

XenServer Release Notes

Version 5.0.0

Published September 2008

1.0 Edition

Table of Contents

About this document	1
New features in XenServer 5.0.0	1
Installation and upgrades	4
Known issues	4
Hardware and installation	4
Virtual Machines	5
Storage	6
Other	7

About this document

This document lists known issues in the current release of XenServer™ 5.0.0. Where possible, workarounds are described.

Release Notes specific to the supported Windows Virtual Machines are present in the "Installing Windows VMs" chapter of the *XenServer Virtual Machine Installation Guide*. Release Notes specific to the supported Linux Virtual Machines are present in the "Installing Windows VMs" chapter of the same document.

New features in XenServer 5.0.0

XenServer 5.0.0 includes a number of new features and ongoing improvements in hardware support, performance and scalability.

New guest support

The following new guest operating systems are supported for installation:

- Windows Server 2008 32-bit and 64-bit support, with WHQL signed para-virtual drivers and initial enlightenment optimizations.
- Windows XP Service Pack 3 and Vista Service Pack 1 support.
- SUSE Enterprise Linux 9 SP4 32-bit and 10 SP1/2 64-bit support.
- Red Hat Enterprise Linux 4.7 32-bit and 5.2 32-bit/64-bit support, including graphical installation.
- SUSE Enterprise Linux 10 installation directly from CD or ISO images as well as network repositories.

Linux guests have also been refreshed to their latest upstream versions, including the bundled XenServer kernels which fix stability issues not addressed in the upstream releases. Direct installation of Red Hat Enterprise Linux 4.1 and 4.4 guests has been deprecated in favor of versions 4.5 and greater, and is now only supported for P2V conversion. Fresh installations of EL5-based distributions should use version 4.5 or greater to receive the benefit of the latest security updates from the upstream vendor of the EL5 distribution as well.

Business continuity

XenServer supports several features to guarantee service uptime in the event of infrastructure failure. Firstly, resource pools can be configured for automated high availability (HA). This deals with individual host failures by restarting VMs that were running on that host on the next available machine in the resource pool. Notable features include:

- Peer-to-peer "self-healing" architecture ensures there is no single point of management failure.
- VM restart priorities can be set individually, to control the order in which services are restarted in the event of host failure.
- Dynamic failure planning algorithms allow administrators to see how many hosts failures can be tolerated without compromising services.
- Redundant storage (via multipath) and network (via NIC bonding) heartbeats eliminate single-points of network failure and permit reliable detection of genuine host failure.
- E-mail and XenCenter alerts are generated in the event of host failures.

Integrated disaster recovery to enable regular backups of virtual machine metadata. When combined with SAN storage replication, this provides a fast way to migrate entire resource pools to another physical site and continue running services with little interruption. It also permits the use of storage repositories which include all of the metadata for VMs installed on them, permitting a "transportable VM" model across resource pools.

Fibre Channel and iSCSI multipath support, configurable from XenCenter. This ensures that redundant storage links to an FC SAN or iSCSI SR permit link failures without loss of service.

Improved network reliability by support active/active NIC aggregation. Existing active/passive NIC bonds on older installations will be upgraded to aggregation mode for active/active usage, permitting full use of all available bandwidth while still maintaining redundant links.

Multiple management network interfaces can be defined in the control domain, and individual ones can be dedicated for use by network storage (e.g. iSCSI or NFS). This helps improve isolation between VMs and storage infrastructure traffic.

One of the most common reasons for system failure is operator error through misconfiguration. XenServer now provides e-mail and XenCenter alerting for potentially dangerous configurations, for example when VM performance will be degraded or a resource pool is overcommitted with respect to high availability.

New bootloader and xsconsole

The Extlinux bootloader is now used rather than GRUB as in XenServer 4.1.0 and earlier. To apply boot-time options for advanced users:

- at boot-time, press shift to get a prompt
- enter **menu.c32** at the prompt
- press the **TAB** key within 5 seconds of the menu appearing, and edit the settings

To make changes permanent, edit the `/boot/extlinux.conf` file.

Once booted, instead of a Linux login prompt, XenServer now displays *xsconsole*, a convenient system configuration console.

