November 2004

Enterprise Internal Network Security in a Box

IBM @server BladeCenter with Mazu Networks

Executive Summary

What has the power to take down critical services and compromise sensitive data?

Worms and insider threats attack critical internal resources, destroy information and compromise or damage sensitive financial records and data. Not only are the number of attacks on the rise, but organizations are sometimes more vulnerable than ever. There has been growing access to networks for remote employees and third parties, wider usage of mobile devices and a rise in Web services. Broadening access to resources within corporate networks has greatly improved the efficiency of many business processes, but also increased organizations' vulnerability to security threats.

It is not simply the growing number of attacks that are of concern, but their growing destructive power, and the shrinking time lag between when vulnerabilities are announced and are exploited. A study of network worms launched over the last 24 months shows the time lag shrinking from 330 days for the Nimda Worm in 2002, to 16 days for Sasser in April of 2004. According to Gartner research, this pattern will only get worse, with a projected 25% increase in the so-called Day 15 attacks over the next several years. (Stay ahead of Software Vulnerabilities – 26-April 2004 – J Pescatore)

Insider threats, or attacks launched by internal credentialed users, are causing even more damage. During the last year, each instance of insider attacks has cost companies more than $500,000 in damage. Credentialed users are not simply employees, but outsourced workers, remote users and even business partners. These attacks are also on the rise. Gartner suggests that over 70% of future attacks will be from insiders.(Stay ahead of Software Vulnerabilities – 26-April 2004 – J Pescatore)

Current strategies for protecting the internal network against these threats are largely focused on perimeter security and patch management. While perimeter security appliances such as intrusion detection systems (IDS) and intrusion prevention systems (IPS) have helped with known threats, they have largely been ineffective against new attacks and threats brought into the network on mobile devices. And while patch management has similarly helped alleviate vulnerabilities to many known threats, it simply cannot keep up with the increasing speed of new exploits. The percentage of organizations reporting compromises is on the rise.

Mazu Profiler™ is the first Behavioral IPS solution to utilize the industry-leading scalability, resilience and integration capabilities of IBM @server® BladeCenter™. As one of its core competencies, Mazu exploited the BladeCenter architecture so that Mazu Profiler delivers the following:

- Scales to meet the needs of the largest, most complex networks
- Collects data from across the network
- Centrally models and analyzes data
- Distributes processing of network information across multiple server blades
- Provides a unified view of networks of hundreds of thousands of hosts in real time

These super dense and simple to install servers house the computing power, I/O flexibility, and memory needed to drive application performance for Mazu Profiler. With BladeCenter's vast level of flexibility, it is rapidly becoming the platform of choice for applications demanding fast scale out infrastructure. The BladeCenter form factor can significantly increase performance density over traditional server form factors, while helping lower administration, power, cooling and installation costs. These cost-savings are achievable because BladeCenter's integrated architecture and its consolidated management solution allow administration of the solution from virtually any location.

Worms

A worm is a self-replicating program that scans for vulnerable hosts on a network and then infects those hosts, which in turn starts the process all over again. Unlike a typical computer virus that requires some human interaction, such as opening an attachment or clicking on a link, worms do not need to attach themselves to a host program, and they require no user interaction. As a result, the speed with which worms propagate across a network is unlike virtually any other form of attack. This makes them very difficult to contain, and even without a malicious payload (the action it takes other than to infect other hosts) worms can disrupt critical services in a network by consuming bandwidth and overloading systems. In an instant your business grinds to a halt.

Worms are made of two main components; a spread algorithm, the method by which it locates other vulnerable hosts, and a payload, the action it takes other than spreading itself.

The Spread Algorithm

Different worms have different spread algorithms, but at a high level one can observe a chain of host scans and port scans. The infected host performs host scans in order to locate other hosts in the network. It then performs some form of port scan in order to identify vulnerabilities. The infected host then infects the vulnerable host by transferring its code, after which the newly infected hosts begin the process all over again. Charting out a worm's progress in a network often looks like a tree diagram, fanning out from a point of entry in the network.

Different worms use different spread algorithms, enabling some to propagate more quickly, elude defenses or even fly under the radar. But often the spread of the worm itself can be devastating. "Worm storms" or the resulting bandwidth surge caused by the escalation in scanning activity can create denial of service (DoS) conditions—grinding the network to a halt.

