Creating Flexible, Highly Available SAP Solutions

Leveraging Oracle9*i* and Linux on Dell Servers and Dell/EMC Storage

Flexible, scalable, and highly available servers and storage systems are required for SAP® software implementations. IT administrators now have an alternative to the expensive, proprietary UNIX® platforms typically used for SAP environments: cost-effective, standards-based Dell™ servers and Dell/EMC storage leveraging Oracle9*i*™ Real Application Clusters and the Red Hat® Enterprise Linux® OS. The key to success is proper configuration of the hardware as well as the Linux, Oracle9*i*, and SAP software stack. Best practices for infrastructure planning, setup, and installation can help ensure that systems function as expected and enable optimal performance, easy scalability, and quick implementation.

BY DAVID DETWEILER, ACHIM LERNHARD, FLORENZ KLEY, THORSTEN STAERK, AND WOLFGANG TRENKLE

Related Categories:

Clustering

Database

Dell PowerEdge servers

Dell/EMC storage

High availability (HA)

Linux

Oracle

Red Hat Enterprise Linux

SAP

Scalable enterprise

Visit www.dell.com/powersolutions for the complete category index.

Whether upgrading existing SAP systems, migrating to the next SAP software release, or making new investments in mySAP™ technology and applications or the SAP NetWeaver® platform, IT organizations invariably face the paramount consideration of how to enhance value with limited resources. Enabling technologies and service-delivery capabilities in both software and hardware infrastructure can help organizations adapt the IT framework to the expanding boundaries set by evolving enterprise priorities.

Enterprise solutions such as the mySAP Business Suite need high availability and easy scalability—regardless of the size of the organization. Users of SAP solutions demand uninterrupted database access even if hardware or software failures occur. As businesses and organizations grow, their IT requirements can change rapidly. Adding capacity or changing workloads for SAP systems may require reconfiguring or replacing the database server to adapt to the new situation. Deploying a database cluster instead of a single, dedicated server can help maximize overall system availability in such situations. In addition, some SAP components such as the SAP central instance—which includes the Message Service and the Enqueue Replication Service—can be made redundant to enhance the robustness of an SAP environment.

Working closely with SAP, Oracle, Intel, and Red Hat, Dell has leveraged Oracle9*i* Real Application Clusters (RAC) technology to demonstrate the technical viability of running mySAP solutions in a robust, highly available, and scalable standards-based environment. Those wanting to implement

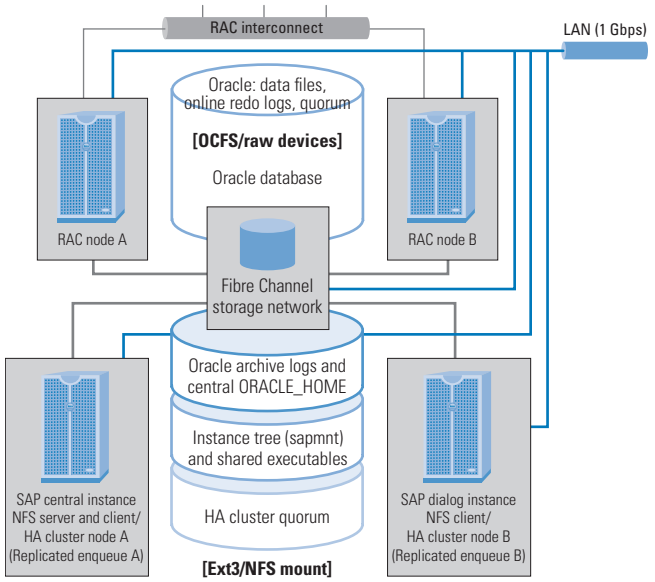

Figure 1. Basic architecture of an SAP central instance using Oracle9*i* RAC databases

an SAP solution on a Linux platform using Oracle9*i* RAC technology have various options for configuring the overall SAP infrastructure to best meet specific business objectives and budgetary requirements. Important aspects of determining the architecture include:

- Whether the SAP system will be used in a production or test/development environment
- The priority and time requirements or constraints of database availability in the overall infrastructure
- The expected growth of the user base or SAP functionality scope
- The importance of continuous SAP central instance availability in the overall SAP system landscape

Figure 1 shows the basic architecture of a fully redundant SAP system that is set up with two Oracle9*i* RAC database nodes and the highly available SAP central instance using the SAP Enqueue Replication Service. This setup can be adapted or scaled to meet various performance and deployment needs, as shown in the example configuration in Figure 2.

