WHITE PAPER: TECHNICAL

symantec.
Confidence in the connected world.

Symantec High Availability and Disaster Recovery Solution for SAP

Keep your SAP application online, all the time with Symantec

Venkata Reddy Chappavarapu
Sr. Software Engineer
Data Center Management Group
Venkatareddy_Chappavarapu@Symantec.com

Table of Contents

Table of Figures

Introduction

Many organizations rely on SAP applications to support vital business processes. Any disruption of these services translates directly into bottom-line losses. As an organization's information systems become increasingly integrated and interdependent, the potential impact of failures and outages grows to enormous proportions.

The challenge for IT organizations is to maintain continuous SAP application availability in a complex, interconnected, and heterogeneous application environment. The difficulties are significant:

- there are many potential points of failure or disruption
- the interdependencies between components complicates administration
- the infrastructure itself undergoes constant change

To gain additional competitive advantage, enterprises must now work more closely together and integrate their SAP environment with those of other organizations, such as partners, customers, or suppliers. The availability of these applications is therefore essential.

There are three main availability classes, depending on the degree of availability required:

- Standard Availability – achievable availability without additional measures
- High Availability – increased availability after elimination of single points of failure within the local datacenter
- Disaster Recovery – highest availability, which even overcomes the failure of an entire production site

Symantec helps the organizations that rely on SAP applications with an integrated, out-of-the-box solution for SAP availability. Symantec's High Availability and Disaster Recovery solution for SAP enhances both local and global availability for business critical SAP applications.

- ***Local high availability*:** By clustering critical application components with application-specific monitoring and failover, Symantec's solution simplifies the management of complex environments. Administrators can manually move services for preventative and proactive maintenance, and the software automatically migrates and restarts applications in the case of failures.

- ***Global availability/disaster recovery*:** By replicating data across geographically dispersed data centers and using global failover capabilities, companies can provide access to essential services in the event of major site disruptions. Using Symantec's solution, administrators can migrate applications or an entire data center within minutes, with a single click at a central console. Symantec's flexible, hardware independent solution supports a variety of cost-effective strategies for leveraging your investment in disaster recovery resources.

The Symantec High Availability and Disaster Recovery solution for SAP utilizes the following products: Veritas Storage Foundation™, Veritas Storage Foundation Cluster File System™, Veritas Cluster Server™ HA/DR, Veritas Volume Replicator™, and Cluster Server agents that are designed specifically for SAP applications. The result is an out-of-the-box solution that you can quickly deploy which protects critical SAP applications immediately from either planned or unplanned downtime.

This paper describes the overall SAP availability environment and Symantec's solution for High Availability and Disaster Recovery. It then offers a technical overview of both local and global availability in a SAP environment, and steps through a number of scenarios using the Symantec solution.

The challenges in protecting critical SAP applications

Symantec recently conducted a global survey of large enterprises regarding their disaster recovery plans. The application types they deemed most critical to their businesses, was their Enterprise Resource Planning (ERP) software, their databases, and their web applications. These are the key components of a SAP environment. This isn't surprising given the close integration that SAP applications have to its customers' critical business processes.

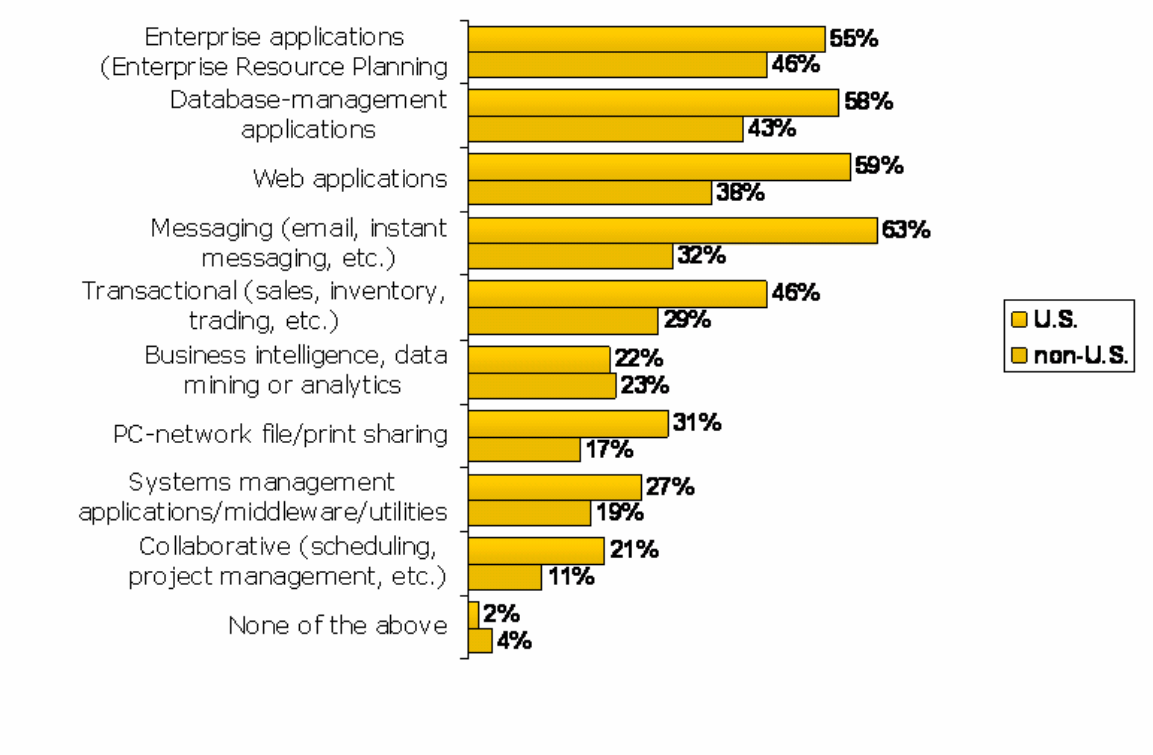

Figure 1: Mission critical applications

SAP applications run vital business processes at companies around the globe. Having made significant investments in the application environment as well as business process changes, these companies now struggle to meet demanding service level agreements (SLAs) for SAP applications. In this same survey, Symantec found that nearly 48% of companies have had to execute their disaster recovery plan in the last year. Thus, explains the urgency that IT organizations have In seeking solutions that will make their SAP environments more available.

Maintaining optimal performance and availability for critical SAP applications is increasingly difficult for many IT departments, because the IT environment itself is complex, dynamic, and inherently unstable. Several factors combine to make the availability of applications both costly and challenging.

The SAP application environment is complex and heterogeneous

An SAP application environment may include the following services: a database server, a central services instance, a central instance, an enqueue replication server, and one or more dialog instances. It is essential that the data center staff understand the dependencies among these application services. For example, the database must be running before the central services instance can be started.

Between the application and the end user, however, there are several other application services, such as web servers, load balancers, and the logical IP address infrastructure. These services may use different operating systems and equipment from multiple vendors.

Because application delivery depends on the availability of the total environment, the operations team must be able to manage and maintain, or at least troubleshoot, application servers, operating systems, storage systems, network infrastructure, and databases. Most often these skills are distributed through a number of different people, possibly in different organizational units. This situation leads to confusion and the "blame game" starts when problems occur. To reduce the administrative costs while improving availability, IT organizations need a single interface for managing all services of the application environment, across different platforms and databases.

N-tier applications have many potential points of failure

To access the application, the user needs the entire application infrastructure to be working. As the number of tiers within an application's architecture increase, the number of points of failure that must be either eliminated or mitigated also increases.

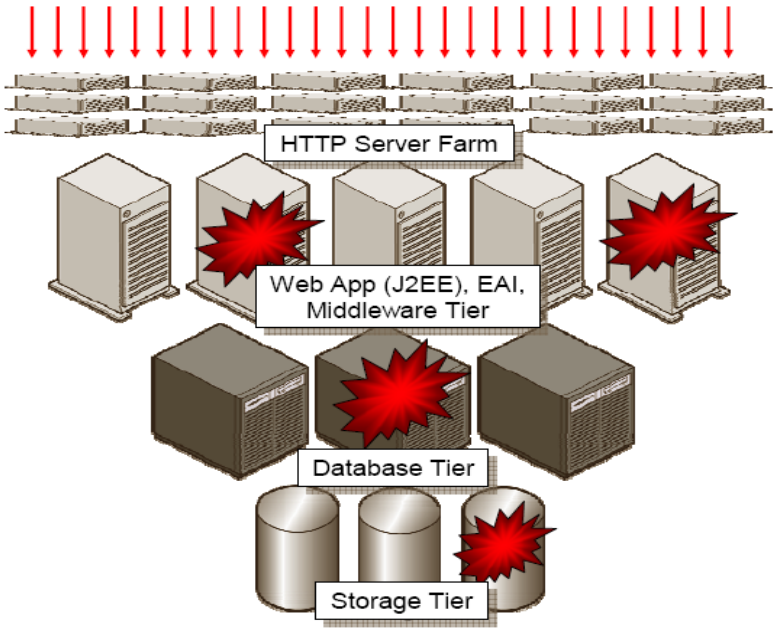

Figure 2: N-tier application architecture

IT organizations typically focus on protecting specific, essential components of the infrastructure, such as the database. But protecting the database alone is not effective if the network router fails. Applications depend not only on the other application services (such as the database and application servers), but also on the infrastructure components, such as the disks and volumes, NFS shared mounts, virtual IP addresses, network cards, and so on. Building redundancy at a component level at each layer increases the complexity of the IT environment – and hence the cost of managing that environment.

IT organizations need an end-to-end approach in protecting SAP application availability that incorporates all of the essential application tiers, services, and components.

The application infrastructure is in constant change

Application availability is further imperiled by the fact that the application infrastructure is in constant change – and change introduces instability.

Most companies carefully manage and monitor change to vital applications themselves, with processes in place to restrict, approve and test any changes to the application and the underlying database. But the infrastructure on which the application depends is made up of so many different levels and components and is essentially a moving target, with:

- Firmware updates
- OS patches
- Capacity upgrades
- Preventative maintenance on storage hardware
- Driver updates

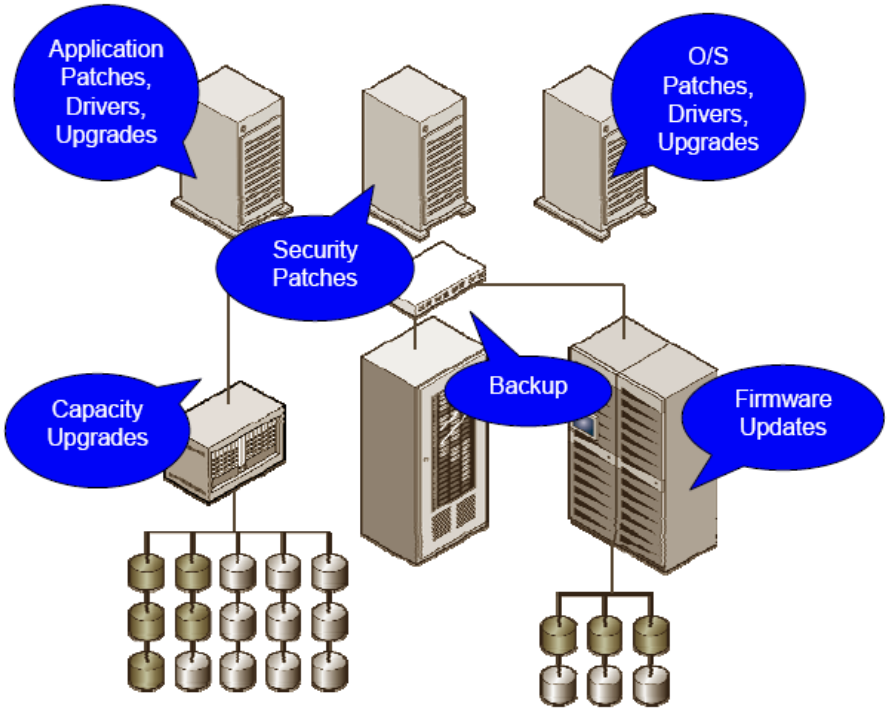

Figure 3: Infrastructure changes diagram

These changes need to be tested before being applied to the production environment. These changes are necessary and usually beneficial. However, every change introduces potential for error and system instability.

