Introduction to VMware vSphere® Replication

Simple and Effective Virtual Machine Protection

TECHNICAL WHITE PAPER

vmware®

Table of Contents

Introduction

A fundamental part of protecting IT is ensuring that the services provided by virtual machines are resilient, and robust at all levels of the compute stack, from hardware through to the application.

vSphere Replication is a feature introduced with VMware vSphere® 5.1 ("vSphere 5.1"). It is designed to augment the recovery capabilities of the VMware vSphere® ("vSphere") platform by providing a built-in capability to continually replicate a running virtual machine to another location. Replication creates a copy of a virtual machine that can be stored locally within a cluster or at another site, providing a data source to rapidly restore a virtual machine within minutes.

vSphere Replication augments offerings in the vSphere availability protection matrix. It provides a solution that enables recovery time better than that of restoring from backup, without introducing the complexity of a complete storage array–based replication configuration. vSphere Replication also enables configuring replication on a per–virtual machine basis and significantly rounds out the capabilities of protection offered by vSphere.

This paper will help you understand what vSphere Replication is and some of the benefits of its features. It will also discuss how it works to protect your virtual machines against failure.

What Is vSphere Replication?

vSphere Replication is a feature of the vSphere platform. It copies a virtual machine to another location, within or between clusters, and makes that copy available for restoration through the VMware® vCenter Server™ Web-based user interface.

vSphere Replication continues to protect the virtual machine on an ongoing basis. It replicates to the copy the changes that are made to the virtual machine. This ensures that the virtual machine remains protected and is available for recovery without requiring restore from backup.

vSphere Replication is provided as a no-charge component of all eligible vSphere licenses, ranging from the VMware vSphere® Essentials Plus Kit through the VMware vSphere® Enterprise Plus Edition™. As with backups through VMware data protection, protecting a virtual machine is a critical function of a hypervisor platform for the datacenter.

Unified management of vSphere Replication is provided via the next-generation VMware vSphere® Web Client. This provides a common and unified screen for all aspects of virtual datacenter management, including many aspects of protection, such as replication, backup and restore.

