



everRun[®]

Product Overview

Release 6.0
October 2010

MARATHON

The Application Availability Experts™

NOTICE

Marathon Technologies Corporation reserves the right to make improvements to this document and the product it describes at any time and without further notice.

COPYRIGHT

© Marathon Technologies Corporation 2008, 2009-2010. All rights reserved.

This document is copyrighted, and all rights are reserved. No part of the document or the products it describes may be reproduced by any means or in any form, without prior consent in writing from Marathon Technologies Corporation. Printed in the U.S.A.

Marathon everRun products are protected by one or more of the following patents:

U.S. Patent Numbers: 5,600,784; 5,615,403; 5,787,485; 5,790,397; 5,896,523; 5,956,474; 5,983,371; 6,038,685; 6,205,565; 6,279,119; 6,473,869; 6,728,898, 7,373,545.

European Patent Numbers: EP0731945; EP0974912; EP0986784; EP0993633; EP1000397; EP1000404; EP1029267; EP1496434; GB2392536; Japanese Patent Numbers: 3679412; 4166939; 4264136; 4472995. Other patents pending.

Marathon Assured Availability, ComputeThru, SplitSite, everRun, and the Marathon logos are registered trademarks or trademarks of Marathon Technologies Corporation. Copyright © 2008, 2009-2010 Marathon Technologies Corporation. All rights reserved.

SOFTWARE COPYRIGHT NOTICE

The software described in this document is covered by the following copyright:

© Marathon Technologies Corporation. 2008, 2009-2010.

GNU PUBLIC LICENSE AND LESSER GNU PUBLIC LICENSE

The product makes use of software covered under the GNU Public License (GPL) and the Lesser GNU Public License (LGPL). For a written copy of these licenses, see `/opt/everRun/current_everRun/licenses` on any XenServer host on which everRun VM is installed. Marathon will make source components available upon request, via email to support@marathontechnologies.com, as required under the terms of the GPL.

TRADEMARK NOTICE

Microsoft® and Windows® are registered trademarks of Microsoft Corporation in the United States and/or other countries; Xen® and Citrix® are registered trademarks and XenServer™ and XenCenter™ are trademarks of Citrix Systems, Inc.; Java® is a registered trademark of Sun Microsystems; Linux® is a registered trademark of Linus Torvalds. Intel® is a registered trademark of Intel Corporation.

All other brands and product names are trademarks of their respective companies or organizations.

SOFTWARE REVISION

The revision of the software that this document supports is Release 6.0.

Marathon Technologies Corporation
295 Foster Street, Littleton, MA 01460
(978) 489.1100 or (888) 682.1142
www.marathontechnologies.com

Contents

Chapter 1 —Welcome to everRun Availability

Overview	1
Why This Approach?	1
Problems with Existing Solutions	2
The everRun Solution	2
FT-Class Availability for Virtual Environments	3
Simple, Centralized Management	4

Chapter 2 —System Setup and Operation

System Setup	5
Configuring the XenServer Hosts and Resource Pool	5
Installing everRun	6
Creating Virtual Machines (VMs) and Adding Applications	6
Integrated Availability Levels	7
Choosing System-Level Fault Tolerance (Level 3)	7
Choosing Component-Level Fault Tolerance (Level 2)	7
Choosing Basic Failover (Level 1)	8
Managing Your Protected Applications	8

Chapter 3 —everRun System Management

At-a-Glance Resource Selection	9
Wizard-Driven Protect/Unprotect Operations	10
Protect Wizard	10

Unprotect Wizard	10
Easy-to-Monitor Operations	12
Accessing Pool and Resource Information	12
Icon-Based Status	13
Easy-to-Access Management Functions	14
Migrating Protected VMs from Host to Host	15
Replacing Hosts or Components	15
Creating Backup Images Using everRun Snapshots	15
Monitoring Applications with System Alerts	15

Chapter 4 —Where to Learn More

Marathon Documentation	17
Information on the Marathon Web Site	19
Citrix XenServer Documentation	19
Useful Web Addresses	20
Reference Books	20

Welcome to everRun Availability

1



Overview

The everRun[®] system represents a new class of availability solutions. This breakthrough product combines the power of Marathon everRun fault tolerant-class availability software with robust server virtualization delivered by Citrix[®] XenServer[™] software.

The new generation of everRun products provide simple, nondisruptive, and cost-effective availability to protect either 32-bit or 64-bit Windows environments from the costly effects of downtime. With everRun availability protection, you can:

- Protect one or more applications running on different hosts in one or more resource pools, and select an appropriate level of protection for each application.
- Manage multiple resource pools simultaneously from a single instance of the everRun graphical user interface.
- Choose to configure workloads with one or more virtual CPUs.

This document describes how everRun enables you to attain the highest levels of virtual server availability in the industry.

Why everRun?

As server virtualization has proven its ability to reduce corporate costs and enhance the scheduling of planned downtime, IT managers are looking for increased protection from unplanned downtime in virtual environments.

everRun provides the required application availability, using techniques that overcome the shortcomings of traditional approaches such as clustering or failover.

