WHITE PAPER: ENTERPRISE SOLUTIONS

Business Continuity:

Symantec Intelligent Application Recovery Solutions Guide

White Paper: Enterprise Solutions

Symantec Intelligent Application Recovery Solutions Guide

Contents

Introduction

Business continuity and disaster recovery (BC/DR) refers to the capability to restore normal (or near-normal) business operations, from a critical business application perspective, after the occurrence of a disaster that interrupts business operations. It requires the ability to bring up mission-critical applications and the data these applications depend on and make them available to users as quickly as business requirements dictate. In cases where downtime is costly, the process will likely involve automation. In other cases, manual processes may be appropriate. For mission-critical applications that demand minimal downtime, the disaster recovery process must be highly automated and resilient—such applications require an intelligent application recovery infrastructure. This document will describe the clustering and replication technologies offered by Symantec that, when used together, deliver intelligent application recovery as part of an overall BC/DR infrastructure. This document also offers guidelines and best practices around intelligent application recovery to ensure the right architecture is matched to the business requirements at hand.

Disaster recovery, in the broad sense, encompasses much more than just recovery of Information Technology (IT) systems and services (for example personnel relocation, power and cooling, etc.); however, in the context of this paper, disaster recovery and intelligent application recovery more specifically will be limited to recovery of mission-critical applications.

Symantec's intelligent application recovery solutions are based in large part on the flexible architecture of Veritas™ Cluster Server. The same architecture that works in a simple two-node cluster can be easily extended to:

- A single cluster spanning two or more locations where the data at each location is a mirror of the storage at the primary location
- A single cluster spanning two or more locations where the primary storage is replicated to alternate locations
- Up to four clusters, with a single cluster at each location where the primary storage is replicated to the alternate locations

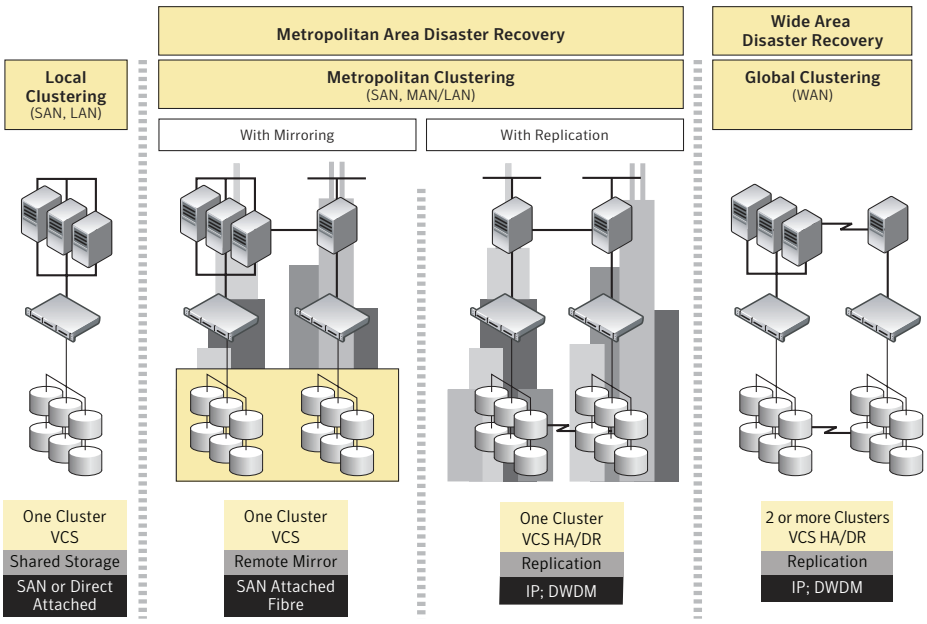

Figure 1. Flexible Veritas Cluster Server architecture

Symantec offers Veritas Cluster Server in two versions:

1. **Veritas Cluster Server**—limited to high availability at a single site as well as Metro Clustering with Veritas Storage Foundation™ mirroring only
2. **Veritas Cluster Server HA/DR**—for all high availability and disaster recovery configurations. This version also supports VCS HA/DR Fire Drill for automated DR testing.

Throughout this guide, the term “Veritas Cluster Server HA/DR” will be used unless the specific goal under discussion can be accomplished using "Veritas Cluster Server" only.

The following factors have influenced the design of Symantec’s intelligent application recovery solutions:

- Standardization across the data center
- Maximization of server utilization
- Simpler installation and management
- Single architecture at any distance
- Dual-use disaster recovery
- Ability to measure and test

As with everything else, determining how quickly you can resume mission-critical operations and the level of automation is always a tradeoff between the business needs of the corporation and the following factors:

- Costs
- Abilities of IT/IS staff to manage the disaster recovery environment
- Existing levels of administrator skill sets
- Concerns regarding complexity of the disaster recovery environment

In making this determination, it is vital that business interests are the driving factor: What level of disaster recovery does the business need and can it be justified in business terms? It should not be a decision on what is thought to be affordable, or one that is made by the IT/IS staff in a vacuum without regard for the needs of the business, or a "let's get the Bentley of disaster recovery when a Morris Mini approach may suffice." The CEO of the company has a fiduciary responsibility to the board of directors and the shareholders (in publicly held companies) and investors (in privately held companies) to properly protect the interests of the business. Government agencies have a moral (and in many cases, legal) obligation to be sure that services are always available. Part of protecting business interests is ensuring that there is a well-defined, documented, and tested Disaster Recovery Plan for the technology-based assets in place that is part of a larger Business Continuity Plan and program based on a thorough and regularly updated Business Impact Analysis. There are many consultants in this field to assist with this kind of work, including Symantec's consulting organization.

The following sections will discuss how to choose the proper Symantec solutions that can sufficiently meet the disaster recovery needs of the corporation without getting too much or too little. Each solution will be examined in depth, highlighting the strengths as well as the challenges of each option.

Intelligent application recovery concepts

The following sections provide background information to insure that common terms are being used.

Definitions

- **Business Continuity Planning—**A proactive and ongoing process that identifies the key functions of an organization and the likely threats to those functions. From this information, plans and procedures are developed, enabling the organization to respond to a business interruption incident and ensure that critical processes can be maintained at a pre-agreed-upon level of functionality.
- **Recovery Requirements Definition/Business Impact Analysis—**A process designed to prioritize business functions by assessing the potential quantitative (financial) and qualitative (non-financial) impact that might result if an organization was to experience a business continuity event.
- **Disaster/Event/Incident—**A sudden, unplanned calamitous event causing great damage or loss as defined or determined by a risk assessment and BIA; Any event that creates an inability on an organization's part to provide critical business functions for some predetermined period of time. The period when company management decides to divert from normal production responses and exercises its disaster recovery plan. Typically signifies the beginning of a move from a primary to an alternate location.

  Significant impacts can be some or all of the following (but not limited to):
  - Revenue loss
  - The inability to recognize revenue
  - Loss of market share
  - Potential litigation
  - Brand damage
  - Reduced customer satisfaction
- **Disaster Recovery Plan—**An integral part of an organization's Business Continuity Plan, which defines the actions to be taken in the event of a disaster. This plan encompasses everything from communications, to locating and occupying temporary space for employees to work out of, to recovery of the business's technology assets (business applications and data, servers, storage, networking, etc.) on which their mission-critical applications run.

- **Mission-critical Application—**An application that is essential to the organization's ability to perform necessary business functions. Loss of a mission-critical application would have a negative impact on the business, as well as potential legal or regulatory impacts.

Recovery point objectives/recovery time objectives

In determining the level of disaster recovery that is needed, it is important to understand the following parameters for each and every current or future mission-critical application. The assessments of these values are business decisions. They are not technical or administrative decisions. The needs of the business must dictate what these parameters are for each mission-critical application.

- **Recovery point objective (RPO)—**The point in time, preceding a business interruption, that data must be available, from an application perspective, at the recovery site. This is essentially the "currency" or "freshness" of the data at the recovery site. How current must the mission-critical data be in order to still provide the services that customers expect? Does the enterprise need data to be current up to the last minute prior to the disaster? Last hour? Last four hours? Last day? The RPO is a business decision relating to how far back in time from the disaster the business data must be current in order to still be effective. How much data loss is acceptable in a disaster? This is generally a tradeoff between how up-to-date the data is at the recovery site and cost.

An example of an RPO that would be used for applications protected by tape backups would be based on the assumption that a full backup to tape is performed nightly, and then shipped off-site the next morning. The worst case for recovering data is when a data loss occurs just before the scheduled backup is to occur. The data on the backup tapes from the previous night is now 24 hours old and has been shipped off-site. The time to retrieve the backup tapes must be included into the RPO. This results in a minimum RPO for tape backups of at least 36 hours, possibly even 48 hours. However, this solution may be more than adequate for internal business functions such as payroll, HR. legal, and so on. External functions that directly impact bottom line revenue may need an RPO closer to 0.