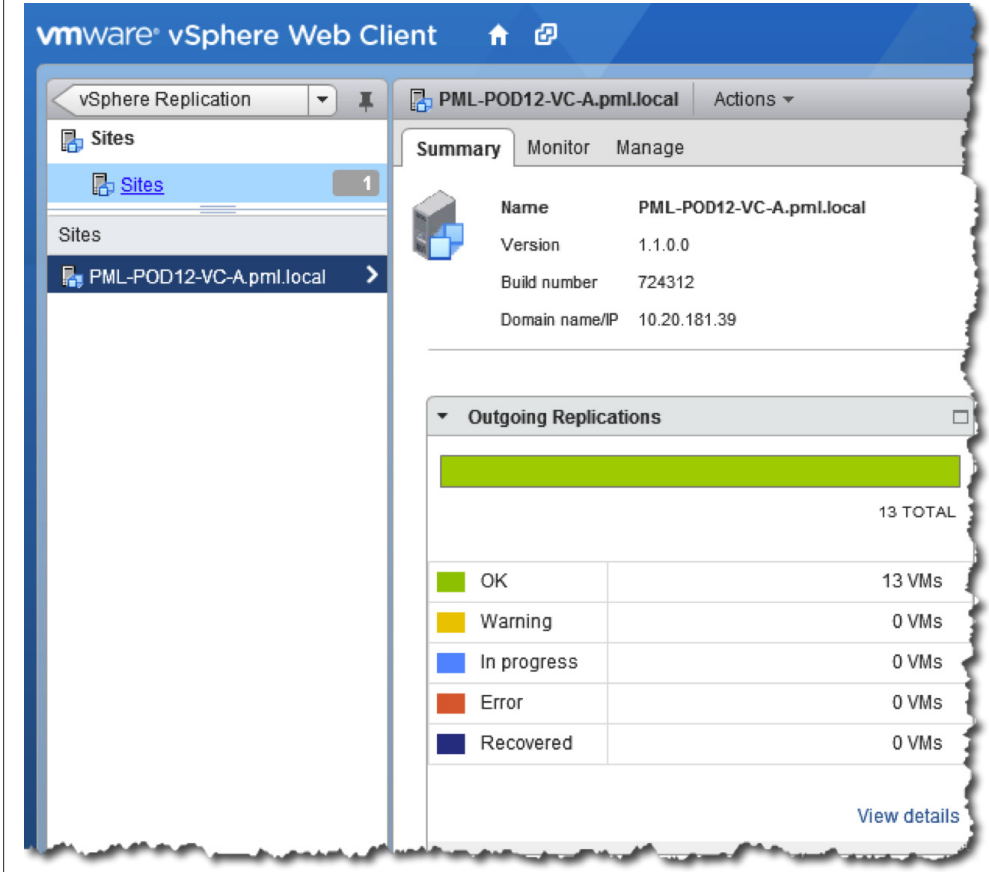

Figure 1. Integrated with vSphere Web Client

Some replication technologies provide mere copies of a virtual machine at a remote site without any consideration for the consistency of the application data within the virtual machine. vSphere Replication can be configured to ensure consistent application data, along with virtual machine data, with one simple click when configuring a virtual machine for replication.

Automatic integration with Microsoft's Volume Shadow Copy Service (VSS) ensures that applications such as Microsoft Exchange or Microsoft SQL Server databases are quiescent and consistent when replica data is being generated. A very quick call to the virtual machine's VSS layer flushes the database writers for an instant to ensure that the data replicated is static and fully recoverable.

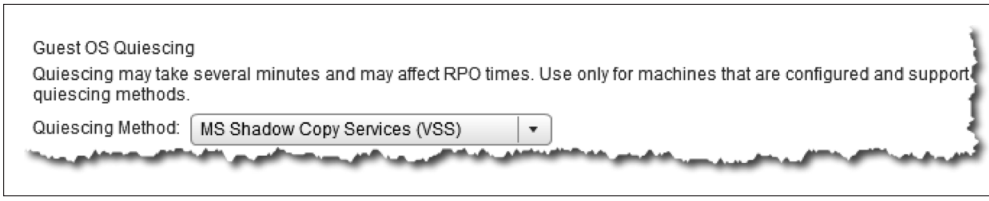

Figure 2. Application Consistency via VSS

There are no application agents or administration required for this process. Application consistency is inherent to the copies made by vSphere Replication.

How Does It Work?

vSphere Replication is a deeply integrated component of the vSphere platform. It is the only truly "hypervisor-level" replication engine available today. Changed blocks in the virtual machine disk(s) for a running virtual machine at a primary site are sent to a secondary site, where they are applied to the virtual machine disks for the offline (protection) copy of the virtual machine.

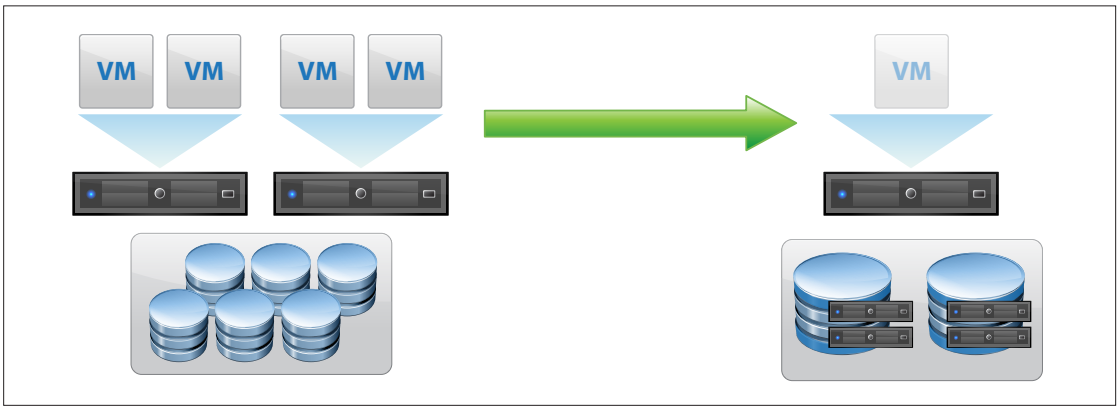

Figure 3. Replicate Virtual Machines' Changed Blocks

vSphere Replication consists of an agent that is part of the core vSphere 5 installation package on each host, and a set of virtual appliances that are deployed from the management interface. From a conceptual perspective, the agent is responsible for sending changed data from a running virtual machine, and the appliance receives the replication at a remote site and applies it to the offline disk files for the virtual machine. The vSphere Replication appliance is also responsible for managing replication, which gives the administrator visibility into the status of protection of the virtual machines, as well as the ability to recover virtual machines with a few simple clicks.

