Enterprise Information System Support Manual

OES Linux Server Templates

Last Revised: February 11, 2010

Revision 1.07

Managing Enterprise Services
Server Administration Templates for Novell OES Linux Server Provisioning

OES Linux Server Checklist

Checklist Item	Completed	Documentation Reference
BIOS/firmware updates and configuration		Consult SA staff
Network address and DNS configuration		Consult SA staff
Disk partitions configured		pages 1 - 2
OES and SLES Software Selections		pages 1 - 4
Install procedures for virtualized servers	Complete or NA[1]	pages 5 - 6
Install procedures for servers that will **not** use eDirectory, NCP, and NSS	Complete or NA[1]	page 6
Install procedures for servers that use eDirectory, NCP, and NSS	Complete or NA[1]	pages 6 - 8
Root user password configuration		page 9
Host name configuration for OES	Complete or NA[1]	page 26
Certificate Authority configuration		page 10
Configure VNC access		page 20
Disable IP version 6		page 16
LUM configuration	Complete or NA[1]	page 26
LDAP service configuration		pages 35 - 37
GRUB configuration		pages 11 - 13
Default Init/Run level configuration		page 10
Disable unnecessary services		pages 15 - 16
Key sequence configuration		pages 10 - 11
NTP configuration		page 10
Disable unnecessary apache modules		page 17
Additional apache hardening configuration		pages 17 - 18
SLP configuration		pages 18 - 19, pages 32 - 33
SMS configuration		pages 18 – 19, pages 30 - 31
LUM post installation configuration	Complete or NA[1]	Pages 27 - 28
SA server access configuration	Complete or NA[1]	page 29
SA home directory configuration	Complete or NA[1]	pages 29 - 30
Proxy and/or service user configuration	Complete or NA[1]	pages 7 – 8, Appendix
Check for and run nssid.sh file	Complete or NA[1]	page 31
LDAP eDirectory configuration	Complete or NA[1]	pages 34 - 35
Supportconfig application installation		page 13
NDS repair for Unix menu wrapper install	Complete or NA[1]	pages 33 - 34
SSH configuration		page 19
Novell Customer Center Registration		page 9
Apply all current online patches		pages 22 - 25

VMware Tools install and configuration	Complete or NA[1]	Consult ESX administrators
Virtual server network card configuration	Complete or NA[1]	page 14
Configure automated fsck intervals		pages 20 - 22
Clustered volume file system configuration	Complete or NA[1]	pages 38 - 39
NSS volume configuration	Complete or NA[1]	pages 31 - 32
Configure agents load and unload scripts	Complete or NA[1]	pages 39 - 40
Verify VMI kernel functionality	Complete or NA[1]	page 6

† Requirement to complete is dependant on software and/or environment chosen for installation

Managing Enterprise Services

Server Administration Templates for Novell OES Linux Server Provisioning

Table of Contents

Managing Enterprise Services

Server Administration Templates for Novell OES Linux Server Provisioning

General server installation models

If eDirectory is your primary, or preferred, user store then a server deployment using a basic OES footprint may be a better deployment option even if the desire is for a pure SLES host.

For example, consider the following installation pattern for an OES 2 Linux server:

- Server Base System
- Novell AppArmor
- Documentation
- Novell Backup/Storage Management Services
- Novell Linux User Management (LUM)
- Novell Remote Manager (NRM)
- Gnome desktop Environment for Server
- X Window System
- C/C++ Compiler and Tools

Effectively the resulting server OS footprint and functionality results in a near native SLES 10 box. However it offers enhanced access and management features.

A summary of benefits resulting from using this installation pattern:

- Using eDirectory accounts and features to authenticate and authorize users
- Using SMS compliant backup services
- Using familiar web based management tools
- Unified patching facilities
- The use of eDirectory certificate services for secure NCP/LDAP authentication services
- Basic NCP connectivity through the use of the Dclient
- A local easily configured firewall for OES Linux services
- Installing the OES service framework now simplifies adding additional OES services later

If you are not using eDirectory as your user store the pure SLES server deployment is obviously a more practical choice. Potential best practice server models are detailed in the following section.

Partition and file system configurations

The referenced disk, or logical drive for RAID configurations, capacities will vary from hardware to hardware instance in some cases. If RAID devices are used then the logical drive hosting the root and supporting partitions, excluding application or end user data partitions, should be initialized in a RAID 1 configuration. The partition size referenced for the root partition should be considered a guideline, but the other partition choices and their placement on the disk should be adhered to unless there is a functional necessity to vary from it. The EXT3 file system is preferred for OES Linux server OS partitions because of its performance and journaling capabilities.

Recommendations for physical servers:

Partition Number	Partition	Partition Size	File System Type
Partition 1	Swap	1 GB	Swap
Partition 2	/	12 GB minimum (36 GB Recommended)	EXT3
Partition 3	/srv	8 GB minimum	EXT3
Partition 4	/var	8 GB minimum	EXT3

Recommendations for virtual servers:

Partition Number	Partition	Partition Size	File System Type
Partition 1	Swap	1 GB	Swap
Partition 2 [1]	/boot	350 MB	EXT2
Partition 3	/	16 GB minimum (20 GB Recommended)	EXT3
Partition 4	/var	8 GB minimum	EXT3

† Due to differences in virtualized guest boot mechanics a non-journalized file system is used

Clustered server configurations

Software Installation (Always use current and supported installation media)

For clustered servers the minimal software that should be installed is:
- Server Base System
- Novell AppArmor
- Documentation
- Novell Backup/Storage Management Services
- Novell Cluster Services (NCS)
- Novell eDirectory
- Novell Linux User Management (LUM)
- Novell NCP Server/Dynamic Storage Technology
- Novell Remote Manager (NRM)
- *Novell Storage Services (NSS)*[1] [2] [3]
- Gnome desktop Environment for Server
- X Window System
- C/C++ Compiler and Tools
- Kernel Source

† Although it is an requirement for OES NetWare systems to use eDirectory, NCP Server, and NSS in NCS implementations these are not requirements for OES Linux NCS implementations. However it is the EISS standard to use them so the eDirectory and NCP Server options should always be included but the NSS option can be excluded if not needed

†† For most file systems verify the primary name space for NSS volumes is set to LONG

††† NSS file systems used for GroupWise stores may perform better using the UNIX name space

** *Verify the fiber HBA driver and firmware versions required by your SAN vendor. This may require and update or even a down grade of the drivers provided with the operation system for your hardware.*

Novell Cluster Services and not using NSS or eDirectory

It is possible for an OES Linux server that does not utilize NSS to participate in a NCS cluster. If the cluster node will not be hosting services that require NSS it can be omitted from the application installation list. eDirectory is a dependency of the NSS service so it could be omitted in this scenario as well. Many Novell services on OES 2 Linux servers do not require a local instance of eDirectory. OES Linux servers without eDirectory installed will use a local service, referred to as the Dclient, to instantiate and manage NCP connections to the servers in the target eDirectory tree(s) as needed.

It is important to note that server LDAP and SLP targets still need to be properly configured when using this directory connectivity model.

Organization and management of your LDAP infrastructure will be of significant benefit when utilizing OES Linux servers without local eDirectory instances.

Non-clustered server configurations

Software Installation (Always use current and supported installation media)

For non-clustered servers where the need for NCP services and NSS file systems has been identified the minimal software that should be installed is:

- Server Base System
- Novell AppArmor
- Documentation
- Novell Backup/Storage Management Services
- Novell eDirectory
- Novell Linux User Management (LUM)
- Novell NCP Server/Dynamic Storage Technology
- Novell Remote Manager (NRM)
- Novell Storage Services (NSS)[1] [2] [3]
- Gnome desktop Environment for Server
- X Window System
- C/C++ Compiler and Tools

† NSS file systems should be placed on a non-system disk device that is not used for the OS
†† For most file systems verify the primary name space for NSS volumes is set to LONG
††† NSS file systems used for GroupWise stores may perform better using the UNIX name space

For non-clustered servers where the need for NCP services is identified the minimal software that should be installed is:

- Server Base System
- Novell AppArmor
- Documentation
- Novell Backup/Storage Management Services
- Novell Linux User Management (LUM)
- Novell NCP Server/Dynamic Storage Technology
- Novell Remote Manager (NRM)
- Gnome desktop Environment for Server
- X Window System
- C/C++ Compiler and Tools

** Organization and management of your LDAP infrastructure will be of significant benefit when utilizing OES Linux servers without local eDirectory instances.

For non-clustered servers where the need for SLES services is identified the minimal software that should be installed is:

- Server Base System
- Novell AppArmor
- Documentation
- Novell Backup/Storage Management Services
- Novell Linux User Management (LUM)
- Novell Remote Manager (NRM)
- Gnome desktop Environment for Server
- X Window System
- C/C++ Compiler and Tools

Virtual server configurations

Software Installation (Always use current and supported installation media)

In addition to any the previous software configurations the following package should be installed for any virtualized 32 bit OES 2 Linux and Suse Linux Enterprise 10 virtualized servers.

- kernel-vmi
- VMwareTools-XXXX-XXXXXX[1]

32 bit SLES based guests will function in a fully virtualized mode using the standard kernel configuration. However the VMware Virtual Management Interface kernel (VMI) will allow for hypervisor specific optimizations to be used resulting in better performance of the guest OS.

The installation and configuration of the VMware Tools software is required for all EISS managed OES Linux and Suse Linux Enterprise Servers hosted on the ESX platform. If the VMI kernel is installed and active for the ESX guest the marriage of the specialized kernel and tool software is also required to realize the comprehensive benefits of the VMI implementation.

It is strongly recommended that all EISS managed ESX hosted Linux servers have the VMwareTools software installed and configured.

† Consult with your ESX Server administrator on how to obtain and install the correct VMwareTools software package for the target server.

Server installation procedures and standards overview

Procedures (Pre-installation):

For all virtualized servers:

- Obtain IP addresses and DNS records for each interface/service the server will host
- Disable serial and parallel ports in the virtual BIOS settings. This will prevent the generation of unnecessary interrupts for the virtual server.

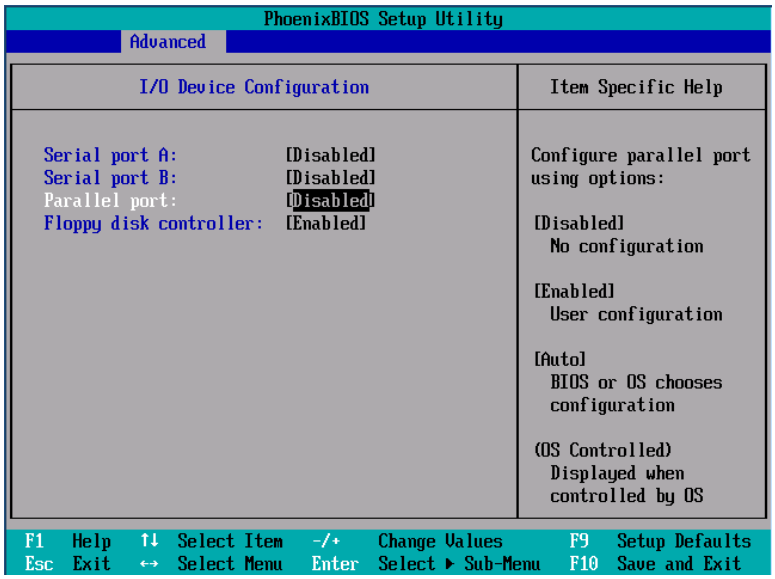

- If the VMI enabled kernel will be installed for the virtual server, enable the VMI para-virtualization features for the server. This is setting is accessible the “Options” section dialog in the Virtual Center management utility.

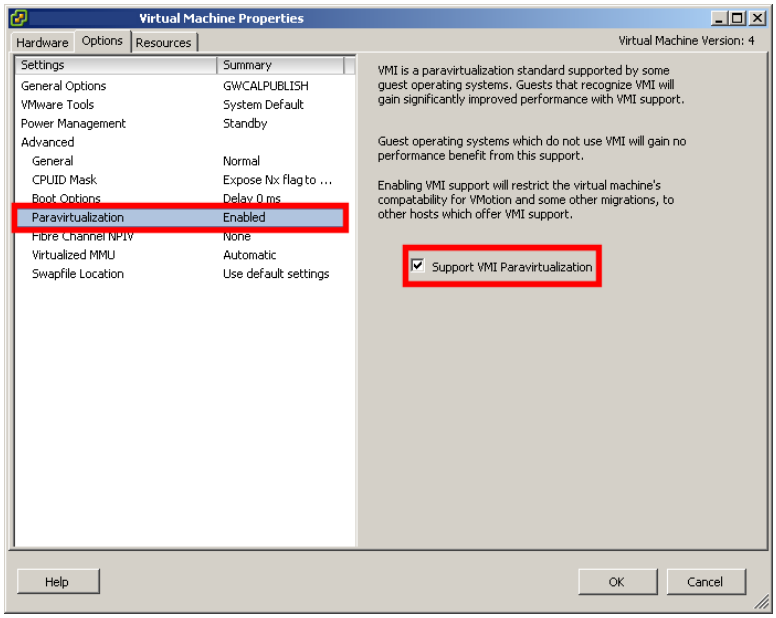

Procedures (Post Installation):

- Verify that the VMI kernel has been implemented successfully on the virtual server.

