

Confidence in the connected world.

# **Using Veritas Storage Foundation Cluster File System to increase availability of SAP NetWeaver**

*Oscar Wahlberg and David Noy  
Storage and Server Management Group*

# Using Veritas Storage Foundation Cluster File System to increase availability of SAP NetWeaver

Content

**Introduction..... 3**

**Business requirements for high availability with SAP NetWeaver ..... 4**

Business Problem ..... 4

Elements of the solution..... 6

Benefits of using Veritas Cluster File System for high availability of SAP NetWeaver ..... 8

**SAP NetWeaver architecture ..... 10**

SAP system architecture..... 10

Typical hardware configurations for SAP NetWeaver ..... 13

**How to implement a SAP NetWeaver in high availability solution with Veritas Storage Foundation Cluster File System..... 14**

Hardware considerations..... 14

Software considerations ..... 15

Setting up the cluster ..... 15

Cluster Volume Manager and Cluster File System Infrastructure resources ..... 19

Summary ..... 27

Where to get more information ..... 27

**Appendix ..... 28**

Veritas Cluster Server configuration file: main.cf ..... 28

## Introduction

SAP NetWeaver suite is a suite of business products that is designed to support complete business management tasks of a corporation, company or institution. SAP NetWeaver is used in business systems to handle invoice payment, production resource management and financial account control. These specific tasks are accomplished by employing application modules, often described with acronyms such as MM (materials management) and SD (sales and distribution).

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 maintaining 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
- Interdependencies between components complicate administration
- The infrastructure itself undergoes constant change

To gain additional competitive advantage, enterprises must now work more closely together and integrate their current IT solutions with those of other organizations, such as partners, customers, or suppliers for SAP. The availability of these solutions is therefore essential. The number of areas that are critical to enterprises is increasing. SAP provides open, scalable, and available solutions to meet these requirements effectively.

This paper describes how Veritas Storage Foundation Cluster File System can be used to implement high availability of SAP deployments on UNIX. This paper focuses on both the business benefits of building SAP on top of Veritas Cluster File System as well as high level coverage on how to configure SAP services and the underlying database with Veritas Cluster File System.

For a more in-depth view on how to ensure availability of a complete SAP environment, it is recommended that the reader reference:

- [Symantec High Availability and Disaster Recovery Solution for SAP](#) by Symantec Corporation
- [Veritas Cluster Server for SAP Solutions](#) by Symantec Corporation
- [Veritas High Availability Agent for SAP NetWeaver Installation and Configuration Guide](#) by Symantec Corporation
- [Technical Infrastructure Guide – SAP NetWeaver](#) by SAP
- [Architecting a high availability SAP NetWeaver Infrastructure](#) by SAP

## **Business requirements for high availability with SAP NetWeaver**

### **Business Problem**

There are many concerns that customers face when trying to deploy SAP NetWeaver in a highly available configuration that can directly be addressed by the Veritas Cluster File System. These include:

#### **High Service Level Requirements**

SAP applications are business critical. SAP NetWeaver deployments require that a clustering solution be in place that can detect application failovers and react accordingly to restore operation. SAP NetWeaver application availability can also be impacted by underlying array or path failures.

When application failures do occur, they need to be detected and handled promptly to ensure high service levels and minimum disruption. If a failure occurs and redundancy of software and hardware is lost, it is important that redundancy be restored once the failed hardware has been replaced or the failed application has been restarted. Furthermore, customers increasingly seek a cluster solution which proactively detects storage array and path failures and adjusts accordingly to prevent service disruptions.

#### **Reduce deployment costs**

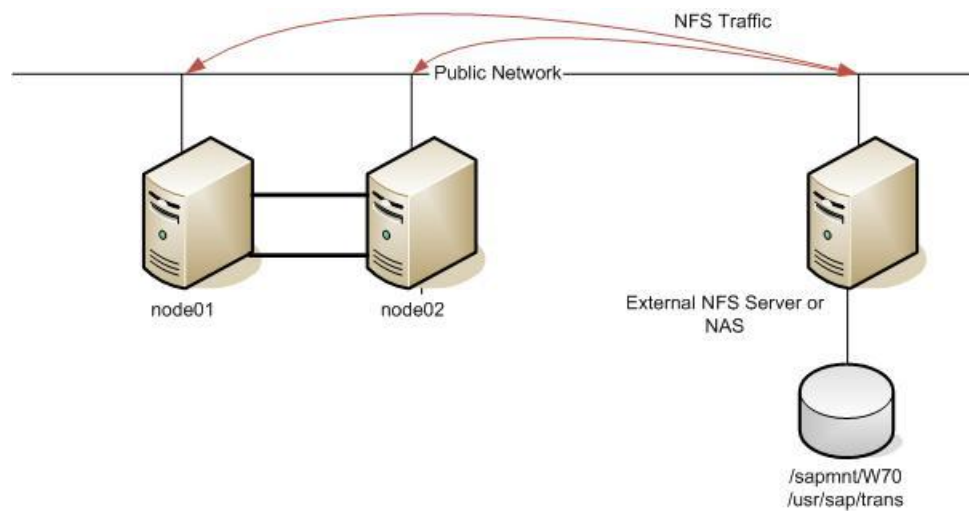
Customers seek to reduce the costs typically incurred as part of solutions around maintaining the high service levels discussed above. Customers typically deploy active/passive clusters which are expensive because 50% of the compute power is then sitting idle as a result. Furthermore, with most clustering tools, customers are locked into specific vendors which limit their ability to leverage existing hardware.

#### **Solutions have to account for the database also**


SAP NetWeaver applications are built around enterprise grade databases. Customers seek a storage solution which provides maximum performance, high availability, integrated backup and specific database interoperability.

### NFS is an insufficient solution

NFS is typically deployed in conjunction with SAP to reduce the total cost of the solution but this has several important drawbacks. Firstly, NFS performance cannot compare with the performance of fiber connected storage. Veritas Cluster File system will show better performance than NFS both for throughput and with regards to failover times. Secondly, NFS has poor reliability relative to Veritas Cluster File System. This is due to the fact that NFS is only as reliable as the network it runs on and because NFS has no knowledge of underlying storage path health.



**Figure 1 SAP NetWeaver cluster with an external NFS server**



**Figure 2 SAP NetWeaver cluster with an internal NFS server**

## Elements of the solution

To address the problems listed above customers require a shared file system solution with the following properties:

### **File system mounted on all nodes simultaneously**

Having the file system mounted and available on all nodes in a cluster reduces failover times by eliminating the extra step required to bring a file system online when failover occurs. Because data is presented to all nodes in a cluster with such a solution, SAP applications can quickly resume operation where a failed application left off. Once the failed node has been restored, a shared storage solution has the capability of presenting files to the now healthy node non-disruptively so that redundancy can be restored.

### **Integration with clustering services**

A cluster file system by itself has some value to SAP, but that value is dramatically increased when the cluster file system integrates with clustered application management. Cluster file systems that are integrated with clustered application management servers can respond to failures automatically without user intervention thereby reducing the time required to resume operation as well as restore redundancy.

Clustered application management servers can detect failures outside the scope of SAP and can proactively respond to problems in the server environment.

When integrated with the underlying cluster file system, clustered application management servers can ensure that SAP application resources are started and available before attempting to start SAP applications so that resource dependencies don't prevent SAP from starting.

### **Active / Active configuration**

A scalable cluster file system provides high performance, scalable access to data across all nodes in a cluster simultaneously. This configuration enables all nodes in a cluster to host active components of the SAP NetWeaver solution reducing the cost of deployment and improving the total availability of the system. Cost is reduced because all nodes can run in active mode and do not have to sit idle waiting for failure to occur. Furthermore, availability is improved because rather than reserving specific nodes as failover nodes, any node in the cluster can serve as a failover node.

### **Multi-Pathing solution**

A highly-available cluster file system solution benefits from an integrated multi-pathing solution because the multi-pathing solution can coordinate path utilization across all nodes accessing the cluster file system. An integrated multi-pathing solution that detects path failures on one node can use that information to coordinate path utilization across all nodes in the cluster to prevent performance degradation and potential trespassing on a storage array.

### **Database integrated and aware solution**

Because SAP NetWeaver applications persists information in a database, a clustering solution is required which can provide maximum performance and availability for a wide array of databases as well. This typical database deployed with SAP is Oracle. As a result, a solution which is targeted at Oracle deployments is well suited to improving the availability of SAP NetWeaver applications.

### **Benefits of using Veritas Cluster File System for high availability of SAP NetWeaver**

The Veritas Cluster File System achieves all of the aforementioned elements of a highly available shared storage solution for SAP across the broadest set of platforms of any competing file system.

Specifically, Veritas Cluster File System provides the following benefits for ensuring high availability of SAP NetWeaver:

#### **Automatic and fast failover**

Veritas Cluster File System (CFS) can be used to mount a file system across all nodes in a cluster to ensure that the file system is already available for use in a failover scenario. When used in conjunction with Veritas Cluster Server (VCS), node failures or application failures are detected by VCS agents and failover is initiated automatically. The failover time is reduced to the time it takes for the application to resume on a failover node. Because CFS is built with a highly scalable architecture, all nodes in the cluster can participate as “active” nodes running SAP or database processes. Similarly, all nodes in the cluster can serve as failover nodes.

#### **Services started in the correct order**

The Veritas Cluster Server ensures that all SAP NetWeaver dependencies are started in the correct order. This is important for automated recovery of a failed node or when new nodes are configured and added to the cluster. By monitoring dependencies outside the scope of SAP, VCS can determine if a failure has occurred in a critical dependency for SAP and take the appropriate response.

#### **Redundancy restored when nodes recover**

The VCS product is able to restore the redundancy of a cluster by bringing a node and its applications back online and rejoining the cluster. Because there is no I/O quiescing on cluster reconfiguration, nodes can rejoin the cluster without disrupting applications already running within the cluster.

#### **POSIX Lock Compliance**

The Veritas Cluster File System complies with POSIX standards for file locking. This standard for file locking is used to ensure that file write operations are transactional and nodes in a cluster are not corrupting each other's data writes.



### Extended file system features

The Veritas Cluster File System is built on top of the VxFS file system which offers several additional features that can improve availability and reduce the time it takes to restore a cluster in the event of failure.

