Windows 2000 Resource Kit utility for dumping registry information.

Overview

Regdmp can be used for troubleshooting registry problems by dumping portions of the registry to standard output (screen) or redirecting output to a file. The tool also can be used for exporting registry information that then can be modified and used by Regini, another Resource Kit tool, for scripted batch copying of registry information into target systems. Like most security tools, however, Regini can also be used for malicious purposes. Specifically, an intruder who has compromised a machine running a version of the Microsoft Windows operating system may be able to use Regdmp for registry enumeration to find additional user and system configuration information.

Locking down registry permissions can help prevent such use, and default permissions restrict remote registry access to administrators only. DumpSec from Somarsoft (www.somarsoft.com) is a registry-auditing tool that can be used to identify possible security vulnerabilities in your registry permissions.

See Also: *enumeration*

registration authority (RA)

A trusted entity that acts as an intermediary between entities requesting digital certificates and a certificate authority (CA) issuing them.

Overview

Registration authorities (RAs) are used in large Public Key Infrastructure (PKI) systems to enhance scalability

by providing an extra layer of functionality between users and CAs. Typically, an RA receives a certificate request submitted by a user, usually a PKCS #10 or X.509 type of request. Once the RA receives a request, it verifies the identity of the user, and then forwards the request to the CA for processing. The CA then issues the certificate and returns it to the RA, which forwards it to the user who requested it. All communication between the RA and the CA is encrypted and signed to ensure maximum protection for the user's identity information. RAs are thus involved in the registration activity of requesting certificates and are usually deployed when large numbers of certificates need to be issued at remote locations where it is impractical to deploy and manage a CA for these purposes.

The term **RA** applies generally to the collection of processes, tools, and personnel needed to perform this kind of function. RAs also may be involved in the revocation of certificates and maintaining certificate revocation lists (CRLs) of lost or compromised certificates.

See Also: *certificate authority (CA), digital certificate, Public Key Infrastructure (PKI)*

regression testing

A quality control technique for ensuring the reliability and security of software.

Overview

Regression testing is performed whenever an operating system or application has been updated by creating a hotfix or patch to address some flaw or vulnerability discovered in the original platform. Regression testing involves comparing every aspect of the functionality of the program before and after applying the patch to make sure that fixing one problem hasn't created a new issue or vulnerability. Regression testing is also performed to ensure new features added to a product don't conflict with the operation of existing features. Regression testing is an essential part of ensuring the security of products and is used by Microsoft and other vendors in testing some patches for products in which vulnerabilities have been identified. Many security patches released by vendors are not regression tested, and

administrators should exercise caution when introducing such patches.

See Also: *hotfix, patch*

remote administration tool (RAT)

Also called a remote administration Trojan, a program used to covertly control a remote host.

Overview

Remote administration tools (RATs) are Trojans that are covertly installed on target systems and used by attackers to secretly control the systems from a remote location. One of the most famous RATs was Back Orifice, a powerful Trojan for the Microsoft Windows 95 and Windows 98 platforms that was developed by Cult of the Dead Cow (cDc) and released at Defcon 6 in 1998. This was soon followed by Back Orifice 2000 (BO2K), which affects the Windows NT and Windows 2000 platforms and can be used either legitimately as a remote administration tool or maliciously as a tool for monitoring and controlling compromised systems. Several other RATs such as Netbus and Hack'a'tack soon followed, including the notoriously powerful and easy-to-use SubSeven Trojan, which appeared in February 1999 and has become the most popular RAT by far among the black hat community.

All of these tools are client/server applications that require some vulnerability to be exploited on the target host so that the server component of the RAT can be secretly installed. Most of them can have their executable files renamed and can be configured to operate on any port number, making their presence difficult to detect using intrusion detection tools. Most of them also can be used as legitimate tools for remotely administering network servers.

See Also: *Back Orifice, Back Orifice 2000 (BO2K), Netbus, SubSeven, Trojan*

Remote Authentication Dial-In User Service (RADIUS)

A security protocol used for centralized authentication, authorization, and accounting of network access.

Overview

Remote Authentication Dial-In User Service (RADIUS) was originally defined in RFCs 2865 and 2866 as a method for authenticating dial-up remote access.

RADIUS is now used for a variety of authentication purposes, however, including authenticating virtual private network (VPN) servers, wireless access points, Ethernet switches, and other platforms.

Implementation

In a typical remote access scenario, a remote client uses Point-to-Point Protocol (PPP) to submit a user's credentials to a remote access server, which is also a RADIUS client. The RADIUS client then forwards the user's credentials and connection establishment information to a RADIUS server, which authenticates and authorizes the request and returns a response to the RADIUS client (the access server), which then accepts the connection attempt from the remote client. Some RADIUS clients also send accounting messages to RADIUS servers to track usage of the access server for billing purposes. Messages between RADIUS clients and servers use User Datagram Protocol (UDP) port 1812 for RADIUS authentication messages and port 1813 for accounting messages.

RADIUS security is provided by a shared secret between the RADIUS client and server, which is used to encrypt messages sent between the client and server. If the secret is weak, an eavesdropper might be able to crack the messages and hijack authentication sessions for malicious purposes. By default, messages sent by RADIUS clients to RADIUS servers are not cryptographically verified, which allows the source addresses of RADIUS messages to be spoofed easily. However, most implementations of RADIUS optionally can be configured to use message digest 5 (MD5) for cryptographically verifying RADIUS messages.

See Also: *authentication, Authentication, Authorization, and Accounting (AAA)*

replay attack

Also called packet replay, an attack based on capturing and resending packets on a network.

See: *packet replay*

repudiation

The ability of a user to deny having performed an action.

Overview

Repudiation is the process whereby a user denies performing an action and other parties involved cannot prove otherwise. For example, if file system auditing is not configured on a system, a user who deletes a file can repudiate (deny) that he or she was the one who did it and no one can prove otherwise. Secure systems generally require a high degree of **nonrepudiation**, the ability to establish the identity of the one who performed an action such as sending a message, deleting a file, or rebooting a system.

See Also: *nonrepudiation*

resource exhaustion attack

Denial of service (DoS) by “starving” a system resource.

Overview

Resource exhaustion (or resource starvation) is a form of DoS attack in which the attacker uses up a resource on the target system, with the result that no resources are available for legitimate users trying to access the system. Examples of types of resources that can be “starved” include central processing unit (CPU) cycles, memory (physical or virtual), network bandwidth, disk space, disk quota, file handles, processes, and threads. An example of a resource exhaustion attack is the Stream.c exploit, which drives up a system's CPU usage by sending streams of malformed Transmission Control Protocol (TCP) ACK packets having random source addresses and sequence numbers to a series of ports on the target machine. A similar exploit is Raped.c, which uses spoofed source addresses instead and has the same effect. These exploits affect a variety of platforms, both UNIX/Linux and Microsoft Windows operating systems, but can be prevented or mitigated by keeping systems up to date with patches issued by vendors.

See Also: *denial of service (DoS)*

restrictive shell

A command shell that limits what users can do.

Overview

Restrictive shells are used on UNIX systems to provide secure environments for users to perform necessary tasks while preventing them from executing commands that could affect the larger system environment. For example, a restrictive shell could allow users to run some programs but not others, work in some directories but not others, and so on. Restrictive shells are often used when users have to run a shell on a host over the insecure Internet to limit the damage that could occur should the session be hijacked.

If restrictive shells are implemented poorly, however, it may be possible for users (or attackers who hijacked the shell) to break out of the shell and run programs to which they normally would not have access. One way of doing this is if the user can start a program that itself includes a shell function. Another way of circumventing shell restrictions is to export files using Ftp or some other command to perform manipulations on them that normally are forbidden.

See Also: *hijacking*

reverse Telnet

Initiating a Telnet session from the host instead of the client.

Overview

Reverse Telnet (sometimes called **direct Telnet**) is a technique sometimes used by administrators for troubleshooting remote hosts, particularly for configuring routers and access servers. However, attackers also can use it to try to compromise a vulnerable system by obtaining an interactive shell for running commands on the system. In a typical scenario, a Web server vulnerable to a malformed Uniform Resource Locator (URL) exploit might be running behind a firewall that blocks all ports except ports 80 and 443, the standard and secure ports for Hypertext Transfer Protocol (HTTP) traffic. The attacker first opens two Netcat windows on the attacker's own machine, one listening to port 80 and the other to port 443. The attacker then performs an

exploit by sending a URL that starts the Telnet client on the target machine to connect to one Netcat window and pipe any commands typed there into a shell running on the other Netcat window, creating a back channel that can be used to run arbitrary commands on the target. Similar exploits can be performed using Xterm or Nc if they are running on the target host and no Telnet client is available on the target.

See Also: *malformed URL attack, Netcat*

reversible encryption

Any form of encryption that also can be decrypted.

Overview

Encryption algorithms may be either one-way or reversible, depending on how they are mathematically designed. Reversible algorithms generally are used for encrypted communications, since you want the receiving party to be able to decrypt encrypted messages when they are received. One-way encryption, which is also called hashing, is used in other contexts such as secure storage of passwords and creating message digests. By storing hashed values of passwords instead of reversibly encrypted versions, intruders are prevented from obtaining actual passwords even if they compromise a system and obtain a copy of the password files. Hashed passwords are also secure from eavesdropping attacks in which an intruder uses a packet sniffer to capture and analyze authentication session traffic. Some authentication schemes such as Challenge Handshake Authentication Protocol (CHAP) and Digest Authentication require reversibly encrypted passwords, however, and Microsoft Windows 2000 operating system supports a password policy setting that can enable reversible encryption for these purposes. Enabling reversible encryption of passwords is not recommended, however, since reversibly encrypted passwords are essentially as vulnerable as plaintext passwords to eavesdropping attacks.

See Also: *authentication, Challenge Handshake Authentication Protocol (CHAP), Digest Authentication, encryption, hashing algorithm, password, password policy*

Rexec

A UNIX utility for executing a command on a remote host.

Overview

Rexec, which stands for “remote execute,” is part of the r-command package of UNIX tools that also includes Rlogin and Rsh. Using the Rexec command on a client machine, a user can run a command on a remote server on which the Rexec daemon (service) is running. Rexec is unable to run most interactive commands, however, so you can’t use it to run Vi or Emacs on a remote host (use Telnet instead for this purpose). Rexec is similar in function to Rsh, but Rexec prompts for a user name and password to be sent to the remote host, while Rsh doesn’t require this. Like other UNIX r-commands, Rexec is not considered secure and now is generally replaced by such tools as Secure Shell (SSH).

See Also: *Rlogin, Rsh, Secure Shell (SSH)*

.rhosts

A file on UNIX systems that specifies remote users who are not required to provide a login password in order to run r-commands.

Overview

The .rhosts file is a hidden file located in a user’s home directory that contains a list of entries of the form *host user*. Here, *host* specifies the name of a remote host (the full Domain Name System, or DNS, name must be used if the remote host is in a different domain). If *user* is unspecified, any user logged on to the remote host can run r-commands such as Rcp, Rlogin, or Rsh on the local host. If *user* is specified, only that user can run these commands (the user account must be the same on both the local and the remote hosts).

Since the UNIX r-commands are notoriously insecure, security best practice suggests that .rhosts should not be used. UNIX admins may even want to consider running a script that periodically clears .rhosts files from their systems since these files often are the target of attackers and can provide a way to gain entry into a system to compromise its security. If .rhosts files must be used,

permissions on them should be set to root ownership and they should not be group or world writable.

See Also: *Rlogin, Rsh*

rights

Authorization to perform an operation that affects an entire system instead of just a specific object on the system.

Overview

On Microsoft Windows platforms, user rights are divided into two categories:

- **Logon rights:** These define the access that users and other security principals have to the system, whether through interactive keyboard commands, across a network connection, as a batch job, or as a network service. The first of the following tables lists the different logon rights available on machines running the Microsoft Windows 2000 operating system.
- **Privileges:** These define which users are authorized to manipulate different types of system resources such as resetting the internal system clock, loading or unloading device drivers, backing up or restoring files, and anything else that affects the system as a whole. The second of the following tables lists the different privileges available on machines running the Windows 2000 operating system, together with the local groups and special identities that are assigned these privileges by default (where applicable).

Rights are different from **permissions**, which are granted by the owner of an object to define who can access the object and what level of access that user can have. User rights are instead assigned through security policy, either Local Security Policy on stand-alone machines in a workgroup environment, or Group Policy in a domain-based scenario using Active Directory directory service.

Notes

The LocalSystem special identity has almost all privileges and logon rights assigned to it by default, and processes running as part of the operating system usually

run within the context of this account since they require a complete set of user rights to have full access to system resources. Most system services included in the Windows 2000 operating system are configured auto-

matically to run as LocalSystem, but on Windows Server 2003 security has been tightened by having many services run under the less privileged Network-Service identity instead.

User rights available in Windows 2000

| <i>Right</i> | <i>Description</i> |
|---|--|
| Access this computer from network | Allows a user to connect to the computer from the network. By default, this right is assigned to Administrators, Everyone, and Power Users. |
| Log on as a batch job | Allows a user to log on by using a batch-queue facility. By default, this right is assigned to Administrators. |
| Log on locally | Allows a user to log on locally at the computer's keyboard. By default, this right is assigned to Administrators, Account Operators, Backup Operators, Print Operators, and Server Operators. |
| Log on as a service | Allows a security principal to log on as a service. Services can be configured to run under the LocalSystem account, which has a built-in right to log on as a service. Any service that runs under a separate account must be assigned the right. By default, this right is not assigned to anyone. |
| Deny access to this computer from network | Prohibits a user or group from connecting to the computer from the network. By default, no one is denied this right. |
| Deny local logon | Prohibits a user or group from logging on locally at the keyboard. By default, no one is denied this right. |
| Deny logon as a batch job | Prohibits a user or group from logging on through a batch-queue facility. By default, no one is denied the right to log on as a batch job. |
| Deny logon as a service | Prohibits a user or group from logging on as a service. By default, no one is denied the right to log on as a service. |

Privileges available in Windows 2000

| <i>Privilege</i> | <i>Description</i> |
|-------------------------------------|--|
| Act as part of the operating system | <p>Allows a process to authenticate like a user and thus gain access to the same resources as a user. Only low-level authentication services should require this privilege.</p> <p>Note that potential access is not limited to what is associated with the user by default; the calling process might request that arbitrary additional privileges be added to the access token. Note that the calling process can also build an anonymous token that does not provide a primary identity for tracking events in the audit log.</p> <p>When a service requires this privilege, configure the service to use the LocalSystem account (which already includes the privilege), rather than create a separate account and assign the privilege to it.</p> |

Privileges available in Windows 2000 (continued)

| <i>Privilege</i> | <i>Description</i> |
|---------------------------------|--|
| Add workstations to a domain | <p>Allows the user to add a computer to a specific domain. For the privilege to be effective, it must be assigned to the user as part of Local Security Policy for domain controllers in the domain. A user who has this privilege can add up to 10 workstations to the domain.</p> <p>In Windows 2000, the behavior of this privilege is duplicated by the Create Computer Objects permission for organizational units and the default Computers container in Active Directory. Users who have the Create Computer Objects permission can add an unlimited number of computers to the domain.</p> |
| Back up files and directories | <p>Allows the user to circumvent file and directory permissions to back up the system. The privilege is selected only when an application attempts access through the NTFS backup application programming interface (API). Otherwise, normal file and directory permissions apply.</p> <p>By default, this privilege is assigned to Administrators and Backup Operators.</p> |
| Bypass traverse checking | <p>Allows the user to pass through folders to which the user otherwise has no access while navigating an object path in any Microsoft Windows file system or in the registry. This privilege does not allow the user to list the contents of a folder; it allows the user only to traverse its directories.</p> <p>By default, this privilege is assigned to Administrators, Backup Operators, Power Users, Users, and Everyone.</p> |
| Change the system time | <p>Allows the user to set the time for the internal clock of the computer.</p> <p>By default, this privilege is assigned to Administrators and Power Users.</p> |
| Create a token object | <p>Allows a process to create an access token by calling <code>NtCreateToken()</code> or other token-creating APIs.</p> <p>When a process requires this privilege, use the <code>LocalSystem</code> account (which already includes the privilege), rather than create a separate user account and assign this privilege to it.</p> |
| Create permanent shared objects | <p>Allows a process to create a directory object in the Windows 2000 object manager. This privilege is useful to kernel-mode components that extend the Windows 2000 object namespace. Components that are running in kernel mode already have this privilege assigned to them; it is not necessary to assign them the privilege.</p> |
| Create a page file | <p>Allows the user to create and change the size of a page file. This is done by specifying a paging file size for a particular drive under Performance Options on the Advanced tab of System Properties.</p> <p>By default, this privilege is assigned to Administrators.</p> |
| Debug programs | <p>Allows the user to attach a debugger to any process. This privilege provides access to sensitive and critical operating system components.</p> <p>By default, this privilege is assigned to Administrators.</p> |

Privileges available in Windows 2000 (continued)

| <i>Privilege</i> | <i>Description</i> |
|--|---|
| Enable computer and user accounts to be trusted for delegation | <p>Allows the user to change the Trusted for Delegation setting on a user or computer object in Active Directory. The user or computer that is granted this privilege must also have write access to the account control flags on the object.</p> <p>Delegation of authentication is a capability that is used by multitier client/server applications. It allows a front-end service to use the credentials of a client in authenticating to a back-end service. For this to be possible, both client and server must be running under accounts that are trusted for delegation.</p> <p>Misuse of this privilege or the Trusted for Delegation settings can make the network vulnerable to sophisticated attacks that use Trojans, which impersonate incoming clients and use their credentials to gain access to network resources.</p> |
| Force shutdown from a remote system | <p>Allows a user to shut down a computer from a remote location on the network. By default, this privilege is assigned to Administrators.</p> |
| Generate security audits | <p>Allows a process to generate entries in the security log. The security log is used to trace unauthorized system access. (See also “Manage auditing and security log” in this table.)</p> |
| Increase quotas | <p>Allows a process that has Write Property access to another process to increase the processor quota that is assigned to the other process. This privilege is useful for system tuning, but it can be abused, as in a denial-of-service (DoS) attack.</p> <p>By default, this privilege is assigned to Administrators.</p> |
| Increase scheduling priority | <p>Allows a process that has Write Property access to another process to increase the execution priority of the other process. A user with this privilege can change the scheduling priority of a process in the Task Manager dialog box.</p> <p>By default, this privilege is assigned to Administrators.</p> |
| Load and unload device drivers | <p>Allows a user to install and uninstall plug and play device drivers. This privilege does not apply to device drivers that are not plug and play; these device drivers can be installed only by Administrators. Note that device drivers run as trusted (highly privileged) programs; a user can abuse this privilege by installing hostile programs and giving them destructive access to resources.</p> <p>By default, this privilege is assigned to Administrators.</p> |
| Lock pages in memory | <p>Allows a process to keep data in physical memory, which prevents the system from paging the data to virtual memory on disk. Assigning this privilege can result in significant degradation of system performance. This privilege is obsolete and is therefore never selected.</p> |
| Manage auditing and security log | <p>Allows a user to specify object access auditing options for individual resources such as files, Active Directory objects, and registry keys. Object access auditing is not actually performed unless you have enabled it in Audit Policy (under Security Settings, Local Policies). A user who has this privilege also can view and clear the security log from Event Viewer.</p> <p>By default, this privilege is assigned to Administrators.</p> |
| Modify firmware environment values | <p>Allows modification of system environment variables either by a process through an API or by a user through System Properties.</p> <p>By default, this privilege is assigned to Administrators.</p> |

Privileges available in Windows 2000 (continued)

| <i>Privilege</i> | <i>Description</i> |
|--|---|
| Profile a single process | Allows a user to run Microsoft Windows NT and Windows 2000 performance-monitoring tools to monitor the performance of nonsystem processes. By default, this privilege is assigned to Administrators and Power Users. |
| Profile system performance | Allows a user to run Windows NT and Windows 2000 performance-monitoring tools to monitor the performance of system processes. By default, this privilege is assigned to Administrators. |
| Remove computer from docking station | Allows the user of a portable computer to undock the computer by clicking Eject PC on the Start menu. By default, this privilege is assigned to Administrators, Power Users, and Users. |
| Replace a process-level token | Allows a parent process to replace the access token that is associated with a child process. |
| Restore files and directories | Allows a user to circumvent file and directory permissions when restoring backed-up files and directories and to set any valid security principal as the owner of an object. By default, this privilege is assigned to Administrators and Backup Operators. |
| Shut down the system | Allows a user to shut down the local computer. In Microsoft Windows 2000 Professional, this privilege is assigned by default to Administrators, Backup Operators, Power Users, and Users. In Microsoft Windows 2000 Server, this privilege is by default not assigned to Users; it is assigned only to Administrators, Backup Operators, and Power Users. |
| Synchronize directory service data | Allows a process to provide directory synchronization services. This privilege is relevant only on domain controllers. By default, this privilege is assigned to the Administrator and LocalSystem accounts on domain controllers. |
| Take ownership of files or other objects | Allows a user to take ownership of any securable object in the system, including Active Directory objects, files and folders, printers, registry keys, processes and threads. By default, this privilege is assigned to Administrators. |

Rijndael

The block cipher at the heart of the Advanced Encryption Standard (AES).

Overview

Rijndael (pronounced “rain doll”) was one of a number of candidates submitted to the National Institute of Standards and Technology (NIST) as a replacement for the aging (and now insecure) Data Encryption Standard (DES), the FIPS 46-9 encryption standard used for many years by the U.S. federal government. Rijndael was the winning candidate and was chosen by NIST as the basis of AES to replace DES as the encryption standard for U.S. government and industry. While Rijndael/

AES officially replaces DES, it will take some years before Rijndael/AES becomes prevalent. This is partly because of the ubiquity of DES, the cost of replacing or upgrading encryption hardware and software, and the fact that Triple-DES (3DES) also is widely used and is likely to remain uncrackable for many years. The main advantage of Rijndael/AES is that it is much faster than 3DES and has a small footprint suitable for embedded and smart card encryption hardware.

Implementation

Rijndael is a symmetric block cipher that was developed by two Belgian cryptographers, Vincent Rijmen and Joan Daemen, and its design is based on an earlier

block cipher named Square. Both the block length and key length in Rijndael are variable, with possible block lengths of 128, 192, or 256 bits and key lengths of the same three values (both lengths can be extended further in multiples of 32 bits). Like DES and the International Data Encryption Algorithm (IDEA), Rijndael uses a series of rounds to transform plaintext blocks into ciphertext, with more rounds used for bigger block sizes and larger keys (the actual number of rounds is equal to six more than the larger of the key and block sizes). Each round combines substitutions, rotations, XOR operations, and mixing of columns in the state table.

The design and operation of Rijndael is freely available and has been described by its authors in detail in the book *The Design of Rijndael*, published by Springer-Verlag. Publishing the operation of an encryption algorithm like this generally enhances its security rather than diminishes it by exposing it to peer review so that cryptanalysis everywhere can try to devise methods for cracking it.

See Also: 3DES, Advanced Encryption Standard (AES), block cipher, Data Encryption Standard (DES), International Data Encryption Algorithm (IDEA), National Institute of Standards and Technology (NIST)

Rinetd

A tool for redirecting Transmission Control Protocol (TCP) connections.

Overview

Rinetd is a UNIX utility that can be used to redirect TCP connections to different Internet Protocol (IP) addresses and/or port numbers. Rinetd can be used for port redirection to allow external users to access hosts inside a firewall that uses Network Address Translation (NAT) or IP masquerading. Like other port redirection tools such as Fpipe, Rinetd can be used for penetration testing and for circumventing firewall restrictions by mapping traffic to open ports. One type of traffic Rinetd cannot redirect, however, is File Transfer Protocol (FTP), which requires an additional socket for data

transfer. Rinetd is open source software released under the GNU Public License (GPL).

For More Information

Visit www.boutell.com/rinetd/ for more information.

See Also: Fpipe, port redirection

RIP spoofing

Forging packets for Routing Information Protocol (RIP) packets.

Overview

RIP is one of several routing protocols used to share routing table information between routers on large internetworks. RIP spoofing involves forging RIP packets to modify routing tables, typically to redirect traffic to a compromised host for eavesdropping purposes. RIP traffic is especially easy to spoof since RIPv1 uses no authentication at all, while RIPv2 authenticates RIP communications using passwords in cleartext, which can be obtained by attackers using packet-sniffing tools. To prevent such attacks you can do either of the following:

- Use static routes for all routers on your network and disable RIP entirely.
- Use static routes for your perimeter routers, block RIP traffic at the perimeter of your network using firewalls, and use RIP only within your firewalled internal network.
- Use Open Shortest Path First (OSPF) instead of RIP as your routing protocol (OSPF has additional features that make it more secure than RIP).

See Also: sniffing, spoofing

risk assessment

Identifying the potential for compromise of a protected system.

Overview

At its most basic technological level, risk assessment involves identifying the likelihood of compromise of a system. Such compromise can be the result of either internal or external threats and involves exploiting

known or potential vulnerabilities in the system's defenses. At a more general business-oriented level, risk assessment involves answering the following questions:

- What valuable resources are important to protect?
- What degree of threat to these resources is likely to exist?
- How vulnerable are the resources likely to be in the face of such threats?
- What would be the business impact of the resources being compromised?

The last question is particularly important since it bridges to the wider area of risk management, which involves comparing the value of the protected resources with the loss that would result should the resources be compromised. This value/loss equation is essentially a business decision: if it costs more to protect a resource than the value of the resource itself, the risk of the resource being compromised is an acceptable or "manageable" risk from a business perspective.

Risk assessments can be performed by companies either using internal assessment teams or by outsourcing the job to independent security-auditing organizations, which may conduct penetration testing to determine the strength of the company's defenses against common forms of attack. Risk assessment is an essential part of business planning in the Internet age because of the ubiquity of the Internet, the complexity of modern computing systems, and the prevalence of hacking and cracking activities across the Internet.

See Also: ISO 17799, penetration testing

Rivest-Shamir-Adleman (RSA)

A widely used public key cryptography algorithm.

Overview

The Rivest-Shamir-Adleman (RSA) algorithm is named after its originators, Ronald Rivest, Adi Shamir, and Leonard Adleman. The algorithm was developed at Massachusetts Institute of Technology (MIT) in 1977 and is the most popular public key encryption algorithm

in use today, with implementations in the Secure Sockets Layer (SSL) protocol, the Secure Electronic Transactions (SET) protocol, Pretty Good Privacy (PGP), and other cryptographic authentication and messaging schemes. RSA can be used for encryption and decryption of information and for the generation and verification of digital signatures. The algorithm was patented and owned by RSA Security, but the patent expired in September 2000 and the algorithm is now in the public domain.

Implementation

RSA is an asymmetric algorithm that supports both variable key length and variable block size. The RSA algorithm requires that the plaintext block being encrypted must be smaller in size than the key used to encrypt it, while the resulting ciphertext block is equal in size to the key itself. The most common key length used for RSA is 512 bits, but a key of this length was cracked in 1999 using hundreds of desktop workstations working together for several months (dedicated hardware could probably crack such a key much faster). As a result, RSA Security now recommends keys of 768 bits for ordinary users, 1024 bits for enterprises, and 2048 bits for certificate authorities (CAs). Even a 768-bit key is still considered uncrackable, though no one knows for how long this will remain so.

The operation of RSA is based on the factoring of large prime numbers, an intractable mathematical problem for which no effective solution is known other than the brute-force approach. The RSA key generation algorithm randomly generates a pair of large primes whose product has the same number of bits as the key, and then it uses modular arithmetic to generate a pair of keys, one of which is the private key and the other, the public key. These keys then are used as follows:

- To encrypt a message and send it to a second party, use the second party's public key, which is available from public sources. When the second party receives your encrypted message, that person can use his or her private key to decrypt it.
- To digitally sign a message and send it to a second party, create a digest of your message and use your

own private key to encrypt the digest to create the signature. Attach the signature to your message and send it to the second party. When the second party receives your message with attached signature, that person can use your public key (which is available from public sources) to decrypt your message digest (MD), then create his or her own MD by hashing the received message, and finally compare the two digests to verify that the signature is valid and the message hasn't been tampered with in transit.

Since RSA is a relatively slow algorithm, it is generally not used to encrypt the messages themselves. Instead, RSA is used to exchange a much smaller session key securely between the two parties, which then use the session key for encrypting messages using a symmetric cipher such as Data Encryption Standard (DES) or International Data Encryption Algorithm (IDEA). Also, when sending encrypted messages that are also signed, best practice is to use different RSA keys for encrypting the session key and for signing the message.

Issues

Because of certain weaknesses in how RSA encrypts messages that make it susceptible to certain kinds of eavesdropping attacks, a series of de facto standards called public key cryptography standards (PKCS) have been developed by RSA to ensure interoperability between different implementations of RSA and to avoid various pitfalls related to padding messages with predictable data.

Notes

Although Rivest, Shamir, and Adleman are generally credited with creating this algorithm, the method was independently discovered four years earlier by Clifford Cocks, a cryptographer working at GCHQ, the British equivalent of the National Security Agency (NSA) in the United States, but the work was kept secret at the time and revealed only in 1997.

For More Information

Visit RSA Security at www.rsasecurity.com for more information.

See Also: *public key cryptography, public key cryptography standards (PKCS), session key*

Rlogin

A UNIX command opening a terminal session with a remote host.

Overview

Rlogin, which stands for “remote login,” is part of the r-command package of UNIX tools that also includes Rexec and Rsh. Rlogin is an Internet standard protocol defined in RFC 1282. It works similarly to Telnet in that it can be used to establish a terminal session with a remote host. Although it supports a larger range of terminal environment semantics than Telnet, it is also a less secure tool since it can be configured to operate without a password being required by defining a list of trusted hosts in the user's .rhosts file. The simplicity of Rlogin resulted in its common use in early UNIX environments, but its lack of security and the ease by which it can be used by attackers for compromising systems means that nowadays the older Rlogin (and other r-commands) usually are replaced by more secure versions that support authentication and encryption, by running them within a wrapper program such as TCP Wrappers, or by using such tools as Secure Shell (SSH) instead.

See Also: *Rexec, .rhosts, Rsh, Secure Shell (SSH)*

Rnmap

An enhanced version of Nmap for centralized port scanning.

Overview

Rnmap, which stands for “Remote Nmap,” is a client/server tool that lets you perform Nmap port scans from a central server instead of from client machines. To use Rnmap, a client connects to an Rnmap server, which does the actual scanning of the remote network. Rnmap is available on SourceForge and is released as open source software under the GNU Public License (GPL). Rnmap works on a variety of UNIX and Linux platforms and is written entirely in Python.

For More Information

Visit rnmap.sourceforge.net for more information.

See Also: *Nmap, port scanning*

role-based access control (RBAC)

A security model used by the Solaris platform developed by Sun Microsystems.

Overview

The traditional security model for UNIX platforms is an “all or nothing” model in which a user either is or is not a superuser, and there is nothing in between ordinary users and the Superuser account called root. This traditional model violates the important security principle of least privilege, which says that entities (users, applications, or devices) should be assigned only the minimum privileges (rights or permissions) they need to fulfill their purposes and nothing more.

Solaris’s role-based access control (RBAC) model enables administrators to separate the rights and privileges of the Superuser account and assign them to different roles such as primary administrator, junior administrator, system administrator, or operator. These roles can then be assigned to individual users and groups to grant them rights for performing specific system tasks such as backing up servers, while at the same time preventing them from being able to perform the full range of tasks that root can.

On Solaris using RBAC, roles can be added, removed, and modified using the Roleadd, Roledel, and Rolemod commands, respectively, while the Role command can be used to display the roles that have been assigned to specific users or groups.

See Also: root

role-based authorization

Authorization that uses roles to determine access rights and privileges.

Overview

A role is a symbolic category that collects together users who share the same levels of security privileges. Role-based authorization is a mechanism that uses roles to assign users suitable rights for performing system tasks and permissions for accessing resources. Role-based authorization is commonly used in business and financial applications to simplify the application of pol-

icy regarding who has what level of access to which resources.

Authorization Manager, a new feature of Microsoft Windows Server 2003, provides support for role-based authorization for the platform. This authorization model has several advantages over traditional low-level (ACL-based, or those methods based on access control lists) authorization methods, including the following:

- Simplifies access control management
- Allows scripts and applications to access authorization information easily
- Provides a mechanism for applying runtime business logic when checking access permissions

See Also: permissions, rights

role-based security

Any general mechanism that controls access to resources using roles instead of user credentials.

Overview

Role-based security is at the heart of many platforms and products including Microsoft Windows operating systems. The architecture of Microsoft Windows NT uses role-based security based on privileges assigned to local groups such as Administrators, Users, and Guests. By simply making a user a member of one of these groups, the user assumes the role of the generic group member and has all the rights to perform system tasks and permissions to access resources that belong to the group. Microsoft Transaction Server (MTS), and later COM+, enhanced this role-based security approach by providing developers with ways of defining their own abstract roles for use with custom-developed applications. Administrators could then assign users to specific roles to define levels of access to distributed applications and resources on a network. The Microsoft .NET Framework extends this model further by including support for role-based authorization within the common language runtime based on Windows accounts or custom identities.

See Also: access control, permissions, rights, role-based authorization

rollup

A cumulative set of hotfixes that can be applied in a single step.

Overview

Rollups (or security rollups) are packages of hotfixes provided by Microsoft Corporation that can patch a number of vulnerabilities in a single operation. Rollups simplify the deployment of security hotfixes and help administrators keep their systems up to date and secure from attack. Rollups generally target specific components or areas of a product's operation and are released from time to time by the Microsoft Security Response Center (MSRC), a team of security professionals at Microsoft responsible for responding to security threats involving Microsoft products.

See Also: *hotfix, Microsoft Security Response Center (MSRC), patch*

root

The superuser on UNIX/Linux platforms.

Overview

The root user (or simply **root**) in UNIX corresponds to the Administrator account on Microsoft Windows platforms and is the all-powerful account with virtually complete control over the system. Because of its extraordinary rights and privileges, root should always be protected with strong passwords, and only trusted individuals should be granted access to this account. The power of the account also makes it a prime target for attackers, however. The “quest for root” is the “holy grail” of cracking since gaining control of this account allows attackers to defeat virtually every aspect of a UNIX system's security, except any protection mechanisms implemented using physical security on the premises, such as lock-and-key access to a network attached storage (NAS) device or token-based smart card authentication. If an intruder is able to compromise root, the intruder also can erase all trail of the exploit by cleaning system logs and other auditing information, though there may be residual evidence that the audit log has been purged.

See Also: *Administrator, password, smart card*

root CA

The certificate authority (CA) at the top of a hierarchical Public Key Infrastructure (PKI).

Overview

A CA is a trusted entity (organization, company, or agency) that issues digital certificates for e-commerce, secure e-mail, and code-signing purposes. CAs are the foundation of PKI systems, both public and private, and most large PKI systems consist of multiple CAs arranged in a hierarchy of trust. At the top of the CA hierarchy is the root CA, which is the ultimate authority for the system. The root CA issues digital certificates to CAs under it to verify their identity, but no one can verify the identity of the root CA except itself, as the chain of trust must stop somewhere. The root CA, therefore, issues and signs its own certificate, called a root certificate. By deciding whether to trust this root certificate, a user decides whether to trust the entire PKI system.

See Also: *CA certificate, CA hierarchy, certificate authority (CA), digital certificate, Public Key Infrastructure (PKI), root certificate*

root certificate

A digital certificate identifying a root certificate authority (CA).

Overview

In a Public Key Infrastructure (PKI), each CA must have its own certificate so that users can trust the CA and verify its identity. In a typical hierarchical PKI model, the certificate for each CA (called a CA certificate) is issued by the CA immediately above it in the hierarchy. At the top of the hierarchy is the root CA, which must issue and sign its own certificate, called a root certificate. Ultimately, the trust that users have toward a particular PKI system depends on the trustworthiness of the root CA and its self-signed root certificate.

Marketplace

In the public arena, CAs act like passport offices to validate the identity of users and Web sites around the world. To facilitate e-commerce, secure e-mail, and secure downloading of code over the Internet, Microsoft Corporation maintains a list of trusted

third-party commercial CAs and preinstalls root certificates for these authorities in the Microsoft Internet Explorer Web browser. To be accepted as a trustworthy authority, a CA must apply to the Microsoft Root Certificate Program, which certifies it through an independent third-party audit by WebTrust for Certificate Authorities. The following table lists CAs certified as trustworthy together with the types of certificates they can issue, namely the following:

- Server certificates to verify the server to the client
- Client certificates to verify the client to the server
- Certificates for sending and receiving secure e-mail
- Certificates for signing program code

The table also shows which CAs time stamp their certificates.

Microsoft Root Certificate Program Trusted Certificate Authorities

| <i>Organization</i> | <i>Secure E-mail</i> | <i>Server Authentication</i> | <i>Client Authentication</i> | <i>Code Sign</i> | <i>Time Stamp</i> |
|---|----------------------|------------------------------|------------------------------|------------------|-------------------|
| Asociacion Nacional del Notariado Mexicano (www.notariadomexicano.org.mx/Asociados/index_asociados.htm) | Yes | Yes | No | No | No |
| Autoridade Certificadora Raiz Brasileira (www.icpbrasil.gov.br) | Yes | Yes | Yes | No | No |
| Baltimore (www.baltimore.com) | Yes | Yes | Yes | Yes | No |
| Belgacom E-Trust (www.e-trust.be/en/) | Yes | Yes | No | No | No |
| Certisign (www.certisign.com.br/) | Yes | Yes | No | No | No |
| CertPlus (www.certplus.com) | Yes | Yes | Yes | Yes | No |
| Correo (http://correo.com.uy/) | No | Yes | No | No | No |
| Deutsche Telekom (www.telekom.de) | Yes | Yes | No | No | No |
| DST (www.digsigtrust.com/) | Yes | Yes | No | No | No |
| Entrust (www.entrust.com/certificate_services/index.htm) | Yes | Yes | Yes | Yes | Yes |
| eSign (www.esign.com.au/) | Yes | No | Yes | No | Yes |
| EUnet International (www.eunet.fi/) | Yes | Yes | No | No | No |
| FESTE (www.feste.org/) | Yes | Yes | No | No | Yes |
| First Data Digital Certificates (www.firstdata.com/index.jsp) | Yes | No | Yes | Yes | Yes |
| FNMT (www.ceres.fnmt.es/) | Yes | Yes | No | No | No |
| Gatekeeper Root CA (www.govonline.gov.au/projects/confidence/Securing/Gatekeeper.htm) | Yes | No | Yes | No | Yes |
| GeoTrust (www.geotrust.com) | Yes | Yes | Yes | Yes | No |
| GlobalSign (www.globalsign.com/) | No | Yes | No | Yes | Yes |
| IPS SERVIDORES (www.ips.es/) | Yes | Yes | Yes | No | No |
| KMD (www.kmd-ca.dk) | Yes | No | Yes | No | Yes |
| NetLock (www.netlock.hu/) | Yes | No | Yes | No | No |
| Post.Trust (www.post.trust.ie/) | Yes | Yes | Yes | No | Yes |
| PTT Post (www.ptt-post.nl) | Yes | Yes | No | No | No |
| RSA (www.rsasecurity.com/) | No | Yes | No | No | No |

Microsoft Root Certificate Program Trusted Certificate Authorities (continued)

| <i>Organization</i> | <i>Secure E-mail</i> | <i>Server Authentication</i> | <i>Client Authentication</i> | <i>Code Sign</i> | <i>Time Stamp</i> |
|---|----------------------|------------------------------|------------------------------|------------------|-------------------|
| Saunalahden Serveri (www.sau-nalahti.fi/) | No | Yes | No | No | No |
| SecureNet (www.securenetasia.com/) | Yes | Yes | Yes | No | No |
| SecureSign (www2.jcsinc.co.jp) | No | Yes | No | No | No |
| SIA (https://ca.sia.it/) | No | Yes | No | No | No |
| TC TrustCenter (www.trustcenter.de/) | Yes | Yes | No | No | No |
| Thawte (www.thawte.com/) | Yes | Yes | Yes | Yes | Yes |
| UserTRUST (www.usertrust.com/) | Yes | Yes | No | Yes | No |
| ValiCert (www.valicert.com/) | Yes | Yes | No | No | No |
| Verisign (www.verisign.com/) | Yes | Yes | Yes | Yes | Yes |
| Wells Fargo Root Certificate Authority (www.wellsfargo.com/certpolicy) | Yes | Yes | Yes | No | No |

See Also: CA certificate, certificate authority (CA), digital certificate, Public Key Infrastructure (PKI), root CA, root rollover

rootkit

A set of tools installed by intruders once a system has been compromised.

Overview

Once an attacker has compromised a target system either by exploiting some known vulnerability, by brute-force attack, or through social engineering, the next step is usually to install a rootkit. A rootkit is basically a set of tools and scripts for automating certain tasks, including the following:

- Installing backdoors to allow compromised systems to be stealthily reentered
- Installing Trojans to capture login credentials and log keystrokes, create covert channels for information leakage, and other nefarious purposes
- Using Mount, Cron, or some other system tool to try to gain root privileges on the system using known exploits
- Installing remote administration tools (RATs) to all compromised systems to be controlled remotely

- Installing packet sniffers to capture passwords and other sensitive information transmitted over the network
- Replacing system files such as Netstat that could be used to detect the presence of the installed rootkit
- Cleaning log files to cover the tracks of the intruder and hide the fact that the system has been compromised

Rootkits first appeared in the early 1990s for SunOS 4.x and other UNIX platforms, when attackers used vulnerabilities in BIND and Xlib to gain a foothold in remote systems and install tools for further compromising the systems. Rootkits for Linux appeared soon afterward and were eventually followed by similar tools for compromising MS-DOS and Microsoft Windows systems. Some examples of Windows rootkits include Hacker Defender, HE4Hook, Null.sys, and Slanret.

The latest evolution of the rootkit is the kernel rootkit, which hides the rootkit within the operating system kernel, making it much harder to detect than traditional rootkits. This type of exploit uses the loadable kernel module (LKM) architecture used by many UNIX/Linux platforms to add rootkit functionality to the system by intercepting system calls instead of replacing existing system files. Knark and Adore are two examples of kernel rootkits targeting Linux platforms.

Marketplace

There are a number of free tools available for detecting the presence of installed rootkits, especially on UNIX/Linux systems, and they work with varying degrees of effectiveness. A utility called Rkdet (vancouver-webpages.com/rkdet/) can run in the background to detect attempts to install rootkits or packet sniffers.

When such an attempt is detected, Rkdet can shut down the system, disable the network connection, or send an alert by e-mail. Tripwire is another useful tool for detecting attempted installation of rootkits, while Chkrootkit can be used to check whether a system has already been compromised with a rootkit. Most commercial intrusion detection systems (IDSs) also include signatures for detecting common rootkit exploits.

Detecting and preventing kernel rootkits is more difficult, but one useful tool is Linux Intrusion Detection System (www.lids.org), which can seal the Linux kernel from being modified and can disable loading of LKMs. Sourceforge (www.sourceforge.net) has an LKM called St. Jude that can provide IDS capability built into the kernel to detect any attempts at installing rootkits or backdoors on Linux systems.

The usual way of recovering a system on which a rootkit has been installed is to remove the machine from the network, wipe its hard drive, and perform a clean install from original read-only installation media. Restoring from backup media is another possibility, but first certainty is required that the backup set was made before the system was compromised.

See Also: backdoor, intrusion detection system (IDS), Knark, Netstat, remote administration tool (RAT), sniffing, Trojan

root rollover

Expiration of the root certificate for a root certificate authority (CA).

Overview

In a Public Key Infrastructure (PKI), each root CA self-signs its own certificate, called a root certificate. This root certificate is typically installed in a user's Web browser or e-mail program so that the user can trust the

CA for purposes of e-commerce and secure e-mail. Because of constant advances in encryption technology, however, root certificates are generally valid only for a period of time before they expire, typically after 5 or 10 years; this process is called root rollover. Once a root certificate expires, users need to accept and install a new root certificate from the CA in their client software to continue using the CA to verify the identity of secure e-commerce sites or to send secure e-mail. Typically, such root rollover affects only users with older browsers or mail clients, since updated root certificates are usually preinstalled by vendors in newer versions of their client software.

See Also: certificate authority (CA), digital certificate, Public Key Infrastructure (PKI), root CA, root certificate

route verification

A packet-filtering technique for blocking spoofed packets.