Configuring replication for up to 500 virtual machines through the same management interface that is used for all vCenter operations is a simple process of right-clicking a virtual machine and selecting the destination for its replica.

Part of this process is to select a recovery point objective (RPO), which provides information to vSphere Replication regarding an age limit for the copy of the virtual machine. It then attempts to replicate data to meet the RPO at all times, ensuring that the virtual machine data is never older than the defined policy for each virtual machine configured for replication.

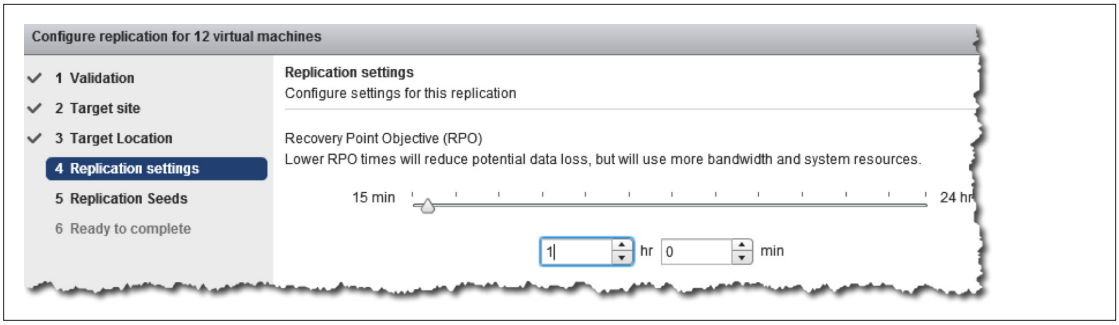

Figure 4. Selecting an RPO for Multiple Virtual Machines

vSphere Replication does an initial full synchronization of the source virtual machine and its replica copy. If it is wanted, a seed copy of data can be placed at the destination, to reduce the amount of time and bandwidth required for the initial replication. A seed copy of a virtual machine consists of a virtual machine disk file that can be placed at the target location through almost any mechanism. Seeding is not a necessary process, and vSphere Replication creates an initial copy at the target location, whether or not a seed is present. If a user has given a seed for replication, the data therein will be used to reduce the amount of replication needed to initially synchronize the primary disk and its replica. A seed can be created manually or copied into place by any mechanism the administrator chooses, such as offline copying, FTP, "sneakernet," or even using an ISO or a clone of a virtual machine.

After baseline synchronization is complete, vSphere Replication switches to transferring only the blocks of data that have changed. The vSphere kernel itself tracks unique writes to the disk files of protected virtual machines and identifies and replicates only those blocks that have experienced unique writes during the configured RPO. This ensures that a minimized amount of data is sent over the network to the target and enables aggressive RPOs. After unique data has been sent, it doesn't need to be sent again. Only changes will be replicated, and the blocks will be sent to the target location's vSphere Replication appliance.

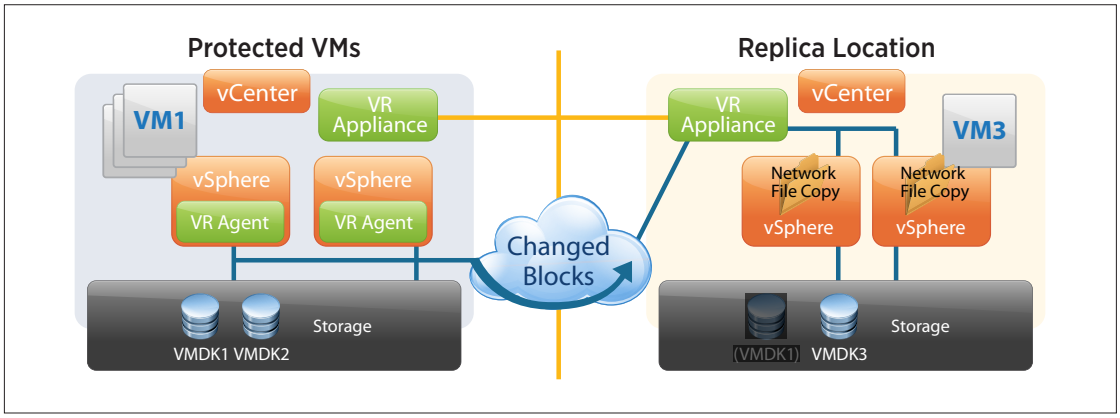

Figure 5. Data Flow of Replication from vSphere Agent to vSphere Replication Appliance

