



VMware ESX Server 3

Best Practices for VMware ESX Server 3

Summary

VMware ESX Server can be deployed in many ways. In this document, VMware recommends specific guidelines for the deployment of VMware ESX Server 3. Following these guidelines will maximize the ease of use of ESX Server in your environment.

What Is a Best Practice?

VMware ESX Server has been shipping since 2001. During that time, VMware and its customers have developed significant operational experience with ESX Server software and have documented a number of techniques that are most likely to result in a highly successful deployment. These techniques are called best practices, and this document describes them.

Please note that these practices are only recommendations, not requirements. Not following these recommendations does not affect whether your system is supported, and some recommendations may not apply to your circumstances. However, VMware believes everyone will benefit from thinking through these recommendations and making a conscious decision to follow or reject them.

Definitions

VMkernel: The core of the ESX Server product, the VMkernel is a proprietary microkernel that separates the virtual machines from the underlying hardware.

Service Console: The service console, sometimes called the "Console OS" or "COS," is the ESX Server command-line interface. It provides the ESX Server management tools, includes a command prompt for direct management of ESX Server, and keeps track of all the virtual machines on the server as well as their configurations. The service console in ESX Server 3 is a customized version of the Linux operating system based on Red Hat Enterprise Linux 3, Update 6 (RHEL 3 U6).

While the service console remains available for advanced administration tasks, most tasks can now be performed through the Virtual Infrastructure Client (VI Client) or using Virtual Infrastructure Web Access (VI Web Access).

Virtual machine configuration file: This is a text file that declares the virtual hardware that composes a virtual machine.

VMFS: The VMware File System. This is a high-performance, dedicated file system, capable of simultaneous access by multiple ESX Server systems, and on which virtual disks for virtual machines are stored as files.



Plan Your Virtual Infrastructure

Before deploying VMware ESX Server 3 and other virtual infrastructure products, it is useful to plan a number of aspects of the installation. This section addresses the following recommendations:

- [Plan Your Deployment Before You Install on page 3](#)
- [Plan Your Licensing Structure on page 5](#)
- [Plan Your User Structure on page 6](#)
- [Plan Your Network Structure on page 7](#)
- [Plan Your Disk Storage Structure on page 8](#)



Plan Your Deployment Before You Install

Before you install the ESX Server software, consider the virtual machines you plan to run on this instance of ESX Server. How many virtual machines will there be, what resources will they use, and what mix of applications will they support?

Load Mix

Assess the load mix you plan to put on this ESX Server system. Is your server hardware sized appropriately? If you are consolidating existing servers, profile the CPU and I/O loads on the systems to be consolidated. Otherwise, if the virtual machines are to be new deployments, examine the utilization of systems in similar roles. Choose server hardware that is capable of sustaining the total average load (if individual virtual machine peak loads tend not to occur simultaneously) or the total peak load (if the peaks do tend to occur simultaneously). VMware's Professional Services Organization can assist you with this task as a consulting engagement. For more information visit <http://www.vmware.com/services/consulting.html> or write to consulting@vmware.com.

Note that a heterogeneous load mix will offer you better system utilization than will a homogeneous load mix. For example, if some of your virtual machines are CPU-intensive and some I/O intensive, you will be able to use more of your total server capacity than if all the virtual machines have the same limiting resource.

Virtual SMP

ESX Server now supports as many as four virtual processors per virtual machine. When a virtual machine will be running a single-threaded application, however, consider configuring that virtual machine as a uniprocessor system, thus avoiding the overhead of an SMP operating system. Even when a virtual machine is configured with multiple virtual processors, configure it with only as many virtual processors as it will actually use. Unused virtual processors still consume some system resources, potentially slowing down other virtual machines on the system.

When configuring virtual machines, remember that the ESX Server software itself also has some overhead, and take care not to overcommit processor resources.

