



Microsoft Exchange 2010 on VMware Availability and Recovery Options

© 2011 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. This product is covered by one or more patents listed at <http://www.vmware.com/download/patents.html>.

VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

VMware, Inc
3401 Hillview Ave
Palo Alto, CA 94304
www.vmware.com

Contents

1. Introduction	5
2. VMware vSphere Platform Advantages.....	5
3. Increase Availability across the Exchange Lifecycle.....	6
3.1 Simplify Upgrades and Reduce or Eliminate Downtime	6
3.2 Performance and Recovery Advantages	6
4. Local Site Availability Options	7
4.1 VMware vSphere HA, DRS, and vMotion	7
4.2 Application Aware vSphere HA.....	8
4.3 Microsoft Database Availability Groups	9
5. Remote Site Availability Options	11
5.1 VMware vCenter Site Recovery Manager	11
5.2 Exchange 2010 DAG with Delayed Log Replay	12
5.3 Third-party Software-based Replication	13
6. Backup and Restore Options	15
6.1 In-Guest Software Solutions	15
6.2 VMware Data Recovery	16
6.3 Array-Based Backup Solutions	17
7. Additional Information	18

1. Introduction

By leveraging the inherent benefits of a VMware-based platform, a Microsoft Exchange Server 2010 deployment on VMware vSphere® offers a choice of several availability and recovery options, each providing varying levels of protection and cost. This solution brief provides a description of the various options available. Topics include:

- VMware vSphere Platform Advantages.
- Increase Availability across the Exchange Lifecycle.
- Local Site Availability Options.
- Remote Site Availability Options.
- Backup and Restore Options.

2. VMware vSphere Platform Advantages

Although application-level clustering has been the prevalent solution for most Exchange implementations, features of the vSphere platform can enhance the overall availability of Exchange by providing options that help to limit both planned and unplanned downtime. In fact, for many organizations, the features provided by vSphere may satisfy the availability requirements of their business without needing to follow traditional clustering approaches. For other organizations that require a greater degree of availability, application-level clustering can be combined with the vSphere features to create an extremely flexible environment, with options for failover and recovery at both the hardware and application levels. Some of the advantages of the vSphere platform include:

- Virtual machines are portable – This means that your Exchange server is no longer bound to a particular piece of hardware, this can enhance availability in several ways:
 - Design decisions are no longer permanent – You can adjust your CPU and memory requirements with a simple reconfiguration and reboot.
 - Easily upgrade to newer hardware – As your Exchange environment grows or changes, simply move the Exchange virtual machine to newer hardware to accommodate increased workloads.
- Virtual machines are hardware independent – Hardware independence means increased flexibility when designing both production and disaster recovery components. Cluster nodes and recovery servers can be virtualized, eliminating the need for identical hardware.
- VMware High Availability (HA) protects your server from hardware failure – If your physical server or any critical component within the server fails for any reason, vSphere HA will automatically reboot the Exchange virtual machine on another physical server, acting as a first line of defense against service outage. By combining vSphere HA with traditional clustering approaches, you can mitigate both hardware and software failures for maximum availability
- VMware vSphere Distributed Resource Scheduler (DRS) can balance workloads and speed recovery. DRS is VMware vSphere® vMotion® with intelligence. As Exchange workloads increase, DRS can move a bottlenecked virtual machine to another host with more available resources, automatically and without downtime. DRS can also help to recover more quickly after server hardware failure. For example, if a physical server fails, HA will reboot the virtual machine on another physical server. When the failed server is replaced, DRS migrates the virtual machine back to its original location with no downtime and no interruption to the end-user.
- VMware offers consolidation opportunities – Underutilized Exchange components, such as Hub Transport Servers, Client Access Servers, and Domain Controllers—and even non-Exchange components such as underutilized file servers or Blackberry Enterprise Servers (link)—can be consolidated onto fewer physical servers for maximum hardware utilization and lower costs.

