



BEA WebLogic Server Virtual Edition

Configuration and User Guide

Version 9.2 v1.1
Revised: January 2008

Contents

1. Configuration Overview and Roadmap

Configuration Overview	1-1
Roadmap for Installing and Configuring WLS-VE	1-3
Comparing Startup, Configuration, and Logging Options to Non-virtualized WLS ...	1-5

2. Understanding LiquidVM

What is LiquidVM?	2-1
Understanding the LiquidVM File System	2-2
Using the Virtual Local Disk	2-3
Determining the Amount of Free Disk Space	2-4
Increasing the Size of the Local Disk	2-4
Using the LiquidVM SSH Service	2-4
SSH Listen Port	2-5
Authenticating with the SSH Service	2-5
Auditing SSH Actions	2-5
Using LiquidVM in Passive Mode	2-6

3. Configuring LiquidVM Connection Parameters

Before You Begin	3-1
Configuring LiquidVM in Graphical Mode	3-3
Configuring LiquidVM in Console Mode	3-5
Understanding the bea.lvm.info File	3-8
Troubleshooting a LiquidVM Configuration	3-9

4. Creating and Copying WLS-VE Domains

Before You Begin.	4-2
Creating WLS-VE Domains Using the Configuration Wizard	4-2
Editing the Start Scripts	4-5
Copying the Domain Using the LiquidVM SSH Service.	4-10
Step 1: Start LiquidVM in Passive Mode	4-10
Step 2: Copy the Domain to the Virtual Machine	4-13
Step 3: Shut Down the LiquidVM Instance	4-14
Step 4: Start the WLS-VE Instance.	4-14
Starting Managed Servers	4-15
Running the Administration Server on a Standard OS.	4-16
Generating the Domain to an NFS server	4-17
Copying a Production License to the Virtual Machine	4-18
Upgrading and Promoting Domains.	4-18
Upgrade from a Non-virtualized WLS 9.2 to WLS-VE 9.2	4-19
Upgrade from an Earlier Version of WLS to WLS-VE 9.2	4-20
Moving a WLS-VE Domain to a Production Environment	4-20
Additional Configuration Tasks	4-20
Tuning LiquidVM	4-20
Deploying an Application to WLS-VE	4-21

5. Working with WLS-VE Using the VMware VI Client

Setting Up VMware and Enabling SSL	5-2
Starting (and Stopping) WLS-VE	5-2
Editing VM Properties	5-2
Pausing a VM.	5-5
Working with the Console Tab	5-5
Inactive Keyboard.	5-5

Console Log.	5-6
Pre-console Log.	5-6

6. Starting and Stopping WLS-VE

Starting WLS-VE	6-1
From the Command-line	6-1
From VirtualCenter	6-2
Starting the Administration Console	6-2
Stopping WLS-VE	6-3

7. Configuring Logging

Understanding the Log Files.	7-1
WLS Logs	7-1
LiquidVM Log.	7-2
VMware Log	7-2
Accessing WLS-VE Log Files	7-2
Copying the Log Files Using SSH	7-3
Configuring LiquidVM to Use Remote syslog	7-3
Storing the Log Files on an NFS Share	7-4

8. Securing Your Production Environment

Securing LiquidVM	8-1
Securing WLS.	8-2
Securing the VMware VirtualCenter	8-3

9. Tuning the WLS-VE LiquidVM Kernel

Tuning the LiquidVM Kernel Startup Options.	9-2
Comparing OS and LiquidVM Kernel Tuning.	9-3

10. Diagnostics and Troubleshooting

Troubleshooting WLS-VE Problems	10-1
Troubleshooting Common WLS-VE Issues	10-2
“Could not find the disk” Error	10-2
The Server Shuts Down Soon After Startup	10-3
“netSend failed: -3” Error	10-4
“Configured IP [...] in use by MAC” Error	10-4
Troubleshooting WLS Issues	10-4
Performance Issues	10-5
Server Failure.	10-5
Clustering Issues	10-5
Other WLS Problems.	10-5
Troubleshooting LiquidVM Issues	10-6
Handling Suspend Files	10-7
Displaying Version Information.	10-7
Reporting a Problem to BEA Support	10-7
Trouble Reporting Process Overview	10-8
Identify Your Problem Type	10-8
Verify That You’re Running a Supported Configuration	10-8
Collect Enough Information to Define Your Issue	10-8

Configuration Overview and Roadmap


The following sections provide an overview of a sample WebLogic Server Virtual-Edition (WLS-VE) configuration, describe the main components in the configuration, and summarize the main steps required to successfully install and configure WLS-VE:

- [“Configuration Overview” on page 1-1](#)
- [“Roadmap for Installing and Configuring WLS-VE” on page 1-3](#)
- [“Comparing Startup, Configuration, and Logging Options to Non-virtualized WLS” on page 1-5](#)

Configuration Overview

To successfully configure and use WLS-VE, it is important to understand the overall architecture and components of a WLS-VE configuration. [Figure 1-1](#) provides a sample configuration consisting of two WLS-VE instances in a virtual environment.

Figure 1-1 Sample WLS-VE Configuration



The following table describes the components of the sample WLS-VE configuration.

Table 1-1 WLS-VE 1.1 Configuration Components

Component	Description
Launcher Machine	<p>The machine, either Windows or Linux, on which you execute the WLS-VE installation program and initiate the creation of virtual machines on the VMware ESX machine. The WLS-VE installation includes WLS 9.2 and LiquidVM 1.1, which provide the tools needed to create, control, and start WLS-VE and LiquidVM instances.</p> <p>The installation also includes an ISO image containing the WLS classes and LiquidVM executables used to run WLS-VE and the applications on the hypervisor host (VMware ESX server). This ISO image must be copied to the hypervisor host from the launcher machine.</p> <p>LiquidVM 1.1 also provides an SSH service which provides a secure mechanism to transfer files, including WebLogic domains, to and from the LiquidVM instance on the hypervisor host.</p>
VMware VirtualCenter	<p>The primary controller for configuring and managing the virtual environment. Users can connect to the VirtualCenter server using the Virtual Infrastructure Client (VI Client). The VI Client is only supported on Windows platforms, therefore there must be at least one Windows machine in your configuration. The BEA tools on the launcher machine interface with the ESX environment through the VirtualCenter server.</p>

Table 1-1 WLS-VE 1.1 Configuration Components

Component	Description
VMware ESX Server	The hypervisor machine with VMware installed and CPUs available for the creation of virtual machines. The LiquidVM launcher accesses the ESX server through the VirtualCenter server.
Virtual Infrastructure	The full infrastructure virtualization suite that provides the hardware and software resources required to support a virtualized environment.
WLS-VE instances	Virtual machines running on the ESX server. In WLS-VE 1.1, LiquidVM includes a virtual local disk, which can be mapped to the SAN, for each virtual machine. You can transfer files to and from the local disk using the SSH service.
SAN Disk	Recommended physical storage device for the ISO image and the local disks on the virtual machines.

Roadmap for Installing and Configuring WLS-VE

The following steps summarize the overall process for installing and configuring BEA WebLogic Server Virtual Edition™. Installation procedures are documented in the [WLS-VE Installation Guide](#). Subsequent sections of this document provide details for configuring and using WLS-VE.

Table 1-2 Roadmap for Installing, Configuring, and Using WLS-VE

Step	Description
1. Install the WLS-VE software.	<p>Before you configure WLS-VE, you need to:</p> <ul style="list-style-type: none"> Plan your virtualized installation Install the WLS-VE software Copy the ISO image to the ESX server <p>Details about installing the software are provided in the WLS-VE Installation Guide.</p>
2. Configure the LiquidVM connection parameters.	<p>Run the LiquidVM Configuration Wizard on the launcher machine to configure the connection between the LiquidVM tools running on the launcher machine and the Virtual Center and ESX server.</p> <p>For details, see Chapter 3, “Configuring LiquidVM Connection Parameters.”</p>

Table 1-2 Roadmap for Installing, Configuring, and Using WLS-VE (Continued)

Step	Description
3. Create the WLS-VE domain on the launcher machine.	<p>Use the WebLogic domain Configuration Wizard to create the WLS-VE domain on the launcher machine. For details, see “Creating WLS-VE Domains Using the Configuration Wizard” on page 4-2.</p> <p>If you are using an NFS share, you can generate the domain directly to the NFS share. For more information, see “Generating the Domain to an NFS server” on page 4-17.</p>
4. Edit the WLS-VE start scripts.	<p>Before you can copy a domain or start a WLS-VE instance, you need to edit the default start scripts that are created when you create the domain. For details, see “Editing the Start Scripts” on page 4-5.</p>
5. Copy the domain to the local disk of the WLS-VE instance.	<p>After you have created the domain and edited the start scripts, you need to copy the domain to the virtual machine in either of the following ways:</p> <ul style="list-style-type: none">• Start the LVM instance in passive mode and transfer the domain to the virtual machine. For details, see “Copying the Domain Using the LiquidVM SSH Service” on page 4-10.• Start the Administration server on the launcher machine and distribute the domain using the WLS distribution mechanism. For details, see “Running the Administration Server on a Standard OS” on page 4-16.
6. Start and Stop WLS-VE.	<p>Start and stop the WLS-VE instances as described in:</p> <ul style="list-style-type: none">• “Starting Managed Servers” on page 4-15• “Starting and Stopping WLS-VE” on page 6-1
7. Administer the WLS-VE environment.	<p>To administer the WLS-VE environment, see the following topics:</p> <ul style="list-style-type: none">• Working with WLS-VE Using the VMware VI Client• Configuring Logging• Copying a Production License to the Virtual Machine• Securing Your Production Environment• Tuning the WLS-VE LiquidVM Kernel• Diagnostics and Troubleshooting

Comparing Startup, Configuration, and Logging Options to Non-virtualized WLS

If you have experience using non-virtualized WLS, you might want to use some of the configuration techniques common to that product. If so, you need to be aware that some of techniques with which you are familiar will not work with WLS-VE. For example:

- Some customers prefer to use a script to configure their environment before running WLS. You can do this only on a standard operating system before starting the WLS-VE instance. You cannot run scripts on the WLS-VE instance itself.
- When you create a domain with the Domain Configuration Wizard, you must provide an administrator password. The Configuration Wizard encrypts the password and stores it in the domain directory in the `boot.properties` and `SerializedSystemIni.dat` files. When WLS-VE is started, it obtains the password from these files. WLS-VE does not provide an interactive prompt so you cannot specify credentials when the server is starting.
- You can track the progress of WLS-VE by viewing the log files. The following types of log files are generated:
 - Server logs
 - LiquidVM console logs

By default, both the server and LiquidVM console log files are created on the local disk of the virtual machine. VMware logs are also created and stored on the ESX server.

To view the log files, you need to copy them to your local machine using SSH, use a remote syslog collector, or put the logs on an NFS share. For more information, see [Chapter 7, “Configuring Logging.”](#)

- If the Administration Server will use a fixed IP address, you need to specify it when you create the domain using the Configuration Wizard. If you do not specify an IP address, DHCP assigns one from the available network addresses. If you do specify an IP address for the Administration Server, the `startWLSVE.sh` start script that is created will specify the IP address for the Administration Server.

Understanding LiquidVM

WLS-VE combines WebLogic Server (WLS) with LiquidVM, a Java Virtual Machine (JVM) that works with hypervisor software and provides only the set of operating system (OS) features that WLS needs to offer its full range of services. Details about LiquidVM are provided in the following topics:

- “What is LiquidVM?” on page 2-1
- “Understanding the LiquidVM File System” on page 2-2
- “Using the Virtual Local Disk” on page 2-3
- “Using the LiquidVM SSH Service” on page 2-4
- “Using LiquidVM in Passive Mode” on page 2-6

What is LiquidVM?

LiquidVM is a virtualization enabled version of the JRockit JVM that can run on a hypervisor without a standard OS, allowing Java applications to run directly on the virtualized hardware. For more information, see [WLS-VE Overview](#).

LiquidVM provides the following features:

- JRockit JVM as the Java runtime component.
- A scaled-down OS kernel that replaces the OS for LiquidVM. It differs from a normal OS in that it is a single-user, single-process environment that is designed to only run a single

JVM. No other processes can be started. It implements the following services that the JRockit JVM needs to run Java:

- Low-level memory management
 - Thread-scheduling
 - Filesystem
 - Networking
 - Interaction with the hypervisor
- Java-based services, started after the JRockit JVM has started, that run in threads that are separate from your application. The LiquidVM services are:
 - LiquidVM SSH-service
 - LiquidVM heap-resizer
 - LiquidVM syslog publisher
 - Tools that are used to create and control LiquidVM instances. The LiquidVM tools run on a standard OS; they do not run inside LiquidVM. LiquidVM tools include:
 - LiquidVM launcher—used to start and stop LiquidVM instances
 - LiquidVM Configuration Wizard—used to configure the connection between LiquidVM and the VMware VirtualCenter
 - A virtual local disk for each virtual machine. The local disk removes the dependence on NFS and provides faster and more secure file transfers

Understanding the LiquidVM File System

The LiquidVM file system is very similar to most Unix-like OSes. There is a single root (/) directory; disks and remote NFS file-shares can be mounted in sub-directories. By default, the virtual disk, if configured, is mounted in the root directory (/).

The directory structure for WLS-VE is described in [Table 2-1](#).

Table 2-1 WLS-VE Directory Structure

Directory	Contents	Notes
/appliance	The contents of the WLS-VE ISO image	<p>This directory is read-only.</p> <p>The WLS-VE ISO image file to be used is specified in the <code>bea.lvm.info</code> file using the <code>vmwareDiskPath</code> parameter. For example:</p> <pre>vmwareDiskPath=[storage1] wlsve/wlsve922.iso</pre> <p>For more information on <code>/appliance</code>, see “The Installation Directory for the WLS-VE ISO Image” in the <i>WLS-VE Installation Guide</i>.</p>
/bea	<ul style="list-style-type: none"> <code>license.bea</code> <code>patch_weblogicnnn</code> 	<p>The default location for the BEA Home directory on the virtual local disk. You can change the location of the BEA Home directory by specifying the <code>-Dbea.home</code> option in the WLS-VE start script.</p> <p>The BEA Home directory contains the <code>license.bea</code> file, and any WLS patches.</p>
/domain	<ul style="list-style-type: none"> <code>server-root</code> directory log files application classes 	The default location of the WLS domain and the current working directory for WLS-VE.
/tmp		Directory for temporary files created when WLS boots. The temporary files are erased each time the LiquidVM instance reboots.

