



An Oracle Technical Article
January 2012

How Dell Migrated from SUSE Linux to Oracle Linux



ORACLE

Introduction

Switching the underlying operating system on a single server is not trivial. Neither is dealing with the related conversion and compatibility issues. Imagine what's involved in switching the operating system on thousands of servers spread globally across an enterprise, like Dell just did.

In June of 2010, Dell made the decision to migrate 1,700 systems from SUSE Linux to Oracle Linux, while leaving the hardware and application layers unchanged. Standardization across the Linux platforms helped make this large-scale conversion possible. The majority of the site-specific operating system and application configuration could simply be backed up and restored directly on the new operating system. Configuration changes were minimal and most could be automated, easing the administration effort required and helping achieve a reliable and consistent transition procedure.

This article describes how Dell planned and implemented the migration, including key conversion issues and an overview of their transition process.

Dell's Deployment Environment

Dell had approximately 1,700 physical systems running SUSE Linux at the start of this migration process. These systems, geographically dispersed around the world, used a mix of eighth-generation (Dell PowerEdge 2850 and 2950 servers) and newer Dell hardware. Fibre Channel SAN storage comprised EMC Symmetrix and CLARiiON devices. The software environment included SUSE Linux 10 Service Pack 1 with multipath I/O (MPIO), Oracle Database 10g Release 2, Oracle Real Application Clusters (Oracle RAC), and Oracle Automatic Storage Management, as shown in Figure 1.

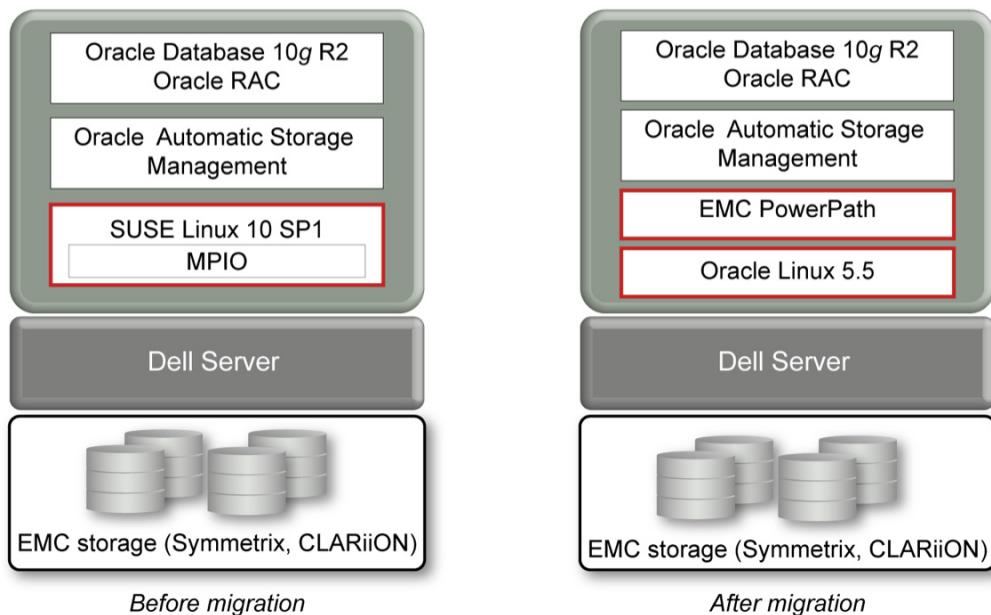


Figure 1. Dell's deployment environment, before and after migration to Oracle Linux

The migration primarily involved the operating system, moving from SUSE Linux 10 to Oracle Linux 5.5. The same physical server and storage hardware was retained during the migration. Similarly, the Oracle

software remained unchanged after the migration. An additional change for Dell was a switch from SUSE Linux's built-in multipath I/O support to EMC PowerPath for automated data path management. (Note: the actual conversion from MPIO to PowerPath is tangential to the operating system migration and is beyond the scope of this document.)

This migration process also served as a time for Dell to re-evaluate its servers running SUSE Linux to determine whether the applications running on these servers could be decommissioned or deployed on an existing MegaGrid environment instead. Dell uses 16-node racks, each capable of hosting 300 databases, for their MegaGrid deployment. In some cases, there was sufficient capacity on an existing MegaGrid infrastructure, and the applications and databases could be migrated to the grid and the SUSE Linux server powered down and decommissioned. This consolidation provided savings in power, cooling, and reduced space requirements. In other cases where consolidation wasn't feasible, the SUSE Linux system was migrated to Oracle Linux, using the processes described in this document.