3. Increase Availability across the Exchange Lifecycle

3.1 Simplify Upgrades and Reduce or Eliminate Downtime

- Traditional, physical environment upgrades and scale-up activities require a great deal of resources, including:
 - Planning and implementation time from engineering resources, including application administration, server administration, and SAN administration.
 - Sizing and acquisition of new hardware.
 - Downtime required to perform an upgrade, which results in higher costs and risks.
- In comparison, scaling your environment using vSphere means simply adding more Exchange virtual machines as the client count increases.
- Physical Exchange server environments are tightly bound to a storage technology and are extremely difficult to scale. Adding more storage capacity to Exchange virtual machines is less complex because vSphere emulates storage to a simple SCSI device. The end result is that the Exchange environment can be upgraded regardless of the underlying storage technologies (iSCSI or Fibre Channel).
- With the VMware Virtual Machine File System (VMFS), the storage capacity serving Exchange environments can be reduced or increased on the fly with the hot add/remove storage functionality in vSphere.
- Virtualized Exchange environments can be serviced (for example, adding more physical memory or CPU) without interruptions due to the shared storage functionality from VMware VMFS and VMware vMotion technology.

3.2 Performance and Recovery Advantages

- Virtualized Exchange environments can recover from:
 - Planned and/or unplanned hardware outages using vSphere HA.
 - Hardware degradation by using DRS capability to automatically balance workloads.
 - Application failure using Microsoft Cluster Service (MSCS) or Failover Clusters within a virtual machine on shared storage.
- With the built-in multipathing capability and advanced queuing techniques available in virtual machine architectures, virtualized Exchange environments can leverage advanced configuration options to:
 - Increase the IOPS/transactions to service more clients.
 - Balance the workloads of multiple Exchange servers sharing the same physical server to utilize multiple SAN paths and storage processor ports.

High I/O, memory- and CPU-intensive applications such as Exchange can better recover from SAN errors because the applications reside on VMFS, which hides SAN errors from guest operating systems.

4. Local Site Availability Options

4.1 VMware vSphere HA, DRS, and vMotion


When deploying Exchange 2010, the high availability options available in physical environment continue to be available in a virtual environment. vSphere features such as VMware vSphere High Availability (HA), VMware vSphere Distributed Resource Scheduler, and VMware vSphere vMotion are available to allow the highest level of performance and recovery in the case of a host failure.

- VMware vSphere HA provides easy-to-use, cost-effective, high availability for applications running in virtual machines. In the event of physical server failure, affected virtual machines are automatically restarted on other production servers with spare capacity. Additionally, if there is an OS-related failure within a virtual machine, the failure is detected by vSphere HA and the affected virtual machine is restarted on the same physical server.
- VMware vSphere Distributed Resource Scheduler (DRS) collects resource usage information for all hosts and virtual machines and generates recommendations for virtual machine placement. These recommendations can be applied manually or automatically. DRS can dynamically load balance all virtual machines in the environment by shifting workloads across the entire pool of VMware ESX®/ESXi™ hosts. This ensures that critical Exchange virtual machines in the environment will always have the CPU and RAM resources needed to maintain optimal performance.
- VMware vSphere vMotion leverages the complete virtualization of servers, storage and networking to move a running virtual machine from one physical server to another. This migration is done with no impact to running workloads or connected users. During a vMotion migration, the active memory and execution state of the virtual machine is rapidly transmitted over the network to the new physical server, all while maintaining its network identity and connections.

4.1.1 Example: Standalone Exchange Mailbox Server Virtual Machine with vSphere HA, DRS, and vMotion

Out of the box, vSphere features can help to protect your standalone Exchange virtual machine from server host failure. vSphere HA automatically reboots your Exchange virtual machine on another server if the current one fails, so your virtual machine can be restored to normal operation in the time that it takes to reboot the operating system and start the Exchange services. After the original server hardware is fixed or replaced, DRS and VMware vMotion can be used to quickly move the virtual machine back to its original ESX host, with no additional downtime.

