



everRun[®]

Configuration Planning Guide

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 — Introduction

Overview	1
Contents of This Guide	1

Chapter 2 — Configuration Fundamentals

Overview	3
Production Deployments	3
Option A: Local Two-Host Pool	4
Option B: Local Multi-Host Pool	5
Option C: SplitSite Two-Host Pool	6
General Deployment Guidelines	7
Summary	8

Chapter 3 — Server Fundamentals

Overview	9
XenServer Hosts	9
Where to start?	9
How many servers?	9
How many CPUs?	10
What CPUs are supported?	11
How much RAM?	11
CPU and Server Rules	12
System Management (Client) Computer	12
quorum computers	13
Summary	13

Chapter 4 — Network Fundamentals

Overview	15
LAN Requirements	15
Management LAN	15
Management LAN Best Practices	16
Providing an Isolation IP Address	16
Private LANs (Availability Links)	16
Private LAN (A-Link) Rules	17
Subnet Configuration Options	17
Private LAN (A-Link) Best Practices	18
Public Production LANs	19
Public Production LAN Rules	19
Public Production LAN Best Practices	19
Additional NICs	20
Multicast Network Traffic	20
Provide TCP/IP Configurations for Production Deployment	20
Summary	20

Chapter 5 — Quorum Service Fundamentals

Overview	21
When is Quorum Service Required?	21
Quorum Service Rules	22
Quorum Service Best Practices	23
Summary	23

Chapter 6 — Storage Fundamentals

Overview	25
everRun Storage Support	25
everRun Data Protection	26
everRun Mirroring.....	26
Storage Subsystem Data Protection	27
Performance and Placement	29
Storage Rules	30
Summary of Storage Options	30

Chapter 7 — Summary of Deployment Options

Overview	31
Configuration Options	31
Meeting Your Availability Goals	31
Protection Level 3 (System Level Fault Tolerance)	32
Protection Level 2 (Component Level Fault Tolerance)	32
Protection Level 1 (Basic Failover)	33
Surviving Management LAN Segmentation Failure	34
Best Practices to Avoid LAN Segmentation Failure	34
Summary of Deployment Guidelines	34
Comparison of everRun Deployment Options	35

Appendix A — Option A: Local Two-Host Pool

Overview	37
Software	37
Servers	38
Networking	39
Storage	41
System Connectivity	41

Appendix B — Option B: Local Multi-Host Pool

Overview	43
Software	43
Servers	44
Quorum Service	45
Networking	46
Storage	48
System Connectivity	48

Appendix C — Option C: Two-Host SplitSite Pool

Overview	51
Software	51
Servers	52
Quorum Service	53
Networking	54

Storage	54
System Connectivity	56

Appendix D — Option D: Two-Host Pool with Level 3 SMP 59

Overview	59
Software	59
Servers	60
Networking	61
Storage	63
System Connectivity	63

Brief Glossary

everRun Terminology	65
XenServer and Industry Terminology	67



Overview

This document provides guidelines to help you select and plan your everRun configuration — and then to provide hardware, network, and storage components for that configuration. These planning guidelines are meant to ensure maximum success during system installation and operation.

For help in understanding the terms used in this guide, please read the *everRun Product Overview* and consult the brief glossary at the end of this document. Before starting your installation, review the *everRun Release Notes* provided with this release. To install, use both the *Read Me First* installation overview and the step-by-step installation details in the *everRun Setup and Installation Guide*.

Contents of This Guide

This guide describes the supported everRun deployments.

- Chapter 2, Configuration Fundamentals, describes the elements of an everRun system. It describes four supported production configurations:
 - Option A: Local, two-host pool
 - Option B: Local, multi-host pool
 - Option C: SplitSite® (disaster tolerant) two-host pool
 - Option D: Local, two-host pool with at least one Level 3 SMP protected VM

Appendixes A to D provide detailed information about the specific software, server, network, quorum service, and storage configuration guidelines for each of these deployments.

- Chapter 3, Server Fundamentals, lists XenServer host requirements.
- Chapter 4, Networking Fundamentals, describes networking hardware and IP

addressing schemes needed to interconnect pool resources.

- Chapter 5, Quorum Service Fundamentals, discusses quorum service, which runs on general-purpose computers and provides continuous communication among servers in an everRun pool. The quorum service computers may be in a local configuration (Option B) or a geographically separated configuration (Option C) using the SplitSite option.
- Chapter 6, Storage Fundamentals, describes storage hardware resources and configuration information.
- Chapter 7, Summary of Deployment Options, lists strategies for meeting your availability goals using everRun Level 3, Level 2, and Level 1 fault tolerance, as defined in the following table. The chapter describes the configuration options as mandatory or not supported, and rates them as Good, Better, or Best choices for fault tolerance or resiliency.

Fault Tolerance Level	Properties
Level 3, System-Level Fault Tolerance	<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
Level 2, everRun Component-Level Fault Tolerance	<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 • 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
Level 1, Basic Failover (Citrix XenServer HA)	<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



Overview

This chapter describes the production deployments supported in this release.

NOTE: This guide does not describe evaluation versions of everRun. The Marathon website at <http://www.marathontechnologies.com/> invites you to try a free 30-day evaluation of everRun.

Production Deployments

The production options detailed in this guide are:

- **Option A: Local two-host pool.** This redundant two-host configuration is located in a single site, such as a computer room. It uses private, directly connected availability link (A-link) networks. Option A is the simplest to configure. It is the configuration encouraged for evaluations — including tests and demonstrations of failures, recoveries, and repair scenarios. For a similar deployment that includes at least one Level 3 PVM running SMP, see Option D.
- **Option B: Local multi-host pool.** This configuration extends the local pool beyond two servers. Marathon has tested up to four servers interconnected within a pool. Because this deployment requires active network devices (switches, routers) for A-link networks, it also requires a quorum service to ensure the integrity of PVMs in the pool against real-world failure scenarios.
- **Option C: SplitSite® (disaster tolerant) two-host pool.** This configuration includes the Marathon SplitSite option, to maintain hardware redundancy and redundancy of physical computer rooms and the buildings containing them. Because of the geographic separation, this configuration requires careful planning of component placement and more complex networking topologies. Quorum service is also required, as in Option B, because a SplitSite configuration exposes the A-link networks to other potential failure scenarios.

- **Option D: Local two-host pool with Level 3 PVMs running on multiprocessors.** This deployment is identical to Option A, except that it includes at least one Level 3 PVM running on multiprocessors (SMP). Because SMP operation imposes different pool requirements, it is described in this document as a separate configuration option.

NOTE: Networking considerations are further described in Chapter 4, quorum considerations in Chapter 5, and storage considerations in Chapter 6. The summary table in Chapter 7 compares the availability of various deployment options and includes additional SplitSite perspectives.

Option A: Local Two-Host Pool

An everRun protected VM (PVM) requires two physical servers, and provides redundancy by connecting a set of resources on a primary host server to a second set of resources on a different physical host, as shown in *Figure 1*.

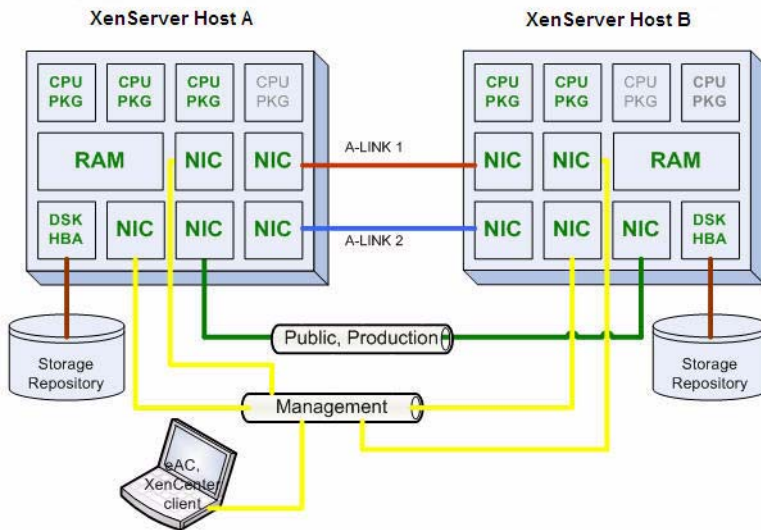


Figure 1 Typical two-host pool installed in a single computer room.

This example uses two servers in a single room with directly connected Ethernet cable/fiber connections. The first everRun host for the protected VM resides on Host A, and the second everRun host is located on Host B. The two A-links provide a redundant connection between the two everRun hosts.

This configuration does not allow you to enable quorum service. Instead, the short, point-to-point A-link LANs minimize the likelihood of simultaneous A-link network failures, in order to avoid an undesirable split-brain scenario.

Variations of this configuration are possible. If you pool includes a Level 3 PVM running SMP, see Option D. If your deployment requires more than two hosts or quorum service, see Option B. If your deployment requires two sites, see Option C.

If Option A meets your deployment needs, review the details in Appendix A before beginning your installation.

Option B: Local Multi-Host Pool

Use Option B (*Figure 2*) if you are planning to deploy a pool of two to four XenServer hosts on a single site — with active, switched A-link networking and quorum service.

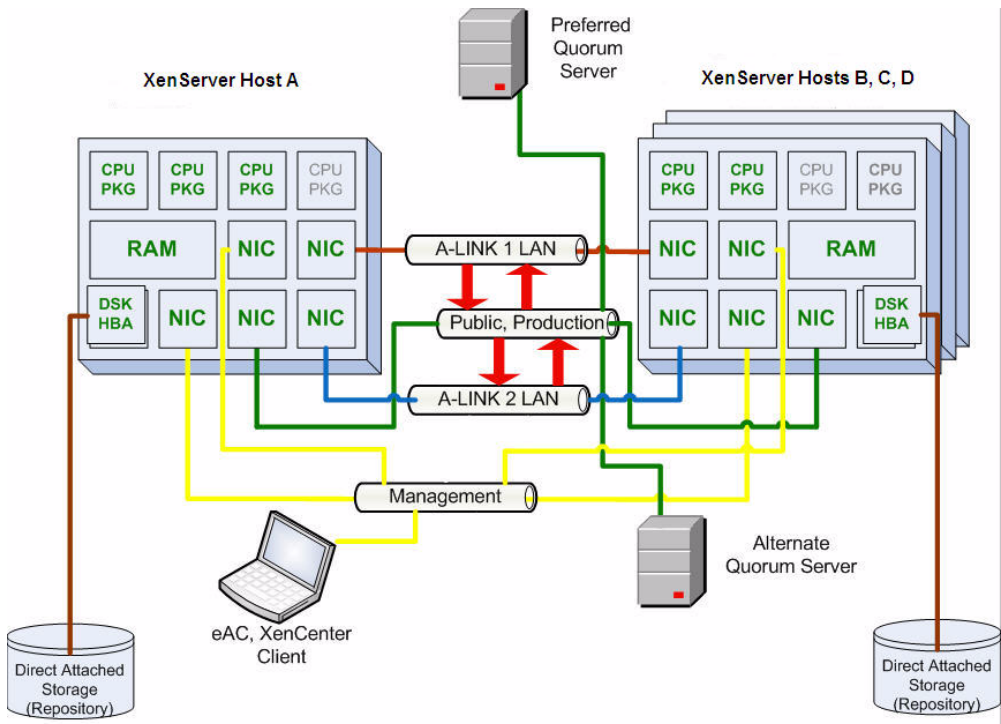


Figure 2 A multi-host pool with two to four hosts on one site, implementing active, switched A-link networks, and storage attached directly to each host. Arrows represent gateways between additional A-link vLANs and vLANs where quorum service clients will be placed in your infrastructure.

This deployment includes two to four XenServer hosts at the same location. To supply the required quorum service, two quorum computers are configured and placed at strategic locations on a production LAN. The networks are interconnected as shown in [Figure 2](#). The A-link LANs are used to access quorum service computers via gateways.

The storage architecture for this option must be selected to meet the availability goals, features, and performance objectives of all the PVMs you plan to deploy on the pool.

If Option B meets your deployment needs, review the details in Appendix B before beginning your installation.

Option C: SplitSite Two-Host Pool

In a SplitSite deployment (shown in [Figure 3](#) on the following page), the two XenServer hosts are located in separate data centers or in geographically separate buildings, with wide area networking configured between them as shown in the drawing. (The red arrows are required, multiple gateways between additional vLANs in your infrastructure.) SplitSite requires quorum service, configured in accordance with best practices. For best fault tolerance, a preferred quorum computer is located in a third facility and an alternate is located in a fourth. However, the alternate quorum computer can also be placed with the preferred quorum computer and still yield satisfactory results.

The storage architecture for this option must be selected to meet the availability goals, features, and performance objectives of all the PVMs you plan to deploy on the pool.

If Option C meets your deployment needs, review the details in Appendix C before beginning your installation.

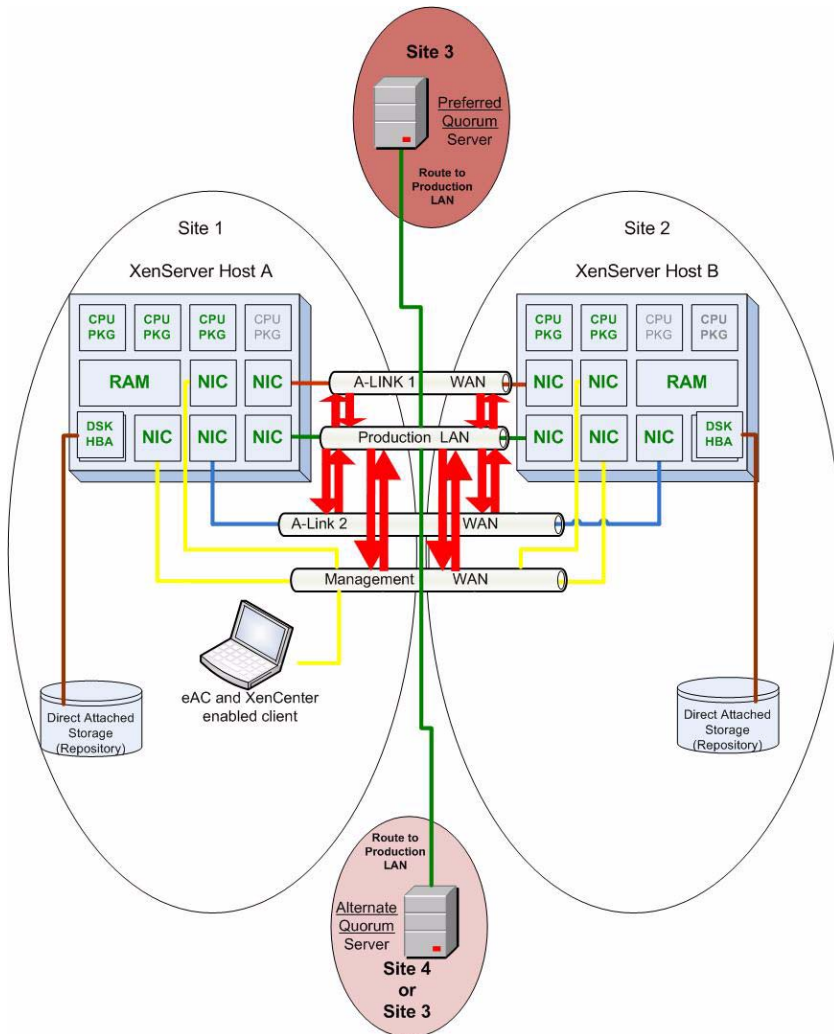


Figure 3 A SplitSite configuration has one XenServer host on each of two sites. The deployment includes active, switched A-link networks and storage attached directly to each host.