  Issue the following command as root at the server console:
  dmesg | grep VMI

  If the VMI kernel was implemented properly and is communicating with the ESX hypervisor a solllar message to the one below will be displayed.

  Detected VMI ROM version 3.0
  VMI Timer active.
  VMI Timer cycles/sec = 2133407000 ; cycles/jiffy = 8533628 ;cycles/alarm = 21334070
- Configure your VMwareTools for the virtual server:

  Issue the following command as root at the server console:
  /usr/bin/vmware-config-tools.pl

  Accept all default configuration values presented by the configuration script unless you have a reason to change them. Verify the configuration process completes successfully and that the ESX host recognizes the VMware tools package is up to date and running.

- It is considered best practice to document configuration notes, installed applications, and service commands specific to an ESX Guest. The Virtual Infrastructure Client used in the ESX environment has a “Notes” field within the “Summary” tab for each ESX guest which should be used to accomplish this.

<u>For servers that will not have eDirectory, NCP, or NSS services installed:</u>

Procedures (Pre-installation):

- Obtain IP addresses and DNS records for each interface/service the server will host
- If using physical hardware, update all server BIOS and firmware code to the latest supported versions and disable serial and parallel ports in the BIOS settings.
- Ensure the desired installation media for selected configuration(s) is available in the enterprise software repository.
- Ensure the enterprise software repository is accessible to the target server over the network for installation. Do not install the server from local media.

<u>For servers that will have eDirectory, NCP, or NSS services installed:</u>

Procedures (Pre-installation):

- Obtain IP addresses and DNS records for each interface/service the server will host
- If using physical hardware, update all server BIOS and firmware code to the latest supported versions and disable serial and parallel ports in the BIOS settings.
- Perform eDirectory Service Health Checks on all replicas of the eDirectory partition that will host the OES Linux server NCP or NCP and NSS objects.
- Perform eDirectory Service Health Checks on all replicas of the eDirectory partition [ROOT].
- Ensure the desired installation media for selected configuration(s) is available in the enterprise software repository.
- Ensure the enterprise software repository is accessible to the target server over the network for installation. Do not install the server from local media.

Procedures (During Installation):

- Ensure NTP is configured properly to allow the new server to successfully join and participate in the target eDirectory tree
- If creating NCP objects in eDirectory use uppercase characters for the Linux server short hostname

Procedures (Post Installation):

- Perform eDirectory Service Health Checks on all replicas of the eDirectory partition that host the OES Linux server NCP or NCP and NSS objects.
- Perform eDirectory Service Health Checks on all replicas of the eDirectory partition [ROOT].
- Populate required OES Linux service user and group class object attributes with standard EISS values.

OES Linux service user and group object and standards:

Minimum user object attributes and value pairs:

User class Attribute	Attribute Value
sn	User
Given Name	Descriptive name representing the target service for the user
Full name	Given Name + Full Name
SSN	Generic
Security Code	Generic
Company	Darkvixen Enterprises
L	Enterprise Information System Support
description	Do not delete, expire, or disable logins

a. Be sure to apply any password restrictions required by EISS standards for generic users
b. If possible apply any network address restrictions required by EISS standards for generic users.
c. Comply with naming and placement standards for OES Linux service user, group, and proxy user objects in place.
d. Configure OES Linux services to use existing configured proxy users where applicable. (See Appendix for more information on key OES system proxy user configuration details)

Minimum group object attributes and value pairs:

User class Attribute	Attribute Value
description	- Descriptive information about the target service - Group Security Authority name and contact information

Supplemental information:

Some system created OES Linux service proxy users allow the installer to name them and configure their placement in the target eDirectory tree. Naming and placement standards are in place for many of these objects and those standards should be applied to new objects if they are created during the OES Linux server installation.

Typical OES Linux service user class objects are:

User class object	Description	LUM enabled
novlwww	OES Linux Tomcat User	yes
novlxregd	OES Linux Xtier Registry User	yes
novlxsrvd	OES Linux Xtier Service User	yes
wwwrun	OES Linux Apache Service User	yes
nssAdmin	OES Linux NSS User	no

Typical OES Linux service group class objects are:

Group class object	Description	LUM enabled
novlxtier	OES Linux Xtier service group	yes
www	OES Linux apache group	yes

Proxy users are utilized for the following OES Linux services:

OES Linux Service	Proxy Scope	Placement or Naming Standards	System Created
Linux User Management	Global proxy user	Refer to standards table 1	no
Novell QuickFinder	Global proxy user	Refer to standards table 1	no
Novell Cluster Services	Global proxy user	Refer to standards table 1	no
Novel Samba	Server specific user	Refer to standards table 1	yes
Novel Samba	Server specific group	Refer to standards table 1	yes
Novell NetStorage	Server specific user	Refer to standards table 1	yes
Novell Storage Services	Server specific user	Refer to standards table 1	yes

Standards table 1:

OES Linux Service User	Naming Standard	Placement Standard
Linux User Management	LUM_PROXY	.PROXYUSR.LUM.SVS.DVC
Novell QuickFinder	QF_ADMIN	.DVC
Novell Cluster Services	NCSConfig	.PROXYUSR.<*CLUSTER*>.SVS.DVC
Novel Samba (Clustered)	<SERVER>_SambaProxy	.PROXYUSR.<*CLUSTER*>.SVS.DVC
Novel Samba	<SERVER>_SambaProxy	.PROXYUSR.<*SERVER_CONTEXT*>
Novell NetStorage (Clustered)	<SERVER>_NetStorage	.PROXYUSR.<*CLUSTER*>.SVS.DVC
Novell NetStorage	<SERVER>_NetStorage	.PROXYUSR.<*SERVER_CONTEXT*>
Novell Storage Services	<SERVER>_NSS_admin	<*SERVER_CONTEXT*>

Information for all SLES and OES Linux server configurations

Backing up configuration files

One of the most important best practices for managing Linux server configuration files is to back up the current file before performing any modifications. The easiest way to do this is to simply save the current file to the same location with a different name.

The current standard being used is to rename the original file, meaning the initial copy of a file that has never been modified, using an ".original" extension.

For example:
"**etc/slp.conf**" is renamed to "**etc/slp.conf.original**"

If the file has been renamed before, meaning a version of the file already exists with the .original extension, then the file is renamed with an extension that reflects the date the file was modified.

For example:
"**etc/slp.conf**" is renamed to "**etc/slp.conf.02142009**"

Root user password configuration

There is a current standard in place for local root user passwords for EISS managed OES Linux and Suse Linux Enterprise servers. The password used is the same across servers and is stored in the Serve Administration protected KeePass application.

It is important this password not be compromised so the transmission of the root user credentials over unsecure channels is not permitted by policy and systemically disabled where possible.

Register server with the Novell Customer Center

All OES Linux and Suse Linux Enterprise servers sould be registered with the Novell Customer Center (NCC) online database. Registering the server with the NCC ensures the server qualifies for online updates to facilitate server patching and even adds the required update servers to the server configuration. Valid registration/licensing codes for OES and SLES products are required.

Procedure:

Server registration with the NCC can be done from the server YaST utility but it is somewhat easier from the command line using the suse_register application.

To register a Suse Linux Enterprise server from the command line use the following syntax:
suse_register -a email=<YOUR_NCC_REGISTERED_EMAIL_ADDRESS> -a regcode-sles=<SLES_SERVER_REGISTRATION_CODE>

To register a OES Linux server from the command line use the following syntax:
suse_register -a email=<YOUR_NCC_REGISTERED_EMAIL_ADDRESS> -a regcode-sles=<SLES_SERVER_REGISTRATION_CODE> -a regcode-oes=<OES_SERVER_REGISTRATION_CODE>

It may take several minutes for the command to complete so do not close your terminal session until the command prompt returns or reports a failure.

Local Certificate Authority and certificate configuration

Even though OES Linux and Suse Linux Enterprise servers utilize eDirectory Certificate Authorities and certificates for HTTPS and LDAPS services they, and standard SLES instances, still utilize a local CA and certificates for Linux specific applications like SSH and Apache.

There is a current standard for local CA and certificate configurations on EISS managed Linux servers. The standards are comprised of naming and password conventions and the certificate expiration periods respectively.

The server NCP name and EISS organizational information are incorporated into the local CA configuration as detailed below.

Field	Field value
CA name	<SERVER_NAME>_CA
CA Common Name	CA (<SERVER_NAME>)
Organization	DarkVixen Enterprises
Organizational Unit	Enterprise Information System Support
Locality	Augusta
State	Georgia
Country	US
Email	postmaster@dvc.edu
CA password[1]:	Different from server root user and the standardized password for EISS managed servers is stored in the KeePass application

† It is important to consider setting the local CA password for your Linux servers to a standardized password other than the root user's password. If your organization regularly changes passwords you may run into management and access issues if the legacy, current, and future values for the root password differ.

The original local server certificate is revoked and a new one that is valid for two years is issued as a post installation task. This practice ensures the local certificates and the eDirectory certificates used for the server expire at the same time.

Init or "Run level" configuration

Currently EISS managed OES Linux servers are configured to start at init 5 or "run level" 5 which starts the server GUI interface and is the default for both OES Linux and Suse Linux Enterprise servers. This is primarily to accommodate remote VNC console access to the servers and to aid administrators in managing the server and its' hosted services. Due to the security issues associated with VNC access to Linux servers this configuration standard may change at a later date. The primary security concern with VNC access to Linux servers centers around administrators not securing sessions properly. Alternatively the primary operational concern would be to return the computational resources consumed by the server GUI application to the back to the server.

Key sequence configuration

A simple configuration change recommended for EISS managed OES Linux and Suse Linux Enterprise servers is to disable the <CTRL> <ALT> <DEL> key sequence. This may seem trivial but can prevent unintentional server reboots when OES servers are in an environment where operators who may be accustomed to the key sequences used in other operating systems have access to server consoles.

Procedures:

To disable the <CTRL> <ALT> <DEL> key sequence launch YaST as root and navigate to the “Security and Users” section and then the “Local Security” option.

Continue to click the “Next” button without changing any settings until the “Boot Settings” screen is reached. Set the <CTRL> <ALT> <DEL> key sequence setting to “ignore”.

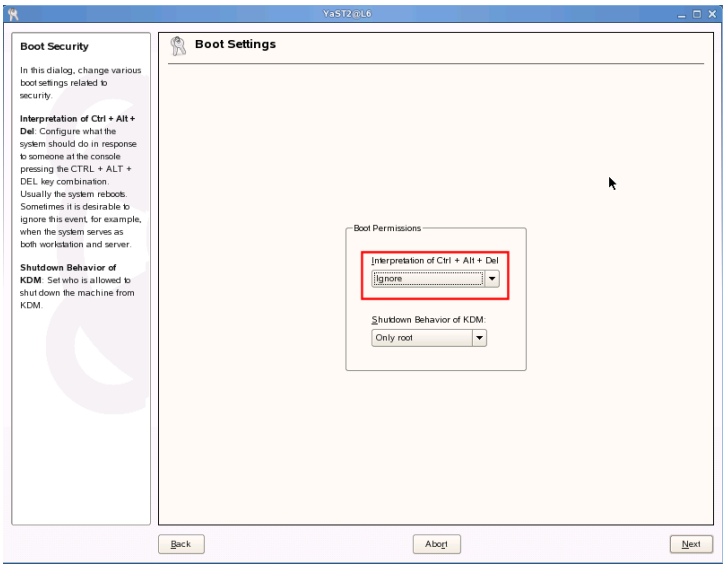

Continue to click the “Next” button without changing any settings until the “Finish” button is reached. Clicking “Finish” completes this procedure.