Spread algorithms are becoming increasingly sophisticated, exploiting methods that enable them to spread quicker and elude modern defenses. Polymorphic worms dynamically change characteristics as they spread, such as changing the ports they propagate over. This makes it very difficult to track and respond.

Another concern is called "Stealth Worms". These are worms that propagate slowly and fly under the radar. It's a common misconception that the faster the worm propagates the more danger it poses. To be certain, a fast propagation vector creates DoS conditions, but it's immediately clear to everyone in the network that something is amiss. By contrast, a worm that propagates very slowly could easily go unnoticed until it's too late.

The Payload

The payload refers to the action or actions the worm takes other than spreading itself. In the worms we have seen to date, payloads have:

- Created backdoors, enabling hackers to take control of the computer at a later time
- Disrupted or defaced Web sites
- Created system zombies, later used to launch distributed denial of service (DDoS) attacks against other sites

To be clear, most of the damage caused by major worms we've seen to date has come from their spread rather than their payload, but this will change. Future worms could contain payloads that cause extensive corruption of data, theft of sensitive data including personal information, intellectual property and critical business data, and launch large scale simultaneous DDoS attacks.

In an era of increasing dependence on the network for key business processes, increasing need to share data and make applications and information more accessible to more people, and increasing regulatory pressures, enterprises and government agencies need to take the worm threat seriously. They are not just nuisances. Worms can be devastating to the organization and the community at large.

However, solving the problem is not easily accomplished with existing solutions and approaches. Most organizations take a two-pronged approach to the problem:

1. **Patching:** Patching systems quicker makes them less vulnerable.
2. **Perimeter Security**: Implement Firewalls, IDS and IPS to limit a worm's ability to enter the network.

Patch Management

Worms need vulnerabilities to exploit. Organizations are deploying increasingly sophisticated patch management processes, but the rate at which new vulnerabilities are being found is astounding. According to CERT, between 70 and 80 new vulnerabilities are being announced every week. That number has been steadily growing over the last five years.

Companies do not simply update all systems automatically as soon as patches are available. A new patch is just as likely to create new problems because of bugs or incompatibilities with existing systems. And so the patch cycle begins: triage, prioritize, test and roll out. A recent study of medium and large enterprises by the Yankee Group showed a mean cycle time from patch announcements to rollout of 60 days.

By contrast, the time it takes worm and virus writers to exploit vulnerabilities is steadily shrinking from almost a year for Nimda to a scant 16 days for Sasser. According to a study by Gartner Research, "Cyberattacks that exploit vulnerabilities where a patch has been available for less than 30 days will increase from 15% of total cyberattacks in 2003 to 30% in 2006."

Through new technologies and processes, companies will continue to improve the patch management process. This will serve to narrow the threat window and make systems less vulnerable to known attacks, but in the foreseeable future this process will never catch up to the hackers creating new exploits. The patch process is not the solution for new attacks.

Perimeter Security: Firewall, IDS and IPS

The second common layer of defense is firewalls, IDS and IPS products.

The most common of these is the firewall. Firewalls help keep the doors to the kingdom from being simply left wide open. The challenge for firewalls is that accessibility is as important as security. Remote employees, contractors, suppliers and other trusted third parties need access to the network. Similarly, hosts inside the network need access to the outside. The problem is an increasing number of exploits bury themselves in common ports and protocols and essentially walk right through the firewall.

IDSs are devices that generate alerts when known exploits pass across their sensors. These products suffer from several problems. First, for these products to catch worms, the signature of the worm must be in their databases. They are therefore incapable of catching any new worm. Second, IDSs are prone to false and duplicate alerts that are often not watched very carefully. Third, they are passive devices, so the worm still infects the network, and the IDS provides no ability to contain and clean up the infection quickly.

Signature-based Intrusion Prevention Systems *(IPSs)* are more promising. IPS products are designed to not only detect attacks, but block malicious activity before damage occurs. As such, IPS products are delivering much higher value than their IDS predecessors. Signature-based IPSs use databases of signatures as well as protocol and application anomaly techniques. They provide greater value than the IDSs in part because they can actually filter the attack if it is detected. There is still concern about accuracy issues, but by and large for known attacks, there is significant value in filtering malicious traffic at the perimeter, regardless of whether you are vulnerable to the exploit.

Unfortunately, signature-based IPSs provide little value for new attacks. In addition, IPSs focus on mitigating attacks and do not provide the same alerting, visibility and analysis that traditional IDSs provide. This reduces the administrator's ability to understand what's going on in the network.