Overview

Route verification is used by packet filters to identify packets that could not have originated from a legitimate host. For example, say the internal network uses the address block 172.16.0.0/16 for host addresses. If a packet originating from outside the network has an address from this same range, the packet must be spoofed because packets with such addresses should enter the packet filter only from the interface connected to the internal network. Normally, a packet like this would be dropped by the filter once its route has been verified as being illegitimate.

See Also: packet filtering

Rpcdump

A *Windows 2000 Resource Kit* tool for displaying services registered with the remote procedure call (RPC) endpoint mapper.

Overview

Rpcdump is a troubleshooting tool that also can be used by attackers to enumerate systems targeted for intrusion. RPC is a standard interprocess communication mechanism used by distributed applications for client/

server communication. When an RPC client wants to communicate with an RPC-enabled service running on a different host, the client first connects to the RPC endpoint mapper on Transmission Control Protocol (TCP) port number 135. The endpoint mapper responds by informing the client of the TCP port number of the desired service, and the client then can connect to the service. A common example is Microsoft Outlook, a Messaging Application Programming Interface (MAPI) client that uses RPCs to connect to Microsoft Exchange Server.

When problems with RPC communications arise, Rpcdump can be used to query the endpoint mapper for information about RPC services listening on the server. Using Rpcdump to troubleshoot RPC connections is similar to using Netstat to troubleshoot TCP connections, with the difference, however, that Rpcdump displays not just which ports are listening but also the names of the registered RPC services listening on these ports.

The main security concern is when RPC services are exposed for use over the Internet. By using Rpcdump to query port 135, an attacker on the Internet can enumerate a detailed list of RPC-based services running on a target machine, which can sometimes be used to identify vulnerabilities to exploit in order to compromise the target. If port 135 is blocked by a packet filter to prevent this, legitimate outside clients will be unable to connect to the host using RPCs. Another type of RPC attack that can be undertaken is a denial of service (DoS) attack against the endpoint mapper, which is accomplished by sending large numbers of RPC queries to port 135. Such an attack can also prevent legitimate RPC clients from connecting to the server. One solution to protect against these kinds of attacks is to use Microsoft Internet Security and Acceleration (ISA) Server, which is able to proxy RPC connections to port 135 while preventing Rpcdump from obtaining information about RPC services running on the protected network.

A similar exploit can be performed against UNIX hosts running the Rpcbind daemon, the RPC portmapper. This is accomplished by using the Rpcinfo utility,

which enumerates listening RPC services by connecting to port 111 (or port 32771 on Solaris hosts). The ubiquitous securing and hacking tool Nmap also can be used for RPC scanning to enumerate target networks. Some UNIX platforms have alternative versions of Rpcbind that can authenticate RPC clients using encryption; an example is SecureRPC from Sun Microsystems.

Notes

BindView's RAZOR research group (*razor.bind-view.com*) also has developed its own enhanced version of Rpcdump.

See Also: *denial of service (DoS), enumeration, Internet Security and Acceleration (ISA) Server, Netstat*

RSA

Stands for Rivest-Shamir-Adleman, a widely used public key cryptography algorithm.

See: *Rivest-Shamir-Adleman (RSA)*

Rsh

A UNIX utility for executing commands on a remote host.

Overview

Rsh, which stands for "remote shell," is part of the r-command package of UNIX tools that also includes Rexec and Rlogin. Rsh connects to the Rshd daemon (service) running on a remote host to run a specified command. If no credentials are specified with Rsh, the command is executed using the credentials of the currently logged on user, which must be specified in the .rhosts file on the remote host. Rsh cannot be used to run interactive commands such as the Vi editor; use Rlogin for this purpose instead. Like other r-commands, Rsh is not considered secure since it uses the .rhosts file, which is a frequent target for attackers trying to gain useful account information for attempting to compromise a target system. A better approach is to use the Secure Shell (SSH) tool for such purposes.

See Also: *Rexec, .rhosts, Rlogin, Secure Shell (SSH)*

rule

In packet filtering, a condition that determines whether to pass or drop a packet.

Overview

Packet-filtering routers use a list of rules to process incoming packets and decide whether to forward them across the interface or block them from being forwarded. These rules are generally processed sequentially until either a condition is met and the rule is applied, or the last rule is encountered and the default rule (usually “deny all”) takes effect. This default rule might be either explicitly included as the last rule in the filter or implicitly applied by the router.

Rules either allow or deny based on whether the packet matches the stated condition. This condition can include packet parameters such as the following:

- Source address
- Destination address
- Source port
- Destination port
- Service (protocol) type
- Various flags and different packet fields, depending on the protocol

Rules can also include such values as “all” or “any” so that all packets of a given type or source address or destination port can be processed similarly. The order of rules is important since reordering the rules can have different results. Depending on the implementation, some filters may apply the first matching rule and ignore the rest, while others may process the entire list and then apply the last matching rule from the list.

See Also: *packet filtering*

Runas

A Microsoft Windows command that allows a user to run an application using different credentials from those used for the current logon session.

Overview

Runas is a command included in Microsoft Windows 2000 and later versions of the operating system as a way of implementing the “least privilege” security principle. Best practice suggests that administrators should normally log on to their workstations using an ordinary user account in order to perform mundane tasks such as checking e-mail or browsing the World Wide Web. This reduces the chance of a malicious e-mail worm or virus being downloaded and gaining administrative privileges, and thus wreaking havoc on the system and possibly the entire network.

In previous versions of Windows, to perform an administrative task administrators would have to log off of their ordinary user account, log on using their Administrator account, perform the task, log off, and then log back on with their ordinary user account. Using Runas, however, administrators can temporarily run a command or application using administrative credentials while still logged on with their ordinary user account. This process is also known as secondary logon because a secondary set of credentials is used to execute the application or command while logged on to a session using a primary set of credentials.

Secondary logon can be performed either from the command line using Runas or from the desktop by holding down the Shift key, right-clicking on the icon or shortcut, and selecting Run As from the context menu. The Runas service must be running in order for secondary logon to work.

Notes

The Windows Runas command is similar in operation to the UNIX Su command.

See Also: *least privilege, Su, Sudo*

Rwho

A UNIX command for displaying information about users logged on to a remote host.

Overview

Rwho is a UNIX command that listens for the Rwhod daemon (service) running on remote hosts. Rwhod periodically broadcasts the user name, host name, and session start time for all users currently logged on to the remote host, and Rwho gathers this information and displays it for all network hosts on which Rwhod is running. While

the related Who command displays such information for a single remote host, Rwho displays such information for all hosts that are running Rwhod. On even a moderate-sized network, the broadcast traffic from such activity is considerable, and as a result best practice usually suggests that Rwhod be disabled on all systems. From a security perspective this is also a good idea since an intruder could use Rwho to gather considerable information for footprinting a target network.

See Also: *footprinting*

SACL

Stands for system access control list, a type of access control list (ACL) used for auditing securable objects.

See: *system access control list (SACL)*

sacrificial lamb

A server placed outside the firewall with the expectation that it may become compromised.

Overview

Sacrificial lambs are used for several reasons related to network security. First, any server placed outside the firewall is a prime target for attackers, and unless it is completely hardened it is likely to become compromised. However, such a server also can deflect the attention of attackers from more valuable servers hidden behind the firewall, with the result that the time intruders spend attacking a sacrificial lamb may possibly equal time not spent trying to compromise less exposed servers. In this way a sacrificial lamb functions similarly to a honeypot, and the two concepts are often used interchangeably.

More proper use of the term **sacrificial lamb**, however, is in the context of public Web servers, which generally need to be well exposed to the Internet so the public can access them easily. By placing your public Web server outside your firewall, you are viewing it as expendable and almost guaranteeing that it will eventually become compromised. But by making this your strategy and frequently backing up content on the server, the idea becomes that whenever the server is “sacrificed,” you can simply restore from backup and quickly get it up and running again. While some administrators follow this practice since it reduces the amount of effort needed to maintain a hardened Web server, many view this practice with suspicion and prefer to locate Web

servers on a screened subnet between two firewalls and expose only those services needed for external users to access them. But like most things concerning computer security, there is always a trade-off between time, cost, and effort on the one hand and degree of security on the other.

See Also: *hardening, honeypot*

Sadmind

A worm that compromises one platform to attack another.

Overview

The Sadmind worm appeared in May 2001 and represented a unique evolution in the development of Internet worms, namely, a worm that infects one operating system platform to infect another, an approach called a combination attack by some virus protection software vendors. Sadmind exploited a well-known buffer overflow vulnerability in the Sadmind service on the Solaris platform from Sun Microsystems, and it propagated across the Internet by taking advantage of the fact that many administrators had failed to patch their systems against this exploit even though the patch had been available for more than two years.

Once a Solaris host is compromised, Sadmind then tries to seek out and attack Microsoft Internet Information Services (IIS) Web servers by exploiting another vulnerability for which a patch had been available for almost seven months. The resulting attack on IIS machines caused Web site defacement and left backdoors open that could grant the attacker the ability to run arbitrary code on the servers to compromise them further.

Sadmind is a good lesson on the importance of administrators keeping their systems up to date with patches issued by vendors, because if the Solaris and IIS

patches had been applied, the worm would have been unable to do the damage it did.

See Also: patch, worm

Safe Harbor Agreement

An international agreement regarding the transfer of personally identifiable information (PII).

Overview

The Safe Harbor Agreement is a framework established in 2000 between the United States and the European Union (EU). The framework was initiated in response to the EU's Directive on Data Protection of 1998, which prohibits the transfer of PII to non-EU nations that fail to meet European standards concerning what constitutes the protection of individual privacy. Because the United States and EU take different approaches to what privacy is and how it is protected, the Safe Harbor Agreement was instituted to ensure the uninterrupted flow of business dealings between the two economies.

The Safe Harbor Agreement is consistent with Fair Information Practices, a series of standards governing collection and use of personal data that derived from the U.S. Privacy Act of 1974. To register for Safe Harbor with the U.S. Department of Commerce, companies must agree to abide by the Safe Harbor Agreement to adequately protect PII transferred from the EU to the United States.

For More Information

Visit www.export.gov/safeharbor/ for more information.

See Also: Fair Information Practices (FIP), personally identifiable information (PII), privacy, Safe Harbor Principles

Safe Harbor Principles

A series of directives for harmonizing privacy protection practices between the United States and the European Union (EU).

Overview

The Safe Harbor Principles are guidelines mandated by the Safe Harbor Agreement between the U.S. and EU

governments. These principles govern the transfer of personally identifiable information (PII) from the EU to the United States for companies that have registered for Safe Harbor under the agreement. The seven principles of Safe Harbor are as follows:

- **Access:** Individuals must have access to their PII that a company has collected to maintain its accuracy, within reasonable expense to the company.
- **Choice:** Companies must provide individuals with the opportunity to opt in or out of having their PII disclosed to third parties.
- **Data integrity:** The only PII that a company may collect from an individual is that which is relevant for the purposes for which it is used, and companies must take reasonable steps to ensure the information remains accurate and current.
- **Enforcement:** Individuals with complaints against a company for how it handles their PII must have affordable, independent mechanisms for recourse to investigate such problems and award damages where applicable by law, and there must also be mechanisms in place for verifying and enforcing the compliance of companies for Safe Harbor.
- **Notice:** Companies must provide individuals with notice concerning the purpose and use of the PII they collect, including disclosure practices to third parties and how complaints are handled.
- **Onward transfer:** Companies must disclose PII only to third parties that abide by Safe Harbor or the EU Directive on Data Protection or agree to abide by similar levels of data protection.
- **Security:** Companies must take reasonable precautions to protect PII collected from individuals against loss or misuse.

See Also: personally identifiable information (PII), privacy, Safe Harbor Agreement

SAINT

Stands for Security Administrator's Integrated Network Tool, a tool for assessing the security of a network.

See: Security Administrator's Integrated Network Tool (SAINT)

salt

Information added to a session key or password to strengthen it.

Overview

Many symmetric encryption schemes use a salt value to make it more difficult for an attacker to mount a brute-force attack against the key. Some authentication protocols use salts as well to increase the length and complexity of user passwords to make authentication more secure from eavesdropping. Adding a salt to a password is also important since two users could have the same password, and the salt combined with the password is then hashed to create a unique challenge string for each user.

Salts are generally random strings generated by pseudo-random number generators (PRNGs), but they also can consist of or include user-specific information such as a user's name or telephone number. When the hash of the password-plus-salt combination is transmitted, the plaintext version of the salt is usually appended since this is needed in order to generate a similar hash at the receiving end for verifying the password.

See Also: hashing algorithm, password, pseudorandom number generator (PRNG), session key

SAM

Stands for Security Accounts Manager, a database of local user accounts on Microsoft Windows NT or later.

See: Security Accounts Manager (SAM)

SAML

Stands for Security Assertion Markup Language, an Extensible Markup Language (XML) dialect for exchanging security information.

See: Security Assertion Markup Language (SAML)

Sam Spade

A site for tracking down spammers and a set of tools for the same purpose.

Overview

Sam Spade (www.samspade.org) is a popular Web site run by Steve Atkins that provides online access to network diagnostic tools useful for tracking the origin of spam. Sam Spade is also the name of a free package of downloadable tools for Microsoft Windows platforms that can be used for network tracking, diagnostics, and spam tracking. Using these tools to track down spammers requires some understanding of how Transmission Control Protocol/Internet Protocol (TCP/IP) works, however.

The name of the Web site originates with the detective played by Humphrey Bogart in the film noir *The Maltese Falcon*.

See Also: spam

sandbox

A protective mechanism used by some programming environments to limit access by programs to system resources.

Overview

Sandboxes are implemented to make it difficult or impossible for a program to damage user data or otherwise affect system security. Typically, a sandbox will restrict the privileges and commands that code can perform by providing a bounded, trusted environment within which the code can run. An example of a programming environment built around the sandbox model of code security is the Java programming language, in which an "applet" downloaded from a Web server to a user machine is "sandboxed" to prevent it from

performing any malicious actions that could harm the user's data. Within the sandbox the actions the applet may perform are limited; for example, an applet may not do the following:

- Read or write to the hard disk
- Spawn a new process
- Load a dynamic-link library (DLL) by directly calling a native application programming interface (API)
- Establish a new network connection

Sandbox models exist for other programming environments such as Python and TCL.

Sandwich Test

A rule of thumb for deciding whether to open e-mail attachments.

Overview

In an age of script kiddies and proliferating spam, e-mail attachments can sometimes contain malicious scripts, Trojans, or other malware that could do harm to the system of the user who opens them. E-mail filtering tools and security patches such as the Microsoft Outlook E-mail Security Update can help protect e-mail clients like Microsoft Outlook from malicious attachments, if they are installed and properly configured. However, sometimes it simply comes down to "Should I open this mail attachment or not?" and the Sandwich Test is a simple and proverbial method for deciding how to answer this question.

The idea is this: if you met a stranger on the street and he offered you a sandwich, would you eat it? Probably not; so don't open attachments from strangers as well. What if your sister handed you the sandwich instead? Well, that depends on what sort of relationship you have with your sister, for instance, where she might have got the sandwich, and perhaps whether you think she's smart enough to tell a nasty sandwich from a healthy one. I don't know about you, but I'd probably thank my sister, put the sandwich in my pocket, and toss it in the trash when she is not looking, unless I was either starving

or actually saw her buy the sandwich from a vendor that had an acceptable level of cleanliness in its operation.

The same principle applies to e-mail received from people you know when they forward you attachments they've received from others on the Internet. Most of the time they go straight into my Deleted Items folder without being opened—what about you? The whole idea here is that it's not technology that keeps our systems and networks secure, it's our brains that really do it—technology is dumb unless the people who use it are smart (Tulloch's "Principle of Least Smartness").

See Also: *spam, Trojan*

sanitized name

A standard format for certificate authority (CA) names.

Overview

A sanitized name is the form of a CA name used for a file name when storing such information in a Public Key Infrastructure (PKI) system. For example, the sanitized version of a common CA name would be used within a certificate revocation list (CRL), a list of revoked certificates maintained by a CA. When a CA name is sanitized, any illegal characters are removed, such as characters that are not allowed in file names, registry key names, Distinguished Names (DNs) or for some other technology-specific reason.

In Microsoft Certificate Services, sanitizing a common CA name causes any illegal characters to be converted into a five-character string of the form `!xxxx`, where `!` is employed as an escape character and `xxxx` represents four hexadecimal integers that uniquely identifies the character that is converted.

See Also: *certificate authority (CA), certificate revocation list (CRL), Public Key Infrastructure (PKI)*

SANS Institute

A cooperative research and education organization devoted to information security research, certification, and education.

Overview

The SANS Institute, established in 1989 and comprising security practitioners from government, business, and academia, is a trusted leader in information security. The institute provides news, security alerts, research papers, training courses, and other resources for professional development of system administrators, network administrators, auditors, and security professionals. Some of the programs and initiatives developed by SANS include the following:

- **SANS Computer & Information Security Training (www.sans.org):** Online and instructor-led courses covering practical steps necessary for protecting systems and networks against common threats
- **SANS/FBI Top 20 List (www.sans.org/top20):** List of current well-known and frequently exploited vulnerabilities together with step-by-step instructions on how to correct them; helps organizations prioritize their efforts to secure their information systems against attack
- **SANS Resources (www.sans.org/resources):** News digests, research summaries, security alerts, and other information useful for system and network professionals and security practitioners
- **SANS Reading Room (rr.sans.org):** Over 1300 articles in 63 different categories relating to information security
- **Global Information Assurance Certification (GIAC) (www.giac.org):** Certification program in intrusion detection, incident handling, firewalls, operating system security, and other security topics
- **Internet Storm Center (isc.incidents.org):** Research center that analyzes data collected from thousands of firewalls and intrusion detection systems (IDSs) in over 60 countries to search for trends and identify potential threats

Notes

The SANS in SANS Institute stands for “SysAdmin, Audit, Network, Security,” reflecting the areas of exper-

tise of its supporting members and scope of the institute’s activities.

See Also: *Global Information Assurance Certification (GIAC)*

SARA

Stands for Security Auditor’s Research Assistant, a tool for auditing the security of a network.

See: *Security Auditor’s Research Assistant (SARA)*

SAS

Stands for secure attention sequence, a special sequence of events that enables a user to log on or off a computer running Microsoft Windows NT or later.

See: *secure attention sequence (SAS)*

SATAN

Stands for System Administrator Tool for Analyzing Networks, a tool for identifying vulnerabilities in networks.

See: *System Administrator Tool for Analyzing Networks (SATAN)*

scanning

Short for port scanning, a method for determining which ports are “listening” (open) on a target system or network.

See: *port scanning*

screened subnet

Another name for demilitarized zone (DMZ), an isolated network segment at the point where a corporate network meets the Internet.

See: *demilitarized zone (DMZ)*

screening router

Another name for a packet-filtering router, a router that blocks packets based on a list of predetermined rules.

See: *packet filtering*

script kiddie

Pejorative term used to describe individuals who use packaged cracking tools widely available on the Internet.

Overview

The term **script kiddie** reflects the disdain that black hat hackers, who are usually highly competent programmers, have toward young individuals who download and use cracking tools to try to “take out” as many systems as they can on the Internet. The media picture of a typical script kiddie (and it’s usually not all that far off) is a lonely, bored, teenage boy who seeks attention by and takes pride in crashing Web servers and other public targets on the Internet. Script kiddies generally have a “wannabe” attitude, and if they perform an exploit, they are likely to get caught afterward because they bragged about their accomplishments on Internet Relay Chat (IRC). Unfortunately, there are widely available scripted toolsets on the Internet that can perform automated attacks by scanning large ranges of Internet Protocol (IP) addresses for vulnerable machines, and the availability of these tools is what makes script kiddies possible since they generally have limited programming skills. Despite the disdain that true hackers and crackers have toward these individuals, some have gained media notoriety through their success in bringing down major commercial Web sites and new media sites, actions that have sometimes landed them in prison—that is, if they’re old enough to be prosecuted.

For More Information

For a humorous look at script kiddies, search www.google.com for “Top 10 Ways to Spot a Script Kiddie.”

See Also: *hacking*

Sechole

A Trojan that exploited an elevation of privileges (EoP) vulnerability in Microsoft Windows NT.

Overview

Sechole was a sophisticated Trojan that appeared on the Internet in 1998. Sechole infected Windows NT machines running Internet Information Services (IIS) through Hypertext Transfer Protocol (HTTP) connections to infected Web sites. The Trojan used debugging application programming interfaces (APIs) to elevate the IUSR_servername anonymous Internet account to administrative privileges, and then added additional user accounts to the Administrators local group on the IIS machine. This allowed attackers to gain full control over the compromised machine without requiring local console access.

See Also: *Trojan*

secondary data uses

Using personally identifiable information (PII) for purposes other than why it was collected.

Overview

When companies collect PII from individuals, they are required to follow certain guidelines concerning how that information is used. Examples of such guidelines include Fair Information Practices (FIP), a set of standards governing the collection and use of personal data derived from the U.S. Privacy Act of 1974, and the Safe Harbor Agreement, an international agreement regarding the transfer of PII between the United States and the European Union (EU). Companies generally express compliance with these guidelines by publishing privacy policies that spell out in detail to individuals what the policy of the company is toward PII they collect from individuals. Any use of PII for purposes other than those for which the information was expressly stated as being collected constitutes secondary data use.

See Also: *Fair Information Practices (FIP), personally identifiable information (PII), privacy, Safe Harbor Agreement*

secondary logon

Another name for the Runas command, a Microsoft Windows command that allows a user to run an application using different credentials from those used for the current logon session.

See: Runas

secret key

A key used in secret key encryption.

Overview

A **secret key** is a key known only to parties engaging in encrypted communications using secret key encryption and is used for both encrypting plaintext and decrypting ciphertext that was encrypted using the same key. Secret keys generally range from 56 to 256 bits in length, and the longer the key, the more secure the encryption scheme. This is because exhaustive key search attacks that employ the brute-force method have to work twice as hard for each bit added to the length of the key, so a 64-bit key is 2 to the power $(64-56) = 256$ times harder to crack than a 56-bit key. For secure communications using secret key encryption, key lengths of at least 128 bits are now recommended, and it is unlikely that such keys will be crackable for at least a decade or more, but of course this depends on advances both in computer technology and in the mathematical science of cryptanalysis.

See Also: private key, public key, secret key encryption

secret key encryption

Encryption based on a shared secret between the parties communicating.

Overview

In secret key encryption, both parties share a common secret key, which they use to secure communications between them. Because secret key encryption systems use symmetric key algorithms, the same key can be used both for encrypting plaintext and decrypting corresponding ciphertext. An example of a symmetric algorithm used for secret key encryption is the Data Encryption Standard (DES), an encryption standard

used for many years by the U.S. government as a Federal Information Processing Standard (FIPS).

The advantage of secret key encryption over public key encryption is that because the key sizes are much smaller with secret key systems, the encryption process is generally faster and more efficient. The main drawback of secret key encryption is that some additional secure method must be used to distribute the shared secret to both parties before it can be used for encrypted communications. In the early days of cryptography and espionage, this was often done using a courier who would deliver the key to the other party by hand. Nowadays, electronic communications systems generally employ public key cryptography to securely exchange a secret key at the start of a session, after which the secret key is used for the remainder of the session for encrypting and decrypting data.

Another drawback of secret key encryption is that it does not support nonrepudiation, the ability to prove who performed an action such as sending a message, deleting a file, or rebooting a system. In a system that uses only secret key encryption, a party possessing a secret key can impersonate any other party sharing the same secret and then claim that “the other guy did it,” and there is no way to prove otherwise. To provide nonrepudiation for secure communications based on secret key encryption, public key cryptography can be employed to generate digital signatures that can be attached to messages to verify the identity of the sender.

See Also: asymmetric key algorithm, Data Encryption Standard (DES), Federal Information Processing Standard (FIPS), key, nonrepudiation, public key cryptography, secret key

secure attention sequence (SAS)

A special sequence of events that enables a user to log on or off a computer running Microsoft Windows NT or later.

Overview

Secure attention sequence (SAS) enables a unique sequence of events to alert the Microsoft Windows

security subsystem (Winlogon.exe) either that a user wants to log on to the computer or that the currently logged on user wants to log off, lock the workstation, or shut down the machine. This provides a secure mechanism for controlling who has interactive control of the console and a protected environment that prevents Trojans or other malware from fooling users into giving up their credentials. On an ordinary personal computer (PC) without special authentication hardware, SAS employs the Ctrl+Alt+Delete keystroke combination to notify Winlogon.exe to display the Windows Security dialog box. If additional authentication hardware such as a smart card reader is employed, SAS uses a different sequence of events such as swiping a card through the card reader.

See Also: *logon*

Secure Electronic Transaction (SET)

A family of specifications for secure credit card transactions over the Internet.

Overview

Work on Secure Electronic Transaction (SET) began in 1996 by a consortium that included Visa, MasterCard, Microsoft Corporation, and Netscape Communications. The idea behind SET was to develop an open technical standard to facilitate secure payment card transactions over the Internet. SET includes features of Public Key Infrastructure (PKI) and uses digital certificates to create a chain of trust to verify the identity of both the cardholder and the merchant.

The SET specification continues to evolve, and interoperability testing for SET-approved software has already resulted in some products being approved for the standard. To manage the evolution of the SET specification, Visa and MasterCard formed the Secure Electronic Transaction LLC (SETCo), which coordinates efforts toward the adoption of SET as a global standard for processing credit card payments.

For More Information

Visit www.setco.org for more information.

See Also: *Public Key Infrastructure (PKI)*

Secure Hash Algorithm-1 (SHA-1)

A hashing algorithm for generating a message digest.

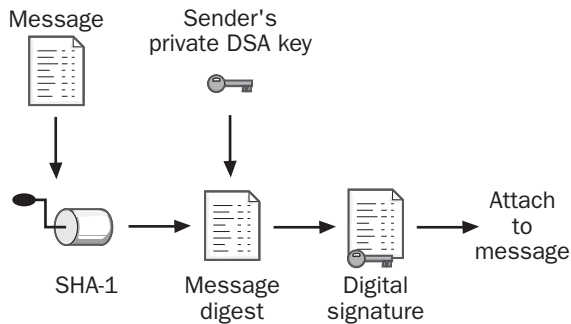
Overview

Secure Hash Algorithm-1 (SHA-1) is defined by the FIPS 180-1 standard, the Secure Hashing Standard (SHS), which was published by the National Institute of Standards and Technology (NIST) in 1993. SHA-1 is a hashing algorithm designed to create message digests that are signed using the Digital Signature Algorithm (DSA) defined in the FIPS 186-2 Digital Signature Standard (DSS). SHA-1 is used for both digitally signing messages and for verifying signatures received from others.

Implementation

The operation of SHA-1 is similar to the message digest 5 (MD5) algorithm defined in RFC 1320. SHA-1 takes a message of arbitrary length (actually, up to 2^{64} bits, which is astronomically huge) and uses a series of five stages to transform the message into a final digest, 160 bits in length (since MD5 uses only four such stages, SHA-1 is likely to be somewhat more secure than MD5). At the end of each stage the scrambled output from that stage is added to the value it had before the stage scrambled it.

To use SHA-1 together with DSA, the message is first fed into SHA-1 to generate a message digest. This digest is then signed using the DSA private key of the sender to create a digital signature for verifying the identity of the sender. When the recipient receives the message and attached signature, SHA-1 is applied to the received message to generate a new message digest, the signature is decrypted using the sender's public key to recover the original message digest, and the two digests are compared. If they are the same, the identity of the sender has been verified; if different, the message may have been intercepted and tampered with in transit.



Secure Hash Algorithm-1 (SHA-1). How SHA-1 and DSA are used to digitally sign a message.

Notes

The 1 in SHA-1 is there because a flaw was discovered in original Secure Hash Algorithm (SHA) after it was published by NIST as FIPS 180. The flaw was corrected by modifying one small step in the algorithm, and the standard was republished as FIPS 180-1, which renamed the algorithm SHA-1.

See Also: *Digital Signature Algorithm (DSA), Digital Signature Standard (DSS), Federal Information Processing Standard (FIPS), hashing algorithm, message digest (MD), message digest 5 (MD5), National Institute of Standards and Technology (NIST), Secure Hash Standard (SHS), SHA-2*

Secure Hash Standard (SHS)

The Federal Information Processing Standard (FIPS) defining the Secure Hash Algorithm-1 (SHA-1).

Overview

Secure Hash Standard (SHS) was originally outlined in the FIPS-180 standard published in 1993 by the National Institute of Standards and Technology (NIST). SHS defined the Secure Hash Algorithm (SHA), a hashing algorithm for creating message digests that then can be digitally signed. SHS was later revised in 1994 and republished as FIPS 180-1, which slightly modified SHA

to make it more secure, and the modified algorithm was called Secure Hash Algorithm-1 (SHA-1).

See Also: *Federal Information Processing Standard (FIPS), hashing algorithm, message digest (MD), National Institute of Standards and Technology (NIST), Secure Hash Algorithm-1 (SHA-1)*

Secure Hypertext Transfer Protocol (S-HTTP)

An extension for Hypertext Transfer Protocol (HTTP) to allow secure transfer of files.

Overview

Secure Hypertext Transfer Protocol (S-HTTP) was developed by Enterprise Integration Technologies and is defined in RFC 2660 as a method for encrypting and digitally signing files transferred using HTTP, the standard protocol of the World Wide Web. S-HTTP performs functions similar to that of Secure Sockets Layer (SSL) with the following differences:

- S-HTTP is an application-layer protocol (a variant of HTTP), while SSL is a transport-layer protocol. As a result, S-HTTP can be used only for securing HTTP traffic, while other kinds of traffic such as Simple Mail Transport Protocol (SMTP) can run on top of SSL to have their security enhanced.
- S-HTTP has a flexible architecture that can support a variety of encryption technologies, while cryptographic support in SSL is more fixed. Also, S-HTTP encrypts individual messages while SSL encrypts entire sessions.

While most browsers support S-HTTP, SSL is the most common method used for securing HTTP traffic, a scheme called HTTP over SSL (HTTPS).

See Also: *Secure Sockets Layer (SSL)*

Secure/Multipurpose Internet Mail Extensions (S/MIME)

An e-mail security standard that uses public key encryption.

Overview

Secure/Multipurpose Internet Mail Extensions (S/MIME) adds security to Simple Mail Transport Protocol (SMTP) messaging systems by allowing multipart Multipurpose Internet Mail Extensions (MIME)-encoded messages and attachments to be encrypted and digitally signed. S/MIME was developed in 1995 and adds privacy, confidentiality, and data integrity to e-mail messaging. S/MIME version 2 is defined by RFCs 3211 and 2312, and a new version 3 of S/MIME is currently an Internet-Draft standard.

Implementation

S/MIME is based on a hierarchical Public Key Infrastructure (PKI) system for issuing digital certificates and public/private key pairs to users. Because of the hierarchical nature of the system, S/MIME scales much better than Pretty Good Privacy (PGP), another popular secure e-mail system that relies on a diffuse “web of trust” model in which users exchange keys with everyone they want to send encrypted messages to. PKI is implemented in S/MIME using the Rivest-Shamir-Adleman (RSA) algorithm, which is used for securely exchanging session keys between users for encrypted messaging. S/MIME supports a variety of symmetric encryption schemes for encrypting messages using session keys, including Data Encryption Standard (DES), Triple DES (3DES), and RC2. The digital certificate format used by S/MIME is the standard X.509 format developed by the International Telecommunications Union (ITU).

S/MIME is supported by most popular e-mail client software, including Microsoft Outlook and Outlook Express.

See Also: *Data Encryption Standard (DES)*, *Pretty Good Privacy (PGP)*, *Public Key Infrastructure (PKI)*, *RC2*, *Rivest-Shamir-Adleman (RSA)*

Secure Shell (SSH)

A popular remote login protocol.

Overview

Secure Shell (SSH) was developed by Tatu Ylonen as a secure replacement for UNIX r-commands such as

Rlogin, Rexec, and Rcp. SSH also can replace Telnet by providing encrypted terminal connections and is more secure than File Transfer Protocol (FTP) for performing file transfers with remote hosts. In the SSH application suite the Ssh utility replaces Rlogin and Telnet, the Scp utility replaces Rcp, and Sftp replaces Ftp. The server-side component of SSH is Sshd.

SSH was originally proposed as an Internet standard but is now licensed by SSH Communications Security. However, several open source versions implementing the protocol are also available, including OpenSSH and others.

See Also: *OpenSSH*, *Rexec*, *Rlogin*

Secure Sockets Layer (SSL)

A protocol for secure communications over the Internet.

Overview

Secure Sockets Layer (SSL) is a protocol developed by Netscape Communications to allow sensitive or private information like credit card numbers to be transmitted securely over a medium that is inherently insecure: the Internet. SSL went through several versions resulting in SSLv3, which formed the basis of the Transport Layer Security (TLS) protocol described in RFC 2246.

SSL operates at the transport layer of the Transmission Control Protocol/Internet Protocol (TCP/IP) protocol suite. As a result, SSL is independent of the application-layer protocol above it and can be used not only for encrypting Web traffic using Hypertext Transfer Protocol (HTTP) but also mail or newsgroup traffic as well. The main use for SSL, however, is for encrypting HTTP traffic, and the combination of HTTP running over SSL is known as HTTPS.

Implementation

The operation of SSL involves a combination of public key cryptography and secret key encryption to provide data confidentiality through encryption. The digital certificates and public/private key pairs used in SSL are generated using the Rivest-Shamir-Adleman (RSA) public key algorithm.

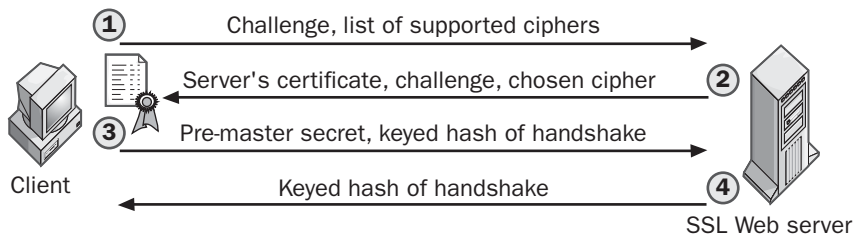
Use HTTPS as an example: when a client Web browser wants to connect to a Web server that uses SSL, the client uses a Uniform Resource Locator (URL) beginning with `https://` to initiate the SSL handshaking process with the server. This handshaking process is used to negotiate the secret key encryption algorithm both parties will use for encrypting information sent between them during the session. The initial information sent by the client to the server includes a list of encryption algorithms the client supports and a random challenge string used later on in the handshake.

Once the client has challenged the server, the server responds by returning a copy of its server certificate, a digital certificate used by the server to authenticate its identity to others. In order for SSL to work, the server earlier must have obtained a server certificate from a certificate authority (CA) such as Verisign. Along with the certificate, the server also includes its own random challenge string and selects an encryption algorithm to use from the list submitted earlier by the client. Examples of secret key encryption algorithms supported by SSL include RC4 and Data Encryption Standard (DES).

The client then validates the certificate sent by the server using the server's public key to ensure it is actu-

ally talking to the server it wants to talk with. The client obtains the server's public key by extracting it from the server's certificate received in the previous step. The client then generates another random string called the **premaster secret**, which will be used in the process for generating the session key for the session. The client also encrypts another value called the premaster secret using the server's public key and sends the encrypted premaster secret to the server, along with a keyed hash of the handshaking messages together with a master secret. This hash helps ensure that the handshaking messages have not been tampered with by an eavesdropper attempting to hijack the session. The key used for the hash is derived from the two random strings sent earlier by each party together with the master secret.

The server finally completes the process by sending the client a keyed hash of all the handshaking messages exchanged up to this point. Both parties then derive the session key from the different random values and keys exchanged using a complicated mathematical operation. All data between the client and server for the duration of the session is encrypted using the session key, which is then discarded when the session terminates or times out.



Secure Sockets Layer (SSL). How the SSL handshake works.

Notes

SSL can also be used to authenticate the client to the server if the client also has obtained a public/private key pair and digital certificate from a CA trusted by the server.

See Also: *public key cryptography, secret key encryption, server certificate, SSL accelerator, Transport Layer Security (TLS)*

Secure Windows Initiative (SWI)

A Microsoft Corporation initiative to ensure the security of its products.

Overview

Secure Windows Initiative (SWI) is part of a broader security initiative at Microsoft that includes the Microsoft Security Response Center (MSRC) and the

Trustworthy Computing Initiative (TCI). The focus of SWI is to help product development teams at Microsoft design and build products that are secure from malicious attack. The SWI acts as a central security consulting arm for developer teams and helps them write secure code by implementing the following:

- Periodically running Security Review Days that provide instruction in different aspects of code security
- Live and online presentations dealing with threat analysis, coding practices, and secure application configuration broadcast over the Microsoft intranet
- Best practice documentation that records any code security issues discovered in existing products and how to fix them
- Improvements to Microsoft Developer Network (MSDN) and the platform software development kit (SDK) to include security issues for function calls

The SWI works together with MSRC to resolve security vulnerabilities when they are discovered in Microsoft products. The SWI becomes involved at the beginning of each incident and helps triage the problem to determine whether it is a security bug or some other issue and helps ensure the problem is addressed properly.

For More Information

Read the book *Writing Secure Code* (Microsoft Press, 2002) by Michael Howard and David LeBlanc, both of whom are security experts with Microsoft.

See Also: *Microsoft Security Response Center (MSRC), Trustworthy Computing Initiative (TCI)*

S

Security+

A vendor-neutral security certification developed by the Computing Technology Industry Association (CompTIA).

Overview

Security+ is a widely recognized certification exam developed in collaboration with IT (information technology) security practitioners from industry, academia, and government. The aim of Security+ is to provide a

way for individuals to demonstrate a basic level of competency in information systems security by passing a standardized exam covering five subject areas:

- general security concepts
- communication security
- infrastructure security
- basics of cryptography
- operational/organizational security

For More Information

Visit www.comptia.org/certification/Security/ for more information.

See Also: *Certified Information Systems Security Professional (CISSP), Global Information Assurance Certification (GIAC)*

Security Accounts Manager (SAM)

The database of local user accounts on Microsoft Windows NT or later.

Overview

The Security Accounts Manager (SAM) database contains security information for local user and group accounts on standalone machines running Windows NT, Windows 2000, Windows XP, and Windows Server 2003. The SAM database is implemented as a registry hive named HKEY_LOCAL_MACHINE\SAM, whose contents are not accessible using normal registry editing tools while Windows is running.

The SAM database is a common target for attackers trying to compromise the security of a Windows machine, and if they can gain access to the database, they will then try to extract password information from it using common password-cracking tools such as L0phtcrack or John the Ripper. Ensuring the security of the database is therefore important, and one way of doing this is by using Syskey, a utility that uses strong encryption for strengthening password security.

The SAM database is mainly used on standalone Windows machines belonging to a workgroup. When

a member server is promoted to a domain controller, all account information stored in the SAM is migrated to the Active Directory directory service. The only time the SAM is used on a domain controller is when an administrator boots the domain controller into Directory Services Restore Mode or uses the Recovery Console.

See Also: *John the Ripper, L0phtcrack, password cracking, Syskey*

Security Administrator's Integrated Network Tool (SAINT)

A tool for assessing the security of a network.

Overview

Security Administrator's Integrated Network Tool (SAINT) is a security auditing and assessment tool that can be used to identify vulnerabilities in networks so that their security can be enhanced. SAINT works by scanning networks to find live Internet Protocol (IP) hosts and identify different services running on the hosts. For each service identified, SAINT launches probes to test the service for a variety of known vulnerabilities that could be exploited by attackers for compromising the host. SAINT not only identifies such vulnerabilities but also categorizes them and displays information concerning how to correct them by applying vendor patches or upgrading to new software versions. SAINT identifies such vulnerabilities according to their names as used by the Common Vulnerabilities and Exposures (CVE) standard from MITRE Corporation, CERT Coordination Center (CERT/CC) security advisories, and other industry-standard vulnerability naming schemes. SAINT is highly configurable and can prioritize found vulnerabilities to allow administrators to address them in a planned, methodical manner.

The SAINT4 Vulnerability Assessment tool includes the SAINT scanning issue with its easy-to-use graphical interface and SAINTwriter for generating customized vulnerability reports. SAINT is available for the Solaris, HP-UX, FreeBSD, OpenBSD, and Linux platforms.

For More Information

Visit www.saintcorporation.com for more information.

See Also: *CERT Coordination Center (CERT/CC), Common Vulnerabilities and Exposures (CVE), port scanning, Security Auditor's Research Assistant (SARA), System Administrator Tool for Analyzing Networks (SATAN), vulnerability*

Security Assertion Markup Language (SAML)

An Extensible Markup Language (XML) dialect for exchanging security information.

Overview

Security Assertion Markup Language (SAML) is an XML language designed to allow Web services platforms from different vendors to interoperate in the area of security. Using SAML, a client can be authenticated and authorized with a Web service using standard XML message formats that any SAML-compliant platform can understand. SAML is designed to facilitate the growth of e-commerce by providing a common language for products and services from different companies to exchange security information. SAML can benefit both the business-to-business (B2B) and business-to-consumer (B2C) marketplaces by enabling suppliers, business partners, and consumers to authenticate with each other in a standard way.

SAML 1 has been approved as a standard by the Organization for the Advancement of Structured Information Standards (OASIS). SAML 1 describes how security and policy domains can communicate using Simple Object Access Protocol (SOAP) messages over Hypertext Transfer Protocol (HTTP) and supports single sign-on (SSO) authentication and authorization.

See Also: *authentication*

Security Auditor's Research Assistant (SARA)

A tool for auditing the security of a network.

Overview

Security Auditor's Research Assistant (SARA) is a tool for performing internal and external audits of a network's security. SARA is based on the System Administrator Tool for Analyzing Networks (SATAN) and scans networks or individual systems for vulnerabilities that could be exploited by attackers. SARA uses an extensible framework that includes support for third-party security plug-ins, integration with Nmap, and integration with Samba for Server Message Block (SMB) scanning. SARA employs an easy-to-use Web interface for scanning targets and analyzing the results and can perform various levels of scanning from light to extreme in which higher levels of scanning are more likely to trigger intrusion detection systems (IDSs), generate system log entries or error messages, or even cause services to fail or systems to crash.

SARA is certified by the SANS Institute and lets administrators scan their networks for vulnerabilities listed in the FBI/SANS Top 20 Vulnerabilities list. SARA also supports scanning for vulnerabilities identified by Common Vulnerabilities and Exposures (CVE) from MITRE Corporation. SARA runs on UNIX platforms including Mac OS X and is updated frequently to reflect the changing vulnerabilities landscape. SARA is free for use under a General Public License (GPL)-like license.

For More Information

Visit www-arc.com/sara for more information.

See Also: *Common Vulnerabilities and Exposures (CVE), Nmap, port scanning, SANS Institute, Security Administrator's Integrated Network Tool (SAINT), System Administrator Tool for Analyzing Networks (SATAN), vulnerability*

Security Configuration and Analysis

A tool for managing security settings on machines running Microsoft Windows 2000 or later.

Overview

Security Configuration and Analysis is a Microsoft Management Console (MMC) snap-in for analyzing and configuring the security of a machine running Windows. The tool can be used to compare a machine's current security settings to those defined by a **security template**, a collection of settings defining security policy for a computer, or to apply the settings defined in a security template to the machine's Local Security Policy or across a domain using Group Policy.

Notes

There is also a command-line version of Security Configuration and Analysis called Secedit.

See Also: *security template*

security context

The security attributes or rules currently in effect.

Overview

On Microsoft Windows platforms, the current user logged on to the computer or the personal identification number entered by the smart card user is an example of the security context for a system. From the perspective of a service running in the background, the LocalSystem account or some other security principal used as a service account is an example of the security context of a process running on the system. From the perspective of a security support provider interface (SSPI), a security context is an opaque data structure containing security data relevant to a connection; for example, a session key.

See Also: *security principal*

security descriptor

A data structure containing security information for a securable object.

Overview

On Microsoft Windows platforms, access control lists (ACLs) alone are not sufficient to describe the complete security attributes of a securable object such as a file or folder on an NTFS volume. A security descriptor is also needed to complete the object's security attributes,

and the security descriptor contains the following four elements:

- The security identifier (SID) of the owner of the object
- The SID of the primary group for the object (used by POSIX applications only)
- The discretionary access control list (DACL) that controls access to the object
- The system access control list (SACL) that controls the logging of attempts to access the object

When you want to display or modify the security attributes of a securable object, you use the object's security descriptor to do this.

