VERITAS™

A Disaster Recovery Planning Guide for Data Availability

www.veritas.com

Table of Contents

The Evaluation Process

It is imperative that any business today keep its data available at all times in order to be successful. But what happens in case of a disastrous event with the potential of data loss? Is your business disaster-tolerant? To what degree is it disaster-tolerant? There are many threatening events that we tend to forget; they include:

- Computer/Internet crime
- Computer viruses
- Power failures
- Telecom/network failure
- Hardware/software failure
- Human error

Disasters can disrupt any chance of achieving your business goals in seconds and destroy your business' viability altogether. Suddenly, your business is negatively impacted in its operation and logistics, revenue generation, customer confidence and competitiveness, among other areas. According to an article in Contingency Planning, the U.S. Labor Department reports that 43% of the companies that have experienced a disaster never reopen. 29% of them that do reopen, close within two years. Recent problems such as viruses spread over the Internet can bring a company down to its knees within a few hours.

The graph below, from a December 1999 report by Gartner/Dataquest, shows a breakdown of the causes of unplanned downtime.

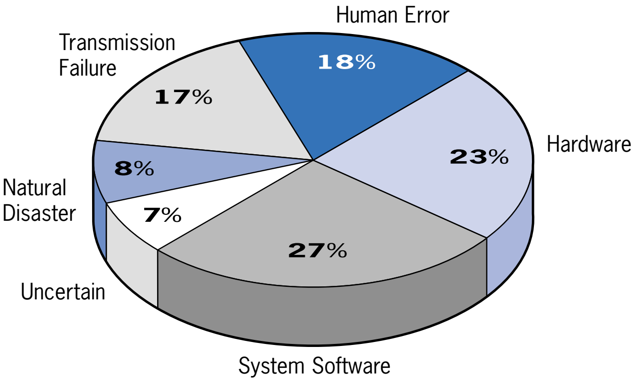

Figure 1: Causes of Unplanned Downtime

Often, when we think of disasters, we automatically bring to mind natural disasters. As you can see, only 8% of unplanned downtime is caused by natural disasters, while the disasters we often forget account for the majority of unplanned downtime: failures of hardware, software, and transmissions, or simple human error. When we think of human error, we often forget common incidents like the database administrator who forgot to add table spaces to the database for a mission-critical application–and suddenly it was too late. How many of these types of disasters are you prepared for?

This guide gives a practical outline of the different technologies available for disaster recovery (DR) of data. It assumes that the reader has some practical experience in data recovery. The guide aims to provide readers with a better understanding of the technologies available today and discusses important questions to consider when putting together a disaster recovery plan for data availability.

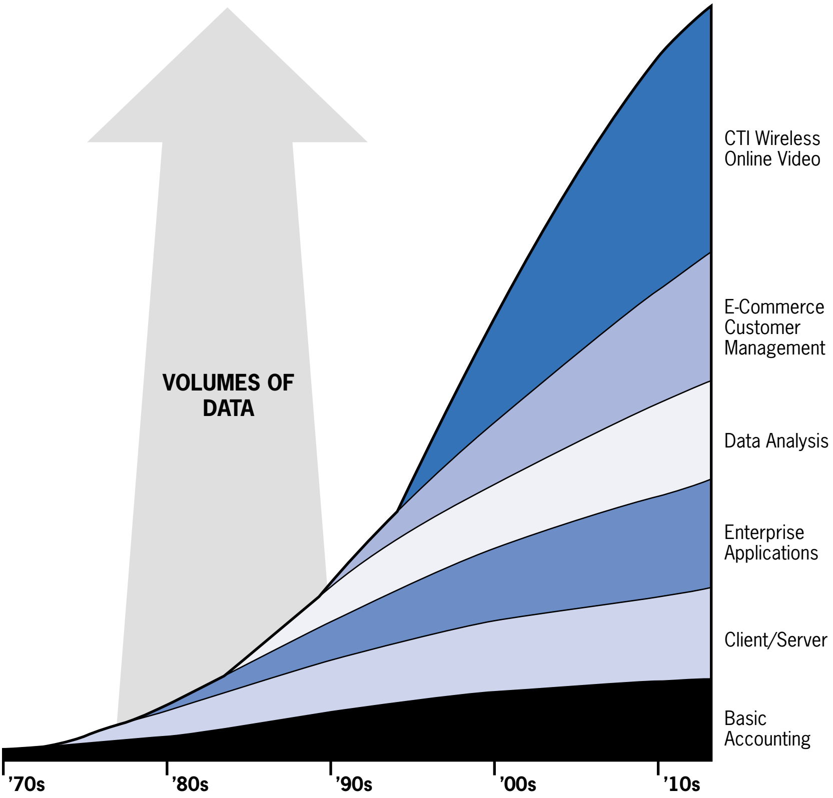

Figure 2: Explosion in Data Volumes

All That Data!

