

# Protecting Virtual Machines: The “Best of VMworld” Approach

Simon Crosby, CTO, Citrix Virtualization and Management Division  
Jerry Melnick, CTO, Marathon Technologies

## VIRTUALIZATION AND AVAILABILITY – THE NEED FOR A NEW APPROACH

*"Although consolidation and legacy migration have been the primary motives for server virtualization, going forward we see high availability and disaster recovery for business resiliency as key reasons why organizations will deploy virtualization software."*

[John Humphreys, IDC]

IT managers and executives have enthusiastically embraced server virtualization because it greatly increases server utilization, reduces capital costs and cuts operating expenses. And a growing number of IT executives are looking beyond server consolidation to use virtualization for simpler, more efficient disaster recovery and application availability.

By enabling IT managers to move virtual machines (and the applications they host) between physical servers, server virtualization gives them a powerful method for reducing planned downtime and speeding recovery. But protecting virtual machines from unplanned downtime and unforeseen outages is a challenge. Traditional approaches to ensuring application availability are complex and costly, and these shortcomings are amplified when applied to virtual environments. Adding cost and complexity can erode the cost savings and efficiency benefits of server virtualization.

IT managers need availability solutions that are simple, non-disruptive and cost-effective so they can extend availability to any virtual machine to ensure business resiliency. This white paper will explore a completely new approach to availability for virtual machines—an approach recognized with the *Best of VMworld 2007 - New Technology* award.



# Protecting Virtual Machines: The Best of VMworld Approach

## VIRTUALIZATION AND UNPLANNED DOWNTIME

Application availability is a major concern for production applications in virtual environments because the server becomes the single point of failure for multiple applications. The consequences of downtime are much greater when all of your eggs are in one basket. Protecting virtual machines in this environment is essential.

*“Best-in-Class” companies used virtualization technologies to reduce unplanned downtime by an average of 53%. [Aberdeen Group: Virtualization and Business Continuity: Are You Protected?]*

Today’s virtualization technologies are particularly useful for protecting applications from planned downtime—outages necessary for administrative purposes. Using *live migration* technologies, companies can move virtual machines and their running applications between physical machines without disruption. Examples of planned migration technologies include Citrix® XenMotion™ and VMware® VMotion™.

But protecting virtual environments from unplanned downtime is a different matter. In many cases, virtual environments employ traditional clustering and failover techniques that were never designed for virtualized environments. These traditional approaches use rudimentary heartbeat pings to check the status of a virtual machine. They suffer from several drawbacks:

- Clustering and failover add cost and complexity to the environment, requiring manual configuration, setup, scripting and testing to define the appropriate actions to take in case of failures. This additional administrative complexity can introduce errors, contributing to availability issues.
- Heartbeat pings are unable to reliably determine the health of a virtual machine and may not distinguish between I/O path failures, server failures, or lack of system resources. In some cases, these limitations may result in unnecessary or false failovers. In other cases, discrete storage or network device outages are not identified as failures and the system does not fail over.
- The failover process is far from certain; it assumes that the administrator has configured the standby system appropriately for the application and has maintained that configuration. If the target system is not configured appropriately, then when a failover does occur, the application or virtual machine is inoperable on the standby system, causing a “failed failover”. Given the sense of uncertainty, some refer to this approach as “ping and pray.”

Marathon Technologies and XenSource (now Citrix Virtualization and Management Division) have worked closely together to offer a new, innovative approach to availability in virtual environments. Marathon **everRun VM** goes beyond clustering

# Protecting Virtual Machines: The Best of VMworld Approach

and failover to create a fault tolerant-class of availability for virtual machines—eliminating downtime from device and even complete server failures. This integrated solution uses Citrix XenServer™, a powerful virtualization environment for Windows clients, to provision multiple virtual machines on a single physical server while everRun™ provides the fault tolerant-class protection.

Unlike existing availability solutions that introduce cost and complexity into the virtual environment, **everRun VM** is both simple and reliable, while offering unprecedented levels of availability for business-critical applications in virtual environments.

## XenServer: THE OPEN, SCALABLE, HIGH-PERFORMANCE VIRTUALIZATION ENVIRONMENT

Citrix XenServer Enterprise Edition 4 (referred to in this paper as XenServer v4) is an enterprise-class virtualization solution based on the open source Xen hypervisor, offering x86-based server virtualization for Windows and Linux clients. As a virtualization platform, XenServer v4 is characterized by performance, openness and simplicity:

- **Performance:** XenServer v4 uses hardware-assisted virtualization on Intel® VT and AMD® Virtualization (AMD-V™) platforms, delivering performance approaching bare-metal levels and supporting both 32-bit and 64-bit environments.
- **Openness:** The open source Xen hypervisor has broad industry support and participation.
- **Simplicity:** XenServer v4 is a lightweight solution that is simple to install and use, and that includes necessary virtualization management utilities.
- **Resiliency:** The XenMotion feature enables live migration of running virtual machines between hosts to achieve resilient, optimal resource utilization for virtual infrastructures.

