

Develop a multi-tiered disaster-recovery strategy

By John Bullitt

When IT managers prioritize upcoming projects, implementing or improving disaster-recovery plans is always high on the agenda. However, because of the high cost of monolithic storage systems and shrinking IT budgets most organizations do little more than make minor updates to disaster-recovery plans, if they're able to do anything at all.

This is particularly true for small to mid-sized organizations that cannot afford an expensive storage system, let alone two systems (primary disk and replication target). The alternative is to implement a host-based disaster-recovery plan using replication and application availability software.

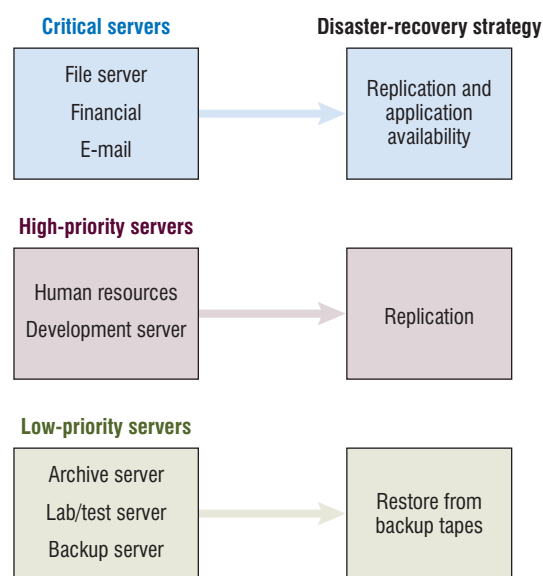
We recommend first reviewing service-level agreements (SLAs) and business needs with disaster recovery in mind. Rather than treating a majority of servers in the same way (e.g., putting those systems in a storage area network [SAN] and replicating from one large disk system to another), companies should look at the particular business needs of individual systems and then logically divide the disaster-recovery strategy based on those needs.

Companies shouldn't look at disaster recovery as a high-cost all-or-nothing purchase. Most companies have a select number of critical systems that must be kept up at all times, a larger number of systems that are high priority but less-critical, and a number of other systems that can go down for longer periods without substantially affecting customers or the company's bottom line.

We recommend establishing several levels of importance for disaster recovery: critical, high priority, and low priority (see figure). By dividing systems into three levels, IT departments can apply host-based software solutions to meet the specific business needs of each level.

Using host-based software for replica-

Establish priority levels



Source: Cambridge Computer

tion and application availability instead of storage-based replication can allow companies to immediately—and more affordably—address disaster-recovery needs. Disaster recovery for less-critical servers can be implemented as IT budgets loosen and at a much lower price point than that of two monolithic storage systems.

BENEFITS OF HOST-BASED SOFTWARE

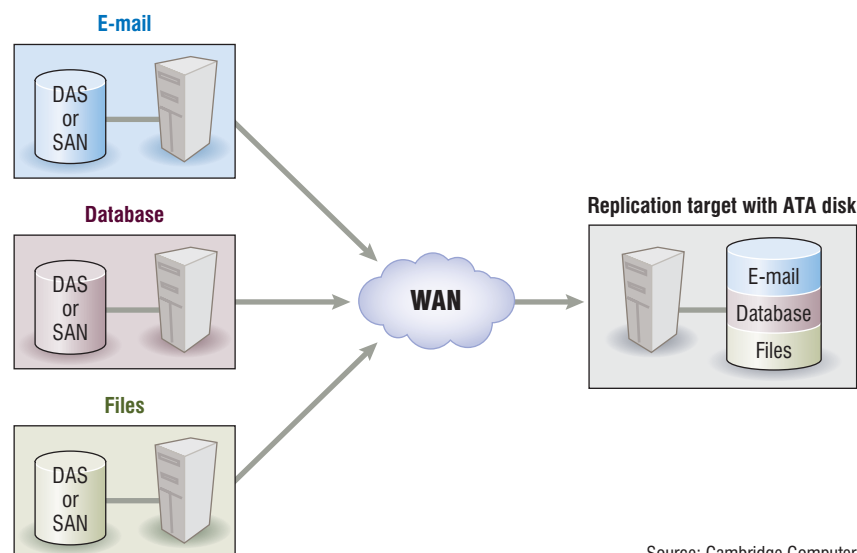
There are a number of benefits to a host-based approach to disaster-recovery management, such as the following:

- **Price savings**—Not only is the initial entry cost for software low enough for most companies to afford but additional savings can be achieved by replicating to less-expensive storage.
- **Tailored solutions**—Companies are able to get a better grasp on their overall disaster-recovery plans and find solutions that are tailored to specific business needs.
- **Immediate action**—Taking the first steps on the road to a comprehensive disaster-recovery plan is better than waiting. Since it is impossible to schedule disasters to correspond with budget availability, it is better for companies to do what they can sooner rather than later.
- **Broader support**—Host-based software gives companies far more freedom in server and storage platforms than do alternative approaches.
- **Flexibility**—Software replication solutions are often more flexible, which means that they lend themselves well to a multi-tiered disaster-recovery strategy.

A variety of host-based replication software products are available, including Fujitsu Softek's TDMF, Legato's Replistor, NSI Software's DoubleTake, and Veritas' Volume Replicator.

Once you determine which product to implement, you need to make three additional decisions: whether to do volume- or

Many-to-one replication



Source: Cambridge Computer

file-level replication; whether to replicate synchronously, asynchronously, or semi-synchronously; and how to design your replication and availability implementation.

FILE-LEVEL VS. VOLUME-LEVEL REPLICATION

A common misperception about host-based replication is that it is less reliable than storage-based alternatives. This probably stems from the fact that many of the cost-effective

replication products are file-system-based, not volume-based.

While host-based file-system replication is fine in many cases, it has some shortcomings when it comes to certain applications and databases. An alternative is host-based volume-level replication. Host-based volume-level replication works similarly to storage-based replication, and the same replication options are supported (i.e., synchronous, asynchronous, and semi-

synchronous). The software monitors activity below the file system and replicates blocks of data as they change.

File-level replication still has some advantages over volume-level replication. First, file-level replication is typically less expensive than block-level replication. Second, because file-level replication monitors the file system itself, administrators can selectively replicate particular files or directories instead of entire volumes. Finally, data from

multiple sources can be replicated to sub-directories of a single directory on one target, which simplifies storage provisioning on the replication target.

ONE TO ONE, MANY TO ONE

The best way for companies to achieve the cost savings of host-based replication is to take advantage of its flexible source and target options. One to one, many to one, many to many, and one to many refer to the number of source servers to target servers. Cost savings can be achieved through establishing the appropriate configuration for each level of disaster-recovery need.

First, critical or high-priority servers can be replicated from the primary location to the disaster-recovery site on a one-to-one basis. While the servers in the disaster-recovery site do not need to be identical to those in the primary site, they should be robust enough to handle the immediate load after a failure. Non-critical systems can use less-powerful hardware, and the load from two applications can run on a single system.

Second, low-priority servers can be replicated on a many-to-one basis (see figure on p. x). Multiple application and file servers in the primary site can replicate data to separate partitions on a large ATA disk array that is attached to a single server at the disaster-recovery site. In a major failure, all data is available in a more accessible form than it would be from tape backup media.

This flexibility ensures that the hardware and software costs for replication can be tailored to a company's specific disaster-recovery needs and budget.

APPLICATION AVAILABILITY

Once you replicate data from one site to another, what do you do next? For critical and high-priority servers, ensuring application availability is the next logical step: It really doesn't do a company any good if it doesn't have access to the data at the disaster-recovery site.

Fail-over products that support multiple operating systems are available from a variety of vendors, including Legato and Veritas. Many of these products also provide a centralized interface for managing larger cluster environments.

By implementing a host-based disaster-recovery solution, IT departments can ensure that their requirements match SLAs and business needs. This type of proactive approach allows companies to implement disaster-recovery plans in phases, in line with budgets. While monolithic storage systems are an excellent foundation for disaster-recovery solutions, host-based replication and application availability software can provide a more flexible and scalable alternative. □

John Bullitt is a senior consultant at Cambridge Computer (www.cambridgecomputer.com) in Waltham, MA.