With the Internet boom, the amount of data continues to grow exponentially. The days where all of a company's data could be managed by a few machines are gone. Today's data is not only growing at a fast rate, but it is also spread throughout organizations. In addition to being able to manage the data, it is critical to have a disaster recovery plan that continues to adapt to the changing conditions of your operation. You need to have the right tools so your plan can keep protecting your data as your business and computing environments develop.

Objectives of a Plan

The objectives of a disaster recovery plan for data availability are the same fundamental objectives as those of any disaster recovery plan.

- Minimize any economic loss
- Ensure a safe recovery
- Minimize any decision making during the recovery plan process
- Reduce the reliance on key individuals

One of the main questions to ask is, what is your service level agreement for disaster? Is it within 48 hours? Is it within 10 minutes? What is it for mission-critical applications compared to non-mission-critical applications? Many organizations today are faced with strict government regulations, particularly in the financial, banking, and healthcare industries. Strategically planning the right disaster recovery solution for high availability can provide companies with an affordable way to meet their service level agreements, comply with government regulations, and minimize their business risk.

Starting a Plan

One of the first steps in DR planning is determining what data assets you have and where they reside. Determine which applications are your most mission-critical and the interruption of which could impact your business most detrimentally. According to a report by Contingency Planning Research, this is the expected financial loss per hour due to unplanned downtime.

Application	Hourly Cost
Retail Brokerage	$6,450,000
Credit Card Sales	$2,600,000
Pay-per-View	$150,000
Home Shopping	$113,000
Catalog Sales	$90,000
Airline Reservations	$89,500

Critical factors in this are planning on the right amount of protection while evaluating and balancing the risk. It is also important to have tools that are flexible enough to adapt to changing needs. Today it is not uncommon for companies to go through drastic changes in their computing environments in just a few months as they accommodate rapid growth. Studies show that business planning cycles have gone from about three to five years to less than a year today. Businesses regularly need to reassess their needs to ensure they are still disaster-tolerant.

In today's world, your web application or ERP system can be one of your most mission-critical applications. However, the most obvious is not necessarily the most mission-critical application in an organization. For example, the salesperson's laptop in a remote office can also contain critical data for your organization. You need to rate how vital your applications are, and determine both the amounts of downtime and data loss that are tolerable. What degree of data loss and how much downtime can you afford without irreparably harming the business?

Often, then, you can look at the amount of tolerable downtime as a good indicator of the type of protection needed. In terms of data availability, you need to determine whether days, a few hours, a few minutes, or only a few seconds are tolerable. According to a Gartner/Dataquest report in December 1999, the top three mission-critical applications are manufacturing, financials, and customer relationship management.

After you've determined what and where your most mission-critical applications are, you must assess the types of disaster you need to protect your data from. A disaster recovery plan often involves several layers. Although this is one of the most critical elements, it is often difficult to determine the most appropriate type of protection.

In creating a disaster recovery plan you must also ask whether there are any single points of failure. There are two reasons why this may be helpful. One is that single points of failure, whatever they are, will sooner or later do just that: they fail. The other is that evaluating this helps you plan for disastrous events you might ordinarily have forgotten.

Disaster recovery planning is the creation of a comprehensive plan for the event of any kind of disaster that can occur to a company. Generally it is true that the more protection, the less the risk, but costs go up as well. Risk assessment and effective protection need to be evaluated as they relate to cost. The highest investments may not necessarily bring the most significant returns for businesses in need of protection. Where and how, then, should disaster recovery planning start?

The Basics

The foundation for any disaster recovery plan for data availability is to ensure you have a backup copy of your data. This holds true as a basic disaster recovery plan for any system–offline storage. Offline storage offers basic recovery and protection of data, but does not necessarily provide for immediate or instantaneous recovery from such things as site disasters.

At a minimum, a copy of all data should be stored offsite for recovery purposes. This protects you in the event of any disaster, including one that halts operations at an entire site. Another copy of your data is available in form of backup tapes that can be used for recovery.

But what if you didn’t experience a site disaster, and all you needed was the backup tape for one of your systems, say, because of a minor problem like a disk or application failure? Obviously, getting that recovery tape from your offsite location can take more time. Ideally, you should make sure that your backup and recovery system can produce duplicate copies of your tapes for onsite and offsite storage. This way you are prepared to recover the data more efficiently by storing the latest copy onsite for minor disasters as well as offsite for site disasters. It is also important to have management capabilities to manage your tapes offsite. This is often called electronic vaulting.