- **Recovery time objective (RTO)—**The maximum acceptable length of time that can elapse before the lack of a business function severely impacts the business entity. An RTO comprises two components:
  - The time before declaration
  - The time to perform the tasks (as documented in their plans) to the point of business resumption

An RTO is based upon the level of projected exposure over time that would be acceptable to senior management from a business perspective.

An example of an RTO that would be used for a mission-critical Web server application assumes that the Web server can be off the Internet for no longer than four hours. This means that the company has decided to accept the risk of not having this mission-critical Web server available for a maximum of four hours. Beyond four hours of unavailability means that the company is incurring losses that it no longer deems acceptable in terms of its business requirements.

As RPO and RTO values decrease (get closer to the actual time of the disaster by approaching values at or near 0), costs for recovering data and bringing applications online go up significantly. Providing a disaster recovery capability is very similar to purchasing insurance. Companies are willing to pay a specific cost to provide a specific amount of coverage. The important factor is that the insurance should not cost more than the potential loss.

Making a decision on RPO/RTO is very similar. If an application costs $1,000 per hour of downtime, providing a comprehensive intelligent application recovery solution offering recovery in seconds at a cost of $100,000 per month may be a bit excessive. Tape backup and off-site tape storage at a cost of $5,000 per month may make more sense. At the same time, using a tape backup and recovery solution for an application that cannot afford to lose any data, at a cost of $100,000 per hour of downtime, leaves the organization under-insured.

A disaster will happen. It is just a question of when and how severe will it be. How much disaster recovery capability is required depends on the needs of each organization. In some cases organizations will need the capability to support an RPO/RTO of near 0 for certain applications, in other words those applications require an intelligent application recovery infrastructure. In other cases, an RPO/RTO of 12 to 24 hours will be sufficient.

A complete disaster recovery plan is not delivered by any one technology, service, or vendor but rather by a combination of products that are implemented in order to provide the needed RPO and RTO of an application. When analyzing a disaster recovery solution, many components must be implemented in order to guarantee data and application availability.

Figure 2 outlines technologies that map to an organization's traditionally accepted RPO and RTO requirements. The burst in the middle represents the point in time that a disaster has occurred. To the left of the burst is the RPO, which outlines technologies that provide data recovery based on the organization's needs. For example, if the business can afford a particular application's data loss of a day or more, then a tape backup approach may be sufficient for backup and recovery of

that application. However, if a day or more worth of data loss will cause substantial impact (an unacceptable loss of revenue, loss of reputation, legal implications, etc.) on the operations of the business, then replication technologies, in addition to the backup strategy, must be implemented in order to protect against substantial data loss. Replication goes beyond the traditional tape backup approach by creating a duplicate copy of data; in real time or near real time, at an alternate facility so that it can be accessed immediately should a complete site outage occur.

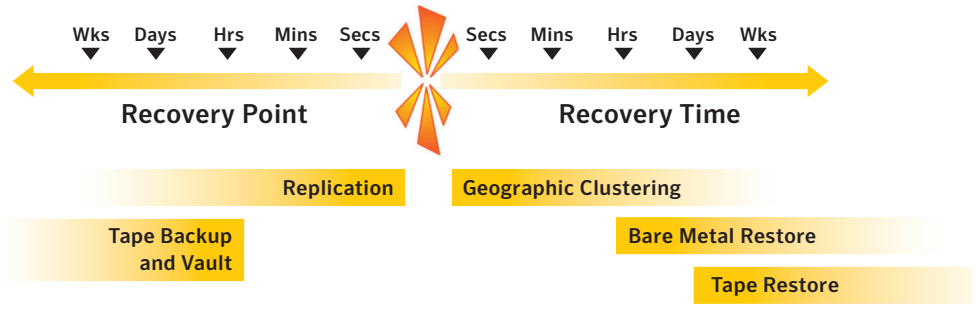

Figure 2. Determining the appropriate technology based on RPOs and RTOs

To the right of the burst is the RTO. If an organization can afford for an application to take several days or more to resume normal activity, then manual backup tape restores will most likely satisfy their needs. The RTO can be improved by using automation such as bare metal restore capabilities (rather than manually building servers prior to loading backup data), assuming that there are servers available for this purpose. However, if an organization cannot afford to have an application unavailable for more than a few hours, manual recovery with traditional tape restore is no longer an option. Unless the organization is willing to immediately restore tapes at the DR site as soon as they arrive to reduce this time when a disaster occurs, an online data mobility solution such as replication or mirroring must be considered so that the data is already “ready” for recovery. In addition, a reliable automation solution must in place to coordinate the startup of applications on top of the replicated data. Application clustering solutions with comprehensive intelligent application recovery capabilities can be used to automate the application and database failover process and reduce downtime to a minimum.

Tables 1 and 2 suggest some appropriate products for a given RPO and RTO value or range.

RPO Objective	Technology	Example Product Offering
1 week or more	Daily or weekly backups and weekly off-site vaulting	IBM® Tivoli® Storage, Manager EMC NetWorker, Veritas NetBackup™ and Vault Option
1 day or more	Daily backups and daily off-site vaulting	IBM Tivoli Storage Manager, EMC NetWorker, Veritas NetBackup and Vault Option
24 hours–1 minute	Asynchronous replication	EMC SRDF/A, Hitachi TrueCopy, Oracle® Data Guard, Veritas™ Volume Replicator
Approaching 0	Synchronous replication	EMC SRDF/S, Hitachi TrueCopy, Oracle Data Guard, Veritas™ Volume Manager, Veritas Volume Replicator

Table 1. Appropriate technologies and products to meet RPOs

Now that the technology has been identified to protect the business data (Table 1), the technology to recover that data must be identified (Table 2).

RTO Objective	Technology	Example Product Offering
3 days or more	Typical tape restore	Veritas NetBackup, IBM Tivoli Storage Manager, EMC NetWorker
12–24 hours	Bare metal restore to rebuild servers quickly	Veritas™ Provisioning Manager, Veritas Bare Metal Restore, Cristie Bare Machine Recovery for IBM Tivoli Storage Manager, EMC NetWorker Recovery Manager
4–12 hours	Only replicate data to secondary site so that all data lives on disk and is able to be restored immediately	Veritas Provisioning Manager (for automation), plus any of Veritas Volume Replicator, Veritas Storage Foundation, EMC SRDF, Hitachi TrueCopy (for data replication)
10 minutes–4 hours	Automatically bring up services at secondary location	Veritas Cluster Server HA/DR (for automation) plus any of Veritas Volume Replicator, Veritas Storage Foundation, EMC SRDF, Hitachi TrueCopy (for data replication)

Table 2. Appropriate technologies and products to meet RTOs

Note: Symantec defines recovery as the ability for users of mission-critical applications and data to have access to those applications and data with as little interruption to them (and/or reconfiguration to be done by them), in as highly an automated manner as possible.

For the purposes of this document, the focus will be on low RTO applications.

Why backup alone is not enough

Many disaster recovery plans rely on data backup and recovery environments as the primary component of their disaster recovery plan. For many years, traditional disaster recovery consisted of no more than the data backup and recovery environment. Typically, tape-based data backup and recovery environments can provide an RPO of about two days (assuming backups are run every night and the tapes are sent off-site) and an RTO of around 12 hours, depending on the amount of data to be restored. The RTO will be longer if very large amounts of data need to be restored. Symantec sells several major data backup and recovery products (NetBackup and Backup Exec™) as a technology that is but one component of a comprehensive disaster recovery plan. Backup and recovery environments can be suitable for applications that have high RPO and RTO values (48 hours or more). However, there is a considerable amount of manual work that needs to be done at the disaster recovery site, before the applications are made available to the users:

- Declare a disaster at the primary site and inform the disaster recovery site of the declaration.
- Obtain the most recent full backup tapes and any subsequent incremental or differential backup tapes. In an ideal world, these tapes may already exist in the backup tape library at the disaster recovery site. If not, they will have to be brought to the disaster recovery site.
- Provision and image backup server(s) to run the backup environment at the disaster recovery site, if they are not already running. This may mean configuring the backup server as it was at the primary data center, down to the host name. It may also mean importing the backup catalog before any restore operations can take place.
- Provision and image target servers to which the data will be restored. Provision and allocate storage for the target servers. This could be problematic in that the servers may not exist or the exact patch levels may be different. A bare metal restore capability would mitigate the issue of patch levels, where a predefined image can be loaded onto the target server. The backup and restore application would have to be part of the image; otherwise it would have to be loaded separately after the operating system and patches have been installed.
- Restore the backup tapes to the target server(s).

- Make the appropriate DNS changes, to allow users to transparently access the applications.
- Validate that the applications are running correctly and that the data they're using is also valid.
- Manually bring up the applications in the correct sequence, and make them available to the users.