Downtime is not an option

The issues described in the preceding paragraphs would not be a serious problem if mistakes did not matter, and if you could take the whole application offline for hours or days to fix any problems.

SAP applications, however, are usually critical to an organization's daily operations, making it difficult to find time to perform disruptive work. Thus, it is essential to respond instantly to any problems that arise.

- With global operations, the few hours that operations staff once had in the middle of the night are now business hours somewhere around the globe, making planned downtime increasingly scarce and expensive.
- Unplanned downtime may interrupt critical business processes and cost companies thousands of dollars or more per hour.

In an inherently volatile environment, IT organizations need the ability to upgrade, test, maintain and deploy infrastructure components without disrupting operations. They need the time to maintain applications, and correct problems as they occur, without interrupting application access.

The growing interdependence of business systems increases risk

With the increasing adoption of enterprise application integration (EAI) initiatives, companies are linking vital applications together with middleware to create a service-oriented architecture (SOA).

The unintended effect of this architecture is the increased risk of an outage in case of an application failure. When a critical application is tightly linked with other applications in a composite architecture, the loss of the application affects wider service availability.

Applications are essential to business continuity

In addition to the moment-to-moment and day-to-day availability concerns, IT organizations have to also figure out ways to provide continued access to critical applications in the case of a disaster. Most critical applications must have the ability to run at an alternate, geographically-separate location from the primary production environment.

Alternate data centers are costly. The money spent on disaster recovery facilities comes from a fixed pool of resources available for IT and infrastructure. For many companies, disaster recovery planning entails identifying a "cold" recovery site that requires an initial investment and ongoing maintenance to remain current. In the case of a disaster, recovery may require significant time and manual intervention. By taking advantage of existing resources and multi-purposing equipment, it is possible for companies to build "hot" recovery sites that can provide nearly instant recovery without manual intervention.

Businesses face escalating risks and costs

In a complex application environment, the risk of application downtime escalates. To mitigate this risk, companies are increasing their expenditures on equipment for redundancy and operational staff. By increasing the complexity of the infrastructure, IT organizations are increasing training and administrative costs.

To reduce the cost and risks of maintaining SAP application environments, companies need an end-to-end solution for application availability that lets them manage the entire application environment with a single interface. They need the ability to manage change proactively, swap or upgrade components without disrupting user access to the application. For business continuity purposes, they need the ability to switch applications to alternate sites, and leverage alternate site investments.

Symantec has an integrated solution for SAP that solves these availability problems. Symantec's High Availability and Disaster Recovery solution reduces the risks and costs of maintaining SAP applications in support of critical business processes.

Symantec's solution for ensuring SAP availability

Symantec offers an end-to-end, fully integrated solution for ensuring highly available SAP environments. Symantec's solution reduces planned[†] and unplanned downtime, simplifies the administration of the complex environment with a single interface, and supports global failover for disaster recovery purposes.

Symantec's High Availability and Disaster Recovery solution for SAP combines Symantec's industry-leading, hardware independent software for storage management and availability with a deep understanding of the SAP application environment and its essential components.

Symantec/Veritas and SAP have worked together since 1997 to optimize the availability, performance and automation of SAP environments. As an SAP Software Partner, Symantec ensures the performance and interoperability of the Symantec solutions in the SAP environment. Symantec's High Availability and Disaster Recovery solution has undergone the SAP documented integration tests, and has been proven in real-world environments with stringent availability requirements.

Symantec's High Availability and Disaster Recovery solution components

Symantec's High Availability and Disaster Recovery solution integrates the following software products:

Veritas Storage Foundation combines Symantec's industry-leading file system and volume management solutions to create a highly available, robust foundation for SAP data. The journal file system restarts in seconds for fast failovers. Logical volumes support highly available, high performance storage configurations. Database-specific components such as direct I/O accelerate database read and write performance while simplifying the manageability of database data. Storage Foundation provides database-specific optimizations for Oracle, DB/2, Sybase, Microsoft SQL Server, and Oracle RAC databases.

Veritas Storage Foundation Cluster File System (CFS) builds upon Symantec's industry-leading file system to provide a solution which allows parallel access to data across all members of a cluster. Because the file system can be mounted on all nodes in a cluster, the time normally required to mount a file system in the event of a failover is eliminated. This can have a dramatic effect on improving failover times. Cluster File System provides cache coherency and POSIX compliance across nodes, so that data changes are atomically seen by all cluster nodes simultaneously. This benefits applications within the SAP environment that require a common set of binaries and/or configuration files.

[†] Planned downtime can be reduced as failover systems could be used to "free" nodes for maintenance operations.

Veritas Cluster Server (VCS) eliminates planned and unplanned downtime by clustering critical applications and the resources they require. Specific agents for SAP, the underlying database, and the file server ensure that all of the critical components of your SAP environment are monitored and managed centrally to ensure maximum application availability.

Veritas Cluster Server HA/DR is a Cluster Server package that monitors and controls multiple, geographically-distributed VCS clusters as well as replication occurring with Veritas Volume Replicator or third-party replication technologies. Veritas Cluster Server HA/DR lets administrators migrate all of the applications in a data center with a single click, enabling companies to survive serious local disruptions without significant interruptions to critical services.

For data replication, the following choices can be made:

Veritas Volume Replicator delivers reliable, storage-independent replication over any IP network, providing a critical component of a rapid disaster recovery configuration. The product replicates data at the logical volume level, and ensures data integrity and reliability during replication.

Or

Veritas Cluster Server agents for Hardware Replication monitor and manage the state of the replicated devices that are attached to VCS nodes. These agents ensure that the system on which the replication resources are online has safe and exclusive access to the configured devices. These agents can be used in single data clusters and multi-cluster environments that are setup using Veritas Cluster Server HA/DR. You can use any of the available hardware based agents from Symantec depending on the array used. Symantec provides agents for most array-based replication technologies including Hitachi TrueCopy, IBM MetroMirror, EMC SRDF, and EMC MirrorView. These agents come with the Veritas Cluster Server HA/DR license.

Veritas Cluster Server agent for SAP NetWeaver[*] starts the SAP application during online, stops the SAP application during offline, monitors the SAP applications for critical processes, and cleans the environment for SAP applications in case of any issue. This agent comes with the Veritas Cluster Server license.

Veritas Cluster Server agent for DNS updates the Domain Name Server with the virtual hostname and IP for SAP applications for DNS based SAP configurations for local or remote cluster failover. This agent comes with the Veritas Cluster Server license.

Complete descriptions of these products can be found at the Symantec web site, www.symantec.com.

[*] SAP NetWeaver provides an open integration and application platform and permits the integration of the Enterprise Services Architecture. You can standardize business processes across technological boundaries, integrate applications for your employees as needed, and access and edit simple information easily and in a structured manner.

How components work together

The different components of the Symantec solution work together and help IT organizations improve SAP application availability on a daily basis while offering significant protection from the loss of service and data in the case of a disaster or regional disruption.

Local Availability

Veritas Storage Foundation and Veritas Cluster Server, which are bundled together as Storage Foundation HA, create a highly available, local data and application availability environment that helps you do the following:

- Manage complexity by providing a single interface for starting, stopping, monitoring and maintaining SAP application services[§].
- Manage change by proactively moving application services to enable dynamic maintenance and testing.
- Improve availability with automated, application-specific monitoring and failover when problems do occur, and fast reconfiguration when the problem is resolved.
- Consolidate servers and make better use of resources through local clustering and the virtualization of critical services.
- Improve storage utilization and administrator productivity through policy-based storage management.

Global Availability/Disaster Recovery

When you add Veritas Cluster Server HA/DR and Veritas Volume Replicator, i.e., Storage Foundation HA/DR with the Volume Replicator option, to the environment, you gain the ability to replicate SAP data to another location and move critical SAP applications within minutes to a data center across the globe with a single click. Because the secondary disaster recovery site does not need to match the primary site exactly and can be used for other purposes, companies can leverage disaster recovery investments and control the costs of supporting disaster recovery by:

- Using lower-cost or lower-capacity components at the off-site recovery location, with the understanding that you will accept degraded application performance at the recovery site.
- Using a single data center as an off-site recovery location for multiple data center locations.
- Using the recovery site to run non-critical services, such as development and testing, that can be interrupted, stopped and replaced with critical applications in case of a global failover.

[§] See Appendix for supported SAP services.

Architecture of SAP systems

Any availability solution for SAP must address the specific availability requirements and constraints of the SAP system environment. This section outlines a few of the SAP-specific issues affecting application availability.

SAP System Components

An SAP application instance has multiple services which are typically deployed across multiple servers. SAP identifies the following services as critical to the application environment, representing potential single points of failure:

- Database instance
- Central Instance
- Messaging service
- Enqueue service
- Enqueue Replication service
- Shared file system (either a cluster file system, Network File System - NFS, or Common Internet File System - CIFS)

Both the Enqueue and Message service typically reside in the SAP central instance (CI). In more advanced configurations, the Enqueue service runs independently and is backed up by a separately operating Enqueue Replication Server (ERS). In most large SAP installations, the database uses a separate server for performance purposes. The file server is typically separate as well.

SAP applications typically have more than one dialog instance in addition to the CI. Because dialog instances are redundant, they do not present a potential single point of failure. However, the loss of a dialog instance in a heavily-loaded system can significantly degrade performance, and the dialog instance is a key part of a global failover scenario.

Dependencies

The different components of SAP have strict dependencies in the order in which they are restarted in case of a failure.

For example:

- In most cases the shared file system should be available before the database and CI start.
- The database server must be available before the CI is started.
- The CI must be online before other dialog instances are brought online.

Database layer

SAP environments may use any of the relational database management systems as the database server. Administering the database itself requires experience in the specific database used, such as DB/2, Oracle, or Microsoft SQL Server.

The database represents a single point of failure for the SAP system, unless the database itself is clustered. The Symantec's High Availability and Disaster Recovery solution has database-specific components for each these databases.

Application layer

With mySAP.com Technology, SAP provides a proven, scalable, fault-tolerant, multi-tier architecture. The individual components can be protected either by horizontal scalability – that is, the use of multiple components that tolerate the failure of individual components – or by cluster and switchover solutions, such as Veritas Cluster Server.

SAP Web Application Server

With SAP Web Application Server (SAP Web AS), SAP enables web applications to be directly supported by the application server and combines a ABAP** and Java execution stack in one infrastructure.