- Storage Checkpoint is a feature which can be used to set a recovery point objective for rolling back files to a specific point in time.
- The File Change Log feature tracks changes to files and directories. Using the FCL can improve application performance significantly if the application is required to perform actions like scanning an entire file system to detect changes since the last scan.
- Multi Volume File Systems is a feature of both the Veritas File System and the Veritas Volume Manager. The integration allows a file system to span multiple volumes. This technology enables high performance by allowing metadata and data to be stored on separate volumes and RAID-levels. The Multi Volume File System is also a key enabler for Dynamic Storage Tiering.
- With the Dynamic Storage Tiering feature, the Veritas File System can leverage multi-volume file systems to present a single namespace of different tiers of storage.

### Veritas Dynamic Multi-Pathing

The Veritas Dynamic Multi-Pathing (DMP) product ships as part of Veritas Cluster File System and provides the widest range of storage array support of any multi-pathing solution on the market. Because it is integrated into CFS and VCS, DMP can detect path and storage failures and coordinate an appropriate response cluster wide to ensure the highest service levels. DMP can also proactively detect potential problems with storage and take action to prevent disruptions before they happen.

### Database support

VCS can be used in conjunction with CFS and database specific agents to ensure high availability of the database used for persistent SAP data. Storage Foundation for Databases contains additional features to enhance the functionality and performance of Oracle, Sybase and DB2.

### One integrated solution

The Veritas Cluster File System High Availability product (CFS HA) is based on the Veritas Cluster File System, Veritas Cluster Server, Veritas Cluster Volume Manager, and Veritas Dynamic Multi-Pathing solutions. All of these products coordinate with one another to ensure the optimal response to changing conditions in a cluster. The result is a strongly integrated end to end clustering solution that monitors from the application layer down to the infrastructure layer of SAP NetWeaver deployments and can respond accordingly when failures occur. Because it is a single integrated solution the likelihood of missing potential problems or incorrectly configuring the HA solution is greatly reduced. Furthermore, administrators will be able to have the confidence in using the same products that they use to ensure high availability of their other clustered deployments.

## **SAP NetWeaver architecture**

This chapter describes the typical architecture of SAP NetWeaver deployments. The discussion includes elements of the SAP NetWeaver product itself, the supporting database, and the underlying storage. This chapter will specifically call out areas of concern for customers seeking high availability solutions for their SAP deployments.

### **SAP system architecture**

The SAP NetWeaver architecture consists of multiple components that usually are spread across multiple physical servers to increase performance and scalability:

- Central Instance
- SAP Central Services (SCS)
- Enqueue Replication Service
- Dialog Instance
- Database Instance
- NFS (Network File System) Service

Of these components, SAP identifies the following components as system-wide, single points of failures in non-clustered SAP NetWeaver environments:

- Central Instance
- SAP Central Services (SCS)
- Database Instance
- NFS (Network File System) Service

Highly available SAP environments need to make all of these services resilient to failures.

### ***Central Instance***

The Central Instance, consisting of the Dispatcher and Work processes and the Enqueue process, can be considered the core. The Central Instance work processes are what actual performs the user requested tasks.

### ***SAP Central Services***

The SAP Central Services instance consists of the Enqueue and Messaging Services. Therefore the corresponding Central Instance does not contain the Enqueue and Message service.

The Enqueue service maintains the logical SAP locks and ensures synchronization within the SAP environment, the Messaging Service is the communication hub within the SAP solution.

### ***Enqueue Replication Service***

The Enqueue Replication Service enables the lock tables to be replicated to a hot standby server (replication server) to prevent client disruption due to a failover of the SAP Central Services. If the Enqueue Replication service is not deployed, all open transactions will be rolled back and user initiated jobs will be aborted. If the Enqueue Replication service is deployed the user will not have open transactions aborted.

### ***Dialog Instance***

The Dialog Instances can be considered to be the application servers. They are responsible for receiving and processing client requests.

### ***Database instance***

The Database Instance is a 3<sup>rd</sup> party relationship database manager service provided for persistence of data used by SAP. SAP can use a variety of different database engines from all major vendors. In this whitepaper we have chosen to use Oracle 10g as an example.

### NFS Service

Generally, enterprise application vendors of highly available, mission critical software tend to be wary of NFS as an underlying storage mechanism. When used alone, NFS does not provide a mechanism for automatic failover of applications. If instances of SAP components or the database it relies on should fail, there is no mechanism in place to restore them. Even when used with application clustering and high availability software, such as Veritas Cluster Server, using NFS as the underlying storage for SAP deployments can cause result in lockups that can take up to minutes to resolve. This is unacceptable for today's mission critical applications where strict Service Level Agreements are the norm. Because NFS is a network protocol, it is also vulnerable to network reliability issues. Finally, NFS has no knowledge of the health of the paths from nodes in a cluster to the file server. Path failure on a node can cause performance degradation when traffic across the paths to storage is not coordinated across nodes. As a result, we recommend that Veritas Cluster File System be used for shared storage.

If the administrator decides to pursue the approach of deploying NFS anyway, the NFS Service provides shared storage for all instances, but the dialog and application servers in particular. The shared data consists of SAP kernel executables, and profiles that need to be available on all participating servers.

The SAP executables and binaries need to be available on all servers, either via NFS, a local binary installation or via a cluster file system. According to the SAP NetWeaver documentation the following directories ought to be shared between all instances of a system.

Directory	Description
/sapmnt/<SID>/profile	Contains the different profiles, to simplify maintenance
/sapmnt/<SID>/global	Contains log files of batch jobs, central Syslog facility
/usr/sap/trans	Contains data and log files for objects transported between different SAP systems (For example between development and production).
/sapmnt/<SID>/exe	Contains the SAP kernel executables. These executables should be accessible locally on all servers running a SAP instances. If NFS is used this can be achieved either by using the <i>sapcpe</i> utility to synchronize the needed files to the instance specific executable directory. If a cluster file system such as Veritas Storage Foundation Cluster File System is used this is not needed.

### **Typical hardware configurations for SAP NetWeaver**

A typical SAP NetWeaver deployment spreads the critical services across multiple servers.

This provides the ability to tune the performance of the underlying storage for the type of service running on each machine. This configuration is ideal for high availability because it also means that a part of the environment can experience a hardware failure without affecting the whole environment.

The most common way to configure the services is to separate the Central Instance and the Database so that they run on different servers. This provides the service partitioning and flexibility for tuning, yet uses the least amount of physical servers.

The NFS Service can be provided either by a server within the environment or an external server or NAS appliance.

When the NFS service is provided from a server within the environment, it is usually from the Central Instance or Database server. If the Central Instance server provides the NFS Service several problems arise:

The server that hosts the NFS Service needs to re-mount the share locally. This is not always supported by the NFS server/client.

Network disruption between the SAP servers can cause access problems and might affect the SAP application(s)

When the NFS Service is configured in a clustered environment, the NFS service can sometimes prevent a failover of the Central Instance to a standby server

When an external NFS server or NAS appliance is utilized, the availability of the SAP environment becomes vulnerable to the network connection to that server or appliance. If the NFS server, NAS appliance, or network connection experience a problem, the SAP environment may become unavailable to users.

SAP environments typically contain one or more Dialog instances. These are by design made to be deployed in a redundant configuration as the SAP architecture allows for multiple parallel instances. It should be noted that if a Dialog instances should fail the performance of the entire environment will drop.

## **How to implement a SAP NetWeaver in high availability solution with Veritas Storage Foundation Cluster File System**

The following chapter describes a 4-node SAP NetWeaver and Oracle 10g configuration based on Veritas Storage Foundation Cluster File System 5.0.

Precautions have been taken during the design phase to allow for future growth. All critical services have been compartmentalized using virtual IP addresses to ensure that they can fail-over independently and so that they are easy to move to other servers if there is a need for further growth.

The intention of this document is not to give a detailed step by step description of how to configure SAP NetWeaver or Oracle, but to give the reader an overview on how to leverage the Veritas Clustered File Systems in these configurations. Please be mindful that there are many different ways to deploy SAP and this document only describe one configuration. It is highly recommended to involve an experienced consultant with knowledge about the Veritas Storage Foundation HA product line during the design phase of any SAP clustering solutions to ensure that the solution meet your requirements.

Previous exposure to Veritas Cluster Server will help the reader to fully understand the concepts discussed in this chapter. If a refresher is needed the Veritas Cluster Server Administrator's Guide (<http://seer.support.veritas.com/docs/286758.htm>) provides a very good introduction to clustering technologies.

### **Hardware considerations**

The configuration described in this document is based on 4 servers, each capable of running SAP and Oracle 10g.

In addition to fulfilling requirements from SAP and Oracle each server must be equipped with:

- 2 Network Interface Cards dedicated for cluster communication

- Fiber-channel Host Bus Adapter card to interface with a Storage Area Network

- A minimum of 2 processors and 2GB RAM per core

The official Hardware Compatibility List for the servers and storage platforms supported by Veritas Storage Foundation Cluster File System can be downloaded from <http://support.symantec.com>.

## Software considerations

Veritas Storage Foundation Cluster File System/HA 5.0 is supported on a wide array of UNIX/Linux platforms with different patch requirements. The latest information for each release is documented in the Release Notes. Late breaking news, or post release information, is always available at the Symantec support site <http://support.symantec.com> as Tech Notes.

The configuration described in this document requires:

- Veritas Storage Foundation Cluster File System High Availability 5.0

- Veritas Cluster Server Agent for SAP 5.0

- Veritas Cluster Server Agent for Oracle 5.0

## Setting up the cluster

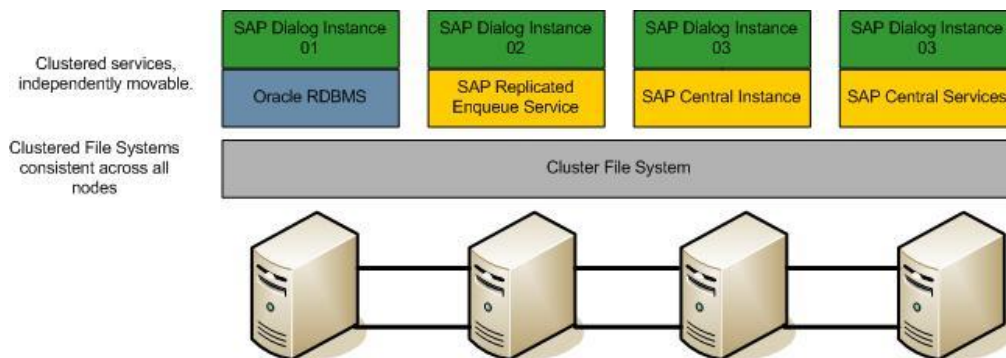
The cluster should be configured so that Cluster File System and Veritas Cluster Server are running on all 4 nodes within the cluster. This allows the file system to be mounted on all nodes simultaneously, so that failover times can be minimized. In the example below the different services (Database, Central Instance etc) are placed on different nodes.

