



BEA WebLogic Server® Virtual Edition

Configuration and User Guide

Version 10 v1.2
Revised: April 2008

Contents

1. Configuration Overview and Roadmap

Configuration Overview	1-1
Roadmap for Configuring a WLS-VE Domain	1-4
Converting a Physical Domain to a Virtual Domain	1-4
Comparing Startup, Configuration, and Logging Options to Non-virtualized WLS ...	1-6

2. Accessing and Using the WLS-VE Distribution Archive

Obtaining the WLS-VE Distribution Archive.....	2-2
Unpacking the WLS-VE Distribution Archive.....	2-2
The BEA Home Directory	2-2
Sharing the BEA Home Directory on the Launcher Machine with Other BEA Products	2-3
Copying the WLS-VE ISO Image.....	2-3
Upgrading and Promoting Domains	2-6
Upgrading from a Non-virtualized WLS 10.0 to WLS-VE 10.0.....	2-6
Upgrading from an Earlier Version of WLS to WLS-VE 10.0	2-6
Downloading and Applying Patches.....	2-7
Running Smart Update on an Ordinary OS and Copy the Patches	2-7
Running Smart Update on a BEA Home on an NFS Share.....	2-7
WLS Service Packs	2-8
What's Next?.....	2-8

3. Understanding and Using LiquidVM

What is LiquidVM?	3-2
Understanding the LiquidVM File System	3-3
Using the Virtual Local Disk	3-3
Determining the Amount of Free Disk Space	3-4
Increasing the Size of the Local Disk	3-4
Using the LiquidVM SSH Service	3-4
SSH Listen Port	3-5
Authenticating with the SSH Service	3-5
Installing a Real Password In Addition to a Public Key	3-6
Auditing SSH Actions	3-6
Using LiquidVM in Passive Mode	3-6

4. Configuring LiquidVM Connection Parameters

Before You Begin	4-2
Configuring LiquidVM in Graphical Mode	4-3
Configuring LiquidVM in Console Mode	4-6
Understanding the bea.lvm.info File	4-9
Troubleshooting a LiquidVM Configuration	4-10

5. Creating Virtual WebLogic Domains

Overview	5-2
Before You Begin	5-3
Preparing an Existing WebLogic Domain for Virtualization	5-3
Out-of-Domain Dependencies	5-4
ListenAddress	5-4
WLST connect()	5-4
Using the P2V Domain Conversion Utility	5-5

Run the P2V-Generated Start Scripts	5-10
Copying Domain Artifacts Using the LiquidVM SSH Service	5-11
Task 1: Start LiquidVM in Passive Mode	5-12
Task 2: Copy the Files to the WLS-VE Instance.	5-14
Task 3: Shut Down the LiquidVM Instance	5-15
Task 4: Re-start the WLS-VE Instance	5-15
Configuring a Virtual Domain to Access an NFS Server.	5-16
Moving a WLS-VE Domain to a Production Environment	5-16
Additional Configuration Tasks	5-17
Tuning LiquidVM	5-17
Deploying an Application to WLS-VE.	5-18

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

Setting Up VMware and Enabling SSL	6-2
Starting WLS-VE Instances	6-2
Editing VM Properties	6-2
Pausing a VM	6-5
Working with the Console Tab.	6-5
Inactive Keyboard	6-5
Console Log.	6-6
Pre-console Log.	6-6

7. Starting and Stopping WLS-VE

Starting WLS-VE	7-2
Starting WLS-VE from a Command-line	7-3
Starting WLS-VE from the VMware VI Client.	7-3
Starting the Administration Console	7-4
Stopping WLS-VE Instances	7-4

8. Configuring Logging

Understanding the Log Files	8-1
WLS Logs	8-1
LiquidVM Log	8-2
VMware Log	8-2
Accessing WLS-VE Log Files	8-2
Copying the Log Files Using SSH	8-3
Configuring LiquidVM to Use Remote syslog	8-3
Storing the Log Files on an NFS Share.	8-4

9. Securing Your Production Environment

Securing LiquidVM	9-1
Securing WLS	9-3
Securing the VMware VirtualCenter	9-3

10. Tuning the WLS-VE LiquidVM Kernel

Tuning the LiquidVM Kernel Startup Options	10-2
Comparing OS and LiquidVM Kernel Tuning.	10-3

11. Diagnostics and Troubleshooting

Troubleshooting WLS-VE	11-1
Diagnosing WLS-VE Issues	11-2
“Could not find the disk” Error	11-2
Server Shuts Down Soon After Startup	11-3
“netSend failed: -3” Error	11-3
“Configured IP [...] in use by MAC” Error	11-4
Diagnosing WLS Issues	11-4
Performance Issues	11-4
Server Failure.	11-5

Clustering Issues	11-5
Other WLS Issues	11-5
Diagnosing LiquidVM Issues	11-5
Handling Suspend Files	11-6
Displaying Version Information	11-7
Reporting a Problem to BEA Support	11-7
Verify That You Are Running a Supported Configuration	11-7
Collect Enough Information to Define Your Issue	11-7

12. Adding Managed Servers to a WLS-VE Domain

Steps for Adding a Managed Server to WLS-VE Domain	12-2
The P2V-Generated Start Script Properties	12-5

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 tasks required to successfully install and configure WLS-VE:

- [Configuration Overview](#)
- [Roadmap for Configuring a WLS-VE Domain](#)
- [Comparing Startup, Configuration, and Logging Options to Non-virtualized WLS](#)

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 on WLS 10.0 MP1



[Table 1-1](#) describes the components of the sample WLS-VE configuration.

Table 1-1 WLS-VE 10 Configuration Components

Component	Description
Launcher Machine	<p>The Windows or Linux machine where the WebLogic domain resides and where you unpackage the WLS-VE distribution archive. It is also the machine on which you run the LiquidVM Configuration Wizard, the P2V Domain Conversion utility, and initiate the creation of WLS-VE instances on the VMware ESX machine.</p> <p>The WLS-VE distribution archive includes LiquidVM 1.2, which provide the tools needed to create, control, and start LiquidVM instances. The archive also includes an ISO image for WebLogic products. For example, the WLS ISO image contains 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 1.2 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>
VMware ESX Server	<p>The hypervisor machine with VMware installed that is available for the creation of virtual machines. The LiquidVM launcher accesses the ESX server through the VirtualCenter server.</p>
Virtual Infrastructure	<p>The full infrastructure virtualization suite that provides the hardware and software resources required to support a virtualized environment.</p>
WLS-VE instances	<p>Virtual WLS instances running on the ESX server. The WLS-VE 10 LiquidVM includes a virtual local disk, which can be mapped to the SAN, for each WLS-VE instance. You can transfer files to and from the local disk using the SSH service.</p>
SAN Disk	<p>Recommended physical storage device for the ISO image and the local disks on the WLS-VE instances.</p>

Roadmap for Configuring a WLS-VE Domain

The following section summarizes the overall process for installing and configuring WLS-VE, depending on whether you are either:

- Installing a WebLogic product from scratch and configuring a domain, and then using WLS-VE to virtualize the domain.
- Using WLS-VE to virtualize an existing WebLogic domain.

The procedures for downloading and unpacking the WLS-VE distribution archive are documented in [Accessing and Using the WLS-VE Distribution Archive](#). Subsequent sections of this document provide details for configuring and using WLS-VE.

Converting a Physical Domain to a Virtual Domain

[Table 1-2](#) describes the main steps required to convert an existing WLS 10.0 MP1 domain to a corresponding virtual domain on the VMware ESX machine.

