

WHITE PAPER

ESX Server™ 2

Mainframe-Class Virtual Machines for the Most Demanding Environments



BEST PRACTICES



Best Practices for VMware ESX Server 2

Summary

VMware ESX Server can be deployed in many ways. In this document, we recommend specific deployment guidelines. Following these guidelines will maximize the ease of use of ESX Server in your environment. This document applies only to VMware ESX Server 2 and Virtual SMP.

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 and have documented what techniques are most likely to result in a highly successful deployment. We call these *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 may not apply to your circumstances. However, we believe 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 micro-kernel that acts as a host for virtual machines.

Service Console: the user interface for the VMkernel. The Service Console supports the Web-based Management User Interface, as well as remote access to a virtual machine's keyboard, display, and mouse. Before ESX Server 2, the Service Console was called the "Console OS". The Service Console in ESX Server 2 is based on RedHat 7.2.

Virtual machine configuration file: a text file that declares the virtual hardware composing a virtual machine. These files are stored in the Service Console's file systems.

VMFS: the VMware File System. Virtual disks for virtual machines are stored as files in this high-performance, dedicated file system.

Plan your load mix before you install

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

Assess the load mix you plan to put on this ESX Server. Is your server hardware sized appropriately? If you are doing a server consolidation project, profile the CPU and I/O loads on the systems to be consolidated; otherwise, if the virtual machines are to be purpose-built, examine the utilization of systems in similar roles. Choose server hardware that is capable of sustaining the total average load (if load peaks do not all occur at the same time) or the total peak load (if the peaks do occur simultaneously). VMware's Professional Services Organization can assist you with this task as a consulting engagement. For more information, write to consulting@vmware.com.

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

The number of virtual machines running under an instance of ESX Server will determine how much memory you allocate to ESX Server's Service Console. See the VMware ESX Server Installation Guide (version 2.0) for specific values.

The amount of memory allocated to the Service Console will in turn determine how much swap space you allocate to the Service Console. This amount is required very early in the ESX Server installation process. Specifically, during the base install, you have the opportunity to size the disk partition used by the Service Console as swap space. Choose a swap-space size that is twice the amount of memory allocated to the Service Console.

Clone virtual machines rather than creating them from scratch

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. VMware's informal term for these master system images is golden masters.

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

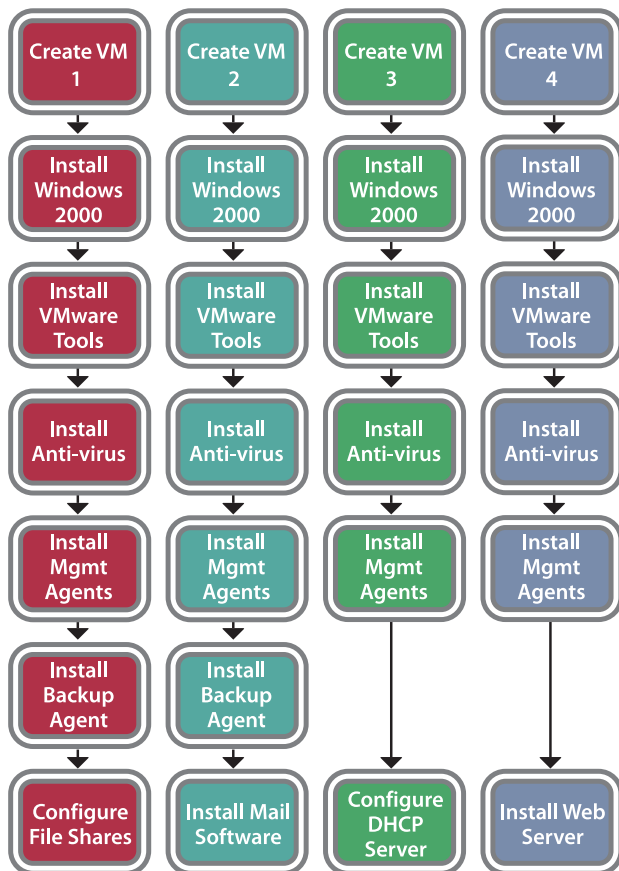
Because the virtual hardware in a virtual machine is uniform and insulated from the physical hardware, you can clone a golden master into service at any time, even if you have changed the underlying physical hardware, or even if you plan to run the new clone on a different ESX Server instance from the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

Before adding a golden master to your library, install in it all the correct tools and applications it will need:

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

Try to have as few of these golden masters as possible, and always use them. The result will be that there are no needless differences among your deployed systems. These differences are a typical cause of operator error.