Using the Virtual Local Disk

LiquidVM provides a virtual local disk for each virtual machine. The local disk can be mapped to a SAN disk attached to the ESX server. You specify the size of the disk by passing the `diskSize` parameter as a startup option to the LiquidVM launcher. To do so, specify the disk size in the WLS-VE start script using the `LVM_DISKSIZE` option. The default is 1GB (1024). For more information, see [“Editing the Start Scripts” on page 4-5](#). If no disk is specified, the local disk is not created.

The first time that you boot LiquidVM, it detects that a virtual hard disk is attached and that it is empty. LiquidVM formats the disk and mounts it in the root directory (/).

To the VMware ESXserver, the virtual hard disk is a vmdk file. The vmdk file is placed in the same directory as the Virtual Machine's configuration file on the ESX server/SAN. You can specify which VMware datastore the VM configuration files and the virtual hard disk should be placed in by specifying the `vmwareVmDatastore` option to the launcher. To do so, add a `vmwareVmDatastore=` entry to the `bea.lvm.info` file.

You can transfer files to and from the local disk by using the LiquidVM SSH service. For more information, see [“Using the LiquidVM SSH Service” on page 2-4](#).

Determining the Amount of Free Disk Space

To determine the amount of free disk space on your virtual machine while the WLS-VE instance is running, press **F1** in the VMware VM console. Details about the running system, including the amount of free disk space, are displayed in the console.

Increasing the Size of the Local Disk

If your disk is full, you can shut-down WLS-VE and specify a larger disk by increasing the value of the `LVM_DISKSIZE` property in the WLS-VE start script.

VMware does not provide a way to increase the size of the disk. Instead, when you restart WLS-VE, the LiquidVM launcher creates a new larger disk and copies the files from the old disk to the new disk. When you increase the size of the disk, the initial restart of WLS-VE will take longer depending on the size of the disk and the amount of files to be copied.

Note: LiquidVM does not provide a mechanism to reduce the size of a disk.

Using the LiquidVM SSH Service

LiquidVM provides a SSH2-compatible service for transporting files to and from LiquidVM. The SSH service does not provide shell services; that is, LiquidVM does not support scripts or editing files from the SSH shell.

The SSH service in LiquidVM provides an encrypted communication channel between the server and the client. The encryption protocol used is AES-128. Unencrypted communication is not supported.

You can transfer files using the `scp` and `sftp` extensions to SSH. The Linux OS includes `sftp` and `scp` clients. On Windows systems, several free SSH2 clients are available for download (e.g., PuTTY, WinSCP, FSecure, and FileZilla.)

Note: LiquidVM does not support SSH1 clients.

You authenticate with the SSH service using either password-based or public/private key authentication. For more information, see [“Authenticating with the SSH Service” on page 2-5](#).

LiquidVM is a single-process, single-user environment, therefore only the `liquidvm` user is supported. Multiple users cannot login into LiquidVM.

SSH Listen Port

The SSH2 server normally listens on the standard SSH port (port 22), but you can change the SSH listen port by setting the Java property `-Dlvm.ssh.port`. You may prefer to change the SSH port to something other than the default (port 22), since most SSH attacks try to attack the default port.

Authenticating with the SSH Service

When you use the LiquidVM launcher to create a new instance, you can specify the type of authentication to be used. To use public/private key authentication, you need to provide your public key as a startup option to the LiquidVM launcher. To do so, specify the location of the public key in the WLS-VE start script using the `LVM_SSH_PUBLIC_KEY` option. For more information, see [“Editing the Start Scripts” on page 4-5](#). When you attempt to authenticate with the SSH service, you will be prompted to provide your private key.

In a development environment, you may not want to bother with keys and secure passwords. In that case, LiquidVM provides a simpler, but unsafe method, of specifying an SSH password in clear-text from the launcher. You can specify the password in the WLS-VE start script using the `LVM_SSH_UNSAFE_PASSWORD` option.

Caution: The password is stored in clear-text. This option should not be used in a production environment. BEA recommends using public/private key authentication in development environments also. Once you have specified a real password or set up an SSH public-key, the unsafe password is no longer valid.

Auditing SSH Actions

The SSH service will send audit messages to the syslog for the following actions:

- Successful or failed log-in attempts. These audit messages include information about where the client tried to login from and the authentication method used.
- Logouts

You can use remote logging facilities to send these message to a remote log-collector for compliancy verification. For more information, see [Chapter 7, “Configuring Logging.”](#)

Using LiquidVM in Passive Mode

To copy files to or from the LiquidVM instance before starting WLS, you can use LiquidVM in passive mode. In passive mode, only the LiquidVM services, including the SSH service, are started. WLS is not started.

To start LiquidVM in passive mode, add the `startMode=passive` option to the launcher start arguments. Once LiquidVM is started in passive mode, you can log in over SSH and transfer your files. When you have finished transferring your files, you can either restart the server or login over SSH and run the start command to resume execution.

If you want the LiquidVM launcher to wait for the SSH service to be started before the launcher exits, you can specify `waitForSSH=true`. This can be useful in a scripting environment where you first start the instance and then you want to copy files from or to the server as soon as SSH is running on the newly started server.

Configuring LiquidVM Connection Parameters

After installing WLS-VE, you need to run the LiquidVM Configuration Wizard to configure the connection between LiquidVM and the VMware VirtualCenter server. You can run the wizard in either a GUI-driven, graphical mode or a command-line-based console mode. Configuring LiquidVM is a critical step because this is where you identify the file system locations of all WLS-VE components. An improperly configured LiquidVM will not run. This procedure is described in the following topics:

- [Before You Begin](#)
- [Configuring LiquidVM in Graphical Mode](#)
- [Configuring LiquidVM in Console Mode](#)
- [Understanding the bea.lvm.info File](#)
- [Troubleshooting a LiquidVM Configuration](#)

Before You Begin

Before you run the LiquidVM Configuration Wizard, be sure that you have access to the configuration data listed in [Table 3-1](#). You will be prompted to provide this information as you step through the configuration wizard.

Table 3-1 LiquidVM Configuration Data

Configuration Data	Description
Virtual Center server	The fully-qualified host name of the server hosting the VirtualCenter to which you need to connect.
VirtualCenter login credentials: <ul style="list-style-type: none"> VC user name VC password 	The user name and password required to access the VirtualCenter server. These should be provided to you by your VMware administrator. Refer to “Securing the VMware VirtualCenter” on page 8-3 for critical security information about your VirtualCenter password.
ESX datacenter name	The VMware datacenter in which your WLS-VE instance is created. The datacenter is the top-level structure in the VirtualCenter server.
ESX compute resource	The IP address or fully-qualified name of your ESX host, or cluster of hosts, in which the WLS-VE instances will run.
ESX resource pool (Optional)	The default VMware resource pool into which LiquidVM should place new VMs. You can override this parameter using the <code>VI_RESOURCE_POOL</code> parameter in your WLS-VE startup scripts. A VMware resource pool is a mechanism provided by VMware that allows you to allocate resources dynamically across a large set of servers. See the VMware documentation for more information on resource pools.
VMware network (Optional)	The VMware virtual network to use. If set to <any>, LiquidVM uses the first available VMware network. Use the drop-down menu to see a list of the available VMware networks.
Location of LiquidVM disk on ESX server	<p>The location of the WLS-VE ISO image on the ESX server. The location consists of two components: the name of the datastore on the ESX server and the pathname to the ISO image on the disk. The datastore name is always enclosed in square brackets; for example, <code>[Storage1]</code>.</p> <p>The storage location and path can be specified as:</p> <pre>[Storage1] myLocalStore/myISO.iso</pre> <p>The ISO image is installed during installation of your BEA software and is copied to the ESX server as described in “Copying the WLS-VE ISO Image” in the <i>WLS-VE Installation Guide</i>.</p>

Configuring LiquidVM in Graphical Mode

Note: If you are a Linux user and don't have access to a GUI, then use the procedure described in [“Configuring LiquidVM in Console Mode” on page 3-5](#).

To run the LiquidVM Configuration Wizard in graphical mode, use the following procedure.

1. Start the LiquidVM Configuration Wizard in graphical mode as shown in the following table.

To start the Liquid VM Configuration Wizard on this platform . . .	Perform the following steps . . .
Windows	<p>From the Start menu:</p> <p>Select Start→All Programs→BEA Products→Tools→LiquidVM Configuration Wizard.</p> <hr/> <p>From the command-line:</p> <ol style="list-style-type: none"> 1. Open a Command Prompt window and navigate to <code>BEA_HOME\tools\virtualization\control_1.1\bin\</code> 2. Enter the following command: <code>lvm_configwizard.cmd</code>
Linux	<ol style="list-style-type: none"> 1. Set the <code>DISPLAY</code> environment variable 2. Open a command shell and navigate to <code>BEA_HOME/tools/virtualization/control_1.1/bin</code> 3. Enter the following command: <code>lvm_configwizard.sh</code>

Note: In these command-lines, `BEA_HOME` represents the BEA Home directory in which you installed WLS-VE. For more information about the BEA Home directory, see [“The BEA Home Directory”](#) in the *WLS-VE Installation Guide*.

The Configuration Wizard starts and the Virtual Center server window appears.

2. Log into the VirtualCenter server by performing the following steps:
 - a. In the [Virtual Center server](#) field, enter the IP address or fully-qualified name of the server on which VirtualCenter is running.

- b. Select **Secure connection to VC** to use SSL for communication with VirtualCenter (recommended). This is selected by default.
- c. Enter the username and password for the VirtualCenter server in the **VC user name** and **VC password** fields.

Note: The first time that you run the LiquidVM Configuration Wizard, the username and password are saved in the `bea.lvm.info` file. If you want to log in as a different user, select **Change VC login credentials** and enter the username and password for the new user.

- d. Click **Next**.

A status window is displayed as you are connected to the Virtual Center. Once you are connected, the datacenter information window of the Configuration Wizard is displayed.

Note: After the datacenter information window initially displays, it may take a few moments for the wizard to obtain the available configuration from the VirtualCenter Server.

- 3. Specify the Datacenter, Host, and Resources by completing the following steps:
 - a. Select the **ESX datacenter name** that contains your WLS-VE instance. You can obtain this name from the datacenter administrator.
 - b. Select the **ESX compute resource** from the list of available hosts. The ESX compute resource is the name of the ESX host as displayed from within VirtualCenter.
 - c. Optionally, select an **ESX resource pool** in which to place your WLS-VE instance.

If you accept the default `<root>`, the VM is placed in the compute resource directly, and not in any resource pool within that compute resource. For more information about resource pools, see the VMware documentation at http://www.vmware.com/support/pubs/vi_pubs.html.

- d. Click **Next**.

The LiquidVM disk location window is displayed.

- 4. Specify the VMware network and location of the disk on the ESX server as follows:
 - a. From the drop-down menu, select the **VMware network** for LiquidVM to use. If you accept the default, `<any>`, LiquidVM uses the first available VMware network.
 - b. In the **Location of LiquidVM disk on ESX server** field, do one of the following:

- Enter the name of the datastore and the path to the LiquidVM iso image on the disk, using the format shown; that is:
`[storage name] path/filename.iso`

You must include the square brackets around the name of the datastore.

- Click **Browse** and navigate to the location of the LiquidVM ISO image on the disk. Select the file and click **Select**.

The field is populated with the datastore name and path to the LiquidVM image in the proper format.

c. Click **Finish**.

The successful configuration confirmation window appears.

5. Click **Close** to close the Configuration Wizard.

When you have successfully configured LiquidVM, the LiquidVM Configuration Wizard creates a file named `bea.lvm.info` in your home directory on your system (e.g., `C:\Document and Settings\username`), which contains all of the information you provided while running the wizard. WLS-VE reads this file when it launches to determine the location of critical files. For more information on the `bea.lvm.info` file, see [“Understanding the bea.lvm.info File” on page 3-8](#).

Configuring LiquidVM in Console Mode

Console-mode installation is an interactive, text-based method for configuring your software from the command-line. This mode is useful for Linux users who don’t have a GUI display or don’t want to otherwise use the graphical configuration mode described in [“Configuring LiquidVM in Graphical Mode” on page 3-3](#). You can also use console mode on a Windows platform.

Be sure you have access to the configuration data provided in [Table 3-1](#), since you will be prompted to supply this information as you step through the wizard.

To complete the console-mode configuration process, respond to the prompts by entering the text representing your choice (filepath, server name, etc.) or by pressing Enter to accept the default. To exit the configuration process, press Ctrl-C in response to any prompt.

To configure LiquidVM using the LiquidVM Configuration Wizard in console mode, follow these steps:

1. Start the LiquidVM Configuration Wizard as shown in the following table.

To start the Liquid VM Configuration Wizard on this platform . . .	Perform the following steps . . .
Windows	<ol style="list-style-type: none"> 1. Open a Command Prompt window and navigate to <code>BEA_HOME\tools\virtualization\control_1.1\bin\</code> 2. Enter the following command: <code>lvm_configwizard.cmd -mode=console</code>
Linux	<ol style="list-style-type: none"> 1. Open a command shell and navigate to <code>BEA_HOME/tools/virtualization/control_1.1/bin</code> 2. Enter the following command: <code>lvm_configwizard.sh -mode=console</code>

Note: In these command-lines, *BEA_HOME* represents the BEA Home directory in which you installed WLS-VE. For more information about the BEA Home directory, see [“The BEA Home Directory”](#) in the *WLS-VE Installation Guide*.

The system responds:

```
LiquidVM Configuration Wizard for VMware ESX (text-mode)
-----
Collecting information VMware Virtual Infrastructure environment...
Virtual center server
```

2. Enter either the IP address or the fully-qualified name (that is, you must include the domain name) of the [Virtual Center server](#). Press **Enter**.