For example the Oracle database runs on node1 and the SAP Central Instance on node3.

Dialog Instances are deployed on all four nodes for scalability and reliability. If the Dialog Instance on any node fails, the other nodes will remain functional.

Each individual failover service group can, if the need arise; fail over to any of the other clustered nodes.

Redundant private interconnects are setup between the four nodes to keep track of cluster membership, control application availability, and provide coordination between the Cluster File System on each node.



**Figure 3 A 4 node Oracle and SAP cluster using CFS**

## Overview of clustered services

The table below describes the different service groups configured within Veritas Cluster Server.

This setup is divided into separate service groups:

Service Group	Description
public_network	A dedicated group for monitoring the external network configuration. This group is configured as a parallel group and runs concurrently on all systems.
cvm	The Veritas Cluster Volume Manager group. This group ensures availability of the Cluster Volume Manager infrastructure and daemons. This group is configured automatically by the installation scripts for Storage Foundation Cluster File System. This group is configured as a parallel group and runs concurrently on all systems.
cfs	This group is responsible for the resources related to the Cluster File System. It ensures that all the needed file systems are mounted on the required nodes. This group is configured as a parallel group and runs concurrently on all systems.
sap-db	The Oracle database instance named <i>W70</i> . Ensures high availability of the database instance and the Oracle Listener.
sap-ci	A SAP Central Instance named <i>JC00</i> . This service group contains the necessary resources to ensure availability of the Central Instance.
sap-rep	SAP Replicated Enqueue service, instance <i>REP02</i> . This group controls the SAP Replicated Enqueue Service.
sap-dia	SAP Dialog servers. One Dialog servers have been configured per physical node for redundancy and load balancing. Each Dialog instance has a unique name D10 to D14. This group is configured as a parallel group and runs concurrently on all systems.
sap-scs	SAP Central Services instance <i>SCS01</i> .

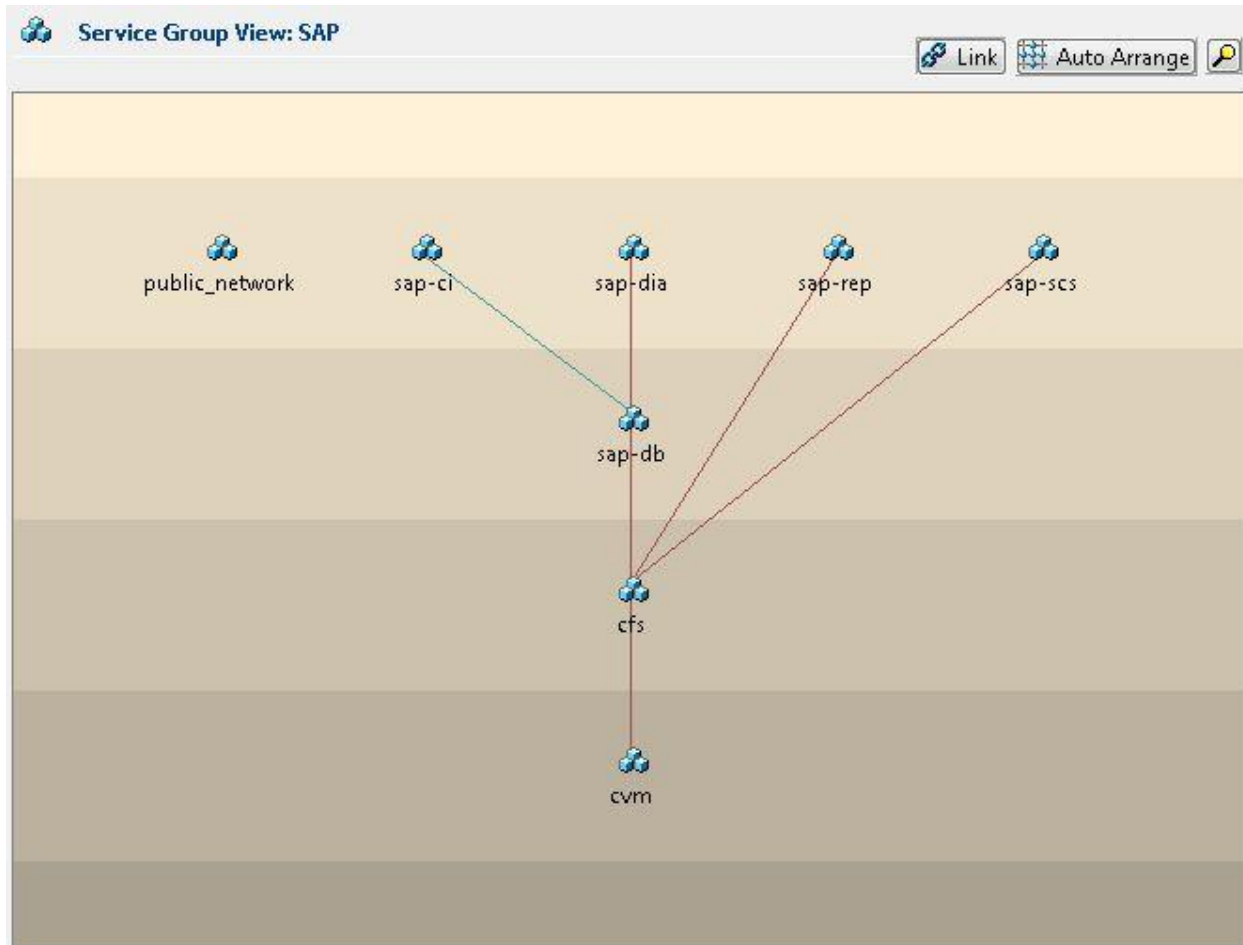


### ***Creating service dependencies***

The SAP architecture creates several dependencies between the different services within the environment. The clustering solution must take these dependencies into consideration. Veritas Cluster Server includes a facility called Service Group Dependencies to ensure that the clustering software take application dependencies into account.

Veritas Cluster Server has ability to use dependencies between each Service Group to ensure that failover happens correctly and applications are able to be restarted on failover nodes. Service groups ensure that entire sets of software components and applications failover as a group. Service group dependencies ensure that these groups are brought offline on a failed node, and then online again on a failover node, in the correct order.

The figure below illustrates how the different service groups within this solution depend on each other. For an in-depth description of the different types of service group dependencies, please see the Veritas Cluster Server Administrator's Guide. (<http://seer.support.veritas.com/docs/286758.htm>)



**Figure 4 Service group dependencies**

## Using Veritas Storage Foundation Cluster File System to increase availability of SAP NetWeaver

At the bottom of the dependencies tree the *cvm*-group provides the infrastructure for the Veritas Cluster Volume Manager and Cluster File System. This group will be the first to start, and the last to stop. The *cvm* group is a parallel group and will be online on all servers within the cluster during normal operations.

The *cfs* service group ensures availability of the clustered file systems and it has an “Online Local Firm” dependency on the *cvm* service group. This is to ensure that all required clustering components are functioning correctly before proceeding.

The database service group (*sap-db*), the SAP central services group (*sap-scs*) and the SAP replicated enqueue service group (*sap-rep*) have an “Online Local Firm” dependency on the *cfs* group to ensure that the file systems are available prior to starting the applications.

The central instance service group (*sap-ci*) and the dialog instance group (*sap-dia*) both have an “Online Global Soft” dependency on the database group (*sap-db*). This is to ensure that the database is online on any node within the cluster.

Each of the application service groups also have a proxy resource configured within the service group. This is done to ensure that all the cluster file systems are available on the local node. This is achieved by having the proxy resource represent the status of the *cfs\_online\_phantom* resource in the *cfs* group.

### **Configuring the public network group**

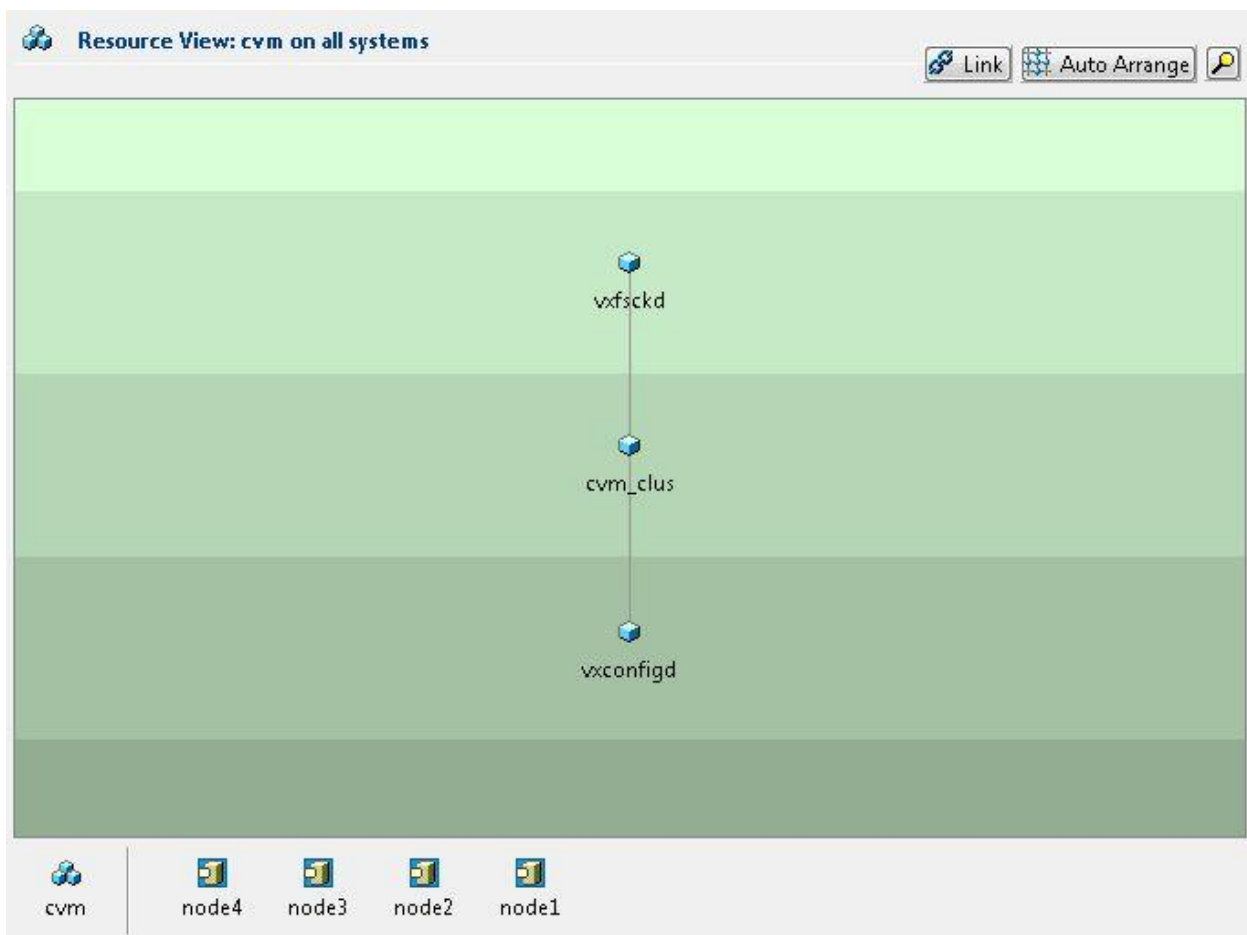
The group *public\_network* is responsible for monitoring the network interfaces used within the cluster. This service group contains the VCS resources that regularly monitor the external network connectivity. In each of the other service groups, a proxy resource points to the *public\_nic* resource within this group, mirroring its status. This is done so that VCS only needs to monitor the network once per node rather than once per service group that utilizes the public network.

