HP StorageWorks Continuous Access XP Journal advanced technology white paper

HITACHI
DATA SYSTEMS

hp invent

Executive summary

Business continuity requirements are changing. When establishing and updating their business continuity plans, enterprises must respond to new business drivers such as round-the-clock operations, higher service-level expectations, closer regulatory scrutiny with the emergence of stringent out-of-region data protection requirements, and increased sensitivity to loss of data and information assets.

The challenge is to reduce risk and increase business resilience, while also reducing costs and increasing efficiency.

HP StorageWorks Continuous Access XP Journal for HP StorageWorks XP12000 Disk Array is an advanced technology for asynchronously replicating data hosted on the XP12000 Disk Array—or on externally attached storage systems from HP and other vendors. Continuous Access XP Journal evolves and improves the HP industry-leading asynchronous remote copy solutions.

Continuous Access XP Journal uses disk-based journaling and a pull-style replication engine to reduce resource consumption and costs, while increasing performance and operational resilience. These features make it more efficient and cost-effective than traditional replication methods.

Continuous Access XP Journal also makes out-of-region replication more cost-effective, and enables wider adoption of advanced three data center (3DC) business continuity strategies that provide "no data loss" replication and business continuity protection over extended distances.

With these advanced technology capabilities, Continuous Access XP Journal and the XP12000 Disk Array enable enterprises to improve business continuity, reduce costs and complexity, and increase business resilience.

Introduction

Risk reduction and business resilience

Many organizations—along with stakeholders and regulatory authorities—are focusing on operational risk assessment, and one key area of focus is information systems and data protection. Companies typically need to enhance system availability, and develop or update a robust disaster recovery strategy.

Compliance mandates. In many industries and geographies, government regulations require companies to have effective business continuity plans that will enable them to protect information assets and maintain their service capabilities in spite of local or regional disaster. The most commonly regulated industries likely to adopt out-of-region strategies worldwide include telecom, transportation, banking and other financial services, government, utilities, healthcare, and e-commerce.

Business resilience. In many cases, IT is still falling short in providing organizations with a fully resilient environment. One key requirement is data replication for business continuity—with guaranteed data integrity and consistency.

Wider scope of protection. While some organizations can define a very narrow scope for data protection and replication, many demanding environments will need to replicate essentially all of their actively changing data sets. And the volume of data is growing exponentially.

Better RTO and RPO. For business continuity planning, two of the key metrics are recovery time objective (RTO) and recovery point objective (RPO). The RTO measures how long it takes to resume essential operations: how long it takes to get back on your feet. The most demanding environments may have a hot site that is fully redundant—with the needed equipment, applications, people, and processes—using technologies such as server clustering to minimize RTO in case of a primary site failure. The RPO is a measure of data currency: how far behind the organization can afford to be when resuming operations after an interruption. Depending on the data location and the nature of the interruption, the copy maintained at the recovery site could be out of date by minutes or even hours.

Many organizations are seeking to improve RTO and RPO, while also protecting their operations and data from a wider range of threats and disaster scenarios.

Cost reduction and efficiency

Cost pressures continue to constrain IT organizations in their business continuity initiatives, so organizations must find ways to meet requirements while reducing costs and increasing efficiency.

Migration and consolidation. Organizations may need to replicate data to and from multiple data centers for testing, upgrades, or load balancing to improve IT operations. Consolidation of computing assets may improve efficiency and reduce costs, but is also akin to putting more eggs in one basket. The consolidated platform or data center requires a higher level of data protection and disaster resilience.

Resource utilization. IT organizations are expected to support higher levels of application performance and responsiveness, while also improving the level of protection and resilience. Business continuity solutions can use remote replication to improve resilience and rapid recovery capabilities, but the replication solutions themselves may consume scarce resources and introduce management complexity.

Management complexity. IT organizations must manage complex and heterogeneous server and storage infrastructures, and this makes business continuity planning and testing more difficult to manage.

Organizations must support local and remote replication using a combination of disjointed replication solutions, in the absence of centralized and universal management tools. This makes business continuity testing a challenge, and increases the risks and costs of failure in a real recovery situation.

This white paper takes a closer look at the strengths and issues of existing approaches to remote replication and business continuance, and introduces a new technology that helps IT organizations meet their data movement and protection needs more simply and efficiently.

Remote replication: Strengths and issues

Data replication is widely used for business continuance, to protect businesses from interruption of critical applications and business processes. Organizations have tried a variety of replication approaches based on server, storage, and network capabilities. Each has its strengths and issues, but storage-based replication is generally considered the best approach for replication of critical data, particularly in heterogeneous server and application environments.

Historically, storage arrays were not designed for replication. However, replication capabilities have been added to storage systems so successfully that storage-based replication is widely used today. It provides the best performance for demanding applications, and has a proven track record for data protection and disaster recovery.

Strengths

For business continuity, disk-based remote replication has greatly improved RTO and RPO, compared with earlier tape-based methods. Customers and analysts agree that storage-based replication is the platform of choice for replication of critical data, generally superior to server-based or fabric-based replication. As one HP customer remarked, "…HP technologies allow us to provide higher redundancy, availability and response for our service-level agreements."

HP customers are quick to note that they have never lost data with their HP storage and HP StorageWorks Continuous Access XP software.

Local disasters. Synchronous replication to an in-region hot site is now widely used for business continuity, to protect businesses from interruption of critical applications and data access at a production data center. However, synchronous replication is limited to relatively short distances—typically less than 50 miles—and is suitable for replication to in-region recovery sites. This approach does not protect from regional disasters that may affect both the production site and the in-region recovery site.