An SMP virtual machine can only be powered on when the number of available physical or logical processors equals or exceeds the number of virtual processors for which that virtual machine is configured. This means that not only does the underlying physical hardware need to have at least as many processors as the virtual machine has virtual processors, but also that the ESX Server software must be able to allocate all of those physical or logical processors simultaneously.

This also means that configuring a virtual machine with more virtual processors than it needs not only can reduce the performance of the over-configured virtual machine (while the ESX Server software delays scheduling that virtual machine until the required number of processors are simultaneously available) but can also reduce the performance of other virtual machines running on the same system (because while the over-configured virtual machine is running, none of the processors assigned to it can be used for other virtual machines).

Each hyper-threaded physical processor appears to contain two logical processors, and therefore counts as two processors for purposes of virtual machine start up. For many workloads, however, a hyper-threaded processor doesn't provide nearly the performance of two discrete processors or two processor cores. Therefore while the ESX Server software may, for example, allow a four-processor virtual machine to be run on a system with two hyper-threaded processors, the performance would likely be extremely poor.



64-Bit Virtual Machines

ESX Server 3 includes limited experimental support for 64-bit virtual machines. If you choose to create 64-bit virtual machines, make sure that your hardware supports them. (See VMware knowledge base article 1901, "Hardware Requirements for 64-Bit Guest Operating Systems," at http://www.vmware.com/support/kb/enduser/std_adp.php?p_faqid=1901.) Also, keep in mind that these virtual machines can be started only in an ESX Server instance running on supported hardware.



Plan Your Licensing Structure

ESX Server 3 and VirtualCenter 2 introduce a new licensing mechanism based on the Macrovision FLEXnet license manager. The new licensing mechanism allows the choice of “server-based” or “host-based” licensing for each ESX Server system.

Server-based licensing is now the default configuration. In this configuration, licenses for multiple VMware products can be stored in a single license file (typically located on the system where VirtualCenter is installed). This allows a central pool of licenses to be used as needed, and allows capacity to be reallocated on the fly. During periods of license server unavailability, VirtualCenter Servers and ESX Server systems using server-based licenses are unaffected, relying on cached licensing configurations, even across reboots.

Host-based licensing allows licenses to be stored in files located on each ESX Server system. Though server-based licensing is recommended, host-based licensing might be suitable for small environments, or for locations in which a decentralized approach is preferred.

You can install, launch, and configure ESX Server version 3 and VirtualCenter version 2 without a software license. However, software licenses are required for useful operation, such as powering on a virtual machine.

When you generate your licenses you need to select whether you want them to be server-based or host-based. Once generated, server-based licenses will work only for server-based licensing configurations and host-based licenses will work only for host-based licensing configurations.

Though the license configuration for an ESX Server instance can be changed after installation, VMware recommends that a licensing model be selected before initial deployment.

More information about licensing can be found in the *VMware Infrastructure 3 Installation and Upgrade Guide*.



Plan Your User Structure

Carefully planning user and group structures before deployment can enhance the benefits of virtual infrastructure while increasing security.

VirtualCenter and ESX Servers each have their own sets of users and groups. VirtualCenter gets its user list from the Microsoft Windows domain or Active Directory, while ESX Servers each have their own internal list of "direct access" users. Most interaction with the ESX Servers is typically performed through the VirtualCenter interface.

In both VirtualCenter and ESX Server users can be members of groups. Permissions (called "roles") can be granted to individual users or to groups (in which case all members of the group inherit those permissions).

Think through how your virtual infrastructure will be managed and used. Determine who will need to grant permissions, create virtual machines, manipulate datastores, etc., then assign roles accordingly.

Each individual should have their own login so that there is an audit trail in the event of problems. Users should be organized into teams so that co-workers can cover for one another during vacations or absences.



Plan Your Network Structure

ESX Server 3 provides networking capabilities that, if planned appropriately, can reduce costs and increase security and availability.

