XenServer 6.2.0 User Security

Published November 2013
1.0 Edition

XenServer 6.2.0 User Security

Version: 6.2.0

Citrix, Inc.
851 West Cypress Creek Road
Fort Lauderdale, FL 33309
United States of America

Disclaimers

Trademarks

Citrix®
XenServer ®
XenCenter ®

CITRIX

Contents

1. Introduction 1

1.1. Architecture 1

2. Host security 3

2.1. Networking Configuration 3

2.1.1. Administration Network 3

2.1.1.1. Resource Pools 3

2.1.1.2. Storage Network 4

2.1.2. Virtual Machine Network 4

2.2. Storage Configuration 4

2.3. Accessing XenServer Hosts 5

2.3.1. Verifying Host Identity 5

2.3.2. Replacing the SSL certificate of a host 5

2.4. Updating the XenServer Host 6

2.4.1. Hotfix Format 6

2.4.2. Updating using XenCenter 6

2.4.3. Updating using the CLI 7

3. Guest Security 8

3.1. Hypervisor Protection 8

3.2. Guest Communication 8

3.3. Guest Consoles 8

3.4. Operating System Recommendations 8

3.4.1. Microsoft Windows 9

3.4.2. Linux 9

CITRIX

Chapter 1. Introduction

XenServer is a server virtualization platform that can run multiple operating systems simultaneously. This consolidation both reduces the operational costs and increases the agility associated with running IT infrastructure. Although XenServer aims to be secure by default and requires little out-of-the-box configuration with respect to security, you should be aware of how it works and how to maintain a secure installation. This document provides an outline of the security architecture of XenServer, and pointers to maintenance procedures and best practices for keeping your XenServer deployment secure and reliable.

References

[XS Guest] *Citrix XenServer 6.2.0 Virtual Machine User's Guide*.

[XS Admin] *Citrix XenServer 6.2.0 Administrator's Guide*.

[XS Install] *Citrix XenServer 6.2.0 Installation Guide*.

[SDK] *Citrix XenServer 6.2.0 Software Development Kit*.

1.1. Architecture

XenServer consists of several different components:

- The *Xen hypervisor* is the first software to load when XenServer boots. It runs in 64-bit mode and virtualizes the CPUs, interrupts and host memory. Xen is a very thin layer of software with minimal I/O functionality and consists of around 150,000 lines of code.
- The *control domain* boots next, which is a 32-bit Linux-based embedded distribution. The control domain is a normal XenServer VM that has additional privileges granted to it which allows it to control host hardware devices and also create further guest domains.
- The *XAPI management stack* runs inside the control domain and manages all resources required for running guest domains. It consists of a distributed database and control software which listens on the administration interface for XenAPI clients that issue control instructions.
- The *Storage Manager* (or SMAPI) runs inside the control domain and provides a consistent interface to a variety of storage backends, such as Fibre Channel, iSCSI, file-based VHD disks, or local storage.
- The *XenCenter* user interface is a Windows application that provides a user-friendly way to manage large numbers of XenServer hosts.
- XenServer can run multiple *Virtual Machines* on the same host, each of which provides isolated computation, storage and networking to the operating system running inside of it.
- Finally, multiple XenServer hosts can be aggregated into a *resource pool* which acts as a single unit of administration across a cluster of machines.