Regional disasters. Use of asynchronous replication is growing, as companies respond to operational risks and compliance requirements. Asynchronous remote replication enables organizations to maintain very current data copies at out-of-region recovery sites. HP is the leader in asynchronous remote replication, based on the success of HP StorageWorks Continuous Access XP Extension.

However, as requirements have evolved over time, and with soaring amounts of replicated data, customers have pushed the limits of current asynchronous replication technology. Furthermore, regulations worldwide are endorsing (if not requiring) out-of-region replication for critical industries such as banking and securities trading. A number of issues have arisen that indicate the need for an improved solution. HP now delivers that improved solution with HP StorageWorks Continuous Access XP Journal on the HP StorageWorks XP12000 Disk Array.

Issues

For all their benefits, current remote replication technologies are not perfect. Enterprise IT managers have identified common issues with remote replication solutions from all vendors. Here are a few typical comments:

- "This stuff uses up all my cache."
- "I know replication is crucial, but I'd like to focus system resources on my applications' transactions."
- "Bandwidth requirements are killing me."
- "I wish I had only one tool for replicating all my data."
- "How can I implement out-of-region disaster recovery without losing data?"

The issues generally involve resource utilization, network bandwidth, management complexity, and the trade-offs between alternative replication strategies.

Resource utilization. Current remote copy technologies can consume tremendous amounts of resources. In server-based approaches, replication consumes server memory to buffer the changed data during replication. In storage-based solutions, replication uses part of the storage system cache to capture changes and transmit them to the other side. It also uses processing cycles on the storage systems—primarily at the originating (production) data center. These resources are, in effect, taken away from production applications. The result is lower application performance, or the increased cost of adding resources to maintain required performance and throughput.

Data growth makes these problems worse as resource consumption increases at the production site. Customers are looking for a better solution. Continuous Access XP Journal, with its replication engine, now returns these resources to where they belong: the application.

Network problems. Replication network bandwidth is another key resource consumed by remote replication technologies. Bandwidth is expensive, and replication requires more of it. Also, since momentary link failures or network congestion can cause cache-based replication to fail, most customers plan and provision bandwidth to meet the maximum peak-load requirement. If they do not, the replication can fail, forcing customers to perform difficult processes to recover and re-establish replication. This can be painful and disruptive.

Replication management complexity. Simplification of management complexity is another requirement from customers. Because of the number of storage vendors and products in a heterogeneous storage environment, many customers must learn and use different replication tools for the various systems. Many customers are asking for a single tool for managing multiple storage-based replication platforms.

Replication strategy choices. In-region replication using synchronous remote copy technology can provide recovery with little or no data loss, but is vulnerable to regional disasters. Out-of-region replication using asynchronous replication protects against regional disasters, but may lose some recent transactions. Multi-data center configurations can combine both approaches. However, existing solutions require storage for multiple copies of the data, as well as complex management and scripting.

Continuous Access XP Journal builds on existing HP StorageWorks XP strengths in remote replication technology, and introduces advanced technology to address these emerging customer issues with a more efficient and cost-effective solution. The next section describes Continuous Access XP Journal features and benefits, and later sections outline improved business continuity approaches that Continuous Access XP Journal makes technically possible and economically feasible.

HP StorageWorks Continuous Access XP Journal: Overview

HP StorageWorks Continuous Access XP Journal for HP StorageWorks XP12000 Disk Array is an advanced technology for asynchronously replicating data hosted on XP12000 Disk Array or on externally attached storage systems from HP and other vendors.[1]

Continuous Access XP Journal uses disk-based journaling and a pull-style replication engine to reduce resource consumption and costs, while increasing performance and operational resilience. These features make it more efficient and cost-effective than traditional replication methods.

Continuous Access XP Journal also enables more resilient and cost-effective business continuity solutions, especially when combined with Continuous Access XP in 3DC replication and recovery strategies.

The following sections describe two key technical innovations that underlie the strengths and benefits of Continuous Access XP Journal:

- Performance-optimized, disk-based journaling
- Pull-style replication engine

[1] For a discussion of the HP StorageWorks XP12000 Disk Array and business continuity applications of its capabilities, including HP StorageWorks Continuous Access XP Journal, see *Business Continuity and the HP StorageWorks XP12000 Disk Array white paper*. Customers should check with HP on supported platforms and configurations.

Journal-based replication

Figure 1 illustrates HP StorageWorks Continuous Access XP Journal's application of performance-optimized, disk-based journaling.

In this illustration, online transaction processing (OLTP) data volumes at the primary data center are being replicated to a second HP StorageWorks XP12000 Disk Array at the remote data center.