Storage and provisioning

Support has been added for QLogic and Emulex 8GB Fibre Channel HBA devices, as well as bug-fixes and stability improvements for other HBA devices.

Diskless virtual machines running on Citrix Provisioning Server now support more advanced PXE configurations, such as DHCP Proxy ARP responses or wide area DHCP relays. Support for PXE servers such as Altiris and Windows Deployment Services (WDS) has been improved.

There is CLI-only support for quiesced, fully consistent VM storage snapshots for Windows VMs. The storage backend must support fast-cloning and snapshot, e.g. the NetApp and Equallogic storage repositories.

The VHD-based storage repositories have undergone significant tuning, improving performance and reducing resource utilization for the local VHD and NFS storage repositories.

This release of XenServer 5.0.0 provides a preview of storage repository support for Dell EqualLogic arrays. This storage repository type supports fast-cloning, snapshots and high-performance LUN-per-VDI access by directly controlling the storage array.

Usability and reliability

Performance statistics are now persisted in a round-robin database on the XenServer pool, with decreasing granularity as time passes. This enables long-term trending and resource optimization for complex data centre deployments. The new scheme is also far more scalable across large resource pools, and so the older mechanism for accessing metrics has been deprecated for this release. In order to activate the legacy metrics support on your resource pool, please refer to the *Persistent XenServer Performance Statistics* in the *XenServer Software Development Guide*.

XenCenter supports connecting to both 4.1 and 5.0.0 hosts, and has several new features:

- Host networking configuration support, including dedicating network interfaces for storage use.
- Permit editing settings on existing SRs, and improved workflow for detaching and reattaching SRs.
- Support user-defined grouping and searching across VMs, hosts and resource pools, including defining custom fields and tags to identify resources.
- New interactive graphing interface to display persistent performance statistics across resource pools.
- Support creation of Fibre Channel and Equallogic storage repositories directly from XenCenter.
- Detect when a 5.0.0 host has been re-installed and warn the administrator.

Enhanced support for keeping hosts up to date with the latest improvements:

- Rolling zero-downtime resource pool upgrades from 4.1 to XenServer hosts.
- XenCenter support for automatically checking for updates and generating alerts for the administrators if any are available.

A menu-driven text console is now present when connecting to the main screen of the XenServer host. This new interface provides an alternative to XenCenter for common operations, or when networking is not otherwise available. The boot process is also much more user-friendly now, with a graphical splash screen and progress indicator.

Performance and hardware

The core Xen hypervisor has been upgraded to a version based off the stable 3.2 releases. Major improvements are present in:

- Improvements to emulating legacy 16-bit code means that a wider variety of older applications and bootloaders will now successfully run. Support for foreign-language versions of Windows have also been improved.
- The Windows paravirtual storage drivers now utilize a SCSI filter interface which can reduce overhead and increase performance when used with multiple virtual disks.
- Citrix XenApp performance has been further optimized with improvements to Xen shadow page-table handling.

- Reduced memory usage in the control domain increases the maximum number of VMs supported.

Many hardware device drivers have been refreshed to their latest stable versions, including MegaRAID SAS, MPT Fusion, Broadcom and 3Ware RAID cards. For 10Gb networking adapters, support has been added for Neterion and Broadcom cards, and Chelsio performance has been significantly improved.

Installation and upgrades

System requirements, preparation, installation, and initial configuration are described in the *XenServer Installation Guide* document.

Note

XenCenter from release 4.0.1 or 4.1.0 does not work with XenServer hosts from release 5.0.0. However, XenCenter from release 5.0.0 *does* work with XenServer hosts from release 4.1.0.

This release fully supports upgrading from XenServer 4.1.0 hosts. Simply install the 5.0.0 installation CD and select the Update option from the on-screen menus. You can also upgrade 4.1.0 resource pools without any VM downtime by using the rolling upgrade procedures described in the *XenServer Installation Guide*. Note that rolling upgrades are *not* supported when upgrading from beta versions of XenServer 5.0.0. For upgrades from beta you should upgrade all the hosts in the pool from the installation CD and then bring up the pool master and slaves.

The Windows Xen VSS provider is not installed by default (unlike in beta releases), and must be specifically installed in order to enable quiesced VM snapshots via the `install-XenProvider.cmd` script. Please refer to the *Virtual Machine Administration Guide* under the Windows section for full details on how to do this.