The Internet Communication Manager (ICM) has also been implemented as another new process in the application server framework. It enables communication between the SAP Web AS and external partners using Internet standard protocols, such as HTTP, HTTPS, SMTP, SOAP, and the Java communication services. The SAP Java Connector (SAP JCo) enables method calls between Java applications and ABAP applications. The following diagram shows how SAP JCo connects an ABAP system and a Java system:

** ABAP – Advanced Business Application Programming is a high level programming language created by SAP.

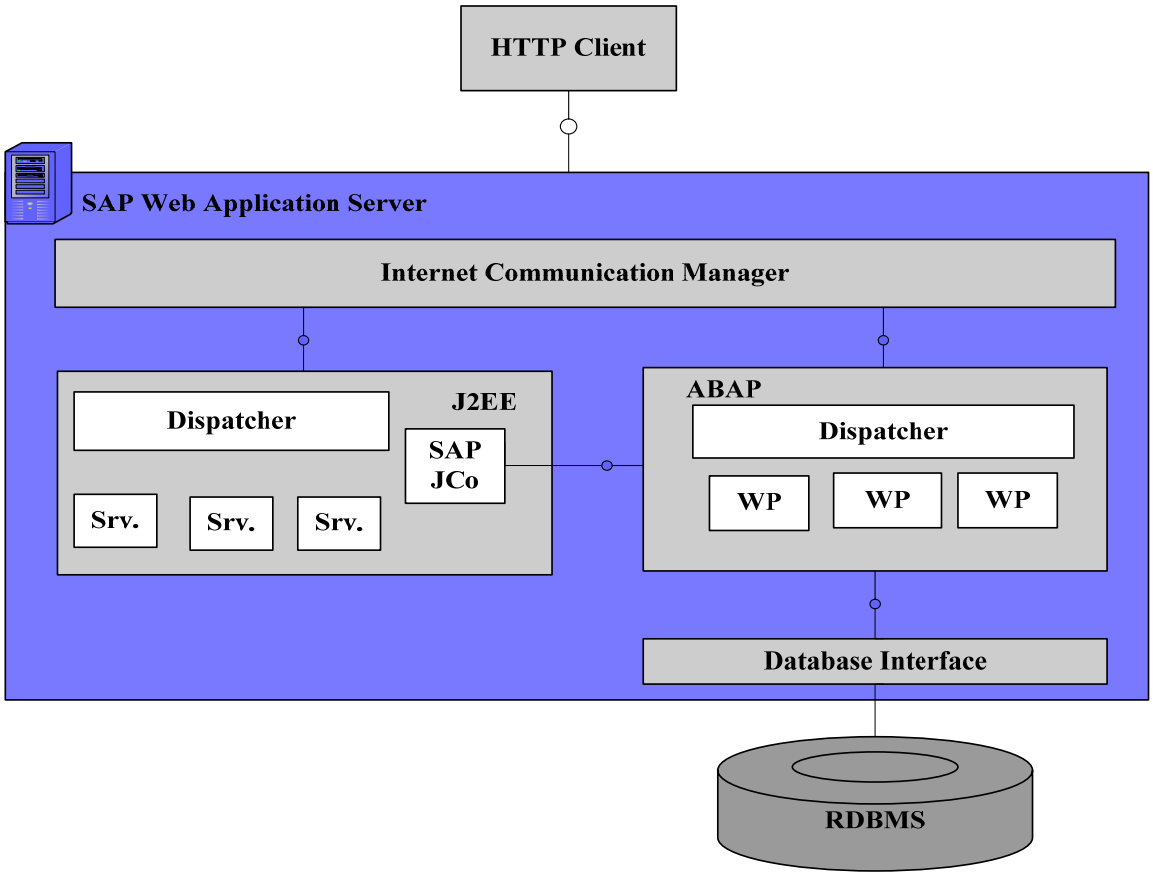

Figure 4: SAP Web Application Server architecture

An SAP application consists of one or more instances of an application server. Each instance can run on a separate server, but it is also possible to operate multiple instances on one host. An SAP instance can provide different service types. The standard SAP services that can be configured on all instances of the SAP component are dialog, batch, update, and spool. The failure of an SAP instance on which only these standard services are configured causes all the transactions processed on it to be terminated and rolled back. Database consistency is guaranteed at all times. Terminated transactions can be repeated on one of the instances still available.

Apart from the standard services, there are two other services that exist only once. They are supplied by the message service and the enqueue service and are potential single points of failure. The SAP instance with both these services is called the central instance (CI).

In the case of the standalone enqueue service, the enqueue service is operated separately from the other SAP services and independently of the other application servers. This technology can be implemented for SAP components using kernels as of Release 4.6D, depending on availability. The message server can also be operated separately from the other application servers. SAP supports separating both these critical components into one instance and can be restarted faster than a complete application server in case of failure.

Message Server

The Message Service keeps a list of dispatchers and server processes. It represents the infrastructure for data exchange (small datasets only) between the participating nodes. The message service also supplies information to the SAP Web Dispatcher about Load Balancing between many ABAP and Java Instances.

The following diagram shows how the message server helps the SAP Web Dispatcher with load balancing information:

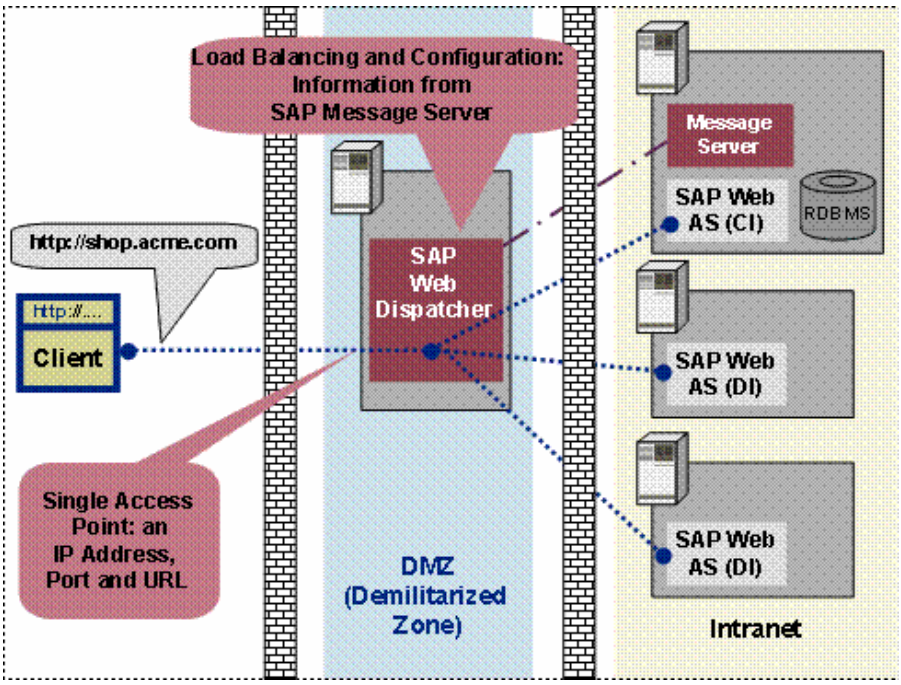

Figure 5: SAP Message Server architecture

To avoid the single point of failure with the message server, the message server must be restarted along with the Enqueue Server quickly. Both can be restarted automatically using cluster software, such as Veritas Cluster Server.

Enqueue Service

The Enqueue service is a critical component of the SAP system. It administers locking using objects within SAP transactions that can be requested by applications to ensure consistency within the SAP system.

Since the lock table is held in the main memory of the Enqueue Server, server failure without additional backup measures results in loss of the locks held. To maintain consistency, all open transactions are rolled back after the Enqueue Server is restarted.

Enqueue Server Failure and Consequences

An Enqueue Server failure has significant implications on the SAP systems. The following steps need to be taken in case of Enqueue Server failure without additional back-up measures:

ABAP Stack:

- Enqueue table contains only locks on data objects, locks are session bound
  - Loss of enqueue service → Enqueue table gets lost
    - → Concerned sessions need to be rolled back (done automatically by SAP)
    - → No restart of whole software cluster necessary
  - Enqueue replication prevents session rollback due to enqueue server failover

Java Stack:

- Enqueue table contains data objects locks plus infrastructure locks (system, not session bound)
  - Loss of enqueue service → Enqueue table gets lost
    - → Concerned sessions need to be rolled back
    - → Due to infrastructure locks in an ambiguous state restart of all J2EE-instances is required!
- Since NetWeaver 04 SPS15 (and also in NetWeaver 2004s SR1):
  - It is strongly recommended to use Enqueue Replication Server (ERS) to replicate the lock table to another cluster node
  - Losing the lock table (e.g. SCS switchover without Enqueue Replication Server) will be detected by the application lock service
  - In this case the cluster state is considered to be nondeterministic
  - A restart of all instances will be enforced by the respective JControl process

The Enqueue Replication Server resolves this challenge.

Enqueue Replication Service

The Enqueue Replication service enables the lock table to be replicated on a second server, the replication server. A copy of the lock table is maintained on this server. If the Enqueue service fails, a new Enqueue service is started on the Replication Server using a failover solution (clustering software from a partner like Symantec) and this replication service creates a new lock table from the copy of the lock table. This enables the Enqueue service, and therefore the whole SAP component, to continue operating almost without interruption. If the Enqueue service fails, transactions are no longer terminated, so that work can be continued transparently.

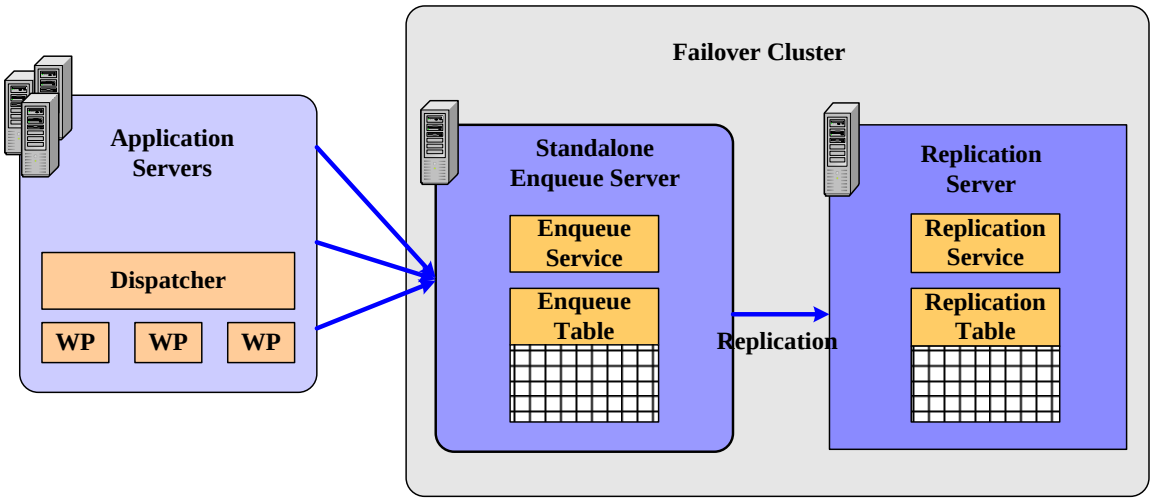

Figure 6: SAP Enqueue and Enqueue Replication Server diagram

Symantec's Veritas Cluster Server provides the cluster technology for the Enqueue Server and its Replication Server, which is required for the Enqueue service to operate without interruption.

SAP Web Dispatcher

The SAP Web dispatcher is recommended when you use an SAP system with several SAP Web Application Servers for Web applications. The SAP Web dispatcher is a program that you can run on the computer that is connected directly to the Internet.

The SAP Web dispatcher lies between the Internet and your SAP System. It is the entry point for HTTP(s) requests into your system, which consists of one or more Web application servers. As a "software web switch", the SAP Web dispatcher can reject or accept connections. When it accepts a connection, it balances the load to ensure an even distribution across the servers.

You can use the SAP Web dispatcher in Add-In systems and in pure Java systems, as well as in pure ABAP systems. It is also beneficial to use the SAP Web dispatcher, if you do not need security functions (entry point in the DMZ, SSL, URL filtering), but you simply want to balance the load between several SAP Web AS instances.

Benefits of Symantec's High Availability and Disaster Recovery solution for SAP

Symantec's High Availability and Disaster Recovery solution supports SAP applications for both local availability and global availability. Symantec's solution supports different deployment configurations for SAP applications. SAP recommends the following HA set-ups for SAP NetWeaver applications[††].