That is a large number of manual steps to be taken. Their successful completion depends heavily on the backup environment and the skills of the backup administrators as well as other administrators (DBAs, network, storage, application, etc.), not to mention the involvement of outside services such as off-site tape archive vendors.

The case for automating disaster recovery, even for applications with RPOs between 24 and 48 hours and RTOs of 12 or more hours, is a strong one. All of the manual steps listed above could easily be avoided if the disaster recovery solution implemented some form of data replication and automated control of applications on a site-wide basis where applications are brought up in the proper sequence, along with properly configuring DNS so that the users may access the applications transparently.

One of the major weaknesses in using a data backup and recovery environment (it doesn't matter what product is used) is that test restores are rarely if ever performed or documented. Each restore of a major environment (not just a few files or directories) is a major undertaking, and is not always successful. The lack of success is generally associated with a process-related failure or the failure to completely backup everything that was required, rather than an out-and-out failure of the backup and recovery environment.

Today's mission-critical enterprise applications require RPO and RTO values that are well beyond the capabilities of any data backup and recovery environment. As such, they require:

- data replication to a facility that can act as a disaster recovery data center
- clustering that can automate much of the disaster recovery process from the primary data center to the disaster recovery data center

Why replication alone is not enough

Many disaster recovery plans rely solely on the replication of data to an alternate site(s). The thinking goes along these lines: If the data is protected—meaning that it has been replicated to one or more sites—the business assets of the organization (the data) have been protected. Disaster recovery built upon products like EMC Symmetrix Remote Data Facility (SRDF) is based on this thinking. EMC SRDF does a perfectly good job at replicating data, as do other products such as Hitachi Data Systems TrueCopy and Veritas™ Volume Replicator from Symantec. It is

essential to understand that disaster recovery is more than just having multiple copies of the data at different locations. It is also critically important to note that often, as new applications are added or storage requirements for particular applications grow, some portion of the mission critical data may not be properly configured for replication. When this happens, the data at the alternate locations is incomplete and applications will not start in the event of a disaster. Organizations are then forced to recover from tape resulting in significantly more downtime. Clearly what is needed is a way to frequently "stress-test" the replicated data to ensure that all of the data that should be configured for replication actually is configured for replication. This is why data replication is just one aspect of a comprehensive and reliable disaster recovery plan.

Another critical aspect of disaster recovery is to be able to launch the business applications that depend on the data that has been copied to the alternate sites. It must be clearly understood that, from a disaster recovery perspective, having the data at a disaster recovery location is essentially useless without the business applications up, running, and being accessed by users. What has been accomplished just by replicating the data to alternate sites is a very expensive backup of that data and nothing else. In order for that data to have value to the business, the mission-critical applications that use that data must be available to the users.

Let's look at a hypothetical example. An online auction service enables its customers to sell and buy items using a bidding process that has been limited by the amount of time that an item is available. This company has two sites: a primary data center where all of the data resides, along with the business applications that users interface to, and a secondary site designated as the disaster recovery site.

In the scenario that only replicates data (using any replication product), the data that drives the business application is replicated, possibly multiple times, to the second site. The mode of replication (synchronous, asynchronous, etc.) is irrelevant for this scenario. A disaster occurs that makes the first site unavailable to the company's users. The type of disaster is also irrelevant for this scenario. The company knows that its data is safe at the disaster recovery site. However, the users cannot use the services of the company if the business application is not up and running. In this case, the company may lose money, its reputation as an online company may suffer as a result, and users may look for a different online auction service to use and never come back to our hypothetical company. In order for this scenario to be less dire than it has been described, several manual steps must occur in a carefully defined sequence:

- A disaster must be declared at the primary data center.
- The replicated copies of the data at the disaster recovery site must be brought online, if they are not already. This could be a multiple-step process.

- DNS changes must be made too, so that users can access the applications at a site different from the primary data center.
- The applications must be brought up in the proper order and sequence—assuming that not just the data was replicated, (and that all of it was replicated, with nothing accidentally omitted) but all of the applications' binaries and configuration files as well. If that is not the case, add the step of loading the binaries and configuration files from the latest backup tapes (which may or may not be at the disaster recovery site). This also assumes that there are servers in place with the proper operating system and patch levels ready to run the applications at the disaster recovery data center.

These are manual steps being performed in a time of crisis (the disaster at the primary data center). There will be varying degrees of documentation describing how to accomplish each step in the process, by people who normally may not be associated with these tasks in the primary data center, whose skills in handling these tasks may not be at the same level as those of the people in the primary data center. These are several additional risk factors that must be considered in designing the disaster recovery plan.

In the scenario that replicates the data and automatically brings up the mission-critical applications at the disaster recovery site, there is a different outcome. The data is replicated as it was in the first scenario. The difference is that the mission-critical applications are brought up automatically in the correct sequence at the disaster recovery site in addition to applying the DNS changes that are necessary for users to access the applications transparently (without having to make any changes to how they access the application). The only real human intervention in this scenario is the initial declaration of the disaster. Once that has been done and the business decision to move operations to the disaster recovery site has been made, everything from this point forward can be done automatically.

This is a very strong argument for automating the disaster recovery of technology-based assets. In a time of disaster, there is a tremendous amount of pressure and stress to get everything back up and running and available to users. In the manual process, mistakes will be made for a variety of reasons. Maybe the documentation for the procedures is out-of-date, poorly written, or incomplete, or it cannot be found or is not available because it was online at the primary data center. Maybe there has been some configuration drift between the primary and disaster recovery data centers. Having an automated disaster recovery capability would eliminate many of these risk factors. In addition, the same disaster recovery infrastructure used in the event

of a disaster can (and should) be used on a regular, frequent basis to "stress-test" recoverability of the replicated data, the server environment, and the application environment. Veritas Cluster Server HA/DR Fire Drill delivers such automated stress-testing out-of-the-box, and is discussed later in this paper.

Expanding the use of replication and automated disaster recovery

The case for automating technology recovery has been made in the previous examples concerning the use of the backup/recovery environment, and the use of replication only as the underlying disaster recovery technology. Does this mean that the backup/recovery environment is no longer needed? Absolutely not. It is still an industry-wide best practice that all data be backed up in a secure and reliable manner with the knowledge that anything from the most trivial file to the most complicated data warehouse can be restored at will. It may impact whether or not duplicate tapes are made. Alternatively, does this mean that replication is of no value? Again, absolutely not. It is vital that mission-critical data, application binaries, configuration files, and user files be “copied” to an alternate site(s) in a manner that is consistent with business requirements.

A cluster-based technology like Veritas Cluster Server HA/DR along with replication technology vastly improves the ability to recover in a well-defined, rehearsed, and highly automated manner from serious site disasters, as well as lesser kinds of interruptions, regardless of RTO/RPO values. Clustering technology in the past was only reserved for a few key applications that had to be highly available. Due to the extensibility of the Veritas Cluster Server HA/DR architecture and its ease of use, clustering can be used to simplify the management of and recovery of a broader range of applications and environments by automating the response to interruptions, major or minor. Replication ensures that mission-critical data is “copied” to the disaster recovery data center in real time or near real time. Clustering ensures that mission-critical applications are managed from the perspective of keeping them highly available, either locally or remotely. Essentially, the traditional ranges of RPO and RTO can be extended for replication and clustering. This is shown in Figure 3 in green. This provides a much greater degree of automation of disaster recovery than using the traditional methods suggested solely by the RPO and RTO values of before.

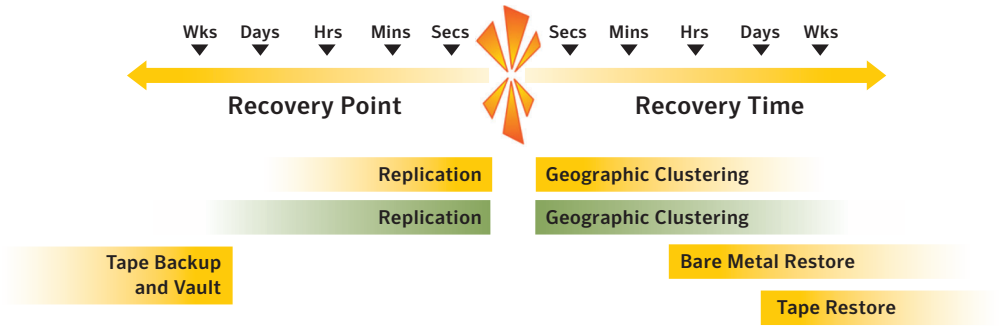

Figure 3. Revised RPO/RTO range for replication and clustering

Take, for example, the following: As the amount of storage increases in the data center, it becomes increasingly difficult to build and maintain a backup and recovery environment that is capable of fully restoring all of the backup data within the recovery time window. In these situations, it may be prudent to use a combination of replication and clustering to maintain a disaster recovery site, even if the RTOs for the business applications suggest that restoring from tape may be adequate. Rarely are backup and recovery environments designed, maintained, and updated to optimize restore performance or keep pace with the amount of storage growth.