Figure 1. HP StorageWorks Continuous Access XP Journal: Asynchronous Remote Copy

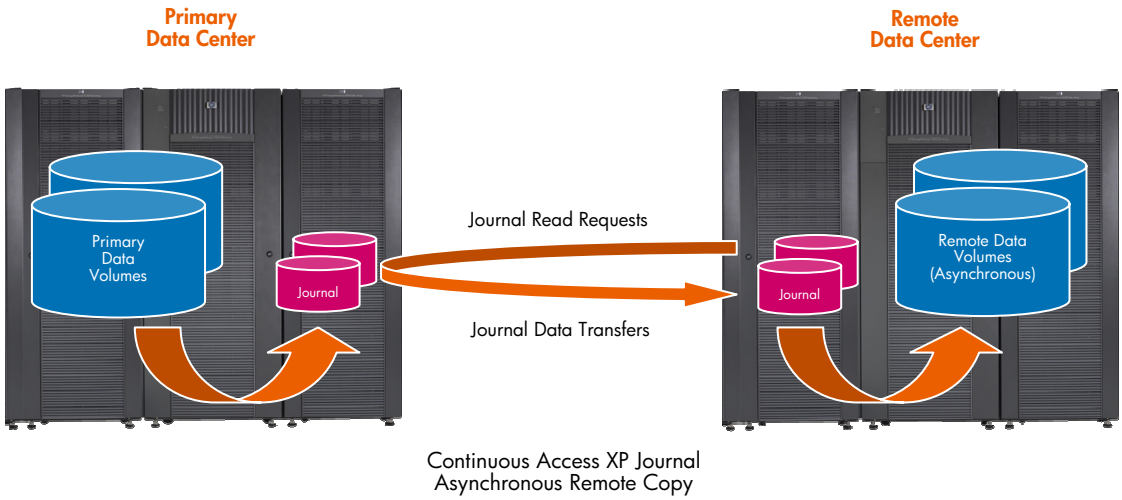

When collecting the data to be replicated, the primary XP12000 Disk Array writes the designated records to a special set of journal volumes. The remote storage system then reads the records from the journal volumes, pulling them across the communication link as described in the next section.

By writing the records to journal disks instead of keeping them in cache, Continuous Access XP Journal overcomes the limitations of earlier asynchronous replication methods. Writes to the journal are cached for application performance reasons, but quickly de-staged to disk to minimize cache usage. The journal disks are specially architected and optimized for maximum performance.

Guaranteed data integrity. In addition to the records being replicated, the journal contains metadata for each record to ensure the integrity and consistency of the replication process. Each transmitted record set includes sequence number information, enabling the replication engine to verify that all the records are received at the remote site and to arrange them in the correct write order for storage. These processes build on the proven algorithms of Continuous Access XP. The journaling and replication processes also support consistency across multiple volumes. Thus Continuous Access XP Journal builds on the proven integrity of existing HP replication solutions, while introducing innovative journaling technology to improve performance and efficiency.

Pull-based replication

In addition to disk-based journaling, HP StorageWorks Continuous Access XP Journal introduces pull-style replication. The primary storage system does not dedicate resources to pushing data across the replication link. Rather, a replication engine on the remote system pulls the data from the primary system's journal volume, across the link, and writes it to the journal volume at the receiving site. The

replication engine then applies the journaled writes to the remote data volumes, using metadata and consistency algorithms to ensure data integrity.

In the default configuration, Continuous Access XP Journal considers replication complete when the data is received in mirrored system cache at the remote system, written to the journal disk, and applied to the remote data volumes.

Since the engine that controls asynchronous replication is located on the remote system, this approach shifts most of the replication workload to the remote site, reducing resource consumption on the primary storage system and improving production application performance. In effect, Continuous Access XP Journal restores primary site storage to its intended role as a transaction processing resource, not a replication engine.

HP StorageWorks Continuous Access XP Journal: Benefits

HP StorageWorks Continuous Access XP Journal significantly reduces resource consumption, and sets a new standard for data protection by enabling replication over any distance for business continuity or improved IT operations, without the need for redundant servers or replication appliances.

Continuous Access XP Journal, running on the HP StorageWorks XP12000 Disk Array, delivers several important benefits:

- Resource optimization
- Mitigation of network problems
- Simplified replication management
- Advanced business continuity configurations

Resource optimization

HP StorageWorks Continuous Access XP Journal significantly reduces consumption of production-site resources, such as storage system cache and processing power. Essentially, it liberates resources and improves performance.

By using local disk-based journaling and a pull-based remote replication engine, Continuous Access XP Journal releases critical resources that are consumed by other asynchronous replication approaches at the primary site, such as disk array cache in storage-based solutions, or server memory in host-based software approaches. Continuous Access XP Journal improves cache utilization, lowering costs and improving performance of production transaction applications. It also maximizes the use of bandwidth by better handling the variations of the replication network resources, enabling enterprises to manage bandwidth cost and RPO more flexibly and intelligently.

The pull-style replication engine also contributes to resource optimization. It controls the replication process from the secondary system and frees up valuable production resources on the primary system.

Mitigation of network problems

In traditional asynchronous replication, typical issues include temporary communication problems, such as communication link failure or insufficient bandwidth for peak-load requirements. These conditions can cause cache-based “push” replication methods to fail. When this happens, traditional replication solutions suspend the replication process and go into bitmap mode, noting changed tracks in a bitmap for future resynchronization. Recovery typically involves a destructive process such as rewriting all the changed tracks, with possible loss of data consistency for ordered writes.

Increased resilience. In contrast, HP StorageWorks Continuous Access XP Journal logs every change to the journal disk at the primary site, including the metadata needed to apply the changes consistently. Should the replication link between sites fail, Continuous Access XP Journal keeps logging changes in the local journal so that they can be transmitted later, without interruption to the protection process or the application. The journal data is simply transferred after the network link failure or bandwidth limitation is corrected, with no loss of consistency. The RPO may be extended a bit during temporary network failures or congestion, but the asynchronous replication process does not fail, and the catch-up process is simple and automatic. Data consistency is preserved, and the user can maintain a more current RPO during outage situations, compared with existing replication methodologies.

Improved bandwidth costs and RPO. With Continuous Access XP Journal, the remote storage system pulls data from the primary journal volumes—over the data replication network—as fast as the bandwidth allows, adjusting to available network conditions. If available bandwidth does not support optimal replication (for example, during peak-load spikes in transaction volume), the primary journal volumes simply buffer the data on disk until more bandwidth becomes available.