Option D: Local Two-Host Pool with Level 3 SMP PVMs

This deployment is identical to Option A, shown in [Figure 1 on page 4](#) except that it includes at least one Level 3 PVM running on up to four vCPUs. Because multiprocessor operation imposes different pool requirements, it is described in this document as a separate configuration option.

The ability to run protected VMs on multiple processors provides important scalability benefits for the protection of critical applications. However, SMP operation imposes certain requirements. For example:

- SMP operation requires specified Intel processors.
- SMP operation has different memory requirements from a conventional everRun pool. For example, each Each Level 3 PVM running in multiprocessor mode (SMP) uses 512 MB incremental RAM rather than 128 MB.
- When getting started with SMP, it is a good idea to deploy Level 3 SMP PVMs incrementally. Start with a single SMP PVM in a pool, assess resource utilization, and then decide whether the remaining resources will support an additional SMP PVM. If you deploy two SMP PVMs, choose a different server as the primary host for each.

If Option D meets your deployment needs, review the details in Appendix D before beginning your installation.

General Deployment Guidelines

You have the flexibility to set up your everRun deployment in a variety of ways. In addition to the supported options outlined in Appendixes A to D, you can plan to protect your virtual machines at Level 3, 2, or 1, or some combination of those protection levels.

Some of the guidelines described in this document apply only to certain deployments. For example, a pool containing Level 3 protected VMs has specific requirements, as mentioned earlier in this chapter.

If you do not plan to include XenServer HA (everRun Level 1) in your pool, some requirements associated with Level 1 operation in this guide are modified. Specifically:

- You do not need to configure shared storage. When using only everRun Level 3 and Level 2, you can reap the benefits of local storage. These benefits include maximum isolation of failure modes and improved I/O performance for Marathon mirror sets.
- There is no need to configure to prevent host fencing as a result of the XenServer HA protection algorithm. You can place Level 3 and Level 2 PVMs on any host in the pool and depend on Marathon fault-handling for all eventualities.
- In addition, you no longer have to meet the XenServer HA requirement of a bonded NIC pair on the management network. *However, Marathon recommends implementing this bonded NIC to provide maximum reliability of the management network.* Because a bonded NIC uses two separate network cards, this configuration consumes a minimum of four, not five NICs.

When XenServer HA is not included in your configuration, an automatic or manual Marathon pool master failover mechanism may be activated. The following points

describe the relationship of quorum service and Marathon automatic/manual master failover.

- If quorum service is enabled, it overrides Marathon's automatic pool master failover. In this configuration, all VMs continue to operate when the pool master fails, but the system administrator manually designates the new pool master host as required.
- Quorum service is not applicable in the case of a local, two-host pool with direct cable or fibre connections. In that configuration, the basic Marathon Management Service (MMS) provides pool master failover.

Summary

When selecting a deployment option, you may find it helpful to review the material in Chapter 7, which lists availability advantages and tradeoffs of each deployment type.

Appendixes A to D provide detailed configuration tables for software, servers, quorum service, networking and storage in the corresponding configuration option.

If you have questions not addressed in this guide, or if you need help to plan a custom configuration that meets your availability needs, consult your Marathon reseller or Marathon professional services.



Overview

Hardware redundancy enables everRun-protected VMs to operate through failures, resulting in very high levels of availability for Windows environments. This chapter describes the basic hardware and software requirements for hosts in an everRun pool, and the general-purpose computers required for system management or quorum service management.

XenServer Hosts

The basic hardware requirement for a single protected VM is two separate XenServer hosts that are members of the same resource pool. Each server should be configured according to the guidelines in this chapter.

Where to start?

Choose two or more servers from those listed on the Citrix XenServer Hardware Compatibility List at the link shown below.

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

To narrow your search, specify “server” and enter your vendor of choice. See “CPU and Server Rules” on page 14 for more information.

How many servers?

Acquire the number of servers you will deploy to meet your capacity and redundancy requirements. See “CPU and Server Rules” on page 14 for more information.

How many CPUs?

Populate each server with enough CPU cores to meet your performance needs. The total number of CPU cores will vary greatly, according to your system design.

This section lists a few factors to consider when planning your CPU requirements. See “CPU and Server Rules” on page 14 for more information.

- The XenServer environment uses one core.
 - Each running VM can use one or more CPUs, and each protected VM represents two VMs on two separate XenServer hosts.
 - Consider the capacity requirements on each host after failover has occurred. For example, Level 3, Level 2 and Level 1 PVMs may require additional resources on the machines they fail over to.
 - **Level 3 protected VMs** — Must use the same number of CPUs on both hosts.
 - **Level 2 protected VMs** — The default configuration uses the same number of CPUs on the active and standby hosts. Performance will be reduced unless free CPU cores are reserved for the worst case scenario, in which a host failure causes all VMs to fail over at the same time. If a PVM must restart on another host using fewer CPUs, online migrations are also inhibited — and availability may be less than it otherwise could be.
 - **Level 1 protected VMs** — Performance will be reduced unless free CPU resources are reserved for the worst case scenario, in which a host failure causes all VMs to fail over at the same time.
 - Various processes running on a pool member host can share CPU cores. However, for best availability and performance, you should provide enough cores on both servers to minimize or eliminate CPU sharing.
-

NOTE: Guidelines and a worksheet for planning your vCPU usage on a XenServer host are available in the KnowledgeBase at the Marathon support website. To review these guidelines, log in to everRun Technical Support with your customer password, click the **Knowledge** button, and choose the **Documents** category in the list at left. Then type "CPU Overcommit Counters" into the **Search This Category** window and click **Search**.

What CPUs are supported?

A list of the processors currently supported for everRun Level 3 SMP, Level 3 uniprocessor, and Level 2 protection is available at http://www.marathontechnologies.com/documents/Hardware_Qualifications.pdf

If you want to evaluate whether your hardware supports everRun Level 3 multiprocessor (SMP), Level 3 uniprocessor, or Level 2 operation, download and run the **everRun Compatibility Check utility** to pre-qualify your hardware. The utility, which is available in the Tools and Utilities section of the Marathon download site, checks CPU vendor, brand, and type; microcode revision; and other features required for everRun operation. It runs from a DVD or USB key attached to the target hardware and does not require XenServer or everRun software to be installed at the time. For instructions, download the *Compatibility Check Readme* document, also available at the Marathon site in the Tools and Utilities category.

How much RAM?

Populate each server with sufficient RAM to support all the VMs you need to run at once. XenServer limits you to the number of VMs the physical RAM can support.

Provide a minimum of 4GB of RAM per server to accommodate four (or more) protected VMs and the XenServer environment. Less RAM will constrain you to have fewer VMs or VMs with smaller amounts of memory.

***Tip:** Make sure both XenServer hosts running protected VM pairs contain enough RAM to support your requirements if a failure of the other host occurs. Meeting this recommendation also provides capacity for online migration in repair and recovery scenarios.*

For best performance, a protected VM should — with few exceptions — have identical RAM on the second everRun host.

- **Level 3 protected VMs** — Require an identical amount of RAM on both hosts. everRun is running Windows simultaneously on both physical hosts *even before a failure*.

Each VM protected at Level 3 (uniprocessor) or Level 2 uses 128 MB of incremental RAM (beyond that configured for Windows) on both hosts. It is mandatory to meet this RAM requirement on each host.

Each Level 3 PVM running in multiprocessor (SMP) mode uses 512 MB incremental RAM rather than 128 MB.

-
- **Level 2 protected VMs** — Unless configured otherwise at the time of protection, Level 2 PVMs reserve identical RAM on both hosts.

RAM allocated to Level 2 PVMs depletes available resources on the host; thus Level 2 PVMs have priority over Level 1 PVMs when failover attempts are processed. In Level 2 protection the amount of RAM configured for the standby side is reserved when the PVM is started (in advance of any host failure).

Online migration (a feature that permits you to maintain PVM operation during software updates or hardware repair scenarios) can be performed only if the RAM is identically configured on both hosts. However, if a PVM is configured to use less RAM on the standby host, it will still be possible to migrate with restart.

- **Level 1 protected VMs** — Requires identical RAM on two hosts, according to your failover needs. It is best to have ample RAM reserved when a single host failure would cause many VMs to fail over. If adequate RAM is not available, some Level 1 VMs cannot be restarted on another XenServer host.
-

CPU and Server Rules

- ▶ *Select a server with approved Intel processors from the lists available in the Read Me First document at the Marathon download site.*
 - ▶ *Minimum of 2 servers per pool; 4 maximum have been qualified.*
 - ▶ *Each CPU must be the same family and model, and otherwise be eligible to join a common XenServer pool.*
 - ▶ *Provide a minimum of 4 CPU cores per server; 32 maximum per server.*
-

System Management (Client) Computer

The client computer(s) used for everRun management can be any general-purpose computers running Windows XP (SP2 or newer), Windows Server 2003, Windows Vista, or Windows 7.

You use the client computer to access the two management interfaces in the everRun system:

- **The everRun Availability Center (eAC)** — The browser-based eAC everRun management GUI can be loaded and run by a Flash-enabled web browser on any computer, including the management computer.
- **The XenCenter interface** — XenCenter is the client application for remote management of XenServer hosts. It can be installed and run on any Windows computer, including the management computer.

Quorum Service Computers

If your deployment uses active, switched networks, quorum service software is required. Quorum service software should be installed on two Windows XP (SP2 or newer), Windows Server 2003 or 2008, Windows Vista, or Windows 7 computers, as described in Chapter 5.

Summary

For more information about everRun server configurations, see Chapter 7, “Summary of Deployment Options”. After you select a deployment option, refer to the relevant appendix (A to D) before procuring hardware and beginning your hardware and software installation.



Overview

Proper attention to network design can make the difference between building a configuration that merely works, and building a configuration that delivers the availability you expect when a hardware or network failure occurs. This chapter describes requirements and best practices for the management LAN, private (A-link) LANs, and public (production) LANs.

NOTE: Static IP assignments are generally accepted practice for server production deployments, and are emphasized for everRun deployments.

When planning your private (A-link) networks, you should adopt an IPv4 network IP address scheme that is compatible with your overall network infrastructure, and has no conflicts with existing addresses. Although everRun provides a default pair of address ranges that are suitable for testing a simple, single pool, Marathon recommends that you provide an IP address scheme that offers maximum reliability in your environment.

LAN Requirements

Four physical LANs (using four or five NICs per host) are recommended, as shown in [Figure 1 on page 4](#). An additional NIC is added when a bonded management LAN is used. A bonded management LAN is strongly recommended when deploying Level 1-protected VMs (called XenServer High Availability [HA] in Citrix documentation).

Management LAN

The XenServer management LAN connects pool members and management applications. To insure against failure of this LAN, Marathon recommends that you use a bonded NIC pair on each host, as recommended by Citrix, to provide management LAN redundancy between hosts and network switches.

Management LAN Best Practices

- ☑ For maximum redundancy, provide independent production and management NICs and LANs (or vLANs) whenever possible. Connect each management NIC through switches and uplinks such that no single switch or uplink failure can impact multiple hosts or both NICs in a bonded pair.

 - ☑ Bonded NICs are strongly recommended in the management LAN, to provide management redundancy in each host. This configuration is required if you plan to enable XenServer HA (deploy Level 1 PVMs) in your pool. Configure the bonded NIC pair on the management LAN, with each NIC cabled to a separate Ethernet switch, before you apply the everRun software to the host servers.

 - ☑ Configure gateway access between your management LAN and other LANs as required for eAC and XenCenter access from clients normally connected to other LANs (including LANs used for remote management purposes).

 - ☑ Configure static IP addresses for management LANs in each pool member host. When necessary, the management LAN connections may place each host on a different IP subnet. Access between hosts can be provided by switched/routed WAN/LAN capability.

 - ☑ To support online migration of a Level 2 protected VM or (re)synchronization of a Level 3 PVM, all NICs, switches, and uplinks must be capable of 155 Mbps or greater, with full duplex connections. 1 GbE components or faster are recommended, although a minimum 155 Mbps uplink capacity is sufficient for SplitSite configurations.

 - ☑ Identify a device with a static IP address for use by Marathon's Management Service (MMS) failover. This device replies to ICMP ping requests initiated by all hosts in the pool, and it should be reachable at all times on the management LAN. MMS failover is defeated when this device is restarted or unavailable for any reason, pending recovery of this device.
-

Providing an Isolation IP Address

As described in the final row of the preceding table, the everRun MMS relies upon a pool-wide isolation IP address to maintain application availability. This mechanism assures that everRun automatically appoints a replacement if the active pool master should fail. During pool configuration in the final step of your system installation, you should supply the IP address of a hardware device that is always reachable from the hosts in the pool.

The pool-wide isolation IP address mechanism is explained in greater detail in the everRun Availability Center (eAC) Online Help.

Private LANs (Availability Links)

VMs protected at everRun Level 3 and Level 2 use two LANs for private availability-link networks (A-link 1 and A-link 2). Within a protected VM, the A-link networks connect the two everRun hosts for the protected VM.

Private LAN (A-Link) Rules

- ▶ *Do not configure NIC bonding for availability-link (A-link) networks.*
 - ▶ *everRun 6.0 or newer allows greater isolation of subnets within the A-links. As in earlier releases, you must configure each A-link on a separate subnet from the other A-link LAN; however, when connecting through a wide area network, you may use additional subnets and routed connections between the endpoints. A maximum of four A-link address ranges can be configured on four distinct subnets. This is described in the following section, “Subnet Configuration Options”.*
-

Subnet Configuration Options

When setting up your network IP addresses, you can elect to assign each A-link to a single IP address range, as in prior everRun releases — or you can assign a separate IP address range to the connection at each end of the A-link, for a total of four A-link address ranges in separate subnets per pool. This section describes the appropriate use of two-subnet and four-subnet approaches for the supported everRun configuration options.

- If you are configuring a **local, two-host pool** (Option A or Option D), there is no need to configure quorum service. A simple two-address-range deployment, where each A-link is on a separate subnet, is adequate.
- Similarly, if you are configuring a **local multi-host pool** (Option B), two address ranges (one for each A-link) are normally adequate, because it is unlikely that the A-link networks would be routed.
- However, in the case of a **SplitSite configuration** (Option C), where the two hosts are in separate locations, it is advantageous to use routed networks for your A-link LANs. Further, it is recommended that you use separate, local gateways to reach the preferred and alternate quorum computers from each site. This enables SplitSite configurations to meet the availability goals without compromising data integrity.

