

Best practices for virtualizing and protecting Exchange Server

Brien M. Posey

Modern Data Protection
Built for Virtualization | **#1 VM Backup**

As is the case with most Microsoft products, Microsoft provides a number of requirements and best practice guidelines for running Exchange Server in a virtualized environment. While there is no denying the importance of adhering to these requirements, it is also important to read between the lines and understand the spirit of the recommendations that Microsoft makes. Doing so will allow you to virtualize Exchange Server in a way that will provide optimal performance, stability and fault tolerance.

Today, the reality is that the resources are in place to virtualize Exchange successfully on vSphere and Hyper-V.

Background on Exchange Server

As you prepare to virtualize Exchange Server, it is important to understand that Exchange Server has been around for a really long time. Although Microsoft has redesigned Exchange Server many times over the years, Exchange Server 2013 still has some things in common with much older versions, such as Exchange Server 5.5.

The reason why this is important is because it plays into the philosophy that you must bear in mind as you prepare to virtualize Exchange Server. Although I am sure there are some who would disagree, Exchange Server was not specifically designed to run in a virtual environment. Yes, Microsoft fully supports running Exchange 2013 on virtual servers. However, the fact that Exchange Server 2013 inherited some of its standards and methodologies from much older Exchange versions means that properly virtualizing Exchange Server can be tricky.

With that in mind, there are two very important things to consider. If you keep both of these things in mind as you virtualize Exchange Server (while also adhering to Microsoft's requirements), you will likely create a virtualized Exchange Server deployment that works very well.

The first thing to keep in mind is that because of its heritage, Exchange Server is really designed for use with physical hardware. As such, any time that you can give Exchange the illusion that it is running on physical hardware, it tends to be a good thing. You will see various examples of this later in this white paper, but one quick example is that Exchange Server tends to perform better if you do not use dynamic memory.

The other important thing to keep in mind is that Exchange Server really doesn't like it if you attempt to manipulate state data. This means that you can run into problems with things like resuming Exchange Server from a saved state or rolling back a snapshot. Again, you will see examples of this as we go along.

The important thing to remember is that as you virtualize Exchange Server, you should work hard to avoid any type of situation that dynamically adjusts the hardware resources that are available to Exchange, overcommits hardware resources or that manipulates state data. These are by far the most important things to keep in mind as you virtualize Exchange Server.

Use a supported hypervisor

Any time that you virtualize a mission-critical resource such as Exchange Server, it is extremely important to do so in a way that is fully supported. As such, you must make sure that the hypervisor that you are using is supported by Microsoft.

Exchange Server 2013 is supported for use on Microsoft Hyper-V 2008 R2 and higher. It is also supported for use on vSphere. That's not to say that these are the only hypervisors that are supported. Microsoft has a Windows Server Virtualization Validation Program, and any hypervisor that is certified by this program can be used with Exchange Server 2013.

There is, however, one exception that seems less than intuitive. Exchange Server 2013 is not certified for use on Windows Azure virtual machines (VMs)—at least not yet. This isn't to say that it doesn't work. It is technically possible to run Exchange Server 2013 on top of Windows Azure. However, doing so should be limited to lab environments only until Microsoft eventually certifies Windows Azure for use with Exchange. There are scores of aggregated best practice resources for virtualizing Exchange, and here are two links to keep handy:

Best practices to virtualize Exchange on VMware:

<http://vee.am/bpexchangeVMware>

Best practices to virtualize Exchange on Hyper-V:

<http://vee.am/bpexchangeHyperV>

Hardware resource allocation

One major key to making Exchange Server perform well in a virtualized environment is to make sure that sufficient hardware resources have been allocated to Exchange and to avoid overcommitting or dynamically adjusting those allocations.

Memory allocation

In a popular TechNet blog, Microsoft states that you should “use dynamic memory on all VMs (unless not supported).” However, Exchange Server 2013 is one of those applications for which the use of dynamic memory is not supported. The Microsoft support guidelines for Exchange Server 2013 state:

Exchange, like many server applications with optimizations for performance that involve caching of data in memory, is susceptible to poor system performance and an unacceptable client experience if it doesn't have full control over the memory allocated to the physical or virtual machine on which it's running. As a result, using dynamic memory features for Exchange isn't supported.

CPU allocation

Memory allocation is not the only hardware resource that must be planned carefully when virtualizing Exchange Server 2013. Administrators must also take care to provide Exchange Server with sufficient CPU resources.

In a Hyper-V environment, each CPU core is treated as a logical processor. However, Hyper-V does not directly allocate logical processors to VMs. Instead, Hyper-V CPU allocations are based on the concept of virtual processors. Virtual processors have no direct real-world correlation to logical processors. As such, a system's CPU cores can easily be oversubscribed by allocating far more virtual processors than the number of logical processors that exist in the system.