The control domain uses slightly more memory in XenServer 5.0.0 relative to the previous releases, and will automatically allocate approximately 100MB of host memory to the control domain upon upgrade. If you were using all available host memory prior to the upgrade, reduce the amount of memory used by one VM by 100MB before powering it on.

Known issues

This section details known issues with this release, and any workarounds which can be applied. Please report any other issues via your Citrix support representative.

Hardware and installation

CA-933

On machines with multiple disks, the XenServer boot loader is only written to the disk selected for installation.

For example, if you set the boot drive to the first disk, and then install the XenServer software on the second drive, and then subsequently change the system BIOS to assign the second drive to be the boot drive, it will fail to boot and properly renumber the disks. Installation to the second disk only works if the first disk has a boot loader already present which lets you boot off second disk, or if the BIOS permits you to boot directly off the second disk. Renumbering disks post-installation will not work. The only workaround is to manually edit `/boot/extlinux.conf`.

CA-4575, CA-5311

On servers with the ICH8 South Bridge chipset, XenServer might not detect the drive at installation due to a problem with the `ata_piix` driver. If you have this problem, set the machine's ATA/IDE mode in the Advanced section of the BIOS to use the Advanced Host Controller Interface (AHCI). This is usually set to off (legacy mode) by default. You might need to set the mode to "Native" to then access further SATA configuration options that allow selecting between AHCI, IDE, or RAID. It might also be available via Advanced->Drive Configuration->Disable Intel® RAID Technology and then Enable SATA AHCI mode. The Dell SAS 5/iR Controller in Dell 490 hosts exhibits this issue.

For more information, see <http://www.intel.com/support/chipsets/imst/sb/cs-015988.htm>.

CA-8767

Motherboards using the Intel 965 chipset with more than 2GB of memory may fail to boot successfully. This has been identified as a BIOS firmware issue, and appears to happen on any 64-bit operating systems (see [Red Hat](#) and [Microsoft](#) related bugs). To workaround this, downgrade your BIOS to version 1669, available from the Intel website.

CA-15337

When doing an upgrade or restoring from an upgrade, users should ensure that removable storage not in use by the XenServer installation is removed, in particular USB flash and hard drives. In particular, remote management tools like HP iLO or Dell RDAC may also cause a virtual USB device to appear in the host, so be sure to disable this functionality before upgrading your server.

CA-20644

When configuring bonds in a pool using XenCenter, the MAC address used may be from an unexpected slave interface of the bond rather than the first, which can cause a different IP to be issued if DHCP is being used. Therefore, it is recommended that a static IP configuration is used in deployments involving NIC bonds.

Virtual Machines

CA-9772

The Windows PV drivers do not send a gratuitous ARP after live relocation if the guest has previously been hibernated. Note that hibernation is not a supported use-case, since direct suspension of a VM is supported instead.

CA-15280

Users should refer to the Administrator Guide when making use of disk QoS settings within XenCenter for information on how these apply to different storage types. Not all storage types fully support setting QoS.

CA-19382

When importing VMs using the `xe` CLI, networks on which the virtual interfaces (VIFs) get connected to are matched by their name on the server the VM was exported from. The default name of the standard pool-wide networks changed between versions 4.0 and 4.1, which means that it may be necessary to re-create the virtual interfaces associated with a VM on the desired network. If using XenCenter, the Import wizard will allow you to remap interfaces as desired as part of the import operation.

CA-22247

In a resource pool, the `xe` CLI commands `vm-copy` and `vm-install` with either `sr-uuid` or `sr-uuid` specified will fail intermittently if some hosts are offline when they are attempted. To work around this issue, repeat the command with the hosts powered on. Alternatively, repeat the command until it succeeds. (The latter works because the virtual disk copy operations invoked by these commands are forwarded by the pool master to any of its member hosts, selected at random.)

CA-22639

On rare occasions the xapi agent might fail to restart after you execute the `xe` CLI commands **pool-emergency-transition-to-master** or **pool-emergency-reset-master**. This will result in XenCenter not being able to connect to the server you have made into the new pool master. Rebooting the server should solve this problem.