Problems with Existing Solutions

Existing availability solutions suffer from a variety of issues:

- Clustering and failover require manual configuration, setup, scripting, and testing to define the appropriate actions to take in case of failure. This adds cost and administrative complexity, which can introduce errors and availability issues.
- Heartbeat pings used by many failover solutions are unable to reliably determine the health of a virtual machine and may not distinguish between I/O path failures, server failures, or lack of system resources. In some cases, these limitations can produce unnecessary or false failovers. In other cases, they may fail to identify outages in discrete storage or network device failures — which means that a needed failover may not occur.
- The failover process itself is far from certain. If the administrator has not configured and maintained the standby environment appropriately, the standby may be unable to support the application or virtual machine after a failover. This is called a “failed failover.”

The everRun Solution

Building on the enterprise-class XenServer virtualization engine, everRun goes beyond clustering or failover to offer new levels of fault tolerant-class availability.

Table 1 Levels of everRun Fault Tolerance

Fault Tolerance Level	Properties
<p>Level 3, everRun Lockstep option</p> <p>System-level fault tolerance for continuous available of mission-critical systems — eliminating transaction loss, data loss, and application downtime</p>	<ul style="list-style-type: none">• Zero downtime for any failure• Retention of application state and memory state during failures• All the benefits of Level 2
<p>Level 2, everRun component-level fault tolerance</p> <p>Eliminates downtime due to I/O component failures and guarantees recovery from system failure</p>	<ul style="list-style-type: none">• Automated fault management: policies handle system, network, and disk I/O failures without IT intervention• Assured recovery of virtual machines• Zero downtime due to I/O failures• Zero data loss• Synchronous data mirroring between hosts; no need for shared storage

Table 1 Levels of everRun Fault Tolerance

Fault Tolerance Level	Properties
	<ul style="list-style-type: none"> • Continuous, active validation of all components assures complete redundancy and recovery at all times • Comprehensive availability — including system, network, and data availability — in a single integrated solution
<p>Level 1, Citrix XenServer HA, based on core Marathon technologies</p> <p>If enabled, offers basic failover and recovery, with restart capabilities for applications in the virtual environment</p>	<ul style="list-style-type: none"> • Basic failover to another host, with resource calculation to determine the number of simultaneous host failures that can be handled • Monitoring the health of hosts within a pool • Shared-storage configuration required

FT-Class Availability for Virtual Environments

everRun transparently protects an application by creating a redundant environment for a virtual machine running across two physical servers in a resource pool.

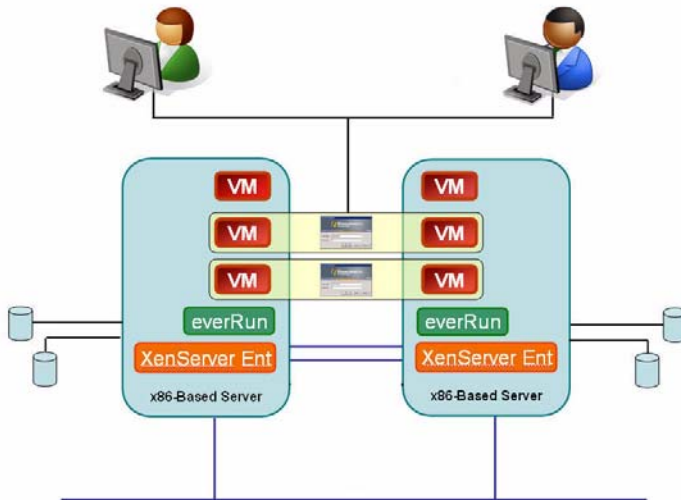


Figure 1 The two parts of a protected VM are located on separate physical servers, but everRun treats them as a single Windows environment.

The everRun management software easily converts an existing virtual machine into an everRun-protected virtual machine. By creating an identical instance of the selected

VM on a second host, it provides FT-class protection of the VM. The system administrator manages this single entity from a separate, browser-based management console, as indicated in Figure 1.

Neither the application nor the user is exposed to the redundant computing resources on the two hosts. There is only one hostname, one IP address, and one MAC address for the application. As the system administrator, you load and configure your applications on the protected VM — just as if you were loading them onto a physical server.

If a fault or failure occurs in a disk or network device, everRun automatically redirects I/O to the paired host to permit continuous operation. While there is a loss of redundancy until the failure is repaired, there is no interruption to client connectivity and no loss of data. The application continues to execute as if nothing had happened. The redundancy, fault detection, isolation, and management are completely transparent to the Windows environment and the application running within it.

Repair is equally transparent and automatic. When a failed component or XenServer host is repaired, everRun automatically incorporates the repaired components into the protected environment and restores redundancy without interrupting the application.

Simple, Centralized Management

The everRun Availability Center (eAC) management console is shown in Figure 2.