Figure 4 is a simplified illustration of a routed four-subnet scheme. Each end of an A-link has a separate path, via a local gateway, to a quorum computer. For simplicity, the drawing shows only the pathway to the preferred quorum computer; a similar pathway would be configured to the alternate quorum computer, using the same approach.

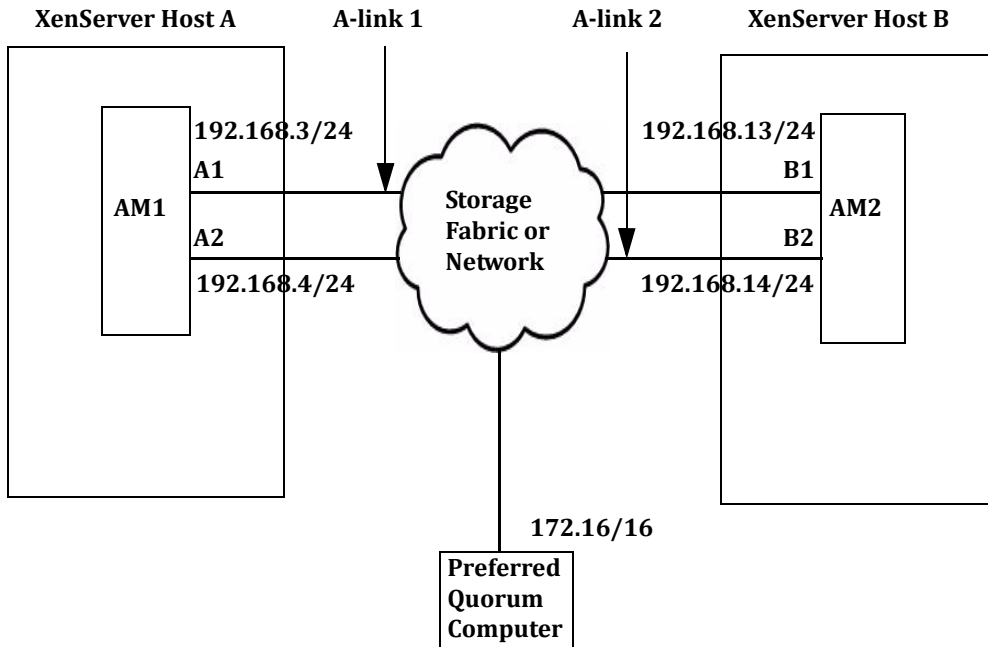


Figure 4 Simplified sketch of a four-subnet scheme that provides separate pathways to each end of each availability link (A-link). IP addresses in the illustration are samples, not examples.

***Private LAN (A-Link) Best Practices
(when deploying SplitSite, or any configuration with switches)***

- ☑ Always provide independent NICs and LANs for A-links in each XenServer host.

- ☑ Connect each A-link NIC through switches, routers, and uplinks such that:
 - No one device or uplink failure can impact the ability of more than one host to contact other hosts.
 - No one device or uplink failure can impact other hosts' access to their respective gateways.
 - No one switch or uplink failure can impact both A-link LANs on the same host.

Private LAN (A-Link) Best Practices (when deploying SplitSite, or any configuration with switches)

- ☑ Identify two or four separate subnets for A-link 1 and A-link 2 that will not conflict with assignments reserved for another purpose. You may be configuring switches and gateways on this LAN to accommodate multiple hosts, sharing the A-link LAN with other everRun pools, and providing access to quorum service computers on other LANs.

Assign subnets for these LANs to avoid potential conflicts with other everRun pools you will install in the future. Avoid the available default address ranges to prevent conflicts with other pools when re-installing hosts during hardware repair/replacement scenarios. (Refer to the appendix that describes your configuration for detailed instructions).

- ☑ Configure a gateway between the A-link1 LAN and the production LAN and another gateway between the A-link 2 LAN and the production LAN for use in accessing quorum service computers. For SplitSite deployments, specify four subnets and configure same-site gateways and routers for each A-Link connection. See Chapter 5 for a discussion of quorum service and quorum service computers.

Exception: Only two subnets can be specified and a gateway cannot be used when two hosts are connected with short, point-to-point cables.

- ☑ All NICs, switches, and uplinks should be capable of 1Gbps (or faster), full-duplex connections. For SplitSite, the minimum speed for all switches, routers, and uplinks must be 155 Mbps.
-

Public Production LANs

Production networks that connect virtual machines and LAN-based clients use a single LAN. VMs protected at everRun Level 3 and Level 2 provide redundant LAN access for the Windows operating system by managing separate physical NICs on two hosts in the pool as a redundant pair. Up to four public production LANs can be configured for each PVM.

Public Production LAN Rules

- ▶ *Although it is possible in some circumstances to combine production and management NICs, such a configuration is not qualified. Further, a bonded management NIC configuration is not supported for use as a public production LAN.*
 - ▶ *Production LANS must pass layer 2 multicasts between all ports connected to the physical NICs on each host in the pool.*
 - ▶ *Do not configure NIC bonding for production networks used by Level 3 or Level 2 PVMs.*
-

Public Production LAN Best Practices

- ☑ Although a single physical NIC and LAN can be deployed for multiple networks, it is not recommended. For maximum redundancy, provide separate production and management NICs and LANs whenever possible.
 - ☑ Connect each production NIC through switches and uplinks such that no one switch or uplink failure can impact the ability of clients to connect to another host (peer or redundant NIC).
-

Additional NICs

everRun supports up to four virtual (redirected) NICs per Windows protected VM. A total of eight redirected devices (NICs and disks combined) per protected VM is supported. For more information about the maximum number of physical NICs supported by XenServer, refer to the Citrix XenServer documentation at the Citrix website (<http://support.citrix.com>) or the Marathon support website (<https://support.marathontechnologies.com/portallogin.asp>).

A XenServer tech note (available at <http://support.citrix.com/article/CTX123158>) describes the recommended configuration of XenServer switch ports.

Multicast Network Traffic

In lightly loaded networks, the absence of multicast network traffic can produce inconsistent failover response. In non-busy, quiet LAN environments Marathon recommends installing a multicast generator to accelerate failover response.

Provide TCP/IP Configurations for Production Deployment

This document does not describe production TCP/IP configurations for the VMs you plan to install on the pool. However, your configuration plan should accommodate each protected VM within your (static) IP address range.

Summary

For more information about everRun network configurations, see Chapter 7, Summary of Deployment Options. After you select a deployment option, use the corresponding appendix to guide your hardware and software installations.

Quorum Service Fundamentals

5



Overview

The quorum service communicates with Level 3 and Level 2 PVMs to ensure integrity when multiple networks (all A-links and production LANs) between hosts in an everRun configuration fail. This failure scenario becomes likely whenever A-Link network complexity exceeds a direct connection of more than, say, three meters. In that event, the fault management software, aided by quorum service managers, selects an appropriate everRun host that will continue. To prevent *split-brain operation* — an undesirable condition in which the two availability managers (AMs) servicing the two halves of a protected VM are partitioned and operating independently — the instance on the other host halts.

This chapter describes deployment options that require quorum service and summarizes the rules and best practices for quorum service configuration.

When is Quorum Service Required?

All configurations, except those where A-links are limited to short, point-to-point cables, require the use of quorum service computers. In configurations where quorum service computers are required, the service is installed on Windows client computers, which are placed strategically and accessed via gateways to one or both A-link LANs. The server requirements for quorum service computers are described in “Quorum Service Computers” on page 15.

VMs protected with everRun Level 3 and Level 2 fault tolerance make appropriate use of quorum service to prevent potential split-brain scenarios. Whenever access to both quorum service computers is lost, impacted PVMs downgrade to guard against the possibility of significant LAN failures that could result in split brain.

Quorum service also improves unattended startup scenarios (for example, recovery from situations such as power outages impacting multiple hosts) that otherwise would require manual intervention. Level 1 basic failover does not make use of quorum service.

Quorum service is required for each of the following deployment scenarios:

- A pool that contains more than two hosts.
- A pool where private (A-link) network(s) are configured with active components; that is, any network connections except short, point-to-point A-link cables.
- A pool that spans multiple sites or computer rooms (SplitSite).

Separate, short, point-to-point A-links (using independent NIC hardware) minimize the likelihood of simultaneous loss of all networks. Quorum service computers can be omitted only for this deployment scenario (called Option A in this guide).

Quorum Service Rules

- ▶ *Provide two quorum service computers for the everRun pool.*
- ▶ *Quorum service installs as a Windows service on XP SP2 or newer, Server 2003, Vista, or Windows 7. Other Windows operating systems may work but are not certified by Marathon.*
- ▶ *Quorum service must not be installed on a VM residing in the pool where everRun is installed.*
- ▶ *Quorum service computers must be configured using static IP addresses.*
- ▶ *Quorum service computers are accessed via gateways having access to A-link 1 or A-link 2.*
- ▶ *everRun protected VMs contact quorum service computers using UDP. By default, firewalls must open port 2189 between A-link NICs and quorum service computer.*
- ▶ *Quorum service computers communicate with AMs operating on everRun-enabled pool member (or master) hosts using UDP. By default, firewalls must open port 2188 between A-link NICs and quorum service computers.*
- ▶ *everRun quorum service installations can coexist on the same computer with other applications and services, including quorum service for earlier versions of everRun FT and everRun HA. That is, quorum service installations from different product families (everRun FT/HA and everRun 2G/VM, for example) operate independently, and may be installed side by side on the same quorum computer.*

If you are running a quorum service from a prior version of everRun 2G/VM, it is recommended that you install the quorum service that is part of the latest everRun product. This version of the quorum service software is backward-compatible with all versions of the everrun 2G/VM product.

Quorum Service Best Practices

- ☑ Configure two quorum computers with minimum common networking between the quorum computers and the pool hosts. In SplitSite configurations, specify four subnets and separate, local gateways associated with each NIC supporting A-link and quorum traffic.
-
- ☑ Quorum computers should not reside on the same site as any XenServer host. If both preferred and alternate quorum computers fail for a common reason, Level 3 and Level 2 PVMs will gracefully downgrade redundancy — then continue to operate using one host, pending recovery of a quorum computer. If a XenServer host and the elected quorum computer fail for a common reason, everRun hosts running on the surviving XenServer shut themselves down.
Exception: When the entire pool is on one site, placing both quorum computers at the same location — where all hosts in the same site can fail — creates little additional liability.
-
- ☑ If the preferred and alternate quorum computers must reside on a common site, power them from separate AC power sources (phases) or configure them on separate UPS devices. Placing both quorum computers on a common third site is preferable to locating either on the same site as a XenServer Host.
-
- ☑ Configure a physical network route to a preferred quorum using A-link 1 and the associated gateway. Make sure that nothing in the route from one everRun host to the quorum service can be impacted by failures in the *other* everRun host's connection path to the preferred quorum computer.
-
- ☑ Configure a separate physical network route to an alternate quorum using A-link 2 LAN/gateway — with no hardware in common with the A-link 1 LAN/gateway. Make sure that nothing in the route from one everRun host to the quorum service can be impacted by failures in the *other* everRun host's connection path to the preferred quorum computer.
-

Summary

Optimum configuration and placement of the quorum service computers may vary somewhat. Key variables are the number of sites you have to work with, as well as the network designs that interconnect the hosts and sites. After choosing a deployment option in Chapter 2, refer to the appendix corresponding to that option for more specific guidelines in implementing quorum service.



Overview

You must configure storage capacity, functionality, and performance sufficient to meet the needs of all PVMs you plan to deploy on the pool. A minimum of 72 GB of direct-attached storage per XenServer host is adequate for evaluation and testing. You should increase storage capacity and functionality as required to accommodate the storage features in your design and the applications you plan to install.

This chapter describes the availability and performance tradeoffs associated with the storage decisions you will make, enabling you to make informed compromises against your storage budget.

everRun Storage Support

everRun supports both types of storage repositories (SRs) described in the XenServer documentation:

- **Host-specific storage** — physical storage directly attached to a host, or networked storage allocated to a single host. In both cases, the storage is not available to other hosts in the pool.
- **Shared storage** — networked physical storage with access to all hosts in the pool. This is a XenServer configuration property, described in the Storage chapter of the *XenServer Administrator's Guide*. It includes a variety of storage options, including iSCSI and FC-SAN. You can configure any of these options as described in the XenServer documentation.

everRun Data Protection

The storage options described in this chapter are based on two methods that can be used for everRun data protection:

- **everRun mirroring** — a process for creating and maintaining a set of identical disk images on separate physical disks. Each everRun-protected VM relies on a

pair of virtual disk images (VDIs) in storage repositories (SRs); this everRun **mirror set** maintains identical information in each image. If a fault occurs on one of the disks in a mirror set and the disk cannot be accessed, the everRun configuration automatically uses the remaining disk in the mirror set to provide continuous access without losing data. The mirror set members should be located in different SRs.

- **Storage subsystem data protection** — Storage subsystem data protection is a path-redundant option, which is available to users whose configuration includes intelligent SAN systems that provide data redundancy within the storage fabric. This option relies exclusively on the storage hardware to protect the data. Both everRun hosts access data on a single VDI. If one pathway to the storage fails, the everRun software activates the alternate pathway on the other server hosting the protected VM.

everRun Mirroring

Host-specific storage repositories using everRun data mirroring (see [Figure 5](#)) provide most reliable protection for Level 3 and Level 2 PVMs.

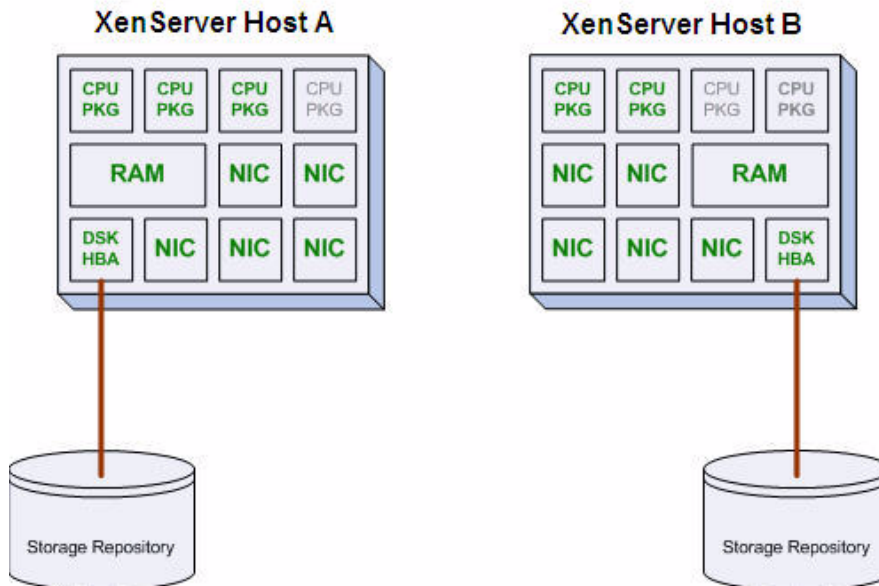


Figure 5 Redundant hosts on one site, with host-specific, independent storage.