Figure 1. Mailbox Server Virtual Machine Protected with HA



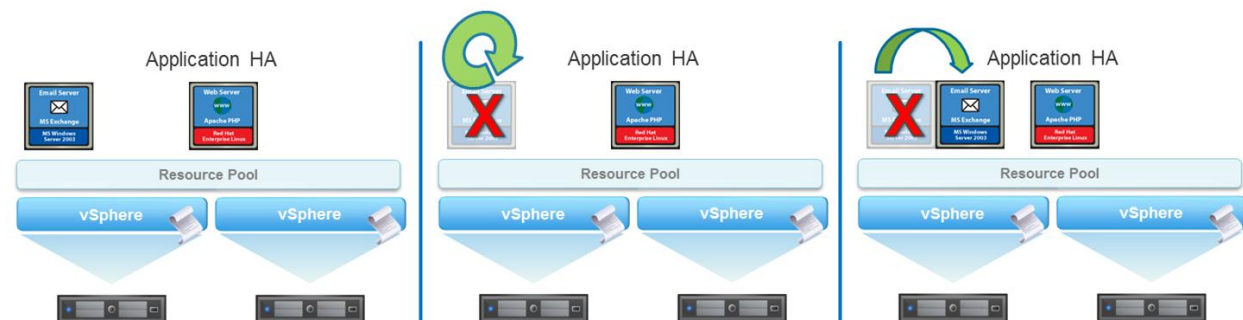
4.2 Application Aware vSphere HA

As of vSphere 4.1 an application programming interface (API) was introduced to provide third-party vendors with the ability to integrate with HA. The capability to allow application monitoring agents to interact with HA is enabled per vSphere cluster with additional configuration options available per virtual machine. When enabled, this feature allows application monitoring agents to send application heartbeats to HA. In the event of an application-level failure the application monitoring agent can take action to either bring the application back online, or stop the application heartbeat, causing HA to initiate a restart of the virtual machine. Prior to vSphere 4.1, only VMware tools heartbeats could trigger HA restarts.

4.2.1 Example: Exchange Mailbox Server Virtual Machine Protected by Application Aware HA

With Application HA a third-party agent, installed in the virtual machine guest OS, monitors the application and its dependencies. When a failure is detected the agent can restart services, mount databases, or invoke HA to initiate a virtual machine reset.

Figure 2. Mailbox Server Virtual Machine Protected with Application Aware HA




4.3 Microsoft Database Availability Groups

Microsoft Database Availability Groups (DAGs) have changed the traditional server-failover model to a database-failover model, where individual databases can fail over to any Exchange server in the DAG. Database Availability Groups provide a non-shared storage failover cluster solution. DAGs use built-in asynchronous log shipping technology to distribute and maintain passive copies of each database on Exchange DAG member servers. Multiple copies of each database can be maintained, to protect from the loss of one or more Exchange DAG member servers. This solution requires the use of Windows Server Enterprise licenses, additional storage capacity to accommodate each additional copy of a database, and a thorough understanding of the additional load passive mailboxes and databases place on mailbox role virtual machines.

4.3.1 Example: Three Exchange Mailbox Server Virtual Machines in DAG

Deploying your database availability groups on a virtualized Exchange 2010 platform allows for consolidation of hardware as well as tiering of your resources, if desired. If a virtual machine experiences operating system or application level corruption the remaining passive databases are activated and continue serving mailbox data to clients. In the event of an ESX/ESXi host failure the remaining passive databases are activated and vSphere HA can power on the affected virtual machines on any remaining hosts to reestablish redundancy in the DAG.


Figure 3. Mailbox Server Virtual Machines in Database Availability Group



4.3.2 Example: vSphere HA Used with DAG Clustering for Faster Recovery

Database level high availability can be achieved through the use of database availability groups. In the event of a server host failure, a passive copy of the affected mailbox databases is activated. Client access servers establish MAPI connectivity to the newly active database copy and client connections are reestablished. In the background, vSphere HA powers-on the failed virtual machine on another server host, restoring the DAG membership and bringing the newly passive database up to date and ready to take over in case of a failure, or to be manually reactivated as the primary active database. The use of database availability groups on top of hypervisor based clustering is supported by Microsoft as of Exchange 2010 SP1. VMware tests have shown that the two technologies can coexist and can be a viable solution to provide maximum recoverability in the case of a host failure. Best practices for deploying HA, vMotion, and DRS with Exchange 2010 DAGs can be found in the whitepaper *Using HA, DRS and vMotion with Exchange 2010 DAGs* (<http://www.vmware.com/files/pdf/solutions/VMware-Using-HA-DRS-vMotion-with-Exchange-2010-DAGs.pdf>).

Figure 4. Database Availability Group with vSphere HA



5. Remote Site Availability Options

Incorporating disaster recovery has become a leading business objective for new deployments and upgrades. Leveraging built-in solutions is ideal for protecting a single business critical application from disaster. Some environments require application and hardware agnostic methods for protecting their mission critical applications and data. Deploying on vSphere provides the flexibility to meet all of these requirements.