In responding to a disaster, the disaster recovery planner strives for repeatability and consistency of operations. This is not the case with the backup-only or replication-only models described above that are completely dependent on so many manual steps in order to effect a recovery. The use of an intelligent application recovery solution combining clustering and data replication by definition provides repeatability and consistency. Recovery can be fully automated. It eliminates the need for people to locate all of the procedures, which may or may not exist in written (or accessible) form at the disaster recovery center. These procedures may or may not be regularly updated. These procedures may or may not be regularly tested. These procedures may or may not be executed in the proper order, by staff members who may or may not have all of the skills that the enterprise needs them to have to recover from a disaster. Veritas Cluster Server HA/DR clustering technology eliminates all these potential areas of increased risk where mistakes can be made in a time of disaster.

The use of Veritas Cluster Server HA/DR clustering technology along with a form of data replication allows the enterprise to maximize its investment in hardware and people by automating the response to minor component failures all the way up to major site-wide disasters using the same set of tools. Veritas Cluster Server HA/DR allows the enterprise to protect its

ability to continue operating in the face of any kind of technology-compromising disaster. It does so with a minimal training footprint for administrators. There is no longer a need to depend on the one person who "knows it all" about every mission-critical application to manage a disaster recovery effort, who themselves are single points of failure.

What about the cost?

One of the challenges often cited with automated disaster recovery is that the cost of a "hot" site is too high. This needs to be put into the proper perspective of preventing the costs that would be associated with a site-wide disaster at the primary data center (loss of revenue and reputation, and legal and regulatory costs as a result of the disaster, etc.). Industry experts have said that companies that do not invest in business continuity plans and disaster recovery technology run a significant risk of going out of business within 5 years after a disaster.

Symantec intelligent application recovery solutions are designed to protect mission-critical applications running at the primary data center from a disaster that no longer allows the mission-critical applications to run at that location. This could be a disaster as simple as a backhoe cutting all of the communications cables outside the primary data center, disconnecting external users from the applications, or it could be as severe as a major natural disaster (earthquake, tsunami, hurricane, tornado, pandemic, etc.) or terrorist activity on the scale of September 11. While there is additional cost associated with configuring automated disaster recovery, it will generally pale in comparison to the costs of not being able to provide services/products to customers. The Symantec intelligent application recovery solutions have been designed to automate the process of recovering from a disaster and to ensure that not only is the mission-critical data protected by using replication, but also that mission-critical applications using that data are highly available locally and remotely in the event of a disaster.

Supported intelligent application recovery configurations

There are basically two tracks of configurations available within the Symantec intelligent application recovery model: Global Clusters in one track, and Metro Clusters in the second track. The major differences between the two tracks depend on how a failover is treated in going from one site to another. Is there a degree of control involved or is it fully automated? Both tracks support local failover within the site.

Global Clusters link two (or more) independent local Veritas Cluster Server HA/DR clusters to form a global failover relationship for specific applications. Failover between local systems in either cluster is fully automatic. Failover between sites (i.e. clusters) requires operator

confirmation. Enabling confirmation gives operators the ability to ask "what should I do?" in the event that local failover protection can no longer protect the mission-critical applications, and a decision (a business decision) has to be made around whether to perform an intelligent application recovery at the disaster recovery location. An operator can thus indicate:

- "Yes," we will incur a site failover. The local disaster is not expected to be resolved in the near future—or more specifically within our RTO limits—and it makes sense from a business perspective to bring all of the mission-critical applications up as quickly as possible at the disaster recovery site. In cases where an asynchronous replication technology is employed, this decision essentially accepts the fact that the company will be coming up on somewhat out-of-date data.
- "No," we believe that the disaster will be resolved shortly, within our RTO limits, and therefore we will not incur a site failover. It is less disruptive to stay at the primary site, even though the site is down, than it is to incur a site failover to the disaster recovery site, and then another one to come back to the primary data center once the disaster has been resolved.

The decision to place an operator in the decision path with global clustering is usually the recommended course of action. It allows the business to assess the severity of the disaster with respect to the RTO goals of the mission-critical applications, and then act accordingly. Manual disaster recovery failover control also eliminates the risk of faulty communications between sites triggering an automated response. In many situations, site-to-site communications are provided by an outside vendor and may not be as reliable as desired.

Metro Clusters extend a single cluster between multiple sites and act at all times like a single cluster. The operator can configure failover ordering to always failover between systems located within the primary site before failing over to systems located at the disaster recovery site. However, when an outage affects all servers at one site, a failover to the second site will occur automatically. The single cluster solution is typically deployed in a metropolitan area, with full synchronous mirroring or replication and very reliable communications links between sites. In such cases, the company is essentially expanding the concept of high availability to encompass more than one data center. Assuming all infrastructure components are very solid, and a business need exists to have full automated failover to a remote site, then this is a very viable solution.

The following sections will provide more detailed information on the various cluster configurations.

Metro Clusters

Metro Clusters are single Veritas Cluster Server clusters that have been extended to more than one site. The Metro Cluster behaves exactly the same way as a Veritas Cluster Server cluster in a single site in terms of failover behavior. The underlying data transport mechanism between the sites provides for slightly different configurations in terms of the storage and how data from the primary site is "copied" to the disaster recovery site. The data can be replicated using application-based, host-based, or array-based replication. Or the data can be mirrored using Veritas Volume Manager, provided as part of Veritas Storage Foundation. Metro Clusters can provide metropolitan area disaster recovery.

Metro Cluster with Replication

A Metro Cluster with Replication is a single Veritas Cluster Server HA/DR cluster that is spread across two physical locations. This configuration eliminates the single point of failure that a single data center represents, as it assumes independent power and communications to each facility. The Metro Cluster essentially provides extended area high availability, which gives it the capability to provide disaster recovery automation at metropolitan distances. The amount of separation between the two data centers depends on the risks that the company wants to protect their environment from, but is limited in total distance supported by the requirement for synchronous replication. This typically is metropolitan in nature (less than 80 kilometers (50 miles).

The nodes within each data center share storage at that data center. There is no shared storage between data centers, nor is there an extended SAN between the data centers. A Metro Cluster uses data replication to assure data access to all nodes at each data center. In a Metro Cluster configuration, if an application or a system fails, the application is restarted by VCS HA/DR on another system within the current primary site or zone. If the entire primary site fails, the application is automatically restarted on a a system at the remote secondary site (which then becomes the new primary). In the event of a storage failure at the primary site, VCS HA/DR will detect that there has been a failure and will perform the operations necessary to prepare the storage at the disaster recovery site for production use, and then restart the application on top of that storage.

Synchronous data replication keeps the copies of data at the two data centers synchronized. Asynchronous replication cannot be used for a Metro Cluster because of the potential of data loss upon an automatic failover between sites. Replication can take place at the application, host, and storage levels. Application-level replication products, such as Oracle Data Guard, maintain consistent copies of data between systems at the SQL or database levels. Host-based replication products, such as Veritas Volume Replicator, maintain consistent storage at the logical volume level. Storage- or array-based replication products such as EMC SRDF or Hitachi Data Systems TrueCopy maintain

consistent copies of data at the disk or RAID LUN level. Supported distances between data centers for synchronous replication are approximately 80 kilometers (50 miles) or less.

The Metro Cluster configuration provides both local high availability and disaster recovery functionality in a single Veritas Cluster Server HA/DR cluster.

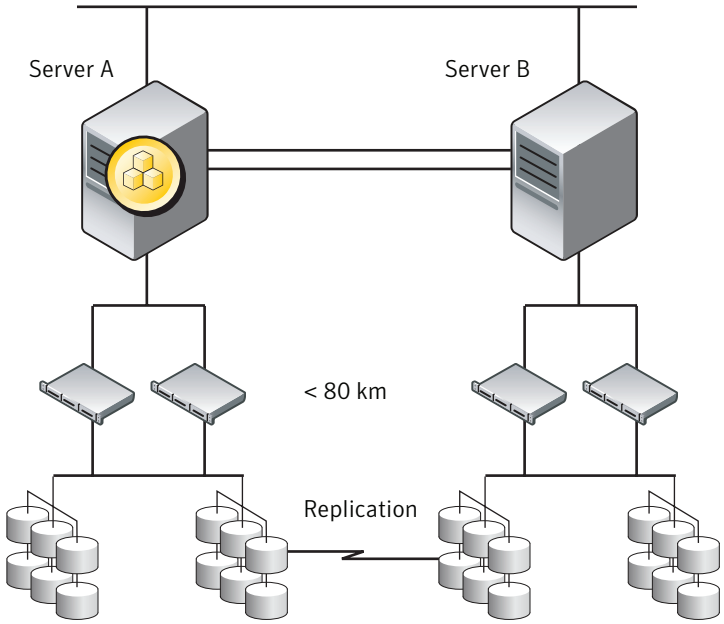

Figure 4. Metro Cluster with Replication

A Metro Cluster configuration is appropriate in situations where dual dedicated cluster interconnects are available between the primary site and the disaster recovery secondary site but lacks shared storage or SAN interconnect between the primary and secondary data centers.