GRUB configuration for physical computers

GRUB is an acronym for “GRand Unified Boot Loader” and is one of several mature Linux boot applications. GRUB is the default boot loader OES Linux servers use to control their boot configuration and behavior. Like most Linux applications, GRUB is managed with text based configuration files. Recognized directives and parameters are passed to the application using the **/boot/grub/menu.lst** configuration file. Only a select few parameters are changed on EISS managed servers. These are displayed in bold below.

default 0
timeout 20
##YaST - generic_mbr
gfxmenu (hd0,1)/boot/message
##YaST - activate

###Don't change this comment - YaST2 identifier: Original name: linux###
title **OES 2 Linux Server – L9**
    root (hd0,1)
    kernel /boot/vmlinuz-2.6.16.46-0.12-bigsmp root=/dev/cciss/c0d0p2 vga=0x317 resume=/dev/cciss/c0d0p1 **splash=0** showopts
    initrd /boot/initrd-2.6.16.46-0.12-bigsmp

###Don't change this comment - YaST2 identifier: Original name: failsafe###
title **Failsafe**
    root (hd0,1)

kernel /boot/vmlinuz-2.6.16.46-0.12-bigsmp root=/dev/cciss/c0d0p2 vga=normal showopts ide=nodma apm=off acpi=off noresume nosmp noapic maxcpus=0 edd=off 3
initrd /boot/initrd-2.6.16.46-0.12-bigsmp

The boot menu time out is changed from 8 seconds to 20.
The default menu item string is changed to use the *Platform_Patch_Level - Server_Name* format.
The boot splash screen display setting is changed to 0.
The Failsafe menu item string is changed to the simple format displayed in the example

Procedure:.

The referenced changes can be made by modifying the **"/boot/grub/menu.lst**" with a Linux text editor, a Linux GUI text editor, or with the YaST Boot Loader module.

The YaST Boot Loader module procedure is detailed here.

Launch YaST as root and navigate to the "System" section and then the "Boot Loader" option.

Click the "Other" button and select 'Edit configuration files" option from the drop down menu.

Select the /boot/grub/menu.lst option from the drop down menu at the top of the dialog and make your changes to the file in the file editor field of the dialog box.

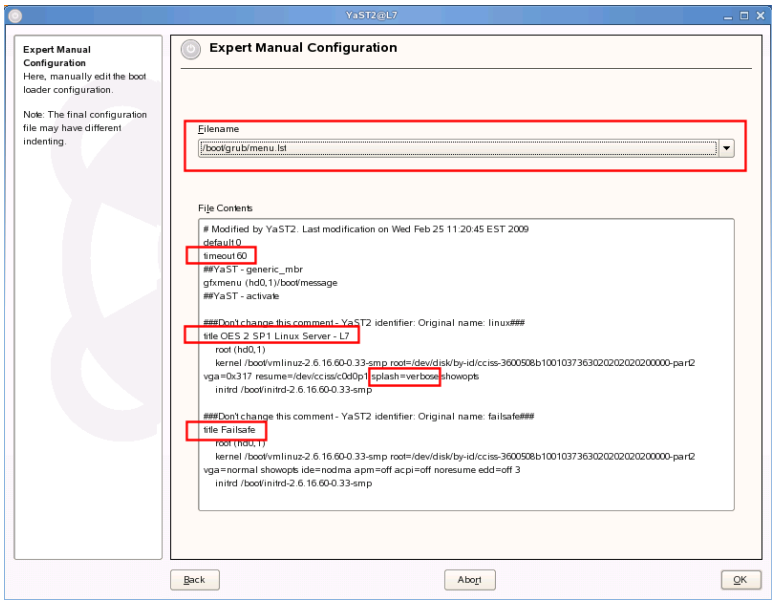

Click the OK button when you are finished to commit your changes.

GRUB configuration for virtual computers

The GRUB configuration for virtual computers is almost identical to that of physical computers with one or possibly two exceptions. The first exception only applies if the VMI kernel is installed.

Procedure:

Issue the following command at server terminal session to determine if the VMI kernel is in use:
/usr/bin/uname –r

If the kernel value returned has the format **<kernel_version_number>-vmi** than this exception applies.

Follow the procedure for GRUB configuration for physical servers but change the title field in the stanza that loads the VMI kernel to the following format and make it the default boot kernel.

*Platform_Patch_Level - Server_Name **(VMI Enabled)***

Leave the contents of the stanza containing the original kernel unchanged.

The second exception relates to the “showopts” directive and applies to all virtualized OES Linux and Suse Linux Enterprise Servers. This parameter configures kernel boot options that are visible on the GRUB boot menu screen. Virtual servers utilize virtual hardware which includes hardware key to accurate time keeping. The “pmtmr” kernel boot option is beneficial to use in virtualized environments.

Add the “clock=pmtmr” value to the showopts directive in the default boot kernel stanza to apply the recommended configuration as in the example below.

```
###Don't change this comment - YaST2 identifier: Original name: linux###
title OES 2 Linux Server – L9
   root (hd0,1)
   kernel /boot/vmlinuz-2.6.16.46-0.12-bigsmp root=/dev/cciss/c0d0p2 vga=0x317
resume=/dev/cciss/c0d0p1 splash=0 showopts clock=pmtmr
   initrd /boot/initrd-2.6.16.46-0.12-bigsmp
```

Installing Novell supportconfig application

Novell Support Services offers an application for download and installation that helps facilitate the transmission of OES Linux and Suse Linux Enterprise server configuration and log information to Novell Support engineers. If you open a Support Request involving an OES Linux or Suse Linux Enterprise server with Novell Support Services you will most likely be asked to provide the support engineer with the output tar archive produced by this application.

Procedure:

Download the most recent copy of the application from the URL listed below and copy it to the target server in a branch of the file system your user has access to.

http://www.novell.com/communities/node/2332/supportconfig-linux

Uninstall the legacy ntsutils RPM package using the following command:
rpm -e ntsutils

Install the supportutils RPM package using the following command:
rpm -Uvh /<RPM_FILE_PATH>/supportutils-<VERSION_NUMBER>.noarch.rpm

Running the utility is simply done by issuing the following command at the server as the root user:
/sbin/supportconfig

The utility writes its archive to the /var/log directory using a naming convention that starts with “nts_” and references the server host name, date, and time the archive was created. This file can be emailed to Novell Support staff or the utility can be instructed to upload the file to Novell Support sytems if the server has FTP service access to the public internet.

Use the following syntax to upload the archive created by the supportconfig application to the Novell Support Services FTP server:
/sbin/supportconfig –ur <SERVICE_REQUEST_NUMBER>

NTP service configuration

All EISS managed OES Linux servers should use the same NTP time source. The standard as of this writing is to configure all OES NetWare and OES Linux servers to consume time from the service available at the "**ntp.dvc.edu**" address. The ntp.dvc.edu service provides time services for servers participating in all EISS production and lab eDirectory environments.

Procedures:

When the server is installed configure the OES Linux or Suse Linux Enterprise Server NTP target as one of your first post installation tasks.

Use the "yast2 ntp-client" command to invoke the yast NTP configuration module. Specify the ntp.dvc.edu address for the time server and use the "Test" option to ensure the server is reachable and responds correctly.

Select the "Finish" option so commit you changes for the service changes.

Virtual server network card configuration

One of the benefits of a virtualized server is the ability to duplicate it quickly to provision a new instance of the guest OS and its' hosted services. VMware ESX services refer to this process as "cloning". Traditionally the process of cloning Linux servers has resulted in anomalies with the network card in the cloned host being unusable as "eth0". Suse OS facilities and YAST uniquely identify network interfaces using ethernet addresses. Having the eth0 identifier associated with the primary network interface is preferred by most system administrators.

Modifying the "FORCE_PERSISTENT_NAMES" directive in the **/etc/sysconfig/network/config** file will allow you to accomplish this in cloned guests. You will have to use YAST to re-setup your network interfaces following the change. However you will be able to setup your interface as eth0 instead of being forced to use eth1.

To implement this change boot the cloned OES Linux or Suse Linux Enterprise guest and follow the procedure detailed below.

Procedure:

- Backup the existing **/etc/sysconfig/network/config** file
- Edit the **/etc/sysconfig/network/config** file and
  change the "FORCE_PERSISTENT_NAMES=no" directive
  to
  FORCE_PERSISTENT_NAMES=yes"
- Reboot the cloned workstation and reconfigure your network card as eth0
  (You may have to delete and recreate the network card)

Disabling unnecessary services

Regardless of the role of the target server it should be deployed with the minimal running application environment possible. What this means in provisioning terms is that it should be deployed with the minimum software footprint required to support the hosted services and possibly any clustering software utilized. There are additional benefits to this deployment model. Primarily these include the conservation of computational resources, increased server security, and reduced management overhead. Most remarkable of these points is that it can result in a OS and/or application instance that is more easily hardened as a result of the omission of the unnecessary services.

The reduction of installed applications is accomplished during the installation of the server itself. However even with a minimal software footprint there are still some standard and default services active on OES Linux installations that can safely be disabled. Examples of these services include sound hardware daemons, modem hardware daemons, and mail handling daemons. A comprehensive reduction of unnecessary services and daemons running on an OES Linux server is often best accomplished post installation.

Below is table of services that can normally be disabled on EISS managed OES Linux or Suse Linux Enterprise servers:

Daemon or service script	Description
alsasound	Sound hardware daemon
apache2[1]	The Apache 2 http daemon
auditd	Linux file system auditing daemon
nfsserver	NFS server daemon
nfsboot	NFS service daemon
nmb	Samba netbios daemon
novell-smdrd	*Novell SMDR daemon and TSA*
novell-tomcat5[1]	Novell implementation of the java application server
novell-zmd[2]	Novell Zenworks Management daemon
smb	Samba service daemon
portmap	NFS/NIS service daemon
postfix	*Mail transport agent (Consider whether or not to disable)*
splash	Splash screen setup
splash_early	Stops animations after the network starts
suseRegister	A start script to execute suseRegister during boot

** *Please note some daemons and services may be disabled to allow the clustering software to control when they start, stop, and restart. Please refer to the relevant services configuration documentation for more information about daemons using this configuration.*

† Consider not loading the apache and tomcat services on most OES Linux cluster servers at all. It may be best to utilize a resource server to run web based management utilities from to manage your OES Linux environment. If you need to run apache or tomcat instances on clustered nodes for services consider configuring discreet instances of them instead of using the primary instance configurations on the server if possible.

†† Consider disabling the Novell ZenWorks management daemon on production OES Linux servers. This application can be resource intensive at times. Loading it manually also allows for a degree of enforcement for patch installations.

Procedure:

For each daemon or service to be disabled (the postfix daemon is the example case):

To check the status of the daemon or service startup configuration
chkconfig postfix

To check the status of the daemon or service init level startup configuration
chkconfig -l postfix

If the daemon or service is being disabled for the first time stop the service first.
/etc/init.d/postfix stop

To disable the daemon or service startup configuration
chkconfig postfix off

Disabling IP version 6

Similar to the reasons for disabling unnecessary services unnecessary service protocols should be disabled as well. IP version 6 has several well know vulnerabilities that can be addressed with patches and configuration changes. However if the protocol is not required in your environment it should be disabled.

IP version 6 is disabled on all EISS managed OES Linux and Suse Linux Enterprise servers.

Procedure:

IP version 6 can be disabled during the server installation or as a post installation task. This procedure details the post installation task steps.

Run Yast as the root user to access the server network configuration:

Issue the following command at the server console:
/sbin/yast network

Then edit the existing network card configuration(s) by selecting EDIT → Then ADVANCED → Then IPv6 and select the option to Disable IPv6.

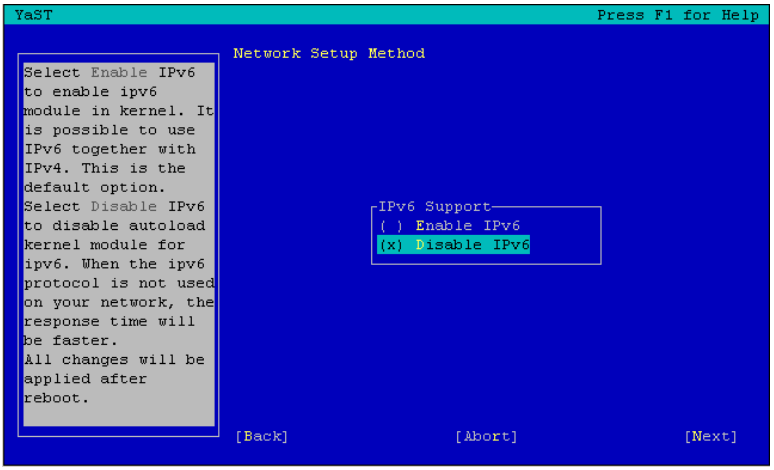

Select NEXT and then FINISH to complete the procedure.