### Cluster Volume Manager and Cluster File System Infrastructure resources

In all Storage Foundation Cluster File System clusters there should be a service group named *cvm*. This group is configured during installation and provides control and monitoring facilities for the Cluster Volume Manager and Cluster File System infrastructure resources. These resources must run on each system participating in a cluster.

The VCS resources *cvm\_vxconfigd*, *cvm\_clus* and *vxfsckd* are mandatory services.

In addition to monitoring and controlling the services required for Veritas Storage Foundation Cluster File System, it is common in smaller configurations to simplify the service groups by using the *cvm* group to host shared file systems resources. In this 4-node configuration we have chosen to not do so and instead a separate service group named *cfs* is used to host the shared file systems used by Oracle and SAP.



**Figure 5 Resources in the *cvm* group. Mandatory group containing infrastructure daemons**

### **Clustered file systems**

The service group *cfs* contain the resources necessary to monitor the disk groups, volumes and file systems utilized by Oracle and SAP in this configuration.

Below is a list of the shared file systems that are configured to be automatically mounted on all participating nodes, please note that this is a simplification of a real SAP and Oracle 10g deployment scenario and that a real deployment is likely to be more complex.

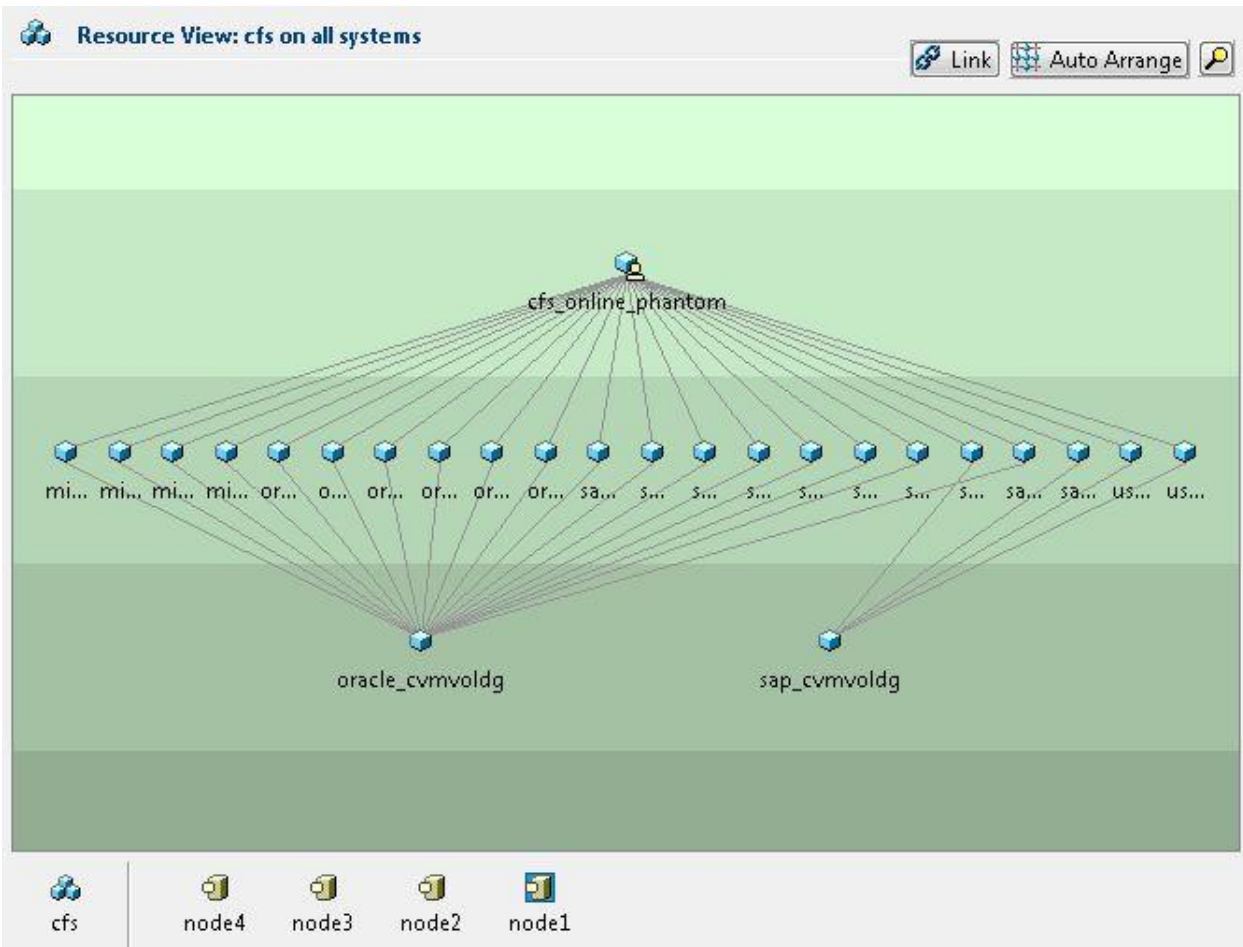
Note that the resource, *cfs\_online\_phantom*, is a Phantom resource that mimics the status of the group (See the [Veritas Cluster Server Bundled Agents Reference Guide](#) for further details). The status of this resource is used to indicate to the application groups that the cluster file systems are available, via proxy resources in each application group.

For SAP the following file systems are mounted on all nodes:

`/usr/sap/W70, /sapmnt/W70, /oracle/W70/saptrace and /usr/sap/trans`

The following file systems are used by Oracle for data storage and they are mounted on all nodes:

`/oracle/W70/102_64, /oracle/W70/oraarch, /oracle/W70/saparch, /oracle/W70/sapbackup,  
/oracle/W70/sapcheck, /oracle/W70/sapreorg,  
/oracle/W70/sapdata1, /oracle/W70/sapdata2, /oracle/W70/sapdata3, /oracle/W70/sapdata4  
/oracle/W70/origlogA, /oracle/W70/origlogB, /oracle/W70/origlogC, /oracle/W70/origlogD,  
/oracle/W70/mirrlogA, /oracle/W70/mirrlogB, /oracle/W70/mirrlogC, /oracle/W70/mirrlogD,`



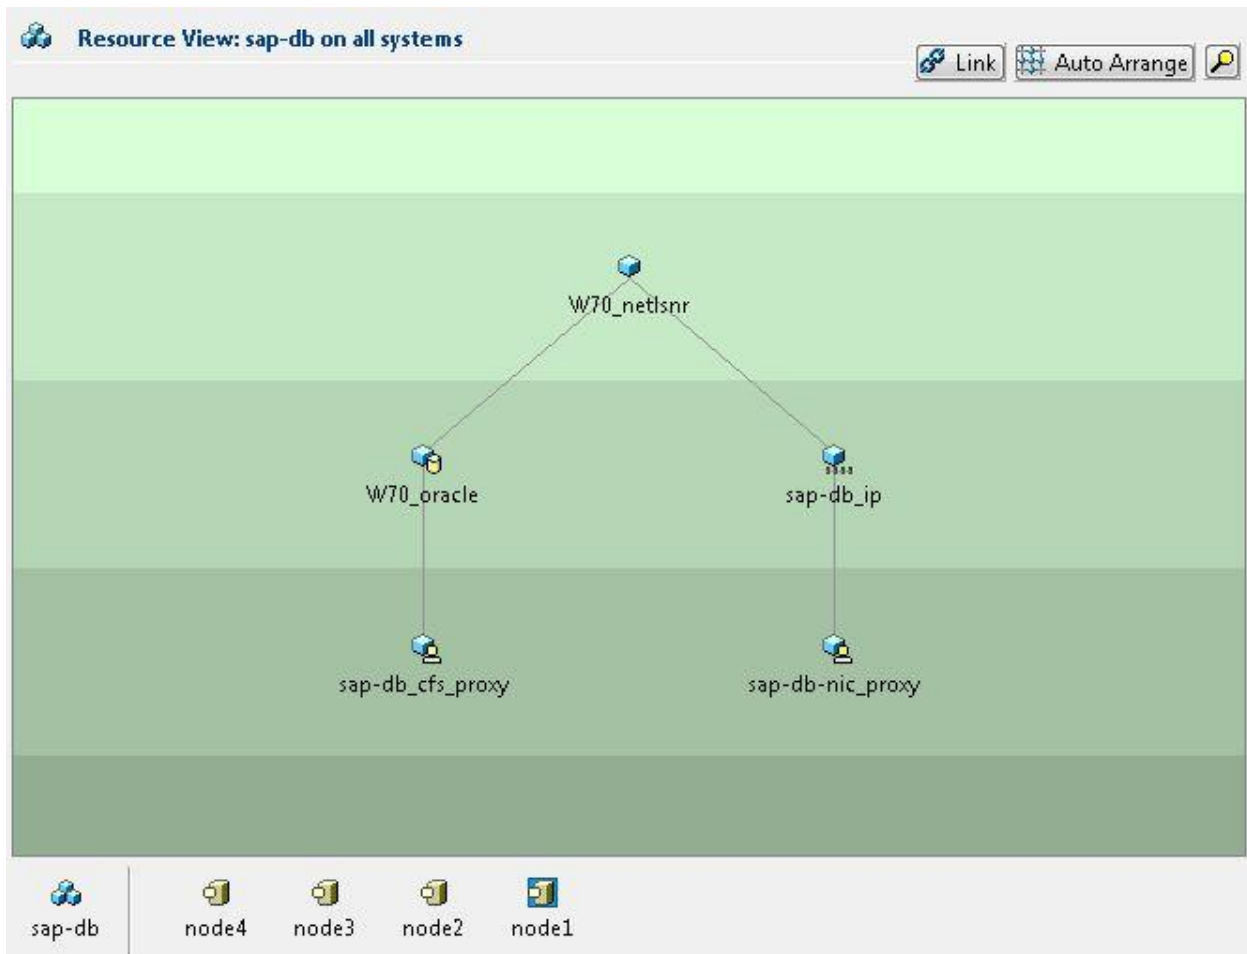
**Figure 6 Resources in the cfs group controlling the cluster file systems.**

### ***Configuring the database for high availability***

The *sap-db* service group contains the resources needed to cluster a single instance Oracle 10g and one database Listener. VCS has the ability to move this service group between the four nodes automatically or manually when a failure is detected.