Migration Process

Given the scale of the migration, planning and automation were essential to the project's success. Aik Zu Shyong, in Core Engineering at Dell, reflects: "We put significant focus on engineering the operating system conversion to make sure we could deliver a simple, reliable, and repeatable automated process. Additionally, by designing the migration to be done in-place instead of using a much slower and cost-prohibitive replacement method, we were able to further reduce downtime and save data center space."

Dell's migration process included three main steps: preparation, reimaging the operating system, and post-installation configuration. First, in the preparation step, Dell saved the existing environment's configuration and safely shut down the applications and database. Next, they reimaged the operating system from SUSE Linux 10 to Oracle Linux 5.5. After the reimaging completed, the post-installation steps configured the new environment and restored the previous data.

Preparation

The following pre-installation steps were used by Dell to prepare for their migration from SUSE Linux to Oracle Linux.

1. First, Dell created a scratch area to save the various configuration files.

For compatibility with the Oracle Linux operating system, Dell created an ext3 file system — not a ReiserFS file system, the default for SUSE Linux 10 — and documented the location of this scratch file system for use after the migration. Dell used a spare volume on the attached Fibre Channel storage device or a secondary drive on the machine, depending on the system configuration, to store the files.

The size of the scratch area varied with the specific system configuration, and it was based on the largest piece of data that needed to be backed up: the `ORACLE_HOME` directory. Sufficient space was reserved to hold this directory plus the various system configuration files that needed to be backed up.

2. Next, administrators shut down applications and services on the system and disabled the `init.d` processes.

Dell followed the Oracle-recommended shutdown order to stop the Oracle Database, Oracle Automatic Storage Management, applications running on the cluster nodes, and Cluster Ready Services (CRS). The `chkconfig` command was used to disable running services:

```
# chkconfig service_name off
```

For Dell's migration, the class of service of the system being migrated affected the shutdown procedure. For non-critical systems, a system maintenance window was requested and the entire cluster was shut down, migrated to Oracle Linux, and then restarted. For systems running business-critical applications, a complete shutdown of services was avoided. In these instances, rolling upgrades were employed. Services were transitioned off a selected cluster node to another node in the cluster, and that selected node was migrated to Oracle Linux. Then, that cluster node was restarted and rejoined the cluster. This process was repeated until all nodes in the cluster had been upgraded.¹

3. Dell confirmed file systems were not in use.

Dell used the `lsof` command to list any open files and make sure any NFS mounts were not in use. They also confirmed the `ORACLE_HOME` directory was free from resource utilization.

```
# lsof
```

4. Dell archived the relevant operating system configuration files and directories.

Using Table 1 as a reference, Dell archived the system configuration and collected the list of operating system files that needed to be retained to restore the site-specific configuration after the Oracle Linux installation. Dell first identified their site-specific configuration files and then created a script that could be used to copy these files to the scratch location created in Step 1.

¹ Although Oracle does not support heterogeneous Oracle RAC clusters, Dell experienced no issues during the transition with nodes running SUSE Linux interoperating with nodes running Oracle Linux. This mixed OS configuration was used only during the migration process, however, and not during normal system operation.

Note: This table is meant as a reference, not a definitive guide to the exact files and file locations. Information may vary based on your site-specific configuration.