5.1 VMware vCenter Site Recovery Manager

VMware vCenter Site Recovery Manager™ (SRM) makes disaster recovery rapid, reliable, manageable, and affordable. SRM leverages VMware vSphere and leading partners' storage replication software to deliver centralized management of recovery plans, automate the recovery process, and enable dramatically improved testing of recovery plans. It transforms the complex hardcopy run books associated with traditional disaster recovery into an integrated element of virtual infrastructure management. VMware vCenter SRM enables organizations to take risk and worry out of disaster recovery—yet another reason the VMware virtualization platform is the safest platform for datacenter applications.

5.1.1 Example: SRM and Exchange 2010 DAGs

Using Exchange 2010 DAGs within the datacenter to provide high availability will meet the requirements of most organizations. With DAGs, a failover can occur automatically at the database or server level and can take place within 15-30 seconds of a detected failure. When designing a disaster recovery solution automated failover is usually not a desirable feature. In many cases a DR facility is designed with a lower SLA and has a slightly delayed version of data than the production facility. Making the choice to activate the DR facility should be a conscious decision that follows an organization's change process. With SRM and a VMware partner's storage replication solution, disaster recovery can be implemented to protect the entire virtual datacenter, including Exchange. Failover testing can be accomplished with no production impact to confirm that the recovery-time and recovery-point objectives are being met. Additionally, customizable recovery plans allow for adding custom scripts, virtual machine power-on priority, and breaks.

Figure 5. VMware vCenter Site Recovery Manager with Database Availability Groups



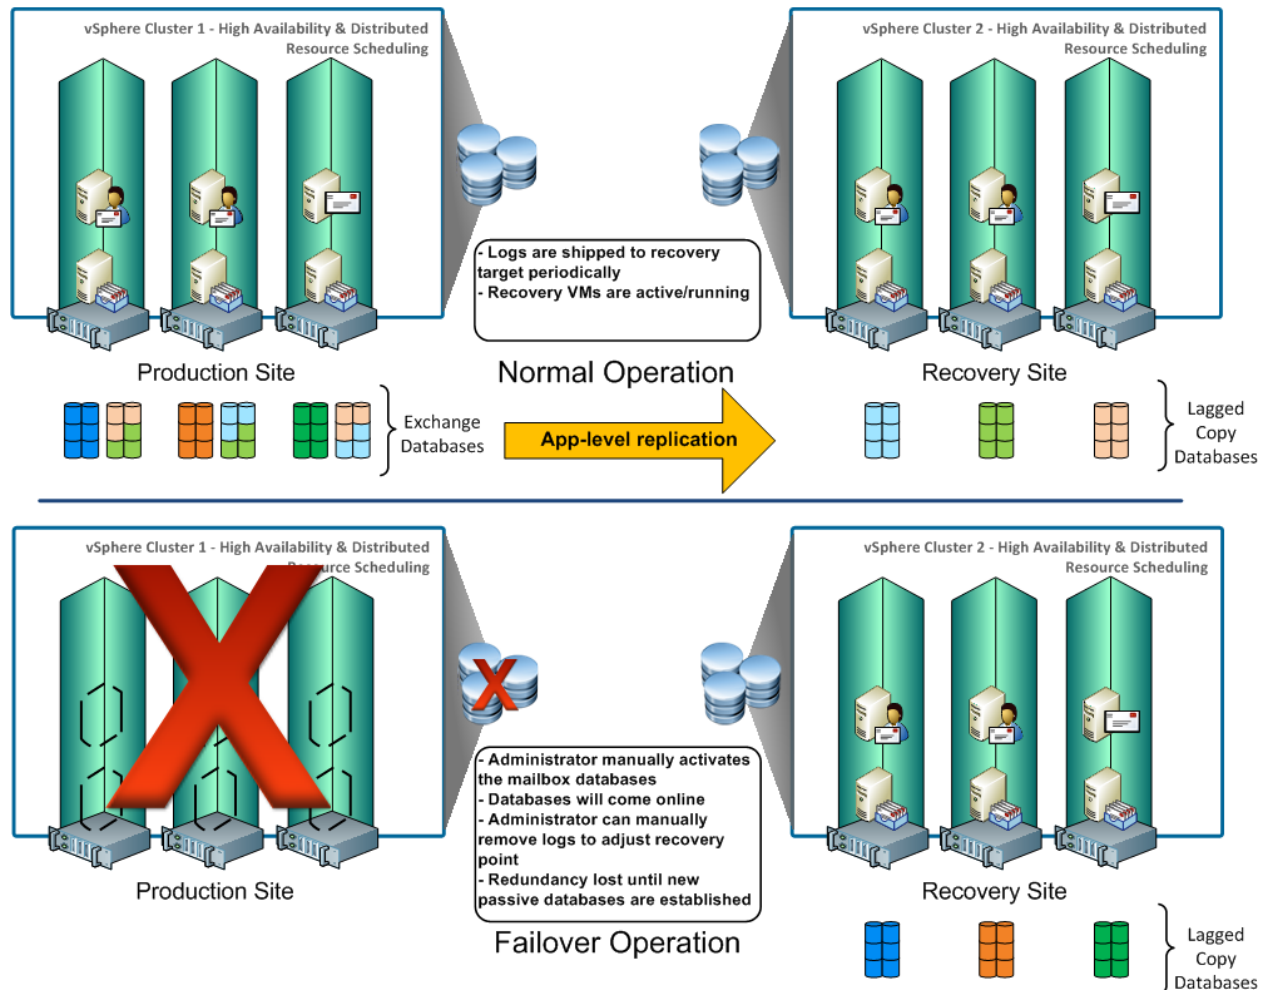
5.2 Exchange 2010 DAG with Delayed Log Replay