The interesting thing about this overcommitment is that Hyper-V servers are rarely CPU bound. I have spoken with engineers for some prominent server manufacturers who have recommended ratios of virtual processors to logical processors of up to 8-to-1. In the case of Exchange Server 2013 however, Microsoft discourages logical processor overcommitment.

If you do plan to overcommit your logical processors, Microsoft recommends a virtual processor to logical processor ratio of no more than 2-to-1. However, they strongly recommend a ratio of 1-to-1.

Another aspect of CPU allocation that is sometimes overlooked is that logical processors must be reserved for use by the host operating system (OS). Microsoft's recommendation is to reserve two logical processors for the host OS, and then use Performance Monitor to keep track of how well the host OS is functioning. If the host OS is suffering from insufficient CPU resources, Microsoft recommends taking away virtual processors from some of the VMs in an effort to free up CPU cores for use by the host OS.

This raises a couple of questions. First, how can you make sure that CPU resources are being reserved for use by the host OS? Second, what types of processes running on the host OS could consume excessive CPU resources?

Unfortunately, the Hyper-V Manager does not provide a mechanism for reserving logical CPUs for use by the host OS. System Center Virtual Machine Manager (SCVMM) does allow you to create host reserves, but the concept can be a little bit misleading.

When you create a host reserve, you can specify the resources such as CPU, disk space and memory that will be required by the host OS. However, these resources do not actually get locked away. SCVMM does not do any sort of metering on the host. Instead, host reserves are used when an administrator deploys a VM. The available hosts are all rated based on their suitability for accommodating the VM. The host reserves are taken into account when calculating this rating so as to avoid depleting a host of resources. However, it is ultimately up to the administrator to determine where they want to place a VM, and if a VM is placed on an unsuitable host, SCVMM does not enforce the host reserves. With that being the case, it is better to think of a host reserve as a prescriptive guideline rather than a true reservation.

So what types of activities can cause the host OS to consume more CPU resources than can be provided by two CPU cores? Under normal circumstances, the host OS should not consume very many CPU resources. However, there are some types of activities that can produce a CPU load. This includes things like compression (which is now an option for Live Migration in Windows Server 2012 R2), post-process deduplication or even an overly aggressive anti-malware scan.

The Microsoft support policy for virtualized Exchange Servers prevents the host OS from running any application aside from those used for day-to-day system maintenance. This helps to limit the host OS CPU consumption. For example, the support policy would allow for antivirus software or for a backup agent, but would prohibit running something like SQL Server within the host OS.

Storage considerations

In the physical world, one of the most difficult aspects of Exchange Server deployment planning is that of planning the storage subsystem. Not only must be storage comply with Microsoft's requirements, but it must also perform well and deliver resiliency. The challenges of storage planning are compounded in a server virtualization environment because the administrator is forced to deal with logical storage (the storage that is visible to Exchange Server) and physical storage (the storage that is available to the hypervisor).

Virtual Hard Disk considerations

When it comes to virtual hard disk (VHD) planning, the first best practice that you should adhere to is to use the new VHDX format whenever possible. VHDX files provide a number of benefits, but are only supported by Windows Server 2012 and Windows Server 2012 R2 with Hyper-V. Comparatively, the VMDK format now supports geometries up to 62 TB with vSphere 5.5.

VHDX-based VHDs increase the maximum disk size from 2 TB to 64 TB, and also offer better performance through improved sector alignment. This format also includes features that make them more resistant to corruption in the event of a power failure. The only situation in which you should avoid using VHDX files is one where the virtual Exchange Server could potentially be moved to an older version of Hyper-V.

Regardless of whether you choose to use the VHD or VHDX format, there are several other considerations you must also take into account. One such consideration is the type of storage provisioning that will be used during the creation of the VHD. As you are probably aware, Hyper-V will allow you to create either a fixed-length VHD, or a thinly-provisioned, dynamically-expanding VHD. This capability is not unique to Hyper-V. VMware and some other competing hypervisors offer similar abilities.

Microsoft and other hypervisor vendors give you a choice of the type of provisioning you want to use because thinly-provisioned VHDs conserve storage space, while fixed-length VHDs perform better.

At the beginning of this white paper, I stated that Exchange Server was not designed with server virtualization specifically in mind and that it is important to try to provide virtualized Exchange Servers with virtual hardware that closely mimics physical hardware. One particularly good example of this is storage provisioning. Microsoft does not support the use of dynamically-expanding VHDs for virtualized Exchange Servers. You must use fixed-length VHDs.

Of course the disadvantage to creating fixed-length VHDs is that you must decide up front how large you want to make them. Windows Server 2012 R2 with Hyper-V supports dynamically-expanding VHDX-based VHDs while a VM is running, but doing so can lead to problems with fragmentation and decreased performance. Furthermore, resizing a VHD does nothing to resize the volumes that exist on that VHD. As such, it is far easier to create a correctly-sized VHD up front rather than creating one that is too large or too small and then trying to modify it later.