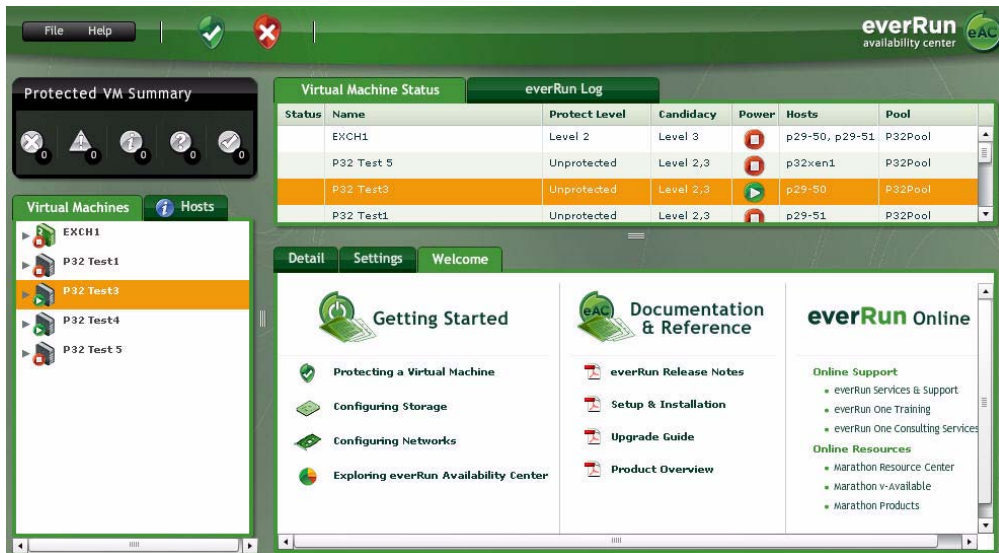


Figure 2 View and manage VMs from the everRun Availability Center.

This easy-to-use workspace reduces complex system management operations to simple pick-and-click routines. You can view and manage the VMs from a single browser-based interface, with at-a-glance insight into events and status.

This chapter describes how the everRun solution combines the reliable virtual environment of Citrix XenServer with the fault-tolerant class of availability solutions from Marathon Technologies.

System Setup

A basic configuration requires two servers that are similarly configured and outfitted with everRun-compatible processors (as determined by the Compatibility Check utility described in Table 2 of Chapter 4).

Configuring the XenServer Hosts and Resource Pool

To create a **resource pool**, you first install XenServer software on two connected **XenServer hosts**. This provides the XenServer **hypervisor** operating environment that supports virtual machines.




Figure 3 A basic everRun configuration contains resource pool of two XenServer hosts and a XenCenter management console.

Next you install a management console, called **XenCenter**, on a separate Windows-based client computer, as shown in Figure 3. XenCenter allows you to perform basic administrative tasks, such as configuring the resource pool and creating virtual machines (VMs) to host your Windows applications.

Installing everRun

To complete your system setup, you install everRun software on the two hosts where the XenServer software is running (Figure 4).




Figure 4 Add the everRun software to complete the installation.

Creating Virtual Machines (VMs) and Adding Applications

Once you have created VMs on the pool master host, you can install applications just as you would on any physical server (Figure 5). Each guest VM is a fully functional operating environment. In addition to local attached storage, everRun supports shared storage such as NFS, iSCSI, or SAN.

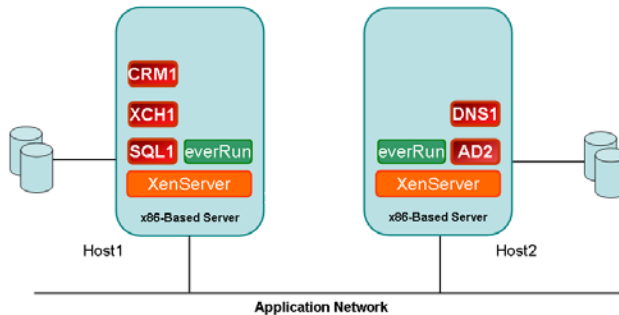


Figure 5 Each XenServer host can contain one or more VMs.

You are now ready to use everRun to create the redundancy that protects your running applications (Figure 6).

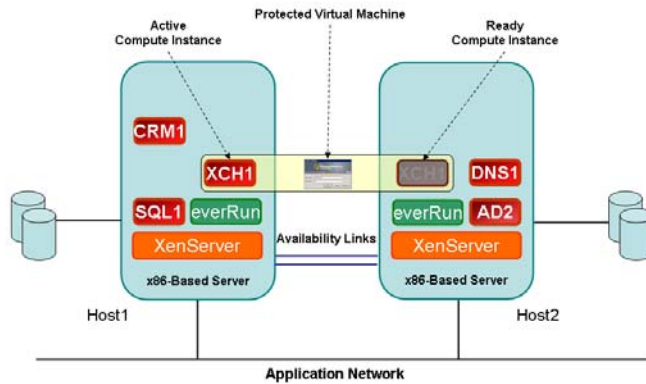


Figure 6 Mirrored everRun instances present a single view to the user, so you can manage a protected VM just as easily as you manage a single application.

Integrated Availability Levels

The tight integration of XenServer and everRun lets you choose the appropriate level of protection for each virtual machine.

Choosing System-Level Fault Tolerance (Level 3)

everRun **system-level fault tolerance** or Level 3 protection provides continuous availability in the face of component or system-wide failure, for zero downtime in mission-critical applications. Use Level 3 protection for systems that must maintain transaction state at all costs.