IT organizations have various options for possible SAP system architectures—and Dell best practices can be used to help guide installation of these architectures. However, some installation procedures discussed in this article are based on SAP and Oracle requirements for SAP systems, and may deviate from Dell-standard procedures for Oracle9*i* RAC and Red Hat Enterprise Linux installations. *Note:* All server names and IP addresses used in this article exemplify an implementation and should be changed to match an organization's specific requirements.

Setting up Dell servers and storage

When an IT organization decides to implement an Oracle9*i* RAC system on a Dell/Oracle-certified configuration, hardware setup and installation can be performed by Dell Services. This can include physically installing the servers in a rack at the organization's site, connecting all necessary cables, installing the OS (Red Hat Enterprise Linux AS 3 is used in the example scenario in this article), and initializing the shared Dell/EMC storage system.

LUN planning

Optimally, logical unit (LUN) planning should be completed before Dell Services sets up the shared storage and the LUNs, enabling the service representative to immediately set up the LUNs according to the specifically planned SAP implementation. Administrators should determine the LUN requirements according to Oracle and SAP

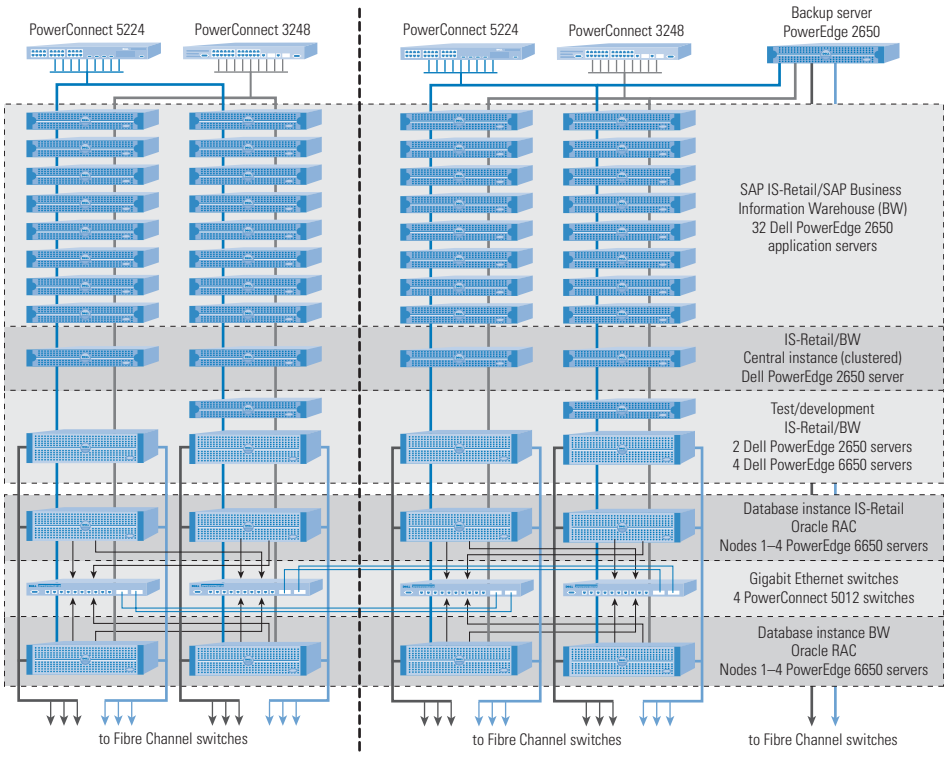

Figure 2. Example deployment of an SAP central instance using Oracle9*i* RAC databases

recommendations for a single-instance Oracle9*i* database for an SAP system. They should then add the LUNs required for RAC communication and the second database node, as well as the LUNs for the Highly Available Network File System (HA NFS) service.