Storage

CA-6966

When using NFS ISO storage repositories, a hard mount is used to communicate with the server. This means that the control domain can hang if the remote NFS server becomes unreachable. The workaround is to use CIFS-based mounts instead.

CA-9208

Citrix has seen data corruption issues using the iSCSI target provided by Adaptec SnapServers. This appears to be a problem with the SnapServer iSCSI implementation, and has been reproduced by Adaptec using a standard (non-XenServer) Linux distribution. We are currently working with Adaptec to find a solution to this problem. Until this issue is resolved, XenServer users are strongly encouraged to use NFS rather than iSCSI storage repositories when using SnapServer products. In general, when using network-based storage hardware, users should ensure the software and/or firmware on the devices being used is up to date, as recommended by the manufacturer.

CA-12866

Users should avoid attaching read-only VDIs to Windows VMs as this may result in unexpected behavior and instability of the virtual machine. This includes NetApp snapshot VDIs; users wishing to attach a snapshot VDI to a Windows VM should first make a read-write copy of the snapshot using the **vdi-copy** CLI command

CA-18965

When a server is ejected from a resource pool, some storage repositories from the ejected server may be shown in XenCenter in a broken state. This is also true if the eject operation was performed before an upgrade from the previous version. This is normal behavior as the storage can no longer be seen by the pool. These storage repositories can be 'forgotten' by right clicking on them and choosing Forget Storage Repository, or by using the **sr-forget** CLI command.

CA-20412

The Dell EqualLogic Storage Repository plugin may cause temporary lockups on the EqualLogic storage array when control plane operations are executed in rapid succession. If this happens, you should slow down the rate of operations being sent to the storage backend. This issue only affects control operations while manipulating storage, and high volume data operations over the iSCSI channel are unaffected. Please contact Dell support for a workaround and a newer firmware version if you experience this issue.

CA-22676

When dedicating a network interface as a storage interface for use with iSCSI or NFS SRs, you must ensure that the dedicated interface uses a separate IP subnet which is not routable from the main management interface. If this is not enforced, then storage traffic may be directed via the main management interface after a host reboot, due to the order in which network interfaces are initialized. If you do require static routing to a shared subnet, then you need to re-plug the storage PBDs after each host restart, or use the **xe pif-forget** command to manually configure up networking rules in the control domain.

CA-22692

Quiesced snapshots taken on Windows Server 2008 guests will not be directly bootable. Instead, you must attach the snapshot disk to an existing Windows Server 2008 VM to access files for restoration purposes.

CA-22632

Snapshots initiated directly from within a Windows guest using the `vshadow` development utility will not be directly bootable. A snapshot taken in this way appears indistinguishable from one taken using the **`xe vm-snapshot-with-quiet`** CLI command, but the latter snapshot taken using the XenServer CLI or XenCenter will boot.

Other

CA-20462

Triggers for alerts are checked at a minimum interval of five minutes (this avoids placing excessive load on the system to check for these conditions and reporting of false positives); setting an alert repeat interval smaller than this will result in the alerts still being generated at the five minute minimum interval.

CA-22686

When configuring e-mail alerts, your mail server may require that the envelope From address of the mail originates from a fully-qualified domain name. Alerts generated by XenServer hosts use the machine hostname, which may not be fully qualified. If this happens, use the **`xe host-set-hostname-live`** CLI command or XenCenter to make it a Fully Qualified Domain Name.

CA-22784

When HA is enabled, users should not attempt to map physical CDs into the virtual CD drive of a protected VM, as this means the VM is no longer agile so cannot be started on another server.

CA-21737

If you use the **`pool-emergency-reset-master`** CLI command and set the `master-address` parameter to the current master address, the master host will fail and the XenCenter connection will be lost. Further attempts to connect through XenCenter will cause an "unknown error" to be reported. To remedy this situation, connect to a console on the affected host and run the **`xe pool-emergency-transition-to-master`** command. Note that you should never have to run the **`pool-emergency-reset-master`** command on a master host, as it is intended to point a slave host at a new master address.

CA-22786

When enabling NIC bonding, you may see a message printed to the host console indicating that the bond MAC address is still in use. This is a transient message shown during the transition from separate interfaces into the bonded state, and can be safely ignored.