

Tivoli Provisioning Manager for Images
Version 7.1.1.14

Troubleshooting and support guide



Tivoli Provisioning Manager for Images
Version 7.1.1.14

Troubleshooting and support guide



Contents

Troubleshooting and support 1

Troubleshooting a problem	1
Searching knowledge bases	3
Getting a fix	4
Contacting IBM Support	4
Error messages	5
System profile error messages	5
OS deployment error messages	5
Exporting the log files in Common Base Event (CBE) format	7
Problems and limitations	8
ISA Compatibility.	8
Windows Service Troubleshooting for OS deployment server	8
Limitations	9
Problems with the kernel-free in BIOS mode	11
Problems with the kernel-free in UEFI mode	11
Problems with the OS deployment server and the web interface extension	12
Problems provisioning Windows operating systems.	15

Problems provisioning Linux and UNIX operating systems	22
Problems provisioning VMWare ESX operating systems.	31
Problems after a deployment	32
Problems while redeploying a profile on the same target	33
Problems with a multiserver architecture	36
Problems with virtual images and migration	39
Problems with booting, network, and file transfers	40
Problems with the database and the database gateway	45
Known limitations applicable to the current release of the Toolkit	48
Miscellaneous problems	49
Conventions	57
Notices	58

Troubleshooting and support

To help you understand, isolate, and resolve problems with your IBM® software, the troubleshooting and support information contains instructions for using the problem-determination resources that are provided with your IBM products.

To resolve a problem on your own, you can find out how to identify the source of a problem, how to gather diagnostic information, where to get fixes, and which knowledge bases to search. If you need to contact IBM Support, you can find out what diagnostic information the service technicians need to help you address a problem.

Troubleshooting a problem

Troubleshooting is a systematic approach to solving a problem. The goal of troubleshooting is to determine why something does not work as expected and explain how to resolve the problem.

The first step in the troubleshooting process is to describe the problem completely. Problem descriptions help you and the IBM Support person know where to start to find the cause of the problem. This step includes asking yourself basic questions:

- What are the symptoms of the problem?
- Where does the problem occur?
- When does the problem occur?
- Under which conditions does the problem occur?
- Can the problem be reproduced?

The answers to these questions typically lead to a good description of the problem, and that is the best way to start down the path of problem resolution.

What are the symptoms of the problem?

When starting to describe a problem, the most obvious question is "What is the problem?" This might seem like a straightforward question; however, you can break it down into several more-focused questions that create a more descriptive picture of the problem. These questions can include:

- Who, or what, is reporting the problem?
- What are the error codes and messages?
- How does the system fail? For example, is it a loop, hang, performance degradation, or incorrect result?
- What is the business impact of the problem?

Where does the problem occur?

Determining where the problem originates is not always easy, but it is one of the most important steps in resolving a problem. Many layers of technology can exist between the reporting and failing components. Networks, disks, and drivers are only a few of the components to consider when you are investigating problems.

The following questions help you to focus on where the problem occurs to isolate the problem layer:

- Is the problem specific to one platform or operating system, or is it common across multiple platforms or operating systems?
- Is the current environment and configuration supported?

Remember that if one layer reports the problem, the problem does not necessarily originate in that layer. Part of identifying where a problem originates is understanding the environment in which it exists. Take some time to completely describe the problem environment, including the operating system and version, all corresponding software and versions, and hardware information. Confirm that you are running within an environment that is a supported configuration; many problems can be traced back to incompatible levels of software that are not intended to run together or have not been fully tested together.

When does the problem occur?

Develop a detailed timeline of events leading up to a failure, especially for those cases that are one-time occurrences. You can most easily do this by working backward: Start at the time an error was reported (as precisely as possible, even down to the millisecond), and work backward using the available logs and information. Typically, you need to look only as far as the first suspicious event that you find in a diagnostic log; however, this is not always easy to do and takes practice. Knowing when to stop looking is especially difficult when multiple layers of technology are involved, and when each has its own diagnostic information.

To develop a detailed timeline of events, answer these questions:

- Does the problem happen only at a certain time of day or night?
- How often does the problem happen?
- What sequence of events leads up to the time that the problem is reported?
- Does the problem happen after an environment change, such as upgrading or installing software or hardware?

Responding to questions like this helps to provide you with a frame of reference in which to investigate the problem.

Under which conditions does the problem occur?

Knowing which systems and applications are running at the time that a problem occurs is an important part of troubleshooting. These questions about your environment can help you to identify the root cause of the problem:

- Does the problem always occur when the same task is being performed?
- Does a certain sequence of events need to occur for the problem to surface?
- Do any other applications fail at the same time?

Answering these types of questions can help you explain the environment in which the problem occurs and correlate any dependencies. Remember that just because multiple problems might have occurred around the same time, the problems are not necessarily related.

Can the problem be reproduced?

From a troubleshooting standpoint, the *ideal* problem is one that can be reproduced. Typically, problems that can be reproduced have a larger set of tools or procedures at your disposal to help you investigate. Consequently, problems that you can reproduce are often easier to debug and solve. However, problems

that you can reproduce can have a disadvantage: If the problem is of significant business impact, you do not want it to recur. If possible, re-create the problem in a test or development environment, which typically offers you more flexibility and control during your investigation.

- Can the problem be recreated on a test system?
- Are multiple users or applications encountering the same type of problem?
- Can the problem be recreated by running a single command, a set of commands, or a particular application, or a stand-alone application?

Searching knowledge bases

You can often find solutions to problems by searching IBM knowledge bases. Learn how to optimize your results by using available resources, support tools, and search methods and how to receive automatic updates.

Available technical resources

In addition to this information center, the following technical resources are available to help you answer questions and resolve problems:

- Tivoli® Provisioning Manager for OS Deployment Support Web site
- Tivoli Redbooks® Domain
- Tivoli support communities (forums and newsgroups)

Searching with support tools

The following tools are available to help you search IBM knowledge bases:

- **IBM Support Assistant (ISA)** is a free software serviceability workbench that helps you resolve questions and problems with IBM software products. Instructions for downloading and installing the ISA can be found on the ISA Web site: www.ibm.com/software/support/isa/
- **IBM Software Support Toolbar** is a browser plug-in that provides you with a mechanism to easily search IBM support sites. You can download the toolbar at: www.ibm.com/software/support/toolbar/.

Search tips

The following resources describe how to optimize your search results:

- Searching the IBM Support Web site
- Using the Google search engine
- Search the information center for more information about error messages that you encounter while using the product.

Receiving automatic updates

You can receive automatic updates in the following ways:

- **My support.** To receive weekly e-mail notifications regarding fixes and other support news, follow these steps:
 1. Go to the IBM Software Support Web site at www.ibm.com/software/support/.
 2. Click **My support** in the upper-right corner of the page under **Personalized support**.

3. If you have already registered for My support, sign in and skip to the next step. If you have not registered, click **Register now**. Complete the registration form using your e-mail address as your IBM ID and click **Submit**.
 4. Click **Edit profile**.
 5. Click **Add products** and choose a product category; for example, **Software**. A second list is displayed.
 6. In the second list, select a product segment; for example, **Data & Information Management**. A third list is displayed.
 7. In the third list, select a product subsegment, for example, **Databases**. A list of applicable products is displayed.
 8. Select the products for which you want to receive updates.
 9. Click **Add products**.
 10. After selecting all products that are of interest to you, click **Subscribe to email** on the **Edit profile** tab.
 11. Select **Please send these documents by weekly email**.
 12. Update your e-mail address as needed.
 13. In the **Documents list**, select the product category; for example, **Software**.
 14. Select the types of documents for which you want to receive information.
 15. Click **Update**.
- **RSS feeds**. For information about RSS, including steps for getting started and a list of RSS-enabled IBM Web pages, visit www.ibm.com/software/support/rss/

Getting a fix

A product fix might be available to resolve your problem.

You can get fixes by following these steps:

1. Obtain the tools required to get the fix.
2. Determine which fix you need.
3. Download the fix. Open the download document and follow the link in the "Download package" section.
4. Apply the fix. Follow the instructions in the "Installation Instructions" section of the download document.

Contacting IBM Support

IBM Support provides assistance with product defects.

Before contacting IBM Support, your company must have an active IBM software maintenance contract, and you must be authorized to submit problems to IBM.

Complete the following steps to contact IBM Support with a problem:

1. Define the problem, gather background information, and determine the severity of the problem.
2. Gather diagnostic information.
3. Submit your problem to IBM Support in one of the following ways:
 - Using IBM Support Assistant (ISA)
 - Online: Click the **Report problems** tab on the IBM Software Support site: <http://www.ibm.com/software/support/probsub.html>

If the problem you submit is for a software defect or for missing or inaccurate documentation, IBM Support creates an Authorized Program Analysis Report (APAR). The APAR describes the problem in detail. Whenever possible, IBM Software IBM Support provides a workaround that you can implement until the APAR is resolved and a fix is delivered. IBM publishes resolved APARs on the IBM Support Web site daily, so that other users who experience the same problem can benefit from the same resolution.

Error messages

System profile error messages

Error messages that can occur when trying to create or to update a system profile.


Sysprep is not installed

This message will appear when you are creating a Windows system profile for cloning. Before creating a system profile, you need to run the Sysprep tool.

Unknown/Unsupported OS

This warning message occurs when creating a system profile for cloning. The operating system installed cannot be recognized or is not supported. The current version of Tivoli Provisioning Manager for Images supports the operating systems listed in the Tivoli Provisioning Manager for Images Installation Guide.

Profiles are too old

The OS deployment server is not able to update Windows Vista/2008/7/2012 64-bit unattended setup profiles created with version 7.1.1.1 and lower of the product if the corresponding WinPE2 64-bit ramdisk software module is not present on the server. In this case the icon of the profile is changed to  , the warning message **Profiles are too old.** is issued , and the profile cannot be deployed anymore. To solve this issue:

1. Create a new system profile.
2. Copy the configurations from the old system profile to the new system profile.
3. Delete the old system profile.

OS deployment error messages

Error messages that can occur on a target during a deployment are displayed in a red panel, in the center of the screen, and are logged to the ODBC database.

SoftwareProfile and SoftwareItem tables must use the same ODBC source

This message should never appear with a standard OS configuration. It will appear if you split the *SoftwareItem* and *SystemProfile* tables into two different ODBC sources.

Invalid destination folder for software copy/system snapshot

The destination folder that you specified during the software module creation does not exist.

Unexpected end of deployment job

One of the required tasks failed during the deployment.

You are not authorized to use this machine (off-line)

This message appears when the process of authentication fails. The cause might be that the network is down.

There is no known OS configuration for this target

You are running without being connected to the network and the database entry for the target to deploy does not contain a valid OS configuration.

This OS configuration was not intended ...

The deployment scheme has the setting **Never edit parameters**. The target is not the same model as the system profile deployed.

No entry found in the BOM for this target

There is no entry in the Bill of Material table (no target definition) matching the target computer MAC address, UUID, and serial number and the deployment settings have been set to disable manual edition of the Bill of Material.

No system partition has been defined ...

This error should never occur unless you have tampered with the definition of a system profile. It results from a system profile definition that has no bootable partition (OSPart is zero in the database).

Invalid Software Item in the database

This error should never occur, unless you have tampered with the definition of software items. It results from an unknown software module type.

Cannot process ... software items in pass zero

A floppy-disk or partition software module is scheduled for use in pass zero, conflicting with the Sysprep process. To avoid this error, either schedule these software items with negative pass numbers (before Sysprep) or with positive pass numbers (after Sysprep).

Cannot process ... software items before pass zero

A software module that involves writing to the operating system partition is used before pass zero, when the operating system partition is formatted. To avoid this error, always schedule these software items with a positive or zero pass number.

Required file has not been enumerated

This AutoCD-specific error message should never occur. A CD-set has not been generated correctly, probably because of an error of the program. If the message occurs, send a report to your reseller, with a copy of the CD-set.

There is not enough space in partition ... to download the images

A system profile partitioning scheme is not compatible with Tivoli Provisioning Manager for Images. The hard disk partition scheme must be created so that the sum of the unpartitioned disk space and of the free space in the last partition is large enough to store all compressed partition and software images.

System setup has not been properly completed

This error results from a previous serious error in the Sysprep process, that has prevented the mini-setup to complete (or even to start).

Connection refused ... in sql.rbc

The TCP to ODBC gateway service that should be running on the OS deployment server is not accepting connections. This service is typically automatically started when the provisioning server service is started. Check (using the service manager) that the TCP to ODBC Gateway service is installed and running on the computer hosting the provisioning server.

Network not initialized

This error results from an abruptly stopped deployment, followed by a hard-disk boot that tries to restart the deployment. However, because the computer has not been started on the network, this is not possible. To restart the deployment, reboot the computer on network boot.

This computer has been interrupted during a deployment

This message appears (on a black background) when you reboot on the hard disk after a stopped deployment (typically because of an error or to the user pressing **Cancel**). The deployment was not completed, and must be restarted because the operating system is not fully installed.

Partitions do not fit on this hard disk

The system profile to be deployed on this target is bigger than the size of its hard disk. Alternately, a protected partition on your disk might not have enough space on the disk for the current profile.

Fatal error, No hard disk detected!

This message appears if you try to deploy a target without hard disk.

Cannot access installation media

This message might appear during the deployment of a SLES 10 profile on a PowerPC target. It indicates that the switch did not reestablish the link to the network card in a timely way. Select **OK** or **Retry** to resume the deployment.

Exporting the log files in Common Base Event (CBE) format

Symptoms

In case you want to export the log files in Common Base Event (CBE) format and analyze them using PD tools, use the following procedure.

Resolving the problem

Export the log files using the following command:

```
rbagent rad-exportlogs destination=destination [mode=mode  
[compress]] [activityid=activityid]
```

where:

destination

Is the path where the common base event must be stored.

mode The following modes are available:

- mode 1: All the logs are placed in destination.(log | trc).xml mode.
- mode 2: All the logs are placed in the destination directory (only one file for all tasks).
- mode 3: All the logs are placed in the destination directory (one file for each task).
- mode 4: All the logs are placed in the destination directory in raw format.

compress

Is the option to compress the resulting file, only available if you choose mode 4.

activityid

Is the option to export only the activity with the specified id.

Note: The modes 1, 2, and 3 export the log files in CBE format.

Problems and limitations

This section provides troubleshooting guidance and information about product limitations for operating system deployment.

ISA Compatibility

From version 5.1.1 onwards, Tivoli Provisioning Manager for OS deployment is compatible with IBM Support Assistant (ISA). ISA can help you collect logs needed for troubleshooting.