This technology approach gives enterprises a more flexible way to manage network bandwidth. They can improve bandwidth utilization and lower their communication costs by provisioning bandwidth for the average traffic level—not for peak usage—and sizing the journal volumes to handle the peak load.

Enterprises can manage their RPO intelligently, in relation to infrastructure and communication costs. The overall result is better replication with improved RTO and RPO, lower costs, and greater resilience.

Simplified replication management

HP StorageWorks Continuous Access XP Journal will simplify replication management by leveraging the HP StorageWorks XP12000 Disk Array's support for externally attached storage from multiple vendors. Continuous Access XP Journal will ultimately support real-time replication for all internal and externally connected disk storage volumes, making it a very powerful solution for remote copy operations.[2] Thus XP12000 Disk Array with Continuous Access XP Journal will support the data replication and movement requirements of disparate applications and business processes, managed from a single pane of glass. It will enable a consistent replication approach and a common set of procedures, tools, and skills to manage diverse, heterogeneous storage and application environments.

Advanced business continuity configurations

HP StorageWorks Continuous Access XP Journal, in combination with synchronous HP StorageWorks Continuous Access XP, supports improved performance and cost reduction in multi-site business continuity configurations. Advanced 3DC configurations intelligently combine synchronous and asynchronous replication to provide a "no data loss" replication capability at any distance.

Advanced configurations can also use fewer intermediate copies than alternative approaches, while placing the copies on less expensive storage and reducing resource consumption at the primary site.

Thus, Continuous Access XP Journal makes the business continuity benefits of 3DC configurations practical and affordable for a wider range of organizations, applications, and business processes.

[2] Customers should check with their HP representatives regarding the scheduled availability of HP StorageWorks Continuous Access XP Journal support for externally connected disk subsystems from HP and other vendors.

Remote replication configurations for business continuity

HP StorageWorks Continuous Access XP Journal improves remote copy solutions, making out-of-region replication more efficient and cost-effective, and thus enabling its wider use.[3] More and more organizations will find that out-of-region replication is both necessary and feasible, so this section provides an overview of the characteristics and benefits of several alternative approaches. It includes a description of several two data center (2DC) and 3DC configurations, followed by a technical comparison and a summary of the relative benefits of each approach. The next section then shows how Continuous Access XP Journal improves these alternatives—particularly the 3DC configurations—and makes them more efficient, affordable, and cost-effective for a wider range of organizations.

Figure 2 illustrates several replication configurations involving two or three data centers.

Figure 2. Remote replication configurations: 2DC and 3DC

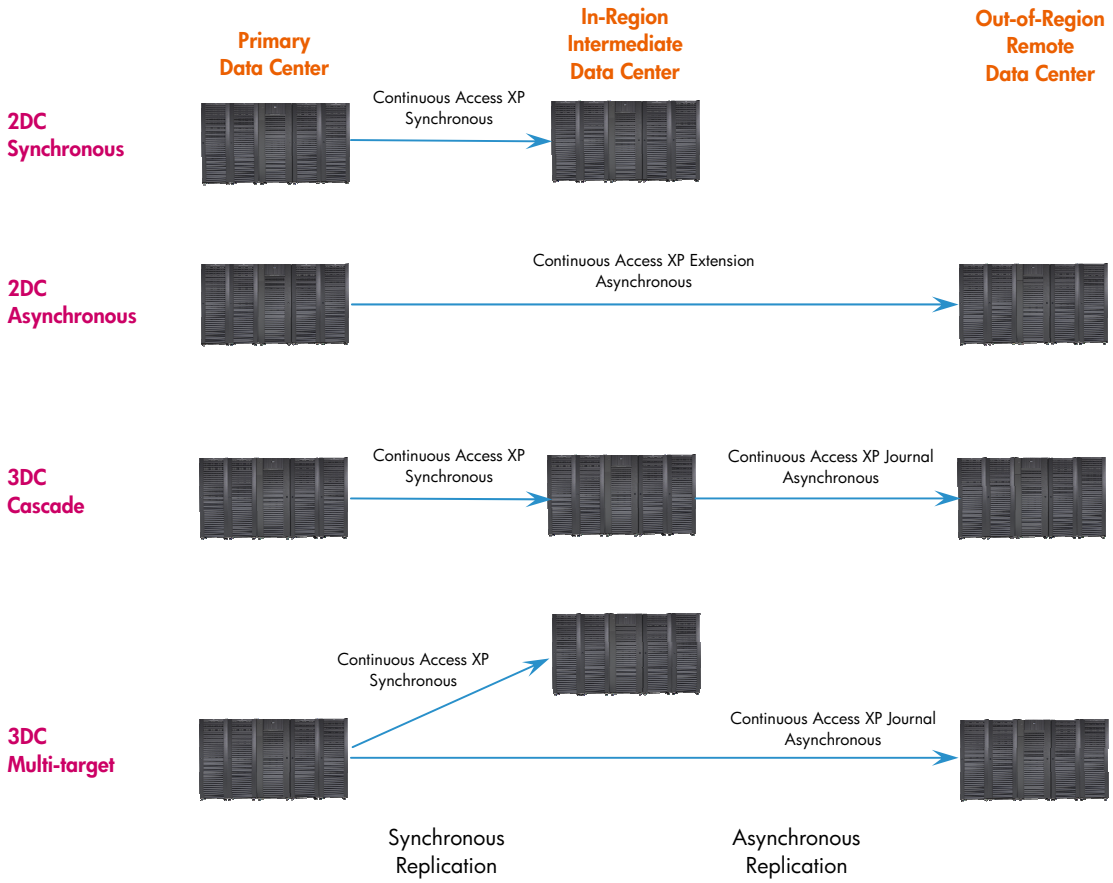

Remote replication strategies may involve two data centers or three data centers and use in-region or out-of-region replication—or both.

