

Introduction

VMware Infrastructure is deployed in data centers for deploying mission critical applications. Deployment of Microsoft Exchange is a very important task for the IT staff. Email system is an extremely critical piece of business operation. Exchange must thus be kept always available. Not only availability, but the Exchange data must be protected so that quick restore and recovery can be done.

Protecting data is an important task for any Exchange deployment from not only operational point of view, but also from the point of view of legal compliance.

Many data protections techniques are available on the market that can effectively protect Exchange database. However, they tend to be complex and expensive to implement. Especially in the physical (non-virtual) implementation of Exchange, servers and applications are protected using complicated “bare metal recovery” systems. These technologies, which are usually adopted from legacy applications are not especially well suited for Exchange environments.

When Exchange is deployed in virtual environment, VMware Infrastructure’s built-in availability and data protection tools can be used to offer superior protection to Exchange deployments. Old, complex methods of providing data protection for servers can immediately be replaced by using VMware Consolidated backup technology.

At the same time, either VMware based or storage based snapshot mechanisms can be used for protecting the Exchange database.

VMware Infrastructure

VMware Infrastructure includes two main components: VMware ESX Server, which is the high performance and production proven hypervisor; and Virtual Center, the management interface. Apart from hypervisor, ESX server also provides clustering services, VMware’s high

performance file system store (Virtual Machine File System, VMFS) as well as networking support. Exchange server is deployed in the guest OS that runs inside the virtual machine. The virtual machines run on the ESX server hypervisor. Virtual machines are stored on the VMFS data stores.

VMFS provides the locking mechanism for virtual machines, used for clustering solution as well as snapshots, which are leveraged for implementing the data protection solution. Using VMware Dynamic Resource Scheduler (DRS) and VMotion technology, customers can increase the utilization of their Exchange deployment at the same time ensuring that Service Level Agreements (SLAs) are met.

Setup of Exchange in a VM

Customers use virtualization for Exchange for a variety of reasons. Some of the most common reasons for deploying Exchange in a virtual infrastructure are:

- Improve server utilization of Exchange servers. Many Exchange servers are over-provisioned and tend to be under utilized in normal times. They are deployed with the peak usage in mind. In non-virtual environments, there is no way to use this spare capacity. With consolidation, VMware Infrastructure allows you to use the spare capacity of Exchange servers at the same time re-provisioning the resources to Exchange when needed.
- Better use of Shared storage: Exchange can utilize advanced capabilities of Storage systems like SAN arrays to implement a faster, efficient solution. With VMware ESX server, customers can use the shared storage to deploy less complex, and high performing clustering and disaster recovery solutions.
- Simplify data protection and disaster recovery. Virtualization enables customers to choose between using newer snapshot technologies or continue to use legacy way of protecting application and data.

- Increase availability and load balancing. Using innovative technology like VMotion, customers can increase utilization of their existing resources. Also VMware HA provides protection against downtime due to server hardware failures, by automatically restarting the virtual machines.

Importance of data protection

Database is the most important part of Microsoft Exchange deployment. It stores all the configuration information about the users, their status in the system as well as all the email and other type of collaboration information like calendars etc. Loss of these records could be catastrophic. Data protection is thus an important task for Exchange administrators.

Exchange administrators routinely are called upon to restore individual mailboxes. Sometimes they are also required to restore the entire database, lost due to some component failure or software corruption etc.

Companies are also required by law to store or archive all the emails for a fixed amount of time, during which time, companies can be asked to furnish the email records.

Many technologies are available on the market to help protect data. In addition to offering the unique set of advantages to customers for protecting exchange application and data, VMware Infrastructure enables customers to select best of breed tools for data protection.

Options in data protection

With virtualized Exchange, customers can use any number of backup solutions for data protection. It is important to note here that in order to have the complete solution, it is essential that both application and data be protected. Protecting both application and data, guarantees faster recovery and lower downtime.

There are three different options when deploying backup for Exchange. Choice of the effective data protection solution depends on user needs, storage device capabilities and resources available.

Storage for Exchange database can be configured in one of three ways to work with virtual machines.