Virtual machines communicate with each other and with the outside world through virtual switches, called vSwitches. Each virtual machine can be connected to one or more vSwitches. Each vSwitch can, in turn, be connected to one or more other virtual machines, to one or more physical network interface cards, or to both virtual machines and network interface cards.

Multiple connections from a single vSwitch to the physical network form a NIC team. Such a NIC team can increase performance by distributing the traffic across those physical network connections and provide passive failover in the event of hardware failure or network outage.

A virtual machine can be multi-homed by connecting it to more than one vSwitch. For example, a public Web server that needs to contact a back-end database might be multi-homed between a vSwitch connected through a physical NIC to the Internet and an internal-only vSwitch on which the database server virtual machine resides. This provides excellent protection for the database server because it cannot be accessed from the Internet.

It is also possible to use multi-homed virtual machines as routers and firewalls. Doing so can allow you the best of both worlds: the improved performance and security of a virtual network and the ease of management of a unified network structure.

ESX Server 3 also implements an IEEE standard networking scheme known as VLANs. VLANs allow a physical or virtual network to be segmented, providing enhanced security.

It is a good idea to isolate the service console by configuring it on its own VLAN. No virtual machine or other system facility (such as VMotion) should use this VLAN.

For security you should also provide an entirely separate Gigabit Ethernet migration network devoted to VMotion traffic. This is because when a virtual machine is migrated with VMotion the contents of the guest operating system's memory are transmitted over the network.



Plan Your Disk Storage Structure

VMware makes recommendations about the layout of your disks, especially about the size and placement of service console storage. Here is a short summary:

- Locate a VMFS partition on locally attached storage (or on LUN0 of a SAN), to be used as a swap partition for the service console. In order to obtain the best performance, using NAS for swap space is not recommended.
- Have a VMkernel core dump partition (sometimes called vmkcore) on locally attached storage (or on LUN0 of a SAN).
- Protect your service console's root file system from becoming full.
- Arrange for storage of virtual machine templates and ISO images, as discussed in the sections [Clone Virtual Machines from Templates on page 10](#) and [Use ISO Images Instead of Physical CD-ROMs on page 12](#).

A detailed discussion of disk allocation issues can be found in the *VMware Infrastructure 3 Installation and Upgrade Guide*.

In addition to Fibre Channel storage, ESX Server version 3 adds support for NAS and for both hardware and software initiated iSCSI storage.

Note: The required VMkernel core dump partition cannot be located on a software initiated iSCSI volume. Additionally, VMware Infrastructure 3 Starter edition does not include support for iSCSI storage.

Most production ESX Server customers concentrate their virtual machine disk usage on external storage, such as a Fibre Channel SAN, a hardware or software initiated iSCSI storage device, or a remote NAS file server (using the NFS protocol).

ESX Server 3 offers support for failover among Host Bus Adapters. You can also designate preferred paths to each logical unit of storage. Doing so allows for the best possible utilization of your bandwidth to the disk array.



Upgrading from Earlier Versions

Consider these factors when upgrading a system from earlier versions of ESX Server to ESX Server 3:

- Though upgrades are supported from most previous customer releases of VMware ESX Server, upgrading from unsupported releases can fail with lost data and unusable configurations.
- Enhancements made to the VMFS file system structure in ESX Server 3 (VMFS-3) make the new file systems incompatible with earlier versions of ESX using VMFS-1 or VMFS-2.
- In some cases virtual machine downtime might be unavoidable, but careful planning can minimize disruptions.

These and other factors make it advisable to plan upgrades carefully. The *VMware Infrastructure 3 Installation and Upgrade Guide* provides detailed guidance for planning and performing an upgrade.



Clone Virtual Machines from Templates

Rather than creating each new virtual machine from scratch, it is possible to create a small number of standard virtual machines and use these as a pool from which to clone each new virtual machine you need.

These standard virtual machines are called “templates” (or sometimes “golden masters”), and their use can dramatically streamline the task of deploying new virtual machines.