Choosing Component-Level Fault Tolerance (Level 2)

everRun **component-level fault tolerance** is suitable for many business-critical applications, for which data loss and downtime are extremely costly. Email, customer relationship management (CRM), enterprise resource planning (ERP), back-end databases, and financial software are all candidates for Level 2 everRun protection.

everRun Level 2 handles failure of components such as disk, storage, controllers, and network devices without interrupting the applications. As shown in Figure 6, everRun mirroring writes to storage disks on both hosts, ensuring that all transactions are saved on both systems. Should one of the storage devices fail, the system I/O is assumed by the mirrored disk on the surviving host. everRun does this automatically and transparently. The application is not aware of the failure, and it continues to operate without interruption. Users and clients are also completely unaware that a failure has occurred.

Recovery of failed storage devices is just as seamless. After you replace the failed component, everRun automatically re-mirrors the repaired component while the application is running — all without administrator intervention or user interruption.

Network failures are handled in a similar fashion. While each virtual machine in the pair has its own network interface, only one network interface is used at a time. The other interface is ready and available to take over if needed. If the active network interface should fail, all network I/O is redirected immediately and transparently to the operational network interface. Again, this is all done automatically and with no interruption of the application or the clients. You repair the problem, then allow everRun to re-enable the device for service.

Choosing Basic Failover (Level 1)

This level is suitable for applications where recovery is not absolutely critical, and where manual intervention — while not desirable — is acceptable. These may include infrastructure applications or test and development systems.

You can enable Level 1 availability when your XenServer software includes the XenServer High Availability (HA) software, and can be managed with XenServer HA alone. If you are managing VMs on multiple availability levels, you can also manage Level 1 availability from the everRun Availability Center (eAC) GUI.

Managing Your Protected Applications

Whatever level of protection you choose for each application, you can use the convenient everRun Availability Center (eAC) to monitor and manage multiple applications from all your resource pools, as described in the following chapter.

This chapter provides a quick tour of the everRun Availability Center (eAC), where you manage and monitor your everRun-protected VMs.

At-a-Glance Resource Selection

The integrated view in the eAC main workspace (Figure 7) lets you quickly identify VMs that are candidates for protection — then easily configure, monitor, and manage everRun-protected VMs, or PVMs.

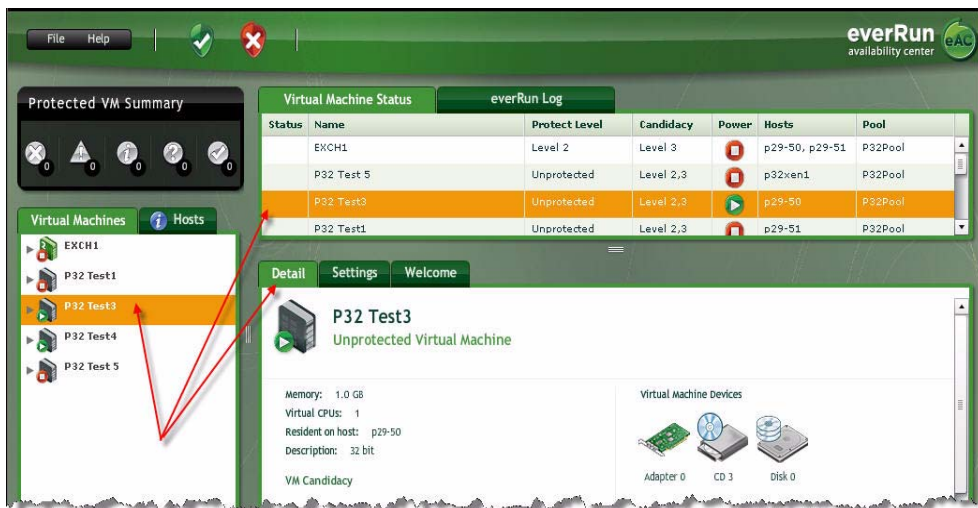


Figure 7 The eAC workspace displays resource information in easy-to-navigate tree views, tables, and configuration tabs.

When you select a resource from the tree view at left, the same resource is highlighted in the Virtual Machine Status table (upper right). Because the interface is context-sensitive, the Detail tab (lower right) shows components configured for that resource.










The Virtual Machine Status table at the top of the workspace provides an orderly list of the VMs and hosts in the pool. It clearly identifies the protection level of any protected VMs, as well as any additional candidacy level(s) available. Status icons quickly portray machine status (such as the Good  icon at the upper left); power status is indicated by the Running  icon or the Halted  icon.


Figure 8 The Status table summarizes virtual machine configuration details.

Virtual Machine Status		everRun Log				
Status	Name	Protect Level	Candidacy	Power	Hosts	Pool
	P29Test1	Level 2 - Componen			p29-50, p29-51	P29-Pool2
	p29test3	Unprotected	Level 2		p29-51	P29-Pool2
	ldev6	Unprotected			p29-50	P29-Pool2
	p29-51					P29-Pool2
	p29-50					P29-Pool2

Wizard-Driven Protect/Unprotect Operations

everRun now supports Level 3, Level 2, and Level 1 protection. Easy-to-use wizards guide you through the Protect and Unprotect processes.