This configuration could easily be modified to include a third node as a fail-over target so that the Central Instance and Oracle do not need share resources if a failure occurs. If availability requirements are great, the single instance Oracle configuration can be upgraded to an Oracle RAC configuration. For further information see the [Datasheet for Veritas Storage Foundation for Oracle RAC](#).

For detailed instructions on how to configure Oracle in Veritas Cluster Server, please read the [Veritas Cluster Server Oracle Agent's Installation and Configuration Guide](#).



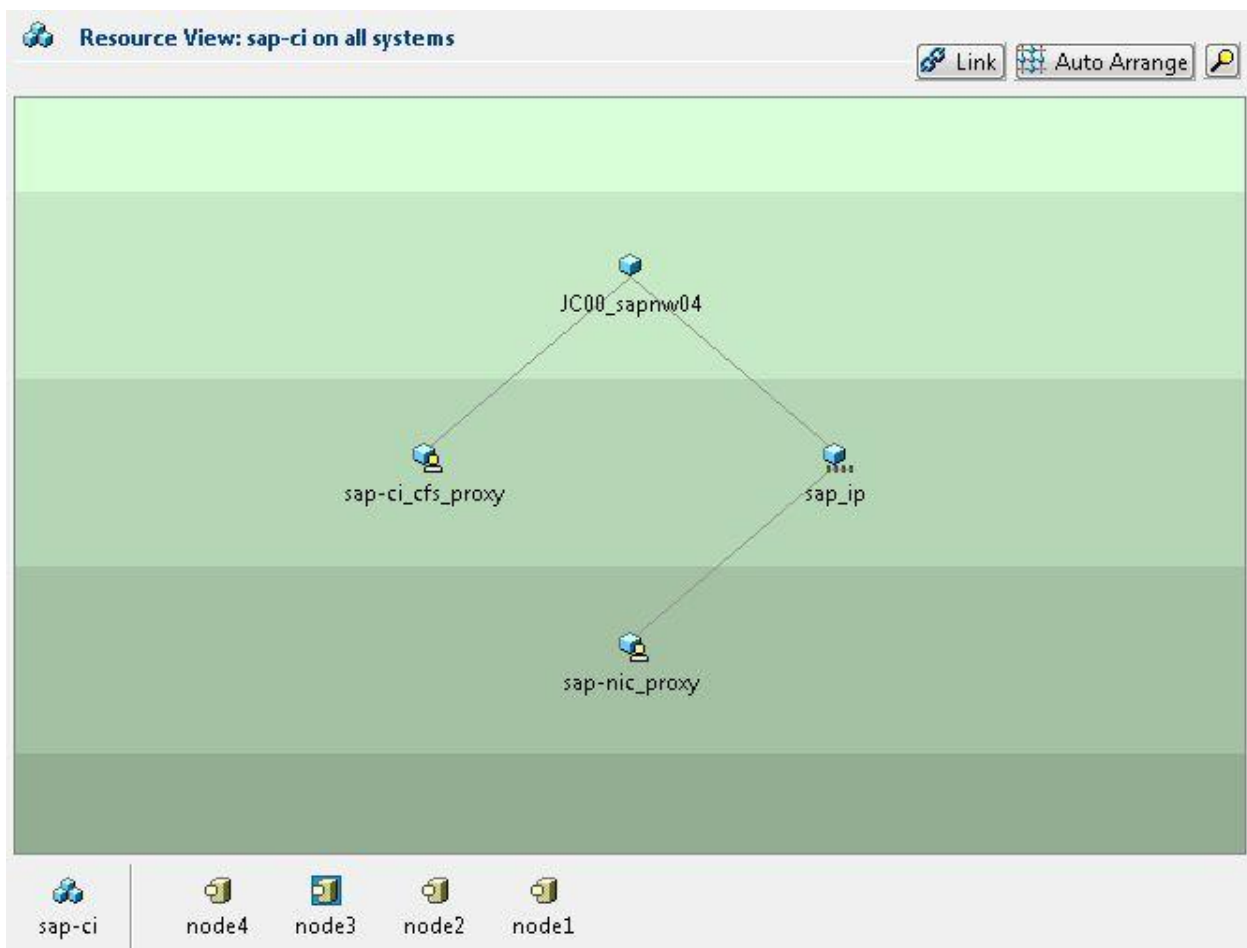
**Figure 7 Resources controlling the Oracle database**

### **SAP Central Instance**

The service group controlling the Central Instance contains only the resources needed for the virtual IP address and the SAPNW04 resource that controls the SAP instance.

The SAPNW04 agent is the VCS agent that is used to monitor the appropriate SAP services, and in the case of the Central Instance the agent will monitor the Dispatcher and Server Processes.

A service group dependency on the *sap-db* group ensures that the database is available and the proxy resource *sap-db\_cfs\_proxy* ensures that the SAP central instance is not started until the file system it relies on has been successfully mounted.



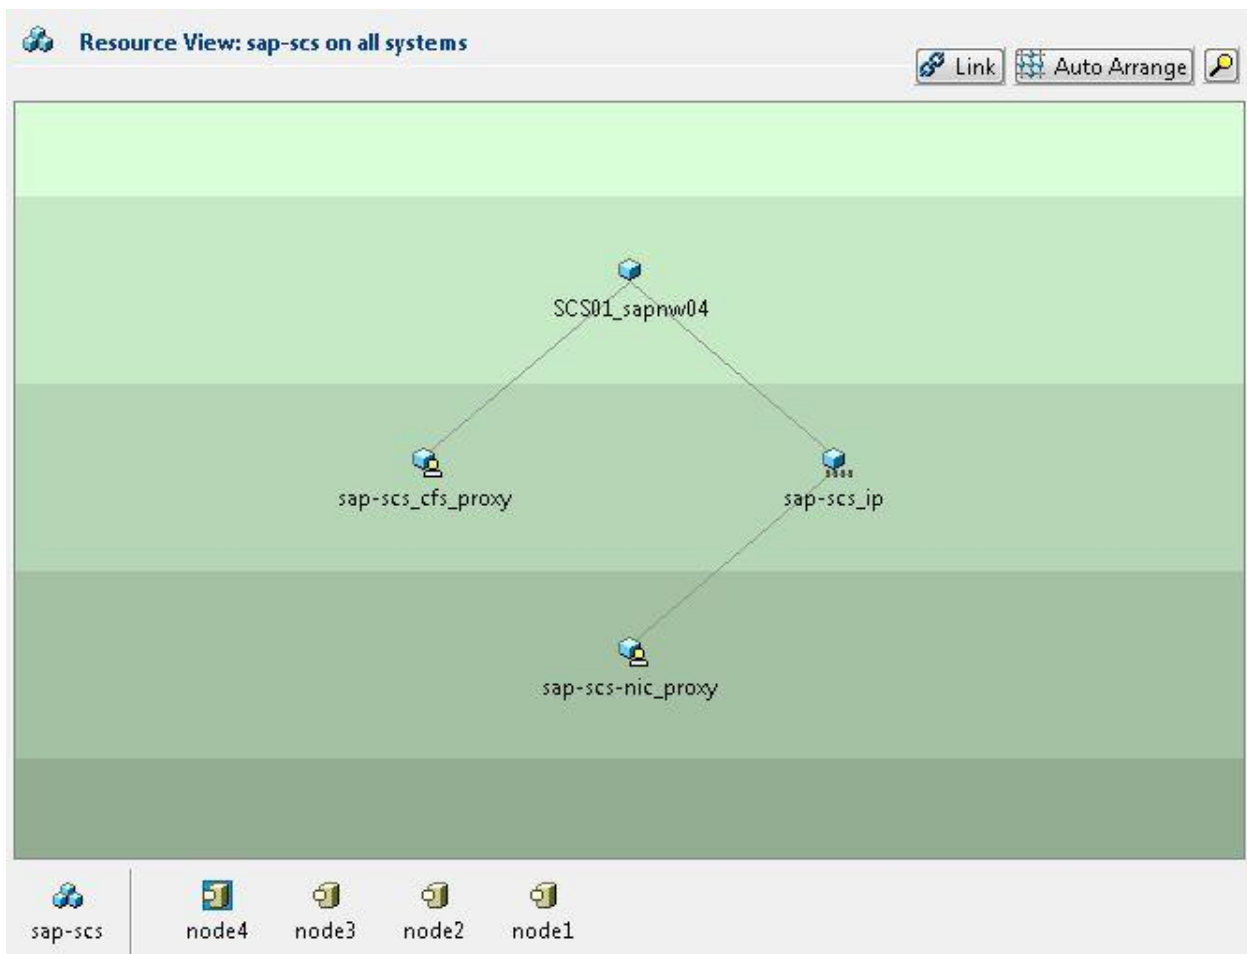
**Figure 8 Resources for the SAP NetWeaver Central Instance**

### **SAP Central Services**

The service group controlling the Central Services contains only the resources needed for the virtual IP address and the SAPNW04 resource that controls the SAP instance.