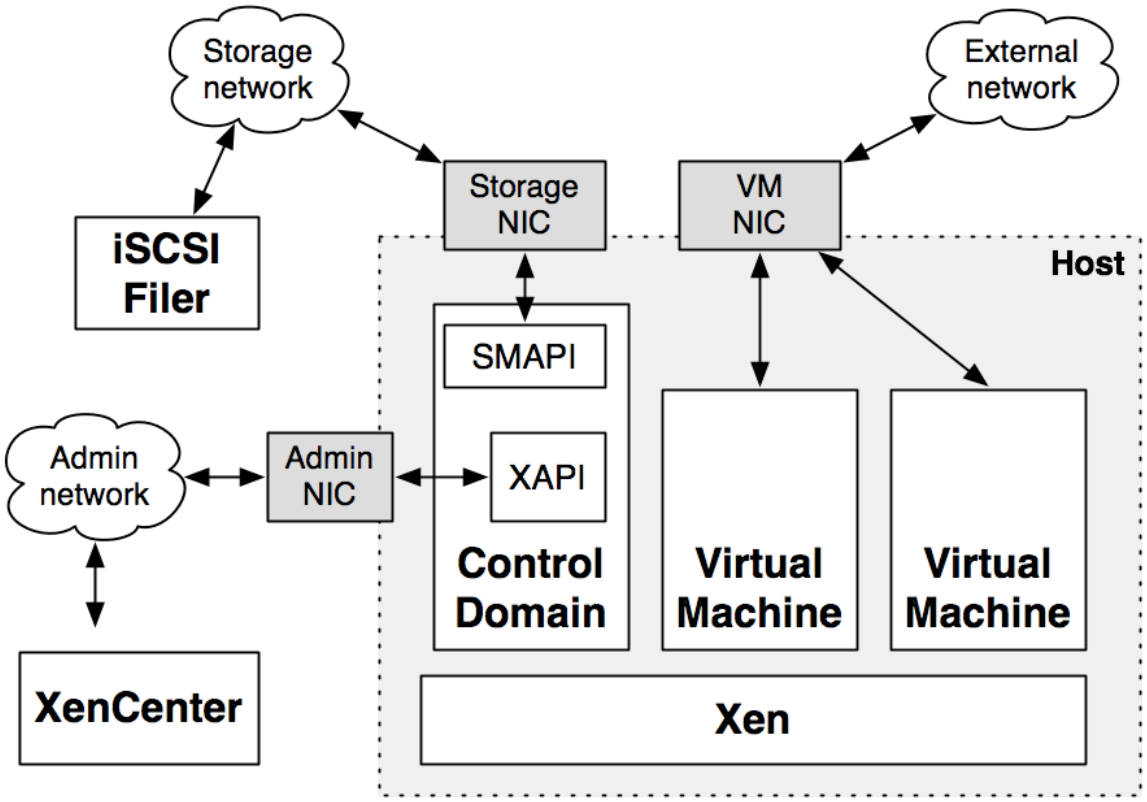

These concepts are illustrated in the above diagram which shows the components in a single physical host. The following chapters describe:

- how to secure the XenServer host itself by ensuring that the control domain is correctly configured and patched
- best practices for storage configuration
- best practices for securing VMs

CITRIX

Chapter 2. Host security

The XenServer control domain provides administrative access to a resource pool. If access to this control interface is compromised, attackers could control virtual machines and the storage and networking layers.

2.1. Networking Configuration

The control domain is initially configured during installation (see Section 3.2, “Installing the XenServer Host” [XS Install]) and subsequently by using the XenCenter GUI. There are three distinct types of networks which may be configured in the control domain: the administration network, the storage network, and the VM network.

The three networks can also run on the same physical NIC, and are initially set up like this after installation. For maximum security, you can assign a set of NICs for dedicated use for administration, storage and VM traffic. These dedicated NICs can also be paired up as NIC bonds to ensure maximum resilience against physical component failure. See Section 4.3.4, “NIC Bonds” in the [XS Admin] for more information.

2.1.1. Administration Network

The administration network is used for the following functions:

- control requests from the XenAPI
- Virtual Machine live relocation (or *XenMotion*)
- Virtual Machine live storage relocation (or *Storage XenMotion*)
- VM import, VM export, hotfix application, and resource pool metadata backups.
- sending e-mail alerts
- intra-resource pool communication

The XAPI tool-stack listens on port 80 (plain-text) and port 443 (SSL encrypted) for XenAPI requests. XenCenter exclusively uses the SSL port to ensure that all control traffic to and from the host is encrypted, but other third-party XenAPI clients may not. To be sure that all traffic is encrypted, add rules to the routers on the administration network to block external requests from port 80.

The XAPI tool-stack itself is written in a high-level, statically type-safe language known as Objective Caml (or *OCaml*). This provides protection against low-level memory corruption issues such as buffer overflows or integer overflows, making it much more robust against malicious attacks over the administration network. The SSL layer uses the popular `stunnel` package to provide industry-standard encryption.

VM live relocation involves transferring the memory image of the VM while it is still running. Since a high-performance transfer will minimize the performance impact on the running VM, and live relocation is only supported between machines on a local network, this transfer occurs unencrypted. Bear in mind that if you do configure XenMotion across WAN links that you will need to use IP-level security (e.g. IPsec) to encrypt the memory image.

VM live storage migration involves transferring the storage of the VM while it is still running. This transfer also occurs unencrypted so if you intend to perform a live storage migration over an insecure network you may wish to configure a secure tunnel across the insecure parts (e.g. by using IPsec).