Protect Wizard

You can start the Protect wizard by clicking the Protect  icon or by right-clicking a resource and choosing **Protect** from the pop-up menu.

The Protect wizard leads you quickly through the multistep process, as shown for a Level 2 PVM in Figure 9. First, you select the VM you want to protect and choose the level of availability required. Next, you choose the secondary host that will provide the redundant resources — making sure you choose a different physical server from the one that is currently hosting the VM. In the remaining screens you configure resources on the secondary host and choose storage resources for the protected VM. Finally, you review your settings and click Protect to start the automatic process.

Depending on the level of protection you choose, the steps in the protection wizard may vary.

Unprotect Wizard

The two-step Unprotect wizard is equally simple to use. To start it, you can click the Unprotect icon () or right-click a resource and choose Unprotect from the pop-up menu. Then you simply specify the target resource and click Unprotect.

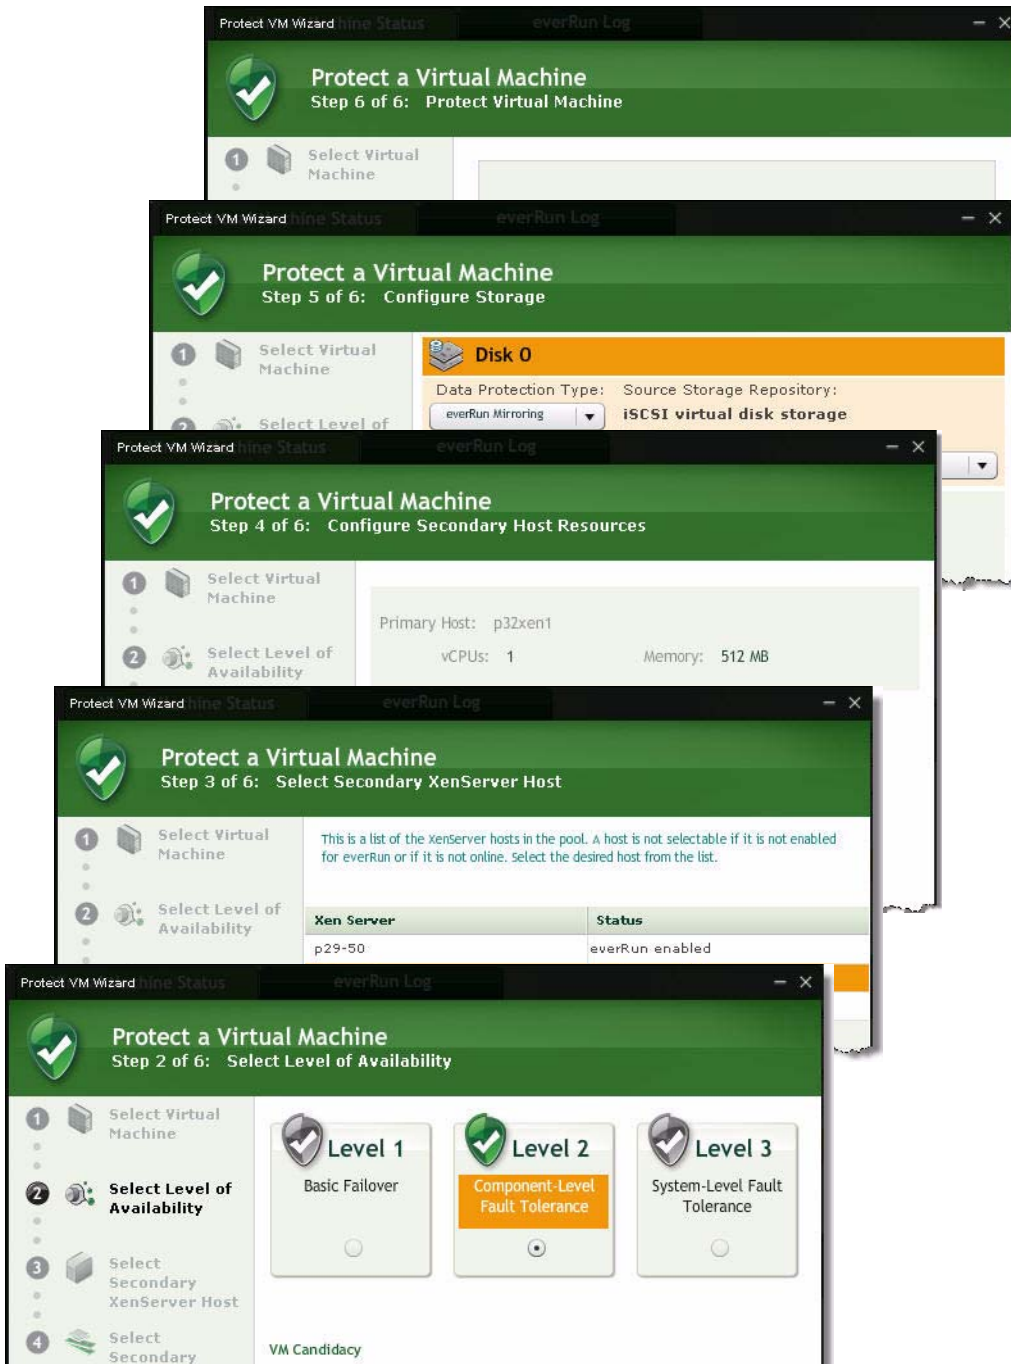


Figure 9 The Protect wizard steps you through the everRun protection process.