TABLE 1. OPERATING SYSTEM CONFIGURATION INFORMATION

ARCHIVE STEP	COMMENTS
Hardware info	Archive hardware info using Dell OpenManage Server Administrator (OMSA) or native Linux commands; save to file (for example, hardware.txt)
Network card info	Archive IP address, subnet and gateway info, MAC address, link speed/duplex info, and network bonding configuration; save to file (for example, network.txt)
Memory info	(Optional) Archive the memory utilization records; use a maximum one-week snapshot, if necessary, to prove equivalent or better performance; save to file (for example, memory.txt)
OS *-release file	/etc/SuSE-release
Kernel modules info	/lib/modules/*, /etc/{modprobe.conf, modprobe.conf.local, modprobe.d/*}
Authentication (PAM), users and groups, and nsswitch.conf	/etc/pam.d/*, /etc/nsswitch.conf, /etc/passwd, /etc/shadow, /etc/group, /etc/sudoers, /etc/security/*
Device manager (udev) rules	/etc/udev/udev.conf, /etc/udev/rules.d/*
Automount file system info	/etc/auto.* (Optional; only needed if you are using automount)
Bootloader configuration files	/boot/grub/*, /etc/grub.conf, /etc/sysconfig/bootloader
/var/log/messages file	/var/log/{boot.msg, boot.ormsg, localmessages, messages}
Runlevel config	/etc/inittab, /etc/init.d/boot.local
rclocal script	/etc/rc.d/rclocal
Cron job config	/etc/cron/{daily, hourly, monthly}/*, /var/spool/cron/tabs/*
MPIO config	/etc/multipath.conf
Network config (NIC, routing, and so on)	/etc/sysconfig/network/ifcfg-*, /etc/sysconfig/network/*, /etc/resolv.conf
NTP config	/etc/ntp.conf
NFS config	/etc/exports, /etc/fstab
Name service config	/etc/nscd.conf
Hosts config	/etc/{hosts, host.conf, hosts.allow, hosts.deny, HOSTNAME}
System configuration (sysconfig)	/etc/sysconfig/* (including all subdirectories), /etc/sys/* (including all subdirectories)
/proc/info files	/proc/* (including all subdirectories)
SSH config	/etc/ssh/*, /etc/sshd.config, /etc/pam.d/ssh
SAR data files	/var/log/sa/sa*
Apache config files	Optional; needed only if running Apache /etc/httpd*
FTP	Optional; needed only if running FTP services
CIFS	Optional; needed only if running CIFS services

Shell/profile information	/etc/{bash.bashrc, csh.cshrc, csh.login, ksh.kshrc}, /etc/profile, /etc/profile.d/*
Pre-login message	/etc/issue
/etc/default directory files	/etc/default/*
PowerPath licensing and config files	/etc/emcp*
Additional software/applications	Back up any third-party non-Oracle software applications

5. Dell converted MPIO to PowerPath.

Dell chose to convert from SUSE Linux's built-in MPIO support to EMC PowerPath for automated data path management because this was the Dell standard for other non-SUSE Linux systems. Using EMC PowerPath also made it easier to copy over LUN mappings after the conversion.

A custom script was written by EMC to perform the conversion from MPIO to PowerPath. Details of this conversion step are beyond the scope of this paper. Readers are referred to EMC or their storage provider for more information on converting data path management, if needed.

6. Dell archived the Oracle-specific configuration information.

Similar to the operating system configuration files, Dell stored these Oracle-specific configuration files on a spare volume on the attached Fibre Channel storage device or on a secondary drive on the machine. Table 2 lists the Oracle-specific configuration files that Dell saved in preparation for the migration to Oracle Linux.

TABLE 2. ORACLE-SPECIFIC CONFIGURATION

ARCHIVE STEP	COMMENTS
Profiles for oracle and svcgrid users	.profile files for oracle and svcgrid users (Note: Oracle Linux files are named .bash_profile)
LUN mapping information	/u02; this directory contained the symbolic links for the LUN mappings in the SUSE Linux environment
Oracle Inventory Pointer (oraInst.loc) and oratab files	/etc/oraInst.loc, /etc/oratab
Oracle inventory file (oraInventory)	/etc/oracle/oraInventory
OCR file	/etc/oracle/ocr.loc
ssh trusted key for oracle user	~oracle/.ssh/*
Database-specific kernel settings	/etc/sysctl.conf
Home directory of Oracle user	Site-specific; /home/oracle for Dell configurations
Home directory of Oracle software (ORACLE_BASE)	Site-specific; /u01/app/oracle for Dell configurations

7. Dell created a backup image of the saved configuration files using the tar utility.

Reimage the Operating System

After configuration information was saved and all essential services were moved to backup servers, the system was ready to have the new Oracle Linux operating system installed. The kickstart installation method was used to automatically perform the installation of the Oracle Linux 5.5 operating system across the network. Using kickstart helped ensure quick, efficient, and consistent operating system installations on the client systems.

The standard kickstart configuration and installation were employed, with a central kickstart server on the network used for the installations. ISO images for Oracle Linux 5.5 were copied to Dell's regional imaging server and made available over the network. A kickstart configuration file was created that specified kickstart options and the packages to be installed. The client machine was booted using a USB flash drive, and the kickstart configuration file was downloaded. Installation proceeded automatically and was completed without requiring user intervention.