See Also: *access control list (ACL), security identifier (SID)*

security identifier (SID)

A string value that uniquely identifies a security principal.

Overview

Security principals are entities such as user accounts or logon sessions that can be authenticated by the Microsoft Windows security subsystem. Each security principal is assigned a unique security identifier (SID) that is implemented as a data structure of variable length and used by Windows processes to uniquely identify the principal to the system or network. There are also a number of well-known SIDs whose values are the same on all Windows platforms and which identify generic users or groups to operating system processes. Knowledge of well-known SIDs can be useful for troubleshooting issues involving security, and they are listed in the following table for reference.

Well-Known SIDs

| <i>SID</i> | <i>Name</i> | <i>Description</i> |
|------------|----------------------|---|
| S-1-0 | Null Authority | An identifier authority. |
| S-1-0-0 | Nobody | No security principal. |
| S-1-1 | World Authority | An identifier authority. |
| S-1-1-0 | Everyone | A group that includes all users, even anonymous users and guests. Membership is controlled by the operating system. |
| S-1-2 | Local Authority | An identifier authority. |
| S-1-3 | Creator Authority | An identifier authority. |
| S-1-3-0 | Creator Owner | A placeholder in an inheritable access control entry (ACE). When the ACE is inherited, the system replaces this SID with the SID for the object's creator. |
| S-1-3-1 | Creator Group | A placeholder in an inheritable ACE. When the ACE is inherited, the system replaces this SID with the SID for the primary group of the object's creator. The primary group is used only by the POSIX subsystem. |
| S-1-3-2 | Creator Owner Server | This SID is not used in Windows 2000. |
| S-1-3-3 | Creator Group Server | This SID is not used in Windows 2000. |
| S-1-4 | Nonunique Authority | An identifier authority. |
| S-1-5 | NT Authority | An identifier authority. |
| S-1-5-1 | Dialup | A group that includes all users who have logged on through a dial-up connection. Membership is controlled by the operating system. |

Well-Known SIDs (continued)

| <i>SID</i> | <i>Name</i> | <i>Description</i> |
|------------------|---------------------------|--|
| S-1-5-2 | Network | A group that includes all users that have logged on through a network connection. Membership is controlled by the operating system. |
| S-1-5-3 | Batch | A group that includes all users that have logged on through a batch queue facility. Membership is controlled by the operating system. |
| S-1-5-4 | Interactive | A group that includes all users that have logged on interactively. Membership is controlled by the operating system. |
| S-1-5-5-X-Y | Logon Session | A logon session. The X and Y values for these SIDs are different for each session. |
| S-1-5-6 | Service | A group that includes all security principals that have logged on as a service. Membership is controlled by the operating system. |
| S-1-5-7 | Anonymous | A group that includes all users that have logged on anonymously. Membership is controlled by the operating system. |
| S-1-5-8 | Proxy | This SID is not used in Windows 2000. |
| S-1-5-9 | Enterprise Controllers | A group that includes all domain controllers in a forest that uses an Active Directory directory service. Membership is controlled by the operating system. |
| S-1-5-10 | Principal Self | A placeholder in an inheritable ACE on an account object or group object in Active Directory. When the ACE is inherited, the system replaces this SID with the SID for the security principal who holds the account. |
| S-1-5-11 | Authenticated Users | A group that includes all users whose identities were authenticated when they logged on. Membership is controlled by the operating system. |
| S-1-5-12 | Restricted Code | This SID is reserved for future use. |
| S-1-5-13 | Terminal Server Users | A group that includes all users that have logged on to a Terminal Services server. Membership is controlled by the operating system. |
| S-1-5-18 | LocalSystem | A service account that is used by the operating system. |
| S-1-5-19 | NT Authority | Local Service. |
| S-1-5-20 | NT Authority | Network Service. |
| S-1-5-domain-500 | Administrator | A user account for the system administrator. By default, it is the only user account that is given full control over the system. |
| S-1-5-domain-501 | Guest | A user account for people who do not have individual accounts. This user account does not require a password. By default, the Guest account is disabled. |
| S-1-5-domain-502 | KRBTGT | A service account that is used by the Key Distribution Center (KDC) service. |
| S-1-5-domain-512 | Domain Admins | A global group whose members are authorized to administer the domain. By default, the Domain Admins group is a member of the Administrators group on all computers that have joined a domain, including the domain controllers. Domain Admins is the default owner of any object that is created by any member of the group. |
| S-1-5-domain-513 | Domain Users | A global group that, by default, includes all user accounts in a domain. When you create a user account in a domain, it is added to this group by default. |

Well-Known SIDs (continued)

| <i>SID</i> | <i>Name</i> | <i>Description</i> |
|-----------------------|-----------------------------|--|
| S-1-5-domain-514 | Domain Guests | A global group that, by default, has only one member, the domain's built-in Guest account. |
| S-1-5-domain-515 | Domain Computers | A global group that includes all clients and servers that have joined the domain. |
| S-1-5-domain-516 | Domain Controllers | A global group that includes all domain controllers in the domain. New domain controllers are added to this group by default. |
| S-1-5-domain-517 | Cert Publishers | A global group that includes all computers that are running an enterprise certification authority. Cert Publishers are authorized to publish certificates for User objects in Active Directory. |
| S-1-5-root domain-518 | Schema Admins | A universal group in a native-mode domain; a global group in a mixed-mode domain. The group is authorized to make schema changes in Active Directory. By default, the only member of the group is the Administrator account for the forest root domain. |
| S-1-5-root domain-519 | Enterprise Admins | A universal group in a native-mode domain; a global group in a mixed-mode domain. The group is authorized to make forest-wide changes in Active Directory, such as adding child domains. By default, the only member of the group is the Administrator account for the forest root domain. |
| S-1-5-domain-520 | Group Policy Creator Owners | A global group that is authorized to create new Group Policy objects in Active Directory. By default, the only member of the group is Administrator. |
| S-1-5-domain-533 | RAS and IAS Servers | A domain local group. By default, this group has no members. Servers in this group have Read Account Restrictions and Read Logon Information access to User objects in the Active Directory domain local group. By default, this group has no members. |
| S-1-5-32-544 | Administrators | A built-in group. After the initial installation of the operating system, the only member of the group is the Administrator account. When a computer joins a domain, the Domain Admins group is added to the Administrators group. When a server becomes a domain controller, the Enterprise Admins group also is added to the Administrators group. |
| S-1-5-32-545 | Users | A built-in group. After the initial installation of the operating system, the only member is the Authenticated Users group. When a computer joins a domain, the Domain Users group is added to the Users group on the computer. |
| S-1-5-32-546 | Guests | A built-in group. By default, the only member is the Guest account. The Guests group allows occasional or one-time users to log on with limited privileges to a computer's built-in Guest account. |
| S-1-5-32-547 | Power Users | A built-in group. By default, the group has no members. Power Users can create local users and groups; modify and delete accounts that they have created; and remove users from the Power Users, Users, and Guests groups. Power Users also can install programs; create, manage, and delete local printers; and create and delete file shares. |

Well-Known SIDs (continued)

| <i>SID</i> | <i>Name</i> | <i>Description</i> |
|--------------|-------------------|--|
| S-1-5-32-548 | Account Operators | A built-in group that exists only on domain controllers. By default, the group has no members. By default, Account Operators have permission to create, modify, and delete accounts for users, groups, and computers in all containers and organizational units of Active Directory except the Built-in container and the Domain Controllers OU. Account Operators do not have permission to modify the Administrators and Domain Admins groups, nor do they have permission to modify the accounts for members of those groups. |
| S-1-5-32-549 | Server Operators | A built-in group that exists only on domain controllers. By default, the group has no members. Server Operators can log on to a server interactively, create and delete network shares, start and stop services, back up and restore files, format the hard disk of the computer, and shut down the computer. |
| S-1-5-32-550 | Print Operators | A built-in group that exists only on domain controllers. By default, the only member is the Domain Users group. Print Operators can manage printers and document queues. |
| S-1-5-32-551 | Backup Operators | A built-in group. By default, the group has no members. Backup Operators can back up and restore all files on a computer, regardless of the permissions that protect those files. Backup Operators also can log on to the computer and shut it down. |
| S-1-5-32-552 | Replicators | A built-in group that is used by the File Replication service on domain controllers. By default, the group has no members. Do not add users to this group. |

See Also: *security principal, Sid2user, User2sid*

security log

An event log on Microsoft Windows platforms used for auditing security events.

Overview

The security log is used for recording success and failure audit events when auditing is configured on machines running Windows NT or later operating systems. Audit entries recorded in the security log display the action performed, the user who performed it, and the date and time of the action to create an audit trail for troubleshooting security-related issues. On machines running Windows 2000 and later, auditing can be enabled and configured using the Audit Policy settings under Local Policies in Group Policy. Once file and object access is enabled, you can specify which files

on NTFS volumes to monitor and for which users and groups.

See Also: *auditing, event logs*

security policy

A written policy outlining the implementation and management of network security.

Overview

A security policy is a document containing guidelines and instructions regarding subjects such as the following:

- Generating and using passwords for authentication purposes
- Protecting the privacy of users' personally identifiable information (PII)
- Defining who has what access rights and privileges to which resources on the network and why

- Performing periodic audits of network security
- Handling incidents in which systems are compromised by intruders
- Establishing expectations for users regarding system availability
- Purchasing policy for security tools, systems, and software
- Limiting physical access to computing resources
- Reporting violations of the policy and enforcing its provisions
- Legal and regulatory issues in which user compliance is required

To develop a security policy for your network, you might follow a procedure like the following:

- 1- Form a team that includes IT (information technology) staff, management, and legal counsel.
- 2- Perform an inventory of your security needs including an audit of your current level of network security.
- 3- Weigh your security needs against their possible cost and how they can affect the ease of use of your computing resources.
- 4- Define the practices needed to meet and maintain your security needs from the perspective of the ordinary user.
- 5- Write down your policy in a clearly understandable fashion.
- 6- Review your policy to ensure it can be implemented and enforced in a practical way.
- 7- Publish your policy so that users can have easy access to it.
- 8- Call attention to it on a regular basis and enforce violations with consistency.
- 9- Revise your policy periodically after careful review.

For More Information

Review the Internet standard document RFC 2196 called “Site Security Handbook” at www.ietf.org/rfc/rfc2196.txt for more information.

See Also: *access control, password, personally identifiable information (PII)*

security principal

An entity that can be authenticated by the Microsoft Windows security subsystem.

Overview

On Windows platforms, a security principal is an account that has a security identifier (SID) assigned to it so that access by the account to resources can be controlled. Examples of different types of security principals include users, groups, computers, and **special identities**, well-known security principals that are managed by the operating system instead of administrators. From the perspective of the security subsystem of the Windows operating system, it doesn’t matter whether the security principal represents a human user or a process running on the computer, both are recognized by the operating system and treated the same way.

In domain-based networks using Windows 2000 or later, security principals are stored in Active Directory directory service. In the earlier Windows NT operating system, they were stored in the SAM database.

See Also: *security identifier (SID), special identities*

security rollup package

Often simply called a rollup, a cumulative set of hot-fixes that can be applied in a single step.

See: *rollup*

security support provider interface (SSPI)

A set of application programming interfaces (APIs) for accessing security services on Microsoft Windows platforms.

Overview

The security support provider interface (SSPI) is a common interface between transport-level applications and security providers and allows a transport application to call the specific security provider it needs to obtain an authenticated connection without requiring detailed knowledge of how the security protocol is implemented. For example, a transport application such as the remote procedure call (RPC) facility could call packages for either Kerberos or LAN Manager authentication depending on whether downlevel computers running on the Windows platform were involved. The SSPI is implemented by the security support provider (SPP), a dynamic-link library (DLL) that makes security packages available to applications. Each package provides mappings between the application's SSPI function calls and the actual functions of the security model.

SSPI provides a wide range of security packages for such services as authentication, message integrity, message privacy, and quality of service. From an architectural perspective, SSPI is Microsoft Corporation's implementation of the Generic Security Service API (GSSAPI) standard described in RFCs 1508 and 1509.

See Also: LAN Manager authentication, Kerberos

security template

In Microsoft Windows 2000 and later, a collection of settings defining security policy for a computer.

Overview

Security templates are *.inf files used for defining policy settings for securing different aspects of a computer running on the Windows platform. Security templates include settings for the following:

- Account policies concerning passwords, lockouts, and Kerberos settings
- Local policies concerning user rights and logging of security events
- Membership restrictions for local groups on the system
- Security for registry keys on the system

- File system security on the system
- Security and startup mode for local services

Using the Security Configuration and Analysis tool, administrators can select a predefined or custom security template and apply it to a system to lock down the system according to the level of security that matches its role. The predefined templates available vary with different versions of the Windows operating system. On the Windows 2000 platform, for example, some of the templates available include the following:

- **Basic:** Default security settings for safe environments
- **Secure:** Settings for secure workstations and servers in mixed Windows 2000/NT environments
- **Highly Secure:** Settings for secure workstations and servers in native Windows 2000 environments
- **Compatible:** Relaxed security settings to resolve application compatibility issues

Security templates also can be imported into Group Policy and applied at the domain, site, or organizational unit level, simplifying the job of administering large numbers of computers on a network. Security templates can be created or modified using the Security Templates snap-in.

See Also: Security Configuration and Analysis

security zone

A security feature implemented by Microsoft Internet Explorer for safer browsing.

Overview

To make browsing the World Wide Web a safer experience, Internet Explorer divides **URLspace** (the abstract space of all possible URLs in the world) into different zones, with each zone having its own security template. The purpose of this approach is to allow users to assign a site to a zone to control which actions can be performed on the browser for sites in that zone. Each zone has its own default security template defining its security policy settings, but each zone's settings also can be

customized if required to prohibit or allow different forms of content. Examples of policy areas in zone security templates include the following:

- Whether ActiveX controls and browser plug-ins may be downloaded, installed, or used
- Whether persistent or per-session cookies can be saved
- Whether files and fonts can be downloaded
- The safety level in which the Java Virtual Machine runs
- Miscellaneous security settings involving certificates, software channels, drag and drop, frames, and so on
- Settings for active scripting and scripting of Java applets
- What forms of authentication are allowed

There are four different security zones to which you can add sites or domains:

- **Local Intranet:** Contains all network connections established using a Universal Naming Convention (UNC) path. It also includes Web sites that bypass a proxy server or have names without periods (such as **http://servername**), provided these sites are not assigned to another zone.
- **Trusted Sites:** Contains sites you explicitly trust as safe. Examples of trusted sites might be Web sites on your company intranet or the site of a business partner you have confidence in. There are initially no sites in this zone until you add them.
- **Restricted Sites:** Contains sites you explicitly choose not to trust. Examples of restricted sites might be sites with malicious scripts or those from which viruses can be downloaded. There are initially no sites in this zone until you add them.
- **Internet:** Contains all Web sites not included in any other zones, which by default is any site on the Internet.

There is also a fifth implied zone called the Local Machine or My Computer zone that applies to any files stored on your local computer. This zone cannot be configured using the Web browser interface and can be modified only using the Microsoft Internet Explorer Administration Kit (IEAK).

The default level of security associated with each zone is as follows:

- **Local Intranet:** Medium for Internet Explorer 4 and medium-low for Internet Explorer 5 and 6
- **Trusted Sites:** Low for all versions of Internet Explorer
- **Restricted Sites:** High for all versions of Internet Explorer
- **Internet:** Medium for Internet Explorer 4, 5, and 6 but high for Internet Explorer 6 on Microsoft Windows Server 2003

Notes

An easy way to tell which zone your current Web page belongs to is by looking at the icon at the right side of the Internet Explorer status bar.

SendIP

A tool for sending arbitrary Internet Protocol (IP) packets.

Overview

SendIP is a command-line program developed by Mike Ricketts, a software engineer working at Hursley Laboratories for IBM UK. Using SendIP, you can create a packet from scratch and assign it any header options you desire. SendIP allows you to create a wide range of different kinds of packets, including the following:

- Internet Protocol version 4 (IPv4) or Internet Protocol version 6 (IPv6) raw packets
- Transmission Control Protocol (TCP) packets
- User Datagram Protocol (UDP) packets
- Internet Control Message Protocol (ICMP) packets
- Routing Information Protocol (RIP) packets

- Border Gateway Protocol (BGP) packets
- Network Time Protocol (NTP) packets

SendIP also allows you to add an arbitrary data payload to the packet and to calculate or forge a checksum to attach to the packet. SendIP is popular with the hacking community, both white and black hat, for its simplicity and power in probing the inner workings of Transmission Control Protocol/Internet Protocol (TCP/IP) stacks on different platforms and systems. Like most security and network troubleshooting tools, SendIP can be used for malicious purposes as well; for example, to generate packet fragments for an IP fragmentation attack or to create spoofed packets for various types of denial of service (DoS) attacks.

SendIP runs on UNIX/Linux platforms and is distributed as open source software under the General Public License (GPL).

For More Information

Visit www.earth.li/projectpurple/ for more information.

See Also: denial of service (DoS), hacking, spoofing

sensitive data

Personally identifiable information (PII) that is protected in special ways by law or policy.

Overview

What constitutes sensitive data is different for different jurisdictions around the world. For example, from the European Union perspective, sensitive data is any PII having to do with race or ethnic origin, political opinions, religious or philosophical beliefs, sexual preference, or trade union membership. From the U.S. perspective, sensitive data also includes other information such as medical and financial data concerning the individual. Online businesses must be aware of these differences when practicing commerce over the Internet lest they violate privacy laws in other jurisdictions and lay themselves open to legal action or lawsuits. The Safe Harbor Agreement between the United States and European Union (EU) is one example of a mechanism

to minimize the impact of the different ways PII is handled by different economies around the globe.

See Also: personally identifiable information (PII), privacy, Safe Harbor Agreement

SERPENT

A block cipher developed by Ross Anderson, Eli Biham, and Lars Knudsen.

Overview

SERPENT is a block cipher that supports variable key length and block size. Using a 128-bit block size and 256-bit key length, SERPENT performs 32 encryption rounds to transform the plaintext into ciphertext. Other supported key lengths are 128 and 192 bits. SERPENT was a candidate for the Advanced Encryption Standard (AES) but lost out to Rijndael. SERPENT is now in the public domain.

For More Information

Visit www.cl.cam.ac.uk/~rja14/serpent.html for more information.

See Also: Advanced Encryption Standard (AES), block cipher, Rijndael

server certificate

A digital certificate installed on a server.

Overview

A server certificate is a digital certificate issued by a certificate authority (CA) so that clients connecting to the server can verify the server's identity. Server certificates are commonly used to validate Web servers running e-commerce sites so clients visiting the site can know they are providing their credit card information to the site they believe they are visiting instead of a malicious Web server impersonating the site. Server certificates are an integral component of the Secure Sockets Layer (SSL) protocol that allows confidential transmission of information over the Internet.

Server certificates for use with SSL can be obtained by submitting a certificate request to a public CA such as Verisign, Entrust, or Thawte, which issues the server the certificate and public/private key pair that makes SSL

communications possible with customers. These public CAs generally charge a fee for issuing a server certificate; the fee generally varies in accordance with the strength of the encryption keys generated.

Server certificates are used in other systems in which public key cryptography is used for authentication and encryption purposes. An example is the Wireless Application Protocol (WAP), in which wireless access points use server certificates to verify their identity to wireless clients accessing the network through them. In an enterprise environment, administrators also can issue their own server certificates to provide authentication, encryption, and data integrity in communications. Certificate Services, a component of Microsoft Windows 2000 Server and Windows Server 2003, can be used to implement the various elements of a Public Key Infrastructure (PKI) in an enterprise networking environment.

See Also: *certificate authority (CA), digital certificate, public key cryptography, Public Key Infrastructure (PKI), root certificate, Secure Sockets Layer (SSL)*

server-gated cryptography (SGC)

An extension of Secure Sockets Layer (SSL).

Overview

Server-gated cryptography (SGC) is an extension to SSL that enables financial institutions with export versions of Internet Information Services (IIS) to employ strong 128-bit encryption. Although SGC capabilities are built directly into IIS, a special SGC certificate is still required to use SGC. If you configure SGC on IIS, however, users trying to establish a secure communications channel with IIS must use a Web browser capable of communicating using a 128-bit session key.

See Also: *Secure Sockets Layer (SSL)*

service account

An account used as a security context for running a service.

Overview

Services are processes that run in the background waiting for other processes to request actions they can perform. An example is the Alerter service, which notifies selected users and computers of administrative alerts on Microsoft Windows NT and later versions of the platform. Since services are processes, they require a security context within which to run, and that context is provided by a service account, which can be either a built-in account called a special identity or an ordinary user account set apart for this purpose. On Windows NT and Windows 2000 platforms, most services run under the context of the LocalSystem account, a highly privileged built-in account that has all privileges for performing operating system tasks. On Windows Server 2003 platform, some services such as the Dynamic Host Configuration Protocol (DHCP) and Domain Name Service (DNS) client services now run using the less-privileged NetworkService account for greater security.

See Also: *special identities*

service pack (SP)

A cumulative set of hotfixes that can be applied in a single operation.

Overview

Service packs (SPs) are periodically released by software vendors as a way of assisting administrators in the process of keeping their systems securely patched and up-to-date. When new vulnerabilities are discovered in software applications, vendors usually quickly release a patch or hotfix that can correct the problem and prevent an exploit from compromising systems running the application. Not all vulnerabilities are serious, however, since some exploits can be performed only under exceptional conditions, which make them highly unlikely to occur. Other times bugs that are discovered are of a less serious nature and affect the performance of certain components under usual conditions without constituting a security threat. Software vendors usually respond to the varying levels of bugs and vulnerabilities by prioritizing hotfixes and publicly releasing only those that correct critical issues that could seriously

impact the security or reliability of their products. The remaining hotfixes not released for general use are then rolled together and incorporated as part of the next SP for the product.

SPs generally are cumulative and contain all critical and noncritical hotfixes since the original release of the product. For example, if a user has not applied SP 1 or 2, the user can still apply SP 3, which fixes everything that SP 1 and 2 did and more. SPs typically include product fixes for issues in the areas of performance, reliability, security, and compatibility. They might also contain new features and enhanced versions of old components to provide users with better functionality and improved manageability. Service packs also sometimes include components for compliance with federal laws; for example, to ensure compliance with revised encryption export controls.

Notes

The terms **service release** and **integrated service pack** refer to the combination of the original product together with the SP released in a single package.

See Also: *hotfix, patch, Software Update Services (SUS), Windows Update*

Service Release (SR)

A service pack (SP) that includes the original product.

Overview

An SP is a cumulative set of hotfixes that can be applied in a single operation. Software vendors typically distribute Service Releases (SRs) to original equipment manufacturers (OEMs) to support new hardware. SRs also might be released to corporate customers to simplify the job of performing clean installs of updated versions of products instead of having to install the original product and then apply an SP to bring it up-to-date. For example, Service Release 1A (SR-1A) for Microsoft Office 2000 was created to allow corporate customers and OEMs to install an updated version of Office 2000 as quickly and easily as possible.

See Also: *hotfix, service pack (SP)*

session hijacking

Short for TCP session hijacking, an exploit in which an attacker takes control of one end of a Transmission Control Protocol (TCP) session.

See: *TCP session hijacking*

session key

A key that is used for a single session and then discarded.

Overview

Session keys are used in most real-world implementations of secret key encryption systems to enhance the security of such systems. In a traditional secret key system, a trusted mechanism such as a courier or public key cryptography is used to distribute a secret key to the parties who need to communicate secretly. Once the parties share this secret, they can use it for encrypting and decrypting messages sent between them.

A weakness with this approach is that if an attacker can obtain a portion of ciphertext through eavesdropping and then somehow also obtain (or guess) the corresponding plaintext associated with the ciphertext, the attacker may be able to crack the key and decode future messages sent between the parties. The problem is compounded if the secret never changes, since the more ciphertext an attacker can obtain through eavesdropping, the more material the attacker has to work with in trying to guess the key. If, however, a new secret is generated and shared between the parties each time they need to communicate, cracking a key for one message is useless for decoding subsequent messages since the key is no longer the same.

Another advantage of using session keys as opposed to long-term secrets is to ensure the security of archived information. If a different session key is used for encrypting each file in an encrypted file storage system, cracking one session key provides access to only one file's contents. One more reason for using session keys is for situations in which the other party is someone you're not sure you would trust with a long-term secret.

Generally, session keys usually are created using a pseudorandom number generator (PRNG), an application for generating a string of apparently random numbers or characters, though some cryptographic systems derive session keys from hashing algorithms instead. For even greater security, some cryptographic systems even change the session key several times during a single communication session, even as far as using a new session key to encrypt each packet in a transmission.

See Also: *ciphertext, hashing algorithm, plaintext, pseudorandom number generator (PRNG), public key cryptography, secret key encryption*

SET

Stands for Secure Electronic Transaction, a family of specifications for secure credit card transactions over the Internet.

See: *Secure Electronic Transaction (SET)*

SGC

Stands for server-gated cryptography, an extension of Secure Sockets Layer (SSL).

See: *server-gated cryptography (SGC)*

SHA-1

Stands for Secure Hash Algorithm-1, a hashing algorithm for generating a message digest.

See: *Secure Hash Algorithm-1 (SHA-1)*

SHA-2

An umbrella designation for variants of the Secure Hash Algorithm-1 (SHA-1).

Overview

SHA-2 is sometimes used to designate any variant of SHA-1 that generates a message digest longer than the 160-bit digest created by applying SHA-1 to an arbitrary message. The most popular variants of SHA-1 are as follows:

- **SHA-256:** Generates a 256-bit message digest

- **SHA-384:** Generates a 384-bit message digest
- **SHA-512:** Generates a 512-bit message digest

There are significant differences in the internal operation of these variant algorithms compared with the original version SHA-1.

See Also: *Secure Hash Algorithm-1 (SHA-1)*

shadow password file

A file used for implementing password shadowing, a technique used on UNIX platforms for hiding the location of passwords.

See: *password shadowing*

shared secret

Another name for a secret key, a key used in secret key encryption.

See: *secret key*

share-level security

Protecting shared resources using only a password.

Overview

Share-level security involves using a password to control access to shared resources on a network. Since any user who knows the password can access the share, share-level security affects all users the same way regardless of what rights or privileges they possess through group or role membership. For example, a share that has Read permission assigned to it can be read by any user who knows the password for the share, and no users are therefore able to write to the share.

Share-level security is different from user-level security, in which permissions are assigned to individual users and the groups to which they belong. User-level security is more granular than share-level security since different users can be assigned different levels of access to the resource. User-level security is a feature of file systems such as the NTFS file system supported by Microsoft Windows NT and later versions of the platform. As a result, user-level security can control access to resources both locally and remotely over the

network, while share-level security can control access only for remote network users.

Share-level security is supported by all versions of Windows and by non-Windows applications such as Samba. When share- and user-level permissions are combined on systems running Windows, the result is sometimes called effective permissions.

See Also: *password, permissions, user-level security*

Shiva PAP (SPAP)

An enhanced version of Password Authentication Protocol (PAP) developed by Shiva Corporation.

Overview

PAP is a remote access authentication protocol used for authenticating Point-to-Point Protocol (PPP) communication sessions. PAP is inherently insecure, however, since it transmits the user's credentials (user name and password) over the connection in cleartext. Shiva PAP (SPAP) is a proprietary version of PAP patented by Shiva that enhances PAP security using two-way encrypted authentication. SPAP is more secure than PAP but not as secure as Challenge Handshake Authentication Protocol (CHAP) and Microsoft Challenge Handshake Authentication Protocol (MS-CHAP), which do not transmit any passwords over a connection but use encrypted challenge strings instead.

See Also: *Password Authentication Protocol (PAP)*

ShowAcls

A *Windows 2000 Resource Kit* tool for displaying NTFS permissions.

Overview

ShowAcls is a security tool that can be used to display the NTFS permissions assigned to files and directories on NTFS volumes. In particular, ShowAcls can be used to enumerate which groups a user belongs to and to match the user's security identifier (SID) and group SIDs to the SIDs for an access control entry (ACE). While ShowAcls is a legitimate security auditing and troubleshooting tool, it also can be used for malicious

purposes if an attacker has gained entry to a system and has sufficient permissions to run the tool.

See Also: *access control entry (ACE), permissions, security identifier (SID)*

ShowPriv

A *Windows 2000 Resource Kit* tool for displaying privileges granted to users and groups.

Overview

ShowPriv is a security tool that can be used to display the user rights and privileges assigned to user and group accounts. ShowPriv uses a command-line syntax that displays the accounts that have a specified privilege such as the right to remotely shut down the system or the right to back up files on the system. While ShowPriv is a legitimate security auditing and troubleshooting tool, it also can be used for malicious purposes if an attacker has gained entry to a system and has sufficient permissions to run the tool.

See Also: *rights*

SHS

Stands for Secure Hash Standard, the Federal Information Processing Standard (FIPS) defining the Secure Hash Algorithm-1 (SHA-1).

See: *Secure Hash Standard (SHS)*

S-HTTP

Stands for Secure Hypertext Transfer Protocol, an extension for Hypertext Transfer Protocol (HTTP) to allow secure transfer of files.

See: *Secure Hypertext Transfer Protocol (S-HTTP)*

Sid2user

A tool for obtaining the user name associated with a security identifier (SID).

Overview

SIDs are an integral part of security on Microsoft Windows platforms and are strings that uniquely

identify security principals on a Windows-based system or network. For example, each user account is assigned a SID when the account is first created, and the SID remains the same even if the account itself is renamed. From the perspective of internal processes and services on machines running on Windows, it is the SID that identifies the user, not the user name of the account.

Since user accounts are one of the main targets of attackers trying to compromise a network, obtaining the names of accounts stored on a machine can provide useful information for exploiting vulnerabilities. The Sid2user utility allows a user to obtain the name of an account based on knowledge of its SID. Using Sid2user together with its companion utility, User2sid, which allows a SID to be obtained for a given user name, an attacker may be able to compromise the security of a Windows system and obtain useful information about user accounts on the system.

Evgenii Rudnyi of Moscow State University developed both Sid2user and User2sid based on published information of Windows application programming interfaces (APIs). Both utilities can be used against remote machines running on Windows NT or later without needing authentication if null sessions can be established with target machines. One way of preventing attacks that use these tools is to block Transmission Control Protocol (TCP) port 139, which is used for NetBIOS session enumeration.

See Also: null session attack, security identifier (SID), security principal

SIIA

Stands for Software & Information Industry Association, a trade association for the software industry.

See: Software & Information Industry Association (SIIA)

single sign-on (SSO)

Any system for authenticating users for a wide range of resources using a single set of credentials for each user.

Overview

Single sign-on (SSO) systems can be implemented in heterogeneous networking environments to make it simpler for users to be authenticated for a wide variety of platforms and resources. For example, instead of having to remember one password for a Microsoft Windows-based network, one for a UNIX network, one for mainframe access, and so on, SSO can allow a user to enter a single user name and password to gain access to all systems on the network. SSO systems also can be implemented on the Internet to help manage the wide range of credentials users require for accessing e-commerce sites and other sites that require users to register before using them.

While SSO makes it simpler for users to manage their credentials and be authenticated by diverse systems, SSO can also represent a single point of failure for user authentication and a target for attack.

Marketplace

Microsoft .NET Passport is one example of a Web services-based SSO system that can be used for both network and Internet authentication purposes. The Liberty Alliance Project is an example of a similar system, and a wide variety of packaged and custom SSO solutions also are available from different vendors.

See Also: authentication, Liberty Alliance Project, .NET Passport

Sircam

A notorious mass-mailer worm.

Overview

The Sircam worm first appeared in 2001 and rapidly spread across the Internet, infecting both home users and business networks with computers running on Microsoft Windows platforms. The worm initially infects machines through an attachment to an e-mail message that has a random subject name and a message body saying, "Hi! How are you?" The attachment has a double extension so that it initially appears to the user to be an image file, Word document, or other harmless file, whereas in reality it is the executable payload of the virus. Once the virus is executed, the worm replicates

itself to several directories on the machine and modifies the registry so that it is executed again upon startup. The worm then extracts e-mail addresses from the Windows address book and other locations and uses its own mailing engine to send out copies of itself to the first 100 addresses it finds.

What made the worm especially troublesome for businesses is that once Sircam infected a machine by e-mail, the worm then scanned the network to enumerate a list of network shares and then tried to write itself to these shares to infect other systems on the network. Once the worm spread across a network, it had to be eradicated from every machine on the network or reinfection easily could have occurred. The worm could have caused even more damage than it did, since it was designed to trigger a payload that could cause the hard drives of some systems to be erased, but fortunately there was a coding error in the worm that prevented this from working.

See Also: worm

site certificate

Another name for server certificates and certificate authority (CA) certificates.

Overview

Digital certificates for servers and CAs are sometimes together called site certificates. A server certificate identifies the server (usually a Web server) presenting the certificate to clients who visit the site. Similarly, a CA certificate identifies the CA that issued the server its own server certificate. Because both types of certificates are involved in using Secure Sockets Layer (SSL) for secure communication with Web sites, they are collectively known as site certificates.

See Also: certificate authority (CA), digital certificate, server certificate, Secure Sockets Layer (SSL)

Six/Four

A technology for circumventing attempts to censor traffic on the Internet.

Overview

Six/Four is a peer-to-peer (P2P) protocol developed by “Mixer,” a German hacker associated with the black hat group Hactivismo. Six/Four allows the creation of virtual private networks (VPNs) that are set up and managed in a decentralized fashion for privately relaying Internet content to users who cannot access it by normal means because of censorship activities of governments and service providers.

Six/Four is described as privacy-enhancing software and uses a generic tunneling protocol that is application independent and thus can relay any form of content. Routing across a Six/Four VPN is anonymized, and the topography of connections is established on an ad hoc basis. Content being transferred is encrypted to ensure its privacy using preauthenticated keys located on trusted peers. When a user wants to access restricted content anonymously, the user uses Six/Four to secretly connect to a trusted peer outside the restricted network, which can use tunneled connections to communicate with other trusted peers until finally content is obtained using normal Internet protocols from the server hosting the requested content.

Notes

The name **Six/Four** is based on the date June 4, which in 1989 was when Chinese authorities began cracking down on activists protesting for democracy in Tiananmen Square.

See Also: Peekabooby Project, Publius Project, virtual private network (VPN)

S/Key

A one-time password (OTP) system developed by Bell Communications Laboratories (Bellcore).

Overview

OTP systems are authentication schemes that require a new password from the user each time authentication must be performed. OTPs are used in high-security environments as a means of preventing possible eavesdropping on open network connections since it is impossible for an attacker to mount a replay attack to

capture and replay authentication traffic and hijack a session if each session requires a unique password.

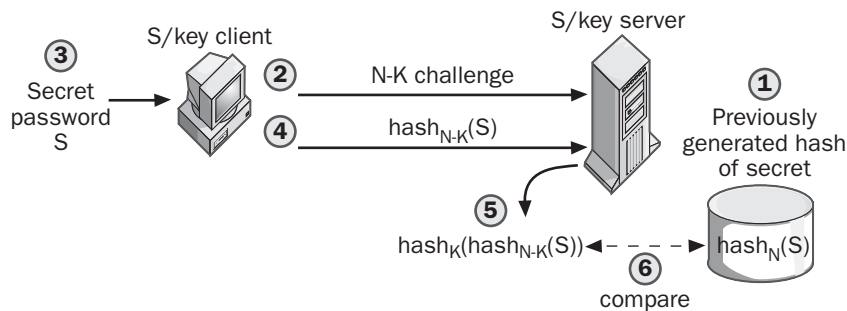
Implementation

To use S/Key the user first securely submits a secret password and a number N to a local S/Key server (that is, not over a network connection). This secret password is typically a phrase or group of words that is easy to memorize and is used to initialize the S/Key system for the user. The S/Key server initializes the S/Key system for the user by taking this secret, adding a random seed value, and applying the message digest 4 (MD4) hashing algorithm to the result N times, and then it securely stores the resulting hash for later authentication of the user.

Now when the user wants to log on to the network, the S/Key server issues the number $(N - 1)$ to the user as an authentication challenge. The user enters his or her secret password to S/Key client software on the computer,

hashes the secret $(N - 1)$ times, and sends the result to the S/Key server. This hash represents a OTP and appears totally random to anyone eavesdropping on the network. When the S/Key server receives the hash, the server hashes it one additional time (since $N - 1 + 1 = N$) and compares the result with the hash it previously stored for use, and if the two match, the user is granted access to the network. The next time the user tries to log on, the challenge $(N - 2)$ is issued and so on until either no more hashes are available for authentication and the user is locked out of the network or the user applies to the server again to generate a new series of OTPs.

S/Key can be used for more than just network logon and can be implemented wherever secure authentication is required. For example, you could secure File Transfer Protocol (FTP) or Telnet using S/Key by developing suitable software. A modified version of S/Key is standardized by RFC 1938.



S/Key. How S/Key authentication works.

Marketplace

S/Key software is available for a variety of platforms, including the Microsoft Windows, UNIX/Linux, and Macintosh operating systems. The FreeBSD operating system includes a variant of S/Key called Onetime Passwords In Everything (OPIE) that uses the stronger message digest 5 (MD5) algorithm.

See Also: *one-time password (OTP)*

Skipjack

A secret key encryption algorithm developed by the National Security Agency (NSA).

Overview

The Skipjack algorithm was designed to be efficient for implementation in tamper-resistant cryptographic hardware devices such as Fortezza cards used for authentication in high-security environments. The algorithm also was implemented in the Clipper Chip technology for providing backdoor access to cryptographic communications. Skipjack was developed in 1993 and initially was classified as SECRET, but in 1998 the algorithm was unclassified and its details were published by the National Institute of Standards and Technology (NIST).

Implementation

Skipjack is a symmetric algorithm based on an electronic codebook (ECB) approach that transforms 64-bit blocks of plaintext into same-sized blocks of ciphertext. The secret key used by the algorithm is 80 bits long, and the transformation process involves 32 rounds of permutations and other mathematical operations. The algorithm is parameterized to operate in any of four different modes.

See Also: *electronic codebook (ECB), National Institute of Standards and Technology (NIST), National Security Agency (NSA), secret key encryption*

Slammer

A notorious worm that affected Microsoft SQL Server.

Overview

Slammer, also known as “SQL Slammer” and “Sapphire,” appeared in January 2003 and spread across the Internet at a phenomenal rate, bringing down large portions of the Internet in less than 10 minutes after it appeared. Slammer is by far the fastest-propagating worm to appear, and it exploited a known vulnerability in SQL Server for which a patch had been issued months before. Unfortunately, large numbers of administrators had failed to apply the patch, and at least 75,000 SQL servers were affected by the worm. Meanwhile, vast amounts of scanning traffic from worms looking for new hosts to infect resulted in problems with airline reservation systems and even crashed banking machines. The mechanism that enabled the worm to propagate so rapidly (almost 100 times faster than CodeRed, an earlier worm) is a scanning engine that selects Internet Protocol (IP) addresses at random to scan for the presence of SQL Server.

See Also: *CodeRed, worm*

Slashdot Effect

A denial of service (DoS) condition that results when too much interest is generated concerning a Web site.

Overview

The term **Slashdot Effect** derives from a series of cases that involved articles posted on Internet news sites such

as Slashdot, a popular site publishing “news for nerds.” What typically happens is that one or more news sites publish articles describing something on some other Web site, which we might call the “victim” site. If the articles are especially timely and interesting, large numbers of readers start visiting the victim site, and its Web server is suddenly overwhelmed with so many requests that it hangs or crashes. The result is the same as if an attacker deliberately targeted the victim with a DoS attack, but in this case the effect is usually incidental and unintentional on the part of those publishing the articles that caused the effect.

In order for the Slashdot Effect to really occur, two things are required:

- The news sites must have a large and eager readership that visits them frequently to check for new stories.
- The news sites must update their sites frequently with unique and interesting stories about other Web sites on the Internet.

Slashdot was probably used for naming this effect because of the intense interest that this site has for many Internet-savvy users, but the effect has been observed to occur as a result of content posted by a wide range of news media sites.

See Also: *denial of service (DoS)*

smart card

A plastic card with an embedded microchip used for access and authentication.

Overview

Smart cards are credit card-sized devices that have embedded memory or microprocessors. They are primarily used as physical authentication tokens for gaining access or logging in to systems and networks. Logging in with a smart card can be as simple as swiping the card through a reader or terminal. A smart card reader is a peripheral that can be attached to a standard PC using a serial or universal serial bus (USB) connection and used to read and/or write to smart cards. A terminal is a stand-alone smart card-swiping device that is connected to a network or telecommunication system

for reading credentials from smart cards. In addition to swiping the card, the user usually has to provide a password or personal identification number (PIN) to complete the authentication process since, otherwise, a stolen card would provide an intruder with a means of accessing the network.

Smart cards are widely used in high-security environments, in the banking and financial industry, in laboratories, and in the transportation industry. Depending on the purpose of their use, smart cards can contain personally identifiable information (PII), user credentials and passwords, encryption keys for public key cryptography, and other information.

Smart cards are governed by several standards including the ISO 7816 family of standards from the International Organization for Standardization (ISO) and the FIPS 140-1 Federal Information Processing Standard (FIPS) from the U.S. National Institute of Standards and Technology (NIST).

See Also: *authentication, Federal Information Processing Standard (FIPS), National Institute of Standards and Technology (NIST), personal identification number (PIN), personally identifiable information (PII)*

SMBRelay

A backdoor Trojan exploiting Server Message Block (SMB) protocol.

Overview

SMBRelay is a Trojan that exploits the SMB port 139 used by NetBIOS. Once installed on a compromised system, SMBRelay is able to take SMB traffic received by the compromised host and relay it to the attacker's host, providing the attacker with access to sensitive information that can be used for further exploits. For example, on Microsoft Windows systems, SMBRelay can cause the target host to try to authenticate with the attacker's machine. This can provide the attacker with an NTLM password hash that the attacker can then try to crack using L0phtcrack or some other password-cracking program.

SMBRelay exists in several variants, and its activity can be blocked by implementing SMB signing or by blocking the following ports on your firewall:

- User Datagram Protocol (UDP) ports 137 and 138
- Transmission Control Protocol (TCP) ports 139 and 445

See Also: *backdoor, L0phtcrack, SMB signing, Trojan*

SMB signing

A secure version of Server Message Block (SMB) protocol.

Overview

SMB signing helps protect network hosts against exploits targeting NetBIOS and other SMB-based services. Traditional SMB authentication is vulnerable to man-in-the-middle (MITM) attacks, but by implementing mutual authentication SMB signing thwarts this kind of attack. SMB signing strengthens SMB authentication by adding digital signatures into SMB packets so that both the client and the server can verify the authenticity of the packets. SMB signing is supported by Windows NT 4.0 Service Pack 3 and later, but the feature is not enabled by default. While SMB signing prevents exploits such as the SMBRelay Trojan from being effective, it does so at an added performance cost because of the increased processing and network traffic required to digitally sign each SMB packet.

See Also: *man-in-the-middle (MITM) attack, SMBRelay, Trojan*

S/MIME

Stands for Secure/Multipurpose Internet Mail Extensions, an e-mail security standard that uses public key encryption.

See: *Secure/Multipurpose Internet Mail Extensions (S/MIME)*

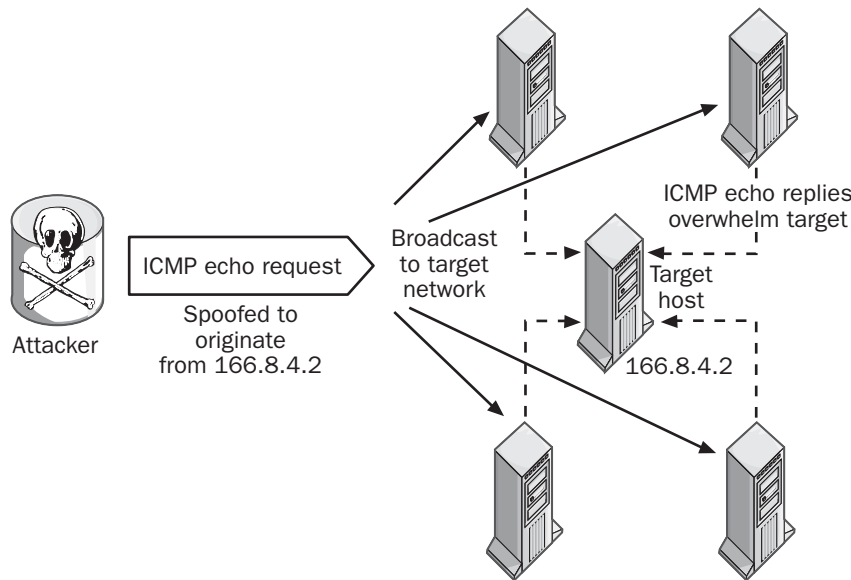
Smurf attack

A classic denial of service (DoS) attack exploiting Internet Control Message Protocol (ICMP).

