

IBM ServerGuide Scripting Toolkit, Linux Edition



Altiris Deployment Solution User's Reference

Version 940

IBM ServerGuide Scripting Toolkit, Linux Edition



Altiris Deployment Solution User's Reference

Version 940

Note:

Before using this information and the product it supports, read the information in "Notices".

Contents

Chapter 1. Introducing the ServerGuide Scripting Toolkit, Linux Edition for Altiris Deployment Solution 1

Chapter 2. Installing the ServerGuide Scripting Toolkit, Linux Edition. 3

Extracting the Linux Scripting Toolkit	3
Adding the Altiris Linux x64 Automation Environment	3
Adding the IBM Linux Pre-installation Environment for x64	4
Adding System Enablement Packs	4
Adding a new supported system using SEP without Internet connectivity	5
Integrating with Altiris Deployment Solution	6
Reinstalling the Linux Scripting Toolkit	7

Chapter 3. Adding and sharing content on the Linux Scripting Toolkit source server 9

Adding UpdateXpress System Packs	9
Adding device drivers	9
Adding Systems Director Agent files	10
Adding operating system installation files	10
Sharing operating system files on your network	11

Chapter 4. Quick start deployment scenarios 13

Validating the Altiris environment.	13
Modular deployment tools	13
Pre-installation tasks	14
Operating system installation	18
Post installation tasks	20
Deployment solutions	21

Chapter 5. Customizing deployment scenarios 23

Customizing sample jobs	23
Sample job definitions	29
Customizing Fibre HBA boot configuration.	29
Capture and Deploy RAID Configuration and Linux image	33
Configure RAID + Install Linux OS	34
Configuring RAID with a policy file	34
Capture and deploy RAID configuration.	34
Server data disposal	35
Linux and VMware unattended scripted installation	35
Capture and deploy a Linux image	36

Linux Pre-Boot Test.	37
------------------------------	----

Chapter 6. Supported hardware and software 39

Operating system support	39
RAID controller support	40
Fibre Channel HBA support	40

Chapter 7. Linux Scripting Toolkit utilities and tools 41

altiris_sgtklinux.cmd	41
acquire_seps.cmd	41
Linux Scripting Toolkit utilities	42
HWDTECT	42
SAVSTAT.	45
PRAID	46
INVRAID	64
VALRAID	67
Tools included with the Linux Scripting Toolkit	70
Advanced Settings Utility	70
QAUCLI	71
septool210.	72
UpdateXpress System Pack Installer	72

Chapter 8. Incorporating the Scripting Toolkit with your existing process . . . 75

Chapter 9. IBM ServerProven compatibility 77

Chapter 10. Hints and tips 79

Known problems and limitations	83
--	----

Appendix A. Getting help and technical assistance 91

Before you call	91
Using the documentation.	91
Getting help and information from the World Wide Web	92
Software service and support	92
Hardware service and support	92

Appendix B. Notices 93

Edition notice	93
Trademarks	93
Important notes	93

Index 95

Chapter 1. Introducing the ServerGuide Scripting Toolkit, Linux Edition for Altiris Deployment Solution

The ServerGuide Scripting Toolkit, Linux Edition for Altiris Deployment Solution (Linux Scripting Toolkit) enables you to tailor and build custom hardware deployment solutions. It provides hardware configuration and Linux operating system (OS) installation for IBM System x, BladeCenter, and iData-Plex hardware.

When integrated with the Altiris Deployment Solution, the Linux Scripting Toolkit simplifies the steps in creating and customizing jobs to deploy hardware configurations and Network Operating System (NOS) deployments. The Linux Scripting Toolkit supports the following types of deployment:

- Policy-based RAID configuration using PRAID
- Cloning of RAID configuration
- Configuration of System settings using the ASU (ASU)
- Configuration of Fibre Channel Host Bus Adapters (HBAs)
- Firmware updates using the UpdateXpress System Pack Installer (UXSPi)
- UpdateXpress System Pack installation integrated with automated deployment of a Network Operating System (NOS)
- IBM Systems Director Agent installation integrated with automated deployment of a NOS
- Automated deployment of the following Network Operating Systems (NOSs):
 - SUSE Linux Enterprise Server 9 32 bit SP4
 - SUSE Linux Enterprise Server 9 x64 SP4
 - SUSE Linux Enterprise Server 10 32 bit SP1/SP2/SP3/SP4
 - SUSE Linux Enterprise Server 10 x64 SP1/SP2/SP3/SP4
 - SUSE Linux Enterprise Server 11 32 bit Base/SP1/SP2/SP3
 - SUSE Linux Enterprise Server 11 x64 Base/SP1/SP2/SP3
 - Red Hat Enterprise Linux 4 AS/ES 32 bit U6/U7/U8
 - Red Hat Enterprise Linux 4 AS/ES x64 U6/U7/U8
 - Red Hat Enterprise Linux 5 32 bit U1/U2/U3/U4/U5/U6/U7/U8/U9
 - Red Hat Enterprise Linux 5 x64 U1/U2/U3/U4/U5/U6/U7/U8/U9
 - Red Hat Enterprise Linux 6 32 bit U1/U2/U3/U4
 - Red Hat Enterprise Linux 6 x64 U1/U2/U3/U4
 - VMware ESX Server 3.5 U4/U5
 - VMware ESX Server 4.0/4.0u1/4.0u2/4.1/4.1u1/4.1u2/4.1u3/4.1u4
- Automated post-installation deployment of the following NOSs:
 - Red Hat Enterprise Linux 4 AS/ES 32 bit U9
 - Red Hat Enterprise Linux 4 AS/ES x64 U9
- Automated deployment of the following NOSs in Native uEFI mode:
 - SUSE Linux Enterprise Server 11 SP1/SP2
 - Red Hat Enterprise Linux 6 x64 U1/U2/U3/U4
- Remote Supervisor Adapter II (RSA II) and BladeCenter Management Module and Advanced Management Module remote disk scenarios
- Integration with Altiris Deployment Solution

- Cloning of NOS images
- Installation of IBM Systems Director Agent integrated with scripted NOS deployment
- Post UXSP installation using the Altiris Agent
- Post IBM Systems Director Agent installation using the Altiris Agent
- Remote deployment via the Integrated Management Module (IMM)

The toolkit provides sample jobs for Altiris Deployment Solution for hardware configuration and operating system deployment using Linux PE. The Linux Scripting Toolkit supports Altiris Deployment Solution 6.9 SP1-SP5 and the Hewlett-Packard Rapid Deployment Pack (RDP) 3.80, 3.81, 3.82, 3.83, 6.0, and 6.2.

Chapter 2. Installing the ServerGuide Scripting Toolkit, Linux Edition

This section describes the steps to install the Linux Scripting Toolkit. It includes information on the prerequisites, instructions for installation, and instructions for removal.

Prerequisites

The ServerGuide Scripting Toolkit, Linux Edition for Altiris Deployment Solution has three components:

1. The IBM Pre-boot environment
2. The IBM Scripting Toolkit Scripts and Utilities
3. The IBM Scripting Toolkit Altiris jobs within the Altiris Deployment Solution console

When you install or re-install the ServerGuide Scripting Toolkit, Linux Edition for Altiris Deployment Solution, these components are updated as a group. After the update, the Altiris Deployment Solution jobs associated with the previous version of the ServerGuide Scripting Toolkit, Linux Edition for Altiris Deployment Solution no longer work.

The Linux Scripting Toolkit requires the following components to install and create the IBM Linux x64 PE environment:

- Altiris Linux Automation Environment for Deployment Solution
BDCgpl_6.9.xxxx.frm, where xxxx is the version number
- IBM Linux Pre-installation Environment boot files for x64
ibm_utl_boot_tools-xxx_anyos_x64-full.zip where xxx is the version number.

The Linux Scripting Toolkit also requires that the target server on which the deployment is to be executed has at least 1 GB of memory.

Extracting the Linux Scripting Toolkit

This section describes the steps for extracting the Linux Scripting Toolkit zip file into the Altiris Deployment Server folder.

Procedure

1. Right-click the Linux Scripting Toolkit zip archive and select **Extract All**.
2. Extract the files to the Altiris Deployment Server directory. By default this directory is C:\Program Files\Altiris\Express\Deployment Server.
3. Verify that the Linux Scripting Toolkit is extracted. If the files extracted properly, they are located in ...\\Altiris\Express\Deployment Server\sgdeploy\sgtklinux\

Adding the Altiris Linux x64 Automation Environment

This section describes how to add the Altiris Linux Automation Environment for Linux x64 if it is not already installed.

Before you begin

Acquire the Linux Automation Environment file. You can acquire this file from the Altiris Deployment Solution or from the Hewlett-Packard Rapid Deployment Pack.

You can get the Altiris Deployment Solution file from <http://www.altiris.com>. The filename is `BDCgpl_6.9.xxxx.frm` where `xxxx` is the version number of the file.

The Hewlett-Packard Rapid Deployment Pack version is located on the DVD at: `\ds\BDCgpl_6.9.xxxx.frm`, where `xxxx` is the version number of the file.

Procedure

There are two ways to install the Altiris Linux x64 Automation Environment:

- Use the Altiris PXE Configuration Utility to install the pre-boot environment manually.
- Copy the file `BDCgpl_6.9.xxxx.frm` to `...\sgdeploy\sgtklinux\boot\`, and retain the original filename. When the `altiris_sgtklinux.cmd` installation script runs, it integrates the environment as part of the installation.

Adding the IBM Linux Pre-installation Environment for x64

If the Altiris Deployment Server is connected to the web and can reach <http://www.ibm.com>, the install process downloads the IBM Linux PE zip file.

Before you begin

This procedure requires access to the Internet. This topic describes the process for acquiring the files when the Altiris Deployment Server does not have Internet connectivity.

Procedure

1. Copy the file `sgdeploy\sgtklinux\altiris\bin\win32\ibm_utl_bomc_x.xx_windows_i386.exe` from the Altiris Deployment Server to a system connected to the Internet.
2. Run the following command to acquire the pre-installation environment zip file and save it in `C:\temp\`:

```
ibm_utl_bomc_x.xx_windows_i386.exe --function=linuxtk -l C:\temp
```

The Bootable Media Creator tool acquires the pre-installation environment zip file `ibm_utl_boot_tools-140_anyos_x86-64-full.zip` and saves it in `C:\temp\` on the workstation.

3. Copy `C:\temp\ibm_utl_boot_tools-140_anyos_x86-64-full.zip` to the Altiris Deployment Server at `..\sgdeploy\sgtklinux\boot\`, preserving the original filename.

Adding System Enablement Packs

System Enablement Packs (SEPs) allow you to add support for hardware released after the current release of the ServerGuide Scripting Toolkit, Linux Edition for Altiris Deployment Solution. This section describes the process for adding SEPs to the Altiris Deployment Server.

Before you begin

If the Altiris Deployment Server is connected to the Web and can reach www.ibm.com, use the **acquire_seps.cmd** command to acquire System Enablement Packs (SEPs) as described here. If the Altiris Deployment Server does not have Internet connectivity, see “Adding a new supported system using SEP without Internet connectivity” for information on using a system other than the Altiris Deployment Server to acquire SEPs.

Procedure

1. Modify the default settings for acquiring SEPs as required by editing the file `acquire_seps.ini` and modifying the appropriate values. The default location of this file is `C:\Program Files\Altiris\express\Deployment Server\sgdeploy\sgtklinux\altiris\acquire_seps.ini`. The following table describes the values you can modify.

Value	Description
TK_SEP_OS	A comma-separated list of operating systems for which to download SEPs. Valid values are: <i>none</i> , <i>all</i> . The default value is <i>none</i>
TK_SEP_Arch	A comma-separated list of processor architectures for which to download SEPs. Valid values are: <i>x86</i> , <i>x64</i> , <i>all</i> .
TK_SEP_PKG_TGT	The target directory on the local system for storing SEPs.
TK_SEPTOOL_OptionalParams	Additional parameters to be used when downloading SEPs.

2. Open a command window and change to the directory containing the **acquire_seps.cmd** file. By default, this file is located at `C:\Program Files\Altiris\express\Deployment Server\sgdeploy\sgtklinux\altiris\acquire_seps.cmd`.
3. Enter the **acquire_seps.cmd** command, followed by a comma-delimited list of machine types. For example, `acquire_seps.cmd 7979,4199`. You can also acquire SEPs for all available machine types by specifying *all* instead of a machine type: `acquire_seps.cmd all`.

Adding a new supported system using SEP without Internet connectivity

System Enablement Packs (SEPs) allow you to add support for hardware released after the current release of the . This section describes the process for adding SEPs to the Toolkit Source Server when the server does not have Internet connectivity.

Before you begin

If the Linux Scripting Toolkit Source Server is not connected to the web, complete the following steps to acquire SEPs.

Procedure

1. Copy the file `septool.zip` file (`win_septoolxxx.zip`, where *xxx* is the version number of the tool) from the Toolkit Source Server to a system with Internet connectivity. The default location for this file is:

2. On the system where you copied the zip file, extract all of the files in the archive.
3. From the directory where you extracted the zip file, run the following command to acquire the SEP and save it in C:\temp:

```
septoolxxx.exe acquire -l C:\temp -m machine_type -o none -a x64
```

where *xxx* is the version of the septool and *machine_type* is the machine type of the system for which you want to download SEPs.

4. Copy the files from C:\temp and place them in the updates folder of the Linux Scripting Toolkit directory tree. The default location is .

Integrating with Altiris Deployment Solution

This section describes the requirements and method for integrating the Linux Scripting Toolkit into your existing Altiris Deployment Solution console.

Installation

To integrate the Linux Scripting Toolkit into the Altiris Deployment Solution console, follow these steps:

1. To suppress prompts or user settings before installing the Scripting Toolkit into the Altiris Deployment Solution console, open the `altiris_sgtklinux.ini` file and modify the appropriate values shown in the table below.

Value	Description
Section: [altiris_SGTKLinux Build Settings]	
TK_Build_SuppressPrompts	<p>Suppress interactive prompts which pause the process to enable the user settings to be read after they are displayed. Valid values are:</p> <ul style="list-style-type: none"> • Yes • No <p>Note: Interactive prompts regarding required removal of jobs during upgrade or reinstall scenarios are always displayed.</p>
TK_Build_DisplayBuildSettings	<p>Display the user settings before installing the Toolkit into Altiris Deployment Solution. Valid values are:</p> <ul style="list-style-type: none"> • Yes • No

2. Open a command window and change the current directory to the directory containing the `altiris_sgtklinux.cmd` file. The default file location is `C:\Program Files\Altiris\express\Deployment Server\sgdeploy\sgtklinux\altiris\altiris_sgtklinux.cmd`.
3. At the command window enter the **`altiris_sgtklinux.cmd`** command. It processes the Toolkit settings in the `altiris_sgtklinux.ini` file.
This process can take between 5 and 15 minutes to complete depending on the server setup.
4. When the `altiris_sgtklinux.cmd` command completes, restart the server and all Altiris services to complete the integration.

After the Toolkit is installed, there is a new folder with the product name, version, and build ID under the Jobs frame in the Altiris Deployment Solution console. For example: **IBM Scripting Toolkit, Linux Edition 9.40**

Reinstalling the Linux Scripting Toolkit

This section describes the process for reinstalling the Linux Scripting Toolkit.

The ServerGuide Scripting Toolkit, Linux Edition for Altiris Deployment Solution has three components:

- The IBM Pre-boot environment
- The IBM Scripting Toolkit Scripts and Utilities
- The IBM Scripting Toolkit Altiris jobs within the Altiris Deployment Solution console

When you install or re-install the ServerGuide Scripting Toolkit, Linux Edition for Altiris Deployment Solution, these components are updated as a group. After the update, the Altiris Deployment Solution jobs associated with the previous version of the ServerGuide Scripting Toolkit, Linux Edition for Altiris Deployment Solution no longer works.

To re-install the Linux Scripting Toolkit, follow these steps:

1. Follow the instructions from the “Installation” on page 6 section of “Integrating with Altiris Deployment Solution” on page 6.
2. The installer displays the following message to indicate that a previous installation has been detected:

The following installed items were detected:

IBM Scripting Toolkit, Linux Edition 2.00

This must be uninstalled before this process can continue. Please select an option below:

- 1) Exit The Reinstaller.
- 2) Automatically Reinstall The Scripting Toolkit, Linux Edition.

Select an above option:

Follow the steps below for the appropriate removal method.

Manually removing the Linux Scripting Toolkit

To manually remove the Linux Scripting Toolkit, complete these steps:

1. Select option number 1 from the prompt to exit the installation process.
2. Remove all jobs associated with the IBM® Linux Toolkit image in the Altiris Deployment Solution console.
3. Open the PXE configuration utility from the Tools menu within the Altiris Deployment Solution Console.
4. Select the IBM Linux Toolkit image.
5. Verify that the column labeled **In Use By DS** does not indicate **Yes** for that image. If it does, there are additional jobs associated with that image. Find and remove them.
6. Click the **Delete** button to remove the PXE image.
7. Click **Save** to save the changes
8. Click **OK** to exit the PXE Configuration utility.

Automatically removing the Linux Scripting Toolkit

To automatically remove the Linux Scripting Toolkit, follow these steps:

1. Select option 2 from the prompt. You receive the following prompt:

```
*****  
You have chosen to reinstall the following:
```

```
IBM Scripting Toolkit, Linux Edition 2.00
```

```
You may keep the jobs associated with the currently installed  
Scripting Toolkit, however, they will be moved to a folder labeled  
IBM Scripting Toolkit, Linux Edition (Archived w/ver. 2.00).  
*****
```

```
Would you like to keep those jobs?<Y/N>
```

2. Select **Y** to remove the current Linux PE image and move the current IBM ServerGuide Toolkit, Linux Edition jobs to a folder named IBM ServerGuide Toolkit, Linux Edition (Archive). Select **N** to remove all previously installed jobs, not including any user-created jobs.

Chapter 3. Adding and sharing content on the Linux Scripting Toolkit source server

This section describes how to add to the source server and share content on your network with the source server.

You can add the following types of files to the source server:

UpdateXpress System Pack (UXSP) files

to the sgdeploy/updates/uxsp/ directory.

Application files, such as the IBM Systems Director Agent

to the sgdeploy/apps/ directory.

Operating system files

to the sgdeploy/os/ directory.

The remaining topics in this section describe the process for adding these files and for sharing them across the network.

Adding UpdateXpress System Packs

To add UpdateXpress System Packs (UXSPs) to the source tree, you must download them to the sgdeploy\updates\uxsp directory on the source server.

About this task

Follow this procedure to add UXSPs to the source server:

Procedure

1. Start the UpdateXpress System Pack Installer, `ibm_utl_uxspi_x.xx_winsrvr_32-64.exe`. The default location is `C:\Program Files\Altiris\Express\Deployment Server\sgdeploy\updates\uxsp\`.
2. Use the UpdateXpress System Pack Installer to download the updates for the machine types and operating systems you need into the source tree. Download the files to `sgdeploy/updates/uxsp`

Note: Do not manually create directories under `sgdeploy/updates/uxsp/`. Doing so prevents the UpdateXpress System Pack Installer from locating the update files.

3. Click **OK** if you are prompted to overwrite existing files.

What to do next

If the Altiris Deployment Server is not connected to the Internet, you can download the updates to an internet-connected machine and copy them to the `sgdeploy/updates/uxsp/` directory.

Adding device drivers

To inject extra device drivers, which are required for some options on the target server, into deployment media, you must first import the device drivers to the Altiris Deployment Server by following this process.

Procedure

1. Start the Toolkit Configuration Utility, **tkconfig.exe**. The default location is C:\Program Files\Altiris\Express\Deployment Server\sgdeploy\sgtklinux\tkconfig\tkconfig.exe.
2. Select **Add Operating System Device Drivers** from the task list. The wizard starts.
3. Select the machine type and OS type that you want to deploy, then click **Next**.
4. Download the device driver package from the link to a local path.
 - If the wizard displays Don't need extra driver packages for the Machine Type and OS type, go to step 7.
 - If the deployment scenario needs extra device drivers, the driver package download link is shown on the wizard. There are two types of device drivers, optional and mandatory. For optional drivers, you can ignore and go to step 7. For mandatory drivers, you must import.
5. Select the location of the package.
6. Click **Next** to import required device drivers to the Altiris Deployment Server.
7. If desired, repeat the above steps to import extra device drivers for other deployment scenarios.
8. Click **Finish** to exit the wizard.

Adding Systems Director Agent files

To include the IBM Systems Director Agent in a deployment, you must first add the files to the Altiris Deployment Server by following this process.

Procedure

1. Download the latest version of the IBM Systems Director Agent installation files from <http://www.ibm.com/systems/management/director/downloads.html>.
2. Unzip the files in a convenient location.
3. Start the Toolkit Configuration Utility, **tkconfig.exe**. The default location is C:\Program Files\Altiris\Express\Deployment Server\sgdeploy\sgtklinux\tkconfig\tkconfig.exe.
4. Select **Add Operating System Application Files** from the task list. The Operating System Application files wizard starts.
5. Select **IBM Director 6.x or IBM Director 6.x.x for Linux**.
6. Select **Core Services Agent (Level 1) or Common Agent (Level 2)**.
7. Browse to the location where you unzipped the Director Agent files, and click **Next** to begin copying the files.
8. When the copy is complete, exit the wizard.

What to do next

You can now include installation of the IBM Systems Director Agent in your deployments.

Adding operating system installation files

To include operating system files in your deployment image, you must make the files accessible to the deployment server. You can share the location of the files with the server, or you can copy the files to the server by following these steps.

Before you begin

If the operating system files you want to use in your deployments are already in a location accessible to the deployment server, you do not have to add them to the deployment server itself.