Easy-to-Monitor Operations

When you start a PVM that is protected with everRun mirroring, you can monitor the status of the disk mirror copy by clicking the Mirror Copy Status tab. When the copy is complete, you can view details of device status, copy rates, and elapsed time by hovering the mouse over the disk icon to display the pop-up message.

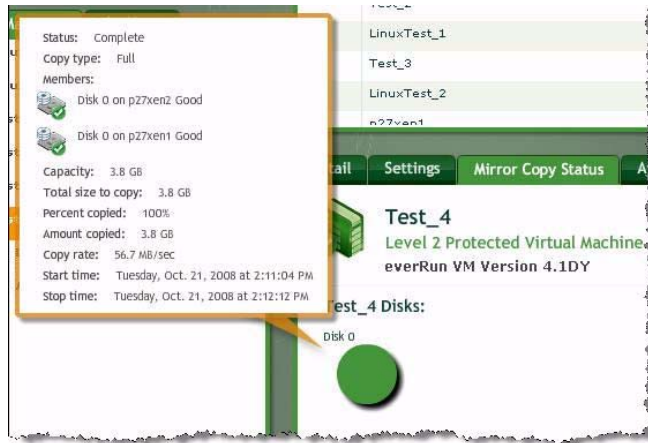


Figure 10 A progress icon turns green as the mirror process completes.

Accessing Pool and Resource Information

To monitor your pool resources, open the Hosts tab in the eAC resource pane, then highlight the line that contains the pool icon in the navigation tree (at the orange highlight in Figure 11). Each host icon in the tree opens to show the components available on that host.



Figure 11 The Hosts tab shows the resource pool icon at the top of the list.

When you highlight the pool name, the Detail tab at lower right shows resource information for the pool (Figure 12). In this example, you see the pool name, pool members, hardware configured, everRun license and version number, and relevant informational messages.

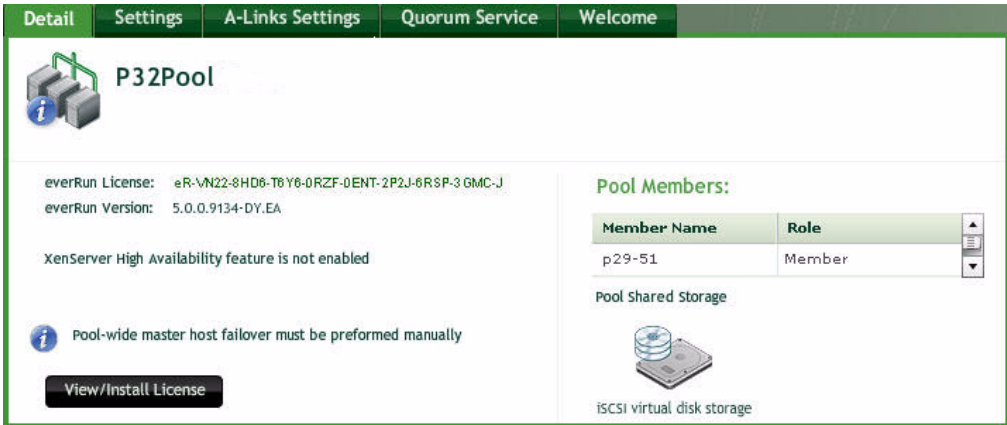


Figure 12 The Detail tab displays basic information about an everRun pool.

To learn more about a component displayed on the pool Detail tab, just hover the mouse over the component icon.

Icon-Based Status

Icon					
Meaning	Good	Warning	Error	Information	Unknown

Figure 13 Status icons allow you to quickly monitor the health of everRun-protected VMs.

You will learn to rely on eAC icons for an instant view of component status. For example, the gray server icon in Figure 11 shows that the resource is unprotected. (A protected resource has a green icon.) The server icon in Figure 11 also includes a small Running icon () to indicate that the server is currently powered on.

As a protected VM starts up, it displays a Starting icon () in the Power column of the Virtual Machine Status table — which changes to the Running () icon when startup is complete.

When a protected VM shuts down, it displays a Stopping icon () in the Power column the icon — which changes to the Halted icon () when shutdown is complete.

A VM protected at everRun Level 3 or Level 2 displays a large icon in the upper-right corner of the Detail tab.

Easy-to-Access Management Functions

Throughout the eAC, easy-to-access commands on context-sensitive menus support all management functions. For example, a single, flexible configuration screen lets you modify your key pool configuration settings.— networks, storage, and vCPU or memory assignments — at any time. .