At the target location, the data is received and checked within the vSphere Replication appliance. Only fully consistent data is then written (via network file copy) to the target cluster's vSphere hosts and thereby to disk. This manner of waiting for a completely consistent block group ensures recoverability of the replica virtual machine at all times, even if data is lost during transit or a crash occurs at any point during the transfer.

From the perspective of the protected virtual machine, this entire process is completely transparent and requires no changes to configuration or ongoing management. The replication is nonintrusive, irrespective of the operating system within the virtual machine.

Simplicity Without Sacrificing Functionality

vSphere Replication is inherently a lightweight replication protocol. By replicating only changed blocks on an ongoing basis, network bandwidth can be saved and commit times for data are minimized.

The vSphere Replication framework consists of a vSphere Replication appliance that is paired with a single vCenter server. This appliance provides the management framework for vSphere Replication and also acts as the target for replicated blocks. Increased scale or remote sites managed by a single vCenter server might lead to the choice of deploying up to 10 vSphere Replication Server (VR Server) virtual appliances. The VR server appliances function strictly as a destination target for replication and enable simple scaling and distribution of the replication framework. As many as 500 virtual machines can be protected with an individual vSphere Replication instance.

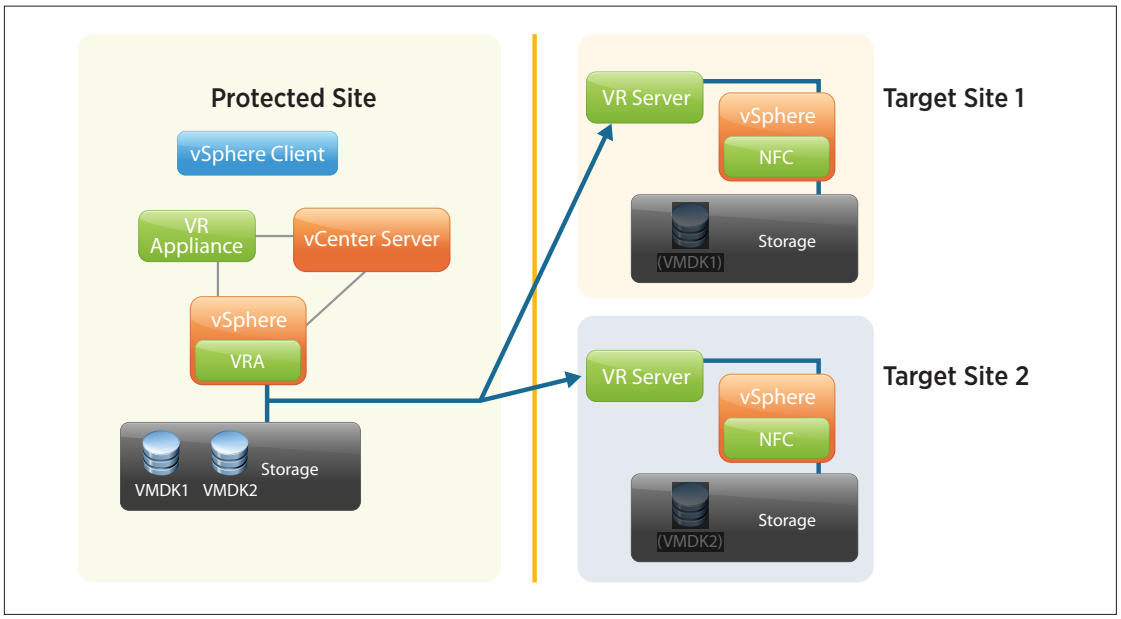

Figure 6. Framework for Replication with Multiple Sites and VR Servers

Within this framework, RPOs for each virtual machine can be defined, ranging from 15 minutes to 24 hours. This RPO can be changed on demand without interruption, giving administrators the ability to fine-tune replication based on dynamic factors such as change rates, bandwidth availability, and so forth.