In addition to the basic requirements of duplicate tapes and offsite management, another backup and recovery key feature is the ability to de-multiplex tapes. Most backup and recovery systems use multiplexing to back up several systems in a shorter amount of time by streaming multiple data streams to one tape device. De-multiplexing tapes during the duplication process helps to produce a more condensed version of your backup tape for faster recovery. The de-multiplexing process should be able to also produce non-proprietary format tapes, like gnu/tar for UNIX systems. For disaster recovery purposes, this is important because it enables the restoration of your files without installing the backup and recovery application first.

In the event of any disaster, it is important to be up and operational as quickly as possible. This includes making sure your plan eliminates any activities you can do without. For example, in case of a host disaster, usually you must first install your operating system and then all of your applications. The optimal backup solution should include what is known as a "bare metal restore." "Bare metal restore" reduces the number of recovery steps by automating both the operating system and data recovery in one process. Usually, it takes much time to recover a system when the operator needs to manually go through each part of the reinstallation process. System parameters and other logistics take time to set up. However, when your application has a "bare metal restore" capability, all of the system settings are also backed up and quickly restored. This will noticeably accelerate the recovery process.

Beyond Backup

Backup provides users with the basic recovery tools for data availability. Before exploring more far-reaching disaster recovery tools, what other basics should one consider?

Easier tools for managing your data are essential. Tools that give you a higher degree of manageability will be very helpful in preventing disasters. They also make you less dependent on a single resource, and require less of a learning curve for the data managers.

For example, let's consider the file system. Are there features and tools included in your file system that can help with the disaster prevention or recovery of your data? A file system is the fundamental element of any information management system. File systems should enable online file system resizing with uninterrupted data access. They should also support more effective data storage management through capabilities such as compatibility with any disk storage hardware platform. In addition to enhancing data availability on heterogeneous hardware platforms, your file system should have an easy-to-use interface, keeping the amount of needed training low.

Remember the disaster where the DBA forgot to add table spaces for the database of a mission-critical application? What data management tools could have helped to avoid this?

While there is no perfect safeguard against the mistakes we make, a software solution like a volume manager can help the system administrator monitor the data growth in this example. Volume managers can allow online storage reconfiguration and database table space growth; they may also come with proactive monitoring and maintenance tools that help you prevent disasters. It's important to be able to perform these tasks without disrupting users and incurring downtime.

Volume managers can also be used to help set up mirroring of your data between multiple disk arrays or storage systems. Mirroring is often done in a local environment, where it is relatively inexpensive and the performance impacts are minimal. Mirroring protects against data loss due to, for example, disk failure. A volume manager should be able to mirror data across multiple disks. With more disks holding redundant data, the likelihood of surviving a failure increases. This protects against an easy-to-pinpoint failure within a site. If an entire site stops operating in a disastrous event, replication technology, discussed below, can be used to continue production.

Beyond the Basics

As mentioned, one of the keys for DR planning is determining what are your most mission-critical applications, the ones that would cause the most significant losses if they were unavailable.

Most mission-critical business applications rely on large databases, the information in which is at the core of what allows the business to function. An interruption to their availability is not acceptable. Traditional backup and recovery methods are often not fast enough to handle the recovery of data used by mission-critical applications. That's why many companies today implement data replication to keep an up-to-date "hot" standby of their most critical data for immediate point of failure recovery. Data replication enables a company to recover mission-critical applications within hours or minutes.

One of the most important aspects of replication is data integrity. Replicated data needs to be consistent, up-to-date, and ready to use at a moment's notice, but also must be transparent to the application. The replication technology should be seamless, so that the application data can be sent from one primary site to multiple secondary sites, or replicated from many primary sites to a secondary location where the data is warehoused. Replicating data to multiple secondary sites is advisable because it offers greater protection and doesn't leave room for a single point of failure, like a network connection.

There are two main methods in which data is replicated, **synchronous** or **asynchronous**. Synchronous data replication writes the data to the replication target system before the write operation even completes on the host system. At any time, data on the target will be in the exact same state as on the primary host. Synchronous replication may cause performance delays on the source system, especially if there is a slow network connection. Asynchronous data replication yields maximum performance with lower overhead than synchronous data replication, but tends to have a small lag in data currency. Journaling can be used to make sure the data are written in the same order on the target as they are on the source, which helps to maintain data integrity. There are also replication solutions that can switch between synchronous and asynchronous, depending on network conditions and the resolution of communications problems.

Many replication technologies depend on a dedicated network. However, this opens up a single point of failure. Replication should not depend on a dedicated network, nor any vendor-specific storage hardware platform. Hardware limitations can lock an organization into a vendor-specific framework without the needed scalability and flexibility for managing growth. Vendor-specific hardware configurations for replication are also often very expensive.