Disabling unnecessary Apache modules

The same logic for disabling and unloading unnecessary server applications and services can be applied to apache modules as well. If the module is not running it does not need to be managed or secured.

With some degree of certainty there is a small subset of apache server modules that can and should be disabled on all EISS managed OES Linux and Suse Linux Enterprise servers. These primarily deal with file and database authentication and authorization services. Generally EISS implementations will employ secure LDAP facilities for apache authentication and authorization services.

Procedures:

Modify the “APACHE_MODULES” directive in the **/etc/sysconfig/apache2** file to remove unnecessary modules from the server configuration.

For example remove the references for the following apache modules:
auth_basic
authn_file
authz_groupfile
authz_user
authn_dbm
suexec
userdir

** *However the required module set for any specific server implementation can vary so do not remove any modules your implementation may require.*

Restart apache using the following command:
/etc/init.d/apache2 stop && /etc/init.d/apache2 start

Additional Apache hardening suggestions

There are many configuration changes that can be implemented to further harden apache instances on OES Linux and Suse Linux Enterprise servers. However the suggestions here are minimum baseline configurations that are determined to be a good fit for the network environment that hosts the servers and services.

A simple change could be used to prevent apache from serving certain types of files as HTTP content or for download. Some of these file types may contain server and service configuration information or other information the organization may not intend to publish. A file named “security.conf” containing the required directives could be created and placed in the **/etc/apache2/conf.d** directory of the target server. Apache on OES and Suse Linux distributions is pre-configured to process all files that end in a “.conf” extension in this directory on service startup. Apache should be restarted after placing the file in the referenced directory.

Restart apache using the following command:
/etc/init.d/apache2 stop && /etc/init.d/apache2 start

The contents of the “security.conf” file are displayed below:

```
<Files ~ "\.(bak|old|~|2|copy|tmp|swp)$">
Order allow,deny
Deny from all
</Files>
```

Some additional security steps can be implemented to instruct apache not to divulge unnecessary information just because a client asks for it. An example of this would be the vendor and version of the operating system hosting the apache instance.

Instructing Apache to not declare the operating system provider used on its host is quite easy to do. Modify the "**APACHE_SERVERTOKENS**" directive in the in **/etc/sysconfig/apache2** file to be something other than "**OS**". It is recommended that you set it to "**ProductOnly**" so apache only declares the apache product and no version or OS info if interrogated by a client.

After making the change restart apache using the following command:
/etc/init.d/apache2 stop && /etc/init.d/apache2 start

Service Location Protocol configuration

The Service Location Protocol (SLP) configuration and role for EISS managed OES Linux servers are usually configured during installation. However the SLP configuration and role for servers can be adjusted, and should be verified, post server installation as well. Generally the only SLP service configuration parameter changed for OES Linux servers are the Directory Agents (DA) specified for the service to use. In the EISS managed environment the primary and secondary DA servers user are slp1.dvc.edu (10.6.10.74) and slp2.dvc.edu (10.6.11.136) respectively. For SLP service configuration the SLP software is biased to prefer the IP address for service configuration.

The relevant section from the **/etc/slp.conf** SLP service configuration file containing the DA server parameters in bold is displayed below:

```
#--------------------------------------------------------------------------
# Static Scope and Static DA Configuration
#--------------------------------------------------------------------------

# Allows administrator to force UA and SA agents to use specific DAs.  If
# this setting is not used dynamic DA discovery will be used to determine
# which DAs to use.  (Default is to use dynamic DA discovery)
net.slp.DAAddresses = 10.6.10.74,10.6.11.136

#--------------------------------------------------------------------------
# UA Specific Configuration
#--------------------------------------------------------------------------
# A 32 bit integer giving the maximum number of results to accumulate and
# return for a synchronous request before the timeout, or the maximum number
# of results to return through a callback if the request results are
# reported asynchronously (default value is 256).
net.slp.maxResults = 768
```

Some of the most common commands an administrator is most likely to need and use to manage the SLP on OES Linux are displayed below:

To determine the scope(s) currently known by the SLP agents on an OES Linux server use the following command:
/usr/bin/slptool findscopes

To determine the providers of a particular type of service currently known by the SLP agents on an OES Linux server use the following command:
/usr/bin/slptool findsrvs service:<SERVICE_TYPE>

To determine the DAs currently known by the SLP agents on an OES Linux server use the following command:

/usr/bin/slptool findsrvs service:directory-agent

To determine the SAs currently known by the SLP agents on an OES Linux server use the following command:
/usr/bin/slptool findsrvs service:service-agent

SSH service configuration

The SSH services running on POSIX systems offer a reasonable level of security in their default configurations. However, most organizations should implement some basic service configuration changes that will enhance the security and usability of their SSH services even if only slightly.

Often institutions and enterprise IT departments limit the ability of the user root to logon to POSIX based server operating systems remotely using SSH. Specifically authentication to one of these systems as user root can be accomplished in only two ways. One is to do so from the physical server console and the other is by using the "su" command when authenticated as another valid user. Implementing this feature is a simple but effective component in hardening POSIX systems and can easily be implemented on OES Linux servers.

A second change limits the use of service to SSH protocol version 2 only as version 1 has many documented security flaws.

Finally the incorrect user login count is decremented for usability reasons however your security model may not agree with this change.

These changes can be implemented for SSH daemons on OES Linux servers by modifying the proper directives in the **/etc/ssh/sshd_config** file.

The referenced configuration file contents are presented below. The majority of the file contents have been omitted for brevity. Lines containing directives and parameters that require modification are displayed in bold.

sshd_config:

```
# The strategy used for options in the default sshd_config shipped with
# OpenSSH is to specify options with their default value where
# possible, but leave them commented.  Uncommented options change a
# default value.

#Port 22
#Protocol 2,1
Protocol 2
#AddressFamily any
#ListenAddress 0.0.0.0
#ListenAddress ::

# Authentication:

#LoginGraceTime 2m
# PermitRootLogin yes
PermitRootLogin no
#StrictModes yes
#MaxAuthTries 6
MaxAuthTries 2
```

VNC service configuration

VNC services allow remote access to Linux and Unix server X-windows desktop environments. Unlike RDP sessions the desktop environment is actually exported to a remote end point instead of allowing remote access to the server desktop environment. This model allows multiple independent desktop sessions to a single server.

VNC requires that the X-Windows desktop manager be running on the target server.

VNC access is not natively secure. It requires HTTPS or SSH frameworks to provide security for the application.

** VNC should be enabled on EISS managed OES Linux and SUSE Linux Enterprise Servers but administrators should always secure the application during use with any of the secure tunneling applications available to them.

Procedures:

Some EISS managed OES Linux or Suse Linux Enterprise Servers may not have a X-Windows Desktop Manager enabled by default. First ensure it is running prior to attempting VNC access to the server.

To enable or verify VNC is available on a server become user root and issue the following command in a SSH session with the server:
/sbin/yast remote

Ensure “Allow remote administration” is selected and select the “Finish” option.

If immediate VNC access is needed issue the following command in a SSH session with the server to start/restart the desktop Manager :
/etc/init.d/xdm stop && /etc/init.d/xdm start

Automated file system check intervals for physical computers

OES Linux and Suse Linux Enterprise Servers that utilize EXT2 and EXT3 file systems have default disk check configurations in place when the server is installed. These automated checks are activated when the file system has been mounted a set number of times and when a specific time interval has elapsed.

It is not recommended to change or disable the default mount count criteria for physical servers. Many of the factors that can cause data loss in these file systems are hardware related and periodic checks of the file systems are an inconvenient necessity.

The default interval for the elapsed time criteria between disk checks is 60 days and can be adjusted to prevent production systems from running disk checks too frequently. This could be very inconvenient for systems with large capacity file systems. The larger the file system the longer the disk check will take to complete. Related systems with very large file systems could also have their periodic disk check events staggered to better accommodate this maintenance.

The **/sbin/tune2fs** application can be used to adjust both of the criteria used to instantiate a disk check events as well as many other file system tuning parameters. This example however will be limited to the disk check event configuration.

Procedure:

- Ensure there are EXT2 and/or EXT3 file systems present that can be configured
  Issue the following command at the server console:
  /bin/mount

You should get a similar output on EISS managed Linux servers on HP DL360 Proliant hardware:

/dev/cciss/c0d0p2 on / type ext3 (rw,acl,user_xattr)
proc on /proc type proc (rw)
sysfs on /sys type sysfs (rw)
debugfs on /sys/kernel/debug type debugfs (rw)
udev on /dev type tmpfs (rw)
devpts on /dev/pts type devpts (rw,mode=0620,gid=5)
/**dev/cciss/c0d0p3 on /srv type ext3** (rw,acl,user_xattr)
/dev/cciss/c0d0p4 on /var type ext3 (rw,acl,user_xattr)
fusectl on /sys/fs/fuse/connections type fusectl (rw)
securityfs on /sys/kernel/security type securityfs (rw)
proc on /var/lib/ntp/proc type proc (rw)

Note the bolded entries which are local EXT3 devices. Only local devices are within the scope of this example. Each device will need to be configured discretely.

EISS managed physical servers have normal capacity (less than 64 GB) file systems configured for a disk check interval of 120 days.

- Issue the following commands at the server console for each normal capacity EXT partition listed:
  /sbin/tune2fs -i 120d /dev/cciss/c0d0p2
  /sbin/tune2fs -i 120d /dev/cciss/c0d0p3
  /sbin/tune2fs -i 120d /dev/cciss/c0d0p4

The output for a successful configuration change should be similar to the following:

/sbin/tune2fs –i 120d /dev/cciss/c0d0p2
tune2fs 1.38 (30-Jun-2005)
Setting interval between checks to 10368000 seconds

Automated file system check intervals for virtual computers

The configuration of default disk check events for EISS managed virtualized Linux servers is very similar to that for physical servers. The only difference is that the mount count based criteria disk check event is disabled. Since hardware failures are much less likely to affect virtualized servers, relying on the elapsed time interval criteria alone should be sufficient to maintain these systems.

Procedure:

- Ensure there are EXT2 and/or EXT3 file systems present that can be configured
  Issue the following command at the server console:
  /bin/mount

You should get a similar output on EISS managed Linux servers:
/dev/sda2 on / type ext3 (rw,acl,user_xattr)
proc on /proc type proc (rw)
sysfs on /sys type sysfs (rw)
debugfs on /sys/kernel/debug type debugfs (rw)

udev on /dev type tmpfs (rw)
devpts on /dev/pts type devpts (rw,mode=0620,gid=5)
/dev/sda3 on /srv type ext3 (rw,acl,user_xattr)
/dev/sda4 on /var type ext3 (rw,acl,user_xattr)
securityfs on /sys/kernel/security type securityfs (rw)

- Issue the following commands at the server console for each normal capacity EXT partition listed:
  /sbin/tune2fs -c 0 -i 120d /dev/sda2
  /sbin/tune2fs -c 0 -i 120d /dev/sda3
  /sbin/tune2fs -c 0 -i 120d /dev/sda4

The output for a successful configuration change should be similar to the following:

/sbin/tune2fs -c 0 -i 120d /dev/sda2
tune2fs 1.38 (30-Jun-2005)
Setting maximal mount count to -1
Setting interval between checks to 10368000 seconds

Installation of server patches

Even when the current supported media is used to install systems there are post release patches, often called “micro-patches”, available from Novell that should be applied. Many of these patches can be downloaded and installed discretely however OES Linux and Suse Linux Enterprise Servers offer online patching services to simplify the patching process. It is required to register the target system with the Novell Customer Center with valid evaluation or purchased license credentials to authorize and configure the online patch channels available to the server. The local applications used to patch the server communicate with update servers managed by Novell to distribute patches. These applications also evaluate which patches to make available using current system software footprint, server platform, and software dependency criteria. This makes it a much better option for successfully patching systems.

Two applications can be used to access the configured online patching services. The Novell ZenWorks Update Agent (**/usr/bin/zen-updater**) and the “Rug” (**/usr/bin/rug**) applications can be used to access the online patches in the X-windows environment and the command line environment respectively.

Both the ZenWorks Update Agent and rug applications require the ZenWorks Management Daemon (ZMD) to be running to access online patch services. The ZMD is disabled on EISS managed OES Linux and SLES systems but it can be loaded and unloaded manually when needed. This allows for some enforcement of how and when patches are applied to servers.