HA-Setup *), **)	SAP Rating SAP NetWeaver AS Java	SAP Rating SAP NetWeaver AS Add-In***)	SAP Rating SAP NetWeaver AS ABAP***)
1) DB only in switch-over group	POSSIBLE	POSSIBLE	POSSIBLE
2) CI, ASCS/JSCS and DB in one switch-over group	POSSIBLE	POSSIBLE	POSSIBLE
3) CI and ASCS/JSCS in one switch-over group, DB in another	POSSIBLE	POSSIBLE	POSSIBLE
4) DB and ASCS/JSCS in one switch-over group	POSSIBLE	POSSIBLE	POSSIBLE
5) DB and ASCS/JSCS, each in its own switch-over group	RECOMMENDED	RECOMMENDED	RECOMMENDED

Figure 7: SAP HA Setups Table

* Prerequisite: Separation of SCS for ABAP and JAVA

** Each HA-Setup should be extended by the Enqueue Replication Server

*** In an ABAP only or Add-In installation, it is recommended to separate an ASCS instance

Symantec's high availability and disaster recovery solution provides:

- **Local failover of SAP systems for high availability**
- **Global failover of SAP systems for disaster recovery**

The solution is discussed in detail in the sections that follow.

[††] SAP NetWeaver is the basis for all SAP solutions on a given hardware. mySAP Business Suite is one example of a solution that uses all key capabilities of SAP NetWeaver.

Local failover for data center high availability

This section describes how you can configure an SAP system for local high availability.

Setting up the cluster

Veritas Cluster Server allows for highly flexible, scalable clustering configurations. For an SAP environment, you will probably want to include the following system components in a VCS cluster:

- SAP central services instance
- The CI server
- Enqueue Replication Service
- Any dialog[##] or application server instances
- The server running the database
- The file system server

Although the dialog instances do not represent single points of failure, the loss of an instance can degrade application performance. If you plan to implement global failover for disaster recovery, you will likely need to include the dialog instances in the VCS cluster configuration even if they do not necessarily failover. Finally, including dialog instances in the local cluster simplifies dialog server administration, allowing operators to use the same interface to start, stop, monitor and manage these additional SAP services.

Symantec strongly recommends configuring the dialog instances under the cluster software for better manageability and performance. Customers can cluster as many dialog instances as needed to service the basic functions after failover and the necessary dialog instances (not so important) for additional connections.

To provide transparent failover, you will need to implement shared data between the cluster nodes, typically with SAN-based storage such as Veritas Cluster File System. The cluster configuration requires redundant private Ethernet connections between nodes to support cluster heartbeat communications.

You can add other servers as well, to consolidate management of a high availability environment and provide spare capacity for critical applications. VCS clusters can support up to 32 nodes, allowing the clusters to be expanded for maximum server efficiency.

Veritas Cluster Server supports both active/passive and active/active configurations, meaning that you do not need to dedicate spare, unused processing capacity for each critical server. Many customer sites maintain an n+1 clustering environment, with one spare or a little-used server providing failover capacity for other servers in the cluster.

It is also possible to have an entirely active/active (n-to-m) cluster, using multiple application groups on multiple servers, each group capable of being failed over to different servers in the cluster. The cluster only needs sufficient spare capacity available on production servers to absorb the applications of a failed system.

[##] Dialog servers do not need to be included in a VCS cluster configuration for local failover. But if you plan to implement global failover, you will likely need to include the dialog instances in the VCS cluster configuration.

Configuring the VCS agent for SAP NetWeaver

Veritas Cluster Server uses application-specific agents to start, stop, monitor, and switch over different applications and infrastructure components. A Veritas Cluster Server environment running SAP uses VCS agents to monitor and track the SAP system, the database, and the file server.

The Veritas Cluster Server agent for SAP NetWeaver starts, stops, and monitors essential SAP application services, including the:

- Enqueue Service (running within a Central Instance or within a Central Services Instance (SCS))
- Message Service (running within a Central Instance or within a Central Services Instance (SCS))
- Enqueue Replication Service (ERS)
- Dispatcher and Worker Processes in both Dialog and Central Instances
- Collector and Sender (CO & SE) processes in both Dialog and Central Instances
- Gateway processes in both Dialog and Central Instances
- Internet Graphics Service (IGS) process in both Dialog and Central Instances
- SAP OS Collector process

The agent provides the ability to start, stop, and monitor each of these SAP services. Additionally, if a service fails, the agent will clean the node for any remaining system processes and shared resources.

Veritas Cluster Server provides multiple levels of monitoring for SAP services:

- The first (default) level confirms the existence of essential processes in the process table.
- The second level, which is optional and additional to the first, runs SAP supplied tools to check the health of the SAP systems.
- The third level, which is also optional, invokes an external monitoring program, allowing the user to provide custom scripts for application monitoring.

The following screen picture shows a SAP Central Services (SCS) Instance configuration in a VCS cluster environment.

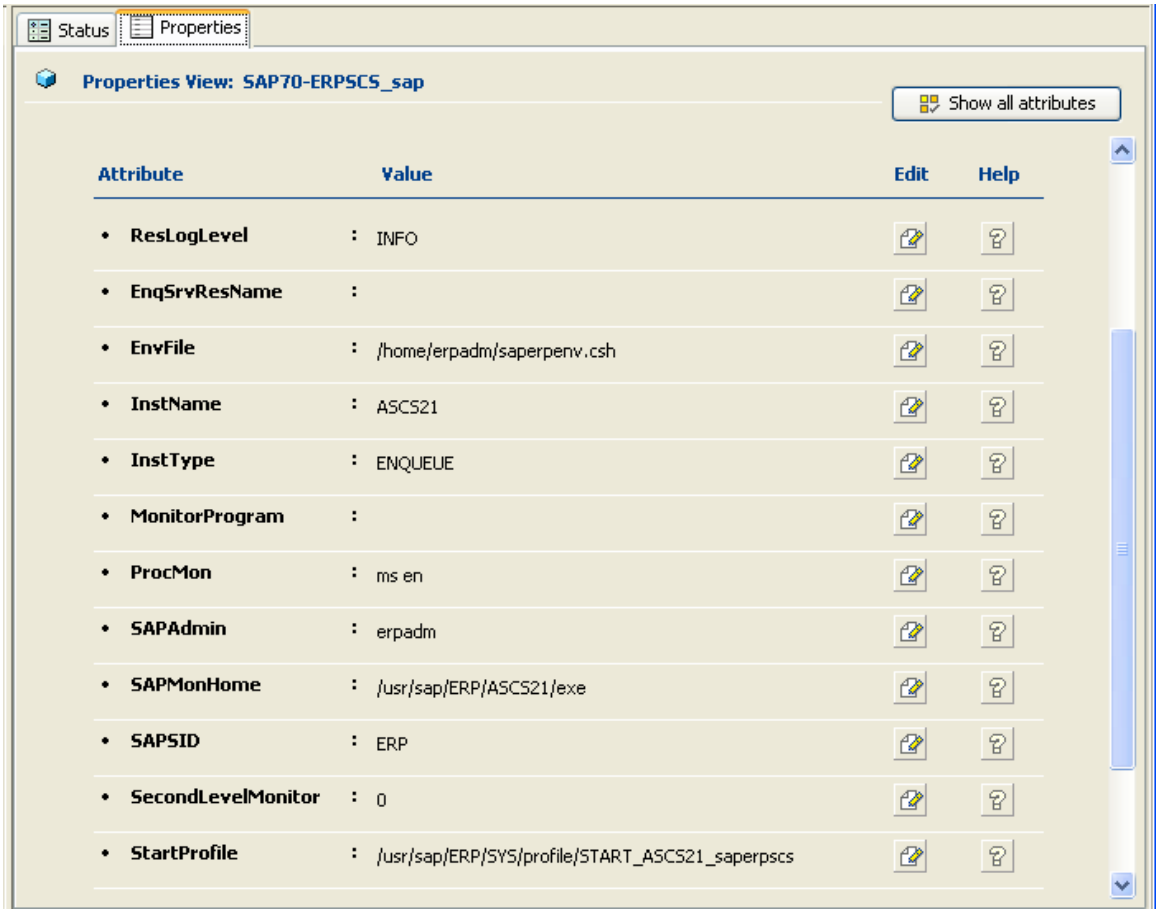

Figure 8: Veritas Cluster Server configuration for SCS Instance

Creating service groups

Veritas Cluster Server provides application failover by encapsulating the resources required for each application into a *service group* – creating virtualized application services that can be moved among cluster nodes. Operations staff can operate on the cluster itself, on the service group (starting, stopping, switching over, etc.), or on the specific resources within the service group. A VCS service group is the smallest unit of failover.

Each SAP service group contains a set of dependent resources – the lower-level components that an application requires to operate successfully. Resources include disk groups, disk volumes, file systems, IP addresses, and dependent application processes. The resources within a service group have dependencies which define the start and stop order that Veritas Cluster Server uses to bring the service group online and offline, respectively.

Veritas Cluster Server starts, stops, monitors and switches service groups on any server in the cluster in response to server or resource faults. In addition, an administrator can proactively move a service group between cluster nodes to perform preventative maintenance or apply patches. The service group includes logic about the dependencies between application components.

For example, the following figure illustrates the relationship between the Cluster Server resources required to support a SAP central services instance (assuming that storage is based on NFS):

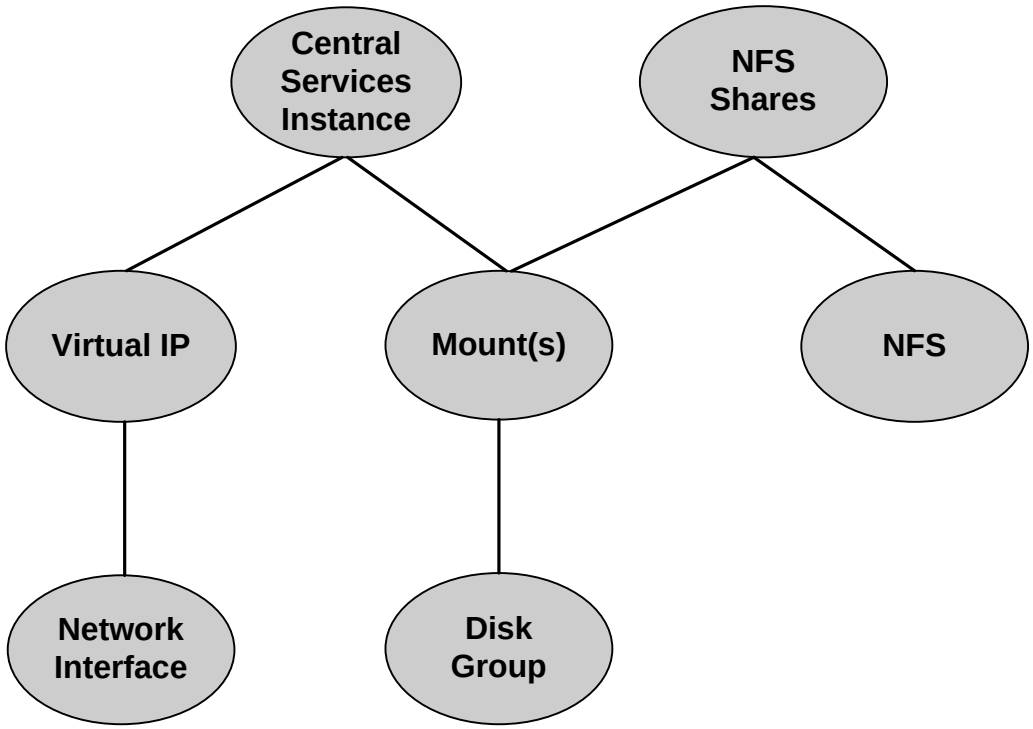

Figure 9: General VCS Service Group diagram for SAP

The resource for the SAP Central Services instance depends on the Virtual IP address and the file system mounted by the mount agent. The mounted file system is shared to allow access to other SAP application servers running on different systems.