One of the most common and well-tested ways of replicating data between two remote sites, besides replicating the primary data between multiple sites, is to add three-way mirroring to a secondary or replicated site. This way you have added protection if your primary site is hit by a site disaster and you are relying on the availability of your secondary site. The second site would be protected from a disk failure, for example, by having two other mirrored copies of your data. The third copy of the mirror can also be used as a break-off of the replicated data to provide a point-in-time representation of data, potentially a very useful additional layer of protection. For example, one company used the third copy of mirrored data to protect themselves against the “ILOVEYOU” virus. They were able, when the virus reached them, to roll back and use their third copy of mirrored data, which was a point-in-time copy of the data before the system was struck.

Data replication is best used for applications that can only tolerate a few minutes of downtime. Inexpensive local data replication can be realized as far apart as 10 kilometers, using fibre channel technology. Data replication can be used for both LAN and WAN environments, depending on the type of protection needed. However, it is primarily used in WAN environments. Many companies replicate their data to multiple sites.

Optimal Recovery

For optimal data protection of a mission-critical application, replication combined with clustering gives the best protection for any disaster.

Clustering is defined as two or more host machines sharing one or more storage resources. Each host is referred to as a node in the cluster. The nodes are connected over a network. If one node fails, the other node(s) can still access the disks or other storage.

Clustering enables applications to fail over from one node in a cluster to another node in the same cluster. This way, mission-critical applications can be restarted after failing over without interruption for the application service. With fibre channel technology, machines as far as 10km apart in a cluster can fail over to each other and protect against local disasters.

Even the simplest cluster offers protection against local disasters, like a machine crash, with immediate recovery. With an application installed on both nodes of a cluster and the data on the shared resource (disks), failover from one node to the other is seamless and data integrity is preserved. However, if the shared resource fails, this again demonstrates the vulnerability operations incur from a single point of failure. If the data is not replicated, recovery will not be immediate, and administrators will have to use backup tapes for recovery.

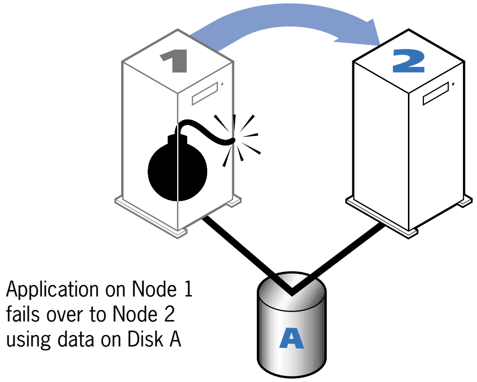
1
2
A
Application on Node 1 fails over to Node 2 using data on Disk A

Figure 3: Application Failover between Two Nodes

The complete solution is to combine replication or mirroring with clustering. Applications should be installed on both nodes with data replicated or mirrored on each node. If one node fails, the other node can restart the application with the primary data. However, if the primary data disk on one node fails, the other node can also restart the application with the replicated or mirrored data. The combination of replication or mirroring with clustering offers stronger protection in the event of a disaster and doesn't leave room for a single point of failure. It also offers an immediate recovery of applications.

To bridge large distances, replication is preferable to simple disk mirroring. Mirroring requires more overhead and is better suited to confined local environments.

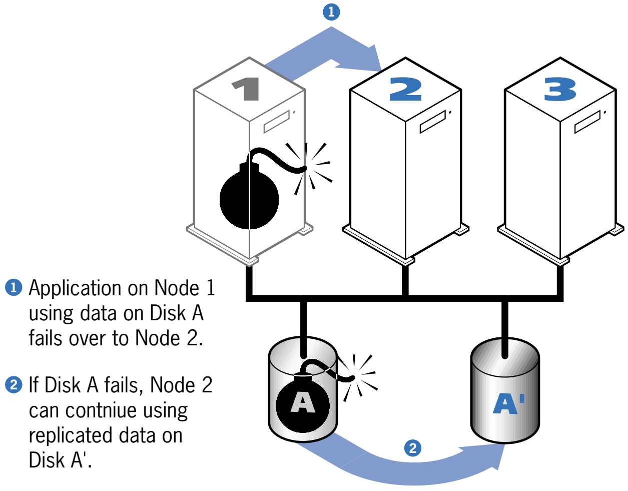
1
1
2
3
A
A'
2
❶ Application on Node 1 using data on Disk A fails over to Node 2.
❷ If Disk A fails, Node 2 can contniue using replicated data on Disk A'.

Figure 4: Failover Using Clustering and Replication

Clustering technology has evolved in recent years. With some clustering packages, failover can only occur between 2-4 nodes in a cluster. It is important to take note of the number of machines that the clustering solution can support. The more nodes available for failover, the better.