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

Step	Description
1. Create a new physical domain or make a copy of an existing domain.	<ul style="list-style-type: none">• If necessary, configure a new domain—The Physical-to-Virtual (P2V) Domain Conversion utility only converts physical WebLogic domains and Managed Servers into virtual instances. Complete details for creating WebLogic domains is provided in Creating WebLogic Domains Using the Configuration Wizard.• Make a copy of the existing domain—For existing WebLogic 10.0 domains, make a copy of the entire domain, and use the back-up copy of the domains' configuration file (<code>config.xml</code>) for the necessary editing and transferring of domain artifacts to the virtual environment.
2. Make sure the physical domain is ready for virtualization.	<p>The P2V utility does not handle out-of-domain dependencies, such as referenced libraries, nor does it modify the domain configuration file (<code>config.xml</code>) or any other files under the domain directory. Therefore, any path references that are valid only for the local physical machine must be modified with respect to the LVM file system path. Also, if server instances in the domain use the <code>ListenAddress</code> attribute, the address must be removed for that server to be virtualized.</p> <p>See Preparing an Existing WebLogic Domain for Virtualization.</p>

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

Step	Description
3. Generate an SSH key-pair	<p>If public key authentication is to be used for SSH connections to LVM instances, then generate an SSH key-pair on the launcher machine using a tool such as <code>ssh-keygen</code>.</p> <p>Note: Only RSA format is supported, so the public key string must begin with <code>ssh-rsa</code>.</p>
4. Download and unzip the WLS-VE distribution archive.	Details about downloading and unpacking the distribution archive are provided in Accessing and Using the WLS-VE Distribution Archive .
5. Copy the ISO image to the ESX server.	Details about the copying the ISO image are provided in Copying the WLS-VE ISO Image .
6. Configure the LiquidVM connection parameters.	<p>On the launcher machine where the WebLogic domain resides, run the LiquidVM Configuration Wizard to configure the connection between the LiquidVM tools running on the launcher machine and the hypervisor environment (Virtual Center and ESX server).</p> <p>For details, see Configuring LiquidVM Connection Parameters</p>
7. Convert the WebLogic domain to a virtualized domain.	<p>On the launcher machine, run the P2V Domain Conversion utility to convert the physical WLS 10.0 MP1-based domain (or selected servers in the domain) to a virtual domain in the hypervisor environment. The P2V utility makes sure that the domain, patches, and connection information is available for the new WLS-VE instances.</p> <p>For details, see Using the P2V Domain Conversion Utility.</p> <p>Note: In many cases, additional steps are required to complete the transformation of physical WebLogic domains into virtual domains.</p>
8. Use the P2V-generated start scripts to complete the configuration and start the WLS-VE instances.	The P2V utility generates WLS-VE start scripts in the domain directory on the launcher machine. You must run these start scripts to complete the configuration process and start the new WLS-VE instances, as described in Run the P2V-Generated Start Scripts .

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

Step	Description
9. If necessary, copy additional files to the WLS-VE instances.	If additional files need to be copied to a WLS-VE instance, such as application files than are stored outside the domain, use the LiquidVM SSH service to securely copy them to the hypervisor environment. See Copying Domain Artifacts Using the LiquidVM SSH Service .
10. Administer the WLS-VE environment.	To administer the WLS-VE environment, see the following topics: <ul style="list-style-type: none">• Working with WLS-VE Using the VMware VI Client• Configuring Logging• Moving a WLS-VE Domain to a Production Environment• 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 10.0 MP1, 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.
- In a physical WebLogic domain, the encrypted administrator user name and password, which are stored in the domain directory in the `boot.properties` file, can be passed on the command-line when booting the Administration Server. WLS-VE does not provide an interactive prompt, however, so you cannot specify credentials when the virtual server is starting. Therefore, if you want to use these credentials for the virtual domain, the `boot.properties` file must be available when using the P2V utility to virtualize the domain. See [Using the P2V Domain Conversion Utility](#).
- 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 WLS-VE instance. VMware logs are also created and stored on the ESX server.

To view the log files on the WLS-VE instance's local disk, the instance must be running in passive mode with SSH enabled. Then you can 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 [Configuring Logging](#).

- If the Administration Server will use a fixed IP address, you need to specify it when you convert the physical domain using the P2V utility. If you do not specify an IP address, DHCP assigns one from the available network addresses. The static IP or the IP obtained using DHCP for the Administration Server will be used as the host address in the `admin_url` parameter set in the Managed Server start scripts that are generated by the P2V utility.
- Note:** If DHCP is selected in the P2V utility for the Administration Server, then no IP address will be set in the generated start scripts and LVM will query the DHCP server to obtain an IP address. If the IP address's time-to-live has expired, then LVM can get a new IP address. Therefore, BEA does not recommend using DHCP for the Administration Server because in such a situation the `admin_url` parameter in the Managed Server's start script will be invalid. If DHCP is used, your network administrators must guarantee that the Administration Server gets a fixed IP from the DHCP server.

Accessing and Using the WLS-VE Distribution Archive

The following sections describe how to access and use the WLS-VE distribution archive contents, and how to prepare a domain for virtualization:

- [Obtaining the WLS-VE Distribution Archive](#)
- [Unpacking the WLS-VE Distribution Archive](#)
- [Copying the WLS-VE ISO Image](#)
- [Upgrading and Promoting Domains](#)
- [Downloading and Applying Patches](#)
- [What's Next?](#)

Obtaining the WLS-VE Distribution Archive

For this release, WLS-VE is only available as a distribution archive and does not include a BEA product installer mechanism. You can download the WLS-VE distribution archive from the [BEA web site](#).

Unpacking the WLS-VE Distribution Archive

After obtaining the WLS-VE distribution archive for WLS 10.0 MP1, follow these steps to unpack the archive.

1. Navigate to the directory containing the distribution archive file: `server1001ve12.tar.gz`.
2. BEA recommends unzipping the archive to the BEA Home directory for your WLS 10.0 MP1 installation. See [The BEA Home Directory](#).
3. Once unzipped, verify the following product directory structure:

Table 2-1 WLS-VE 10.0 Product Installation Directory Structure

Directory . . .	Description . . .
<code>server1001ve12</code>	The WLS-VE Home directory.
<code>iso</code>	The WLS-VE ISO image (<code>wlsve1001.iso</code>) contains the LiquidVM and WebLogic Server classes that you use to host your business applications. Each physical machine that hosts an instance of WLS-VE needs access to this ISO image file.
<code>tools</code> directory	Contains the P2V Domain Conversion Utility, the LiquidVM Configuration Wizard, and the WLS-VE startup utilities.
<code>tools/lib</code> directory	Contains the WLS-VE runtime utilities.
<code>WLSVE_VERSION</code> file	Contains the WLS-VE version number.

The BEA Home Directory

The BEA Home directory serves as a repository for your files that facilitate any future upgrades or installation of patches. If you already have a BEA Home directory on the local disk of the launcher machine, unzip the distribution archive in the existing directory.

Note: Optionally, you can also create your BEA Home directory on an NFS file server that is accessible to the WLS-VE installer through an NFS mount. For instructions on using an NFS file server, see the [Creating and Sharing Directories](#) section in the WLS-VE v1.0 *Installation and Configuration Guide*.

If the BEA Home is not on an NFS share, to install update patches on the VM instance, you first need to install them on the local disk of the launcher machine, and then copy them to the local disk of each VM. For more information, see [Moving a WLS-VE Domain to a Production Environment](#) and [Downloading and Applying Patches](#).

Note: The copy process is only necessary for installing update patches.

Sharing the BEA Home Directory on the Launcher Machine with Other BEA Products

The BEA Home directory can be considered a *central support directory* for all the BEA products installed on your system. For example, if you use WLS, WLW, and WLS-VE, you can maintain a single BEA Home directory for all products on the launcher machine. However, each WLS-VE instance has its own BEA Home directory, `/bea`, by default.


Copying the WLS-VE ISO Image

When you unpack the WLS-VE distribution archive on the launcher machine, the WLS 10.0 MP1 ISO image (`wlsve1001.iso`) is copied to the `/server1001ve12` directory. The ISO image contains the LiquidVM and WebLogic Server classes that run on hypervisor software and host your Java applications. To give the ESX server access to these classes, you need to copy the ISO image to a datastore on each ESX server that will run WLS-VE. (See [Figure 2-1](#).) A recommended best practice is to copy the ISO image to a SAN that can be accessed from each ESX server host.

Notes: Copying the ISO image to a physical disk that is local to the ESX server may disable some VMware functionality, such as VMotion.

If you are using an NFS share, datastores will experience degraded performance when booting WLS-VE.

