Data Replication Considerations

WHITE PAPER

This white paper provides an overview of data replication alternatives and a review of the considerations enterprises must weigh in determining the appropriate use of replication technology.

By understanding the differences between server-based replication, controller-based replication, and SAN-based replication, you can determine the best approach for your environment. Also discussed are the various replication modes (synchronous, asynchronous, and point-in-time), as well as the factors involved in choosing the right combination of replication mode and platform.

Table of Contents

Introduction 2

Data Replication Alternatives 2

Server-Based Replication 3

Controller-Based Replication 3

SAN-Based Replication 4

Replication Mode 4

Choosing a Data Replication Solution 6

Conclusion 8

Introduction

Data replication technology is finding its way into every aspect of the IT application infrastructure. Moving data between servers, a task that once meant physically carrying tapes between locations, can now be streamlined by transmitting data directly over a network. Enterprise, metro-area, and wide-area networks, available at relatively low cost, offer the bandwidth and speed necessary to efficiently transfer data across the data center, or across the world, for disaster recovery, content distribution, hardware migration or system consolidation.

Data replication has revolutionized enterprise disaster preparedness. Organizations large and small rely increasingly on software applications for every facet of business activity. Any interruption of access to application data can be devastating. Using data replication technologies, organizations are now able to safeguard the IT infrastructure from disaster.

For IT staff involved in configuring replication, staying abreast of business process vulnerability to data loss and changes in replication technology is essential. This paper aims to provide an overview of data replication alternatives and a review of the considerations enterprises must weigh in determining the appropriate use of replication technology.

Data Replication Alternatives

The term data replication covers a broad spectrum of methods for creating duplicate copies of operational data, either locally or at a remote data center. Like other applications of data replication technology, disaster recovery planning requires comprehensive and universally applicable data copy functionality that offers high-performance movement of data across significant distances.

The storage infrastructure provides multiple platforms for hosting enterprise data replication processing. Whether operating from an application server, the controller of a high-end storage array, or from a network-resident dedicated device, the data replication application creates mirror images of primary system data on secondary volumes located anywhere on the network. The backup copies can then be made available to business applications if an outage disrupts access to the primary source of data.

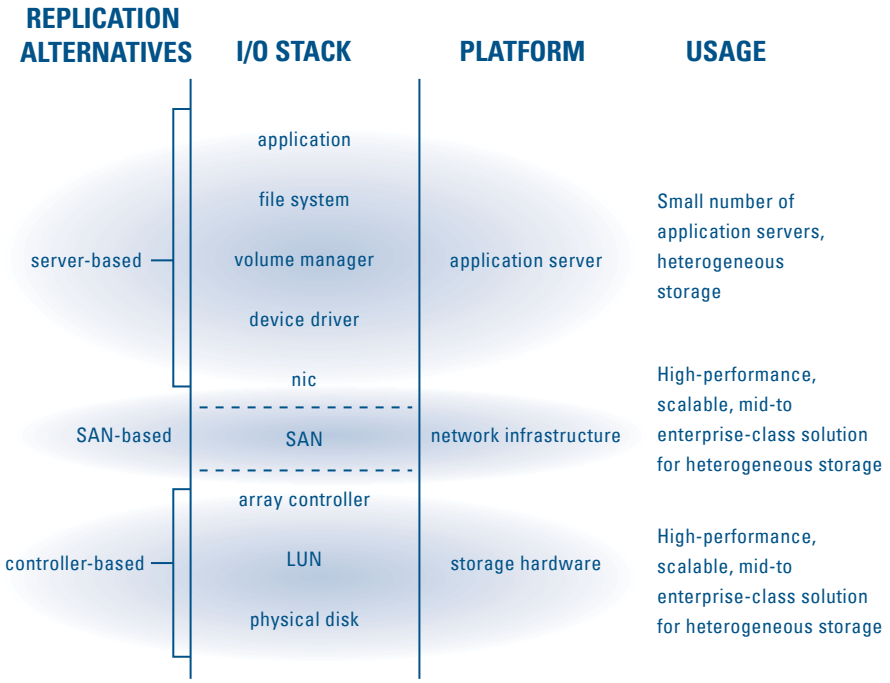

Figure 1: Analyzing data replication alternatives

Server-Based Replication

Server-based data replication, referred to by some vendors as host-based or application-based, is installed on, and operates from, an application server platform. Often running on the same server as the business application, server-based solutions are frequently co-located with volume management products—Volume Replicator from VERITAS offers an example. The replication software intercepts application I/O and duplicates write activity to a secondary volume. Replicated volumes can be mounted to either the application server or to a backup application server located elsewhere on the network. Although server-based solutions consume valuable CPU and memory resources at the application server, they provide the storage administrator with significant flexibility when determining the topology of data replication. Any storage hardware supported by the application server can be configured to receive replicated data.