Database copies support lagged transaction log replay. This is an improvement to the Exchange 2007 feature standby-continuous replication (SCR). SCR uses a hard-coded reply lag of 50 transaction log files. With Exchange 2010 and DAGs a mailbox copy can be configured with an administrator-defined lag time, in minutes, to delay the replaying of log files into the passive database copy. Replay lag provides protection against logical database corruption by providing the ability to recover up to the last copied and inspected log file, or to a specific point-in-time (PIT) within the lag window by manipulating the log files and running `eseutil`.

5.2.1 Example: DAG with Delayed Log Replay

To maintain high availability within the datacenter, and also provide disaster recovery protection, a third database copy can be established at a DR datacenter. By implementing delayed log replay you can avoid logical corruption by requiring manual activation of the DR databases. Testing of the disaster recovery plan should be performed to familiarize the operations team with the `eseutil` tool, which is required for log replay. When testing disaster recovery plans, take into account any impact disaster recovery testing may have on active client connections.

Figure 6. Recovery Process using DAG and Delayed Log Replay



5.3 Third-party Software-based Replication

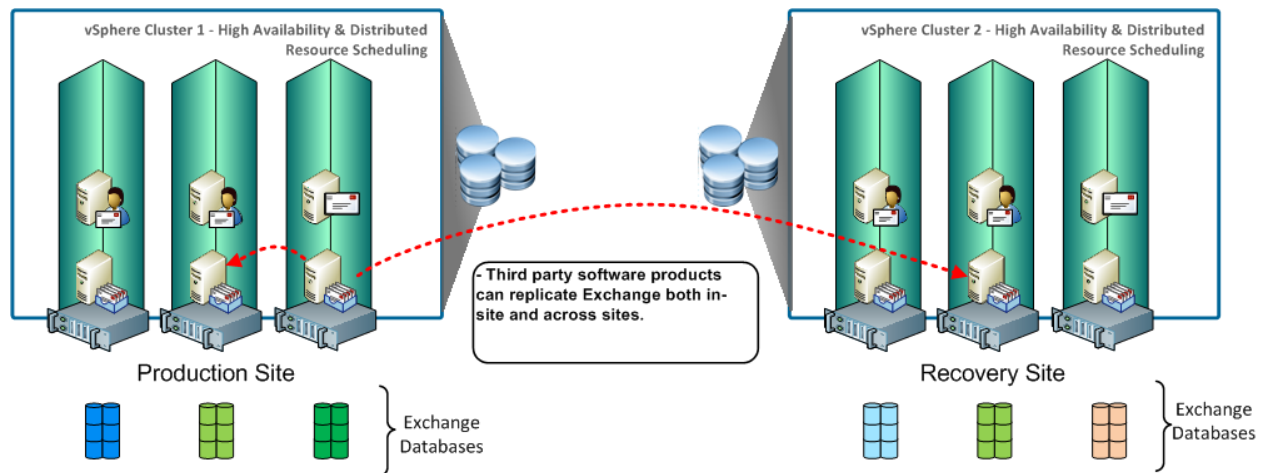
Software-based replication products from third-party vendors provide high availability from within the guest OS. This is similar to deploying a DAG, but has some additional benefits. By not deploying DAGs you can eliminate the need for enterprise-based Windows Server licenses, and have near zero-downtime planned migrations and failovers. Software-based replication typically depends on an agent running in the guest OS to replicate transactions to the remote location. These in-guest agents allow for application-aware replicas similar to those offered with DAGs.