Disk Groups

Each service group requires a dedicated file system, volume and disk group to store the service group's data and programs. By importing and deporting this set of storage objects on different servers in the cluster without affecting other service groups, Veritas Cluster Server allows the service groups themselves to be independent of the underlying architecture and mobile across the cluster.

For example, when Veritas Cluster Server shuts down a service group, the resources are shutdown in the following hierarchy:

- As the file system resource is shutdown, the file system is un-mounted

- As the volume is shutdown, the volume is stopped
- As the disk group is shutdown, the disk group is deported.

Veritas Cluster Server initiates startup of the service group on another system in the cluster where each resource is started in dependent order. On the new system, the disk group will be imported, the volume started and the file system will be mounted. This entire process happens automatically in the event of a failure, significantly reducing the downtime associated with a failure or outage. This sequence could not be accomplished if the SAP application component was installed on a local system disk on one node in the cluster.

By deploying the Veritas Storage Foundation Cluster File System, the steps outlined above can be eliminated. While each service group still requires a dedicated file system, volume, and disk group, by deploying Cluster File System, the file system, disk, and volume groups will already be started on all nodes in the cluster. Therefore, there is no need to unmount the file system, shutdown the volumes, and deport the disk groups. Additionally, there is no need for Veritas Cluster Server to start the corresponding groups on the other node, as they will have already been started. Cluster Server will still have the role of ensuring Cluster File System has been started and is running on all of the nodes in a cluster.

Network Resources

In addition to disk groups, applications also require specific network resources, such as Network Interface Cards (NICs) and IP addresses.

The following screenshot shows the resource dependencies for a SAP Central Services service group in a Veritas Cluster Server configuration with NIC, IP, DiskGroup, Mount and SAPNW04 resources.

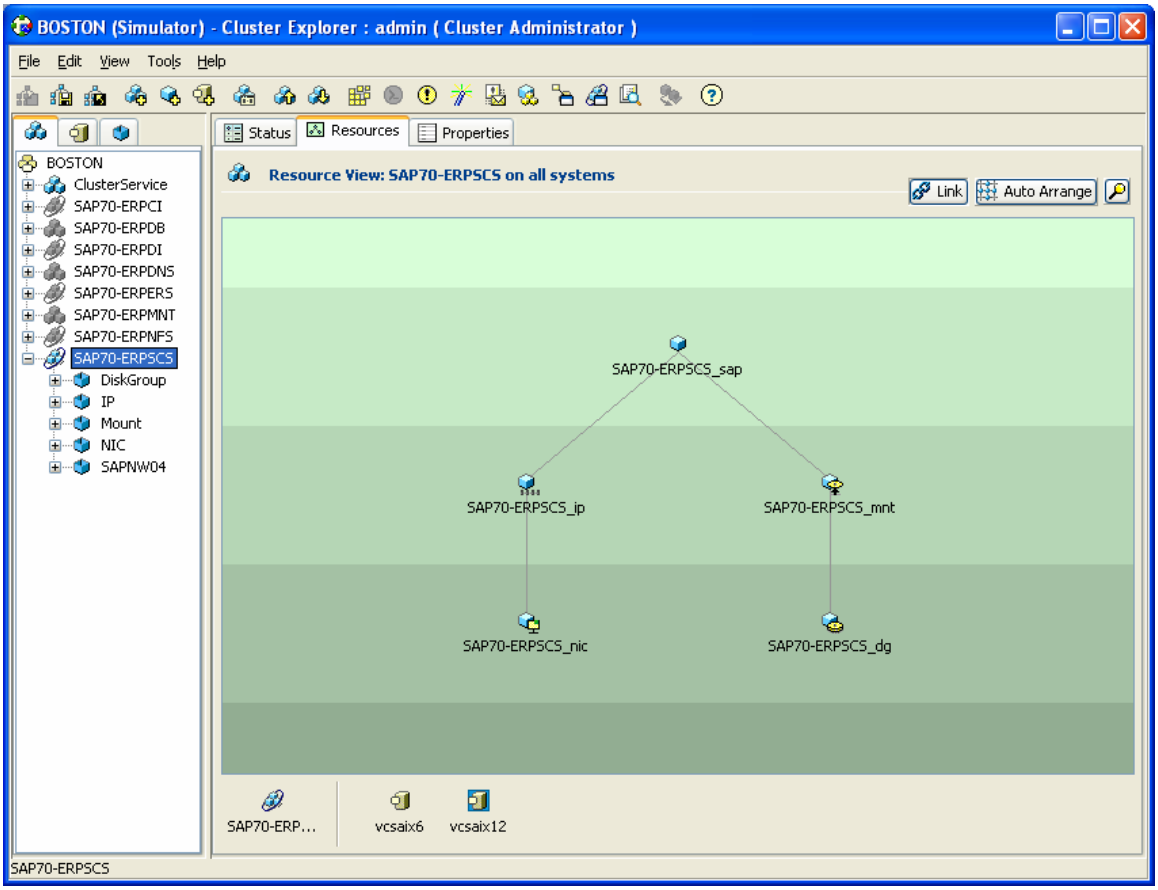

Figure 10: Resource dependency diagram for SCS Instance

The service group includes a Veritas Cluster Server resource for the Network Interface (NIC). It also includes a resource for the IP. This resource plumbs an IP address to the host when the resource is brought online and un-plumbs the IP when the resource is taken offline.

The service group includes a resource to manage the Disk Group and another resource to mount the file system (Mount) for the SAP application. And, finally the service group includes a resource to manage the SAP application (SAPNW04). This resource starts the SAP application when it is brought online and shuts down the SAP application when it is taken offline.

To support automated failover, SAP application components should be configured with virtual IP addresses. If a service group becomes unavailable, Veritas Cluster Server frees the virtual IP address so it can be reconfigured on the failover host. In this way, users connect to the application without regard for its physical location.

Local failover scenarios

The Symantec solution works in an SAP environment as illustrated in the following scenario.

The sample environment is a four-node cluster, with a web server, two dialog servers, and one server running both the Central Instances and an Oracle database to support it. All the four nodes access storage from a shared disk array.

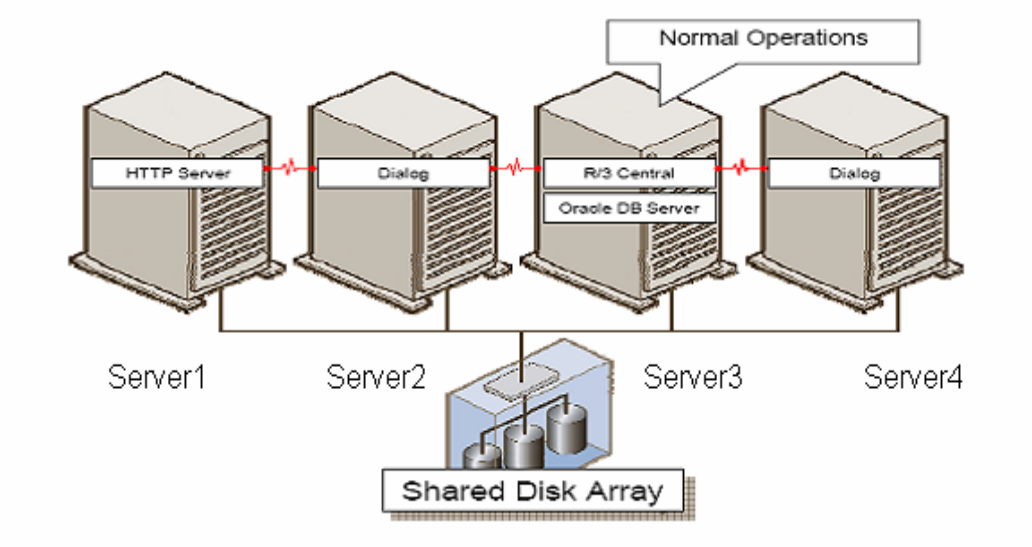

Figure 11: Sample local cluster environment for SAP

Veritas Cluster Server monitors each component with application-specific agents, and maintains “heartbeat” connections between cluster nodes to determine system availability. With this infrastructure in place, Cluster Server supports both proactive application switchover and automated failover processes.

Proactive switchover for planned downtime purposes

In a planned downtime scenario, the administrator initiates the switch of the service group to another node in the cluster. When this happens, the following processes occur:

- Appropriate application services are stopped in a clean and orderly manner.
- The virtual IP address is unconfigured on the current node and reconfigured on the new node.
- If Cluster File System is not used, then file systems are unmounted, then remounted on the new node.
- The application services are started on the new node.

The shutdown and startup processes follow a given order.

In the first example, Server 3 (which runs both the CI and the Oracle database) needs preventative maintenance. There is no "spare" server in the cluster, so the application services on Server 3 need to be switched to servers performing other tasks. Since Server 3 is a high-end server and the tasks are resource-intensive, the services on Server 3 are split between other servers in the cluster.

The administrator first switches the Oracle database server to Server 2, which is also running a dialog instance. The SAP DB Reconnect feature ensures that users are reconnected to the relocated database instance transparently, without interruption of their current sessions. Previously open transactions are rolled back in the DB restart phase.

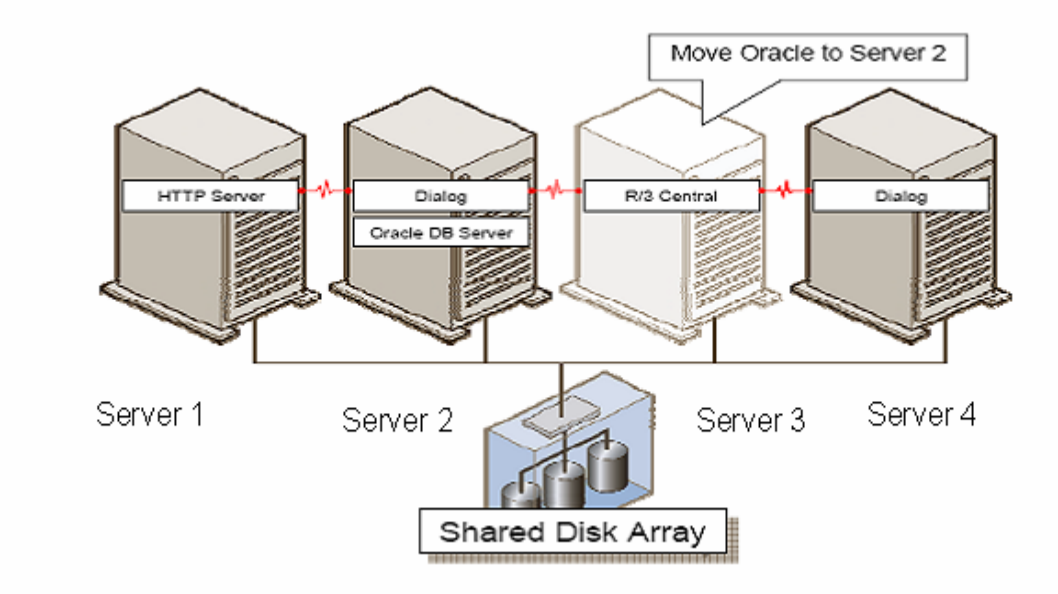

Figure 12: Sample local cluster environment for SAP after database switch

Next, the administrator chooses to switch the Central Instance to Server 4. The Central Instance will be shutdown on server 3 and started on server 4 where the dialog instance is also running. However if you configure the Central Instance not to run on the same physical server as a Dialog Instance, Veritas Cluster Server automatically shuts down the dialog instance, because it knows about the service group dependency between the Central Instance and the Dialog Instance. (The dialog instance could also be switched to a server running a web server process or to another server in the cluster with extra capacity.)