Here is an example. 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.

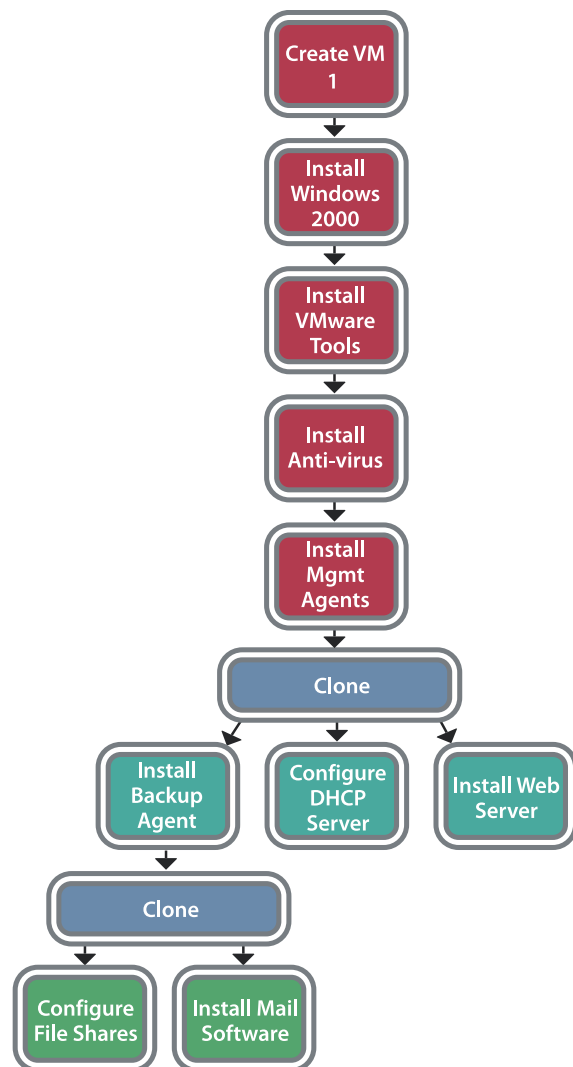


One way to create these virtual machines would be to build each from scratch. The diagram on the next page illustrates this process.

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 this example, we do each package installation only once, and we clone the virtual machine at appropriate junctures.



Make sure that your cloning process does not result in virtual machines with duplicate system attributes, such as Windows SIDs. To avoid duplicate system attributes in Windows virtual machines, run a tool such as Microsoft's `sysprep` before powering off the virtual machine and copying its virtual disks. To avoid duplicate system attributes in Linux virtual machines, make the virtual machine into a DHCP client before copying it.

Plan for storage of cloned images. For example, in the above scenario, you would want to retain the images at each point of cloning for future use as golden masters. If your ESX Server 2 systems are connected to a Storage Area Network (SAN), you can use a public VMFS as a shared repository of golden masters. Be sure to choose public mode, not shared mode. This capability is new in ESX Server 2. Note, however, that VMFS file systems do not support subdirectories, and that a VMFS may hold no more than 255 files at a time. So, if you wish to use VMFSes as a shared repository, you might decide to have several VMFSes of moderate size rather than one large one.

You may also use a local Service Console file system as a repository. Finally, you may make the Service Console 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: a few large applications stop working if the NetBIOS name of the computer they run on changes, as happens during a cloning operation. Be sure to test your cloned system images.

Use ISO files rather than physical CD-ROMs

VMware virtual machines can use a Service Console disk file as if it were a CD-ROM. These files are called "ISO files" because they contain 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 disk files is faster than that to CD-ROM drives, especially when more than one person is accessing them at the same time.
- Disk files can be mounted and unmounted through the VMware Remote Console application, unlike physical CD-ROMs, which must be manually handled.
- Having a repository of CD-ROM files in your Service Console encourages personnel to use them consistently.