Overview

The Smurf attack is based on spoofing ICMP echo request packets, the packets used by the ping network-troubleshooting utility. The attacker begins by forging ICMP echo request packets with source Internet Protocol (IP) addresses the same as the target's IP address. These packets then are broadcast onto the network where the target resides, and other hosts that receive them respond by sending a flood of ICMP echo reply

packets to the target host. The result can overwhelm the target, creating a DoS condition that can prevent legitimate users from accessing services on the target. Some target systems might even hang up or crash, and the flood of ICMP traffic can cause network congestion that can affect other hosts on the network as well. If the attack is directed toward a service provider network, it can cause degradation of network performance and affect services provided to the provider's clients.



Smurf attack. How a Smurf attack works.

The Smurf attack first appeared in 1997 and gained media attention when Yahoo! Web servers experienced three hours of outage as targets of the attack. Part of the insidiousness of a Smurf attack is that the attacker doesn't need to generate a large amount of traffic to initiate the attack, since, because of the amplification factor of using an attack based on broadcast packets, even one spoofed ICMP echo request can generate a large amount of ICMP echo replies if the target network has a sufficient number of hosts. In addition to Web servers, when it first appeared the attack also targeted Internet Relay Chat (IRC) servers.

Some of the methods used for countering the Smurf attack include the following:

- Disabling IP-directed broadcast traffic on perimeter routers
- Disabling hosts from responding to ICMP packets sent to broadcast addresses
- Ingress filtering on perimeter routers to prevent spoofed traffic from entering the network

Notes

A related DoS attack called Fraggle is similar to Smurf but uses User Datagram Protocol (UDP) packets instead of ICMP.

See Also: *denial of service (DoS)*

sniffing

Capturing and analyzing network traffic.

Overview

The term **sniffing** commonly is used to describe protocol analysis, the process of viewing and analyzing the contents of packets on a network. Tools used for this purpose are properly called protocol analyzers, but the term **sniffer** also is widely used for such tools despite the fact that “Sniffer” is a trademark of Network Associates, a leading provider of tools for managing and troubleshooting networks.

Sniffing networks can be done for legitimate reasons to profile the flow of network traffic or troubleshoot communication problems. It is also commonly performed for malicious reasons to capture passwords or credit card information, enumerate services available on targets, and other purposes. Sniffing is easy on a shared network where segments are connected by hubs. If you have physical access to a network, you can connect a protocol analyzer to the network and capture everything happening on the segment. For remote sniffing you can install a sniffing program on a network host whose network card is configured to run in promiscuous mode, and then have the program capture traffic and send it to your console, a method often used by attackers. Sniffing switched networks is more difficult since each port is its own network segment, but most modern switches include a monitor port to which you can connect a sniffer to capture traffic on the switch’s backplane.

The best way to defeat sniffing is to encrypt network traffic using Internet Protocol Security (IPSec) or some other mechanism. This won’t stop attackers from sniffing your network, but it will make it harder for them to gain anything useful from their efforts.

Marketplace

Some popular sniffers frequently used by security professionals and black hat hackers alike include Dsniff, Tcpdump, Ngrep, Ethereal, Sniffit, and Snort for UNIX/Linux platforms, and Windump, Ethereal, Etherpeek, Network Associates’s Sniffer, and Microsoft Network Monitor for Microsoft Windows platforms.

Notes

Another name for sniffing is **packet snooping**.

See Also: *Dsniff, Internet Protocol Security (IPSec), protocol analyzer, Tcpdump, Windump*

Snort

A popular open source intrusion detection tool.

Overview

Snort is a simple but effective network intrusion detection system (NIDS) that can be used to analyze Internet Protocol (IP) traffic in real time, looking for evidence of intrusion. Snort can detect a wide range of exploits, including stealth scans, Server Message Block (SMB) probes, stack fingerprinting, and more. The tool includes a flexible rule-based architecture and an extensible detection system that can incorporate custom modules for enhancing its capabilities. Snort also can analyze the flow of Hypertext Transfer Protocol (HTTP) traffic and perform stateful pattern matching.

Snort is available for a wide range of platforms, including most versions of UNIX/Linux and Microsoft Windows platforms. The current version, Snort 2, is licensed under the General Public License (GPL).

For More Information

Visit www.snort.org for more information.

See Also: *network-based intrusion detection system (NIDS)*

social engineering

Using persuasion or deception rather than technology to compromise information system security.

Overview

Human psychology often is more effective at compromising network security than technological tools and exploits. A **social engineer** is a malicious hacker who uses persuasion, deception, fraud, and espionage to obtain access to information systems and overcome network security protection measures. Social engineering basically can take place by several means:

- **Physical means:** By “Dumpster diving” (sifting through company trash), an attacker often can obtain useful information that can help compromise a network or support further social-engineering efforts. Discarded system manuals and other technical documentation often can help the attacker map the network for targeted intrusion. Memos and other communications can provide clues concerning company security policy and how it might be circumvented. Some employees even print out their e-mail in order to read it, and discarded e-mail can provide clues for user names for logging on to networks. And company organizational charts and phone directories can provide social engineers with identities they can impersonate to engineer physical intrusion.
- **Psychological means:** Posing as frustrated users who are having problems with their computers, social engineers can contact help desk personnel and take advantage of the fact that such personnel have been trained to be helpful in such situations and may provide new passwords or other information useful for network intrusion. Walking through the front door dressed as maintenance workers, for instance, social engineers often can persuade receptionists and security personnel to give them access to work areas; once entry is gained, social engineers can check under keyboards for stickies with passwords written on them, try to find a desk that a hapless employee has momentarily left without locking the workstation, or watch over users’ shoulders as they enter passwords to access their computers. As a result of a social engineer’s use of impersonation, persuasion, ingratiation, and friendliness, it is amazing how many people will violate company

security policy and provide sensitive information to total strangers.

For More Information

Kevin Mitnick, a master of social engineering who spent four years in prison for malicious hacking exploits, has written an instructive account of what social engineering is about in his book *The Art of Deception: Controlling the Human Element of Security* (John Wiley & Sons, 2002).

See Also: *hacking*

SOCKS

A generic proxy protocol for Internet Protocol (IP) networks.

Overview

SOCKS, also known as Authenticated Firewall Transfer (AFT), is a protocol used in proxy servers and firewalls and for virtual private networks (VPNs). SOCKS is implemented as a client/server protocol to allow hosts outside a firewall to access hosts inside the firewall without needing a direct IP connection between the two hosts. Since SOCKS works at the transport layer, its operation is independent of what happens at the application layer, so SOCKS can proxy new applications without having to reconfigure the firewall or install new software.

SOCKS comes in two versions:

- **SOCKSv4:** Establishes proxy circuits and connection requests and relays application data between proxied hosts
- **SOCKSv5:** Extends SOCKSv4 by adding authentication for greater security

SOCKS is an Internet standard protocol defined in RFCs 1928, 1929, and 1961.

For More Information

Visit www.socks.permeo.com for more information.

See Also: *firewall, virtual private network (VPN)*

Software & Information Industry Association (SIIA)

A trade association for the software industry.

Overview

The Software & Information Industry Association (SIIA) is an alliance of over 1200 software vendors worldwide that engages in activities related to ensuring the growth and expansion of the software industry. One of its most important efforts is trying to prevent software piracy, a rapidly growing problem that costs the software industry worldwide tens of billions of dollars each year. The SIIA works to protect the rights of its members and is an advocate for legislation to stop piracy and protect intellectual property rights. The SIIA runs an antipiracy hotline (800-388-7478) that individuals can use to report real or suspected cases of software piracy in retail, corporate, or Internet environments.

For More Information

Visit www.siiia.net for more information.

See Also: *software piracy*

software piracy

Using software in any fashion that violates its license agreement.

Overview

Software piracy is a violation of intellectual property rights. When you purchase software, you agree to comply with the licensing agreement included with the software, and violations of this agreement can result in criminal and civil penalties. Software piracy can take many forms, including the following:

- Illegally copying software by burning a copy for a friend or colleague
- Installing single-use software on more than one computer (sometimes called “softlifting”)
- Posting files for commercially licensed software on the Internet for others to download
- Underreporting the number of computers on which commercial software has been installed
- Duplicating software and selling it fraudulently as legitimate software
- Selling not-for-resale samples of software received for testing and evaluation purposes
- Selling Original Equipment Manufacturer (OEM) “backup copies” of software included with preinstalled systems (called “OEM unbundling”)
- Renting software for temporary use by others (violates the U.S. Software Rental Amendments Act of 1990)

Microsoft provides businesses and consumers with a number of tools to ensure they are not willingly or unwillingly using pirated software or violating licensing agreements, including special identification features on legitimate product CDs, the Microsoft Software Inventory Analyzer, and various licensing guides explaining different licensing programs available. Software piracy is a growing problem that costs the United States and other national economies billions of dollars each year.

For More Information

Visit www.microsoft.com/piracy for more information.

See Also: *Software & Information Industry Association (SIIA)*

Software Update Services (SUS)

A tool for keeping critical software updates up-to-date on Microsoft Windows 2000, Windows XP, and Windows Server 2003.

Overview

Software Update Services (SUS) is designed to help administrators ensure their networks are secure and reliable by automatically deploying critical updates, security updates, and security rollups to computers on a network. SUS can distribute patches to both servers and desktop computers, and it consists of a server component that runs on a computer implementing Windows 2000 Server or Windows Server 2003 and Automatic Updates agent software that runs on the target computers

that receive the updates. SUS servers periodically download the latest critical updates from Windows Update and then distribute them according to a schedule defined by administrators. SUS can run in either a workgroup or domain scenario, and multiple SUS servers can be used to scale for enterprise networking environments.

SUS can be used only to apply hotfixes and rollups for Windows operating systems. It cannot be used to distribute service packs or device drivers, and it cannot be used to provision updates for other Microsoft server products such as Microsoft Exchange Server or Microsoft SQL Server. For a more powerful and scalable tool to manage updates for the complete range of Microsoft products across an enterprise, use Microsoft Systems Management Server (SMS) instead.

For More Information

Visit www.microsoft.com/windows2000/windowsupdate/sus/ for more information.

See Also: *hotfix, Windows Update*

source routing

A method for specifying ahead of time the route a packet should take when traversing a routed internetwork.

Overview

Normally, when an Internet Protocol (IP) packet is sent toward a remote host on an internetwork, routers along the way make the decisions concerning which route or path the packet should take to reach its destination. Source routing is a feature built into the IP protocol that allows the packet's sender to include information in the packet that determines the path the packet takes to reach its target. Source routing can be either of the following:

- **Strict source routing:** The path the packet should take is completely predetermined by the sender.
- **Loose source routing:** Only portions of the path the packet should take are determined by the sender.

Source routing was built into the IP protocol for legitimate purposes, including these:

- Forcing traffic to take an alternate path to avoid network congestion or a downed link
- Troubleshooting problems with the flow of traffic between two remote networks

Source routing also can be exploited by attackers, however, for mapping networks in preparation for an attack and for impersonating hosts on a network in order to hijack sessions. To prevent malicious use of source routing, most operating systems allow you to disable this feature, and firewall filters can usually be configured to block source-routed traffic as well.

Source routing is described in RFC 791.

See Also: *Tcp_wrapper, Traceroute*

SP

Stands for service pack, a cumulative set of all hotfixes that can be applied in one step, for example SP1, SP2, and so on.

See: *service pack (SP)*

spam

No longer just a form of luncheon meat.

Overview

Spam has become the bane of the Internet, and still there is no real solution in sight. Spam is usually defined as “unsolicited e-mail” and resembles the flyers from stores that clog your postal mailbox each morning, but it's much more than that. Spam, depending on who is discussing the topic, can include the following:

- Unsolicited e-mail from legitimate businesses trying to market their products to a wider audience
- Bulk e-mail from organizations announcing new services of various kinds that recipients might be interested in
- Fraudulent e-mail trying to perpetrate various get-rich schemes that actually make the perpetrators rich instead of the recipient

- Offensive e-mail inviting readers to purchase memberships in pornography sites
- Sneaky e-mail that includes scripts or programming code so that, when the recipient opens it, the sender is notified, allowing the sender to target the sender's address for an even greater volume of e-mail

A more formal way of defining **spam** is any form of e-mail that tries to hide its originating e-mail address to make it hard to trace the sender or that uses deception in the subject line to try to induce the recipient to open the message. A simple way of defining it with which almost any e-mail user will agree is, "I never asked to be on their mailing list!"

Issues

The problem of spam is partly an issue of free speech and partly the nature of the Internet: a distributed system that no one really controls and which developed through a series of consensus decisions on technical issues. Spam crosses national borders easily and often is difficult to distinguish from legitimate e-mail. The result is that spam-filtering software is only limited in its effectiveness, while the number of "spammers" and their ingenuity seems to be growing at an exponential rate. In 1999, a survey suggested that almost 25% of most business e-mail messages were spam; in 2003, that number shot up to 75%. In only a few years e-mail may become a useless tool for legitimate communications unless something is done to stop the trend.

Domain Name System (DNS) blacklists are one way of fighting spam. If a mail server on the Internet is configured as an open mail relay so that it forwards anything it receives, the mail server quickly gets placed on one of many blacklists available across the Internet. The problem is that occasionally legitimate mail servers get placed on these blacklists and their mail suffers havoc while the administrator tries to convince the blacklist to let them off the hook.

Another way of fighting spam is to hide your e-mail address from public display so that it won't be added to spammers' mailing lists when they use various "address-harvesting" techniques to cull valid e-mail addresses from newsgroups, Web pages, and other

sources. The problem is, you usually want your e-mail address to be public so legitimate people can contact you, and using myname@NOSPAM.mydomain.com in your e-mail address can prevent unsophisticated users from knowing how to contact you.

User education is another aspect of fighting spam, and by using online tools such as SamSpade.org, technically savvy users can track down the origin of unsolicited e-mail messages and contact their service providers or a law enforcement agency to complain. If enough individuals did this, it might make a difference, but the vast majority of e-mail users don't know how to use Trace-route or Whois and likely never will.

Law enforcement may be the only answer, or more specifically, lawsuits. Antispam legislation has begun to appear in different jurisdictions, and using these laws, attempts may be made to prosecute spammers in other locales since spam crosses all boundaries. Almost 9 out of 10 business users in one survey supported legislation against spam.

For More Information

Visit www.unc.edu/courses/pre2000fall/law357c/cyberprojects/spring99/spam/name.html to find out how spam got its name (and for a bit of a laugh as well)

See Also: mail relaying, Sam Spade

SPAP

Stands for Shiva PAP, an enhanced version of Password Authentication Protocol (PAP) developed by Shiva Corporation.

See: Shiva PAP (SPAP)

Spar

A free tool for auditing processing accounting on UNIX platforms.

Overview

Spar, which stands for "show process accounting records," is a free utility available for most UNIX platforms that is used to parse process-accounting records on systems for which process accounting has been

enabled. Spar can be used to enhance the security of UNIX systems by providing administrators with a way of determining whether rouge processes have been running on their machines. The typical approach to using Spar is to first use it to base line the processes running on a newly hardened system that has not yet been connected to the network or Internet. Once this process baseline has been established, regular use of Spar and comparison with the baseline can help administrators recognize behaviors that can indicate a compromised machine. Spar also can be used to identify discrepancies in the process reporting of other UNIX tools such as Ps and can verify security events by displaying the process and user associated with the event.

See Also: *auditing*

special identities

Well-known security principals managed by the operating system instead of administrators.

Overview

Microsoft Windows platforms include a number of built-in security principals called special identities. These principals behave similarly to groups since user accounts can assume their roles; the difference, however, is that administrators assign users to groups, but the operating system assigns users to special identities. Some of the more important special identities on Microsoft Windows 2000 and later include the following:

- **Anonymous Login:** Any user currently accessing a specific resource but who hasn't had his or her user account authenticated
- **Authenticated User:** All users who have had their accounts authenticated on the network using a valid account (does not include anonymous users)
- **Creator Owner:** The user who created or took ownership of the resource under consideration
- **Dialup:** All users currently connected to the network over a dial-up connection

- **Everyone:** All currently logged on network users, including guests and users from other domains
- **Interactive:** All users currently logged on to a specific computer and accessing a resource on that computer as opposed to over the network
- **Network:** All users currently accessing a specific resource over the network as opposed to interactively from the local console

See Also: *security principal*

spoofing

Forging packets so they appear to originate from a trusted host.

Overview

In a general sense, spoofing is any method for making a transmission appear to have come from someone other than the one who originated the transmission. Spoofing is a threat to information security because it violates the trust hosts use for establishing communication between themselves on a network. Some of the common forms of spoofing used by malicious hackers include these:

- **Internet Protocol (IP) address spoofing:** Falsifying the source addresses of IP packets
- **Address Resolution Protocol (ARP) spoofing:** Falsifying the Media Access Control (MAC) addresses of Ethernet frames
- **Domain Name System (DNS) spoofing:** Impersonating name servers by falsifying information in DNS packets.

See Also: *ARP spoofing* , *DNS spoofing*, *IP address spoofing*

spyware

A form of malware that installs itself for information leakage purposes.

Overview

Spyware is a term used to describe certain forms of "adware," software that installs itself on your system

without your knowledge and displays advertisements when you browse the Internet. While some adware simply might display advertisements in your browser periodically, spyware is more insidious because it collects information about you and surreptitiously sends it to the originator using your Internet connection, in effect creating a “back channel” for covert communications. Such information often is simply Web-browsing habits and usually not actual personally identifiable information (PII), but nevertheless it is viewed by users mostly as an invasion of privacy and misuse of their Internet bandwidth.

For More Information

Visit Opt Out at grc.com/optout.htm for lists of known and suspected spyware.

See Also: *adware, malware, personally identifiable information (PII), privacy*

SR

Stands for Service Release, a service pack (SP) that includes the original product.

See: *Service Release (SR)*

SSCP

Stands for System Security Certified Practitioner, a security certification from International Information Systems Security Certification Consortium (ISC)².

See: *System Security Certified Practitioner (SSCP)*

SSH

Stands for Secure Shell, a popular remote login protocol.

See: *Secure Shell (SSH)*

SSL

Stands for Secure Sockets Layer, a protocol for secure communications over the Internet.

See: *Secure Sockets Layer (SSL)*

SSL accelerator

Hardware to speed up processing of Secure Sockets Layer (SSL) encryption.

Overview

SSL is the de facto standard for secure communication over the Internet and is used by businesses for encrypting sensitive information exchanged with consumers and partner businesses. SSL is a slow protocol, however, and using SSL on a Web server often can reduce the speed of transactions 10-fold or more. As a result, many businesses enhance the performance of their Web servers by moving SSL processing off the Web server and onto special hardware called an SSL accelerator. Such hardware can be implemented either as a rack-mountable box or a card that can be inserted into the server, and it offloads all SSL processing from the server’s central processing unit (CPU) to the accelerator device.

Marketplace

Examples of the many popular SSL acceleration hardware devices available include the Alteon SSL Accelerator from Nortel Networks, the CertainT 100 Secure Sockets Layer Accelerator from Radware, the BIG-IP SSL Accelerator 800 from F5 Networks.

See Also: *Secure Sockets Layer (SSL)*

SSPI

Stands for security support provider interface, a set of application programming interfaces (APIs) for accessing security services on Microsoft Windows platforms.

See: *security support provider interface (SSPI)*

Stacheldraht

A tool used for launching distributed denial of service (DDoS) attacks.

Overview

Stacheldraht, which means “barbed wire” in German, first appeared in February 2000 when several large-scale DDoS attacks occurred on the Internet. Stacheldraht is similar to TFN2K in that it can be used to launch a variety of types of DDoS attacks, but

includes support for encrypted communication between the client controlling the attack and the compromised hosts or “zombies” actually employed to generate the packets used in the attack. Stacheldraht employs the typical master/slave architecture of other DDoS tools and includes support for scripting automated attacks against individual hosts, networks, or multiple networks.

See Also: *distributed denial of service (DDoS), zombie*

stealth scanning

Any type of port scanning that doesn't actually establish connections with ports on target hosts.

Overview

Port scanning is a set of methods for determining which ports are open and listening on a target system and is used commonly by attackers to seek out vulnerable hosts to attack. Some forms of port scanning form Transmission Control Protocol (TCP) connections to ports on target servers and are easy to detect by an intrusion detection system (IDS) set up to protect the remote network. More difficult for an IDS to detect is a stealth scan, any type of scan in which a TCP connection is not established with the remote host. Some examples of different types of stealth scans include ACK, FIN, NUL, SYN, and XMAS scans. ACK and FIN scans are especially stealthy and often can circumvent firewalls and sneak in under the radar of an IDS, but they generally work only with older operating systems that have flaws in how their Transmission Control Protocol/Internet Protocol (TCP/IP) stack is implemented.

Other stealth-scanning techniques sometimes used by attackers include the following:

- **Slow scanning:** The attacker sends only one or two packets a day from different source addresses to avoid the threshold trigger level of an IDS protecting the remote network.
- **User Datagram Protocol (UDP) scanning:** Since UDP is connectionless, it is intrinsically “stealthy” as a scanning method but is more difficult to

implement than TCP scans and usually relies on bugs in the implementation of targeted TCP/IP stacks.

See Also: *intrusion detection system (IDS), port scanning, SYN scan*

stream cipher

A cipher that encrypts data one bit at a time.

Overview

Stream ciphers are symmetric key algorithms that generally are much faster for encryption information than block ciphers, which encrypt data in discrete chunks called blocks. Stream ciphers generally work by generating a string of bits called a keystream and then XORing this with the stream of data one bit at a time. There are two general kinds of stream ciphers based on how the keystream is generated:

- **Synchronous stream cipher:** The process for generating the keystream is independent of both the plaintext data and its corresponding ciphertext. Synchronous stream ciphers are the most common form of stream cipher, and some synchronous ciphers use a keystream called a one-time pad (OTP) that is completely random and must be generated anew each time the cipher is applied. The advantage of OTPs is that they are almost uncrackable, because even if an attacker could somehow obtain a copy of the OTP for one message, it would be useless for decrypting other messages. One advantage of a synchronous stream cipher is that if a bit of plaintext becomes corrupted, only the single corresponding ciphertext bit will be corrupted as well. A disadvantage, however, is that if a bit of plaintext is somehow dropped, all ciphertext after this bit will become garbage.
- **Self-synchronizing stream cipher:** The keystream is generated from the plaintext and possibly the corresponding ciphertext as well. The advantage of a self-synchronizing stream cipher over a synchronous one is that if a bit of plaintext is dropped, only

a finite amount of ciphertext after it will be garbage, and after that the cipher will correct itself and correctly encrypt data again.

The best-known stream cipher is RC4, a stream cipher with variable key length developed by Ron Rivest.

Other stream ciphers include A5, PANAMA, SEAL, and SOBER. PKZIP, an algorithm used for compressing files, is also a form of stream cipher and is easily cracked.

Notes

Many block ciphers also can operate as stream ciphers, for example, Data Encryption Standard (DES) operating in cipher feedback (CFB) mode; but they are still much slower than true stream ciphers.

See Also: *block cipher, cipher feedback (CFB), one-time pad (OTP), RC4*

STPP

Stands for Microsoft Strategic Technology Protection Program, an initiative launched by Microsoft Corporation in October 2001 to help protect its customers against threats from the Internet.

See: *Microsoft Strategic Technology Protection Program (STPP)*

strong encryption

Encryption with a key long enough to make cracking it unfeasible.

Overview

Strong encryption is a term generally used to describe secret key encryption schemes in which the key is at least 128 bits in length. Any scheme with smaller keys such as 40 or 56 bits is described as “weak encryption” since such keys now can be cracked using off-the-shelf technology in a not unreasonable amount of time, whereas 128-bit encryption schemes commonly are viewed as likely to be uncrackable for the foreseeable future (probably at least a decade or more). Strong encryption is important to understand since export of technology using such encryption often is restricted under U.S. federal law.

See Also: *encryption, secret key encryption*

Su

A UNIX command that allows a user to run an application using different credentials from those employed for the current logon session.

Overview

Su is the UNIX equivalent of the Runas command on Microsoft Windows platforms, though Su is actually the older of the two. **Su** stands for “super user” and allows users to become another user (even root, if they have the credentials) to perform tasks without logging off from their currently logged on account. When you use Su, you supply a password for the second user, unless you are already logged on as root. Once the password is verified, a new shell is opened using the credentials of the second and allowing you to run commands using these credentials. By pressing Ctrl+D, you can exit the shell and return to your current user session.

Su allows administrators to log on using their ordinary user account and then “su” whenever they need to perform an administrative task requiring root privileges.

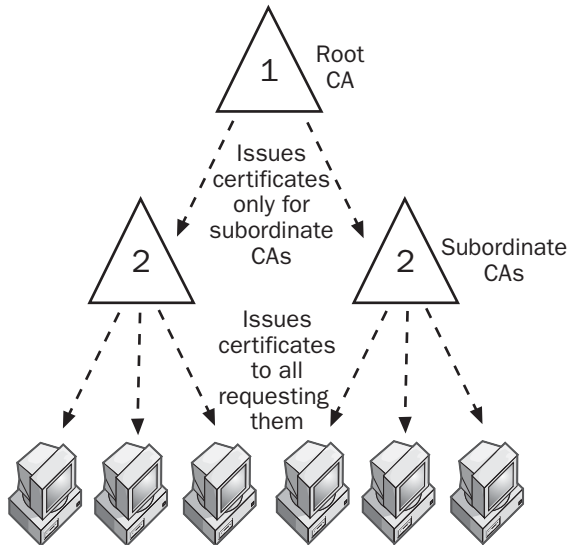
See Also: *Runas*

subordinate CA

A certificate authority (CA) at a level beneath the root CA.

Overview

In a hierarchical Public Key Infrastructure (PKI) system, the top of the hierarchy is held by the root CA, and other CAs beneath it are called subordinate CAs. The job of the root CA is to certify the identity of subordinate CAs by issuing them digital certificates. The subordinate CAs then are responsible for issuing certificates to users who request them; the root CA generally is not involved in servicing such user certificate requests at all. In some implementations certain subordinate CAs also might be assigned other roles such as issuing certificates for applications or smart cards. Because a subordinate CA issues certificates in response to certificate requests, it is sometimes called an “issuing CA” instead.



Subordinate CAs. Subordinate CAs in operation.

See Also: certificate authority (CA), digital certificate, Public Key Infrastructure (PKI), root CA

SubSeven

A notorious remote administration tool (RAT) and Trojan.

Overview

SubSeven appeared in February 1999 as a tool for stealthily gaining control over machines running Microsoft Windows. Although the tool is not as well known as Back Orifice, its power and ease of use resulted in its becoming the most popular RAT used by malicious attackers. For example, at the click of a button an attacker can reboot an infected machine, capture video or sound from a camera or microphone connected to it, record screen shots, copy or delete files, or run arbitrary programs. An infected computer also can be used as a launching pad for performing attacks on other systems using port scanning and redirection. SubSeven also continues to evolve; various strains are available at black hat sites across the Internet, and it remains a dangerous threat that is best guarded against by a defense-in-depth approach that detects and prevents all forms of intrusion.

See Also: Back Orifice, remote administration tool (RAT), Trojan

Sudo

A UNIX command that allows administrators to grant partial root privileges to other users.

Overview

Sudo, which stands for “superuser do,” allows administrators to grant selective privileges to users and groups for running different commands. When a user has been assigned privileges to “sudo” some command, the user simply types *sudo* followed by the command. Sudo then checks to see whether the user has suitable privileges for running the specified command, and if so, performs the command. Sudo also can be configured to prompt the user for a password for more security.

Sudo is a useful command in enterprise environments in which delegating limited privileges to individuals can simplify the job of the administrator. By using Sudo to grant users the privilege of running only certain commands as root and not others, Sudo can create a more secure network environment than one in which root privileges are assigned indiscriminately.

Sudo was developed by Todd Miller and is free software distributed under a Berkeley Software Distribution (BSD)–style license. Sudo is available for most versions of UNIX/Linux.

For More Information

Visit www.courtesan.com/sudo for more information.

See Also: root

SUID root

A process on UNIX platforms that executes with root privileges regardless of its owner.

Overview

SUID is a special type of UNIX permission that allows an ordinary user to execute a process using root privileges. By setting the SUID bit using the `Chmod u+s` command, the owner of a file can assign it SUID permission, and if root is the owner of the file, this procedure results in a SUID root file. The resulting UNIX permissions for a SUID file are `-rwsr-xr-x`.

Although the purpose of SUID is to allow ordinary users to run certain system tasks, SUID can be misused

easily. For example, by running a SUID shell script, an ordinary user could execute arbitrary commands on the system using root privileges. Intruders who compromise UNIX systems often look for vulnerabilities resulting from poor use of SUID, such as SUID files in users' home directories, to perform a root exploit and gain full control of the system. As a result, best practice for administrators normally is to remove SUID root files whenever possible to harden the system against such vulnerabilities.

Notes

Another type of special UNIX permission called SGID (for "set group id") sometimes can be exploited by hackers for performing elevation of privileges (EoP) attacks.

See Also: *elevation of privileges (EoP), root*

superuser

Another name for root, the all-powerful user account on UNIX/Linux platforms.

See: *root*

SUS

Stands for Software Update Services, a tool for keeping critical software updates up-to-date on Microsoft Windows 2000, Microsoft Windows XP, and Microsoft Windows Server 2003 operating systems.

See: *Software Update Services (SUS)*

Swatch

A tool for log file monitoring on UNIX platforms.

Overview

Swatch, which stands for "Simple WATCHer and filter," is an active monitor for the Syslog daemon on UNIX platforms. Swatch can monitor a wide range of log files, watching for specific rules or "triggers" and generating alerts as a result of system activity. Swatch can help administrators avoid missing important Syslog messages that might indicate an intrusion or other malicious activity on their systems. Swatch was developed

by Todd Atkins at Stanford University and is now distributed through SourceForge.

For More Information

Visit swatch.sourceforge.net for more information.

See Also: *log analysis software, Syslog*

SWI

Stands for Secure Windows Initiative, a Microsoft initiative to ensure the security of its products.

See: *Secure Windows Initiative (SWI)*

symmetric key

Another name for secret key, a key used in secret key encryption.

See: *secret key*

symmetric key algorithm

A mathematical algorithm used in secret key encryption.

Overview

Symmetric key algorithms are algorithms that use shared secrets for both encrypting and decrypting information. This shared secret is called a secret key, but if the secret exists only during the lifetime of the communication session and is then discarded, it is also called a session key. There are two main categories of symmetric algorithms:

- **Block ciphers:** These are algorithms that encrypt whole blocks of data at a time and include the Data Encryption Standard (DES), Triple DES (3DES), RC2, and the Advanced Encryption Standard (AES).
- **Stream ciphers:** These are algorithms that encrypt a stream of data one bit at a time and include A5, RC4, SEAL, and others.

See Also: *3DES, Advanced Encryption Standard (AES), block cipher, Data Encryption Standard (DES), RC4, secret key, secret key encryption, stream cipher*

symmetric key encryption

Another name for secret key encryption, which is encryption based on a shared secret between the parties communicating.

See: secret key encryption

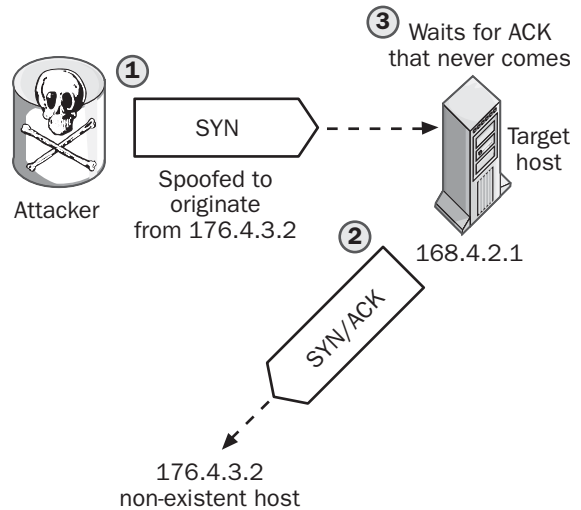
SYN flooding

A type of denial of service (DoS) attack using SYN packets.

Overview

Transmission Control Protocol (TCP) SYN packets are used to initiate connections between two hosts and are sent by the initiating host to the target as the first step of a TCP three-way handshake. In a SYN flood, an attacker sends TCP SYN packets to listening ports on a target host. These SYN packets are spoofed so that they have source addresses that do not correspond to actual systems. When the target receives a spoofed SYN packet, it responds with a SYN/ACK packet directed toward the address from which the SYN packet originated and waits for an ACK packet in reply to complete the connection. Since, however, the source address is spoofed, the ACK packet never comes and the targeted port simply waits until the connection attempt times out. If a listening port receives multiple SYN packets, the port responds with SYN/ACK as to many of them as it can buffer within the memory resources allocated to it by the operating system.

The number of TCP connection attempts a host can buffer varies with different platforms, but is usually no more than several hundred. By sending a flood of such SYN packets to listening ports on the target host, the connection buffers can become full and the target will be unable to respond to additional connection attempts until time outs expire and buffers have room for more attempts. Some operating systems even might hang or crash when connection buffers become full and then need to be rebooted. The result in either case is that connection attempts from legitimate users cannot be accepted and users experience denial of the service they are trying to connect to on the server.



SYN flooding. How a SYN flood works.

There are several ways of preventing or mitigating the effect of SYN flooding attacks:

- **Increase the size of the TCP connection buffers to allow more simultaneous connection attempts:** Unfortunately, the attacker also might be able to increase the rate of SYN packets to compensate.
- **Decrease the time out value for TCP connection attempts:** Unfortunately, this also might make it more difficult for legitimate clients to connect over slow or busy connections.
- **Implement ingress filtering on service provider routers:** This blocks all attacks that use source address spoofing but is really effective only if all Internet service providers (ISPs) agree to implement it.
- **Monitor firewalls and reconfigure them to block SYN attacks when they occur:** This approach, and the use of intrusion prevention systems (IPSs), is the most common method for thwarting SYN floods but requires careful monitoring to ensure legitimate clients aren't blocked out.

See Also: denial of service (DoS), firewall, ingress filtering, SYN scan

SYN scan

A type of stealth scan that makes use of SYN packets.

Overview

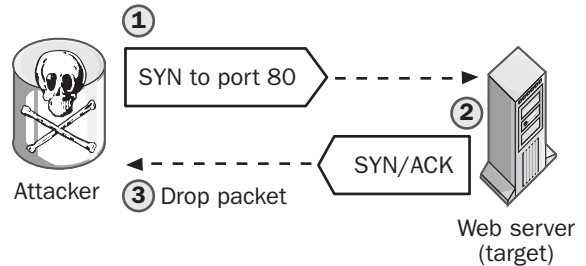
Transmission Control Protocol (TCP) uses a three-way handshake process to establish a connection between two hosts, for which the following steps take place:

- 1- The host wishing to establish the connection sends a SYN packet to the target host to request a socket connection.
- 2- The target host responds with a SYN/ACK that acknowledges receipt of the original SYN packet and sends its own SYN to request a socket.
- 3- The originating host replies with an ACK, and a connection between the two hosts is established.

In a SYN scan, an attacker sends a SYN packet to a port on a target host to see how the host responds. If the host responds with a SYN/ACK packet, this means the targeted port is listening (open) and may be targeted for further attack. Meanwhile, the attacker simply drops the received SYN/ACK packet instead of acknowledging it, which means a connection is not established with the target host. Alternatively, the attacker might respond with an RST packet, which can sometimes help prevent the remote host from logging the connection attempt. If the target port on the remote host is not listening, the remote host responds with an RST packet instead (or possibly provides no response, if a firewall blocks RST packets from leaving the network).

Notes

Because a SYN scan fails to complete a TCP connection that the attacker tries to initiate with the target, it is sometimes called a “half-open” scan.



SYN scan. How a SYN scan works.

See Also: port scanning, stealth scanning

Syskey

A Microsoft Windows NT utility for strengthening password security.

Overview

Syskey first was released as a post–Service Pack 2 (SP2) hotfix for Windows NT and later was included as part of Service Pack 3. Syskey helps protect Windows NT passwords by implementing strong 128-bit encryption for password hashes instead of the previous 40-bit level of encryption. Should an attacker compromise a system and extract password hashes from the SAM database, Syskey makes cracking these hashes much more difficult. However, implementing Syskey is an irreversible step, and the encryption key must be safely stored since if it is lost or corrupted, the system will be unbootable. To provide administrators with flexibility in protecting this key, Syskey provides three key management options:

- **Store the startup key locally on the system:** The disadvantage is that if the system is compromised and the startup key is obtained, an attacker could crack stored passwords.
- **Store the startup key on a floppy disk:** The disadvantage is that the floppy disk must be inserted each time the system needs to be booted, and if the floppy is lost, the system will be unbootable. Managing large numbers of such floppies also can be an administrative headache if there are many servers

that use Syskey. This approach also poses a dilemma, since administrators often disable floppy drives on servers to increase physical security.

- **Use a password entered at startup to derive the encryption key:** Making sure that only authorized personnel know the password is, of course, the main vulnerability here.

In a domain environment, Syskey must be applied to every domain controller for its security to be effective.

See Also: *key, password*

Syslog

A UNIX feature for logging system activity.

Overview

Syslog is the de facto standard for logging system events on UNIX platforms. The feature uses the Syslogd daemon (service) and `/etc/syslog.conf` configuration file to record system messages in specified log files. By default, Syslog logs information locally, but it also can be configured to log to a remote system or to e-mail messages to administrators. Remote logging using Syslog consumes additional network bandwidth but makes it more difficult for intruders to cover their tracks by log cleaning on compromised systems.

Syslog classifies system messages according to seven different levels of severity:

- emerg (emergency)
- alert (alert)
- crit (critical)
- err (error)
- info (informational event)
- debug (debugging information)
- none (no action needed)

Syslog messages also include information about the source that generated the message, such as kern (kernel), daemon (service), or user (user action).

Marketplace

Syslog messages are unauthenticated and therefore vulnerable to spoofing on compromised systems. As a result, replacements for Syslog have been developed, such as the SDSC Secure Syslog project from the San Diego Supercomputer Center.

See Also: *log cleaning, Swatch*

system access control list (SACL)

A type of access control list (ACL) used for auditing securable objects.

Overview

Microsoft Windows platforms support two kinds of ACLs: discretionary access control lists (DACLS) and system access control lists (SACLs). DACLS are the common kind and are used whenever a user configures permissions to control access to an object such as a file or folder. SACLs are used only for the special purpose of controlling the generation of audit messages resulting from attempts to access a securable object. Normally, only administrators have the right to configure SACLs, which is done by enabling auditing on securable objects.

SACLs contain access control entries (ACEs) just like DACLS do, only for SACLs each ACE specifies the types of access attempts by a specified trustee that will result in the system logging the occurrence in the security event. ACEs in SACL can generate such events when attempts at access fail, when they succeed, or both.

See Also: *access control list (ACL), auditing, discretionary access control list (DACL), security log*

System Administrator Tool for Analyzing Networks (SATAN)

A tool for identifying vulnerabilities in networks.

Overview

System Administrator Tool for Analyzing Networks (SATAN) can be used for scanning remote systems and networks to identify known vulnerabilities that could be exploited by an attacker trying to compromise network security. For each vulnerability that SATAN identifies, information also is provided on how to correct the underlying problem, typically through system reconfiguration or applying patches available from vendors. SATAN runs on UNIX platforms and is available free under a General Public License (GPL)-like license.

SATAN was developed in 1995 by Dan Farmer and Wietse Venema and was one of the earliest attempts to develop an integrated platform for protecting network security by automatically scanning networks for known vulnerabilities. SATAN is not a cookbook or toolkit for breaking into networks, but like any security tool it can be used for illegitimate purposes; when SATAN first appeared, there was some concern in the security community that the tool could make the life of crackers easier. Since that time, however, SATAN and related tools such as Security Administrator's Integrated Network Tool (SAINT) and Security Auditor's Research Assistant (SARA) have become essential tools for network administrators in their efforts to cope with the flood of threats against networks connected to the Internet.

See Also: port scanning, Security Administrator's Integrated Network Tool (SAINT), Security Auditor's Research Assistant (SARA), vulnerability

System Security Certified Practitioner (SSCP)

A security certification from International Information Systems Security Certification Consortium (ISC)².

Overview

System Security Certified Practitioner (SSCP) is one of two rigorous information security certifications from (ISC)² that requires candidates to have years of field experience and passing scores on difficult exams to acquire; the other is the respected Certified Information Systems Security Professional (CISSP) certification. SSCP is targeted toward individuals involved in the security administration of systems and networks and those who develop information security standards, policies, and procedures for organizations. SSCP covers the following seven topic areas:

- Access controls
- Administration
- Audit and monitoring
- Risk, response, and recovery
- Cryptography
- Data communications
- Malicious code/malware

SSCP-certified individuals also must recertify every three years to maintain their credentials in good standing and ensure they are knowledgeable concerning the latest tools and techniques for information security.

See Also: Certified Information Systems Security Professional (CISSP), International Information Systems Security Certification Consortium (ISC)²

TACACS

Stands for Terminal Access Controller Access Control System, a security protocol for Authentication, Authorization, and Accounting (AAA).

See: Terminal Access Controller Access Control System (TACACS)

TACACS+

An enhanced version of the Terminal Access Controller Access Control System (TACACS) security protocol.

Overview

TACACS is a security protocol used for Authentication, Authorization, and Accounting (AAA) that was developed in the 1980s by the Defense Data Network (DDN) for MILNET, the U.S. military part of the Internet. TACACS is similar in operation to the industry standard Remote Authentication Dial-In User Service (RADIUS) protocol used by service providers for authenticating users for remote access or Internet connectivity.

TACACS+ is an enhanced version of TACACS developed by Cisco Systems that has enhanced security features, including support for up to 16 different privilege levels and a wide-range of authentication methods. TACACS+ is not compatible with the original version of TACACS and is used mainly for AAA servers used by Internet service providers (ISPs).

Notes

TACACS+ is sometimes referred to as **tac_plus** or **T+**.

See Also: Authentication, Authorization, and Accounting (AAA), Remote Authentication Dial-In User Service (RADIUS), Terminal Access Controller Access Control System (TACACS)

TCPA

Stands for Trusted Computing Platform Alliance, an industry consortium dedicated to improving trust and security on computing platforms.

See: Trusted Computing Platform Alliance (TCPA)

Tcpdump

A UNIX tool for monitoring network traffic.

Overview

Tcpdump is a free sniffing tool for “dumping” (displaying) traffic on a Transmission Control Protocol/Internet Protocol (TCP/IP) network. TCP is a powerful tool that operates from the command line and has numerous options. Tcpdump lets you capture packets whose headers match a specified Boolean expression and then display the packets or save them to a data file for analysis later. Tcpdump also can be used to parse or filter previously saved data files for offline analysis of network traffic. Tcpdump can continue capturing traffic in the background until an interrupt signal is sent or a specified number of packets have been processed.

Tcpdump is available for a wide range of UNIX/Linux platforms and is popular with both security professionals and black hat hackers. The current version of the utility is Tcpdump 3.7.2, and it continues to evolve and be enhanced with new features.

For More Information

Visit www.tcpdump.org for more information.

See Also: sniffing

Tcp_scan

A popular UNIX tool for port scanning.

Overview

Tcp_scan is a free tool for scanning hosts to see which Transmission Control Protocol (TCP) ports are listening for incoming connection attempts. For example, an Apache Web server normally listens on TCP port 80 for Hypertext Transfer Protocol (HTTP) requests issued by Web browsers running on client machines, and Tcp_scan would identify that port 80 is in a LISTENING state on the Apache machine.