Figure 2-1 Copy the ISO Image to Datastores for Each ESX Server



To copy the WLS-VE ISO image to a datastore on a local disk or SAN, you can use the following secure copy (`scp`) command syntax:

```
scp -p source-file username@esxhost:/vmfs/volumes/datastore/path/
```

where:

- `-p` preserves modification times, access times, and modes from the original file.
- `source-file` is the relative or absolute path and file name of the WLS-VE ISO image file.
- `username` is the name of a user in your network who has write privileges on the hypervisor host machine.
- `esxhost` is the name of an ESX server host.
- `/vmfs/volumes/` is the directory within the VMware file system under which ESX Server stores datastores.
- `datastore` is the name of a datastore.
- `path` is one or more optional directory levels.

Notes:

- BEA recommends using `scp` to copy ISO images. The Linux operating system includes an `scp` client. For Windows, you can install (or download for free) third-party utilities that include `scp` clients. Also note that `scp` is just one of many ways to copy the disk, depending on how you have configured your VMware environment. For example, in a Windows environment, the SAN disk could be an additional drive that is accessible to the ESX server; therefore, you can copy the ISO image to that drive using Windows Explorer.

- Copying an ISO using `scp` requires that the ESX server has been configured to allow `user-login` over SSH (which is off by default).

Be sure to record the pathname that you specify. When you use the LiquidVM Configuration Wizard to configure virtual machines, you will need to provide this location. Note that the syntax in ESX Server for specifying this path is:

```
[datastore] /path/file
```

For example, you can use the following `scp` command to copy the WLS-VE ISO image from the current directory to the default `storage1` datastore (VMware creates the `storage1` datastore when you install ESX Server):

```
scp -p wlsve1001.iso myusername@myESXHost:/vmfs/volumes/storage1/
```

With the above example, the pathname that you specify in the LiquidVM Configuration Wizard is:

```
[storage1] /wlsve1001.iso
```

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.

Upgrading from a Non-virtualized WLS 10.0 to WLS-VE 10.0

When you upgrade an application from a non-virtualized implementation of WLS 10.0 to WLS-VE 10, you need to make some modifications to the application code to ensure successful operation. See [Preparing an Existing WebLogic Domain for Virtualization](#).

Upgrading from an Earlier Version of WLS to WLS-VE 10.0

You cannot migrate an application directly to WLS-VE 10.0 MP1 v1.2 from an earlier version of WLS-VE (for example, version 9.2 v1.1). Instead, you need to follow the upgrade procedures outlined in [Upgrading WebLogic Application Environments](#) to upgrade your domain to a WLS 10.0 MP1 domain. Once you have done so, prepare the domain for virtualization as described in [Preparing an Existing WebLogic Domain for Virtualization](#).

Downloading and Applying Patches

WLS patches can be installed on each WLS-VE instance separately without shutting down the entire domain or cluster. How WLS-VE is patched depends on where the issue occurs and how the WLS-VE instance is configured in your environment, as follows:

- If there is a LiquidVM issue, either in the OS-layer or in the JVM layer, then a new CD-ISO-image must be acquired from BEA Product Support and uploaded to the ESX Server. Also, the WLS-VE instance has to be configured to use the new ISO.
- If there is a WLS issue, then you can use the normal WLS patch mechanism. Follow the instructions provided in the following sections that are appropriate to your WLS-VE configuration.

For information about upgrading your software with maintenance patches and service packs, if available, see [Installing Maintenance Updates and Service Packs](#).

Running Smart Update on an Ordinary OS and Copy the Patches

If you created your BEA Home on the local disk of the launcher machine, follow these steps:

1. Run Smart Update on the OS using a BEA Home directory on that machine, as explained in [Starting Smart Update](#). This procedure will create a `patch_weblogicNNN` directory under the BEA Home, where *NNN* specifies the WLS release (for example, 1001 for 10.0 MP1).
2. Follow steps 1 and 2 in the [Moving a WLS-VE Domain to a Production Environment](#) section, which explains how to use SSH to copy files to your WLS-VE instances.
3. Using SSH, copy the entire `patch_weblogicNNN` directory into the instance's `/bea` directory (where `/bea` is the BEA Home directory on your WLS-VE instance).

Running Smart Update on a BEA Home on an NFS Share

If your server has configured a BEA Home on an NFS share, you can run the Smart Update tool from an ordinary OS that also has access to the same NFS share with the BEA Home on it. Smart Update will apply the patches and put them in the BEA Home on the NFS share. The next time you restart your WLS-VE instance the patches will get automatically loaded.

For more information on using NFS shares, refer to the [Creating and Sharing Directories](#) section in the WLS-VE v1.0 *Installation and Configuration Guide*.

WLS Service Packs

If you want to apply a WLS service pack, a new CD-ISO-image must be acquired from BEA Product Support and uploaded to the ESX Server. You must also reconfigure your WLS-VE instance to use the service pack ISO instead of the previous CD-ISO.

What's Next?

After you have unpackaged the software and copied the ISO to the hypervisor host, you need to configure your environment, and create WLS-VE instances.

- Configure the LiquidVM connection parameters, as described in [Configuring LiquidVM Connection Parameters](#).
- Prepare the WLS-VE domain on the launcher machine, use the P2V utility to stage the WLS-VE instances on the ESX server, and run the necessary P2V-generated start scripts, as described in [Creating Virtual WebLogic Domains](#).
- Start and administer the WLS-VE environment (for example, security, logging, and troubleshooting), as described in:
 - [Starting and Stopping WLS-VE](#)
 - [Securing Your Production Environment](#)
 - [Configuring Logging](#)
 - [Tuning the WLS-VE LiquidVM Kernel](#)
 - [Diagnostics and Troubleshooting](#)

For specific information about this WLS-VE release, see the WLS-VE [Release Notes](#).

Understanding and Using LiquidVM

The following topics provide an overview of LiquidVM and describe how to use its main features:

- [What is LiquidVM?](#)
- [Understanding the LiquidVM File System](#)
- [Using the Virtual Local Disk](#)
- [Using the LiquidVM SSH Service](#)
- [Using LiquidVM in Passive Mode](#)

What is LiquidVM?

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

LiquidVM provides the following features:

- JRockit JVM as the Java runtime component.
- Java run-time class libraries.
- 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 WLS-VE instance. 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 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 3-1](#).

Table 3-1 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/wlsve1001.iso</pre>
/bea	<ul style="list-style-type: none"> patch_weblogicnnn 	<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 any WLS patches.</p>
/domain	<ul style="list-style-type: none"> server-root 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 WLS-VE instance. The local disk can be mapped to a SAN disk attached to the ESX server. You specify the size of the disk in the WLS-VE start script using the P2V Domain Conversion utility. The default is 500 MB. For more information, see [Using the P2V Domain Conversion Utility](#). 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](#).

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 (for example, 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](#).

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 [The P2V-Generated Start Script Properties](#). 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.

Installing a Real Password In Addition to a Public Key

If you should happen to lose your private key, it is good idea to have a secure *real* password as well, so that you can still log in and access the files on the local disk. You can do password-based authentication that is not clear-text, follows:

1. Using either the public/private-key authentication or the unsafe clear-text password, log into the virtual machine using SSH.
2. Type `passwd`. You will be asked for the existing password. If you are using a public key, leave it blank and press `Enter`.
3. When prompted for new password, enter a secure real password.
4. Confirm the new password.

This installs a real password; therefore, the unsecure clear-text password will no longer work. The public key authentication will continue to work if you prefer to use that method.

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 [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 4-1](#). You will be prompted to provide this information as you step through the configuration wizard.

Table 4-1 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 nameVC 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 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.

Table 4-1 LiquidVM Configuration Data (Continued)

Configuration Data	Description
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, [Storage1].</p> <p>The storage location and path can be specified as:</p> <pre>[storage1] myLocalStore/myISO.iso</pre> <p>The ISO image is installed when you unpack the WLS-VE distribution archive and is then manually copied to the ESX server as described in Copying the WLS-VE ISO Image.</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](#).