Finally, there is an additional flaw in all these methods. Mobile devices such as laptops move in and out of the network every day. They circumvent perimeter defenses altogether, and the internal network is a big place to watch. Simply deploying IPSs on every link will not solve the problem. On the contrary, it would be enormously expensive, impossible to manage, miss all new attacks and put companies at serious risk of accidentally disrupting critical services.

What's Missing

To truly solve the worm problem enterprises need a solution that is capable of:

- Catching new and even zero-day attacks. It cannot rely on signatures.
- Catching attacks that circumvent the perimeter on mobile devices. It needs to watch over an entire network—not just a link.
- Surgically mitigating attacks: It's not enough to see the attack; it needs to help teams mitigate the attack. But this is particularly challenging both because the attack can be anywhere in the network and because a great deal of internal traffic is highly critical. A solution that enables teams to block internal traffic must be highly accurate in detection, surgically precise and provide visibility into the impact of the mitigation prior to implementing it.

In addition, we need something that will help harden the internal network beyond simply patching systems. This would include:

- Optimizing segmentation and access policies
- Eliminating rogue services such as e-mail servers, Web servers and wireless access points
- Enforcing usage policies

Insider threats

Insider threats, while garnering a lot less attention from the media, can be equally devastating. Insider threat refers to both potential and actual attacks by credentialed users or "insiders". These attacks can result in theft or contamination of sensitive data, or disruption of highly critical services. These incidents are far more common than the media coverage would suggest, in large part because many of these attacks are not publicly disclosed. Gartner Research suggests that the internal trusted zone is where the motivated and potentially knowledgeable hacker resides. They also suggest that inside of companies' networks is where 70% of their future threats are going to come.

According to the findings of the Secret Service and the CERT Coordination Center in a study of insider attacks against financial organizations, it doesn't take a person with a lot of technical knowledge to steal or wreak havoc on an IT system from the inside, and these inside attackers do not fit any common profile. There are no signatures that can be created to stop insider attacks.

Several trends seem to be at the heart of this increased threat. The first is the growth in credentialed users. In the past, these threats were characterized by employees or former employees of the company. Today, however, with companies increasingly giving access to corporate resources to contractors, remote employees and business partners, there are often significantly more credentialed users than employees.

The second trend is the move to wider access and Web services. As applications move from mainframes to client servers to Web-based services, it has become increasingly easier and more cost-effective to give a wider group of users access to applications and data. This has several advantages, not the least of which is far greater efficiencies in key business processes. It also creates several security challenges.

A Web browser is a highly effective hacking tool. This creates a very big security challenge. Traditional solutions such as firewalls, signature IDS and signature IPS are infective in dealing with this problem. To begin with, these devices are typically deployed at the perimeter. Furthermore, the internal network has far too many links to watch, making it exceedingly expensive with these link-based solutions. Lastly, threats such as unauthorized access or sabotage of custom applications largely defy signatures.

To solve the insider threat problem, many of the same attributes are required as described in the worm solution outlined above. Enterprises need a solution that is capable of:

- Catching unauthorized access and disruption to critical services. It cannot rely on signatures.
- Catching attacks launched from anywhere in the internal network. It needs to watch over an entire network—not just a link.
- Surgically mitigating attacks: It is not enough to see the attack; it needs to help teams mitigate the attack. Once again, because of the sensitive nature of internal services, the solution must be highly accurate in detection and characterization of the attack, surgically precise and provide visibility into the impact of the mitigation prior to implementing it.

Hardening the internal network is once again essential in preventing these attacks. As with the worm solution, this requires a solution capable of:

- Optimizing segmentation and access policies
- Eliminating rogue services such e-mail servers, Web servers and wireless access points
- Enforcing usage policies

Behavioral Intrusion Prevention System (Behavioral IPS)

Behavioral IPS is a new breed of intrusion prevention solution. It not only offers a more accurate and comprehensive means of detecting threats, but uniquely provides accurate detection of new and zero day threats, provides the ability to characterize or model threats, and provides visibility into the impact of changes in policy. Unlike traditional IPS solutions, behavioral IPSs detect threats by observing how systems on your network are behaving. For example, hosts infected with a worm typically begin performing host scans and then port scans, and systems they connect to begin similar behavior. Not only are hosts behaving differently than they normally do, but the chain of scanning activity represents a “pattern” of behavior indicative of a worm.