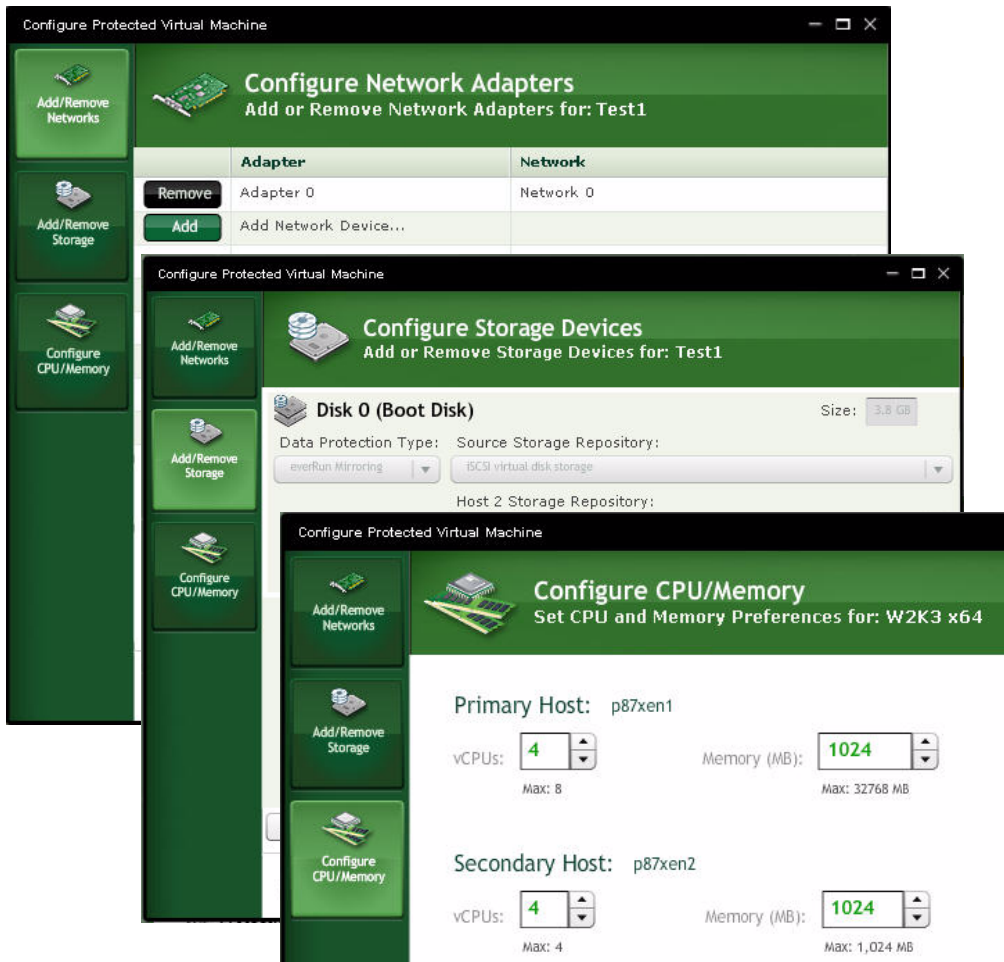


Figure 14 Click an icon on the left edge of a configuration screen to choose the pool element you want to configure: networks, storage, or CPU/memory.

Migrating Protected VMs from Host to Host

everRun supports two types of migration for protected VMs: migrate with reboot or online migration. In other words, you can either shut down a PVM and then move it to a different host (reboot migration), or you can move it without shutting it down (live or online migration).

Replacing Hosts or Components

When a host or component fails, you can quickly swap in a new one: just disable the host that contains the failed component, connect the new or repaired component, then re-enable the host.

Creating Backup Images Using everRun Snapshots

An everRun **snapshot** is a read-only image that represents the disk image and configuration of a protected VM at a certain point in time. It includes everRun drivers and configuration details, and maintains an association with the PVM from which it was created.

By maintaining stored images of active PVMs, you can capture an up-to-date record of your application state. If the need arises, a recent image can be used to recreate and deploy a new PVM that contains configuration and disk information from the original PVM.

Monitoring Components and States with System Alerts

everRun now includes automatic Simple Network Management Protocol (SNMP) system alerts, delivered to an email address. These messages provide information about state changes within the protected virtual machine (PVM) that help you monitor the health of each application.

Where to Learn More

4



This chapter tells you where to find additional information about Marathon everRun and Citrix XenServer software,

Marathon Documentation

Marathon provides you with a variety of documentation and online help sources to give you information you need to implement an everRun configuration quickly and efficiently. Please log into the Marathon Customer Portal at:

<https://support.marathontechnologies.com/portallogin.asp>

Click the Download button and navigate to the everRun section for links to the relevant documentation. Except as noted, the documentation listed in Table 2 is also accessible from the everRun Availability Center (eAC) management console.

Table 2 everRun Component and Feature Documentation

Documentation	Contents
<i>everRun Configuration Planning Guide</i>	Describes supported configurations with recommended best practices for servers, networks, and storage. Read before purchasing and installing, to plan your system deployment; use during deployment to verify or modify configuration details. (Available from the Marathon Customer Portal only.)
<i>everRun Compatibility Check utility and Readme</i>	Enables you to check the everRun compatibility of your servers' processors. It is available in the "Tools and Utilities" section of each everRun download area.