The system responds:

```
Use secure connection (https) to virtual center? [Y/n]
```

3. Enter **y** (or press **Enter**) to use SSL to connect to the VirtualCenter server.

The system responds:

```
Virtual center username
```

4. Enter the appropriate VirtualCenter [VC user name](#). This should be provided to you by your VMware administrator. Press **Enter**.

The system responds:

Do you want to provide the password for your virtual center user? if you do the password will be stored in the configuration file encrypted, if you don't you will be asked for the password every time you launch a LiquidVM. [Y/n]

5. If you enter **x**, the system responds:

Virtual center password (you will not see what you type)

6. Enter the **VC password** you want to use to control access to the VirtualCenter server. This should be provided to you by your VMware administrator.

Note: If you have already set up a password for the VirtualCenter server and want to use that one, simply press **Enter**.

The system responds:

Connecting to Virtual Center. May take 30 seconds or more...

Looking up datacenters...

VMware Datacenter

[*numbered list of available datacenters*]

Please select one of the above numbers

7. Enter the number that corresponds to the name of your VMware datacenter (**ESX datacenter name**). Press **Enter**.

The system responds:

Looking up compute resources (hosts) in datacenter [*datacenter name*]...

Default VMware Compute Resource (ESX Host or Cluster)

[*numbered list of available resources*]

Please select one of the above numbers

8. Enter the number that corresponds to the name of your **ESX compute resource**. Press **Enter**.

The system responds:

Looking up resource pools in [*ESX host name*]...

VMware Resource Pool (or type any for default resource pool)

[*numbered list of available resource pools*]

Please select one of the above numbers [default: <root>]

9. Enter the name of the **ESX resource pool**, if specified.

The system responds:

Looking up VMware Networks available to [*ESX host name*]...

VMware Network (or type any to use any available)

[*numbered list of available virtual networks*]

Please select one of the above numbers [default: <any>]

10. Enter the number that corresponds to the [VMware network](#) to use.

The system responds:

```
ISO-image datastore  
[numbered list of available datastores]  
Please select one of the above numbers [default: storagel]
```

11. Enter the ISO image datastore. Press **Enter**.

The system responds:

```
Now you should provide the path on storagel where to find the wlsve.iso  
An example of a path is wlsve/wlsve921.iso  
ISO-image path:
```

12. Enter the path to the ISO image file and press **Enter**.

The system responds:

```
Checking path...
```

The system responds:

```
Datastore for new VMs  
[numbered list of available datastores]  
Please select one of the above numbers [default: storagel]
```

13. Enter the datastore name where the WLS-VE VMware configuration files should be stored.
Press **Enter**.

The system responds with this confirmation message:

```
The LiquidVM configuration has now completed successfully.  
Configuration data has been stored in the 'bea.lvm.info' file
```

When you have successfully configured LiquidVM, the LiquidVM Configuration Wizard creates a file named `bea.lvm.info` in your home directory on your system (e.g., `C:\Document and Settings\username`), which contains all of the information you provided while running the wizard. WLS-VE reads this file when it launches to determine the location of critical files. For more information on the `bea.lvm.info` file, see [“Understanding the bea.lvm.info File” on page 3-8](#).

Understanding the bea.lvm.info File

After you have run the LiquidVM Configuration Wizard and connected to the VirtualCenter server, the Configuration Wizard creates a file named `bea.lvm.info` and stores it in your home directory (for Windows users, that's `\\Documents and Settings\yourHome`; for example `C:\Documents and Settings\jtsmith`). This file contains all of the configuration

information you entered while running the wizard. [Listing 3-1](#) shows an example of a `bea.lvm.info` file.

Listing 3-1 Sample `bea.lvm.info` File

```
#BareMetal ESX-launcher configuration information
#Thu Jan 03 21:08:15 EST 2008
vmwareUsername=user1
vmwareDiskPath=[storage1] wlsve/wlsve922.iso
vmwareVcHost=vmwarevc.bea.com
vmwareKeystore=C:\\Documents and
Settings\\jtsmith\\bm_vmwarevc.bea.com.keystore
vmwarePassword=d90f1423925849c78e6dd9100d162f3f
vmwareVmDatastore=storage2
vmwareComputeResource=esx.bea.com
vmwareKeystorePassword=07300ca783a0a3e92f8fc6121e2d14aa
LiquidVM.config.version=5
vmwareDatacenter=TESTCENTER
```

Note: Refer to [Securing the VMware VirtualCenter](#) for critical security information regarding the `vmwarePassword=` property in the `bea.lvm.info` file.

The LiquidVM launcher reads the `bea.lvm.info` file at startup to obtain your LiquidVM configuration specifics. The WLS-VE launcher looks for `bea.lvm.info` in the location specified by the `LVM_INFO` environment variable (your home directory by default). This file contains information about VirtualCenter and default information about the ESX Server on which to start new WLS-VE instances. Typically, none of this information in this file is specific to the machine on which you ran the Configuration Wizard so you can copy it between different launching machines. However, since the launcher machine searches for this file in your home directory, if you move it to another location (or rename it), you need to set the `LVM_INFO` environment variable in the start script to point to the new location of the file.

Troubleshooting a LiquidVM Configuration

For information about how to troubleshoot problems that may result from the LiquidVM instance configuration, see [“Troubleshooting WLS-VE Problems” on page 10-1](#).

Creating and Copying WLS-VE Domains

To use WLS-VE, you need to create a WebLogic domain from the launcher machine, and then copy the domain and/or the servers to the virtual machine. A sample WLS-VE configuration is illustrated and described in [“Configuration Overview” on page 1-1](#).

There are several ways that you can create domains and make them accessible to the virtual machine:

- Generate the domain on the launcher machine and copy it to the WLS-VE instance using an SSH client.
- Generate the domain and start the Admin Server on the launcher machine, and use the WLS domain distribution mechanism to send the Managed Servers to the virtual machine.
- Generate the domain to a NFS server.

The following topics describe the procedures to create and copy WLS-VE domains, and also provide additional configuration information:

- [“Before You Begin” on page 4-2](#)
- [“Creating WLS-VE Domains Using the Configuration Wizard” on page 4-2](#)
- [“Editing the Start Scripts” on page 4-5](#)
- [“Copying the Domain Using the LiquidVM SSH Service” on page 4-10](#)
- [“Starting Managed Servers” on page 4-15](#)
- [“Running the Administration Server on a Standard OS” on page 4-16](#)

- [“Generating the Domain to an NFS server” on page 4-17](#)
- [“Copying a Production License to the Virtual Machine” on page 4-18](#)
- [“Upgrading and Promoting Domains” on page 4-18](#)
- [“Additional Configuration Tasks” on page 4-20](#)
- [“Deploying an Application to WLS-VE” on page 4-21](#)

Before You Begin

Before attempting to configure and run WLS-VE, ensure the following:

- ESX artifacts (clusters, disk mount, etc.) exist.
- Load balancers and Web servers have been configured.
- The WLS-VE ISO image is available from the VMware Virtual Infrastructure. BEA recommends that the ISO image reside in a datastore on the local disk for each ESX Server, or on a SAN. Normally, you copy the ISO image to the datastore during the installation process. For more information, see [“Copying the WLS-VE ISO Image”](#) in the *WLS-VE Installation Guide*.
- You have access to the IP addresses for your virtual machines, if required.

Creating WLS-VE Domains Using the Configuration Wizard

Use the steps in the following procedure to create a WLS-VE domain and servers (“virtualized servers”) using the Configuration Wizard. Complete details for creating WebLogic domains is provided in [Creating WebLogic Domains Using the Configuration Wizard](#).

Note: The following procedure describes how to create a domain using the Configuration Wizard in graphical mode. If you start the wizard in console mode, equivalent prompts are provided in text format. Use the guidelines provided in the console prompts to create the domain using console mode.

1. Start the domain Configuration Wizard on the launcher machine using one of the procedures described in [Table 4-1](#).

Table 4-1 Configuration Wizard Start-up Modes

To start the Configuration Wizard in this mode . . .	Do this . . .
Graphical	On the Windows desktop: <ul style="list-style-type: none"> • Select Start→Programs→BEA Products→Tools→Configuration Wizard
	On Linux: <ol style="list-style-type: none"> 1. Set the <code>DISPLAY</code> environment variable. 2. Navigate to the <code>/common/bin</code> subdirectory of the <code>WEBLOGIC_HOME</code> directory (for example, <code>/bea/weblogic92/common/bin</code>) 3. Enter <code>./config.sh</code>
Console	<ol style="list-style-type: none"> 1. Open a command-line shell and navigate to the <code>\common\bin</code> subdirectory of the <code>WEBLOGIC_HOME</code> directory (for example, <code>c:\bea\weblogic92\common\bin</code>). 2. From the command-line, enter: <code>./config.sh -mode=console</code> (or <code>config.cmd -mode=console</code> on Windows)

The Configuration Wizard launches and displays the **Welcome** window.

2. Select **Create a new WebLogic domain** and click **Next**.

The **Select Domain Source** window is displayed.

3. Select **Generate a domain configured automatically to support the following BEA products** and click **Next**.

Note: The **WebLogic Server** and **WebLogic Server VE Extension** options are selected by default and cannot be unchecked. These components are required to create a WLS-VE domain.

The **Configure Administrator Username and Password** window is displayed.

4. Enter a username and password to be used by the WebLogic Administrator and click **Next**.

These credentials are encrypted and stored in a boot identity file (`boot.properties`) in the domain directory. When you start the Administration Server, it refers to this file for

user credentials. WLS-VE does not provide an interactive prompt, so you cannot specify credentials when the server is starting.

The **Configure Server Start Mode and JDK** window is displayed.

5. Select the startup mode for the domain and click **Next**.

Note: You cannot select an alternate JDK. By default, the domain uses LiquidVM, which includes the JRockit JDK.

The **Customize Environment and Services Settings** window is displayed.

6. Do one of the following:
 - Select **Yes** to customize your domain configuration per the requirements of your implementation. Note the following:
 - If the Administration Server will use a fixed IP address, you need to enter it in the **Listen address** field of the **Configure the Administration Server** window. If you accept the default, **All Local Addresses**, DHCP assigns an IP address from the available network addresses.
 - If your configuration includes Managed Servers, you need to create them in the **Configure Managed Servers** window. For each Managed Server specify a name and, optionally, a listen address, listen port, SSL listen port, and whether SSL is enabled. If you accept the default, **All Local Addresses**, DHCP assigns an IP address from the available network addresses.
 - WLS-VE does not support the Node Manager, therefore there is no need to configure any Node Manager-related options.
 - Select **No** to accept the default settings defined in the domain source.

When you have completed defining your domain, click **Next**.

The **Create WebLogic Domain** window is displayed.

7. Enter a domain name and a location on the launcher machine in which to create the domain. The default location is `BEA_HOME\user_projects\domains`.

Be sure to note the directory in which you created the domain and the names you assigned to the Managed Servers.

8. Click **Create**.

The **Creating Domain** window displays status messages during the domain creation process. When the process is complete, the new domain is ready for use.

Note: If you are going to boot the Administration Server on a virtual machine, the **Start Admin Server** option on the **Creating a Domain** window should remain unchecked. However, if you are going to start the Administration Server on the launcher machine, you can select this option if you want to start the server immediately. For more information, see [“Running the Administration Server on a Standard OS” on page 4-16](#).

9. Click **Done** to exit the Configuration Wizard.

What's Next?

As shown in [“Roadmap for Installing and Configuring WLS-VE” on page 1-3](#), the next step in the process is to edit the start scripts as described in [“Editing the Start Scripts” on page 4-5](#):

Editing the Start Scripts

To start WLS-VE and LiquidVM, you need to provide start options to the LiquidVM launcher. When you create a WLS-VE domain using the Configuration Wizard, a start script containing a list of the options required to start WLS-VE on the virtual machine is created in the following directory:

- `DOMAIN_NAME\bin\startWLSVE.cmd` (Windows)
- `DOMAIN_NAME/bin/startWLSVE.sh` (Linux)

where `DOMAIN_NAME` is the name of the directory in which you located the domain, typically `BEA_HOME\user_projects\domains\DOMAIN_NAME`.

Note: The `bin/startWebLogic.cmd` and `startWebLogic.sh` files are used to start a WLS instance on a standard OS and cannot be used to start a WLS-VE instance on a virtual machine.

Before starting WLS-VE or LiquidVM, you need to create copies of the `startWLSVE` script and edit the copies to define the properties for your configuration, or to start the LiquidVM in passive mode. To do so, follow these steps:

1. In your local file system, navigate to `DOMAIN_NAME/bin/` and locate the `startWLSVE.sh` file (or `.cmd` on Windows).
2. Create a copy of the start script for each server that you intend to start, as defined in [Table 4-2](#).

Table 4-2 Creating Start Scripts

For . . .	Do this . . .
The Administration Server	Make a copy of the <code>startWLSVE</code> script and rename it <code>startWLSVE_Admin.sh</code> (or <code>.cmd</code>).
Each Managed Server	<p>Make separate copies of <code>startWLSVE.sh</code> (or <code>.cmd</code>) for each Managed Server in your configuration and name them <code>startWLSVE_Managed_nn.sh</code> (or <code>.cmd</code>), where <i>nn</i> is a sequential number identifying the individual Managed Server; for example, <code>startWLSVE_Managed__01</code>, <code>startWLSVE_Managed__02</code>, and so on.</p> <p>Note: To create a Managed Server in your WLS-VE domain, it is not enough to create a start script; you must also have created the Managed Server when you created the domain.</p>

3. Edit the start scripts `startWLSVE_Admin.sh` (or `.cmd`) and `startWLSVE_Managed_nn.sh` (or `.cmd`) by setting the properties required for your implementation. The start script properties are defined in the [Table 4-3](#).