[3] To support business continuity objectives, many enterprises send copies of production data to remote sites for use in disaster recovery processes. Early disaster recovery approaches used local replication or backup software to create copies on removable media, and then shipped the media to an offsite recovery provider. Such approaches support RPOs and RTOs of several days, and may be appropriate for less critical applications and some fixed content. However, many enterprises would suffer substantial economic damage if critical business applications were shut down for several days after a data center disaster. Electronic replication strategies involving interconnected data centers can provide faster recovery, less risk of data loss, and protection from a wider range of disruptive events.

2DC strategies involve replication from the primary data center to a second data center. The second data center might be an in-region recovery site, such as a hot site serving a campus-level or metro-level server cluster. Or the second site might be an out-of-region recovery site, located hundreds of miles away.

- **2DC Synchronous:** Synchronous replication can provide a very fast recovery time (low RTO) and good data currency (low RPO) when recovering at an in-region hot site. But in-region replication does not protect from regional disasters (such as power grid failures, earthquakes, and floods) that can impact critical business operations and recovery capabilities at both primary and secondary data centers.
- **2DC Asynchronous:** To reduce business risks from regional disasters, many organizations have established remote data replication to out-of-region data centers. Because synchronous replication over long distances can impact application performance, out-of-region replication typically uses asynchronous replication.[4] As already shown, HP StorageWorks Continuous Access XP Journal makes 2DC asynchronous replication more resilient in response to network issues, and more cost-effective in terms of resource utilization.

3DC strategies provide the best of both worlds: fast recovery and data currency for local site failures, combined with good protection from regional disasters. Figure 2 illustrates several configurations that are introduced in general terms and then described in more detail.

- **3DC Cascade:** Synchronous replication supports a current copy of production data at an in-region site that does provide processing capabilities for recovery. The in-region site also forwards the data to a remote recovery site, using asynchronous replication. 3DC cascade is the traditional 3DC replication configuration, and of course it benefits from Continuous Access XP Journal's more resilient and cost-effective approach to asynchronous replication.
- **3DC Multi-target:** Synchronous replication supports a current copy of production data at an in-region site that provides processing capabilities for recovery. A separate asynchronous replication process copies data from the primary site to the out-of-region recovery site, over a separate replication network. Until now, 3DC multi-target solutions have been deployed only in the most critical business continuity environments. However, Continuous Access XP Journal makes 3DC multi-target configurations feasible and practical for a wider range of organizations and applications.

Technical comparison of 2DC and 3DC strategies

This section presents a more detailed comparison of the previously introduced 2DC and 3DC replication configurations.

[4] Synchronous replication requires the application to wait for the remote site to confirm each write operation before sending the next write operation. As the replication distance increases, the time lag between synchronous write operations gets longer and longer, and these delays become intolerable for high-volume, write-intensive transaction processing applications. Asynchronous replication allows the production application to continue, and keeps track of the pending writes until they are complete. A well-designed asynchronous replication solution maintains an I/O-consistent copy at the remote site. It may deliver a somewhat lengthened RPO and RTO, but it does protect the enterprise from major data loss in case of a regional disaster.

Table 1. Technical comparison: 2DC and 3DC replication configurations

Data center strategy	1DC*	2DC		3DC	
Remote replication configuration	**Local**	**Sync with Failover**	**Async out of Region**	**Cascade**	**Multitarget**
Technology					
Synchronous replication	N/A	Continuous Access XP	N/A	Continuous Access XP	Continuous Access XP
Asynchronous replication	N/A	N/A	Continuous Access XP Extension	Continuous Access XP Journal for XP12000 Disk Array	Continuous Access XP Journal for XP12000 Disk Array
Characteristics					
Recovery locations supported	0	1	1	2	2
In-region site failover	No	Yes	No	Yes	Yes
Out-of-region site failover	No	No	Yes	Yes	Yes
"No data loss" available	No	Yes	No	Yes	Yes
Minimum number of logical full-volume copies	1	2	2	3	3
Primary site failure					
RPO	24–168 hrs	0–2 mins	0–5 mins	0–2 mins	0–2 mins
RTO	48–168 hrs	5–60 mins	1–8hrs	5–60 mins	5–60 mins
Regional disaster					
RPO	24–168 hrs	24–168 hrs	0–5 mins	0–5 mins	0–5 mins
RTO	48–168 hrs	48–168 hrs	1–8hrs	1–8hrs	1–8hrs
Impact of single failure outside primary site					
Recovery/intermediate site failure (primary still up)	N/A	No disaster recovery	N/A	No disaster recovery	OK
Remote out-of-region site failure (primary still up)	N/A	N/A	No disaster recovery	OK	OK
Link failures only—primary to in-region	N/A	No disaster recovery DR	N/A	No disaster recovery	OK
Link failures only—primary to out-of-region remote	N/A	N/A	No disaster recovery	N/A	OK
Link failures only—intermediate to remote	N/A	N/A	N/A	OK	N/A

Note: Alternative replication strategies provide different levels of protection in terms of data currency and recovery time, for various failure or disaster recovery scenarios. The RPO and RTO ranges in this table are HP estimates based on customer experience in real-world enterprise application environments.