Planners need to consider how the failover can occur. Can it have different applications running on one node fail over and restart on different nodes in the cluster? What if the secondary node fails after the application has failed over from the primary node? Can the secondary node fail over to a third node? These scenarios are called, respectively, **distributed** and **cascading** failover.

The rules for failover from one application to another depend on the business requirements of the application. Configurable policies need to be defined in the clustering software so that they can dynamically determine where and how after a failover recovery should occur. Failover rules should determine both policy- and resource-based failover strategies.

Using configurable policies, distributed failover is defined as the ability to fail over specific applications on one node to many other nodes. Imagine a primary node that is running ftp, an Oracle application, and a Sybase application. Suddenly the machine crashes. The optimal failover scenario would be for ftp to fail over and restart on one node, the Oracle application to fail over and restart on another node, and the Sybase application to fail over and restart to yet another node. This is known as distributed clustering. Since it does not force all applications to fail over to the same node, it allows for a flexible, logical disaster recovery.

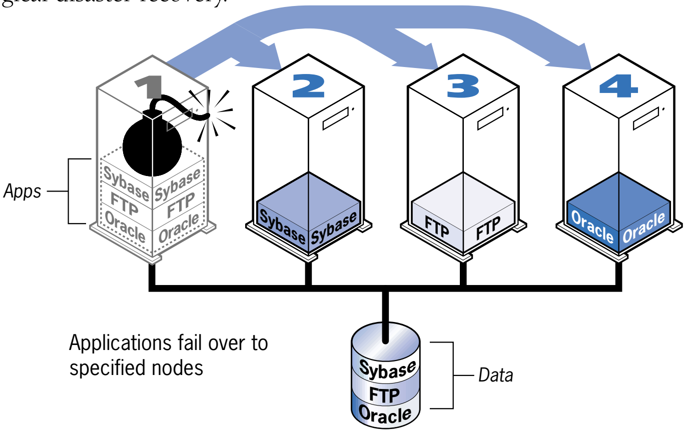

Figure 5: Distributed Failover

Cascading failover in clustering is another way of adding more protection by ensuring automatic recovery from consecutive failures. For example, if a primary node has failed over to a secondary node, and then the secondary node fails as well, a third node should be able to recover and restart the application that failed on the secondary node. This is different from distributed failover, because it supports recovery from failures after the initial one.

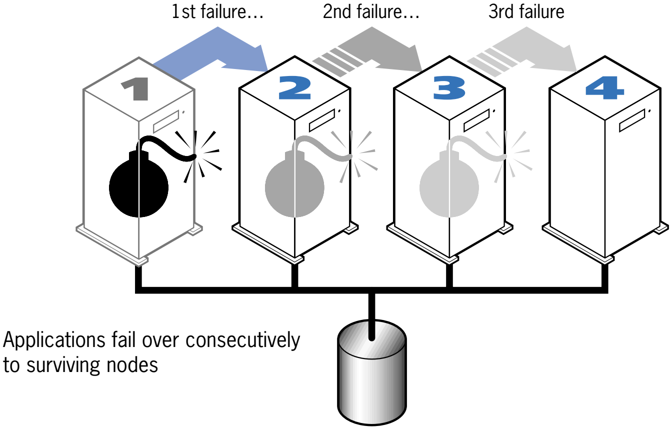

Figure 6: Cascading Failover

All of the clustering solutions mentioned above protect a local site environment, which is where you typically find single clusters. However, for a true WAN solution that can offer immediate recovery and recovery from a complete site disaster, failover capabilities between clusters and replicated data over a WAN is a must. One of the newest clustering technologies is the ability to manage multiple clusters. Management of multiple clusters adds functionality with increased failover capabilities so that a cluster at one site can fail over to another cluster elsewhere. A company can ensure almost immediate recovery of its mission-critical application by having a replicated site anywhere in the world, with clustering capabilities that allow a node on a cluster at one site to fail over to another node on the other clustered site with the replicated data. This yields optimal data protection for total site disasters. Business requirements will need to determine the way in which multiple clusters react to various scenarios, just as they do in a single cluster environment. Fully scalable, global data protection is the optimal protection for any enterprise.