The SAPNW04 agent is the VCS agent that is used to monitor the appropriate SAP services, and in the case of the Central Services the agent will monitor the Enqueue and Messaging Services.

A service group dependency on the *cfs* group ensures that the SAP central instance is not started until the file system it relies on has been successfully mounted.



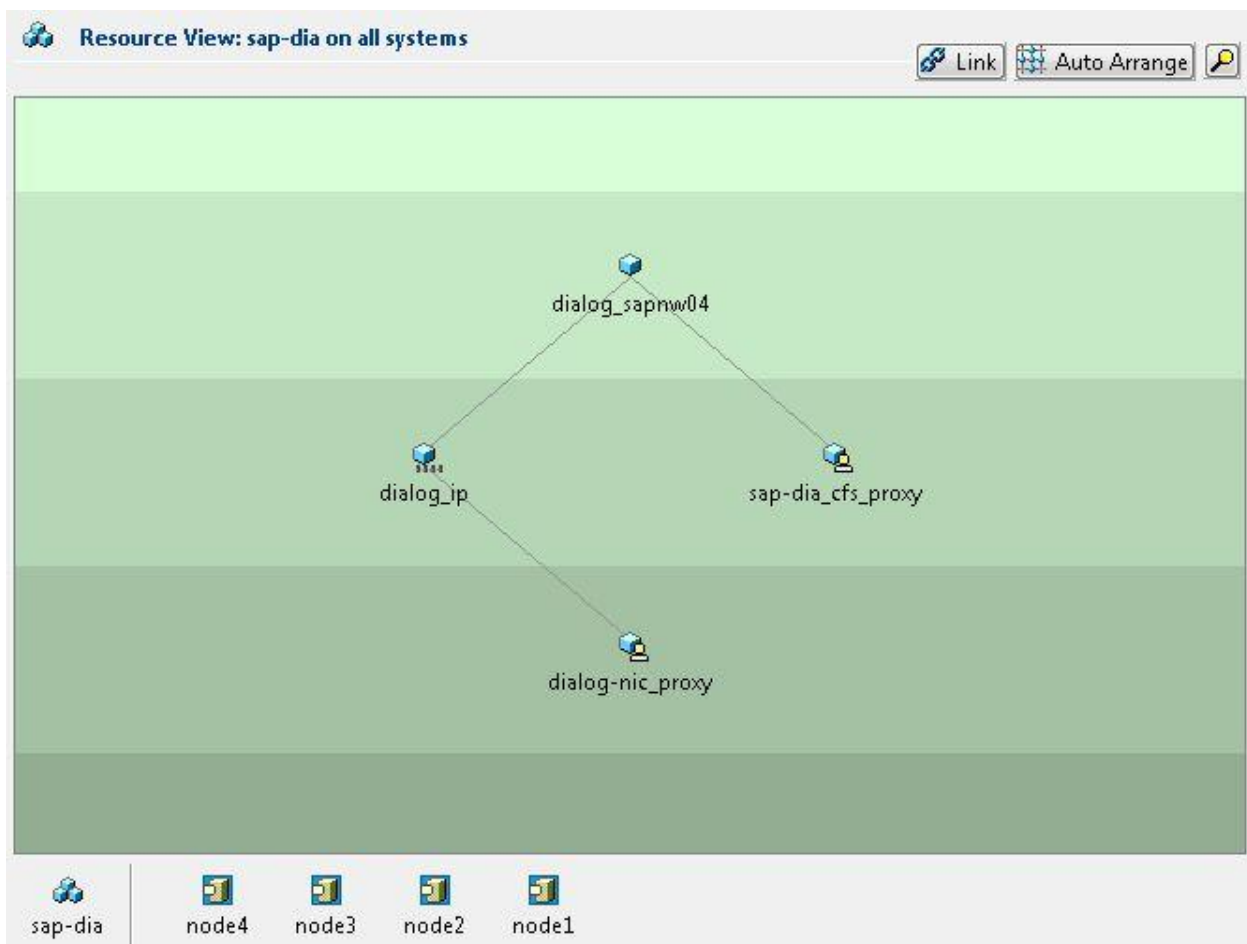
**Figure 9 Resources for the SAP NetWeaver Central Services**



### Dialog Instances

The SAP Dialog Instances have been configured as virtual servers to facilitate easy expansion. By configuring the Dialog server to use a virtual IP address instead of the host IP address the Dialog servers can easily be migrated to another node if more are added. This is primarily done to allow for future configuration changes and not to increase the availability as these service groups are not configured for failover. The SAP architecture allows for multiple Dialog servers, so each one does not necessarily need to be highly available as long as there are others available.

A service group dependency on the *sap-db* group ensures that the database is available. The proxy resource *sap-dia\_cfs\_proxy* ensures that the Dialog instances are not started until the cluster file systems have been successfully mounted.



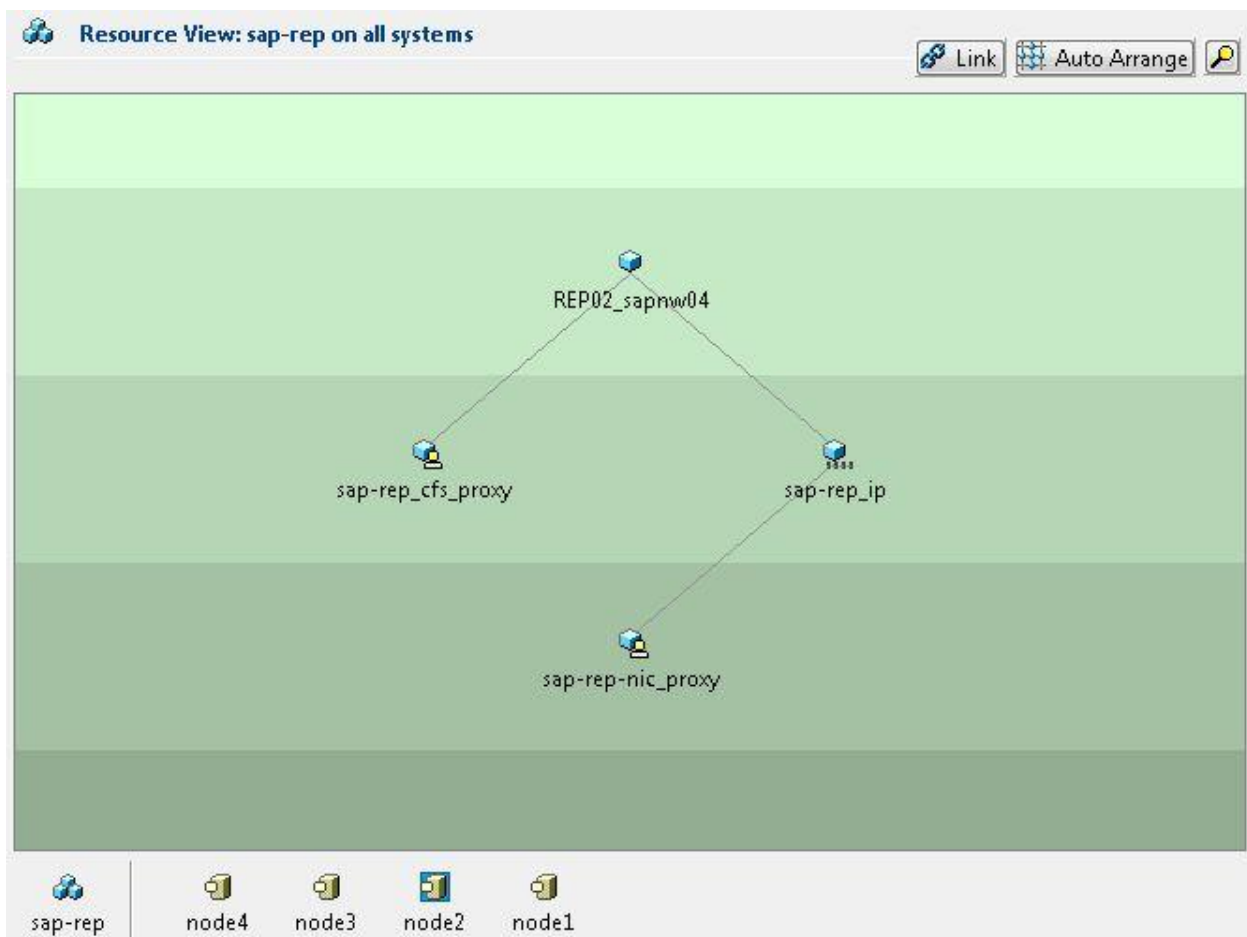
**Figure 10 Resources for the SAP NetWeaver Dialog Instances**

### **Replicated Enqueue Service**

The replicated Enqueue Service is strongly recommended by SAP for HA deployments. The Replicated Enqueue Service enables a transparent failover of the Enqueue Service without impacting open transactions or requiring Java instances to be restarted.

To facilitate high availability the Replicated Enqueue Service has been configured as a virtual service using a movable IP address instead of the hosts IP address, allowing the cluster server to migrate the service to another cluster node if a problem is detected.

A service group dependency to the *cfs* group ensures that the Replicated Enqueue Service is not started until the file systems it relies on have been successfully mounted.



**Figure 11 Resources for the SAP NetWeaver Replicated Enqueue Service**

## Using Veritas Storage Foundation Cluster File System to increase availability of SAP NetWeaver

### Summary

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. From time to time, faults will occur within the data center. The unparalleled integration of file system and clustered application management of CFS and VCS ensures that faults which arise are detected and handled proactively. The Veritas Storage Foundation Cluster File System, in combination with the Veritas Cluster Server, provides dependable high availability of enterprise applications, protecting against unplanned downtime and minimizing exposure during planned maintenance as well.

The combination of SAP NetWeaver with Veritas Storage Foundation Cluster File System and Veritas Cluster Server from Symantec delivers significant benefits for organizations.