NOTE: Two *separate* shared storage devices, such as iSCSI, can provide separate storage repositories (SRs) for everRun mirrored storage.

everRun mirrored storage requires separate virtual disks (XenServer virtual disk images, or VDIs) for redundancy. Storage repositories provide the greatest fault tolerance when implemented in *physically separate* storage subsystems.

However, host-specific storage can not support everRun Level 1 protected VMs, nor can it support the storage subsystem (path-redundant) feature for everRun Levels 3 and 2.

Storage Subsystem Data Protection

Storage subsystem data protection (the path-redundant alternative to everRun mirrored protection) relies exclusively on features of the storage hardware to provide fault resiliency (see [Figure 6](#)). In this type of data protection, everRun provides a redundant access path to common, shared storage, but everRun does not provide redundancy of the storage itself.

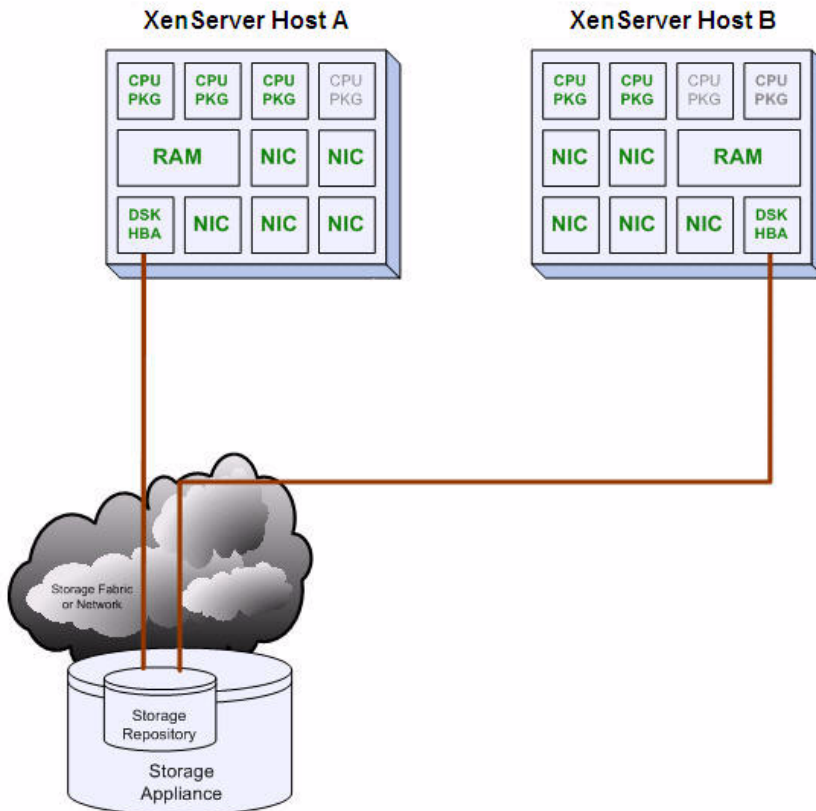


Figure 6 Redundant hosts at one site, using storage subsystem data protection.

Shared storage accessible to *all* hosts in the pool supports:

- Level 3 and Level 2 storage subsystem data protection, with path redundancy.
- Level 1 basic failover protection *requires* shared storage (see [Figure 6](#)).

everRun mirrored data protection requires separate virtual disks, and performs best when implemented in physically separate storage subsystems (see [Figure 7](#)).

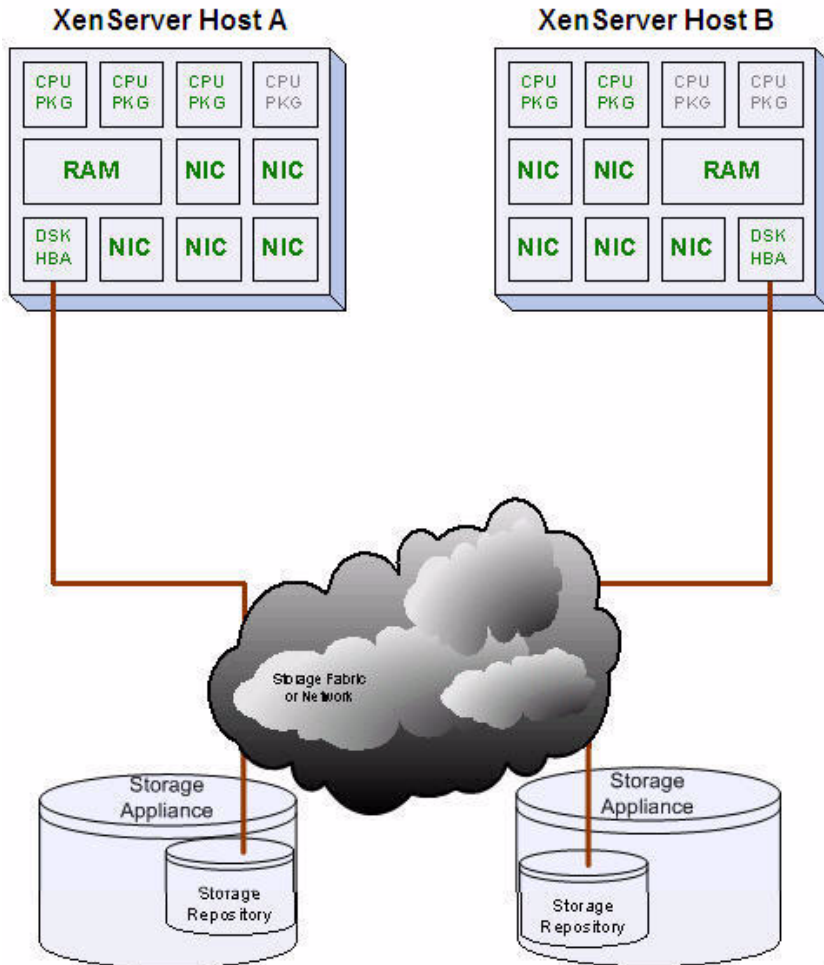


Figure 7 Redundant hosts with shared storage-based local storage repositories (SRs) in separate storage systems (can be NFS-based, iSCSI, or a SAN fabric).

Performance and Placement

You should consider performance needs and placement requirements when configuring SAN(s). Performance considerations include:

Performance Issue	Considerations
Resource contention	I/O intensive loads compete for the performance capabilities of host bus adapter (HBA) caches, requirements of the battery backup unit (BBU), pathways to the storage system, and the storage system itself.
Storage system	Controllers, caches, and physical access characteristics are shared among all hosts competing for common physical resources. Random patterns of access can be accommodated better than I/Os that are synchronized, such as dual reads/writes against everRun mirrored storage. For best performance with everRun mirror sets, plan to locate the mirror set members in physically different storage systems.
Storage fabric	As in storage systems, common switches and physical pathways perform best when loads are distributed randomly. For best performance with everRun mirror sets, plan to configure separate pathways through the storage network for each host (to access associated mirror set members).

Placement considerations include:

Placement Issue	Considerations
Single points of failure	Storage vendors minimize single points of failure. However, problems with firmware or storage networks can render storage systems inaccessible to all hosts at the same time. Separate storage systems and storage network configurations add insurance against inadvertently configuring a single point of failure.

Storage Rules

- ▶ *Do not configure any (iSCSI) storage access via the A-link networks.*
 - ▶ *When enabling XenServer HA, the state file (the file that maintains the pool heartbeat) cannot reside on NFS servers. Only FC-SAN and iSCSI are supported for this purpose. This limitation applies only to the XenServer state file. Storage for Level 3 and Level 2 PVMs can reside on NFS servers.*
-

Summary of Storage Options

- The number of storage systems provided, and their location, creates tradeoffs of availability, performance, and functionality. You may configure FC-SAN, iSCSI, or NFS servers as host-specific storage directly attached to a host, or as shared repositories available to all pool members. (EXCEPTION: An NFS server cannot host the XenServer state file in Level 1 PVMs.)
- In XenServer releases since version 5.0, iSCSI and NFS servers must not be configured on the management LAN or any IP subnet that the management LAN can route to. Therefore, a NIC must be added for the dedicated use of LAN-based storage architectures.
- It is recommended that you use two NICs for storage access — either bonded or configured as path-redundant (multipath) storage. When adding NICs, be aware of how many active NICs are accumulating within each host in this pool. An active/standby bonded pair counts as a single network toward the XenServer maximum of six active NICs.
- Independent storage systems for each pool member (based either on host-specific storage using everRun mirroring or on path-redundant storage subsystems providing data redundancy within the storage fabric) provide:
 - Highest levels of redundancy (data availability)
 - Higher performance (if using everRun mirror sets)

NOTE: Always configure large caches in RAID controllers and provide battery backup (BBU). The absence of BBU (or battery failure) will result in significant performance loss.

- XenServer VMs support storage volumes up to 2 TB in capacity, and the same limit is imposed on everRun-protected VMs.

Summary of Deployment Options

7



Overview

This chapter recaps the configurations suitable for everRun production deployments. It lists considerations for applying Level 3, Level 2, and Level 1 everRun protection on each configuration type. It then contrasts the supported configurations in a summary table, which rates availability features as Good, Better, and Best.

Configuration Options

All production deployments require hardware redundancy to deliver the value of everRun protection. Redundant physical hosts, storage, networking (and quorum service) are appropriately configured and deployed in support of the availability goals for your production environment.

The table in “Comparison of everRun Deployment Options” on page 37 compares the deployments in terms of their fault tolerance features. After selecting a deployment, refer to the corresponding appendix to develop a detailed plan for hardware configuration and integration with your infrastructure.

Meeting Your Availability Goals

everRun protection Levels 3 and 2 can be deployed on each of the documented configurations. Level 1 depends on shared storage and bonded NICs, and is compatible with configuration Options A and B.

NOTE: Option C bridges multiple sites. Attempts to configure shared storage across multiple sites, even if possible, cannot uphold the objectives of site-level fault tolerance.

The following sections list factors to consider when designing a system to protect VMs at selected levels of fault tolerance.

Protection Level 3 (System-Level Fault Tolerance)

- Choose Level 3 for applications that cannot be permitted to fail as a consequence of a (single) hardware failure in the pool, storage, or networks.
- Level 3 PVMs continue to run through failures of storage or networks — and even after a complete failure of a physical host.

Option A or D Local Two-Host*	Option B Local Multi-Host	Option C SplitSite Two-Host
<ul style="list-style-type: none"> • Level 3 protected VMs have no single point of failure — not the servers, nor the storage, nor the network devices. • You can deploy less critical, Level 2 PVMs on the same two hosts without compromising uptime of more important Level 3 PVMs. • In Option A, you can deploy less critical, Level 1 PVMs on the same two hosts without compromising uptime of more important Level 3 PVMs. ** 	<ul style="list-style-type: none"> • If you “misconfigure” Level 3 PVMs (by placing them across more than two hosts and enabling XenServer HA) the survival of all cannot be guaranteed in case of a management LAN segmentation. • You can deploy less critical, Level 2 PVMs to any hosts in the pool without impacting the ability to preserve hosts containing Level 3 PVMs. • You can deploy less critical, Level 1 PVMs on any hosts in the pool without impacting the ability to preserve hosts containing Level 3 PVMs.* 	<ul style="list-style-type: none"> • Level 3 Protected VMs have no single point of failure — including the room (site) where the servers, storage and network devices physically reside. • You can deploy less critical, Level 2 PVMs on the same two hosts without compromising uptime of Level 3 PVMs.

* If your deployment uses any Level 3 PVMs with two or more CPUs (SMP), 1Gbps of bandwidth for both A-link networks is required. Deployments using only uniprocessor Level 3 PVMs can operate with a minimum of 155 Mbps A-link network bandwidth.

** An exception occurs when failures already exist (for instance, when repairs are not completed) before a management LAN segmentation failure. This is a double failure scenario.

Protection Level 2 (Component-Level Fault Tolerance)

- Choose Level 2 for applications that must tolerate storage and network failures, even in cases when some XenServer host failures may require (protected) Windows migration (with reboot) to the other host.
- Level 2 is the next highest level of availability to Level 3 and supports multiprocessor Windows installations with a lower A-link bandwidth requirement.

Option A or D Local Two-Host	Option B Local Multi-Host	Option C SplitSite Two-Host
<ul style="list-style-type: none"> You can migrate applications to another host when the online host hardware fails, with only a momentary interruption for restart. You can deploy less critical, Level 1 PVMs on the same two hosts without compromising uptime of more important Level 3 or Level 2 PVMs.* 	<ul style="list-style-type: none"> If you "misconfigure" Level 2 PVMs by placing them on hosts that do not have the greatest number of Level 3 PVMs and by enabling XenServer HA, the Level 2 PVMs cannot be guaranteed if a management LAN segmentation occurs. You can deploy less critical, Level 1 PVMs on the same hosts without compromising uptime of more important Level 3 or Level 2 PVMs.* 	<ul style="list-style-type: none"> Failure is limited to an immediate restart when you lose an online host. All components (including the room where the servers, storage, and network devices physically reside) are protected against outage.

* An exception occurs when failures already exist (for instance, when repairs are not completed) before a management LAN segmentation failure. This is a double failure scenario.

Protection Level 1 (Basic Failover)

- Choose Level 1 protection for applications that can restart on another host (with reboot) if the host where it is running fails.
- Level 1 depends on shared storage to make its disks available to an alternate host.
- Network failures that disconnect a Level 1 PVM from a production LAN require manual intervention to online migrate the VM to an alternate host.

Option A Local Two-Host	Option B Local Multi-Host	Option C SplitSite Two-Host
<ul style="list-style-type: none"> Level 1 PVMs can be combined on the same two hosts without compromising uptime of more important Level 3 or Level 2 PVMs.* 	<ul style="list-style-type: none"> Level 1 PVMs can be combined on various hosts without compromising uptime of Level 3 or Level 2 PVMs.* 	<ul style="list-style-type: none"> Generally speaking, Level 1 PVMs cannot be deployed across multiple sites. Shared storage is typically located in a single physical computer room.

* An exception occurs when failures already exist (for instance, when repairs are not completed) before a management LAN segmentation failure. This is a double failure scenario.

Attempts to achieve conflicting goals within a single pool can lead to unanticipated compromise. Multiple small pools offer a better model than a single, larger pool that attempts to meet a broad range of availability requirements.

Surviving Management LAN Segmentation Failure

The likelihood of a segmented management LAN (failure) depends on the configuration of NICs, switches, and uplinks in the management LAN.

Use the guidelines in this section to ensure that applications protected at Level 3 and Level 2 will be able to survive potential management LAN segmentation failures.