Tcp_scan is a command-line tool developed by Wietse Venema and runs on various UNIX/Linux platforms that support the use of raw Internet Control Message Protocol (ICMP) sockets. Like most security tools, it can be used for good or bad purposes. For example, administrators might use the tool for testing firewall configurations or auditing services running on a network. A malicious hacker likely would use it to enumerate which services are running on a target host in preparation for trying to exploit known vulnerabilities associated with such services.

Tcp_scan is included as one of the tools that make up System Administrator Tool for Analyzing Networks (SATAN), a comprehensive package developed in 1995 by Dan Farmer and Wietse Venema for auditing the security of networks, and a similar tool called Security Administrator's Integrated Network Tool (SAINT).

Notes

Wietse Venema also wrote a similar tool called Udp_scan for scanning User Datagram Protocol (UDP) ports.

See Also: port scanning, Security Administrator's Integrated Network Tool (SAINT), System Administrator Tool for Analyzing Networks (SATAN)

TCP session hijacking

Taking control of a Transmission Control Protocol (TCP) session between two hosts.

Overview

TCP session hijacking is a term used to describe a variety of techniques used by attackers to “break into” a TCP session and impersonate one or both of the parties communicating. TCP session hijacking requires that the

attacker first be able to eavesdrop on the session, typically using a sniffer to capture traffic over a poorly secured network connection or by compromising a host on a remote network and installing sniffing software to monitor the network. Once the session can be monitored, the attacker uses spoofing tools to forge Internet Protocol (IP) packets and introduce them into the data stream as part of the session.

There are several methods by which TCP session hijacking can be done, including the following:

- Using source routing to redirect IP packets to a host controlled by the attacker where sniffing and spoofing tools are used to hijack the session.
- Trying to guess or predict TCP sequence numbers for the session and blindly introducing packets, hoping one of the hosts will acknowledge a packet and start directing packets to the attacker.

If TCP session hijacking efforts are directed against only one end of a connection, the attacker may use a denial of service (DoS) attack to temporarily take the host at the other end of the connection while trying to perform the exploit. If the attacker tries to capture both ends of the session, the attack is generally referred to as a man-in-the-middle (MITM) attack instead.

See Also: denial of service (DoS), hijacking, man-in-the-middle (MITM) attack, sniffing, source routing, spoofing, TCP three-way handshake

TCP SYN flooding

Another name for SYN flooding, a type of denial of service (DoS) attack using SYN packets.

See: SYN flooding

TCP three-way handshake

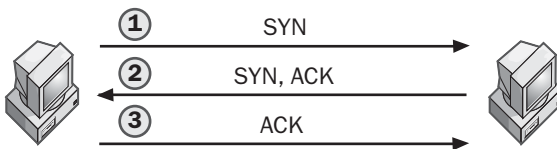
The procedure used by Transmission Control Protocol (TCP) for establishing a session.

Overview

TCP is a connection-oriented protocol for reliable transport of Internet Protocol (IP) packets between hosts on a network. For two hosts to start communicating with each other, they first must establish a session

between them, and this is accomplished using a procedure called a TCP three-way handshake. The three steps in this process are as follows:

- 1- The initiating host sends a TCP SYN packet (a TCP packet with its SYN flag set) to a port on the target host, indicating that it desires to establish a session and to see whether the target is “listening” for connection attempts.
- 2- If the target host is listening, it responds by sending a TCP SYN ACK packet (a packet with both ACK and SYN flags set) back to the initiating host to acknowledge that it is ready to establish a connection and to verify that the initiating host is also listening.
- 3- The initiating host responds with a TCP ACK packet to acknowledge it is listening, and the connection has now been established between the two hosts, allowing a communication session to take place in which IP packets are reliably exchanged between them.



TCP three-way handshake. How a TCP three-way handshake works.

Knowledge of how TCP sessions are established is used by attackers for several purposes, including these:

- Scanning hosts to determine which services are running on them in order to exploit known vulnerabilities for compromising systems
- Hijacking TCP sessions and spoofing packets to impersonate one or both of the hosts in a session

For More Information

For more information about TCP, see the *Microsoft Encyclopedia of Networking, Second Edition*, available from Microsoft Press.

See Also: port scanning, sniffing, spoofing, TCP session hijacking

Tcp_wrapper

A UNIX tool for monitoring and filtering incoming requests for common network services.

Overview

Tcp_wrapper is popular a packet-filtering tool for UNIX systems commonly used for enhancing the security of Transmission Control Protocol/Internet Protocol (TCP/IP) networks. The tool was developed by Eindhoven University of Technology in the Netherlands for monitoring attempts by malicious hackers to compromise UNIX systems. Tcp_wrapper works by substituting a normal call a client makes to a daemon (service) with a call the program makes, thus wrapping Transmission Control Protocol (TCP) connection attempts with an added layer of security. This adds almost no overhead to how the client accesses the service and can be used to report the name of the client and the service it is trying to access. Tcp_wrapper can be used for both monitoring and filtering of TCP connection attempts and can control which Internet Protocol (IP) addresses are allowed to access common TCP ports.

With Tcp_wrapper, an administrator could monitor the source IP address of all connection attempts and filter out connection attempts with clients with specified addresses; an e-mail message is sent to the administrator when such a client tries to connect. Used together with the Ident daemon described by RFC 931, Tcp_wrapper also can be used for monitoring and filtering users who try to remotely log on to systems or otherwise connect to network services.

For More Information

Visit ftp.porcupine.org/pub/security for more information.

See Also: packet filtering

TCT

Stands for The Coroner’s Toolkit, a package of tools for forensic analysis of compromised UNIX systems.

See: The Coroner’s Toolkit (TCT)

Teardrop attack

One of the earliest types of denial of service (DoS) attacks.

Overview

The Teardrop attack appeared in 1997 and exploited weaknesses in the implementation of the Transmission Control Protocol/Internet Protocol (TCP/IP) stacks on the Linux, Microsoft Windows 95, and Windows NT platforms. Teardrop exploited the fact that the routine used for reassembling fragmented packets did not work properly if the fragments were overlapping, something that normally doesn't happen with legitimate network traffic but which is a situation that easily could be created by an attacker using a tool for forging Internet Protocol (IP) packets.

Implementation

The original Teardrop attack, which was named after a C program called Teardrop.c that could be compiled into a tool for performing it, used overlapping User Datagram Protocol (UDP) packets to perform the exploit. UDP port 53, the Domain Name System (DNS) port, was selected for the exploit because it is frequently open on firewalls. Teardrop forged two UDP packets, usually with spoofed source addresses to hide the host performing the attack, and sent these packets to a target machine running on the Windows or Linux platforms and connected to the Internet. When the target received the packets and tried to reassemble them, memory violations occurred and the target either hung or crashed. In the initial phase of the attack, hosts in the .edu and .gov domains were usually targeted, but this expanded to a much wider range of targets as time went on.

Vendors soon released patches to resolve the TCP/IP stack issues that made Teardrop possible, but new variants of the attack appeared rapidly, including Teardrop2, Newtear, Bonk, and Boink, all of which used some variation of the Teardrop approach and exploited other issues with TCP/IP stack operations. Eventually, the bugs in these stacks were flushed out and the attacks dried up.

See Also: Boink attack, Bonk attack, denial of service (DoS), fragmentation, IP fragmentation attack

Temporal Key Integrity Protocol (TKIP)

The replacement for Wired Equivalent Privacy (WEP) in the 802.11i specification for wireless network security.

Overview

Temporal Key Integrity Protocol (TKIP) is designed as a replacement for WEP, the flawed security protocol that is part of the 802.11 standards for wireless networking. TKIP (pronounced “tee-kip”) adds several additional features to WEP, including the following:

- Hashing of the initialization vector that is added to the WEP key to create the session key used to encrypt traffic. This feature helps TKIP protect wireless networks against sniffing exploits that can allow attackers to eavesdrop connections and impersonate legitimate stations.
- A message integrity code (MIC) used to guarantee the integrity of packets and determine when an attack has captured and modified packets. This feature helps protect against key-cracking attacks based on packet injection.
- A mechanism to generate dynamic keys to replace the more easily cracked static keys used by WEP. This feature helps protect against key-cracking attacks based on brute force.

Support for TKIP is also included in Wi-Fi Protected Access (WPA), an interim protocol developed by the Wi-Fi Alliance as a solution to WEP problems until the 802.11i specification is completed.

See Also: 802.11i, hashing algorithm, message integrity code (MIC), Wi-Fi Protected Access (WPA), Wired Equivalent Privacy (WEP)

Terminal Access Controller Access Control System (TACACS)

A security protocol used for Authentication, Authorization, and Accounting (AAA).

Overview

Terminal Access Controller Access Control System (TACACS) was developed in the 1980s by the Defense

Data Network (DDN) for MILNET, the U.S. military part of the Internet. TACACS is similar in functionality to the industry standard Remote Authentication Dial-In User Service (RADIUS) protocol for authenticating users for remote access or Internet connectivity. TACACS is more flexible than RADIUS, however, since it separates the AAA components and allows them to be used independently of one another. For example, in a typical scenario a service provider might use RADIUS for authentication purposes while using TACACS for authorization and accounting.

See Also: *Authentication, Authorization, and Accounting (AAA), Remote Authentication Dial-In User Service (RADIUS)*

TFN

Stands for Tribal Flood Network, a type of distributed denial of service (DDoS) attack developed by “Mixer.”

See: *Tribal Flood Network (TFN)*

The Coroner’s Toolkit (TCT)

A package of tools for forensic analysis of compromised UNIX systems.

Overview

Computer forensics is the process of analyzing compromised systems to obtain evidence for prosecuting criminal activity. In general, computer forensics involves the application of both computer technology and legal expertise and can be a complex and difficult task when systems have been rendered unbootable and data stolen or destroyed. One tool that can help with identifying the exploits of intruders is The Coroner’s Toolkit (TCT), a set of free UNIX tools that takes a “snapshot” of a damaged system to allow forensic analysis to extract as much useful information as possible that might indicate the course of the attack. TCT includes several programs for forensic analysis, including the following:

- **Grave-robber:** Used to capture system information for forensic analysis
- **lls and Mactime:** Used to display access patterns for files

- **Unrm and Lazarus:** Used to make copies of swap files and deleted disk space and then to try to recover data from them

TCT was written by Dan Farmer and Wietse Venema, who also developed System Administrator Tool for Analyzing Networks (SATAN), a comprehensive package of tools for auditing the security of networks. TCT is available for a number of UNIX platforms, including the Berkeley Software Distribution (BSD) family and Solaris, and for Linux platforms.

For More Information

Visit www.porcupine.org/forensics for more information.

See Also: *computer forensics*

threat

Also called an attack, any method used to try to breach the security of a network or system.

See: *attack*

ticket

In Kerberos authentication, a data structure used to provide access to resources.

Overview

A ticket is a set of identification data for a security principal (user or application) issued by a ticket-granting service (TGS), a Kerberos service running on a key distribution center (KDC). Tickets contain information about the identity of the principal and are used for authenticating the principal within a Kerberos realm or domain.

There are two types of Kerberos tickets:

- **Ticket-granting ticket (TGT):** Issued to a user by the authentication service (AS), another Kerberos service running on the KDC in their local realm, after the user submits his or her logon credentials to the network. Once a user has a TGT, the user can present the TGT to the TGS to request a service ticket.
- **Service ticket:** Issued to a user by the TGS in response to the user submitting his or her TGT.

Once the user has a service ticket, the user can present this to a network service in order to authenticate with the service and establish a session.

Implementation

The structure of a service ticket follows a standard pattern and includes the following fields:

- **Message type:** Tickets are used in several kinds of Kerberos messages.
- **Protocol version number:** This is 5 for Kerberos v5 protocol.
- **Sname and Realm:** The name and Kerberos realm of the party to which the ticket is being presented; for example, a network service running on a server.
- **Flags:** A series of options used for specifying how the ticket might be used by different parties.
- **Key:** The session key given to the holder of the ticket for encrypting communication when authenticating with other parties.
- **Cname and Crealm:** The name and realm of the holder of the ticket (a security principal).
- **Transited:** The names of any realms that must be crossed in order for the ticket holder to present it to the target party.
- **Time stamps:** Values describing when the ticket was issued and when it expires.
- **Caddr:** An optional set of addresses from which the ticket must be presented to be accepted as valid by the target party.
- **Authorization data:** Information limiting the rights of the ticket holder (varies with application being used).

Note that the first three fields are in plaintext, while the data in the remaining fields is encrypted using the master key of the target party.

See Also: Kerberos, key distribution center (KDC)

TKIP

Stands for Temporal Key Integrity Protocol, the replacement for Wired Equivalent Privacy (WEP) in the 802.11i specification for wireless network security.

See: Temporal Key Integrity Protocol (TKIP)

Tlist

A tool for displaying running processes on machines running on Microsoft Windows NT or later versions of the operating system.

Overview

Tlist is a Resource Kit tool that displays a “task tree” of running processes on local or remote computers. Tlist can search for processes specified using regular expressions and can match the processes against task names or the names displayed in window titles. Tlist also can display the active services for each process and return the process id (PID) for each process. One common use for Tlist by security professionals is to look for “rogue processes” on a system that might indicate the system has been compromised with a Trojan.

Notes

Another tool called Pulist provides the same functionality as Tlist together with information concerning the owner of the task or process.

See Also: Pulist, Trojan

TLS

Stands for Transport Layer Security, an Internet standard version of Secure Sockets Layer (SSL), Netscape’s protocol for secure communications over the Internet.

See: Transport Layer Security (TLS)

Traceroute

A UNIX tool for displaying the path taken by packets on a network.

Overview

Traceroute is a useful tool for troubleshooting Transmission Control Protocol/Internet Protocol (TCP/IP) networks by displaying the route packets take as they

are forwarded by routers across an internetwork. Traceroute employs the Time To Live (TTL) field in Internet Protocol (IP) packets to try to elicit Internet Control Message Protocol (ICMP) Time Exceeded responses from each router along the path packets travel to a specified remote host. The remote host may be specified either by IP address or fully qualified domain name (FQDN).

Traceroute often is used by malicious hackers as well for two purposes:

- For footprinting a large network to gain a better understanding of the possible targets that can be attacked.
- For identifying special hosts such as firewalls that could represent problems for mounting an attack.

Not all routers or IP hosts respond to Traceroute by sending replies, either because ICMP has been disabled on the host or is blocked by a firewall. As a result, the output of Traceroute is not an infallible guide to the structure of the network being examined. Some service providers block Traceroute's operation entirely by blocking all incoming ICMP traffic on their routers.

Marketplace

There are a number of free tracerouting services available on the Internet, and one use for these tools is for mapping the path packets travel to reach your network from www.traceroute.org, www.traceroute-gateways.com, and many other places. Universities in particular often provide traceroute gateways for public use.

See Also: *footprinting*

Tracert

The Microsoft Windows version of Traceroute, a UNIX tool for displaying the path taken by packets on a network.

See: *Traceroute*

Transport Layer Security (TLS)

An Internet standard version of Secure Sockets Layer (SSL), Netscape's protocol for secure communications over the Internet.

Overview

Transport Layer Security (TLS) is almost identical to Secure Sockets Layer version 3 (SSLv3) and is standardized in RFC 2246 and later RFCs. TLS differs from SSL in only the following ways:

- TLS more clearly separates the handshaking process from the record layer mechanism.
- The protocol can be extended by adding new authentication methods to its operation.
- It improves performance over SSL by using session caching.

The current version of the TLS standard is TLS 1.

See Also: *Secure Sockets Layer (SSL)*

trapdoor

A hidden entry point in a program or system.

Overview

Trapdoors are sometimes coded into applications or operating systems to provide the designers with a secret way of entering the system by circumventing normal security requirements such as authentication and access control. Trapdoors are generally not implemented for malicious reasons but to simplify the task of debugging code during the development process, and often they are simply forgotten about and left in as part of the final released code. Obviously, such trapdoors can be exploited by attackers as well if malicious users can discover them, and trapdoors in general constitute a security risk or vulnerability that may be exploited.

Notes

The terms **trapdoor** and **backdoor** often are used interchangeably, but in the context of information security the term **backdoor** now generally means any mechanism by which an attacker can stealthily reenter a compromised system without needing to repeat the exploit that originally provided access. The term **trapdoor** also is used sometimes to represent key escrow or some other mechanism that enables an authority with legal permission to bypass the security of a system and obtain access to data without the permission of users on the system.

See Also: *backdoor*

Trash2

A denial of service (DoS) exploit that uses Internet Control Message Protocol (ICMP) packets to hang or crash targeted systems.

Overview

Trash2 appeared in the wild in 1999 and was written in C code to execute on UNIX/Linux platforms. Trash2 targets limitations in the Transmission Control Protocol/Internet Protocol (TCP/IP) buffer memory in Microsoft Windows platforms in Windows 95 and later versions of the operating system. Trash2 works by generating large numbers of ICMP packets with random information in their Type and Code fields. These packets are then sent to a target host connected to the Internet and rapidly exhaust the memory resources of the ICMP buffer on the host, usually causing the host to crash and require rebooting. To prevent the target's administrator from simply filtering out the flood, Trash2 also spoofs the source Internet Protocol (IP) address for each packet by assigning it a randomly generated address.

Trash2 is one of many DoS exploits that use ICMP to deny services to legitimate users on a network; some others are Gin, ICMPEX, Smurf, PapaSmurf, and Twinge. The usual way of handling such exploits is to block ICMP traffic on your firewall or router.

Notes

There was indeed a Trash1 exploit by the same author "Misteri0," but it wasn't used much since it didn't spoof source addresses and therefore would allow targeted networks to easily trace the origin of the attack.

See Also: denial of service (DoS), Smurf attack, spoofing

Tribal Flood Network (TFN)

A type of distributed denial of service (DDoS) attack developed by "Mixer."

Overview

Tribal Flood Network (TFN) is similar to the earlier Trin00 in its general operation and runs on UNIX/Linux platforms. TFN can use a collection of compromised

"zombies" to launch several types of denial of service (DoS) attacks against a target host, including Smurf attacks, Internet Control Message Protocol (ICMP) floods, User Datagram Protocol (UDP) floods, and SYN floods. To start the attack, the attacker uses the command-line TFN client program running in a root shell on a compromised "master" system to send an ICMP echo reply message to the TFN server program running on each zombie. Using ICMP messages to launch the attack adds to the stealthy nature of TFN and makes it more difficult to trace compromised master machines. The actual attack command is hidden away in the ID field of the ICMP packets. To make it difficult to shut down an attack, TFN typically uses multiple masters located on different networks connected to the Internet.

Later versions of TFN that use Blowfish encryption further hide their nature and activities.

See Also: denial of service (DoS), distributed denial of service (DDoS), Smurf attack, Stacheldraht, SYN flooding, Tribal Flood Network 2000 (TFN2K), Trin00

Tribal Flood Network 2000 (TFN2K)

A distributed denial of service (DDoS) tool based on the earlier Tribal Flood Network (TFN) exploit.

Overview

Tribal Flood Network 2000 (TFN2K) is similar in operation to TFN (q.v.) but differs in some respects, namely the following:

- While TFN clients on compromised "master" machines use Internet Control Message Protocol (ICMP) echo reply packets to communicate with TFN servers on "zombies," TFN2K clients also can use Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) packets with random port numbers as well for this purpose.
- TFN2K attack commands sent by masters to zombies are repeated 20 times to ensure they get through because of the unreliable delivery mechanism of UDP and ICMP.

- TFN2K sends out dummy attack command packets to random hosts as decoys to make it more difficult for administrators to determine which machines are actually zombies.
- TFN2K encrypts all attack commands using CAST-256 encryption, and the commands are binary and not strings.
- TFN2K zombies are completely silent and do not issue any responses to attack commands they receive from masters.

See Also: denial of service (DoS), distributed denial of service (DDoS), Smurf attack, Stacheldraht, SYN flooding, Tribal Flood Network (TFN), Trin00

Trin00

The original distributed denial of service (DDoS) attack that appeared in 1999.

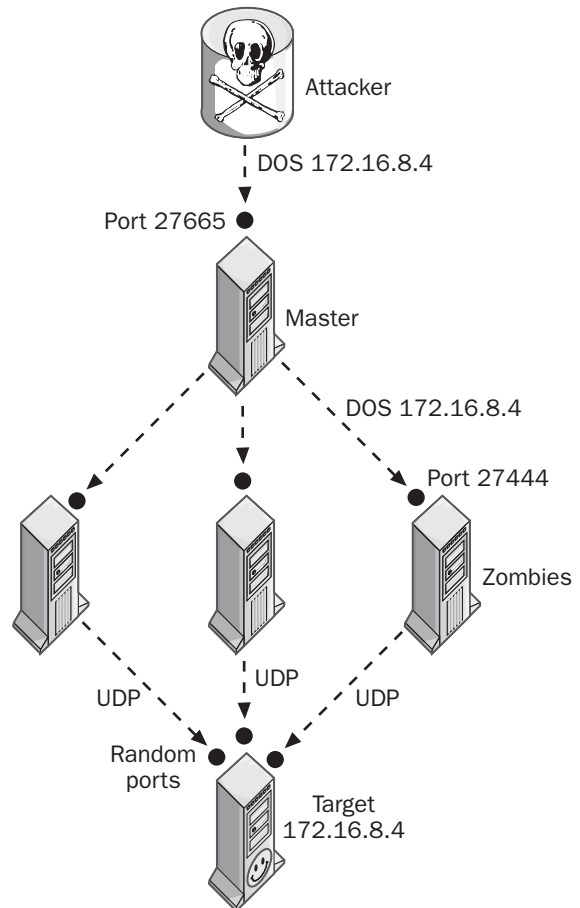
Overview

Trin00 (or Trinoo) was the first DDoS attack to appear in the wild and the prototype of later similar exploits. Trin00 is based on a compromised “master” host used by the attacker to control compromised “zombie” systems that perform the actual attack. The master host runs the Trin00 client program, and the zombies have the Trin00 daemon (service) installed. The attacker is on another machine somewhere else, making it difficult to trace back to the source of the attack.

To launch a Trin00 attack, the attacker first uses Telnet to connect to the master on Transmission Control Protocol (TCP) port 27665. Once connected to the master, the attacker can issue a command to the Trin00 client, which then relays the command to User Datagram Protocol (UDP) port 27444 on the zombies. Some of the commands that can be issued to masters include the following:

- **Dos:** Launch a DoS attack against the host with the specified address
- **Mtimer:** Set the duration of the attack
- **Msize:** Set the size of DoS packets to use

When a command to launch a DoS attack is relayed to a zombie, the zombie starts sending a flood of UDP packets to random ports on the target. If enough zombies and available network bandwidth are involved, the target is soon overwhelmed and legitimate users are unable to access its services. Both the client and server programs are password-protected to prevent anyone else from gaining control of the exploit.



Trin00. How Trin00 works.

Trin00 attacks were relatively easy to block since the source Internet Protocol (IP) addresses of packets sent by zombies were not spoofed. As a result, by configuring the firewall to block the addresses of all zombies (a tedious job if there are thousands of them), the attack

effectively is stopped. Later DDoS attacks such as Tribal Flood Network (TFN) use source address spoofing to make it more difficult to block attacks.

See Also: *denial of service (DoS), distributed denial of service (DDoS), Smurf attack, Stacheldraht, SYN flooding, Tribal Flood Network (TFN), Tribal Flood Network 2000 (TFN2K)*

Trinoo

Properly known as Trin00, the original distributed denial of service (DDoS) attack that appeared in 1999.

See: *Trin00*

Trinux

A security toolkit for Linux that runs from a floppy disk or CD-ROM.

Overview

Linux is a popular platform for many security professionals and black hat hackers because of the wide variety of UNIX security tools that run on the platform. Trinux is a toolkit that provides a popular selection of security tools and a small-footprint ramdisk-based distribution of Linux that gives you access to these tools without the need to install and configure a full-blown Linux system. Trinux includes security tools such as Nmap, Tcpdump, Ngrep, OpenSSH, and others in a format that will run on even a 486 machine with only 12 MB of memory. Trinux is open source software released under the General Public License (GPL) and is maintained by Matthew Franz.

For More Information

Visit trinux.sourceforge.net for more information.

See Also: *Ngrep, Nmap, OpenSSH, Tcpdump*

Triple-A

More commonly known as AAA or Authentication, Authorization, and Accounting, a security framework for controlling access to network resources.

See: *Authentication, Authorization, and Accounting (AAA)*

Triple DES

Usually called 3DES, a secret key encryption algorithm based on repeated application of the Data Encryption Standard (DES).

See: *3DES*

Tripwire

A popular file integrity checker.

Overview

Tripwire can be used to monitor a system looking for attempts to modify or replace important files such as critical operating system files and log files. Tripwire monitors attributes of files that normally shouldn't change, including file size, binary signature, and more. Tripwire is popular with security professionals and can be used for intrusion detection, data integrity assurance, and testing policy compliance.

Tripwire was originally created by Dr. Eugene Spafford and Gene Kim at Perdue University in 1992. Tripwire exists in two forms: open source and commercial. The differences between them are as follows:

- **Commercial version:** Development of the commercial version is managed by Tripwire, Inc. (www.tripwire.com) and has evolved into a number of editions for business, enterprise, and academic use. The current version, Tripwire 3, includes a manager and server component and is available for Microsoft Windows NT 4, Windows XP Professional, Windows 2000, and Solaris.
- **Open source version:** The free version of Tripwire runs only on Linux and evolved from the source code for version 2.2.1 of commercial Tripwire, which was released under the General Public License (GPL) in October 2000.

For More Information

Visit www.tripwire.org for more information.

See Also: *file integrity checker, intrusion detection system (IDS)*

Trojan

A form of malware that often can do considerable damage to a system or network.

Overview

Trojans are distinguished from other malware such as viruses and worms in two ways:

- They generally are stealthy in operation and often masquerade as legitimate programs, while viruses and worms usually have a more obvious effect such as corrupting files or causing a system to crash.
- They usually don't replicate like viruses and worms do.

The terms **Trojan** and **Trojan horse** originally meant malicious code hidden inside a legitimate, useful program, much as the original Trojan horse hid enemy soldiers within an innocent-looking sculpture. If a Trojan attaches itself to some legitimate program to modify that program's behavior, it is sometimes called a Trojan virus. Most Trojans nowadays are self-sufficient executable (*.exe) files that malicious hackers insert in compromised systems to gain control over the system or steal users' data. Trojans can also infect systems when users download applications such as games from untrusted sources on the Internet. Such games may have Trojan code embedded in them that can give a cracker a foothold in a system and threaten the integrity of data.

Some examples of different kinds of Trojans include the following:

- **Password stealers:** These may search for stored passwords on your system and e-mail them to the attacker. Alternatively, they may masquerade as legitimate login screens and wait for you to enter your password and then steal it. A notorious example is the Passfilt Trojan, which masquerades as the Passfilt.dll file used to add strong password security to Microsoft Windows NT. An administrator downloading this file from an untrusted source would not only not be strengthening password security, but actually allowing an attacker to capture passwords!
- **Keystroke loggers:** These monitor everything you type and e-mail it to the attacker or save it in a file for later retrieval.
- **Remote administration tools (RATs):** These allow attackers to gain complete control over your system and do anything they want with it from their remote location. Popular examples include Back Orifice and SubSeven.
- **Zombies:** These are used for launching distributed denial of service (DDoS) attacks against hosts targeted by the attacker.

Detecting the presence of Trojans on a system is not easy. One way is to see what ports are listening on your system by running Netstat or some other utility from the command line. Since many Trojans run services in the background so they can receive commands from the attacker's remote station, any unusual Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) ports in a LISTENING state are a good indication that there might be a Trojan. Another way to detect Trojans is to inventory all executable files on your system and compare your list with an earlier list made when the system was first installed and configured in a "clean room" situation (no network connectivity). Any unusual files, especially if they are named similarly to legitimate files ("service.exe" instead of "services.exe," for example) could be a Trojan.

If you do find a Trojan on your system, you should either do a complete restore from a backup known to be uninfected or reinstall everything from scratch. You might be tempted simply to delete the Trojan, but its presence often can mean the system has been compromised and other exploits were therefore likely to have been performed against it. Rather than spending hours or days tracking down every single change on your system, just reinstall and move on.

Marketplace

There are a number of products available in the market for scanning your system for Trojans and eliminating them. Some of these products include Anti-Trojan, Digital Patrol, Pest Patrol, TrojanHunter, TrojanShield, and The Cleaner.

See Also: Back Orifice, remote administration tool (RAT), SubSeven, virus, worm

Trojan horse

Now commonly known as just **Trojan**, a form of malware that often can do considerable damage to a system or network.

See: Trojan

trust

A relationship between domains that allows a user in one domain to be authenticated for access to the other.

Overview

Trusts are used to create authentication paths between domains in Microsoft Windows NT and later versions of the operating system. Trusts are built on the foundation of two authentication protocols supported by Windows platforms:

- **NTLM:** Supported by Windows NT, Windows 2000, Windows XP Professional, and Windows Server 2003
- **Kerberos:** Supported by Windows 2000 and Windows Server 2003

A trust is always a relationship between two domains with the following characteristics:

- The trusted domain is the domain in which the user has logged on.
- The trusting domain is the domain in which the resource that the user wants to access resides.

Windows NT only supported one-way, nontransitive trusts. This meant the following was true:

- If A trusts B, then B does not trust A unless a second trust in the opposite direction is also created.
- If A trusts B and B trusts C, then A does not trust C unless a trust from A to C is explicitly created.

In Windows 2000 and later, two-way, transitive trusts are the default within a forest, while one-way, nontransitive trusts can be explicitly created for cross-domain trusts and trusts with MIT Kerberos v5 realms on UNIX

networks. In a two-way transitive trust the following are true:

- If A trusts B, then B automatically trusts A as well.
- If A trusts B and B trusts C, then A trusts C also.

See Also: Kerberos, NTLM

Trustbridge

An emerging Microsoft technology for federated identity management among businesses.

Overview

Trustbridge is the code name for a new technology developed by Microsoft Corporation that will allow businesses to share user identity information between applications and organizations. Trustbridge is built on a foundation that includes the trusted Kerberos authentication protocol and the emerging Web Services Security (WS-Security) standard for adding security to Extensible Markup Language (XML) Web services.

Trustbridge is designed to support interoperability between both proprietary and standards-based identity management systems used by different vendors. Using Trustbridge, for example, a Microsoft Windows-based network running Active Directory directory service will be able to recognize and share user identities in organizations using UNIX networks supporting Kerberos authentication. Trustbridge leverages the power of Simple Object Access Protocol (SOAP) running over Hypertext Transfer Protocol (HTTP) with the security of WS-Security to connect businesses and make it easier for them to build deeper and more dynamic relationships with customers, partners, and suppliers, and it also helps mobile employees increase their productivity. Trustbridge also can federate with the Microsoft .NET Passport single sign-on (SSO) identity management system and other WS-Security-based authentication systems.

For More Information

Visit www.microsoft.com/presspass/press/2002/jun02/06-06TrustbridgePR.asp for more information.

See Also: Kerberos, Liberty Alliance Project, .NET Passport, Web Services Security (WS-Security)

TRUSTe

A nonprofit organization that monitors how participating online businesses comply with their privacy policies.

Overview

TRUSTe is an independent third-party organization whose goal is to generate confidence in e-commerce by providing standard mechanisms for monitoring how privacy policies are enforced. When an online business becomes a licensed member of TRUSTe, it agrees to follow the practices of disclosure and informed consent promoted by TRUSTe, to comply with ongoing oversight, and to submit privacy disputes to TRUSTe's dispute resolution process.

TRUSTe acts as an "assurance broker" to ensure companies do what they promise consumers regarding the privacy of personally identifiable information (PII) they collect from them. Should a consumer feel a member company is misusing his or her PII, that person can report the complaint to TRUSTe, which will investigate the violation and, if necessary, take some form of action against the member. Such actions could include enforcing a penalty, performing an audit, or even revoking TRUSTe privileges. In extreme cases TRUSTe could even report the violation to the Federal Trade Commission (FTC) for further action.

For More Information

Visit www.truste.org for more information.

See Also: personally identifiable information (PII), privacy, privacy policy

Trusted Computer System Evaluation Criteria (TCSEC)

A set of security classifications for computer systems developed by the U.S. Department of Defense.

Overview

Trusted Computer System Evaluation Criteria (TCSEC), commonly known as the Orange Book because of the color of its cover, is a standard for computer security developed in 1983 and still important today. TCSEC outlines procedures and methods for evaluating the security of both stand-alone and network operating systems and classifies systems according to the level of security they provide. These classification levels are as follows:

- **D Minimal Protection:** Systems that fail to comply with any other classification
- **C Discretionary Protection:** Systems that use file, directory, or device protection
 - **C1 Discretionary Security Protection:** Systems that use discretionary access control (DAC). Some early UNIX platforms were certified as C1.
 - **C2 Controlled Access Protection:** Same as C1 with the addition of per-user object security through access control lists (ACLs). This is the most common security level, and examples of C2 systems include Microsoft Windows NT, NetWare 4.11, Oracle 7, and IBM OS/400.
- **B Mandatory Protection:** Systems that use mandatory access control (MAC)
 - **B1 Labeled Security Protection:** Same as C2 with the addition of labeling of files, processes, and devices. Examples of B1-certified platforms include Trusted IRIX and HP-UX BLS.
 - **B2 Structured Protection:** Same as B1 with the addition of separation of critical and non-critical elements and protection against covert entry. Examples of B2 platforms include Multics and Trusted XENIX.
 - **B3 Security Domains:** Same as B2 with the addition of reference monitoring of all object access and system recovery procedures. The only B3 platform is Getronics/Wang Federal XTS-300.

- **A Verified Protection:** The highest level of computer system security
 - **A1 Verified Protection:** Same as B3 with the addition of formal proof of integrity and verified design. The only A1 platforms are Boeing MLS LAN, Gemini Trusted Network Processor, and Honeywell SCOMP.
 - **A2 and above:** TCSEC provides for higher levels but these have not been formally defined.

See Also: discretionary access control (DAC), mandatory access control (MAC)

Trusted Computing Platform Alliance (TCPA)

An industry consortium dedicated to improving trust and security on computing platforms.

Overview

The Trusted Computing Platform Alliance (TCPA) was formed in 1999 by Compaq, Hewlett-Packard (HP), IBM, Intel, and Microsoft Corporation and has grown to over 150 participating companies. The goals of the TCPA are to develop advanced hardware and software technologies that will incorporate trust directly into hardware platforms, operating systems, and applications. The TCPA is developing a set of specifications for achieving these goals, which include the following:

- Enhancements in how cryptographic technologies are incorporated into hardware and software
- Improved mechanisms for measuring platform integrity

For More Information

Visit www.trustedcomputing.org for more information.

See Also: Next-Generation Secure Computing Base for Windows, privacy, software piracy, Trustworthy Computing

Trustworthy Computing

An idea that Microsoft Corporation is developing to make computer technology more secure and reliable.

Overview

Trustworthy Computing is an umbrella term for a wide range of technological advances that must be made in order for ordinary people to feel as safe and comfortable about using computers as they do about trusting the lights will come on when they flip a switch on the wall. Trustworthy Computing is emerging as a pervasive distinction of Microsoft corporate culture and is defined by its goals, its means, and its execution.

The goals of Trustworthy Computing consider trust from the user's point of view and include four key goals:

- **Security:** Customers can expect that systems are resilient to attack and that the confidentiality, integrity, and availability of the system and its data are protected.
- **Privacy:** Customers are able to control data about themselves, and those using such data adhere to fair information principles.
- **Reliability:** Customers can depend on the product to fulfill its functions when required to do so.
- **Business integrity:** The vendor of a product behaves in a responsive and responsible manner.

The means of meeting these goals consider trust from industry's point of view and include the following:

- **Secure by design, by default, and in deployment:** Steps have been taken to protect the confidentiality, integrity, and availability of data and systems at every phase of the software development process—from design to delivery to maintenance.
- **Fair information principles:** End-user data is never collected and shared with people or organizations without the consent of the individual. Privacy is respected when information is collected, stored, and used consistent with Fair Information Practices (FIP).
- **Availability:** The system is present and ready for use as required.
- **Manageability:** The system is easy to install and manage, relative to its size and complexity.

(Scalability, efficiency, and cost-effectiveness are considered to be part of manageability.)

- **Accuracy:** The system performs its functions correctly. Results of calculations are free from error, and data is protected from loss or corruption.
- **Usability:** The software is easy to use and suitable to the user's needs.
- **Responsiveness:** The company accepts responsibility for problems and takes action to correct them. Help is provided to customers in planning for, installing, and operating the product.
- **Transparency:** The company is open in its dealings with customers. Its motives are clear, it keeps its word, and customers know where they stand in a transaction or interaction with the company.

The execution of these means is the way in which an organization conducts its operations to deliver the components required for Trustworthy Computing; it includes the following:

- **Intents:** These comprise the following:
 - Company policies, directives, benchmarks, guidelines
 - Contracts and undertakings with customers, including Service Level Agreements (SLAs)
 - Corporate, industry, and regulatory standards
 - Government legislation, policies, and regulations
- **Implementation:** This includes the following:
 - Risk analysis
 - Development practices, including architecture, coding, documentation, and testing
 - Training and education
 - Terms of business
 - Marketing and sales practices
 - Operations practices, including deployment, maintenance, sales and support, and risk management
 - Enforcement of intents and dispute resolution
- **Evidence:** This involves the following:
 - Self-assessment

- Accreditation by third parties
- External audit

For More Information

Visit www.microsoft.com/presspass/exec/craig/10-02trustworthywp.asp for more information.

See Also: *privacy*

TSEnum

A tool for scanning for the presence of Microsoft Windows terminal servers.

Overview

TSEnum is a security tool that can scan the local network to locate any terminal servers running on it. The tool identifies such servers even if they aren't using their default listing Transmission Control Protocol (TCP) port 3389 and works by querying the Browser service, which builds a list of services running on the network. TSEnum does not require special privileges to run and can be used by attackers for enumeration purposes.

For More Information

Visit www.hammerofgod.com/download.htm for more information.

See Also: *enumeration*

tunneling

Transporting packets from a private network over a public network.

Overview

Tunneling is used in virtual private networks (VPNs) to provide remote access connectivity over the Internet. To create a tunnel for a remote client to connect to a VPN server, a tunneling protocol is used to encapsulate network traffic at the client into packets that can be transmitted over the Internet. Examples of common tunneling protocols include Point-to-Point Tunneling Protocol (PPTP) and Layer 2 Tunneling Protocol (L2TP).

In a PPTP scenario, Internet Protocol (IP) packets are first encapsulated in Point-to-Point Protocol (PPP) packets typically used for remote access over dial-up connections. Then these PPP packets are encapsulated a

second time into IP packets that are then sent over the Internet between the remote client and VPN server as a logical point-to-point connection or tunnel.

Tunneling and encapsulation in and of themselves do not add security to transmission over an intrinsically insecure medium like the Internet, but PPTP also encrypts traffic to ensure it can't be read by eavesdroppers. L2TP does not include encryption but is usually used together with Internet Protocol Security (IPSec), which adds encryption to the tunnel to ensure confidentiality.

See Also: *Internet Protocol Security (IPSec), Layer 2 Tunneling Protocol (L2TP), Point-to-Point Tunneling Protocol (PPTP), virtual private network (VPN)*

two-factor authentication

Authentication that uses two separate items or tasks to verify a user's identity.

Overview

Password-based authentication is used widely as a standard mechanism for enforcing network security. When users want to log on to a network, they enter their user name and secret password and the security provider grants or denies them access according to whether their credentials are legitimate or not. The problem with passwords is that they can be lost, stolen, or guessed.

An example of a different approach to authentication would be swiping a smart card through a card reader. A smart card can contain in tamper-proof casing encrypted credentials of the user who owns it. Smart cards also have their security problems, however, since they, too, can be lost, stolen, or even manufactured if the cracker has the right skill and technology.

By combining smart card authentication with a secret password known only to the owner of the card, namely,

a personal identification number (PIN), a two-factor authentication scheme is established that requires two steps for users to be authenticated: swipe the card and enter the secret PIN. This scheme is inherently more secure than the single-factor schemes described earlier, since the smart card and PIN are secured differently: the smart card is kept in the pocket or attached to a wrist band, while the PIN is kept in the user's mind.

Two-factor authentication systems are common in high-security environments such as banks and government institutions. For even more security, a three-factor (or higher) scheme can be devised easily.

See Also: *authentication, password, smart card, social engineering*

Twofish

A block cipher that was one of five candidates for the Advanced Encryption Standard (AES).

Overview

Twofish was developed by a team of cryptographers at Counterpane Labs and was proposed for AES but lost out to Rijndael, the winner of the competition run by the National Institute of Standards and Technology (NIST). Twofish employs a fixed block size of 128 bits and variable key length of 128, 192, or 256 bits. The cipher utilizes 16 rounds for converting plaintext into ciphertext and supports all the standard block cipher modes. Twofish is not patented, not copyrighted, and free for anyone to use.

For More Information

Visit www.counterpane.com for more information.

See Also: *Advanced Encryption Standard (AES), National Institute of Standards and Technology (NIST), Rijndael*

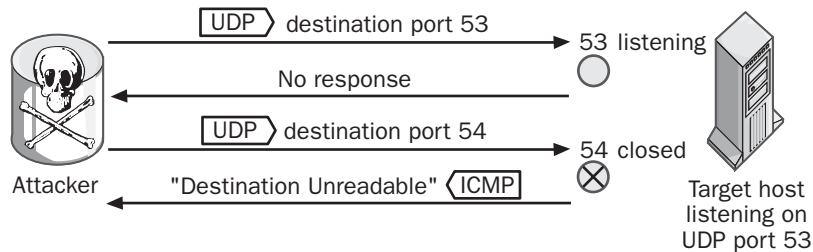
UDP scanning

Enumerating services on a target by using User Datagram Protocol (UDP) packets.

Overview

UDP scanning is a stealthy but somewhat unreliable way of scanning a remote host to see which services are running on it. Some examples of common network services that use UDP include the following:

- Domain Name System (DNS)
- Dynamic Host Configuration Protocol (DHCP)
- Network Time Protocol (NTP)
- Remote Authentication Dial-In User Service (RADIUS)
- Routing Information Protocol (RIP)



UDP scanning. How UDP scanning works.

At least that's how it works in theory. In practice, however, Transmission Control Protocol/Internet Protocol (TCP/IP) protocol stacks for some operating systems are not completely compliant with the RFC 1122 standard, which means that UDP scanning works with some platforms and not with others. Other factors also can affect the reliability of this approach, including these:

- Packet loss because of network congestion

- Simple Network Management Protocol (SNMP)
- Trivial File Transfer Protocol (TFTP)

Implementation

UDP scanning is based on RFC 1122, which indicates that a system should generate Internet Control Message Protocol (ICMP) error messages when a UDP packet is targeted to a closed port; that is, a port that is not in a LISTENING state.

In a typical exploit, an attacker uses a port scanner such as Nmap to send UDP packets to every possible port on the target host. Ports for services that are running generate no response, while ports that are closed generate ICMP Destination Unreachable messages. The attacker can thus determine which ports are listening on the target and then uses this info to test for common vulnerabilities in these services.

- Packet filtering at the firewall
- Unusual configuration of the host

Marketplace

Nmap and Netcat are two popular port scanners used by security professionals and black hat hackers that support UDP scanning. Two tools designed only for UDP scanning include ScanUDP and Udp_scan.

See Also: Nmap, port scanning

UDP tunneling

A method of using User Datagram Protocol (UDP) to establish a covert channel.

Overview

Covert channels are communications channels that hide illicit information flow within a normal communications stream. One method of establishing a covert channel on an Internet Protocol (IP) network is to hide data in packets that normally don't carry payloads. An example is UDP tunneling, which hides data in UDP packets used by such services as the Domain Name System (DNS) service. If firewalls are configured to pass such traffic, information can be leaked from the system without being detected by a firewall or intrusion detection system (IDS).

A common use of covert channels is communication with backdoors. Once an attacker has compromised a system and installed a backdoor, a covert channel allows the attacker to control the system or leak information from it using innocuous-looking UDP packets. One tool that attackers can use for this purpose is Loki, a program first published in *Phrack* magazine. The best way of preventing UDP tunneling is to block unnecessary UDP traffic at the firewall and disable on your host unnecessary services that use UDP.

See Also: *covert channel, Phrack*

URLScan

A tool for securing Microsoft Internet Information Services (IIS).