Table 2 everRun Component and Feature Documentation (Continued)

Documentation	Contents
<i>Read Me First</i>	<i>Please read before installing everRun.</i> Includes instructions for accessing installation materials and overall guidelines for installation. (Available from the Marathon Customer Portal only.)
<i>everRun Release Notes</i>	Provides release-specific information, including late-breaking information about known issues or limitations and suggested workarounds. This information is not included elsewhere in the everRun documentation set.
<i>everRun Setup and Installation Guide</i>	Explains how to install and configure the system software.
<i>everRun Upgrade Guide</i>	Provides instructions for upgrading to the current release, as well as earlier versions of everRun.
<i>everRun SNMP Guide</i>	Includes instructions for enabling and interpreting everRun SNMP alerts.
<i>everRun Messages Guide</i>	Interprets everRun messages displayed in Availability Manager (AM) log files and Windows event logs, and suggests appropriate actions.
<i>everRun Online Help</i>	Instructions, tips, and guidance include: <ul style="list-style-type: none">• How to use the everRun Availability Center to monitor and manage protected virtual machines• How to use the various administration tools and utilities• Troubleshooting and Tips for Success• Alphabetical glossary of everRun and XenServer terms• Summary of everRun commands (CLI) used for scripts and system administration

Table 2 everRun Component and Feature Documentation (Continued)

Documentation	Contents
Event reporting	<p>The everRun Log tab in the eAC provides feedback on task-based events:</p> <ul style="list-style-type: none">• Task detail describes issues that may affect completion• Detail suggests steps for remediation <p>The everRun harvest utility, <code>mtc_log-collector.sh</code>, collects log information about the operation of everRun-protected VMs. If you require Marathon support, your Marathon representative will ask you to email the output of the utility for diagnostic purposes.</p>

Information on the Marathon Web Site

Visit the Support web site at <http://support.marathontechnologies.com> for updates to information about the release, including:

- Supported Windows Service Packs, as well as operating systems that support everRun-protected VMs
- Tools and utilities available for problem diagnosis and resolution.
- Information about packaging of event logs and error log files for transmission to Marathon for analysis and problem resolution
- Technical Information Knowledgebase

Citrix XenServer Documentation

Open the XenServer corporate product site at <http://support.citrix.com>

Search for the current version of XenServer software, then click **Product Documentation**.

You can also download the XenServer documentation from the Marathon website at <https://support.marathontechnologies.com/portallogin.asp>

- *XenServer Release Notes*

The release notes cover any known issues with the named release of XenServer. Please read them before installing the product.

- *XenServer Installation Guide*

The installation guide walks you through how to install the software on a physical machine to convert it into a XenServer virtualization platform.

- *XenServer Virtual Machine Installation Guide*
Once you have installed a XenServer system, this guide shows you how to install virtual machines that run guest operating systems such as Windows or Linux.
- *XenServer Administrator's Guide (XenServer Reference Guide)*
The comprehensive reference manual for all the powerful features provided by your XenServer system, including the command-line interface documentation.
- *XenServer Software Development Kit Guide*
Tutorial for using the XenAPI Software Development Kit (SDK) to develop your own applications to control XenServer hosts via various programming languages and shell scripts.

Useful Web Addresses

<http://hcl.xensource.com/?showall=no&subtab=systems>

A list of hardware and software components that have been reported to work with XenSource products, either by Citrix, by the individual vendors, or by the Xen community. The preferred products are those listed by Citrix; those listed by individual vendors are also recommended.

Reference Books

Chisnall, Dave. *The Definitive Guide to the Xen Hypervisor*. Prentice Hall: November, 2007.

“An in-depth exploration of the architecture, interfaces, device support, management tools, and internal workings of the Xen hypervisor — including key information for developers who want to optimize applications for virtual environments.” (from publisher’s description)

Citrix White Paper: *The Three Levels of High Availability*.

A collaborative white paper, authored by Citrix Systems, Inc. and Marathon Technologies Corporation, which describes a range of availability solutions -- from core failover/restart through high and continuous availability. Use this white paper to balance priorities and cost, and to identify capabilities appropriate for your business application and budget. Available at http://www.marathontechnologies.com/resource_center.html

Matthews, Jeanna N., Eli M. Dow, Todd Deshane, Wenjin Hu, Jeremy Bongio, Patrick F. Wilbur, Brendan Johnson. *Running Xen: A Hands-On Guide to the Art of Virtualization*. Prentice Hall PTR: April 16, 2008.

A primer on Xen virtualization, along with insights, best practices, and case studies based on real-world environments. Endorsed by Ian Pratt, Xen Project Leader and VP Advanced Technology, Citrix Systems.

von Hagen, William. *Professional XEN Virtualization*. Wrox: January, 2008.

“A complete foundation in virtualization, Xen, and competing approaches with in-depth discussions of installing, configuring, and using Xen, including detailed examples of Xen configuration files, system configuration files, and system-level configuration information” (publisher’s description).

Wolf, Chris and Erick M. Halter, *Virtualization: From the Desktop to the Enterprise*. Apress: May, 2005.

Covers all aspects of virtualization, including virtual machines, virtual file systems, clustering, and virtualized storage solutions.