Plan to build a library of master system images. A key advantage of virtual machines is that they are easily cloned. Cloning a virtual machine is mostly a matter of copying its virtual disks (possibly resetting system attributes before doing so) and then building virtual machine configurations around the copies. Once created, these system images can be designated as templates, and can be prevented from being powered on.

Note: You are responsible for complying with all operating system and application vendors’ license agreements even if you clone a virtual machine.

Because the virtual hardware in a virtual machine is uniform and insulated from the physical hardware, you can clone a template into service at any time, even if you have changed the underlying physical hardware or you plan to run the new clone on a different ESX Server instance from the original.

Before adding a template to your library, install in it all the correct tools and applications it will need. These will typically include:

- Guest OS service packs and patches
- VMware Tools
- Management agents
- Anti-virus software
- Backup agents
- Applications

Try to have as few of these templates as possible, and always use them when creating new virtual machines. The result will be that there are no needless differences among your deployed systems. These differences are a typical cause of provisioning errors.

Here is an example of how templates can be used. Suppose you wish to create four virtual machines, all running Windows 2000 Server. One is to be a file server, one a mail server, one a DHCP server, and one an intranet Web server. The first two have locally stored data, and the last two do not.

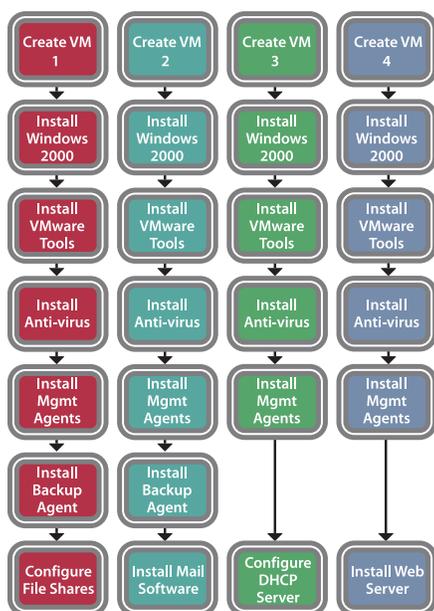
One way to create these virtual machines would be to build each from scratch. Figure 1, part A illustrates this process.

This approach is very labor intensive. It is also highly error-prone, because many steps must be repeated. The risk is high that some steps will be missed, done incorrectly, or done in the wrong order.

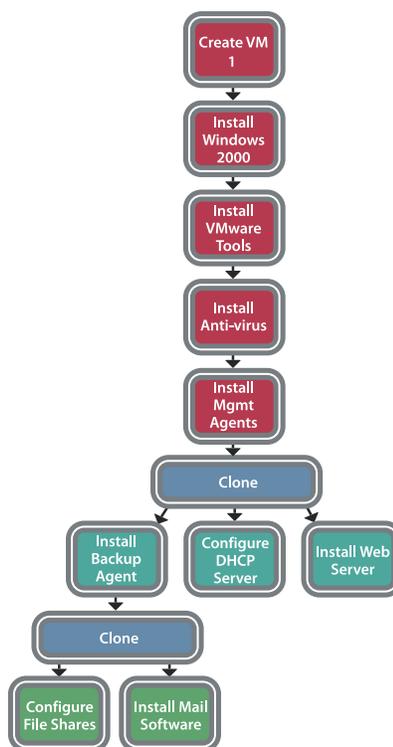
A better approach uses cloning. In the example illustrated in Figure 1, part B, each package is installed only once, and the virtual machine is cloned at appropriate junctures.



A: Building four machines from scratch



B: Building one machine and then cloning it

**Figure 1** Creating Virtual Machines From Scratch versus Cloning

To make sure that your cloning process does not result in virtual machines with duplicate system attributes, such as Microsoft Windows security identifiers (SIDs), you can use the “Deploy Virtual Machine from Template” Wizard. Alternately, for systems running Microsoft Windows you can use Microsoft’s System Preparation Tool (`sysprep`), and for systems running the Linux operating system you can configure the virtual machine as a DHCP client before cloning it.