Best Practices to Avoid LAN Segmentation Failure

- ☑ To ensure that *all Level 3 protected applications* survive a potential LAN segmentation:
 - 1.) Protect all Level 3 PVMs **on the same two hosts** (in pools of more than two hosts.) Protect VMs that have less critical uptime requirements on different hosts or in different pools.
 - OR–
 - 2.) Do not enable XenServer HA on the pool. *

- ☑ To ensure that *all Level 2 protected applications* survive a potential LAN segmentation:
 - 1.) Protect Level 2 PVMs **on the two hosts that contain the largest number of Level 3 PVMs**. Assign VMs that have less critical uptime requirements (e.g., Level 1 protection) to other hosts or to another pool.
 - OR–
 - 2.) Do not enable XenServer HA on the pool. *

*When XenServer HA is enabled and a segmented management LAN occurs, the hosts that contain the largest number of Level 3 PVMs will continue to run. Hosts connected to other segments of the management LAN will be reset.

Summary of Deployment Guidelines

Use the following guidelines when choosing your deployment.

Deployment Guidelines

- ▶ Comply with all best practices. Provide ample CPU cores and RAM to avoid overcommitting the resources. Overcommitting can produce configurations that cannot run or will not deliver to your availability expectations.

 - ▶ To insure all Level 3 and Level 2 PVMs survive segmentation of the management LAN, protect all Level 3 and the most important Level 2 PVMs on a common pair of hosts. Assign PVMs with less critical uptime requirements (Level 1 PVMs) to other hosts.

 - ▶ Take all measures available to minimize the likelihood of a segmented management LAN. Then, deploy PVMs under the assumption that it still can happen.

 - ▶ A bonded NIC configuration for the management LAN in each XenServer host is strongly recommended, and is required if XenServer HA is enabled.
-

- ▶ At times, you may best meet conflicting availability objectives by splitting the deployment into multiple pools (that is, by consolidating Level 1 PVMs into a dedicated pool or by reducing the number of PVMs to avoid overcommitting pool resources).

Comparison of everRun Deployment Options

The following chart contrasts the availability and performance considerations of the everRun deployment options described in this document. The legend describes the symbols used. Good, Better, and Best indicate relative fault tolerance and resiliency ratings.

Legend			
✓	Mandatory	✘	Not supported or not applicable
❖	Good	●	Better
		☺	Best

	Option A Local Two-Host Pool	Option B Local Multi-Host Pool	Option C Two-Host SplitSite Pool
General Redundancy	See Chapter 2 to review pool configuration fundamentals.		
2 Hosts	●	●	☺
More than 2 Hosts	✘	●	✘
everRun Protection Level 1	❖	❖	✘
everRun Protection Level 2	●	●	● ⁺
everRun Protection Level 3*	☺	☺	☺ ⁺

⁺ Adds physical site protection

* Level 3 adds virtual OS redundancy (in addition to I/O redundancy in Level 2)

	Option A Local Two-Host Pool	Option B Local Multi-Host Pool	Option C Two-Host SplitSite Pool
Networking	See Chapter 4 to review networking fundamentals.		
Private LAN Redundancy			
Direct-attached	✓	✘	✘
Two routable A-links	✓	✓	✓

Public (Production) LAN Redundancy			
everRun Protection Level 1	✘**	✘**	✘
everRun Protection Level 2	😊	😊	😊+
everRun Protection Level 3	😊	😊	😊+
Management LAN Redundancy			
XenServer NIC bonding	😊	😊	😊
Storage: Host-Specific SR(s)	See Chapter 6 to review storage fundamentals.		
Data Availability/ Integrity***			
everRun mirroring (Protection Level 1)	✘	✘	✘
everRun mirroring (Protection Level 2)	●	●	😊+
everRun mirroring (Protection Level 3)	●	●	😊+
Storage subsystem data protection (Protection Level 1)	✘	✘	✘
Storage subsystem data protection (Protection Level 2)	✘	✘	✘
Storage subsystem data protection (Protection Level 3)	✘	✘	✘
Performance***	😊	😊	●
Storage: Shared SR(s)	See Chapter 6 to review storage fundamentals.		
Data Availability/Integrity***			
everRun mirroring (Protection Level 1)	✘	✘	✘
everRun mirroring (Protection Level 2)	●	●	✘
everRun mirroring (Protection Level 3)	●	●	✘
Storage subsystem data protection (Protection Level 1)	✘	✘	✘
Storage subsystem data protection (Protection Level 2)	❖	❖	✘
Storage subsystem data protection (Protection Level 3)	❖	❖	✘
Performance***	●	●	✘
+ Adds physical site protection			
** Requires manual detection of the LAN issue and migration to another host (or restart of Windows and protected applications on another host)			
*** Storage configuration choices can alter ratings			

Option A: Local Two-Host Pool

A



Overview

Use this appendix to deploy Option A — a two-host pool in a single room with static IP addressing and point-to-point (private) A-link networking. If you plan to deploy Level 3 PVMs with SMP (also called FT-SMP) in a two-host pool, see Appendix D.

The XenServer hosts must be interconnected with short, crossover (or patch) Ethernet cables or fiber connections for A-link networks as shown in [Figure 8 on page 44](#). By following these guidelines, you minimize the possibility of simultaneous A-link network failures and reduce the likelihood of split-brain operation.

Choose a storage architecture for this configuration to meet the availability features and performance objectives of all VMs planned for deployment on the pool.

This appendix lists requirements for Option A. The best practice for each item is classified as either “Required” or “Recommended.” You can use these tables as worksheets to ensure that you have appropriate resources configured for Option A.

Software

Local Two-Host Pool

	Host A		Host B	
	Required	Recommended	Required	Recommended
Software*				
XenServer 5.5 Enterprise Essentials with Update 2 and trial license; use free XenServer 5.5 with Update 2 if you do not require XenServer HA.	✓		✓	
everRun pool-wide license (enables all hosts in pool)	✓			

* everRun Release Notes and Marathon Technical Support web portal lists officially required updates. (Log in at <http://support.marathontechnologies.com>).

Servers

Local Two-Host Pool

	Host A		Host B	
	Required	Recommended	Required	Recommended
Servers (Listed on Citrix XenServer HCL (http://hcl.xensource.com/?showall=no&subtab=systems))				
Approved servers on XenServer HCL	✓		✓	
Tested with XenServer 5.5, update 2	✓		✓	
CPUs				
Intel VT-capable CPUs (1.5 GHz minimum; virtualization enabled in BIOS; (same family, stepping, model in all hosts))*	✓		✓	
1 core for XenServer Dom0	✓		✓	
1 Add'l core for each Level 3 PVM		✓		✓
1 Add'l core for each Level 2 PVM		✓		✓
Add'l cores to meet requirements for all VMs planned. (Include non-protected VMs.)**		Qty: ____		Qty: ____
32 cores maximum per XenServer host				
RAM				
1 GB for XenServer	✓		✓	
Add'l RAM to meet full requirement of each Level 3 PVM ***	✓		✓	
Add'l RAM to meet full requirement of each Level 2 PVM***	✓			Qty: ____
128 MB add'l RAM for each Level 3 uniprocessor or Level 2 PVM (everRun AM)	✓		✓	
Add'l RAM to meet full requirement of each Level 1 PVM***	✓			Qty: ____
Add'l RAM to meet needs of each non-protected VM		Qty: ____		Qty: ____
128 GB RAM maximum per XenServer host				
<p>* A list of the currently supported processors for everRun VM Level 3 (both uniprocessor and SMP) and Level 2 protection is available at http://www.marathontechnologies.com/documents/Hardware_Qualifications.pdf. The everRun Compatibility Check utility, described in "What CPUs are supported?" on page 13, is a simple-to-use Marathon tool to prequalify your system hardware before installing XenServer or everRun software.</p>				
<p>** Allocate base core requirements, plus redundancy requirements on secondary host. The virtual CPU requirement must be no greater than 1.5 times the quantity of physical cores provided. Note that a 1-to-1 ratio is recommended, as specified in "CPU and Server Rules" on page 8 of this guide.</p>				
<p>*** Allocate identical RAM on the secondary host (the host that a Level 3 PVM is protected on, or that a Level 2 PVM will someday migrate to, or that a Level 1 PVM will someday fail over to). Level 2 PVMs cannot migrate online unless the identical RAM requirement is met on its standby host.</p>				

Networking

Local Two-Host Pool

	Host A		Host B	
	Required	Recommended	Required	Recommended
A-link 1 Network				
Crossover cable/fiber (as short as possible)	✓		✓	
Isolate from A-link 2 (physical cable route)	✓		✓	
1 Gbit, full-duplex NIC minimum	✓		✓	
10 Gbit, full-duplex NIC		✓		✓
Use default eth1		✓		✓
Default static IP assignments in subnet 192.168.3/24*		✓		✓
Avoid using common card (multiport NIC) with A-link 2		✓		✓
A-link 2 Network				
Crossover cable/fiber (as short as possible)	✓		✓	
Isolate from A-link 1 (physical cable route)	✓		✓	
1 Gbit, full-duplex NIC minimum	✓		✓	
10 Gbit, full-duplex NIC		✓		✓
Use default eth2		✓		✓
Default static IP assignments in subnet 192.168.4/24*		✓		✓
Avoid using common card (multiport NIC) with A-link 1		✓		✓
Management Network				
Configure on management vLAN		✓		✓
1 Gbit, full-duplex NIC minimum**	✓		✓	
10 Gbit, full-duplex NIC (optional)**		✓		✓
Use default eth0		✓		✓
Bonded eth0 with eth4 (use default)***		✓		✓
Assign a static IP configuration (address, gateway, DNS)		✓		✓
Separate NIC and IP subnet from any iSCSI or NFS servers	✓		✓	
Connect each NIC to separate switch(es) from other pool members		✓		✓
Connect bonded NIC to separate switch(es) from the other LAN in the bonded pair, and to same management vLAN		✓		✓
*Defaults will work for all direct-connected installations; however, it is preferable to assign your own local subnet in order to avoid LAN conflicts that could occur if the pool if you plan to integrate with Option B or C in the future.				
**All network components (switches, uplinks, etc.) must have >155 MBps minimum capacity end-to-end.				
***For maximum redundancy of the management LAN, Marathon strongly recommends using a bonded NIC on this network. If you plan to deploy Level 1 protection (XenServer HA) in your pool, a bonded NIC is required.				

Networking (cont.)

Local Two-Host Pool

	Host A		Host B	
	Required	Recommended	Required	Recommended
Configure gateway to production LAN for remote management		✓		✓
Specify an isolation IP address for automatic pool failover*	✓			
Production Network**				
Use default eth3		✓		✓
Configure on production vLAN		✓		✓
Same layer 2 multicast domain as other host(s) in the pool; DST MAC=01E009000005; DIX packet type 0x8806	✓		✓	
Same layer 2 multicast domain as other host(s) in the pool; DST MAC=01E009000001; DIX packet type 0x8805	✓		✓	
Same IP subnet as other host(s) in the pool	✓		✓	
Configure each NIC on separate switch(es) from the other pool member host		✓		✓
Configure gateway to management LAN for remote management		✓		✓
* Isolation IP is a pool-wide concept. Allocate only one address per pool, not an address for each XenServer host in the pool.				
** Second, third, fourth production LANs are optional. Repeat these requirements for each LAN configured.				

Storage

Local Two-Host Pool

	Host A		Host B	
	Required	Recommended	Required	Recommended
Storage				
Minimum 4 GB per XenServer host	✓		✓	
Shared SR*				
Add'l shared storage to meet full requirement of each VM protected at Level 3**	Qty: ____			Qty: ____
Add'l shared storage to meet performance, availability, and capacity of each VM protected at Level 2**	Qty: ____			Qty: ____
Add'l shared storage to meet full requirement of each VM protected at Level 1**	Qty: ____			
Host-Specific SR				
Add'l host-specific storage to meet performance, availability, and capacity of each VM protected at Level 3***	Qty: ____		Qty: ____	
Add'l host-specific storage to meet performance, availability, and capacity of each VM protected at Level 2***	Qty: ____		Qty: ____	
<p>* Allocate storage on a shared SR once only (not on both hosts). Please note: Do not choose storage subsystem data protection for this release; see the Release Notes for details.</p> <p>** Allocate storage on a shared SR or a host-specific SR — not both. IF configuring storage subsystem data protection on a shared SR as an alternative to an everRun mirrored disk, do not allocate storage for the second host. IF configuring host-specific everRun mirrored disks, allocate storage on two hosts and configure to meet performance needs.</p> <p>*** Allocate storage on a shared SR or a host-specific SR — not both. To configure everRun disk mirroring, allocate storage on two separate hosts.</p>				

System Connectivity

Figure 8 on page 44 shows one way to comply with best practices and set up a local two-host pool.

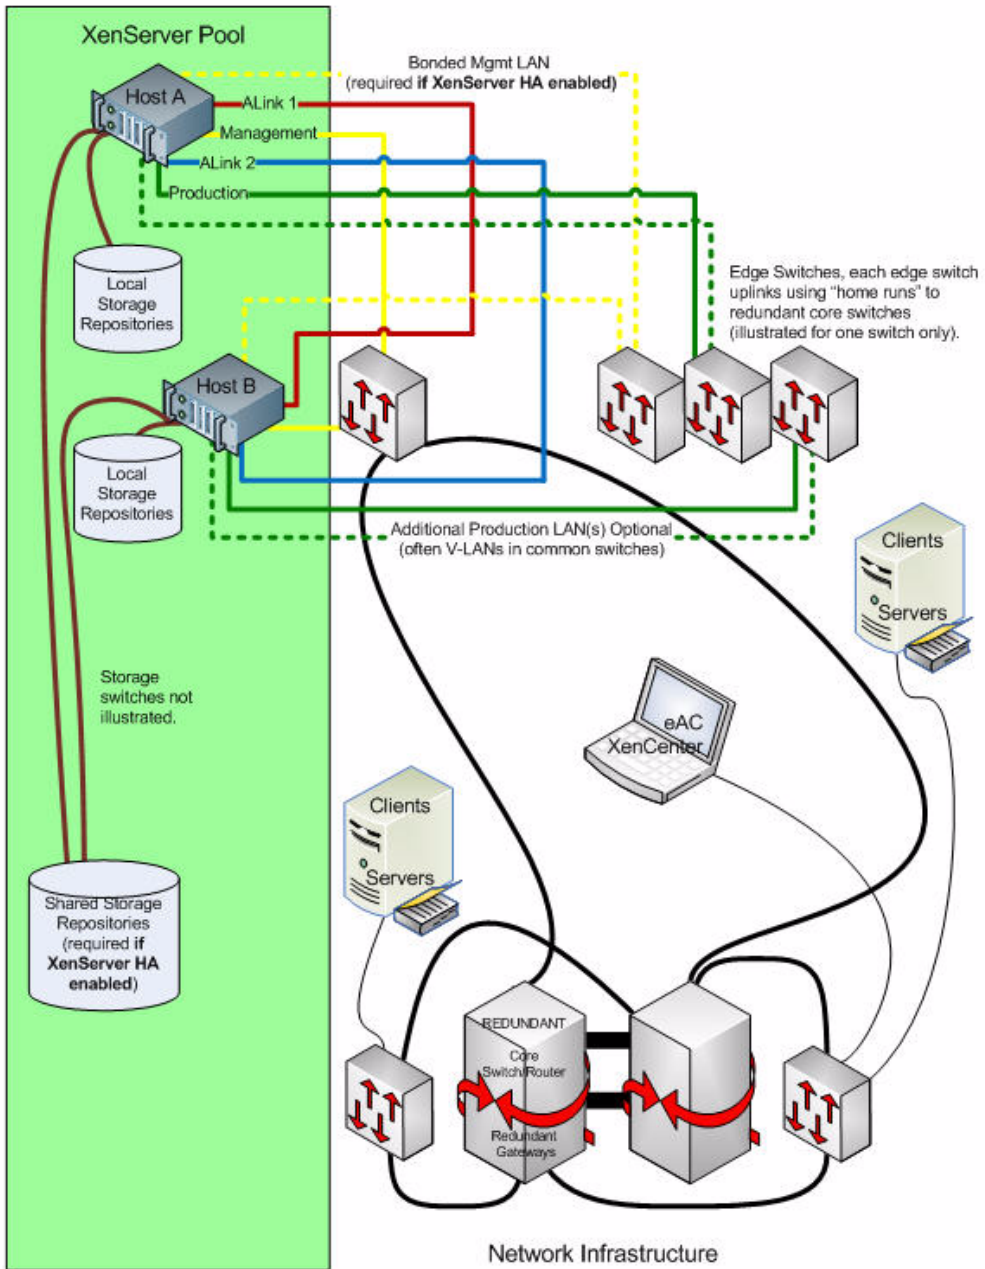


Figure 8 Connectivity for local two-host pool.