The following LUNs are required:

- A set of LUNs for a single-node Oracle9*i* database for SAP
- A set of LUNs for the redo thread of each additional database instance (online redo logs and the undo/rollback tablespace)
- Two (small) LUNs for the quorum of the Red Hat Enterprise Linux cluster
- A LUN for the shared executables ($ORACLE_HOME and SAP executables)

The LUNs should be organized into two separate storage groups governing access. One storage group should contain all database-relevant LUNs; the other storage group should contain all LUNs belonging to the Red Hat Enterprise Linux cluster.

If installing the database on raw devices and not on OCFS, administrators should be aware that what can be a file in an OCFS implementation must be configured on a separate raw device, and thus must be on a separate LUN or partition:

- A small LUN for the RAC cluster quorum (to ensure exclusive access on the SCSI device level, this must be a separate LUN and must not be a partition on a LUN shared with other partitions)
- A small LUN for the shared configuration file (this can be one of several partitions on a LUN)
- Three or more small LUNs for the database control files (these can be distributed among several partitions on a LUN)

Visual check of the hardware

To verify that systems are wired correctly, administrators should start by visually checking the overall system. The Dell/EMC storage system, the two Dell PowerEdge™ servers for the database cluster, and the two PowerEdge servers for the SAP central instance cluster should each be connected to the Fibre Channel switch. Additional SAP application servers may not need storage area network (SAN) connectivity.

The primary network card in each server and the network interface of the Dell/EMC storage system should be connected to the data center LAN, over which these servers will be accessed from the outside. The network should satisfy SAP requirements for the connection of SAP database servers and SAP application servers.

The secondary interface of the database and the SAP central instance cluster servers should be connected to a private LAN, over which these systems connect only to the other nodes in the respective cluster. In the case of a two-node Oracle9*i* RAC configuration with a two-node SAP central instance cluster, each secondary interface can be connected with a cross-wired Ethernet cable (also referred to as a cross-connect or X-connect) to its clustered counterpart. Alternatively, administrators can connect the clustered servers over a dedicated Gigabit Ethernet switch (a requirement in Oracle9*i* RAC configurations with more than two database nodes). The only fabric currently supported for the database interconnect is Gigabit Ethernet.[1]

ONLINE EXTRA: LOGIN AND MANAGEMENT TOOLS FOR SAP ON ORACLE/LINUX PLATFORMS

To learn more about remote login capabilities, the Linux desktop, and versioning tools for SAP on Oracle/Linux platforms, visit *Dell Power Solutions* online at www.dell.com/powersolutions. This special supplement to the print article also provides example LUN configurations for Oracle9*i* RAC database nodes and high-availability NFS clusters in an SAP environment.

Installing the Linux, Oracle9*i* RAC, and SAP software

When installing Red Hat Enterprise Linux, Oracle9*i*, and SAP software, administrators should heed the following recommendations:

- Install a valid hardware/software combination that is certified by Dell, Red Hat, Oracle, and SAP. Dell-certified hardware for SAP and Linux can be found in SAP Note 300900 or at www.dell.com/sap. Because SAP specifically requires running only SAP-certified Linux kernel and glibc library versions, administrators should make sure to use only these versions and to exclude the Linux kernel and glibc library from software that automatically installs other versions of these packages. The latest information on certified Linux and glibc versions is available at www.sap.com/linux. The latest information on platform availability for Oracle9*i* RAC can be found in SAP Note 527843 (a username and password for the SAP Service Marketplace is required).
- Unless otherwise noted, use the original installation media provided by Red Hat, Oracle, and SAP to help ensure uniformity. *Note:* The Dell Deployment CD for Oracle9*i* RAC can be used to install Red Hat Enterprise Linux and Oracle on Dell PowerEdge servers. However, this CD currently does not support specific SAP requirements on RAC setups for SAP and therefore should not be used for RAC installations in SAP environments.
- Plan and configure storage systems carefully, optimally with the help of a Dell Services representative. Storage plays a

[1] This term does not connote an actual operating speed of 1 Gbps. For high-speed transmission, connection to a Gigabit Ethernet server and network infrastructure is required.

vital role in the success of any SAP project—particularly for implementations that include Oracle9*i* RAC technology.
- Use standard SAP installation procedures and utilities for Oracle9*i*. A single-instance database and the SAP software should be installed and then additional nodes can be added, activating the RAC feature of the database.
- Pay careful attention to the installation and the high-availability aspects of the implementation when configuring Oracle9*i* RAC for SAP software.