- Database in a virtual disk. Exchange database can be implemented in a virtual disk (.vmdk file), stored on a virtual machine file system (VMFS). The database is accessed by the Exchange application running inside the VM as a raw disk. This .vmdk file can be stored on a VMFS volume, separate from the virtual machines system disk, in order to get maximum storage throughput.
- Database stored on RDM volume (physical compatibility mode): Raw Device Mapping is a special technology built into ESX server that allows users to access storage LUNs directly, thereby bypassing VMFS file system. The RDM can be configured to entirely bypass VMFS, so all writes go directly from the VM to the device. This mode of operation is called physical compatibility mode RDM.
- Database stored on RDM volume (virtual compatibility mode). Like above, virtual compatibility mode RDM volume bypasses VMFS. However, in virtual compatibility mode, VMFS can temporarily cache the writes going from the VM to the device. That is the only difference between the physical and virtual compatibility modes of RDM operation.

Protecting Exchange Application

Whereas Exchange data can be stored on the VMFS volume inside a vmdk file or on an independent LUN as an RDM, the Exchange application benefits by being installed in the virtual machine. Exchange application can reside on the system drive for the virtual machine.

VMware


Data protection for the virtual machine with Exchange application installed can be implemented using VMware Consolidated Backup (VCB).

VCB is a framework technology that enables data protection software to create snapshot copies of virtual machines without first off-lining the VM. The image is then transported to secondary medium (like tape library or disk library) over SAN.

This methodology has the advantage that the backup of the virtual machine can be performed at any time. A new full backup of the virtual machine can be created any time there is some configuration changes to the OS or the application or when the software patches are applied. This allows for faster recovery in case of the failure or loss of data.

There is no need to backup the server or the application, every time database backup is performed.

Figure below shows how VCB works.



Protecting Exchange Database

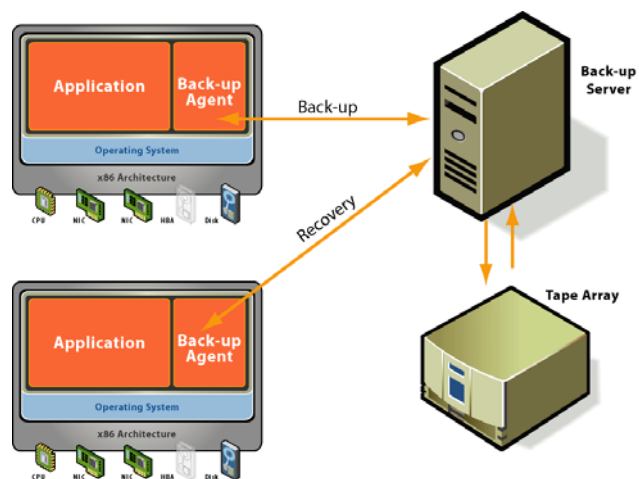
Backing up Exchange database is very important for the businesses. Exchange database is critical to operation of the messaging system. Database back up is done frequently with an eye towards restorability of the backup. Recovery point objective for mission critical Exchange database is rather stringent. Exchange deployment in VMware virtual environment allows backup administrators flexibility of data protection unmatched in non-virtualized environments. The

following subsections describe each option in detail.

Option 1: Database in VMDK

When Exchange database is kept on a volume that is stored on VMFS file system as a VMDK file, it can be protected using a backup agent installed inside the virtual machine. The backup agent helps in backing up the database for easy restore of the entire database or even individual mailboxes. Some backup agents allow the restore of individual messages.

If you already have exchange protected using backup agents on non-virtualized servers, your backup methodology does not change at all for virtualized exchange implementation. Figure below shows how backups work with agents inside the VM.



The advantages of using backup agent inside the VM to protect Exchange database are:

- Stream based backup to protect the database works exactly same for non-virtualized and virtualized environments. The agents are also application aware, so application consistent backup can be done
- Mailbox level backup can be performed for better control of backup data. Since the agent is aware of exchange data structures, it can backup database incrementally. This helps in reducing the amount of data backed up.