Overview

URLScan is used for screening incoming requests to the server and filtering them based on rules you specify. URLScan helps protect IIS Web servers by blocking potentially harmful Uniform Resource Locators (URLs) that have the following characteristics:

- Are excessively long
- Request an unusual action to be performed
- Are encoded in an alternate character set
- Include sequences of characters that are rarely found in legitimate requests

URLScan 2 is part of the IIS Lockdown Tool, a wizard-based security tool.

For More Information

Visit www.microsoft.com/downloads/ for more information and to download the IIS Lockdown Tool.

See Also: *IIS Lockdown Tool*

User2sid

A tool for obtaining the security identifier (SID) associated with a user name.

Overview

SIDs are an integral part of security on Microsoft Windows platforms and are strings that uniquely identify security principals on a Windows-based system or network. For example, each user account is assigned a SID when the account is first created, and the SID remains the same even if the account itself is renamed. From the perspective of internal processes and services running on Windows-based machines, it is the SID that identifies the user, not the user name of the account.

Since user accounts are one of the main targets of attackers trying to compromise a network, obtaining the SIDs of accounts stored on a machine can provide useful information for exploiting vulnerabilities. The User2sid utility allows a user to obtain the SID of an account based on knowledge of its user name, which often can be guessed or obtained in other ways; for example, from e-mail addresses. Using User2sid together with its companion utility, Sid2user, which allows a user name to be obtained for a given SID, an attacker might be able to compromise the security of a Windows-based system and obtain useful information about user accounts on the system.

Both User2sid and Sid2user were developed by Evgenii Rudnyi of Moscow State University based on published information of Windows application programming interfaces (APIs). Both utilities can be used against remote machines running on Windows NT or later without needing authentication if null sessions can be established with target machines. One way of preventing attacks that use these tools is to block Transmission

Control Protocol (TCP) port 139, which is used for NetBIOS session enumeration.

See Also: null session attack, security identifier (SID), security principal

UserDump

A tool for enumerating user accounts on Microsoft Windows-based systems.

Overview

UserDump is a tool created by Hammer of God for obtaining the security identifiers (SIDs) of user accounts on a targeted system running Windows NT or later. UserDump does this by using Server Message Block (SMB) null session enumeration combined with “SIDwalking,” a technique that starts by determining the remote system SID and then increments it with expected values to find user accounts, starting with the default Administrator account. Knowing such information can be of use to attackers trying to compromise a target system and gain control over it. UserDump connects to Transmission Control Protocol (TCP) port 445 for performing enumeration, and the simplest defense against this exploit is to block this and other NetBIOS ports at the firewall.

Notes

Don't confuse this tool with UserDump, a tool for cracking the Novell NetWare 3.x bindery, or with Userdump, a tool for capturing a crash dump that occurs when a computer running Windows has a blue-screen error.

For More Information

Visit www.hammerofgod.com/download.htm for more information.

See Also: null session attack, security identifier (SID)

user-level security

Protecting shared resources using user authentication.

Overview

User-level security involves assigning permissions to users to define the level of access they can have for a securable resource such as a file or folder. User-level security requires that users be authenticated when they want to access the resource, and depending on the network model (workgroup or domain), this can be local authentication by the SAM database on the local machine or network authentication using Active Directory directory service in a network based on Microsoft Windows 2000 or later.

User-level security is different from share-level security, which involves using a password to control access to shared resources on a network. Since any user who knows the password can access the share, share-level security affects all users the same way regardless of what rights or privileges they possess through group or role membership. With user-level security, however, different users can have different permissions on a resource. User-level security also can secure resources both for local use and over the network, while share-level security applies only to resources shared on the network.

User-level security is a feature of file systems such as the NTFS file system supported by Windows NT and later versions of the platform. User-level security is more granular than share-level security and provides a greater degree of control over what users are allowed to do with a secured resource. When share- and user-level permissions are combined on Windows-based systems, the result is sometimes called effective permissions.

See Also: permissions, share-level security

victim host

Another name for sacrificial lamb, a server placed outside the firewall with the expectation that it may become compromised.

See: sacrificial lamb

virtual private network (VPN)

A method for establishing secure remote access connections over the Internet.

Overview

Traditional remote access solutions range from slow dial-up connections for occasional access using modems to fast but expensive leased lines such as T1, typically used to connect branch offices to company headquarters over dedicated wide area network (WAN) links. The standard WAN protocol typically used in most remote access scenarios is the Point-to-Point Protocol (PPP), which encapsulates Internet Protocol (IP) packets into a format suitable for sending over WAN connections.

With the ubiquity of the Internet nowadays, however, companies can save the cost of leased lines by sending their IP traffic over the Internet instead. They also can save money by having remote dial-up clients use local Internet service providers (ISPs) instead of having them dial long-distance and connect to a modem pool at the office. The problem with using the Internet for remote access, however, is that the Internet is a notoriously insecure network for sensitive information to travel over. That's where a virtual private network (VPN) comes in handy.

A VPN is essentially a kind of tunnel between two hosts on the Internet that encrypts data and securely sends it from one host to the other. This tunnel is created using a tunneling protocol such as Layer 2 Tunneling Protocol

(L2TP) or Point-to-Point Tunneling Protocol (PPTP), which serves to encapsulate the PPP frames, which are encapsulating the actual IP packets that the hosts are communicating. Using these protocols and suitable firewall configuration on the company network, a remote client can transparently and securely communicate with a server on the company network as if it were actually on the network itself. And instead of paying costly leased-line fees or long-distance charges to phone companies, the only cost required is the cost of being connected to the Internet.

Issues

VPNs have exploded in popularity in recent years as businesses seek ways to cut costs to improve their bottom line. Since VPNs usually use encryption, most administrators assume they are secure, but this is not always the case. Of the two tunneling protocols, PPTP is probably the least secure. The original implementation of PPTP from Microsoft Corporation was found to have some vulnerabilities associated with it involving reuse of session keys and a control channel that was unauthenticated and unencrypted. Microsoft fixed most of these issues in *Service Pack 4 for Windows NT 4.0*, and PPTP has been enhanced significantly in Microsoft Windows 2000 and later platforms. One feature that was not enhanced, however, is the fact that session keys are derived from passwords, and since passwords usually are much shorter than keys, the result is diminished entropy of the key space, making keys easier to crack.

The other tunneling protocol, L2TP, does not provide encryption services, so it usually is used in conjunction with Internet Protocol Security (IPSec), a collection of security extensions for IP that provides end-to-end encryption. Although IPSec is becoming ubiquitous in both operating systems and VPN appliances, its security is considered questionable by some cryptanalysts. The reason for this is the high degree of complexity of

the standards that define IPSec, which make the protocol both difficult to analyze from a cryptographic standpoint and complex to implement from a vendor standpoint. Cryptanalysts Bruce Schneier and Niels Ferguson, who were part of the team that developed the Twofish algorithm, one of the candidates that was considered for the Advanced Encryption Standard (AES), tried analyzing IPSec from a cryptographic standpoint and concluded that the protocol was too complex. They stated that because “complexity is security’s worst enemy,” IPSec could be insecure in some hidden way, and they suggested a number of recommendations for simplifying it to make it more secure and easier to implement. Another cryptanalyst, Steve Bellovin, also pointed out that since IPSec traffic typically contains a significant amount of known plaintext in the form of Transmission Control Protocol/Internet Protocol (TCP/IP) packet header data, it could be susceptible to certain kinds of known plaintext attacks. Of course, if IPSec turns out to be insecure, so do VPNs based on L2TP/IPSec.

For More Information

See www.counterpane.com/ipsec.html for Schneier and Ferguson’s paper on IPSec.

See Also: Internet Protocol Security (IPSec), Layer 2 Tunneling Protocol (L2TP), Point-to-Point Tunneling Protocol (PPTP)

virus

Malicious code that infects files on your system.

Overview

Viruses are code designed specifically to infect files and cause mischief, loss of data, or system failure. Viruses can enter a system several ways, including the following:

- An infected floppy disk received from a friend
- A malicious e-mail attachment
- Infected shareware downloaded from the Internet

Some of the common types of viruses include these:

- **Boot sector viruses:** These infect the master boot record of your hard drive and are loaded into memory each time your system starts.

- **File viruses:** Also called program or parasitic viruses, these attach themselves to executable programs and are loaded into memory when such programs are run.
- **Macro viruses:** These are written in macro languages used by applications such as Microsoft Word and typically infect systems by e-mail. Since macro viruses use applications, they can be multiplatform, infecting computers running different operating systems but sharing the same applications.
- **Multipartite viruses:** These are viruses that are a combination of boot sector and file viruses.
- **Polymorphic viruses:** These might be any of the preceding types of viruses but include the capability of mutating their own code to make them more difficult for signature-based virus protection software to detect and remove.

There are thousands of viruses currently “in the wild,” and more appear each week, so protecting your systems against infection by viruses is an essential part of information security. Protecting yourself against viruses involves the following safeguards:

- Installing and regularly updating virus protection software on all your systems
- Educating users not to open e-mail attachments from strangers
- Making regular full backups of systems that have been scanned for viruses and are clean of infection

If a virus somehow is found on a system on your network, you should do the following:

- Scan all your systems for evidence of the virus.
- Disconnect any infected systems immediately from your network.
- Restore the infected systems from a clean backup.
- Notify your antivirus vendor so it can ensure its signature database is up-to-date.

For More Information

Visit the Symantec Security Response Center at www.sarc.com for information about the latest news concerning virus threats.

See Also: *hoax, Trojan, virus protection software, worm*

virus protection software

Software for detecting and removing viruses.

Overview

Virus protection (or antivirus) software are applications that can determine when a system has been infected with a virus. Typically, such software runs in the background and scans files whenever they are downloaded from the Internet, received as attachments to e-mail, or modified by another application running on the system. Most virus protection software employ one of following methods:

- Signature-based detection:** This is the traditional approach and searches for “signatures,” or known portions of code of viruses that have been detected and cataloged in the wild. Signature-based products are fast and reliable in detecting previously known viruses but generally cannot detect new viruses until the vendor has updated its signature database
- Behavior-blocking detection:** This is a newer approach borrowed from intrusion detection system (IDS) technologies and uses policies to define which kinds of system behaviors might indicate the presence of a virus infection. Should an action occur that violates such a policy, such as code trying to access the address book to mass mail itself through e-mail, the software steps in and prevents this from happening and can also isolate the suspect code in a “sandbox” until the administrator decides what to do with it. The advantage of behavior-blocking detection is that it can detect new viruses for which no signatures are known. The disadvantage is that, like IDSs, such detection systems can generate false positives if the detection threshold is set too low or can miss real infections if it is set too high. A few newer virus protection products include behavior-blocking technology, but most still operate using signature databases.

Marketplace

There are numerous virus protection software vendors in the marketplace, and the following table lists some for reference.

Vendors of Virus Protection Software

| <i>Vendor</i> | <i>Web site</i> | <i>Products</i> |
|---------------------------|--|-------------------|
| Aladdin Knowledge Systems | www.aks.com | eSafe |
| Alwil Software | www.asw.cz | Avast |
| Anyware Software | www.helpvirus.com | Anyware AntiVirus |
| AVG AntiVirus | www.grisoft.com | AVG Anti-Virus |
| Cat Computer Services | www.quickheal.com | Quick Heal |
| Central Command Software | www.centralcommand.com | Vexira Antivirus |
| Command Software, Inc. | www.commandsoftware.com/index.cfm | Command AntiVirus |
| Computer Associates | www.ca.com/virusinfo/ | eTrust |
| Cybersoft | www.cyber.com | waVe Antivirus |
| DialogueScience | www.dials.ru | SpIDer Guard |
| Eset Software | www.nod32.com | NOD32 |
| Frisk Software | www.f-prot.com | F-Prot Antivirus |

Vendors of Virus Protection Software (continued)

| <i>Vendor</i> | <i>Web site</i> | <i>Products</i> |
|------------------------------|--|--------------------------------|
| F-Secure | www.fsecure.com | F-Secure Anti-virus |
| GeCAD | www.rav.ro | RAV AntiVirus |
| H+BEDV Datentechnik | www.antivir.de | AntiVir |
| HAURI | www.hauri.co.kr | ViRobot, DataMedic, Live-Call |
| Hiwire Computer and Security | www.hiwire.com.sg | WinProof and ExcelProof |
| Ikarus | www.ikarus-software.at | Die Klinik |
| Kaspersky Labs | www.kaspersky.com | Kaspersk Anti-Virus (AVP) |
| Leprechaun Software | www.leprechaun.com.au | VirusBUSTER II |
| MessageLabs | www.messageLabs.com/virusEye/ | email scanning services |
| MicroWorld Software | www.microworldtechnologies.com | eScan |
| MKS | www.mks.com.pl | MKS Vir |
| Network Associates | www.mcafee.com or www.nai.com | McAfee Anti-Virus |
| NetZ Computing | www.invincible.com | InVircible AV |
| Norman Data Defense Systems | www.norman.no | Norman Virus Control |
| Panda Software | www.pandasoftware.com | Panda AntiVirus |
| Per Systems | www.persystems.com/antivir.htm | Per AntiVirus |
| Proland Software | www.pspl.com | Protector Plus |
| Safetynet | www.safe.net | VirusNet PC and VirusNet LAN |
| Softwin | www.bitdefender.com | BitDefender |
| Sophos | www.sophos.com | Sophos Anti-Virus |
| Sybari Software | www.sybari.com | Antigen for Microsoft Exchange |
| Symantec | www.symantec.com | Norton Antivirus |
| TREND Micro | www.trendmicro.com | Trend Virus Control System |
| VirusBuster Ltd. | www.vbuster.hu | VirusBuster |

See Also: intrusion detection system (IDS), sandbox, virus

VLAD

An open source tool for scanning systems for critical vulnerabilities.

Overview

VLAD was developed by Bindview's RAZOR security team for checking systems for critical security vulnerabilities referenced in the Top Ten list of the SANS Institute.

These vulnerabilities are listed at www.sans.org/top20/top10.php and are considered by many security professionals to be among the most common security issues that administrators need to guard against to protect their networks. VLAD is especially good at identifying security issues with Common Gateway Interface (CGI) programs, which commonly are used on UNIX Web servers for providing dynamic content to Web pages. VLAD runs on Linux, OpenBSD, and FreeBSD, and its current version is v0.9.2.

For More Information

Visit razor.bindview.com/tools/vlad/ for more information.

See Also: SANS Institute, *vulnerability*

VPN

Stands for virtual private network, a method for establishing secure remote access connections over the Internet.

See: *virtual private network (VPN)*

vulnerability

Anything that gives an attacker the opportunity to perform an exploit.

Overview

To compromise a system an attacker wants to target, a malicious hacker begins by looking for vulnerabilities. A vulnerability can be many things, including the following:

- An error in configuring a service running on the target

- A flaw or bug in the operating system or an application running on the target
- A weakness in an underlying protocol used by a service on the target

When new vulnerabilities are identified on popular platforms and products, incident response centers such as the CERT Coordination Center (CERT/CC) generally issue advisories to notify administrators of the vulnerability along with information about how they work, the impact they can have against a network, and how to protect networks by reconfiguring firewalls, installing vendor patches, modifying the registry, or performing some other action.

See Also: *exploit*

vulnerability scanner

Another name for a port scanner, a program that can determine which ports are “listening” (open) on a target system or network.

See: *port scanning*

wardialing

Cracking the telephone system and networks that use the system for remote access.

Overview

Wardialing refers to the practice of automatically dialing large numbers of telephone numbers in hope of finding a modem or modem bank used by a corporate network for dial-up remote access. Once a modem is found, the attacker can try different passwords to break into the network and compromise security. **Wardialing** also can refer to identifying special carrier equipment such as private branch exchanges (PBXs) and taking control of them to obtain unlimited free long-distance or redirect calls to the attacker and perform social-engineering exploits. Some “wardrivers” even target voice mail systems in hope of taking control and listening to messages that might divulge sensitive information.

Wardialing is essentially a footprinting technique used by attackers to identify potential targets for attacking. Wardialing is one of the oldest activities of the hacking and cracking communities and is less popular now since the Internet has replaced many traditional dial-up remote access solutions with virtual private networks (VPNs) instead. Wardialing is illegal in some jurisdictions, and a “phreaker” can incur legal penalties as a result of abusing the phone system this way.

Marketplace

Some free tools that can be used for wardialing include ToneLoc, one of the earliest DOS-based tools for wardialing, and THC-Scan, a tool developed by the German black hat group called The Hacker’s Choice. Commercial wardialing tools also are available and generally are easier to use than free tools, examples being PhoneSweep from Sandstorm Enterprises and TeleSweep Secure from Secure Logix.

Notes

The analogous activity of footprinting 802.11 wireless networks is called wardriving.

See Also: footprinting, phreaking, virtual private network (VPN), wardriving

wardriving

Eavesdropping on wireless networks.

Overview

Wardriving refers to the practice of driving around downtown with a laptop, 802.11 wireless network card, antenna, and packet sniffer to “sniff out” poorly protected wireless networks to target for compromise. The term derives from the classic “wardialing” attack in which a cracker tries dialing a range of phone numbers in search of modems that can give access to corporate networks. Of course, you don’t have to drive around in a car to practice wardriving—you can walk instead. And you don’t need a laptop, either; a wireless Personal Digital Assistant (PDA) or PocketPC will do just as well and be even less conspicuous.

The motivation for wardriving is mainly the weak security of Wireless Equivalent Privacy (WEP), the security protocol for 802.11 wireless networks. Vulnerabilities in WEP and the fact that some network administrators don’t even turn WEP on or don’t configure it properly make wireless networks particularly easy to eavesdrop on and perform exploits against to compromise their security.

Marketplace

Several sniffing tools are popular with “wardrivers,” including NetStumbler for the Microsoft Windows platform and Kismet and Dstumbler for UNIX/Linux. These tools are not general-purpose network sniffers but rather tools to sniff out the service set identifier

(SSID), Media Access Control (MAC) address, WEP key, and other information from wireless traffic. With this information, attackers quite often can impersonate a wireless client, hijack a session, or compromise an access point. Other tools such as Prism2dump, a general-purpose sniffer for displaying the details of 802.11 frames, and AirSnort, a set of scripts for cracking WEP encryption, also are commonly used by wardrivers. Many wardriving tools can interface with a global positioning system (GPS) device to facilitate mapping wireless “hot spots” as wardrivers snoop around a neighborhood. Other tools such as GPSMap and StumbVerter can take GPS information from a wardriving session and convert it into actual maps using Microsoft MapPoint or some other service.

Like most hacking and cracking tools, wardriving tools also can be used for legitimate network security purposes; for example, to audit the security of a company’s wireless network by performing penetration testing.

See Also: *footprinting, penetration testing, sniffing, wardialing, Wireless Equivalent Privacy (WEP)*

Wassenaar Arrangement

An international agreement on export controls for conventional arms and dual-use goods and technologies.

Overview

The Wassenaar Arrangement was ratified in 1996 by 33 nations to provide an international mechanism for controlling the proliferation of arms and dual-use technologies that could be used for military purposes. From an information security perspective, the agreement is important because it includes restrictions concerning both hardware (advanced materials, materials processing, electronics, and computers) and software (cryptographic systems and technologies). The agreement is reviewed periodically by the nations involved to take into account changes in technology and the geopolitical scene. The agreement acts as an umbrella for national policies and laws regarding arms and dual-use goods.

For More Information

Visit www.wassenaar.org for more information.

See Also: *cryptography*

weak key

A cryptographic key that can be cracked easily because of its unique mathematical properties.

Overview

Some cryptosystems that are normally considered extremely strong have a small number of possible keys with unique properties that make it undesirable to use them for cryptographic purposes. An example is the Data Encryption Standard (DES), an encryption standard used since 1977 by U.S. federal agencies for protecting the confidentiality and integrity of sensitive information. DES has 2^{56} (or approximately 72 quadrillion) different keys, and almost all these keys were considered secure until only a few years ago when DES was first cracked. Because of the internal operation of the cryptographic algorithm DES is based on, however, a small number of these keys are actually trivial to crack. Specifically, 4 keys are considered weak and 12 semiweak because of the simple values they express at certain points in the cipher process, namely, blocks of all ones or all zeros. Cryptanalysts do not consider the existence of such keys as significant in relation to the security of DES, however, since the chance of one of these keys being randomly generated is about 4 quadrillion to 1.

See Also: *Data Encryption Standard (DES), encryption algorithm, key*

Web anonymizer

Any tool for anonymous Web browsing, any method for browsing the World Wide Web anonymously.

See: *anonymous Web browsing*

Web bug

An invisible graphic embedded in a Web page and used to monitor a user visiting the page.

Overview

Web bugs are usually tiny image files one pixel in size and the same color as the background of the Web page on which they reside. The image file doesn't reside on the site itself—the IMG tag on the page simply uses a Uniform Resource Locator (URL) to load the bug from a different site, usually one belonging to a marketing or advertising company. When the page is loaded, the image is called from the remote site and a log is created that can be used to collect information about the user loading the page, such as the Internet Protocol (IP) address of the computer, the URL of the page visited, and the time the visit occurred. The word **bug** in Web bug originates because such covert monitoring activity is called “planting a bug” in the espionage trade, and many users find Web bugs and similar technologies a violation of their privacy if such practices are not explicitly mentioned in the online privacy policy for the site users are visiting.

Web bugs can be used for a number of purposes, including the following:

- Track visits to a site to better target marketing efforts
- Track the browsing habits of users as they traverse links across the site
- Exchange information between sites concerning visitors who visit both sites

Marketplace

The hard way of detecting Web bugs is to view the Hypertext Markup Language (HTML) source of each Web page you visit, looking for an embedded URL that references a tiny invisible image on another site. A free tool called Bugnosis from Privacy Foundation (www.bugnosis.org) can detect Web bugs embedded in Web pages and alert you when they are found, giving you details about the bug and making it visible on the page. To block Web bugs entirely, use ad-blocking utilities such as AdSubtract or WebWasher.

See Also: *privacy*

web of trust

The approach used by Pretty Good Privacy (PGP) for managing trust between users.

Overview

PGP differs from most other public key cryptography systems in that instead of using a hierarchy of certificate authorities (CAs) for issuing keys and verifying digital certificates, it allows each user to decide which keys of other users to trust. This trust model is called direct trust and results in a complex mesh or “web” of trust relationships between users of PGP. PGP also allows a user to be a “trusted introducer” for other users and to act as a “mini-CA” for others. Trust also can be granted at various levels, including implicit, complete, marginal, and no trust. Certificate revocation can be performed either by the owner of a certificate or someone the owner designates as a “trusted revoker.” Such a system is simple to manage but does not scale as well as a hierarchical Public Key Infrastructure (PKI) system does.

See Also: *certificate authority (CA), Pretty Good Privacy (PGP), Public Key Infrastructure (PKI)*

Web permissions

Special permissions for configuring access to Web content in Microsoft Internet Information Services (IIS).

Overview

Although NTFS permissions are the primary method for controlling access to Web content on IIS Web servers, another set of permissions called Web permissions also is involved in the process. These Web permissions are similar to shared folder permissions for network shares in that they affect all users the same way, as opposed to the user-level security implemented with NTFS. The Web permissions available on IIS 5 and 6 are these:

- **Read:** Allows users to read the page or download files in the directory
- **Write:** Allows users to modify the page or upload files to the directory

- **Directory browsing:** Allows users to view a list of files in the directory when there is no default document present
- **Script source access:** Allows users to access source code of the page if either Read or Write is allowed
- **Execute:** Determines which types of executable content can run on the page, which can be one of the following:
 - Scripts and Executables
 - Scripts only
 - None

See Also: permissions

Web Services Security (WS-Security)

An emerging standard for adding authentication, confidentiality, and data integrity to Web services.

Overview

Web Services Security (WS-Security) is a specification developed by IBM, Microsoft, and Verisign to enhance the Simple Object Access Protocol (SOAP), a message-passing protocol used in distributed Web services scenarios. WS-Security extends SOAP in the following areas:

- Provides support for integrating different security models and encryption technologies into SOAP
- Allows SOAP messages to be associated with security tokens using a variety of formats
- Specifies procedures for encoding digital certificates and Kerberos tickets into SOAP messages

WS-Security does not add security to SOAP but rather provides an extensible framework for incorporating any security model or encryption technology into SOAP for securing distributed Web services applications.

For More Information

For more information about SOAP, see the *Microsoft Encyclopedia of Networking, Second Edition*, available from Microsoft Press.

See Also: digital certificate, Kerberos

WEP

Stands for Wired Equivalent Privacy, the security protocol for 802.11 wireless networking.

See: Wired Equivalent Privacy (WEP)

WFP

Stands for Windows File Protection, a mechanism for preventing critical system files from being modified on Microsoft Windows platforms.

See: Windows File Protection (WFP)

Whisker

A tool for scanning a network for hosts running Common Gateway Interface (CGI) applications.

Overview

CGI was developed in the UNIX networking environment to enable Web browsers to execute “gateway” applications on Web servers. These gateway programs are typically written in PERL script and sometimes C programming language and are used to enhance the functionality of Web sites by providing handlers for forms, database access, and other dynamic content features.

Whisker, a tool developed by Rain Forest Puppy of SecuriTeam, can be used to scan networks for the presence of Web servers running CGI applications. Whisker can detect CGI code running on both Apache and Internet Information Services (IIS) machines, and it has a number of advanced features, including the ability to scan for CGI directories other than the default cgi-bin directory, scan applications in virtual subdirectories, and use an Nmap results file as the basis for directing a scan.

Like most security tools, Whisker can be used for good or bad. A network administrator could use it to audit CGI applications running on the network, while an attacker could use it to footprint a network in preparation for launching an attack.

For More Information

Visit www.securiteam.com/tools/3R5QHQAPPY.html for more information.

See Also: port scanning

white hat

Euphemism for a security professional who performs hacking activities for (possibly) legitimate purposes.

Overview

The term **white hat** generally means an expert in hacking and cracking who seeks out vulnerabilities in platforms and either reports them immediately to the vendor or publicizes them on a security advisory site or list such as CERT Coordination Center (CERT/CC) or Bugtraq. Unlike their black hat counterparts, the motives of the white hats are not to damage systems, steal data, or gain notoriety, but instead to improve the security of vendors' products and of the Internet overall.

Most vendors prefer information concerning vulnerabilities in their products reported first to them so they can then publicize it when appropriate (for example, after a patch is ready to be issued). This issue of "vulnerability disclosure" is a hot topic, and there are two sides to consider:

- What if someone reports a vulnerability to the vendor and the vendor chooses to sit on it by not fixing the problem and not publicizing it? Some white hats see such a "security through obscurity" approach as dangerous and self-serving because black hats are likely to discover (or may have already discovered) the same vulnerability and use it for malicious purposes against the vendor's customers. In this sense, the white hat may see him- or herself as the Lone Ranger (who actually wore a white hat) riding to the rescue of these companies and protecting them from the Big Bad Vendor.

- On the other hand, what if the white hat publicizes the vulnerability immediately before reporting it to the vendor and before a patch is available to correct it? Many vendors, businesses, and law enforcement agencies frown (or worse) upon such activities and liken them to putting a gun in the hands of the Bad Guy, the black hat community, who may rush with glee to start exploiting the problem, causing chaos and costing businesses money. To those on this side of the issue, white hats look more like a darker shade of gray.

See Also: black hat, Bugtraq, CERT Coordination Center (CERT/CC), gray hat, hacking, vulnerability

Whois

A protocol and tool for looking up information about registered Internet names and numbers.

Overview

Whois is a tool provided by Internet name registrars for searching databases of Domain Name System (DNS) names, Internet Protocol (IP) addresses, Autonomous System (AS) numbers, and other identifiers upon which the operation of the Internet depends. Whois is based on an Internet standard protocol defined in RFC 954 and allows registrar databases to be searched for owner and contact information for Internet names and numbers.

An example of Whois is the service provided by Network Solutions, a company owned by Verisign. By entering a domain name, you can look up the name and address of the individual who registered the domain name, the administrative and technical contacts that maintain the DNS records for the domain, when the name was registered and when it expires, and the names of the primary and secondary DNS name servers for the domain. Other domain name registrars provide similar Whois functionality for performing domain name lookups.

The American Registry for Internet Numbers (ARIN) provides a somewhat different Whois tool. Using the ARIN Whois database, you can look up registration and contact information for IP addresses and AS numbers registered with ARIN. ARIN's Whois can't be used for

looking up addresses for U.S. military networks, however, because such information is managed by the Defense Information Systems Agency (DISA) instead.

Whois lookups are used by network administrators for troubleshooting Internet addressing and naming issues, but they also are used by malicious hackers for footprinting target networks prior to commencing an attack. The following table lists some of the more common or

important sites for performing Whois lookups to find various information. Since Whois is not a distributed database but a series of disjointed databases managed by different companies and organizations, performing a Whois lookup first involves knowing which database to search. There are some tools available on the Internet, however, that act as metaquery agents that enable you to query multiple Whois sources in a single step, including Allwhois, BetterWhois, and others.

Popular Whois Sites on the Internet

| <i>Organization</i> | <i>Link</i> |
|--|--|
| American Registry for Internet Numbers (ARIN)• | www.arin.net |
| RIPE Network Coordination Center• | www.ripe.net |
| Asia Pacific Network Information Center (APNIC)• | www.apnic.net |
| U.S. Government Domain Registration and Services• | www.nic.gov |
| U.S. Department of Defense Network Information Center• | whois.nic.mil |
| Network Solutions• | www.networksolutions.com |
| Allwhois• | www.allwhois.com |
| BetterWhois• | betterwhois.com |

See Also: *footprinting*

Wi-Fi Protected Access

A technology developed by the Wi-Fi Alliance as a solution to problems with Wireless Equivalent Privacy (WEP), the security protocol for 802.11 wireless networks.

Overview

Wi-Fi Protected Access is designed to resolve security issues with WEP by providing a standards-based solution that is backwardly compatible with existing 802.11 wireless networking hardware. Wi-Fi Protected Access consists of a subset of 802.11i, the emerging standard specifying security enhancements for 802.11 wireless networking. Although implementing the full 802.11i specifications would require changes in hardware design of access points and other wireless devices, Wi-Fi Protected Access provides many of the advantages of 802.11i while requiring only a software upgrade on wireless hardware.

Wi-Fi Protected Access improves WEP security in two ways:

- Improvements in data encryption through the Temporal Key Integrity Protocol (TKIP) to ensure greater confidentiality of communications
- Support for strong authentication through 802.1x and Extensible Authentication Protocol (EAP) together with centralized Remote Authentication Dial-In User Service (RADIUS) servers

For More Information

Visit www.weca.net for more information.

See Also: *802.1x, 802.11i, Extensible Authentication Protocol (EAP), Wireless Equivalent Privacy (WEP)*

Windows File Protection (WFP)

A mechanism for preventing critical system files from being modified on Microsoft Windows platforms.

Overview

In early versions of the Windows platform, the operating system sometimes could be damaged by applications that overwrote executable files or dynamic-link libraries (DLLs). The effects of such damage could range from unpredictable system instability to the inability to boot the system.

Windows File Protection (WFP) is a feature included in Windows 2000 and later to prevent this type of problem from occurring. WFP maintains a cache of copies of *.exe and *.dll files critical to the operating system, and should any system file be somehow overwritten or modified, WFP detects the occurrence and automatically replaces the modified file with its pristine version from the Dllcache directory. Should a pristine copy of the file not be found in the cache, WFP prompts you to specify the location of the installation media, which is typically either the product CD or a network distribution share. WFP also creates an audit trail by logging the event in the system event log.

See Also: *file integrity checker*

Windows NT Challenge/Response

Another name for NTLM, the authentication protocol used by Microsoft Windows NT.

See: *NTLM*

Windows Product Activation (WPA)

A Microsoft technology aimed at reducing software piracy.

Overview

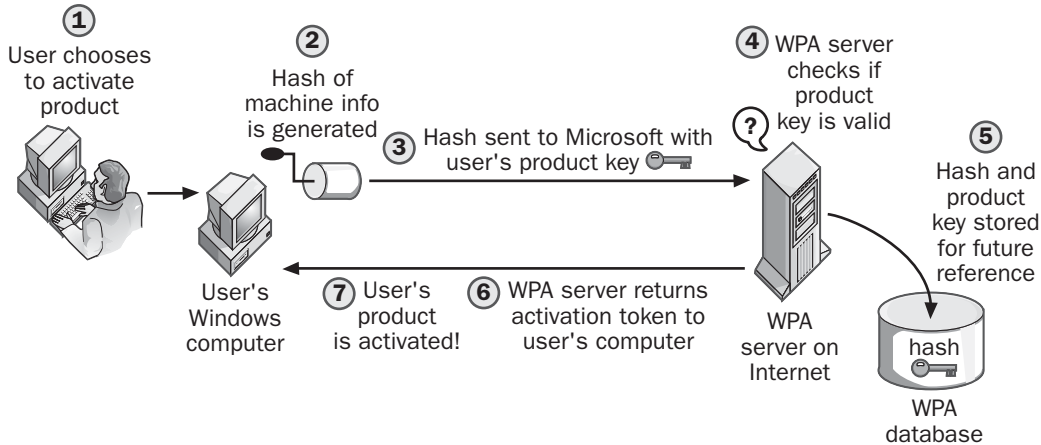
Software piracy has become in recent years a billion-dollar industry and a drain on the high-technology economy. To combat software piracy, Microsoft Corporation developed Windows Product Activation (WPA) and has incorporated it into Microsoft Windows XP,

Windows Server 2003, and Office XP product lines. WPA is designed to make it harder to pirate these products and helps ensure customers that software they've purchased or obtained is legitimate and of the quality they expect.

Implementation

When a product such as Windows XP is installed on a system, the user is prompted at the end of Setup to activate the product, and the user can either choose to do so immediately or delay until a later time. Different products have different grace periods for activation, and with Windows XP, for example, you have 30 days from installation to activate it or its functionality will be temporarily locked down until you choose to comply with activation. If you decide to activate your product, you can do so in one of two ways:

- **Internet:** If you have a live Internet connection on your machine, activation is as simple as clicking a few times and takes less than a minute. During the activation process, the Windows operating system determines certain features of your machine's hardware configuration and anonymously sends this information to Microsoft, which maintains the association between your product key and machine information in its online activation database. This association or binding of product key to machine information prevents others from using your product CD to install another copy of the Windows operating system on a different machine.
- **Telephone:** If you require or prefer phone activation, you can call a Microsoft customer representative, provide that person with your product key; then you will receive further instructions on how to activate your product.



Windows Product Activation (WPA). How WPA works over the Internet.

Activation is needed only for commercial packaged retail products. Enterprise customers with volume licensing schemes do not have to activate their products, and original equipment manufacturers (OEMs) can preactivate products on systems they assemble and sell to make things easier for their customers.

Notes

Note that activation is not the same as product registration, which uses personal information for registering your product with Microsoft to provide you with timely information and offers on other products. Activation is entirely anonymous and no personally identifiable information (PII) is transmitted to Microsoft during the process. This information concerning your machine's hardware configuration collected by Microsoft is hashed to protect it from eavesdropping while it is transmitted, and it is stored on Microsoft's activation servers in hashed form as well. The algorithm used to hash the information is a one-way function that prevents even Microsoft from determining the exact details of your system's hardware and devices, so everything is totally anonymous.

See Also: software piracy

Windows Rights Management (WRM)

A new technology from Microsoft Corporation for secure content management.

Overview

Windows Rights Management (WRM) is an extension for Microsoft Windows Server 2003 that allows companies to control access to documents they have created, allowing sensitive internal documents such as financial statements or e-mail messages to be managed more securely and preventing them from being "leaked" or stolen. WRM allows users to control access to documents they have created, deciding whether such documents may be copied, forwarded, or printed, and whether recipients can hold on to them indefinitely or whether they expire after a certain time.

WRM is scheduled for release sometime in 2003 and is part of Microsoft's ongoing Trustworthy Computing initiative to enhance the security of the Windows platform.

See Also: Trustworthy Computing

Windows Update

An online extension of the Microsoft Windows platform for keeping your system secure and up-to-date.

Overview

Windows Update is a Web site run by Microsoft that allows users of different Windows platforms to find and download software updates. These updates can be used to fix security vulnerabilities found in Windows, to keep Windows running smoothly and reliably, and to provide users with the latest features and enhancements in Windows functionality.

Microsoft Windows 2000 with Service Pack 3, Microsoft Windows XP, and Microsoft Windows Server 2003 also include a feature called Automatic Updates that allows your system to contact Windows Update automatically on a specified schedule to download and install updates on your computer. For corporate users, Microsoft Software Update Services (SUS) allows administrators to manage distribution of updates across a network of computers running on the Windows 2000, Window XP, and Windows Server 2003 platforms.

Implementation

To use Windows Update, you simply open Microsoft Internet Explorer and visit the Windows Update Web site at v4.windowsupdate.microsoft.com. Once there, you use the following procedure:

- 1 Choose to allow Windows Update to scan your system for new updates to download and install.
- 2 Browse through the various recommended updates in each category and select the ones you want to download.
- 3 Download the updates and install them on your system (may require a reboot).

The scanning process protects your privacy and does not collect any information that can be used to identify you. The information collected is used only to determine what updates your system needs or supports, and includes only the following:

- Operating system version number
- Internet Explorer version number
- Version numbers of other software for which Windows Update provides updates
- Plug and Play ID numbers of hardware devices
- Region and Language setting

In addition, the Product ID and Product Key confirm that you are running a validly licensed copy of Windows operating system.

The different categories of updates you can choose from include the following:

- **Critical Updates:** Security fixes and other important updates to keep your computers current and your network secure
- **Recommended Downloads:** The latest Windows and Internet Explorer service packs and other important updates
- **Windows Tools:** Utilities and other tools provided to enhance performance, facilitate upgrades, and ease the burden on system administrators
- **Internet and Multimedia Updates:** The latest Internet Explorer releases, upgrades to your Windows Media Player, and more
- **Additional Windows Downloads:** Updates for your desktop settings and other Windows features
- **Multi-Language Features:** Menus and dialog boxes, language support, and Input Method Editors for a variety of languages

The various Windows platforms that support Windows Update include these:

- Windows 98
- Windows 98 Second Edition
- Windows 2000 Professional
- Windows 2000 Server
- Windows 2000 Advanced Server
- Windows Millennium Edition (Windows Me)
- Windows XP
- Windows Server 2003 (including 64-bit versions)

See Also: *privacy, Software Update Services (SUS)*

Winnuke

A classic denial of service (DoS) attack.

Overview

Winnuke was one of the first DoS attacks devised against computers running Microsoft Windows, and it appeared on the Internet in 1997. Winnuke worked by exploiting a vulnerability in the Transmission Control Protocol/Internet Protocol (TCP/IP) stack of Microsoft Windows 95 involving User Datagram Protocol (UDP) port 139, the NetBIOS port. To perform the attack, an attacker crafted a packet with the out of band (OOB) flag set and garbage data as payload. When a computer running Windows 95 received such a packet, it crashed and had to be rebooted, which resulted in logged-on users losing their work. Later versions of Winnuke also worked against the Windows NT platform until patches were released to resolve the issue.

Winnuke spread widely across the Internet and particularly targeted Internet Relay Chat (IRC) servers, almost bringing IRC to a halt. Winnuke was the first of a long line of DoS exploits that use nonstandard packets to hang or crash systems. A system affected by such an attack is said to have been “winnuked” or simply “nuked.”

See Also: denial of service (DoS)

WinTrinoo

A tool for launching a distributed denial of service (DDoS) attack.

Overview

WinTrinoo is a Microsoft Windows–based version of Trin00, an early UNIX tool for performing DDoS attacks. WinTrinoo installs a Trojan named service.exe on compromised intermediary hosts or “zombies” and uses these hosts for overwhelming a target system or network. The Trojan is 23,145 bytes in size and is not to be confused with services.exe, the Service Control Manager (SCM) on Microsoft Windows platforms. The WinTrinoo Trojan communicates on port 34555 with the control station used by the attacker for starting the attack, so if you find systems on your network that are listening on this port, they likely have been compro-

mised and should be disconnected from the network and cleaned immediately.

See Also: distributed denial of service (DDoS), Trin00, Trojan horse

Winux

A virus that infects both Microsoft Windows– and Linux-based systems.

Overview

Winux appeared in March 2003 and represented a significant step in the evolution of viruses; namely, it is a multiplatform virus that infects two different operating systems. There have been previous viruses that could produce this kind of infection, but these depended on cross-platform applications such as Microsoft Office instead of the operating systems themselves. Winux is not completely platform-independent, however, because it first must infect a machine running Windows before it can infect Linux systems. Once a Windows-based system has been infected, the virus spreads by enumerating network shares on both Windows and Linux servers and by infecting executables and other types of files.

Winux does not do any real damage and is more of a proof of concept than a threat. But its very existence indicates new directions in the black hat community that one day might result in viruses that hop platforms easily, which makes them more difficult to guard against and eliminate.

See Also: virus

Wired Equivalent Privacy (WEP)

The security protocol for 802.11 wireless networking.

Overview

Wired Equivalent Privacy (WEP) is an Institute of Electrical and Electronics Engineers (IEEE) standard for encrypting traffic between 802.11 access points and wireless clients. WEP was designed to protect wireless networks from eavesdropping and to provide confidentiality for wireless communications. WEP uses the RC4

stream cipher to encrypt traffic. The session key is a 64-bit key comprising a 24-bit initialization vector that changes periodically and a 40-bit WEP key. One weakness of this scheme is the small size of the WEP key, which makes the system susceptible to brute-force cracking attacks. Another weakness is the low entropy of the initialization vector, which provides only 16,777,216 different session keys for a given WEP key, which can result in reuse of an initialization vector during a long session.

Other weaknesses of WEP include the following:

- Only the data portion of 802.11 packets are encrypted and not the headers.
- The initialization vector is sent over the network in cleartext.
- No authentication of client or access point is provided.

Because of these weaknesses, WEP is considered to be insufficient as a solution for securing wireless networks. Several popular tools such as AirSnort and DWEPCrack, available for “wardriving” or eavesdropping on wireless networks, can exploit vulnerabilities in WEP to give attackers access to networks.

As a result of these issues, wireless vendors and standards bodies have been working on several replacements for WEP including 802.11i, WEP2, and others. Wi-Fi Protected Access, an interim solution developed by the Wi-Fi Alliance, has gained some popularity. It requires only a software upgrade for 802.11 devices to support it, and it is forward-compatible with 802.11i, which is likely to be the future standard for wireless security. Another possible solution is to implement virtual private network (VPN) tunnels over wireless networks using Internet Protocol Security (IPSec) to encrypt traffic end to end.

See Also: 802.11i, Internet Protocol Security (IPSec), RC4, virtual private network (VPN), wardriving, Wi-Fi Protected Access

workaround

A temporary solution to a problem.

Overview

When a vulnerability that affects its security is discovered in a software product, usually the vendor quickly tries to develop, test, and issue a patch to correct the problem. Sometimes developing a patch is not practical or necessary, and in such cases the vendor instead issues a workaround for the problem. A workaround generally consists of performing steps to configure certain features of the product, and it may involve delving into advanced configuration aspects such as editing the registry or modifying configuration files.

Workarounds usually are issued instead of patches for the following situations:

- When the patch will take some time to develop because of the complexity of the issue and the nature of the threat
- When a patch is impractical to develop since the problem is intrinsic to the underlying architecture of the product and cannot be changed until a new version of the product is developed
- When the chances of the vulnerability being exploited are remote because of the unusual conditions under which it can occur, which makes it simpler to issue a workaround than develop a patch
- When the number of users likely to be affected is only a small percentage of those using the product, which makes it uneconomical for the vendor to develop a patch

See Also: patch, vulnerability

world-writable

In UNIX, permissions on a file that allow any user to modify the file.

Overview

World-writable files are a common security issue with UNIX platforms and can lead to vulnerabilities which crackers can use to perform a root exploit to gain complete control of your systems. The problem is not the fact that files can be world-writable, but which files have this property, since inexperienced administrators often grant this privilege to files for the convenience of

remotely managing various aspects of UNIX over the network. In particular, the following types of files should never be made world-writable:

- System initialization files
- System configuration files
- Startup scripts

Best practice is to use the Find command to search for world-writable files on all your UNIX systems and remove such permissions unless they are absolutely necessary; for example, for certain device files in the /dev directory. Also, be sure to test your systems each time you remove such permissions from a system file.

Notes

A similar problem can occur with SUID root files.