VMware has a comprehensive list of ISV partners that provide software-based replication of Exchange servers to a DR site.

5.3.1 Example: Exchange Mailbox Virtual Machines Using Third-party Software-Based Replication

A stand-alone Exchange mailbox virtual machine can be configured with third-party software-based replication solution to provide disaster recovery protection. An agent installed in the guest OS replicates transactions using the third-party software vendor's proprietary software to provide consistency, and provides monitoring for automated failover in the case of a guest OS or application level failure. The target location can be within the same vSphere HA cluster or across a WAN link to your disaster recovery site.

Figure 7. Exchange Mailbox VMs Using Third-Party Software-Based Replication



6. Backup and Restore Options

The feature set available to an application, when deployed in a virtual environment, is no different than what is available with a physical deployment. In fact, there are more options available for protecting entire virtual machines. This is especially useful for applications that require extensive configuration. For Exchange 2010 the standard methods for backup are supported. These tend to be deployed using a third-party backup agent that uses a VSS requestor to coordinate with the VSS writer to prepare the database files for backup. Regardless of the backup solution required, VMware and VMware partners have provided solutions for most situations.


6.1 In-Guest Software Solutions

Many organizations have dedicated backup support teams or requirements that may not allow them to integrate the backup solution to the level that is available with vSphere. In these situations, traditional backup methods are desirable, and a virtualized environment allows for their use. Many of the leading backup software providers are VMware partners and provide full support for using their backup solutions within a virtualized guest operating system. Backup administrators can continue to deploy and manage the backup agents, jobs, and restores as though they were running on physical systems.

6.1.1 Example: In-guest Exchange-aware backup solution

Centralized backup management software controls the backup schedule, save set, and target location for all systems, both virtual and physical. Backup agent software loaded within the guest operating system allows the virtual machine guest OS to be managed in the same way as all other systems. Additional plug-ins from backup software vendors provides application aware support.

Figure 8. In-Guest, Agent-Based Backup Process



6.2 VMware Data Recovery

VMware Data Recovery (VDR) protects your data at the virtual machine level capturing application and system data as a full virtual machine image. VMware Data Recovery runs at the ESX/ESXi host level as a virtual appliance to provide streamlined deployment and full integration with vCenter Server. VDR stores multiple restore points for each virtual machine using *deduplication* technology to not only provide point-in-time restore capabilities, but efficiently use available disk space.

As most Exchange administrators know, the use of an Exchange-aware backup agent provides database health checking and log truncation. While VDR can use the VSS framework to back up Windows guest OS virtual machines it does not contain the required VSS components to properly restore an Exchange mailbox database, although the backups created do use VSS and are application consistent. Exchange peripheral roles such as the client access, hub, and edge transport roles are relatively stateless and may benefit significantly from being protected using VMware Data Recovery.

6.2.1 Example: Using VDR with Exchange Client Access and Hub Transports

The Client Access and Hub Transport roles do not store end-user data, but they do contain data that is critical to the functionality of the roles. You can quickly recover a failed Client Access or Hub Transport virtual machine by using VMware Data Recovery to back up the entire contents of the virtual machine. In the case of guest OS corruption or unsuccessful changes to the Exchange configuration the entire virtual machine can be restored in minutes using the latest VMware Data Recovery back up image.