## XenServer 4: ELIMINATING THE SINGLE POINT OF FAILURE MANAGEMENT TOOL

VMware Virtual Center is a traditional, heavyweight management application with a database that owns authoritative state for the managed hosts. As such, it is a single point of failure that is vulnerable to catastrophic failure. However, in a XenServer v4

*By 2009, there will be three competitive hypervisor architectures: VMware ESX Server, Xen and Microsoft's hypervisor (0.8 probability). [Gartner]*

# Protecting Virtual Machines: The Best of VMworld Approach

resource pool, management state for all hosts and virtual machines is stored redundantly in every host in the pool. At any time, one host is the pool leader, offering a powerful XenAPI interface to the XenCenter management console and third party applications. The XenCenter console delivers event notifications and performance data, and takes instructions for the pool to execute (such as managing VM lifecycle). If the pool leader fails, any other server in the resource pool can simply take over the job with no loss of state information.

## ADDING FAULT TOLERANT-CLASS AVAILABILITY TO THE VIRTUAL ENVIRONMENT

Fault tolerance, or the ability to remain operational in the event of a fault, was once the domain of high-end, mainframe-class servers running specialized operating systems. Marathon makes fault tolerant-class computing available to everyone with its software-based solution providing the highest levels of availability for virtual machine environments.

*everRun is installed in over 1,500 companies and institutions including stock exchanges, television networks, casinos, pharmaceutical companies, and manufacturing plants throughout the world.*

everRun transparently combines and manages the resources of two virtual machines running on different servers in a XenServer v4 resource pool to create a single protected virtual machine environment. The protected virtual machine appears and is managed just like a standard Windows server. Disk data is mirrored synchronously to redundant storage and network and server operations are protected from failure. The administrator loads and configures the application in the protected virtual machine as though it was being loaded onto a physical server and then simply walks away. There is nothing else to do!

If a fault or failure occurs in a disk, network device or even a XenServer host, everRun's ComputeThru™ capability automatically reconfigures resources to permit the application to continue operating without interruption or loss of client connectivity.

Repair is transparent and automatic as well. When a failed component or XenServer host is repaired, the administrator boots the repaired server into the environment and everRun automatically incorporates the virtual machine and components into the protected environment, restoring redundancy without interrupting application operations.

The two XenServer hosts can be located in separate computer rooms, campus buildings or even across a WAN to protect applications from site-wide disasters.

The everRun solution eliminates the complexity and cost of traditional clustering

# Protecting Virtual Machines: The Best of VMworld Approach

solutions. Failure detection is automatic and recovery policy is embedded. There is no need to create and test complex failover scripts. everRun works equally well with all classes of storage devices, unlike clusters that require expensive shared storage. Configuration is simple; everRun presents a single protected virtual machine to the administrator and transparently manages redundant resources for availability, removing opportunities for operator-induced errors.

Failure recovery is always reliable —with active redundancy the system is constantly validating itself because the two virtual machines comprising the everRun protected environment continuously execute the same application operations. This unique architecture assures the availability and correct operation of system resources when they are needed in the event of a failure.