Note: The Linux Scripting Toolkit supports only Service Pack 4 for SuSE Linux Enterprise Server (SLES) 9. To add Service Pack 4 to the source server, you must first add all CDs of base SLES 9. After adding these files, you can add the service pack following the same procedure. You must select the same target directory you used for the base when adding the service pack.

About this task

If the operating system files are not already in a location accessible to the deployment server, add them to the server by following these steps.

Procedure

1. Start the Toolkit Configuration Utility, **tkconfig.exe**. The default location is C:\Program Files\Altiris\express\Deployment Server\sgdeploy\sgtklinux\tkconfig\tkconfig.exe.
2. Select **Add Operating System Application Files** from the task list. The Operating System Application files wizard starts.
3. Select the operating system you want to add from the top panel.
4. Select the version of the operating system to import from the lower panel.
5. Select the location of the files.
6. Click **Next** to copy the files to the source server.
7. Repeat steps 3 through 6 for each operating system you want to add.

Sharing operating system files on your network

The operating system files you are using for your deployments must be shared over the network to be accessible to the target server during deployment. Follow these steps to share the files.

About this task

In order to complete a deployment, the operating system files used for the deployment must be shared with the target system using one of the following protocols:

- FTP
- HTTP
- NFS

In order to share the files from the Altiris Deployment Server, you must configure the server to share the sgdeploy/ directory.

The operating system installation jobs provided by the Linux Scripting Toolkit use the FTP protocol. Follow the steps below to configure the an Altiris Deployment Server running Windows Internet Information Server (IIS) to share the sgdeploy/ directory.

Note: Please consult the IIS documentation for more options and security considerations.

Procedure

1. Start the Internet Information Services Manager.
2. Select **Default FTP Site**.
3. Right click and select **New > Virtual Directory...**
4. Set the alias name to **sgdeploy**.
5. Set the home directory to **C:\Program Files\Altiris\Deployment Server**.
6. Allow Read access.
7. Start the FTP Site Service.

What to do next

You can now use the operating system files in the sgdeploy directory in your deployments.

Chapter 4. Quick start deployment scenarios

This section contains basic information on deployment scenarios to allow you to begin using the Linux Scripting Toolkit with Altiris Deployment Solution as quickly as possible.

The Linux Scripting Toolkit provides a collection of sample jobs for use by the Altiris Deployment solution. The following topics provide additional information on the most commonly used jobs. For more information on the sample jobs provided by the Scripting Toolkit, see “Sample job definitions” on page 29

The Linux Scripting Toolkit for use with Altiris Deployment Solution provides two types of jobs:

- “Modular deployment tools”
- “Deployment solutions” on page 21

For more information on tailoring deployments to your needs, see Chapter 5, “Customizing deployment scenarios,” on page 23.

Validating the Altiris environment

After installing the Linux Scripting Toolkit for use with Altiris Deployment Solution, you should validate that the environment is operating properly. This topic describes the process for doing so, using the sample job provided.

Before you begin

Complete the installation method appropriate for your environment as described in Chapter 2, “Installing the ServerGuide Scripting Toolkit, Linux Edition,” on page 3.

About this task

When the installation steps are complete:

Procedure

1. Open the Altiris Deployment Solution console.
2. Run the job **Linux Pre-Boot Test**, located in the Pre-Boot OS Connectivity Test folder, in the **IBM Scripting Toolkit, Linux Edition 9.40** folder. The job boots into the Linux PE x64 environment on the target server. Once connectivity is established, it performs a **wait** task.

Results

If the task is successful, you can then right-click on the computer icon and select **Power Option > Shutdown**. Execution of any subsequent job will use Wake on LAN to wake the system.

Modular deployment tools

The sample jobs provided in the **Modular Deployment Tools** folder are designed to be used individually to complete common deployment tasks. This section describes how and in what order to use the modular deployment tools.

The modular deployment tools provided by the Linux Scripting Toolkit are each designed to complete one step in the deployment process. The following sections will provide information on using the sample jobs for:

- Pre-installation hardware configuration
- Operating system installation
- Post-installation configuration

Pre-installation tasks

This section provides information on the tasks available before you install an operating system using the modular tools.

Before installing an operating system using the modular tools, you can perform the following tasks:

- Replicate system settings
- Configure RAID
- Replicate RAID
- Configure Fibre HBAs
- Update firmware using UpdateXpress System Packs
- Server disk analysis
- Server data disposal

The following sections describe how to perform these tasks using the modular tools.

Replicate system settings

Before installing an operating system, you can capture system settings from a compatible server to a file and deploy them to your new server using the sample jobs provided by the Linux Scripting Toolkit and the Advanced Settings Utility.

The Toolkit provides sample jobs to capture and deploy settings using ASU. These jobs are located in the **IBM Scripting Toolkit, Linux Edition 9.40 > Modular Deployment Tools > Step 1 - Pre-installation Hardware Configuration > System Settings Replication** folder.

Capture system settings

The **Capture System Settings** sample job captures the system settings of the server and stores them in a file named `ASU_Capture`. By default all captured System Settings are stored in `sgdeploy/sgtklinux/ini/asu/PlatformID` where *PlatformID* is the machine type of the captured system.

Deploy system settings

The **Deploy System Settings** sample job deploys the system settings from the `ASU_Capture.ini` file located in the specified directory.

You can customize the following variables for these jobs from the Altiris Console. The values shown are the defaults.

- `TK_ASU_File=sgdeploy/sgtklinux/ini/asu/\${BIOS_Build_Level:0:2}/asu_capture`
- `TK_ASU_Mode=save`

RAID configuration

Before installing an operating system, you can configure RAID on the target server using the Linux Scripting Toolkit and the PRAID utility.

Linux Scripting Toolkit provides sample jobs to perform a number of RAID configurations. These jobs are located in the **IBM Scripting Toolkit, Linux Edition 9.40 > Modular Deployment Tools > Step 1 - Pre-installation Hardware Configuration > RAID Configuration** folder, and are named according to the RAID configuration they will create. For example, **Configure RAID [RAID 0]**, and **Configure RAID [RAID 1 + Hot Spare]**.

To perform RAID configuration using the modular tools:

1. Start the Altiris Deployment Solution console.
2. Navigate to the **IBM Scripting Toolkit, Linux Edition 9.40 > Modular Deployment Tools > Step 1 - Pre-installation Hardware Configuration > RAID Configuration** folder.
3. Run the appropriate job against the target system to create the RAID configuration you want.

RAID replication

If you have an existing RAID configuration that you want to replicate to other systems, you can use the sample jobs **RAID Capture** and **RAID configuration** to capture and deploy the settings to a new target server.

The sample jobs are located in the **IBM Scripting Toolkit, Linux Edition 9.40 > Modular Deployment Tools > Step 1 - Pre-installation Hardware Configuration > RAID Replication** folder in the Altiris Deployment Solution console.

To capture an existing RAID configuration, follow these steps:

1. Start the Altiris Deployment Solution console.
2. Navigate to the **IBM Scripting Toolkit, Linux Edition 9.40 > Modular Deployment Tools > Step 1 - Pre-installation Hardware Configuration > RAID Replication** folder.
3. Run the **Capture RAID configuration** job against the system whose RAID configuration you want to replicate.

You can customize the following variables for these jobs from the Altiris Deployment Solution console. The values shown are the defaults:

- **TK_pRAID_File**=sgdeploy/sgtklinux/ini/praid/*Machine_Type*/raid_capture.ini where *Machine_Type* is the machine type of the system being captured.
- **TK_pRAID_Flags**=

When you have captured the RAID configuration, you can deploy it using the **Deploy RAID configuration** job in the **IBM Scripting Toolkit, Linux Edition 9.40 > Modular Deployment Tools > Step 1 - Pre-installation Hardware Configuration > RAID Replication** folder.

Configure a Fibre HBA

You can use this pre-installation task to configure your Fibre Channel Host Bus Adapter (HBA) to boot from your Storage Area Network (SAN).

Before you begin

In order to boot from the SAN, it must be configured with a logical drive mapped to a Logical Unit Number (LUN). The HBA you are configuring must have permission to access the logical drive and the LUN.

About this task

Note that the default settings configure the first drive mapped to LUN 0 as a boot device (for QLogic Fibre HBAs only. For information Emulex Fibre HBAs, see “Known problems and limitations” on page 83.). To avoid errors, ensure that your SAN configuration includes at least one logical drive mapped to LUN 0.

Procedure

1. Start the .Altiris Deployment Solution console
2. Navigate to the **IBM Scripting Toolkit, Linux Edition 9.40 > Modular Deployment Tools > Step 1 - Pre-installation Hardware Configuration > Fibre Configuration** folder.
3. Run the **Configure Fibre HBA for Boot** job against the target system. You can customize the following variables for these jobs from the Altiris Deployment Solution console. The default values are shown.
 - TK_FIBRE_COUNT="1"
 - TK_FIBRE_1_HBA_ID="0"
 - TK_FIBRE_1_BOOT_DISABLE="No"
 - TK_FIBRE_1_BOOT_PRIM="0 0 0"
 - TK_FIBRE_1_BOOT_ALT1=""
 - TK_FIBRE_1_BOOT_ALT2=""
 - TK_FIBRE_1_BOOT_ALT3=""

Update Firmware

You can use the **Update Firmware (UXSP)** job to update all of the firmware on a target server from the IBM Preboot environment.

About this task

To update the firmware on a server using the **Update Firmware (UXSP)** job, follow these steps:

Procedure

1. Acquire UpdateXpress System Packs (UXSPs) for SUSE Linux Enterprise Server 10 for all of the systems to be updated.
2. Start the .Altiris Deployment Solution console
3. Navigate to the **IBM Scripting Toolkit, Linux Edition 9.40 > Modular Deployment Tools > Step 1 - Pre-installation Hardware Configuration > Update Firmware (UXSP)** folder.
4. Run the **Update System Firmware** job against the target system. You can customize the following variables for these jobs from the Altiris Deployment Solution console. The default values are shown.
 - TK_UXSP_UpdateXpressSystemPacks="/sgdeploy/updates/uxsp"
 - TK_UXSP_ApplyLatest="No"
 - TK_UXSP_UXSPIUpdateFlags="update --unattended --firmware"

Server disk analysis

This sample job captures the disk information for all disks attached to a given server. You can use this job after performing RAID configuration to ensure that when you install a supported operating system, you install to disk /dev/sda on the target server.

Before you begin

You must perform RAID configuration before running this job, as the disk information collected during analysis is changed during RAID configuration.

About this task

To capture disk information, follow these steps:

Procedure

1. Start the Altiris Deployment Solution console.
2. Navigate to the **IBM Scripting Toolkit, Linux Edition 9.40 > Modular Deployment Tools > Step 1 - Pre-installation Hardware Configuration > Server Disk Analysis** folder.
3. Run the **Capture Disk Data** job against the target system. You can customize the following variables for these jobs from the Altiris Deployment Solution console. The default values are shown.
 - TK_DISK_File=/mnt/ds/Temp/\\$ID/disk_info.txt

Server data disposal

The Linux Scripting Toolkit provides sample jobs to destroy data from drives and optionally to delete or reset RAID configuration. These jobs are located in the **IBM Scripting Toolkit, Linux Edition 9.40 > Modular Deployment Tools > Step 1 - Pre-installation Hardware Configuration > Server Data Disposal** folder.

The following topics describe the process for using these sample jobs.

Server data disposal with RAID reset:

Use this job to remove data from the drives and reset the RAID configuration of the target system.

Procedure

1. Start the Altiris Deployment Solution console.
2. Navigate to the **IBM Scripting Toolkit, Linux Edition 9.40 > Modular Deployment Tools > Step 1 - Pre-installation Hardware Configuration > Server Data Disposal** folder.
3. Run the **Server Data Disposal and Reset RAID** job against the target system. You can customize the following variables for these jobs from the Altiris Deployment Solution console. The default values are shown.
 - TK_Wipe_Disk=1
 - TK_Wipe_Repeat_Number=00
 - TK_Wipe_Level=quick

Server data disposal without RAID reset:

Use this job to remove data from the drives without resetting the RAID configuration of the target system.

Procedure

1. Start the Altiris Deployment Solution console.

2. Navigate to the **IBM Scripting Toolkit, Linux Edition 9.40 > Modular Deployment Tools > Step 1 - Pre-installation Hardware Configuration > Server Data Disposal** folder.
3. Run the **Server Data Disposal non-RAID** job against the target system. You can customize the following variables for these jobs from the Altiris Deployment Solution console. The default values are shown.
 - TK_Wipe_Disk=1
 - TK_Wipe_Repeat_Number=00
 - TK_Wipe_Level=quick

Operating system installation

After you have configured RAID on the target server, you can use the jobs provided by the Linux Scripting Toolkit to install a supported operating system. This section describes that process.

The Linux Scripting Toolkit provides two types of jobs for deploying operating systems:

- Operating system cloning
- Scripted operating system installation

The following sections describe how to perform each type of installation.

Operating system cloning

The Linux Scripting Toolkit provides sample jobs to capture and deploy supported Linux operating systems.

Before you begin

In order to run the **Capture Linux Image** sample job, the Altiris agent must first be installed on the donor system. The scripted operating system installation jobs provided by the Linux Scripting Toolkit install the Altiris agent automatically. If you are cloning an operating system that was not installed using the Toolkit, you must install the Altiris agent manually.

About this task

To capture a supported operating system from one server and deploy it to another, follow these steps:

Procedure

1. Start the Altiris Deployment Solution console.
2. Navigate to the **IBM Scripting Toolkit, Linux Edition 9.40 > Modular Deployment Tools > Step 2 - Operating System Installation > Operating System Imaging** folder.
3. Open the **Capture Linux Image** job.
4. Select the **Create Disk Image** task and click **Modify**.
5. Change the path and filename for the captured image and click **Finish** to save your changes.
6. Open the **Deploy Linux Image** job.
7. Select the **Distribute Disk Image** task and click **Modify**.
8. Change the path and filename to match the ones you entered for the captured image and click **Finish** to save your changes.

9. Run the **Capture Linux Image** job against the donor server.
10. When the capture job is complete, run the **Deploy Linux Image** job against the target server.

Scripting operating system installation

If you do not want to deploy a clone of an existing operating system installation, the Linux Scripting Toolkit provides sample jobs to deploy each of the supported operating systems. Use the following procedure to install a supported operating system using the sample jobs provided:

Before you begin

Before you can install an operating system, the operating system installation files must be available to the source server. See “Adding operating system installation files” on page 10 for more information.

Procedure

1. Start the Altiris Deployment Solution console.
2. Navigate to the **IBM Scripting Toolkit, Linux Edition 9.40 > Modular Deployment Tools > Step 2 - Operating System Installation > Scripted Operating System Installation** folder.
3. Select the job that corresponds to the operating system you want to install and run it against the target server.

Scripted operating system installation in uEFI

Linux Scripting Toolkit provides sample jobs to deploy supported operating systems in uEFI mode. This topic describes the process for performing these deployments.

Before you begin

Before you can install an operating system, the operating system installation files must be available to the source server. See “Adding operating system installation files” on page 10 for more information.

About this task

Use the following procedure to install a supported operating system in uEFI mode using these sample jobs:

- **Combined Deployment Solutions > Scripted Deployments:**
 - **Configure RAID + Install RHEL 6 x64 in uEFI**
 - **Configure RAID + Install SLES 11 x64 in uEFI**
- **Modular Deployment Tools > Scripted Operating System Installation:**
 - **Install RHEL 6 x64 in uEFI**
 - **Install SLES 11 x64 in uEFI**

Procedure

1. Start the Altiris Deployment Solution console.
2. Navigate to the **IBM Scripting Toolkit, Linux Edition 9.40 > Modular Deployment Tools > Step 2 - Operating System Installation > Scripted Operating System Installation** folder.
3. Select the job that corresponds to the operating system you want to install and run it against the target server.

Note: Only SLES 11 x64 and RHEL x64 are supported for native uEFI installation. The *TK_DISTRO_CFG* job variable determines whether the installation is a native uEFI installation or a legacy installation. The Linux Scripting Toolkit provides two sample answer files for native uEFI installations:

- *sles11x64_efi.xml* for native uEFI installation of SLES 11 x64.
- *rhel6x64_efi.xml* for native uEFI installation of RHEL 6 x64.

If you want to customize other answer files for native uEFI installation, ensure that the file contains an entry for */boot/efi*. This entry can be commented out, but must remain visible in the file.

Post installation tasks

When you have completed RAID configuration and operating system installation, you can perform post installation tasks. This section describes the methods for installing the IBM Systems Director Agent and applying UpdateXpress System Pack (UXSP) updates.

The Linux Scripting Toolkit provides sample jobs for installing the IBM Systems Director Agent and applying UpdateXpress System Pack updates to a target system on which RAID has been configured and a supported operating system installed. Please note that in order to complete these jobs, the Altiris agent must be installed on the target system. The Altiris agent is installed automatically by the operating system installation tasks provided by the Linux Scripting Toolkit. If the operating system on the target system was not installed using these jobs, you might have to install the Altiris agent manually. See “Operating system installation” on page 18 for more information on using the modular deployment tools to install operating systems.

The following sections describe the process for using the modular deployment tools to install the IBM Director Agent and apply UXSP updates to a target system.

Installing the IBM Systems Director Agent

After you have configured RAID and installed a supported operating system, use the Linux Scripting Toolkit to install the IBM Systems Director Agent.

Procedure

1. Ensure that the IBM Systems Director Agent source files have been added to the source tree on the source server. For more information, see “Adding Systems Director Agent files” on page 10.
2. Open the Altiris Deployment Solution console and navigate to the **IBM Scripting Toolkit, Linux Edition 9.40 > Modular Deployment Tools > Step 3 - Post-Installation Configuration > Application Installation** folder.
3. Run the **IBM Director Agent for Linux nstall** job against the target system.

Installing UpdateXpress System Pack updates

You can ensure that your target server has current firmware fixes by using the Linux Scripting Toolkit to install UpdateXpress System Pack (UXSP) updates.

Before you begin

Before you can install UXSPs, you must have added the latest updates for your target server machine type to the source server. See “Adding UpdateXpress System Packs” on page 9 for more information.

About this task

Follow these steps to apply UXSPs to your target system:

Procedure

1. Open the Altiris Deployment Solution console.
2. Navigate to the **IBM Scripting Toolkit, Linux Edition 9.40 > Modular Deployment Tools > Step 3 - Post-Installation Configuration > IBM UpdateXpress System Packs** folder.
3. Run the **Run UpdateXpress System Pack Installer** job against the target server.

Deployment solutions

The sample jobs provided in the **Combined Deployment Solutions** folder are designed to consolidate multiple deployment elements in a single Altiris job. This section describes how to use these jobs to perform complete deployments.

The deployment solutions provided by the Linux Scripting Toolkit are designed to perform a combination of deployment steps in a single job. The following sections will provide information on using the sample jobs for:

- Image based deployment (image cloning)
- Scripted deployment

Image based deployment

The Linux Scripting Toolkit provides a set of tools to capture and deploy an existing system configuration, including RAID configuration, operating system, and installed applications and updates. These jobs require the Altiris agent to be installed on the donor system. The Altiris agent is installed automatically as part of the Scripting Toolkit operating system installation jobs. If the donor system operating system was not installed using the Scripting Toolkit, you might need to install the Altiris agent manually.

Follow these steps to capture and deploy a system configuration:

1. Open the Altiris Deployment Solution console and navigate to the **IBM Scripting Toolkit, Linux Edition 9.40 > Combined Deployment Solutions > Image-based Deployment** folder.
2. Run the **Capture RAID Configuration and Linux Image** job against the donor system.
3. When that job completes, run the **Deploy RAID Configuration and Linux Image** against the target system.

Scripted deployment

The scripted deployment sample jobs provided by the ServerGuide Scripting Toolkit are designed to integrate all of the steps in deploying a system configuration in a single Altiris job. To use a scripted deployment job to configure RAID, install an operating system, install the IBM Director Agent, and install UpdateXpress System Pack updates on the target system, follow these steps:

1. Ensure that the necessary operating system, Director Agent, and UXSP files are properly included in the source tree. See Chapter 3, “Adding and sharing content on the Linux Scripting Toolkit source server,” on page 9 for more information.

2. Open the Altiris Deployment Solution console and navigate to the **IBM Scripting Toolkit, Linux Edition 9.40 > Combined Deployment Solutions > Scripted Deployment** folder.
3. In this folder, select the job to deploy the operating system of your choice. Modify the following job variables:

TK_DirAgent_DirectorAgent=dirlin

To install the IBM Systems Director Agent. Uncomment this variable and assign the folder under /sgdeploy/apps/ where the IBM Systems Director Agent files are located.

TK_UXSP_UpdateXpressSystemPacks=uxsp

To install UpdateXpress System Pack updates.

See Chapter 5, “Customizing deployment scenarios,” on page 23 for more information on configuring job variables.

4. Run the job against the target server.

Chapter 5. Customizing deployment scenarios

This section provides information on customizing deployment scenarios.

You can customize your deployment scenarios in the following ways:

- Customize source server settings
- Add PRAID policy files
- Add ASU files
- Customize Fibre HBA boot configuration
- Customize your Linux installation
- Add installation of the IBM Systems Director Agent to your deployment
- Add installation of UpdateXpress System Packs to your deployment

Customizing sample jobs

This section provides information on customizing the sample Altiris jobs provided with the ServerGuide Scripting Toolkit, Linux Edition.

Each sample job provided within the Altiris Deployment Solution console that contains customizable settings begins with a task that sets all variables for the job. This task is labeled **Customize Job Variables**. The variables that you can modify are contained within a block that is similar to the following:

```
[User_Customizable_Variables]
```

```
Variables
```

```
[End_User_Customizable_Variables]
```

You cannot modify variables outside this block. When you modify a job, it is a good practice to create a copy of the job and modify the copy, rather than modifying the original job.