By investigating anomalous activity, analyzing patterns of behavior and correlating anomalous behavior with other activity on the network, behavioral IPSs are able to accurately detect insider threats and attacks, and other dangerous activities such as misuse of network services or assets. For example, a host connecting to a server for the first time is not necessarily a case of unauthorized access, but if no other hosts in the group ever connect to the server (e.g. marketing desktops connecting to a finance server), the activity is a little more suspicious. Furthermore, if

the port being used is not typically used by that server, this arouses further suspicion. By correlating enough information, the system is able to determine if activity represents an operationally relevant threat.

Equally important, behavioral IPSs provide enormous visibility into how the network and assets in the network are actually being used. This helps companies become more proactive about security. Hardening the internal network becomes more efficient and effective with a clear understanding of how the network is being used and by whom.

Mazu Profiler

Mazu Profiler is the leading Behavioral IPS for protecting internal networks against worms and insider threats. Mazu Profiler allows companies to detect and surgically mitigate attacks, harden the internal network against future attacks, and monitor and audit sensitive assets and services.

Surgical Mitigation System: Detect and surgically mitigate worms

- Detect – Detect and characterize new attacks
- Mitigate – Provide detailed impact analysis and recommendations and perform precise mitigation
- Remediate – Identify and prioritize compromised hosts and their dependencies

Detect

Without the correct information and the proper research, companies are faced with the risk of losing critical services on their networks. The challenge is that during an attack, every second counts. These decisions must not only be informed decisions, but made very quickly.

The Surgical Mitigation System is the first technology that provides IT departments and administrators with detailed and real-time information about their networks as well as any and all worms and threats that are inside the network.

Mitigate

Provide detailed impact analysis and recommendation

Surgical Mitigation System is the first technology to provide organizations with instant, detailed impact analysis regarding the effect on business activities of mitigating worms and insider threats. With traditional solutions, the cure is often worse than the disease. Stepping on traffic without knowing the consequences is a much more significant problem in the core of the network than at the perimeter. Mitigating techniques can cost companies if the impact of the mitigation is not completely understood. Surgical Mitigation System allows companies to instantly and completely understand the effect of mitigating threats while providing detailed mitigation strategies and recommendations.

Recently, a large financial institution was being attacked by the Sasser worm. The worm began to propagate over a certain port not used by standard applications. A logical reaction would be to have administrators shut down that port across the network. Fortunately, impact analysis showed that the port in question was in heavy use by a custom application in the call center. The action would have shut down the call center and potentially cost the company millions. A more surgical response was chosen. The key is not simply being precise, but knowing exactly where to cut. The goal, after all, is business continuity. If critical services are taken down, it is not relevant if the cause was the worm or the actions taken to eliminate the worm. In some cases, millions of dollars can hang in the balance of these decisions.

Perform precise mitigation

In the previous example, the Surgical Mitigation System recommended and performed mitigation on the host-to-host connections and avoided shutting down the port entirely. This precision allowed the company to continue using the port and the call center application while mitigating the worm.

Other mitigation technologies and solutions cannot perform and provide detailed impact analysis based on real-time network information and data. Surgical Mitigation System is perhaps the first technology capable of surgically mitigating threats without disturbing or disrupting business operations.

Remediate

Profiler's recovery reports provide security and network teams with up to the minute reports of all compromised hosts, prioritized by group, and a list of all their dependencies. This enables a focused, efficient and effective recovery process.

Network Immunization: Harden the internal network against future attacks

- Eliminate – Rogue services such as rogue e-mail and Web servers or wireless access points
- Enforce – Monitor and enforce policies across the network
- Optimize – Access policies for firewalls and routers

Detailed and dynamic information about actual network usage helps security teams be more proactive and take steps that prevent damaging attacks. Mazu Profiler shows which individuals and groups within an organization use exactly which network resources. This gives security teams an efficient way to optimize access policies for firewalls and other systems. As a result, enterprises can avoid problems associated with access policies that are either too tight or too loose. In addition, visibility into the internal traffic provides a highly efficient way to monitor the network for usage policy violations without using agent-based technologies.

Also, detailed information on network usage can be used to identify the most critical applications and services in a network environment. This feature often identifies network resources that are underrated in terms of their importance to key business processes. By highlighting their importance, security teams can prioritize their protection, and avoid disruption of critical business activities.