Table 4-3 WLS-VE Start-up Options

Property	Description
LVM_INFO	The location of the <code>bea.lvm.info</code> file created by the LiquidVM Configuration Wizard on your local machine. This file contains the default settings for LiquidVM for your virtualization environment. By default, the <code>bea.lvm.info</code> file is created in your user home directory. If you change the location of the <code>bea.lvm.info</code> file, you must set this property to the new location. For more information about the <code>bea.lvm.info</code> file, see “Understanding the bea.lvm.info File” on page 3-8 .
SERVER_NAME Required for all servers.	The name of the server to be started. In VMware, <code>SERVER_NAME</code> is called <code>WLS-SERVER_NAME</code> . You must set this property to correspond to the server name that you specified when you created the domain.
LVM_IP_ADDRESS	The IP address this LVM should use. If left unset, the LVM uses DHCP to dynamically obtain an IP address

Table 4-3 WLS-VE Start-up Options

Property	Description
ADMIN_URL Required for Managed Servers.	<p>The listen address (host name or IP address) and port number of the Administration Server for the domain.</p> <p>This value must be set if you are configuring a Managed Server. For the Administration Server, ADMIN_URL must be left blank.</p>
LVM_CPUS	The number of CPUs to use (1,2,...).
LVM_MEMORY	The amount of memory allocated to the virtual machine.
LVM_DISKSIZE	<p>The size of the local disk allocated to the virtual machine. The default is defined in MB as 1024.</p> <p>For details about the LiquidVM local disk, see “Using the Virtual Local Disk” on page 2-3.</p>
LVM_NETMASK	The subnet mask for your network. You need to set this value if you are not using DHCP or you are not using default settings for netmask. The default netmask is 255.255.255.0.
LVM_GATEWAY	<p>The octet (###.###.###.###) for the gateway between your current network and the one you want to access. You need to set this value if you are not using DHCP or you are not using default settings for netmask. The standard gateway is the static IP address masked with the set netmask, with a 1 in the lowest octet; for example if the netmask is the standard 255.255.255.0 and the static IP is 172.23.80.102, then the default gateway is 172.23.80.1. If the netmask is 255.255.0.0 and the static IP address is the same (172.23.80.102), then the gateway is 172.23.0.1.</p>
LVM_DNS_SERVER	The DNS server the LVM should use.
LVM_DOMAIN_NAME	The network domain name for the LVM instance.
LVM_START_MODE Required if starting the server in passive mode.	<p>The mode in which to start the LiquidVM instance (active passive).</p> <p>Passive mode allows you to start only the LVM services, included the SSH service. WLS is not started. Once the LVM is started in passive mode, you can log into the LVM instance using SSH and transfer files from the launcher machine, such as domains that you have created, to the local disk of the virtual machine. When you have finished transferring your files, you can restart the server.</p> <p>Active mode, the default, starts both LVM and WLS.</p>

Table 4-3 WLS-VE Start-up Options

Property	Description
LVM_SSH	Determines whether the SSH service is to be used (<code>on off</code>). The default is <code>off</code> . If you set LVM_START_MODE to <code>passive</code> , SSH is used by default.
LVM_SSH_PUBLIC_KEY	The local path to the SSH public key. Specify a value for this property if you are using public/private key authentication for the SSH service (recommended).
LVM_WAIT_FOR_SSH	Specifies whether the LiquidVM launcher should wait for the SSH-server to be up and running before the launcher returns (<code>true false</code>). The default is <code>false</code> .
LVM_SSH_UNSAFE_PASSWORD	A clear-text password to be used to log into the LiquidVM instance over SSH. This password is not secure because it is stored in clear text. Do not use the unsafe password option in a production environment. BEA recommends that you use public/private key authentication.
LVM_CONSOLE_LOG	The path inside LiquidVM to the console log. The default is <code>/domain/<WLS=SERVER_NAME>.lvm.out</code> .
LVM_SYSLOG_RECEIVER	A hostname or IP-address of a remote syslog receiver.
BEA_HOME_MOUNT	<p>The NFS storage location for the BEA Home directory.</p> <p>BEA recommends that you use the local disk on the LVM instance. Generally, input and output with a local disk is faster and more secure than with a file server. Do not set this property if you are using the local disk on the LVM instance.</p>
WLS_DOMAIN_MOUNT	<p>The NFS storage location of the domain directory.</p> <p>BEA recommends that you use the local disk on the LVM instance. Generally, input and output with a local disk is faster and more secure than with a file server. Do not set this property if you are using the local disk on the LVM instance.</p>
TMP_MOUNT	<p>The NFS storage location for tmp files.</p> <p>BEA recommends that you use the local disk on the LVM instance. Generally, input and output with a local disk is faster and more secure than with a file server. Do not set this property if you are using the local disk on the LVM instance.</p>

Table 4-3 WLS-VE Start-up Options

Property	Description
PRODUCTION_MODE	The mode in which the server is started. If this variable is set to <code>true</code> , the server starts in production mode; if set to <code>false</code> , the server starts in development mode.
WLS_USER Required for Managed Servers.	<p>The username to be used by the WebLogic administrator to access WLS. This value is set by the domain Configuration Wizard in the <code>boot.properties</code> file and does not need to be set to start the Administration Server. To use the values specified in <code>boot.properties</code>, do not set this property.</p> <p>You must set this property in the Managed Server start scripts if you are running the Administration Server on a standard OS. For more information, see “Running the Administration Server on a Standard OS” on page 4-16.</p>
WLS_PW Required for Managed Servers.	<p>The password to be used by the WebLogic administrator to access WLS. This value is set by the domain Configuration Wizard in the <code>boot.properties</code> file. To use the values specified in <code>boot.properties</code>, do not set this property.</p> <p>You must set this property in the Managed Server start scripts if you are running the Administration Server on a standard OS. For more information, see “Running the Administration Server on a Standard OS” on page 4-16.</p>
PRE_CLASSPATH	Classpath to prepend to the standard classpath.
POST_CLASSPATH	Classpath to append to the standard classpath.
JAVA_OPTIONS	<p>Java options passed directly to the JVM. By default, the maximum Java heap size is set to <code>Xmx256m</code>.</p> <p>For more information about passing Java options to the JVM, see “Tuning LiquidVM” on page 4-20.</p>
JAVA_PROPERTIES	Other Java properties on the form <code>-Dproperty=val</code> that will be passed straight to the JVM. This is similar to the <code>JAVA_OPTIONS</code> , but should only be used for properties. For example, to specify an alternate BEA Home directory on the virtual machine, set <code>-Dbea.home=dir</code> using this property.

Note: The start scripts do not and should not include a path reference to the WLS classes. WLS-VE uses the WLS classes that are located in the WLS-VE ISO image.

What's Next?

As shown in [“Roadmap for Installing and Configuring WLS-VE” on page 1-3](#), the next step in the process is to copy the domain using one of the following procedures:

- Copy the domain to the local disk of the LVM instance on the ESX server using the SSH service provided by the LVM. For details, see [“Copying the Domain Using the LiquidVM SSH Service” on page 4-10](#).
- Start the Administration Server on the launcher machine and use the WLS domain distribution mechanism to distribute the servers to the virtual machine. For details, see [“Running the Administration Server on a Standard OS” on page 4-16](#).

Copying the Domain Using the LiquidVM SSH Service

After you have created the domain on the launcher machine, you can copy the domain to the virtual machine using the SSH service provided with LiquidVM.

Note: BEA recommends that you follow these steps for each of the servers that you created in your domain. However, if you choose not to copy the domain for Managed Servers, you can still start them on the VM from an empty domain as described in [“Starting Managed Servers” on page 4-15](#).

To copy a domain using the LiquidVM SSH service, follow these steps.

Step 1: Start LiquidVM in Passive Mode

When you start LiquidVM in passive mode, only the LiquidVM services, including the SSH service, are started. To start LiquidVM in passive mode, you need to set the `startMode=passive` option in the launcher start arguments and start LiquidVM as described in the following steps.

1. Modify the `startWLSVE.cmd/sh` start script for the LVM instance as described in [“Editing the Start Scripts” on page 4-5](#).

The following values are required:

- `SERVER_NAME`—Specify the name of the server that you defined when you created the domain.
- `LVM_START_MODE`—Set this parameter to `passive` to start LiquidVM in passive mode. The default is `active` mode. Note that when you set the start mode to `passive`, SSH is used by default.

Optionally, you can set the following values:

- `LVM_SSH`—Set this parameter to `on`. SSH is required in passive mode. However, if you do not set this parameter, SSH is used by default.
- `LVM_IP_ADDRESS`—If you specified an IP address for the server when you created the domain, enter it here. If you did not specify an IP address, do not define one for this parameter. If you leave it unset, LiquidVM uses DHCP to dynamically get an IP address.
- `LVM_SSH_PUBLIC_KEY`—If you are using public/private key authentication (recommended), specify the local path to the SSH public key.
- `LVM_SSH_UNSAFE_PASSWORD`—If you working in a development environment, you can provide a clear text password as an alternative to using public/private key authentication.

Note: Do not use the unsafe password option in a production environment.

Note: If you do not specify any type of SSH authentication, a temporary password, valid for 300 seconds, is displayed in the LiquidVM Console window in the VI Client when you attempt to authenticate to the SSH service. You can use this time-limited password to login, but you need to change it immediately to a permanent, secure password by executing the `passwd` command.

If you are starting an LVM instance for a Managed Server, verify that the `ADMIN_URL` is pointing to the IP address of the Administration Server (for the Administration Server, `ADMIN_URL` must be blank).

2. Start LiquidVM in passive mode as follows:

- a. Navigate to the `DOMAIN_NAME\bin\` directory, where `DOMAIN_NAME` is the name of the directory in which you located the domain, typically `BEA_HOME\user_projects\domains\DOMAIN_NAME`.
- b. Execute the start command by entering the name of the server start script at the prompt. For example, if you named your start scripts `startWLSVE_ADMIN.sh/cmd` or `startWLSVE_Managed_nn.sh/cmd` as described in [“Editing the Start Scripts” on page 4-5](#), you enter the name of that start script at the prompt:

```
startWLSVE_Admin.sh (.cmd on Windows)
```

or

```
startWLSVE_Managed_nn.sh (.cmd on Windows)
```

Output similar to the following is displayed in the Command Prompt window:

```
LiquidVM instance name = WLS-AdminServer
```

```
LVM_INFO = C:\Documents and Settings\myhome\bea.lvm.info
#BareMetal ESX-launcher configuration information
#Mon Jan 07 16:31:32 EST 2008
vmwareUsername=user1
vmwareDiskPath=[storage2] wlsve/wlsve.iso
vmwareVcHost=vmwarevc.bea.com
vmwareKeystore=C:\Documents and
Settings\myhome\bm_vmwarevc.bea.com.keystore
vmwarePassword=d90f1423925849c78e6dd9100d162f3f
vmwareComputeResource=sthx6414.jrpg.bea.com
vmwareKeystorePassword=07300ca783a0a3e92f8fc6121e2d14aa
LiquidVM.config.version=5
vmwareDatacenter=JRPG-BM
```

```
## LiquidVM configuration for instance: WLS-AdminServer
name=WLS-AdminServer
diskSize=1024
startMode=passive
ssh=on
cwd=/domain
mount= ''
```

```
JAVA_OPTIONS      : -Xmx256m -Xverify:none
WLS_CLASSPATH     :
:/bea/patch_weblogic922/profiles/default/sys_manifest_clas
spath/weblogic_patch.jar:/appliance/java/lib/tools.jar:/appliance/bea/w
eblogic92
/server/lib/weblogic.jar:/appliance/bea/weblogic92/server/lib/webservic
es.jar:
LICENSE_DIR       : /bea
PRODUCTION_MODE   :
JAVA_PROPERTIES   : -Dweblogic.management.discover=true
```

```
Launching WLS-VE instance WLS-AdminServer ...
Starting WLS-AdminServer connect...lookup...create...start...booting...
    Formatting new disk (may take some time)...
Initial log from LiquidVM instance follows:
```

```
-----
Found empty mass storage device dev1, initializing filesystem
MKFS block device dev1 (1040384k) |Baremetal hostname: "172.23.82.203"
IP address: 123.45.67.890
.....| Done
LiquidVM R1.1.4.0-92759 (BareMetal 4.0.4.0-92746-229)
-----
```


See the console log-file for further data
 LiquidVM IP-address: 123.45.67.890

Note: Be sure to take note of the LiquidVM IP address. You will need to provide it when you log into the SSH service, as described in [“Step 2: Copy the Domain to the Virtual Machine” on page 4-13](#).

By default, the LiquidVM output is recorded in `\domain\WLS-servername.lvm.out`, unless you have specified a different location using the `LVM_CONSOLE_LOG` property in the start script.

3. Optionally, verify that the WLS-VE instance started in passive mode as follows:

a. Log into the VMware VI Client.

The server name, prefixed with `WLS-`, should be listed in the Host and Clusters list in the left navigation pane. For example, if you named your Admin Server `QA_AdminServer`, the name displayed in the navigation pane is `WLS-QA_AdminServer`.

b. Select the **server**, and then select the **Console** tab. The console displays the IP address of the server, and the following message:

```
INFO: LiquidVM SSH-Server running on port 22
```

Step 2: Copy the Domain to the Virtual Machine

Use an SSH-2 compatible file transfer client of your choice, such as `scp` or `sftp`, to login to the SSH service and transfer the domain. Note the following:

- You need to specify the IP address of the LiquidVM instance noted in [“Step 1: Start LiquidVM in Passive Mode” on page 4-10](#) as the Hostname.
- If you provided an SSH public key in the start script, you will be prompted to provide the private key to login to the SSH service.
- Login as the user `liquidvm`. LiquidVM is a single-process, single-user environment; only the user `liquidvm` is supported.
- Copy the domain directory files from the `DOMAIN_NAME` directory on the launcher machine to the `/domain` directory inside the WLS-VE instance on the ESX server. (Do not create a `DOMAIN_NAME` subdirectory under the `/domain` directory on the virtual machine; the files must be copied directly under the `/domain` directory.) By default, `/domain` is defined as the current working directory within the LiquidVM instance.

Note: Before you can start a WLS-VE instance, a valid `license.bea` file must be installed in the `/bea` directory of the virtual machine. By default, an evaluation license is provided in the ISO image and is copied to the `/bea` directory the first time the server is started. However, if you want to use a production license, you will need to copy it to the virtual machine. You can copy the production license file during this session, or at a later time. For more information, see [“Copying a Production License to the Virtual Machine” on page 4-18.](#)

Step 3: Shut Down the LiquidVM Instance

You can shut down the LiquidVM instance using the VMware VI Client as follows:

1. If you have not already done so, log into the VMware VI Client.

The server name, prefixed with `WLS-`, should be listed in the Host and Clusters list in the left navigation pane. For example, if you named your Admin Server `QA_AdminServer`, the name displayed in the navigation pane is `WLS-QA_AdminServer`.