To collect log files with the ISA, you need to have ISA and the Tivoli Provisioning Manager for OS deployment plugin installed on the provisioning server.

For more information about ISA and to download it, visit <http://www.ibm.com/software/support/isa/>.

During the data collection process, ISA asks you for the Tivoli Provisioning Manager for OS deployment installation path to know where the log files are located. The collected log files are then compressed into a jar file. The path and name of the compressed file are given by ISA at the end of the process. On Windows, it is located at C:/Program Files/ISA and ESA/IBM Support Assistant/space/logs/. The file name starts with isaCollector_SS3HLM.

Windows Service Troubleshooting for OS deployment server

If your OS deployment server does not work correctly, or you suspect that something is wrong, you have several options to collect debugging information from the OS deployment server:

- If the service works and the server is reachable, use the web interface and click **Server status/installation check**. If you can read the summary, the OS deployment server is running and you can check here for possible errors.
- If your service works but you want more debugging information, you can also check the **Server log files** where all log types are displayed. The log verbosity level can be increased if necessary using **Server parameters/Configuration**.
- If the web interface cannot contact the OS deployment server, check that the Tivoli Provisioning Manager for Images service and processes are running:
 - Windows: Use the service manager
 - Linux/FreeBSD/OS X: type `ps aux | grep rembo`
 - Solaris: type `ps -elf | grep rembo`
- Check the Windows Event Viewer. The OS deployment server logs unrecoverable errors messages into *event manager*.
- If the service does not start, run `rembo.exe` from the command-line with the following options: `rembo -d -v 4`. This will run the OS deployment server as a console application, with all debugging output redirected to your command window. You can increase the debug level (the `-v` parameter) to 6 for maximum detail. See “OS deployment server command-line options” on page 9 for more on command line arguments. If Tivoli Provisioning Manager for Images complains about an error in the configuration, reinstall a fresh copy of Tivoli Provisioning Manager for Images. If the error message is related to your network configuration, try to fix your network configuration and run the server again. In particular, you must change the Interfaces global parameter if your computer has more than one network interface.

If it still does not work, contact your IBM Software Support representative.

OS deployment server command-line options

rembo [-d] [-v loglevel] [-c configfile] [-cert rembokey]

- -d prints debug info to the standard output, does not run as daemon (do not detach)
- -v sets the verbosity level (default: 2)
- -c specifies the config file name (the default is rembo.conf)

The verbosity levels are defined as:

- 0 : no output
- 1 : log error messages only
- 2 : log error and warning messages
- 3 : log error, warning and info messages
- 4 : same as 3, but also log notice messages
- 5 : same as 4, with debug output
- 6 : same as 5, with network trace

Limitations

A list of known limitations and problems.

Tivoli Provisioning Manager for Images limitations

Note: Refer to the product release notes or most recent readme file provided with each fix pack for the most up-to-date information related to problems or limitations.

- When performing direct downloads of generated ISO or CAB files, the download might not work correctly with Internet Explorer 9 installed on a Windows 2008 host workstation. The same downloads will work with Internet Explorer 9 installed on a Windows 7 host workstation.
- When deploying a system profile against a Sony model, the last digit of the serial number is not reported in the Target Monitor. The serial number gets truncated to 15 characters.
- The **Boot on OS** function does not choose the operating system to run but only boots from the disk and is not able to understand if you deployed a multi-boot system.
- The deployment on KVM guests is not supported using the "Virtio" disk. You must change the device type to "IDE Disk" by removing the "Virtio" device.
- On some RAID controllers, the automatic RAID configuration might fail. In these cases, use the custom configurations.
- For workstations with an amount of memory slots bigger than four, the size of the memory reported on the inventory tab of the target details might be inaccurate.
- VMWare ESX 3.5 u5 profiles might cause network issues on HP DL585G2 hardware.
- The network might not be configured properly on HP DL580 G7 hardware after a successful Windows 2003 deployment even if the network driver is bound to the profile.
- SuSE Linux Enterprise Server (SLES) 11 SP1 operating system deployments are not supported on HP DL580 G7 workstations.
- 32-bit profiles are not supported on IBM x3850 x5 servers.

- Red Hat Enterprise Linux (RHEL) 5.5 32-bit operating system is not supported on IBM x3850 x5 servers. This operating system is supported only on x3850A servers.
- For software modules in pre-partitioning and pre-OS phases the following applies:
 - The current Web interface enables you to add a reboot within the "Software application order" popup dialog, even if it does not make sense to add a reboot between these software stages because, after a reboot, the deployment state is lost in both cases.
 - Pre-partitioning software is applied into the RAM. The deployment state is lost after the reboot.
 - For pre-OS software there is currently no cache and the deployment state is also lost after the reboot.
- To create one of the following system profiles, you must use the browser on a Linux workstation with the web interface extension installed:
 - SLES 10 SP3 32-bit
 - SLES 10 SP3 64-bit
 - SLES 10 SP4 32-bit
 - SLES 10 SP4 64-bit
- After the target is booted on the OS deployment server, if the target has two NIC cards, the inventory details for that target only show one entry for one NIC card.
- In the operating system deployment, LVM is not supported, if it spans multiple disks.
- Software snapshots are deprecated. They are not supported on Linux, and you can only restore previously created software snapshots on Windows.
- For Windows cloning, because of Sysprep limitations, it is not possible to change the administrator password during the deployment if the system profile contains a non-empty administrator password.
- For Windows golden master image, because of Sysprep limitations, it is not possible to change the administrator password during the deployment if the system profile contains a non-empty administrator password. This limitation only applies if you run Sysprep yourself manually.
- targets only supports incremental images on the primary OS partition. If you want to install software on a secondary partition, you must use software update packages with an unattended setup command line.
- targets supports only up to 3 primary partitions. If you have more than 3 primary partitions, transform the extra ones into extended or logical partitions.
- The following options are not available when deploying Windows 2000 on targets:
 - Advanced IP settings
 - Join a Windows domain on site
 - Add domain user to local admin group
 - Force user to put a new admin password
 - Create local account for user
- VMware ESXi 5.0 and 5.1 do not support Boot On OS and Green banner end action. They will default to other end action. Green banner performs a shutdown and Boot On OS performs a reboot.
- When performing Direct Migration or while capturing a profile or virtual image, the Microsoft System Reserved partition is captured but it is empty.

Database-related limitations

The hardware inventory records up to four memory slots only.

Problems with the kernel-free in BIOS mode

Inventory limitation in kernel-free mode

Symptoms

CD-DVD information is not captured in the inventory details when the target boots in kernel-free mode.

Resolving the problem

No workaround is currently available to solve this issue.

Boot on OS on IBM x3850 fails in kernel-free mode

Symptoms

You perform a deployment on an IBM x3850 hardware in kernel-free mode. The deployment is successful and it displays the green banner. But after clicking **Boot on OS**, the target does not boot on the operating system, it boots on the network, again in kernel-free mode.

Resolving the problem

To perform the **Boot on OS** on an IBM x3850 workstation, it is necessary to update its BIOS with the latest version 1.32.

Boot fails on some targets when deploying in kernel-free mode

Symptoms

When deploying on multiple targets (more than 20) in kernel-free mode, some targets might not be able to boot during the deployment. The following message is displayed on the target for several minutes:

Boot failed: press a key to retry, or wait for reset....

Resolving the problem

As a workaround, specify the network before the hard disk in the boot list. After the reboot, the task is able to be automatically recovered and no manual steps are needed.

Problems with the kernel-free in UEFI mode

Incomplete inventory in UEFI mode on IBM HS22

Symptoms

When you PXE boot in UEFI mode on IBM HS22, the PCIIO protocol is not initialized and the inventory is incomplete.

Resolving the problem

To perform the PXE boot in UEFI mode, it is necessary to update the BIOS with the version P9E155B or later.

Some UEFI targets fail to PXE boot

Symptoms

While PXE booting UEFI targets, you can encounter problems if the UEFI firmware version is out of date.

Resolving the problem

With the current UEFI firmware version, some UEFI targets do not process correctly DHCP option 43, they require option 66 and 67 to be set. For more information check the DHCP server configuration section in the Installation Guide.

Problems with the OS deployment server and the web interface extension

Many tasks suddenly fail

When many tasks are suddenly failing, there is a simple check that you can easily perform to check the state of your OS deployment server.

Symptoms

Many of the tasks you try to perform fail without apparent reason.

Causes

Some of the services or daemons of the OS deployment server are not properly started.

Diagnosing the problem

Go to **Server > OS deployment > Target Monitor**. . You should have at least the following two entries:

unknown

This entry refers to the OS deployment server itself.

IP of the OS deployment server

This IP address is present when the Web interface extension is running on the OS deployment server.

One of these entries, or both, does not appear in the Target Monitor anymore.

Resolving the problem

- If **unknown** has disappeared, restart the OS deployment server.
- If the IP address of the OS deployment server has disappeared, make sure the web interface extension is installed on the OS deployment server and restart it.

After these steps, both **unknown** and the IP address of the OS deployment server should be present and your OS deployment server should work again.

Failures in creating or deploying a profile or software module

When you cannot create or deploy profiles or software modules, OS deployment server might have a full temporary directory.

Symptoms

You cannot create or deploy profiles or software modules.

Causes

OS deployment server might have a full temporary directory.

Resolving the problem

Clean up the temporary directory inside the OS deployment server installation directory on a regular basis. The temporary directory is `TPMfOSD_INST_PATH\temp`.

OS deployment server stops responding

The OS deployment server might stop responding if too many ports are already in use and it does not have any left for communication.

Symptoms

Under Windows, the OS deployment server stops responding without any apparent cause.

Causes

This might be due to the limited number of port and sockets available by default on Windows operating systems. Use of the Java™ API might cause to reach this limit.

Resolving the problem

You can try to solve this problem by allowing TCP to assign higher port numbers than the default 5000 and providing a smaller waiting time, in seconds, before TCP can release a closed connection. To do so:

1. Stop the OS deployment server
2. Edit the following registry keys and provide the values suggested

```
"HKLM/system/CurrentControlSet/Services/TcpIp/Parameters/MaxUserPort" = 65534
"HKLM/system/CurrentControlSet/Services/TcpIp/Parameters/TCPTimedWaitDelay" = 30
```
3. Restart you OS deployment server.

Unable to connect to server with KVM hypervisor

Cannot establish a server connection when using an environment that includes KVM machines.

Symptoms

When working with KVM machines, a server connection cannot be established.

Causes

When working with KVM machines, ensure you have a correctly bridged NIC.

Diagnosing the problem

If the following message is logged to the trace file, rbagent.trc, on the hypervisor, then the rbagent is unable to work with the IP:

```
#####
...
bind address already in use
HTTP SERVICE DISABLED on XXXX (ip)
...
#####
```

Resolving the problem

Ensure your physical NIC (eth0) is correctly bridged, usually to br0, otherwise, the rbagent running on the hypervisor cannot contact the OS deployment server during virtual infrastructure tasks.

Corrupted character display on AIX during setup

Symptoms

The characters displayed during product setup on AIX® are corrupted. Messages are unreadable.

Causes

Your locale is not UTF-8 compliant.

Environment

- Your operating system is AIX.
- Your locale is not UTF-8 compliant.

Diagnosing the problem

You are installing Tivoli Provisioning Manager for OS deployment on an AIX computer using `./setup`. The messages displayed by `./setup` contain corrupted characters which render the content of the messages unreadable. The language you are using to install the product has several locales and the one you currently using is not UTF-8 compliant.

Resolving the problem

There are two possible workarounds:

- Change the character encoding of your terminal session to UTF-8 and run setup again.
- Change your sessions parameters to an UTF-8 compliant language encoding and run setup again.

Server and web interface extension

The OS deployment server and the web interface extension cannot start properly on a Linux target which has its locale set to Turkish.

Symptoms

The OS deployment server or the web interface extension fails to start on Linux targets when the locale is set to Turkish.

An error with the following text might appear:

Unknown error (ERROR: column not found: @@id)

Causes

Tivoli Provisioning Manager for Images cannot work in a Turkish locale.

Environment

The OS deployment server or the web interface extension is installed on a target running Linux in a Turkish locale.

Resolving the problem

Set the locale to something else than Turkish, for example US English.

Tasks appear with erroneous dates and times in the web interface

Problem description

When using an Apache Derby database with an OS deployment server installed on Linux, tasks may appear with erroneous dates and times in the web interface.

Problem resolution

- Set the TZ variable. To do so
 1. Use the `tzselect` tool to help you find your appropriate time zone.
 2. Set the TZ environment variable to the appropriate value.
 3. Export the TZ variable in the Apache Derby environment.

Example for steps 2 and 3 in Switzerland:

```
TZ='Europe/Zurich'; export  
TZ
```

- Modify your Apache Derby startup script, if you have one. Include the `-Duser.timezone` command line option line in your script. For example, if you want to use UTC, type `-Duser.timezone=UTC`.

The web interface extension is not detected

Sometimes, the web interface extension is not correctly detected on UNIX and Linux OS deployment servers.

Symptoms

The web interface extension is correctly installed and is running, but it is not detected by the OS deployment server. A red icon for the web interface extension is displayed.

Causes

The OS deployment server is not listening to the correct interface and cannot therefore detect the web interface extension.

Resolving the problem

1. Open `rbagent.log` and find the last occurrence `Connect xx.xx.xx.xx -> yy.yy.yy.yy` where `xx.xx.xx.xx` and `yy.yy.yy.yy` are both IP addresses of the OS deployment server.
2. Take note of the `yy.yy.yy.yy`.
3. Edit `/etc/hosts`.
4. Move or add the line with the `yy.yy.yy.yy` before the line with the `localhost` interface (`127.0.0.1`).
5. Restart the OS deployment server daemon.

Now, the OS deployment server can resolve the host name properly and detect the web interface extension.

Problems provisioning Windows operating systems

Deploying Windows fails due to missing storage driver

Symptoms

When you deploy Windows operating system, an incorrect built-in mass storage driver is used by the operating system, and a blue screen displays when booting on the operating system.

Causes

In the operating system is included an incorrect driver that matches the storage device. Even if you bind the correct software module driver, the incorrect driver is taken.

Resolving the problem

To successfully deploy Windows, perform these steps:

1. Locate the `stpwndxxxx.1` file under the `TPMfOSD_DATA_DIR/global/partition` directory. This is the image file of the operating system you are deploying.
2. Rename it into `stpwndxxxx.wim`.
3. Mount it using Windows Automated Installation Kit (WAIK).
4. Remove the driver using WAIK.
5. Unmount the `stpwndxxxx.wim` file and rename it back into `stpwndxxxx.1`.