To understand how a Metro Cluster configuration works, let us take the example of an Oracle database configured in a Veritas Cluster Server HA/DR Metro Cluster. The configuration has two system zones:

- Primary zone (zone 0) comprising nodes located at the primary site and attached to the primary storage
- Secondary zone (zone 1) comprising nodes located at the secondary site and attached to the secondary storage

Oracle is installed and configured on all nodes in the cluster. Oracle data is located on shared disks within each Metro Cluster zone and is synchronously replicated across Metro Cluster zones to ensure data concurrency. The Oracle service group (a service group within Veritas Cluster Server is the smallest unit of failover—it contains all of the resources that a mission-critical application needs to come online) is online on a system in the current primary zone and is configured to fail over in the cluster.

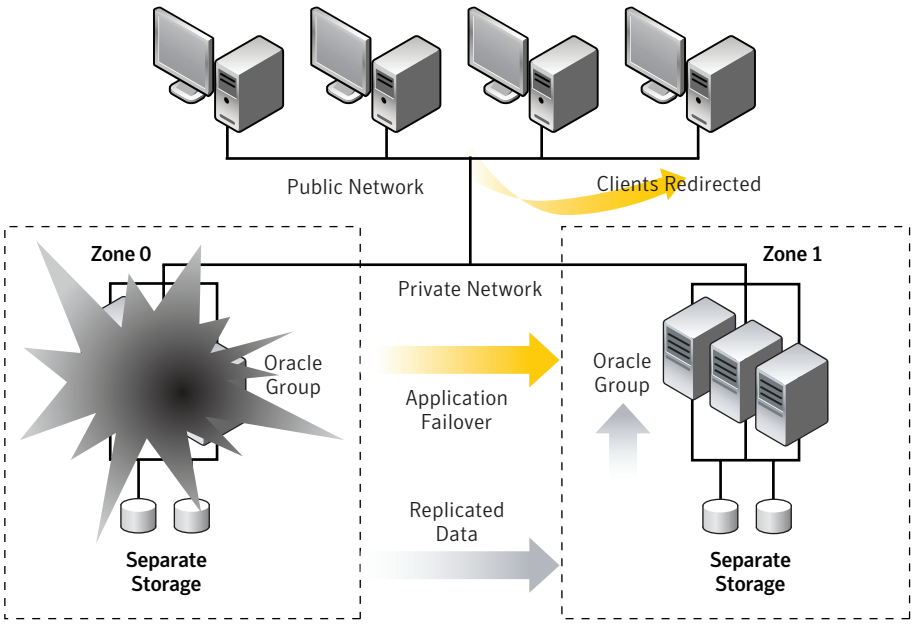

Figure 5. Metro Cluster with Replication failover example

In the event of a system or application failure, Veritas Cluster Server HA/DR attempts to fail over the Oracle service group to another system within the same Metro Cluster zone. However, in the event that Veritas Cluster Server HA/DR fails to find a failover target node within the primary Metro Cluster zone, Veritas Cluster Server HA/DR automatically switches the service group to a node in the current secondary Metro Cluster zone (zone 1). Veritas Cluster Server HA/DR also redirects clients by bringing up the application IP address(es) at the zone 1 data center. The Metro Cluster with Replication configuration requires Veritas Cluster Server HA/DR installed on each server at each location, and supports the use of VCS HA/DR Fire Drill for automated testing.

Metro Cluster with Mirroring

The Metro Cluster with Mirroring is a single Veritas Cluster Server cluster that spans two or more physical locations similar to the Metro Cluster. The storage at one location is mirrored to the other location using Veritas Storage Foundation. Mirroring is a synchronous process. Unlike the Metro Cluster with Replication, the storage at both data centers is connected using an extended SAN. For all intents, the user is configuring a RAID 1 mirror between two arrays at different sites. All data is written simultaneously to both mirrors. Reads can be configured to be serviced by the site where the application is currently running. The separation between the data centers has a practical limit in terms of performance of approximately 80 kilometers (50 miles) as the fiber is pulled, not as the crow flies. A Metro Cluster with Mirroring can also be configured using the concept of zones to set failover ordering to always attempt failover locally first.

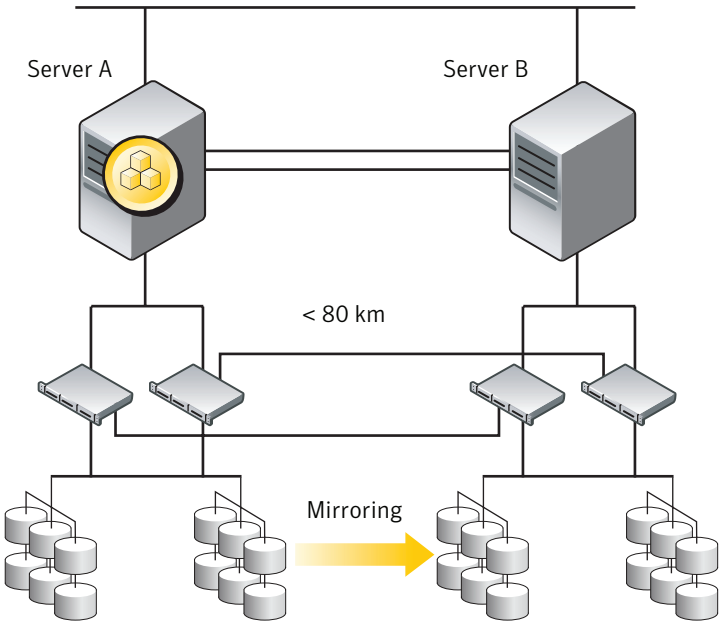

Figure 6. Metro Cluster with Mirroring

A Metro Cluster configuration is appropriate in situations where dual dedicated cluster interconnects are available between the primary site and the disaster recovery secondary site and there is an extended SAN connecting the shared storage between the two sites. The storage at one site is configured so that it will be mirrored to the second site. Veritas Cluster Server and Veritas Storage Foundation will handle the management of the disk plexes in the mirrored storage environment.

For failures not involving storage, a Metro Cluster with Mirroring will have the exact same failover behavior as the Metro Cluster with Replication. The Metro Cluster with Mirroring configuration requires Veritas Cluster Server installed on each server at each location. VCS HA/DR Fire Drill is not supported in this configuration.

Global Clusters

A Global Cluster is a collection of two or more Veritas Cluster Server HA/DR clusters (up to four) at separate locations that are linked together to enable intelligent application recovery over any distance, i.e. globally. There is no shared storage between the clusters in a Global Cluster. Each cluster within the Global Cluster is connected to its own shared storage. A Global Cluster has a single primary cluster (i.e. site), and up to three secondary clusters (sites). The storage within the Global Cluster is replicated, either synchronously or asynchronously, from the primary cluster to each of the other secondary clusters. Typically, asynchronous replication over the wide area network (WAN) connecting the data centers is used, but synchronous replication can be used for shorter distances (less than 80 kilometers (50 miles)).

Local clustering provides local failover for each site or building. Metro Cluster configurations offer protection against disasters affecting very small geographic regions. Large-scale disasters such as major floods, hurricanes, earthquakes, and acts of terrorism can cause outages for an entire city or region. In such situations, a company can ensure global availability by migrating applications to sites located considerable distances apart. Over the past few years, the minimum best practice distance separating the primary data center and disaster recovery data center(s) has grown from 50 miles to over 200 miles.

In a Global Cluster, if an application or a system fails, the application is migrated to another system within the same cluster. If the entire cluster fails, the enterprise can make a business decision on whether or not to move operations to one of the alternate disaster recovery sites. If the decision made is to move to a specific disaster recovery site, the application is automatically restarted on a system in the cluster at that disaster recovery site(s).