2. Select the **server**, and then select the **Console** tab.
3. Click inside the Console window of the VI Client and press Ctrl-C.

Note: When you click inside the Console window, most keys on your keyboard are disabled and your mouse pointer disappears. However, certain keyboard functions, including Ctrl-C still function.

For more information about using the VI Client, see [Chapter 5, “Working with WLS-VE Using the VMware VI Client.”](#)

Step 4: Start the WLS-VE Instance

1. In your local file system on the launcher machine, navigate to `DOMAIN_NAME/bin/` and open the start script that you edited in [“Step 1: Start LiquidVM in Passive Mode” on page 4-10.](#)
2. Edit the following property in the start script and save the file:

- `LVM_START_MODE`—Set this parameter to `active` to start the WLS-VE instance in active mode. In active mode, both the LVM services and the main WLS classes are started.

Note: To continue to use SSH in active mode, set the `LVM_SSH` parameter to `on` if you did not specifically set this parameter in passive mode.

3. Execute the start command by entering the name of the server start script at the prompt; for example:

```
startWLSVE_Admin.sh (.cmd on Windows)
```

or

```
startWLSVE_Managed_nn.sh (.cmd on Windows)
```

4. Optionally, verify that the WLS-VE instance started as follows:

- a. Log into the VMware VI Client.

The server name, prefixed with WLS-, should be listed in the left navigation pane. For example, if you named your Admin Server QA_AdminServer, the name displayed in the navigation pane is WLS-QA_AdminServer.

- b. Select the **server**, and then click the **Summary** tab.

In the General pane, the State field indicates **Powered On**.

What's Next?

As shown in [“Roadmap for Installing and Configuring WLS-VE” on page 1-3](#), after you have successfully started your WLS-VE instance from the command-line, you can start and stop the servers, and administer the WLS-VE environment as required.

Starting Managed Servers

If you have not copied the domain to the VM for each Managed Server as described in [“Copying the Domain Using the LiquidVM SSH Service” on page 4-10](#), there is no domain configuration information on the WLS-VE instance for the Managed Server. In this case, you can start the Managed Server on the VM as described in this section and the Managed Server will obtain its domain configuration from the Administration Server. However, using this method is less secure because you need to specify the username and password for the Administration Server in clear text in your start script.

To start a Managed Server from an empty domain, follow these steps.

1. As described in [“Editing the Start Scripts” on page 4-5](#), create a Managed Server start script for each Managed Server in the domain using the `startWLSVE.cmd/sh` script as a template.

Be sure to set the following values in the start script:

- `SERVER_NAME`—Specify the name of the Managed Server that you defined when you created the domain.

- `ADMIN_URL`—Specify the listen address (host name or IP address) and port number of the Administration Server for the domain.
- `WLS_USER`—Specify the Administrator username. This value is required to log into the Administration Server.
- `WLS_PW`—Specify the Administrator password. This value is required to log into the Administration Server.

Optionally, you can set the `LVM_SSH` parameter to `on` to use the SSH service.

2. Start each Managed Server by entering the name of the server start script at the prompt. For example, if you named your start script `startWLSVE_Managed_01.sh/cmd` as described in [“Editing the Start Scripts” on page 4-5](#), you enter the name of that start script at the prompt:

When the server starts, it obtains its domain configuration information from the Administration Server.

Running the Administration Server on a Standard OS

An alternate approach to copying the domain to the virtual machine is create a domain that consists of an Administration Server and Managed Servers. You can then run the Administration Server on an ordinary OS and let the WLS built-in domain distribution mechanism send out the domain configuration to the WLS-VE instances for the Managed Servers. This method should be used only in development environments because it is less secure than using SSH.

To use this approach, do the following:

1. Create the domain as described in [“Creating WLS-VE Domains Using the Configuration Wizard” on page 4-2](#). Be sure to create at least one Managed Server.
2. Start the WLS Administration Server for the domain on the launcher machine in one of the following ways:
 - While you are creating the domain, select **Start Admin Server** in the **Creating a Domain** window of the Configuration Wizard.
 - After the domain is created, navigate to the `DOMAIN_NAME\bin\` directory, where `DOMAIN_NAME` is the name of the directory in which you located the domain, typically `BEA_HOME\user_projects\domains\DOMAIN_NAME`. Enter `startWebLogic.cmd/sh` at the prompt.

Note: In this case, you are using the `startWebLogic.cmd/sh` start script because you are starting the Administration Server on a standard OS on the launcher machine. The `startWLSVE.cmd/sh` scripts are used to start servers on the virtual machine.

3. Using the `startWLSVE.cmd/sh` script as a template, create a Managed Server start script for each Managed Server in the domain as described in [“Editing the Start Scripts” on page 4-5](#).

Be sure to set the following values in the start script:

- `SERVER_NAME`—Specify the name of the Managed Server that you defined when you created the domain.
- `ADMIN_URL`—Specify the listen address (host name or IP address) and port number of the Administration Server for the domain.
- `WLS_USER`—Specify the Administrator username. This value is required to log into the Administration Server.
- `WLS_PW`—Specify the Administrator password. This value is required to log into the Administration Server.

Optionally, you can set the `LVM_SSH` parameter to `on` to use the SSH service.

4. Start each Managed Server by entering the name of the server start script at the prompt. For example, if you named your start script `startWLSVE_Managed_01.sh/cmd` as described in [“Editing the Start Scripts” on page 4-5](#), you enter the name of that start script at the prompt:

When the WLS-VE instance boots up, it will automatically obtain its domain by contacting the Administration Server running on the standard OS.

The potential drawback of this approach is that you have to have a mixed environment with one non-virtualized WLS acting as the Administration Server.

What’s Next?

As shown in [“Roadmap for Installing and Configuring WLS-VE” on page 1-3](#), after you have successfully started your WLS-VE instance from the command-line, you can start and stop the servers, and administer the WLS-VE environment as required.

Generating the Domain to an NFS server

If both your ordinary OS and your WLS-VE instance have access to the same NFS-share, you can run the domain Configuration Wizard on the launcher machine, which is running a standard OS, and generate the domain to the NFS directory. When WLS-VE is started it will see and use the newly generated domain.

For details about generating a domain to an NFS share, see [“Configuring and Starting WLS-VE Domains”](#) in version 1.0 [Installation and Configuration Guide](#).

Note: Creating a domain on an NFS share is less secure than using a local disk. BEA recommends using a local disk on the virtual machine instead.

Copying a Production License to the Virtual Machine

Before you can start a WLS-VE instance, a valid `license.bea` file must be installed in the `/bea` directory of the virtual machine. By default, an evaluation license is provided in the ISO image and is copied to the `/bea` directory the first time the server is started. To update the evaluation license with a production license, you need to copy the production license from the BEA Home directory on the launcher machine to the BEA Home directory, `/bea`, on the local disk of the WLS-VE instance.

Use the following procedure to copy a production `license.bea` file to the virtual machine:

1. Edit the start script for the WLS-VE instance to start the LVM in passive mode as described in [“Step 1: Start LiquidVM in Passive Mode” on page 4-10](#).

Note that `/bea` is the default BEA Home directory on the virtual machine. If you want to use a directory other than `/bea` as the BEA Home directory, set it in the start script using the `JAVA_PROPERTIES="-Dbea.home= dir"` property.

2. Use an SSH-2 compatible file transfer client of your choice to log into the SSH service. Note the following:
 - If you provided an SSH public key in the start script, you will be prompted to provide the private key to login to the SSH service.
 - Login as the user `liquidvm`. LiquidVM is a single-process, single-user environment; only the user `liquidvm` is supported.
3. Copy the `license.bea` file from the BEA Home directory on the local machine to the BEA Home directory, `/bea` by default, on the virtual machine. Note that the default working directory is `/domain`, so you will need to navigate to the `/bea` directory.
4. Shut down the LVM instance as described in [“Step 3: Shut Down the LiquidVM Instance” on page 4-14](#).

Upgrading and Promoting Domains

Generally, upgrading and promoting virtualized domains requires the same steps used for upgrading and promoting non-virtualized domains. The main steps in this process are:

1. Plan the upgrade.

In this step, you need to inventory the application environment, verify supported configuration information, review the compatibility information, and create an upgrade plan.

2. Prepare to upgrade.

In this step, you undeploy any deployed applications, shut down all servers in the application environment, back up the application environment, install any required BEA products, prepare the remote Managed Server domain directories, and set up the environment.

3. Upgrade your application environment.

4. Complete post-upgrade procedures.

For detailed instructions on these steps, see [Upgrading WebLogic Application Environments](#).

Due to its virtualized nature, when you upgrade to WLS-VE, depending on your required upgrade scenario, you will need to modify the standard upgrade procedure to address important virtualization issues. The changes you need to make are described in the following sections.

Upgrade from a Non-virtualized WLS 9.2 to WLS-VE 9.2

When you upgrade an application from a non-virtualized implementation of WLS 9.2 to WLS-VE 9.2, you need to make some modifications to the application code to ensure successful operation.

- Use the Domain Configuration Wizard to create an empty WLS-VE domain, customize it for your environment, and transfer the domain to the virtual machine.
 - Change any Java commands used to boot a server from `java` to `java_esx`.
 - If the application has hardcoded `localhost` as the host name, you must change the host name to the LiquidVM IP address.
- Note:** If your application uses Pointbase, you will need to change `localhost` to the IP address of Pointbase.
- If your application uses default values for the host in the WLST `connect()` command, you will need to change it so that the hostname can be passed in.

Upgrade from an Earlier Version of WLS to WLS-VE 9.2

You cannot migrate an application directly from an earlier version of WLS (for example, version 8.1) directly to WLS-VE 9.2. Instead, you need to follow the upgrade procedures outlined in [Upgrading WebLogic Application Environments](#).

Moving a WLS-VE Domain to a Production Environment

You use the same procedures to move a WLS-VE domain from a development environment to a production environment that you use for standard WLS domains. Ensuring your environment is secure is critical in a production environment. For important security recommendations, see [Chapter 8, “Securing Your Production Environment.”](#)

Additional Configuration Tasks

Since WLS-VE contains both a JVM and a virtualized WLS instance, you can configure both devices by using the same configuration flags used by their non-virtualized editions. Usually, you can do this from the WLS Administration Console. Refer to [System Administration for BEA WebLogic Server 9.2](#) for complete information on how to:

- Configure a WLS environment
- Configure server security
- Configure system resources
- Configure and deploy applications
- Configure WLS environments for high availability

Tuning LiquidVM

The JVM should already be well-tuned for most WLS applications but you can configure and tune the Java behavior of a machine by setting the necessary Java options in the start-up script for the domain in question. Simply enter the standard J2SE start-up options or BEA JRockit's non-standard `-X` and `-XX` options at the `JAVA_OPTIONS=` statement.

[Listing 4-1](#) shows a snippet of the Administration Server start-up script, `startWLSVE.cmd`, with `JAVA_OPTIONS=` highlighted.

Listing 4-1 startWLSVE.cmd Code Snippet

```
@ECHO OFF
SETLOCAL

.
.
.
set PRE_CLASSPATH=
set POST_CLASSPATH=
set JAVA_OPTIONS=
set JAVA_PROPERTIES=
```

For example, suppose you want to start the machine so that LiquidVM uses a garbage collector (that is, a memory management system) optimized for application throughput. You would do this by setting `JAVA_OPTIONS` as follows:

```
JAVA_OPTIONS="-xgcprio:throughput"
```

You can string together as many valid options as you need; however, you must place them within quotation marks and separate them with a single space. For example, the following code:

```
JAVA_OPTIONS="-xgcprio:throughput -xgcreport -Xss:512k"
```

tells the JVM to:

- Start WLS-VE with its JVM using a garbage collector optimized for application throughput (`"-xgcprio:throughput"`).
- Generate an end-of-run report that shows garbage collection statistics (`"-xgcreport"`).
- Set the thread stack size (memory areas allocated for each Java thread for their internal use) to 512 KB (`"-Xss:512k"`).

See the BEA JRockit [Command Line Reference](#) for a list of valid LiquidVM start-up options and instructions for using them. For LiquidVM tuning and configuration guidelines, see [“Profiling and Performance Tuning”](#) in the BEA JRockit *Diagnostics Guide*.

Deploying an Application to WLS-VE

Deploy applications on WLS-VE the same way you deploy them on non-virtualized WLS. Application deployment generally involves the following tasks:

- Preparing applications and modules for deployment
- Configuring applications for production deployment
- Exporting an application for deployment to new environments
- Deploying applications and modules with `weblogic.Deployer`
- Redeploying applications in a production environment
- Managing deployed applications

These tasks are detailed in [*Deploying Applications on BEA WebLogic Server 9.2*](#).

Working with WLS-VE Using the VMware VI Client

The VMware VI Client provides a graphical view to WLS-VE through the VMware VirtualCenter, a component of VMware Infrastructure. It allows you to provision virtual machines and monitor performance of physical servers and virtual machines. VirtualCenter can optimize resources, ensure high availability to all applications running on virtual machines and improve the responsiveness of your IT environment with virtualization-based distributed services.

This section provides an overview of the VI Client to help you familiarize yourself with its components and some of its uses. This section is not intended as a VI Client user guide beyond functionality that is specific to WLS-VE. For complete information on this product, we strongly recommend that you refer to the [VMware Infrastructure Documentation](http://www.vmware.com/support/pubs/vi_pubs.html) at http://www.vmware.com/support/pubs/vi_pubs.html

This section contains information on these topics:

- “Setting Up VMware and Enabling SSL” on page 5-2
- “Starting (and Stopping) WLS-VE” on page 5-2
- “Editing VM Properties” on page 5-2
- “Pausing a VM” on page 5-5
- “Working with the Console Tab” on page 5-5


Setting Up VMware and Enabling SSL

For information about installing VMware VirtualCenter, see the VMware installation documentation at http://www.vmware.com/support/pubs/vi_pubs.html.