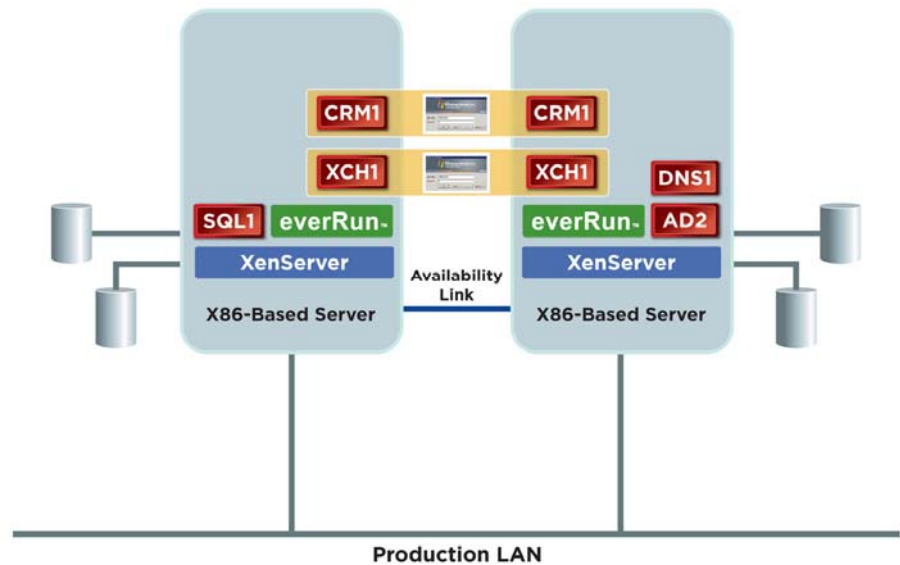
## HOW IT WORKS

We'll examine how the solution works in the simplest case, with two identical Intel or AMD-based physical servers. These servers are configured with Ethernet adapters for application use and either local (direct-attached) or shared storage. Redundant Ethernet paths, known as Availability Links, provide a private communications path used by everRun to maintain synchronization of storage and application operations for the virtual machines that everRun protects.

Each server runs the XenServer Enterprise Edition virtualization software as the host environment. The everRun technology is loaded as a virtual appliance (the everRun Availability Manager). The Availability Manager runs within the XenServer environment as a purpose-built appliance that provides availability for other virtual machines on the host. The Availability Manager establishes a tightly-coupled relationship to the virtual machines it protects. It resides in the data path between the virtual machine and the control domain, which handles I/O for the virtual machines.

*everRun VM showcase demo won a "Best of VMworld – New Technology" award at VMworld 2007*

# Protecting Virtual Machines: The Best of VMworld Approach



## everRun Availability Manager Protects Virtual Machines

This advanced availability architecture provides a number of benefits over traditional clustering and failover technologies.

- It offers superior insight into the health and activity of the protected virtual machine. With access to the data path, everRun can rapidly identify errors in specific devices and recover operations in a discrete and responsive manner, rather than incurring a complete failover process for every class of error.
- everRun controls virtual machine operations transparently, redirecting I/O away from failed devices and correctly operating devices without losing I/O transactions or interrupting the application.
- The Availability Manager replicates disk data and maintains application synchronization between the two host servers to provide the resource redundancy required to support the highest levels of virtual machine availability.

The everRun Availability Center management console provides a consolidated view of the entire environment with at-a-glance insight into events and status. You are able to view and manage the entire everRun availability environment from a browser-based interface.

# Protecting Virtual Machines: The Best of VMworld Approach

To protect a virtual machine, the administrator uses the everRun Availability Center to indicate which virtual machine to protect with everRun and then selects the desired protection level and which host to use for protection. That's it. There is no manual configuration, no script writing or failover testing required.

everRun creates a paired virtual machine environment that appears as a single virtual machine environment for both administration and application operations. The environment presents a single IP address and hostname to the application network so users never have to make client-side changes when failures occur.

With everRun ComputeThru technology, protection can be configured so that component and even server failures go unnoticed. You simply keep computing through the failure, rather than failing over between servers.

Of course, the two-server scenario is simplistic. You can have many servers in the resource pool, and choose different configurations of servers for the protection of virtual machines.

## THE FUTURE OF RESILIENT VIRTUALIZATION

Citrix and Marathon Technologies have worked closely to develop and test this integrated solution for fault tolerant-class availability in a virtual environment. Yet different applications have distinct availability requirements. For some, data replication alone may be sufficient, while others need remote recoverability for disaster recovery purposes and others fault tolerance. Moving forward, Marathon and Citrix will expand the integrated solution to support different levels of availability, giving administrators the ability to set the right level of availability for application requirements.

By combining high-performance, easy-to-use virtualization with the fault tolerant-class availability of everRun, Citrix and Marathon Technologies enable new levels of availability and resilience for virtual environments. By increasing both the availability and simplicity of virtual environments, Citrix and Marathon Technologies are driving broader adoption of virtualization with new, innovative uses for the technology across a broad spectrum of applications.

Notice: The information in this publication is subject to change without notice. THIS PUBLICATION IS PROVIDED "AS IS" WITHOUT WARRANTIES OF ANY KIND, EXPRESSED OR IMPLIED, INCLUDING ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT. CITRIX SYSTEMS, INC. ("CITRIX") AND MARATHON TECHNOLOGIES CORPORATION ("MARATHON"), SHALL NOT BE LIABLE FOR TECHNICAL OR EDITORIAL ERRORS OR OMISSIONS CONTAINED HEREIN, NOR FOR DIRECT, INCIDENTAL, CONSEQUENTIAL OR ANY OTHER DAMAGES RESULTING FROM THE FURNISHING, PERFORMANCE, OR USE.