Table 1 provides a list of the variables in the sample Altiris jobs that you can customize. Note that all values for variables must be enclosed in quotation marks.

Table 1. Customizable variables in Altiris sample jobs

Variable	Description
[PRAID Settings]	
TK_pRAID_File	Specifies the pRAID policy file to be used for the job. Default: Varies per job, see next section.
TK_pRAID_Flags	Specifies any additional flags to be passed to the pRAID utility. Default: -r
[Fibre Settings]	

Table 1. Customizable variables in Altiris sample jobs (continued)

Variable	Description
TK_FIBRE_COUNT	<p>Specifies the number of HBA ports to configure.</p> <p>Valid values are 1–<i>n</i>, where <i>n</i> is the number of HBA ports available.</p> <p>This variable affects the use of the following variables:</p> <ul style="list-style-type: none"> • TK_FIBRE_N_HBA_ID • TK_FIBRE_N_BOOT_DISABLE • TK_FIBRE_N_BOOT_PRIM • TK_FIBRE_N_BOOT_ALT1 • TK_FIBRE_N_BOOT_ALT2 • TK_FIBRE_N_BOOT_ALT3 <p>Where N is the HBA number to be configured.</p> <p>Note: You must complete one of each of these variables for every HBA port you configure. So if TK_FIBRE_COUNT=2, you must complete one set of these variables for the first port and one for the second.</p>

Table 1. Customizable variables in Altiris sample jobs (continued)

Variable	Description
TK_FIBRE_N_HBA_ID	<p>Identifies the Qlogic/Emulex HBA to be configured, where <i>N</i> is the HBA number to be configured.</p> <p>Valid values are:</p> <p><i>hba_instance</i></p> <p>The instance number of an HBA port. Valid values are integers from 0 to <i>n-1</i>, where <i>n</i> is the number of HBAs in the system.</p> <p>For example, to configure HBA instance 0:TK_FIBRE_1_HBA_ID=0</p> <p><i>hba_wwpn</i></p> <p>the World Wide Port Name of an HBA port, in the format xx-xx-xx-xx-xx-xx-xx-xx or xxxxxxxxxxxxxxxx.</p> <p>For example, to configure HBA: 90-87-AA-BB-65-34-BB-E0:</p> <p>TK_FIBRE_1_HBA_ID=90-87-AA-BB-65-34-BB-E0</p> <p>Default: 0</p>
	<p>Identifies the Brocade HBA to be configured, where <i>N</i> is the HBA number to be configured.</p> <p>Valid values are:</p> <p><i>hba_instance</i></p> <p>the instance number of an HBA port. Valid format is <i>N/P</i>, where <i>N</i> is the adapter number from 1 to <i>N</i>, and <i>P</i> is the port number from 0 to <i>p-1</i>.</p> <p>For example, to configure HBA instance 0: TK_FIBRE_1_HBA_ID=1/0</p> <p><i>hba_wwpn</i></p> <p>the World Wide Port Name of an HBA port, in the format xx-xx-xx-xx-xx-xx-xx-xx or xxxxxxxxxxxxxxxx.</p> <p>For example, to configure HBA: 90-87-AA-BB-65-34-BB-E0:</p> <p>TK_FIBRE_1_HBA_ID=90-87-AA-BB-65-34-BB-E0</p> <p>Default: 0</p>

Table 1. Customizable variables in Altiris sample jobs (continued)

Variable	Description
TK_FIBRE_N_BOOT_DISABLE	<p>Disable the selected current boot device settings on the specified HBA port, where <i>N</i> is the HBA number to be configured.</p> <p>Valid values are</p> <p>No Does not clear or disable any boot settings.</p> <p>All Disables the primary and all alternate boot settings - Prim, Alt1, Alt2, and Alt3.</p> <p>Prim Disables only the primary boot setting.</p> <p>Alt1 Disables the Alternative 1 boot setting.</p> <p>Alt2 Disables the Alternative 2 boot setting.</p> <p>Alt3 Disables the Alternative 3 boot setting.</p> <p>Default: No.</p>
TK_FIBRE_N_BOOT_PRIM = <i>target_wwnn target_wwpn lun_id</i>	<p>Defines the primary boot target settings, where <i>N</i> is the HBA number to be configured, and:</p> <ul style="list-style-type: none"> <i>target_wwnn</i> - is the World Wide Node Name of a device, in the format xx-xx-xx-xx-xx-xx-xx-xx or xxxxxxxxxxxxxxxx. <i>target_wwpn</i> - is the World Wide Port Name of a device, in the format xx-xx-xx-xx-xx-xx-xx-xx or xxxxxxxxxxxxxxxx. <i>lun_id</i> - is the Logical Unit Number of a device. <p>Default: 0 0 0</p> <p>Example:</p> <p>TK_FIBRE_1_BOOT_PRIM= BB-CC-AA-BB-65-34-BB-F1 BB-CC-AA-BB-FF-34-BB-F1 9</p>
TK_FIBRE_N_BOOT_ALT1 = <i>target_wwnn target_wwpn lun_id</i>	<p>Configures the operating system to use the indicated target as the first alternate boot device, where <i>N</i> is the HBA number to be configured, and</p> <ul style="list-style-type: none"> <i>target_wwnn</i> - is the World Wide Node Name of a device, in the format xx-xx-xx-xx-xx-xx-xx-xx or xxxxxxxxxxxxxxxx. <i>target_wwpn</i> - is the World Wide Port Name of a device, in the format xx-xx-xx-xx-xx-xx-xx-xx or xxxxxxxxxxxxxxxx. <i>lun_id</i> - is the Logical Unit Number of a device. <p>Default: blank.</p> <p>Example:</p> <p>TK_FIBRE_1_BOOT_ALT1= BB-CC-AA-BB-65-34-BB-FD BB-CC-AA-BB-FF-40-BB-F1 5</p>

Table 1. Customizable variables in Altiris sample jobs (continued)

Variable	Description
TK_FIBRE_N_BOOT_ALT2 = <i>target_wwnn target_wwpn lun_id</i>	<p>Configures the operating system to use the indicated target as the second alternate boot device, where <i>N</i> is the HBA number to be configured, and</p> <ul style="list-style-type: none"> <i>target_wwnn</i> - is the World Wide Node Name of a device, in the format xx-xx-xx-xx-xx-xx-xx-xx or xxxxxxxxxxxxxxxx. <i>target_wwpn</i> - is the World Wide Port Name of a device, in the format xx-xx-xx-xx-xx-xx-xx-xx or xxxxxxxxxxxxxxxx. <i>lun_id</i> - is the Logical Unit Number of a device. <p>Default: blank.</p> <p>Example:</p> <p>TK_FIBRE_1_BOOT_ALT2= BB-CC-AA-BB-65-34-BB-FD BB-CC-AA-BB-FF-40-BB-F1 5</p>
TK_FIBRE_N_BOOT_ALT3 = <i>target_wwnn target_wwpn lun_id</i>	<p>Configures the operating system to use the indicated target as the third alternate boot device, where <i>N</i> is the HBA number to be configured, and</p> <ul style="list-style-type: none"> <i>target_wwnn</i> - is the World Wide Node Name of a device, in the format xx-xx-xx-xx-xx-xx-xx-xx or xxxxxxxxxxxxxxxx. <i>target_wwpn</i> - is the World Wide Port Name of a device, in the format xx-xx-xx-xx-xx-xx-xx-xx or xxxxxxxxxxxxxxxx. <i>lun_id</i> - is the Logical Unit Number of a device. <p>Default: blank.</p> <p>Example:</p> <p>TK_FIBRE_1_BOOT_ALT3= BB-CC-AA-BB-65-34-BB-FD BB-CC-AA-BB-FF-40-BB-F1 5</p>
[ASU Settings]	
TK_ASU_File	<p>Specifies the file to be used for a Capture System Settings or Deploy System Settings operation. If you are capturing system settings, they will be written to a file with this name. If you are deploying system settings, they will be read from this file.</p> <p>Default: /sgdeploy/sgtklinux/ini/asu/ \$(BIOS_Build_Level0:2)/asu_capture</p>
TK_ASU_Mode	<p>Specifies the ASU mode being used. This value is used with the asu command and the value of the TK_ASU_File variable to create a complete ASU command.</p> <p>Valid values are: save, replicate, and batch.</p>
TK_ASU_Flags	<p>Specifies additional flags for the current mode.</p> <p>The default value is -group-bios.</p>
[NOS Installation Settings]	

Table 1. Customizable variables in Altiris sample jobs (continued)

Variable	Description
TK_DISTRO_TYPE	Specifies the Linux distribution to be installed. Valid values are: RHEL, SLES, VMWARE
TK_DISTRO_VERSION	Specifies the version of the operating system indicated by TK_DISTRO_TYPE. Valid values correspond to the supported versions of the supported operating systems. Note: The version number must be a valid and supported version of the Linux distribution specified by TK_DISTRO_TYPE.
TK_DISTRO_PATH	Specifies the fully qualified URL for the location of the network share where the operating system files are located. Supported protocols are FTP, HTTP, and NFS. Default: ftp://\$TK_Altiris_IP/sgdeploy/os/osfolder
TK_DISTRO_CFG	Specifies the unattended answer file to be used for the operating system installation. This file must be formatted as a kickstart or autoyast file, depending on the distribution being deployed. Default: ./sgdeploy/sgtklinux/altiris/answerfiles/answerfile
TK_DISTRO_KPARAM	Specifies additional kernel parameters to be used for the deployment, such as VNC parameters or a static IP address. Default: none
TK_DISTRO_NIC_PORT	Specifies a Network Interface Card (NIC) to be used for the installation, if the NIC is different than the one used for the pre-installation execution environment (PXE) boot. If no NIC is specified, the NIC used for PXE boot will be used for the installation, using the bootif method. Note: This variable is not valid for VMware ESX Server. To indicate the network device for VMware ESX Server, edit the network section of the VMware kickstart file.
[Director Agent Settings]	
TK_DirAgent_DirectorAgent	Specifies the location of the IBM Director Agent application files on the source server. Valid values are: <ul style="list-style-type: none"> • dirlin • A user-supplied directory in the sgdeploy/apps directory Default: dirlin
[UXSP Settings]	

Table 1. Customizable variables in Altiris sample jobs (continued)

Variable	Description
TK_UXSP_UpdateXpressSystemPacks	<p>Specifies the location of the UpdateXpress System Packs on the source server.</p> <p>Valid values are:</p> <ul style="list-style-type: none"> • uxsp • a user-specified directory in the sgdeploy/updates directory <p>Default: uxsp</p>

Sample job definitions

This section provides descriptions of the sample Altiris jobs provided by the Linux Scripting Toolkit, including a description of the variables that can be customized for each job.

All of the sample jobs except imaging jobs without RAID configuration and the Linux Pre-Boot test add logging information to the ToolkitProcess.log file.

Customizing Fibre HBA boot configuration

You can use Toolkit variables to customize the configuration of Fibre HBAs on the target system, allowing them to boot from SAN targets.

By default, the ServerGuide Scripting Toolkit will configure the first HBA on the system to boot from the first available SAN target (for QLogic Fibre HBAs only. For information Emulex Fibre HBAs, see “Known problems and limitations” on page 83). The BIOS configures the first disk drive that it finds that is also a LUN 0 as a boot device. The ServerGuide Scripting Toolkit uses the following variables to configure Fibre HBAs:

Note: Please note that while some examples are broken across multiple lines for formatting reasons, when using these settings, you must include all the information for each variable on a single line.

Table 2. Fibre HBA boot configuration variables

Variable	Description
TK_FIBRE_COUNT	<p>Specifies the number of HBA ports to configure.</p> <p>Valid values are 1–<i>n</i>, where <i>n</i> is the number of HBA ports available.</p> <p>This variable affects the use of the following variables:</p> <ul style="list-style-type: none"> • TK_FIBRE_N_HBA_ID • TK_FIBRE_N_BOOT_DISABLE • TK_FIBRE_N_BOOT_PRIM • TK_FIBRE_N_BOOT_ALT1 • TK_FIBRE_N_BOOT_ALT2 • TK_FIBRE_N_BOOT_ALT3 <p>Where N is the HBA number to be configured.</p> <p>Note: You must complete one of each of these variables for every HBA port you configure. So if TK_FIBRE_COUNT=2, you must complete one set of these variables for the first port and one for the second.</p>

Table 2. Fibre HBA boot configuration variables (continued)

Variable	Description
TK_FIBRE_N_HBA_ID	<p>Identifies the Qlogic/Emulex HBA to be configured, where <i>N</i> is the HBA number to be configured.</p> <p>Valid values are:</p> <p><i>hba_instance</i></p> <p>The instance number of an HBA port. Valid values are integers from 0 to <i>n-1</i>, where <i>n</i> is the number of HBAs in the system.</p> <p>For example, to configure HBA instance 0:TK_FIBRE_1_HBA_ID=0</p> <p><i>hba_wwpn</i></p> <p>the World Wide Port Name of an HBA port, in the format xx-xx-xx-xx-xx-xx-xx-xx or xxxxxxxxxxxxxxxxx.</p> <p>For example, to configure HBA: 90-87-AA-BB-65-34-BB-E0:</p> <p>TK_FIBRE_1_HBA_ID=90-87-AA-BB-65-34-BB-E0</p> <p>Default: 0</p>
	<p>Identifies the Brocade HBA to be configured, where <i>N</i> is the HBA number to be configured.</p> <p>Valid values are:</p> <p><i>hba_instance</i></p> <p>the instance number of an HBA port. Valid format is <i>N/P</i>, where <i>N</i> is the adapter number from 1 to <i>N</i>, and <i>P</i> is the port number from 0 to <i>p-1</i>.</p> <p>For example, to configure HBA instance 0: TK_FIBRE_1_HBA_ID=1/0</p> <p><i>hba_wwpn</i></p> <p>the World Wide Port Name of an HBA port, in the format xx-xx-xx-xx-xx-xx-xx-xx or xxxxxxxxxxxxxxxxx.</p> <p>For example, to configure HBA: 90-87-AA-BB-65-34-BB-E0:</p> <p>TK_FIBRE_1_HBA_ID=90-87-AA-BB-65-34-BB-E0</p> <p>Default: 0</p>

Table 2. Fibre HBA boot configuration variables (continued)

Variable	Description
TK_FIBRE_N_BOOT_DISABLE	<p>Disable the selected current boot device settings on the specified HBA port, where <i>N</i> is the HBA number to be configured.</p> <p>Valid values are</p> <p>No Does not clear or disable any boot settings.</p> <p>All Disables the primary and all alternate boot settings - Prim, Alt1, Alt2, and Alt3.</p> <p>Prim Disables only the primary boot setting.</p> <p>Alt1 Disables the Alternative 1 boot setting.</p> <p>Alt2 Disables the Alternative 2 boot setting.</p> <p>Alt3 Disables the Alternative 3 boot setting.</p> <p>Default: No.</p>
TK_FIBRE_N_BOOT_PRIM = <i>target_wwnn target_wwpn lun_id</i>	<p>Defines the primary boot target settings, where <i>N</i> is the HBA number to be configured, and:</p> <ul style="list-style-type: none"> <i>target_wwnn</i> - is the World Wide Node Name of a device, in the format xx-xx-xx-xx-xx-xx-xx-xx or xxxxxxxxxxxxxxxx. <i>target_wwpn</i> - is the World Wide Port Name of a device, in the format xx-xx-xx-xx-xx-xx-xx-xx or xxxxxxxxxxxxxxxx. <i>lun_id</i> - is the Logical Unit Number of a device. <p>Default: 0 0 0</p> <p>Example:</p> <p>TK_FIBRE_1_BOOT_PRIM= BB-CC-AA-BB-65-34-BB-F1 BB-CC-AA-BB-FF-34-BB-F1 9</p>
TK_FIBRE_N_BOOT_ALT1 = <i>target_wwnn target_wwpn lun_id</i>	<p>Configures the operating system to use the indicated target as the first alternate boot device, where <i>N</i> is the HBA number to be configured, and</p> <ul style="list-style-type: none"> <i>target_wwnn</i> - is the World Wide Node Name of a device, in the format xx-xx-xx-xx-xx-xx-xx-xx or xxxxxxxxxxxxxxxx. <i>target_wwpn</i> - is the World Wide Port Name of a device, in the format xx-xx-xx-xx-xx-xx-xx-xx or xxxxxxxxxxxxxxxx. <i>lun_id</i> - is the Logical Unit Number of a device. <p>Default: blank.</p> <p>Example:</p> <p>TK_FIBRE_1_BOOT_ALT1= BB-CC-AA-BB-65-34-BB-FD BB-CC-AA-BB-FF-40-BB-F1 5</p>

Table 2. Fibre HBA boot configuration variables (continued)

Variable	Description
TK_FIBRE_N_BOOT_ALT2 = <i>target_wwnn target_wwpn lun_id</i>	<p>Configures the operating system to use the indicated target as the second alternate boot device, where <i>N</i> is the HBA number to be configured, and</p> <ul style="list-style-type: none"> • <i>target_wwnn</i> - is the World Wide Node Name of a device, in the format xx-xx-xx-xx-xx-xx-xx-xx or xxxxxxxxxxxxxx. • <i>target_wwpn</i> -is the World Wide Port Name of a device, in the format xx-xx-xx-xx-xx-xx-xx-xx or xxxxxxxxxxxxxx. • <i>lun_id</i> - is the Logical Unit Number of a device. <p>Default: blank.</p> <p>Example:</p> <p>TK_FIBRE_1_BOOT_ALT2= BB-CC-AA-BB-65-34-BB-FD BB-CC-AA-BB-FF-40-BB-F1 5</p>
TK_FIBRE_N_BOOT_ALT3 = <i>target_wwnn target_wwpn lun_id</i>	<p>Configures the operating system to use the indicated target as the third alternate boot device, where <i>N</i> is the HBA number to be configured, and</p> <ul style="list-style-type: none"> • <i>target_wwnn</i> - is the World Wide Node Name of a device, in the format xx-xx-xx-xx-xx-xx-xx-xx or xxxxxxxxxxxxxx. • <i>target_wwpn</i> -is the World Wide Port Name of a device, in the format xx-xx-xx-xx-xx-xx-xx-xx or xxxxxxxxxxxxxx. • <i>lun_id</i> - is the Logical Unit Number of a device. <p>Default: blank.</p> <p>Example:</p> <p>TK_FIBRE_1_BOOT_ALT3= BB-CC-AA-BB-65-34-BB-FD BB-CC-AA-BB-FF-40-BB-F1 5</p>

Capture and Deploy RAID Configuration and Linux image

The ServerGuide Scripting Toolkit, Linux Edition provides sample Altiris jobs to perform operating system imaging tasks using Altiris. Imaging tasks include capturing a copy of an installed operating system and deploying that image to a target server. You might need to configure the target server before you deploy the image.

Note: The donor and target systems must be configured identically in order to deploy the captured RAID configuration and Linux image.

These jobs capture a RAID Configuration along with the OS image, so that you can use both to redeploy the image. In order to run the capture RAID configuration and Linux image sample job, you must first install the Altiris agent on the donor system. The Toolkit's scripted OS install jobs automatically install the Altiris agent.

The capture sample job captures the RAID configuration of the server, and then captures an OS image from the target server and stores both on the Altiris server.

Each deployment sample job deploys the captured RAID configuration, and OS image to the target server. The sample job configures RAID, calls a reboot, and then performs the image deployment.

The following variables may be customized for these jobs within the Altiris Console:

- TK_pRAID_File="/sgdeploy/sgtklinux/ini/praid/\\$Machine_Type/raid_capture.ini"

Configure RAID + Install Linux OS

This sample job performs a RAID configuration, restarts the system, and then performs a scripted operating system installation. You must set up the operating system being deployed in the Linux Scripting Toolkit source tree on the Altiris server. See "Adding operating system installation files" on page 10 for details.

The following variables may be customized for this job within the .Altiris Deployment Solution console. The values shown are the defaults:

- TK_pRAID_File="/sgdeploy/sgtklinux/ini/praid/default_raid.ini"
- TK_pRAID_Flags="-r"
- TK_Altiris_IP
- TK_DISTRO_TYPE
- TK_DISTRO_VERSION
- TK_DISTRO_PATH=ftp://\$TK_Altiris_IP/sgdeploy/os/rhel5x64
- TK_DISTRO_CFG
- TK_DirAgent_DirectorAgent=dirlin
- TK_UXSP_UpdateXpressSystemPacks=uxsp
- TK_DISTRO_KPARAM
- TK_DISTRO_NIC_PORT="link"

Configuring RAID with a policy file

You can configure RAID controllers using standalone jobs that only configure the RAID controllers and then stop, and jobs integrated with scripted and image-based deployment. The jobs that specify a specific RAID level in them use a pre-configured policy file to create a RAID array of that type on the target machine. The other jobs either deploy a previously captured RAID configuration or create the default RAID configuration if no policy file is specified.

Follow these steps to use a user defined policy file:

1. Create the policy file in the sgdeploy/sgtklinux/ini/praid directory.
2. Copy a pre-existing Toolkit job with RAID configuration.
3. Modify the TK_pRAID_File variable in the copied job to point to your policy file. In this example: TK_pRAID_File="/sgdeploy/sgtklinux/ini/praid/UserPolicy1.ini".

Capture and deploy RAID configuration

The Linux Scripting Toolkit provides sample Altiris jobs to capture and deploy a RAID Configuration, allowing the RAID configuration of a single server to provide a configuration basis for any number of servers. The capture sample job captures the RAID configuration of the server and stores it in a file named RAID_Configuration.ini.

By default, all captured RAID configurations are stored in the following directory on the Deployment Server:

AltirisPath\sgdeploy\sgtklinux\ini\praid\machinetype

where:

- *AltirisPath* is the fully qualified path to the Altiris Deployment Solution directory, for example C:\Program Files\Altiris\eXpress\Deployment Server.
- *machinetype* is the machine type of the captured system.