Alternatively, modify the image before creating the system profile.

Windows 2003 deployment failure on target with two hard disks

During Windows 2003 deployment, WinPE boots from a disk different than the one used by Preboot Execution Environment (PXE).

Symptoms

Unattended deployment of Windows 2003 fails with the following error:
Rembo MBR detected, and no partition is marked with the bootable flag

Resolving the problem

From BIOS on the target machine change the boot order of the hard disks to have the same boot order for WinPE and PXE.

Windows 2000/2003/2008/XP/Vista/7 reports that it has discovered a new device

Symptoms

In some cases, Windows 2000/2003/2008/XP/Vista/7 might report the detection of a new device and ask for a reboot after restoring an image using Tivoli Provisioning Manager for Images, even if the image was made on the same hardware.

Resolving the problem

There are two causes of a Windows redetection of hardware:

- Restoring an image on the exact same kind of hardware but on another computer. Some components (including the hard disk drive) include a unique serial number. This is not visible when deploying an image in Sysprep because Sysprep handles the redetection silently, but this can affect a restoration if Sysprep mini-setup had not been used during the creation of the image.
- A change in the partition size. Windows 2000/2003/2008/XP/Vista/7 stores information regarding the operating system partition layout in the registry, and might need a reboot if the partition has changed. This is typically not visible when deploying an image in Sysprep because Sysprep handles the redetection silently, but it can affect a plain restoration. In some cases, it can affect a typical deployment, if the operating system partition goes to the end of the disk, because Tivoli Provisioning Manager for Images needs to resize it temporarily to store its image files during the deployment. The workaround is to have another partition after the operating system partition, so that the operating system partition itself is not resized during the deployment.

Accessing target devices during deployment

Symptoms

During deployment on an IBM system x3350 the following message is displayed:
Please insert disk into drive.

Resolving the problem

Verify that all your devices can be accessed.

WinPE does not have the driver to handle the NIC

Symptoms

Windows Preinstallation Environment (WinPE) starts on the target and does not have the driver to handle the NIC.

Resolving the problem

Ensure that you install the driver, which is WinPE certified, for handling the NIC. Create the software module, related to the WinPE driver, for the target. From the driver grid, bind the driver to the software module. For more details, see Binding drivers to a WinPE deployment engine.

Windows Vista/2008/7/2012 ask for a product key during deployment

Problem description

In some circumstances, Windows Vista/2008/7/2012 ask for a product key during deployment of a cloning or cloning WIM system profile created with Sysprep in audit mode.

This problem occurs only if

- You are deploying a cloning or cloning WIM system profile of Windows Vista/2008/7/2012.
- The profile was created with Sysprep in audit mode.
- You do not have a Volume Licensing version of Windows Vista/2008/7/2012.

Problem resolution

This is known behavior of both Windows Vista/2008/7/2012 operating systems.

If the product key was provided in the OS configuration, you can leave the field empty and click to continue.

Windows Vista/2008/7/2012 prompts you for an Administrator user name during deployment

Problem description

Windows Vista/2008/7/2012 prompts you for an Administrator user name during deployment.

Problem resolution

If Windows Vista/2008/7/2012 prompts you for an Administrator user name, it is because Windows Vista/2008/7/2012 requires a new local account when starting for the first time. If you want to avoid being prompted for an Administrator user name, provide it in the OS configuration.

Drive letter incorrectly set

When deploying a virtual image or migrating a Windows operating system, the W drive is assigned a new drive letter.

Symptoms

On Windows operating system, one of your drives is set to letter W. After direct migration or the deployment of a virtual image, this drive is randomly assigned another letter on the receiving target. All other drives keep their original letter.

Causes

Drive letter W is not available on Windows operating systems installed by the product.

Domain name server is not set after a Windows XP 64-bit deployment

Symptoms

Domain name server (DNS) information is not correct after the deployment of a Windows XP 64-bit platform.

Causes

On Windows XP 64-bit platforms, if you create a profile, modify the target settings with advanced network settings and complete the following fields: IP address, gateway, DNS server and DNS domain. After the deployment all settings are correctly set, except for the DNS domain.

Deploying Windows on Linux partitions fails

Your Windows system profile deployment fails.

Symptoms

Your deployment of a Windows system profile containing Linux partitions (for example, Linux swap and EXT2) fails.

The error message indicates <ERR> Command [diskpart /s X:\windows\TEMP\4ECCC1F83AC358F3_1\partscript.txt] failed with exit code 5.

Causes

Since version 7.1.1, the product leverages WinPE for Windows operating system installation. WinPE does not recognize Linux partitions. This causes the deployment to fail.

Resolving the problem

The only solution is to edit the system profile to remove the Linux partitions.

Windows deployment fails with a stop error 7B on a blue screen

Your Windows deployment fails with a stop error 7B on a blue screen because of a missing disk driver.

Symptoms

Your Windows deployment fails with a stop error 7B on a blue screen, as shown in Figure 1 on page 19.



Figure 1. Stop error 7B in a blue screen

Causes

A disk driver is missing from your system profile.

Resolving the problem

You must figure out which driver is missing, bind it to the system profile, and deploy the system profile again.

To discover which driver is missing:

1. Go to **Server > OS deployment > System profiles > Profile details**.
2. Scroll down to the **Drivers handling** section.
3. If you are using the regular software rules, click **Switch to driver specific bindings mode**.
4. Click **Edit profile's driver bindings**.
5. Use the grid to locate missing mass storage device drivers and to bind suitable drivers.
6. Start your deployment again.

Timeout limit can be reached during OS Probe of Windows OS

The timeout limit can be reached during OS Probe of Windows OS with slow networks or slow machines.

Symptoms

When WinPE is downloaded on targets from server, it can take a long time for the image to be transferred on the network. If the transfer is very slow, some tasks like OS probe (detection of OS to capture) could hit a timeout error since default timeout for this task is 10 minutes.

Resolving the problem

To speed up the transfer of WinPE image from server to targets, the following parameters must be customized for WinPE task creation:

- TftpBlockSize
- TftpWindowSize

To customize these parameters, go to the table **Settings**, look for the record where Type = "WPECREAT" and get the value corresponding to the field **IniFile**, for example, radset-sttnsfrRTPSPEC9.ini. Edit the INI file, located in the <DATADIR>/global/schemes by changing the value assigned to the TftpBlockSize and TftpWindowSize , as shown in the following example:

```
[Task]
type="WPECREAT"

[Settings]
StartEnv="*"
TftpBlockSize="1432"
TftpWindowSize="4"
ScratchSpace="32"
```

It is not necessary to restart the server or the Web Interface extension, you can just recreate WinPE engines.

WinPE scratch-space is insufficient

The timeout limit can be reached during OS Probe of Windows OS with slow networks or slow machines.

Symptoms

While running the command startnet.cmd in the WinPE environment, the following message could appear: **There is not enough space on the disk.** After this message is displayed, a reboot could be performed. It is very likely that WinPE scratch-space is insufficient. Scratch space setting represents the amount of writable space available on the Windows PE system volume when booted in RAM disk mode. By default, Windows PE allocates 32 MB of memory to allocate to the Windows PE work space, valid values are 32, 64, 128, 256, and 512. The WinPE scratch space might provide insufficient space to copy all the required drivers software modules on the specific target model. To use the WinPE driver injection by the Binding Grid for a specific target model to inject a set of software modules requires an amount of space greater than 32Mb. The WinPE involved must be created with an appropriate quantity of scratch space. You can find examples about WinPE scratch space at <http://technet.microsoft.com/en-us/library/dd799244%28v=ws.10%29.aspx> and <http://technet.microsoft.com/en-us/library/hh824936.aspx>

Resolving the problem

To customize scratch-space, go to the table Settings, look for the record where Type = "WPECREAT" and get the value corresponding to the field **IniFile**, for example, radset-sttnsfrRTPSPEC9.ini. Edit the INI file, located in the <DATADIR>/global/schemes by changing the value assigned to the ScratchSpace , as shown in the following example :

```
[Task]
type="WPECREAT"

[Settings]
StartEnv="*"
TftpBlockSize="1432"
TftpWindowSize="4"
ScratchSpace="32"
```

It is not necessary to restart the server or the Web Interface extension, you can just recreate WinPE engines.

Update fails on WinPE engines with maintenance levels earlier than Fix Pack 12

Update fails on WinPE engines with maintenance levels earlier than Fix Pack 12

Symptoms

Starting from Fix Pack 12, the update of some old Windows PE engines might fail with the following warning message:

WinPe analysis detected a RISKYALLOC prone engine

. This message is reported in the `rbagent.trc` file if the server tried to update all engines at startup time, or displayed in the web interface if you used the “Update it” link in the **Advanced Features > Deployment Engines > Engine details** panel.

Causes

The update failure is caused by incorrect handling of the old format of Tivoli Provisioning Manager for OS Deployment Windows deployment engines when the update is automatically scheduled after you upgrade the product version. This results in memory allocation problems.

Resolving the problem

To force the update of WinPE proceed as follows:

1. Create the file '`<Base dir for OSD files>\global_NEVER_RISKYALLOCSAFE_`' to unblock the unsafe WinPE update.
2. Open a DOS Shell and set the following variable to modify the allowed threshold of risky memory allocations:
`set RBOHARMLESSSIZE=<new allowed threshold of risky memory allocation in bytes>`

The current value is set to 64MB. This value can be raised, just for the update procedure, to a higher value, for example 240MB (1024*1024*200 bytes):

```
set RBOHARMLESSSIZE=209715200
```

3. If the Web Interface Extension service is running, stop it via service manager.
4. Manually launch the following command in the same DOS shell of the Web Interface Extension from the `<TPMfOSD install directory>`, usually `C:\Program Files\Common Files\IBM Tivoli`:
`rbagent -d -v 4`
5. Run the WinPE update process using one of the following methods:
 - If your OSD server is installed on a Windows system, restart the OSD Server service and the local IBM Tivoli Web Interface extension service. The server detects which WinPE engines are not updated and runs the update again.
 - If your OSD server is installed on other operating systems, Select the **Advanced Features > Deployment Engine** page on the web interface, move the line focus on each WinPE engine to identify which engine displays, at the bottom of the panel, the warning message: “This deployment engine is not updated”. Then click on the **update it** link.

Messages similar to the following could be displayed in the console traces:

```
Crash dump: detected RBOHARMLESSIZE environment variable (value=209715200 bytes)
Crash dump skipped: requested 104857600 bytes
```

These messages are normal.

6. After the upgrade completes, stop the rbagent that you launched manually, and restart the Web interface Extension service via the service manager.

Problems provisioning Linux and UNIX operating systems

RPM packages are missing from the profile

Symptoms

A Red Hat Enterprise Linux deployment might fail due to missing RPM packages. This might happen because a "PRUNE" phase is performed, when creating the Linux system profile, to reduce the profile size by excluding unnecessary RPM packages. The selection by the user of different combinations of software groups when creating the profile might cause the "PRUNE" process to discard some needed RPM packages.

Resolving the problem

To solve this issue, create on the OS deployment server an empty file named:

```
RH_PROFILE_INCLUDE_ALL_RPM
```

under the *Base directory for OS deployment server files/global* directory; then recreate the system profile.

Deploying Linux on IBM T410 laptop fails

Symptoms

When you perform HTTP deployments of Linux operating system on IBM T410 laptop, the deployment fails.

Causes

Linux HTTP deployments, on IBM T410 laptop, using static IP address fails because the NIC wakes up slowly, after kickstart reconfigures it, and the installer starts before the NIC is completely ready.

Resolving the problem

To solve this issue, set the TCP/IP setting to Dynamic IP (DHCP).

Installing Linux on a target with multiple NIC fails

Symptoms

When you install Linux on a workstation with multiple NIC, the deployment might fail because the target cannot connect to the server.

Causes

If you do not switch to Advanced IP network settings, only the first NIC will be configured. The deployment fails if the target tries to connect to the server with a different NIC.

Resolving the problem

To solve this issue, switch to Advanced IP network settings when installing the Linux operating system on a workstation with multiple NIC, by selecting **Target Details > Common networking info > Switch to Advanced IP setting mode**.

Issue on Linux target with multiple NIC

Symptoms

On a Linux workstation with multiple network cards, if the network card connected to the OS deployment server is not the first one as described in "Installing Linux on a target with multiple NIC fails," the unattended Linux installation loops booting.

Resolving the problem

To solve this issue, select **Target Monitor > Target Details > Advanced network setting** of the network card connected to the OS deployment server, and specify the correct **Connection name** for that network card. It is the logical name that the Linux operating system assigns to the network device, such as eth3 or eth5.

Adding a new partition to RHEL unattended system profile fails

Symptoms

When you install an unattended Red Hat Enterprise Linux (RHEL) system profile, if the system profile has no mount point provided for a partition, during the HTTP deployment, the following error is displayed:

The following error was found while parsing the kickstart configuration file.
The following problem occurred on line 8 of the kickstart file:
the mount point "none" is not valid.

Resolving the problem

To solve this issue, always provide a mount point for partitions in the system profile partition layout or use non-HTTP deployments.

Deploying VMWare ESX/Linux fails

Symptoms

When you deploy Linux or VMWare ESX operating systems, the deployment might fail and the following error occurs:

Unable to boot Linux with the requested parameters

because the Linux kernel or initial ramdisk cannot be found.

Causes

The default device mapping is not working.

Resolving the problem

To solve this issue, manually install Linux or VMWare ESX using the CD/DVD on the workstation to identify the devices used and to use the correct device name in the profile.

Deploying Linux in kernel mode fails

Symptoms

When you deploy Linux using the standard installation (HTTP mode, kernel mode), during the deployment the following error occurs:

Unable to boot Linux with the requested parameters

Causes

The issue is caused by the fact that the /boot directory is on an ext4 file system.

Resolving the problem

To successfully deploy Linux using the /boot directory on an ext4 file system, perform the deployment in kernel-free mode.

Graphics Manager is not correctly configured after deploying SLES 11.1

Symptoms

When installing SuSE Linux Enterprise Server (SLES) 11.1 on an hyper-v guest with a graphical display manager (gnome, kdm), the installation ends successfully but you might experience some problems when loading the GUI panel, after the deployment is completed.