Red Hat Enterprise Linux 3

Before installing Red Hat Enterprise Linux 3, administrators should upgrade the Fibre Channel adapter firmware to the most recent version listed in the EMC compatibility matrix (www.emc.com/interoperability/index.jsp). During the installation, administrators should disconnect the Fibre Channel cables from the host bus adapters (HBAs) and set up access to the shared storage with the OS in place. Administrators should install the Red Hat Enterprise Linux OS on each Dell PowerEdge server according to the documentation provided with the server and the Red Hat installation guide. Best practices recommend using English as the default language.

Additional recommendations include the following:

- Partition the hard drives manually to properly reserve space for the mySAP application system data.
- Select "No firewall" when asked in the dialog box to configure the firewall. If a firewall is required, configure it manually after the system installation. In this case, ensure that the required ports on the mySAP servers are available from the other servers. The used ports on the mySAP servers can be ascertained after the installation in /etc/services.
- Select the time zone to be Coordinated Universal Time (UTC) and check the "System clock uses UTC" box. Ensure that the system's BIOS clock is set to UTC.
- Select "Customize the set of packages to be installed" in the Package Default dialog box. In the Package Group Selection dialog box, select the following as a minimum (additional items can be chosen): X Window system, Editors, Development Tools, Legacy Software Development, Administration Tools, and Printing Support.

After the first system boot, a configuration tool should appear. If a Network Time Protocol (NTP) Time Server is available at the site, Dell best practices recommend using this NTP server as a common time source for all servers in the SAP environment. Administrators should select "Enable Network Time Protocol," and under "Date and Time" they should enter the address of the time server.

The system should be registered with Red Hat Network so that update packages are automatically made available. To exclude the Linux kernel and important libraries from this auto-update, administrators should issue the `up2date --configure` command and configure the "Package Exceptions" to exclude the following: kernel*, glibc*, nscd, and nptl-devel.

Post-installation tasks for SAP environments

Once the initial installation and configuration on each server has been completed, administrators must conduct various tasks before proceeding with the Oracle and SAP installation. First, the Linux kernel parameters must be changed to meet the requirements for SAP software. The following lines should be added to the /etc/sysctl.conf file (or those already there should be edited to match the following):

```
#SAP settings
kernel.shmmax=2313682943
kernel.msgmni=1024
kernel.sem=1250 256000 100 1024
fs.file-max=8162
```

Once the preceding lines have been added, administrators should activate these settings with the `sysctl -p` command.

A temporary file system (tmpfs) needs to be configured on each SAP system. Space for the tmpfs is allocated in the system's virtual memory, which is composed of the system's physical memory (RAM) and the configured swap space. Tmpfs holds the SAP system's shared memory. The default setting for tmpfs is generally half of the total size of system memory. This setting can be changed in the /etc/fstab:

```
tmpfs   /dev/shm   tmpfs   size=size 0 0
```

To determine the proper size of tmpfs, administrators must consider the amount of virtual memory available and the memory requirements of the SAP system itself. If the system has 4 GB of RAM and 8 GB of configured swap space, the tmpfs can be limited to 6 GB. Administrators would then edit the line in /etc/fstab so that *size* is "6G" or "6144M":

```
tmpfs   /dev/shm   tmpfs   size=6G 0 0
```

However, administrators should note the following:

- Tmpfs must always be smaller than the amount of virtual memory available.
- Tmpfs should not be larger than twice the system's physical memory.