The ZenWorks Update Agent requires root user access to configure it and to enable non-root users to use it to apply patches. The rug utility is accessible using a terminal session and requires root user access to use it to patch servers successfully.

Procedure:

** *This procedure is best performed as a post installation task. Patches should not be applied as part of the initial installation and configuration tasks.*

First the ZenWorks Management Daemon must be started. This can be done from a terminal session at the target server. If the ZenWorks Update Agent will be used open a terminal session within an existing X-windows session.

Start the ZenWorks Management Daemon using the following command:
/etc/init.d/novell-zmd start

Wait a few minutes for the ZMD to initialize after loading successfully.

The preferred method to update EISS managed servers is to use the rug client. It is often quicker and does not require a local persistent user profile to be created.

To use the rug client to update the OES Linux or Suse Linux Enterprise Server:

Issue the following command:
/usr/bin/rug refresh

This will refresh all services and update channel catalogs

Issue the following command:
/usr/bin/rug service-list

This will display the local and remote update channels the server is subscribed to. Ensure the remote channel **https://nu.novell.com** is available or re-register the server with the Novell Customer Center.

Issuing the following command:
/usr/bin/rug catalogs

Will display the patch catalogs the server is subscribed to and the subscription status

Issuing the following command:
/usr/bin/rug list-updates

Will display available updates in the configured update channel catalog(s)

or

Issue the following command:
/usr/bin/rug list-updates <CATALOG_NAME_1> <CATALOG_NAME_2>

This will display available updates for a specific catalog(s)

Actually patching the server:

Issue the following command:
/usr/bin/rug update -t patch <CATALOG_NAME_1> <CATALOG_NAME_2>

For example, if the server should have both OES 2 and SLES 10 patches installed:
/usr/bin/rug update -t patch OES2-SP2-Updates SLES10-SP3-Updates

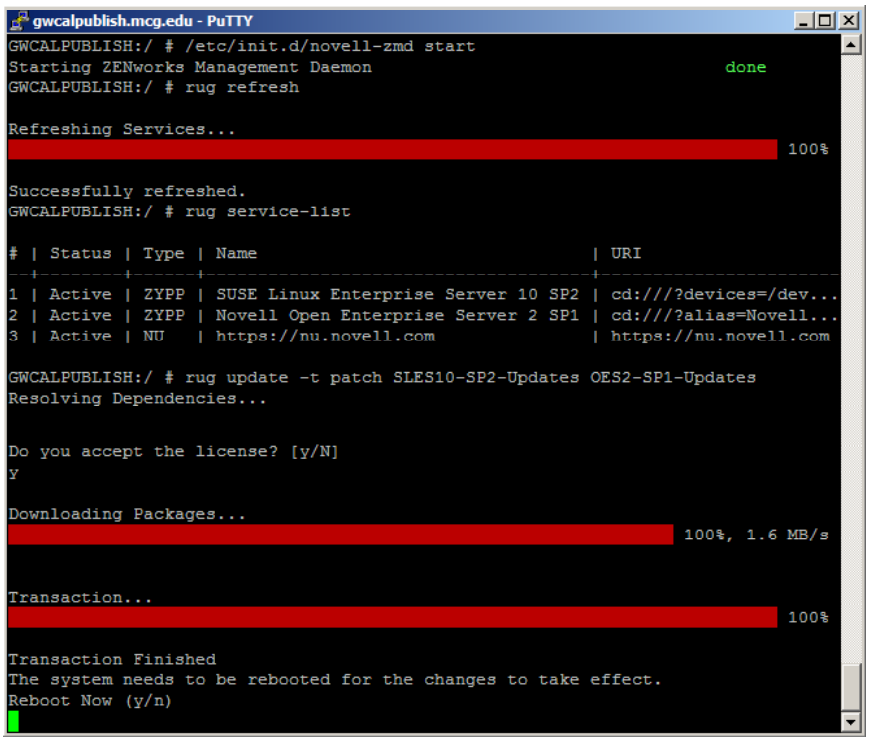

This will install all available patches for the target server. Answer all questions presented by the update service to begin the update process. Reboot the server, restart the ZMD, and repeat the process until all server patches have been applied.

Stop the ZenWorks Management Daemon using the following command:
/etc/init.d/novell-zmd stop
To use the ZenWorks Update Agent to update the OES Linux or Suse Linux Enterprise Server:

Start the ZenWorks Management Daemon by issuing the following command in a terminal session within the X-windows session:
/etc/init.d/novell-zmd start

Issue the following command:
/usr/bin/zen-updater &

** The “&” argument forces the application to become a background process so the terminal instance can be used to issue additional commands.

It may take several minutes for the ZenWorks Update Agent to initialize. When it does an icon will appear in the system tray area of the desktop indicating whether or not updates are available. Double click the icon with the exclamation point to begin the patch installation process. You can also right click the icon to force an agent refresh operation, configure agent properties, or to shut down the agent.

Select the patches to apply and click the “Update” button.

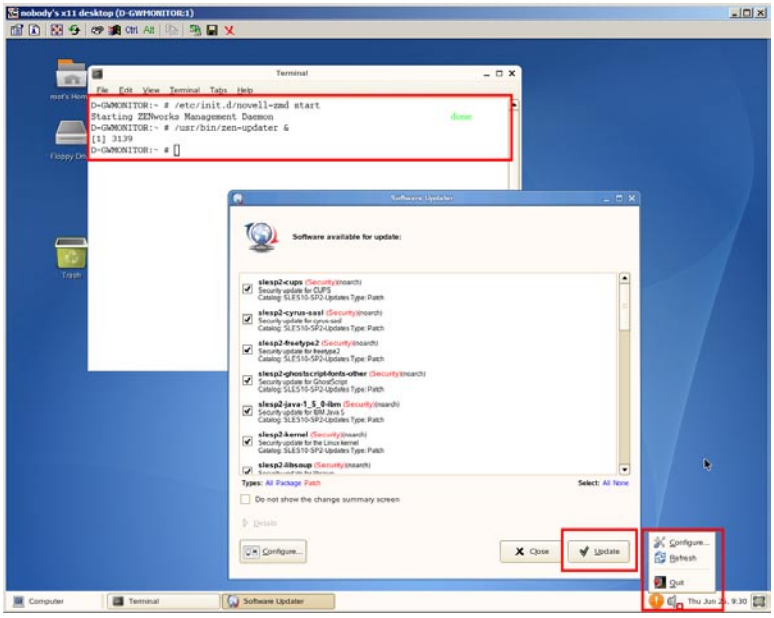

Answer all questions presented by the update service to complete the update process. Shut down the ZenWorks Updater Agent by right clicking the application icon and choosing “Quit”

Stop the ZenWorks Management Daemon by issuing the following command:
/etc/init.d/novell-zmd start

Reboot the server, restart the ZMD, and repeat the process until all server patches have been applied.

Advanced installation of server patches

It should be noted that both applications used to patch OES Linux and Suse Linux Enterprise Servers have some useful advanced features. These features allow administrators to filter patches using view criteria, interrogate patch packages for additional information, apply discrete patches, and establishing check points that can be used to facilitate patch roll back.

Information for all OES 2 server configurations

Special configuration notes for OES systems with NCP installed

It is the standard for EISS NCP object names to utilize all uppercase characters. This is especially true for NCP server object names. When the NCP server component is installed on OES Linux servers the short DNS name and its' existing case is used for the NCP server object name.

Since the network stack is configured prior to OES services during an OES Linux installation the host name portion of the fully qualified DNS name should be entered in all uppercase characters.

For example:
L9.dvc.edu

** *It is advised that the host name on existing SLES servers be modified to match this format if the NCP server component is installed on them at a later date.*

Special configuration notes for OES systems with LUM installed

It is important to understand what Linux User Management (LUM) brings to the table over using Suse Linux Enterprise server LDAP based authentication facilities. An OES Linux server with LUM installed uses special binaries and configuration settings that implement many eDirectory authentication and authorization features on the Linux server. Services like Universal Password support, eDirectory specific account restrictions, file system access and authorization using local Linux applications and even eDirectory password management are supported with LUM.

The LUM implementation in the DVC eDirectory environment has been designed with scalability in mind. Understanding that LUM is a server centric technology that has limitations in a directory enabled environment is an important concept when designing an enterprise service. Although a comprehensive understanding of LUM is of benefit to administrators, if the standards and recommendations in this document are implemented it is not required.

EISS managed servers utilize the following LUM object standards:

OES Linux LUM service objects:

LUM service object	Description	Placement or Naming Standards
Unix Config	Maintains unique UIDs and GIDs	Refer to standards table 2
Unix Workstation	Provides access to Linux servers	Refer to standards table 2

Standards table 2:

LUM service object	Naming Standard	Placement Standard
Unix Config[1]	Unix Config	.DVC
Unix Workstation[2]	Unix Workstation - <SERVER_NAME>	.LUM.SVS.DVC
Unix Workstation[3]	Unix Workstation - <SERVER_NAME>	.<SERVICE_OU>.LUM.SVS.DVC

† Only one Unix Config object exists in EISS managed eDirectory trees
†† Unix Workstation objects for servers used for enterprise wide services are located here
††† Unix workstation objects for servers used for specific application or non-enterprise wide Services are located in a dedicated eDirectory container in the .LUM.SVS.DVC context. If the target container does not exist it is required to be created before the LUM service configuration is attempted.

EISS managed servers utilize the following LUM configuration standards:

- A standard service proxy user[1]
- Multiple secure LDAP targets to access eDirectory user stores
- An increased application thread count
- eDirectory synchronization refresh every 60 minutes
- User home directories are not created (unless this is a requirement for dependant services)
- Persistent LDAP connections to eDirectory targets are not used
- Does not require matching character case for authentication users

All of these configuration changes can be implemented by modifying the **/etc/nam.conf** file.

† Currently if the LUM proxy user is specified at installation time the password for the existing configured proxy user is changed to an unknown value. This can be an issue when a LUM infrastructure is currently in place. LUM can be installed successfully without specifying the proxy user during the installation. Simply configure the proxy user credentials when the nam.conf file is configured post installation.

For example, consider the following example where the bolded entries reflect the referenced configuration options:

base-name=o=DVC
admin-fdn=cn=admin,o=DVC
preferred-server=a2.dvc.edu
alternative-ldap-server-list=a1.dvc.edu
proxy-user-fdn=cn=LUM_PROXY,ou=PROXYUSR,ou=LUM,ou=SVS,o=DVC
proxy-user-pwd=<PROXY_USER_PASSWORD>
num-threads=10
schema=rfc2307
enable-persistent-cache=yes
user-hash-size=211
group-hash-size=211
persistent-cache-refresh-period=3600
persistent-cache-refresh-flag=all
create-home=no
type-of-authentication=2
certificate-file-type=der
ldap-ssl-port=636
ldap-port=389
support-alias-name=no
support-outside-base-context=yes
cache-only=no
persistent-search=no
case-sensitive=no
convert-lowercase=no

Procedure:

- Stop the LUM service daemon by using the following command at the server console: **/etc/init.d/namcd stop**
- Backup the existing **/etc/nam.conf** file on the server being deployed.
- Configure the **/etc/nam.conf** file on the server being deployed to match the one displayed in the example.
- Issue the following command at the server console: **/usr/bin/namconfig –k**

Provide the target eDirectory tree admin password and ensure the target LDAP server certificate imports successfully

- Start the LUM service daemon by using the following command at the server console:
  /etc/init.d/namcd start
  Ensure the service starts successfully and without errors
- Configure the LUM service to use multiple LDAP targets using the information presented in the next section.

Configuring LUM to utilize multiple LDAP servers

You will need to ensure the LDAP server currently configured using the "preferred-server directive in the **/etc/nam.conf** file has its' certificate imported into the into the LUM certificate store and the LUM daemon (namcd) can use the certificate.

Then you will need to temporarily change the "preferred server" directive in the **/etc/nam.conf** file to contain the address(es) of any LDAP server(s) specified by the "alternative-ldap-server-list" directive(s) in turn and import those server certificates into the LUM certificate store.

After stopping and starting the LUM daemon to put changes made to the **/etc/nam.conf** file into effect the **/usr/bin/namconfig** application is used to import the preferred-server certificate. Running the **/usr/bin/namconfig** utility with the **-k** option as root imports the LDAP server certificate for the configured preferred-server into the LUM certificate store.

For example:
Open two terminal sessions as root on the same server. One to monitor logs and one to manage the LUM service components on.