Monitor and Audit Sensitive Assets and Services

Because Mazu Profiler maintains baselines of how resources are used, it provides an enormously effective means of monitoring and auditing the usage of financially sensitive data and other sensitive assets in the network. This is critically important for a range of regulatory standards including Sarbanes Oxley and HIPPA.

How it works

Mazu Profiler collects network traffic statistics from multiple Mazu Sensors and/or NetFlow sources (such as routers). Mazu Sensors can tap links or collect mirror traffic from switches to monitor network traffic.

Profiler collects this information from sources deployed across the entire network and de-duplicates the data to create an accurate, centralized view of all network activity. For the first week, this data is used to create a baseline of the network, which includes the following information:

- Which hosts communicate with each other
- What services do they consume or provide on the network
- How much traffic (connections, bytes, packets) is sent or received for each of those services

The system can be configured to have multiple baselines for different periods of the week. For example, companies can configure Profiler to create one baseline for weekdays, one for weeknights,and one for weekends. Companies can even create baselines for special periods like end of quarter or market open. After the baselines are generated, Profiler's heuristics are activated to compare current network traffic to the relevant baseline to look for malicious behaviors and immediately alert on the following:

- Worms (rapid or stealthy)
- Unauthorized Access
- Host Scans (rapid or stealthy)
- Port Scans (rapid or stealthy)
- DDoS
- New Service or Host
- Silent Host

Users can also specify policies or rules, and Profiler will alert when it sees traffic that violates those policies. For example, many companies have a policy that desktop systems should not run Web servers because they are not managed by the IT group that will keep the systems patched, protecting them from vulnerabilities. If a company would like to monitor for this, it can specify a rule that will create an alert if any systems in the "desktop" group starts providing services over Web server ports. This can be useful to catch activity immediately after it occurs rather than waiting for the next weekly or monthly vulnerability scan.

In addition to these detection and alerting capabilities, Profiler has a very rich set of reporting tools. The system can report on information from two primary traffic data sets:

1) **Baselines:** As mentioned above, for a particular time frame, the baselines indicate what behaviors are "typical" for any host on the network.
2) **Flow log:** This is a log of all time-stamped flow records sent to the Profiler by Mazu Sensors or by NetFlow. This is useful for forensic reporting, analysis and troubleshooting

Next Generation Internal Network Security

To secure the internal network against worms and insider threats, companies need to turn to Mazu Profiler. Profiler is the leading behavioral IPS for the enterprise. It has several technological advancements that make it unique among other solutions:

1. Adaptive Profiling Engine
2. Blade Architecture
3. Power Search and/or audit and verify (reports)

1. Adaptive Profiling Engine

The Mazu Profiler is the only behavioral IPS with an Adaptive Profiling Engine. This unique technology enables the system to maintain more granular baselines and maintain its content automatically on an ongoing basis. The result is the only solution that can maintain accuracy in complex and dynamic environments.

Dynamic Baselining

Most Behavioral IPSs use static baselining techniques. They snapshot the network, creating white lists of traffic patterns and relationships between hosts. The problem is that networks are dynamic. New data centers are brought online, older applications are decommissioned. New employees start all the time, a group of contractors leave, and an acquisition is merged into the network. Static baselines simply go increasingly out of date, causing an increasing amount of false alerts. Eventually, systems need to be taken offline and the network must be rebaselined.

All of this creates more work, generates more false alerts and creates more exposure for the network.

Mazu Profiler's dynamic baseline adjusts itself over time to changes in how the network is used. Activity that disappears from the network automatically ages out of the model over time. New hosts or services in the network are identified, analyzed for potential threats and then worked into the model. The result is a baseline that is always current, so that the system always knows what "typical behavior" is. There are fewer false alerts, less maintenance effort and fewer windows of vulnerability typical of systems that need to periodically rebaseline the network.

Multi-Baselining

Another example of changes is the way networks are used in different cycles; weekdays versus weekends versus end of quarter. Most behavioral IPSs use a single baseline no matter where they are in the business cycle. The problem is that what looks odd for a typical Monday morning may be completely normal for an end of quarter. What looks normal for a Monday morning may be very odd for Sunday at 2am.

