

Clustering and Mirroring.

Processes offer failover and replication capabilities proving critical to business-continuity and disaster-recovery strategies.





IT disasters can cause system downtime or make data unavailable. They can be as small and specific as a hard-drive failure, loss of power to a rack or a cooling system malfunction. They can also be larger in scope, for example a tornado, flood or earthquake.

Even if your year was free of IT disasters, your company's IT systems and cooling facilities still need scheduled downtimes. This may be for routine maintenance, general service and system upgrades. Regardless, these downtimes can impact applications and data.

So how does your company allow key IT data and functions to be available during planned and unplanned outages? And does your firm have a strategy for continuing with business or recovering from disasters?

Server clustering and data mirroring can play an important role in implementing high system, application and data availability. They can also prove cornerstone to an effective Business Continuity (BC) and Disaster Recovery (DR) strategy.

Clustering is done to achieve high-availability, high-performance computing (HPC) or a combination of the two.

"Clustering is the number one solution today for high availability," Lamorena adds. "Fault-tolerance — or duplicate components in a server — is extremely expensive."

Clustering for high availability uses automated failover to other servers in the cluster. A high-availability cluster uses detailed monitoring of an application and all its components, including operating system, hardware, networking and storage.

The clustering software determines when to perform a "failover" by continually checking for each application's "heartbeat" signal. If one system has a problem, a copy of the application on another system in the cluster takes over.

Clustering nodes help provide a highly reliable environment for critical services like databases," comments Bob Zuber, IBM worldwide product marketing manager. "You don't want to take a hit on one of these, because restarting it and applying the redo log can take hours. So businesses look for high redundancy."

Clustering Benefits —

- **High Availability:** Designed to avoid a single point-of-failure
- **Scalability:** Computing power can be increased by adding more processors or computers
- **Manageability:** Appears as a single-system image with a single point-of-control

"Business Continuity means all of your employees can use IT tools whenever and wherever they are with no disruptions," says Chris McCall, director of product marketing, LeftHand Networks.

"Disaster Recovery is the ability for an IT environment to survive catastrophic failure with no data loss or degradation to information," he adds.

Server clustering is all about providing high availability, thereby coinciding nicely with BC efforts. And data mirroring offers a form of data replication and recovery, thereby aligning with both BC and DR efforts.

Server Clustering — Failover Assurance

"Server clustering is about making sure the application is up and running," says Dan Lamorena, senior product marketing manager, Symantec. "Mirroring and data replication are the data availability piece. Combined with clustering, it's the whole solution."

Server clustering uses two or more "nodes" — systems — connected and configured so they are sharing storage.

Depending on how the cluster is configured — active/active or active/passive — the other system(s) may provide load-sharing when the primary system is up.

"An active/active system, with each node doing half the work, can be tricky," notes Zuber. "You have to make sure your utilization doesn't go above 100 percent."

Equally, though, with an active/passive cluster configuration, Zuber points out, "You have to make sure the scripts are done right, so you have access when you failover. You need to have the same driver levels and the software has to be the same."

A cluster may have more than two nodes. For example, a four-node cluster may have three active nodes and one passive node. "Your average clustering product is probably one-to-one with one server doing the work, and the other is a spare," Lamorena says.

"However, that doesn't scale very well," he notes. "Symantec Veritas products offer the ability to have larger clusters, for example, up to 32 servers in a cluster, of which three to four may be spares. You don't need as many spare servers taking up space and power."



Let CDW design, implement, host and maintain a high-availability business-continuity solution for your firm.

Data Mirroring — Storage Replication

Having an application available is only half of the IT requirement. The data applications created and accessed must be equally available in order for business to continue.

Disk mirroring is the recording of redundant data on two partitions of the same disk or two separate disks for fault-tolerant operation.

With mirroring, "The application waits to 'hear' that the mirrored writing is done before it does the next transaction," says Symantec's Lamorena.

"So if one storage device goes down, or something happens to the network, the second one stays up and you've got an offsite mirror with a recovery point objective of zero," Lamorena adds. "That is, no data loss, and applications can obtain the data from the second device."

Mirroring is a key component in the highest level of data protection and disaster recovery. It means there's a guaranteed exact copy of data in two places. Mirroring is different from backups, which are saves of the data state at a point in time, or replication, where a complete volume is replicated, often for use in testing.

Within an enclosure, mirroring is often done using Level 1 Redundant Array of Independent Disks (RAID) features. RAID can be provided through the motherboard or a controller card, or built into a dedicated disk array.

"Mirroring tends to be used in a local or campus environment," Symantec's Lamorena says. "You would mirror the data to two separate Logical Unit Numbers (LUNs) or Storage Area Networks (SANs). The same data being written on the one site would also be written on the other site."

Mirroring Benefits —

- **Protects Against Data Loss:** Added redundancy offers backup in case of hardware failure
- **Disaster Protection:** Offers quick recovery against site- and region-wide incidents
- **Individual Disk Access:** Each disk can be accessed separately for reading purposes

Although mirroring is essential to ensuring high availability of data, it's not a complete data protection solution by itself. "If the data is corrupted, mirroring doesn't help," Lamorena cautions. For example, a virus might scramble data, or a user accidentally delete data. "This is why data protection, in the form of backup, is also necessary."

Clustering and Mirroring Vendors

Many servers include clustering and mirroring features. Third-party software is also available from companies including CA and Symantec. Microsoft offers versions of Windows with support for clustering, including Windows Server 2003 R2 Enterprise Edition — Server Cluster. Windows Server 2008 also includes failover cluster features.