Causes

Resolving the problem

To solve this issue, choose one of the following options:

- Use the following custom `autoinst.xml` file:

```
<profile xmlns="http://www.suse.com/1.0/yast2ns";;  
xmlns:config="http://www.suse.com/1.0/configs">;  
  <x11>  
    <color_depth config:type="integer">16</color_depth>  
    <display_manager>gdm</display_manager>  
    <enable_3d config:type="boolean">true</enable_3d>  
    <enable_xgl config:type="boolean">true</enable_xgl>  
    <monitor>  
      <display>  
        <max_hsync config:type="integer">97</max_hsync>  
        <max_vsync config:type="integer">160</max_vsync>  
        <min_hsync config:type="integer">30</min_hsync>  
        <min_vsync config:type="integer">43</min_vsync>  
      </display>  
      <monitor_device>RYI 109-B5</monitor_device>  
      <monitor_vendor>RYI</monitor_vendor>  
    </monitor>  
    <resolution>1280x1024 (SXGA)</resolution>  
    <window_manager>gnome</window_manager>  
  </x11>  
</profile>
```

- Use the following custom `autoinst.xml` file:

```
<profile xmlns="http://www.suse.com/1.0/yast2ns";;  
xmlns:config="http://www.suse.com/1.0/configs">;  
  <x11>  
    <resolution>1280x1024 (SXGA)</resolution>  
  </x11>  
</profile>
```

- Run from the terminal the following command:

```
sax2 -ra; /usr/sbin/rcmxdm restart
```

to reconfigure the GUI.

fdisk utility fails during deployment

Symptoms

When you perform a deployment task, the `fdisk` utility might fail.

Causes

The issue is very likely related to the RAID disk partitioning.

Resolving the problem

To solve this issue, you have one of the following options:

- Perform a disk blanking and then run again the deployment.
- Log into Linux deployment engine and run the following command:

```
dd if=/dev/zero of=/dev/cciss/c0d0 count=1
```

Reboot the system and try again the deployment.

Deployment of SLES using Kernel Free method and http technique might fail

Symptoms

When deploying the SuSE Linux Enterprise Server (SLES) operating system, using the Kernel Free method and the HTTP technique, the deployment might fail and display the following error message:

Unable to find installation source.....

Causes

The client is unable to access the OS deployment server due to a missing network driver.

Resolving the problem

Ensure that you have the driver for your network card copied onto a CD-ROM or a USB drive. When the screen displays the error, press ENTER. Then insert the CD-ROM or the USB pen. From the red screen click **OK** -> **Kernel modules** -> **Add driver update** and select the USB pen or the CD-ROM.

You should receive a confirmation that the driver update was successful.

At the error screen press <alt> F9> to access the command window and configure the network card with your network specific parameters, for example:

```
ifconfig eth0 inet 192.168.1.1 broadcast 192.168.1.255 netmask 255.255.255.0 up
```

Verify if the configuration was successful by pinging the OS deployment server.

Return to the Web interface dialog where the error popped up by pressing <alt> <f1> or <f3>. Press ENTER to start the manual installation. Select **Start Installation or System**. Specify the installation options such as type of installation "network" , technique "http", and provide either the static or the dynamic IP address. You can now proceed with the installation.

Issue when deploying Linux profile on KVM guest

Symptoms

When deploying a Linux profile on a KVM guest, the deployment might fail and display the following error message:

```
autoinst.xml cannot be found
```

Resolving the problem

Ensure that the disk mapping of the profile has no other mappings checked except for /dev/hda.

SLES deployment on PowerPC switches to interactive

A SuSE Linux Enterprise Server deployment on a PowerPC® with multiple network cards switches to interactive installation when it is not registered with its first network card in the OS deployment server.

Symptoms

You are deploying a SuSE Linux Enterprise Server system profile on a PowerPC target with more than one network card. The deployment starts, searches for a DHCP address for its first network card and does not find it. The message *Sending DHCP request for <firstnetworkcard>* is displayed, where <firstnetworkcard> is the name of the first network card of the target. The target switches to interactive installation.

Causes

You have registered you target with a network card which is not the first.

Resolving the problem

1. Go to **OS Deployment > Profiles > Profile details > OS configuration details** .
2. Select **Unix**.

3. Click **Edit** to edit the **Fixed UNIX-specific prop..**
4. Update the field **Net boot device** to reflect the network card which was used when registering the target in the OS deployment server.
5. Once your configuration is updated with the appropriate network card, you can start the deployment again.

Firmware error during Linux deployment on PowerPC

SuSE Linux Enterprise Server (SLES) 10 deployment fails on PowerPC, after the files are copied on the target and the target is restarted. The installation process cannot continue because a firmware exception is caught.

Symptoms

When deploying a SLES 10 on PowerPC, the deployment starts correctly. The YAST installer installs files on the target and correctly restarts the target to continue the installation process with target configuration. However, when the target restarts once more, the installation process stops because a firmware exception is caught. The target hangs.

When using the same media to install the target manually, the installation proceeds smoothly.

Causes

This seems to happen only on old firmware.

Resolving the problem

Update the firmware of the target and try deploying again.

Linux deployment of unattended setup image fails with space error

Symptoms

When deploying a Linux profile, the deployment fails with the message No space left on device.

Resolving the problem

To work around this, increase the size of the swap partition on the hard disk drive of your target computer.

LVM partitions are not supported with RedHat 4.8 unattended setup

After successful deployment of a RedHat 4.8 unattended setup system profile with an LVM partition in which Kickstart needs to write data, starting on the operating system fails.

Symptoms

You have created a RedHat 4.8 unattended setup system profile with LVM partitions. You have successfully deployed the profile, but you are not able to start the operating system.

Causes

Kickstart of RedHat 4.8 is not able to write on LVM partitions.

Resolving the problem

Do not use LVM partitions in your RedHat 4.8 unattended setup system profile.

SLES setup from USB cannot access installation media

Symptoms

A problem occurs when deploying the SLES setup profile from USB. If you create a SLES11 x64 setup profile with additional package selection (during the setup creation wizard) and deploy the profile from USB drive (offline), the deployment stops twice and requires user interaction to proceed. Then, if you skip the missing packages, the deployment completes successfully.

Resolving the problem

As a workaround create and deploy a SLES11 x64 setup profile without any additional package and add any required software with Tivoli Provisioning Manager for OS Deployment software modules.

User interaction required when deploying SLES 10 on PowerPC

Depending on your network environment, deploying SuSE Linux Enterprise Server 10 on PowerPC might require user interaction.

Symptoms

When deploying a SuSE Linux Enterprise Server 10 system profile on a PowerPC target, the deployment might stop with a blue screen and the following message:

```
Cannot access installation media
SuSE Linux Enterprise Server SP2.
Check that the server is accessible
```

Environment

This is an environmental issue. In some cases, the network switch does not reestablish the link with the network card in a timely fashion and a timeout occurs.

Resolving the problem

For an immediate solution, select **OK** or **Retry**.

For a long term solution, configure your switches or modify your network infrastructure to avoid long delays when switches must reestablish links to network cards.

PowerPC does not reboot on hard disk at the end of a deployment

Symptoms

At the final reboot of a PowerPC deployment, the target sometimes reboots either in the SMS menu or in the Open Firmware prompt instead of on the hard disk.

Causes

The origin of the problem seems to reside in the version of the firmware and in the operating system which was previously deployed on the target.

Resolving the problem

- If the target boots into the SMS menu at the end of the deployment:
 1. Select **Boot options**
 2. Select **Boot device**
 3. Select **Hard drive**and your target will boot on the hard disk drive.
- If the target boots into the Open Firmware prompt at the end of the deployment, run boot disk and your target will boot on the hard disk drive.

Incomplete logs for Solaris unattended setup system profiles

On Solaris, some of the logging information appears on the target rather than in the log files of the OS deployment server.

Symptoms

When creating a Solaris unattended setup system profile or deploying it, the activity logs seem incomplete.

Causes

This logging information cannot be both captured for inclusion in the log files of the OS deployment server and displayed on the target.

A design choice was made to display the information about the target.

Missing RPM files

There is no verification of RPM dependencies when RPM packages are added through `autoyast.xml`.

Symptoms

When deploying a Linux system profile in which you have added your own RPM packages through `autoyast.xml`, you get an error indicating that some RPM files are missing.

Causes

The product does not resolve RPM dependencies and you have not included all the necessary RPM packages in your system profile.

Resolving the problem

There are two options to solve this problem. You need to perform only one.

- Edit the content of Linux installation media to build a custom Linux distribution.
- After you have created your system profile, edit it and add the missing RPM packages.
 1. Go to **Server > OS deployment > System Profiles**.
 2. Double-click on a system profile to open the **Profile details** page.
 3. Double-click on an OS configuration to open the **OS configuration details** page.
 4. Select **Disks** to view the partition layout.
 5. Click **Browse image of primary partition** to obtain access to the files.
 6. In the contextual menu, select **Upload file** and indicate the file location.

Note: File upload is limited to 16 MB.

7. Repeat step 4 for each RPM package you want to add.

Linux deployment does not end on a Blade

While deploying a Linux system profile on a Blade, Linux deployment engine loading goes on indefinitely when the target is not connected to a remote console.

Symptoms

- You are deploying a Linux system profile on a Blade.
- The target is not connected to a remote console.
- The deployment does not progress anymore when Linux deployment engine is being loaded.

The deployment is successful when a remote console is connected to the target.

Resolving the problem

1. Go to **Server > OS deployment > Target Monitor**. Double-click on a target to view its details.
2. Scroll to the **Boot Settings** section and click **Edit**.
3. Select **Disable USB** and click **Save**.
4. Start your deployment again.

The keyboard layout cannot be changed during offline deployment

You are deploying a Linux operating system from a deployment media and you cannot change the keyboard layout to enter target values.

Symptoms

You are deploying a Linux operating system from a deployment media. Your target has booted into Linux deployment engine and you are required to enter some target information. The keys you press on the target keyboard do not correspond to what is shown on the display, because of a different keyboard layout. Your attempt to change the keyboard layout fails.

Resolving the problem

Due to Linux deployment engine limitations, you cannot change the keyboard layout on the target while you are deploying a Linux operating system from a deployment media. Use the US English keyboard layout to enter the required values.

Deployment of a SLES system profile with more than two logical partitions fails

Your deployment of a SLES 10 SP2 or a SLES 11 system profile with more than two logical partitions fails.

Symptoms

Your SLES 10 SP2 or SLES 11 system profile contains more than two logical partitions.

While you are deploying your system profile, you encounter one of these issues:

- The deployment stalls during partition creation
- The YaST installer stops with the message
Error during configuration of the partition. Pls try again

Causes

The cause is located in the YaST installer which cannot handle more than two logical partitions from the partition layout information generated by the OS deployment server.

Resolving the problem

For SLES 11, a fix provided by NOVELL was validated. You should request from your SLES vendor a driver update. The reference is Bugzilla # 620212 (SUSE 11) .

For SLES 10 SP2, there is no workaround because NOVELL does not provide fixes for this version anymore.

Accessing target devices during deployment

Symptoms

During deployment on an IBM system x3350 the following message is displayed: Please insert disk into drive.

Resolving the problem

Verify that all your devices can be accessed.

SLES deployment fails on HP workstations

On HP workstations, SLES deployment hangs when starting the Linux installer.

Symptoms

You are trying to deploy a SLES system profile on an HP Blade HPBL460c G6 or HP DL580 G7. The deployment fails when starting the Linux installer.

Resolving the problem

To work around this issue, you can enable the kernel-free mode on the target. For more details about the kernel-free mode, see Types of PXE network boot.

Note: If you work with a multiple server architecture, you must be aware that the \TPMf0S Files\tftp folder is not replicated from parent to child servers. Therefore, if a task is triggered on the parent server but the target boots on the child, you must use an external replication mechanism to ensure that the parent and child versions of this folder are synchronized.

Linux unattended setup deployment fails

While trying to deploy a Linux unattended system profile, when starting the Linux kernel, the target hangs. A black screen is sometimes displayed.

Symptoms

The target hangs in the middle of the Linux deployment engine-based unattended setup deployment, when trying to start the Linux kernel. It may show a black screen.

Note: HTTP deployment is disabled.

Causes

Some low level hardware incompatibilities on recent BIOSes may cause the pre-boot engine to hang when loading the Linux kernel during an unattended setup Linux deployments.

Resolving the problem

To work around this issue, you can rely on the Syslinux native Linux loader:

1. Download the Syslinux open source tool available on the web.
2. On the OS deployment server used to create the network boot media, and on the OS deployment server used for deployment, you need to copy some Syslinux files into the <DATADIR>/global/updates folder, for example C:\TPMf0S Files\global\updates on a Windows operating system.
 - a. From Syslinux 3.63, copy the files extlinux from the extlinux folder, and mbr.bin from the mbr folder, into <DATADIR>\global\updates\extlinux on the server.
 - b. From Syslinux 4.02 or above, copy the files isolinux.bin from the core folder, and lua.c32 from the com32/lua/src folder, into <DATADIR>\global\updates\isolinux on the server.
3. Restart your OS deployment servers.
4. Create a new deployment scheme, or use a deployment scheme, where **Download files with a network share or Linux HTTP when applicable** is set to **No**.
5. Generate a network boot CD.

- a. Go to **Server > OS deployment > Deployment schemes**.
- b. Click **Generate Media**.
- c. Select **Generate a network boot CD/DVD Isolinux based** and click **Next**.

Note: This option is available only when the appropriate content has been copied into <DATADIR>/global/updates.

- d. Follow the remaining instructions of the wizard to generate your ISO network boot CD/DVD.
6. On the target, boot from the CD/DVD.
7. At the isolinux prompt, type **LinuxDE** to start Linux deployment engine from the network boot CD/DVD.
8. When the green banner is displayed, to boot from the hard drive on which Linux has been installed, click **Reboot** on the target.

Error during Linux 64-bit deployment or migration when running mkinitrd

When deploying or migrating a Linux 64-bit operating system, you may encounter an error when running the **mkinitrd** script. This error sometimes results in a failed deployment or migration.

Symptoms

You are deploying a Linux 64-bit operating system.

You see the following error message in the logs, although the deployment continues:

```
<ERR> mkinitrd failed, you may experience problem during deployment!
```

Causes

Linux deployment engine had generated a 32-bit **mkinitrd**, that sometimes causes problems for 64-bit deployment and migration.

UEFI targets fail to PXE boot on VMWare virtual machines with Linux operating system

VMWare virtual machines with Linux guest operating system need e1000 Network Card Interface (NIC) for UEFI to PXE boot.

Symptoms

When the product server is running on a VMWare virtual machine with Linux guest operating system, the UEFI targets fails to PXE boot.

Resolving the problem

Make sure the Network Card Interface (NIC) type on the product server is e1000.