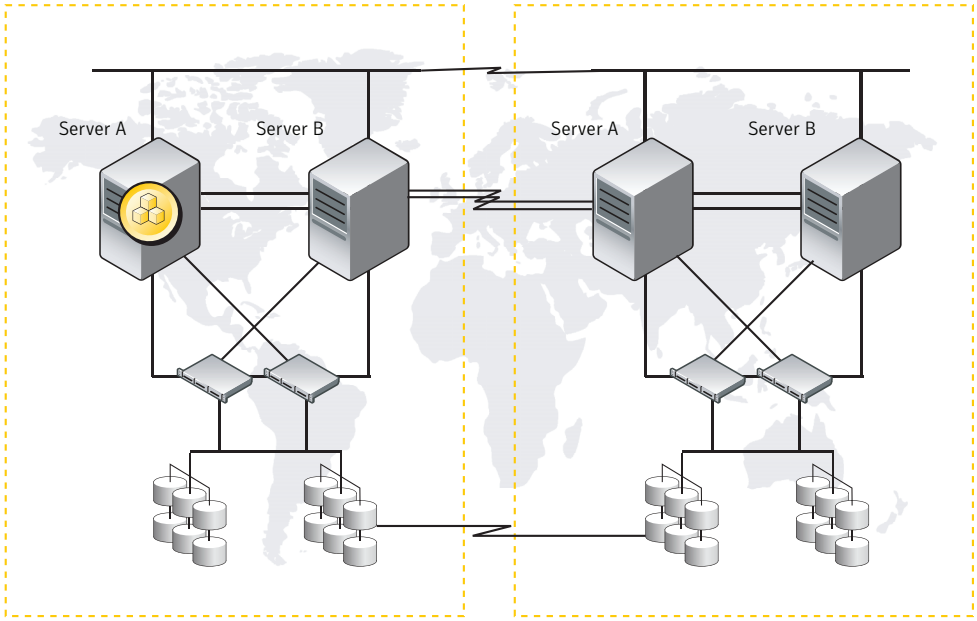

Figure 7. Global Cluster

Let us take the example of an Oracle database configured in a Veritas Cluster Server HA/DR Global Cluster that connects two clusters together. Oracle is installed and configured in both clusters. Oracle data is located on shared disks within each cluster and replicated across clusters to ensure data concurrency. The Oracle service group is online on a system in cluster A and is configured to fail over globally, on clusters A and B.

Veritas Cluster Server HA/DR continuously monitors and communicates events between clusters. Intercluster communication (ICMP-based) ensures that the Global Cluster is aware of the state of the global service group at all times.

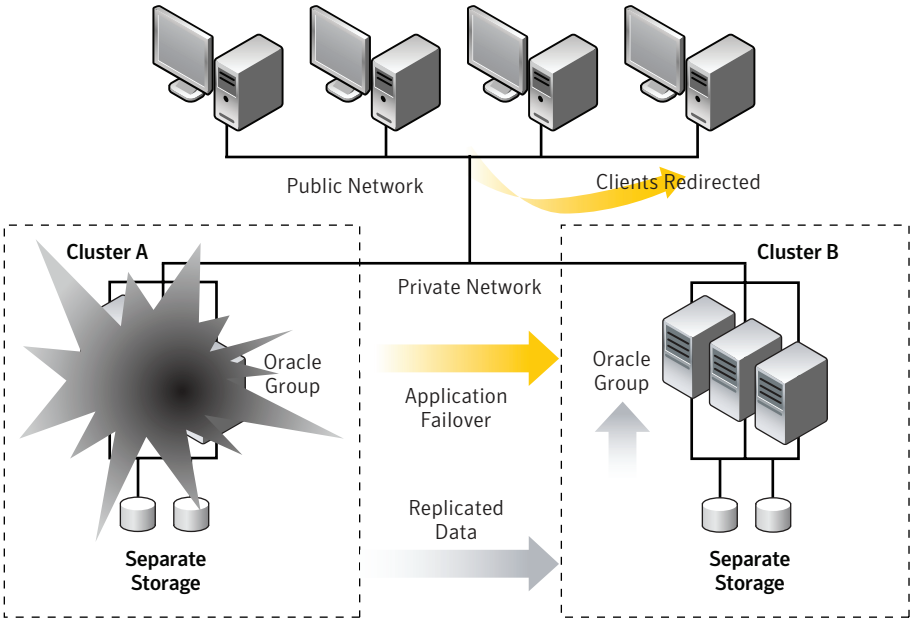

Figure 8. Global Cluster failover example

In the event of a system or application failure, Veritas Cluster Server HA/DR fails over the Oracle service group to another system in the same cluster. If the entire cluster fails, Veritas Cluster Server HA/DR will alert the operator and provide the opportunity to take action. The operator may declare a "*disaster*," which indicates that the primary data center has been lost, (at least as far as the application is concerned—this may be a result of a localized power outage) in which case operations are migrated to the disaster recovery data center automatically. Or the operator may declare an "outage" and decline to allow failover in cases where local restoration of service will happen in a short period of time. In either case, a business decision must be made on whether or not a failover is in the enterprise's best interests, considering RTO limits and what is known about the particular disaster. The Global Cluster configuration requires Veritas Cluster Server HA/DR installed on each server at each location, and supports the use of VCS HA/DR Fire Drill for automated testing.

Replication support in Metro and Global Clusters

In addition to supporting Veritas Volume Replicator, an option for Veritas Storage Foundation, Veritas Cluster Server HA/DR supports a wide variety of third-party replication products. Veritas Cluster Server HA/DR includes a number of Agents for third-party replication that completely automate the process of replication management and application startup at the remote site without the need for complicated manual recovery procedures involving storage and application administrators. The Veritas Cluster Server HA/DR replication agents provide all the necessary logic to completely control the underlying replication configuration, whether that replication operates:

- At the storage array level (e.g., EMC SRDF)
- At the database level (e.g., Oracle Data Guard)
- Synchronously or asynchronously

Depending on the type of failure, this control may involve reversing the direction of the replication (otherwise known as role reversal, role swap, dynamic swap, or personality swap), or simply moving the data and applications back when the original site comes back online. This comprehensive solution also includes the capability to select automatic or operator-confirmed site-to-site failover.

By deploying a comprehensive and integrated disaster recovery solution, the IT department is assured that end users will be able to access their critical data and applications, even in the event of a disaster, without the need for highly trained administrators to manage each component of the disaster recovery action. By placing all storage, application, and network components under control of Veritas Cluster Server HA/DR, organizations can create a completely automated solution designed to carry out the same tested and approved intelligent application disaster recovery actions regardless of time of day or skill set of the personnel present. For the latest in Veritas Cluster Server HA/DR Replication support, please visit:
http://eval.veritas.com/mktginfo/products/Sales_Docs/High_Availability/agent_list_partner.xls

Design considerations

There are a few design considerations that apply generally to the cluster types that have been discussed. The first issue concerns the number of servers or nodes within the cluster at the disaster recovery data center. The second issue concerns updating DNS once the mission-critical applications have been migrated to the disaster recovery data center.

Clusters at the disaster recovery site

The purpose of the DR1 site (disaster recovery data center) pictured in Figure 9 is to handle a massive disaster that disables/destroys the P1 site (primary data center). Another site could be built at some point after the disaster, if it appears that P1 is not salvageable. The Global Cluster capability of Veritas Cluster Server HA/DR would handle the site-to-site failover from the P1 environment to DR1, as well as the necessary DNS updates to allow users to get to the mission-critical applications.

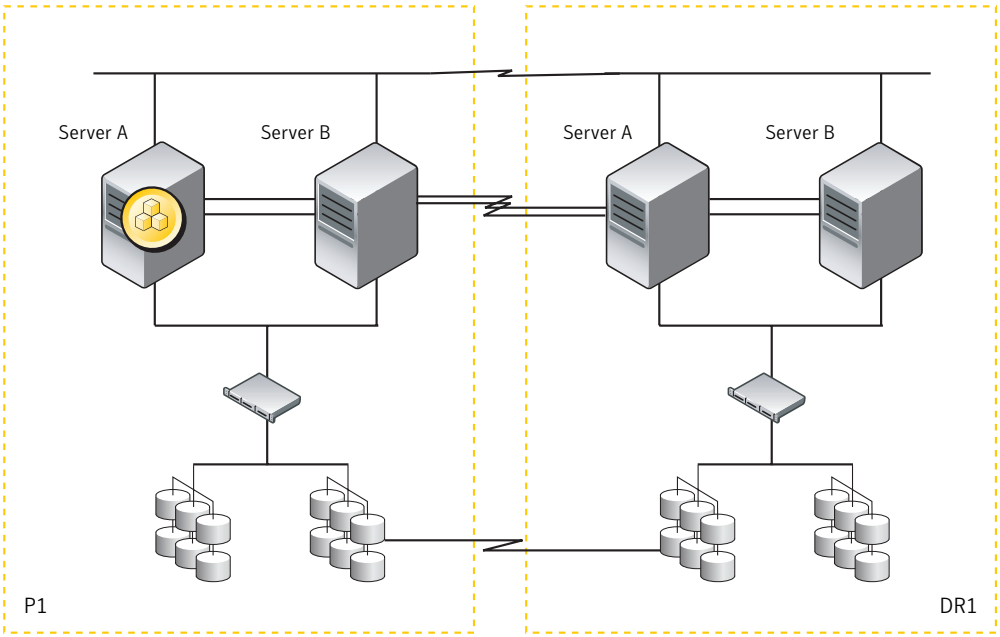

Figure 9. Multi-node cluster at the disaster recovery site

Notice that the cluster depicted at the DR1 site is a multi-node cluster. Veritas Cluster Server HA/DR does support the use of a single-node cluster as the disaster recovery site cluster, and many companies do in fact have a single-node cluster at their disaster recovery site. This would be a slightly less expensive alternative in terms of initial outlays; however, the drawback to this approach is that if all the mission-critical processes are now at the disaster recovery data center, they are not highly available locally within that data center. If the single-node cluster experiences a failure that cannot be cleared/resolved by rebooting, there is no place for the mission-critical applications to run.

