

Article Source: <http://www.enterprisestorageforum.com>

May 2002

Farewell to Data Loss: Understanding Data Replication

Data protection and continuous availability are top priorities of IT managers and C-level executives alike, but as companies allocate additional resources to ensuring 24x7 access to data and applications, the question of what strategy and products to employ emerges. With that in mind, this article will help you better understand the role of off-site replication technology in a data protection solution, the difference between synchronous replication and asynchronous replication and how asynchronous replication works.

The Transition

Relying on tape back-up alone is no longer adequate. Despite increases in megabytes-per-minute speed, tape back-up technology has not kept up with the growth in server size causing further increases in restoration times. For example, years ago, with traditional tape-backup technology a typical 50-gigabyte server could be restored in about eight hours at a rate of around 10 MB/min. Even with the best tape-backup available today, restoration for a 1-terabyte server at a 900 MB/min or 52 GB/hour rate will still take approximately 19 hours, meaning the application and server will be unavailable to users during that time. Such a situation could potentially threaten service level agreements (SLA) and customer relationships.

Recognizing the restrictions of only relying on tape back-up, companies today are integrating replication technology to maintain real-time copies of data and applications at one or more off-site locations. A company headquartered in New York, for example, may transfer data to both a New Jersey location that employees can reach relatively quickly and to a rural desert location hundreds of miles away. Validating this evolving protection strategy, analyst firm Gartner Group forecasts that by 2003 75 percent of large enterprises will be combining disk-based data replication and tape-based technology for rapid application recovery.

Technology Alternatives

As companies research and evaluate products to augment their tape-back up strategies with replication technology, it is important to understand the level of protection a business's applications require. If Automated Teller Machine (ATM), stock or other critical transactions cannot be lost under any circumstances, and cost is not a concern, synchronous replication is the route to go. It is vital to note, however, that the "zero loss" associated with synchronous replication is expensive and can reduce overall application performance.

Maintaining "zero loss" through synchronous replication requires a two-phase commit approach where each write to a disk block is written to and acknowledged by the target drive and then written to the source drive, and finally committed to both before any subsequent read or write input/output (I/O), the transfer of information between devices, can be processed by the disk subsystem. The performance penalties that may emerge can become proportionately greater as the distance between the systems increase because communication speed is limited by the speed of light, in the best case. Often times once network protocol and routing latency are factored in, it is much slower.

One alternative to avoid the round-trip delay overhead associated with synchronous mirroring or two-phase commit is to buffer then transmit the changes as fast as available bandwidth allows. Providing the available bandwidth is equal to or greater than the rate of data change, data will be transmitted and applied nearly instantaneously providing "near zero" data loss. With this buffering alternative, if the rate of data change temporarily exceeds available bandwidth, seconds or even minutes of changes could be queued, waiting to be transmitted. Since the changes are still on-site, they could be lost in the event of a disastrous failure. Losing the changes would be impossible with a synchronous system since the transactions would never occur because to allow the mirroring to keep pace everything would have been slowed down to the rate of data transmission.

Enter Asynchronous Replication Technology

Past replication technologies worked either within the application (such as SQL transaction replication), where the level of protection was limited to a single application engine and typically caused overhead to the production environment, or at the hardware layer, which often caused latency to the production disk and/or significant and cost-prohibitive wide area network (WAN) usage. Today, asynchronous replication technologies capture file system changes within the operating system, eradicating the aforementioned challenges.

Asynchronous replication offers the advantage of significantly reducing the need for network bandwidth, particularly for large files such as databases and activity logs. It is also quite flexible, allowing users to continuously replicate only the files or directories they deem business critical and worthy of immediate recovery. This approach significantly reduces the cost of off-site protection, compared with disk mirroring applications that take an "all or nothing" approach.

Asynchronous replication captures changes to any files managed by the server Operating System (OS) at a byte level by installing a File-System Filter Driver, which filters all transactions sent to the file system. Through a few simple rules (e.g. "ignore reads"), the filter driver captures a copy of each transaction and sends it to a system service or daemon. The system service or daemon then transmits it via TCP/IP to the target server.

Specifically, the data flows from the application layer, the software located in Layer 7 of the network, to the File System in virtually "real time." Next, data moves to the hardware as the storage solution is ready to begin transmitting across the network as bandwidth becomes available. Irrespective of which application is creating the data change (i.e. Oracle, SQL, Exchange, Web-Services or File-Sharing), the file system write appears the same when the OS views it. This approach ensures that the data replication is completely independent of the application.

Replication solutions are generally hardware independent, so it does not matter whether a Windows 2000 operating system is storing data on a SAN or its own storage drives. While the asynchronous nature of replication may lead one to think the replicated data is not as current as the production data, this is inaccurate. In many environments, particularly large databases, the I/O demands on the production disk are significantly higher in reads than writes. As a result, the small percentage of write I/O is actually replicated with insignificant latency to the target. However in the limited case that there is a constant flow of writes to the production disks and the amount of actual bytes changed is greater than the bandwidth of the connecting pipe, the replicated data may be a few minutes behind the production data.

Finally, data-integrity and delivery must be highlighted when discussing asynchronous replication technology. To ensure the highest levels of data-integrity it is important to ensure the sequential transmission as well as the asynchronous replication of data. Given that asynchronous replication implies the potential of latency and it is possible to lose a packet across wide area networking, it is absolutely crucial to ensure that packets arrive at the target in the order that they were transmitted. The reason being that if an interruption were caused after only two blocks were written to physical disk, the source could recover from "A-B" (losing C), but the target data may be invalid if only "C-A" exists (losing B).

Nearly all replication technologies use standard networking protocols for delivery, with the one differentiator being whether they use a dedicated "cross-over cable" versus "standard-IP." Cross-over connections require that the servers be in relatively close proximity (same room or building), negating their viability for disaster recovery (DR). Using standard IP networking, in conjunction with the asynchronous replication, however, companies can deploy a cost-effective DR plan that transfers data without geographical limitations to an off-site location.

In conclusion, asynchronous replication technology allows companies of all sizes to cost-effectively protect a wide-array of data and applications, many of which were once left vulnerable due to financial resources and/or distance limitations. While synchronous replication technology is ideal for ATM or stock-related transactions, where "zero loss" regardless of price or performance is required.

Author Bio: David J. Demlow, vice president of product management, joined NSI in March 1997. He is responsible for defining the company's product roadmap and positioning. This includes identifying market opportunities and requirements from existing and potential customers, performing competitive analysis and monitoring overall industry trends and events. Demlow has more than 10 years of product marketing and management experience in the networking and storage industry.