If the pool of application servers requiring replication is relatively homogenous, a single software product will likely satisfy all data copy needs. However, the more varied the pool of servers becomes, the more likely that multiple replication products will be required to accommodate all operating system software and hardware configurations. Storage administrators, tasked with managing these solutions, must learn the ins and outs of each product and address the inevitable product compatibility issues.

The close integration between existing storage management applications and same-vendor replication software has definite advantages. Replication utilities will probably have a familiar look and feel, and product support will be integrated, speeding resolution of potential problems.

The most significant drawback to server-based replication is limited scalability. Software must be installed and managed at each server with data volumes to be replicated, adding to replication costs and to the administration burden of storage managers. For small-scale configurations, however, server-based replication presents a cost effective approach.

Controller-Based Replication

High-end storage arrays offer a wide range of proprietary storage management functionality that leverages the array controller as an operating environment. Logical Unit Number (LUN) management, performance monitoring and management, and bulk data movement are just some of the functions commonly built in to the storage array. The tight integration between hardware and software allows vendor developers unmatched scope for optimizing the speed and functionality of storage management applications.

All vendors of enterprise class storage arrays now offer data replication as a controller-based—sometimes called disk-based—application. Products include: SRDF from EMC; PPRC from IBM; TrueCopy from Hitachi Data Systems; and ContinuousAccess from HP.

Functionally, controller-based replication applications are identical to their server-based cousins. Each write I/O generated by a business application results in two downstream writes; one write to the primary storage volume and one to a volume defined on a secondary storage array. Controller-based replication products generally insist on like-to-like data copies, which means that the target storage array must be identical to the primary. Although understandable from the perspective of the storage vendor, who has high-performance promises to keep, for IT managers, the single-vendor storage solution invariably results in a more costly replication configuration.

As with server-based replication solutions, deployed alongside volume management software from the same vendor, controller-based replication delivers single-vendor support and common look-and-feel benefits. And it is difficult to beat the performance of an application specifically optimized to the hardware platform. For mission-critical applica-

tions in high-end enterprise infrastructures, where cost is secondary to performance, controller-based replication is an obvious first choice.

SAN-Based Replication

Storage management applications are gradually finding their way into the intelligent storage networks being rapidly deployed across all levels of the enterprise. Whether running on a dedicated network appliance or an intelligent network switch—technology that has been announced but not yet delivered—SAN-based storage management solutions, and data replication in particular, offer tremendous benefits to the storage administrator.

Hosting data replication in the network provides administrators with a huge productivity boost. Rather than manage replication at each application server and storage array, all replication is controlled from a central platform in the network. SAN-based replication is also application server and storage vendor agnostic, offering significant cost advantages when purchasing and deploying a heterogeneous pool of storage devices.

SAN-based replication solutions are available in two distinct topological flavors: in-band —also know as symmetric—and out-of-band—also known as asymmetric.

In-Band

In-band replication solutions place the hardware running replication logic in the path of application I/O. When a host server initiates I/O it passes through the replication appliance. Any write activity generates two downstream I/Os—one to the primary storage volume and one to the secondary.

In-band replication solutions are non-invasive, having no need for a software footprint on the application server. The configuration does, however, add a "hop" in the data path, with implications for performance and scalability.

Out-of-Band

Out-of-band SAN-based replication requires a small software component, known as an agent, to be installed on the application server. The agent communicates with a metadata server in the network to obtain replication control information. After determining the replication configuration the agent issues the appropriate I/O to the primary and secondary storage volumes.

Out-of-band replication solves the latency and scalability issues associated with in-band solutions by eliminating the extra "hop" in the application I/O path. However, the agent software consumes processor, memory, and I/O resources at the application server.

Replication Mode

Whether choosing server-based, controller-based, or SAN-based replication storage managers, application architects must determine what mode of replication will best meet the needs of the application. Replication mode influences the probability of data loss in the event of a failure, and the impact of the replication task on application performance. And because vendors interpret replication mode in slightly different ways, understanding how each product functions, with regard to mode, is essential.