Use the following command on one terminal instance to monitor last 15 lines of the **/var/log/messages** in real time:
tail -f –n 15 /var/log/mssages

Use the <CTRL> <C> key sequence to close the tail application process
Configure the "**preferred-server**" and "**alternative-ldap-server-list**" directives in the **/etc/nam.conf** file so they contain the desired secure LDAP server targets. The "**alternative-ldap-server-list**" directive can contain a comma delimited list of LDAP targets. Secure LDAP connections are enforced by setting the "**type-of-authentication**" directive to a value of "**2**".

Then stop and restart LUM using the following commands on the other terminal instance:
/etc/init.d/namcd stop && /etc/init.d/namcd start

Verify that the messages log file output reflects the current preferred-server LDAP target can be contacted successfully and securely.

If not import the LDAP server certificate into the local LUM certificate store using the following command:
/user/bin/namconfig –k
Provide the LDAP server admin user credentials when prompted

Start and stop the LUM service again and verify that the messages log file output reflects the current preferred-server LDAP target can be contacted successfully and securely.

Configure each LDAP server specified in the "**alternative-ldap-server-list**" directive in the "**preferred-server**" directive and repeat this process until all LDAP server certificates are imported and can be used by the LUM service instance successfully.

Configuring global OES Linux server access for administrators

Configure all new LUM service instances so Server Administration staff that belong to the SA_LUM_G LUM enabled eDirectory group. This group is used by server administration staff to authenticate to all EISS managed OES Linux servers using their eDirectory credentials.

Add all new Linux Workstation objects to “uamPosixWorkstationList” attribute of the .SA_LUM_G.LUM.SVS.DVC LUM enabled eDirectory group.

This is best accomplished using the “Modify Unix Workstation Object” option within the Linux User Management iManager snapins. Adding the Linux Workstation object(s) to this attribute gives LUM access to the new server(s) for all EISS Server Administrators that are members of the group.

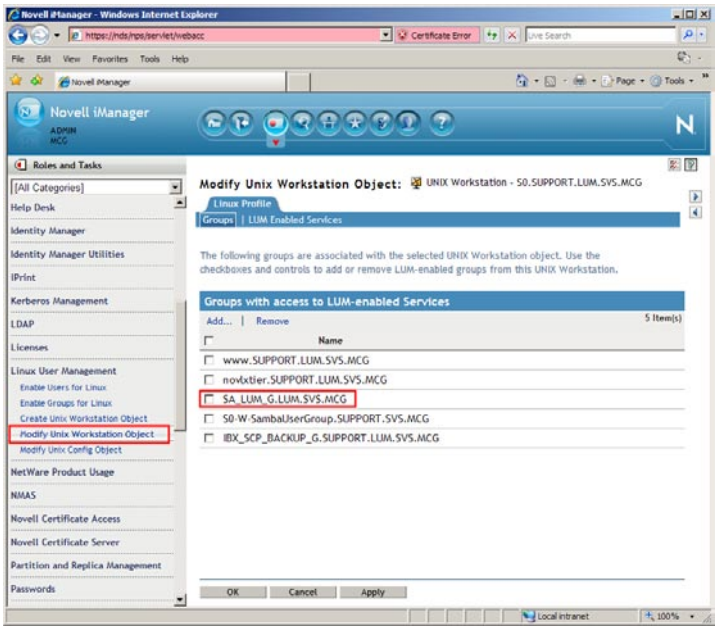

Configuring home directories for OES Linux server administrators

Most EISS managed OES Linux servers are provisioned to provide services and not personalized file storage for users. However it is often useful for server administrators to have some storage unique to them available on the servers they manage. Since the automated creation of home directories is disabled in the standard EISS Linux server deployments the home directories for server administrators must be created manually.

A script file named “makehome.sh” is maintained on EISS deployed Linux servers for this purpose and is placed in the root user’s home directory by administrators. Copy the file to the root user’s home directory on the new server file system from any other Linux server where LUM has been deployed and modify it if necessary so it reflects the current active server administration staff and run the script.

For example, consider the following makehome.sh example:

```
#!/bin/bash
mkdir /home/admin
mkdir /home/BOB
mkdir /home/ED
mkdir /home/WENDY
mkdir /home/GERRY
mkdir /home/JASON
sleep 3
chmod 750 /home/admin
chmod 750 /home/BOB
chmod 750 /home/ED
chmod 750 /home/WENDY
chmod 750 /home/GERRY
chmod 750 /home/JASON
sleep 3
chown ADMIN:SA_LUM_G /home/admin
chown AHK:SA_LUM_G /home/BOB
chown EDH:SA_LUM_G /home/ED
chown FRL:SA_LUM_G /home/WENDY
chown GAL:SA_LUM_G /home/GERRY
chown JSA:SA_LUM_G /home/JASON
```

Ensure there are no errors thrown by the script and all home directories are created successfully.

Information for OES 2 server configurations using NSS

These Novell NSS and SMS components implemented on OES Linux all have their own respective text based configuration files.

OES configuration file	File system path	Used by
nssstart.cfg	/etc/opt/novell/nss/	Novell Storage Services
smdrd.conf	/etc/opt/novell/sms/	Novell SMDR daemon
tsafs.conf	/etc/opt/novell/sms/	Novell OES Linux TSA

Common configuration settings that enhance the functionality and performance of these services are displayed below.

Recommended settings are in bold.

nssstart.cfg:
/ListXattrNWMetadata
/CtimeIsMetadataModTime

*** These settings are required for several SMS compliant backup services to successfully backup and restore data from NSS volumes hosted on OES Linux servers. Consult you backup service documentation to see if they should be implemented in your environment.*

smdrd.conf: **(Can also be modified using iManager)**
hosts: enable
slp: enable
ip: 0.0.0.0 **(Can be set to a specific IP address if necessary)**
autoload: tsafs **(tsafs should auto-load with the SMDR service)**
priority: slp, hosts
uds: enable

*** It may be useful to restart the SMDR daemon in NCS load and unload scripts when NCP volumes are migrated between nodes to ensure the SMS state is refreshed on both nodes involved in the migration.*

tsafs.conf:	**(Can also be modified using iManager)**
cluster=enable	**(Verify backup service compatibility for a cluster node)**
cachememorythreshold=10	
readaheadthrottle=2	
readbuffersize=65536	
readthreadallocation=100	
readthreadsperjob=4	
tsamode=Dual	**(Allows TSA to process POSIX, NCP, and NSS targets)**
cachingmode=enable	

Checking for the presence of an nssid.sh file

Check the "**/opt/novell/oes-install**" directory on servers where NSS is installed, whether during the initial server installation or as a post installation task, for the presence of a system generated script file named "**nssid.sh**".

This script will most likely be generated for the following reasons:

You have installed an OES 2 Linux server using NSS into an eDirectory tree that currently has or has ever had OES 1 Linux servers using NSS.

You have installed an OES 1 Linux server using NSS into an eDirectory tree that currently has or has ever had OES 1 Linux or OES 2 Linux servers using NSS.

If this script is present you should run it as the root user as one of your first post NSS installation tasks. The purpose of this script is to synchronize key file ownership information for specific OES system users. Not running the script could result in server and service functionality and security issues. When the script is run successfully the script file is removed from the system.

NSS volume configuration

It is often advisable to disable the use of the file and directory access time meta-data as it is rarely of benefit for NSS file systems. Maintaining access time meta-data can actually negatively impact volume performance for large or frequently accessed volumes.

To disable access time management enter the NSS console using the "**nsscon**" command and then enter the following NSS command:

nss /noatime=<VOLUME_NAME>

When implementing NSS volumes on OES 1 Linux systems the name space for the volume defaulted to "Unix". The Unix name space default allowed for some compatibility benefits for clients when implementing NSS on OES Linux systems but injected some performance issues. If the NSS volume will be used exclusively for NCP aware services or clients the default name space for the volume should be changed to "Long".

When implementing NSS pools and volumes on OES 2 Linux server the default name space is set to "Long" but you should verify this during new volume deployments.

Use the following command to invoke the NSS Management Utility at the OES 2 Linux server console hosting the target NSS volume(s):

/sbin/nssmu

From the menu select “Volumes” and press <ENTER>. Then select each mounted NSS volume to review its configured settings.

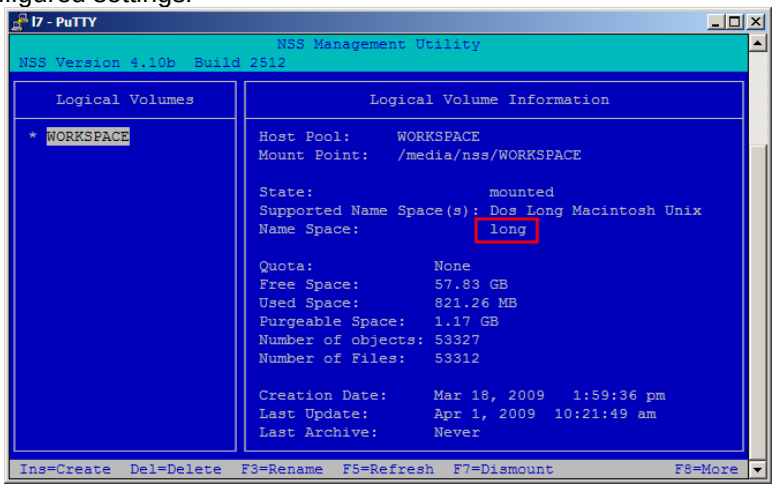

If the volume name space is not set to “Long” press the <F8> key until the name space option is available in the lower status bar. Then press the <F9> key to change the name space to the desired default.

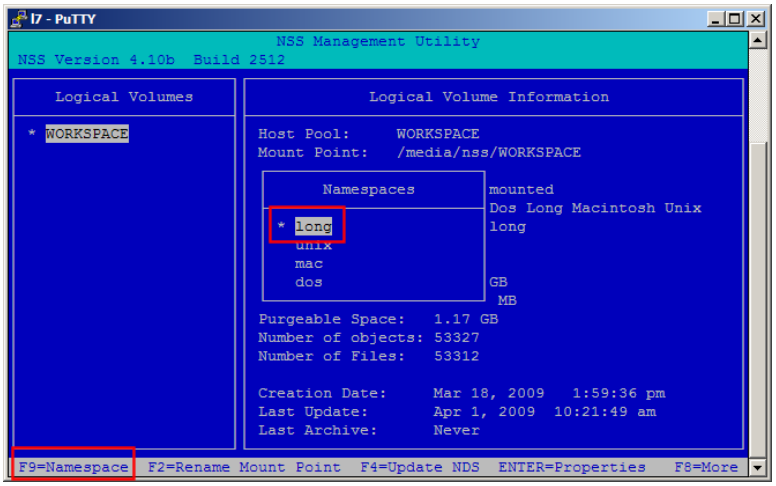

*** You can also make the desired volume name space changes using the storage plugins in the Novell iManager application.*

eDirectory management configuration

Many configuration options can benefit your eDirectory implementations on the OES Linux platform. However two should always be configured without fail when EISS managed OES Linux servers are deployed. They are the Service Location Protocol service (SLP) and the Novell DSRepair Menu wrapper script.

SLP configuration is detailed in an earlier section but will be presented again.

The Service Location Protocol (SLP) configuration and role for EISS managed OES Linux servers are usually configured during installation. However the SLP configuration and role for servers can be adjusted, and should be verified, post server installation as well. Generally the only SLP service configuration parameter changed for OES Linux servers are the Directory Agents (DA)

specified for the service to use. In the EISS managed environment the primary and secondary DA servers user are slp1.dvc.edu (10.6.10.74) and slp2.dvc.edu (10.6.11.136) respectively. For SLP service configuration the SLP software is biased to prefer the IP address for service configuration.