- Parallel access to binaries and configuration files
- Failover times are reduced because volumes and file systems do not need to be brought online after a recovery as they are already available to all nodes in the cluster
- Availability of SAP is improved in three areas:
  - CFS file lock management is significantly enhanced compared to NFS, which makes failover more reliable
  - CFS integrates up and down the clustering software stack from the Veritas Cluster Server to the Veritas Dynamic Multi-Pathing product for comprehensive fault detection
  - The redundancy of the SAP solution can be automatically restored after a failover occurs
- CFS has database accelerators for optimal database performance

### Where to get more information

For more information please visit us at

[www.symantec.com](http://www.symantec.com)

or visit the Symantec Storage Foundation Cluster File System product page at

[www.symantec.com/business/products/overview.jsp?pcid=2245&pvid=209\\_1](http://www.symantec.com/business/products/overview.jsp?pcid=2245&pvid=209_1)

## Appendix

### Veritas Cluster Server configuration file: main.cf

The configuration file included below can be used to rebuild or study the 4-node Veritas Storage Foundation Cluster File System described previously in this document.

```
include "types.cf"
include "CFSTypes.cf"
include "CVMTypes.cf"
include "OracleTypes.cf"
include "SAPNW04Types.cf"

cluster SAP (
    UserNames = { admin = ElmElgLimHmmKumGlj }
    Administrators = { admin }
)

system node1 (
)

system node2 (
)

system node3 (
)

system node4 (
)

group cfs (
    SystemList = { node1 = 0, node2 = 1, node3 = 2, node4 = 3 }
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { node1, node2, node3, node4 }
)

    CFSMount sapdata1_cfsmount (
        MountPoint = "/oracle/W70/sapdata1"
        BlockDevice = "/dev/vx/dsk/oracle/sapdata1"
    )

    CFSMount sapdata2_cfsmount (
        MountPoint = "/oracle/W70/sapdata2"
        BlockDevice = "/dev/vx/dsk/oracle/sapdata2"
    )

    CFSMount sapdata3_cfsmount (
        MountPoint = "/oracle/W70/sapdata3"
        BlockDevice = "/dev/vx/dsk/oracle/sapdata3"
    )

    CFSMount sapdata4_cfsmount (
        MountPoint = "/oracle/W70/sapdata4"
        BlockDevice = "/dev/vx/dsk/oracle/sapdata4"
    )

    CFSMount saparch_cfsmount (
        MountPoint = "/oracle/W70/saparch"
        BlockDevice = "/dev/vx/dsk/oracle/saparch"
    )

    CFSMount origlogA_cfsmount (
        MountPoint = "/oracle/W70/origlogA"
        BlockDevice = "/dev/vx/dsk/oracle/origlogA"
    )

    CFSMount origlogB_cfsmount (
        MountPoint = "/oracle/W70/origlogB"
    )
```

```
        BlockDevice = "/dev/vx/dsk/oracle/origlogB"
    )

    CFSMount origlogC_cfsmount (
        MountPoint = "/oracle/W70/origlogC"
        BlockDevice = "/dev/vx/dsk/oracle/origlogC"
    )

    CFSMount origlogD_cfsmount (
        MountPoint = "/oracle/W70/origlogD"
        BlockDevice = "/dev/vx/dsk/oracle/origlogD"
    )

    CFSMount mirrlogA_cfsmount (
        MountPoint = "/oracle/W70/mirrlogA"
        BlockDevice = "/dev/vx/dsk/oracle/mirrlogA"
    )

    CFSMount mirrlogB_cfsmount (
        MountPoint = "/oracle/W70/mirrlogB"
        BlockDevice = "/dev/vx/dsk/oracle/mirrlogB"
    )

    CFSMount mirrlogC_cfsmount (
        MountPoint = "/oracle/W70/mirrlogC"
        BlockDevice = "/dev/vx/dsk/oracle/mirrlogC"
    )

    CFSMount mirrlogD_cfsmount (
        MountPoint = "/oracle/W70/mirrlogD"
        BlockDevice = "/dev/vx/dsk/oracle/mirrlogD"
    )

    CFSMount sapmnt_W70_cfsmount (
        MountPoint = "/sapmnt/W70"
        BlockDevice = "/dev/vx/dsk/sapdg/sapmnt_W70"
    )

    CFSMount usrsap_W70_cfsmount (
        MountPoint = "/usr/sap/W70"
        BlockDevice = "/dev/vx/dsk/sapdg/usrsap_W70"
    )

    CFSMount usrsaptrans_cfsmount (
        MountPoint = "/usr/sap/trans"
        BlockDevice = "/dev/vx/dsk/sapdg/usrsaptrans"
    )

    CFSMount saptrace_W70_cfsmount (
        MountPoint = "/oracle/W70/saptrace"
        BlockDevice = "/dev/vx/dsk/sapdg/saptrace_W70"
    )

    CFSMount orahome_cfsmount (
        MountPoint = "/oracle/W70/102_64"
        BlockDevice = "/dev/vx/dsk/oracle/orahome"
    )

    CFSMount oraarch_cfsmount (
        MountPoint = "/oracle/W70/oraarch"
        BlockDevice = "/dev/vx/dsk/oracle/oraarch"
    )

    CFSMount sapbackup_cfsmount (
        MountPoint = "/oracle/W70/sapbackup"
        BlockDevice = "/dev/vx/dsk/oracle/sapbackup"
    )

    CFSMount sapcheck_cfsmount (
        MountPoint = "/oracle/W70/sapcheck"
        BlockDevice = "/dev/vx/dsk/oracle/sapcheck"
    )

    CFSMount sapreorg_cfsmount (
        MountPoint = "/oracle/W70/sapreorg"
```

```

        BlockDevice = "/dev/vx/dsk/oracle/sapreorg"
    )

    CVMVolDg oracle_cvmvoldg (
        CVMDiskGroup = sapdg
        CVMVolume = { orahome, oraarch, sapdata1, sapdata2, sapdata3, sapdata4,
            saparch, sapbackup, sapcheck, sapreorg,
            origlogA, origlogB, origlogC, origlogD,
            mirrlogA, mirrlogB, mirrlogC, mirrlogD }
        CVMActivation @node1 = sw
        CVMActivation @node2 = sw
        CVMActivation @node3 = sw
        CVMActivation @node4 = sw
    )

    CVMVolDg sap_cvmvoldg (
        CVMDiskGroup = sapdg
        CVMVolume = { sapmnt_W70, usrsap_W70, usrsaptrans, saptrace_W70 }
        CVMActivation @node1 = sw
        CVMActivation @node2 = sw
        CVMActivation @node3 = sw
        CVMActivation @node4 = sw
    )

    Phantom cfs_online_phantom (
        Critical = 0
    )

    requires group cvm online local firm
    sapdata1_cfsmount requires oracle_cvmvoldg
    sapdata2_cfsmount requires oracle_cvmvoldg
    sapdata3_cfsmount requires oracle_cvmvoldg
    sapdata4_cfsmount requires oracle_cvmvoldg
    saparch_cfsmount requires oracle_cvmvoldg
    origlogA_cfsmount requires oracle_cvmvoldg
    origlogB_cfsmount requires oracle_cvmvoldg
    origlogC_cfsmount requires oracle_cvmvoldg
    origlogD_cfsmount requires oracle_cvmvoldg
    mirrlogA_cfsmount requires oracle_cvmvoldg
    mirrlogB_cfsmount requires oracle_cvmvoldg
    mirrlogC_cfsmount requires oracle_cvmvoldg
    mirrlogD_cfsmount requires oracle_cvmvoldg
    sapmnt_W70_cfsmount requires sap_cvmvoldg
    usrsap_W70_cfsmount requires sap_cvmvoldg
    usrsaptrans_cfsmount requires sap_cvmvoldg
    saptrace_W70_cfsmount requires sap_cvmvoldg
    orahome_cfsmount requires oracle_cvmvoldg
    oraarch_cfsmount requires oracle_cvmvoldg
    sapbackup_cfsmount requires oracle_cvmvoldg
    sapcheck_cfsmount requires oracle_cvmvoldg
    sapreorg_cfsmount requires oracle_cvmvoldg
    cfs_online_phantom requires mirrlogA_cfsmount
    cfs_online_phantom requires mirrlogB_cfsmount
    cfs_online_phantom requires mirrlogC_cfsmount
    cfs_online_phantom requires mirrlogD_cfsmount
    cfs_online_phantom requires oraarch_cfsmount
    cfs_online_phantom requires orahome_cfsmount
    cfs_online_phantom requires sapbackup_cfsmount
    cfs_online_phantom requires sapcheck_cfsmount
    cfs_online_phantom requires sapreorg_cfsmount
    cfs_online_phantom requires usrsap_W70_cfsmount
    cfs_online_phantom requires sapmnt_W70_cfsmount
    cfs_online_phantom requires saptrace_W70_cfsmount
    cfs_online_phantom requires usrsaptrans_cfsmount
    cfs_online_phantom requires sapdata2_cfsmount
    cfs_online_phantom requires sapdata3_cfsmount
    cfs_online_phantom requires sapdata4_cfsmount
    cfs_online_phantom requires origlogA_cfsmount
    cfs_online_phantom requires origlogB_cfsmount
    cfs_online_phantom requires origlogC_cfsmount
    cfs_online_phantom requires origlogD_cfsmount
    cfs_online_phantom requires saparch_cfsmount
    cfs_online_phantom requires sapdata1_cfsmount