* 1DC strategy represents a traditional tape-based disaster recovery approach.

Table 1 shows a number of quantitative parameters of each 2DC and 3DC configuration. For comparison, the traditional tape-based disaster recovery approach is included as the "one data center" (1DC) configuration.[5] The RPO and RTO ranges in this table are HP estimates based on customer experience in real-world enterprise application environments.

Technology. The chart shows that synchronous remote replication is supported by HP StorageWorks Continuous Access XP, and asynchronous replication is supported by HP StorageWorks Continuous Access XP Extension (on all HP StorageWorks XP arrays) and by HP StorageWorks Continuous Access XP Journal (on the HP StorageWorks XP12000 Disk Array only).

Characteristics. The 2DC configurations each support one recovery location: either an in-region recovery site or an out-of-region remote site, as illustrated in Figure 2. The 3DC cascade and multi-target configurations both support two recovery locations.

The "no data loss" capability provided by synchronous replication (in 2DC synchronous, 3DC cascade, and 3DC multi-target configurations) is best thought of as excellent data currency—within zero to two minutes—rather than absolutely no transactions lost.[6]

Primary site failure. The data currency is excellent, with target RPO levels of zero to two minutes, for 2DC Sync and for all of the 3DC configurations because they can all recover from a synchronous copy at the recovery site. In contrast, 2DC asynchronous replication can lag farther behind the primary-site data, with target RPO values up to five minutes or more, depending on available bandwidth and other factors.

In terms of RTO, the 2DC synchronous and the other two 3DC configurations provide the best RTO performance.

Regional disaster. In the case of a regional disaster, the 2DC synchronous copy at the in-region site can be lost along with the primary-site copy, forcing the organization to recover from backup tapes—with RPO and RTO comparable to traditional 1DC approaches. In contrast, the 2DC asynchronous and all the 3DC approaches provide similar RPO and RTO during a regional disaster: they all require recovery from an asynchronous copy, at the out-of-region data center.

Impact of single failure outside primary site. Even if the production site remains up and running, a failure or disruption that impacts one of the remote sites or communication links can become a problem. Such a failure can eliminate the disaster recovery capability that the primary site depends on, leading to longer outages and more data loss in the event of a subsequent failure of the primary site.

[5] Note on 1DC strategy: Organizations with one data center need to make other arrangements for disaster recovery, such as a warm site with capability to recover data from tapes and restart applications. Estimates in this table assume a traditional disaster recovery strategy with an offsite recovery site provider, and that current tape backups exist (including journal tapes if needed). This is the best case scenario.

[6] Real-world synchronous replication configurations—especially for transaction-processing databases—involve settings of "criticality" parameters that impact application performance as well as data currency. In practice, they are rarely set for "critical" write ordering. The technical details are beyond the scope of this paper.

For example, the 3DC cascade configuration provides ongoing protection if the out-of-region site fails, so the chart shows this scenario as "OK." In contrast, failure of the intermediate site (or either communication link) prevents ongoing replication to the remote site as well, leaving no active remote replication for primary-site data. Thus the chart shows "No Disaster Recovery" for this failure scenario.

In this respect, the 3DC multi-target configuration provides a significant benefit: either of the remote data centers can fail, and the other one will continue providing disaster recovery protection for the production site. If the out-of-region site fails, the in-region site continues to provide synchronous remote copy and excellent RPO and RTO in the event of a subsequent failure at the primary site. If the in-region recovery site fails, the out-of-region site provides ongoing recovery capabilities with asynchronous replication.

Considering all these different disaster scenarios and impacts, 3DC multi-target provides the broadest protection and the best combination of capabilities. The price can be higher as well, and each company must assess its risks and requirements when evaluating these advanced business continuity configurations. However, as seen in the next section, Continuous Access XP Journal makes 3DC configurations more cost-effective, thus enabling more enterprises and applications to implement improved protection at an affordable cost.

Benefits comparison

Table 2 summarizes the benefits of these 2DC and 3DC remote-replication configurations, compared to a traditional 1DC approach that relies on physical tape movement.

Table 2. Benefits comparison summary

Data center strategy	**1DC**	**2DC**		**3DC**	
Replication configuration	**On-site**	**Sync Near**	**Async Far**	**Cascade**	**Multitarget**
Primary site failure/failover					
Speed of recovery (RTO)	Bad	Better	Good	Better	Better
Data currency (RPO)	Bad	Better	Good	Better	Better
Regional disaster (RTO)	Bad	Bad	Good	Good	Good
Protection after failure outside primary site	N/A	Bad	Bad	Depends	Better

Note: Alternative replication strategies provide different levels of protection.

The chart shows the relative benefits of these configurations in terms of recovery speed and data currency after a primary site failure, as well as recovery speed after a regional disaster. The chart also illustrates the impact of a failure outside the primary site, which could affect the ongoing level of protection and recovery capability in case of an additional site failure.

The scope of protection increases from left to right in Table 2:

- **The 1DC strategy** offers the least protection from data loss during disasters, and can require days to recover data from tape copies. The best-case scenario for recovery may be 24 to 168 hours, depending on the availability of data tapes at offsite recovery services.
- **The 2DC strategies** can provide good recovery speed and data currency for primary site failures, with certain trade-offs: synchronous in-region replication is the best data currency solution, but provides poor protection against a regional disaster. Asynchronous out-of-region replication is the best regional disaster protection, but the I/O-consistent copy at the remote site may not be 100 percent current (RPO > 0).
- **The 3DC strategies** can provide "zero data loss" at any distance after a primary site failure, with different levels of resilience in case of additional failures. The most familiar approach, 3DC cascade, provides superior recovery speed and data currency after primary site failures, along with good protection from regional disasters. However, a failure at the intermediate site can leave the primary site without ongoing disaster recovery protection. 3DC multi-target offers all the benefits of 3DC cascade, and also preserves ongoing replication and recovery capability if the intermediate data center or the remote data center fails while the production site survives.