Caution: Make sure that the installation process does not erase the backup disk that is used to store the archived system information. Dell's kickstart process specifically touched only the /dev/sda disk, leaving /dev/sdb available for safely archiving the backup information.

Post-installation

The following key steps were part of the post-installation process in Dell's migration to Oracle Linux.

1. Dell restored/converted the operating system configuration files from SUSE Linux to Oracle Linux.

Dell restored the operating system configuration information that was saved (see Table 1 in the previous section) to enable the transition from SUSE Linux to Oracle Linux. The majority of the configuration files could be restored directly from the backup copy and did not require any conversion. The settings that were not directly restored include the following:

- **Ioscheduler information.** Because the `grub.conf` file is different for Oracle Linux, the equivalent SUSE Linux file could not be copied directly. Instead, an entry for the preferred ioscheduler was added to the new Oracle Linux `/boot/grub/grub.conf` configuration file, for example:

```
kernel KERNEL_PARAMETERS elevator=deadline
```

- **Password information.** Because the SUSE Linux environment used Blowfish and the new Oracle Linux environment used MD5 cryptographic hash functions, the encrypted password information in the `/etc/passwd` and `/etc/shadow` files could not be copied directly. Instead, the passwords for the few local user accounts (for example, the `oracle` user) were manually restored.
- **Host keys for ssh.** After installing the new Oracle Linux operating system, the host keys returned by the `ssh` daemon changed. Therefore, new `known_hosts` key files, needed for `ssh` client access, were regenerated for the hosts in the cluster.
Note: While Dell chose to generate new client keys, it would also be possible to restore the old host keys from backup.

- **Network bonding configuration.** SUSE Linux directly loads the bonding kernel modules through `ifcfg-bondN` configuration files. In contrast, Oracle Linux uses the `/etc/modprobe.conf` file to load the bonding kernel module and its options. Therefore, entries were added to the new Oracle Linux `/etc/modprobe.conf` file to load the bonding kernel modules and set options, for example:

```
alias bond0 bonding
options bond0 mode=active-backup miimon=100 downdelay=100
updelay=200
```

Both SUSE Linux and Oracle Linux store the network device information in `ifcfg-bondN` and `ethN` files. However, SUSE Linux stores these files in the `/etc/sysconfig/network` directory, and Oracle Linux uses the `/etc/sysconfig/network-scripts` directory. Table 3 shows an example `ifcfg-bondN` file for both the previous SUSE Linux environment and the new Oracle Linux environment.

TABLE 3. EXAMPLE IFCFG-BONDN FILES

SUSE LINUX	ORACLE LINUX
<code>/etc/sysconfig/network/ifcfg-bond1</code>	<code>/etc/sysconfig/network-scripts/ifcfg-bond0</code>

DEVICE=bond1
 BOOTPROTO='static'
 BROADCAST='192.168.255.255'
 IPADDR='192.168.0.190'
 NETMASK='255.255.0.0'
 NETWORK='192.168.0.0'
 REMOTE_IPADDR=''
 MTU=''
 STARTMODE='onboot'
 BONDING_MASTER='yes'
 BONDING_SLAVE_0='eth2'
 BONDING_SLAVE_1='eth3'
 BONDING_MODULE_OPTS='mode=active-backup miimon=100 downdelay=100 updelay=200'

DEVICE=bond0
 BOOTPROTO=none
 ONBOOT=yes
 IPADDR='192.168.0.190'
 NETMASK=255.255.0.0
 NETWORK=192.168.0.0
 USERCTL=no

Table 4 shows an example `ifcfg-eth2` file for both the previous SUSE Linux environment and the new Oracle Linux environment.

TABLE 4. EXAMPLE IFCFG-ETH2 FILES

SUSE LINUX	ORACLE LINUX
<code>/etc/sysconfig/network/ifcfg-eth2</code>	<code>/etc/sysconfig/network-scripts/ifcfg-eth2</code>

DEVICE=eth2
 STARTMODE='onboot'
 BOOTPROTO='none'
 MASTER='bond1'
 SLAVE='yes'

DEVICE=eth2
 HWADDR=00:15:17:97:CD:4E
 BOOTPROTO=none
 ONBOOT=yes
 MASTER=bond0
 SLAVE=yes
 USERCTL=no

Refer to the Oracle Linux system administration documentation for more complete details on setting up network bonding on an Oracle Linux system.

- Dell restored the Oracle configuration settings and files.