See Also: SUID root

worm

Autonomous code that propagates across a network.

Overview

Viruses are malicious code that infects files on your system, but worms do this one better—they can spread to other systems on the network as well. Many worms simply replicate, filling up memory and eating up network bandwidth until legitimate users are denied access to network services and servers have to be rebooted. Some worms are more malicious and include viral code that can corrupt files, steal documents and mail them away, or render systems inoperable.

Some examples of famous (or infamous) worms include the following:

- **ADMw0rm:** Exploited a buffer overflow in Berkeley Internet Name Domain (BIND)
- **Code Red:** Exploited a buffer overflow vulnerability in Microsoft Internet Information Services (IIS) 4 and 5
- **LifeChanges:** Exploited a vulnerability in Microsoft Windows that allowed scrap shell (*.shs) files to be used to execute arbitrary code
- **LoveLetter:** Exploited Visual Basic Script (VBScript) to mass mail itself to everyone in the Windows address book
- **Melissa:** Exploited a vulnerability in Microsoft Outlook and Outlook Express to mass mail itself to everyone in the Windows address book
- **Morris:** Exploited a vulnerability in the debug mode of Sendmail
- **Nimda:** Exploited the Hypertext Markup Language (HTML) IFRAME tag to automatically execute e-mail attachments in HTML messages
- **Slapper:** Exploited a buffer overflow vulnerability on the Apache Web server platform
- **Slammer:** Exploited a buffer overflow in unpatched machines running Microsoft SQL Server

See Also: Trojan, virus

WPA

Stands for Windows Product Activation, a Microsoft technology aimed at reducing software piracy.

See: Windows Product Activation (WPA)

WRM

Stands for Windows Rights Management, a new technology from Microsoft for secure content management.

See: Windows Rights Management (WRM)

WS-Security

Stands for Web Services Security, an emerging standard for adding authentication, confidentiality, and data integrity to Web services.

See: Web Services Security (WS-Security)

WWWhack

A tool for cracking password-protected Web sites.

Overview

Password cracking is the science (and art) of guessing at the password for an application or system until the correct one is found. While popular tools such as L0phtCrack and John the Ripper exist for cracking passwords used to authenticate users to Microsoft Windows– and UNIX-based systems and networks, other tools such as WWWhack can be used to crack a different kind of password, one that is used to control access by users to a Web site.

WWWhack can use both a dictionary attack or brute-force approach to cracking Web site credentials, and like most security tools, can be used either for good (auditing the security of password-protected sites) or bad (cracking passwords to steal sensitive data) purposes.

For More Information

Visit packetstormsecurity.org/Crackers/ for more information.

See Also: *John the Ripper, L0phtCrack, password cracking*

X.509

A standard format for digital certificates.

Overview

X.509 is a recommendation from the International Telecommunications Union (ITU) defining the type of information contained in a digital certificate in a Public Key Infrastructure (PKI) system. Such certificates are used for several purposes in a PKI system, including the following:

- Distributing public keys for encrypting messages to provide confidentiality.
- Signing messages to provide data integrity and authenticity.

There are three versions of the X.509 standard:

- **X.509v1:** The original X.509 certificate standard defined in 1988 and still widely used in many applications
- **X.509v2:** Adds features to deal with possible reuse of subject and issue names, but superseded by X.509v3 and not widely used
- **X.509v3:** Current X.509 standard defined in 1996 that supports arbitrary extensions to certificates based on RFC 2459 and other work of the Public-Key Infrastructure (PKIX) Working Group from the Internet Engineering Task Force (IETF)

X.509 is employed by many security protocols that use PKI systems, including Secure Sockets Layer (SSL), Secure Electronic Transaction (SET), and Secure/Multipurpose Internet Mail Extensions (S/MIME).

Implementation

The main elements of an X.509 digital certificate are the following:

- **Version:** X.509v3 is identified as version 2 since versions start at 0.
- **SerialNumber:** An integer that uniquely identifies the certificate to the certificate authority (CA) issuing it.
- **Signature:** The encryption algorithm used to generate the signature of the certificate.
- **Issuer:** The name of the CA that issued the certificate, expressed in X.500 naming format.
- **Validity:** The date/time when the certificate was issued and the date/time it expires.
- **Subject:** The name of the entity being issued the certificate, expressed in X.500 naming format.
- **SubjectPublicKeyInfo:** The encryption algorithm used to generate the public key for the entity being issued the certificate, and also the public key itself.
- **Encrypted:** The digital signature for the certificate.
- **Extensions:** Any additional fields recommended by PKIX or used by the PKI vendor for proprietary purposes. Examples may include key identifiers, key usage, certificate policies, and various constraints.

In addition to defining the information fields found in a certificate, X.509 also defines the data format used for encoding the certificate. This format is a combination of two standards:

- **Abstract Syntax Notation One (ASN.1):** Defines the detailed data format

- **Definite Encoding Rules (DER):** Defines how to store and transfer the data

See Also: digital certificate, Public Key Infrastructure (PKI), Public-Key Infrastructure (X.509) (PKIX), Secure Electronic Transaction (SET), Secure/Multipurpose Internet Mail Extensions (S/MIME), Secure Sockets Layer (SSL)

XACML

A standard for secure information access using Extensible Markup Language (XML).

Overview

XACML, which stands for Extensible Access Control Markup Language, is a standard that specifies how to express XML policies for controlling access to information over the Internet. XACML was developed by the Organization for the Advancement of Structured Information Standards (OASIS) and is the result of efforts by IBM, Sun, Entrust, and others to advance e-business by standardizing Web services security operations using XML. XACML is part of a broader secu-

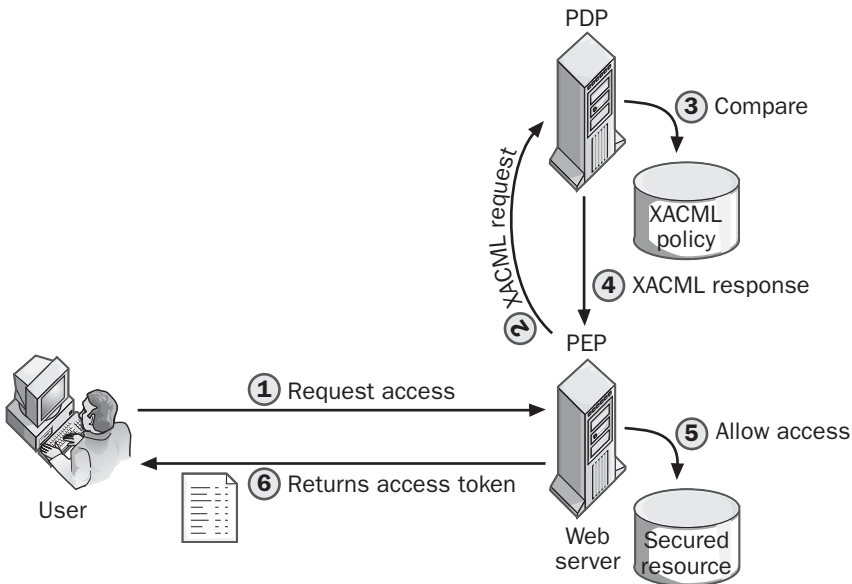
urity standards initiative from OASIS that includes Security Assertion Markup Language (SAML) and Web Services Security (WS-Security).

The advantage of XACML is that it represents an industry standard that can provide interoperability between diverse platforms and systems using different proprietary access control mechanisms. XACML also is an extensible framework that can evolve to support new access control policy needs and mechanisms. The current version of the standard is XACML 1.

Implementation

XACML defines two kinds of XML dialects for security purposes:

- A control language for expressing access control policies defining who has what kind of access to which resources
- A request/response language for expressing requests for access to resources and responses that allow or deny access



XACML. How XACML uses policy to control access to resources.

When a user wants to access a resource secured using XACML, the user first submits an access request to the Policy Enforcement Point (PEP), which is the system on which the resource resides; for example, a Web server. The PEP formulates an XACML request message that includes the security attributes of the user, the attributes of the resource, and the action the user wants to perform. The PEP forwards this message to the Policy Decision Point (PDP), an XACML authentication server that compares the XACML request message to stored security policies written using the XACML control language, which then makes a decision concerning whether to grant or deny the request. The PDP returns an XACML response message to the PEP, which then allows or denies access to the resource accordingly.

See Also: *Security Assertion Markup Language (SAML)*, *Web Services Security (WS-Security)*

Xauth

A UNIX tool for facilitating secure access to X Window System–based hosts.

Overview

Certain versions of X Window System can employ a token-based scheme for authenticating local or remote users. The token, called a magic cookie, is a randomly generated number stored in a file named `.Xauthority` and is used by clients to authenticate them for X connections. The resulting authentication scheme is transparent to users and is facilitated by Xauth, a program that can be used to display and edit these magic cookies. To log on to an X Window System–based host from a remote machine, the user first changes the magic cookie on the client machine and then uses Rsh or some other tool to copy the cookie to the X host. Using Xauth, X Window System authentication is user-based as opposed to the less secure host-based security implemented by Xhost.

See Also: *Xhost*

Xhost

A UNIX tool for specifying which hosts can make connections to a host running X Window System.

Overview

Xhost is a notoriously insecure way of managing security for X Window System–based systems. Using Xhost, an administrator can configure host-based authentication for X Window System, specifying which hosts are allowed to connect and which are not.

Though Xhost usually is recommended only for single-user workstation environments, administrators sometimes use Xhost as a simple method for allowing them to run X programs on remote systems while displaying them on the local machine. For example, by entering the command “`xhost +name`” (where *name* is the name of the remote system), you are actually allowing anyone who is logged on to the remote system to view the contents of your screen and log your keystrokes. If an attacker has compromised the remote host, this gives the attacker an opportunity to break into your system as well. Unfortunately, attackers do not require root privileges to modify access to X Window System–based machines using Xhosts.

A more secure (though less convenient) approach for managing X Window System security is using the Xauth tool instead.

See Also: *Xauth*

XKMS

Stands for XML Key Management Specification, an emerging standard for encrypting information based on Extensible Markup Language (XML).

See: *XML Key Management Specification (XKMS)*

XMAS scan

A port-scanning technique based on some obscure aspects of RFC 793, the Transmission Control Protocol (TCP) standard.

Overview

The XMAS scan is an extension of the FIN scan, a technique used to determine listening ports on a target host by sending a FIN packet to every port on the host. On some platforms, closed ports respond to FIN with RST packets, while listening ports make no response. This technique allows the attacker to enumerate which services are running on the target system.

The XMAS scan takes this approach further by sending a TCP packet with the FIN, PSH, and URG flags set. For target hosts whose operating system has an RFC-compliant implementation of TCP, closed ports respond with RST, while open ports are silent. Unfortunately (or fortunately, depending on your point of view), not all platforms comply with every detail of RFC 793. Therefore, for some platforms, such as Microsoft Windows, the results of this scan are unreliable.

See Also: *port scanning*

XMLDSIG

Stands for XML Signatures, an emerging standard for creating and managing digital signatures using Extensible Markup Language (XML).

See: *XML Signatures (XMLDSIG)*

XMLENC

Stands for XML Encryption, an emerging standard for encrypting information based on Extensible Markup Language (XML).

See: *XML Encryption (XMLENC)*

XML Encryption (XMLENC)

An emerging standard for encrypting information based on Extensible Markup Language (XML).

Overview

XML Encryption (XMLENC) is an emerging standard from the Internet Engineering Task Force (IETF) that can be used for encrypting information and represent-

ing the result using XML. Any kind of information can be encrypted this way, including whole XML documents, an element within an XML document, the content of an XML element, or arbitrary data such as a binary attachment.

The advantage of XMLENC is that it can use end-to-end encryption to facilitate secure e-commerce and e-business transactions over the Internet. This does not mean that XMLENC is intended as a replacement for other Web encryption schemes such as Secure Sockets Layer (SSL) or Transport Layer Security (TLS). Instead, XMLENC is meant to complement these technologies and be implemented mainly in business-to-business (B2B) environments for secure exchange of sensitive information expressed in XML syntax when more than two parties are involved. To perform encryption using XMLENC, a mechanism is required for securely exchanging the session keys used for encrypting and decrypting data. This is accomplished using XML Signatures (XMLDSIG), another emerging XML security standard.

The IETF is developing XMLENC in conjunction with the World Wide Web Consortium (W3C), which developed the core XML standards.

For More Information

Visit www.w3.org/TR/xmlenc-core/ for more information.

See Also: *encryption, Secure Sockets Layer (SSL), Transport Layer Security (TLS), XML Key Management Specification (XKMS), XML Signatures (XMLDSIG)*

XML Key Management Specification (XKMS)

An emerging standard for encrypting information based on Extensible Markup Language (XML).

Overview

XML Key Management Specification (XKMS) is an emerging standard from the Internet Engineering Task Force (IETF) that can be used for registering and distributing keys for any Public Key Infrastructure

(PKI) system. These keys can then be used with the following specifications:

- **XML Encryption (XMLENC):** Used for encrypting and decrypting information exchanged between several parties to provide confidentiality
- **XML Signatures (XMLDSIG):** Used for digitally signing information to provide data integrity and nonrepudiation

XKMS includes two main subspecifications that involve the Simple Object Access Protocol (SOAP) and Web Services Definition Language (WDSL):

- XML Key Registration Service Specification (X-KRSS)
- XML Key Information Service Specification (X-KISS)

The IETF is developing XMLENC in conjunction with the World Wide Web Consortium (W3C), which developed the core XML standards.

For More Information

Visit www.w3.org/TR/xkms/ for more information.

See Also: *Public Key Infrastructure (PKI), XML Encryption (XMLENC)*

XML Signatures (XMLDSIG)

An emerging standard for creating and managing digital signatures using Extensible Markup Language (XML).

Overview

Digital signatures are used in Public Key Infrastructure (PKI) systems to provide data integrity and nonrepudiation for information transmitted between parties. XML Signatures (XMLDSIG) is an emerging standard from the Internet Engineering Task Force (IETF) that can use XML syntax to digitally sign any kind of document, including e-mail messages, binary attachments, Web pages, or XML documents. XMLDSIG can use several

standard hashing algorithms for generating signatures, including both the Digital Signature Standard (DSS) and hash-based message authentication code (HMAC) used in conjunction with Secure Hash Algorithm-1 (SHA-1).

The advantage of XMLDSIG is that it can facilitate secure e-commerce and e-business transactions over the Internet. The IETF is developing XMLDSIG in conjunction with the World Wide Web Consortium (W3C), which developed the core XML standards, and has published the standards track RFC 3275 *XML-Signature Syntax and Processing*.

For More Information

Visit www.w3.org/TR/smlsig-core/ for more information.

See Also: *digital signature, Digital Signature Standard (DSS), hash-based message authentication code (HMAC), Public Key Infrastructure (PKI), XML Encryption (XMLENC), XML Key Management Specification (XKMS)*

Xscan

A tool used to scan for hosts running X Window System.

Overview

UNIX systems running X Window System can be vulnerable to attack if they are not configured and administered properly. For example, administrators who use the `Xhost +` command to grant themselves easy access to their X server from any other host on the network are leaving their X server wide open to attack. Attackers also can exploit the Xterm tool used to open a local command shell on an X server together with buffer overflows or other vulnerabilities to grant themselves shell access to the server. Because of such issues, attackers often use Xscan to scan remote networks for X servers they can try to compromise and use as launching points for further network intrusion.

Xscan works by scanning a specified subnet for X hosts. If an X host is found that has the Xhost + command allowing any remote host to connect to it, Xscan also logs all keystrokes typed on the console, showing the actions performed by administrators on the remote server in real time, which can quickly lead to root passwords being compromised and the attacker taking total control of the target.

See Also: Xhost, Xterm

Xterm

A UNIX tool for opening a local command shell on a host running X Window System.

Overview

Administrators commonly use Xterm for running commands to administer UNIX systems. Xterm is also

popular with attackers, however, since it can be used to provide them with shell access that allows them to run arbitrary code on the target system. In a typical exploit, an attacker uses a known vulnerability such as a buffer overflow in a Web server to execute a command on the remote target to run the Xterm command with the -Display option. This allows the attacker to direct the resulting shell to the attacker's own X server, providing shell access with complete control over the target system. Removing X entirely from your Web server is a simple way of eliminating this threat, but other methods such as reverse Telnet also can be used to try to gain shell access to a remote server.

See Also: reverse Telnet

Ypgrab

A tool for extracting password tables from Network Information System (NIS) hosts.

Overview

NIS is a protocol used on some UNIX platform network naming and directory services, and it functions as a kind of telephone book for locating resources on a Transmission Control Protocol/Internet Protocol (TCP/IP) network. Ypgrab is a tool used by crackers to obtain password tables from NIS and requires only that the attacker guess the name of the NIS domain. After obtaining the tables, the attacker can run common password-cracking tools to guess the passwords and use them to log on to the domain and try to perform further exploits.

The original version of NIS is relatively insecure, and Ypgrab is an older exploit that may not work on newer systems. Ypgrab is also ineffective against the more secure NIS+, which was developed by Sun Microsystems and includes added security features compared to NIS.

For More Information

For more information about NIS, see the *Microsoft Encyclopedia of Networking, Second Edition*, available from Microsoft Press.

See Also: *password cracking*

Zap

A tool for cleaning log files on UNIX systems.

Overview

Zap is a tool often used by attackers for covering their tracks once they've compromised a vulnerable system. Zap can zero out a number of log entries, including UTMP, WTMP, LASTLOG, and ACCT entries used to record login activity on UNIX platforms. Zap is only 2 KB in size and is often included as part of rootkits such as Knark. A newer version called Zap2, or Z2, is also available on many black hat sites.

See Also: *Knark, log cleaning, rootkit*

zombie

A compromised system used as an intermediary in a distributed denial of service (DDoS) attack.

Overview

A DDoS attack is a type of denial of service (DoS) attack that leverages the power of multiple intermediary hosts to overwhelm the target. The intermediary hosts generally are poorly secured systems connected to the Internet, which the attacker compromises and on which the attacker installs special DDoS agent software. This agent software, typically known as a **zombie agent**, is installed in a “sleeping” state and waits until the attacker issues a command telling it which host to target, whereupon it “wakes up” and begins sending spoofed packets to the target (the type of packet depends on the specific DDoS attack tool used).

The name **zombie** is used to describe such compromised hosts because in voodoo terminology a zombie is

a dead body into which a supernatural host has entered, causing it to become a “living dead” person. Other common names for such hosts are **drone** and **slave**.

Using large numbers of zombies is the key to a DDoS attack and provides the amplification factor that makes them so much more effective than traditional DoS attacks.

Notes

The origin of the term **zombie** actually predates DDoS attacks and refers to any program secretly installed on a compromised system and then remotely triggered later on for malicious purposes.

See Also: *denial of service (DoS), distributed denial of service (DDoS), Zombie Zapper*

Zombie Zapper

A free tool for stopping distributed denial of service (DDoS) attacks.

Overview

Zombie Zapper is an open source tool from Bindview Corporation's RAZOR security team that can be used as a countermeasure to a DDoS attack. Zombie Zapper works by impersonating the master program used by the attacker to control the zombies. To prepare for a DDoS attack, an attacker seeks out vulnerable hosts connected to the Internet and installs “zombie agent” software on them, leaving these “zombie systems” in a sleeping state until the attacker is ready to launch the attack. When the time is ripe, the attacker uses a “zombie master” program to issue a command to the zombies, which begin flooding the target with network traffic.

Zombie Zapper works by issuing commands telling the attacking zombies to go back to sleep. Zombie Zapper works only if the default password for launching the attack has not been changed by the attacker, which is often the case with attacks performed with the DDoS tools TFN, Trinoo, and Stacheldraht. With the TFN2K exploit, however, the attacker is forced to change the default password before launching the attack, and Zombie Zapper fails to stop the attack in this case.

Notes

Some other tools that work as countermeasures against certain DDoS exploits include the following:

- **Find_ddos:** Available from the National Infrastructure Protection Center (NIPC) (www.npic.gov)
- **DDOSPing:** Available from Foundstone (www.foundstone.com)

For More Information

Visit razor.bindview.com/tools/ for more information.

See Also: *distributed denial of service (DDoS), Stacheldraht, Tribal Flood Network (TFN), Tribal Flood Network 2000 (TFN2K), Trin00, zombie*

zone

Another name for security zone, a security feature implemented by Microsoft Internet Explorer for safer browsing.

See: *security zone*

Appendix I-

Applying Key Principles of Security-

Note The following document first appeared in the *Microsoft Windows Security Resource Kit*. To learn more about this book, visit <http://www.microsoft.com/MSPress/books/6418.asp>.

Managing information security is difficult. To do it well requires a combination of technical, business, and people skills, many of which are not intuitive. The foundation of information security is risk management.

Without a good understanding of risk management, it is impossible to secure any large modern network. More often than not, the failure of network administrators and managers to build a secure network results in the organization's most closely held information being as secure as the lunch menu. Thus, either the lunch menu will be very secure, or the security of important information will be very weak. Neither situation is workable in the long run.

Not every network administrator is a security expert, and most need not be. However, all network administrators must understand the basics of security. There are several key principles that you can follow to secure your networks and applications. By acting on these key principles when completing your day-to-day tasks, you can secure your network—even without being a security expert. And if you are a security specialist or want to become one, you must master these key principles.

Understanding Risk Management

The first key principle of security is that no network is completely secure—information security is really about risk management. In the most basic of terms, the more

important the asset is and the more it is exposed to security threats, the more resources you should put into securing it. Thus, it is imperative that you understand how to evaluate an asset's value, the threats to an asset, and the appropriate security measures. In general, without training, administrators respond to a security threat in one of three ways:

- Ignore the threat or acknowledge it but do nothing to prevent it from occurring
- Address the threat in an ad hoc fashion
- Attempt to completely secure all assets to the utmost degree, without regard for usability or manageability

None of these strategies takes into account what the actual risk is, and all of them will almost certainly lead to long-term failure.

Learning to Manage Risk

Managing security risks can be an incredibly daunting task, especially if you fail to do so in a well-organized and well-planned manner. Risk management often requires experience with financial accounting and budgeting as well as the input of business analysts. Building a risk assessment of an organization's security can take months and generally involves many people from many parts of the company. You can follow this simple process for assessing and managing risk:

- 1 Set a scope.** If you try to assess and manage all security risks in your organization, you are likely to be overwhelmed and certain to miss critical details. Before starting the risk assessment, set the scope of the risk assessment project. This will enable you to

better estimate the time and cost required to assess the security risks in the project and to more easily document and track the results.

2 Identify assets and determine their value. The first step in assessing risk is to identify assets and determine their value. When determining an asset's value, take these three factors into account:

- The financial impact of the asset's compromise or loss
- The nonfinancial impact of the asset's compromise or loss
- The value of the asset to your competitors

The financial impact of an asset's compromise or loss includes revenue and productivity lost because of downtime, costs associated with recovering services, and direct equipment losses. The nonfinancial impact of an asset's compromise or loss includes resources used in shaping public perception of a security incident, such as advertising campaigns, and loss of public trust or confidence, known as *goodwill in accounting*. The value of the asset to your organization should be the main factor in determining how you secure the resource. If you do not adequately understand your assets and their value, you might end up securing the lunch menu in the cafeteria as stringently as you secure your trade secrets.

1 Predict threats and vulnerabilities to assets. The process of predicting threats and vulnerabilities to assets is known as *threat modeling*. Through the exercise of modeling threats, you will likely discover threats and vulnerabilities that you did not know about or had overlooked, and you will document the more well-known threats and vulnerabilities. You can then proactively mitigate risk rather than having to react to it after a security incident.

2 Document the security risks After completing the threat model, it is essential that you document the security risks so that they can be reviewed by all relevant people and addressed systematically. When documenting the risks, you might want to rank them. You can rank risks either *quantitatively* or *qualitatively*. Quantitative rankings will use actual

and estimated financial data about the assets to assess the severity of the risks. For example, you might determine that a single incident of a security risk will cost your organization \$20,000 in financial losses while another will cost the organization only \$5,000. Qualitative rankings use a system to assess the relative impact of the risks. For example, a common qualitative system is to rank the product of the probability of the risk occurring and the value of the asset on a 10-point scale. Neither quantitative nor qualitative risk assessment is superior to the other; rather, they complement each other. Quantitative ranking often requires acute accounting skills, while qualitative ranking often requires acute technical skills.

3 Determine a risk management strategy. After completing the risk assessment, you must determine what general risk management strategy to pursue and what security measures you will implement in support of the risk management strategy. The result from this step is a risk management plan. The risk management plan should clearly state the risk, threat, impact on the organization, risk management strategy, and security measures that will be taken. As a security administrator, you will likely be responsible for or involved in implementing the security measures in the risk management plan.

4 Monitor the assets. Once the actions defined in the risk management plan have been implemented, you will need to monitor the assets for realization of the security risks. As we've alluded, realization of a security risk is called a *security incident*. You will need to trigger actions defined in contingency plans and start investigating the security incident as soon as possible to limit the damage to your organization.

5 Track changes to risks. As time progresses, changes to your organization's hardware, software, personnel, and business processes will add and obsolete security risks. Similarly, threats to assets and vulnerabilities will evolve and increase in sophistication. You will need to track these changes and update the risk management plan and the associated security measures regularly.

Risk Management Strategies

Once you have identified an asset and the threats to it, you can begin determining what security measures to implement. The first step is to decide on the appropriate risk management strategy. The rest of this section will examine the four general categories of risk management that you can pursue:

- Acceptance
- Mitigation
- Transference
- Avoidance

Accepting Risk

By taking no proactive measures, you accept the full exposure and consequences of the security threats to an asset. Accepting risk is an extreme reaction to a threat. You should accept risk only as a last resort when no other reasonable alternatives exist, or when the costs associated with mitigating or transferring the risk are prohibitive or unreasonable. When accepting risk, it is always a good idea to create a contingency plan. A contingency plan details a set of actions that will be taken after the risk is realized and will lessen the impact of the compromise or loss of the asset.

Mitigating Risk

The most common method of securing computers and networks is to mitigate security risks. By taking proactive measures to either reduce an asset's exposure to threats or reduce the organization's dependency on the asset, you are mitigating the security risk. Generally, reducing an organization's dependency on an asset is beyond the scope of a security administrator's control; however, the former is the primary job function of a security administrator. One of the simplest examples of mitigating a security risk is installing antivirus software. By installing and maintaining antivirus software, you greatly reduce a computer's exposure to computer viruses, worms, and Trojan horses. Installing and maintaining antivirus software does not eliminate the possibility of a computer being infected with a virus because there will inevitably be new viruses that the antivirus software cannot yet protect the computer against. Thus,

when a risk is mitigated, you still should create a contingency plan to follow if the risk is realized.

When deciding to mitigate risk, one of the key financial metrics to consider is how much your organization will save because of mitigating the risk, less the cost of implementing the security measure. If the result is a positive number and no other prohibitive factors exist, such as major conflicts with business operations, implementing the security measure is generally a good idea. On occasion, the cost of implementing the security measure will exceed the amount of money saved but will still be worthwhile—for example, when human life is at risk.

Transferring Risk

An increasingly common and important method of addressing security risks is to transfer some of the risk to a third party. You can transfer a security risk to another party to take advantage of economies of scale, such as insurance, or to take advantage of another organization's expertise and services, such as a Web hosting service. With insurance, you are paying a relatively small fee to recuperate or lessen financial losses if the security risk should occur. This is especially important when the financial consequences of your security risk are abnormally large, such as making your organization vulnerable to class action lawsuits. When contracting a company to host your organization's Web site, you stand to gain sophisticated Web security services and a highly trained, Web-savvy staff that your organization might not have afforded otherwise. When you engage in this type of risk transference, the details of the arrangement should be clearly stated in a contract known as a *service level agreement (SLA)*. Always have your organization's legal staff thoroughly investigate all third parties and contracts when transferring risk.

Avoiding Risk

The opposite of accepting risk is to avoid the risk entirely. To avoid risk, you must remove the source of the threat, exposure to the threat, or your organization's reliance on the asset. Generally, you avoid risk when there are little to no possibilities for mitigating or transferring the risk, or when the consequences of realizing

the risk far outweigh the benefits gained from undertaking the risk. For example, a law enforcement agency might want to create a database of known informants that officers can access through the Internet. A successful compromise of the database could result in lives being lost. Thus, even though many ways to secure access to the database exist, there is zero tolerance of a security compromise. Therefore, risk must be avoided by not placing the database on the Internet, or perhaps not storing the information electronically at all.

Understanding Security

The most fundamental skill in securing computers and networks is understanding the big picture of security. By understanding the big picture of how to secure computers and networks as well as the limitations of security, you can avoid spending time, money, and energy attempting impossible or impractical security measures. You can also spend less time resecurating assets that have been jeopardized by poorly conceived or ineffective security measures.

Granting the Least Privilege Required

Always think of security in terms of granting the least amount of privileges required to carry out the task. Excess privileges serve no useful business or technical purposes and can lead to users, administrators, or attackers taking advantage of them.

Defending Each Network Layer

Imagine the security of your network as an onion. Each layer you pull away gets you closer to the center, where the critical asset exists. On your network, defend each layer as though the next outer layer is ineffective or nonexistent. The aggregate security of your network will increase exponentially if you defend vigilantly at all levels.

Reducing the Attack Surface

Attackers are functionally unlimited and thus possess unlimited time, while you have limited time and resources. (The concept of being functionally unlimited is detailed in Appendix II, “Understanding Your Enemy.”) An attacker needs to know of only one vulnerability to attack your network successfully, while you must pinpoint all your vulnerabilities to defend

your network. The smaller your attack surface, the better chance you have of accounting for all assets and their protection. Attackers will have fewer targets, and you will have less to monitor and maintain.

Avoiding Assumptions

Making assumptions will generally result in you overlooking, prematurely dismissing, or incorrectly assessing critical details. Often these details are not obvious or are buried deep within a process or technology. That is why you must test everything! You might also want to hire a third party to assess the security of your network or applications. Some organizations might even have legal or regulatory compliance statutes that require them to undergo this type of evaluation.

Protecting, Detecting, and Responding

When you think about securing a computer or a network, think about how you can protect the asset proactively, detect attempted security incidents, and respond to security incidents. This is the security life cycle. By looking at security from this perspective, you will be better prepared to handle unpredictable events.

Securing by Design, Default, and Deployment

When you design networks, ensure that the following criteria are met:

- Your design is completed with security as an integral component.
- Your design is secure by default.
- The deployment and ongoing management of the implementation maintains the security of the network.

By accomplishing these three goals, you can address security proactively and natively rather than reactively and artificially.

The 10 Immutable Laws of Security

In 2000, Scott Culp of the Microsoft Security Response Center published the article “The Ten Immutable Laws of Security” on the Microsoft Web site, which you can read at <http://www.microsoft.com/technet/treeview/default.asp?url=/technet/columns/security/essays/10imlaws.asp>. Even though the Internet and computer

security are changing at a staggering rate, these laws remain true. These 10 laws do an excellent job of describing some of the limitations of security:

- 1 If a bad guy can persuade you to run his program on your computer, it's not your computer anymore.** Often attackers attempt to encourage the user to install software on the attacker's behalf. Many viruses and Trojan horse applications operate this way. For example, the ILOVEYOU virus succeeded only because unwitting users ran the script when it arrived in an e-mail message. Another emerging class of applications that attackers prompt a user to install are spyware applications. Once installed, spyware monitors a user's activities on his computer and reports the results to the attacker.
- 2 If a bad guy can alter the operating system on your computer, it's not your computer anymore.** A securely installed operating system and the securely procured hardware that it is installed on are referred to as a *Trusted Computing Base (TCB)*. If an attacker can replace or modify any of the operating system files or certain components of the system's hardware, the TCB can no longer be trusted. For example, an attacker might replace the file `Passfilt.dll`, which is used to enforce password complexity with a version of the file that also records all passwords used on the system. If an operating system has been compromised or you cannot prove that it has not been compromised, you should no longer trust the operating system.
- 3 If a bad guy has unrestricted physical access to your computer, it's not your computer anymore.** Once an attacker possesses physical access to a computer, you can do little to prevent the attacker from gaining administrator privileges on the operating system. With administrator privileges compromised, nearly all persistently stored data is at risk of being exposed. Similarly, an attacker with physical access could install hardware or software to monitor and record keystrokes that is completely transparent to the user. If a computer has been physically compromised or you cannot prove otherwise, you should not trust the computer.
- 4 If you allow a bad guy to upload programs to your Web site, it's not your Web site anymore.** An attacker who can execute applications or modify code on your Web site can take full control of the Web site. The most obvious symptom of this is an attacker defacing an organization's Web site. A corollary to this law is that if a Web site requests input from the user, attackers will use bad input. For example, you might have a form that asks for a number between 1 and 100. While normal users will enter numbers within the specified data range, an attacker will try to use any data input he or she feels will break the back-end application.
- 5 Weak passwords trump strong security.** Even if a network design is thoroughly secure, if users and administrators use blank, default, or otherwise simple passwords, the security will be rendered ineffective once an attacker cracks the password.
- 6 A machine is only as secure as the administrator is trustworthy.** One constant on all networks is that you must trust the network administrators. The more administrative privileges an administrator account has, the more the administrator must be trusted. In other words, if you do not trust someone, do not give him or her administrator privileges.
- 7 Encrypted data is only as secure as the decryption key.** No encryption algorithm will protect the ciphertext from an attacker if he or she possesses or can gain possession of the decryption key. Encryption alone is not a solution to a business problem unless there is a strong component of key management and unless users and administrators are vigilant in protecting their keys or key material.
- 8 An out-of-date virus scanner is only marginally better than no virus scanner at all.** New computer viruses, worms, and Trojan horses are always emerging and existing ones are always evolving. Consequently, antivirus software can become outdated quickly. As new or modified viruses are released, antivirus software is updated. Antivirus software that is not updated to recognize a given virus will not be able to prevent it.

- 9 Absolute anonymity isn't practical, in real life or on the Web.** Two issues related to security that are often confused are *privacy* and *anonymity*. Anonymity means that your identity and details about your identity are completely unknown and untraceable, while privacy means that your identity and details about your identity are not disclosed. Privacy is essential, and technology and laws make achieving it possible. On the other hand, anonymity is not possible or practical when on the Internet, or when using computers in general.
- 10 Technology is not a panacea.** Although technology can secure computers and computer networks, it is not—and will never be—a solution in and of itself. You must combine technology with people and processes to create a secure computing environment.

The 10 Immutable Laws of Security Administration

Microsoft's Scott Culp wrote "The Ten Immutable Laws of Security Administration," which you can find at <http://www.microsoft.com/technet/treeview/default.asp?url=/technet/columns/security/essays/10salaws.asp>. These 10 laws address the security issues that network administrators must contend with, issues entirely separate from the day-to-day security concerns of users:

- 1 Nobody believes anything bad can happen to them, until it does.** Because attacks on computer networks often cannot be seen, felt, or heard, it is easy for users and administrators to place them out of their minds. With attacks far from one's mind, it is difficult to see the need for security. Unfortunately, after a security incident takes place, the need for security is frequently still dismissed and the breach regarded as a one-time incident. Attackers *will* attempt to compromise the security of your network. It is not a question of if or when—it is a question of how frequently they will do it. You must protect your networks against attackers, detect their attempts to compromise your network, and respond when security incidents do occur.
- 2 Security works only if the secure way also happens to be the easy way.** For most users and administrators, the more difficult or invasive a security measure is, the more likely they are to ignore it, forget it, or subvert it. Ideally, security should be transparent to users and administrators. When the security measure requires a user or an administrator to change his or her behavior, you should create clear and easy-to-follow procedures for completing the task in question and explain your rationale for implementing the security measure.
- 3 If you don't keep up with security fixes, your network won't be yours for long.** After a security update is announced and the vulnerability is explained, a race begins between attackers attempting to exploit the vulnerability and administrators attempting to apply the security update. If you do not keep up with applying security updates, an attacker will exploit one of the known vulnerabilities on your network.
- 4 It doesn't do much good to install security fixes on a computer that was never secure to begin with.** Although installing security updates will prevent exposure to newly discovered vulnerabilities, installing security updates in and of itself will not result in a secure computer. For a computer to be secure, it is essential that the base operating system be securely configured.
- 5 Eternal vigilance is the price of security.** Security is an ongoing effort. The security administrator must remain vigilant to attacks and attackers who constantly strive to increase the level of sophistication of their attacks. An infinite number of potential attackers exists, and they have infinite time on their hands to crack your network. Attackers have little to lose and need to know only one exploit. Security administrators, on the other hand, have a finite amount of time and resources to defend their organization's network. A security administrator is defeated when a single attack is successful against the network.

- 6 There really is someone out there trying to guess your passwords.** Because of the mythic qualities surrounding attackers—much like the monster under the bed—it is easy to push the possibility of attackers out of one’s mind. Unlike the monster under the bed, attackers do exist and they do attack networks. In movies, attackers break powerful encryption algorithms; in real life, they guess simple passwords and exploit mundane, known vulnerabilities.
- 7 The most secure network is a well-administered one.** Although a security expert can secure a network, it will not remain secure if it is not well managed—from the CIO to the security administrator to the user.
- 8 The difficulty of defending a network is directly proportional to its complexity.** The more complex a network is, the greater the chance for administrators to misconfigure computers, lose track of the configuration of computers, and fail to understand how the network really works. When in doubt, keep it simple.
- 9 Security isn’t about risk avoidance; it’s about risk management.** You will never avoid all security risks. It would be too costly and impractical. Claims of unbreakable security stem from ignorance or arrogance.
- 10 Technology is not a panacea.** Although it is essential to ensure the bits and bytes on your network are configured securely, doing so will not prevent rogue administrators, poor processes, careless users, or apathetic managers.

Appendix II-

Understanding Your Enemy-

Note The following document first appeared in the *Microsoft Windows Security Resource Kit*. To learn more about this book, visit <http://www.microsoft.com/MSPress/books/6418.asp>.

Knowing Yourself

In respect to information security, knowing yourself and your enemy is not necessarily a straightforward endeavor—if it were, networks would be much more secure than they are today. To know yourself, you must do the following:

- Accurately assess your own skills.
- Possess detailed documentation of your network.
- Understand the level of organizational support you receive.

Accurately Assessing Your Own Skills

The skill set of a network administrator should include formal training on operating systems and applications; experience designing, installing, and configuring networks and network services; and the ability to predict problems before they occur and solve them when they do. To prevent design and configuration mistakes that can lead to security breaches, you must be able to accurately assess your network management skill set. Overestimating your knowledge of a network, operating system, or application can easily lead to vulnerabilities that attackers can exploit. Accurately assessing your skill set enables you to be proactive in obtaining training and acquiring the services of experienced consultants if the situation requires it.

For example, you might be asked to install and configure an Internet Web server for customers to access their order history on a Web application that your organization is deploying. Although you might be an experienced MCSE who has installed and configured intranet Web servers, you might not have any knowledge or experience with Internet Web applications or configuring servers that have direct Internet connectivity. By not accurately assessing your skills, you could easily and unwittingly expose customer information to attackers and not realize it until the information has already been compromised.

Possessing Detailed Documentation of Your Network

A key requirement of securing your organization's network is maintaining detailed documentation about the physical infrastructure of your network, complete and up-to-date network diagrams, documentation of the configuration of computers, applications, and the audit log. Without this documentation, network administrators might overlook resources that must be secured, and the network will almost certainly be inconsistent in its level of security. Without baseline performance and security information, it is difficult to detect attacks, regardless of whether they are successful. For example, your network might have a direct connection to the Internet that is no longer used but is still connected to a router. Over time, a router's access control list (ACL) can become outdated and can present a significant security risk. This is because the outdated ACL can enable an attacker to compromise your organization's network by using tactics that did not exist when the router was

secure and was monitored. Although such situations might seem obscure, they are quite common for organizations that have grown by being acquired by another organization. When consolidating IT resources, you can easily overlook these types of details. Similarly, it often takes such organizations a long time to create detailed documentation on their newly formed networks.

Understanding the Level of Organizational Support You Receive

The level of support that you receive from your organization—from management to your end users—greatly determines how you will secure network resources. This is often called your organization's *security position* or *security posture*. The security position of your organization includes the level of executive sponsorship for security policies and procedures, security requirements mandated by industry or government regulations, end user compliance with security policies and procedures, and training for end users and administrators. Your organization's security policies and procedures are central to the level of organizational support that you have.

In general, the completeness and clarity of an organization's security policies and procedures can indicate the support network administrators will receive for securing a network. Failing to understand your organization's security position can result in you oversecuring network resources to the point that end users will work around security measures and cause security vulnerabilities. For example, your organization might create policies that greatly restrict the types of Web applications that can be installed on a Web server, causing departments to purchase and deploy their own Web servers. Because the IT department does not know of the rogue Web servers, it cannot manage the application of security updates and service packs to those servers.

Identifying Your Attacker

Knowing your enemy is as complicated as knowing yourself—maybe even more so. Too often, network administrators know their enemies only through stereo-

types of attackers, and like most stereotypes, these are generally not accurate and rely on fear. For example, when you see movies that portray computer crime, more often than not, the penetration of the computer systems involves breaking an encryption key. The movie attacker fiercely pounds at his or her keyboard to break the encryption key by guessing it, which usually happens within a matter of seconds. Or he or she quickly writes a program with a well-designed user interface featuring big numbers that crack each character in the encryption key one by one. Although both attacks add drama to these movies, they are not only mathematically absurd and impossible, they also are not an accurate depiction of how networks are attacked. If this is all you know about the people who will attack your network, your network will be compromised.

In reality, breaking an industry-standard encryption key such as the 25-year-old Data Encryption Standard (DES) algorithm takes special hardware, significant computer programming skills, and plenty of time. To prove the insecurity of the DES algorithm, the Electronic Frontier Foundation (EFF) spent more than a year building a computer, using custom-built hardware and software, that could crack a 56-bit DES key. It took three days to crack the key.

You could design and build a network more secure than a government currency vault, but it would take only one computer that does not have the latest service pack installed for an attacker to compromise the network. A computer network looks very different from the attacker's point of view than from your viewpoint, as the defender. For example, you might think applying a security update for a known vulnerability to all but one computer on your network will be successful. To the attacker, this lone computer without the security update is the key to compromising the network.

By understanding (or “knowing”) the attacker, you can think like an attacker when designing security for your network. For example, many organizations complete vulnerability assessments on their networks. But you might want to consider training members of your organization's IT staff or hiring external experts to attempt to break into the network from the outside. In fact, there

are those in the field of computer security who boldly assert that you cannot secure a computer network without being able to attack one.

When most people think about attackers, or “hackers,” they generally think of a know-it-all, 14-year-old boy who wears a black T-shirt every day and is pale as a vampire as a result of all the hours he spends in front of his computer or videogame console. Although this stereotypical attacker certainly exists, he represents only a small portion of the attacker population. For convenience, we’ll group attackers into two general categories: external attackers (those outside your organization) and internal attackers (those within it).

Understanding External Attackers

The majority of attackers that you hear about in the media work outside the organizations they attack. These attackers include everyone from teenagers to professional hackers employed by governments and rogue nations. In addition to the attackers who are outright malicious, there exist groups of self-styled “white hat,” or nonmalicious, attackers. Although these attackers might not have malicious intentions, they also present significant dangers to networks. For example, a “harmless” attacker might break into a network for the challenge, but while attempting to compromise a server, might render it inoperable, resulting in a denial of service (DoS) condition. When examining attackers, it can be helpful to think about the dangers they present in terms of their skill level—be it novice, intermediate, or advanced.

Novice Attackers

Novice attackers generally possess only rudimentary programming skills and basic knowledge of the inner workings of operating systems and applications. These attackers represent the majority of attackers. Although this group of attackers might not possess significant skills, they are a threat to networks primarily because of the number of them out there and the knowledge they lack. For example, a novice attacker is much more apt to destroy information (either intentionally or accidentally) even though it will reveal his or her compromise of the network and quite possibly result in being arrested. Although secure networks will rarely be com-

promised by novice attackers, networks that are not vigilantly secured are extremely vulnerable to this type of attacker because of the sheer number of them. Novice attackers exploit known vulnerabilities with tools created by more experienced attackers, and thus are often called *script kiddies*. They also present a serious threat to obvious security vulnerabilities, such as weak passwords. Novice attackers who are also employees (making them internal attackers) often present the same level of danger as external attackers because they already possess valid network credentials from which they can launch attacks and they have access to network documentation.