Procedures:

Modify the /etc/slp.conf file as instructed below:

The relevant section from the **/etc/slp.conf** SLP service configuration file containing the DA server parameters in bold is displayed below:

```
#---------------------------------------------------------------------------
# Static Scope and Static DA Configuration
#---------------------------------------------------------------------------

# Allows administrator to force UA and SA agents to use specific DAs.  If
# this setting is not used dynamic DA discovery will be used to determine
# which DAs to use.  (Default is to use dynamic DA discovery)
net.slp.DAAddresses = 10.6.10.74,10.6.11.136

#---------------------------------------------------------------------------
# UA Specific Configuration
#---------------------------------------------------------------------------
# A 32 bit integer giving the maximum number of results to accumulate and
# return for a synchronous request before the timeout, or the maximum number
# of results to return through a callback if the request results are
# reported asynchronously (default value is 256).
net.slp.maxResults = 1024
```

Some of the most common commands an administrator is most likely to need and use to manage the SLP on OES Linux are displayed below:

To determine the scope(s) currently known by the SLP agents on an OES Linux server use the following command:
/usr/bin/slptool findscopes

To determine the providers of a particular type of service currently known by the SLP agents on an OES Linux server use the following command:
/usr/bin/slptool findsrvs service:<SERVICE_TYPE>

To determine the DAs currently known by the SLP agents on an OES Linux server use the following command:
/usr/bin/slptool findsrvs service:directory-agent

To determine the SAs currently known by the SLP agents on an OES Linux server use the following command:
/usr/bin/slptool findsrvs service:service-agent

NDS repair for Unix Menu Wrapper

The NDS Repair for UNIX Menu Wrapper is implemented with a shell script file named "**dsrmenu.sh**". The dsrmenu.sh script file is simply a front end to simplify the use of the eDirectory maintenance and management command line utility **/usr/bin/ndsrepair**. However it provides a menu with the same look and feel as the DSREPAIR.NLM application on OES NetWare servers.

Administrators new to eDirectory management on the UNIX, Linux, and Solaris platforms may find the command line tools challenging to use at first. One nuance of using the script file is that it displays the command line that would be used to complete a task selected from the menu to the console or terminal operator. This should help administrators learn how to use the command line utility.

Procedures:

Installation of the NDS Repair menu wrapper is a simple as placing the dsrmenu.sh script file in the **/usr/sbin** directory on EISS managed OES Linux servers. The current version used is 1.04.01. This version includes support for all current eDirectory versions and is available from Novell's public download site or it can simply be copied from an OES Linux server that already has it installed.

To use the utility run the "/usr/sbin/dsrmeu.sh" command as the root user then select the eDirectory version and instance from the subsequent menus. You will then be presented with the familiar menu shown below.

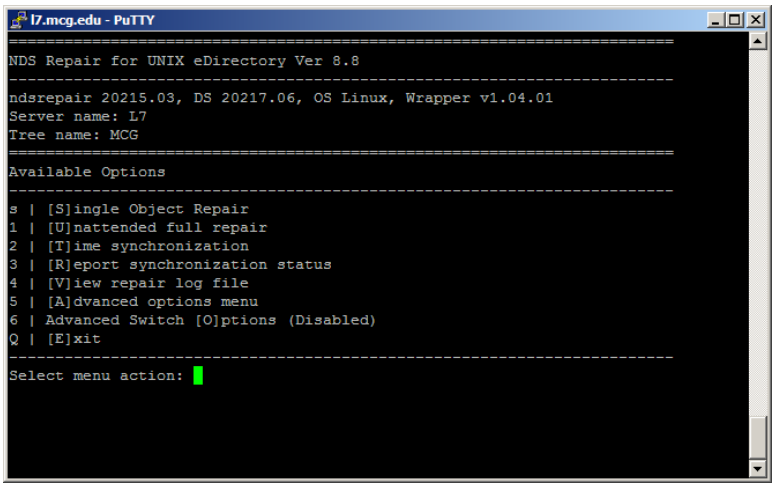

LDAP eDirectory configuration

Novell services running on OES Linux servers are heavily dependant on a properly configured LDAP service infrastructure. Key to any managed LDAP infrastructure is the security implemented. Many security components simply take the form of policies and standards. This may seem unremarkable initially but these types of standards and policies are a common server and service hardening tool.

The LDAP services hosted on EISS managed OES Linux and OES NetWare servers are managed using dedicated service proxy accounts and network address restrictions. The LDAP service proxy user is an eDirectory user object in the DVC eDirectory tree that has read only eDirectory rights to a sub-set of eDirectory user object attributes and values. The marriage of this rights model for the LDAP service proxy account with the TCP/IP network address restrictions greatly enhances the security of the OES LDAP implementation. This means that for a server or service to offer LDAP services for consumption by LDAP clients in the DVC eDirectory tree the following must be true:

1. The server LDAP service must be configured to use the DVC eDirectory tree LDAP service proxy user.

2. One of the server's TCP/IP network addresses must be added to the DVC eDirectory tree LDAP service proxy user's "Address Restriction" attribute.

** *You should not install OpenLDAP on an OES Linux server intended to host Novell services.* Just as with OES NetWare based servers, server specific LDAP eDirectory objects are created for OES Linux servers at the time of server installation. The two objects created are the LDAP SERVER - <SERVER_NAME> and the LDAP Group - <SERVER_NAME> and they are placed in the same eDirectory context of the host server.

The EISS standards for LDAP services are implemented through the configuration of these objects. The configured object parameters and their values are presented in the diagrams below.

LDAP object	Setting implemented or configured
LDAP Group - <SERVER_NAME>	LDAP service proxy user **(cn=LDAP_PROXY.O=DVC)**
	Require TLS for simple binds with passwords **(Enabled)**
LDAP SERVER - <SERVER_NAME>	Require TLS for all operations **(Disabled)**
	Enable persistent search **(Disabled)**
	Enable event monitoring **(Disabled)**

** *It is important to note that disabling TLS for all operations on OES servers can expose user credentials to interception efforts. It is possible to deny unsecure bind attempts but the credentials used are still sent over the network unencrypted. It may be best to prevent this scenario as well.*

** *Consider disabling the persistent search and event monitoring services for your OES Linux servers and their eDirectory targets. These services can negatively impact your LDAP environment performance if you have not engineered the infrastructure for load considerations or manage the behavior of the clients or applications consuming the provided LDAP services.*

LDAP service configuration

OES 2 Linux servers offer enhanced features for configuring the server LDAP connectivity when compared to OES 1 Linux servers. Most remarkable is the ability to configure a hierarchy of multiple LDAP targets.

EISS managed OES Linux servers do not currently hold eDirectory partition relicas. Because they will always defer any LDAP searches to remote servers always place the primary enterprise remote LDAP service targets first in the server list when configuring the LDAP services for the server. Conversely the local server should be specified last.

Procedures:

Access the Yast utility as user root on the target server. Access the OES LDAP configuration by selecting the "Open Enterprise Server" module, then select the "OES Install and Configuration" option.

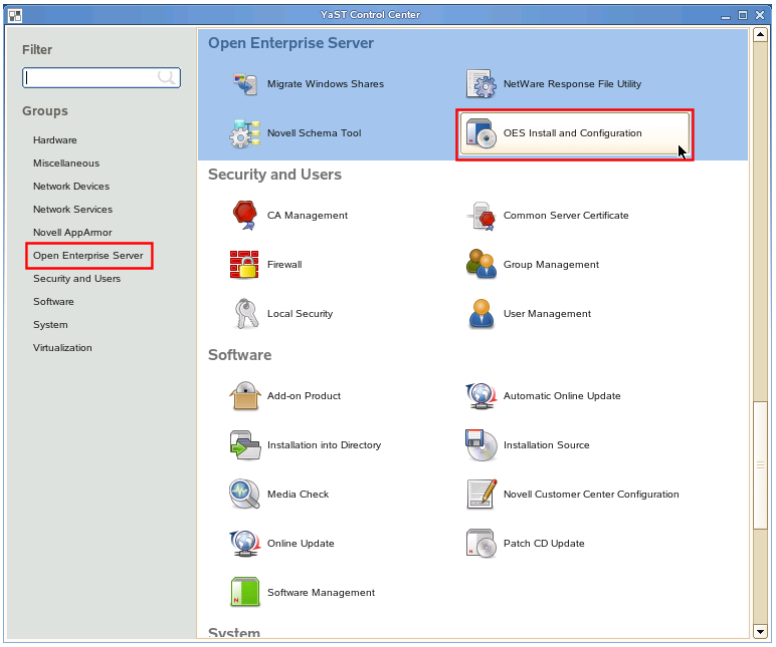

Do not make any software package changes on the subsequent dialog screen and click the "Accept" button.

Click the "enable" link on the "LDAP Configuration for Open Enterprise Server" section of the form and then select the LDAP configuration option from the "Change" drop down button menu.

Configure your LDAP service targets as recommended providing any required credentials.

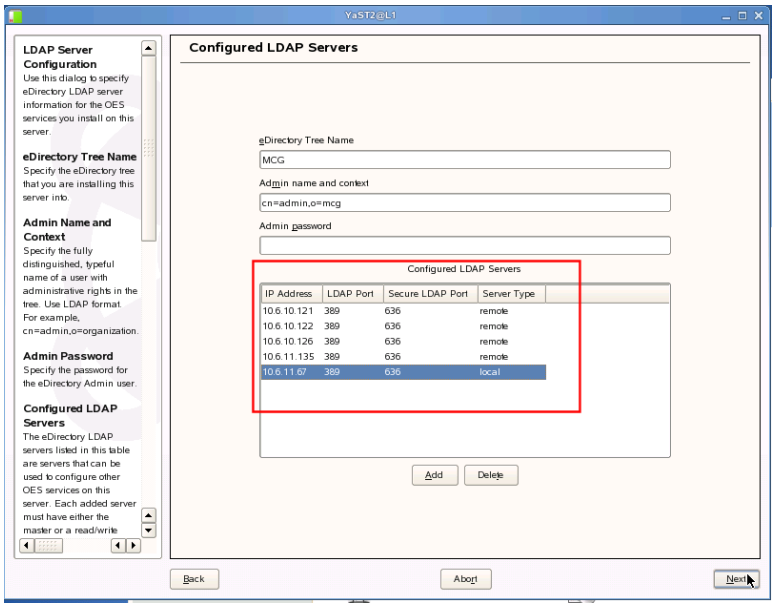

Click the "Next" button and then accept or confirm your selections on any subsequent dialog screens.

*** If the target server does have local eDirectory partition replicas and could benefit from using them you could of course adjust your LDAP server configuration accordingly.*

Light Directory Access Protocol (LDAP) on EISS managed OES Linux servers is implemented using the Novell provided **/opt/novell/eDirectory/sbin/nldap** application. The **/etc/init.d/nldap** script or the actual binary file can used to start and stop the nldap application on OES Linux servers. The nldap application is not configured to automatically load at server start. It is started dynamically by other OES applications, like eDirectory, that have LDAP dependencies.

Some of the most common command lines to manage the nldap application an administrator is most likely to need and use are displayed below:

To manually stop the Novell nldap application use the following command(s):
nldap -u or
/etc/init.d/nldap stop

To manually start the Novell nldap application using the following command(s):
nldap -l or
/etc/init.d/nldap start

To check the status of TCP and LDAPS port for the Novell nldap application use the following command:
/etc/init.d/nldap -s

To refresh the Novell nldap application to instate configuration changes use the following command and provide the password when prompted:
ldapconfig -R -a .<admin_user>.<context>

Information for clustered OES 2 server configurations

Clustered file system and Novell Cluster Service script configuration

Novell Clustering Services (NCS) uses system scripts to predictably transition the states of managed cluster resources by passing structured commands to the cluster software. NCS cluster resource load and unload scripts have a 1024 character limit. Since this limit can be easily exceeded it is considered best practice to “batch” required service commands in specialized script files. Those script files can then be referenced in NCS cluster resource load and unload scripts using standard Linux path and file name conventions.

Using sets of scripts for different types of services injects some organizational and management benefits into NCS configurations. For example, scripts to load and unload application services, web services, and miscellaneous services or agents could be implemented. All of these specified scripts are processed by the NCS software automatically but if there are issues it may be helpful to load and unload them discreetly.

The SAN attached storage can be used as the central repository for the scripts referenced by NCS load and unload scripts. The primary benefits of doing this are that there is a single set of scripts to maintain and the scripts can be invoked or modified manually if NCS management facilities are unavailable. A sort of “out of band” NCS management option results from this practice.

The cluster resource file system should be standardized for this model. The root of the volume should be used to store global configuration data such as that used by NCS or enterprise backup applications. Subdirectories should be created for each application managed by NCS and these should serve as the root for their specific configuration, application data and/or files.

For example consider the following volume used by NCS to host the service script files used by NCS, Tivoli Storage Manager, and the Novell Teaming application and configuration files.