Problems provisioning VMWare ESX operating systems

Deploying VMWare ESX/Linux fails

Symptoms

When you deploy Linux or VMWare ESX operating systems, the deployment might fail and the following error occurs:

```
Unable to boot Linux with the requested parameters
```

because the Linux kernel or initial ramdisk cannot be found.

Causes

The default device mapping is not working.

Resolving the problem

To solve this issue, manually install Linux or VMWare ESX using the CD/DVD on the workstation to identify the devices used and to use the correct device name in the profile.

HS20 hangs when VMWare ESX 3.5 is being deployed

You cannot deploy a VMWare ESX 3.5 on an HS20 target without interaction.

Symptoms

You are deploying a VMWare ESX 3.5 system profile on an HS20 target. The deployment hangs when a reboot is needed.

Causes

This issue is specific to the HS20 which cannot be rebooted by the ESX installer.

Resolving the problem

If you reboot the HS20 target manually, the deployment resumes and ends successfully.

Problems after a deployment

Booting a target with two disks after a successful deployment fails

After a successful deployment, you cannot boot your target which contains two hard disks.

Symptoms

Your system profile contains only one disk, but the target on which you are deploying it contains two disks. The deployment is successful. However, you are not able to start the operating system on your target.

Causes

You are trying to boot on one disk while the operating system was installed on the other disk.

Resolving the problem

You can try to change the boot order of the target to put the second disk first.

Boot on OS fails after a successful deployment

After a successful deployment of a system profile, clicking **Boot on OS** does not start the operating system on the target.

Symptoms

You have created a system profile. You have performed a successful deployment of this system profile. When the green banner is displayed, you click on **Boot on OS** to start the operating system. There is an error and the operating system does not start.

Resolving the problem

After the successful deployment, do not click **Boot on OS**. Click **Reboot** instead.

RAID deployment task fails

Symptoms

When you perform a RAID deployment task using the RAID command, the console displays that the RAID configuration command exits with an error.

Resolving the problem

If the RAID hardware configuration fails, check the activity .log and .trc files for detailed error messages and advanced troubleshooting. For example, if you receive the following error message:

```
ERROR: (2829) Cannot create Array, Disk Selection Error:  
Drive "1I:1:3" is not valid
```

it means that you selected a wrong disk in the hardware configuration.

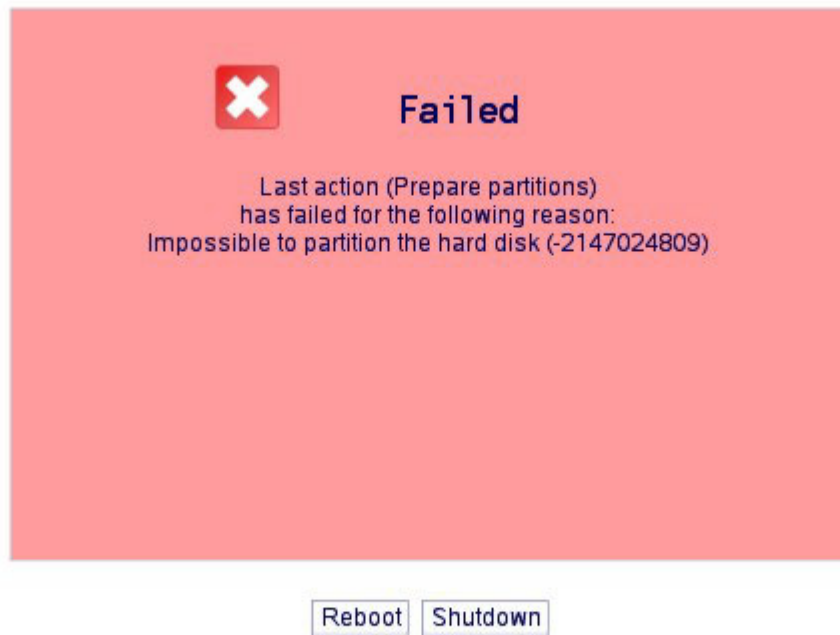
Problems while redeploying a profile on the same target

Deploying on a Windows target while keeping the existing partition

Symptoms

When you try to deploy again just one partition of a system profile that has a partition layout with two partitions, for example, C and D, using simple deployment by setting only the option *Must be deployed* to **no** for the partition that you do not want to deploy, for example, deploy again C while leaving D untouched, the deployment fails with the following error:

Impossible to partition the hard disk



Waiting for the OS deployment server to unblock this target

Resolving the problem

To solve this issue, follow these steps:

1. Make sure the following prerequisites are met:

- In the partition layout, 100% of the disk must be assigned. In this example, there are two partitions, one occupying 85% of the disk space and the other occupying 15%.
 - The second partition must not be completely full to allow the shrinking operation.
2. Create a software module pre-OS (with a stage equal to 1, before using the hard disk) to shrink the partition that you do not want to deploy again. Here is a sample of a script for this type of software module:


```
REM processing disk 0 partition 2 shrinking
select disk=0
select partition=2
shrink DESIRED=2048 MINIMUM=1024
```

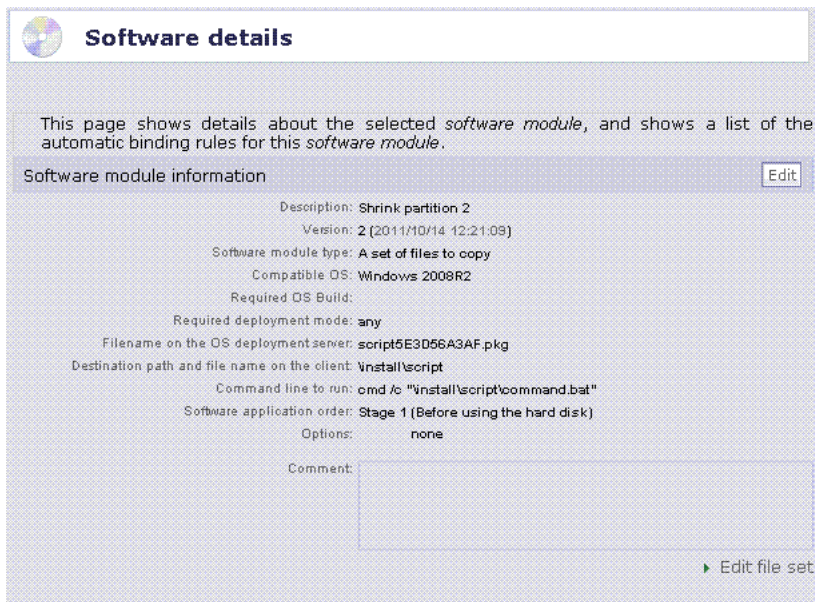


Figure 2. Software Module Details

3. Before performing the second deployment, make sure that you have changed the must be deployed option of D: partition to no and then select the software module defined in the previous step.

When partitions C and D are primary, you have an alternative solution: from the first deployment, the scenario works only if in the disk at least 1 or 2 GB of unpartitioned space is left on the disk.

Redeploying partition D with the BitLocker feature enabled

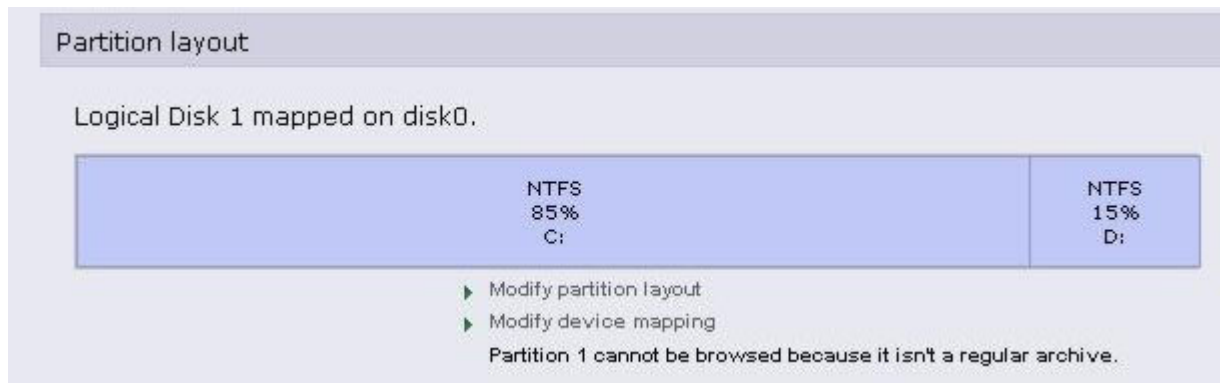
This topic describes how redeploy partition D when the BitLocker feature is enabled.

In the first part of this procedure (steps 1 to 3) you create a system profile to be used in deployments and redeployments. In the second part of the procedure (steps 4 to 5) you perform the redeployment. The prerequisite for this procedure is to decrypt partitions C and D.

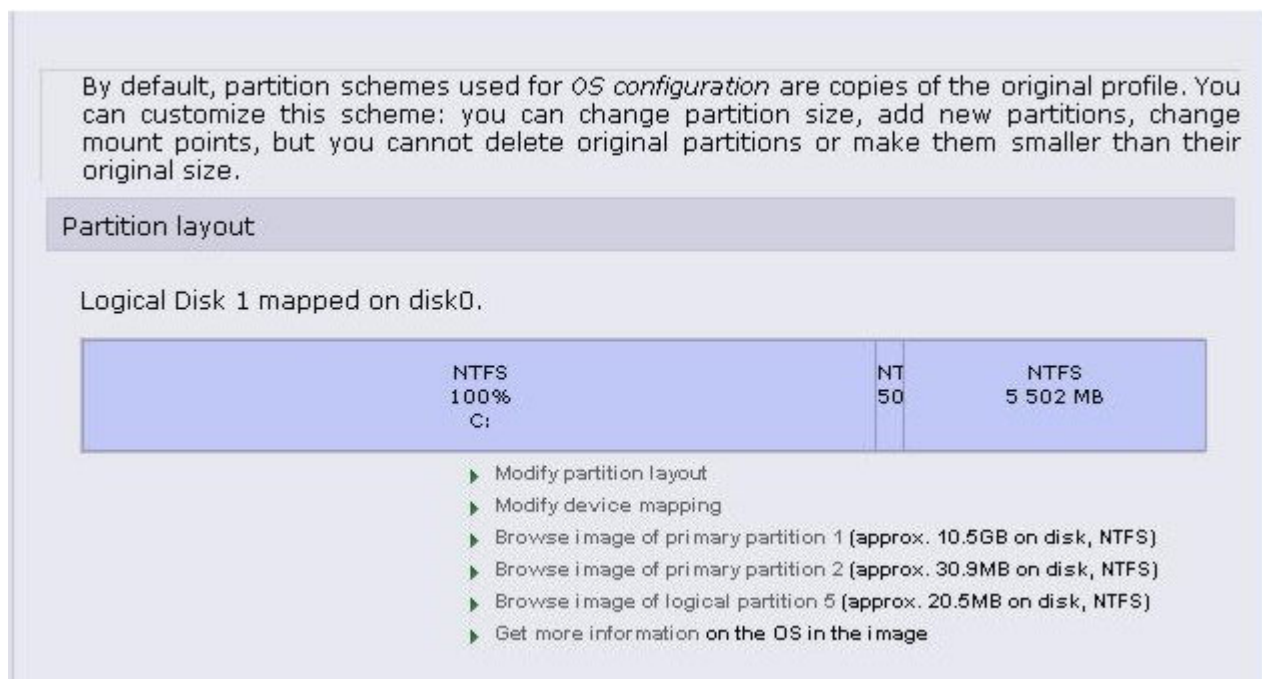
To create a system profile to be used for deployment and redeployment with the BitLocker feature enabled, perform the following steps:

1. Deploy a Windows 7 Enterprise with a partition layout that has two partitions, for example C and D, where D is not primary, and the option Must be deployed

is set to yes for both C and D, as shown in the following figure:



2. After successful deployment, log in into Windows and follow the procedure to enable the BitLocker feature. For more information, see <http://www.howtogeek.com/howto/6229/how-to-use-bitlocker-on-drives-without-tpm/>. The BitLocker setup is divided into two parts:
 - a. Prepare your drive for BitLocker
 - b. Encrypt the driveTo prepare your drive, a system and hidden partition of 300 MB is created between C and D. This partition is left decrypted and is used to boot the target. After the reboot, you can encrypt both partitions C and D. At this stage, skip the encryption part. To proceed with next steps, both C and D partitions must be decrypted.
3. Create a new system profile by cloning the target. Launch the command sysprep before creating the profile. After cloning, you see the partition layout of the cloned profile as shown in the following figure:



The hidden partition is now present between C and D. Use this clone for a deployment of a new target, as the base image, or for a redeployment on a target previously installed with the same image, saving partition D.

4. Create some files on partition D, to verify that after the redeployment they still exist on the partition.
5. Decrypt both partitions C and D (if they were encrypted).
6. Modify the cloned profile by changing the setting of the option Must be deployed of the partition D to no.
7. Create a pre-OS software module (with a stage equal to 1, before using the partition) to shrink the partition that you do not want to deploy again. This is a script sample for this type of software module:

```
REM processing disk 0 partition 3 shrinking
select disk=0
select partition=3
shrink DESIRED=2048 MINIMUM=1024
```

This page shows details about the selected *software module*, and shows a list of the automatic binding rules for this *software module*.

Software module information

Edit

Description: Shrink partition 3

Version: 3 (2011/11/23 12:32:19)

Software module type: A set of files to copy

Compatible OS: Windows 7

Required OS Build:

Required deployment mode: any

Filename on the OS deployment server: scriptD65256A3AF.pkg

Destination path on the client: \install\script

Command line to run: cmd /c "\install\script\command.bat"

Software application order: Stage 1 (Before using the hard disk)

Options: none

Comment:

Edit file set

Software binding rules

8. Make sure that you select the software module specified in step 7, before performing the second deployment.

This second deployment reinstalls partition C in the same state when it was cloned, and leaves untouched partition D.

Problems with a multiserver architecture

Errors in the vm.trc log when replicating with flag h in AutoSync

There can be error messages when replicating for the first time a server with the flag h in AutoSync.

Symptoms

- You are working in a multiserver multiple database environment.
- You use a config.csv file to replicate your servers.
- You have set flag *h* in **AutoSync**.
- Error messages appear in vm.trc log.
- The error messages all occur inside **RADRefreshHosts**.