```

```
// resource dependency tree
//
//   group cfs
//   {
//   Phantom cfs_online_phantom
//   {
//     CFSMount mirrlogA_cfsmount
//     {
//       CVMVolDg oracle_cvmvoldg
//     }
//     CFSMount mirrlogB_cfsmount
//     {
//       CVMVolDg oracle_cvmvoldg
//     }
//     CFSMount mirrlogC_cfsmount
//     {
//       CVMVolDg oracle_cvmvoldg
//     }
//     CFSMount mirrlogD_cfsmount
//     {
//       CVMVolDg oracle_cvmvoldg
//     }
//     CFSMount oraarch_cfsmount
//     {
//       CVMVolDg oracle_cvmvoldg
//     }
//     CFSMount orahome_cfsmount
//     {
//       CVMVolDg oracle_cvmvoldg
//     }
//     CFSMount sapbackup_cfsmount
//     {
//       CVMVolDg oracle_cvmvoldg
//     }
//     CFSMount sapcheck_cfsmount
//     {
//       CVMVolDg oracle_cvmvoldg
//     }
//     CFSMount sapreorg_cfsmount
//     {
//       CVMVolDg oracle_cvmvoldg
//     }
//     CFSMount usrsap_W70_cfsmount
//     {
//       CVMVolDg sap_cvmvoldg
//     }
//     CFSMount sapmnt_W70_cfsmount
//     {
//       CVMVolDg sap_cvmvoldg
//     }
//     CFSMount saptrace_W70_cfsmount
//     {
//       CVMVolDg sap_cvmvoldg
//     }
//     CFSMount usrsaptrans_cfsmount
//     {
//       CVMVolDg sap_cvmvoldg
//     }
//     CFSMount sapdata2_cfsmount
//     {
//       CVMVolDg oracle_cvmvoldg
//     }
//     CFSMount sapdata3_cfsmount
//     {
//       CVMVolDg oracle_cvmvoldg
//     }
//     CFSMount sapdata4_cfsmount
//     {
//       CVMVolDg oracle_cvmvoldg
//     }
//     CFSMount origlogA_cfsmount
//     {
//       CVMVolDg oracle_cvmvoldg
//     }
//     CFSMount origlogB_cfsmount
```

```
//      {
//      CVMVolDg oracle_cvmvoldg
//      }
//      CFSSMount origlogC_cfsmount
//      {
//      CVMVolDg oracle_cvmvoldg
//      }
//      CFSSMount origlogD_cfsmount
//      {
//      CVMVolDg oracle_cvmvoldg
//      }
//      CFSSMount saparch_cfsmount
//      {
//      CVMVolDg oracle_cvmvoldg
//      }
//      CFSSMount sapdata1_cfsmount
//      {
//      CVMVolDg oracle_cvmvoldg
//      }
//      }
//      }

group cvm (
    SystemList = { node1 = 0, node2 = 1, node3 = 2, node4 = 3 }
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { node1, node2, node3, node4 }
)

CFSfsckd vxfsckd (
    ActivationMode @node1 = { sapdg = sw, oracledg = sw }
    ActivationMode @node2 = { sapdg = sw, oracledg = sw }
    ActivationMode @node3 = { sapdg = sw, oracledg = sw }
    ActivationMode @node4 = { sapdg = sw, oracledg = sw }
)

CVMCluster cvm_clus (
    CVMClustName = sap
    CVMNodeId = { node1 = 0, node2 = 1, node3 = 2, node4 = 3 }
    CVMTransport = gab
    CVMTimeout = 200
)

CVMVxconfigd vxconfigd (
    Critical = 0
    CVMVxconfigdArgs = { syslog }
)

vxfsckd requires cvm_clus
cvm_clus requires vxconfigd

// resource dependency tree
//
//      group cvm
//      {
//      CFSfsckd vxfsckd
//      {
//      CVMCluster cvm_clus
//      {
//      CVMVxconfigd vxconfigd
//      }
//      }
//      }
//      }

group public_network (
    SystemList = { node1 = 0, node2 = 1, node3 = 2, node4 = 3 }
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { node1, node2, node3, node4 }
)

NIC public_ce0 (
```



```

        Device = ce0
    )

    Phantom public_network_phantom (
        Critical = 0
    )

    // resource dependency tree
    //
    //     group public_network
    //     {
    //         NIC public_ce0
    //         Phantom public_network_phantom
    //     }

group sap-ci (
    SystemList = { node1 = 0, node2 = 1, node3 = 2, node4 = 3 }
    AutoStartList = { node3 }
)

    IP sap_ip (
        Device = ce0
        Address = "192.168.0.100"
        NetMask = "255.255.255.0"
    )

    Proxy sap-nic_proxy (
        TargetResName = public_ce0
    )

    Proxy sap-ci_cfs_proxy (
        TargetResName = cfs_online_phantom
    )

    SAPNW04 JC00_sapnw04 (
        EnvFile = "/home/w70adm/env.vcs"
        InstName = JC00
        ProcMon = "dw jc"
        SAPAdmin = w70adm
        SAPMonHome = "/usr/sap/W70/SYS/exe/run"
        StartProfile = "/sapmnt/W70/profile/START_JC00_sap70ci"
    )

    requires group sap-db online global soft
    sap_ip requires sap-nic_proxy
    JC00_sapnw04 requires sap-ci_cfs_proxy
    JC00_sapnw04 requires sap_ip

    // resource dependency tree
    //
    //     group sap-ci
    //     {
    //         SAPNW04 JC00_sapnw04
    //         {
    //             Proxy sap-ci_cfs_proxy
    //             IP sap_ip
    //             {
    //                 Proxy sap-nic_proxy
    //             }
    //         }
    //     }

group sap-db (
    SystemList = { node1 = 0, node2 = 1, node3 = 2, node4 = 3 }
    AutoStartList = { node1 }
)

    IP sap-db_ip (
        Device = ce0
        Address = "192.168.0.50"
    )

```

```

        NetMask = "255.255.255.0"
    )

    Netlsnr W70_netlsnr (
        Owner = oracle
        Home = "/oracle/W70/102_64"
        TnsAdmin = "/oracle/W70/102_64/network/admin"
        Listener = LISTENER
    )

    Oracle W70_oracle (
        Sid = W70
        Owner = oracle
        Home = "/oracle/W70/102_64"
        Pfile = "/oracle/W70/102_64/dbs/initW70.ora"
    )

    Proxy sap-db-nic_proxy (
        TargetResName = public_ce0
    )

    Proxy sap-db_cfs_proxy (
        Enabled = 0
        TargetResName = cfs_online_phantom
    )

    requires group cfs online local firm
    sap-db_ip requires sap-db-nic_proxy
    W70_netlsnr requires W70_oracle
    W70_netlsnr requires sap-db_ip
    W70_oracle requires sap-db_cfs_proxy

    // resource dependency tree
    //
    //     group sap-db
    //     {
    //         Netlsnr W70_netlsnr
    //         {
    //             Oracle W70_oracle
    //             {
    //                 Proxy sap-db_cfs_proxy
    //             }
    //         }
    //         IP sap-db_ip
    //         {
    //             Proxy sap-db-nic_proxy
    //         }
    //     }
    //
    // }

group sap-dia (
    SystemList = { node1 = 0, node2 = 1, node3 = 2, node4 = 3 }
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { node1, node2, node3, node4 }
)

IP dialog_ip (
    Device = ce0
    Address @node1 = "192.168.0.111"
    Address @node2 = "192.168.0.112"
    Address @node3 = "192.168.0.113"
    Address @node4 = "192.168.0.114"
    NetMask = "255.255.255.0"
)

Proxy dialog-nic_proxy (
    TargetResName = public_ce0
)

Proxy sap-dia_cfs_proxy (
    TargetResName = cfs_online_phantom
)

```

```

SAPNW04 dialog_sapnw04 (
    EnvFile = "/home/w70adm/env.vcs"
    InstName @node1 = D10
    InstName @node2 = D11
    InstName @node3 = D12
    InstName @node4 = D13
    InstType = DIALOG
    ProcMon = "dw jc"
    SAPAdmin = w70adm
    SAPMonHome = "/usr/sap/W70/SYS/exe/run"
    SAPSID = W70
    StartProfile @node1 = "/usr/sap/W70/SYS/profile/START_D11"
    StartProfile @node2 = "/usr/sap/W70/SYS/profile/START_D12"
    StartProfile @node3 = "/usr/sap/W70/SYS/profile/START_D13"
    StartProfile @node4 = "/usr/sap/W70/SYS/profile/START_D14"
)

requires group sap-db online global firm
dialog_ip requires dialog-nic_proxy
dialog_sapnw04 requires sap-dia_cfs_proxy
dialog_sapnw04 requires dialog_ip

// resource dependency tree
//
//      group sap-dia
//      {
//      SAPNW04 dialog_sapnw04
//      {
//      Proxy sap-dia_cfs_proxy
//      IP dialog_ip
//      {
//      Proxy dialog-nic_proxy
//      }
//      }
//      }

group sap-rep (
    SystemList = { node1 = 0, node2 = 1, node3 = 2, node4 = 3 }
    AutoStartList = { node2 }
    PreOnline = 1
)

IP sap-rep_ip (
    Device = ce0
    Address = "192.168.0.102"
    NetMask = "255.255.255.0"
)

Proxy sap-rep-nic_proxy (
    TargetResName = public_ce0
)

Proxy sap-rep_cfs_proxy (
    TargetResName = cfs_online_phantom
)

SAPNW04 REP02_sapnw04 (
    EnqSrvResName = SCS01_sapnw04
    EnvFile = "/home/w70adm/env.vcs"
    InstName = REP02
    InstType = JENQREP
    ProcMon = enr
    SAPAdmin = w70adm
    SAPMonHome = "/usr/sap/W70/SYS/exe/run"
    SAPSID = W70
    StartProfile = "/usr/sap/W70/SYS/profile/START_REP02"
)

requires group cfs online local firm
sap-rep_ip requires sap-rep-nic_proxy
REP02_sapnw04 requires sap-rep_cfs_proxy
REP02_sapnw04 requires sap-rep_ip

```

```
// resource dependency tree
//
//      group sap-rep
//      {
//      SAPNW04 REP02_sapnw04
//      {
//      Proxy sap-rep_cfs_proxy
//      IP sap-rep_ip
//      {
//      Proxy sap-rep-nic_proxy
//      }
//      }
//      }

group sap-scs (
  SystemList = { node1 = 0, node2 = 1, node3 = 2, node4 = 3 }
  AutoStartList = { node4 }
  PreOnline = 1
)

  IP sap-scs_ip (
    Device = ce0
    Address = "192.168.0.101"
    NetMask = "255.255.255.0"
  )

  Proxy sap-scs-nic_proxy (
    TargetResName = public_ce0
  )

  Proxy sap-scs_cfs_proxy (
    TargetResName = cfs_online_phantom
  )

  SAPNW04 SCS01_sapnw04 (
    EnvFile = "/home/w70adm/env.vcs"
    InstName = SCS01
    InstType = JENQUEUE
    ProcMon = "en ms"
    SAPAdmin = w70adm
    SAPMonHome = "/usr/sap/W70/SYS/exe/run"
    SAPSID = W70
    StartProfile = "/usr/sap/W70/SYS/profile/START_SCS01"
  )

  requires group cfs online local firm
  sap-scs_ip requires sap-scs-nic_proxy
  SCS01_sapnw04 requires sap-scs_cfs_proxy
  SCS01_sapnw04 requires sap-scs_ip

// resource dependency tree
//
//      group sap-scs
//      {
//      SAPNW04 SCS01_sapnw04
//      {
//      Proxy sap-scs_cfs_proxy
//      IP sap-scs_ip
//      {
//      Proxy sap-scs-nic_proxy
//      }
//      }
//      }

```

## 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](http://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](http://www.symantec.com)

Copyright © 2007 Symantec Corporation. All rights reserved. Symantec and the Symantec logo are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

10/07