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

1. Depending on your OS, start the LiquidVM Configuration Wizard in graphical mode as shown here.

To start the Liquid VM Configuration Wizard on this platform . . .	Perform the following steps . . .
Windows	<p>From the command-line:</p> <ol style="list-style-type: none"> 1. Open a Command Prompt window and navigate to <code>server1001ve12\tools\</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>server1001ve12/tools/</code> 3. Enter the following command: <code>lvm_configwizard.sh</code>

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](#).
 - 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 (for example, `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](#).

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](#). You can also use console mode on a Windows platform.

Be sure you have access to the configuration data provided in [Table 4-1](#), 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, and so on) 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. Depending on your OS, start the LiquidVM Configuration Wizard as shown here:

To start the Liquid VM Configuration Wizard on this platform . . .	Perform the following steps . . .
Windows	<p>From the command-line:</p> <ol style="list-style-type: none">1. Open a Command Prompt window and navigate to <code>server1001ve12\tools\</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>server1001ve12/tools/</code>2. Enter the following command: <code>lvm_configwizard.sh -mode=console</code>

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 **y**, 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/wlsve1001.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 (for example, `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](#).

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 is `\\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 4-1](#) shows an example of a `bea.lvm.info` file.

Listing 4-1 Sample bea.lvm.info File

```
#BareMetal ESX-launcher configuration information
#Mon Mar 24 21:08:15 EDT 2008
vmwareUsername=user1
vmwareDiskPath=[storage1] wlsve/wlsve1001.iso
vmwareVcHost=vmwarevc.bea.com
vmwareKeystore=C:\\Documents and
Settings\\jtsmith\\bm_vmwarevc.bea.com.keystore
vmwarePassword=d90f1423925849c78e6dd9100d162f3f
vmwareComputeResource=esx.bea.com
vmwareResourcePool=TESTPOOL
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 issues that may result from the LiquidVM instance configuration, see [Troubleshooting WLS-VE](#).

Creating Virtual WebLogic Domains

The following topics describe the procedures to create virtual WebLogic domains and also provide additional configuration information:

- [Overview](#)
- [Before You Begin](#)
- [Using the P2V Domain Conversion Utility](#)
- [Copying Domain Artifacts Using the LiquidVM SSH Service](#)
- [Configuring a Virtual Domain to Access an NFS Server](#)
- [Moving a WLS-VE Domain to a Production Environment](#)
- [Moving a WLS-VE Domain to a Production Environment](#)
- [Additional Configuration Tasks](#)
- [Deploying an Application to WLS-VE](#)

Overview

The Physical-to-Virtual (P2V) Domain Conversion utility for WebLogic helps simplify the conversion of WLS Managed Servers in a domain to WLS-VE instances in the hypervisor environment, as follows:

- Stages WLS-VE instances on the ESX server for selected servers in a domain, and makes sure that the domain, patches, and connection information is available for the new WLS-VE instances. You create one such WLS-VE instance for each Managed Server (and optionally the Administration Server) in a domain.
- Generates start scripts in the domain directory on the launcher machine (the host on which the tool is run). These scripts are used to complete the configuration process and start the newly generated WLS-VE instances.

The P2V utility packages up the following WLS artifacts:

- Product artifacts — WLS patches and security files. The P2V utility will either copy the `boot.properties` file from the `/AdminServer` directory or create such a file from the user-specified name and password. It will also either copy/install the SSH public key when specified or set the password when this option is selected.
- Domain artifacts — All shared files in the entire domain directory (excluding logs and tmp, and only the specific folder (under the `/servers` directory) for each server being virtualized.
- Application artifacts — EARs deployed in the domain (does not include referenced applications or libraries).

The P2V utility also generates start scripts in the domain directory on the launcher machine to complete the configuration and start the newly generated WLS-VE instances. A sample WLS-VE configuration is illustrated and described in the [Configuration Overview](#).

Note: If you are using an NFS share, you must configure any NFS connections after using the P2V utility to create a virtualized domain.

Before You Begin

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

- Java SE 5.0 or above must be installed on the launcher machine where the P2V utility will be run.
- ESX artifacts (datacenter, host or cluster, optional resource pool, datastore, and so on) exist.
- 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 WLS-VE configuration process. See [Copying the WLS-VE ISO Image](#).
- 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. See [Configuring LiquidVM Connection Parameters](#).
- Depending on the networking options you want to use for LVM, confirm your configuration information:
 - DHCP — Uses DHCP for WLS-VE network configuration.

Note: If DHCP is used for the Administration Server, your network administrators must guarantee that the it gets a fixed IP from the DHCP server.
 - Static IP and DHCP — Uses a static IP address but DHCP will be used for other networking attributes.
 - Static IP — Obtains a set of IP addresses for WLS-VE instances. Also obtain networking information (for example, netmask, gateway, DNS servers, and domain name) if static configuration is to be used for all networking attributes.
- If public key authentication is to be used for SSH connections to LVM instances, then generate SSH key-pair on the launcher machine using a tool like `ssh-keygen`. You should also review the [Using the LiquidVM SSH Service](#) documentation.

Preparing an Existing WebLogic Domain for Virtualization

When you convert from a non-virtualized implementation of a WLS 10.0 MP1 domain to a WLS-VE domain, you may need to make some modifications to the domain's configuration file (`config.xml`) and/or application code to ensure successful operation.

Tip: BEA recommends making a back-up copy of your entire domain, and then editing and using this copy when virtualizing your domain so your original domain remains intact.

Out-of-Domain Dependencies

The P2V utility does not handle out-of-domain dependencies, such as referenced libraries, nor does it modify the WLS domain configuration file (`config.xml`) or any other files under the domain directory. Therefore, any path references in the `config.xml` that are valid only for the local physical machine must be modified with respect to the LVM file system path using the UNIX-style path convention, as follows:

- Change the path to one that is relative to the domain directory, using the UNIX-style path convention (for example, `./applications/myapp/myapp.jar`)

In this case, the files must be present in the domain directory before the P2V utility is run to transfer the domain.

- Change the path to an absolute path with respect to LVM file system, using UNIX style path convention (for example, `/usr/apps/myapp/myapp.jar`)

In this case, the P2V utility is not going to transfer the files to LVM file system. Instead, the files should be copied using a `scp` or `sftp` client after the LVM instance is created and then restarted in passive mode. See [Copying Domain Artifacts Using the LiquidVM SSH Service](#).

ListenAddress

If server instances that are to be virtualized use the `ListenAddress` attribute (including the `ListenAddress` attribute in any configured `NetworkAccessPoint`), you must remove the address for those servers so that WLS automatically binds to the IP address of LVM.

WLST connect()

If an 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.

Using the P2V Domain Conversion Utility

The following steps explain how to use the P2V Domain Conversion utility to convert the physical WLS instances in the domain to WLS-VE instances in the hypervisor environment. You create one such instance for each Managed Server (and optionally the Administration Server) in the domain.

Complete details for creating WebLogic domains is provided in [Creating WebLogic Domains Using the Configuration Wizard](#).

1. Depending on your OS, start the P2V Domain Conversion utility on the launcher machine as shown in [Table 5-1](#):

Table 5-1 Starting the P2V Utility

To start the P2V utility on this platform . . .	Perform the following steps . . .
Windows	<p>From the command-line:</p> <ol style="list-style-type: none"> 1. Open a Command Prompt window and navigate to <code>server1001ve12\tools\</code> 2. Enter the following command: <code>p2v_wizard.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>server1001ve12/tools/</code> 3. Enter the following command: <code>p2v_wizard.sh</code>

The P2V Domain Conversion utility launches.

2. On the **Select Domain & BEA-home** window, specify the BEA Home and WebLogic domain to virtualize:
 - **BEA-home:** Enter the BEA Home on the launcher machine that contains the patches for the domain that you want virtualize.
 - **WL-domain to virtualize:** Enter a domain name and a location on the launcher machine that you want to virtualize. The default location is `BEA_HOME\user_projects\domains`. At a minimum, the domain directory must contain a `\config` directory with a `config.xml` file.

Click **Next**.

3. On the **Provide WLS user & password** window, enter a username and password to be used by the WebLogic Administrator and click **Next**.

Note: This option is only displayed if these credential are not stored in a boot identity file (`boot.properties`) in the physical domain, specifically, in the Administration Server's security directory (for example, `domain\servers\AdminServer\security\boot-properties`). When you start the Administration Server, it refers to the boot identity file for the encrypted user credentials. WLS-VE does not provide an interactive prompt, so you cannot specify these credentials when the WLS-VE instance is booting.

Click **Next**.

4. On the **SSH information** window, select whether SSH should be on by default or not, and if used, then select the type of authentication that should be configured for the ISO image:
 - **off** (the default)
 - **Public Key Authentication** – in the **SSH Public Key** field, specify the location of the SSH Public Key (which should be in a file in OpenSSH-format). Only RSA format is supported, so the public key string must begin with `ssh-rsa`.
 - **Password Authentication** – in the **SSH Password** field, specify the password.

Note: When enabled, the same authentication will be installed on each virtual server. If you want to have different authentication keys or passwords on different virtual servers, you have to manually make those changes post-installation.

Click **Next**.

Caution: If the SSH private key is somehow lost, there is no way to exchange the old key that is installed on the WLS-VE instance. For a workaround to this situation, see [Installing a Real Password In Addition to a Public Key](#).

5. On the **Select servers to virtualize** window, define which elements of the domain that you want to be virtualized. For example, you may want to virtualize all the servers in the domain or just the Managed Servers.
 - a. Select the tab for the server instance in the domain that you want to virtualize, and then specify the properties described in [Table 5-2](#):

Table 5-2 Create VE Server Properties

Property	Description
Convert to WLS-VE	Select this check box to enable the virtual server selection properties on the tab.
Server name	Displays the selected WLS name in the specified domain.
Server type	Displays the type of server, either an Administration Server or a Managed Server.
LVM name	The name for the WLS-VE instance.
VM size (MB)	<p>The amount of memory allocated to the WLS-VE instance.</p> <ul style="list-style-type: none"> • Default: 1,024 MB • Minimum: 500 MB • Maximum: 3700 MB
Disk size (MB)	<p>The disk size allocated to the WLS-VE instance.</p> <ul style="list-style-type: none"> • Default: 500 MB • Minimum: 200 MB • Maximum: dependent on server capacity
Number of processors	The number of CPUs to use (1,2, or 4).
Log Receiver address	A hostname or IP address of a remote syslog receiver.

Table 5-2 Create VE Server Properties (Continued)

Property	Description
Network configuration	<p>The network configuration defaults to DHCP. If DHCP is selected for the Administration Server, then no IP address will be set in the generated start scripts and LVM will query the DHCP server to obtain an IP address. If the IP address's time-to-live has expired, then LVM can get a new IP address.</p> <p>To use a different network configuration, click the Change button to specify other options:</p> <ul style="list-style-type: none"> • Network Configuration Type: Select either <code>STATIC_IP_AND_DHCP</code> or <code>STATIC_IP</code>. <p>Note: To avoid an invalid IP address in the start script files, BEA recommends that the Administration Server uses a static IP address. If DHCP is used, your network administrator must guarantee that the Administration Server gets a fixed IP from the DHCP server.</p> • IP-address: If using <code>STATIC_IP_AND_DHCP</code> or <code>STATIC_IP</code>, the IP address this LVM should use. If left unset, the LVM uses DHCP to dynamically obtain an IP address. • Network-mask: 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 <code>255.255.255.0</code>. • Gateway: The octet (<code>###.###.###.###</code>) 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. 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 <code>255.255.255.0</code> and the static IP is <code>172.23.80.102</code>, then the default gateway is <code>172.23.80.1</code>. If the netmask is <code>255.255.0.0</code> and the static IP address is the same (<code>172.23.80.102</code>), then the gateway is <code>172.23.0.1</code>. • DNS Servers: The DNS server the LVM should use. • Domain name: The network domain name for the LVM instance. <p>Click OK when finished.</p>
VMware Compute Resource	Select the target ESX server within the VMware Datacenter.
VMware ResourcePool	Select the VMware resource pool on the ESX server to use.
VMware VM Datastore	Select the VMware datastore to store the virtual server on.

Table 5-2 Create VE Server Properties (Continued)

Property	Description
VMware Network	Select the VMware network to use.
WebLogic ISO file	<p>Define the local storage on the ESX server where you copied the WLS-VE ISO image, making sure to use the following naming convention (including the brackets):</p> <p><code>[storage name] path/*.iso</code></p> <p>Note: The P2V utility cannot verify the contents of a remote ISO image, so you must verify that you are selecting the right ISO-image. Further, the P2V utility does not upload a local ISO image to the VI environment; therefore, it is up to you to make sure the appropriate ISO images have been uploaded before the utility is run.</p>

- b. Click the tab for other server instances in the domain that you want to virtualize and repeat the property selection process.
 - c. Click **Next** when finished.
6. The **Confirm server generation** window lists the virtual servers that will be generated. Click **Next** to start the process.

The **Creating servers and generating scripts** window displays the progress during the domain creation process. At this point, the ISO is being created, and then an SSH session is opened with the VMware server and the virtual servers and their corresponding start scripts are created.

When successful, you should see a message similar to the following:

```

Creating LVM WLS-AdminServer for WL server AdminServer
Started LVM. Waiting for SSH to start
Connecting to SSH server at 172.23.80.102 on port 22
Copying server domain files and patches
Successfully created LVM WLS-AdminServer on esx.bea.com in demo_2
resource pool

Creating LVM WLS_MS1 for WL server MS1
Started LVM. Waiting for SSH to start
Connecting to SSH server at 172.23.81.105 on port 22
Copying server domain files and patches
Successfully created LVM WLS_MS1 on esx.bea.com in demo_2 resource pool

```

Generating scripts for starting WebLogic servers
Generated scripts in C:\wls1001\user_projects\domains\dev_domain\virtual

7. Click **Done** to exit the utility.
8. Use the VMware VI Client to verify that the new WLS-VE instance is listed in the Host and Clusters list in the left navigation pane.

What's Next?

As shown in [Roadmap for Configuring a WLS-VE Domain](#), the next step in the process is to start new virtual servers using the WLS-VE start scripts generated by the P2V utility, as described in [Run the P2V-Generated Start Scripts](#):

Run the P2V-Generated Start Scripts

As shown in [Roadmap for Configuring a WLS-VE Domain](#), the next step to complete the configuration process is to run the server start scripts in the /virtual directory generated by the P2V utility, as shown in [Table 5-3](#):

Table 5-3 P2V-generated Start Scripts

Platform	Start Script Names
Windows	<i>DOMAIN_NAME</i> \virtual\ <ul style="list-style-type: none">• <i>startservername.cmd</i>• <i>commonVEStart.cmd</i>
Linux	<i>DOMAIN_NAME</i> /virtual/ <ul style="list-style-type: none">• <i>startservername.sh</i>• <i>commonVEStart.sh</i>

where *DOMAIN_NAME* is the directory name of the domain that was selected for virtualization, typically *BEA_HOME\user_projects\domains\DOMAIN_NAME*.

Initially, the scripts must be run from a command-line to complete the P2V configuration process, as follows:

1. Open a command-line shell and navigate to the *DOMAIN_NAME/virtual* directory containing the start scripts.

2. Execute the start command by entering the name of the server start script at the prompt. Sample start script commands are shown in [Table 5-4](#).

Table 5-4 Running the Start Scripts

To . . .	Enter . . .
Start a WLS-VE Administration Server	<ul style="list-style-type: none"> • Windows — <code>startAdminServerName.cmd</code> • Linux — <code>startAdminServerName.sh</code>
Start a WLS-VE Managed Server	<ul style="list-style-type: none"> • Windows — <code>startManagedServerName.cmd</code> • Linux — <code>startManagedServerName.sh</code>

Tip: For domains with Managed Servers, start the Administration Server first, followed by the Managed Servers. This way the Managed Servers will obtain their domain configuration from the Administration Server.

The startup scripts also call the `commonVEStart.cmd` startup script in the `/virtual` directory. Therefore, when you start a WLS-VE instance, LiquidVM boots (in active mode) and simultaneously boots the associated server.

Note: 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 [Task 1: Start LiquidVM in Passive Mode](#).

For more information about starting and stopping WLS-VE instances after the initial startup, see [Starting and Stopping WLS-VE](#).

What's Next?

As shown in [Roadmap for Configuring a WLS-VE Domain](#), the next step, if necessary, is to copy domain-related files to the local disk of the LVM instance on the ESX server using the SSH service provided by the LVM. For details, see [Copying Domain Artifacts Using the LiquidVM SSH Service](#).

Copying Domain Artifacts Using the LiquidVM SSH Service

After you have created and started the new WLS-VE instances, you can copy any application files not stored in the domain directory (that is, referenced applications) to the instances using the SSH

service provided with LiquidVM. This may be necessary because the P2V utility does not handle out-of-domain dependencies, such as referenced libraries.

Note: BEA recommends that you follow these steps for each of the virtual servers that you created in your domain.

For a description of all the P2V-generated start script parameters, see [The P2V-Generated Start Script Properties](#).

To copy files to a virtual server using the LiquidVM SSH service, follow these tasks.

Task 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 `LVM_startMode=passive` option in the start script arguments and start LiquidVM as described in the following steps.

1. Modify the P2V-generated start script(s) for the LVM instance as follows:

- `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.

Note: If `LVM_SSH=on`, the SSH server will be available in `active` mode as well as `passive` mode.

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:

- Navigate to the `DOMAIN_NAME\virtual\directory`, where `DOMAIN_NAME` is the name of the directory in which you located the domain, typically `BEA_HOME\user_projects\domains\DOMAIN_NAME`.
- Execute the start command by entering the name of the server start script at the prompt. For example, if your start scripts are named `startADMIN.sh/cmd` or `startManaged_nn.sh/cmd`, as described in [Run the P2V-Generated Start Scripts](#), you enter the name of that start script at the prompt:

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

or

`startManaged_nn.sh (.cmd on Windows)`

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