Mazu Profiler is the only behavioral IPS that builds and maintains multiple baselines in order to independently track different parts of the business cycle. A configurable setting allows administrators to set up multiple baselines such as weekday, weekend, weeknight, end of quarter, etc. The result is far greater accuracy in assessing deviations from typical behavior. This helps reduce or even eliminate false positives from traffic that seems odd for a Monday, but normal given that it's end of quarter. It alleviates false negatives such as flagging behavior that is typical for that same Monday morning, but very unusual for Sunday at 2am.

Adaptive profiling helps Mazu Profiler be more accurate in complex and dynamic environments.

2. Blade Architecture

The Mazu Profiler is the only behavioral IPS designed to exploit the blade server architecture, making it the only solution capable of scaling to meet the needs of even the largest, most complex networks.

Mazu Profiler is unique in its ability to accurately detect and characterize new and zero-day attacks and simulate policy changes or mitigation actions before they are executed. However, securing the internal network has the added challenge of being a distributed problem. Unlike the perimeter, where there are fewer links to watch, the internal network has a multitude of links. And given that an infected laptop can be inserted anywhere in the network, it is critical to watch over the entire network.

Yet, the problem cannot be solved by simply deploying more boxes all over the network, each analyzing its own link and correlating only its own alerts. Many threats are difficult to detect and understand from the single-link perspective. In addition, this approach leads to duplicate alerts and the inability to correlate all hosts associated with an event to the event itself. Lastly, with no unified networkwide model of all network activity, it is very difficult to assess how a mitigation action will impact normal operations of the network. As a result, you are more likely to disrupt critical services accidentally when mitigating attacks.

Solving these problems requires the ability to collect data from all across the network and centrally model and analyze the data. Higher-end behavioral IPSs take this approach, leveraging distributed devices such as routers, switches and probes and analyzing their data on a single system. The challenge is scaling the solution for a large network. The amount of data to analyze can be staggering.

The Mazu Profiler runs only on IBM @server® BladeCenter™. It distributes processing of network profile information across multiple “server blades” that analyze traffic data in parallel. An administrator can scale the system as the needs grow simply by adding analysis blades. It provides a unified, coherent view of networks of hundreds of thousands of hosts in real time. BladeCenter is the key to providing this unparalleled scalability with its superior scale out architecture. It is an industry-leading superior implementation of the blade server concept of physical consolidation of servers into a smaller, more manageable environment to achieve efficiencies of operation.

The BladeCenter design brings the customers’ computing resources into a cost-effective, highly reliable, modular new form factor at up to twice the density of conventional comparable 1U Intel® processor-based servers. Coupled with the fastest Intel Xeon® processors, modular Fibre Channel and Ethernet switches (Cisco® layer 2, and Nortel® NetworksLayer 2-7) built into the BladeCenter chassis and advanced management of storage, networking, servers and applications through IBM Director, organizations can take control of the computing environment and potentially reduce costs. Physical costs alone can potentially be reduced with a smaller footprint for multiple servers (14, 2-Way in 7U) and up to an 83% reduction in cabling. BladeCenter supports IBM’s TotalStorage and networking solution in a common fully managed architecture. Additionally, BladeCenter often takes less time to install, can require fewer people to manage and maintain, provides modular scalability and provides an environment with almost no single points of failure.

BladeCenter – Mazu Profiler Base Configuration

BladeCenter Chassis

mBlade

Manager Blade

Base configuration supports up to 50,000 hosts.

The advantages of this approach are:

- **Greater accuracy** – There are fewer false positives, fewer missed events and fewer duplicate alerts.
- **Lower management cost** – Multiple devices and user interfaces in separate locations require more maintenance and trained staff. A single, centralized device minimizes operational costs.
- **Better network visibility** –
    - The centralized BladeCenter always has a coherent, up-to-the-minute view of the entire network. Distributed devices may have stale data or only partial views.

    - Some vendors sell distributed systems that only aggregate events, not network data. This level of aggregation gives little insight into actual network behavior, and often requires laborious manual event correlation.
- **Better attack response** – The Profiler Blade System coordinates enterprisewide attack responses from a single location, decreasing both labor and response time.
- **Low network bandwidth utilization** – Mazu uses advanced data-compression algorithms to minimize the bandwidth required to transfer traffic statistics from Sensors to the Profiler.