The deployment sample job deploys the captured RAID configuration from RAID_Configuration.ini located in the above directory by default. You can customize the following variables for these jobs within the Altiris Console:

- TK_pRAID_File=/sgdeploy/sgtklinux/ini/praid/*machinetype*/raid_capture.ini
- TK_pRAID_Flags="-r"

Server data disposal

These sample jobs perform a server data disposal. The first performs a disposal and resets the RAID Configuration. The second performs only a server data disposal.

The following variables may be customized for these jobs within the Altiris Console:

Variable	Values
TK_Wipe_Level	<ul style="list-style-type: none">• "quick" - performs a quick wipe of the disks.• "dod" - performs a multipass wipe of the disks that conforms to DOD standards.
TK_Wipe_Repeat_Number	<i>number</i> - indicates the number of passes to complete.
TK_Wipe_Disk	<i>disk_number</i> - indicates the number of the disk to wipe. By default, "1" is the first disk in the system.

Note: Depending on the size of the drives involved and the wipe level, these jobs can take up to several hours to complete.

Linux and VMware unattended scripted installation

These sample jobs perform scripted installation of the supported Linux distributions. You must setup the distribution to be installed on the Altiris server before running these jobs. See "Adding operating system installation files" on page 10 for more information.

You can customize the following variables for these jobs:

- TK_Altiris_IP
- TK_DISTRO_TYPE
- TK_DISTRO_VERSION
- TK_DISTRO_PATH
- TK_DISTRO_CFG
- TK_DirAgent_DirectorAgent

- TK_UXSP_UpdateXpressSystemPacks

Capture and deploy a Linux image

The Linux Scripting Toolkit provides sample jobs to perform OS imaging tasks using Altiris. Operating system imaging tasks include capturing a copy of an installed OS and deploying that image to a target server. You might need to perform RAID configuration on the target server before you can deploy the OS image to it. In order to run the **Capture Linux Image** sample job, the Altiris agent must first be installed on the donor system. The Toolkit's scripted OS install jobs automatically install the Altiris agent.

Note: The **Capture Linux Image** job is not supported for VMware ESX Server.

To modify the filename and location of the captured image, perform the following steps:

1. Double-Click on the **Capture Linux Image** job.
2. Select the **Create Disk Image** task and click **Modify**.
3. Change the path and filename for the captured image.
4. Click **Finish** to save.
5. Double-click on the **Deploy Linux Image** job.
6. Select the **Distribute Disk Image** task and click **Modify**.
7. Change the path and filename to the same values used above for the captured image.
8. Click **Finish** to save the changes.

Note:

- When you deploy a target server using the Image Install job, the image will include the Altiris Agent, because the donor server is required to have the Altiris Agent installed.
- Because the Linux imaging sample jobs use a sector-based imaging technique, the donor and target servers must have identical hardware.
- Altiris RapiDeploy, which is used for the cloning jobs, does not support cloning of multiple disks in a single pass. To work around this issue, you can place all of the Linux partitions on a single disk during installation by modifying the sample kickstart files to include the **-ondisk** parameter in the partition configuration *part* lines. Refer to the operating system documentation for additional information.
- To specify the disk to be cloned, you must add the **-d** parameter, specifying the disk number, to the command line options for the Capture Linux Image and Deploy Linux Image jobs. For example, to clone the operating system image on disk two, you would add: **-d2** to the command line options for the job.
- If the capture or deploy fails, please refer the rdeploy log file for more information. By default, the log file location is `$AltirisInstalledPath\RDeploy\Linux\x64\rdlog\` where `$AltirisInstalledPath` is the fully qualified path to the Altiris Deployment Solution directory. For example, `C:\Program Files\Altiris\Express\Deployment Server\RDeploy\Linux\x64\rdlog\`.

Linux Pre-Boot Test

This sample job performs a basic task to test connectivity between the target server and the Altiris Deployment Server. The job boots into the IBM Linux PE x64 environment on the target server.

Chapter 6. Supported hardware and software

This section lists the operating systems, adapters, and RAID controllers supported by the Linux Scripting Toolkit, as well as systems that support BIOS and firmware updates using the ASU.

The Linux Scripting Toolkit supports deployment of Linux operating systems on IBM System x[®] and BladeCenter[®] servers. In general, the Linux Scripting Toolkit provides support for ServerProven[®] IBM or third-party adapters in the following categories:

- Ethernet
- Fibre Channel
- IDE and IDE RAID
- SAS and SAS RAID
- SATA and SATA RAID
- SCSI and SCSI RAID, including Ultra-SCSI

This section contains the following information about specific hardware and software support for deployment scenarios:

- Supported operating system and server combinations
- RAID and Fibre channel HBA support by server
- Network device driver support by server
- Limitations of support for applicable servers

The most up-to-date support information is contained in the `readme.htm` file. You can download the latest version of `readme.htm` file from the ServerGuide Scripting Toolkit Web page. See IBM deployment resources on the World Wide Web for information.

Operating system support

This section lists operating system deployment/server combinations supported by the Linux Scripting Toolkit.

You can use the Linux Scripting Toolkit to deploy supported Linux distributions to any IBM System x, BladeCenter, or iDataPlex server that supports that distribution. To determine what distribution/server combinations are supported, see IBM ServerProven.

The Linux Scripting Toolkit supports these Linux distributions:

- SUSE Linux Enterprise Server 9 32 bit SP4
- SUSE Linux Enterprise Server 9 x64 SP4
- SUSE Linux Enterprise Server 10 32 bit SP1/SP2/SP3/SP4
- SUSE Linux Enterprise Server 10 x64 SP1/SP2/SP3/SP4
- SUSE Linux Enterprise Server 11 32 bit Base/SP1/SP2/SP3
- SUSE Linux Enterprise Server 11 x64 Base/SP1/SP2/SP3
- Red Hat Enterprise Linux 4 AS/ES 32 bit U6/U7/U8
- Red Hat Enterprise Linux 4 AS/ES x64 U6/U7/U8

- Red Hat Enterprise Linux 5 32 bit U1/U2/U3/U4/U5/U6/U7/U8/U9
- Red Hat Enterprise Linux 5 x64 U1/U2/U3/U4/U5/U6/U7/U8/U9
- Red Hat Enterprise Linux 6 32 bit U1/U2/U3/U4
- Red Hat Enterprise Linux 6 x64 U1/U2/U3/U4
- VMware ESX Server 3.5 U4/U5
- VMware ESX Server 4.0/4.0u1/4.0u2/4.1/4.1u1/4.1u2/4.1u3/4.1u4

RAID controller support

You can use the Linux Scripting Toolkit to configure any RAID controller supported by the IBM System x, Blade Center, iDataPlex, or PureFlex server in which it is installed. For information on supported RAID controller/server combinations, see Storage Controllers and IBM ServerProven.

Fibre Channel HBA support

You can use the Linux Scripting Toolkit to configure any FC HBA supported by the IBM System x, Blade Center, iDataPlex, or PureFlex server in which it is installed.

For information on supported FC HBA/server combinations, see Shared Storage Adapters and IBM ServerProven.

Chapter 7. Linux Scripting Toolkit utilities and tools

This section contains information about the utilities that are included in the Linux Scripting Toolkit and the tools that are shipped with it. For each utility there is a description of parameters, along with examples.

For each included tool there is a brief description of the tool and instructions on using it with the Linux Scripting Toolkit, as well as pointers on where to get more information on the tool and its use.

altiris_sgtklinux.cmd

Purpose

This script builds the IBM-specific Linux Pre-installation Environment for Altiris Deployment Solution and integrates the Toolkit into Altiris Deployment Solution

Note:

- This script must be executed from the directory in which it resides.
- The Altiris Deployment Solution and Linux Scripting Toolkit must be installed in a path that does not contain parentheses.

altiris_sgtklinux.cmd

Parameters

There are no parameters.

Sample

Altiris_SGTKLinux

acquire_seps.cmd

This script acquires the System Enablement Packs (SEP) for the specified machine types given and copies the packs to a location where Linux Toolkit can use them during deployments.

acquire_seps.cmd

Parameters

machine_type

Specifies the machine type for which to acquire SEPs. You can specify multiple machine types by separating them with commas.

All Acquires all available SEPs.

Sample

acquire_seps.cmd 7979

acquire_seps.cmd 7979,4199

acquire_seps.cmd all

Linux Scripting Toolkit utilities

This section contains information about the utilities that are included in the Linux Scripting Toolkit. For each utility there is a description of the parameters with examples. These utilities are included in the Linux Scripting Toolkit directory structure. The default location is C:\Program Files\Altiris\eXpress\Deployment Server\sgdeploy\sgtklinux\tk\bin.

The command-line syntax examples in this documentation use the following conventions:

- Variables are shown in *italics*
- Required parameters are shown within <> brackets
- Optional parameters are shown within [] brackets
- Required or optional parameters from which you must make a unique choice are separated by a vertical bar (|) character

You must enter all parameters for a utility on a single command line even when the information in this documentation is shown on multiple lines.

HWDETECT

HWDETECT is used to perform basic hardware detection functions that are typically obtained using SMBIOS and a PCI scan. This utility contains options that can be used to dump all of the hardware information to an output file, or it can be used to query hardware information and return values that set the *errorlevel* environment variable or the return code, for example \$?.

HWDETECT has basic hardware scan functions and more complex PCI device detection options. The basic hardware scan functions can only be used singularly. The PCI device detection functions can be used with each other to produce a query based on multiple restrictions. You can only use the hwdetect utility basic hardware scan functions one at a time. The PCI-device detection functions can be combined or used more than once on the same command line.

Usage:

hwdetect [-s|-i|-p|--m=machinetype] -f=filename

Parameter	Description	Example
-s	Determines if the target server is an IBM System x, xSeries®, or BladeCenter server. The return values are: <ul style="list-style-type: none">• 0 for an IBM system• 1 for a non-IBM system	<pre>./hwdetect -s if [\$? -eq 1]; then echo "Perform non-IBM equipment specific steps here." else echo "Perform IBM equipment specific steps here." fi</pre>
-i	Dumps all available information about the system hardware to the screen in an ini file format. You can use the -f parameter to output this information to a file. A return code of zero indicates success. All other return codes indicate an error.	<pre>./hwdetect -i</pre>

Parameter	Description	Example
<code>-f=filename</code>	Directs the output to the indicated file. This parameter can be used in conjunction with the <code>-i</code> or <code>-p</code> parameters. A return code of 254 indicates that <code>hwdetect</code> was unable to open the specified file.	<pre>./hwdetect -i --f=hwdetect.out cat hwdetect.out grep "Bus_Number.21 = 41"</pre>
<code>--m=machinetype</code>	Compares the machine type of the current system to the specified machine type. Return codes: <ul style="list-style-type: none"> 0 indicates that the machine types do not match. 1 indicates a match 	<pre>./hwdetect --m=8676 if [\$? -eq 8676]; then echo "It is an IBM system." else echo "It is not an IBM system."</pre>

You can also use **HWDETECT** to inventory PCI devices on the target system.

Usage:

```
hwdetect [--vid=vendor_id|--did=device_id|--svid=sub-vendor_id|
--sdid=sub-device_id|--bn=bus_number|--dn=device_number|--add=number]
```

Parameter	Description	Example
<code>--vid=vendor_id</code>	Searches for PCI devices with the indicated hexadecimal vendor ID.	<pre>./hwdetect --vid=40 echo "Found \$? matches..."</pre>
<code>--did=device_id</code>	Searches for PCI devices with the indicated hexadecimal device ID.	<pre>./hwdetect --did=41 echo "Found \$? matches..."</pre>
<code>--svid=sub-vendor_id</code>	Searches for PCI devices with the indicated hexadecimal sub-vendor ID.	<pre>./hwdetect --svid=42 echo "Found \$? matches..."</pre>
<code>--sdid=sub-device_id</code>	Searches for PCI devices with the indicated hexadecimal sub-device ID.	<pre>./hwdetect --sdid=43 echo "Found \$? matches..."</pre>
<code>--bn=bus_number</code>	Starts the search at the indicated decimal bus number.	<pre>./hwdetect --bn=44 echo "Found \$? matches..."</pre>
<code>--dn=device_number</code>	Starts the search at the indicated decimal device number.	<pre>./hwdetect --dn=45 echo "Found \$? matches..."</pre>
<code>--add=number</code>	Adds the specified decimal value to the return value before exiting.	<pre>./hwdetect --vid=46 --add=1 echo "Found \$? - 1 matches..."</pre>

Below is an example of the `hwdetect.out` file created by the `-i` flag:

```
[System]
Machine_Type=8674
Model_Number=42X
Serial_Number=78Z9506
Product_Name=eserver xSeries 330
BIOS_version=1.04
BIOS_Build_Level=EME112A
BIOS_DATE=06/28/2002
BIOS_Manufacturer=IBM
BIOS_Language=US
Number_Of_Enclosures=1
Enclosure_Type.0=23
Processor_Slots=2
Active_Processors=1
Processor_Family.0=17
Processor_Speed_MHz.0=1400
```

```

Processor_X64 = TRUE
Total_Enabled_Memory_Mb=256
ROM_Diagnostics_Build_Level=EME112A
ISMP_Build_Level=BR8T30A
RSA_Build_Level=GEE834A
System_UUID = 8030E01060F010B010605090D0A020F0
Blade_Chassis_UUID = 0F020A0D0900F00F020A0D0900F00F02
Blade_Slot = 02

```

```

[PCI]
Total_Number_Devices=10
Bus_Number.0=0
Device_Number.0=1
Function_Number.0=0
Class_Code.0=0000
Revision.0=0
Header_Type.0=0
Vendor_ID.0=5333
Device_ID.0=8A22
Subvendor_ID.0=1014
Subdevice_ID.0=01C5
Bus_Number.1=0
Device_Number.1=2
Function_Number.1=0
Class_Code.1=0000
Revision.1=0
Header_Type.1=0
Vendor_ID.1=8086
Device_ID.1=1229
Subvendor_ID.1=1014
Subdevice_ID.1=105C

```

Using the `-p` flag produces the same output with the exception that the section names are tacked onto the beginning of each keyword:

```

System_Machine_Type = 8674
System_Model_Number = 42X
System_Serial_Number = 78Z9506
...
PCI_Bus_Number.0 = 0
PCI_Device_Number.0 = 1
...

```

Notes:

1. The BIOS_DATE value is listed in mm/dd/yyyy format.
2. The Enclosure_Type.0=23 is based on SMBIOS 2.3 spec. 23 = Main chassis.
3. There is an entry for Processor_Family and Processor_Speed_MHz for each microprocessor in the server.
4. The ROM_Diagnostics_Build_Level is empty for servers that do not support ROM diagnostics.
5. PCI devices are listed in the order they are scanned.
6. PCI devices are listed in the *Value.n* format, where *Value* is the variable name and *n* is the nth PCI device scanned.
7. The header_type field is not available for versions of hwdetect running on Windows 32 or 64-bit operating systems.
8. The vendor, device, subvendor, and subdevice values are in hexadecimal notation.

SAVESTAT

The Savestat utility allows you to store and retrieve up to twenty values to persistent storage. The utility is designed to help you remember where you left off in an installation script even when a system reboot is required. This utility is designed to return values that set the `?` environment variable so that you can branch in a script (sh) file based on the result of the utility's execution.

Savestat uses the persistent storage capability of **ASU**. Therefore the following files must be available for the script to work:

- ASU package (ibm_utl_asu_asut69*_linux_x86-64.tgz)
- savestat.sh script
- savestat.def

Usage

The savestat utility that comes with the ServerGuide Scripting Toolkit has the following command-line syntax:

```
SAVESTAT [/q] -set1=value [...-set2=value ... -set21=value]
SAVESTAT [/q] -getn
SAVESTAT [/q] -validate
SAVESTAT [/q] -signature
```

Parameter	Description	Usage
-setn=value	Saves an integer value, <i>value</i> , to the <i>n</i> th location in persistent-storage memory, where <i>n</i> is an integer from 1-21. Return codes: <ul style="list-style-type: none">• 0 if successful• 1 if not successful	<code>./savestat.sh -setn=value</code> Where: <ul style="list-style-type: none">• <i>n</i> is an integer from 1–21• <i>value</i> is an integer from 0–254
-getn	Retrieves a value currently set in the <i>n</i> th location in persistent-storage memory. Return codes: <ul style="list-style-type: none">• The value stored at the location specified by <i>n</i>, if successful.• 255 if not successful.	<code>./savestat.sh -getn</code> Where <i>n</i> is the location of a previously-stored value.
-signature	Verifies that the persistent storage contains the savestat signature. Return codes: <ul style="list-style-type: none">• 0 if storage contains the signature• 1 if storage does not contain the signature	<code>./savestat.sh -signature</code>
-validate	Verifies that the system is supported by savestat . Return codes: <ul style="list-style-type: none">• 0 if the system is supported• 1 if the system is not supported	<code>./savestat.sh -validate</code>
-q	Invokes the quiet mode, which suppresses prompting. This parameter is optional and can be used with any other savestat parameter.	<code>./savestat.sh -q -set1=100</code>

Note: The help for **savestat.sh** indicates that the **--reset** parameter is supported. **Savestat.sh** does not currently support the **--reset** parameter. To reset all of the storage locations to zero, use the **savestat.sh --set** command as shown here:

```
savestat.sh --set1=0 --set2=0 --set3=0 --set4=0 --set5=0 --set6=0 --set7=0
--set8=0 --set9=0 --set10=0 --set11=0 --set12=0 --set13=0 --set14=0
--set15=0 --set16=0 --set17=0 --set18=0 --set19=0 --set20=0 --set21=0
```

Examples

The following examples illustrate savestat utility usage.

Example	Description
<code>./savestat.sh -set2=100</code>	Stores the value 100 in the second persistent-storage memory location
<pre>./savestat.sh -get2 if [\$? -eq 100]; then echo "The value 100 was found successfully." else echo "The value 100 was not found." fi</pre>	Retrieves the value of the second persistent-storage memory location and branches in the script file according to the value returned

PRAID

PRAID is a scriptable utility that offers a single user interface for both configuring and replicating all RAID controllers supported by the Linux Scripting Toolkit.

PRAID has three modes of operation:

- **Deploy mode** – for scripted configuration of RAID controllers.
- **Capture mode** – for replicating RAID controller settings.
- **Restore-defaults mode** – for resetting RAID controllers to factory-default settings only.

Deploy mode

Used in Deploy mode, PRAID offers the following features:

- Configures all RAID controllers in a server with a single call to the program.
- Automatically resets all RAID controllers to factory-default settings before configuring.
- Uses customizable logic to decide which configuration (policy) is applied to a server based on system hardware. The logic can involve:
 - Machine type of the server
 - Serial number of the server
 - Number of drives connected to the RAID controller
 - RAID controller type
 - Controller number (order) of the RAID controller
- Can be highly customized for specific RAID configurations or generalized to handle many different RAID configurations.
- Provides a default or AUTO mode for automatically creating arrays and logical drives using default settings. This mode requires no knowledge of the number, size, or location of the drives connected to the RAID controllers.

- Automatically applies default values for any RAID configuration parameters that you do not supply. You supply only the parameters that you want to change.
- Default values for each configuration parameter are equivalent to the default settings of the ServeRAID Manager express configuration method where applicable.
- Allows up to 50 policies for configuring RAID controllers to be specified in a single policies file.

Note:

When using PRAID in Deploy mode, the **-r** parameter is required.

To delete RAID configuration on all controllers, specify **-r**. To delete RAID configuration on a specific controller, specify **-r#** where # is the controller number.

For example, `praid -f:policy.ini -r -y`.

Capture mode

Used in Capture mode, PRAID offers the following features:

- Captures the RAID configurations of all supported controllers to a text file, the policies file, with a common format.
- Captured RAID configurations can be immediately used with PRAID in deploy mode to easily replicate the RAID configuration to many servers.
- Allows customizable logic when saving the captured parameters to determine when each captured configuration must be deployed.
- Saves useful information about each captured configuration, including the system machine type, date, and time when the configuration was captured.
- Allows you to edit any RAID configurations that you capture before deploying them to other systems.

Restore-defaults mode

Used in Restore-defaults mode, PRAID offers the following features:

- Deletes all arrays and logical drives on all RAID controllers.
- Sets other RAID controller settings back to factory defaults.

Environment requirements

The table below provides the supported RAID adapter information by PRAID. PRAID works by parsing the output of other RAID configuration utilities. To accomplish this, the utilities must be in the system search path.