The Central Instance is then restarted on Server 4. The dialog servers cache transactions while waiting for the Central Instance to return. Once the CI instance is fully restored, the dialog servers will re-acquire critical database connections and will re-issue transactions that did not complete when Central Instance services were interrupted.

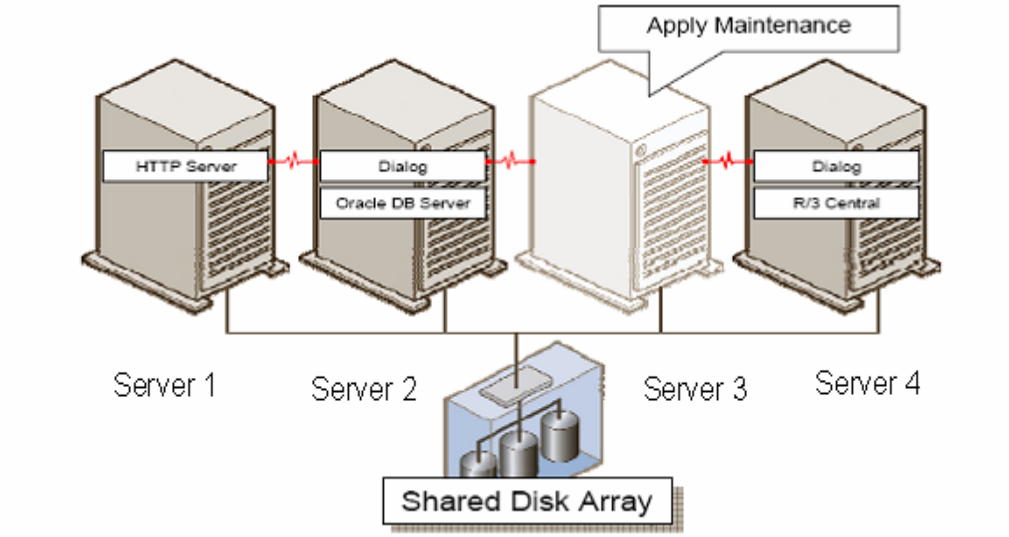

Figure 13: sample local cluster environment after SAP Central Instance switch

At this point, Server 3 is available for maintenance. When the maintenance is complete, the administrator uses Veritas Cluster Server to switch the applications back to their original configuration on the newly updated Server 3, and to restart the dialog instance on Server 4 if needed.

In a configuration using a standalone Enqueue Server, the Enqueue Server would failover to the cluster node running the Enqueue Replication Server. The Enqueue server would read the Replication Server's shared memory copy of the lock table and reinitialize the locks from this table. Once fully restored, the Enqueue server would accept connections from the dialog instances, which could resume normal operations and re-acquire locks without having to redo transactions in process. This configuration is better for stateful recovery.

Automated failover in response to unplanned downtime

The primary difference between a switchover and a failover is that in a failover, Veritas Cluster Server detects the failure and initiates the switch. If an orderly shutdown is not possible (i.e. in the case of a kernel panic[§§], memory fault, or operating system anomaly) then the agent performs a more aggressive and immediate shutdown of the application (i.e. kill -9). Once the application is shut down, the agent completes the process by "cleaning" the SAP instance environment, which includes removing SAP IPC resources (shared memory segments, semaphores and message queues), removing lock files, and removing the kill.sap program. At this point, the SAP instance can safely migrate to an alternate cluster node or can restart on the same node depending on the desired behavior.

Veritas Cluster Server automatically restarts the service group on other nodes of the cluster. Cluster Server uses a workload management feature to make intelligent decisions about system capacity and resource availability so that applications are hosted on the best server available within the cluster. You can also design failover policies to determine where applications are restarted.

[§§] A kernel panic is an action taken by an operating system upon detecting an internal fatal error from which it cannot recover.

Site migration for Disaster Recovery

Creating a highly available local clustering environment protects critical SAP applications from a wide range of component failures as well as supporting preventative and proactive maintenance without service disruption. However, it still leaves companies vulnerable to site-wide disruptions. Many companies are giving their business continuity planning efforts renewed emphasis in recent years, and critical SAP applications are a logical place to start planning.

By adding data replication and global clustering to your highly available SAP environment, you gain the ability to switch over an entire SAP application between geographically-distributed data centers, quickly and accurately.

Primary and Secondary sites

Global clustering depends on the existence of one or more alternate computing sites. The cost of alternate sites is a major concern for many organizations – how many can afford to have a duplicate data center, somewhere distant from the primary data center, ready to accept production processing at a moment's notice in case of a disaster?

Symantec's High Availability and Disaster Recovery solution puts a high level of application continuity within reach of many organizations by offering:

- ***Server savings*:** The secondary site need not be identical in computing resources to the primary site, provided degraded performance is acceptable. For example, the primary site might use four very large 12-way SMP machines, while the DR site might use six smaller, 4-way machines. The DR site might have less capacity, but still provide enough computing resources to run the SAP application should a disaster strike.

- ***Storage savings*:** The secondary site can use the same storage from a single vendor, different storage from a single vendor, or storage from multiple vendors, as long as the overall capacity is sufficient. For example, the production site might use very expensive, fast SAN disk arrays providing a total capacity of 5 terabytes. While the DR site also needs 5 terabytes to support the full data replication role, the DR site could use less expensive SAN disk to mitigate the total cost of ownership of the system. Veritas Storage Foundation and Volume Replicator do not require the same disk array on both sites of the replication link, and work effectively in a heterogeneous storage environment.

- ***Alternate uses for the DR environment*:** Many organizations actively use the DR site for other purposes that could be stopped in the case of a true disaster. For example, the secondary site could host development and QA processes for the SAP application or other applications. You would then configure, as part of the failover or switchover process, the graceful shutdown of development processes before bringing the production instance online. In this way, companies can leverage their investment in DR capacity to provide new computing resources for development, quality assurance, and non-essential production processes.

- ***Flexible global configurations*:** Primary and secondary cluster designation is relative, not absolute. For example, the primary site for a HQ instance of SAP is located at Site A. Site B serves as the primary site for a separate "production" instance of SAP in a multi-national company where the company has decided to maintain a separate and distinct instance to support a wholly-owned subsidiary. In this case, Site B might serve as the secondary site for the HQ SAP instance while Site A might serve as the secondary site for the "subsidiary" instance. The data center architect could configure a Site C to serve as a single DR site for both Sites A and B. The high availability architect has a number of very good choices open to him or her and, in many cases, is able to leverage existing resources to achieve a level of protection not previously available at a manageable cost.

Veritas Volume Replicatior

To support site migration at the click of a button, the production data must already be resident, and up-to-date, at the secondary site. Veritas Volume Manager, found in Veritas Storage Foundation, can be used to synchronously mirror data over an extended SAN infrastructure (up to 80km) to provide protection for metropolitan area networks. Veritas Volume Replicator (VVR), the IP replication option to Veritas Volume Manager, offers the ability to consistently and reliably replicate data across IP networks.

Veritas Volume Replicator replicates the contents of each volume across a wide area network to the secondary site. It is completely transparent to the application components. Unlike database-oriented solutions that replicate transactions or database blocks across distances, Veritas Volume Replicator can also manage other, essential file-based data, such as database configuration files, necessary for a complete site migration. And, unlike traditional block-based approaches, Volume Replicator replicates write I/O's instead of disk tracks to ensure the data is always replicated in a consistent fashion thereby guaranteeing the recoverability of the SAP application.

Veritas Volume Replicator supports both synchronous and asynchronous replication.

- Using synchronous replication, the initial write is not committed until the data has been replicated successfully. This solution guarantees that there is no data loss in case of a site failure, but there will be an application performance impact because the application at the primary site is waiting for the transaction to travel to the secondary site and back before the transaction is committed at the primary site. Over long distances synchronous replication may introduce unacceptable write delays in production systems.

- Using asynchronous replication, Veritas Volume Replicator commits the data at the primary site immediately and then queues replication operations for network availability. Asynchronous replication does not impact the application performance at the primary site but there may be some potential for data loss. Typical data loss between sites over long distance can be measured in milliseconds. Veritas Volume Replicator is unique in the market because it enforces write order on the replicated site, ensuring data integrity and consistency, thereby guaranteeing the data will be recoverable at the secondary site.

Most organizations select asynchronous replication for long-distance, global failover scenarios.

Hardware Based Replication

As discussed in the Veritas Volume Replication section to support the site migration, the production data must already be resident and up-to-date at the secondary site. The Veritas Cluster Server agents for hardware based replication offer the ability to support failover and recovery in environments where array based replication is used to replicate the data between arrays.

The following diagram shows a typical hardware based replication set-up. The set-up has a total of four servers out of which two servers are connected to a source array at the primary site and the remaining two servers are connected to a target array at the remote or secondary site.

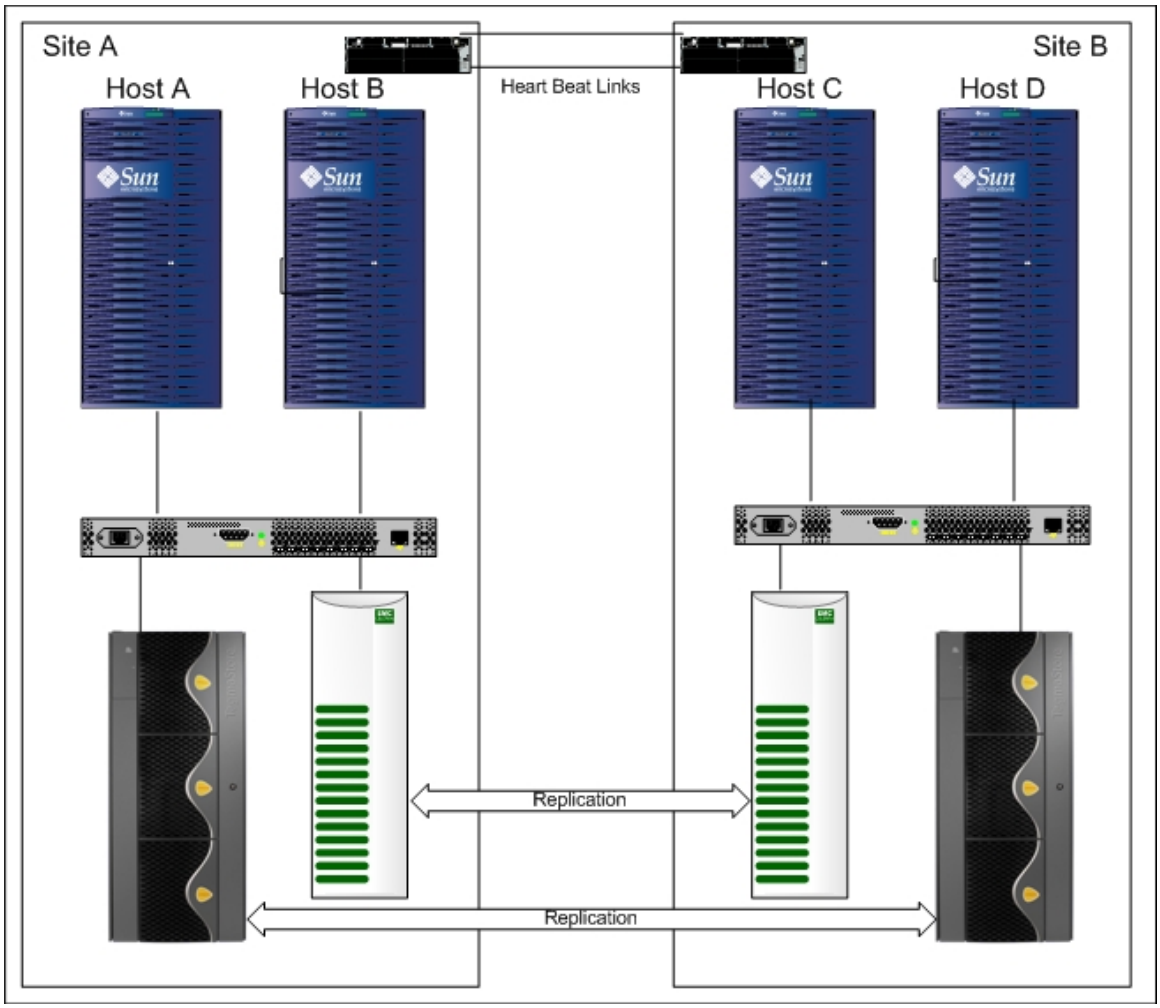

Figure 14: Hardware based replication diagram

Clustering in a hardware replication environment typically consists of the following hardware infrastructure:

- The primary array, comprising one or more hosts directly attached by SCSI or Fibre Channel to a primary array containing source volumes/devices.
- The secondary array, comprising one or more hosts directly attached by SCSI or Fibre Channel to a second array containing target volumes/devices. These devices pair with the source devices in the primary array. These hosts and the array must be at a significant distance from the primary site to survive a potential disaster.
- Network heartbeats, Low Latency Transport (LLT) or TCP/IP between the two data centers to determine the health of the other site.
- In a replicated data cluster environment, all hosts are part of the same cluster. You must connect them with dual dedicated networks that support Low Latency Transport (LLT).
- In a global cluster environment, you must attach all hosts in a cluster to the same array.