To avoid a situation in which swap partitions cannot be found, administrators should check that the order of the HBAs in each

system is correct. After the initial installation, they should check in the /etc/modules.conf file to help ensure that the SCSI RAID controller for the internal hard drives comes before the Fibre Channel HBA. For example, on a Dell PowerEdge 2650 server, the /etc/modules.conf file should look similar to the following:

```
alias scsi_hostadapter aacraid
alias scsi_hostadapter1 qla2300
```

If this is not the case, administrators should change the file and run `mkinitrd` to help ensure the order is kept. If using the kernel 2.4.21-15.ELsmp, administrators should issue the following command:

```
mkinitrd initrd-2.4.21-15.ELsmp.img
   2.4.21-15.ELsmp -f
```

Alternatively, administrators can use the device file system (devfs) or devlabel.[2]

Checking the network connections

Administrators must ensure that all servers are properly connected and integrated into the network. The servers each require a public IP address, and name resolving must function among the servers through the use of either host files on all participating systems or Domain Name System (DNS). Because reliance exclusively on host files is waning, best practices recommend using DNS. However, to prevent DNS name resolving from being a point of failure or a source of instability in the cluster, administrators should set up host files on all servers in the database cluster and the SAP central instance cluster, as well as on additional SAP application servers, to resolve names. Also, administrators should have procedures in place to keep those hosts files current—for example, after changing the DNS records.

Additionally, all database nodes must have an internal IP address for the database interconnect. This internal IP address should be reachable and resolvable from the database nodes exclusively, and does not have to be connected to the enterprise's network (see the "Visual check of the hardware" section in this article for information about cable connectivity).

Next, administrators should check the host names of the servers and test whether they can be reached over the network and whether name resolving works. Administrators should issue the `hostname` command, which should return the node name only—not the fully-qualified domain name (FQDN). Although the definition of a host name can vary—for example, whether it should include the domain name—for Oracle and SAP applications, the servers must be set up according to SAP note 722273 ("Red Hat Enterprise Linux 3 and 4: Installation and upgrade"):

```
# hosts for acme.com
127.0.0.1     localhost.localdomain  localhost
172.16.42.31  snoopy.acme.com        snoopy
```

Administrators should check and, if necessary, adapt the /etc/hosts file on each server. The entry for "localhost" must not contain an alias for the host name. Administrators should decide which name-resolving strategy will be used and test the name-resolving capability. They should then perform a similar test for connectivity with the `ping` command. If DNS is used, administrators should test both DNS and file resolution with the appropriate tools, such as `dig` and `getent`.

Checking the Dell/EMC shared storage

For Dell/EMC Fibre Channel storage, administrators should verify that the EMC® Navisphere® software agent is installed on each node. This installation is usually performed by Dell Services during the storage setup. Using the Red Hat Package Manager (RPM™), administrators can verify whether the Navisphere agent is already installed:

```
[root@rac1 root]# rpm -q naviagent
```

The expected response should look like the following:

```
naviagent-6.6.0.3.8-1
```

Version numbers may differ because software is upgraded over time.

Administrators should check whether LUNs have been created according to Oracle and SAP requirements (as discussed in the "LUN planning" section in this article). They also should check that each node is assigned to the correct storage group in the EMC Navisphere software; more information about this can be found in the documentation that accompanies the Dell/EMC Fibre Channel storage system. If not using EMC PowerPath software, administrators should use devlabel to create persistent device names.

Next, administrators should check in /proc/partitions for the shared storage LUNs. A list of the LUNs or logical disks that are detected by the node is displayed, as well as a list of the partitions that have been created on the external devices. Using this information, administrators can create partitions for OCFS volumes or raw

[2] For more information about devlabel, visit linux.dell.com/devlabel or see "Resolving Device Renaming Issues in Linux" by Gary Lerhaupt, *Dell Power Solutions*, February 2003, www1.us.dell.com/content/topics/global.aspx/power/en/ps1q03_lerhaupt.

devices. Administrators should ensure that the same shared disks are visible from all nodes.[3]

The listed devices vary depending on how the storage is configured. The primary SCSI drive or RAID array on each node will be listed as sda and will be partitioned. If any other SCSI disks or RAID arrays are on the node, they will be listed as sdb, sdc, and so on. The LUNs on the Fibre Channel storage system or SCSI enclosure should also be visible as SCSI devices. For example, if the node has one RAID array and the storage system has three logical disks, the node should identify the RAID array or internal disk as sda and the logical disks as sdb, sdc, and sdd. If three LUNs are on the Fibre Channel storage, the node should recognize the RAID array as sda and the Fibre Channel LUNs as emcpowera, emcpowerb, and emcpowerc (assuming EMC PowerPath® software is being used).