To use SSL with VirtualCenter (which is strongly recommended), you must install the VMware Web Service. The VMware Web Service is installed by default when you choose the typical installation for VirtualCenter. For information about how to set up the VMware Web Service and verify that it is operating correctly, see the VMware *Installation and Upgrade Guide*. See also the *Developing Client Applications* chapter of the *VMware Infrastructure SDK Programming Guide*.

Starting (and Stopping) WLS-VE

You can start an existing WLS-VE instance from the VI Client as follows:

1. Select the machine in the Host and Clusters list in the left navigation pane.
2. Do one of the following:
 - Click the **Power On** button .
 - Select the Summary tab and select **Power On** in the Commands pane.
 - Right-click a machine on the Hosts and Clusters list and select **Power On** from the context menu that appears.

If you stop WLS-VE, you risk losing any underlying server connections, with no guarantee that these connections will be restored when the machine is restarted. For this reason, **stopping WLS-VE using VirtualCenter is not recommended**. If you must stop WLS-VE, follow the instructions for doing so in “[Stopping WLS-VE](#)” on page 6-3. For more detailed information on starting (and stopping) a machine from VirtualCenter, refer to “[Starting WLS-VE](#)” on page 6-1

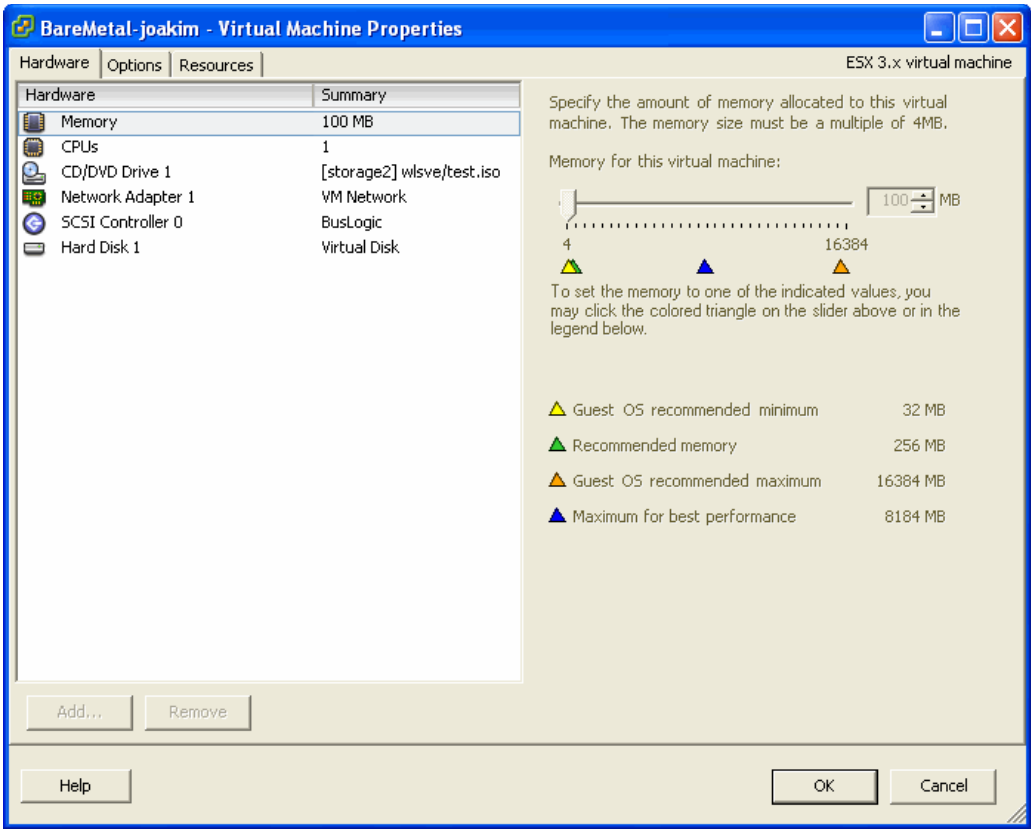
Editing VM Properties

You can edit the properties of a running virtual machine while it is running from within the VI Client. To do so, either:

- Select the Summary tab and select **Edit Settings** in the Commands pane.
or
- Right-click a machine on the Hosts and Clusters list and select **Edit Settings** from the context menu that appears.

The Virtual Machine Properties window appears (Figure 5-1).

Figure 5-1 Virtual Machine Properties Window



This window is divided into three tabs, each providing access to certain groups of properties that you can edit, as described in Table 5-1. Unless otherwise specified, you can edit any of the values on these tabs as long as those properties are enabled for your VM.

Table 5-1 Virtual Machine Properties Tabs

Tab	Properties
Hardware	<p>These properties can be changed:</p> <ul style="list-style-type: none"> • Memory: Amount of memory allocated to the VM • CPUs: the number of virtual processors used by this machine <p>These properties cannot be changed:</p> <ul style="list-style-type: none"> • CD/DVD Drive: The path and filename of the machine's <code>.iso</code> file (there can be more than one drive) • Network Adapter: the device status, adapter label, and device type used by this machine (there can be more than one network adapter) • SCSI Controller: information about the SCSI controller, if any is used • Hard Disk: information about the virtual hard disk used by this machine, including disk size, disk file name, device node, and operational mode (there can be more than one hard disk)
Options	<ul style="list-style-type: none"> • General: information about the VM in use, including its name, configuration file, working location, and guest OS • VMware Tools: identifies which VMware power controls are available to this machine and allows the user to select any scripts to run with the machine • Power Management: allows the user to define how the VM should behave when the guest OS is placed on standby • Advanced: allows the user to set advanced properties such as whether or not to expose the Nx flag, enable logging, and set advanced configuration parameters
Resources	<ul style="list-style-type: none"> • CPU: allows the user to allocate resources to the CPU • Memory: allows the user to allocate memory resources • Disk: allows the user to reallocate hard disk resources • Advanced CPU: allows the user to permit or deny sharing of physical CPU cores when the host supports hyperthreading and to select physical processor affinity for the VM

For details on how to use the Virtual Machine Properties window to edit VM settings, refer to the VMware document, *Basic System Administration*, in the VMware Infrastructure documentation at http://www.vmware.com/support/pubs/vi_pubs.html.

Pausing a VM

Pausing, or hibernating, a VM causes it to stop running while allowing all other processes to continue. You can pause a VM either by clicking the Suspend button or by right-clicking a running machine in the Hosts & Clusters list and selecting Suspend from the context menu that appears.

You should be very careful when pausing a VM; only use this function when absolutely necessary. When you pause a VM instance while you have open connections to the server, you risk losing these connections when you resume the VM. The resultant unpredictable behavior might cause these connections to be reset.

Working with the Console Tab

While most of the VI Client tabs are fairly straightforward in their use, the Console tab does support certain functions specific to WLS-VE. This section describes those functions.

Inactive Keyboard

When you select the Console tab and click inside the Console window, most keys on your keyboard are disabled and your mouse pointer disappears. You can use the key sequences listed in [Table 5-2](#) to perform certain functions when the Console is open:

Table 5-2 VI Client Console Tab Key Sequences

Press . . .	To . . .
Ctrl-Break	Force a stack trace.
Ctrl-C	Shut down LiquidVM.
Ctrl-Alt	Reactivate your mouse pointer.
F1	Display environmental information.
F10	Collect application profile data. First, ensure that you have applied a load to the application. To use this key: <ul style="list-style-type: none"> • Press F10 once to start a data collection. • Press F10 a second time to display the information you collected.

Console Log

Text written to the console also gets written to a log file. This log file is normally stored in the current working directory (`/domain` by default), as seen from inside the WLS-VE instance. By default, the LiquidVM output is recorded in the `\domain\WLS-servername.lvm.out` file, unless you have specified a different location using the `LVM_CONSOLE_LOG` property in the start script. For details about setting properties in the start script, see [“Editing the Start Scripts” on page 4-5](#).

Pre-console Log

Before the LiquidVM has a network connection, it cannot write to the log file because it needs a network connection to do so. If something fails during initialization, you may want to see what happened without having to use the Virtual Infrastructure Client. WLS-VE supports such monitoring with the help of the WLS-VE launcher. After the WLS-VE instance is launched, the launcher continues to monitor the boot process until the JVM has started successfully. If anything fails during this early stage, the error messages are displayed in your launcher. If the failure happened early in the boot process, this message may not be shown in the LiquidVM console log because the log file may not have been created at the time of the failure. This request uses the VMware communication channels that don't require networking to be set up, therefore the WLS-VE launcher can get machine log information even if networking initialization has failed.

Starting and Stopping WLS-VE

The following topics describe how to start and stop virtualized servers in a WLS domain.

- [“Starting WLS-VE” on page 6-1](#)
- [“Stopping WLS-VE” on page 6-3](#)

Starting WLS-VE

When you start a WLS-VE instance, LiquidVM boots (in `active` mode) and simultaneously boots the associated server. You can also start a LiquidVM instance in `passive` mode. In `passive` mode, only the LiquidVM services are started; the server is not started. For instructions about starting LiquidVM in passive mode, see [“Step 1: Start LiquidVM in Passive Mode” on page 4-10](#).

You can start a WLS-VE instance either from the command-line or from the VMware VI Client running on the ESX host.

Note: You *cannot* start a WLS-VE server instance from the WLS Administration Console, nor can you use the Administration Console to resume or suspend a WLS-VE server instance.

However, you can, and should, use the Administration Console to shut down a WLS-VE server instance.

From the Command-line

To start WLS-VE from the command-line:


1. Open a command-line shell and navigate to the `DOMAIN_NAME/bin/` directory.
2. Execute the start command by entering the name of the server start script at the prompt. For details about setting the properties in the server start scripts, see [“Editing the Start Scripts” on page 4-5](#). Sample start script commands are shown in [Table 6-1](#).

Table 6-1 WLS-VE Start Commands

To ...	Enter ...
Start a WLS-VE Administration Server	<code>startWLSVE_Admin.sh</code> (or <code>.cmd</code> on Windows)
Start a WLS-VE Managed Server	<code>startWLSVE_Managed_nn.sh</code> (or <code>.cmd</code> on Windows) where <i>nn</i> is a sequential number identifying the individual Managed Server

From VirtualCenter

After you have started WLS-VE from the command-line at least once, you can then start it from the VMware VirtualCenter:

1. Log into the Virtual Infrastructure Client.
2. In the Inventory panel of VirtualCenter, scroll down to select the virtual machine you want to start.
3. Do one of the following:
 - Click the **Power On** button .
 - Select the Summary tab and select **Power On** in the Commands pane.
 - Right-click a machine on the Hosts and Clusters list and select **Power On** from the context menu that appears.

As the virtual machine starts, a progress meter appears in the Recent Tasks pane.

Once WLS-VE has successfully started, the task status changes to **Completed**.

Starting the Administration Console

After your WLS-VE instance is started, you can start the Administration Server on the launcher machine to administer and shut-down WLS-VE instances.

To start the Administration Console, open a supported Web browser and open the following URL:

```
http://hostname:port/console
```

where *hostname* is the DNS name or IP address of the Administration Server and *port* is the address of the port on which the Administration Server is listening for requests (7001 by default).

If you start the Administration Server using Secure Sockets Layer (SSL), you must add an “s” after http as follows:


```
https://hostname:port/console
```

For detailed information about using the Administration Console, see the WLS [Administration Console Online Help](#).

Stopping WLS-VE

Stopping WLS-VE should be carefully considered as the ripple effects of an improper shutdown can cause unexpected results. Ideally, you should only shut down WLS-VE by using one of following methods:

- Use the WLS Administration Console for the machine and request the server to gracefully shutdown.
 - For information on shutting down a server, see [Shut Down a Server Instance](#) and [Control Graceful Shutdowns](#) in the *Administration Console Online Help*.
 - For information on gracefully shutting down the Managed Servers in a cluster, see [Shut Down Servers in a Cluster](#) in the *Administration Console Online Help*.
- Use WLST scripts to request the server to gracefully shutdown. For more information, see [Shutdown](#) in the [WLST Command and Variable Reference](#).

Note: Pressing Ctrl-C or clicking  in the VI Client results in a forced WLS shutdown, similar to pressing Ctrl-C in a standard OS.

Configuring Logging

LiquidVM-based WLS-VE instances have their own local disks where logs are stored. Many organizations use third-party log management products to collect logs from all running machines. This is typically required for SOX-compliance (Sarbanes-Oxley).

The following topics describe the logs that are created when using WLS-VE and how to access them:

- [“Understanding the Log Files” on page 7-1](#)
- [“Accessing WLS-VE Log Files” on page 7-2](#)

Understanding the Log Files

There are three basic types of log files created when using WLS-VE:

- WLS (Java) logs
- LiquidVM log
- VMware log

WLS Logs

WLS-VE creates the same WLS log files, such as the server log and the domain log, as non-virtualized WLS. In WLS-VE, the log files are stored on the virtual local disk of the WLS-VE instance under the `/domain` directory. For example:

- Server logs are stored in `/domain/servers/SERVER_NAME/logs/SERVER_NAME.log`
- Domain logs are stored in `/domain/servers/ADMIN_SERVER_NAME/logs/DOMAIN_NAME.log`

In this pathname, *ADMIN_SERVER_NAME* is the name of the Administration Server for the domain and *DOMAIN_NAME* is the name of the domain that you provided to the Configuration Wizard.

For more information about the WLS logs, see [“Understanding WebLogic Logging Services”](#) in *Configuring Log Files and Filtering Log Messages*.

LiquidVM Log

By default, the LiquidVM output is written to the console and, by default, is recorded in the following log file:

```
/domain/vmname.lvm.out
```

where `/domain` is the current working directory and *vmname* is the server name you assigned to the server when you created the domain, prefixed with `WLS-`, for example

```
/domain/WLS-AdminServer.lvm.out.
```

VMware Log

The VMware log file, `vmware.log`, is not available directly from LiquidVM, but it is stored on the ESX server in the same directory as the VMware configuration files and local disk. VMware logs record important and critical events from VMware, and warnings and errors reported from LiquidVM.

Accessing WLS-VE Log Files

There are three basic ways to collect WLS and LiquidVM log file information from LiquidVM-based instances:

- Copy the logs using SSH
- Configure LiquidVM to use remote syslog (not applicable to Java logs)
- Put the log-files on an NFS share

The following sections describe these options in more detail.