"Almost all of our StorageWorks products, including disk systems, have a complete portfolio of disaster-recovery and business-continuity solutions attached to them," notes Jeff Day, senior marketing manager, HP StorageWorks division.

CA's XOssoft products support protection of a cluster, according to Frank Jablonski, director, product marketing for Recovery Management software, CA. "CA XOssoft Replication will replicate data in real time to another location, and can fail over to that location," says Jablonski.

"We do asynchronous replication, which can be to a different location. Full mirroring tends to be within the array," he adds. "And this typically costs less than full mirroring."

According to the company, CA XOssoft Replication integrates replication, continuous data protection (CDP) and near-instantaneous failover and restoration in the event of loss or disruption. "And our CA XOssoft High Availability provides automatic failover," Jablonski adds.

CA XOssoft features include "push-button" failover and fallback, which simplifies switching between production and replication servers and CA XOssoft Assured Recovery, CA's automated disaster-recovery testing option, which helps IT ensure readiness.

In cases where clustering technology is too expensive, or a business just wants some increased protection without the extra hardware, CA XOsoft software can provide a level of capability through a software-only solution. "XOsoft lets you replicate information to a remote site," notes Jablonski. "If something happens the main site, it will switch you over."

Symantec clustering and mirroring products for high availability and disaster recovery include Storage Foundation, Veritas Cluster Server (VCS), Veritas Storage Foundation HA (which combines Storage Foundation and VCS), and Veritas Volume Replicator, which provides long-distance data replication.

Clustering Challenges

To make certain that all systems or nodes are on the same page, server clustering can have an administrative overhead.

"A clustered environment can be very complicated, especially with immature solutions," Symantec's Lamorena comments. "If IT doesn't have tools that can do basic checks, like whether a patch is applied or isn't made on a spare server, this can cause problems."

Clustering can create other concerns as well. "If you're using a Service-Oriented Architecture (SOA), where applications are working with each other, you want solutions that understand the dependencies," Lamorena says. "A lot of solutions don't let you do much to a cluster once it's up and running."

Clusters and iSCSI SANs

Geographic separation is one way that companies isolate IT facilities from events — weather, earthquakes, local accidents — that could cause downtime.

One technology making it easier to have data mirrors in multiple locations is Internet Small Computer System Interface (iSCSI) according to Mike Karp, senior analyst at industry consulting firm Enterprise Management Associates (EMA).

"Since iSCSI works over Ethernet, it has the ability to do the kind of long-haul data transfer and management that Fibre Channel hadn't," Karp says. "The main value of iSCSI is that it's much easier to implement and IT already understands it."

Companies like HP offer iSCSI in infrastructure and SAN products, making it easier to implement business-continuity and disaster-recovery solutions.

One company with iSCSI-based SANs, that can be used to provide business continuity and disaster recovery, is LeftHand Networks. The firm offers this through its iSCSI-based NSM storage nodes.

"Instead of having software on the server to do mirroring to multiple volumes, we do 'network RAID,' using our embedded storage virtualization to manage multiple copies of the data within a single volume," LeftHand Networks' McCall says.

"This lets us aggregate storage," he adds. "We can control how data is written across nodes, which provides not only load balancing and high performance, but also availability. So we can manage multiple copies of data, which lets us survive a system failure."

This approach "provides cost-effective high availability, because we can put half of a SAN's nodes into a different geographic location," according

to McCall. "Because our products are using iSCSI and Ethernet, there is no constraint on location or distance other than link latency."

Server Clustering and Virtualization

Server virtualization — running one or more "virtual machines" on one system — is finding a place in clustering solutions for business continuity and disaster recovery.

IT managers are implementing clusters within the same physical node, using virtualization. Should a failure occur, applications, that aren't tied to the hardware on which they run, can easily be moved from one virtual server to another or within a cluster.

Virtualization also fits in well in the recovery site, CA's Jablonski points out. "The data replicated has to be a close match, but the physical hardware no longer needs to be a match.

"Virtualization helps bring the cost of that down tremendously by having one piece of hardware, and virtual servers on it," he adds. "You get more flexibility."

"When things go right, virtualization is very simple," says EMA's Karp. However, Karp cautions, "The challenge is, when you back up a virtual machine, there's no granularity, you're backing up the whole thing. So you need backup and recovery tools that understand virtualization."

Advice for IT

Systems that provide data protection and recovery in an hour, day or week are less expensive than ones that deliver business-critical service, where you want to make sure there's close to no downtime.

"From a business standpoint, you have to decide how much business risk and impact can you handle and accept," says HP's Day. "Take a look at all the business functions and processes that are dependent on IT. Then ask, 'What is the impact profile on each of these services if IT goes down?'"

It's important to understand the level of service that will be required. Service Level Agreements (SLAs) aren't about IT, they're the decisions that drive IT decisions. Or if you have a set of equipment, it defines what service levels you can achieve.

"We view business continuity as more broad and high-level, relating to people, processes and technology," comments Symantec's Lamorena. "It's what you need to do in your organization to keep the business up and running. You'll probably want to hire a management consultant to analyze your business processes."

Also, remember that a plan is only as good as its execution, cautions EMA's Karp. "There are many examples of plans that work well on paper, but won't, if, for example, all the key people are out of the building when the problem occurs.

"Or if no one has printed lists of landline and cell numbers, or even something as simple but important as what if the flashlight batteries are dead? And key to disaster recovery is an ability to test your disaster-recovery plan — under conditions that simulate a disaster," he adds. ◇