The Veritas Cluster Server agent for array based replication starts the replication, stops the replication and monitors the replication link between the source and destination arrays. Symantec supports various hardware based replication technologies. Some of the hardware based replication solutions that Symantec supports includes Hitachi True Copy, EMC MirrorView, and IBM MetroMirror, among others. See the Veritas Cluster Server website for the comprehensive support list. http://go.symantec.com/vcs

Global Clusters for SAP

Veritas Cluster Server HA/DR enables the linking of clusters from separate locations together and connecting SAP applications across clusters. This connection provides complete service level protection against an entire site failure by providing SAP application failover to the remote site.

Veritas Cluster Server HA/DR continuously monitors and communicates relevant SAP application events between clusters. Inter-cluster communication ensures that the remote cluster is aware of the state of the SAP application at all times. In the event of a system or application service failure, Veritas Cluster Server HA/DR fails over the affected services to another system in the same cluster. If the entire cluster or a site fails, Veritas Cluster Server HA/DR fails over the complete SAP application to the remote cluster at the DR site. VCS HA/DR also redirects clients once the application is online to the new location.

Note that the Cluster Server service groups must include *all of the components* necessary to run the SAP application – not just those that represent a single point of failure in a local clustering environment. For example, you might choose to not cluster dialog instances for local failover if you already run redundant dialog servers in your environment, but if you are implementing global failover you likely need to include all SAP components in the cluster, including dialog instances[***], web servers, etc.

[***] Symantec strongly recommends clustering dialog instances for local as well as global availability for easier management and operation of SAP applications.

Here is an example of Cluster Server service group diagram which shows the resource dependencies for a SAP application with global clustering.

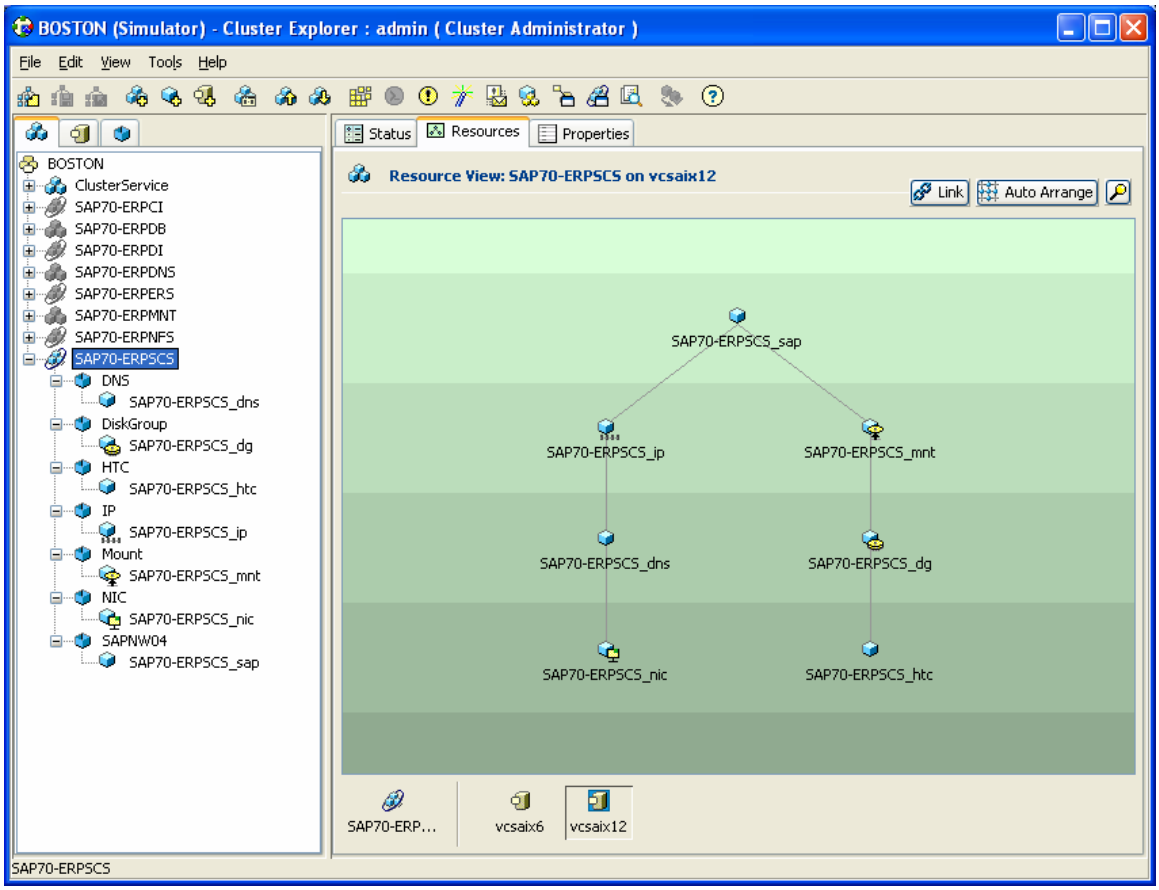

Figure 15: Resource dependencies for a SAP disaster recovery environment

The service group includes a Veritas Cluster Server resource for a Network Interface Card (NIC). The service group includes a resource to update the Domain Name Server (DNS) with a hostname and Virtual IP for the SAP application during global failover. It also includes a resource for IP. This resource plumbs the IP address to the host when the resource is brought online and un-plumbs the IP when the resource is taken offline.

The service group includes a resource for Hitachi True Copy (HTC) for data replication across sites for the SAP application. The service group also includes a resource to manage the Disk Group and another resource to mount the file system (Mount) for the SAP application. Finally, the Service group includes a resource to manage the SAP application, SAPNW04. This resource starts the SAP application when it is brought online and shuts down the SAP application when it is taken offline.

The Veritas Cluster Server service group should be configured as a Global Service Group (GSG). A global service group can be brought online on any node in the global cluster (listed in the ClusterList attribute). The GSG can also failover to any node in any cluster, even across sites.

The following screenshot shows the status of the SAP Central Services global service group in the local cluster, Boston, and the remote cluster, Chicago.

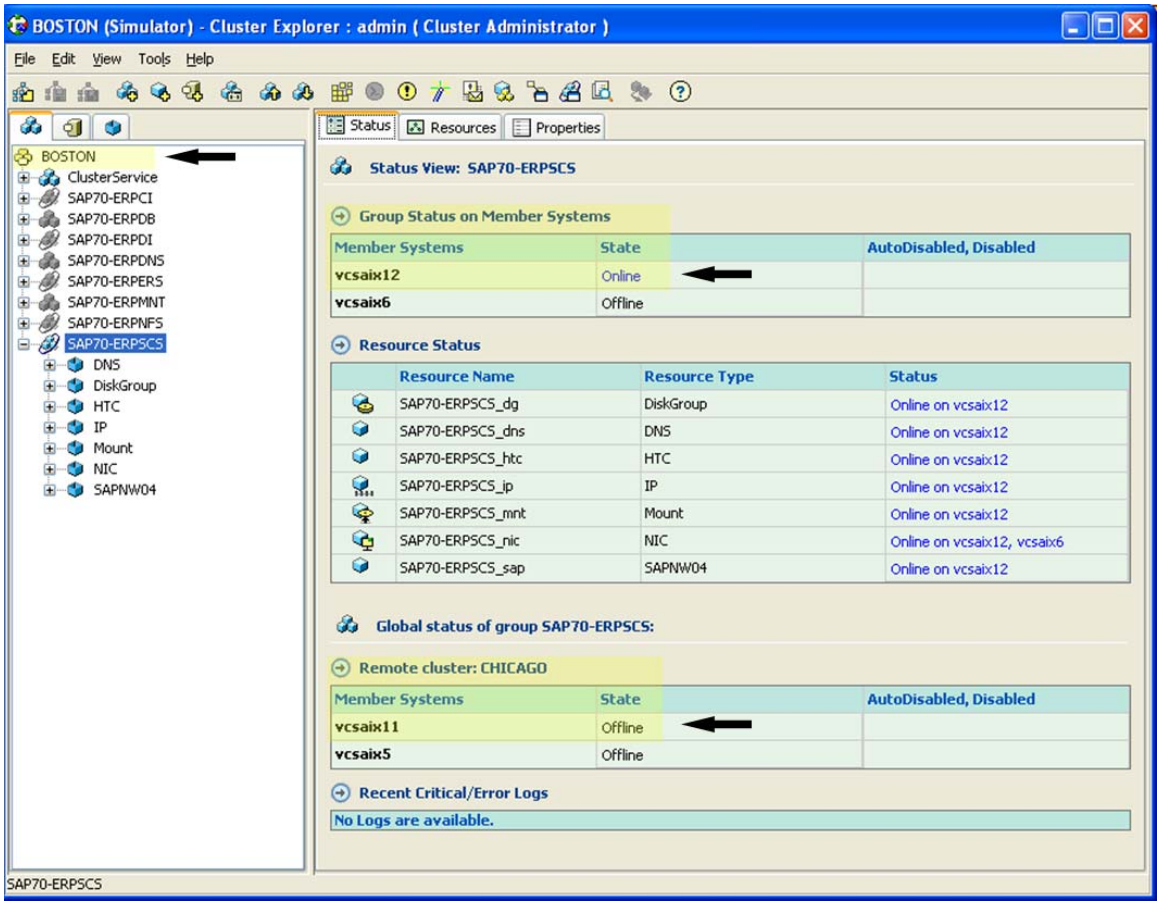

Figure 16: Global Service Group status for the SAP Central Services Instance

The SAP Central Services failover service group is online on system vcsaix12 in the local cluster, Boston, and offline on all the nodes in the remote cluster, Chicago.

Disaster Recovery Scenario

To understand how the integrated Symantec solution can switch a critical SAP application environment across sites for disaster recovery, it's worth walking through a disaster recovery scenario. As with local clustering, the difference between a planned switchover and a failover is whether or not you can shut down the production instances gracefully and initiate the switchover yourself, or whether a hard failure causes the automated failover.

Even in some "disaster" situations, you may have time to perform the switchover proactively, ensuring a smoother transition. For example, in the case of a power failure you may have a short period of time in which power is available from the Uninterruptible Power Supply system, which should be more than adequate to perform the site switch. Weather disasters often come with some advanced warning as well. Validation tests performed by Symantec show that application switchover can occur in only minutes. Actual switchover time largely depends on how long the critical applications take to start.