Figure 9. Using VDR with Exchange Client Access and Hub Transports

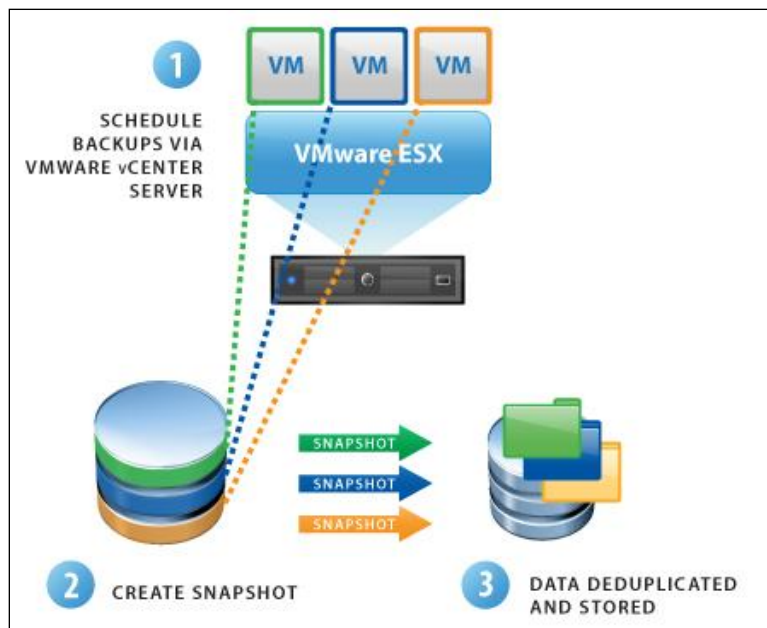
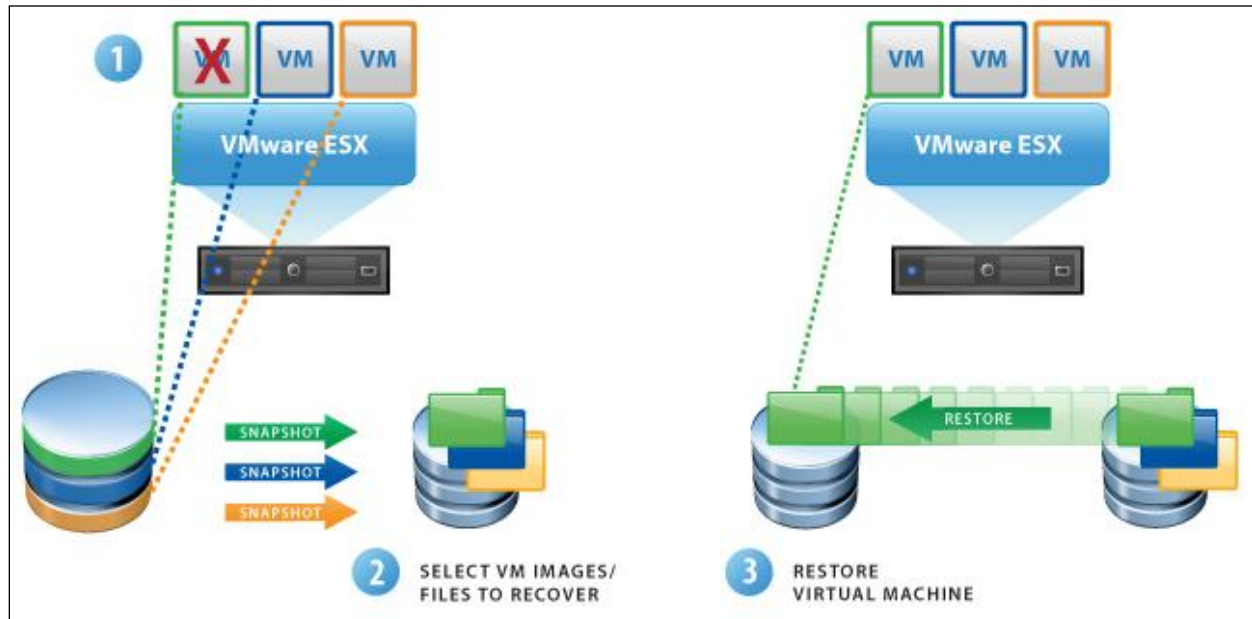