Plan for storage of cloned images. For example in the above scenario you would want to retain the image prior to each cloning operation for future use as a template. If your ESX Server systems are connected to a Storage Area Network (SAN), or to hardware or software initiated iSCSI storage, you can use a public VMFS partition as a shared repository of virtual machine templates. Be sure to choose public mode, not shared mode.

Finally, you can make the VMkernel a client of a file server and site the repository on the file server, but be careful to store virtual disks only in exported format. (Many popular utilities do not properly handle files bigger than two gigabytes.)

Caution: Some applications stop working if the NetBIOS name of the computer they run on changes, or if the MAC addresses of the NICs change, both of which can happen during a cloning operation. Be sure to test your cloned system images.



Use ISO Images Instead of Physical CD-ROMs

Virtual machines within ESX Server can use certain files as if those files were CD-ROMs. These files are called "ISO images" because they contain an image of an ISO-9660 file system, which is the standard format of a CD-ROM. To create one of these files, use a CD ripping tool to copy a disk image from a CD-ROM to a disk file.

Use this technique for operating system and application installation CDs. This is advantageous in several ways:

- Access to ISO images is faster than to CD-ROM drives, especially with multiple simultaneous users.
- ISO images can be mounted and unmounted through the VMware Remote Console application, unlike physical CD-ROMs, which must be manually handled.
- Having a repository of ISO images in your service console encourages personnel to use them consistently.

ISO images, however, do not take the place of virtual machine templates. You should install from the original boot media, whether physical or in ISO-image form, only when no template is available.

Note: You are responsible for complying with all operating-system and application vendors' license agreements.

ESX Server no longer supports browsing the service console file system for ISO images. Instead, it allows browsing for ISO images located on both VMFS-based and NFS-based datastores.

It is now possible to mount NAS volumes through the VMkernel. These VMkernel-based NAS volumes appear as datastores in the VI Client. The virtual machine configuration editor allows you to browse these NAS datastores, as well as VMFS-based datastores, for ISO images to be used as virtual CD-ROM devices.



Keep Your Contact Information Updated

When you sign up for VMware's support and subscription programs, you register an email address with VMware. Make sure that this email address is always correct in VMware's records. VMware will use this address to notify users of the availability of patches, upgrades, and any newly discovered security issues.

A wise technique is to register an email address that points to an internal mailing list. For instance, you might create within your company a mailing list named **esx-contact** and subscribe to that list all employees who need to receive this type of bulletin from VMware. Then register **esx-contact@yourcompany.com** as your support contact. This way if employees of your company go on vacation, leave, or change roles, you need only update the membership of your esx-contact mailing list to continue to receive these notices.

Conclusion

What all VMware best practices have in common is the following rule: **virtualization pays**. The key objectives of any ESX Server deployment are to reduce management costs and to make more efficient use of resources. Look for ways to use ESX Server features towards these aims.

VMware's Professional Services Organization offers hands-on ESX Server training courses in which both the product and these best-practice recommendations are explored in detail. For more information, visit <http://www.vmware.com/services/consulting.html> or write to education@vmware.com.

VMware, Inc., 3145 Porter Drive, Palo Alto, CA 94304

Tel 650-475-5000 Fax 650-475-5001 www.vmware.com

Copyright © 1998-2006 VMware, Inc. All rights reserved. Protected by one or more of U.S. Patent Nos. 6,397,242, 6,496,847, 6,704,925, 6,711,672, 6,725,289, 6,735,601, 6,785,886, 6,789,156 and 6,795,966; patents pending. VMware, the VMware "boxes" logo and design, Virtual SMP and VMotion are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions. Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation. Linux is a registered trademark of Linus Torvalds. All other marks and names mentioned herein may be trademarks of their respective companies. Revision: 20060605 Item: ESX-ENG-Q206-236