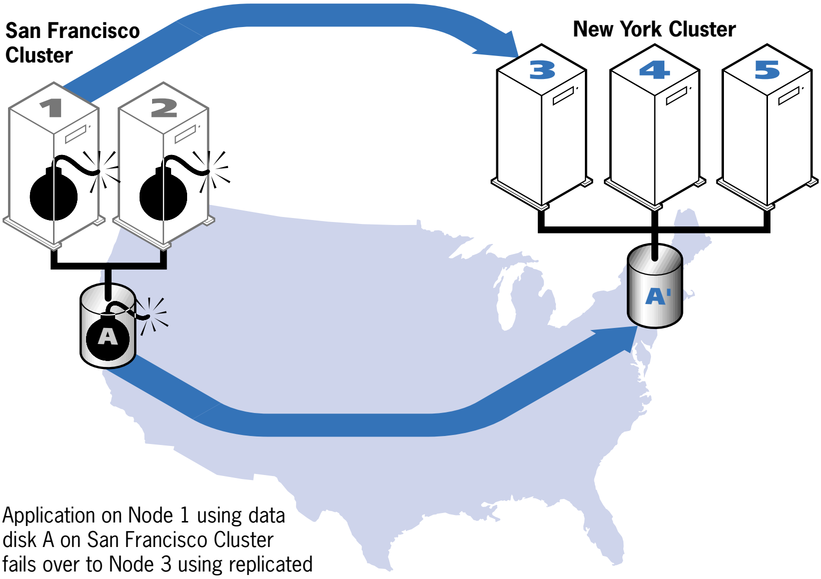

Application on Node 1 using data disk A on San Francisco Cluster fails over to Node 3 using replicated data disk A' on New York Cluster

Figure 7: Global Cluster Failover

Pulling It All Together

Building a disaster recovery plan for data availability is certainly no easy task. It is important to ensure that the plan needs to have a solid framework that enables growth. The phenomenal data growth in our business environments increases the urgency of reviewing and adjusting DR plans to ensure protection for all of the data, no matter what happens. Your disaster recovery plan needs to be flexible enough to evolve with your business.

Your DR plan will involve much thorough planning. Disaster recovery planning needs to encompass people, processes, and technology. Hiring the right consultants that understand the issues and technology involved in disaster recovery and high availability is important for the development of a good disaster recovery plan. When you evaluate consultants, in addition to references, you may want to learn about their:

- Certification for disaster recovery planning
- Experience with proven solutions
- Expertise in system and data availability
- Understanding of complex system and application environments

Hiring the right consultants to assist with your disaster recovery plan can help you to:

- Minimize the complexity of disaster recovery procedures
- Reduce technical infrastructure costs
- Increase effectiveness and efficiency of the plan
- Accelerate implementation

Consultants should be able to help with your planning from assessment through implementation and testing.

Assessment by the consultant should be at the beginning of DR plan design. A business impact analysis assesses the impact of a disruption to your organization's operations. It identifies critical business functions; defines acceptable recovery time frames; and determines recovery priorities. By identifying the costs of a disaster, this assessment also allows you to minimize impacts with prevention and mitigation measures.

The next step is the design process. Once you have determined your business requirements, a recovery strategy should be designed to achieve your goals. To be successful, a strategy needs to address all types of risk. It should be reliable and cost-effective and should make intelligent use of your resources.

As the DR plan develops further, your strategies become actions. Practice and implementation helps your organization understand everybody's roles and responsibilities, and should result in comprehensive recovery procedures that fully meet your needs. Plan testing plays an important role in recovery by assuring the accuracy and completeness of procedures and by building awareness and familiarity in recovery team members.

Based on a full survey of your business requirements, the DR plan is likely to include the right infrastructure to help your business remain viable in the event of any potential disaster. That includes a projection of business developments into the future, so you arrive at a plan that will grow with your organization. The consultants can play a significant role in implementing and testing your disaster recovery plan, then rehearsing it with your staff. Without all of these steps, your plan will be unpredictable and incomplete.

Disaster Recovery Solutions from the Data Availability Leader

VERITAS Software provides essential data availability software solutions that enable customers to protect and access their business-critical data for Business Without Interruption. As a trusted supplier of heterogeneous data availability software including solutions for file and volume management, data protection, and clustering, VERITAS Software ensures integrated, out-of-the-box solutions by fostering premier strategic partnerships with all major industry drivers, including industry pacesetters from the device level to high-end UNIX vendors. VERITAS Software delivers continuous availability of business-critical information to companies of all sizes, from some of the world's best known blue chip corporations to business managers and IT professionals.

Visit our website at **www.veritas.com** for more detailed information on the products and services listed below, as well as other products that bring data availability to your business.

VERITAS Foundation Products

VERITAS File System™ for UNIX

VERITAS File System for UNIX is a high-performance, quick-recovery, and standards-compliant file system. VERITAS File System for UNIX augments UNIX file management with high availability, increased bandwidth, and up-to-date reliable structural integrity. VERITAS File System supports fast recovery following a system crash or reboot. The system completes a file system check (fsck) in seconds, regardless of file system size. In addition, VERITAS File System allows online backup, online resizing, and online defragmentation. VERITAS File System is tuned to allocate contiguous disk space for every file, yielding superior performance over traditional file systems.