If a node does not detect the external storage devices, administrators should perform the following steps:

1. Reload the HBA driver on all nodes to synchronize the kernel's partition tables on the nodes by entering (for QLogic HBAs):

```
rmmod qla2300
modprobe qla2300
```

2. Confirm that all nodes detect the external storage devices by entering:

```
cat /proc/partitions
```

Partitions on shared storage

Even if creating partitions and file systems on shared storage from any of the attached servers is possible, best practices recommend limiting these administrative tasks to only one of the nodes, usually the first node. Because no database or shared file system is running at this stage of the installation, the nodes do not yet have access to the shared storage and consequently do not need to be disconnected from it. However, once a RAC database, an HA NFS server, or a similar service accesses the shared storage from multiple nodes, all other nodes must stop I/O operations and cede control of the storage device on a software level, so that the management node can alter the shared device configuration. Usually, ceding control is accomplished by unloading the Fibre Channel driver module, which is possible only when the database, cluster manager, or a similar service is shut down. Taking into account these organizational measures, administrators should continue creating partitions on the management node as they would on locally attached disks.

Partitions on local disks in Oracle9*i* RAC database nodes

Administrators should create partitions in the local drives on each of the RAC database nodes for the Oracle database software. On the first server, node A, administrators should perform the following steps:

1. Create one partition of 10,000 MB and a partition for the remainder of the disk space using the following command:

```
fdisk /dev/sda
```

2. Reboot the server.
3. After the reboot, format the partitions:

```
mke2fs -j /dev/sda7
mke2fs -j /dev/sda8
```

4. Verify that the entries in the file /etc/fstab are as follows:

```
/dev/sda7   /oracle   ext3   defaults   1 2
/dev/sda8   /sapcd    ext3   defaults   1 2
```

5. Create the /oracle directory:

```
mkdir /oracle
```

6. Create the /sapcd directory:

```
mkdir /sapcd
```

7. Mount the directories according to /etc/fstab:

```
mount -a
```

On the second server, node B, administrators should perform the following steps:

1. Create one partition of 10,000 MB and a partition for the remainder of the disk space using the following command:

```
fdisk /dev/sda
```

2. Reboot the server.
3. After the reboot, format the partition:

```
mke2fs -j /dev/sda6
```

4. Verify that the entries in the file /etc/fstab are as follows:

```
/dev/sda6   /oracle   ext3   defaults   1 2
/dev/sda8   /sapcd    ext3   defaults   1 2
```

[3] For more information about how to ensure that all machines see the same devices under the same name, see "Resolving Device Renaming Issues in Linux" by Gary Lerhaupt, *Dell Power Solutions,* February 2003, www1.us.dell.com/content/topics/global.aspx/power/en/ps1q03_lerhaupt.

5. Create the /oracle directory:

```
mkdir /oracle
```

6. Mount the directory according to /etc/fstab:

```
mount -a
```

Installing OCFS on each RAC database node

Administrators should place Oracle tablespaces and log files on OCFS. To install OCFS on Dell/EMC storage systems, the following packages are required:

- coreutils-debuginfo-4.5.3-35.i386.rpm
- coreutils-4.5.3-35.i386.rpm
- ocfs-2.4.21-EL-smp-1.0.12-1.i686.rpm
- ocfs-support-1.0.10-1.i386.rpm
- ocfs-tools-1.0.10-1.i386.rpm
- ocfs-tools-debuginfo-1.0.10-1.i386.rpm
- tar-debuginfo-1.13.25-16.i386.rpm
- tar-1.13.25-16.i386.rpm