Copying the Log Files Using SSH

The most straight forward approach to collecting logs is to use an SSH-based file transfer client, log into the WLS-VE instance while it is running, and transfer the log files of interest to another machine for inspection. For more information about SSH, see [“Using the LiquidVM SSH Service” on page 2-4](#). To copy the log files using SSH, follow these steps:

1. Make sure the LiquidVM SSH service is enabled and running. The easiest way to do this is to specify `ssh=on` to the WLS-VE launcher by setting the `LVM_SSH=on` property in the `WLS_VE` start script. For more information, see [“Editing the Start Scripts” on page 4-5](#).
2. Use an SSH-2 compatible file transfer client of your choice to log into the SSH service on the WLS-VE instance. Note the following:
 - If you provided an SSH public key in the start script (recommended), you will be prompted to provide the private key to log into the SSH service.
 - Login as the user `liquidvm`. LiquidVM is a single-process, single-user environment; only the user `liquidvm` is supported.
3. Transfer the log files from the WLS-VE instance to a directory on your local machine. You can then view them using any text editor. The logs on the WLS-VE instance are located by default in `/domain`, as described in [“Understanding the Log Files” on page 7-1](#).

Configuring LiquidVM to Use Remote syslog

LiquidVM provides a syslog compliant interface for syslog events. Most third-party log management tools can collect log information from syslog compliant devices. LiquidVM implements the syslog standard (RFC3164) as a service that you can configure to publish syslog information to a remote syslog collector. Note that the syslog is subset of all logs. In particular, it does not contain the WebLogic logs. However, it does contain the events that a generic log management product typically collects.

To enable remote syslog in LiquidVM:

1. Ensure that the host that you want to receive the logs is running a syslog collector. On a Linux machine, you typically enable remote syslog collection by adding `-r` to syslog when starting the syslog daemon. If you are using a log management product, such as rSA Envision, review the product manual for configuration requirements.
2. Specify the hostname or IP address of the receiving host using `logReceiver=hostname` as an argument to the WLS-VE launcher. You can do this by setting `LVM_SYSLOG_RECEIVER`

property in the WLS-VE start script. Setting this property guarantees that WLS-VE will send syslog messages to the specified host.

Storing the Log Files on an NFS Share

You can also configure your environment to store all log files on an NFS share instead of on the local disk. By doing so, the log files are accessible from any other non-virtualized OS machine that can access the NFS share also. BEA does not recommend storing the log files on an NFS share for two primary reasons:

- NFS is less secure and more vulnerable than using the local disk and should be avoided for sensitive data. It is very likely that the application log files may contain sensitive information.
- Performance is worse when files are stored on an NFS share than when they are stored on local disk.

For information about how to configure and use an NFS share, see “Preparing for the Installation” in WLS-VE version 1.0 *[Installation and Configuration Guide](#)*.

Securing Your Production Environment

Before you attempt to use WLS-VE, you need to establish a level of security to protect the integrity of your data and the safety of your transactions. This section describes the most critical security measures you should take before working with WLS-VE. These are:

- “Securing LiquidVM” on page 8-1
- “Securing WLS” on page 8-2
- “Securing the VMware VirtualCenter” on page 8-3

Securing LiquidVM

WARNING: The following information is of critical importance. Please read this section in its entirety.

BEA recommends that you follow these essential guidelines to secure LiquidVM in your production environment:

- Do not store sensitive data on an NFS server. LiquidVM does not encrypt the communication with the NFS-server. Therefore, the data can be snooped on the local network. In general, storing sensitive data on NFS-servers greatly increases the security threats to your system. If you are using an NFS file server, see “NFS Security Measures in the WLS-VE”, version 1.0 [Installation and Configuration Guide](#) for additional NFS security guidelines.

- LiquidVM provides a very secure runtime environment for the Java application out of the box, but the Java application has full access to its files. Therefore, it is important to ensure that the Java application running on LiquidVM is also secure.
- Use a firewall to protect LiquidVM instances running on a local network from external access. In particular SSH, DHCP, ARP and ICMP traffic should not be allowed to reach LiquidVM from an external point.
- Be sure to secure the VMware ESX servers so that no unauthorized users can gain root-access to these servers. LiquidVM is unable to protect itself if unauthorized root access to the ESX-servers is possible.
- Configure the VMware Virtual Infrastructure such that only users that are trusted to modify the runtime state of LiquidVM are given control and console access to the VM. LiquidVM cannot protect itself from VM shutdown and other VM related attacks if this policy is not maintained.
- Choose SSH passwords that are not obvious and store them securely. BEA recommends that your password contain a minimum of 8 characters, and consist of a combination of numbers, signs, and letters. It is critical that you follow this guideline because LiquidVM does not provide any kind of strength validation of the passwords.
- If you use SSH private keys, be sure to store them in a secure fashion on the client machine so that other users cannot gain access to the private key.
- SSH is disabled by default. If SSH is enabled, be sure to install a public key or set a password immediately. Once a password is set, temporary clear-text passwords and console displayed time-limited passwords will no longer work.
- Store the start scripts used by a remote launcher in a secure manner so that unauthorized users do not have access to them. If you do not store these start scripts safely, the LiquidVM startup arguments can be compromised and confidentiality can be breached.

Securing WLS

To ensure the most secure environment for running WLS-VE, BEA recommends that you take the basic security measures required for a non-virtualized implementation of WLS. These measures are:

- Secure the WLS host
- Secure network connections
- Secure your database

- Secure the WebLogic Security Service
- Secure any applications you plan to run

Refer to [Securing a Production Environment](#) for complete information on setting up basic WLS security. Also see the manufacturer's security documentation for any applications you plan to run on WLS-VE.

Securing the VMware VirtualCenter

If you plan to use VMware's VirtualCenter, you should follow all of the security practices recommended by VMware. See the [VMware Infrastructure Documentation](#) for more information.

You should use SSL to connect to VirtualCenter, as described in ["Setting Up VMware and Enabling SSL" on page 5-2](#).

In addition to taking the security measures recommended by VMware, you should also secure your VirtualCenter password by removing it (actually, the encrypted representation of it) from the `bea.lvm.info` file. While the password is stored in an encrypted form to provide a high level of security, you still run the risk of it being compromised. To remove it from the `bea.lvm.info` file, do the following:

1. Go to your home directory (or `//Documents and Settings/myDirectory` on Windows) and open the `bea.lvm.info` file.
2. Locate the statement `vmwarePassword=`.
3. Delete the string of characters following the `=`.

Once the password is removed from the `bea.lvm.info` file, you will need to supply it every time you try to create or start a WLS-VE instance.

Tuning the WLS-VE LiquidVM Kernel

With non-virtualized WLS you can tune your OS to improve the performance of your application. For information about tuning a standard OS for WLS, see [“Operating System Tuning”](#) in *WebLogic Server Performance and Tuning*.

The LiquidVM kernel was designed and developed to provide an optimized runtime for executing Java on the JRockit JVM, or more specifically, Java EE applications deployed on WLS on the JRockit JVM. The LiquidVM kernel is not a general purpose OS, and therefore, contains very few performance related configuration options.

However, there are certain system defaults which, for a given deployment, may be:

- Too conservative and therefore detrimental to performance
- Overzealous, resulting in (perhaps fatal) resource starvation

For the most part, such system defaults are set dynamically depending on a given load. Although, in certain circumstances, dynamic settings may be too slow to scale or may introduce unwanted indeterminism.

The following topics describe how the LiquidVM kernel can be tuned to improve the performance of WLS-VE applications:

- [“Tuning the LiquidVM Kernel Startup Options”](#) on page 9-2
- [“Comparing OS and LiquidVM Kernel Tuning”](#) on page 9-3

Tuning the LiquidVM Kernel Startup Options

The LiquidVM startup options described in [Table 9-1](#) can have an effect on performance and can be tuned as required.

Table 9-1 Tunable LiquidVM Startup Options

Option	Default	Description
<code>netRcvBufSz</code>	40KB	<p>System default <code>SO_RCVBUF</code>.</p> <p>Describes the maximum amount of received data a socket may buffer. The TCP receive window is based upon the free buffer size. Individual sockets may set this number using <code>java.net.Socket.setReceiveBufferSize(int)</code> (see the API Javadoc, Java™ 2 Platform Standard Edition 5.0 API Specification, for more information).</p> <p>Generally speaking, the higher the better, the cost is that the amount of memory used potentially limits the maximum number of sockets.</p>
<code>netSndBufSz</code>	Dynamic, max = 64KB	<p>Static system default <code>SO_SNDBUF</code>. Describes the maximum amount of transmit data a socket may buffer. Individual sockets may set this number using <code>java.net.Socket.setSendBufferSize(int)</code> (see API Javadoc, Java™ 2 Platform Standard Edition 5.0 API Specification, for more information). Normally, the LiquidVM kernel dynamically sets this number based on memory pressure. Explicitly setting this option disables the dynamic configuration, and statically sets the system default. Generally speaking, the higher the better; the cost is that the amount of memory used potentially limits the maximum number of sockets.</p>
<code>netTcpAto</code>	true	<p>Delayed ACK feature.</p> <p>Delays lone ACK packets with the expectation that the user will reply to the received data, piggybacking the ACK on the next transmission. Disabling this feature results in ACK being scheduled for immediate transmission, meaning better latency at the cost of "badput" (needless transmission, costing CPU).</p>

Comparing OS and LiquidVM Kernel Tuning

[Table 9-2](#) provides comparisons for basic OS tuning concepts between non-virtualized WLS and the LiquidVM kernel. For detailed information about OS tuning with non-virtualized WLS, see “[Operating System Tuning](#)” in *WebLogic Server Performance and Tuning*.

Table 9-2 Basic OS Tuning Concept Comparisons

Non-Virtualized WLS	LiquidVM Kernel
The default settings for the Windows OS are usually sufficient and do not need to be tuned.	This is the same for the LiquidVM kernel.
Most error conditions are TCP tuning parameter related and are caused by the OS's failure to release old sockets from a <code>close_wait</code> call. Common errors are "connection refused", "too many open files" on the server-side, and "address in use: connect" on the client-side. In most cases, these errors can be prevented by adjusting the <code>TCP wait_time</code> value and the TCP queue size.	Dynamically estimates effective MSL for any given connection to be a function of RTT.

[Table 9-3](#) lists common OS tunable parameters and their counterparts if they exist in the LiquidVM kernel. In many cases, these tunable parameters are not necessary for the LiquidVM kernel because it is a single-process, single-user operating environment designed to run one Java application most efficiently. Therefore, some of the tunable parameters in normal OSes can be completely eliminated. Tunable parameters that do not have counterparts in the LiquidVM kernel are listed as N/A.

Table 9-3 Tunable OS Parameters in the LiquidVM Kernel

OS Parameter	Relevance to LiquidVM Kernel
Solaris Tuning Parameters	
<code>/dev/tcp tcp_time_wait_interval</code>	N/A. Dynamically estimate effective MSL for each socket. For more information, see Table 9-2 .

Table 9-3 Tunable OS Parameters in the LiquidVM Kernel (Continued)

OS Parameter	Relevance to LiquidVM Kernel
/dev/tcp tcp_conn_req_max_q	N/A. LiquidVM kernel respects the accept backlog given by the user <code>listen(2)</code> . Java JDK provides a default of 50.
/dev/tcp tcp_conn_req_max_q0	N/A. LiquidVM kernel respects the accept backlog given by the user <code>listen(2)</code> . Java JDK provides a default of 50.
/dev/tcp tcp_ip_abort_interval	N/A.
/dev/tcp tcp_keepalive_interval	LiquidVM Kernel boot option "netKeepAlive"
/dev/tcp tcp_rexmit_interval_initial	N/A.
/dev/tcp tcp_rexmit_interval_max	N/A.
/dev/tcp tcp_rexmit_interval_min	N/A.
/dev/tcp tcp_smallest_anon_port	N/A. Smallest anonymous port is 1025.
/dev/tcp tcp_xmit_hiwat	LiquidVM Kernel boot option "netSndBufSz", else dynamically guided by memory pressure.
/dev/tcp tcp_rcv_hiwat	LiquidVM Kernel boot option "netRcvBufSz"
/dev/ce instance	N/A.
/dev/ce rx_intr_time	N/A.
set rlim_fd_cur	N/A. LiquidVM has no such limitation; memory is the only restriction.
set rlim_fd_max	N/A. LiquidVM has no such limitation; memory is the only restriction.
set tcp:tcp_conn_hash_size	N/A.
set shmsys:shminfo_shmmax	N/A.
set autoup	N/A.
set tune_t_fsflushr	N/A.

Table 9-3 Tunable OS Parameters in the LiquidVM Kernel (Continued)

OS Parameter	Relevance to LiquidVM Kernel
Linux Tuning Parameters	
<code>/sbin/ifconfig lo mtu</code>	N/A. There are no other processes to talk to, so if loopback performance "in-process" is important to your application, you are probably facing much bigger issues
<code>kernel.msgmni</code>	N/A. SystemV IPC configuration. There are no other processes to talk to, so if loopback performance "in-process" is important to your application, you are probably facing much bigger issues.
<code>kernel.sem</code>	N/A. SystemV IPC configuration. There are no other processes to talk to, so if loopback performance "in-process" is important to your application, you are probably facing much bigger issues.
<code>kernel.shmmax</code>	N/A. SystemV IPC configuration. There are no other processes to talk to, so if loopback performance "in-process" is important to your application, you are probably facing much bigger issues.
<code>fs.file-max</code>	N/A. LiquidVM Kernel has no such limitation; memory is the only restriction.
<code>net.ipv4.tcp_max_syn_backlog</code>	N/A. LiquidVM Kernel has no such restriction, we abide by the user's accept backlog setting (<code>listen(2)</code>)
HP-UX Tuning Parameters	
<code>tcp_conn_req_max</code>	N/A.
<code>tcp_xmit_hiwater_def</code>	LiquidVM Kernel boot option "netSndBufSz", else dynamically guided by memory pressure.
<code>tcp_ip_abort_interval</code>	LiquidVM Kernel boot option "netRcvBufSz"
<code>tcp_rexmit_interval_initial</code>	N/A
<code>tcp_keepalive_interval</code>	LiquidVM Kernel boot option "netKeepAlive"
Windows Tuning Parameters	

Table 9-3 Tunable OS Parameters in the LiquidVM Kernel (Continued)

OS Parameter	Relevance to LiquidVM Kernel
MaxUserPort	N/A
TcpTimedWaitDelay	N/A. Dynamically estimate effective MSL for each socket.