Table 3. Supported RAID adapter information

Adapter	Controller type	Utility
ServeRAID 7t	ServeRAID-7t	arconf
ServeRAID 8i	ServeRAID-8i	
ServeRAID 8k	ServeRAID-8k	
ServeRAID 8k l	ServeRAID-8k-l	
ServeRAID 8s	ServeRAID-8s	
ServeRAID B5015	ServeRAID-B5015	brcli

Table 3. Supported RAID adapter information (continued)

Adapter	Controller type	Utility
LSI SAS 1078 IR	LSI-SAS-1078-IR	cfggen
LSI SAS (1064/1064E/1068/1078)	LSI-SAS-RAID	
LSI SCSI (1020/1030)	LSI-SCSI-RAID	
ServeRAID BR10i	ServeRAID-BR10i	
ServeRAID BR10il	ServeRAID-BR10il	
ServeRAID 7e SATA	ServeRAID-7e-SATA	hrconf
ServeRAID 7e SCSI	ServeRAID-7e-SCSI	
ServeRAID 8e SAS	ServeRAID-8e-SAS	
ServeRAID 8e SATA	ServeRAID-8e-SATA	
ServeRAID 6M	ServeRAID-6M	ipssend
LSI MegaRAID 8480	LSI-MegaRAID-8480	megacli
ServeRAID C105	ServeRAID-C105	
ServeRAID C100	ServeRAID-M100	
ServeRAID C100 R5	ServeRAID-M100-R5	
ServeRAID M1xxx Series	ServeRAID-M1xxx	
ServeRAID M1xxx Series R5	ServeRAID-M1xxx_R5	
ServeRAID M5014	ServeRAID-M5014	
ServeRAID M5014 R6/R60	ServeRAID-M5014-R6-R60	
ServeRAID M5015	ServeRAID-M5015	
ServeRAID M5015 R6/R60	ServeRAID-M5015-R6-R60	
ServeRAID M5025	ServeRAID-M5025	
ServeRAID-M5025-R6-R60	ServeRAID M5025 R6/R60	
ServeRAID M51xx Series	ServeRAID-M51xx	
ServeRAID M51xx Series R5	ServeRAID-M51xx_R5	
ServeRAID M51xx Series R5/R6	ServeRAID-M51xx_R5_R6	
ServeRAID M51xx Series R6	ServeRAID-M51xx_R6	
ServeRAID MR10i	ServeRAID-MR10i	
ServeRAID MR10ie	ServeRAID-MR10ie	
ServeRAID MR10il	ServeRAID-MR10il	
ServeRAID MR10is	ServeRAID-MR10is	
ServeRAID MR10k	ServeRAID-MR10k	
ServeRAID MR10M	ServeRAID-MR10M	
ServeRAID H1110/H1135	SAS2004	sas2ircu

Usage

Each of the modes supported by PRAID requires a specific syntax, but they all share some common parameters, described in Table 4 on page 49.

Table 4. PRAID parameters common to multiple modes

Parameter	Description	Usage
<p><code>-r:n</code></p> <p>Restore-defaults mode</p>	<p>Restores the RAID controller with the controller number specified by <i>n</i> to factory-default settings and then returns immediately. No RAID configuration is done if you use this parameter.</p> <p>If no value is specified for the controller number, all RAID controllers are reset to factory-default settings.</p> <p>Used alone, the parameter provides Restore-defaults mode. You must use this parameter in conjunction with Deploy mode parameters to reset controllers to the factory default settings before deploying a new configuration.</p>	<p><code>praid -r</code></p> <p>Restores all controllers to factory-default settings.</p> <p><code>praid -r:3</code></p> <p>Restores controller three to factory-default settings. No other controllers are affected.</p> <p><code>PRAID -f:policies.ini -r -v:5 -e1</code></p> <p>Configures the RAID controllers in the system using the policies file <code>policies.ini</code>, sets the verbose mode to maximum, and returns an error code if there were no matching policies for any controllers.</p>
<p><code>-f:policies_file</code></p> <p>Specifies the policy file</p>	<p>The policy file name. This parameter is required for capture mode and for deploy mode unless the <code>-d</code> parameter is used.</p> <p>In deploy mode, this points to the policies that you would like PRAID to use when configuring the RAID controllers. You cannot use this parameter with the <code>-d</code> parameter.</p> <p>In capture mode, this points to the file where you would like the captured configurations to be written. If the file does not exist, PRAID will create it. If the file does exist, PRAID appends to the end of it.</p> <p>The <code>-f</code> parameter is valid in deploy and capture modes.</p>	<p><code>praid -f:myfile.ini</code></p> <p>Uses the policies file, <code>myfile.ini</code>, to configure all RAID controllers.</p> <p><code>praid -c -f:myfile.ini</code></p> <p>Captures the RAID configuration of all controllers to the policy file, <code>myfile.ini</code>.</p>
<p><code>-y</code></p> <p>Suppresses prompting</p>	<p>Suppresses the confirmation prompt. This parameter is optional.</p> <p>If you select the <code>-y</code> parameter, PRAID does not prompt you before resetting controllers to factory-default settings. PRAID always resets all controllers to factory-default settings before configuring them.</p> <p>If you do not supply this parameter, PRAID will pause to warn you before resetting the RAID controllers to factory-default settings.</p> <p>The <code>-y</code> parameter is valid in deploy and restore-defaults modes.</p> <p>This parameter is optional.</p>	<p><code>praid -f:myfile.ini -y</code></p> <p>Uses the policies in <code>myfile.ini</code> to configure the RAID controllers and does not prompt before resetting all controllers to factory-default settings.</p>

Table 4. PRAID parameters common to multiple modes (continued)

Parameter	Description	Usage
-e2 Error code 2 if no supported controllers found	Returns an error code of 2 if there were no supported RAID controllers found in the system. By default, PRAID does not return an error if no controllers are found in the system. This parameter is valid in all modes. This parameter is optional.	<pre>praid -c -f:myfile.ini -e2</pre> Captures the RAID configuration of all RAID controllers to myfile.ini, and returns an error if no controllers are found in the system.
-e3 Error code 3 if no supported drives found	Returns an error code of 3 if at least one controller was found with no drives attached. By default, PRAID does not return an error if no drives are attached to a RAID controller. This parameter is valid in any mode. This parameter is optional.	<pre>praid -d -e3</pre> Configures all RAID controllers with default settings and returns an error if one or more controllers has no drives attached.
-v:n Verbose level	Sets the verbosity level, where n is: <ul style="list-style-type: none"> • 0 - quiet • 3 - default • 5 - maximum This parameter is valid in any mode. This parameter is optional.	<pre>praid -d -v:5</pre> Configures all RAID controllers with default settings, and sets the verbose level to maximum.

Deploy mode

The syntax for Deploy mode is:

```
PRAID.EXE -f:policies -r -d -p:path -e1 -e2 -e3  
-v:n -y -b
```

The parameters unique to Deploy mode are described below.

Table 5. PRAID Deploy mode parameters

Parameter	Description	Usage
-d Configure with defaults	<p>Configure all controllers in the system using default settings instead of using a policies file. The default settings used are the same as the default settings for the policies file.</p> <p>You cannot use this parameter with the -f parameter. See “Default RAID levels” on page 63 for the default values that will be assigned for each RAID controller based on the number of drives attached to the controller.</p> <p>This parameter is required unless the -f parameter is specified.</p>	<p><code>praid -d -r</code></p> <p>Configures all RAID controllers in the system using default settings.</p>
-e1 Error if no policy found	<p>Returns an error code of 1 if one or more controllers are not configured due to the fact that there was no policy found to configure them.</p> <p>This parameter is optional.</p>	<p><code>praid -f:policy.ini -r -e1</code></p> <p>Configures all RAID controllers using the policies file, policy.ini, and returns an error if no matching policy was found.</p>

Capture mode

The syntax for Capture mode is:

`PRAID.EXE -c[:p] -f:policies -e2 -e3 -v:n`

The parameters unique to Capture mode are described below.

Table 6. Capture mode parameters

Parameter	Description	Usage
<p>-c[:p]</p> <p>Capture mode</p>	<p>Indicates capture mode. The :p portion is optional. If you do not include the optional portion, :p will assume the default value: "t,d".</p> <p>You can use :p to provide a list of parameters describing the AppliesTo that is created when capturing the parameters to a policy. See "AppliesTo.n" on page 56.</p> <p>:p is a list containing any of the following:</p> <ul style="list-style-type: none"> • t – use the type of the RAID controller in the AppliesTo.1 entry for the policy. • c – use the controller number (scan order relative to all other RAID controllers in the system) in the AppliesTo.1 entry for the policy. • d – use the number of drives connected to the RAID controller in the AppliesTo.1 entry for the policy. <p>Note: You must specify the name of the policies file using the -f parameter when using the -c parameter.</p> <p>The policy or policies created are appended to the end of the file if the file exists. If the file does not exist, a new file is created. If there are multiple RAID controllers in the system, their configurations are placed in the file in scan order.</p>	<p>praid -c:m,t -f:myfile.ini</p> <p>Captures the configuration of all RAID controllers to myfile.ini using the machine type of the server and the RAID controller type as the AppliesTo.1 entry.</p>

Restore-defaults mode

The syntax for Restore-defaults mode is:

```
PRAID.EXE -r:n -e2 -v:n -y
```

Usage examples

Deploy mode examples

```
PRAID -r -d -y
```

This example is useful for unattended scripted installations.

- Configures all RAID controllers in the system using default settings.
- Does not prompt before setting controllers to factory-default settings.
- Performs drive synchronization without prompting, when required.

```
PRAID -f:policies.ini -r -v:5 -e1
```

- Configures the RAID controllers in the system using the policies file: policies.ini.
- Sets the verbose mode to maximum.

- Returns an error code if there were no matching policies for one or more controllers.

Capture mode examples

```
PRAID -c -f:policies.ini
```

Captures the configuration of all RAID controllers into the file: `policies.ini`.

```
PRAID -c:m,t -f:policies.ini
```

- Captures the configuration of all RAID controllers into the file: `policies.ini`.
- Uses the system machine type and RAID controller type as the `AppliesTo.1` entry in the policies file for each captured configuration.

Restore-defaults mode examples

```
PRAID -r -v:0 -y
```

- Restores all RAID controllers to factory default settings.
- Operates in silent mode; no messages are printed to the screen.
- Does not prompt you before restoring factory-default settings.

Return codes

- **0** - Success.
- **1** - Execution was successful, but the `-e1` parameter was supplied and at least one controller was not configured because there was no matching policy.
- **2** - Execution was successful, but the `-e2` parameter was supplied and no controllers were found in the system.
- **3** - Execution was successful, but the `-e3` parameter was supplied and at least one controller was not configured because no drives were attached.
- **4** - Syntax error on the command line.
- **5** - Syntax error in the policies file or the policy file could not be opened.
- **6** - Reserved
- **7** - Error resetting a controller to the default settings.
- **8** - Error gathering information about a controller.
- **9** - Error in the policy file.
- **10** - Error during processing.
- **11** - Error during deployment.

Policies file

When used in configure mode, the policies file directs how PRAID configures the RAID controllers in a system using keywords and values that you can customize. In capture mode, PRAID creates or appends to the end of a policies file the parameters that can configure other RAID controllers identically to the ones in the current system.

A policies file can be created using any of the following methods:

1. Run PRAID in capture mode to create a policies file from an already-configured RAID controller.
2. Use one of the example policies files provided with the ServerGuide Scripting Toolkit, and customize it to configure your RAID controllers.
3. Use an ASCII text editor to create a new policies file.

The policies file is an ASCII text file that is organized in INI-file format. Each INI-file section name indicates the start of a new policy for configuring RAID controllers.

The policies file must contain one or more uniquely-named sections using the format [Policy.*name*] where *name* is a unique user-assigned name that is used to identify the policy. *name* can be any combination of letters, numbers, underscores, periods, or dashes.

Some examples of legal section names are: [Policy.1], [Policy.mypolicy], and [Policy.My-RAID5-config]. Each section in the policies file represents a single policy for configuring RAID controllers. You can have up to 50 policies in a single policies file.

How PRAID selects a policy: Each section in the policies file represents a single policy for configuring the RAID controllers. In configure mode, each RAID controller is configured using a single policy, but a single policy can be used to configure multiple controllers. Each policy in a policies file contains one or more *AppliesTo.n* entries, where *n* is the number of the AppliesTo parameter within the policy. This entry is required in each section, so every section must contain at least an AppliesTo.1 entry. See “Policies file parameters” for a full description of the AppliesTo.n entry.

These entries are followed by a list of hardware parameters including machine type, number of drives connected to the RAID controller, and scan order, that are evaluated against the current system hardware. If all of the hardware parameters of an AppliesTo.n entry match the hardware being evaluated, this policy is used to configure the hardware. For each policy in the policies file, the AppliesTo.n entries for that policy are evaluated in order starting with AppliesTo.1. If none of the AppliesTo.n entries match the current hardware then the policy is not applied and the AppliesTo.n entries in the next policy are evaluated. This continues until either a match is found or no more policies exist in the file. If the end of the file is reached without a match then the controller is not configured. Because the policies are evaluated in order, you should place more specific policies at the beginning of the policies file.

Policies file parameters: This section describes the parameters used in the policies file. The Policy.*name* header and AppliesTo.1 entry are the only parameters required. All values are case-insensitive.

If you do not specify a value for any of the other parameters, they will be assigned their default value when applicable. If a parameter is not valid for a RAID controller, it is ignored.

In addition to this reference, the ServerGuide Scripting Toolkit also provides two example policies files that you can modify for your own use.

- RAID1-5.ini creates a RAID-1 array using the first two drives, and a RAID-5 array using the remaining drives. Valid for ServeRAID-6M and 8i.
- RAID5HSP.ini creates a single RAID-5 array with a single hot-spare drive using all available drives. Valid for ServeRAID-6M and 8i.
- template.ini provides a policies file template containing all parameters with details about each parameter.
- syntax.txt provides a syntax specification for the policies file.

Table 7. Policy file parameters

Keyword	Required?	Default	Description
<code>Policy.name</code>	Yes	None	This header designates the start of a new policy. See “ <code>Policy.name</code> ” on page 56 for additional information.
<code>AppliesTo.n</code>	Yes	None	Use this parameter to describe when the current policy should be chosen to configure the RAID controllers. See “ <code>AppliesTo.n</code> ” on page 56 for additional information.
<code>ReadAhead</code>	No	<ul style="list-style-type: none"> ADAPTIVE (for ServeRAID 6M) ON (for ServeRAID-7t 8i, 8k, and 8k-l) 	Specifies the read ahead setting that should be applied to the RAID controller. See “ <code>ReadAhead</code> ” on page 57 for additional information.
<code>RebuildRate</code>	No	HIGH	Specifies the rebuild rate that should be applied to the RAID controller. See “ <code>RebuildRate</code> ” on page 57 for additional information.
<code>StripeSize</code>	No	<ul style="list-style-type: none"> 8 (for ServeRAID 6M) 64 (for ServeRAID-7t, 8i, 8k, 8k-l) 	Specifies the stripe-unit size in KB that the controller should use for its arrays. See “ <code>StripeSize</code> ” on page 57 for additional information.
<code>Array_Mode</code>	No	AUTO	Defines the array-creation policy to use when selecting physical disk drives to include in an array. See “ <code>Array_Mode</code> ” on page 57 for additional information.
<code>Array_Defaults</code>	No	<ul style="list-style-type: none"> 0%:1 for ServeRAID-8e-SATA and 8e-SAS, LSI-SCSI-RAID when at least 3 drives are available 0%:1 for ServeRAID-6M, when one or more arrays has 4 or more physical drives 0%:0 for all other cases 	Defines the default values to use for the variance and number of hot-spare drives when AUTO is specified for <code>Array_Mode</code> . See “ <code>Array_Defaults</code> ” on page 58 for additional information.
<code>Array.letter</code>	No	None	Lets you specify exactly how many arrays are created and the exact physical drives that you would like in each array. See “ <code>Array.letter</code> ” on page 58 for additional information.
<code>Hotspares</code>	No	None	Defines a list of specific physical drives to designate as hot-spare drives. See “ <code>Hotspares</code> ” on page 59 for additional information.
<code>Logical_Mode</code>	No	AUTO	Defines the logical-drive creation policy to use when creating logical drives. See “ <code>Logical_Mode</code> ” on page 59 for additional information.

Table 7. Policy file parameters (continued)

Keyword	Required?	Default	Description
Logical_Defaults	No	FILL:AUTO:AUTO	Defines the default logical drive settings that should be used when creating logical drives. See “Logical_Defaults” on page 60 for additional information.
Logical.num	No	None	Lets you specify how many logical drives are created and the specific parameters for each logical drive. See “Logical.num” on page 60 for additional information.

Policy.name:

Description

This header designates the start of a new policy. You can specify *name* using any combination of letters, numbers, underscores, periods, or dashes. There is no maximum length for *name*, but the maximum length for a single line in the policies file is 256 characters. You can have up to 50 policies in a single policies file.

Examples

[Policy.RAID-5-Hotspare]

AppliesTo.n:

Description

Use this parameter to describe when the current policy is chosen to configure the RAID controllers. You can define up to 20 AppliesTo.n entries per policy. You must have an AppliesTo.1 entry for each policy, and AppliesTo.n is the only required parameter of a policy.

AppliesTo.n includes a comma delimited list containing one or more of the following parameters:

- *m:mtpe*, where *mtpe* is the four digit machine type of an IBM eServer™ or xSeries server.
- *s:serial*, where *serial* is the serial number of an IBM eServer or xSeries server.
- *c:contn*, where *contn* is the controller number (scan order) of the RAID controller with respect to all other RAID controllers in the system.
The number assigned to a particular controller is dependent on the controller's physical PCI slot and the order in which your system scans its PCI slots.
- *t:ctype*, where *ctype* is the type of the controller. The type is not case-sensitive, and must be one of the controller types listed in the table of RAID adapters supported by PRAID.
- *d:drives*, where *drives* is an integer value specifying the number of drives connected to the controller. Only drives in a **Ready** state after resetting the controller to factory-default settings are counted.
- ALL. Indicates that this policy must be used for all RAID controllers. This parameter is good to use if you declare a default policy that is not covered by any of the other policies.

Examples

Example using the m,s,c,t, and d parameters:

AppliesTo.1 = m:8865,t:ServeRAID-7t
AppliesTo.2 = c:1,d:15,s:87R478U

Example using the ALL parameter:

AppliesTo.1 = ALL

ReadAhead:

Description

Specifies the read ahead setting that must be applied to the RAID controller. If this parameter is not applicable for a RAID controller, it is ignored. See “Supported settings for RAID controllers” on page 61 for the list of ReadAhead settings supported by PRAID for each RAID controller. Possible settings are:

- Adaptive
- On
- Off

Examples

ReadAhead = On

RebuildRate:

Description

Specifies the rebuild rate that is applied to the RAID controller. If this parameter is not applicable for a RAID controller, it is ignored. See “Supported settings for RAID controllers” on page 61 for the list of RebuildRate settings supported by PRAID for each RAID controller.

- High
- Medium
- Low

Examples

RebuildRate = High

StripeSize:

Description

Specifies the stripe-unit size in KB that the controller uses for its arrays. If this parameter is not applicable for a RAID controller, it is ignored. See “Supported settings for RAID controllers” on page 61 for the list of StripeSize settings supported by PRAID for each RAID controller. Possible values are any stripe size supported by the controller.

Examples

StripeSize = 32

Array_Mode:

Description

Defines the array-creation policy to use when selecting physical disk drives to include in an array. Possible values are:

Auto Creates arrays using drives that have the same size in MB. This is the default. Each set of drives of the same size are combined into a single array. The maximum number of drives allowed per array is determined by

the limits of the RAID controller. Only drives in a **Ready** state after resetting the controller to factory-default settings are used in arrays. Hot-spare drives are created based on the rules supplied with the `Array_Defaults` parameter.

The `Array_Defaults` parameter allows you to modify the default behavior of the AUTO mode for arrays.

Custom Allows you to specify the exact physical disk drives to use in the array. If you specify this value, you must specify the `Array.letter` parameter with a list of drives for each array that you want to create. If you want hot-spare drives to be created, you must use the `Hotspares` parameter to list the hot-spare drives.

Examples

```
Array_mode = CUSTOM
```

Array_Defaults:

Description

Defines the default values to use for the variance and number of hot-spare drives when AUTO is specified for `Array_Mode`. This parameter is not valid if `Array_Mode` is set to CUSTOM.

The value of `Array_Defaults` is expressed in the format: *variance:hotspares*, where:

variance specifies the percentage variance to use when selecting drives to add to the array. This parameter is useful when you are using drives that vary slightly in size. Variance is based on a percentage of the drive size in MB. The valid values are:

- 0% - Combine only drives with equal size in MB into a single array.
- 5% - Combine all drives within 5% size in MB into a single array.
- 10% - Combine all drives within 10% size in MB into a single array.
- 100% - Combine all drives, regardless of size in MB, into a single array.

and *hotspares* is an integer that specifies the total number of hot-spare drives to create. The largest drives are chosen as hot-spare drives first. If not enough drives are available to create hot-spare drives, PRAID does not create any hot-spare drives.

Examples

```
Array_Defaults = 5%:1
```

Array.letter:

Description

Allows you to specify exactly how many arrays are created and the exact physical drives that you would like in each array. You can specify the physical drives using any of the following methods:

- The channel number and SCSI ID (for SCSI) or bus number and target ID (for SATA/SAS) of each drive. The channel number or bus number is always 1-based. The SCSI ID or target ID is always 0-based.
- A list of integer values indicating that the *n*th drive should be included in the array.

- The keyword ALL to indicate that all remaining drives attached to the controller that are not specified in previous arrays must be included in the current array.

The first array must be labeled `Array.A`. Additional arrays are labeled sequentially, `Array.B`, `Array.C`, and so on. The maximum number of arrays allowed per controller is determined by the limits of the specific RAID controller.

Examples

Example using channel number and SCSI ID:

```
Array.A = 1:1,1:2
Array.B = 1:3,1:4,1:5,2:1,2:2,2:3,2:4,2:5,2:6
Array.C = ALL
```

Example using integer values:

```
Array.A = 1,2,3
Array.B = ALL
```

Hotspares:

Description

Defines a list of specific physical drives to designate as hot-spare drives. You can specify the physical drives using any one of these methods:

- The channel number and SCSI ID (for SCSI) or bus number and target ID (for SATA/SAS) of each drive. The channel number or bus number is always 1-based. The SCSI ID or target ID is always 0-based.
- A list of integer values indicating that the *n*th drive must be included in the array.
- The keyword ALL to indicate that all remaining drives attached to the controller that are not specified in previous arrays must be included in the current array.

Examples

Example using channel number and SCSI ID:

```
Hotspares = 1:12,2:14
```

Example using integer value:

```
Hotspares = 12, 13
```

Logical_Mode:

Description

Defines the logical-drive creation policy to use when creating logical drives. Possible values are:

AUTO Indicates that defaults must be used for all parameters. Default parameters are:

- One logical drive is created on each array using all available space.
- The RAID level is set using the AUTO (default) scheme.
- Write-cache mode is set using the default value for the controller.