If there has been a disaster at the primary data center necessitating a site-to-site failover, no assumptions can be made as to how long they'll stay there. Therefore it is prudent that the mission-critical processes have the same level of local high availability protection at the disaster recovery site as they once did at the primary data center site. This generally means a multi-node cluster at the disaster recovery site as well. While a multi-node cluster at the disaster recovery site will be slightly more expensive initially, that expense will pale in comparison to the costs of a disaster that wasn't prevented. Another way to view this is that if a mission-critical application and its data were deemed important enough to the business to implement local high availability at the primary data center, replication of that data to a disaster recovery data center, and the ability to fail the application over to the disaster recovery data center, why isn't local high availability at the disaster recovery center equally important?

DNS updates

In a Global Cluster or a Metro Cluster (assuming that the data centers are on different networks), when the mission-critical applications are migrated from the primary data center to the disaster recovery center, the network environment changes as a result. The virtual IP addresses with which users access the business applications at the primary data center will not exist at the disaster recovery data center. DNS will have to be updated at the disaster recovery data center to redirect clients using the primary data center virtual IP addresses for the business applications to the virtual IP addresses that exist at the disaster recovery data center.

The Veritas Cluster Server DNS Agent performs DNS updates automatically to redirect clients from a failed group or cluster that has gone offline to its replacement that has come online in a different subnet. This process supports replication but does not require it. The Veritas Cluster Server DNS agent is an interface between DNS servers and Veritas Cluster Server for use in failing IP addresses to different subnets—typically the case when data centers are spread across a wide area. The DNS agent can also be used for local HA scenarios as well if the networking environment demands it.

DNS name servers are used to resolve Internet names, such as www.symantec.com, to IP addresses recognized by routers. Canonical names, such as www.symantec.com, are assigned to real machines within a domain and can be reassigned to backup machines when disasters occur. For example, www.symantec.com runs in Sunnyvale, California. There is a disaster recovery data center in Roseville, Minnesota. Normally, www.symantec.com would point to svlbigip1.symantec.com (fictional name of the Sunnyvale BigIP server). In the event of a site disaster that requires site failover, the Veritas Cluster Server DNS agent would re-point www.symantec.com to rosbigip1.symantec.com (fictional name of the Roseville BigIP server)

transparently to the users. The Veritas Cluster Server policy mechanism can be configured to make these changes, and the DNS agent implements the policy.

When should each model be used

Given all of the information presented so far, which configuration should be chosen? As always, the answer is, "it depends." To a large degree it will depend on the hardware choices that may have already been made.

Symantec has developed a rich set of Veritas Cluster Server HA/DR replication agents that support a wide variety of replication products. Veritas Cluster Server HA/DR tightly integrates with and manages these replication products.

A major factor in choosing between Global Clusters and Metro Clusters is the failover mechanism. In a Global Cluster, the choice of failing over to the disaster recovery site can be made manually, and is suggested as a best practice. Site-wide failure should typically be a business decision, not a software-based decision. Once the business decision to fail over to the disaster recovery site has been made and the "GO" button pressed in the Veritas Cluster Server HA/DR console, the rest of the failover process will proceed in a fully automated manner. In Metro Clusters, a single cluster is extended across sites and handles all failover in full automatic fashion. Metro Clusters are only recommended in configurations with rock-solid site-to-site communications and a need for full auto failover. In all other situations, Symantec recommends a Global Cluster configuration.

The other major choice depends on the distance between the primary production data center and the disaster recovery data center. There will be a subset of companies that will need to have a significant amount of distance between their sites, on the order of hundreds of miles. This may be for actual risk reasons or regulatory reasons. For these companies, asynchronous replication is the only option. In these situations, operator confirmation before failover to an asynchronous copy is required. Global Clusters would be the best practice recommendation for these companies.

Table 3 provides intelligent application recovery recommendations given the parameters of how far apart the data centers are, whether or not there is an extended SAN in place, whether third-party replication will be used, whether local failover within the zone or local cluster is desired, and whether synchronous replication is used.

Distance Between Data Centers	Extended SAN	Third-Party Replication	Failover within Local Zone/Cluster	Synchronous Replication	Symantec Recommended Intelligent Application Recovery Solution
> 80 km	No	No	Yes	No	Global Cluster with Veritas Volume Replicator
> 80 km	No	Yes	Yes	No	Global Cluster with third-party replication
< 80 km	No	No	Yes	Yes	Global or Metro Cluster with Veritas Volume Replicator
< 80 km	No	Yes	Yes	Yes	Global or Metro Cluster with third-party replication
< 80 km	Yes	No	No	Mirroring	Metro Cluster with Storage Foundation mirroring
< 80 km	Yes	No	Yes	Yes	Metro Cluster with Veritas Volume Replicator
< 80 km	Yes	Yes	Yes	Yes	Metro Cluster with third-party replication

Table 3. Symantec intelligent application recovery recommendations

These are guidelines only. For example, if business requirements dictate that data centers be more than 80 kilometers (50 miles) apart and a company wants to run synchronous replication, Veritas Cluster Server HA/DR will not interfere with that. The concern would be for application performance, with an application having to wait on the latency that is inherent in synchronous replication (the acknowledgement of every write), and whether or not users/customers could tolerate this from a usability perspective.

Another example might be the case where data centers do not have to be greater than 80 kilometers (50 miles) apart, but the company wants to run asynchronous replication. That is an acceptable configuration although it may result in the data at the disaster recovery data center being slightly behind the production data, which would result in some data loss in the event of a disaster at the primary site.

Testing the intelligent application recovery environment

Having a disaster recovery capability will only be of real use if it is systematically and frequently tested, since continuous change in today's typical data center often results in mis-configurations, expired application licenses, missing network connections etc. The IT organization must continually ensure that, even in the face of unrelenting change, the intelligent application recovery infrastructure is resilient and ready to do its job. Unfortunately, disaster recovery testing

has always been a bane of IT/IS organizations as testing tends to be extremely *expensive*, disruptive to normal production operations, not to mention that the planning involved in a test can swamp IT/IS managers and their staffs for several weeks prior to the test.

Veritas Cluster Server HA/DR Fire Drill

Veritas Cluster Server's Fire Drill allows verification of the intelligent application recovery environment to be done without disrupting normal production operations. The VCS HA/DR Fire Drill procedure tests the fault-readiness of a VCS HA/DR configuration by mimicking a failover from the primary site to the secondary site. This procedure is done without stopping the application at the primary site and disrupting user access, interrupting the flow of replicated data, or causing the secondary to need resynchronization.

The initial step is to create a Fire Drill service group on the secondary site that closely follows the configuration of the original application service group. The Fire Drill service group will not directly use the replicated storage at the disaster recovery data center. Instead, the Fire Drill service group is configured to create and use a point-in-time snapshot (e.g. Business Continuity Volume, or BCV) of that replicated data, using the snapshot capabilities of the underlying replication technology. For example, Fire Drill utilizes EMC TimeFinder for SRDF snapshots, Hitachi ShadowImage for TrueCopy snapshots etc. Bringing the Fire Drill service group online at the secondary site using the point-in-time snapshot demonstrates the ability of the mission-critical application to fail over and come online at the secondary site, should the need arise. If the test fails for any reason, a detailed entry is logged so that the cluster administrator can investigate and correct the problem. When the test is complete the snapshot is reset. Fire Drill service groups do not interact with outside clients or with other instances of resources, (for example, Fire Drill service groups do not invoke any DNS changes) so they can safely come online even when the application service group is online. A Fire Drill should be conducted only at the secondary site; the Fire Drill service group should not be brought online on the node hosting the original application.

Virtual Fire Drill

Configuring high availability for a database or an application requires several infrastructure and configuration settings on multiple systems. However, local cluster environments are subject to change (configuration drift) after the initial setup, just like any other environment. Administrators add disks, create new diskgroups and volumes, add new cluster nodes, or add new NICs to

upgrade and maintain the infrastructure. Keeping the cluster configuration updated and synchronized with the changing infrastructure is critical.