BladeCenter System Profiler – Scaling Up

- **Each additional blade supports 50,000 hosts**
- **One system holds up to 10 analyzer blades to support up to 500,000 hosts**

3. Power Search

Power Search provides the industry's most powerful and flexible querying and reporting capabilities.This tool enables deeper investigation of events, real-time analysis of the network to dive deeper into an attack or simulate a policy change, sophisticated forensic analysis, and proactive analysis to identify critical assets and harden the internal network.

Macro and Micro

In the traditional reporting structure, reports are based in two categories, management reporting for internal decision-making and production reporting to generate operational documents. Many companies need both levels of detail in order to understand and make good business decisions. However, many traditional companies provide only the production or the management level of reports. Power Search in contrast, instantly provides both levels of information. This allows the managers to understand what is going on in their networks and the IT administrators to see how individual hosts are behaving.

Instant Reports

Power Search's Instant Reports give users real-time access to reports on the top talkers in the network, top services, groups, etc. These reports, created by Profiler within seconds, serve to identify critical assets, processes and conversations, and provide a quick assessment of the network. Power Search also integrates Crystal Reports, enabling customers and third parties to create a virtually unlimited set of reports. Third-party Power Search reports currently include packages for firewall tuning, group identification and management-level reporting.

Power Search also offers an array of analytical and reporting tools that leverage the system's detailed insight into how the network is actually used. These tools can be used to reduce risky guesswork, rough approximations and manual processes often associated with network analysis, planning and operational initiatives. Power search leverages its statistical model of the network and grouping algorithms to create a dependency map.

Dependency mapping and automated role grouping can greatly improve both the process and results of projects such as the following:

- Access policy setup, maintenance and auditing
- Application inventory and usage assessment
- Server consolidation
- Disaster recovery planning
- Network segmentation
- Regulatory compliance

Power Search can help companies with compliance requirements associated with regulations such as Sarbanes Oxley. Power Search provides the detailed information and documentation that corporations need for the CEOs and CFOs to certify that their financial records have been controlled and are uncompromised. In addition, Power Search allows administrators to show auditors detailed analysis and operating reports. This information demonstrates that the company can monitor, secure and audit the security of their financial records.

Conclusion

Worms and insider threats represent a clear and present danger to your business. They have the power to disrupt critical services, drain critical resources, compromise sensitive data and put companies at regulatory risk. Isolating your network is not an option. On the contrary, business pressures are forcing companies to make applications and data inside the corporate network more accessible to remote workers, business partners and contractors.

Patch management and perimeter security appliances such as firewalls and IPSs are helping to secure networks, but do little to help protect against new attacks, insider threats and malicious code that circumvent the perimeter on mobile devices.

Enterprises and government organizations need to accurately mitigate these forms of attacks. Companies and organizations need a solution to block malicious traffic without blocking legitimate traffic, simulate and understand the impact of mitigation, provide networkwide coverage, and mitigate attacks and exploits rapidly.

For a growing number of worldwide enterprises, that solution is the Mazu Profiler. Now on it’s fifth generation, Mazu Profiler enables organizations to leverage their existing infrastructure of routers, switches and probes to detect and surgically mitigate new attacks, harden the internal network against future attacks, and audit how sensitive assets are used and by whom. It leads the behavioral IPS market in its ability to scale, remain accurate in highly dynamic environments, surgically mitigate threats and report on virtually any process or activity. The result is the most comprehensive protection against worms and insider threats.

Mazu Profiler in combination with the BladeCenter technology platform delivers a leadership solution to deal with the challenge of worms and internal threats.

Produced in the USA
November 2004

Visit *www.ibm.com/pc/safecomputing* periodically for the latest information on safe and effective computing.
Warranty Information: For a copy of applicable product warranties, write to: Warranty Information, P.O. Box 12195, RTP, NC 27709, Attn: Dept. JDJA/B203. IBM makes no representation or warranty regarding third-party products or services including those designated as ServerProven or ClusterProven.

The examples given in this paper are hypothetical examples of how a customer can use the products described herein and examples of potential cost or efficiency savings are not based on any actual case study. There is no guarantee of comparable results. Many factors determine the sizing requirements and performance of systems architecture. IBM assumes no liability for the methodology used for determining the configurations recommended in this document nor for the results it provides. Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements quoted in this presentation may have been made on development-level systems. There is no guarantee these measurements will be the same on generally-available systems. Some measurements quoted in this presentation may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information in this presentation concerning non-IBM products was obtained from the suppliers of these products, published announcement material or other publicly available sources. IBM has not tested these products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.