Figure 10. VDR Backup and Recovery Processes



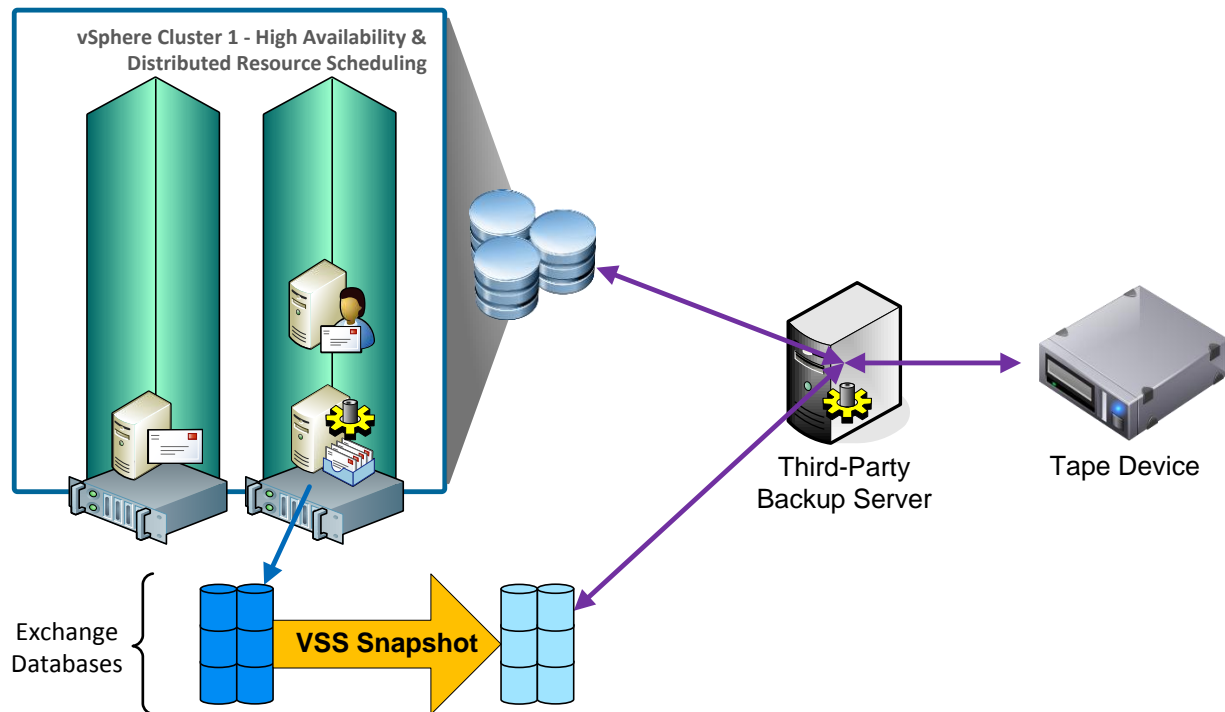
6.3 Array-Based Backup Solutions

As is the case with the in-guest solutions, array-based solutions provided by many of the leading storage vendors continue to work with vSphere deployments of Exchange. Array-based backup solutions for Exchange use the Volume Shadow Copy Service (VSS) available with Exchange 2010 to produce near-instant, application-aware clones or snapshots of Exchange databases. These local clones or snapshots can then be backed up to disk, tape or cloned offsite for disaster recovery purposes. Guidance on proper deployment methods and any additional considerations when running in a virtualized environment must be provided by the storage vendor. VMware has a comprehensive list of ISV partners that provide array-based replication of Exchange servers for backup and restore operations.

6.3.1 Example: Exchange Mailbox Server Virtual Machine with Array-based Backup Solution

An array-based backup solution provides integration with the Exchange application and the underlying storage solution. A software agent provided by your backup vendor coordinates with the Exchange VSS writers to create a supported backup image of your Exchange databases. These databases can later be streamed to tape as flat files for compliance or archive requirements with no IO impact to the production data.

Figure 11. Exchange Mailbox Server Virtual Machine with Array-based Backup Solution



7. Additional Information

VMware vSphere offers many tools and features to increase the availability of Exchange 2010. vMotion, HA, and DRS can help reduce downtime and improve flexibility in your Exchange 2010 architecture while lowering costs. In some cases VMware and Exchange features can be combined to improve overall availability. However, there are a few things to keep in mind when architecting a solution:

- vMotion, HA, and DRS are supported for all Exchange 2010 roles including DAG nodes.
- Exchange 2010 DAG nodes can use Fibre Channel or iSCSI attached storage for Exchange data files. This pertains to storage handled by ESX hosts (RDMs or virtual disks on VMFS volumes).
- Network Attached Storage is not currently supported for Exchange regardless of whether the Exchange server is physical or virtual.
- Use of software iSCSI initiators within guest operating systems configured with MSCS or DAGs, in any configuration supported by Microsoft, is transparent to ESX hosts and there is no need for explicit support statements from VMware.

Visit www.vmware.com/go/exchange for more information.