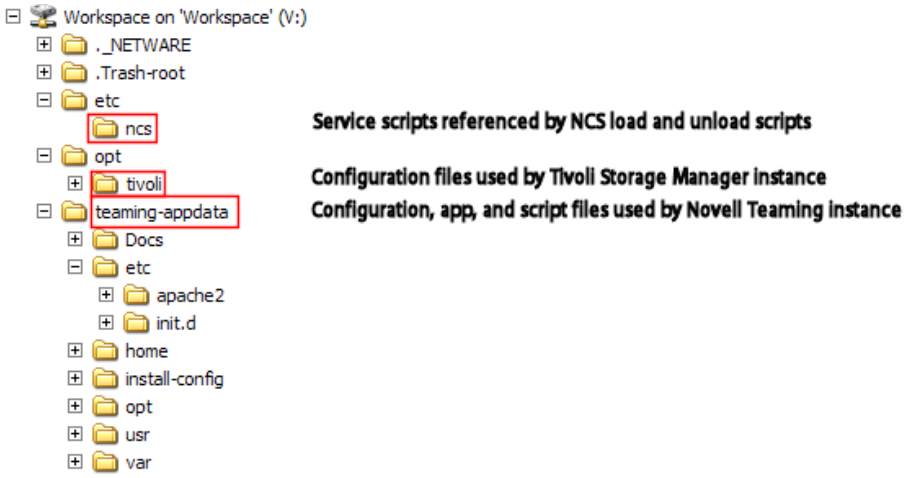

The NCS cluster resource load script that references the example file system depicted is displayed below:

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs
exit_on_error nss /poolact=WORKSPACE
exit_on_error ncpcon mount WORKSPACE=252
exit_on_error add_secondary_ipaddress 10.6.11.243
exit_on_error ncpcon bind --ncpservername=WORKSPACE --ipaddress=10.6.11.243
/media/nss/WORKSPACE/etc/ncs/agents-load.sh
/media/nss/WORKSPACE/etc/ncs/teaming-load.sh
exit 0
```

Using discrete scripts to start and stop services keeps the NCS load and unload scripts orderly and easy to use and understand. The script entries in bolded text are examples of this implementation model.

Required NCS service scripts

Currently the only NCS service scripts administrators are required to provision on clustered OEs Linux file systems are:

/media/nss_or posix>/<MOUNT_POINT>/etc/ncs/agents-load.sh
/media/nss_or posix>/<MOUNT_POINT>/etc/ncs/agents-unload.sh
/media/nss_or posix>/<MOUNT_POINT>/etc/ncs/agents-load.sh.original
/media/nss_or posix>/<MOUNT_POINT>/etc/ncs/agents-unload.sh .original

The **"<nss_or_posix>"** path component reflects the EISS standards for different types of clustered volumes. NCS technology on OES Linux servers allows you to cluster both NSS and POSIX shared disk devices so this path component reflects the storage type used.

The **"<MOUNT_POINT>"** path component reflects the name of the configured NSS pool or the EVMS device name. Again this path component is an EISS standard and helps reflect the storage type used.

The **agents-load.sh** and **agents-unload.sh** files are used minimally to restart the SMDR daemon when a clustered volume changes states or migrates to another cluster node. Many file system backup and restoration services are SMS aware and restarting the SMDR daemon when a file system is mounted or dismounted is a good precaution to take. Doing so should help the reliability of your backup and restoration services.

The **agents-load.sh.original** and **agents-unload.sh.original** files are simply backups done as a best practice precaution.

The agents load and unload scripts can and should be used to start and stop generic agents or services. Generic in this context refers to services native to the OS or common across servers.

An example of an agents-load.sh service script that loads generic services:

```
#!/bin/bash
# Restarts the Novell SMDR daemon after a volume mount state change
/etc/init.d/novell-smdrd restart

# Starts the Tivoli Storage Manager instance configured for this file system
/media/nss/WORKSPACE/opt/tivoli/tivoli-start.sh
```

This script restarts the SMDR daemon and starts the Tivoli Storage Manager instance for the file system.

Procedures:

These scripts should only be composed with a native Linux text editor, or an editor known to be compatible with the Linux text file symantecs.

It may also be prudent to place a text file in the script which warns users not to edit the files with an incompatible editor.

- Compose the following agents load and unload scripts in the **/media/nss_or posix>/<MOUNT_POINT>/etc/ncs/**directory:

  agents-load.sh
  #!/bin/bash
  # Restarts the Novell SMDR daemon after a volume mount state change
  /etc/init.d/novell-smdrd restart

  agents-unload.sh
  #!/bin/bash
  # Restarts the Novell SMDR daemon after a volume mount state change
  /etc/init.d/novell-smdrd restart
- Copy these files to files of the same name with .original file name extensions in the same directory.
- Create an empty text file named "**Do not modify these files with a windows text editor.txt**" in the same directory
- Modify the cluster resource NCS load and unload script to call the appropriate agent load and unload script on resource load and unload events.

Managing Enterprise Services

Server Administration Templates for Novell OES Linux Server Provisioning

System proxy user configurations

- NCS

-- Address restrictions should be used.

NCS is not dependent upon the tree admin user. It's possible to create a cluster admin proxy user, and enter that user's credentials during cluster configuration. The cluster admin user needs create rights to whatever container you intend to create the cluster object in.

-- Assigned to proxy user at the organizational level

[Entry Rights]	[B] Inheritable
CN	[CR] Inheritable

-- Assigned to proxy user at the Cluster object organizational unit level

[Entry Rights]	[BC] Inheritable
[All Attribute Rights]	[CR] Inheritable

** Create containers that contain cluster objects low in the tree to avoid security issues with the [All Attributes] assignment or adjust it accordingly.

- Samba

-- Address restrictions should be used.

-- Assigned to proxy user at the organizational level

Assign the following eDirectory object and attribute rights to the context which is to contain the samba domain objects.

The “Create” right to the parent context is required to allow the proxy user to create new samba domain objects only. Minimally the inheritance flag should be removed or the “Create” right can be added and removed when a new samba domain object needs to be created.

[Entry Rights]	[BC] Inheritable
ObjectClass	[CRW] Inheritable
sambaAcctFlags	[CRW] Inheritable
sambaAlgorithmicRidBase	[CRW] Inheritable
sambaBadPasswordCount	[CRW] Inheritable
sambaBadPasswordTime	[CRW] Inheritable
sambaDomainName	[CRW] Inheritable
sambaHomeDrive	[CRW] Inheritable
sambaHomePath	[CRW] Inheritable
sambaKickoffTime	[CRW] Inheritable
sambaLMPassword	[CRW] Inheritable
sambaLogoffTime	[CRW] Inheritable
sambaLogonHours	[CRW] Inheritable
sambaLogonScript	[CRW] Inheritable
sambaLogonTime	[CRW] Inheritable

sambaMungedDial	[CRW]Inheritable
sambaNTPassword	[CRW]Inheritable
sambaNextGroupRid	[CRW]Inheritable
sambaNextRid	[CRW]Inheritable
sambaNextUserRid	[CRW]Inheritable
sambaPasswordHistory	[CRW]Inheritable
sambaPrimaryGroupSID	[CRW]Inheritable
sambaProfilePath	[CRW]Inheritable
sambaPwdCanChange	[CRW]Inheritable
sambaPwdLastSet	[CRW]Inheritable
sambaPwdMustChange	[CRW]Inheritable
sambaSID	[CRW]Inheritable
sambaUserWorkstations	[CRW]Inheritable

- Samba user groups
-- Have no assigned eDirectory rights but do have Linux service restriction attributes in addition to being LUM enabled

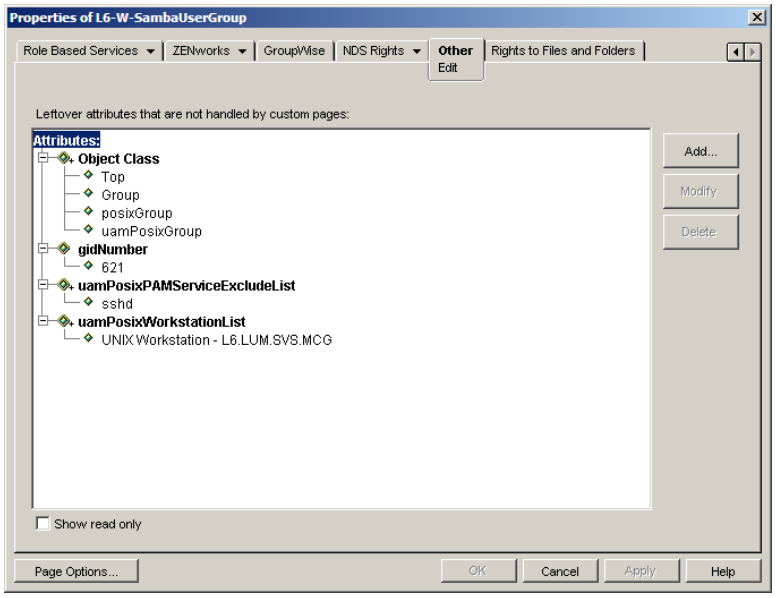

Samba attributes indexing

Deploying the Samba service in large eDirectory environments often requires some backend service optimization. Most often performance will degrade or seem noticeably poor during authentication routines when the potential for object, attribute, and attribute value searches are the most demanding. Indexing services can be configured to overcome this particular bottleneck.

The “Object Class” and “SambaSID” attributes are indexed on the LDAP targets used for samba authentication services. Indexing these attributes greatly decreases search times when authenticating users to the Samba service.

- LUM
-- Address restrictions should be used.

(Also used by the NCS configuration daemon when managing a cluster through iManager)

-- Assigned to proxy user at the LUM organizational unit level

-- .LUM.SVS.DVC

[Entry Rights]	[BC] Inheritable
[All Attribute Rights]	[CR] Inheritable
Group Membership	**[CRW] Inheritable**

-- Assigned to proxy user at the organizational level

[Entry Rights]	[B] Inheritable
[Class: User]	
gecos	[CRW] Inheritable
homeDirectory	[CRW] Inheritable
loginShell	[CRW] Inheritable
primaryGroupID	[CRW] Inheritable
uidNumber	[CRW] Inheritable
uamPosixSalt	[CRW] Inheritable
uniqueID	[CR] Inheritable
Login Disabled	[CR] Inheritable
Login Expiration Time	[CR] Inheritable
Password Allow Change	[CR] Inheritable
Password Expiration Interval	[CR] Inheritable
Password Expiration Time	[CR] Inheritable
Password Minimum Length	[CR] Inheritable
Password Required	[CR] Inheritable
Password Unique Required	[CR] Inheritable
[Class: Group]	
uamPosixWorkstationList	[CRW] Inheritable
[Common Attributes]	
CN	[CR] Inheritable
Description	[CRW] Inheritable
ObjectClass	[CRW] Inheritable
gidNumber	[CRW] Inheritable
Group Membership	**[CR] Inheritable**
memberUid	[CRW] Inheritable
uamPosixPAMServiceExcludeList	[CRW] Inheritable
uamPosixWorkstationContexts	[CRW] Inheritable
uamPosixGidNumberDeletedMap	[CR] Inheritable
uamPosixGidNumberEnd	[CR] Inheritable
uamPosixGidNumberLastAssigned	[CRW] Inheritable
uamPosixGidNumberReuse	[CR] Inheritable
uamPosixGidNumberStart	[CR] Inheritable
uamPosixUidNumberDeletedMap	[CR] Inheritable
uamPosixUidNumberEnd	[CR] Inheritable
uamPosixUidNumberLastAssigned	[CRW] Inheritable
uamPosixUidNumberReuse	[CR] Inheritable
uamPosixUidNumberStart	[CR] Inheritable
uamPosixWorkstationList	[CRW] Inheritable

- GroupWise Messenger

-- Address restrictions should be used.

-- Assigned to proxy user at target organizational unit level(s)

[Entry Rights]	[BC]	Inheritable
Description	[CR]	Inheritable
Internet Email Address	[CR]	Inheritable
L	[CR]	Inheritable
OU	[CR]	Inheritable
Title	[CR]	Inheritable
nnmBlocking	[CRW]	Inheritable
nnmBlockingAllowList	[CRW]	Inheritable
nnmBlockingDenyList	[CRW]	Inheritable
nnmClientSettings	[CRW]	Inheritable
nnmContactList	[CRW]	Inheritable
nnmCustomStatusList	[CRW]	Inheritable
nnmLastLogin	[CRW]	Inheritable

- QuickFinder

-- Address restrictions should be used.

- A valid user with Supervisor file system rights to the directory(s) that contain the site configuration and their respective index files.

Administrative users are authenticated via eDirectory (or PAM on Linux) and authorized access if they have write rights to the configuration file in the product directory (SYS:\qfsearch on NetWare® and /var/lib/qfsearch on Linux).

Access to the administrative interface is restricted to valid users that have write rights to the configuration file in the product directory.