You can adjust these default values using the `Logical_Defaults` parameter.

CUSTOM Indicates that you want to specify all of the parameters for each logical drive that is created. If you specify CUSTOM, you must specify the parameters for each logical drive using the `Logical.num` parameter.

Examples

Logical_Mode = CUSTOM

Logical_Defaults:

Description

Defines the default logical drive settings that must be used when creating logical drives. This parameter is only valid when AUTO is specified for Logical_Mode. Values for this parameter are expressed in the format: *size:raidlevel:writemode*, where:

Size specifies the size of each logical drive. One logical drive is created on each array using the given size. *Size* can be in any of the following formats:

- A positive integer – specifies the size in MB.
- A percentage – specifies that a percentage of the total space must be used.
- FILL – indicates that all available space on the array must be used.

Raidlevel specifies the RAID level for the logical drive. See “Supported settings for RAID controllers” on page 61 for the list of RAID level settings supported by PRAID for each controller.

Writemode is an optional parameter that specifies the write-cache mode for each logical drive. If the write-cache mode cannot be set for a specific configuration, this parameter is ignored. See “Supported settings for RAID controllers” on page 61 for the list of write_cache mode settings supported by PRAID for each RAID controller.

Valid values are:

- ON
- OFF
- AUTO uses the default write-cache mode for the controller. (Recommended for most users.) This value is the default value if writemode is not specified.

Examples

Logical_Defaults = 50%:5EE:AUTO

Logical.num:

Description

Allows you to specify how many logical drives are created and the specific parameters for each logical drive. You can set the array letter where the logical drive is located, logical drive size, RAID level, and write-caching mode for each logical drive. The first logical drive must be labeled Logical.1. Additional logical drives are numbered Logical.2, Logical.3, and so on. You must specify at least one logical drive for each array. The maximum number of drives allowed per array and the maximum total number of logical drives allowed is determined by the specific RAID controller.

Values for this parameter are expressed in the format:

array:size:raidlevel:writemode where *array* specifies the array letter, and *size*, *raidlevel*, and *writemode* are as described in “Logical_Defaults.”

Examples

Logical.1 = A:50%:0

Logical.2 = A:50%:5EE

Logical.3 = B:FILL:1:ON

Logical.4 = C:4096:AUTO:AUTO

Supported settings for RAID controllers: Table 8 lists the supported settings for each RAID controller when using PRAID.

In some cases, the list of supported settings when using PRAID might differ from the supported settings of the RAID controller. These known cases are indicated in the table.

Table 8. Supported settings for each RAID controller when using PRAID. Bold settings are defaults.

RAID adapters	Read policy	Write policy	RAID Levels ¹	Stripe Size (KB)
ServeRAID-B5015	<ul style="list-style-type: none"> ON OFF 	[n/a]	R1, R5	4, 8, 16, 32, 64, 128 , 256, 512, 1024
LSI-IDEal-RAID	[n/a]	[n/a]	R0, R1	32, 64 , 128, 256, 512, 1024, 2048, 4096
LSI-MegaRAID-8480	[n/a]	[n/a]	R0, R1, R10, R5, R50	4, 8, 16, 32, 64 , 128
LSI-SAS-1078-IR	[n/a]	[n/a]	R0, R1	[n/a]
LSI-SAS-RAID	[n/a]	[n/a]	R0, R1, R1E	[n/a]
LSI-SCSI-RAID	[n/a]	[n/a]	R1	[n/a]
ServeRAID-7t	<ul style="list-style-type: none"> ON OFF AUTO 	<ul style="list-style-type: none"> ON OFF 	RVOLUME, R0, R1, R10, R5	16, 32, 64
ServeRAID-8i	<ul style="list-style-type: none"> ON OFF AUTO 	<ul style="list-style-type: none"> ON OFF 	RVOLUME, R0, R1, R10, R1E, R5, R50, R5EE, R6, R60	16, 32, 64, 128, 256 , 512, 1024
ServeRAID-8k	<ul style="list-style-type: none"> ON OFF AUTO 	<ul style="list-style-type: none"> ON OFF 	RVOLUME, R0, R1, R10, R1E, R5, R6	16, 32, 64, 128, 256 , 512, 1024
ServeRAID-8k-l	<ul style="list-style-type: none"> ON OFF AUTO 	<ul style="list-style-type: none"> ON OFF 	RVOLUME, R0, R1, R10	16, 32, 64, 128, 256 , 512, 1024
ServeRAID-8s	<ul style="list-style-type: none"> ON OFF AUTO 	<ul style="list-style-type: none"> ON OFF 	RVOLUME, R0, R1, R10, R1E, R5, R50, R6	16, 32, 64, 128, 256 , 512, 1024
ServeRAID-BR10ie	[n/a]	[n/a]	R0, R1, R1E	[n/a]
ServeRAID-BR10il	[n/a]	[n/a]	R0, R1, R1E	[n/a]
ServeRAID-M1015	[n/a]	[n/a]	R0, R1, R10	8, 16, 32, 64
ServeRAID-M1015-R5	[n/a]	[n/a]	R0, R1, R10, R5, R50	8, 16, 32, 64
ServeRAID-M1xxx	[n/a]	[n/a]	R0, R1, R10	8, 16, 32, 64
ServeRAID-M1xxx-R5	[n/a]	[n/a]	R0, R1, R10, R5, R50	8, 16, 32, 64

Table 8. Supported settings for each RAID controller when using PRAID (continued). Bold settings are defaults.

RAID adapters	Read policy	Write policy	RAID Levels ¹	Stripe Size (KB)
ServeRAID-M5014	<ul style="list-style-type: none"> ON OFF AUTO 	<ul style="list-style-type: none"> ON OFF 	R0, R1, R10, R5, R50	8, 16, 32, 64 , 128
ServeRAID-M5014-R6-R60	<ul style="list-style-type: none"> ON OFF AUTO 	<ul style="list-style-type: none"> ON OFF 	R0, R1, R10, R5, R50, R6, R60	8, 16, 32, 64 , 128
ServeRAID-M5015	<ul style="list-style-type: none"> ON OFF AUTO 	<ul style="list-style-type: none"> ON OFF 	R0, R1, R10, R5, R50	8, 16, 32, 64 , 128
ServeRAID-M5015-R6-R60	<ul style="list-style-type: none"> ON OFF AUTO 	<ul style="list-style-type: none"> ON OFF 	R0, R1, R10, R5, R50, R6, R60	8, 16, 32, 64 , 128
ServeRAID-M5025	<ul style="list-style-type: none"> ON OFF AUTO 	<ul style="list-style-type: none"> ON OFF 	R0, R1, R10, R5, R50	8, 16, 32, 64 , 128, 256, 512, 1024
ServeRAID-M5025-R6-R60	<ul style="list-style-type: none"> ON OFF AUTO 	<ul style="list-style-type: none"> ON OFF 	R0, R1, R10, R5, R50, R6, R60	8, 16, 32, 64 , 128, 256, 512, 1024
ServeRAID-M5xxx	<ul style="list-style-type: none"> ON OFF AUTO 	<ul style="list-style-type: none"> ON OFF 	R0, R1, R10, R5, R50, R6, R60	8, 16, 32, 64 , 128, 256, 512, 1024
ServeRAID-M51xx	<ul style="list-style-type: none"> ON OFF AUTO 	<ul style="list-style-type: none"> ON OFF 	R0, R1, R10	8, 16, 32, 64 , 128, 256, 512, 1024
ServeRAID-M51xx_R5	<ul style="list-style-type: none"> ON OFF AUTO 	<ul style="list-style-type: none"> ON OFF 	R0, R1, R10, R5, R50	8, 16, 32, 64 , 128, 256, 512, 1024
ServeRAID-M51xx_R6	<ul style="list-style-type: none"> ON OFF AUTO 	<ul style="list-style-type: none"> ON OFF 	R0, R1, R10, R6, R60	8, 16, 32, 64 , 128, 256, 512, 1024
ServeRAID-M51xx_R5_R6	<ul style="list-style-type: none"> ON OFF AUTO 	<ul style="list-style-type: none"> ON OFF 	R0, R1, R10, R5, R50, R6, R60	8, 16, 32, 64 , 128, 256, 512, 1024
ServeRAID-MR10i	<ul style="list-style-type: none"> ON OFF AUTO 	<ul style="list-style-type: none"> ON OFF 	R0, R1, R10, R5, R50, R6, R60	8, 16, 32, 64 , 128
ServeRAID-MR10il	<ul style="list-style-type: none"> ON OFF AUTO 	<ul style="list-style-type: none"> ON OFF 	R0, R1, R10, R5, R50, R6, R60	8, 16, 32, 64 , 128

Table 8. Supported settings for each RAID controller when using PRAID (continued). Bold settings are defaults.

RAID adapters	Read policy	Write policy	RAID Levels ¹	Stripe Size (KB)
ServeRAID-MR10is	<ul style="list-style-type: none"> ON OFF AUTO 	<ul style="list-style-type: none"> ON OFF 	R0, R1, R10, R5, R50, R6, R60	8, 16, 32, 64 , 128, 256, 512, 1024
ServeRAID-MR10k	<ul style="list-style-type: none"> ON OFF AUTO 	<ul style="list-style-type: none"> ON OFF 	R0, R1, R10, R5, R50, R6, R60	16, 32, 64 , 128, 256, 512, 1024
ServeRAID-MR10M	<ul style="list-style-type: none"> ON OFF AUTO 	<ul style="list-style-type: none"> ON OFF 	R0, R1, R10, R5, R50, R6, R60	8, 16, 32, 64 , 128
ServeRAID-C100	[n/a]	[n/a]	R0, R1, R10	64
ServeRAID-C100-R5	[n/a]	[n/a]	R0, R1, R10, R5	64
ServeRAID-C105	[n/a]	[n/a]	R0, R1	64
SAS2004	[n/a]	[n/a]	R0, R1, R10, R1E	[n/a]

1. RAID levels 5E and 5EE support only one logical drive per array.

Default RAID levels are described in “Default RAID levels.”

Default RAID levels: The default RAID level that is applied to a logical drive depends on the number of drives in the array and the controller type. These default values are designed to match the default values of the express configuration method in ServeRAID Manager where applicable. The following table shows the default RAID values that PRAID uses when AUTO is specified for *raidlevel*.

Table 9. Default RAID levels

Controller	Drives in array				
	1	2	3	4	5 or more
ServeRAID-B5015	[n/a]	RAID 1	RAID 5	RAID 5+Hotspare	RAID 5+Hotspare
LSI-IDEal-RAID	[n/a]	RAID 1	[n/a]	[n/a]	[n/a]
LSI-MegaRAID-8480	RAID 0	RAID 0	RAID 0	RAID 0+Hotspare	RAID 0+Hotspare
LSI-SAS-1078-IR	[n/a]	RAID 1	RAID 1+Hotspare	RAID 1+Hotspare	RAID 1+Hotspare
LSI-SAS-RAID	[n/a]	RAID 1	RAID 1E+Hotspare	RAID 1E+Hotspare	RAID 1E+Hotspare
LSI-SCSI-RAID	[n/a]	RAID 1	RAID 1+Hotspare	RAID 1+Hotspare	RAID 1+Hotspare
ServeRAID-7t	RAID 0	RAID 1	RAID 5	RAID 5+Hotspare	RAID 5+Hotspare
ServeRAID-8i	VOLUME	RAID 1	RAID 5	RAID 5+Hotspare	RAID 5+Hotspare
ServeRAID-8k	VOLUME	RAID 1	RAID 5	RAID 5+Hotspare	RAID 5+Hotspare
ServeRAID-8k-l	VOLUME	RAID 1	RAID 1+Hotspare	RAID 10	RAID 10+Hotspare
ServeRAID-8s	VOLUME	RAID 1	RAID 5	RAID 5+Hotspare	RAID 5+Hotspare
ServeRAID-BR10ie	[n/a]	RAID 1	RAID 1E	RAID 1E + Hotspare	RAID 1E + Hotspare

Table 9. Default RAID levels (continued)

Controller	Drives in array				
	1	2	3	4	5 or more
ServeRAID-BR10il	[n/a]	RAID 1	RAID IE	RAID IE + Hotspare	RAID IE + Hotspare
ServeRAID-M1015	RAID 0	RAID 0	RAID 0	RAID 0 + Hotspare	RAID 0 + Hotspare
ServeRAID-M1015-R5	RAID 0	RAID 0	RAID 0	RAID 0 + Hotspare	RAID 0 + Hotspare
ServeRAID-M1xxx	RAID 0	RAID 0	RAID 0	RAID 0 + Hotspare	RAID 0 + Hotspare
ServeRAID-M1xxx_R5	RAID 0	RAID 0	RAID 0	RAID 0 + Hotspare	RAID 0 + Hotspare
ServeRAID-M5014	RAID 0	RAID 0	RAID 0	RAID 0 + Hotspare	RAID 0 + Hotspare
ServeRAID-M5014-R6-R60	RAID 0	RAID 0	RAID 0	RAID 0 + Hotspare	RAID 0 + Hotspare
ServeRAID-M5015	RAID 0	RAID 0	RAID 0	RAID 0 + Hotspare	RAID 0 + Hotspare
ServeRAID-M5015-R6-R60	RAID 0	RAID 0	RAID 0	RAID 0 + Hotspare	RAID 0 + Hotspare
ServeRAID-M5025	RAID 0	RAID 0	RAID 0	RAID 0 + Hotspare	RAID 0 + Hotspare
ServeRAID-M5025-R6-R60	RAID 0	RAID 0	RAID 0	RAID 0 + Hotspare	RAID 0 + Hotspare
ServeRAID-M5xxx	RAID 0	RAID 0	RAID 0	RAID 0 + Hotspare	RAID 0 + Hotspare
ServeRAID-M51xx	RAID 0	RAID 0	RAID 0	RAID 0 + Hotspare	RAID 0 + Hotspare
ServeRAID-M51xx_R5	RAID 0	RAID 0	RAID 0	RAID 0 + Hotspare	RAID 0 + Hotspare
ServeRAID-M51xx_R6	RAID 0	RAID 0	RAID 0	RAID 0 + Hotspare	RAID 0 + Hotspare
ServeRAID-M51xx_R5_R6	RAID 0	RAID 0	RAID 0	RAID 0 + Hotspare	RAID 0 + Hotspare
ServeRAID-MR10i	RAID 0	RAID 0	RAID 0	RAID 0+Hotspare	RAID 0+Hotspare
ServeRAID-MR10il	RAID 0	RAID 0	RAID 0	RAID 0+Hotspare	RAID 0+Hotspare
ServeRAID-MR10is	RAID 0	RAID 0	RAID 0	RAID 0+Hotspare	RAID 0+Hotspare
ServeRAID-MR10k	RAID 0	RAID 0	RAID 0	RAID 0+Hotspare	RAID 0+Hotspare
ServeRAID-MR10M	RAID 0	RAID 0	RAID 0	RAID 0+Hotspare	RAID 0+Hotspare
ServeRAID-C100	RAID 0	RAID 0	RAID 0	RAID 0+Hotspare	RAID 0+Hotspare
ServeRAID-C100-R5	RAID 0	RAID 0	RAID 0	RAID 0+Hotspare	RAID 0+Hotspare
ServeRAID-C105	RAID 0	RAID 0	RAID 0	RAID 0+Hotspare	RAID 0+Hotspare
SAS2004	[n/a]	RAID 0	RAID 0	RAID 10	RAID 10

INVRAID

Use this program to dump all of the RAID controller configuration information to an output file. Supported RAID controllers by INVRAID are listed in the table, Table 3 on page 47.

Environment requirements

INVRAID works by parsing the output of other RAID configuration utilities. To accomplish this, the utilities used by INVRAID must be located in the system search path.

Usage

invraid [-I | -P] -F

Table 10. INVRAID parameters

Parameter	Description
-I	Displays information about all host adapters in the system in an INI-file format.
-P	Dumps information about all host adapters in a system in a keyword=value format.
-F :filename	Directs the output of invraid to the specified file.

Return values

Table 11 lists the values returned by INVRAID.

Table 11. Values returned by INVRAID

Return Value	Description
0	Success
1	Syntax Error
2	Program Error

Examples

To dump the information about all RAID controllers in a system to a file in INI file format with the name myraid.ini, use the -I parameter as shown here:

```
invraid.exe -i -f:myraid.ini
```

Returns:

```
[System]
Machine_Type = 7233
Serial_Number = 23A0075
Total_Number_Of_Controllers = 2
```

```
[RAIDController.1]
Model = LSI-SAS-1078-IR
BIOSVersion = 6.22.00.00
FirmwareVersion = 1.25.82.00
DriverVersion =
RebuildRate = HIGH
StripeSize =
ReadAhead = ADAPTIVE
PCI = 4:0:0:1000:0062:FFFF:FFFF
```

```
[RAIDController.1.Array]
Total_Number_Of_Arrays = 1
ID.1 = A
Members.1 = 1,2
```

```
[RAIDController.1.Hotspares]
Total_Number_Of_Hotspares = 0
```

```
[RAIDController.1.Logical]
Total_Number_Of_Logicals = 1
```

```
Array.1 = A
Size.1 = 139236
```

```

Raid_Level.1 = 1
WriteCache.1 = AUTO
State.1 = Okay (OKY)
Derived_State.1 = GOOD

[RAIDController.1.Physical]
Total_Number_Of_Physicals = 4

Channel.1 = 1
ID.1 = 0
Size.1 = 140013
Type.1 = SAS
Serial_Number.1 = 3NM2SQED0000980322JB
State.1 = Online (ONL)
Derived_State.1 = GOOD

Channel.2 = 1
ID.2 = 1
Size.2 = 140013
Type.2 = SAS
Serial_Number.2 = 3NM223CV0000974732Y9
State.2 = Online (ONL)
Derived_State.2 = GOOD

Channel.3 = 1
ID.3 = 2
Size.3 = 140013
Type.3 = SAS
Serial_Number.3 = 3NM2000900009746H8BY
State.3 = Ready (RDY)
Derived_State.3 = GOOD

Channel.4 = 1
ID.4 = 3
Size.4 = 140013
Type.4 = SAS
Serial_Number.4 = 3NM23J1J00009746XNSB
State.4 = Ready (RDY)
Derived_State.4 = GOOD

[RAIDController.2]
Model = ServeRAID-MR10M
BIOSVersion = 2.02.00
FirmwareVersion = 1.40.12-0551
DriverVersion =
PCI = 30:0:0:1000:0060:1014:0379

[RAIDController.2.Array]
Total_Number_Of_Arrays = 0

[RAIDController.2.Hotspares]
Total_Number_Of_Hotspares = 0

[RAIDController.2.Logical]
Total_Number_Of_Logicals = 0

[RAIDController.2.Physical]
Total_Number_Of_Physicals = 0

```

Using the -p parameter returns the same information, but the section title from the properties file is shown for each value:

```
invraid -p -f:myfile.ini
```

Returns:


```

System_Machine_Type = 7233
System_Serial_Number = 23A0075
RAIDController.1.Model = LSI-SAS-1078-IR
RAIDController.1.BIOSVersion = 6.22.00.00
RAIDController.1.FirmwareVersion = 1.25.82.00
RAIDController.1.DriverVersion =
RAIDController.1.RebuildRate = HIGH
RAIDController.1.StripeSize =
RAIDController.1.ReadAhead = ADAPTIVE
RAIDController.1.PCI = 4:0:0:1000:0062:FFFF:FFFF

RAIDController.1.Array.ID.1 = A
RAIDController.1.Array.Members.1 = 1,2

RAIDController.1.Logical.Array.1 = A
RAIDController.1.Logical.Size.1 = 139236
RAIDController.1.Logical.Raid_Level.1 = 1
RAIDController.1.Logical.WriteCache.1 = AUTO
RAIDController.1.Logical.State.1 = Okay (OKY)
RAIDController.1.Logical.Derived_State.1 = GOOD

RAIDController.1.Physical.Channel.1 = 1
RAIDController.1.Physical.ID.1 = 0
RAIDController.1.Physical.Size.1 = 140013
RAIDController.1.Physical.Type.1 = SAS
RAIDController.1.Physical.Serial_Number.1 = 3NM2SQED0000980322JB
RAIDController.1.Physical.State.1 = Online (ONL)
RAIDController.1.Physical.Derived_State.1 = GOOD

RAIDController.1.Physical.Channel.2 = 1
RAIDController.1.Physical.ID.2 = 1
RAIDController.1.Physical.Size.2 = 140013
RAIDController.1.Physical.Type.2 = SAS
RAIDController.1.Physical.Serial_Number.2 = 3NM223CV0000974732Y9
RAIDController.1.Physical.State.2 = Online (ONL)
RAIDController.1.Physical.Derived_State.2 = GOOD

RAIDController.1.Physical.Channel.3 = 1
RAIDController.1.Physical.ID.3 = 2
RAIDController.1.Physical.Size.3 = 140013
RAIDController.1.Physical.Type.3 = SAS
RAIDController.1.Physical.Serial_Number.3 = 3NM2000900009746H8BY
RAIDController.1.Physical.State.3 = Ready (RDY)
RAIDController.1.Physical.Derived_State.3 = GOOD

RAIDController.1.Physical.Channel.4 = 1
RAIDController.1.Physical.ID.4 = 3
RAIDController.1.Physical.Size.4 = 140013
RAIDController.1.Physical.Type.4 = SAS
RAIDController.1.Physical.Serial_Number.4 = 3NM23J1J00009746XNSB
RAIDController.1.Physical.State.4 = Ready (RDY)
RAIDController.1.Physical.Derived_State.4 = GOOD

RAIDController.2.Model = ServeRAID-MR10M
RAIDController.2.BIOSVersion = 2.02.00
RAIDController.2.FirmwareVersion = 1.40.12-0551
RAIDController.2.DriverVersion =
RAIDController.2.PCI = 30:0:0:1000:0060:1014:0379

```

VALRAID

VALRAID is a utility program that can be used to validate policy files against inventory files generated by the INVRAID utility.