Option B: Local Multi-Host Pool

B



Overview

Use this appendix only to deploy Option B — a pool of two to four hosts on a single site, with active, switched A-link networking and quorum service. In this option, up to four XenServer hosts are located in a common computer room environment.

NOTE: Option B does not support PVMs protected with Level 3 SMP in everRun 6.0.

This deployment requires quorum service configured in compliance with best practices. Ideally, a preferred quorum service computer is located at a third facility and an alternate is located at a fourth site (or carefully placed at the third site). The networks are interconnected as shown in [Figure 9 on page 51](#).

This appendix lists requirements for Option B. The best practice for each item is classified as either “Required” or “Recommended.” You can use these tables as worksheets to ensure that you have appropriate resources configured for Option B.

Software

Local Multi-Host Pool

	Host A		Hosts B, C, and D	
	Required	Recommended	Required	Recommended
Software*				
XenServer 5.5 Enterprise Essentials with Update 2 and trial license; use free XenServer 5.5 with Update 2 if you do not require XenServer HA.	✓		✓	
everRun pool-wide license (enables all hosts in pool)	✓			

* everRun Release Notes and Marathon Technical Support web portal list officially required updates. (Log in at <http://support.marathontechnologies.com>).

Servers

Local Multi-Host Pool

	Host A		Hosts B, C, and D	
	Required	Recommended	Required	Recommended
Servers (Listed on Citrix XenServer HCL (http://hcl.xensource.com/?showall=no&subtab=systems))				
Approved servers on XenServer HCL	✓		✓	
Tested with XenServer 5.5	✓		✓	
CPUs				
Intel VT-capable CPUs (1.5 GHz minimum; virtualization enabled in BIOS; (same family, stepping, model in all hosts)*)	✓		✓	
1 core for XenServer Dom0	✓		✓	
1 Add'l core for each Level 3 PVM		✓		✓
1 Add'l core for each Level 2 PVM		✓		✓
Add'l cores to meet requirements for all VMs planned. (Include non-protected VMs.)**		Qty: ____		Qty: ____
32 cores maximum per XenServer host				
RAM				
1 GB for XenServer	✓		✓	
Add'l RAM to meet full requirement of each Level 3 PVM ***	✓		Qty: ____	
Add'l RAM to meet full requirement of each Level 2 PVM ***	✓			Qty: ____
128 MB add'l RAM for each Level 2 or Level 3 uniprocessor PVM (everRun AM)	✓		✓	
Add'l RAM to meet full requirement of each Level 1 PVM ***	Qty: ____			
Add'l RAM to meet needs of each non-protected VM	Qty: ____			
128 GB RAM maximum per XenServer host				
<p>* A list of the currently supported processors for everRun VM Level 3 (both uniprocessor and SMP) and Level 2 protection is available at http://www.marathontechnologies.com/documents/Hardware_Qualifications.pdf</p> <p>The everRun Compatibility Check utility, described in "What CPUs are supported?" on page 13, is a simple-to-use Marathon tool to prequalify your system hardware before installing XenServer or everRun software.</p>				
<p>** Allocate base core requirements, plus redundancy requirements on secondary host. The virtual CPU requirement must be no greater than 1.5 times the quantity of physical cores provided. However, a 1-to-1 ratio is recommended, as specified in "CPU and Server Rules" on page 8 of this guide.</p>				
<p>*** Allocate identical RAM on the secondary host (the host that a Level 3 PVM is protected on, or that a Level 2 PVM will someday migrate to, or that a Level 1 PVM will someday fail over to). Level 2 PVMs cannot migrate online unless the identical RAM requirement is met on its standby host.</p>				

Quorum Service

Local Multi-Host Pool

	Preferred Quorum Computer		Alternate Quorum Computer		Pool Configuration	
	Req'd	Recom'd	Req'd	Recom'd	Req'd	Recom'd
Latency < 500ms RTT to all everRun hosts	✓		✓		✓	
UDP port 2189 open from all hosts to quorum computer	✓		✓		✓	
UDP port 2188 open from quorum computer to all hosts	✓		✓		✓	
Computer running Windows Server 2003, Windows XP (SP2 or newer), Windows Vista or Windows 7; always powered on	✓		✓			
Static IP configuration (quorum computers and A-link connections from all AMs in all PVMs)	✓		✓		✓	
Connect to production LAN (connectivity representative of production clients)		✓		✓		
Place at third site (preferred quorum), fourth site or third site (alternate quorum). Different network switch and power source from any host (or other quorum)*		✓		✓		
Access via A-link 1 (gateway)**		✓				✓
Access via A-link 2 (gateway)**				✓		✓
Minimize common switches and uplinks between quorum computer and multiple host sites		✓		✓		✓
Maximum PVMs serviced: one pool		✓		✓		
Static IP configuration	✓		✓		✓	
Connect to production LAN (connectivity representative of production clients)		✓		✓		

*Quorum computers should have as little in common as possible. If both must be located at a common site with the XenServer hosts, care must be taken to minimize dependencies on common power sources or common Ethernet switches to reach the XenServer pool member (and master) hosts. See Chapter 5 for an overview of quorum considerations.

**Physical connectivity between the quorum computers must not route through a third site if deployed with remote locations of quorum computers.

NOTE: Ports 2189, 2188 may be reconfigured from the Marathon defaults in the rare event that you need to accommodate your network infrastructure. Use the instructions in the *everRun Setup and Installation Guide*.

Networking

Local Multi-Host Pool

	Host A		Hosts B, C, and D	
	Required	Recommended	Required	Recommended
A-link 1 Network				
Connected to separate switch from A-link 2 and quorum computers		✓		✓
1 Gbps, full-duplex NIC minimum*	✓		✓	
10 Gbps, full-duplex NIC optional		✓		✓
Max latency RTT 10 ms	✓		✓	
Min bandwidth 155 Mbps per Level 3 uniprocessor or Level 2 PVM*	✓		✓	
Configure your IP address range instead of using the everRun default (92.168.3/24) for the pool**		✓		✓
Same IP subnet as other A-link 1 connections		✓		✓
Different IP subnet from A-link 2 connections	✓		✓	
Use default eth1		✓		✓
Avoid using common card (multiport NIC) with A-link 2		✓		✓
A-link 2 Network				
Connected to separate switch from A-link 1 and quorum computers		✓		✓
1 Gbps, full-duplex NIC minimum*	✓		✓	
10 Gbps, full-duplex NIC optional		✓		✓
Max latency RTT 10 ms	✓		✓	
Min bandwidth 155 Mbps per Level 3 uniprocessor or Level 2 PVM*	✓		✓	
Configure your IP address range instead of using the everRun default (92.168.3/24) for the pool**		✓		✓
Same IP subnet as other A-link 2 connections*		✓		✓
Different IP subnet from A-link 1 connections*	✓		✓	
Use default eth2		✓		✓
Avoid using common card (multiport NIC) for A-link 1		✓		✓
Management Network				
Configure on management vLAN		✓		✓
1 Gbps, full-duplex NIC minimum*	✓		✓	
10 Gbps, full-duplex NIC (optional)		✓		✓
* All network components (switches, uplinks, etc.) must have 155 MBps or greater minimum capacity end-to-end.				
** Assigning your own local subnet is recommended, to avoid LAN conflicts possible with Marathon's default values.				

Networking (cont.)

Local Multi-Host Pool

	Host A		Hosts B, C, and D	
	Required	Recommended	Required	Recommended
Management Network (cont'd)				
Use default eth0		✓		✓
Bonded eth0 with eth4 (use default) *		✓		✓
Assign a static IP configuration (address, gateway, DNS)		✓		✓
Separate NIC and IP subnet from any iSCSI or NFS servers	✓		✓	
Connect each NIC to separate switch(es) from the other NIC in the bonded pair, and to the same management vLAN		✓		✓
Connect each NIC to separate switch(es) from other pool members and to same management vLAN		✓		✓
Configure gateway to production LAN for remote management		✓		✓
Specify an isolation IP address for automatic pool failover **	✓			
Production Network***				
Use default eth3		✓		✓
Configure on production vLAN	✓		✓	
Same layer 2 multicast domain as other host(s) in the pool; DST MAC=01E009000005; DIX packet type 0x8806	✓		✓	
Same layer 2 multicast domain as other host(s) in the pool; DST MAC=01E009000001; DIX packet type 0x8805	✓		✓	
Same IP subnet as other host(s) in the pool	✓		✓	
Configure each NIC on separate switch(es) from the other pool member host; same production/client vLAN		✓		✓
Configure gateway to management LAN for remote management		✓		✓
Configure each NIC on separate switch(es) from the other pool member host; same production/client vLAN		✓		✓
* For maximum redundancy of the management LAN, Marathon strongly recommends using a bonded NIC on this network. If you plan to deploy Level 1 protection (XenServer HA) in your pool, a bonded NIC is required.				
** Isolation IP is a pool-wide concept. Allocate only one IP address per pool; do not allocate an address for each XenServer host in the pool. When quorum service is enabled in this configuration, it overrides pool master failover.				
*** Second, third, and fourth production LANs are optional. Repeat these requirements for each LAN configured.				

Storage

Local Multi-Host Pool

	Host A		Hosts B, C, and D	
	Required	Recommended	Required	Recommended
Storage*				
Minimum 4 GB for XenServer	✓		✓	
Shared SR				
Add'l shared storage to meet full requirement of each VM protected at Level 3**	Qty: ____			Qty: ____
Add'l shared storage to meet performance, availability, and capacity of each VM protected at Level 2**	Qty: ____			Qty: ____
Add'l shared storage to meet performance, availability, and capacity of each VM protected at Level 1**	Qty: ____			
Host-Specific				
Add'l shared storage to meet performance, availability, and capacity of each VM protected at Level 3***	Qty: ____		Qty: ____	
Add'l shared storage to meet performance, availability, and capacity of each VM protected at Level 2***	Qty: ____		Qty: ____	
<p>* Allocate storage on a shared SR once only (not on both hosts). Please note: Do not choose storage subsystem data protection for this release; see the Release Notes for details.</p> <p>** Allocate storage on a shared SR or a host-specific SR — not both. IF configuring storage subsystem data protection on a shared SR as an alternative to an everRun mirrored disk, do not allocate storage for the second host. IF configuring host-specific everRun mirrored disks, allocate storage on two hosts and configure to meet performance needs.</p> <p>*** Allocate storage on a shared SR or a host-specific SR — not both. To configure everRun disk mirroring, allocate storage on two separate hosts.</p>				

System Connectivity

Figure 9 on page 51 shows one way to comply with best practices and set up a local multi-host pool.

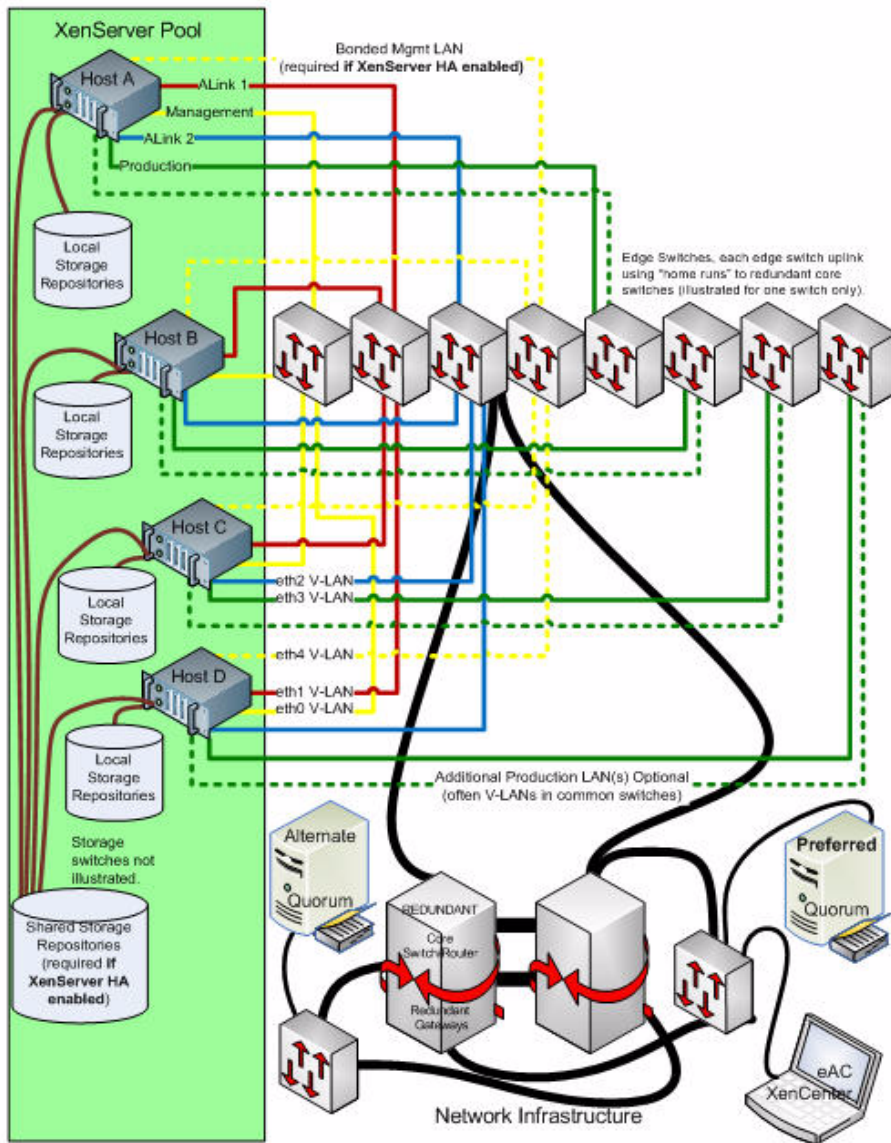


Figure 9 Connectivity for switched A-link, multi-host pool.