```
LVM_INFO:
C:\wls1001\user_projects\domains\dev_domain\virtual\AdminServer_lvm_vmware
.info

LVM_ARGS: name=WLS-AdminServer cwd=/domain cpus=1 memory=1024 diskSize=500
startMode=passive ssh=off ip=123.45.67.890 netMask=255.255.248.0
gateway=172.18.128.1 dns=10.40.0.86 networkDomainName=bea.com

CLASSPATH:
:/bea/patch_wls1001/profiles/default/sys_manifest_classpath/weblogic_patch
.jar:/appliance/java/lib/tools.jar:/appliance/bea/weblogic/server/lib/weblo
gic_sp.jar:/appliance/bea/weblogic/server/lib/weblogic.jar:/appliance/bea
/modules/features/weblogic.server.modules_10.0.1.0.jar:/appliance/bea/modu
les/features/com.bea.cie.common-plugin.launch_2.1.2.0.jar:/appliance/bea/w
eblogic/server/lib/webservices.jar:/appliance/bea/modules/org.apache.ant_1
.6.5/lib/ant-all.jar:/appliance/bea/modules/net.sf.antcontrib_1.0b2.0/lib/
ant-contrib.jar:/appliance/bea/weblogic/common/eval/pointbase/lib/pbclient
51.jar:/appliance/bea/weblogic/server/lib/xqrl.jar:

JAVA_PROPERTIES: -Dbea.home=/bea -Dweblogic.Name=AdminServer
-Dweblogic.ProductionModeEnabled=true -Dweblogic.management.discover=true
startMode=passive was requested but for that to work ssh must also be turned
on.
Turning ssh on!
Starting WLS-AdminServer connect...lookup...configure...start...booting...
Initial log from LiquidVM instance follows:
-----
Baremetal hostname: "LVM14.BEA.COM" IP address: 123.45.67.890
LiquidVM R1.2.4.0-95617 (BareMetal 4.1.4.0-95612-24)
-----
See the console log-file for further data
LiquidVM IP-address: 123.45.67.890
```

Note: If the IP was obtained via DHCP, make 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 [Task 2: Copy the Files to the WLS-VE Instance](#).

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
```

Task 2: Copy the Files to the WLS-VE Instance

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 necessary files. Note the following:

- You need to specify the IP address of the LiquidVM instance noted in [Task 1: Start LiquidVM in Passive Mode](#) 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 necessary files as follows:
 - When copying files that are outside of local `DOMAIN_NAME` directory, you can copy them anywhere in LVM file system. The only restrictions are that you *cannot* write under the `/appliance` directory and you *should not* write under system directories (such as `/etc`). You can write to directories such as `/apps` or `/usr/apps`. These directories are not present by default, but can be created using a `scp` or `sftp` session.
 - When copying domain-related 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 WLS-VE instance; the files must be copied directly under the `/domain` directory.) By default, `/domain` is defined as the current working directory within the LiquidVM instance.

Task 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 [Working with WLS-VE Using the VMware VI Client](#).

Task 4: Re-start the WLS-VE Instance

1. In your local file system on the launcher machine, navigate to `DOMAIN_NAME/virtual/` directory and open the start script that you edited in [Task 1: Start LiquidVM in Passive Mode](#).
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: When SSH is enabled, you can remove the `LVM_SSH` line since this parameter does not control active or passive mode.