VERITAS Volume Manager™ for UNIX

VERITAS Volume Manager for UNIX enables easy-to-use online storage management for enterprise computing environments. Through the support of RAID redundancy techniques, VERITAS Volume Manager protects against disk and hardware failures, while being powerful and flexible enough to extend the capabilities of existing hardware. By creating a logical volume management layer, VERITAS Volume Manager overcomes the physical restrictions imposed by hardware disk devices; for example, by allowing volumes to span multiple spindles.

VERITAS Volume Manager™ for Windows NT
VERITAS Volume Manager™ for Windows 2000

Until now, Windows systems have lacked the sophisticated storage management tools of their enterprise counterparts. VERITAS Volume Manager for Windows NT and Windows 2000 provide easy-to-use, online disk and storage management for the enterprise customer. They allow storage to remain online and available, and feature flexible performance monitoring and optimization, and centralized management from a single, MMC-compliant GUI. Microsoft selected VERITAS to develop disk management software for Windows 2000. The Volume Manager products are built on this foundation and enhance the basic capabilities offered in Windows 2000. This combination of availability, performance optimization, and ease of use make Volume Manager the premier volume management solution for the Windows environment.

VERITAS Backup Products

VERITAS NetBackup™

VERITAS NetBackup delivers enterprise backup and recovery for organizations that require a high-performance, flexible, and easily managed solution. VERITAS NetBackup's central management console and highly scalable architecture enable the backup and recovery solution to continually meet the changing need of the modern IT environment. By ensuring the reliability of backup images, VERITAS NetBackup becomes a critical part of any disaster recovery solution for data availability. NetBackup supports a wide variety of platforms, including all of the major UNIX systems, Windows NT, Novell NetWare, PC workstations, and VMS. VERITAS NetBackup can be used as a centralized solution for large heterogeneous environments scaling from a few to many thousands of nodes, all managed from a single console. VERITAS NetBackup, the chosen backup and recovery solution for many Global 2000 organizations, is also offered to enterprise customers by major hardware vendors such as Sun, Hewlett Packard, and Bull S.A.

VERITAS NetBackup Business Server™

VERITAS NetBackup BusinesServer is an easy-to-use solution that provides a simple way to protect valuable data stored on a small number of UNIX and NT systems. Derived from VERITAS Software's industry-leading NetBackup, BusinesServer offers customers the ability to configure, monitor, and control backup and recovery jobs from a familiar-looking Java GUI interface. An ideal solution for workgroups, remote organizations, or growing businesses, BusinesServer can easily be upgraded to NetBackup, a fully automated, policy-driven, company-wide storage management solution.

VERITAS NetBackup Professional™

VERITAS NetBackup Professional brings automated data protection to desktops and mobile laptops by delivering transparent backup and comprehensive recovery capabilities for Windows clients. It is an easy-to-deploy, easy-to-manage solution that lets customers back up a large number of desktops and laptops using the small hardware capacity. VERITAS NetBackup Professional is integrated with Microsoft Explorer, so it uses an intuitive client interface for recovering individual files and directories. VERITAS NetBackup Professional can recover the complete hard drive image via bootable CDs, balance backup workloads using a common console that supports multiple servers, smoothly move clients from one server to another, and easily copy profiles between servers.

VERITAS Backup Exec™ Desktop Edition

VERITAS Backup Exec Desktop Edition is the solution of choice for reliable, automated Windows 98, Windows 95, and Windows NT Workstation 4.0 data protection. Integrated Emergency Recovery rebuilds your entire system without reinstalling the operating system or the backup software. Peer-to-peer network support protects data on a small network. You can schedule unattended backups with an advanced scheduling system. Backup Exec supports most tape devices; CD-R and CD-RW; DVD-RAM; Zip and Jaz drives; magneto optical; PD/CD; SuperDisk; and even floppy disks, hard disks, and network drives mapped as a device using the File Specification feature.

VERITAS Backup Exec for Windows NT/2000

VERITAS Backup Exec is the de facto industry-standard backup solution with 100% compatibility for all Windows NT and Windows 2000 environments. It makes fully reliable backup and recovery of Windows 2000 Active Directory, Distributed File System, and the System State components possible. Built-in virus protection and an intuitive user interface make for a reliable and easy-to-use data protection solution. Combined with the many available options, Backup Exec is a complete scalable solution for any size network and any level user. Built-in wizards, Agent Accelerator technology, and Advanced Device and Media Management (ADAMM) deliver the ease of use, high performance and flexibility to manage data protection effectively with minimal administrative overhead.