The following packages should be installed with the `-ivh` option:

- coreutils-debuginfo-4.5.3-35.i386.rpm
- ocfs-2.4.21-EL-smp-1.0.12-1.i686.rpm
- ocfs-support-1.0.10-1.i386.rpm
- ocfs-tools-1.0.10-1.i386.rpm
- ocfs-tools-debuginfo-1.0.10-1.i386.rpm
- tar-debuginfo-1.13.25-16.i386.rpm

Packages that should be installed using the `-Uvh` option include coreutils-4.5.3-35.i386.rpm and tar-1.13.25-16.i386.rpm. These two packages contain tools that are OCFS aware; the standard coreutils and tar packages should not be used on OCFS file systems. Once the packages have been installed on both node A and node B, the ocfs.conf file must be edited. On node A, the file should look similar to the following:

```
node_name=nodeA
node_number=0
ip_address=1.0.102.65
ip_port=7000
```

Then, administrators should issue the command `ocfs_uid_gen -c`. This generates a unique node identifier from the information in /etc/ocfs.conf and appends that information to the file. On node B, the ocfs.conf file should look similar to the following:

```
node_name=nodeB
node_number=1
ip_address=1.0.102.66
ip_port=7000
```

As on the first node, administrators should issue the command `ocfs_uid_gen -c` on node B. Next, administrators should enter the `load_ocfs` command on both nodes. At this point, administrators can verify on the system log whether OCFS loaded successfully. Once that is completed, the OCFS partitions must be formatted. Figure 3 shows example formatting commands (SID is the SAP system ID).

```
mkfs.ocfs -b 128 -L saparch -m /oracle/SID/saparch -u 1046 -g 502 /dev/sdc1
mkfs.ocfs -b 128 -L sapdata1 -m /oracle/SID/sapdata1 -u 1046 -g 502 /dev/sdd1
mkfs.ocfs -b 128 -L sapdata2 -m /oracle/SID/sapdata2 -u 1046 -g 502 /dev/sde1
mkfs.ocfs -b 128 -L sapdata3 -m /oracle/SID/sapdata3 -u 1046 -g 502 /dev/sdf1
mkfs.ocfs -b 128 -L sapdata4 -m /oracle/SID/sapdata4 -u 1046 -g 502 /dev/sdg1
mkfs.ocfs -b 128 -L sapdata5 -m /oracle/SID/sapdata5 -u 1046 -g 502 /dev/sdh1
mkfs.ocfs -b 128 -L sapdata6 -m /oracle/SID/sapdata6 -u 1046 -g 502 /dev/sdi1
mkfs.ocfs -b 128 -L origlogA -m /oracle/SID/origlogA -u 1046 -g 502 /dev/sdj1
mkfs.ocfs -b 128 -L origlogB -m /oracle/SID/origlogB -u 1046 -g 502 /dev/sdk1
mkfs.ocfs -b 128 -L mirrlogA -m /oracle/SID/mirrlogA -u 1046 -g 502 /dev/sdl1
mkfs.ocfs -b 128 -L mirrlogB -m /oracle/SID/mirrlogB -u 1046 -g 502 /dev/sdm1
mkfs.ocfs -b 128 -L clusterquorum -m /clusterquorum -u 1046 -g 502 /dev/sdn1
mkfs.ocfs -b 128 -L racquorum -m /racquorum -u 1046 -g 502 /dev/sdo1
```

Figure 3. Commands for formatting OCFS partitions

The next step is to create on all nodes the matching directories as mount points for OCFS:

- /oracle/SID/sapdata1
- /oracle/SID/sapdata2
- /oracle/SID/sapdata3
- /oracle/SID/sapdata4
- /oracle/SID/sapdata5
- /oracle/SID/sapdata6
- /oracle/SID/saparch
- /oracle/SID/origlogA
- /oracle/SID/origlogB
- /oracle/SID/mirrlogA
- /oracle/SID/mirrlogB
- /racquorum

```
LABEL=sapdata1        /oracle/SID/sapdata1   ocfs   _netdev   0 0
LABEL=sapdata2        /oracle/SID/sapdata2   ocfs   _netdev   0 0
LABEL=sapdata3        /oracle/SID/sapdata3   ocfs   _netdev   0 0
LABEL=sapdata4        /oracle/SID/sapdata4   ocfs   _netdev   0 0
LABEL=sapdata5        /oracle/SID/sapdata5   ocfs   _netdev   0 0
LABEL=sapdata6        /oracle/SID/sapdata6   ocfs   _netdev   0 0
LABEL=saparch         /oracle/SID/saparch    ocfs   _netdev   0 0
LABEL=origlogA        /oracle/SID/origlogA   ocfs   _netdev   0 0
LABEL=origlogB        /oracle/SID/origlogB   ocfs   _netdev   0 0
LABEL=mirrlogA        /oracle/SID/mirrlogA   ocfs   _netdev   0 0
LABEL=mirrlogB        /oracle/SID/mirrlogB   ocfs   _netdev   0 0
LABEL=racquorum       /racquorum             ocfs   _netdev   0 0
```

Figure 4. Storage devices in the /etc/fstab directory

Figure 4 shows the entries in the file /etc/fstab. The option `_netdev` is designed to ensure that, during system startup, the OCFS volumes are mounted after the network becomes active, enabling OCFS cluster communication. These devices should be activated with the `mount -a` command. When all the devices on node A have been mounted, administrators should edit the /etc/fstab directory as shown in Figure 4 and issue the `mount -a` command. The ownership for these directories must be changed to the database and user group:

```
chown -R ora<sid>:dba /oracle
```

By completing the steps described in this article, administrators can help ensure that the shared storage for both Oracle9*i* RAC nodes is ready to function as expected.

Planning a highly available, flexible environment for SAP software

SAP software implementations require a robust IT infrastructure to support business-critical processes. Cost-effective, standards-based Dell servers and Dell/EMC storage leveraging Oracle9*i* RAC databases and the Red Hat Enterprise Linux OS can provide a highly available, flexible platform for SAP environments. However, proper planning and configuration—based on best-practices recommendations from Dell, Oracle, Red Hat, and SAP—are key to helping ensure that all systems function optimally.

David Detweiler is the Dell SAP Alliance Manager in Europe, the Middle East, and Africa (EMEA) and a member of the Dell SAP Competence Center in Walldorf, Germany. The Dell SAP Competence Centers help ensure that current and future Dell technologies work together with SAP solutions and provide customers with the architecture, functionality, reliability, and support expected of mission-critical applications.

Achim Lernhard has worked at the Dell SAP Competence Center in Walldorf, Germany, for three years as part of the SAP LinuxLab. He assisted the Oracle9*i* RAC on Linux pilot customer from installation to productivity and worked on the hardware certifications.

Florenz Kley is a consultant for SAP Technology Infrastructure. He has worked for five years at the Dell SAP Competence Center in Walldorf, Germany, as part of the SAP LinuxLab. He conducted performance benchmarks to help prove the scalability and performance of Oracle9*i* RAC for SAP on Linux and helped build the architecture for Dell's Oracle9*i* RAC on Linux pilot customer.

Thorsten Staerk is a consultant at the Dell SAP Competence Center in Walldorf, Germany, as part of the SAP LinuxLab. He has worked extensively on Oracle9*i* RAC technologies for SAP, researches new SAP technologies and functionality, and certifies Dell platforms for SAP on Linux.

Wolfgang Trenkle is a senior consultant at the Dell SAP Competence Center in Walldorf, Germany, and is also a member of the Dell EMEA Enterprise Solutions Center team in Limerick, Ireland. In addition to serving as a consultant and supporting proof of concepts, Wolfgang provides training materials and tools to Dell's global SAP community.

FOR MORE INFORMATION

Red Hat Enterprise Linux 3:
www.redhat.com/docs/manuals/enterprise
www.redhat.com/docs/errata

SAP LinuxLab:
www.sap.com/linux

Dell and Linux:
linux.dell.com/projects.shtml
linux.dell.com/projects.shtml#devlabel