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

```
startAdmin.sh (.cmd on Windows)
or
startManaged_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**.

Note: The **Powered On** state does not always indicate active vs. passive mode. For active mode, the WLS main class will be executed, so the VMware console will show messages from WLS.

What's Next?

As shown in [Roadmap for Configuring a WLS-VE Domain](#), 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. See [Starting and Stopping WLS-VE](#).

Configuring a Virtual Domain to Access an NFS Server

After using the P2V utility to convert a physical domain to a virtualized domain, you must manually configure any NFS connections. Before doing so, make sure that both your ordinary OS and your WLS-VE instance have access to the same NFS-share.

For details about generating a WebLogic 9.2 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 WLS-VE instance instead.

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 [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 WLS 10.0](#) 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 5-1](#) shows a snippet of the Administration Server start-up script, `startAdminServername.sh/cmd`, with `JAVA_OPTIONS=` highlighted.

Listing 5-1 `startAdminServername.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` or the Administration Console
- Redeploying applications in a production environment
- Managing deployed applications

These tasks are detailed in [Deploying Applications on BEA WebLogic Server 10.0](#).

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, BEA strongly recommends that you refer to the [VMware Infrastructure Documentation](#).

This section contains information on these topics:

- [Setting Up VMware and Enabling SSL](#)
- [Starting WLS-VE Instances](#)
- [Editing VM Properties](#)
- [Pausing a VM](#)
- [Working with the Console Tab](#)

Setting Up VMware and Enabling SSL

For information about installing VMware VirtualCenter, see the [VMware Infrastructure Documentation](#).

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 WLS-VE Instances

You can start an existing WLS-VE instance from the VI Client, as described in [Starting and Stopping WLS-VE](#).

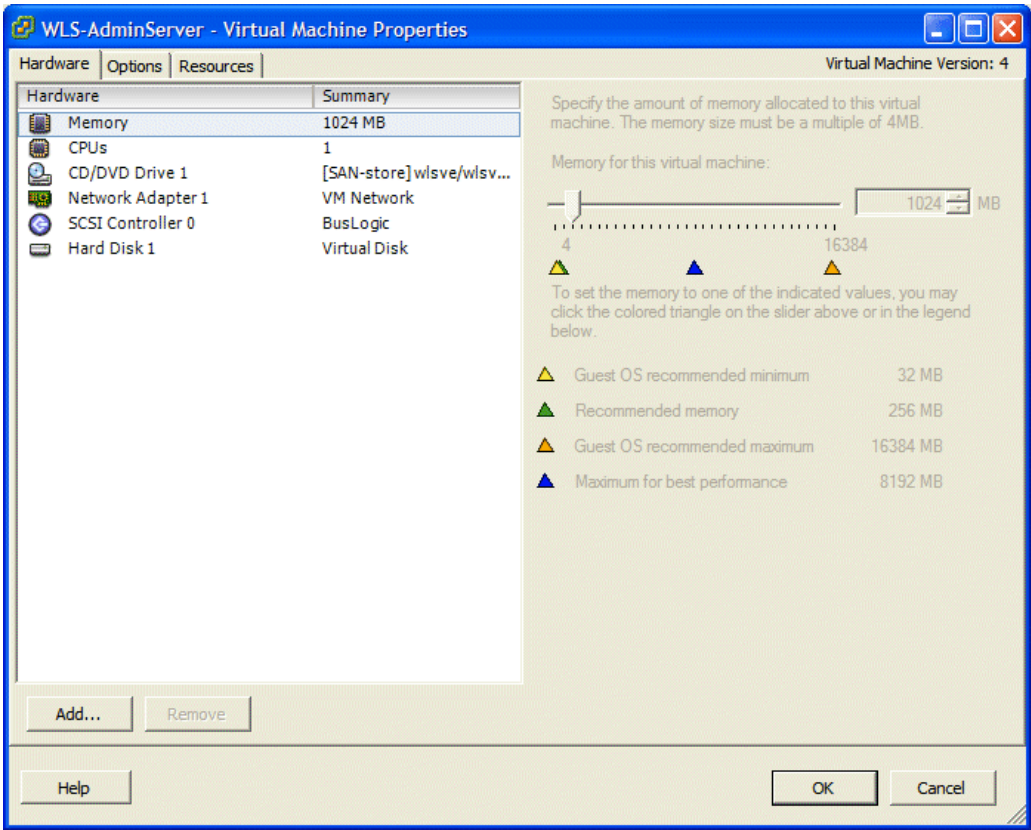
Editing VM Properties

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 6-1](#)).

Figure 6-1 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 6-1](#). Unless otherwise specified, you can edit any of the values on these tabs as long as those properties are enabled for your VM.

Table 6-1 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: Number of virtual processors used by this machine <p>These properties cannot be changed:</p> <ul style="list-style-type: none"> • CD/DVD Drive: Path and filename of the machine's <code>.iso</code> file (there can be more than one drive) • Network Adapter: 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](#).

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 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 6-2](#) to perform certain functions when the Console is open:

Table 6-2 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 [The P2V-Generated Start Script Properties](#).

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 WebLogic domain.

- [Starting WLS-VE](#)
- [Starting the Administration Console](#)
- [Stopping WLS-VE Instances](#)

Starting WLS-VE

When you convert a physical WebLogic domain to a WLS-VE domain using the P2V utility, scripts to start WLS-VE instances on the virtual machine are created in the `/virtual` directory generated by the P2V utility, as shown in [Table 7-1](#):

Table 7-1 P2V-generated Start Scripts

Platform	Start Script Names
Windows	<i>DOMAIN_NAME</i> \virtual\ <ul style="list-style-type: none"><code>startservername.cmd</code><code>commonVEStart.cmd</code>
Linux	<i>DOMAIN_NAME</i> /virtual/ <ul style="list-style-type: none"><code>startservername.sh</code><code>commonVEStart.sh</code>

where *DOMAIN_NAME* is the directory name of the domain that was selected for virtualization, typically `BEA_HOME\user_projects\domains\DOMAIN_NAME`.

The startup scripts also call the `commonVEStart.cmd` startup script in the `/virtual` directory. Therefore, when you start a WLS-VE instance, LiquidVM boots (in active mode) and simultaneously boots the associated server. For more information about LiquidVM, see [Configuring LiquidVM Connection Parameters](#).

Note: 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 [“Task 1: Start LiquidVM in Passive Mode” on page 5-12](#).

Initially, the scripts must be run from a command-line to complete the P2V configuration process. After you have started a WLS-VE instance from the command-line at least once, you can restart it 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.

Starting WLS-VE from a Command-line

To start the WLS-VE instances and the LiquidVM launcher, follow these steps:


1. Open a command-line shell and navigate to the `DOMAIN_NAME/virtual` directory containing the start scripts.
2. Execute the start command by entering the name of the server start script at the prompt. Sample start script commands are shown in [Table 7-2](#).

Table 7-2 Running the Start Scripts

To ...	Enter ...
Start a WLS-VE Administration Server	<ul style="list-style-type: none"> • Windows — <code>startAdminServerName.cmd</code> • Linux — <code>startAdminServerName.sh</code>
Start a WLS-VE Managed Server	<ul style="list-style-type: none"> • Windows — <code>startManagedServerName.cmd</code> • Linux — <code>startManagedServerName.sh</code>

Starting WLS-VE from the VMware VI Client

After you have started a virtual server from the command-line at least once, you can then start it from the VMware VI Client:

1. Log into the Virtual Infrastructure Client.
2. Use the Inventory panel to select the WLS-VE instance 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 WLS-VE instance 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 you have started a virtual Administration Server from the command-line at least once, you can start the Administration Console on the launcher machine to remotely administer and shut-down its managed 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 Instances

Stopping WLS-VE instances should be carefully considered as the ripple effects of an improper shutdown can cause unexpected results, such as losing any underlying server connections, with no guarantee that these connections will be restored when the instance is restarted. For this reason, **stopping** WLS-VE instances **using VirtualCenter is not recommended**. If you must stop WLS-VE, use 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 the **Power Off** button  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](#)
- [Accessing WLS-VE Log Files](#)

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 P2V Domain Conversion utility.

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)
- Store 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](#). 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 [Copying Domain Artifacts Using the LiquidVM SSH Service](#).
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](#).

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 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](#)
- [Securing WLS](#)
- [Securing the VMware VirtualCenter](#)

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,” v1.0 [Installation and Configuration Guide](#) for additional NFS security guidelines.

- LiquidVM provides a 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.

Caution: If the SSH private key is somehow lost, there is no way to exchange the public key that is installed on the WLS-VE instance. This means you will not be able to remotely access the content on the WLS-VE instance because you won't be able to connect via `ssh` or `sftp`. Therefore, you may also want to create a secure *real* password as well, so that you can still log in and access the files on the local disk, as explained in [Installing a Real Password In Addition to a Public Key](#).

- SSH is disabled by default. If SSH is enabled, be sure to install a public key or set a secure real 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](#).

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 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](#)
- [Comparing OS and LiquidVM Kernel Tuning](#)

Tuning the LiquidVM Kernel Startup Options

The LiquidVM startup options described in [Table 10-1](#) can have an effect on performance and can be tuned as required.

Table 10-1 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 10-2](#) 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 10-2 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 OSes 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 10-3](#) 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 10-3 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 10-2 .

Table 10-3 Tunable OS Parameters in the LiquidVM Kernel (Continued)

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

Table 10-3 Tunable OS Parameters in the LiquidVM Kernel (Continued)

OS Parameter	Relevance to LiquidVM Kernel
Linux Tuning Parameters	
/sbin/ifconfig lo mtu	N/A.
kernel.msgmni	N/A. SystemV IPC configuration.
kernel.sem	N/A. SystemV IPC configuration.
kernel.shmmax	N/A. SystemV IPC configuration.
fs.file-max	N/A. LiquidVM Kernel has no such limitation; memory is the only restriction.
net.ipv4.tcp_max_syn_backlog	N/A. LiquidVM Kernel has no such restriction, and abides by the user's accept backlog setting (<code>listen(2)</code>)
HP-UX Tuning Parameters	
tcp_conn_req_max	N/A.
tcp_xmit_hiwater_def	LiquidVM Kernel boot option "netSndBufSz", else dynamically guided by memory pressure.
tcp_ip_abort_interval	LiquidVM Kernel boot option "netRcvBufSz"
tcp_rexmit_interval_initial	N/A
tcp_keepalive_interval	LiquidVM Kernel boot option "netKeepAlive"
Windows Tuning Parameters	
MaxUserPort	N/A
TcpTimedWaitDelay	N/A. Dynamically estimate effective MSL for each socket.

Diagnostics and Troubleshooting

This section describes how to deal with issues that might occur in both WLS and the JVM, as well as issues that are specific WLS-VE. 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](#)
- [Handling Suspend Files](#)
- [Displaying Version Information](#)
- [Reporting a Problem to BEA Support](#)

Troubleshooting WLS-VE

This section provides information you will find helpful in solving issues that might occur with WLS-VE. Generally, you handle WLS and LiquidVM (the BEA JRockit component) issues 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 issues and collect relevant information about runtime activity.

Diagnosing WLS-VE Issues

Issues 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](#)
- [Server Shuts Down Soon After Startup](#)
- [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 WLS-VE instance was created, VMware could not find the ISO image needed to boot up WLS-VE.

Solution:

Check that you have uploaded the WLS-VE ISO image to the ESX server.

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

(`server1001ve12\tools\lvm_configwizard.cmd` or `.sh`), as described in [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 `wlsve1001.iso` is available in your datastore.

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 10.0 MP1 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 did not 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. WLS-VE does not support normal keyboard input, so you cannot enter a user name 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 *Managing Server Startup and Shutdown*) or add a user name 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 are not, 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 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 is using your IP address. Because finding out who that person is might be difficult, contact your system administrator to obtain another IP address.

Diagnosing WLS Issues

WLS-VE can encounter the same type of server-related issues that can occur when running non-virtualized WLS. This section provides an overview of the kinds of WLS issues you should watch for when running WLS-VE.

Performance Issues

Often, a problem with WLS is the result of poor tuning. For example, pool sizes (such as pools for JDBC connections, Stateless Session EJBs, and MDBs) that do not 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 tuning and performance guidelines in [WebLogic Server Performance and Tuning](#). Specific tuning guidelines for WLS-VE are provided in [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 issues can affect the performance of WLS. These issues can occur for many reasons, including versioning errors, multicast addressing issues, errors or misspellings in start-up commands, and even a poorly-tuned memory management systems. You can find guidelines for troubleshooting cluster issues in [Troubleshooting Common Problems](#) in *Using WebLogic Server Clusters*.

Other WLS Issues

Other, non-specific issues can also occur with WLS. When these issues 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 WLS error conditions.

Diagnosing LiquidVM Issues

Issues that do not originate with WLS may occur in LiquidVM and are typical to the kinds of issues you might encounter in non-virtualized JVMs. Issues such as these are documented in the [BEA JRockit Diagnostics Guide](#) (BEA JRockit is the JVM component of LiquidVM). This topic 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 issues 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 OSes 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 resuming execution, because you might lose critical information that would be

helpful in diagnosing the issues causing the crash. You should also be aware that suspend files are large and so it might not be efficient to copy from the ESX server.

Displaying Version Information

A critical piece of information that Support will need to help diagnose any issues you report to them is the version number. You can find this number in the file `WLSVE_VERSION`, which is located in the `server1001ve12/` directory.

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