Option C: Two-Host SplitSite Pool

C



Overview

Use this appendix only to deploy Option C — a pool of two hosts on two sites (one per site), with active, switched A-link networking and quorum service. In this SplitSite[®] option, two XenServer hosts are located in separate buildings with geographic separation and LANs over wide area networking configured between them.

NOTE: Option C does not support PVMs protected with Level 3 SMP in everRun 6.0.

Quorum service is required and configured in compliance with best practices. A preferred quorum service computer is (ideally) located in a third facility and an alternate is located in a fourth site (or carefully placed in the third). The networks are interconnected as shown in [Figure 10 on page 59](#).

Locally attached storage and Marathon mirrored disks are configured, because the benefits of site level redundancy are realized only when storage is local to each site.

The best practice for each item is classified as either “Required” or “Recommended.” Use these tables to ensure that you have appropriate resources configured for Option C.

Software

Two-Host SplitSite Pool

	Host A		Host B	
	Required	Recommended	Required	Recommended
Software*				
XenServer 5.5 Enterprise Essentials and trial license; use free XenServer 5.5 with update 2 if you do not require XenServer HA.	✓		✓	
everRun pool-wide license (enables all hosts in pool)	✓			

* everRun Release Notes and Marathon Technical Support web portal list officially required updates. (Log in at <http://support.marathontechnologies.com>)

Servers

Two-Host SplitSite Pool

	Host A		Host B	
	Required	Recommended	Required	Recommended
Servers (Listed on Citrix XenServer HCL (http://hcl.xensource.com/?showall=no&subtab=systems))				
Approved servers on XenServer HCL	✓		✓	
Tested with XenServer 5.5	✓		✓	
CPUs				
Intel VT-capable CPUs (1.5 GHz minimum; virtualization enabled in BIOS; (same family, model in all hosts)*)	✓		✓	
1 core for XenServer Dom0	✓		✓	
1 Add'l core for each Level 3 PVM		✓		✓
1 Add'l core for each Level 2 PVM		✓		✓
Add'l cores to meet requirements for all VMs planned. (Include non-protected VMs.)**		Qty: ____		Qty: ____
32 cores maximum per XenServer host				
RAM				
1 GB for XenServer	✓		✓	
Add'l RAM to meet full requirement of each Level 3 PVM***	✓		✓	
Add'l RAM to meet full requirement of each Level 2 PVM***	✓			Qty: ____
128 MB add'l RAM for each Level 3 uniprocessor or Level 2 PVM (everRun AM)	✓		✓	
Add'l RAM to meet needs of each non-protected VM		Qty: ____		
128 GB RAM maximum per XenServer host				
<p>* A list of the currently supported processors for everRun VM Level 3 (both uniprocessor and SMP) and Level 2 protection is available at http://www.marathontechnologies.com/documents/Hardware_Qualifications.pdf</p> <p>The everRun Compatibility Check utility, described in "What CPUs are supported?" on page 13, is a simple-to-use Marathon tool to prequalify your system hardware before installing XenServer or everRun software.</p>				
<p>** Allocate base core requirements, plus redundancy requirements on secondary host. The virtual CPU requirement must be no greater than 1.5 times the quantity of physical cores provided. However, a 1-to-1 ratio is recommended, as specified in "CPU and Server Rules" on page 8 of this guide.</p>				
<p>*** Allocate identical RAM on the secondary host (the host that a Level 3 PVM is protected on, or that a Level 2 PVM will someday migrate to, or that a Level 1 PVM will someday fail over to). Level 2 PVMs cannot migrate online unless the identical RAM requirement is met on its standby host.</p>				

Quorum Service

Two-Host SplitSite Pool

	Preferred Quorum Computer		Alternate Quorum Computer		Pool Configuration	
	Req'd	Recom'd	Req'd	Recom'd	Req'd	Recom'd
Latency < 500ms RTT to all everRun hosts	✓		✓		✓	
UDP port2189 open from all hosts to quorum computer	✓		✓		✓	
UDP port2188 open from quorum computer to all hosts	✓		✓		✓	
Computer running Windows Server 2003, Windows XP (SP2 or newer), Windows Vista or Windows 7; always powered on	✓		✓			
Static IP configuration (quorum computers and A-link connections from all AMs in all PVMs)	✓		✓		✓	
Connect to production LAN (connectivity representative of production clients)		✓		✓		
Place at third site (preferred quorum), fourth site or third site (alternate quorum). Different network switch and power source from any host (or other quorum)	✓			✓		
Access via A-link 1 (gateway)**		✓				✓
Access via A-link 2 (gateway)**				✓		✓
Minimize common switches and uplinks between quorum computer and multiple hosts		✓		✓		✓
Maximum PVMs serviced: one pool		✓		✓		
Static IP configuration	✓		✓		✓	
Connect to production LAN (connectivity representative of production clients)		✓		✓		

* Quorum computers should have as little in common as possible. If both must be located at a common (third) site, care must be taken that they don't depend on common power sources or common Ethernet switches to reach the locations (sites) of the pool member hosts. Review Chapter 5 for an overview of quorum considerations.

** Physical connectivity between the quorum computers must not route through either production site to reach the other.

NOTE: Ports 2189, 2188 may be reconfigured from the Marathon defaults in the rare event that you need to accommodate your network infrastructure. See the *everRun Setup and Installation Guide* for instructions.

Networking

Two-Host SplitSite Pool

	Host A		Host B	
	Required	Recommended	Required	Recommended
A-link 1 Network				
Connected to separate switch from other host (same pool) and quorum computers		✓		✓
1 Gbps, full-duplex NIC minimum*	✓		✓	
10 Gbps, full-duplex NIC optional*		✓		✓
Max latency RTT 10 ms	✓		✓	
Min bandwidth 155 Mbps per Level 3 uniprocessor or Level 2 PVM**	✓		✓	
Configure your IP address range instead of using the everRun default (192.168.3/24) for the pool		✓		✓
Different IP subnet from other A-link 1 connection**		✓		✓
Different IP subnet(s) from A-link 2 connection**	✓		✓	
Use default eth1		✓		✓
Avoid using common card (multiport NIC) with A-link 2		✓		✓
A-link 2 Network				
Connected to separate switch from other host (same pool) and quorum computers		✓		✓
1 Gbps, full-duplex NIC minimum*	✓		✓	
10 Gbps, full-duplex NIC optional*		✓		✓
Max latency RTT 10 ms	✓		✓	
Min bandwidth 155 Mbps per Level 3 uniprocessor or Level 2 PVM**	✓		✓	
Configure your IP address range instead of using the everRun default (192.168.3/24) for the pool		✓		✓
Different IP subnet from other A-link 2		✓		✓
Different IP subnet(s) from A-link 1	✓		✓	
Use default eth2		✓		✓
Avoid using common card (multiport NIC) with A-link 1		✓		✓
* All network components (switches, uplinks, etc.) must provide 155 MBps or greater capacity per PVM.				
** Assigning your own local subnet is recommended, to avoid LAN conflicts possible with Marathon's default values.				

NOTE: Switched/routed A-link configurations often preclude configuration of Level 3 SMP due to bandwidth limitations of switches and uplinks.

Networking (cont.)

Two-Host SplitSite Pool

	Host A		Host B	
	Required	Recommended	Required	Recommended
Management Network				
Configure on management vLAN		✓		✓
1 Gbps, full-duplex NIC minimum*	✓		✓	
10 Gbps, full-duplex NIC optional*		✓		✓
Use default eth0		✓		✓
Bonded eth0 with eth4 (use default)**		✓		✓
Assign a static IP configuration (address, gateway, DNS)		✓		✓
Separate NIC and IP subnet from any iSCSI or NFS servers	✓		✓	
Connect each NIC to separate switch(es) from other pool members and to same management vLAN		✓		✓
Configure gateway to production LAN for remote management		✓		✓
Specify an isolation IP address for automatic pool failover***	✓			
Production Network****				
Use default eth3		✓		✓
Configure on production vLAN		✓		✓
Same layer 2 multicast domain as other host(s) in the pool; DST MAC=01E009000005; DIX packet type 0x8806	✓		✓	
Same layer 2 multicast domain as other host(s) in the pool; DST MAC=01E009000001; DIX packet type 0x8805	✓		✓	
Same IP subnet as other host(s) in the pool	✓		✓	
Configure each NIC on separate switch(es) from the other pool member host; same production/client vLAN		✓		✓
*All network components (switches, uplinks, etc.) must have 155 MBps or greater capacity per PVM.				
** For maximum redundancy of the management LAN, Marathon strongly recommends using a bonded NIC on this network. If you plan to deploy Level 1 protection (XenServer HA) in your pool, a bonded NIC is required.				
*** Isolation IP is a pool-wide concept. Allocate only one IP address per pool; do not allocate an address for each XenServer host in the pool. When quorum service is enabled in this configuration, it overrides pool master failover.				
****Second, third, fourth production LANs are optional. Repeat these requirements for each LAN configured.				

Storage

Two-Host SplitSite Pool

	Host A		Host B	
	Required	Recommended	Required	Recommended
Storage				
Minimum 4 GB for XenServer	✓		✓	
Host-Specific*				
Add'l local storage to meet performance, availability, and capacity of each VM protected at Level 3**	Qty: ____		Qty: ____	
Add'l local storage to meet performance, availability, and capacity of each VM protected at Level 2**	Qty: ____		Qty: ____	
* Do not configure shared storage. Host-specific SRs created from separate storage appliances (direct-attached or FC-SAN) are supported for a SplitSite configuration. See Chapter 6 to review overall storage requirements and considerations.				
** Allocate storage on two hosts (configure as an everRun mirrored disk).				

NOTE: If you need help in planning a SplitSite configuration to meet your availability needs, consult your Marathon representative.

System Connectivity

Figure 10 on page 59 shows one way to comply with best practices and set up a two-host SplitSite pool.

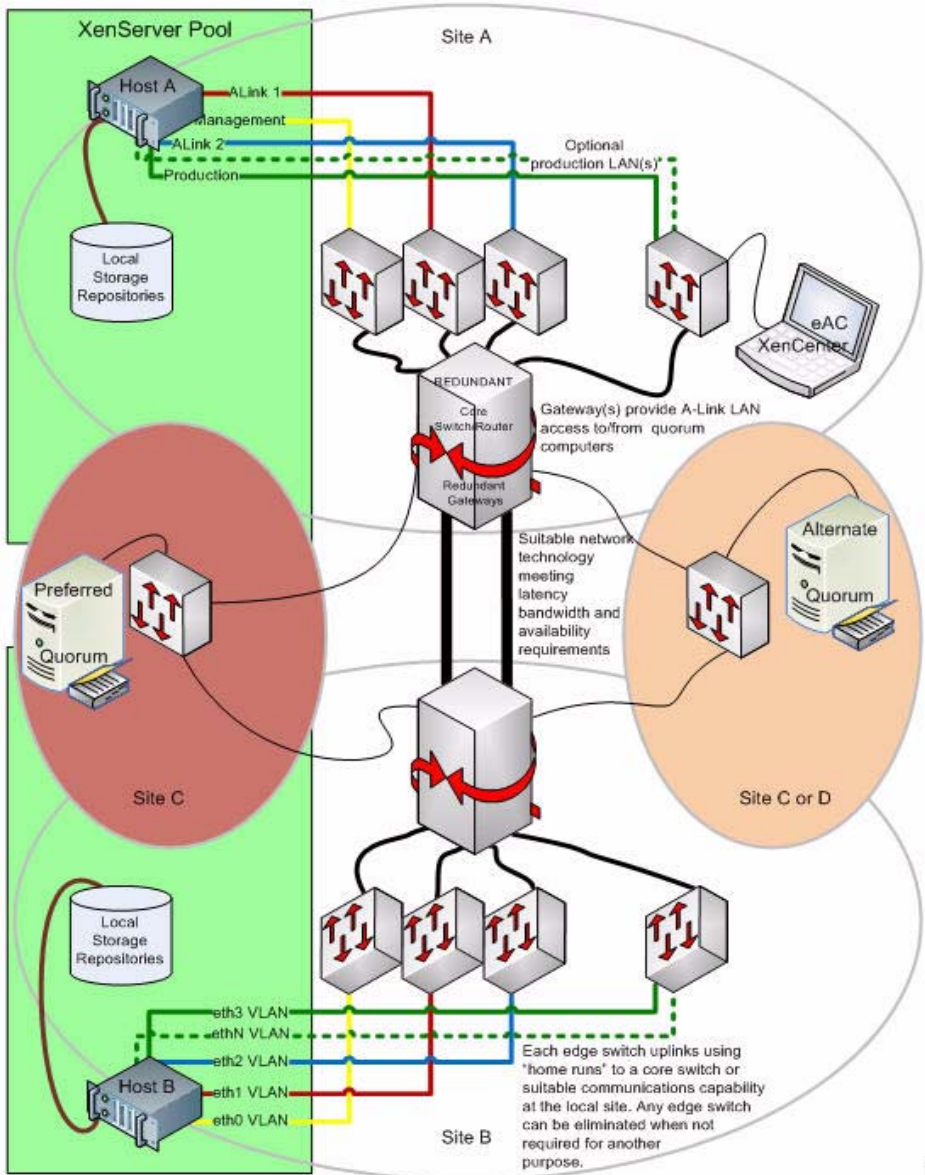


Figure 10 Connectivity for two-host SplitSite pool.

Option D: Two-Host Pool with Level 3 SMP

D



Overview

Use this appendix to deploy Option D — a two-host local pool that contains Level 3 PVms running in symmetric multiprocessor (SMP) mode (also called FT-SMP). The hosts must be interconnected using short, crossover (or patch) Ethernet cables or fiber connections for A-link networks, as shown in Appendix A, [Figure 8 on page 44](#).

NOTE: For Option D, choose applications with process scalability, not network latency, requirements. Please contact your Marathon technical representative if you have configuration questions.

This appendix lists requirements for Option D. The best practice for each item is classified as either “Required” or “Recommended.” You can use these tables as worksheets to ensure that you have appropriate resources configured for Option D.

Software

Local Two-Host Pool with Level 3 SMP PVM(s)

	Host A		Host B	
	Required	Recommended	Required	Recommended
Software*				
XenServer 5.5 Enterprise Essentials with Update 2 and trial license; use free XenServer 5.5 with Update 2 if you do not require XenServer HA.	✓		✓	
everRun pool-wide license (enables all hosts in pool)	✓			

* everRun Release Notes and Marathon Technical Support web portal lists officially required updates. (Log in at <http://support.marathontechnologies.com>).