Because the RPO is unique to each virtual machine, more critical virtual machines can be given a more aggressive replication target than others. Groups of virtual machines, however, can be selected en masse to enable large changes. Each disk within a virtual machine can be independently configured to be replicated or not replicated, offering further opportunity to save on bandwidth and time for recovery. For instance, a database server might consist of multiple disks, one of which is dedicated to a temporary scratch location. This disk is not required for replication. Similarly, sometimes a swap file or paging file is redirected to a dedicated disk that can also be excluded from replication to save bandwidth for transient data. The replicated disks are independent of format, layout and snapshots of the primary copy. A virtual machine disk file that is thick provisioned on a Fibre Channel SAN at the primary location can be thin provisioned to local disk for recovery if wanted. vSphere Replication eliminates the requirement that replication have identical storage hardware. Similarly, because the virtual machine replica will be “cold booted” when recovery takes place, server hardware need not be identical if replicating between clusters.

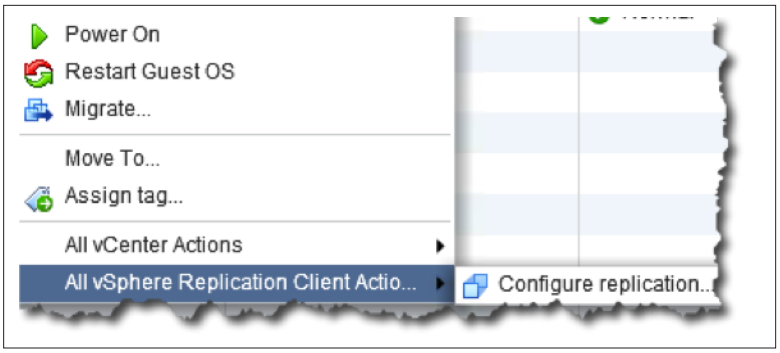

Figure 7. Replication Is a Policy of the Virtual Machine

Essentially, vSphere Replication enables policy-driven protection of a virtual machine—because the configuration of a virtual machine's replication is attached as a property of the virtual machine itself—while also enabling virtual machines to continue operating without any change, overhead or interruption. The tracking mechanism that identifies changed blocks of a virtual machine resides above the storage layer, enabling replication that is fully independent of disk location, disk format, thin or thick provisioned disks and whether or not snapshots are in use by the virtual machine.

Recovering a Virtual Machine

Recovery of a virtual machine might be necessary for a number of reasons, ranging from testing, impending outages or even recovery from a disaster. vSphere Replication is designed to give administrators the ability to manually recover an individual machine with a small number of clicks in the vSphere Web Client.

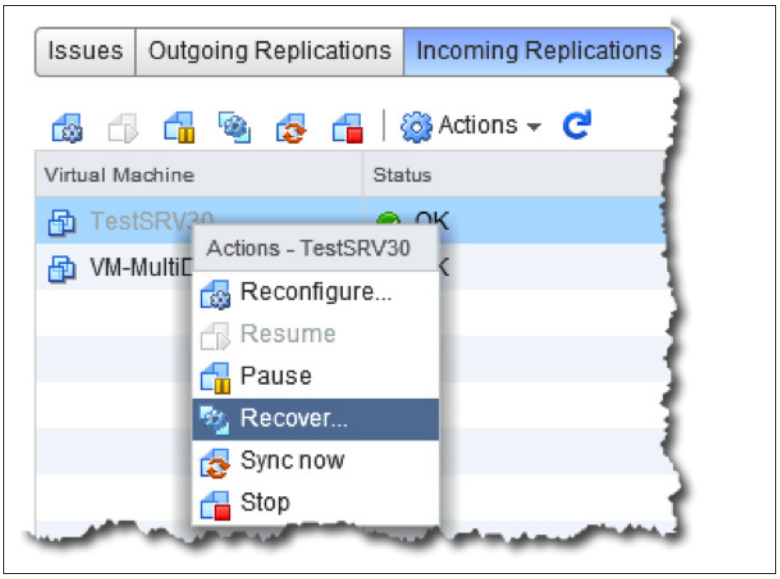

Figure 8. Recover a Virtual Machine by Right-Clicking Its Replication Status

NOTE: It is not the goal of vSphere Replication to provide complex orchestration of disaster recovery in the manner of VMware® vCenter™ Site Recovery Manager™. Instead, vSphere Replication offers replication on a per–virtual machine basis to enable fine-grain recovery configurations for environments that need it.