```
10.0.0.1ve1.2-b16
```

Reporting a Problem to BEA Support

If your machine crashing, running slowly, or returning unpredictable results you may have a problem with WLS-VE. 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.

If you determine that you need to open a case with BEA support, 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.

Verify That You Are 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.

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 guidelines provided in [Troubleshooting WLS-VE](#). 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.

Adding Managed Servers to a WLS-VE Domain

The following topics describe how to create and add Managed Server instances to an existing WLS-VE domain created with the P2V utility.

- [Steps for Adding a Managed Server to WLS-VE Domain](#)
- [The P2V-Generated Start Script Properties](#)

Steps for Adding a Managed Server to WLS-VE Domain

After using the P2V utility to convert a physical domain to a virtual domain to the hypervisor environment, you may need to add server instances to the WLS-VE domain. This section explains how to add a Managed Server to a WLS-VE domain.

1. Use the VMware VI Client to start the Administration Server in the WLS-VE domain, as described in [Starting WLS-VE from the VMware VI Client](#).
2. Start the WLS administration console, as described in [Starting the Administration Console](#).
3. Create a new Managed Server but do not attempt to start it.

Note: If you try to start the Managed Server you will get the following error message:
“The server does not have a machine associated with it. All of the servers selected are currently in a state which is incompatible with this operation or are not associated with a running node manager or you are not authorized to perform the action requested.”

4. On the launcher system, change to the virtual directory of the domain that has been transferred. For example:

```
bea_home/user_projects/domains/testdomain/virtual
```

5. Make a copy of the LVM information file for the newly created Managed Server. For example:

```
cp AdminServer_lvm_vmware.info to Managed1Server_lvm_vmware.info
```

6. Using an existing WLS-VE start script as a template, create a new startup script for the Managed Server, as shown in [Listing 12-1](#).

Listing 12-1

```

export WL_SERVER_NAME=managed1
export ADMIN_URL=http://100.90.80.201:7001
export LVM_NAME=managed1
export LVM_CPUS=1
export LVM_MEMORY=1024
export LVM_DISKSIZE=500
export LVM_IP_ADDRESS=100.90.80.202
export LVM_NETMASK=255.255.248.0
export LVM_GATEWAY=100.90.128.1
export LVM_DNS_SERVERS=10.10.0.86
export LVM_DOMAIN_NAME=domain.com
export LVM_SSH=on
export LVM_SSH_PUBLIC_KEY=
export LVM_SSH_UNSAFE_PASSWORD=
export LVM_SYSLOG_RECEIVER=
export LVM_START_MODE=passive
export
LVM_INFO=/usr/local/boa_home/user_projects/domains/testdomain/virtual/M
anaged1Server_lvm_vmware.info
/usr/local/boa_home/user_projects/domains/testdomain9/virtual/commonVES
tart.sh

```

Note: If the LVM_SSH_PUBLIC_KEY parameter is set to the public_key file, the LiquidVM launcher will install the public_key on LVM authorized_keys.

For a description of all the start script properties, see [The P2V-Generated Start Script Properties](#).

7. Run the newly created startup script to start the Managed Server in passive mode (LVM_START_MODE=passive, as shown in [Listing 12-1](#)), as follows:


```
> startManaged1Server.sh (or .cmd on Windows)
```
8. Use the LiquidVM SSH service to copy the following files to the local disk for the Managed Server:
 - boot.properties
 - security .ldift files
 - SerializedSystemIni.dat
 - BEA patches

Note: If you are using public key authentication for SSH, and if the `LVM_SSH_PUBLIC_KEY` parameter was not set to the `public_key` file, then you also need to copy your public key file to the `/etc/ssh/authorized_keys` directory.

After copying is complete, the local disk should contain the following files:

```
/domain/servers/managed1/security/boot.properties  
/domain/security/*.ldift  
/domain/security/SerializedSystemIni.dat  
/bea/patch_*
```

For more information on using the SSH service, see [Using the LiquidVM SSH Service](#).

9. Power off the Managed Server.
10. Edit the startup script to change from passive mode to active mode set, as follows:
`LVM_START_MODE=active` or `'LVM_START_MODE='`
11. Run the startup script again to start the Managed Server in active mode.
`> startManaged1Server.sh` (or `.cmd` on Windows)

The P2V-Generated Start Script Properties

[Table 12-1](#) describes all the properties in the Administration Server and Managed Server start scripts generated by the P2V Domain Conversion utility.

Table 12-1 WLS-VE Start-up Options

Property	Description
WL_SERVER_NAME	The name of the WLS instance in the domain.
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_NAME Required for all servers.	The name of the virtual server to be started. In VMware, SERVER_NAME is called WLS-SERVER_NAME. This property corresponds to the server name that you specified when you created the domain.
LVM_CPUS	The number of CPUs to use (1,2, or 4).
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.</p>
LVM_IP_ADDRESS	The IP address this LVM should use. If left unset, the LVM uses DHCP to dynamically obtain an IP address
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	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. 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.

Table 12-1 WLS-VE Start-up Options (Continued)

Property	Description
LVM_DNS_SERVER	The DNS server the LVM should use.
LVM_DOMAIN_NAME	The network domain name for the LVM instance.
LVM_SSH	The SSH service mode to be used (<code>on</code> <code>off</code>). The default is <code>off</code> . If you set <code>LVM_START_MODE</code> 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_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_SYSLOG_RECEIVER	A hostname or IP address of a remote syslog receiver.
LVM_START_MODE Required if starting the server in passive mode.	<p>The mode in which to start the LiquidVM instance (<code>active</code> <code>passive</code>).</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>
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 .