Servers

Local Two-Host Pool with Level 3 SMP PVMs

	Host A		Host B	
	Required	Recommended	Required	Recommended
Servers (Listed on Citrix XenServer HCL (http://hcl.xensource.com/?showall=no&subtab=systems))				
Approved servers on XenServer HCL	✓		✓	
Tested with XenServer 5.5, update 2	✓		✓	
CPUs				
Intel E55xx (Nehalem) 2.26 GHz or better OR Intel E56xx (Westmere) 2.4 GHz or better (same processors on both hosts)	✓		✓	
Two processor configuration with four cores per processor minimum	✓		✓	
Start with one FT-SMP PVM per pool; depending on results, add a second*		✓		✓
1 core for XenServer Dom0	✓		✓	
4 Add'l cores for each Level 3 PVM		✓		✓
1 Add'l core for each Level 2 PVM		✓		✓
Add'l cores to meet requirements for all VMs planned. (Include non-protected VMs.)**		Qty: ____		Qty: ____
32 cores maximum per XenServer host				
RAM				
1333 MHz Dual-Ranked RDIMMS configured for 3-way interleaving (6 DIMMS per server), or as needed to achieve 3-way interleaving	✓		✓	
1 GB for XenServer	✓		✓	
Add'l RAM to meet full requirement of each Level 3 PVM ***	✓		✓	
Add'l RAM to meet full requirement of each Level 2 PVM***	✓			Qty: ____
128 MB add'l RAM for each Level 3 uniprocessor or Level 2 PVM (everRun AM)	✓		✓	
512 MB add'l RAM for each Level 3 SMP (everRun AM)	✓		✓	
Add'l RAM to meet full requirement of each Level 1 PVM***	✓			Qty: ____
Add'l RAM to meet needs of each non-protected VM	Qty: ____		Qty: ____	
128 GB RAM maximum per XenServer host				

* If two Level 3 (FT-SMP) PVMs are running in the same two-host pool, it is recommended that you choose a different host as the primary (active) host for each PVM. Do not make the same host the primary side for both PVMs.

** Allocate base core requirements, plus redundancy requirements on secondary host. The virtual CPU requirement must be no greater than 1.5 times the quantity of physical cores provided. Note that a 1-to-1 ratio is recommended, as specified in "CPU and Server Rules" on page 8 of this guide.

*** Remember that each Level 3 SMP (FT-SMP) PVM uses 512 additional RAM. Allocate identical RAM on the second host where a Level 3 PVM is protected; the host that a Level 2 PVM might someday migrate to; or the host that a Level 1 PVM may fail over to. Level 2 PVMs cannot migrate online unless an identical amount of RAM is allocated on its standby host.

Networking

Local Two-Host Pool with Level 3 SMP PVMs

	Host A		Host B	
	Required	Recommended	Required	Recommended
A-link 1 Network				
Crossover cable/fiber (as short as possible)	✓		✓	
Isolate from A-link 2 (physical cable route)	✓		✓	
1 Gbit, full-duplex NIC minimum	✓		✓	
10 Gbit, full-duplex NIC*		✓		✓
Use default eth1		✓		✓
Default static IP assignments in subnet 192.168.3/24**		✓		✓
Avoid using common card (multiport NIC) with A-link 2		✓		✓
A-link 2 Network				
Crossover cable/fiber (as short as possible)	✓		✓	
Isolate from A-link 1 (physical cable route)	✓		✓	
1 Gbit, full-duplex NIC minimum	✓		✓	
10 Gbit, full-duplex NIC*		✓		✓
Use default eth2		✓		✓
Default static IP assignments in subnet 192.168.4/24**		✓		✓
Avoid using common card (multiport NIC) with A-link 1		✓		✓
Management Network				
Configure on management vLAN		✓		✓
1 Gbit, full-duplex NIC minimum***	✓		✓	
10 Gbit, full-duplex NIC (optional)***		✓		✓
Use default eth0		✓		✓
Bonded eth0 with eth4 (use default)****		✓		✓
Assign a static IP configuration (address, gateway, DNS)		✓		✓
Separate NIC and IP subnet from any iSCSI or NFS servers	✓		✓	
Connect each NIC to separate switch(es) from other pool members		✓		✓
<p>* 10 Gbps NICs and switches are highly recommended if the deployment includes Level 3 SMP protected VMs. Contact your Marathon technical representative for more information about specific components recommended.</p> <p>** Defaults will work for all direct-connected installations; however, it is preferable to assign your own local subnet in order to avoid LAN conflicts that could occur if the pool if you plan to integrate with Option B or C in the future.</p> <p>*** All network components (switches, uplinks, etc.) must have 1 Gbps end-to-end for Level 3 SMP.</p> <p>**** For maximum redundancy of the management LAN, Marathon strongly recommends using a bonded NIC on this network.</p>				

Networking (cont.)

Local Two-Host Pool with Level 3 SMP PVMs

	Host A		Host B	
	Required	Recommended	Required	Recommended
Connect bonded NIC to separate switch(es) from the other LAN in the bonded pair, and to same management vLAN		✓		✓
Configure gateway to production LAN for remote management		✓		✓
Specify an isolation IP address for automatic pool failover*	✓			
Production Network**				
Use default eth3		✓		✓
Configure on production vLAN		✓		✓
Same layer 2 multicast domain as other host(s) in the pool; DST MAC=01E009000005; DIX packet type 0x8806	✓		✓	
Same layer 2 multicast domain as other host(s) in the pool; DST MAC=01E009000001; DIX packet type 0x8805	✓		✓	
Same IP subnet as other host(s) in the pool	✓		✓	
Configure each NIC on separate switch(es) from the other pool member host		✓		✓
Configure gateway to management LAN for remote management		✓		✓
* Isolation IP is a pool-wide concept. Allocate only one address per pool, not an address for each XenServer host in the pool.				
** Second, third, fourth production LANs are optional. Repeat these requirements for each LAN configured.				

Storage

Local Two-Host Pool

	Host A		Host B	
	Required	Recommended	Required	Recommended
Storage				
Minimum 4 GB per XenServer host	✓		✓	
Shared SR*				
Add'l shared storage to meet full requirement of each VM protected at Level 3**	Qty: ____			Qty: ____
Add'l shared storage to meet performance, availability, and capacity of each VM protected at Level 2**	Qty: ____			Qty: ____
Add'l shared storage to meet full requirement of each VM protected at Level 1**	Qty: ____			
Host-Specific SR				
Add'l host-specific storage to meet performance, availability, and capacity of each VM protected at Level 3***	Qty: ____		Qty: ____	
Add'l host-specific storage to meet performance, availability, and capacity of each VM protected at Level 2***	Qty: ____		Qty: ____	
<p>* Allocate storage on a shared SR once only (not on both hosts). Please note: Do not choose storage subsystem data protection for this release; see the Release Notes for details.</p> <p>** Allocate storage on a shared SR or a host-specific SR — not both. IF configuring storage subsystem data protection on a shared SR as an alternative to an everRun mirrored disk, do not allocate storage for the second host. IF configuring host-specific everRun mirrored disks, allocate storage on two hosts and configure to meet performance needs.</p> <p>*** Allocate storage on a shared SR or a host-specific SR — not both. To configure everRun disk mirroring, allocate storage on two separate hosts.</p>				

System Connectivity

Figure 8 on page 44 shows one way to comply with best practices and set up a local two-host pool.

Brief Glossary



*Here are some common terms you'll encounter in this guide. everRun terms and XenServer/Industry terms are listed separately. You'll find a more extensive glossary in the Online Help system, available from the **everRun Availability Center (eAC)**.*

everRun® Terminology

availability links (A-links)

The connectivity path(s) between two availability managers (normally located on separate XenServer hosts) in an everRun-protected virtual machine (PVM).

Availability Manager (AM)

A purpose-built virtual appliance containing the everRun software that applies protection to a target virtual machine (VM). The AM, which exists on each everRun-enabled XenServer host, implements the logic and context required to create the protected VM.

bonded network interface card (NIC)

An everRun “bonded” NIC aggregates two separate network interface cards into a single logical interface. everRun uses a bonded NIC on the XenServer management network to provide a hot standby in case one NIC should fail.

Compute Instance (CI)

A component of the protected virtual machine (PVM) that represents a single instance of the everRun PVM on a XenServer host. A duplicate instance is cloned on another XenServer host.

everRun Availability Center (eAC)

The browser-based everRun graphical user interface (GUI) used to manage and monitor everRun-protected virtual machines.

everRun mirroring

A process for creating and maintaining a set of identical disk images on separate physical disks. Each everRun protected VM relies on a pair of storage repositories (SRs), known as a **mirror set**, which maintain identical information.

everRun SplitSite®

An everRun deployment located in two separate physical locations, in order to provide protection against outage at one location. Used for disaster-tolerance or data recovery.

host-specific storage

Physical storage directly attached to a host, or networked storage allocated to a single host. In both cases, the storage is not available to other hosts in the pool.

Compare to **shared storage** under XenServer and Industry Terminology.

isolation IP address

The IP address of an independent network device on the XenServer management network. everRun software pings this device to detect the health of the host's management network connection. The results of the ping help isolate master/member failures and determine corrective action. You must configure this IP address; suggested values are that of the network's gateway device or the preferred quorum computer configured for the pool.

Level 1 protected VM

A VM protected using XenServer HA.

Level 2 protected VM

A VM protected using everRun Level, component-level protection, and configured using one to four CPUs. Continue to run through a failure of storage or network components, or a host failure that requires migration to the alternate host.

Level 3 protected VM

A VM protected using everRun Level, system-level protection, traditionally configured with a single CPU (uniprocessor). Continue to run through a complete failure of a physical host.

Level 3 SMP (or FT-SMP) PVM

A VM protected using everRun Level, system-level protection, and configured with two to four CPUs (multiprocessor). Offers scalability advantages over uniprocessor-based Level 3 PVMs.

protected VM (PVM)

A virtual machines (VM) that has Marathon availability software applied. For optimum availability, a PVM should

run on two physical XenServer hosts. See also Level 2 or Level 3 protected VM.

quorum service

An everRun communication service, installed on one or more Windows clients on the LANs, that prevents host servers from operating independently in a non-communicating, partitioned mode of operation (commonly known as *split brain*). Also enables unattended restarts in failure modes that impact all networks (such as loss of connectivity between the two AMs in a PVM).

quorum service computer

One of two network-accessible PCs that run the everRun quorum service when quorum service is included in the system deployment. The eAC provides a dialog box where you can specify an IP address for the *preferred* quorum computer, as well as an *alternate* quorum computer. All PVMs in a resource pool use the same quorum service computers.

split brain

An undesirable condition in which the two availability managers (AMs) servicing the two halves of an everRun protected VM are partitioned and operating independently.

storage subsystem data protection

An everRun data protection option, which is available to users whose configuration includes intelligent SAN systems that provide redundant paths to the data within the storage fabric. Both everRun hosts access data on a single VDI. If one pathway to the storage fails, the everRun software activates the alternate pathway on the other server hosting the protected VM data.

XenServer™ and Industry Terminology

domain

In XenServer terminology, refers to the execution context that contains a running virtual machine (VM). The relationship between VMs and domains in XenServer is similar to that between programs and processes in an operating system: like a program, a VM is a persistent entity that resides on disk. When it is loaded for execution, a VM runs in a domain with a distinct domain ID. See **Domain0**, **DomainU**.

Domain0 (Dom0)

A privileged virtual machine (VM) that provides the platform management and drives the physical I/O devices and platform resources. Dom0 (rhymes with Mom) is the first domain started by the XenServer hypervisor at boot, running a Linux OS. Dom0 is sometimes referred to as the “control domain.”

DomainU (DomU)

An unprivileged domain on a XenServer, running a guest operating system that has been ported to XenServer; for example, each Windows guest VM runs in a DomU, or user domain.

FC-SAN

In storage virtualization, fibre-channel (FC) storage area networks (SAN) allow you to aggregate disks and logical units (LUNs) from disparate systems into shared storage pools.

iSCSI

Internet SCSI, an IP-based standard for linking data storage devices over a network and transferring data by carrying SCSI commands over IP networks. For an overview, see wikipedia.org.

ISO image

A disk image (archive file) of an optical disc using a conventional ISO format supported by many software vendors. The name "ISO" is taken from the ISO 9660 file system used with CD-ROM media.

logical unit number (LUN)

An address for an individual disk drive. The term is used in the SCSI protocol as a way to differentiate individual disk drives within a common SCSI target. In everRun VM, the LUN provides address information for individual storage volumes within a mirror set.

logical volume

In computer storage, logical volume management (LVM) is a method of allocating space on mass storage devices that is more flexible than conventional partitioning schemes. Volume management is one of many forms of storage virtualization. A volume manager can concatenate, stripe together, or otherwise combine partitions (volumes) into larger virtual ones that can be resized or moved.

management computer

In a XenServer configuration, the general-purpose Windows computer that is used to manage the XenServer software as well as the everRun software. Also called **client computer**.

master host

The physical node in a XenServer resource pool that contains an administration interface and forwards commands to individual member nodes.

member host

A physical node in a XenServer resource pool that has no administration interface and receives commands from the pool's master node. Also called **pool member**.

NFS

Network File System, an industry-standard client/server application that allows network users to access shared files stored on computers of different types. Users can manipulate shared files as if they were stored locally on the user's own hard disk.

PATA

Parallel Advanced Technology Attachment. One type of standard interface for connecting storage devices in personal computers. For an overview, see AT Attachment at wikipedia.org.

resource pool

A connected group of homogeneous XenServer hosts that, combined with shared storage, provide a platform on which virtual machines (VMs) run. Servers in a resource pool monitor the state and availability of their peers. Also called **pool**.

SATA

Serial Advanced Technology Attachment. One type of standard interface for connecting storage devices in personal computers. For an overview, see AT Attachment at wikipedia.org.

SCSI

A type of hard disk drive used in high-performance workstations and servers. For an overview see SCSI at wikipedia.org.

shared storage

Networked physical storage with access to all hosts in the pool. This is a XenServer configuration property, described in the Storage chapter of the XenServer Administrator's Guide. It includes a variety of storage options, including iSCSI and FC-SAN. *Compare to **host-specific storage** under everRunTerminology.*

storage repository

In XenServer, any storage target that contains physical disks.

virtual disk image (VDI)

A disk abstraction that contains the contents of a virtual disk.

virtual machines (VMs)

File-based abstractions of dedicated machines, which provide the environment in which a hosted operating system (OS) runs. A VM is also known as a guest domain or "DomU" in XenServer terminology.

Xen bridge (xenbr)

In XenServer network configuration, the bridge parameter defines a method of connecting objects. For example, if a physical interface (PIF) object connects a network object (n) to a host object (h), the network corresponding to n is bridged onto a physical interface specified by the fields of the PIF object. Thus, if a NIC is named **eth0**, the associated bridge is **xenbr0**.

XenCenter

A Windows client (GUI) application that enables you to manage XenServer hosts, resource pools, and shared storage. It also lets you deploy, manage, and monitor VMs.

XenServer API (xapi)

The Citrix XenServer management API, the applications programming interface that supports both XenCenter GUI and the `xe` command line interface (CLI).

XenServer host

As defined by Citrix, a host is a homogeneous system — aggregated into a single management entity — that can contain a number of virtual machines. See the XenServer documentation for more information.