It is possible to avoid having an IP address assigned to the administration network interface, which will mean that none of the administration functions will work from outside of the local console on the XenServer host. However, you should be aware that in this configuration you will not be able to create resource pools, import/export VMs, or otherwise take advantage of features such as remote logging.

2.1.1.1. Resource Pools

Multiple XenServer hosts can be clustered together into *resource pools*. These resource pools are assigned a *pool master* which controls all the other hosts. All communication between hosts within resource pools is done over

SSL, and hosts authenticate themselves to each other using a randomly generated symmetric key that is created at the time of pool creation.

A common way to isolate this administration traffic is to use Ethernet VLANs to segregate it from other non-administration traffic. You can configure your routers to tag all traffic from the administration NICs with an administration VLAN tag. This VLAN can also be used for other appliance control traffic in your server farm, such as Citrix Provisioning Server or Citrix NetScaler. Note that remote clients such as XenCenter will need to connect to the administration network VLAN. If you choose to use this feature of your routers, it is advisable to check that they maintain VLAN separation under load and/or fault conditions

2.1.1.2. Storage Network

If you are using IP-based storage backends, such as an NFS or iSCSI SR, the storage traffic also flows through the control domain. Storage traffic is not encrypted by XenServer, and so it is important to isolate this traffic from other administration traffic for both reliability and security.

Dedicated storage NICs require an IP address that must be on a different IP subnet to the main administration interface. You can use XenCenter to dedicate a storage NIC and bond multiple NICs together for resilience.

In the case of iSCSI traffic, the software OpenISCSI initiator is used to connect to iSCSI targets. This initiator supports CHAP to authenticate to the remote target, and you should configure this on your iSCSI filer if other third-parties should also be able to connect to the filer and access the VM data.

The file-based NFS storage repository mounts the specified NFS repository on the control domain. Each SR is represented as a directory on the remote NFS server, with each virtual disk being a file in that directory stored in the VHD file format. VHD is not an encrypted file format, and so you should ensure that only the XenServer hosts and authorized administrators can mount the filesystem. For further security, ensure that the storage interfaces of the remote filer are not visible outside of the dedicated storage network.

2.1.2. Virtual Machine Network

The VM network does not require an IP address in the control domain. Instead, VM network packets are *bridged* at the Ethernet layer over the host NIC assigned to the virtual network interface in the VM. You may use XenServer's Port Locking features to restrict the ability of VMs to masquerade using different addresses. See Section 4.4.10.8, “ Using Switch Port Locking” [XS Admin] for more information.

A common configuration is to have different virtual network interfaces inside a VM for "front end" traffic (for example, a web server) and "back end" traffic (for example, a database). This traffic may be isolated by using VLANs, which will tag the Ethernet traffic separately but still go over the same physical NIC on the host (see Section 4.3.3, “VLANs” [XS Admin]).

2.2. Storage Configuration

Storage on the control domain is automatically configured during installation and subsequently by using the XenCenter GUI and for some actions using the CLI. XenServer supports several types of storage repositories (SRs) and virtual disk image (VDI) types (see Chapter 5, *Storage* [XS Admin] for a complete list of SR and VDI types).

Most of the publicly supported SR-types available in XenServer 6.2.0 are designed so that, by default, data stored in a deleted VDI won't be available in a new VDI, even when they occupy the same physical space on the underlying storage media. This includes the following SR-types:

- Local LVHD (*lvm*)
- LVHD over FC (*lvmohba*)
- LVHD over iSCSI (*lvmoiscsi*)
- Local EXT3 VHD (*ext3*)
- NFS VHD (*nfs*)

- iSL NetApp Adaptor (*netapp* or *cslg*)
- iSL EqualLogic Adaptor (*equal* or *cslg*)

Note that iSL relies on the array security to ensure that data in a previous LUN is not re-exposed into a new LUN.

The contents of USB disks and ISO images are not deleted when their corresponding SRs are destroyed.

2.3. Accessing XenServer Hosts

Each XenServer host generates a set of random cryptographic keys when it is installed. Each key is split into two portions: a *public* key which is displayed to clients, and a *private* key which is only known to the host and can be used to prove that it is the owner of the public key.

The two classes of keys generated by a XenServer host are for the Secure Shell (SSH) protocol, and for the XenAPI SSL network service. SSH is only used for advanced configuration of the control domain, and should not be needed in normal use. The SSL communication is used much more often (e.g. by XenCenter).