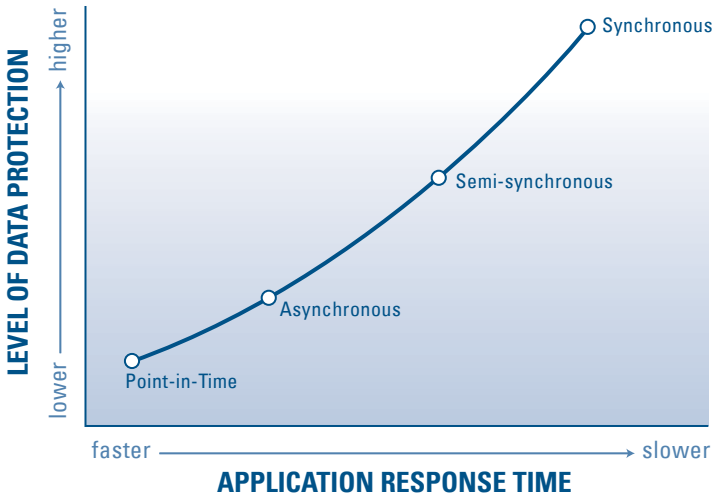

Figure 2: Mapping replication mode to data protection and performance

Point-In-Time

Point-in-time replication generates a snapshot image of a storage volume as it existed at the time the replication utility was run. Database agents are available to work with the application to get a good image in preparation for the replication. The snapshot is replicated to a secondary storage device to create a duplicate image of the primary volume at the remote location.

The impact on application performance from a point-in-time copy is negligible. Although a quiesce of application databases may be necessary, I/O activity between the application server and primary storage device is never impeded. However, in an outage event, point-in-time replication has the highest risk for loss of data. If access to the primary volume is interrupted, data written to this volume since the last snapshot copy was transferred to the secondary storage device will be lost. Storage administrators maintain some control over the amount of data loss incurred, and can increase the frequency of snapshots and make certain that snapshots are run at a transaction boundary, ensuring proper data integrity.

Asynchronous

Asynchronous replication significantly reduces the window for data loss. Each write I/O on the primary storage volume is immediately replicated to a secondary volume, creating a close to real-time duplicate copy at the remote site. The primary and secondary write I/Os are completely independent of each other, eliminating the possibility of a performance drag on the application. This independence introduces the possibility of lost in-flight I/Os during an outage event. In-flight I/Os—those writes successfully committed to the primary volume but not yet written to the secondary—are unlikely to occur conveniently at a transaction boundary. This means that storage administrators must address data integrity problems before the secondary volume data is available to applications.

Unique to asynchronous mirroring is the challenge of ordered writes—keeping the write I/Os in the order they were sent from the application. To ensure data consistency between the local and remote copies, storage administrators can leverage technologies that guarantee ordered writes. As long as the writes are sent in order to the remote location, the data is valid and most transactional applications will recover just as they would with synchronous.

Synchronous

Synchronous replication closes the window on data loss completely. Application write I/Os are simultaneously sent to both primary and secondary storage volumes. The application waits for both writes to be successfully completed before continuing. Should an outage occur, the secondary volume reflects the state of the primary volume at the time of the problem.

Of the available alternatives, synchronous replication has the greatest potential impact on application performance. Excessive latency during primary or secondary I/O will translate directly into longer I/O elapse time for the application.

Variations on Replication Mode

In a bid to improve performance and reduce the risk of data loss, developers of replication software have subjected the basic modes of replication to numerous twists and tweaks. Two common examples include: buffering writes to the secondary replication volume during asynchronous replication to minimize data loss from in-flight I/O; and limiting in-flight I/O to a single write per application, the aim being to reduce the risk of data loss to a single uncommitted write.

Choosing a Data Replication Solution

Establishing the right combination of replication mode and platform requires a detailed understanding of the business application. A business impact analysis can determine how dependent the organization is on an application and how well the application tolerates data loss. The analysis will provide two key pieces of business continuity information: the recovery time objective (RTO); and the recovery point objective (RPO). The RTO indicates how quickly an application must be resurrected after a failure in order to avoid business losses, and is important in the design of failover mechanisms. Of more importance to considerations of data replication, the RPO indicates how much data an application can risk losing before the business begins to suffer losses.

Distance, Data Loss, and Application Performance

Choosing between the principal data replication modes—point-in-time, asynchronous, and synchronous—involves a trade-off between risk of data loss and application performance. Guaranteeing zero data loss imposes the greatest potential impact on application performance. Synchronous replication, the only mode that promises to eliminate data loss completely, telegraphs any latency, from either primary or secondary I/O, directly to the application. The result is slower application I/O response time.

Real-world data replication processes frequently require the geographic separation of primary and secondary volumes. As the distance between the replication platform and the secondary storage volume increases so too does the round-trip latency of the secondary write. And because synchronous replication forces the application to wait for completion of both I/Os before continuing, the speed of the secondary write sets the I/O response time for the entire transaction. Many factors contribute to I/O latency including: propagation delay within the network; the requirements of the network connection protocol—for example, if every send requires acknowledgement, round-trip latency is effectively doubled; and equipment latency, the time it takes for a request to enter and exit each physical device encountered along the path to the remote secondary volume. Each of these factors can be tuned to optimize performance but the physical distance an I/O must travel presents limits that can be difficult to overcome.