VALRAID has two modes of operation:

- **Simulation mode** simulates the effect a policy file would have on a controller.

- **Check mode** determines if the policy file matches the configuration represented in the inventory file.

Simulation mode

Used in simulation mode, VALRAID simulates the effect that a policy file has on a RAID configuration if it is applied using the PRAID utility. This capability can be used when creating PRAID policy files. The policy files can be tested without running PRAID on the target system.

Check mode

Used in check mode, VALRAID determines if the policy file specified matches the RAID configuration represented in the inventory file. Use this capability in operating system deployment scripts to bypass the RAID configuration step if the controller is already configured with the required RAID configuration. This process lets you avoid restarting the system before installing the operating system. VALRAID sets the return code to 20 to indicate that the policy file does not match the configuration represented by the inventory file.

Usage

The two modes of operation share most parameters, but the syntax is mode-specific.

The simulation mode syntax is:

```
valraid -ini:input_inventory_file -inp:input_policy_file -out:output_inventory_file
-outp:output_policy_file -raid:inifiles
```

The check mode syntax is:

```
valraid -c -ini:input_inventory_file -inp:input_policy_file -raid:inifiles
```

Table 12. VALRAID parameters

Parameter	Description	Example
-ini:input_inventory_file	Specifies the input inventory file. Generate the inventory file by running INVRAID against a target system.	valraid -ini:myfile.inv -inp:policy.ini -out:myfile.inv -outp:newpolicy.ini -raid:/inifiles
-inp:input_policy_file	Specifies the input policy file.	valraid -ini:myfile.inv -inp:policy.ini -out:myfile.inv -outp:newpolicy.ini -raid:/inifiles
-out:output_inventory_file	Specifies the filename for the output inventory file. This is an inventory file representing the RAID configuration that would result from using the PRAID utility to apply <i>input_policy_file</i> to the system described in <i>input_inventory_file</i> . This option is valid only for simulation mode.	valraid -ini:myfile.inv -inp:policy.ini -out:myfile.inv -outp:newpolicy.ini -raid:/inifiles

Table 12. VALRAID parameters (continued)

Parameter	Description	Example
<code>-outp:output_policy_file</code>	Specifies the filename for the output policy file. This file can be applied to a target system using the PRAID utility. This option is valid only for simulation mode.	<code>valraid -ini:myfile.inv -inp:policy.ini -outi:newfile.inv -outp:newpolicy.ini -raid:/inifiles</code>
<code>-raid:inifiles</code>	Specifies the directory that contains the RAID ini files.	<code>valraid -ini:myfile.inv -inp:policy.ini -outi:newfile.inv -outp:newpolicy.ini -raid:/inifiles</code>
<code>-c</code>	Specifies check mode. Check mode compares the configuration from <i>input_inventory_file</i> to the configuration represented in <i>input_policy_file</i> . The default is simulation mode.	<code>valraid -c -ini:myfile.inv -inp:policy.ini -raid:/inifiles</code>

Return codes

VALRAID uses the following return codes:

- 0 – Success
- 1 – Error parsing input policy file
- 2 – Error parsing input inventory file
- 3 – Controller is not supported
- 4 – Raid level is not supported
- 5 – Stripesize is not supported
- 6 – Number of arrays not supported
- 7 – Number of drives in array not supported
- 8 – Number of logical volumes in array is not supported
- 9 – Not enough drives to create hotspare
- 10 – Not enough drives of the same size
- 11 – Error opening input policy file
- 12 – Error opening input inventory file
- 13 – Error opening output inventory file
- 14 – Error writing to output inventory file
- 15 – Error opening output policy file
- 16 – Error writing output policy file
- 17 – Partial drive sizing not supported
- 18 – Command line syntax error
- 19 – No policy match
- 20 – Controller not configured, does not match policy file

Tools included with the Linux Scripting Toolkit

The Linux Scripting Toolkit includes several additional tools to make the Toolkit more efficient. This section describes the additional tools provided by this release of the Linux Scripting Toolkit:

- Advanced Settings Utility
- QAUCLI
- UpdateXpress System Pack Installer

Advanced Settings Utility

For convenience, the Linux Scripting Toolkit includes the Advanced Settings Utility (ASU). You can use ASU to modify firmware settings from the command line on multiple operating-system platforms.

The Linux Scripting Toolkit uses a subset of the ASU function to capture and deploy firmware settings as part of your scripted deployments.

Usage

This section describes the ASU functions used by the Linux Scripting Toolkit.

Table 13. ASU functions in Linux Scripting Toolkit

Command	Description
asu show all	Is used to display and capture BIOS settings. You can use redirection to store this output in a file as shown here: <code>asu.exe show bios > bios_settings.ini</code>
asu save <i>filename</i>	Is used to apply CMOS settings from a file. ASU looks for the filename specified by <i>filename</i> , and reads the contents. If the contents are valid CMOS settings, they are applied, one line at a time, to the server. This example applies the settings captured above: <code>asu save bios_settings.ini</code> Note: Only settings captured from an identical model can be replicated, due to difference in BIOS settings and valid values between models.
asu set IMM.HostIPAddress <i>IP address</i>	Sets the external IP address in the Integrated Management Module (IMM) to the specified address. This setting is part of the IMM group.
asu set IMM.LanOverUsb <i>enabled \ disabled</i> -kcs	Enables or disables the IMM LAN over USB interface. Note: When you enable or disable this setting, you must use the KCS interface to ensure that the asu command completes correctly and returns status.

You can use the QAUCLI utility to configure Fibre Host Bus Adapters (HBAs). A 32-bit version of this utility comes with the Linux Scripting Toolkit. You can download this utility from QLogic at <http://www.qlogic.com>.

Usage

Table 14. QAUCLI usage

Command	Description
<code>qaucli -pr fc -e (view ?)</code>	Shows the current boot device information on all HBAs
<code>qaucli -pr fc -e (hba_instance hba_wwpn target_wwnn target_wwpn lun_id [prim alt1 alt2 alt3])</code>	<p>Configures the Operating System to boot from a particular target, where:</p> <p><i>hba_instance</i> The HBA instance number of an HBA port.</p> <p><i>hba_wwpn</i> The World Wide Port Name of an HBA port.</p> <p><i>target_wwnn</i> The World Wide Node Name of a target device, in the format <i>nn-nn-nn-nn-nn-nn-nn-nn</i> or <i>nnnnnnnnnnnnnnnnnnnnnn</i>.</p> <p><i>target_wwpn</i> The World Wide Port Name of a target device, in the format <i>nn-nn-nn-nn-nn-nn-nn-nn</i> or <i>nnnnnnnnnnnnnnnnnnnnnn</i>.</p> <p><i>lun_id</i> The Logical Unit Number of a LUN.</p> <p><i>prim</i> The primary boot port name.</p> <p><i>altn</i> The name of the alternate boot port. You can specify up to three alternate boot ports.</p>
<code>qaucli -pr fc -e (hba_instance hba_wwpn) (view ?)</code>	Shows the current boot device information for the specified HBA port.
<code>qaucli -pr fc -e (hba_instance hba_wwpn) (enable 0 0 0)</code>	Configures the operating system to boot from the first target found by the BIOS. The default LUN is 0.
<code>qaucli -pr fc -e (hba_instance hba_wwpn) disable [prim alt1 alt2 alt3]</code>	Clears the selected boot device settings on the indicated HBA port.
<code>qaucli -pr fc -l (hba_instance hba_wwpn)</code>	Displays information about the LUNs attached to the specified HBA port.

Examples

The following examples illustrate `qaucli` utility usage.

Note: Some of these examples are broken across multiple lines however, when using QAUCLI, you must enter all of the parameters on a single line.

Example	Description
<code>qaucli -pr fc -e view</code>	Displays the current boot device information on all HBAs.
<code>qaucli -pr fc -e E0-FF-EE-DE-CD-34-56-30 E0-00-ED-DE-CD-34-56-30 E0-10-ED-DE-CD-34-56-30 1 prim</code>	Configures HBA E0-FF-EE-DE-CD-34-56-30 E0-00-ED-DE-CD-34-56-30 E0-10-ED-DE-CD-34-56-30 to boot from the primary target.
<code>qaucli -pr fc -e E0-FF-EE-DE-CD-34-56-30 view</code>	Displays the current boot setting information for HBA port E0-FF-EE-DE-CD-34-56-30.
<code>qaucli -pr fc -e E0-FF-EE-DE-CD-34-56-30 disable prim</code>	Clears the selected boot device setting on HBA port E0-FF-EE-DE-CD-34-56-30.
<code>qaucli -pr fc -l E0-FF-EE-DE-CD-34-56-30</code>	Displays information about the LUNs attached to HBA port E0-FF-EE-DE-CD-34-56-30.

septool210

For convenience, the Linux Scripting Toolkit, includes the SEPTool. You can use SEPTool to acquire any additional System Enablement Packs (SEP)s.

Usage

This section describes the septool210 functions.

Table 15. septool functions in Linux Scripting Toolkit

Command	Description
<code>septool210 acquire</code>	downloads SEPs from the IBM web site.

UpdateXpress System Pack Installer

For convenience, the Linux Scripting Toolkit includes the UXSPi to help you acquire updates for inclusion in your deployment scenarios. The UpdateXpress System Pack Installer is located at: `...sgdeploy\updates\uxsp`.

The UpdateXpress System Pack Installer can perform these functions:

- Acquire firmware and driver updates for supported machine type/operating system combinations from a remote location, such as the IBM Support website.
- Inventory a system to be updated and compare the inventory to the list of available updates, then recommend and deploy a set of updates for the system.
- Create bootable media on CD-ROM, DVD, or USB key to use in applying firmware to supported systems.

For more information on running the UpdateXpress System Pack Installer, change directory to the UXSPi directory and run the UXSPi executable shown below:

```
./ibm_utl_uxspi_x.xx_anyos_x86-64.bin -update -help
```

Usage

The Linux Scripting Toolkit uses the UXSPi in the update mode to acquire and deploy device drivers and firmware as part of Linux Scripting Toolkit deployments. This section details the command-line options for the `uxspi -update` mode.

Table 16. UXSPi update mode options

Option	Description
<code>-firmware</code>	The firmware option forces UXSPi to install only firmware updates.
<code>-drivers</code>	The driver option forces UXSPi to install only driver updates.
<code>-f update_ids, -force=update_ids</code>	Specifies that UXSPi use the <code>unattendedForcedInstallCommandLine</code> field in the update XML rather than the <code>unattendedInstallCommand</code> field.
<code>-s update_ids -select=update_ids</code>	The select option deploys the specified set of updates to the target system even if the system version is newer than the version in the update package. Use this option to roll-back firmware and driver levels where necessary.
<code>-l update_xml_path, -local=update_xml_path</code>	Specifies the filename of a local UXSP XML file or the path to search for one.
<code>-n, -new</code>	Selects all updates that are newer than the current system versions or not currently installed on the system.
<code>-e update_ids, -exclude=update_ids</code>	Excludes the specified update IDs. You can provide multiple IDs in a comma-separated list.
<code>-i update_ids, -include= update_ids</code>	Includes the updates specified in the list of update IDs. You can provide multiple IDs in a comma-separated list.
<code>-ignore-undetected=update_ids</code>	Specifies not to apply the indicated update IDs. You can provide multiple IDs in a comma-separated list.
<code>-L, -latest</code>	The default behavior of UXSPi is to apply the latest UXSPi update pack found in the UXSPi directory. This option forces UXSPi to install the latest updates whether they are from an update pack, are individual updates, or a combination of the two.
<code>-remote=remote_address</code>	Runs the update command on the remote server specified by <code>remote_address</code> .
<code>-remote-user=remote_user</code>	Specifies the remote user ID to use when connecting to a remote system specified with <code>-remote</code> .
<code>-remote-password=password</code>	Sets the password for the user ID specified by <code>-remote-user</code> .
<code>--remote-dir=directory</code>	Specifies the staging or working directory on the remote system.
<code>-noinventory</code>	Causes UXSPi to gather only the machine type and operating system information without performing an inventory of existing updates.
<code>-nouxsp</code>	Do not deploy UXSPs.
<code>-r, -report</code>	Displays a summary report of updates used in the compare step.

Examples

The following example can be used to specify an UpdateXpress System PackXML file named `uxsp.xml` located in the same directory as the UXSPi executable.

```
./ibm_utl_uxspi_x.xx_anyos_x86-64.bin update -l uxsp.xml
```

Chapter 8. Incorporating the Scripting Toolkit with your existing process

To incorporate Scripting Toolkit procedures into an existing deployment process, use the HWDETECT utility to determine if the combined process is being executed on Scripting Toolkit supported hardware. You can then add appropriate branches in the batch files to use the existing process or the Scripting Toolkit process.

For example, you might use the -s option of HWDETECT to determine if the current system is an IBM server:

```
hwdetect -s
if errorlevel 1 goto NONIBM
if errorlevel 0 goto IBM

:NONIBM
rem Perform non-IBM equipment specific processing here.

:IBM
rem Perform IBM eServer or xSeries equipment specific processing here.
```

Chapter 9. IBM ServerProven compatibility

The IBM ServerProven website provides valuable information on selected products for compatibility with IBM System x, BladeCenter, and xSeries servers. Check the following link for up-to-date compatibility with operating systems, configuration, and hardware options <http://www.ibm.com/servers/eserver/serverproven/compat/us/>.

Chapter 10. Hints and tips

This section contains information on known problems and limitations, best practices, and hints and tips for using the Linux Scripting Toolkit.

Installing network operating systems using HTTP

To install network operating system using HTTP, Linux Toolkit requires that you enable directory browsing on the web server. If enabling browsing presents a security risk, you can provide the information needed to perform installations by adding and setting the following environment variables in the Altiris jobs:

```
TK_DISTRO_HTTP_KERNEL_PKG  
TK_DISTRO_HTTP_RELEASE_PKG  
TK_DISTRO_HTTP_ISOLINUX  
TK_DISTRO_HTTP_INITRD  
TK_DISTRO_HTTP_VMLINUX
```

This example shows the values for these variables for a Red Hat Enterprise Linux 5, Update 4 x64 installation:

```
TK_DISTRO_HTTP_KERNEL_PKG=Server/kernel-2.6.18-160.el5.x86_64.rpm  
TK_DISTRO_HTTP_RELEASE_PKG=Server/redhat-release-5Server-5.4.0.2.x86_64.rpm  
TK_DISTRO_HTTP_ISOLINUX=isolinux/isolinux.cfg  
TK_DISTRO_HTTP_INITRD=isolinux/initrd.img  
TK_DISTRO_HTTP_VMLINUX=isolinux/vmlinuz
```

Improving PXE boot transfer speed

To speed up the TFTP transfer of the preboot environment, increase the packet size for the TFTP server of Altiris. To increase the packet size, complete the following steps:

1. Open the PXE configuration utility.
2. Click the **Multicast** tab.
3. Change the packet size to 1420 (the default packet size is 768).

If the system is uEFI-based, see “Disabling uEFI PXE to decrease network boot time” for more ways to increase performance.

Disabling uEFI PXE to decrease network boot time

To improve the time it takes to start the network for uEFI-based systems, complete the following steps:

1. Start the system.
2. Press **F1** to display the menu options.
3. Navigate to **System Settings > Network > PXE Configuration**.
4. Select Port %MAC1%.
5. Select **Enable PXE** and press **Enter**.
6. Select **Legacy Support** and press **Enter**.
7. Select **Save Changes** and press **ESC**.
8. Select Port %MAC2%.
9. Select **Enable PXE** and press **Enter**.
10. Select **Legacy Support** and press **Enter**.

11. Select **Save Changes** and press **ESC**.

Performing Clone/Deploy of SLES 11

To perform a clone of SUSE Linux Enterprise Server(SLES) 11 with Altiris, you must use Altiris. Deployment Solution 6.9 SP3. To deploy a cloned image of SLES11 from Altiris, refer to Altiris Knowledgebase article 48984 to enable SLES11 support in Altiris.

After enablement, complete the following steps to deploy the SLES11 cloned image:

1. On the Altiris console, navigate to **Deploy Linux TK**.
2. In the Distribute Disk Image Task of the Deploy Linux TK job, uncheck the option **Automatically perform configuration task after completing this image task**.

Note: If you do not deselect the check box, the network cannot be configured.

Using RapiDeploy with SLES 10 and VMware

RapiDeploy provides native imaging support for EXT2 and EXT3 file systems. By default, SuSE Linux Enterprise Server 10 uses the reiserfs file system when installing. To clone and redeploy a SLES 10 image using the Toolkit jobs, you must perform the original deployment using an EXT3 file system.

To change the file system prior to installation with the Toolkit, edit the sles10.xml AutoYast file provided by the Toolkit.

Replace this partitioning section of the AutoYast file:

```
<partitioning config:type="list">
  <drive>
    <initialize config:type="boolean">true</initialize>
    <use>all</use>
  </drive>
</partitioning>
```

with the following partitioning section:

```
<partitioning config:type="list">
  <drive>
    <initialize config:type="boolean">true</initialize>
    <use>all</use>
    <partitions config:type="list">
      <partition>
        <filesystem config:type="symbol">ext3</filesystem>
        <format config:type="boolean">true</format>
        <mount>/boot</mount>
        <size>100mb</size>
      </partition>
      <partition>
        <filesystem config:type="symbol">swap</filesystem>
        <format config:type="boolean">true</format>
        <mount>swap</mount>
        <size>auto</size>
      </partition>
      <partition>
        <filesystem config:type="symbol">ext3</filesystem>
        <format config:type="boolean">true</format>
        <mount>/</mount>
        <size>auto</size>
      </partition>
    </partitions>
  </drive>
</partitioning>
```

```
</partition>
</partitions>
</drive>
</partitioning>
```

Linux X server considerations

There are special considerations for using Linux X server with a Remote Supervisor Adapter II (RSA II) port. If you are using this configuration, consider the following items:

If the Remote Supervisor Adapter II-EXA is installed on a server that is running either the Red Hat Linux or SuSE Linux operating system, make sure that the Linux operating system is selected in the Remote Supervisor Adapter II settings in the server BIOS. To set Linux as the operating system in the server BIOS, complete the following steps :

1. Start or restart the server.
2. When prompted, press **F1** to display the configuration menu.
3. Click **Advanced Setup > ASM Settings**.
4. In the **OS USB** field, select Linux.
5. Select **Save Values and Reboot ASM**.

Note: If you run an automated X Windows system configuration utility, repeat these configuration changes.

Install the operating system in text mode. Set the color depth to 16-bit and the screen resolution to 1024 x 768.

If SuSE Linux or Red Hat Linux is already installed and configured to run in text mode, and will never use the X Window system, no additional configuration is required for the RSA II-EXA to function correctly.

The Remote Supervisor Adapter II-EXA requires USB mouse support from the operating system. If you install a supported Linux operating system on your server using a USB mouse, the installation process automatically establishes USB mouse support; no further action is required. If you install a supported Linux operating system using a PS/2 mouse, you must modify the Linux files after installation to add USB mouse support. Follow the instructions in this document to add USB mouse support

The Remote Supervisor Adapter II-EXA requires a Video Electronics Standard Association (VESA) device driver. The VESA video device driver enables the remote control screen and the local screen to display the same information (clone mode).

When using power management, the video output might not return correctly from some power saving states. To correct this problem, use the **xset** command to disable DPMS:: `xset -dpms`

For more information, see the *IBM Remote Supervisor Adapter II-EXA Technical Update for Linux Operating Systems* available from ftp://ftp.software.ibm.com/systems/support/system_x_pdf/88p9275.pdf.

Special considerations for BladeCenter Blades and Linux X server configuration

After installing BladeCenter blade and Linux X servers, do not change the monitor configuration or any other graphical settings. If you must change the graphics settings, enter the following command to start the configuration utility:

```
sax2 -m 0=fbdev
```

When using power management, the video output might not return correctly from some power saving states. To correct this problem, use the **xset** command to disable DPMS as shown: `xset -dpms`

IPv6 compliance

Beginning in version 2.20, the Linux Scripting Toolkit provides support for locating operating system files on IPv6 network shares. To use operating system files located in a share on an IPv6 network, you must update the variable **TK_DISTRO_PATH** in the appropriate Altiris job.

Performing network based installations of SLES11 SP1 using static IPv6 addresses

When performing a network-based installation of SLES11 SP1 in a static IPv6 environment, you must use either **ipv6=1**, which accepts both IPv4 and IPv6 addresses, or **ipv6only=1**, which accepts only IPv6 addresses, as a boot parameter.

For example, to configure static IPv6 addresses for an IPv6-only network, use these boot parameters:

```
ipv6only=1 netdevice=eth0 hostip=2000::2dae:2390/64
```

During the **Perform installation** stage of the deployment, the system will restart. You might receive the error message: Download (curl) error ... /media.1/media after the system starts. To avoid this problem, follow these steps:

1. Navigate to another console that has a command prompt.
2. Configure the network by issuing this command:

```
ifconfig eth0 inet6 add 2000::2dae:2390/64
```
3. Return to the console with the error.
4. Click **Retry** to continue the installation.