Intermediate Attackers

Attackers with intermediate skills are less numerous than novice attackers but generally possess programming skills that enable them to automate attacks and better exploit known vulnerabilities in operating systems and applications. This group of attackers is capable of penetrating most networks if given enough time, but they might not be able to do so without being detected. These attackers frequently port attacks from other operating systems and conduct more sophisticated attacks than novice attackers. Attackers with an intermediate skill level often launch such attacks as an attempt to increase their notoriety or boost their skill level by creating tools to attack networks and publishing information that helps other attackers break into networks.

Advanced Attackers

Attackers with advanced skills usually are not only accomplished programmers but also have experience breaking into networks and applications. These attackers discover vulnerabilities in operating systems and applications and create tools to exploit previously unknown vulnerabilities. Advanced attackers are generally capable of compromising most networks without being detected, unless those networks are extremely secure and have well-established incident response procedures.

Understanding Internal Attackers

Contrary to what you might hear in the media, the majority of attacks on networks are conducted by attackers who have company badges—in other words, your fellow employees. Attackers who are employees of the organization they're attacking present a unique danger to networks for several reasons. Such attackers have the following in their favor:

- Higher levels of trust
- Physical access to network resources
- Human resources protections

Higher Levels of Trust

Almost all networks place a much higher level of trust in users and computers accessing resources on the local area network (LAN) than on publicly available network resources, such as servers connected to the Internet. Many networks allow authentication methods and unencrypted data transmissions on LANs that they would never consider using on the Internet. It is also much easier for attackers to enumerate information about the configuration of computers and applications when they have valid credentials on the network. Employees have valid credentials to the network, which also gives them greater initial access to network resources than external attackers might initially have. It can be very difficult to discern whether an employee is using his or her credentials legitimately or illegitimately—especially when the person is a network administrator.

Physical Access to Network Resources

Employees have much greater physical access to network resources—namely, the computers of their coworkers. In general, when an attacker has physical control of a computer, that computer can no longer be protected from the attacker; rather, it is only a matter of time and computing power before the attacker can recover all data on the computer. Similarly, employees have much greater access to documentation on the network, which can be a critical resource for attacking it.

Human Resources Protections

Employees, even those who attack network resources, are often protected by employment laws and HR policies that can greatly hinder their employer from detecting them or preventing them from doing further damage once detected. For example, local laws might prohibit an organization from inspecting the Internet usage of its employees without a court order. An employee could take advantage of this by attacking internal Web resources.

What Motivates Attackers?

Attackers attempt to break into computer networks for many reasons. Although all attackers present a clear and present danger to networks, the motivation of the attacker will greatly determine the actual threat posed. By understanding what might motivate potential attackers to attempt to compromise your organization's network, you can predict what type of threats the network faces. Armed with this knowledge, once you detect an attack, you might be more able to prevent further damage or better equipped to identify who the attacker is.

Many attackers are motivated by more than one factor. Here are the reasons that attackers attempt to break into computer networks, in ascending order of the danger they present:

- Notoriety, acceptance, and ego
- Financial gain
- Challenge
- Activism
- Revenge
- Espionage
- Information warfare

Notoriety, Acceptance, and Ego

An attacker's quest for notoriety, desire for acceptance, and ego comprise one of the most common motivations for attempts to break into computer networks and applications. Attackers motivated by notoriety often are naturally introverted and seeking a way to gain acceptance

in the electronic hacker community; thus, their exploits are very public. Examples of such attacks include defacing Web sites and creating computer viruses and worms.

By breaking into a network of a major company or government agency and defacing its Web site, an attacker is virtually guaranteed national and international publicity and enshrined in the electronic hacker community. For example, Attrition.org runs a Web site that catalogs nearly all Web site defacements in recent years. Querying any major search engine for the phrase Web site defacement invariably returns thousands of accounts of an organization's Web site being defaced, including those of most major corporations and government agencies.

Although not normally regarded as attackers, people who create and release computer viruses and worms cause billions of dollars of damage each year. In 1991, the Michelangelo virus opened a Pandora's box of sorts for computer viruses. Although the Michelangelo virus did little actual damage, the coverage that it received in the mainstream media, including newspapers, magazines, and television news, brought computer viruses into the popular consciousness and opened the door for other malicious publicity seekers. Since then, many other computer viruses have created similar media frenzies, such as Fun Love, I Love You, Melissa, and most recently, Code Red and NIMDA.

Popular media and antiauthoritarian romanticism transformed outlaws of the American western frontier—such as Jesse James and Billy the Kid—from common criminals who robbed banks and murdered people into cult heroes. Similarly, several attackers have gained cult hero status in the hearts and minds of computer geeks. Two recent examples include Kevin Mitnick and Adrian Lamo. Other attackers and prospective attackers seek the attention of the media and hacker communities that Mitnick and Lamo received and are envious, if not worshipful. The cult following of these two hacker legends is particularly strong with impressionable teenagers who have not fully developed their own sense of morality and rarely understand the true consequences their

actions have on business continuity and information technology.

In all these examples and in many similar incidents, the exploits of the attackers received international publicity. Attackers motivated by notoriety, acceptance, and ego look at these incidents as proof that they too can become famous. You can probably imagine the sense of accomplishment an attacker might feel, seeing his handiwork in the headlines of major newspapers and discussed on television news programs by political pundits. Often attackers know that their actions are illegal but consider their behavior harmless because there is no clear victim, no one physically harmed, and no tangible goods stolen or destroyed. Thus, in the minds of many attackers, they are not doing anything discernibly wrong. Certainly this is not the case. For example, although the direct financial consequences of Web site defacements are often low, the loss of public confidence in how well the organization can ensure the confidentiality and privacy of their employee, business partner, and customer information can be severe. This can result in indirect financial losses from customer distrust and defection.

Financial Gain

We can separate attackers motivated by monetary gain into two categories: those motivated by direct financial gain, and those motivated by indirect financial gain.

Attackers motivated by direct financial gain are little more than common criminals, akin to bank robbers with computer skills. These attackers break into computer networks or applications to steal money or information. In the past few years, there have been several high-profile thefts of credit card information from the databases of companies that conduct online commerce. These attackers did one of three things with the credit card information that they stole: they used the credit cards to purchase products or make cash withdrawals, sold the credit card numbers to other criminals, or attempted to extort money from the companies from which they stole the credit cards. In nearly every case, the attacker was apprehended, but not before causing significant damage. For example, in 1994, a Russian attacker broke into Citibank and transferred roughly

\$10 million to accounts in several countries. He was captured, and all but \$400,000 was recovered. But the real damage to Citibank was in their customers' loss of trust because of Citibank's inability to secure their customers' bank accounts. The attacker was sentenced to three years in prison and fined \$240,000, whereas U.S. Federal Sentencing Guidelines call for a minimum 6–10 year sentence for someone with no prior criminal record who robs a bank in person.

Another way that attackers seek financial gain from attacking networks and applications is to successfully break into an organization's network and then offer to help the organization secure the network. Although many of these attackers maintain the position that they are "good guys" wanting only to help the target organization, in reality, they are little more than extortionists demanding "protection money," like a 1920s gangster in cyberspace.

Some attackers are motivated by financial gain but in an indirect manner. A researcher or computer security company might make a large effort to discover vulnerabilities in commercial software applications and operating systems, and then use their discovery and the publication of such previously unknown vulnerabilities as a marketing tool for their own security assessment services. The publicity that a company or individual receives from unearthing a serious vulnerability in a commercial software application, especially a widely used application, can be priceless. For example, most significant vulnerabilities discovered in a widely used software application will be reported on the front page of major news and computer industry Web sites and in the technology or business sections of major newspapers. The discoverer of such a vulnerability might even receive airtime on the cable news television networks. For most small computer consulting companies, obtaining this type of publicity normally would be out of the question.

There is a critical point in the process of discovering commercial software vulnerabilities when one leaves the realm of ethical behavior and becomes an attacker: the reporting of that vulnerability to the general public without the software company's knowledge or consent.

Most commercial software companies are more than willing to work with researchers who have discovered security vulnerabilities to ensure that a software patch is available before the vulnerability is announced. Many software companies will also give credit to the person and company that discovers the vulnerability, thus balancing the interests of their software users with the public recognition earned by the person and company reporting the vulnerability. However, many researchers not only publish the vulnerability without notifying the software vendor, they also create code to exploit the vulnerability. Further complicating this issue are laws such as the 1998 Digital Millennium Copyright Act (DMCA), which prohibits individuals from exposing vulnerabilities in certain software and hardware encryption techniques used for digital rights management. The bottom line is this: although discovering vulnerabilities for indirect financial gain can be done illegitimately via extortion, it can also be done legitimately to advance the mutual business goal of software vendors and researchers—protecting consumers.

Challenge

Many attackers initially attempt to break into networks for the mere challenge. In many ways, attackers view networks as a game of chess—a battle of minds that combines strategic and tactical thinking, patience, and mental strength. However, chess has precisely defined rules, and attackers clearly operate outside the rules. Attackers motivated by the challenge of breaking into networks often do not even comprehend their actions as criminal or wrong. Attackers motivated by the challenge are often indifferent to which network they attack; thus, they will attack everything from military installations to home networks. These attackers are unpredictable, both in their skill level and dedication.

Activism

One newer type of attacker is the *hactivist*, an attacker who breaks into networks as part of a political movement or cause. This type of attacker might break into a Web site and change the content to voice his own message. The "Free Kevin Mitnick" hactivists frequently did this in an attempt to get Mitnick released from U.S. federal custody after he was arrested on multiple counts

of computer crime. Attackers motivated by a specific cause might also publish intellectual property that does not belong to them, such as pirated software or music. They might carry out sophisticated DoS attacks, called *virtual sit-ins*, on major Web sites to call attention to a particular cause.

Revenge

Attackers motivated by revenge are often former employees who feel they were wrongfully terminated or hold ill will toward their former employers. These attackers can be particularly dangerous because they focus on a single target and—being former employees—often have intricate knowledge of the security of the networks. For example, on July 30, 1996, employees of Omega Engineering arrived at work to discover that they could no longer log on to their computers. Later they discovered that nearly all their mission-critical software had been deleted. The attack was linked to a logic bomb planted by an administrator who had been fired three weeks earlier. The attack resulted in more than \$10 million in losses, prompting the layoff of 80 employees. In early 2002, the former administrator was sentenced to 41 months in prison, which pales in comparison to the financial and human damages that he caused.

Espionage

Some attackers break into networks to steal secret information for a third party. Attackers who engage in espionage are generally very skilled and can be well funded. Two types of espionage exist: industrial and international. A company might pay its own employees to break into the networks of its competitors or business partners, or the company might hire someone else to do this. Because of the negative publicity associated with such attacks, successful acts of industrial espionage are underreported by the victimized companies and law enforcement agencies. A widely publicized industrial espionage incident using computers recently took place in Japan. In December 2001, an engineer at Japan's NEC Toshiba Space Systems broke into the network of the National Space Development Agency of Japan. This engineer illegally accessed the antenna designs for a high-speed Internet satellite made by Mitsubishi in an

attempt to help NEC gain business from the space agency. As a result, the Japan Space Agency prohibited NEC from bidding on new contracts for two months, but no criminal charges were filed.

Attackers who engage in international espionage attempt to break into computer networks run by governments, or they work for governments and rogue nations to steal secret information from other governments or corporations. The most famous case of computer-related international espionage is documented in Cliff Stoll's book *The Cuckoo's Egg: Tracking a Spy Through the Maze of Computer Espionage* (Pocket Books, 2000). In 1986, Stoll, an astronomer by trade, was working as a computer operator at Lawrence Berkeley Lab when he discovered a 75-cent discrepancy in an accounting log from the mainframe computer. One thing led to another, and eventually Stoll discovered that German attackers being paid by the KGB were breaking into both military and nonmilitary computers to steal secret information.

Information Warfare

Information warfare is another motivation for attacking computer networks that is becoming increasingly dangerous as people around the world rely on them for mission-critical services. Major wars have been marked by the evolution of weapons systems—the machine gun changed the nature of combat in World War I, the tank changed the nature of combat in World War II, and airpower changed the nature of combat in Vietnam. Behind the scenes, each war also marked the evolution of electronic combat. From intercepted telegrams broken by hand, to radar jamming, to satellite transmissions that could be broken only by stealing the encryption keys (despite the power of many supercomputers)—electronic combat and intelligence has become a deciding factor in modern warfare. Although no widely reported incidents of cyberterrorism exist, you can be certain that these attempts have been made. There is evidence of information warfare in China, Israel, Pakistan, India, and the United States. The U.S. President's Critical Infrastructure Protection Board was formed in 2001 specifically to

address countering the threat of cyber-terrorism and information warfare against the United States.

Why Defending Networks Is Difficult

In traditional combat, defenders enjoy a distinct advantage over their attackers. However, in information technology, several factors give attackers the advantage:

- Attackers have unlimited resources.
- Attackers need to master only one attack.
- Defenders cannot take the offensive.
- Defenders must serve business goals.
- Defenders must win all the time.

Attackers Have Unlimited Resources

At any given time, defenders must protect their network against both attackers around the globe and their own employees. This accumulation of attackers, as a group, limits a defender's resources. Many attackers can spend all day systematically attempting to break into your network. Attackers can collaborate to develop new and more sophisticated attacks. As a network administrator, you have other duties besides defending the network, and unlike attackers, you go home at night, take sick days, and go on vacations. Over time, some attackers will cease attempting to break into your network, but new ones will take their place. Defending networks against unrelenting hoards of attackers with much more time than you gives attackers an advantage.

Attackers Need to Master Only One Attack

As a network administrator, you have to secure many servers and applications. You must learn how all your operating systems, applications, and network devices work, as well as how to secure and manage them. You must determine the threats to each component of your network and keep current with newly reported vulnerabilities. Attackers, on the other hand, need to master attacking only a single application or operating system feature in order to compromise it and break into your network.

Defenders Cannot Take the Offensive

Although attackers can attack networks with a certain amount of impunity, defenders can retaliate only through litigation, which is expensive and time-consuming. Attacking an attacker is not only illegal in most countries, it is impractical. This is because attackers often use previously compromised third-party computers, called *zombie systems*, or many zombie systems acting in unison to attack networks. By using zombie systems to carry out or amplify an attack, an attacker can protect her identity. Frequently attackers use the networks of colleges and universities as an attack vector because of their openness, computing power, and bandwidth. An attack can also originate from another legitimate organization whose employee has attacked your network, or whose network has already been compromised by an intruder. In any of these cases, retaliating against an intruder can result in your organization illegally attacking an unwitting individual, company, or organization. Thus, legally and practically, you cannot retaliate against attackers.

Defenders Must Serve Business Goals

Although network administrators are responsible for securing their organizations' networks, they also must install and configure operating systems and applications that help employees meet the goals of the business. In some situations the pursuit of company business goals conflicts with maintaining the security of the network.

For example, company executives might travel with laptops that contain sensitive information about the company. The executives might be unwilling to comply with security policies that require long and complex passwords. Knowing this, a network administrator might supply the executives with smart cards that they must use to access their laptops. This security measure will better protect the information on the laptop, but it also introduces other potential problems, such as the loss or misplacement of smart cards. To mitigate this situation, a network administrator might decide to create a second account for local computer users that could be used without the protection of a smart card, granting certain trusted employees the new account password,

which could result in a serious security vulnerability. Another situation in which the pursuit of business goals can interfere with the protection of the network occurs when your organization has a business rule that conflicts with the security of the network. For example, your organization might have a business rule that requires network traffic to the payroll server to be encrypted. This security measure will make data transmission of employee compensation secure, but makes it impossible for you to monitor network traffic to determine whether traffic is legitimate or illegitimate. It also prevents you from using any type of network intrusion detection software. In both scenarios, having to serve business goals jeopardizes your ability to protect the network.

Defenders Must Win All the Time

An attacker needs only one successful attack to compromise a network, while a network administrator must prevent *all* attacks to succeed in his role. These are ominous odds for ill-equipped or under-resourced network administrators. Given all the other problems defenders of networks face, it is inevitable that the security of your network will be compromised at some point. As a network administrator, you must ensure that these compromises are detected early and happen infrequently.

Is defending a network impossible? Not at all. But one thing is certain: it is impossible to defend a network without trained, skilled, and knowledgeable network administrators. By applying the key principles of security to the information this book presents on securing computers running Microsoft Windows 2000 and Windows XP, you can build a strong foundation for defending your networks.

Appendix III

Threats and Risk Assessment

Note The following content is from *Secure Messaging with Microsoft Exchange 2000* (Microsoft Press, 2003) by Paul Robichaux.

You can go a long way with a smile. You can go a lot farther with a smile and a gun.

– Attributed to Al Capone

Risk and threat assessment is something humans are notoriously bad at. Examples abound: try asking 10 of your coworkers whether it's more dangerous to fly or drive from Seattle to Denver and see how many of them correctly identify air travel as less risky. Then ask the same group whether the risk of dying in a commercial airline crash is greater or less than the risk of being struck by lightning. Sometimes our inability to properly assess risks is based on a lack of solid objective data about what the risks are, and sometimes the cause is an unwillingness to fully evaluate the threat and the corresponding risks.

This appendix helps you begin to understand the process of threat and risk assessment. This is normally the domain of skilled security practitioners, and you won't necessarily be able to completely evaluate your messaging system risks when you are done reading this book. However, you will be much better prepared to understand what risks you actually face (as opposed to the ones you think will give you trouble), and you'll have a better understanding of how to go about mitigating them.

First, a brief vocabulary lesson. A *threat* is something bad that can happen. Common threats include virus

attacks, internal or external network penetrations, theft of data, eavesdropping, and server failure. A *risk* is the product of two things: the likelihood that a particular threat will occur and the expected damage if it does. For example, my car might be stolen from the airport parking lot. That's a threat. My personal risk is low, though, because my auto insurance will replace the car if it's stolen; I've essentially transferred that risk to someone else. On the other hand, the risk that I'll have to wash my car when I return home is high. The threat (mostly posed by bird droppings) is likely to occur (that is, birds are very likely to fly around and over the car), and the expected effect (that is, bird droppings on the sunroof) is predictable. Professional risk assessors also factor in the frequency of the threat; something that is guaranteed to happen every year and causes moderate damage might be a bigger risk than something that might only happen every 50 years but causes more damage. For a real-world perspective on risks and frequency, consider mudslides and earthquakes in California, hurricanes in the Carolinas or Florida, and tornadoes in Kansas and northern Alabama.

Although statistical risk assessment is a rigorous process that requires a disciplined approach, you can do your own risk assessments. For every risk you identify, you need to do one of four things:

- **Avoid the risk.** This is the simplest (and often the least feasible) approach. If something seems risky, don't do it. If you're worried about e-mail-borne viruses, you can disconnect your servers from the Internet—a measure that would give you pretty good protection, if not good communications. If you're concerned about hackers attacking your factory-floor control systems through your Internet

connection, you might choose to isolate them on a self-contained internal network with no direct or indirect connectivity to other networks.

- **Mitigate the risk.** You can do this by either reducing the associated loss or blocking the associated threat. Installing a good-quality antivirus product on your workstations and servers would mitigate the risk of a virus infection; using covered parking at the airport would mitigate the risk of a bird-dropping attack.
- **Transfer the risk to someone else.** That's what insurance does: you pay someone to assume the risk of loss for you. You generally can't buy computer-security insurance, but you can use a variety of outsourced services that assume some degree of operational or security risk if you feel it worthwhile.
- **Accept the risk.** Some risks are either so unlikely, or so hard to avoid, mitigate, or transfer that you're stuck with them. Most of us accept some risks by default. For example, we generally don't insist on riding around in armored cars, even though that would drastically lessen the risk of bodily injury in a car crash; we might choose to mitigate the risk by choosing safer vehicles or by driving less, but ultimately most of us accept some degree of this particular risk. Once you've done everything you can to reduce, remove, or redirect the risk, you have to accept the degree of risk that's left over. You must be very careful to ensure that you have explicitly identified the risks that you're accepting as part of your messaging security environment.

Types of Security Threats

A complete discussion of all of the possible risks to your network and computers could fill several books. Some of these threats, of course, are much more likely than others—the risk that copper will suddenly stop conducting electricity, although real, is pretty remote, whereas the risk that your network will be attacked by a worm or virus is regrettably large. It's helpful to have a

system to categorize threats in several ways, including by target, type, and severity.

What Makes a Target?

Everyone knows something confidential. Likewise, every company, no matter how small, has at least some data that it would prefer to keep confidential. Some companies (particularly those in the financial services or defense manufacturing industries) have data that is well worth stealing. Other companies might find themselves targeted because of what they do, who they employ, or where they're located. However, because most attacks are initiated by worms and viruses, most victims are randomly targeted. Targets can be grouped into three general categories:

- **Opportunistic targets** are just that; they get attacked simply because they're there. Many attackers are looking for any system to attack, not a particular one. This is especially true for springboard attacks, in which an intruder compromises a machine solely to use it as a launch point for attacks on another (and probably better defended) target. Port scans and Domain Name System (DNS) zone transfers are common ways to map potential targets on a network. Follow-up probes can check for specific vulnerabilities that can be exploited.
- **Incidental targets** end up getting attacked as part of an attack on another system. For example, one variant of the CodeRed worm was programmed to attempt a distributed denial of service (DDoS) attack on the <http://www.whitehouse.gov> Web site. Machines that were compromised by this worm were incidental targets because the real purpose of the attack was to flood the Internet Protocol (IP) address of the White House Web site with traffic. (Fortunately, the designer made a simple implementation mistake that made the attack easy to prevent!)
- **Targeted systems** are attacked because of the data they contain or the role they play. Critical infrastructure systems, like emergency dispatch centers, telephone switches, or public-utility control systems, are frequently (if unsuccessfully) targeted

because disruptions to these systems cause great upset. More ominously for Microsoft Exchange administrators, messaging or storage systems at particular businesses are often targeted for penetration or denial of service (DoS) attacks. Potential attackers include current or former employees (one large financial services firm that I know of loses more than five times more each year to internal thefts than to external ones) or people seeking monetary gain, revenge, or prestige in the hacker community. For instance, *USA Today* recently suffered an attack that was undertaken because it's a well-known national publication. During the 1990s, a little-publicized series of attacks stole several millions of dollars from Citibank (although the attacker was eventually caught). Of course, e-mail systems are often targeted as part of attacks on other systems because an attacker can use the system to monitor the security staff's efforts to catch them by reading their e-mail!

You might think that no one would ever intentionally target your systems because your organization is too small to bother with, or none of your data or resources are valuable enough to attack. You might even be correct in thinking that (although, as I pointed out earlier, even small, unknown companies generally have information of value to dishonest employees or competitors). However, because most attacks are incidental or opportunistic, it's well worth taking good protective measures just in case.

Attack vs. Defense

In war, the advantage typically goes to the defense because in infantry and armor combat the defender can prepare defensive positions that play to the strengths of the defenders' equipment and terrain. Regrettably for us, the opposite principle is true of computer security: the attacker has significant advantages that we cannot always counteract. Michael Howard of Microsoft has outlined a set of four principles that neatly sum up the problem we administrators face:

- 1 The defender must defend all points of vulnerability, including workstations, servers, stored passwords, communication links, and network access

devices. The attacker can choose which point, or points, he or she attacks.

- 2 The defender can defend only against vulnerabilities he or she knows about. The attacker is free to study the systems and networks to find new vulnerabilities and exploits for them. That means that you must stay alert to new classes of attacks and new vulnerabilities as they emerge.
- 3 The defender must be constantly vigilant. The attacker can strike at will. Prime times for attacks are Sunday nights, anytime during long weekends, or major holidays like Christmas or New Year's—all times when administrators are less likely to be vigilantly watching for signs of an attack.
- 4 The defender has to play by the rules, but the attacker can fight dirty. In particular, attackers can use specialized hardware or software; they can attempt to trick employees into giving them passwords, network addresses, or other useful bits of information, and they can gang up on a target.

As you read the material on classifying threats and on applying the two threat models covered in this appendix to your own work, remember these principles—forgetting them can cost you dearly!

Classifying Threats

In his famous speech in the Book of Mormon, King Benjamin says, "I cannot tell you all the things whereby ye may commit sin; for there are divers ways and means, even so many that I cannot number them." (See <http://scriptures.lds.org/mosiah/4/29>.) So it is with security threats: clever attackers are continually finding new vulnerabilities in software, systems, and communications protocols, so it's very difficult to come up with a comprehensive list of potential attacks that will remain useful over time. Rather than a checklist of attack methods, it's more useful to classify threats into general categories, with a few specific examples of each:

- **DoS attacks** are designed to keep legitimate users from using a resource. If someone blocks my car into a parking space, that's an effective DoS attack because I can't move my car until the obstructing

vehicle moves. One common network-based DoS is flooding, in which a target system's Internet connection is overwhelmed with meaningless traffic. It is not uncommon to see DoS attacks that involve sending specially malformed data to a service on the target machine. That data exploits programming flaws that cause the service to crash or hang, or to consume all of the CPU resources or RAM on the target server.

- **DDoS attacks** are a pernicious variant of ordinary DoS attacks. Imagine if a crew of miscreants simultaneously used every pay phone in New York City to dial 911—the 911 dispatch center would quickly be overwhelmed, and legitimate calls couldn't get through. That's the evil genius behind DDoS attacks; they leverage many compromised machines that focus their efforts on a single target. Participating machines are typically compromised either by a worm or a Trojan; once compromised, they can attack the target on a coordinated schedule or when directed to by the original attacker.
- **Penetration attacks** involve gaining surreptitious access to a network. For example, an increasingly common penetration tactic is using software like NetStumbler (<http://www.netstumbler.org>) to find poorly protected wireless local area networks (WLANs), and then to attack them. Penetration attacks are usually a prerequisite to other types of attacks; sometimes an attacker's only goal is to penetrate a particular network so that it can be used as a jumping-off point for attacking a different network. Most penetration attempts are never reported to law enforcement, and I would venture to say that until recently, the majority of attempts went completely undetected by the target—not exactly a comforting thought.
- **Spoofing attacks** are those in which some kind of data is falsified. If you ever get spam, you've probably seen at least one kind of spoof, in which the e-mail headers for the sending domain are falsified. Other spoofs include the injection of fake DNS records or rogue DNS or Dynamic Host Configuration Protocol (DHCP) servers into a network,

as well as more obvious attacks like falsifying or modifying data in databases, file shares, or messages.

- **Escalation of privilege attacks** are quite serious. Thanks to the access control mechanisms that Windows implements, an ordinary user doesn't have privileges to do really destructive or dangerous things. Administrators, however, do. Privilege escalation attacks depend on flaws in the operating system that let an ordinary user gain administrative privileges. These flaws can be exploited by unscrupulous users, attackers who can gain physical access to the machine, or attackers who trick legitimate users into running Trojans.
- **Information disclosure attacks** attempt to steal useful or interesting information. They range from the exotic, like using high-gain wireless antennae to sniff 802.11 signals in the parking lot, to the mundane, like rifling through the company dumpster looking for incriminating documents. This kind of thing sometimes happens to security companies, too; try doing a Web search for "Mykotronx dumpsters" and see what you find! Most commonly, these attacks are accomplished by using privilege escalation or penetration attacks to get the attacker into the system where the target data are stored.
- **Information compromise attacks** aim to carry out the covert modification, substitution, or creation of data. As with disclosure attacks, these attacks are usually undertaken after a successful penetration or privilege escalation attack gives the malefactor access to the needed systems. These are similar to spoofing attacks, but the distinction between them is that spoofing attacks concentrate on falsifying identities or services (for example, the address of a legitimate DNS server for a domain), whereas compromise attacks target data stored on a system (for example, the value of an oil and gas drilling lease or the amount of revenue a company has booked in the year to date).
- **Virus or worm attacks** are usually opportunistic, but they can be quite damaging. These attacks could

lead to other types of attacks; I already mentioned the CodeRed DDoS payload, and it is not uncommon for worms or viruses to carry Trojan payloads that allow remote compromise and exploitation of an infected system.

It's important to note that for some of these attacks, there's no practical distinction between network-borne attacks and those that arrive through other means. Of course, penetration, DoS, and DDoS attacks are dependent on network connectivity, but the other types discussed here are just as feasible from a local workstation as they are from some far corner of the Internet.

Models for Risk Assessment

There are a number of models for identifying and quantifying information systems risks. Most of these models require a fair amount of specialized training to be useful because performing a strict risk assessment involves a number of fine points that are well beyond the scope of this book. However, before you rush off to hire a Certified Information Systems Security Professional (CISSP) to do your risk assessment (not that doing so is a bad idea by any means, as long as you hire one with a background in risk assessment), wouldn't it be helpful if you could evaluate your risks yourself? You can do so to a good degree using the two risk assessment models I present in this section; the models help you begin to identify and quantify various threats and risks to your information systems well enough to start fixing the most serious ones. In fact, I suggest you apply the two models in combination.

The STAVE Model

This model doesn't really have a formal name; I made up STAVE because it's pronounceable. The key elements of risk assessment described in the CISSP curriculum and in Krause and Tipton's *Information Security Management Handbook*, however, revolve around the five elements in STAVE, so it's worth presenting them here to give you a stronger conceptual framework. The five elements of STAVE are simple to understand: safeguards, threats, assets, vulnerabilities, and exploits. However, for maximum understanding, let's talk about the STAVE elements in a slightly different order,

beginning with the assets and working our way through the elements that indicate what risks those assets face and how we can fix them.

Assets

Assets are something valuable that you have; they can be tangible or intangible. If you don't have any valuable assets, it's probable that no one will attack you on purpose. Of course, being asset-poor doesn't mean that you won't be attacked, merely that you won't be an intentional target. The more numerous, valuable, or irreplaceable your assets are, the higher the likelihood that they'll be attacked. Notice that in this context, asset doesn't just mean a physical object or a juicy piece of information; some of the most valuable assets of the sites of organizations such as CNN and the *New York Times* are their perceived trustworthiness and integrity.

Assets have a value associated with them. In the case of a physical asset like a building, a stack of gold ingots, or a fighter plane, the value is pretty easy to calculate. For an intangible asset, like the value of a complete database of your company's customers over the last 15 years, the value might be much more difficult to pin down. Having said that, getting relatively accurate asset values will help you clearly identify where the biggest potential risks are. A small risk of losing a highly valuable asset might be more important than a larger risk to a less valuable asset.

Threats

A threat is something bad that can happen. The exact set of threats you should worry about varies from asset to asset. For example, one of my clients is a large law firm located in a downtown area that occasionally floods. Because the company is located on the 37th floor of the building, the primary concern isn't the physical computer assets; it's the value of the company's data and of its reputation as a trustworthy guardian of the legal records it maintains.

Along with identifying the threats themselves, you need to be able to prioritize them in some way. This could be done by severity (for example, if this threat occurs, how bad will the effects be?), by likelihood, by frequency, or by some other criterion that's specific to your business.

Brainstorming with a list of assets is a great way to develop a prioritized threat list. Draw all of your assets on a chart and then start listing threats to each of them. There might be more threats out there than you realize!

Vulnerabilities

A vulnerability is something that allows a threat to apply to an asset. In other words, a vulnerability is a weak spot that, if not mitigated, allows an attacker to use a specific threat to damage or gain control of a particular asset. Vulnerabilities can be anything that an attacker can exploit: unlocked doors, unpatched workstations, users who keep passwords in plain view, and flaws in installed software are all-too-common vulnerabilities.

Identifying vulnerabilities can be tricky. Some are obvious (like the notes with passwords on them), but some are much more subtle. In fact, attackers can, and do, expend large amounts of time and effort finding new vulnerabilities and using them before you, or the vendors who make the products you use, can find them. For that reason, you cannot always count on being able to eliminate vulnerabilities; in some cases, the best you can do is mitigate the ones you know about and try to proactively protect yourself against known classes of vulnerabilities. This logic gave birth to the modern antivirus software field.

Hold It Right There!

A brief pause in our discussion of the STAVE model is necessary because there's a simple but subtle point to make here: eliminating assets, threats, or vulnerabilities reduces or removes any particular risk. Let's say that you do such a good job of securing your Exchange systems that you get a fat bonus, which you use to buy one of those fancy plasma-screen TVs. You install it in your living room in such a way that it can be seen from the street. You are in the habit of leaving your front door open, with only the glass storm door keeping intruders out. Let's analyze your risk based on what we know:

- The asset is your spiffy new TV, valued at about \$7,000. It's even more valuable to you because your spouse is very unlikely to let you ever buy another one, so you want to hang on to it.
- The threat is having someone steal the TV by entering through your open front door or by breaking the cheap lock on the door while you're not at home.
- There are several vulnerabilities: your habit of leaving the door open, the open door itself, the flimsy locks on the front and back doors, and the lack of an alarm system. The reason I cite the door-open habit separately is simple: without proper attention to processes and education, the best safeguards in the world will be ineffective because administrators and users will fall back on their old, insecure habits.

You can mitigate this risk by doing one of three things. First, you could fasten the TV to the wall, making it more difficult to steal even if a thief manages to break into your house. Second, you could move to a town with a lower occurrence of theft. Finally, you could address the vulnerabilities by changing your security procedures, remembering to close the door, closing your curtains, and upgrading your locks. Doing any of these things greatly reduces the risk to your asset; doing several of them virtually eliminates it. Notice that not all these proposed mitigating measures are really practical, though—clearly, addressing the vulnerabilities is the best place to start.

What does your fancy TV set have to do with information systems security? Using risk assessment to drive security choices acknowledges the compromise between risk and cost. The cost and trouble associated with moving is far too great just to reduce the risk of someone stealing your TV. However, it probably is worthwhile to close your windows and upgrade your locks. You cannot wave a magic wand and make all potential threats disappear, although you can (and should) work to minimize any threats over which you have control or influence. The biggest win for administrators is to clamp down on vulnerabilities.

Exploits

A vulnerability by itself isn't particularly interesting. An exploit, alas, is a different story; it's a piece of code or behavior that takes advantage of a particular vulnerability. The difference between an exploit and a vulnerability is slight but significant. If you forget to

close the telnet port on your firewall, that's a vulnerability. If someone uses it to hack your server, that's an exploit.

Unless you're the one doing the hacking, you probably won't have any control over exploits aimed at your machines. However, every time you fix or remove a vulnerability, you're rendering useless all the exploits that use that particular security gap.

Safeguards

Safeguards are just what their name implies: they are procedures, devices, or programs designed to safeguard assets against threats and exploits. Some safeguards are preventative, whereas others are designed to limit the potential damage from a known or suspected vulnerability. Safeguards are all around in the systems we use today. Banks use safeguards like armed guards, surveillance cameras, big steel vaults, and serial number tracking. Computer systems use safeguards like strong password policies, smart cards, and security auditing logs. A good risk assessment clearly identifies what safeguards are currently in place and what they safeguard against. A better one also points out new or modified safeguards that can help reduce the danger from the risks identified in the assessment.

The STRIDE Model

Microsoft uses a different, more specific model to guide their internal security processes, including design and security reviews. The STAVE model is a good framework for general security concepts; Microsoft's model, called STRIDE, is normally used by developers and designers to identify and resolve security issues in their application code. STRIDE is useful for us too because we can use it to evaluate the potential risks to a messaging system with only slight modification. The six letters in STRIDE each represent a particular risk. Those risks, and their effects on Exchange, are as follows:

- **Spoofing user identity** I've already mentioned spoofing, but the STRIDE model talks about it in the specific context of an attacker who can impersonate another user. Spoofing rears its head in two ways within Exchange. One is usually legitimate; Exchange allows users to delegate access to their mailbox so that another user can send mail on the mailbox owner's behalf. The other, Simple Mail Transfer Protocol (SMTP) spoofing, is difficult to guard against because SMTP wasn't (and still isn't) designed to offer any significant degree of security.
- **Tampering with data** An attacker who maliciously changes data is often much harder to detect, and does much more damage, than a smash-and-grab Web site defacer or disk reformatter. Why? First, you might not find the modified data until some time has passed; once you find one tampered item, you'll have to thoroughly check all the other data on your systems to ensure that nothing else was tampered with. Modifying data in Exchange requires the correct privileges, which means that this particular threat is usually coupled to privilege-escalation attacks.
- **Repudiation** The R in STRIDE stands for repudiation, and it represents the risk that a legitimate transaction will be disowned by one of the participants. This isn't a direct threat to Exchange messaging systems; however, if the systems are used for business transactions, the risk that a transaction will be repudiated exists. Without digital signatures, it's trivial to forge transactions, and most people understand this. Unfortunately, that widespread understanding means that the unscrupulous have a ready-made claim for repudiation: "I never sent that e-mail! Someone must have forged it."
- **Information disclosure** In the STRIDE model, information disclosure means that an attacker can gain access, without permission, to data that the owner doesn't want him or her to have. This is a very broad definition; one of your jobs is to refine it by specifying which kinds of information disclosure worry you. For example, most sites aren't worried about the fact that the Exchange SMTP server identifies itself as such, but some are. On the other hand, almost no company is willing to allow individual users or administrators to have unfettered access to each others' mailboxes.

- **DoS** DoS attacks make it impossible to use a resource. They're tricky to defend against because they involve the overuse of legitimate resources. You can stop all such attacks by removing the resource used by the attacker, but then real users can't use the resource either. In addition to common DoS attacks against the network or Windows components, Exchange is vulnerable to DoS attacks that attempt to consume all of the disk space on the drives where the information store databases are located.
- **Escalation of privilege** Exchange itself depends on the underlying Windows authentication system. You might think that this makes Exchange less vulnerable to privilege escalation attacks; the truth is that this dependency doesn't make it any less vulnerable. An attacker who can successfully escalate his or her privileges might be able to use those privileges to gain elevated access to Exchange, depending on how you've assigned Exchange permissions. The STRIDE model is quite useful as a way to help build a taxonomy of threats. As you list the risks to your Exchange systems (as we'll do in the next section), you can pigeonhole them into one of the six STRIDE categories, giving you a convenient roadmap of the threats and risks you're most likely to face.

Asset and Threat Assessment for Exchange (or, What Would You Not Like to Lose Today?)

Part of risk assessment is identifying the assets you have to lose. Even a quick, back-of-the-envelope inventory is better than nothing, but the more time and effort you put into your inventory the more useful it will be to you. Of course, you have to put a commensurate amount of effort into identifying the threats you face, too, so that you can adequately assess the risks to your assets. The particulars will vary according to your operations, but the overall principle is the same.

Asset Inventory

Make a comprehensive list of your assets, informational or otherwise. Some will be obvious, like the server and

network hardware that hosts your Exchange infrastructure, or the stored message data in your mailbox and public folder stores. Some might be less obvious: Have you considered the value of data that is on your backup tapes? What about the information value of message headers? Here are some specific questions to ask to guide you in this process:

- Which assets, if lost, would result in a cessation of normal business activities? How long would that cessation last, and how could you recover from it? For Exchange, this usually includes stored mailbox data, but it can also include documents or records stored using Exchange's Web Storage System. It might also include permissions or security data.
- For any particular asset, what would happen if you lost access to it for 15 minutes? For an hour? For a workday? For a work week? For an entire month? Your answers to these questions must be specific, and should include objective cost or loss figures. Knowing exactly how much it costs to be without e-mail for a day is a wonderful way to figure out which availability and security measures make sense.
- Which assets could potentially be compromised without you knowing it? The best example is probably communication links, because they can generally be monitored without either endpoint becoming aware of the monitoring. Other examples include wireless local area networks (WLANs) and traffic to and from cell phones or other devices that can wirelessly send and receive e-mail.
- If they were compromised or leaked, which assets could provide useful, nonpublic information to competitors? Which of those assets can be restricted so that their competitive value is eliminated or reduced? Any message-related data that leaves the company falls into this category. That includes messages that have to transit networks (including the Internet) that you don't control.
- Which assets, if disclosed, would make attacks on other assets possible? User credentials and passwords are one obvious example; so are unprotected

backup tapes, DNS data, and other seemingly mundane materials like organizational charts. (There's a reason why the old Soviet Union classified their phone books!)

- Which assets have intrinsic value? These are the assets of greatest interest to an attacker. Companies that deal with high-value transactions like oil and gas leases, commodities futures trades, and the like obviously have lots of these, but so do companies that make expensive manufactured goods or components.
- Which assets are you legally or contractually required to maintain? What happens if you fail to maintain them? Health care companies face stiff penalties if data about their patients is compromised, and many other industries have similar regulations. If compromised, what other assets might make it impossible to meet those requirements?
- Which assets lose value if their integrity is compromised? If a newspaper prints a false story, the costs go far beyond the cost of the paper and the ink. If a stock exchange occasionally forgets trades, the reputation of the stock exchange is permanently damaged, costing far more than the value of the transactions. Damages to some assets can be more complex than total destruction.

Once you've identified the assets and the potential loss associated with each of them, you're ready to start asking some harder questions. For each asset, ask these questions:

- What threats exist to that particular asset? (Do you even know which threats might exist? If not, this is a terrific time to find out!) Are there threats that might cover multiple assets, or even classes of assets?
- What vulnerabilities enable the threats you just identified? Be sure to include vulnerabilities caused by poor security processes or lack of user and administrator education; it's tempting to blame every vulnerability on the software vendors, but that misses some of the biggest, juiciest weak spots.
- How frequently can you expect these vulnerabilities to be exploited? This is a hard judgment call to make because the frequency of exploit will depend on a number of variable, hard-to-determine factors: How public is your site? Is there any special reason people might target you? Do you have particularly valuable, sensitive, or controversial data?
- What safeguards can you apply to mitigate the threat; avoid, reduce, or transfer the risk; or block the vulnerability? Obviously, closing off vulnerabilities is the most effective safeguard in most cases, but you might not always be able to address every vulnerability in that manner.

Summary

Risk and loss assessment are complicated, and there are many subtleties that only a certified disaster planner or CISSP can help you fully explore. However, there's a lot you can do to help yourself: inventory your most valuable assets, identify threats that might cause loss to those assets, and estimate the likelihood that a particular threat will cause loss or damage. Taking these three steps will give you a much better idea of what's actually at risk in your messaging system.

Additional Reading

- The Association for Computing Machinery (ACM) maintains an extremely educational discussion of computer-related risks to the public, the RISKS Digest. The digest is available on Usenet (look in comp.risks) or on the Web; I normally read it at <http://catless.ncl.ac.uk/Risks>.
- Microsoft Press has two good programming security books. Michael Howard and David LeBlanc's *Writing Secure Code* (2001) and Howard's *Designing Secure Web-Based Applications for Microsoft Windows 2000* (2000) describe how programmers and designers can apply the STRIDE model to make their products and services more secure. They're good general-security references, although they don't discuss Exchange in any detail.

- Microsoft has a security Web site (<http://www.microsoft.com/security>) that does a good job of posting information on newly discovered vulnerabilities for Windows systems.
- Microsoft Security Operations Guide for Windows 2000 Server and the companion guide for Exchange 2000 make terrific detailed guides of security practices and settings. These guides cover both policies and practices, and Microsoft has made them freely downloadable from <http://msdn.microsoft.com/practices>. Get them and read them.

About the Author

Mitch Tulloch, MCSE, Cert. Ed., is a trainer, consultant, and author of the award-winning *Microsoft Encyclopedia of Networking* (Microsoft Press). Mitch has written numerous books about administering Microsoft platforms and products, including *Windows Server 2003 in a Nutshell*, *Windows 2000 Administration in a Nutshell*, *Microsoft Exchange Server in a Nutshell*, *Administering Exchange 2000 Server*, *II6 Administration*, *Administering IIS 5*, and *Administering IIS 4*. Mitch has also written feature articles for industry magazines such as *NetworkWorld* and is a contributor to *myITforum.com* and other online communities for IT professionals.



The manuscript for this book was prepared using Microsoft Word 2000 and submitted to Microsoft Press in electronic form. The contents was structured using XML. The majority of the pages were composed using Adobe FrameMaker+SGML 6.0, with text type in Times and display type in ITC Franklin Gothic. Composed pages were sent to the printer as electronic prepress files. The CD was created using XML output from FrameMaker.

Interior Graphic Design: James D. Kramer
Production Information Analyst: Barbara Norfleet

For nSight, Inc. (www.nSightWorks.com)

Project Manager: Susan H. McClung
Copy Editor: Christina Palaia
Technical Editor: Thomas Keegan
Proofreaders: Virginia Carroll, Katie O'Connell
Desktop Production Specialist: Mary Beth McDaniel
Indexer: Jack Lewis