The final section of this paper discusses the benefits of Continuous Access XP Journal advanced technology in each of the 3DC strategies.

HP StorageWorks Continuous Access XP Journal: Improving out-of-region replication

Out-of-region replication is increasingly important for business continuity, compliance, and business resiliency. Strategies for 3DC combine in-region and out-of-region replication to provide the strongest protection: fast recovery and data currency for local site failures, combined with good protection from regional disasters. However, multiple data centers and data copies increase costs, so robust 3DC strategies have typically been limited to large organizations with extremely critical business continuity needs.

HP StorageWorks Continuous Access XP Journal makes remote replication more efficient, enabling more organizations and applications to benefit from 3DC replication strategies. This section summarizes the advantages and benefits of Continuous Access XP Journal in each type of 3DC configuration.

3DC cascade

The most familiar 3DC configuration is 3DC cascade. Figure 3 illustrates a 3DC cascade configuration that uses synchronous HP StorageWorks Continuous Access XP to maintain a current copy of the production data at an in-region data center. As noted earlier, 3DC cascade configurations make sense when the in-region hot site provides processing capabilities for recovery.

Figure 3. 3DC cascade with HP StorageWorks Continuous Access XP Journal

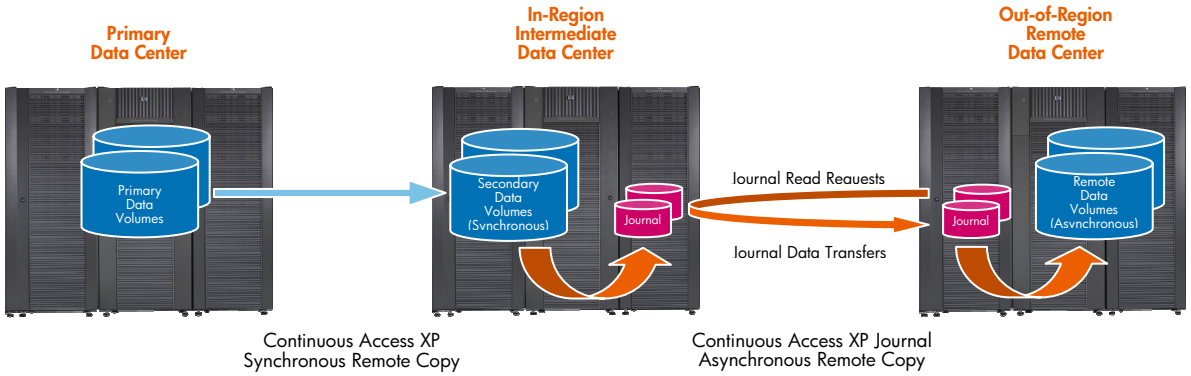

3DC cascade operation provides full volume copies at the in-region data center.

The HP StorageWorks XP12000 Disk Array at the in-region site also cascades the data to an out-of-region recovery site, using Continuous Access XP Journal asynchronous replication.

In comparison with other asynchronous replication technologies, Continuous Access XP Journal does not require an additional point-in-time copy of the data volume at the intermediate site. Continuous Access XP Journal stages the data to the journal disk, which is relatively small compared with a complete data copy. This feature saves physical disk space and reduces the cost of the 3DC configuration.

Continuous Access XP Journal also provides a more robust asynchronous replication technology that enables organizations to reduce bandwidth costs for the remote replication network. It can ride through temporary network outages and peak-load congestion without suspending replication or requiring human intervention or complex recovery procedures. In case of production-site failure, processing can resume at the intermediate site, using a current replica of production data. The in-region hot site can also support planned failover when needed for maintenance, upgrades, or business continuity testing. Meanwhile, Continuous Access XP Journal provides ongoing replication to the out-of-region site, maintaining robust business continuity protection. In case of a regional disaster, the out-of-region site enables rapid recovery with a consistent copy of production data.

3DC multi-target

The multi-target strategy is similar to the 3DC cascade approach in its support of full data copies at the in-region recovery data center. The multi-target configuration is appropriate when the in-region site provides processing capabilities for recovery, and when the organization must maintain business continuity protection despite failure of the in-region recovery site or one part of the replication network.