- Best for recovery of single message/mailbox for ease of restore. Customers can restore not only entire databases, but also individual mailboxes or even individual messages. This is important if users accidentally delete a message or their mailboxes corrupt. In many organizations, this service is limited to only executives and senior management, rather than deployed generally. That is because of the overheads attached with backing up and restoring individual mailboxes.

Option 2: Database on Physical RDM LUN

In a few installations, customers want to dedicate an entire storage LUN to exchange database. This implementation is driven by the rationale that storage arrays or storage devices can most effectively handle the data protection or replication services.

This type of implementation is especially used when customers have invested in buying high-end storage for their Exchange implementation. VMware Virtual Infrastructure supports this type of Exchange implementation. In VMware infrastructure, raw LUNs can be connected to the VMs directly using one of two modes. One is "physical compatibility" mode Raw Device Mapping LUN.

When RDM LUN is attached to the VM in physical compatibility mode, data traffic bypasses VMFS to go directly from the VM to the LUN. In those cases, VMFS does not quiesce the VM and no snapshots can be taken using VMFS. All snapshot and cloning activity must happen inside the storage device.

If the administrator chooses to connect a RDM LUN in physical compatibility mode, backup agents inside the VM may be configured to work with the storage device to take a snapshot of the entire LUN. Those snapshots can then be moved to the secondary storage as backup. Some backup vendor's solutions are also capable of creating clones or keeping track of snapshots of quick recovery within the storage array. Following are the advantages of using this method for protecting exchange data:

- Enables array-snapshot based backup for quick backup. Array based snapshots can be quick and efficient. This method usually does not involve any downtime for application and is done without any significant performance impact.
- VSS integration through backup agent inside the VM is possible. Backup agent inside the VM is exchange aware. As such, the application can be quiesced before the snapshot is taken. The resulting image of the data (snapshot) is application crash consistent.
- Online Snapshot based, off-host backup can also be supported. Some backup application vendors can work with array vendors to move the data directly over SAN from storage array to the backup media.
- Complete DB restore for instant recovery can be implemented. By keeping the snapshots over storage and if backup software can keep track of those snapshots, some vendors can implement instant recovery for the entire Exchange database.

Option 3: Database on Virtual RDM LUN

Customers have a choice when connecting an entire LUN to a single VM for use as Exchange database storage. They can opt to connect the Raw Device Mapping LUN in "virtual compatibility" mode. When connected in this mode, it enables customers to use VMFS snapshots to backup entire exchange database. Especially when used with storage that is not licensed for LUN snapshots, this method gives customers the ability to do snapshot of databases. The following list describes the advantages of this method:

- VCB can be used for creating the snapshot of the RDM. Since VCB uses VMFS snapshot mechanism, RDMs connected in virtual compatibility mode can be snapshot using VCB.

VMware

- Off-host and online backup of the database is possible using this method. VCB snapshots can be migrated to proxy server and backups can take place on the backup servers.
- Complete DB restore can be done. Backup software on proxy can restore the database to the LUN for complete database restore.
- VCB cannot yet quiesce the exchange database before taking snapshot. It is possible to write pre- and post-snapshot scripts for maintaining database consistency. Using these scripts, it's possible to quiesce and un-quiesce the database for crash consistent backup. Barring the scripts, VCB backup must be done for either cold or warm database only. Hot consistent backup (online) can only be done using either option 1 or option 2.
- Secondary storage on same SAN is the best and fastest way to restore the data. Since the data is transferred over SAN rather than network, its faster to restore the data from secondary storage or primary storage.
- Depending on the restore needs for the Exchange database, it may make sense to keep a few snapshots on primary storage for instant recovery.

Recovery and Restore

When protecting Exchange database, it's important to understand the restore limitations of each type of backup. When an entire database is backed up in "database backup" mode, only full database restore can be performed. It is often not possible to restore single mailboxes from database backup of the Exchange database. This function is dependent on your backup software and works the same way for virtualized and non-virtualized environment.

Keep the following tips about restore and recovery of database in mind when backing up the database:

- No single Mailbox restores/recover with database-level backup. No matter which option you utilize to backup the database, you generally cannot restore single mailboxes from database level backups.
- Complete database restore is more efficient than restoring single mailboxes, when restoring and recovering from large data or storage data losses.