If application performance demands prove incompatible with synchronous replication, some risk of data loss must be anticipated. Asynchronous and point-in-time modes reduce the impact of replication on the application and change the decision making

process to one of data timeliness and cost. Asynchronous mode generates a more current copy of the primary volume data, but requires considerable, and continuous, network bandwidth. Point-in-time copy allows flexible scheduling of the data transfer to optimize available network resources, but also introduces the greatest risk of data loss.

Application Server Overhead

Replication mode impacts performance at the microscopic level of individual I/O response times. The replication platform, on the other hand, has just as much impact on the application at the macroscopic level.

Server-based and SAN-based out-of-band replication solutions require software components on the application host that consume processor, memory, and I/O resources. This additional load on the host puts replication in competition with business applications for the limited resources of the server.

Running replication software on the application server also complicates capacity management. An expanding storage pool will increase the amount of server resources needed to manage data replication. Capacity planners must balance the demands of the business application and replication tasks and size the server platform accordingly. Unanticipated growth in demand from either task may require complex and time-consuming upgrades to the server, impacting end user access.

The Cost of Data Replication

Of the many factors that have a bearing on data replication costs, support for heterogeneous storage, licensing fees, and network bandwidth requirements are the most significant.

Replication products that take a vendor agnostic approach to storage can significantly reduce costs by allowing less expensive storage hardware, and redeployed legacy devices, to be used as replication targets. Heterogeneous support also gives IT managers, charged with purchasing and deploying storage, greater leverage when negotiating with vendors. Healthy price competition between multiple vendors vying for an account benefits the customer.

Although controller-based solutions lock IT managers into a single-vendor relationship, replicating between duplicate high-end storage arrays can provide tangential benefits. Close integration between replication and storage platforms often reduces support costs, and gives IT administrators better access to support personnel. A single-vendor solution also reduces the finger pointing that often happens when problems occur with multi-vendor configurations.

Server-based replication solutions often provide the most economic model for small-scale replication tasks. In fact, replication functionality may already be available, at little or no additional cost, from the vendor providing storage management software for the application server. As the number of servers involved in replication grows, however, costs can quickly escalate. Server-based replication requires licensing at each server, and often a duplicate server at the receiving end of the replication.

Keep in mind that vendor solutions that support data compression can significantly offset bandwidth costs. However, even with very high compression ratios, there will always be some expense to transmit information across a network. Both synchronous and asynchronous modes require continuous use of network connections, with enough capacity to support peak write rates. Point-in-time replication, however, is less demanding of the network, and consequently less expensive. The transfer of snapshot copies to a remote location can be scheduled and prioritized to optimize available network bandwidth.

Conclusion

High-speed networks are making it possible to host near real-time copies of production data at what would once have been considered unimaginably distant data centers. Wide availability of high-bandwidth networks, at relatively low-cost, has opened up the market for replicated data within IT. Whether used for content distribution, system consolidation, data migration, or disaster recovery, replication technologies are now supplanting magnetic tape for data movement.

Unlike tape-based data transfer, choosing replication is no longer the end of the decision making process. Myriad choices face the administrator and architect designing an effective data replication configuration. The type of replication (point-in-time, asynchronous, synchronous), platform for replication (server-based, SAN-based, or controller-based), and physical and operational capabilities of the hardware involved must all be carefully correlated to meet the needs of each business application.

In reality, production IT environments support a diverse array of business applications and will require a correspondingly diverse set of data replication solutions. Understanding the needs of each application, and the costs and limitations of each replication approach, is the first step to making an informed decision.

CNT is one of the world's largest providers of comprehensive storage networking solutions. For over 20 years, our experts have analyzed, designed, and built enterprise storage networks.

Visit **www.cnt.com** to learn about our solutions, products, partnerships, career opportunities, and more.

© 2003 by Computer Network Technology Corporation (Nasdaq: CMNT). All rights reserved. Any reproduction of these materials without the prior written consent of CNT is strictly prohibited. CNT, the CNT logo, Channelink, and UltraNet are registered trademarks of Computer Network Technology Corporation. All other trademarks identified herein are the property of their respective owners. CNT is an equal opportunity employer. CNT corporate headquarters' QMS is registered to ISO 9001: 2000. Certificate #006765.

USA: 1-800-638-8324 ▪ Canada: 905-595-1500
UK: 44-1753-792400 ▪ France: 33-1-4130-1212
Australia: 61-2-9540-5486 ▪ Germany: 49-6131-62304-0
Brazil: 55-11-5509-1504 ▪ Japan: 813-5403-4858
Other locations: 1-763-268-6000

PL631 | 0703