Dell restored the Oracle-specific configuration information that was saved (see Table 2 in the previous section) to enable the transition from SUSE Linux to Oracle Linux. Like the operating system configuration files, the majority of the Oracle-specific configuration files were able to be restored directly from the backup copy and did not require any conversion. The two exceptions were the profile files and the Oracle startup scripts in the `inittab` file.

- Profile files.** SUSE Linux uses a `.profile` file, and Oracle Linux uses a `.bash_profile` file. Therefore, the `.profile` files for the `oracle` and `svcgid` users were copied to `.bash_profile` files in the new Oracle Linux environment.
- The `inittab` file.** The `inittab` file is different for the two operating systems. Therefore, entries for the three startup scripts for Oracle software were copied into the new `inittab` file rather than directly copying the `inittab` file in its entirety.

The three relevant lines in the original SUSE Linux `inittab` file — entries for the Event Manager daemon (`evmd`), Oracle Cluster Services Synchronization daemon (`cssd`), and Oracle Cluster Ready Services daemon (`crsd`) — were copied from the archived file and added to the end of new Oracle Linux `inittab` file, for example:

```
# Run xdm in runlevel 5
x:5:respawn:/etc/X11/prefdm -nodaemon
h1:35:respawn:/etc/init.d/init.evmd run >/dev/null 2>&1 </dev/null
h2:35:respawn:/etc/init.d/init.cssd fatal >/dev/null 2>&1 </dev/null
h3:35:respawn:/etc/init.d/init.crsd run >/dev/null 2>&1 </dev/null
```

- Dell rebooted the server.
- Dell restarted the database and confirmed operation. In addition, third-party software products were also verified.

Note: As a best practice, the Oracle product executables might need be relinked after installing the new operating system. For more information, refer to [How to Relink Oracle Software on Unix \[ID 131321.1\]](#) on My Oracle Support (requires a valid Customer Support Identifier (CSI) to view).

Final Thoughts

Migrating 1,700 servers from SUSE Linux to Oracle Linux was an aggressive IT decision, one deemed necessary by Dell to gain better stability and support, easier administration, and lower costs. Extracting the underlying operating system layer and replacing it, while leaving the application layer intact, was possible only because of standardization across the Linux platforms. The bulk of the site-specific operating system configuration could simply be backed up and restored directly on the new operating system. Similarly, Oracle Database and other applications required only minor configuration changes to transition from SUSE Linux to Oracle Linux.

At the time this document was written in December of 2011, Dell was approximately halfway through the migration process, with an anticipated June 2012 completion date. Careful planning before starting the migration to itemize the needed site-specific configuration files and identify the files that required conversion was key to Dell's success. Automation via scripts and kickstart installations, plus attention to detail via checklists during the actual conversion process, reduced risk and provided consistency during the migration process.

Dell is confident that the decision to migrate these servers from SUSE Linux to Oracle Linux was the right decision for their business. According to Jon Senger, Enterprise Architect at Dell, "We took a risk doing a migration of such a large scale with such a challenging infrastructure, but it's really paid off. Not only do we lower our TCO for the environment, but since we have been able to standardize on Oracle Linux, we have achieved the stability and support that our customers demand."

Resources

The following resources are available for Oracle Linux:

- Oracle Linux Web site: <http://oracle.com/linux>
- Oracle Unbreakable Linux Network: <https://linux.oracle.com>
- Oracle Linux blog: <http://blogs.oracle.com/linux>
- Oracle Linux discussion forum: <https://forums.oracle.com/forums/forum.jspa?forumID=822>
- Oracle Linux on Oracle Technology Network: <http://www.oracle.com/technetwork/server-storage/linux/whatsnew/index.html>

Download Oracle Linux:

- Oracle Software Delivery Cloud (free download): <http://edelivery.oracle.com/linux>
- Public yum server for Oracle Linux: <http://public-yum.oracle.com>

Follow Oracle Linux:

- Twitter: [@ORCL_Linux](https://twitter.com/ORCL_Linux)
- Facebook: [http://facebook.com/OracleLinux](https://facebook.com/OracleLinux)
- YouTube: [http://YouTube.com/OracleLinuxChannel](https://YouTube.com/OracleLinuxChannel)

How Dell Migrated from SUSE Linux
to Oracle Linux
January 2012, Revision 1.0
Author: Jon Senger
Contributing Authors: Aik Zu Shyong,
Suzanne Zorn

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200
oracle.com

Copyright © 2012, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark licensed through X/Open Company, Ltd. 0611