Example:

```
[2009/11/18 20:12:08] <ERR> Error raised by RADGetMasterBomIds in console.rbc,
line 2047 [VM:3250]
[2009/11/18 20:12:08] <ERR> Not an array
[2009/11/18 20:12:08] <WRN> (called from RADRefreshHosts (console.rbc:2837))
[2009/11/18 20:12:08] <WRN> (called from RADKeepAlive (console.rbc:3797))
[2009/11/18 20:12:08] <WRN> (called from RADInstaller (console.rbc:8188))
[2009/11/18 20:12:08] <WRN> (called from RunPackageInstallers (server.rbc:19980))
[2009/11/18 20:12:08] <WRN> (called from --toplevel-- (server.rbc:21505))
[2009/11/18 20:12:09] <ERR> Error raised by RADRefreshHosts in console.rbc,
line 2863 [VM:3250]
[2009/11/18 20:12:09] <ERR> Not an array
[2009/11/18 20:12:09] <WRN> (called from RADKeepAlive (console.rbc:3797))
[2009/11/18 20:12:09] <WRN> (called from RADInstaller (console.rbc:8188))
[2009/11/18 20:12:09] <WRN> (called from RunPackageInstallers (server.rbc:19980))
[2009/11/18 20:12:09] <WRN> (called from --toplevel-- (server.rbc:21505))
[2009/11/18 20:12:14] <INF> Faulty query (dbdb.rbc:2049): INSERT INTO "BOMUpd"
("Dep1Count","Description","SN","UUID","MAC","Model","Platform","LogicalName",
"HostName","EnableIPv6","IPSettings","IP","MgmtIF","NetMask","Gateway","DNSServer1",
"DNSServer2","DNSServer3","DNSDomain","DNSOrder","WINSServer1","WINSServer2",
"BitsPerPel","Xresolution","Yresolution","Vrefresh","IdentScope","ScopeName","OrgUnit",
"JoinDomUser","JoinDomPass","ProductKey","AdministratorName","NameService",
"NameServer","LDAPProfile","KerberosRealm","KAdminServer","KDCServer1","KDCServer2",
"KDCServer3","TimeServer","TermInfo","SystemID","Dep1Set","GroupID",
"DeployGroupID","UserID","EnsembleID","Status","StatusDate","IdentDate",
"MonitorLocation","MonitorLayout","RemboServer","RemboOptions","RemboLockout",
"PXEBootMode","PXEBootSrv","RedirectSrvAlt","MgnParam","SrvHostID","Hypervisor",
"InstanceID") VALUES (0,NULL,' 40160191D8B460',NULL,'00:14:5E:54:9E:F2',NULL,
NULL,NULL,NULL,'0',NULL,'9.167.62.167',NULL,NULL,NULL,NULL,NULL,NULL,NULL,
NULL,NULL,NULL,NULL,NULL,NULL,NULL,NULL,NULL,NULL,NULL,NULL,NULL,NULL,
NULL,NULL,NULL,NULL,NULL,NULL,NULL,NULL,NULL,NULL,'Srv5E549EF2-0',0,NULL,
'3174709/12 06:38:34','3174709/12 06:38:34','Default',NULL,'9.167.62.167',0,0,0,
NULL,NULL,NULL,0,0,NULL)
[2009/11/18 20:12:15] <ERR> Error raised by DbRbcExec in dbdb.rbc, line 4174 [:0]
[2009/11/18 20:12:15] <ERR> Unknown error (Problem in a SQL query, check logs for
details[USRMSG]ERROR 1 (22007): DB2 SQL Error: SQLCODE=-180, SQLSTATE=22007,
SQLERRMC=null, DRIVER=3.53.70[dbdb.rbc:2049][I18NMSG]Problem in a SQL query,
check logs for details)
[2009/11/18 20:12:15] <WRN> (called from DbCreateRecord (dbdb.rbc:2090))
[2009/11/18 20:12:15] <WRN> (called from RADForwardEvents (console.rbc:2361))
[2009/11/18 20:12:15] <WRN> (called from RADRefreshHosts (console.rbc:3136))
[2009/11/18 20:12:15] <WRN> (called from RADKeepAlive (console.rbc:3797))
[2009/11/18 20:12:15] <WRN> (called from RADInstaller (console.rbc:8188))
[2009/11/18 20:12:15] <WRN> (called from RunPackageInstallers (server.rbc:19980))
[2009/11/18 20:12:15] <WRN> (called from --toplevel-- (server.rbc:21505))
```

Resolving the problem

These error messages do not compromise server functionality and can be ignored.

Errors in the file.trc log when replicating Symptoms

When performing a replication through sync.pak, a failure occurs and the following error message is written in file.trc:

Insufficient system resources exist

Resolving the problem

To solve this issue, apply to the child server the procedures described in the following Microsoft technote:

<http://support.microsoft.com/kb/312362>

Errors in the NBP log after switching parent and child server roles

Reversing the roles of a parent and a child server in a single database architecture creates errors in the NBP log.

Symptoms

You have reversed the roles of a parent server and a child server in a multiserver single database architecture, using the links provided in the web interface. There are errors in the NBP logs of both the new parent server and the new child server after this process.

The error on the new parent server is similar to

```
<ERR> [9.167.62.244] No session found with identifier 1. Resetting connection.
```

The errors in the new child server are similar to

```
<ERR> [TCPWRITE.9.167.62.246.69bc] No matching connection
<ERR> [9.167.62.246] Error during TCPWRITE (69bc)
<ERR> [TCPCLOSE.9.167.62.246.69bc] No matching connection
<ERR> [9.167.62.246] Error during TCPCLOSE (69bc)
```

Resolving the problem

You can safely ignore these errors in the NBP logs.

Errors in the vm.trc log when changing server role from standalone to child

Error messages can be logged to the trace file when restarting the server after a server has changed role from standalone to child.

Symptoms

- You are working in a multiserver multiple database environment.
- You use a config.csv file to change the role of a server from standalone to child.
- You copy the config.csv file to the RAD directory.
- You restart the server.
- Unexpected error messages appear in vm.trc log.

Sample error message:

```
[2010/02/19 15:30:20] <ERR> Error raised by DbRbcExec in dbdb.rbc, line 4178 [:0]
[2010/02/19 15:30:20] <ERR> Unknown error (Problem in a SQL query, check logs for
details[USRMSG]ERROR 1 (22007): DB2 SQL Error: SQLCODE=-180, SQLSTATE=22007,
SQLERRMC=null, DRIVER=3.53.70[dbdb.rbc:2051][I18NMSG]Problem in a SQL query,
check logs for details)
[2010/02/19 15:30:20] <WRN> (called from DbCreateRecord (dbdb.rbc:2092))
[2010/02/19 15:30:20] <WRN> (called from RADForwardEvents (console.rbc:2426))
[2010/02/19 15:30:20] <WRN> (called from RADRefreshHosts (console.rbc:3201))
[2010/02/19 15:30:20] <WRN> (called from RADKeepAlive (console.rbc:3862))
[2010/02/19 15:30:20] <WRN> (called from RADInstaller (console.rbc:8542))
[2010/02/19 15:30:20] <WRN> (called from RunPackageInstallers (server.rbc:19999))
[2010/02/19 15:30:20] <WRN> (called from --toplevel-- (server.rbc:21524))
```

Resolving the problem

These error messages do not compromise server functionality and can be ignored.

Problems with virtual images and migration

Physical to physical operating system migration of Linux might stop

Symptoms

When performing a physical to physical operating system migration of Linux, the migration might stop due to ACPI issues.

Resolving the problem

When performing a physical to physical operating system migration for Linux platforms, ensure that the base architectures of the source computer and target computer are similar as much as possible. In particular the ACPI / APIC architecture of both computers must be the same to ensure that the migration ends successfully.

The ProgramData folder is not hidden

After direct migration or virtual image deployment, the ProgramData folder is not hidden anymore.

Symptoms

You are deploying a virtual image or performing direct migration of a Windows operating system. On the source, the ProgramData folder is hidden. After task completion, this folder is not hidden on the receiving target.

Resolving the problem

There is currently no workaround.

Rerunning a virtual image export or replication shortly after a first failed attempt might also fail

Symptoms

When you attempt to perform a virtual image export or replication for the second time, shortly after the first attempt was unsuccessful, the second attempt might also fail, with the following error:

Invalid Runtime Environment, current is AGENT but SERVER was expected.

Causes

This is a known problem, and it only occurs when the second task is queued for the OS deployment server's sequencer while the server is trying to save the virtual image to another location.

Resolving the problem

Allow the server some time to clean up the activity, and resubmit the task when the server is less busy.

Rebooted target computer fails to load a new task

Symptoms

An image deployment task failed on the target computer due to a missing device driver. After fixing the problem, a new deployment task was submitted and the target was rebooted. After the reboot, instead of picking up the new task, the target loads the previous failed task.

Causes

Tivoli Provisioning Manager for Images is currently unable to detect CMOS in the target memory, therefore the failed task cannot be cleaned up.

Resolving the problem

The suggested workaround is:

1. From the console window, reboot the target system.
2. When the target loads WinPE and displays the following message, Cache Selected according to CMOS, press CTRL+ALT+DEL to cancel, press **Ctrl+Alt+Del** to cancel the previous CMOS before the next Rembo window is loaded.

Deployment failure caused by different disk partition sizes

Symptoms

You created a virtual image from a machine with multiple partitions and deployed to a hypervisor. The deployment failed.

Causes

You created an image on a large disk and then tried to deploy it on a smaller disk. During the deployment, Tivoli Provisioning Manager for Images tries to resize the target disk by having all the partitions keep the same size except for the largest one that occupies the rest of the disk.

Resolving the problem

Before deploying a virtual image created on a machine with multiple partitions, ensure that the target disk is larger than the source machine disk.

Problems with booting, network, and file transfers

Unknown MAC address during boot on PXE

Symptoms

When you boot an Hyper-V target with two NIC configured, for the secondary NIC the MAC address is set to "Unknown NIC".

Causes

When booting on PXE a target and checking its network card, only the network card used to boot on PXE will display the MAC address.

Resolving the problem

There is no current workaround to solve this issue.

Hang on splash screen

Symptoms

The OS deployment server might be hanging on splash screen when booting on a target.

Resolving the problem

This issue can be due to a missing default gateway information among the options to be set for the DHCP server (option 3). Also verify that any firewall does not block the traffic to the OS deployment server, UDP port 4012.

PXE bootrom not detected

Symptoms

During the boot process, there is no message about the PXE bootrom, and the computer boots normally (on the floppy, hard disk or CD).

Resolving the problem

Check that your network card is correctly installed, and that a PXE bootrom is installed on the network card. To verify that the network works, run Windows or

Linux, and configure the operating system so that you are able to ping other computers (or you are able to see other computers in the network neighborhood).

On certain network cards, the PXE bootrom is not be activated by default. Read the product documentation to find the key combination to press to enter the PXE setup menu at boot time. On Intel EPRO100, the key combination is Ctrl-S, or Shift-Shift (press both Shift keys). These keys must be pressed during the boot process, when the computer is powered on. Some cards do not have a configuration menu.

Enter your BIOS setup during boot time (DEL, or F2 key on most systems), and configure the BIOS boot process so that the network card is the first entry in the boot list. In some BIOS, there is an option to enable boot on network. On other BIOS, you must manually set the LAN (also called NET, or Other) as the first device of the boot order.

If all of these steps fail, try to obtain a flash memory upgrade from your network card vendor, and flash the network card rom with the newest upgrade. If the flash process fails, there is a chance that no bootrom is installed on your network card.

If you are still not seeing the PXE messages, ask for support from your network card manufacturer.

Alternatively, you can create a network boot media to enable your target to connect to your OS deployment server.

The bootrom displays DHCP... and times out

Symptoms

The following message is displayed:

The bootrom does not receive enough information to proceed further.
Either the DHCP server or the
OS deployment server is not correctly configured.

Resolving the problem

Check that your DHCP server is correctly configured as explained in DHCP server OS configuration. In particular, check that option 60 is set to PXEClient if you are running the DHCP server and the PXE server on the same target only.

If the DHCP server and the OS deployment server are on the same target, try to stop both servers, and restart the two servers in the following order: DHCP server first, then the OS deployment server. If the OS deployment server is started first, it might reserve the DHCP port, thus preventing the DHCP server to start.

Check your DHCP configuration: run Windows or Linux on your remote-boot target, and configure the network to use dynamic configuration instead of fixed IP address. If this works (run winipcfg or ipconfig on a Windows computer, ifconfig on a Linux computer), then the DHCP server is correctly configured for this target. Otherwise, check your DHCP server OS configuration, so that the remote-boot target is assigned an IP address, a netmask and a default gateway.

If your server is correctly configured (including option 60), and the target still displays DHCP... followed by an error, check your OS deployment server OS configuration. Stop the OS deployment server, then run rembo.exe -d -v 6, and start the remote-boot target. When starting, the server displays a line saying whether it is acting as a DHCP Proxy or a BINL Proxy. If the DHCP server and the OS deployment server are on the same target, the OS deployment server acts as a BINL proxy. If they are on different hosts, the server acts as a DHCP proxy. If the

server displays a message saying it acts as a BINL proxy, but the two servers are not on the same target, it means that there is a DHCP server installed on the computer where you have installed Tivoli Provisioning Manager for Images.

When a target starts, and DHCP is correctly configured, the OS deployment server (in debug mode) displays Valid discovery from... followed by target... found in group.... If the server displays the first line, but displays target... not found in any group instead of the second line, it means that your configuration file does not contain a default group, and that the remote-boot target is not declared in any group (the target must be declared with its hardware address, not its IP address).

If the server does not display the message Valid discovery request from..., then option 60 on the DHCP server is not correctly set, or the OS deployment server and the DHCP server are not on the same subnet. If you have installed Tivoli Provisioning Manager for Images on a multi-homed target (a computer with more than one network card, or with a dialup adapter), use the Interfaces option to specify which network interface to use.

If it still does not work, send a report to your IBM Software Support representative with the following information:

- All the files from the logs directory of the OS deployment server
- The OS configuration information for your DHCP server
- The OS configuration information for your OS deployment server
- A memory dump of the network traffic between the servers and the remote-boot target (use the MS Network Monitor on Windows NT/2000)

The bootrom displays MTFTP..., and an error message

Symptoms

The following error message is displayed:

The bootrom was unable to receive the Tivoli Provisioning Manager for OS Deployment bootstrap from the server.

Resolving the problem

If the delay between the MTFTP.. message and the error message is short, and the message explains that a file was not found, then the target you have installed Tivoli Provisioning Manager for Images on already runs a TFTP server (and this TFTP server answers request for the OS deployment server). If you are using Windows 2000/2008/XP/Vista on the server, check the list of services, and disable services related to TFTP or Boot protocols (including Intel LCM and Microsoft PXE,).