Of course, ISO files do not take the place of golden masters. You should only install from the original boot media, whether physical or in ISO-file form, if no golden master is available.

Note also that you are responsible for complying with all operating-system and application vendors' license agreements.

Plan your user structure

Each virtual machine under ESX Server has an owner. The owner of a virtual machine is the Service Console user that owns its configuration file. Each virtual machine is affiliated with a group, also the group affiliation of its configuration file. The Service Console's administrative user is called **root**. This is analogous to a Windows system's **Administrator** user.

Avoid having virtual machines owned by **root**. To administer such a virtual machine, you will have to log in as **root**, during which time your errors can have system-wide consequences. A virtual machine gets no special power, privilege, or priority from being owned by **root**.

Design a user and group structure that allows your IT staff to manage the virtual machines they own. Those personnel should have individual logins so that there is an audit trail in the event of problems. Personnel should be organized into teams so that co-workers can cover for one another during vacations or absences.

There is no need to create Service Console logins for users of virtual machines, only for personnel who will be creating or administering those virtual machines.

By default, Service Console passwords are stored locally. Many VMware customers tie their Service Console logins to an external password store, such as a Windows domain. The advantage of doing so is that personnel have one fewer password to remember. The disadvantage is that there is now an extra dependency: if the external password store goes down, some administrators will be unable to log in.

Plan your network structure

Each physical network interface card ("NIC") installed on your ESX Server is allocated either to the Service Console or to the VMkernel. Virtual machines can only use NICs allocated to the VMkernel.

Each virtual Ethernet interface of each virtual machine can be patched to a different virtual switch. Some of these virtual switches are wholly private within the ESX Server itself; there is no hardware associated with them. These are called vmnets. Others are patched through to one of the ESX Server's physical NICs. These are called vmnics.

Look for opportunities to use vmnets for support of back-end processing among virtual machines. The horsepower needed to deliver packets on vmnets comes from the pool of CPU time not used directly by virtual machines. Therefore, so long as the server system's CPUs are not saturated, these vmnets frequently offer faster performance than physical LANs, because no network hardware is involved.

Another advantage of vmnets is their improved physical security. Because they have no physical ports, it is impossible for persons to patch a system to them without knowing a Service Console username and password.

To communicate with the outside world, virtual machines need to use vmnics. A new feature of ESX Server 2 is the ability to group several vmnics together into a logical team or bond. Doing this adds reliability as well as performance, because ESX Server 2 implements fail-over among the individual interfaces in the bonded interface.

Virtual machines can be multi-homed. For example, a public Web server that needs to contact a back-end database might be multi-homed between a vmnic facing the Internet and a vmmnet 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 vmmnets, and the ease of management of a unified network structure.

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.

- Have a VMFS located on locally-attached storage. Use this for swap space for the VMkernel. (Note: VMkernel swap space is completely independent of, and needed in addition to, Service Console swap space.)
- Have a VMkernel core dump partition on locally-attached storage.
- Arrange for storage of golden masters and ISO images, as discussed above.
- Protect your Service Console's root file system from filling.

See ESX Server knowledgebase article 904 for a more detailed discussion. The VMware knowledge base can be accessed at: <http://www.vmware.com/kb>

Most production ESX Server customers concentrate their virtual machine disk usage on external storage, such as a SCSI RAID array or a Fibre Channel Storage Area Network. ESX Server 2 offers improved support for fail-over 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.

Keep your support contact up-to-date

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 at an internal mailing list. For instance, you might create a mailing list inside your company called **esx-contact**; subscribe all persons to that list 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.

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 cost 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, write to education@vmware.com.



VMware, Inc. 3145 Porter Drive Palo Alto CA 94304 USA Tel 650-475-5000 Fax 650-475-5001 www.vmware.com
Copyright © 2003 VMware, Inc. All rights reserved. Protected by one or more of U.S. Patent Nos. 6,397,242 and 6,496,847; patents pending. VMware, the VMware "boxes" logo, GSX Server and ESX Server are trademarks of VMware, Inc. 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.