Diagnostics and Troubleshooting

This section describes how to deal with problems that occur in both the server and the JVM and problems endemic to WLS-VE specifically. It also describes how to obtain information about your WLS-VE instance and provide that information to BEA Support.

This section includes information on the following subjects:

- [“Troubleshooting WLS-VE Problems” on page 10-1](#)
- [“Handling Suspend Files” on page 10-7](#)
- [“Displaying Version Information” on page 10-7](#)
- [“Reporting a Problem to BEA Support” on page 10-7](#)

Troubleshooting WLS-VE Problems

This section provides information you will find helpful in solving problems that might occur with WLS-VE. Generally, you handle WebLogic Server and LiquidVM (the BEA JRockit component) problems the same way you would for their non-virtualized versions. You should follow BEA Support’s instructions for information collection, augmented with those in [Reporting a Problem to BEA Support](#). For BEA JRockit, you can use the standard tools available with BEA JRockit Mission Control—such as the JRockit Runtime Analyzer and Memory Leak Detector—to help you diagnose problems and collect relevant information about runtime activity.

This section contains information on the following subjects:

- [Troubleshooting Common WLS-VE Issues](#)

- [Troubleshooting WLS Issues](#)
- [Troubleshooting LiquidVM Issues](#)

Troubleshooting Common WLS-VE Issues

Problems with WLS-VE not specifically associated with WLS or with LiquidVM, can probably be traced to configuration errors. This section will help you identify the problem and figure out what caused it and how to resolve it. If you cannot find the solution here, collect the necessary information about your system, as described in [Reporting a Problem to BEA Support](#), and open a case with BEA Support.

Note: If you have configured your installation using an NFS share, you may encounter certain error conditions related to the NFS configuration. Common NFS error conditions are described in the WLS-VE version 1.0 [Installation and Configuration Guide](#).

The most common error conditions you might encounter are:

- [“Could not find the disk” Error](#)
- [The Server Shuts Down Soon After Startup](#)
- [The Server Shuts Down Soon After Startup](#)
- [“netSend failed: -3” Error](#)
- [“Configured IP \[...\] in use by MAC” Error](#)

“Could not find the disk” Error

Symptom: When you launch your instance and you get the following output in your OS console window:

```
Starting WLS-MyServer. connect...configure...create...
Could not find the disk: [storage2] wlsve/isoName.iso
```

Problem: When the virtual machine was created, VMware could not find the ISO image that you need to boot up WLS-VE.

Solution:

Check that you have uploaded the WLS-VE ISO image to the ESX server (see [Copying the WLS-VE ISO Image](#) in the *WLS-VE Installation Guide* for more information).

Verify that the `bea.lvm.info` file in your home directory points to the correct location on the ESX server. You can do this either by manually editing the `bea.lvm.info` file in your favorite

editor or by rerunning the LiquidVM configuration wizard (tools\virtualization\control_1.1\bin\lvm_configwizard.cmd or .sh), as described in [Chapter 3, “Configuring LiquidVM Connection Parameters.”](#)

Confirm that the ISO image exists, using the LiquidVM configuration wizard:

1. Select VM Host.
2. Select the **Configuration** tab and note the Datastore Name.
3. Select **Browse Datastore** and confirm that wlsve922.iso is available in your datastore.

The Server Shuts Down Soon After Startup

Symptom: The server shuts down soon after startup and a LiquidVM log file named WLS-*<servername>*.lvm.out appears in the domain directory on the local disk. In that file, you find the following:

```
<Jan 9, 2008 6:50:23 PM EST> <Info> <Management> <BEA-141107> <Version: WebLogic
Server 9.2 MP2 Fri Dec 7 01:21:28 EST 2007 1023546 >
<Jan 9, 2008 6:50:30 PM EST> <Info> <Security> <BEA-090065> <Getting boot
identity from user.>
Enter username to boot WebLogic server:The application tried to read from
keyboard (stdin). However,
reading from keyboard is not possible when running LiquidVM.
Please be aware that this could result in unexpected behaviour
if the application really depends on keyboard input working
<Jan 9, 2008 6:50:30 PM EST> <Error> <Security> <BEA-090783> <Server is
Running in Development Mode and Native Library(terminalio) to read the
password securely from commandline is not found.>
<Jan 9, 2008 6:50:30 PM EST> <Notice> <WebLogicServer> <BEA-000388> <JVM
called WLS shutdown hook. The server will force shutdown now>
<Jan 9, 2008 6:50:30 PM EST> <Alert> <WebLogicServer> <BEA-000396> <Server
shutdown has been requested by <WLS Kernel>>
<Jan 9, 2008 6:50:30 PM EST> <Notice> <WebLogicServer> <BEA-000365> <Server state
changed to FORCE_SHUTTING_DOWN>
```

Problem: You didn't provide a user name and password either in the security directory (no boot.properties file) of the Administration Server's root directory or in the start script using the WLS_USER and WLS_PW properties. Keep in mind that WLS-VE does not support normal keyboard input, so you cannot enter a username and password on the keyboard.

Solution: Add a boot.properties file to the security directory of the Administration Server's root directory (see [Creating a Boot Identity File for an Administration Server](#) in the WLS 9.2 document, *Managing Server Startup and Shutdown*) or add a username and password in the start script and relaunch your WLS-VE server.

“netSend failed: -3” Error

Symptom: You receive the following error message on the Virtual Center console:

```
000000 [rpcconn    WRN] netSend failed: -3
000000 [rpcconn    WRN] Rpc call failed
000000 [rpc        WRN] Rpc request failed: 3
000000 [rpc        WRN] rpcDoRegeust returned 3
000000 [rpc        WRN] rpcCall 3 returned 8549398
```

Problem: Your network configuration is incorrect.

Solution: In the start-up script, check your static IP address, your gateway, and your netmask and verify that they are correct. If they aren't, obtain the correct information and enter it in the respective property.

“Configured IP [...] in use by MAC” Error

Symptom: When you attempt to start a server, you receive this message:

```
000000 [net    WRN] Configured IP [172.18.134.55] in use by MAC: 00:50:56:a0:
06:96
000001 [net    WRN] Network stack initialization FAILED: 98
```

Problem: Someone else is already using the IP address you have specified.

Solution: Another running VM might be using the same IP address. Do the following:

- Verify that you typed the IP address correctly.
- Verify that none of your running VMs already use that IP address.

If neither of the above is the case, someone else is using your IP address. Since finding out who that person is might be difficult, contact your system administrator to obtain another IP address.

Troubleshooting WLS Issues

WLS-VE can encounter the same type of server-related problems that can occur when running non-virtualized WebLogic Server. This section provides an overview of the kinds of WebLogic Server problems you should watch for when running WLS-VE. It includes information on these subjects:

- [Performance Issues](#)
- [Server Failure](#)

- [Clustering Issues](#)
- [Other WLS Problems](#)

Performance Issues

Often, a problem with WebLogic Server is the result of poor tuning. For example, pool sizes (such as pools for JDBC connections, Stateless Session EJBs, and MDBs) that don't maximize concurrency for the expected thread utilization can adversely affect performance. Similarly, applications that handle large amounts of data per request will experience a boost in performance if the chunk size—that is, a unit of memory that the WLS network layer uses to read data from and write data to sockets—on both the client and server sides can be increased, a process called tuning the chunk size.

You can find well-tested tuning and performance guidelines in [WebLogic Server Performance and Tuning](#).

Specific tuning guidelines for WLS-VE are provided in [Chapter 9, “Tuning the WLS-VE LiquidVM Kernel.”](#)

Server Failure

A server instance can fail and different events can lead to this failure. Often one failure condition leads to another. Loss of power, hardware malfunction, OS crashes, network partitions, and unexpected application behavior can all contribute to the failure of a server instance. Even in a clustered environment, server instances may fail periodically and you must be prepared for the recovery process. See [Avoiding and Recovering From Server Failure](#) in *Managing Server Startup and Shutdown* for information on dealing with server failure.

Clustering Issues

A number of cluster problems can affect the performance of WebLogic Server. These problems can occur for many reasons, including licensing and versioning errors, multicast addressing problems, errors or misspellings in start-up commands, and even a poorly-tuned memory management systems. You can find guidelines for troubleshooting cluster problems in [Troubleshooting Common Problems](#) in *Using WebLogic Server Clusters*.

Other WLS Problems

Other, non-specific problems can also occur with WLS. When these problems occur, they usually generate an error message with an associated error code. The [Index of Messages by Message](#)

[Range](#) provides descriptions, possible causes, and corrective actions for all WebLogic Server error conditions.

Troubleshooting LiquidVM Issues

Problems that don't originate with WLS may occur in LiquidVM and are typical to the kinds of problems you might encounter in non-virtualized JVMs. Problems such as these are documented in the [BEA JRockit Diagnostics Guide](#) (BEA JRockit is the JVM component of LiquidVM). This document provides information for either resolving the problem yourself or mining the necessary information required to open a case with BEA Support.

The types of LiquidVM problems you might encounter when running WLS-VE are:

- System crashes occur when the entire system shuts down involuntarily and usually without warning. See [The System is Crashing](#).
- System freezes occur when the application stops answering requests but the process is still there. See [BEA JRockit is Freezing](#).
- Slow startups usually occur when BEA JRockit's optimizing compiler must run extensively to ensure that the most efficient code possible is compiled. See [BEA JRockit Starts Slowly](#).
- Poor performance usually occurs when your application experiences poor throughput. This usually indicates that the memory management system has not been tuned for optimal performance. See [Low Overall Throughput](#).
- Occasional slow response times usually indicate that transactions are taking too long to execute, a bottleneck most often caused by garbage collection pause times lasting too long. See [Long Latencies](#).
- Performance degrading after the application has been running is characterized by your application's behavior: it may be working fine early in its run, but after a while it may report the wrong results, or throw exceptions where it shouldn't, or it simply crashes or hangs at roughly the same time each time you run it. See [BEA JRockit's Performance Degrades Over Time](#).

Note that in most UNIX OSs there is a file descriptor limit that limits the number of files and sockets you can have open. LiquidVM does not have such limits so there is no need (and no way) to set a file descriptor limit.

For complete information on BEA JRockit problem determination and resolution, see [BEA JRockit Tools](#) in the [BEA JRockit Diagnostics Guide](#).

Handling Suspend Files

When WLS-VE crashes, the VM goes into a state of suspension. A pause button will appear on the VirtualCenter and information about the crash will be written to the console. When a suspend file is created, do the following:

1. `tar-gzip` the suspend file. You will find it on the VM's home directory on the ESX server; it will have a filetype of `.vmss`.
2. Copy the `tgz` file from the ESX server to your normal environment (for example, your `My Documents/` folder).
3. Upload the `tgz` file to BEA Support.

Be aware that you might not realize that your machine has actually crashed when it suspends. You should avoid the temptation to resume execution, as you might lose critical information that would be helpful in diagnosing the problems causing the crash. You should also be aware that suspend files are huge and might not be easy for you to copy from the ESX server.

Displaying Version Information

A critical piece of information that Support will need to help diagnose any problems you report to them is the version number. You can find this number in the file `LVM_VERSION`, which is located in the `BEA_HOME/tools/virtualization/control_1.1/` directory.

Open this file to find the version number; for example:

```
LiquidVM R1.1.7.0-93732
```

Reporting a Problem to BEA Support

If you determine that you need to file a trouble report, this section discusses what you need to do before opening the case to ensure that you supply the support personnel assigned to your issue as complete picture of what is wrong as possible. The more information you can provide, the more quickly will the support staff be able to resolve your issue. This section includes information on these subjects:

- [Trouble Reporting Process Overview](#)
- [Identify Your Problem Type](#)
- [Verify That You're Running a Supported Configuration](#)

- [Collect Enough Information to Define Your Issue](#)

Trouble Reporting Process Overview

When you encounter a problem with WebLogic Server Virtual Edition and can't resolve it using the information provided in the relevant BEA documentation, you need to collect the information that best describes your problem and open a case with BEA Support. If you have a service agreement with BEA, the normal process is to contact your Level 1 service provider, who will make the initial attempts to correct the problem. If the case cannot be solved by the Level 1 staff, it is escalated to the Level 2 staff, who will draw on their particular expertise to get your JVM running again. For serious problems, the issue will be handled by the Level 3 staff (the WebLogic Server Virtual Edition developers).

Identify Your Problem Type

Is your machine crashing? Is it running slowly or returning unpredictable results? These are the kind of symptoms that indicate a problem with WebLogic Server Virtual Edition. Being able to identify what kind of problem you are experiencing will help you know what kind of information you need to include when you open the trouble report.

Verify That You're Running a Supported Configuration

Before submitting a bug, verify that the environment in which the problem was found is a supported configuration. See the [Supported Configurations](#) page for WLS-VE 1.1.

Collect Enough Information to Define Your Issue

In addition to testing with the latest update release, use the following guidelines to prepare for submitting a trouble report:

1. Collect as much relevant data as possible. For example, generate a thread-dump in the case of a deadlock, or locate the core file (where applicable) and `hs_err` file in the case of a crash. In all cases it is important to document the environment and the actions performed just before the problem is encountered.
2. Where applicable, try to restore the original state and reproduce the problem using the documented steps. This helps to determine if the problem is reproducible or an intermittent issue.
3. If the issue is reproducible, try to narrow the problem. In some cases, a bug can be demonstrated with a small standalone test case. Bugs demonstrated by small test cases will

typically be easy to diagnose when compared to test cases that consists of a large complex application.

4. Search the bug database to see if the bug, or similar bugs, have been reported. If the bug has already been reported, the bug report may have further information. For example, if the bug has already been fixed, it will indicate the release that the bug was fixed in. The bug may also contain information, such as a workaround or include comments in the evaluation that explain, in further detail, the circumstances that cause the bug to arise.

If you conclude that the bug has not already been reported, then it is important to submit a new bug.