If the delay between the MTFTP.. message and the error message is long, the multicast TFTP datagrams sent by the provisioning server are not being received by the remote-boot target. If you have installed Tivoli Provisioning Manager for Images on a multi-homed computer, use the Interfaces parameter to specify which network interface to use for multicast packets.

If it still does not work, send a report to your IBM Software Support representative with the following information:

- All the files from the logs directory of the provisioning server
- The OS configuration information for your DHCP server
- The OS configuration information for your OS deployment server
- A memory dump of the network traffic between the servers and the remote-boot target (use the MS Network Monitor on Windows 2000/2008/XP/Vista)

Occasional MTFTP timeout (on multihomed server)

Symptoms

When a server network connection is lost and then recovered, the targets report an MTFTP timeout after receiving their DHCP lease.

Resolving the problem

This is because Windows 2000 automatically closes all sockets when a network connection is lost. A workaround is to restart the provisioning server after the network is up again. A long-term fix is to disable the Windows 2000 media sensing on the network card, on the server.

More information about this topic can be found in Microsoft knowledge base, under the title *How to Disable Media Sense for TCP/IP in Windows 2000*.

Subnet not allowing multicast

In some cases, a subnet incorrectly refuses multicast connections.

Symptoms

While you have selected multicast, packets are sent in unicast in a given subnet.

Causes

It is possible that the first multicast probe sent by the OS deployment server failed and that a redetection is needed for the OS deployment server to ascertain that multicast is possible on this subnet.

Resolving the problem

If you believe that your subnet should be able to use multicast, you can try to redetect the subnet settings.

1. Go to **Server > OS deployment > Target Monitor**.
2. Select the problematic subnet.
3. Click **Force redetection of subnet settings**

The multicast option is not respected during deployment

While deploying a system profile with a multicast option in the deployment scheme, the deployment is performed using unicast.

Symptoms

While trying to deploy a system profile with a deployment scheme which request the use of multicast file transfer, the deployment succeeds but unicast was used instead of multicast.

Resolving the problem

Multicast is available only if all the criteria which follow are respected.

- the targets have an Intel x86 or x86-64 architecture
- multicast is selected in the deployment scheme
- the subnet supports multicast traffic
- multicast is not disabled in the boot options of the target
- the target is not a VMWare 3.x guest.

Make sure all the criteria are met before trying to use multicast.

Deployment fails on some Broadcom network adapters

On some Broadcom network adapters, a firmware defect prevents successful deployment or redeployment.

Symptoms

Deployment or redeployment fails during network transfers on some targets with either one of the following Broadcom chips:

- BCM5700
- BCM5701
- BCM5702
- BCM5703
- BCM5704

These Broadcom chips can be potentially included in the following IBM servers, among other IBM and non-IBM servers:

- HS20
- LS20
- x335
- x355
- x3655

Causes

This is a firmware defect which cannot be fixed in Tivoli Provisioning Manager for Images.

Diagnosing the problem

On the target, you see a message starting with **BROKEN FIRMWARE DETECTED FOR YOUR NETWORK ADAPTER**.

Resolving the problem

This issue can only be fixed by an update of your firmware with the correct PXE level. Table 1 indicates which PXE level must be updated and the minimal PXE level to be reached.

Table 1. PXE levels

Problematic PXE level	Correct PXE level
9.0.x	9.0.13 or above
10.0.x	10.4.10 or above
10.4.x	10.4.10 or above
11.0.x	11.4.0 or above

Contact your hardware support to obtain a new firmware version with the required PXE level.

The deployment does not start on a target with multiple network interface cards

You have manually registered a target with multiple network cards, providing the MAC address to register. You now want to deploy the target, but the deployment never starts.

Symptoms

- You have a target with more than one network interface card. Each card has its own MAC address, for example, AA:AA:AA:AA:AA:AA and BB:BB:BB:BB:BB:BB.
- Your OS deployment server works in closed mode.
- You have registered your target using the first MAC address, for example AA:AA:AA:AA:AA:AA.

- On the web interface, you select the target and select **Deploy now** to start a deployment.
- The deployment does not start on the target with MAC address AA:AA:AA:AA:AA:AA.
- However, a separate entry is added in the Target Monitor, with the IP address corresponding to the second MAC address (BB:BB:BB:BB:BB:BB). You now have two entries for the same target in the Target Monitor.

Causes

The MAC address you used to register the target (AA:AA:AA:AA:AA:AA) is not the one from which the target boots by default.

Resolving the problem

There are two options to solve this issue:

- Use the UUID, serial number, or IP address to manually register your target.
- If you must use the MAC address to register your target, make sure that the MAC address that you provide is the one the target boots from by default.

Problems with the database and the database gateway

ODBC Data Source Administrator reports an error when manually removing or configuring a system DSN on Windows 64-bit operating systems

Symptoms

When you install Tivoli Provisioning Manager for Images with the installer on a Windows 64-bit operating system, the installer creates a system DSN called **AutoDeploy** with a Microsoft Access Driver. This system DSN cannot be removed or configured using the ODBC Data Source Administrator. When trying to do so, you get the following error message:

The setup routines for the Microsoft Access Driver (*.mdb) ODBC driver could not be found. Reinstall the driver.

Resolving the problem

The system DSN is removed automatically when you uninstall Tivoli Provisioning Manager for Images. To remove or configure the system DSN manually, you must do so directly in the registry with the registry editor. To open the registry editor:

1. Open a command prompt
2. Type regedit and press **Enter**.

The key and value to remove or configure are

- The name and driver of the source, located at My Computer/HKEY_LOCAL_MACHINE/SOFTWARE/ODBC/ODBC.INI/ODBC Data Sources/ There is an entry with AutoDeploy in the Name column and Microsoft Access Driver (*.mdb) in the Data column
- The source parameters, located at My Computer/HKEY_LOCAL_MACHINE/SOFTWARE/ODBC/ODBC.INI/AutoDeploy

Database connectivity with Apache Derby on a UNIX server

Problem description

You either want to check your database connectivity before installing your provisioning server or you encounter database connectivity errors in the vm.log file of your provisioning server. You work with Apache Derby and a UNIX computer.

Problem resolution

To check the connectivity of your provisioning server installed on UNIX with an Apache Derby database

1. `dbgw.jar` is a Java implementation of the TCP-to-ODBC gateway, using JDBC instead of ODBC for database access. You have two alternatives to start the database gateway.
 - - a. Start the database gateway by copying the `derbyclient.jar` file into the directory where `dbgw.jar` is located. The `dbgw.jar` file is part of the Tivoli Provisioning Manager for OS deployment installation files, and should have been extracted into `/usr/local/tpmfos`.
 - b. Run the following command:

```
java -cp dbgw.jar:derbyclient.jar \ -Djdbc.drivers=org.apache.derby.jdbc.ClientDriver com.rembo.dbgw.Dbgw
```
 - Alternatively, if you do not want to copy the `derbyclient.jar` file, you can start the Java database gateway with the following command:

```
java -cp dbgw.jar:$DERBY_HOME/lib/derbyclient.jar \ -Djdbc.drivers=org.apache.derby.jdbc.ClientDriver com.rembo.dbgw.Dbgw -d
```
2. In another terminal, verify the database connectivity using:

```
telnet localhost 2020 use derby://127.0.0.1:1527/tpmfosddb,root,<password>
```
3. Stop the `dbgw` process. If you do not stop it, you will encounter errors when running the setup or when restarting your provisioning server.

Verifying database connectivity

Problem description

If you suspect there are connectivity problems between the OS deployment server and the database, perform the following checks to isolate the problem.

Problem resolution

To verify the connectivity of the database gateway, the database, and the OS deployment server, perform the following checks:

1. Verify the database gateway.

The database gateway, used by the server, runs a process named **dbgw** on the computer running the OS deployment server. Stop the OS deployment server and attempt to contact the database gateway using the following telnet command:

```
telnet localhost 2020
```

If you successfully run a session then the database gateway is functioning correctly. The following banner is displayed:

```
--> "HELLO [127.0.0.1] - Rembo TCP-to-ODBC gateway / $Revision: #2 $"
```

Note: The **dbgw** service is provided in each TPM server, but in the integrated environment the child server installed by TPM does not use this service.

2. Verify the connection between the database and the database gateway.

Contact the database using the command "**use <connString>**", where the value of `<connString>` depends on the way the server has been configured.

- On Windows: this string could be the name given in the "ODBC administrator".

- On Linux or Windows: this string could be whatever the DB needs.

For example, submit the following command for a default installation of the OS deployment server and DB2 database software on a Linux machine:

```
use db2://127.0.0.1:50000/tpmfosd,db2inst1,<password>
```

If the connection is successful, then you are in a classical SQL evaluator and the following commands are transmitted to the database, evaluated, and the results are displayed in the same window:

```
info
select * from BOM
select * from Servers where ServerIdent
```

When the connection string is not "AutoDeploy", you can find the connection string in the following file: `net://global/rad/radb.ini`. Use the parameters in the [Settings] section to submit the command as follows:

```
use <ODBC_Source>,<ODBC_Username>,<ODBC_Password>
```

If you can successfully verify both of these steps then the database is working properly.

Java exception thrown by dbgw process

Symptoms

On z/Linux platforms, after you install Tivoli Provisioning Manager for OS Deployment, the following exception is thrown by the dbgw process:

Exception: Communications link failure due to underlying exception:

```
** BEGIN NESTED EXCEPTION **
```

```
java.io.EOFException
```

```
STACKTRACE:
```

```
java.io.EOFException
at com.mysql.jdbc.MysqlIO.readFully(MysqlIO.java:1934)
at com.mysql.jdbc.MysqlIO.reuseAndReadPacket(MysqlIO.java:2380)
at com.mysql.jdbc.MysqlIO.checkErrorPacket(MysqlIO.java:2909)
```

Causes

The problem seems to be caused by the database engine which takes long time to start.

Resolving the problem

The product recovery code works around the issue without any side effect.

Error in Target Monitor page when using User Authentication in ODBC

Symptoms

After you install the Microsoft SQL Server database, if you start ODBC config using User Authentication, you might have an error opening the Tivoli Provisioning Manager for Images interface.

Resolving the problem

Ensure that the credentials you use to log on to the Tivoli Provisioning Manager for Images interface and to access the OS deployment server database are the same as those used for the SQL server connection.

Known limitations applicable to the current release of the Toolkit

The following limitations are applicable to the current release of the Toolkit:

- The rbagent process running offline on operating system exits if no pending task is found and no bindings for a binding menu are available.
- In the offline menu for kernel-free targets, if a task is selected having one action with prepare_nextaction=true, the action is performed but not tracked into the database.
- A target remains in idle state while polling the server for a scheduled task. The target runs the task when the administrator has scheduled it.
- Target UI customization is limited in rebranding. For example, due to the new architecture, when a target PXE boots and there are no pending task, the product skin is shown and cannot be changed for now (custom skin is not available in kernelfree and rembokernel).
- When the base flow customization plug-in implements the target machine redirection, if the target does the PXE boot on the server using the kernel mode, it will be redirected on the secondary server in the same kernel mode and the server must be accordingly configured. If the target does the PXE boot in kernel-free mode, it will boot in either kernel-free or kernel mode on the secondary server depending on the current server configuration.

On Linux FAT32 must be created with a hidden parameter for Toolkit offline flows

Problem description

On Linux, if the cache partition is used to store Rembo Kernel as loader, the partition must be formatted using the -h (hidden-sectors) flag.

Problem resolution

You must retrieve the offset for the cache partition by using the following fdisk -ul command:

```
pc-000C29C47389:~# fdisk -ul
```

Disk /dev/sda: 21.5 GB, 21474836480 bytes

255 heads, 63 sectors/track, 2610 cylinders, total 41943040 sectors

Units = sectors of 1 * 512 = 512 bytes

Disk identifier: 0x00051d14

Device Boot	Start	End	Blocks	ID System
/dev/sda1	2048	2080767	1039360	82 Linux swap / Solaris

Partition 1 does not end on cylinder boundary.

Device Boot	Start	End	Blocks	ID System
/dev/sda2	2080768	3112959	516096	83 Linux

Partition 2 does not end on cylinder boundary.

Device Boot	Start	End	Blocks	ID System
/dev/sda3	3112960	29419519	13153280	83 Linux
/dev/sda4	29419520	41873407	6226944	5 Extended
/dev/sda5	29421568	38174719	4376576	b W95 FAT32

Assuming that /dev/sda5 is the partition created to be the cache, then its Start sector is 29421568. This value must be given as input as follows

```
pc-000C29C47389:~# mkfs.vfat -h 29421568 /dev/sda5
```

Miscellaneous problems

rbagent sync-srvradget command fails due to timeout

Symptoms

Performing the **rbagent sync-srvradget** command, an error similar to the following might occur:

```
<ERR> Stopping Web extension
Command error: Timeout, RPC: unable to connect
Error raised by ExternalRPCcall in ibm_utils.rbc, line 356 [:0]
Timeout (RPC: unable to connect)
RbAgent command sync-srvradget has failed.
```

when you export a big amount of data.

Causes

The HTTP answer timeout of the agent during the commands sent to the server is four minutes. Depending from the amount of data to export to the .rad file, four minutes might not be enough for the server to complete the export and send the HTTP answer to the agent, so the HTTP timeout error on the agent might occur for big amount of data. For example 10 GB or more.

Resolving the problem

You can increase the timeout value by setting an environment variable named **RBAGT_HTTP_TIMEOUT** in the command shell where the **rbagent sync-srvradget** command is launched. This makes the agent modify the default value with this time (in seconds). The value depends from the data to export. For example, a value of 1200 can be reasonable for 20 minutes.

Agent behavior when the Reboot on Fatal Error option is specified

Symptoms

When the target operates in kernel-free mode and the following conditions are met:

- WinPe is missing the network drivers
- The **Reboot on Fatal Error** option is selected when the agent is rebooted
- The Automatic Driver Binding mode is specified for WinPe

the product automatically attempts to bind the drivers.

Resolving the problem

A new option was added to the **rad-makewpe** command as follows:

```
rad-makewpe [update=EngineItemID] [drivermode=auto|none]
```

If the driver mode is not specified, the default value (none) is used.

Deploying on KVM 6.1 hypervisor fails

Symptoms

The deployment on a KVM 6.1 hypervisor might fail when the NIC model is "Hypervisor Default". The deployment works when the NIC device model is "E1000". During the deployment the KVM guest MAC address changes from XX:XX:XX:XX:XX:XX to 00:00:00:XX:XX:XX.

Resolving the problem

To solve this issue, a DHCP configuration is required.