Consider the case of a company with a production facility running a SAP application in Boston and a DR/failover site in Chicago. The Chicago site could be running other services, which would then be shut down in case of a DR event. For the purposes of this example, we'll only look at the SAP application.

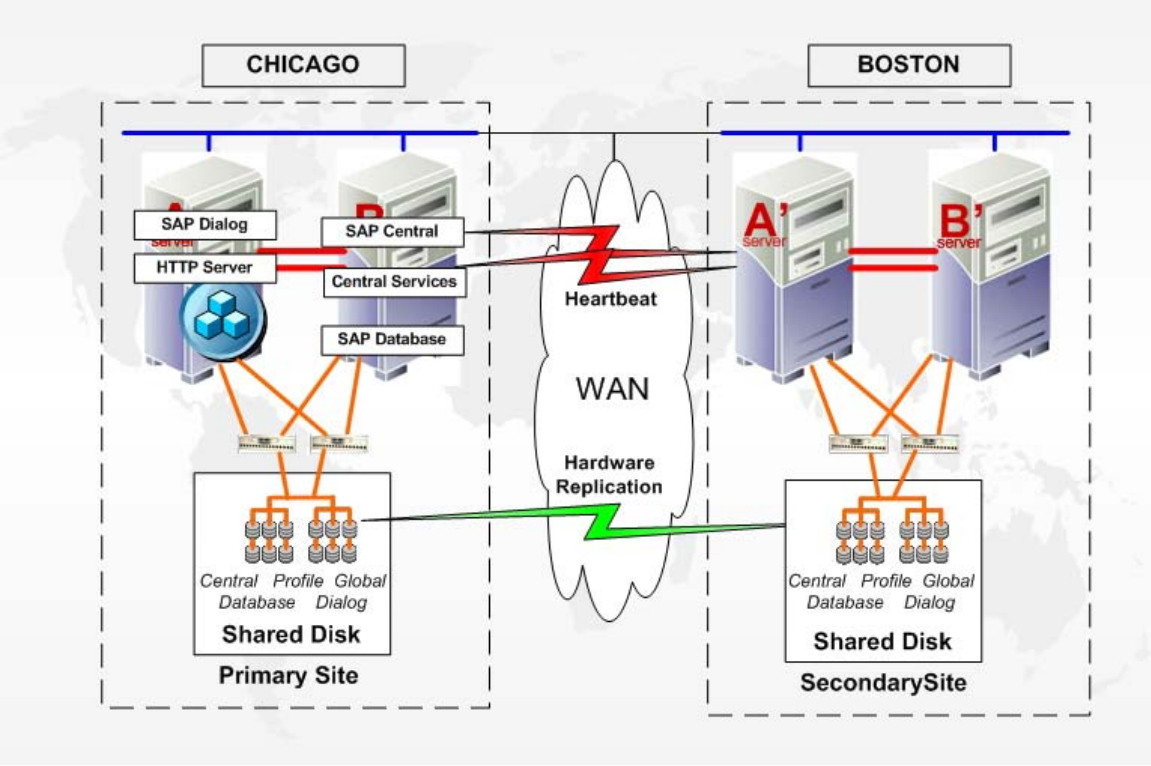

Figure 17: Example Disaster Recovery Setup

If the Boston site experiences a failure, such as a power failure or flooding, an operator would be notified that the site in Boston was down. He would then log on to the Cluster Server console from anywhere in the world to verify the site is actually not available. At that point the operator initiates switching the SAP application to the Chicago data center through a series of quick and easy point-and-click operations.

After Veritas Cluster Server stops any remaining SAP application services, it would perform the following steps:

1. Stop replication: First, Cluster Server will stop any replication between the primary and secondary sites. (If the primary site has experienced a dramatic failure, then replication has already stopped.)

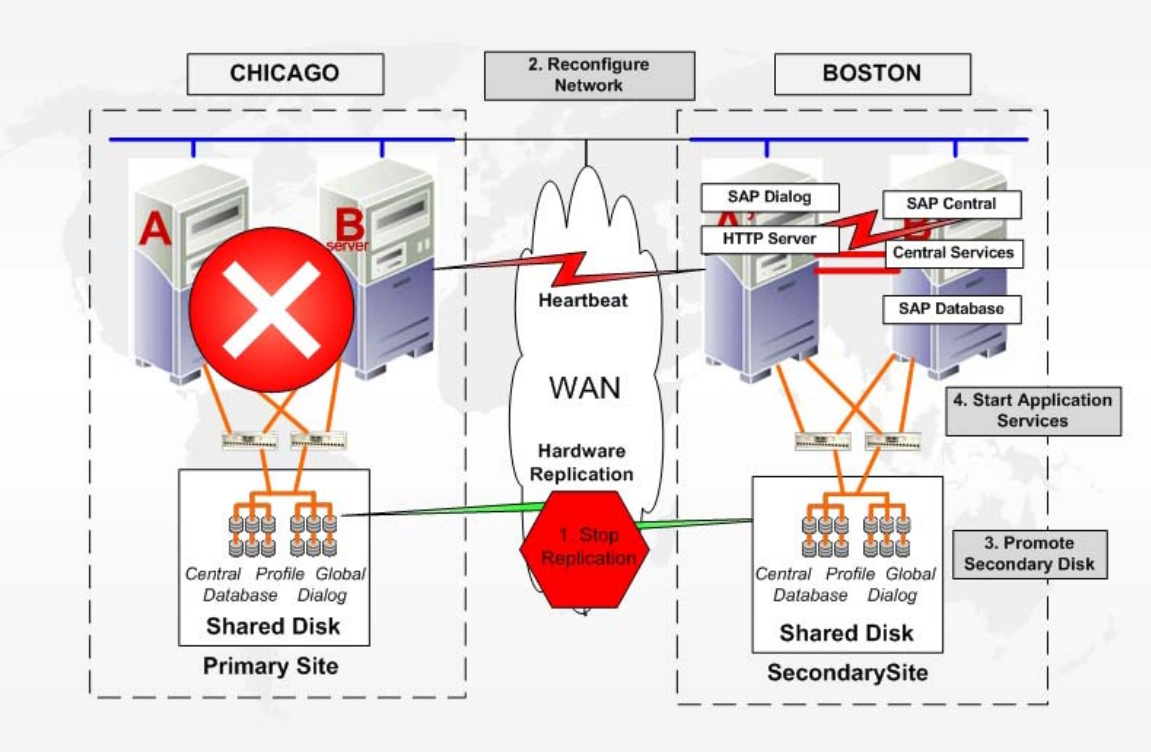

Figure 18: Disaster recovery set-up after disaster at primary site

2. Reconfigure the network: Next, Veritas Cluster Server updates the Domain Name Server (DNS) so that client network traffic is routed to the secondary site in Chicago.
3. Promote the replicated data: Then, Veritas Cluster Server promotes the secondary data to primary status. If a secondary site is still available, the Veritas Agent for hardware based replication or Veritas Volume Replicator can replicate to another site to maintain data availability.
4. Restart services: Finally, Cluster Server starts the SAP application services in the appropriate order, bringing the SAP instance back online.

At this point, the SAP application is now running successfully at Chicago – typically in a matter of minutes. All SAP instances would be restarted, without maintaining their transaction state. As a result, users would have to re-establish and re-authenticate their connections by refreshing the IP address, and restart all transactions not completed during the time of the switch. Otherwise, users would access the application in Boston without being aware that the application was running in a secondary site, Chicago.

When the disaster has passed and the primary data center is running again, you can update the Boston storage, restart replication, and then switch the application back to its original location, using the Cluster Server Management Console. The only data sent back to the Boston site is the data that changed while operations were occurring in Chicago, so that a complete re-initialization does not have to occur between the sites.

Compare this to the processes most companies manage today. First, companies need to maintain the disaster recovery site at the same update and patch levels as the primary site – difficult when the site is physically removed from the primary data center. Coordinating change control for multiple sites for each production application presents a significant challenge. Then, when it's time to switch an application to the disaster recovery site, IT has to replicate their data and applications at the disaster recovery site, and have no automated start-up of the applications. Application start-up has to be done in a manual way, without the documentation being used has every item in the correct order. The process can be very time consuming and expensive, and subject to possible human error.

Symantec's High Availability and Disaster Recovery solution eliminates these change control problems, as all changes to the primary application (and its subsidiary components) are replicated to the disaster recovery site automatically[†††]. The failover is automated, reducing risks while improving application availability.

[†††] This is true for applications that are replicated. This is not true for system level, O/S level patches and firmware, which are generally not replicated. You can use Veritas Configuration Manager to keep track of the changes to the applications and use the Fire Drill feature in Veritas Cluster Server HA/DR to simulate the DR failover scenario before performing the actual failover.

Conclusion

In many ways, the problems faced by IT organizations to maintain SAP application availability are endemic to all kinds of enterprise applications. Application complexity, high service level expectations, and constant change are factors nearly universal in today's IT environments.

These challenges have given rise to the concept of utility computing (or real-time infrastructure, on-demand computing or other like terms). Utility computing is all about reducing complexity, freeing applications from physical resource dependencies and creating a dynamic IT infrastructure that aligns with IT needs. For most vendors, however, utility computing is still a "horizon" technology – they can offer only a blueprint, a few pieces, and a map to future architectures.

In contrast, the Symantec Data Center Availability Solution for SAP provides tangible, demonstrable results in a matter of days, helping reduce the complexity of the SAP application environment while improving service level delivery.

Using the Symantec solution, the SAP application is not tied to a specific server or set of hardware components – it can be relocated manually or automatically when needed. Administrators gain time to do preventative maintenance and infrastructure work, and time to respond to crises appropriately while users continue to access the critical application. Companies gain improved service levels and reduced risk from outages. IT departments can better leverage resources, consolidating capacity in clusters and multi-purposing equipment used in disaster recovery facilities.

Symantec offers a real-world solution that solves immediate needs while supporting the long-term objectives of utility computing.

Appendix

Supported platforms and Veritas Cluster Server versions

Symantec supports the following combination of platforms and SAP versions for UNIX:

Veritas Cluster Server for SAP (UNIX)	
Application versions	SAP R/3 4.6C with 4.6D Kernel, 4.6D, 4.7 Enterprise, and SAP NetWeaver applications based on the Web Application Server 6.20, 6.40 (NW04), 7.00 (NW04s)
Platforms	Solaris™, AIX®, HP-UX, Linux®, VMware Linux guests
Services supported	SAP Central Services (SCS containing Standalone Enqueue Server) Enqueue Replication Server (ERS) Central Instance (CI) Application (Dialog) Servers (DI)

Symantec supports the following combination of platforms and SAP versions for Windows:

Veritas Cluster Server for SAP (Windows)	
Application versions	SAP R/3 4.6C with 4.6D Kernel, 4.6D, 4.7 Enterprise, and SAP NetWeaver applications based on the Web Application Server 6.20, 6.40 (NW04), 7.00 (NW04s)
Platforms	Microsoft® Windows® 2000 (SP4) Microsoft Windows Server® 2003 Microsoft Windows Server 2003 (x64) Microsoft Windows Server 2003 (IA-64)
Services supported	SAP Central Services (SCS containing Standalone Enqueue Server) Enqueue Replication Server (ERS) Central Instance (CI) Application (Dialog) Servers (DI)

For the latest information on supported platforms and Veritas Cluster Server versions, see the Veritas Cluster Server page on the Symantec website at: http://go.symantec.com/vcs.

Additional references

Symantec's High Availability and Disaster Recovery solution for SAP works with an updated Veritas Cluster Server DNS agent for global failover and DNS updates. Please download the updated agent and follow the instructions in the install document to install the agent.

http://support.veritas.com/docs/297120

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

02/08 13803401