Virtual Fire Drills detect discrepancies between the Veritas Cluster Server configuration and the underlying infrastructure on a host server—discrepancies that might prevent a service group from going online on a specific host, without having to create a point-in-time copy of application data. Virtual Fire Drill is for local cluster configuration validation only—they are are not for disaster recovery testing. In addition, Virtual Fire Drill tests are "virtual" since they do not actually try and start applications as part of the test.

Virtual Fire Drill performs a series of infrastructure checks, verifying that the resources defined in the Veritas Cluster Server configuration file (main.cf) properly reflect the required infrastructure to fail over on another node. For example, an infrastructure check for the Mount resource verifies the existence of the mount directory defined in the MountPoint attribute for the resource, and this verification happens for each node in the cluster.

An infrastructure check can be run only against a service group that is online—this ensures that the the cluster configuration is indeed correct. If this were not the case, the service group would not have been able to come online in the first place. The check verifies that the specified node (or nodes) is a viable failover target capable of hosting the service group.

Virtual Fire Drill provides an option to fix specific errors detected during the infrastructure check (e.g. create a mount point on a host if it is missing).

Frequently asked questions

What will be the typical minimum disaster recovery configuration, in terms of the size of the clusters at the primary and secondary sites as well as the number of primary/secondary sites?

There is no minimum cluster size for the primary. It will depend on what the business needs of the company are for application availability and disaster recovery. In some cases there will be a single-node cluster at the primary. In these cases, the company is not interested in local failover, but *is* interested in application disaster recovery; however, this could have implications for data currency at the disaster recovery site. If there is no local failover at the primary data center (a single-node Veritas Cluster Server HA/DR cluster), a failover to the disaster recovery site will be required for every outage that cannot be resolved with rebooting the server. Having said that, it is not unusual to see a single-node Veritas Cluster Server HA/DR cluster at the disaster recovery center. See the discussion above on single-node clusters at the secondary site. The size of the cluster will depend heavily on the business needs of the enterprise.

The number of secondary sites will also depend on the business needs.

Does Symantec recommend disaster recovery with or without clustering at either the primary or secondary sites?

The Symantec intelligent application recovery model depends heavily on the robust, easily configurable, and inherently testable automation provided by Veritas Cluster Server HA/DR. This is a differentiating factor for Veritas Cluster Server HA/DR. Additionally, Veritas Cluster Server HA/DR at the primary site will try to fail over service groups locally before going to an alternate site. These are all characteristics that a robust disaster recovery solution should entail. By not having any clustering at the primary site, the company is depending on manual and/or scripted responses in the event of a disaster, or worse, hoping that a site disaster does not occur. Monitoring must be set up manually as well. Instead, an intelligent application recovery solution from Symantec will utilize Veritas Cluster Server HA/DR to deliver a consistent and repeatable result is exactly what the disaster recovery planner wants in their solution.

What replication products does Symantec recommend?

Symantec offers the capability to use host-based replication such as Veritas Volume Replicator, application-based replication such as Oracle Data Guard, or array-based replication such as EMC SRDF or Hitachi TrueCopy. Veritas Cluster Server has agents that support all of the major replication products. As a best practice, Symantec recommends either host-based or array-based replication as opposed to application-based replication. Host and array-based solutions are easily scalable across many applications and provide a more suitable infrastructure layer for enterprise-class replication.

What is the recommended replication mechanism, synchronous or asynchronous?

This choice will depend on several things. Can the application and the user tolerate waiting for each and every acknowledgement of a write? Is the recovery point objective so small that asynchronous cannot be used? If asynchronous is used, can the application/users accept the fact that, at a given point in time, all of the data that has been written at the primary may not actually be on disk at the secondary? What is the distance between the primary and the secondary data centers? How many writes is the application performing on a normal day, peak day, and slow day? What is the performance of the network link? What other traffic is on the network link? How resilient is the link?

What servers and storage platforms does Symantec recommend?

Symantec does not make specific server hardware recommendations. Symantec does recommend that there be sufficient performance headroom for all of the work that the server has

to perform. In the cases where the more sophisticated Veritas Cluster Server failover models (N+1, N:M, etc.) are used, it is important that the servers are sized to handle the expected loads in event of a failover. Servers, storage, SAN, and network equipment should be built to eliminate single points of failure.

Symantec does not make specific storage recommendations either. Clearly, if a company is considering an intelligent application disaster recovery solution, Symantec would suggest that they invest in a high-performance tier 1 class of storage hardware that offers many high availability and data protection features at the storage level, but this is not required.

What does Symantec suggest for application recovery in the event of disasters? Is it completely automatic?

For companies whose business needs can justify a second data center for the purposes of disaster recovery, Symantec would recommend a combination of Veritas Cluster Server HA/DR, and some kind of data replication. As a best practice, Symantec recommends operator intervention when failing over from the primary data center to the disaster recovery data center. Completely automated failover is advisable only in specific Metro Cluster configurations. Most companies want to confirm that they do indeed have to move their operations from the primary to the disaster recovery data center. Once that confirmation has been given, the rest of the failover operation is automatic.

What does Symantec suggest/recommend in terms of validating and testing the disaster recovery environment without disrupting the production environment?

Veritas Cluster Server HA/DR Fire Drill allows companies to validate/verify their intelligent application disaster recovery environment. The recovery environment should function as expected, once everything comes up clean in the Fire Drill test. An important feature of Veritas Cluster Server HA/DR Fire Drill is that it is non-disruptive to production operations, so that it is easy to initiate or even schedule regular Fire Drill tests.

What does Symantec recommend to insure that the disaster recovery environment is operating as expected?

Disaster recovery best practices say that the disaster recovery plan must be tested, revised, and updated regularly. Symantec recommends that companies take this very seriously and test their disaster recovery environment regularly in accordance with the needs of their business. There is no sense in investing the time, money, and effort to build a disaster recovery capability, if it is not routinely verified. With Veritas Cluster Server HA/DR Fire Drill, organizations can easily test the recoverability of mission critical applications on a weekly basis.

Conclusion

The intelligent application recovery solutions available from Symantec have been described in this paper from a high-level perspective. Veritas Cluster Server has been acknowledged within the industry as a leader in high availability and disaster recovery because of the flexibility that has been highlighted. As such, it has become the standard in many data centers around the world.

This document has not focused on how to maximize server utilization, but Veritas Cluster Server allows high availability designers to do just that by building multi-node clusters with appropriate failover models. The days of active/passive two-node clusters can be a thing of the past, with the proper Veritas Cluster Server failover model (active/active, N:1, N:N).

The proposed intelligent application recovery options are all based on Veritas Cluster Server and Veritas Cluster Server HA/DR. This means that they do not require any additional specialized training or skills if Veritas Cluster Server is already running in-house. For new users, Veritas Cluster Server is very easy to set up and run. Initial clusters can be built within a matter of a few hours, once the server hardware, storage, and network infrastructure are in place. Administrators and operators already trained in Veritas Cluster Server can manage all of the cluster types described above, including Global Clusters.

The intelligent application recovery solutions presented here are based on the single flexible multi-site architecture of Veritas Cluster Server. This architecture can handle the local high availability needs within a data center, all the way up to intelligent application recovery between major data centers in different parts of the world. The architecture is dual-use in nature, providing application availability over a wide range of distances.

Veritas Cluster Server HA/DR even provides a mechanism to test and verify the intelligent application recovery infrastructure, without disturbing the production environment. With this ability, customers will have more confidence when they actually perform disaster recovery tests, which should be done regularly as a best practice.

The majority of the options described take advantage of data replication, which is an integral part of an intelligent application recovery solution. Veritas Cluster Server HA/DR tightly integrates with third-party replication offerings, just as it does with Veritas Volume Replicator, protecting an organization's investment in those solutions. The most up-to-date information on Veritas Cluster Server HA/DR replication agents may be found at http://eval.veritas.com/mktginfo/products/Sales_Docs/High_Availability/agent_list_partner.xls.

Symantec looks forward to working with your organization's disaster recovery team to evaluate your disaster recovery needs, plan your response to the risks that your enterprise faces, implement the proper intelligent application recovery solutions for your business needs, and test your recovery environment.

About Symantec

Symantec is a global leader in infrastructure software, enabling businesses and consumers to have confidence in a connected world. The company helps customers protect their infrastructure, information, and interactions by delivering software and services that address risks to security, availability, compliance, and performance. Headquartered in Cupertino, Calif., Symantec has operations in 40 countries. More information is available at www.symantec.com.

For specific country offices and contact numbers, please visit our Web site. For product information in the U.S., call toll-free 1 (800) 745 6054.

Symantec Corporation
World Headquarters
20330 Stevens Creek Boulevard
Cupertino, CA 95014 USA
+1 (408) 517 8000
1 (800) 721 3934
www.symantec.com

03/07 12114889