The main use of these keys is to confirm that you are connecting to the correct host. Although the SSL communication is encrypted, you also need to ensure that the host you are communicating with is actually the one you think you are connected to. For this reason, XenCenter has special support to warn you when you are connecting to a brand-new XenServer host, and will alert you if a previously-seen host's SSL identity suddenly changes.

2.3.1. Verifying Host Identity

Follow these steps to find the SSL public key:

1. At the physical host console, activate the menu-driven console, if not already running, by typing:

   ```
   xsconsole
   ```

2. Navigate to **Status Display** and note the key fingerprint for HTTPS.

Follow these steps when you first connect to the XenServer host from XenCenter:

1. XenCenter displays a message informing you that the host has not been seen before.
2. Compare the entirety of the key fingerprint displayed in XenCenter to the one you noted down at the host console. The key fingerprint may be obtained from the **New Security Certificate** dialogue by selecting **View Certificate** and examining the Thumbprint value under the Details tab
3. If they are the same, then the host you are connecting to is the correct one and you can safely continue.
4. If the keys are different, decline the dialog box and investigate why the host XenCenter is connecting to is not the one you think it is.

Once you have connected to a XenServer host for the first time, XenCenter caches the public key and associates it with the host record. If the host key changes (for example, it is upgraded or re-installed), a warning is displayed showing the new host key. Follow the same steps as for the initial connection to ensure that you are not under a *man in the middle* attack by an attacker trying to impersonate a XenServer host in order to obtain your authentication credentials.

2.3.2. Replacing the SSL certificate of a host

Citrix recommends that you replace the default certificates on hosts by certificates adhering to the standards of the organization where XenServer is deployed. The replacement file must contain the certificate and private key in PEM format.

To install the new certificate, first move the current SSL certificate:

```
mv /etc/xensource/xapi-ssl.pem /etc/xensource/xapi-ssl.pem_orig
```

Then copy the new certificate into its place:

```
cp <cert>.pem /etc/xensource/xapi-ssl.pem
```

After this reboot the XenServer host.

2.4. Updating the XenServer Host

XenServer hosts are updated by applying *hotfixes*. These hotfixes are cryptographically signed code bundles which update portions of the control domain and/or the hypervisor. The control domain is an embedded Linux distribution with many customizations and modified packages.

Installing upstream Linux packages directly into the control domain will conflict in many cases, and may result in a non-functional product. The sole supported update mechanism for XenServer is through hotfixes issued by Citrix. Configuration changes which are not explicitly documented or approved by Citrix Technical Support may not have been tested and are therefore not supported. In addition, configuration changes may not persist after installation of a hotfix or upgrade, and could also cause a hotfix or upgrade to fail. In particular, Citrix strongly recommends that you do not modify the control domain's internal firewall configuration without consulting Citrix Technical Support

2.4.1. Hotfix Format

XenServer hotfixes are cryptographically signed by Citrix before being issued. When they are uploaded to a XenServer host, this signature is checked and any invalid hotfixes are immediately rejected. Any third-party modifications to a hotfix will result in the signature being invalidated, thus providing protection against the introduction of malware through a hotfix.

The hotfix also includes the following metadata embedded in it:

1. A unique UUID which identifies the hotfix.
2. A pre-check function which ensures that the hotfix is relevant to the XenServer host it has been uploaded to. This function commonly checks the version number and checksums of important files to ensure that it will apply cleanly.
3. Post-install *guidance* which represents actions which need to be taken to activate the hotfix after it has been applied, if any. These include host rebooting, restarting some guests (e.g. all Windows guests), or restarting the management tool-stack.

2.4.2. Updating using XenCenter

If your XenServer is licensed, you can apply updates using XenCenter. XenCenter includes an Install Update wizard which automates the process of applying hotfixes across multiple physical hosts in a resource pool. This wizard works in two modes: automatic and manual. The automatic mode live migrates virtual machines away from hosts sequentially, emptying the host of VMs before applying a hotfix. Once the host has been updated and rebooted, the original VMs are migrated back onto it, and the next host is upgraded. This mechanism means that VMs will not have any downtime if a hotfix is applied across the pool. There is also a manual mode which allows the administrator to control which hotfixes are applied by hand.

XenCenter can be configured to automatically check for new hotfixes. It performs this operation by regularly polling the Citrix XenServer updates web page for an XML file which contains a list of all the publicly issued hotfixes. It then compares these hotfixes to the versions installed in the resource pools connected to XenCenter, and generates an administrator alert if any out of date hosts are found.

XenCenter is also able to apply previously applied hotfixes to new hosts without further downloads from the XenServer updates web page. This means that hosts added to XenCenter in the future may be updated using the update manager.