VERITAS Backup Exec for NetWare

VERITAS Backup Exec for NetWare v8.5, with NDS support, is the latest backup and restore technology for quickly, reliably, and simply protecting server data. The Multi-Server Edition includes agents for Windows NT, Macintosh, UNIX, and Linux servers. Intelligent Disaster Recovery™ and Open File backup options allow for quick point-in-time recovery and protection of data in use on both local and remote NetWare servers. Add the Shared Storage Option (SSO) to create storage area networks, share tape libraries, and protect NetWare clusters. Features for creating duplicate archive tapes, simplified holiday scheduling, ARCserve tape read, and online tape audit make Backup Exec the complete backup solution for NetWare operations.

VERITAS Replication Products

VERITAS Volume Replicator™

VERITAS Volume Replicator provides the foundation for seamless wide area availability, site migration, and disaster recovery. Whether motivated by natural disaster, site failure, or simply a planned site migration, the ability to seamlessly move computing operations with minimal (or even zero) downtime has become a necessity to compete in today's market. Based on the defacto industry-standard VERITAS Volume Manager, the VERITAS Volume Replicator reliably, efficiently, and consistently mirrors data to remote locations over any IP network for maximum business continuity. Volume Replicator provides an elegant, flexible, storage-independent solution to deliver true disaster recovery when data currency and availability are paramount.

VERITAS File Replicator™

VERITAS File Replicator is the perfect tool to keep multiple Web and/or file servers fully synchronized and mounted for maximum availability and load balancing. With VERITAS File Replicator, all replicated file systems are mounted and available for use at secondary systems at all times. Fully synchronous replication ensures that all participating nodes are totally up-to-date at all times. VERITAS File Replicator is the only product of its kind that can offer true bi-directional replication. This means that in a two-node configuration, both systems could both read and write to the replicated data set with complete integrity. The File Replicator can also be configured to have a single master node with many read-only targets.

VERITAS Storage Replicator™

VERITAS Storage Replicator delivers robust, reliable data replication to Windows NT environments. Whether the requirement is for automated data distribution, disaster protection, or many-to-one backup centralization, VERITAS Storage Replicator will meet the need for even the most demanding replication jobs on the Windows NT platform. The need to have files identically replicated to multiple systems is found in web load balancing environments, software distribution, or simply file servers that must be accessed from different physical locations. And for sites that cannot tolerate the data loss or availability restrictions inherent in a traditional tape backup-and-restore scenario, a real-time replication solution like Storage Replicator is a perfect tool to ensure maximum data availability even if an entire site is inaccessible.

VERITAS Clustering Products

VERITAS Cluster Server™

VERITAS Cluster Server (VCS) is an architecture-independent, heterogeneous, availability-management solution focused on proactive management of service groups (application services). It is equally applicable in simple shared disk, shared nothing, or SAN configurations of up to 32 nodes. Policy-based, cascading and multi-directional application failover is supported, and application services can be manually migrated to alternate nodes for maintenance or load-balancing purposes. VCS provides a comprehensive availability management solution designed to minimize both planned and unplanned downtime.

VERITAS Global Cluster Manager™

VERITAS Global Cluster Manager (GCM) is an availability management solution focused on proactive management of the distributed clustered computing environment. Global Cluster Manager leverages existing clustering and replication technology in the enterprise environment to minimize both planned and unplanned downtime. Building on the existing enterprise infrastructure, it extends to customers an availability management growth path for handling one or more site failures or disasters.

VERITAS ClusterX™ for MSCS

VERITAS ClusterX for MSCS is a software tool designed specifically for Microsoft clusters. ClusterX co-exists with and extends the functionality of, Microsoft Cluster Service (MSCS) to increase availability Microsoft clusters and clustered applications in Windows 2000 and Windows NT 4.0 environments. Through powerful functionality, ClusterX makes it easy to install clustered applications, configure clusters and clustered applications, and perform ongoing administration and management.

VERITAS Enterprise Consulting Services

VERITAS Disaster Recovery/Business Continuity Services

A market research firm, the META Group, reports that the cycle time for implementing new business processes has gone from five to seven years in the ’70s and ’80s to less than 12 months today. Through delivery of a comprehensive mix of technical expertise, industry knowledge, and experience, VERITAS Enterprise Consulting helps ensure rapid implementation of your disaster recovery plan. Our highly qualified consultants offer services ranging from needs assessment and planning to installation, implementation, troubleshooting, and custom application design.

VERITAS Software Corporation
Corporate Headquarters
1600 Plymouth Street
Mountain View, CA 94043
650-335-8000

For additional information about VERITAS, its products, or the location of an office near you, please call our corporate headquarters or visit our Web site at www.veritas.com

90-00341-399 • VRTS07-DRPG-0000