With vSphere Replication, recovery of a virtual machine can be started easily by one of three methods. After navigating to **Replication Status** in the **vSphere Replication** area of the vSphere Web Client, the administrator first selects the virtual machine to be recovered. Selecting **Recover...** as an action can be done by right-clicking, then clicking the **Recovery** icon, then selecting the **Recover...** action from the **Actions** drop-down menu.

The replica cannot be powered on and recovered if the original virtual machine is still reachable and is itself still powered on. To continue, the primary copy of the virtual machine must be unreachable by vCenter Server or powered off.

The next steps are simply to choose a destination for the virtual machine, a folder to hold it, and a cluster or resource pool in which it will be powered on.

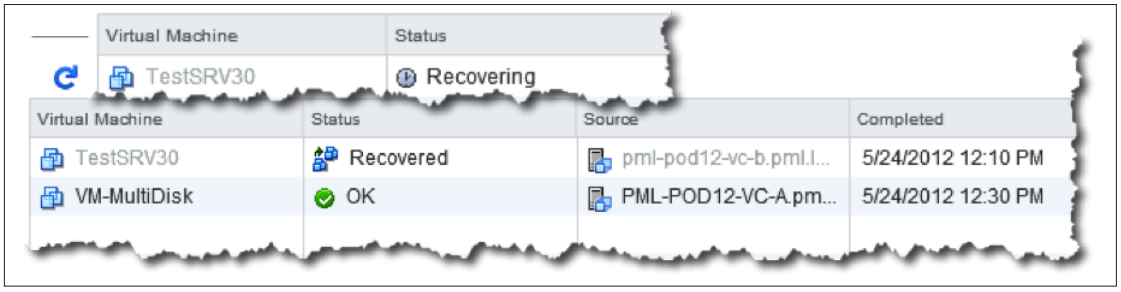

Figure 9. Recovery of a Virtual Machine Is Complete

After these items are selected, recovery will begin. This process creates a powered-on copy of the virtual machine connected to the replica disk but does not connect any of the virtual network cards to any port groups. This can help avoid situations where the same virtual machine might be active on the network in two locations simultaneously, thereby preventing broadcast collision and routing issues involving duplicates.

After the virtual machine is fully booted, the administrator reviews the recovery and status of the booted replica. To complete recovery, the administrator simply attaches the virtual machine to the appropriate networks.

Recovery of a virtual machine is an easy process that can be accomplished with a few clicks in the GUI in a safe and reliable fashion.

Next Steps

Protect Your Virtual Machines from Loss and Downtime with vSphere Replication

Make vSphere Replication a part of your vSphere 5.1 deployment and improve your virtual machine availability without introducing risk.

Register today for a trial of vSphere and enjoy the benefits of replication of your critical virtual machines as an integrated part of your IT platform.

Additional Documentation

For more information about vSphere, visit the product pages at http://www.vmware.com/products/vSphere.

Here are some links to online vSphere and vCenter Server documentation:

- Product documentation
- Support knowledge base

VMware Contact Information

For additional information or to purchase VMware vSphere 5.1, the VMware global network of solutions providers is ready to assist. If you want to contact VMware directly, you can reach a sales representative at 1-877-4VMWARE (650-475-5000 outside North America) or by email at sales@vmware.com. When emailing, include the state, country and company name from which you are inquiring.

Providing Feedback

VMware appreciates your feedback on the material included in this guide. In particular, we would be grateful for any guidance on the following topics:

How useful was the information in this guide?

What other specific topics would you like to see covered?

Please send your feedback to tmfeedback@vmware.com, with "Introduction to vSphere Replication" in the subject line. Thank you for your help in making this guide a valuable resource.

About the Author

Ken Werneburg is a senior technical marketing manager at VMware. He works on business continuity and disaster recovery solutions in the Cloud and Infrastructure Management group. Ken has been with VMware since 2006.

- Follow Ken's blogs at http://blogs.vmware.com/vsphere/uptime.
- Follow Ken on Twitter: @vmKen.

vmware®

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com