Software module wizard does not resolve .lnk files contained in folder

Symptoms

When creating a new software module, the wizard asks you to specify the folder containing the required files. The control is then passed to the rbagent, which creates the requested software module. If the folder contains some link files, having the .lnk extension, they will not be resolved by the rbagent.

Resolving the problem

No workaround is currently available for this issue.

How to use hot RAID disks on HP servers

Symptoms

To correctly use RAID disks on HP servers, you must enter the correct syntax for "Physical Drives". To know the syntax to use for the "Physical drives" field of a RAID hardware configuration, first launch a RAID capture. To launch the RAID capture:

1. Right-click the target.
2. Select **Additional features > Capture hardware parameters > RAID capture**
3. From the capture trace, select **Server history > Tasks > Capture hardware parameters**
4. Expand the section.
5. Right-click the specific target entry.
6. Select **Show trace file**.
7. Click **Full log**.

The raidcapt.rbc output is displayed between:

- start of data in raidcapt.ini and
- end of data in raidcapt.ini

Here you find the information about existing physical drives and the syntax to apply for a RAID configuration:

ex:

```
Drive= 1I:1:1, 1I:1:2, 1I:1:3, 1I:1:4, 2I:1:5,  
      2I:1:6, 2I:1:7, 2I:1:8
```

For example, to create an array with three disks, the syntax is as follows:

```
"Physical disks" --> 1I:1:1, 1I:1:2, 1I:1:3
```

Resolving the problem

To correctly use RAID hot spare disks on HP servers, you must enter the information about which array the hot spare drive belongs to, by specifying the

letter of the array at the beginning of every hot spare drive in the list. For example, if you want to add the disk 3 as a hot spare of the first array (A) defined as follows:

Physical drives: "1I:1.1, 1I:1:2"

the correct syntax is the following:

Hot-spare drives: "A:1I:3.1"

Password must be changed in rembo.conf to log in on Solaris Symptoms

When you install Tivoli Provisioning Manager for Images or Tivoli Provisioning Manager for OS Deployment on a Solaris platform, after a successful installation, you cannot log on to the UI.

Causes

The credentials provided during the installation process are not valid.

Resolving the problem

To solve this issue, perform the following steps:

1. Encrypt the desired password using a md5 algorithm.
2. Replace the value for Netpassword in the rembo.conf file located in the installation directory with the encrypted string:
ex: NetPassword "674861BB588C7D56702E52ED86E7E56C"
3. Replace the value for REMBO_PWD in the tpmfosdvars file located in the /etc directory with the encrypted string:
ex: REMBO_PWD="674861BB588C7D56702E52ED86E7E56C"
4. From the installation directory, manually restart the OS Deployment server by running the command:
./rembo -d -v 4 -c ./rembo.conf
5. From the *install_dir* /scripts/solaris directory, restart the rbagent by running the command:
./3_rbagent.sh start

HP ProLiant DL 585 G2 requires BIOS updates

Symptoms

To successfully perform RAID deployments and hardware capture tasks, the HP ProLiant DL 585 G2 hardware requires some BIOS updates.

Resolving the problem

To resolve this issue, perform the following steps on the HP DL585 G2 ProLiant server:

1. Boot the server and go to the BIOS settings.
2. Select the **Advanced Options**.
3. Select the **Linux x86_64 HPET Option** and disable it.
4. Save and close the BIOS settings.

Minimum BIOS version required on some HP laptops

Symptoms

To successfully deploy in kernel mode on the following HP systems:

- HP EliteBook 8470p Notebook PC
- HP EliteBook 8470w Mobile Workstation
- HP EliteBook 8570p Notebook PC

- HP EliteBook Folio 9470m Notebook PC
- HP ProBook 6470b Notebook PC
- HP ProBook 6570b Notebook PC

a minimum BIOS version is required.

Resolving the problem

To resolve this issue, the BIOS update Version F.40 (released in February 2013) , or any later version, must be applied in advance.

Incorrect fonts on the target screen

Names appearing correctly in the OS deployment server cannot be properly viewed on the screen of the target.

Symptoms

Underscores _ appear on the screen of the target instead of the expected name which displays correctly on the OS deployment server.

Causes

The correct fonts are not loaded to the target because the language of the OS deployment server does not use these fonts. The problematic characters are shown as underscores _.

Resolving the problem

Use names with fonts that are compatible with the language of the OS deployment server.

BIOS is too slow to boot from USB disks

Symptoms

The USB drive is not recognized as a boot device by the BIOS during its first boot sequence, and the target does not boot from the USB device.

Causes

Target computers might be too slow before identifying a USB drive, and the USB drive is not taken into account in the boot sequence. This problem only occurs if the target was completely shut down before starting the deployment.

Resolving the problem

If the disk is not recognized, you are required to perform a warm reboot of the target computer (ctrl+alt+del).

Limited partition removal task is not processed by a target booted from USB

Symptoms

Removing the limited partition on a target computer booted from USB fails.

Causes

If you select a target booted from USB and perform the following actions:

1. Select **Additional features** in the contextual menu.
2. Select **Destroy hard disk content** in the additional feature wizard.
3. Select **Limited Partition removal** as **Disposal method**.
4. Follow the remaining instructions from the wizard.

The task is created and is also visible on the task page, but it does not start.

Resolving the problem

When booting from the network, the task is started and ends successfully.

Troubleshooting Java API

If you are experiencing problems, try looking in the following places for errors:

Java API exceptions

Many problems are reported with an exception at the Java API level. For example, improper use of the API or problems contacting the provisioning server.

Exception status in RBActivityTarget

When there is a problem with a task, it is reported in the status field with the method `getError()`.

Task traces

Traces are stored on the provisioning server. By default, the file path is `c:\tpmfos files\global\hosttasks`. The `.ini` files contain parameters for the task and for each target. The log files contain a copy of the console of the target. The target are identified by their BOMID, an identifier used internally by the product. If you know only the IP address, MAC address or serial number (SN) of the target, you have to search for it in all `.ini` files of target activities. Then you can retrieve the corresponding log file. Accessing these traces from the Web interface is more convenient.

Server traces

Server traces are stored on the provisioning server. By default, they are located in the file path, `c:\tpmfos files\logs`. The default detail level is 3, which can be raised to 4 without impacting the performance, or even to 5, but this slows down the server. The easiest way to change this debug level is in the Web interface, or with `RBServerConfig.setGlobalDebugLevel()`. To change the debug level using the Web interface click on **Server parameters**, then click **Configuration** and under the heading **Base parameters** there is a Global debug level which can be edited. If the parameter is present in the `config.csv` file, the value changes next time the server reloads its OS configuration, for example, when it is restarted. The reference for `config.csv` is at <http://www-1.ibm.com/support/docview.wss?uid=swg21247013>. These traces can also be accessed from the Web interface.

Web Interface Extension traces

These traces are located by default in file path, `c:\program files\common files\ibm tivoli\rbagent.log`. Errors occurring when the target of a task is `rbagent` can be found there.

Database traces

Tracing is disabled by default for the product Database Gateway (= `dbgw.exe` = `remboodbc` service). It can be very useful to turn these traces on because all database requests performed by the provisioning server and all task targets pass using this process. Tracing of the SQL queries is only enabled if the process is run from the command line with `"dbgw -d -v 4 -f <log-path>"`. Tracing of the SQL queries and results are enabled if the verbosity parameter is 5 instead of 4. Both slow down significantly the overall performance of the server. The same command line parameters can be used to run the `remboodbc` service with traces: from Windows service manager user interface to start the service manually with Start parameters `"-d -v 4 -f <log-path>"`. The JDBC version of the database gateway (`com.rembo.dbgw.Dbgw`) traces queries and results when started with option `"-d"`.

Database content check

The easy way to inspect database contents is by opening the file `c:\program files\ibm\tpmfosd\autodeploy.mdb` with Microsoft Access. If Microsoft Access is not installed on your provisioning server and you cannot reach the file from another computer, you can still look at database contents directly on the provisioning server with a simple "telnet localhost 2020". Port 2020 is the TCP port used to communicate with our database gateway. This is documented in <http://www.rembo.com/rbo/docs/SQLOpenEx.html>. An updated description of all tables and columns in the product SQL database can be downloaded directly from the server using the URL `http://<server-name-or-ip>:8080/virtual/AutoDeployDistrib.ini`

Server OS configuration

The main OS configuration files are `rembo.conf`, `config.csv` and `radb.ini`.

Offline deployment fails on virtual machines

Deployment fails on a virtual machine when it is started from an OS deployment CD/DVD.

Symptoms

You are using an OS deployment CD/DVD to deploy a virtual machine. The deployment fails.

Causes

The boot order of your virtual machine lists the CD drive before the hard disk.

Resolving the problem

Change the boot order on your virtual machine to ensure that the CD drive is after the hard disk.

Uploading logs to a USB key does not work

While performing a deployment from a deployment media, you cannot upload the deployment logs to a USB key

Symptoms

You are deploying a system profile from a deployment media. You want to upload the deployment logs to a USB key, but it does not work.

Resolving the problem

If you need to send the logs to the support team, use the show trace feature, scroll down and take pictures with a digital camera of the logs. Send the pictures to the support team.

Customizing the timeout duration while cloning from a reference Windows image file

Symptoms

You are creating a new system profile starting from a WIM file. In the last "Profile Wizard" you are asked to wait for the task to be complete. By default, the timeout is 600 seconds. If the time needed to complete the task exceeds this limit, you receive the following error message:



Figure 3. Profile Wizard

Resolving the problem

To solve this issue, customize the timeout duration. To do this, perform the following actions:

1. In the database, go to the table Settings, look for the record where Type = "IFPROBE" and get the value corresponding to the field "IniFile", for example, radset-prbngnst0B1AS9RI.ini.
2. Edit the .INI file, located in the Base directory/global/schemes by changing the value assigned to the timeout, as shown in the following example:

```
[Task]
type="IFPROBE"

[Settings]
StartEnv="AGENT"
timeout="600"
```

The line `timeout="600"` might not be present, for example, in the case of an upgrade. In this case, you can specify a timeout value after the line `StartEnv="AGENT"`. The value is expressed in seconds. For a fresh installation, instead, the INI file is updated.

It is not necessary to restart the server or the Web Interface extension, you can just restart the task.

Customizing the timeout duration while cloning from a reference machine

Symptoms

You are creating a new system profile starting from a reference machine. In the last "Profile Wizard" you are asked to wait for the task to be complete. By default, the timeout is 600 seconds. If the time needed to complete the task exceeds this limit, you receive the following error message:



Figure 4. Profile Wizard

Resolving the problem

To solve this issue, customize the timeout duration. To do this, perform the following steps:

1. In the database, go to the table Settings, look for the record where Type = "OSPROBE" and get the value corresponding to the field "IniFile", for example, radset-dtctprtn0V47XQD6.ini.
2. Edit the .INI file, located in the Base directory for OSD files/global/schemes by changing the value assigned to the timeout, as shown in the following example:

```
[Task]
type="OSPROBE"

[Settings]
timeout="600"
```

The line timeout="600" might not be present, for example, in the case of an upgrade. In this case, you can add the Settings section with the timeout field to specify your preferred value, which is expressed in seconds. For a fresh installation, instead, the INI file is updated.

It is not necessary to restart the server or the Web Interface extension, you can just restart the task.

The user interface for BIOS and UEFI is not localized

Problem description

The user interface present on the target screen for both BIOS and UEFI switcher is not translated into national languages, it is displayed only in English.

Deployment fails on PowerEdge R720 when configured to boot with SD card

Symptoms

When the machine PowerEdge R720 is configured to boot with a internal SD card, the deployment on that machine fails.

Resolving the problem

To solve this issue, remove the internal SD card or to disable the internal SD port via BIOS. For more information, see <http://www.dell.com/downloads/global/products/pedge/en/poweredge-idsdm-whitepaper-en.pdf>.

Conventions

Cue graphic indicate information that is pertinent only for specific operating systems.

Windows

Specific to all supported Microsoft Windows operating systems.

Windows 7

Specific to Windows 7 operating system.

2008

Specific to Windows Server 2008 operating system.

Vista

Specific to Windows Vista operating system.

2003

Specific to Windows Server 2003 operating system.

XP

Specific to Windows XP operating system.

2000

Specific to Windows 2000 operating system.

UNIX

Specific to all supported UNIX-like operating systems.

Linux

Specific to Linux operating systems.

Red Hat

Specific to RedHat operating system.

SUSE

Specific to SuSE operating system.

AIX

Specific to AIX operating system.

HP-UX

Specific to HP-UX operating system.

Solaris

Specific to Solaris operating system.

Notices

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Corporation
274A/101
11400 Burnet Road
Austin, TX 78758
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Canada Ltd.
Office of the Lab Director
8200 Warden Avenue
Markham, Ontario
L6G 1C7
Canada

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information may contain sample application programs, in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Each copy or any portion of these sample programs or any derivative work must include a copyright notice as follows:

(C) (*your company name*) (*year*). Portions of this code are derived from IBM Corp. Sample Programs. (C) Copyright IBM Corp. _ *enter the year or years* _. All rights reserved.

Trademarks

IBM, the IBM logo, and ibm.com[®] are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (or [™]), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at Copyright and trademark information at www.ibm.com/legal/copytrade.shtml

Adobe is either a registered trademark or trademark of Adobe Systems Incorporated in the United States, other countries, or both.

Intel and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States, and/or other countries.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

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

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

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other product and service names may be trademarks of IBM or other companies.

Copyrights

© Copyright IBM Corporation 2012. All rights reserved.

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

IBM web site pages may contain other proprietary notices and copyright information which should be observed.

Portions of third-party software included in this IBM product is used with permission and is covered under the following copyright attribution statements:

- Copyright (c) 1998-2005, The OpenSSL Project. All rights reserved.
- Copyright (c) 1995-2005 Jean-loup Gailly and Mark Adler, the ZLIB data compression library.
- Copyright 1994-2006, The FreeBSD Project. All rights reserved.

The MD5 Message-Digest Algorithm was developed by Ron Rivest. The public domain C language implementation used in this program was written by Colin Plumb in 1993. Permission to use, copy, modify, and distribute this software and its documentation for any purpose and without fee is hereby granted, without any conditions or restrictions. This software is provided "as is" without express or implied warranty.

Portions include cryptographic software written by Eric Young (<eay@cryptosoft.com>). This product because include software written by Tim Hudson (<tjh@cryptosoft.com>).

Notices for tianocore EDK II UDK2010

Copyright © 2004, Intel Corporation All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of the Intel Corporation nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.



Printed in USA