Enabling Linux Scripting Toolkit PXE images to work with other PXE images

The Linux Scripting Toolkit uses a customized `pxelinux.0` file rather than the default file that comes with `syslinux`. If you already have a PXE server in your network and want to use PXE images generated by the Linux Scripting Toolkit with other PXE images, you must implement a PXE chain. To implement a PXE chain, complete the following steps:

1. Download `syslinux 3.72` or higher from <http://www.kernel.org/pub/linux/utils/boot/syslinux/>.
2. Copy the file `core/pxelinux.0` from the `syslinux` directory structure to your `tftpboot` directory.
3. Extract the file `pxechain.com` from a PXE image created using the Linux Scripting Toolkit. The `pxechain.com` file is located in `/tftpboot/lnxtoolkit/`

image_name/tc.zip. For example, if you have created a PXE image called *PXE_test* using the Linux Scripting Toolkit, you can extract *pxechain.com* using the following command:

```
unzip /tftpboot/lntoolkit/PXE_test/tc.zip
```

4. Copy the *pxechain.com* file to your tftpboot directory.
5. Copy */tftpboot/lntoolkit/pxelinux.0* to your tftpboot directory.
6. Create a subdirectory of your tftpboot directory called *ibm*.
7. Copy the PXE files created by the Linux Scripting Toolkit into the *tftpboot/ibm* directory.
8. Create a subdirectory of tftpboot called *pxelinux.cfg*.
9. Create the file *tftpboot/pxelinux.cfg/default*. This sample default file includes the PXE image created by the Linux Scripting Toolkit. You can add other existing PXE images as shown:

```
prompt 0
default ibmchain
timeout 100
label ibmchain
kernel pxechain.com
append ::ibm/pxelinux.0
label your_other_pxe
kernel pxechain.com
append ::your_other_pxe/pxelinux.0
```

When you have completed these steps, the tftpboot file structure looks like this:

```
-- ibm
|  -- img2a
|  -- img3a
|  -- pxelinux.0      <- IBM's modified pxelinux.0
|  -- pxelinux.cfg
|  |  -- PXE_test    <- The default file for Linux Scripting Toolkit created PXE Image
|  |  -- tc.zip
|  |  -- tcrootfs
|  -- your_other_pxe
|  |  -- vmlinuz
|  |  -- initrd.gz
|  |  -- pxelinux.0  <- your_other_pxe's pxelinux.0
|  |  -- pxelinux.cfg
|  |  |  -- default  <- your_other_pxe's default
|  -- pxechain.com   <- pxechain.com from tc.zip (Step 1)
|  -- pxelinux.0     <- pxelinux.0 from syslinux 3.72 (or later)
|  -- pxelinux.cfg
|  |  -- default     <- default file for pxechain.com
```

Known problems and limitations

This section provides information and alternative solutions for known problems and limitations of the Linux Scripting Toolkit.

Integration script *altiris_sgtklinux.cmd* fails with error code 16

When running the *altiris_sgtklinux.cmd* command, it fails with the following error:

```
[Altiris_PerformUpgradeSteps.cmd]
Running
"C:\Program Files (x86)\Altiris\Express\Deployment Server\pxe\pxeconfig.exe"
-delete "IBM Linux Toolkit 1.00"
"C:\Program Files (x86)\Altiris\Express\Deployment Server\pxe\pxeconfig.exe"
-delete "IBM Linux Toolkit 1.00" returned a non-zero error code
      ReturnCode = 16
      The Altiris_PerformUpgradeSteps.cmd script was aborted.
```

To correct this problem, restart the Windows operating system on which the Altiris Deployment Server and the PXE server are running.

Integration script `altiris_sgtklinux.cmd` fails with error code 10

When running the `altiris_sgtklinux.cmd` command, it fails with the following error:

```
[Altiris_CreatePXEImage.cmd]
Running
"E:\Program Files\Altiris\Express\Deployment Server\pxe\pxeconfig"
-create -linux "IBM Linux Toolkit 1.00" -x64 -defaultLinux -oem "ibm"
-save -notmanaged
      "C:\Program Files\Altiris\Express\Deployment Server\pxe\pxeconfig"
-create -linux "IBM Linux Toolkit 1.00" -x64 -defaultLinux -oem "ibm"
-save -notmanaged returned a non-zero error code
      ReturnCode = 10
      The Altiris_CreatePXEImage.cmd script was aborted.
```

This problem can occur when the Altiris Linux x64 Automation Environment is not properly installed by the `altiris_sgtklinux.cmd` script.

To correct this problem, manually install the Altiris Linux x64 Automation Environment (BDCgpl_6.9.xxxx.frm) then try running the script again.

To manually install the Altiris Linux x64 Automation Environment:

1. Start the PXE Configuration Utility from the Altiris Console.
2. Select **Boot > New...**
3. Select the **Linux** radio button and click **Add pre-boot"**.
4. Select the row **Linux X64** from the table.
5. Enter the path to the BDCgpl_6.9.xxxx.frm file.
6. Click **Next** and follow the on-screen instructions.

Operating system installation halts after reboot when using LSI SAS RAID controller

Some combinations of LSI SAS RAID controllers and operating systems might experience a system halt after rebooting during operating system installation. The affected operating systems are:

- SLES 10
- SLES 11
- RHEL 5
- VMware 4

in combination with one of these RAID controllers:

- LSI-SAS-1078-IR
- LSI-SAS-(1064,1068)
- ServeRAID-BR10i
- ServeRAID-BR10ie

This problem occurs when the server has a drive that is not part of a RAID array and not configured as a hot spare. The problem is caused by the ordering of Linux mptsas devices.

For example, consider a system that has four drives with two configured in a RAID 1 array, one configured as a hot spare, and one outside the array. The BIOS sees the drive outside the array, `/dev/sda`, as HDD1. The RAID array, `/dev/sdb`, is treated as HDD0. The operating system installation puts the boot files on `/dev/sda`, the drive outside the array, but after the reboot, the installation looks to HDD0 for the boot files.

To work around this problem, use one of these options:

- Do not configure RAID.
- Change the RAID configuration so that all drives are included in a RAID array.
- Remove the drive outside the RAID array from the controller.
- Modify the boot order of the system to point to the drive outside the array instead of the array.

UpdateXpress System Pack Installer returns errors when supported hardware is not present

Deployment tasks that include installation of UpdateXpress System Packs (UXSPs) will return errors if hardware supported by the UXSPs is not present in the target system. These errors can be safely ignored.

Savestat will not save to location 9 on xSeries 226 with BIOS PME170CUS

On the xSeries 226 with BIOS Level PME170CUS, **savestat** cannot save a value to byte nine in persistent storage.

Unattended Linux installation requests network device

When performing unattended Linux operating system installs, the process might pause to ask which network device to use if there are multiple devices available. To avoid this problem, you can add a kernel parameter to specify the desired network device during the Workflow creation process.

In the **OS install** section of the workflow, a field is provided for optional kernel parameters.

The kernel parameter varies by operating system:

- For Red Hat and VMware: `ksdevice=eth`, where *eth* is the network device to use. For example `eth0`, `eth1`, and so on.
- For SUSE Linux: `netdevice=eth` where *eth* is the network device to use. For example `eth0`, `eth1`, and so on.

RAID configuration fails for ServeRAID-8E SATA

RAID configuration is not supported for the ServeRAID-8E SATA disk controller. Before you can install an operating system, you must create one or more simple volumes using the BIOS-based utility.

This limitation affects the following systems:

- System x3400, types 7973 and 7974
- System x3550, types 1913 and 7978

Unattended file not found during installation of SLES on uEFI systems

When using Linux Scripting Toolkit to install SLES on a uEFI based system, the installation task might be unable to find the answer file, causing the installation to attempt to continue in manual mode.

To resolve this issue, perform these steps:

1. Edit the workflow for your installation.
2. In the OS install section of the workflow, add `brokenmodules=usb_storage` to the optional kernel parameters.
3. Save the workflow.
4. Create bootable media from the workflow and perform the installation.
5. After the installation is complete, edit the file `/etc/modules.d/blacklist`, it is recommended that you make a copy of this file before editing it.
6. Remove the line `blacklist usb_storage`.

This limitation affects the following systems:

- System x3400 M2, types 7836 and 7837
- System x3500 M2, type 7839
- System x3550 M2, types 7946 and 4198
- System x3650 M2, types 7947 and 4199
- System x iDataPlex dx360 M2 types 7321, 7323 and 6380
- BladeCenter HS22, types 7870 and 1936

Altiris Clone of Image Fails on SLES 11

Altiris RapiDeploy does not support the cloning of SLES 11 for capturing or deploying images, which causes this job to fail. To avoid this problem, use the `-raw` parameter for the **Capture Linux Image** and **Deploy Linux Image** jobs.

You must add `-raw` to the **Additional Parameters** line of both the **Capture Linux Image** and **Deploy Linux Image** jobs. When using this parameter, the operation will take significantly longer. You must also ensure that enough free space is available on the target location to save the images.

VMware 4 installations using Altiris fail

The `computername` variable used in VMware installations must contain only letters, digits, and hyphens. If any other characters are detected, the installation will halt.

By default, the Altiris console supplies the `compname` variable used by the console as the `computername` variable for VMware installations. Therefore to successfully install VMware using Altiris, you must ensure that the `compname` variable used by the Altiris console is allowable as the VMware `computername` variable.

ServeRAID BR10i adapter not supported on iDataPlex dx360 M2 with 12 Bay Storage Chassis (Machine type 7321)

The ServeRAID BR10i adapter is not supported on the iDataPlex dx360 M2 with 12 Bay Storage Chassis, machine type 7321.

RAID configuration fails for LSI SATA RAID

When performing RAID configuration to configure an LSI 1064/1064e SATA controller, you might receive error code 7 or 11. This error is caused when the `cfggen` utility is unable to remove or create a configuration on SATA drives larger than 250GB.

To avoid this problem, remove any logical volumes including RAID arrays on the adapters using the Ctrl+C menu on system POST prior to using Linux Scripting Toolkit.

Default Fibre Configurations not supported on Emulex Fibre HBAs

The Target WWNN, Target WWPN and LUN number on the Fibre HBA Toolkit variables need to be set to configure the Primary, Alternate 1, Alternate 2 and Alternate 3 boot device settings. The default settings will not work on Emulex Fibre HBA adapters.

All values are case sensitive. You must ensure that the configured values are identical to the adapter values with regard to case.

ASU configuration fails for Load Defaults

When performing ASU Configuration to load the system defaults, you might receive an error code of 45. This error is caused when the ASU utility is unable to load defaults for the `ISCSI.InitiatorName` setting. This limitation affects the following systems:

- System x3200 M3, types 7327 and 7328
- System x3250 M3, types 4251, 4252, and 4261
- System x3400 M2, types 7836 and 7837
- System x3500 M2, type 7839
- System x3550 M2, types 7946 and 4198
- System x3650 M2, types 7947 and 4199
- System x iDataplex dx360 M2 types 7321, 7323 and 6380
- BladeCenter HS22, types 7870 and 1936

To avoid this problem, create a new `asu.ini` file with these contents:

```
loaddefault uefi
loaddefault SYSTEM_PROD_DATA
loaddefault BootOrder
loaddefault IMM
```

VMware ESX 4 installation requires a minimum of 4 GB of memory

When performing an installation of VMware ESX 4, ensure that the target system has a minimum of 4 GB of memory.

PXE will not bind to port 4011 on a Windows 2008 server

When Altiris Deployment Server is installed on a Windows 2008 server, you might encounter this problem with PXE-E55 on the target system. Instructions for avoiding the problem are included in this Altiris knowledge base article.

PXE will not bind to port 4011 on a Windows 2008 server

When Altiris Deployment Server is installed on a Windows 2008 server, you might encounter this problem with PXE-E55 on the target system. Instructions for avoiding the problem are included in this Altiris knowledge base article.

Altiris clone of image only supports Ext2/Ext3 file system

Altiris RapiDeploy only supports the cloning of the ext2/ext3 file system, which means that the job "Capture Linux Image" or "Capture RAID Configuration and Linux Image" may fail for a Linux operating system with a non-ext2/ext3 file system.

Altiris deployment of clone image fails on SLES 10.x and 11.x

When using the Toolkit for Altiris to run either "Deploy Linux Image" or "Deploy RAID Configuration and Linux Image" to deploy SUSE Linux Enterprise Server 10.x or 11.x, you may see an error message similar to the following:

```
Waiting for device /dev/disk/by-id/xxxxxxx to appear:
..... not found
--exiting to /bin/sh
```

To avoid the problem, before running "Capture Linux Image" or "Capture RAID Configuration and Linux Image" to capture the SUSE Linux Enterprise Server 10.x or 11.x, you need to boot into the system that you want to capture and perform the following steps:

1. Call up **YaST2 > System > Partitioner**.
2. Type `yast bootloader` (or `yast2 bootloader` for the graphical version).
3. For every boot menu entry, change the root device to match what was put into `/etc/fstab` by YaST Partitioner.
4. Reboot the system to see if it boots properly. If so, the system should be ready for cloning.

If you would rather make the changes manually instead of by using YaST, follow these steps:

1. Open the files `/etc/fstab` and `/etc/zipl.conf` in a text editor. You must do this from the root directory.
2. Replace all entries, such as `root=/dev/disk/by-id/[partition]` with `root=/dev/disk/by-path/[partition]`. If you are unsure about the existing partitions, use the `ls -l` command to view the directory entries in `/dev/disk/by-path`.
3. After making the changes to `/etc/fstab` and `/etc/zipl.conf`, run the commands `mkinitrd` and `zipl`. If you do not, the system will not boot.

For more details, refer to the Novell Knowledgebase document 3580082 at <http://www.novell.com/support/kb/doc.php?id=3580082>.

SLES11.3 Deployment in uEFI mode not supported

The current ServerGuide Scripting Toolkit, Linux Edition for Altiris supports the deployment of SUSE Linux Enterprise Server 11 SP 3 (SLES11.3) in legacy mode only; it does not support the deployment of SLES11.3 in uEFI mode.

VMWARE ESX requires that NUMA system memory be balanced

VMware installations may fail to load the VMkernel when Non-Uniform Memory Access (NUMA) is enabled and each processor does not have memory in its adjoining memory banks. For more information on this problem, see RETAIN tip H1974.

VMware ESX Server 4.1 installation hangs at "Starting vmkernel initialization"

When installing VMware ESX Server 4.1 on a system with a MAX5 memory expansion module, the installation might hang on this screen. This issue can occur on the following systems:

- BladeCenter HX5, type 7872
- System x3690 X5, types 7148, 7149
- System x3850 X5, type 7145

To avoid this problem, add the kernel parameter **allowInterleaveNUMAnodes=TRUE** to the kickstart file and the Altiris job.

To update the kickstart file, follow these steps:

1. Navigate to `..sgdeploy\sgtklinux\altiris\answerfiles\` and make a copy of the existing `esx4.ks` kickstart file.
2. Edit `esx4.ks` to modify the line:

```
bootloader --location=mbr
```

to be:

```
bootloader --location=mbr --append="allowInterleavedNUMAnodes=TRUE"
```

To modify the Altiris job, follow these steps:

1. In the **Customize Job Variables** task of the Altiris job, click **Modify** to edit the task.
2. Edit the line:

```
#TK_DISTRO_KPARAM=
```

by uncommenting it and adding the kernel parameter

allowInterleavedNUMAnodes=TRUE so that the line is now:

```
TK_DISTRO=KPARAM="allowInterleavedNUMAnodes=TRUE"
```

3. Save your changes.

uEFI operating system installations not booting from hard drive

During native uEFI operating system installations, the target system might fail to boot from the hard drive after Linux Scripting Toolkit processes are complete. This can occur if the target system does not automatically boot the efi file (`bootx64.efi` for RHEL6, or `elilo.efi` for SLES11) from the drive.

The solution to this problem is dependent on the operating system. Please consult the operating system information for instructions on adding a new boot option entry for the efi file.

For example, to correct this problem on most IBM systems, you can create a new boot entry for efi file and continue the installation using that option. Follow the steps below to create a new boot entry for the efi file:

1. Power on the system, and, press **F1** to enter setup.
2. Select **Boot Manager**.
3. Select **Add Boot Option**.
4. Select the boot entry which includes string "*.efi"
5. Input the description as `OS_Install` and select **Commit Changes**.

Follow the steps below to continue the installation:

- Power on the system, and press **F1** to enter setup.
- Select **Boot Manager**.
- Select **Boot from File**.
- Select the GUID Partition Tables (GPT) System Partition with the name `OS_Install`.
- Select **EFI**.
- Select **Boot**.
- Select **efi file**.

Note: If the installation completes and the system does not boot to the operating system, go to the **Start Options** section of the setup menu and select the boot entry for the operating system

Appendix A. Getting help and technical assistance

If you need help, service, or technical assistance or just want more information about IBM products, you will find a wide variety of sources available from IBM to assist you. This appendix contains information about where to go for additional information about IBM and IBM products, what to do if you experience a problem with your System x or IntelliStation® system, and whom to call for service, if it is necessary.

Before you call

Before you call, make sure that you have taken these steps to try to solve the problem yourself:

- Check all cables to make sure that they are connected.
- Check the power switches to make sure that the system is turned on.
- Use the troubleshooting information in your system documentation, and use the diagnostic tools that come with your system. Information about diagnostic tools is in the *Hardware Maintenance Manual and Troubleshooting Guide* on the IBM xSeries Documentation CD or in the *IntelliStation Hardware Maintenance Manual* at the IBM Support Web site.
- Go to the IBM Support website at <http://www.ibm.com/pc/support/> to check for technical information, hints, tips, and new device drivers or to submit a request for information.

You can solve many problems without outside assistance by following the troubleshooting procedures that IBM provides in the online help or in the publications that are provided with your system and software. The information that comes with your system also describes the diagnostic tests that you can perform. Most xSeries and IntelliStation systems, operating systems, and programs come with information that contains troubleshooting procedures and explanations of error messages and error codes. If you suspect a software problem, see the information for the operating system or program.

Using the documentation

Information about your IBM System x or IntelliStation system and preinstalled software, if any, is available in the documentation that comes with your system. That documentation includes printed books, online books, readme files, and help files. See the troubleshooting information in your system documentation for instructions for using the diagnostic programs. The troubleshooting information or the diagnostic programs might tell you that you need additional or updated device drivers or other software. IBM maintains pages on the World Wide Web where you can get the latest technical information and download device drivers and updates. To access these pages, go to <http://www.ibm.com/pc/support/> and follow the instructions. Also, you can order publications through the IBM Publications Ordering System at <http://www.elink.ibm.link.ibm.com/public/applications/publications/cgibin/pbi.cgi>.

Getting help and information from the World Wide Web

On the World Wide Web, the IBM website has up-to-date information about IBM System x and IntelliStation products, services, and support. The address for IBM System x information is <http://www.ibm.com/systems/x/>. The address for IBM IntelliStation information is <http://www.ibm.com/systems/intellistation/pro/index.html>.

You can find service information for your IBM products, including supported options, at <http://www.ibm.com/support><http://www.ibm.com/support>.

Software service and support

Through IBM Support Line, you can get telephone assistance, for a fee, with usage, configuration, and software problems with xSeries servers, IntelliStation workstations, and appliances. For information about which products are supported by Support Line in your country or region, go to <http://www.ibm.com/services/sl/products/>.

For more information about Support Line and other IBM services, go to <http://www.ibm.com/services/>, or go to <http://www.ibm.com/planetwide/> for support telephone numbers. In the U.S. and Canada, call 1-800-IBM-SERV (1-800-426-7378).

Hardware service and support

You can receive hardware service through IBM Services or through your IBM reseller, if your reseller is authorized by IBM to provide warranty service. Go to <http://www.ibm.com/planetwide/> for support telephone numbers, or in the U.S. and Canada, call 1-800-IBM-SERV (1-800-426-7378).

In the U.S. and Canada, hardware service and support is available 24 hours a day, 7 days a week. In the U.K., these services are available Monday through Friday, from 9 a.m. to 6 p.m.

Appendix B. Notices

This book contains the following notices designed to highlight key information:

- **Note:** These notices provide important tips, guidance, or advice.
- **Important:** These notices provide information or advice that might help you avoid inconvenient or difficult situations.
- **Attention:** These notices indicate possible damage to programs, devices, or data. An attention notice is placed just before the instruction or situation in which damage could occur.

Edition notice

© COPYRIGHT INTERNATIONAL BUSINESS MACHINES CORPORATION, 2013. All rights reserved.

U.S. Government Users Restricted Rights — Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Trademarks

The following terms are trademarks of International Business Machines Corporation in the United States, other countries, or both:

BladeCenter	IntelliStation
e-business logo	ServeRAID
eServer	ServerGuide
IBM	ServerProven
IBM (logo)	xSeries
TotalStorage	System x

Intel and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a trademark of The Open Group in the United States, other countries, or both.

Other company, product, or service names might be trademarks or service marks of others.

Important notes

When referring to processor storage, real and virtual storage, or channel volume, KB stands for approximately 1000 bytes, MB stands for approximately 1 000 000 bytes, and GB stands for approximately 1 000 000 000 bytes.

When referring to hard disk drive capacity or communications volume, MB stands for 1 000 000 bytes, and GB stands for 1 000 000 000 bytes. Total user-accessible capacity might vary depending on operating environments.

IBM makes no representation or warranties regarding non-IBM products and services that are ServerProven, including but not limited to the implied warranties of merchantability and fitness for a particular purpose. These products are offered and warranted solely by third parties.

IBM makes no representations or warranties with respect to non-IBM products. Support (if any) for the non-IBM products is provided by the third party, not IBM.

Some software can differ from its retail version (if available) and might not include user manuals or all program functionality.

Index

C

customizing deployment scenarios 23

D

deployment scenarios

 customizing 23

 quick start 13

N

notes, important 93

P

policy file parameter

 AppliesTo.n 56

 Array_Defaults 58

 Array_Mode 57

 Array.letter 58

 Hotspares 59

 Logical_Defaults 60

 Logical_Mode 59

 Logical.num 60

 policy.name 56

 ReadAhead 57

 RebuildRate 57

 StripeSize 57

Q

quick start 13

T

toolkit utilities, using 41, 42, 70

trademarks 93

U

using, toolkit utilities 41, 42, 70



Printed in USA