Figure 4. 3DC multi-target with HP StorageWorks Continuous Access XP Journal

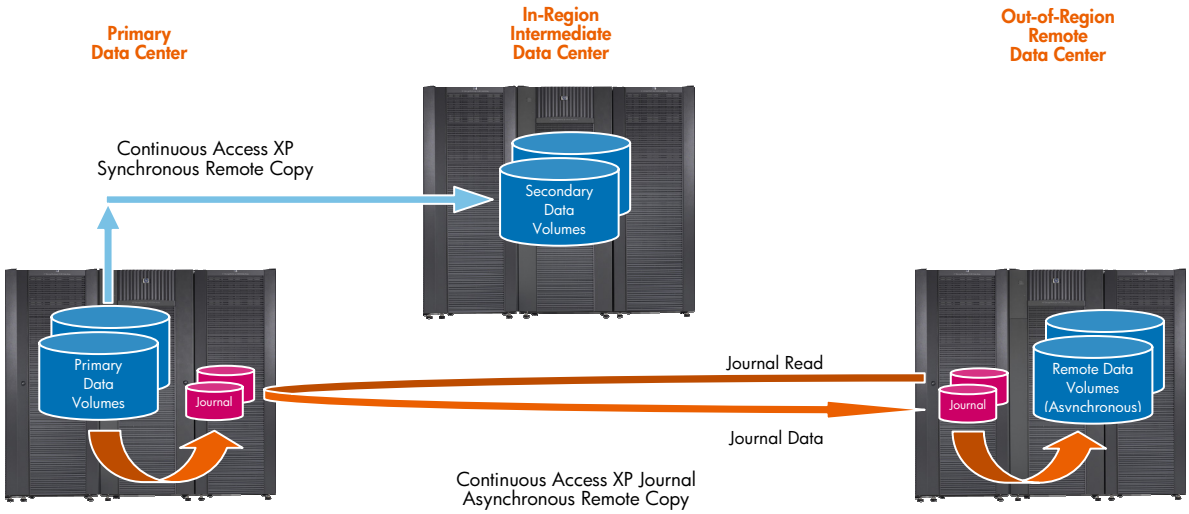

3DC multi-target provides independent replication paths to two recovery sites, and delivers the best protection against a wide range of potential outages at any site.

Figure 5 illustrates a 3DC multi-target configuration, in which data is replicated to two remote sites in parallel. Continuous Access XP maintains a current copy of the production data at an in-region recovery data center. Simultaneously, the XP12000 Disk Array at the primary site replicates the data to an out-of-region recovery site, using Continuous Access XP Journal asynchronous replication across a separate replication network.

In case of production site failure, processing can resume at the in-region recovery site, using a current replica of production data. The in-region hot site can also support planned failover when needed for maintenance, upgrades, or business continuity testing. Meanwhile, Continuous Access XP Journal provides ongoing replication to the out-of-region site, maintaining robust business continuity protection.

In case of a regional disaster, the out-of-region data center can recover rapidly with a slightly older but fully consistent copy of production data.

As noted earlier, 3DC multi-target offers all the benefits of 3DC cascade, and also preserves ongoing replication and recovery capability if the intermediate data center or the remote data center fails while the production site survives. Thus 3DC multi-target delivers the best overall business continuity protection. However, until now, 3DC multi-target solutions have been deployed only in the most critical business continuity environments.

Continuous Access XP Journal makes 3DC multi-target configurations feasible and practical for a wider range of organizations and applications by reducing resource consumption and bandwidth costs and by providing a more resilient technology for out-of-region replication. Because the primary data center must support synchronous and asynchronous replication simultaneously, Continuous Access XP Journal is particularly beneficial for a multi-target strategy because it reduces primary-site resource consumption for the asynchronous copy function.

Compared to a cascade approach, the multi-target configuration eliminates the potential of having no disaster recovery should the in-region recovery site fail. The trade-off is that a failure of the primary site will require recovery using a slightly older copy of production data at the out-of-region site.

In all of these 3DC configurations, XP12000 Disk Array with Continuous Access XP Journal enables enterprises to maintain "no data loss" replication and out-of-region recovery capability with less cost, complexity, scripting, and overhead than previous solutions.

Conclusion

Enterprises are faced with the challenges of meeting business applications' service levels, including increased resilience and protection from local and regional disruptions, while dealing with complex infrastructures and tight budgets.

Within this robust framework, HP StorageWorks Continuous Access XP Journal for HP StorageWorks XP12000 Disk Array delivers an advanced technology for asynchronously replicating data hosted on the XP12000 Disk Array or on externally attached storage systems from HP and other vendors.

Continuous Access XP Journal evolves and improves industry-leading asynchronous remote copy solutions from HP. It improves data currency and recovery time for minor outages. It also enables improved utilization of storage system cache and replication network bandwidth, reducing the cost of asynchronous replication between XP12000 Disk Arrays.

Continuous Access XP Journal also makes out-of-region replication more cost-effective. And it enables wider adoption of advanced 3DC business continuity strategies that provide "no data loss" replication and business continuity protection over extended distances.

With these advanced technology capabilities, Continuous Access XP Journal and the XP12000 Disk Array enables enterprises to improve business continuity, reduce costs and complexity, and increase business resilience.

Appendix: Complementary solutions

HP StorageWorks Business Copy XP

HP StorageWorks Business Copy XP is a local mirroring product that maintains one or several copies of critical data through a split-mirror process. Asynchronous copy volume updates ensure I/O response time for primary applications is not adversely affected. It provides full-copy/clone or snapshot/space-efficient local replication. Nine simultaneous clone copies or 32 simultaneous snapshot copies can be concurrently maintained.

HP StorageWorks Continuous Access XP, HP StorageWorks Continuous Access XP Extension, and HP StorageWorks Continuous Access XP Journal

HP StorageWorks Continuous Access XP, HP StorageWorks Continuous Access XP Extension, and HP StorageWorks Continuous Access XP Journal are high-availability data and disaster recovery solutions that deliver host-independent real-time remote data mirroring between XP disk arrays. Continuous Access XP provides synchronous replication. Continuous Access XP Extension and Continuous Access XP Journal extend capabilities to include asynchronous replication. A variety of remote connectivity infrastructures can be deployed to facilitate array-to-array connections—ESCON, Fibre Channel, ATM, and IP.

For more information

For more information on the HP StorageWorks Continuous Access XP Journal advanced technology, visit: http://www.hp.com/go/enterprisestorage

5983-1290EN, 07/2005