Please refer to the XenCenter help for more information on how to update hosts and pools using XenCenter.

2.4.3. Updating using the CLI

The XE command-line interface also provides full support for applying hotfixes, which is a good way to integrate XenServer hotfix management with any existing configuration management software you may be using for other infrastructure.

The hotfix must be installed on all Citrix XenServer hosts, including each host in a resource pool. For more details, please refer to Chapter 8, *Applying Updates and Hotfixes to XenServer* [XS Install].

CITRIX

Chapter 3. Guest Security

The primary rule for guest operating systems running within XenServer is to ensure that you follow normal security practice on those guests. Install virus scanners within Windows, activate packet filters and firewalls, and keep your packages up-to-date in Linux distributions.

This chapter describes the protection that XenServer uses to isolate VMs from each other.

3.1. Hypervisor Protection

XenServer contains a 64-bit hypervisor that runs guests in two modes: para-virtualized and hardware assisted. In both modes, the hypervisor provides strong isolation against CPU instructions running in one guest affecting the state, including memory, of another guest.

Paravirtual (PV) guests run kernels which have been specially modified to be aware of the XenServer hypervisor. A para-virtualized kernel runs in a deprivileged mode: a 32-bit PV kernel runs in hardware ring1 while a 64-bit PV kernel runs in ring3. All of the pagetables are owned by the hypervisor which provides strict checks when guest kernels request an alteration (using the *hypercall* interface provided by the hypervisor).

Hardware assisted mode (HVM) uses the Intel VT-x or AMD-V hardware extensions to offload the effort of virtualization onto the hardware. Guests in HVM mode see what they believe to be a complete physical machine. When hardware detects that the HVM domain has attempted to perform an operation which requires external help to be serviced correctly, it switches to the hypervisor which applies appropriate safety checks and actions.

HVM mode does not automatically provide block and network devices; instead, these are initially emulated via a `qemu-dm` helper process which runs in the control domain. To protect against bugs in this program compromising the control domain, XenServer runs each `qemu-dm` inside a Linux chroot with a unique unprivileged process ID. Later on in the boot process, the high-performance paravirtual drivers are activated which switch away from the emulated devices to high-speed virtual channels which use a similar mechanism to para-virtualized guests to communicate with the outside world with minimal overhead.

3.2. Guest Communication

Guests communicate control information and flags to and from the control domain via a tree known as *Xenstore*. The Xenstore process runs in the control domain and has a comprehensive Access Control List (ACL) mechanism which limits read, write and access permissions to various portions of the tree. When VMs are started, the control domain writes some control values into Xenstore which are read-only to the guest, and the guest can write into a sub-tree of its namespace in order to communicate information back to the control domain (e.g. its IP address, or performance metrics). Access control is used to make sure that the guest's sub-tree isn't readable by other guests.

3.3. Guest Consoles

Access to VM consoles is provided over the VNC protocol. The exact mechanism varies depending on whether the guest is running in PV or HVM mode. In the case of PV guests, a `vncterm` process in the control domain makes the text console available over VNC for rendering in XenCenter. The `vncterm` process runs in a Linux chroot with a unique process ID to protect against bugs in the console emulation resulting in a compromise of the control domain.

For all guests, the VNC server is exposed via a TCP port bound to the `localhost` interface in the control domain. This precludes direct external access to the port, which is also firewalled within the control domain. The XAPI toolstack then exposes this interface via the XenAPI securely over SSL, ensuring only authorized users can connect to the guest. Note that any user logged into the control domain who can access the `localhost` interface will be able to access all the VM consoles, so do not grant shell access to the control domain to untrusted users.

3.4. Operating System Recommendations

This section covers security recommendations specific to a particular guest operating system.

3.4.1. Microsoft Windows

XenServer includes a set of WHQL-certified disk and network drivers for all supported versions of Windows. Ensure that you have the most up-to-date drivers installed in the guest.

It is easy to determine if you are up-to-date, since XenCenter will display a warning in the General tab if the drivers are missing or out-of-date. If this is the case, simply insert the `xs-tools.iso` CD image into the guest and install the Windows tools to refresh them to the latest versions.

Citrix recommends that you follow Microsoft's advice for securing your Microsoft Windows guest operating systems.

3.4.2. Linux

Most Linux distributions contain kernel support for running on XenServer and so Citrix recommends that you follow your chosen distributions' advice for securing your Linux guest operating systems.