Of course, storage capacity planning is something of an art form. The amount of storage space that will be required can vary greatly depending on how a particular VHD is going to be used. For instance, a VHD that contains the guest OS or the Exchange Server binaries will almost always need far less space than a VHD that contains a mailbox database.

One common mistake that is sometimes made with regard to provisioning a virtualized Exchange Server system volume is not allocating sufficient space. Obviously the VHD will need enough space to accommodate the OS files. This volume will also typically be used to store the Exchange binaries. However, the VHD will likely need quite a bit more space than what these elements collectively consume.

One of the big reasons for this is that the system volume is typically responsible for storing the pagefile. In a physical environment, the pagefile is typically equivalent to the amount of memory that is installed in the server. The same basic concept also applies to virtual servers, but they can be a bit more problematic.

The reason for this is that virtualization administrators are often anxious to achieve the highest possible VM density. They accomplish this by making efficient use of hardware resources and by allocating the fewest resources possible. This can be a problem because an Exchange Server may need to be provisioned with additional hardware resources as its workload increases.

Suppose you were to add an extra 50 mailboxes to a virtualized mailbox server. In order to ensure the virtual server continues to run efficiently, you might need to provision it with some additional physical memory. Upon doing so, however, the size of the VM's pagefile will also increase. If the VHD has been designed to use an absolute minimum amount of space, it might not have sufficient capacity for an increased pagefile.

As you can see, it is a good idea to provision virtualized Exchange Servers with larger VHDs than they really need. Not only can doing so help to make the virtual server more scalable, but the extra space may also come in handy if you ever decide to perform an in-place upgrade to a new OS or to a new version of Exchange Server. Of course some would be quick to point out that in an Exchange Server environment, migrations are usually preferable to in-place upgrades. Even so, the same basic principle of saving some extra space for maintenance still applies. After all, the VM will need sufficient storage space to accommodate things like service pack upgrades.

While I am on the subject of VHDs, I also want to point out that VMs consume storage space outside of the VHDs. A small amount of storage space is used for maintaining the VM configuration, but larger amounts of storage space are used for things like snapshots (which are not supported for use with Exchange Server) and for temporary memory storage.

I already explained that the Windows pagefile usually exists within a VHD, but there are at least two ways in which Hyper-V uses temporary memory storage outside of the VHDs. The storage must be planned for so that it can be properly accommodated.

The first way in which temporary memory storage is used is in the form of a BIN file. The BIN file is equal in size to the amount of memory that has been allocated to the VM.

The second temporary memory storage mechanism only comes into play in Windows Server 2012 and Windows Server 2012 R2 with Hyper-V environments. In Windows Server 2012, Microsoft made a change to the dynamic memory feature that allowed a VM to release some of its memory after the routing and run on a minimal amount of memory. The problem is that if a VM is operating with less physical memory than is required to boot, a reboot could result in a situation where insufficient physical memory is available to complete the boot process. In these types of situations, Hyper-V uses a smart paging file to make up for shortages in physical memory and allow the boot process to complete.

As previously mentioned, dynamic memory is not supported for use with Exchange Server 2013 VMs. It might, therefore, seem that accommodating the smart paging file would be a moot point. However, it is common for multiple VM configurations to share a common storage repository which means an unrelated VM could potentially consume storage on a volume that contains the virtualized Exchange Server's configuration files.

Physical storage considerations

Entire books have been written about Exchange Server storage planning, and virtualization adds an extra layer of complexity on top of an already complicated situation. With that being the case, my goal is not to make you an Exchange Server storage expert, but rather to give you some important things to think about as you plan to virtualize Exchange.

Hyper-V supports a wide variety of storage types. Since the physical storage is not directly exposed to a virtualized Exchange Server, Exchange is unaware of what kind of physical storage is actually being used. Even so, there are advantages and disadvantages to each storage type, and there are often corresponding limitations that impact functionality.

The first thing that you need to understand about Exchange is that it depends on block storage. As such, the use of NAS storage is generally unsupported. The exception to this rule is that SMB 3.0 storage is supported if the virtual Exchange Server exists on Windows Server 2012 or Windows Server 2012 R2 with Hyper-V. Older versions of SMB are not supported.

If you are considering the use of SMB storage, Microsoft does not support providing a virtualized Exchange Server with direct access (pass-through access) to the storage. Instead, you must create a VHD on the SMB-accessible storage and link it to the virtual Exchange Server.

Another storage limitation that you need to be aware of relates to the use of SCSI pass-through disks. Pass-through disks are fully supported for use with virtual Exchange Servers. However, backups can be problematic when pass-through disks are used. Hyper-V backups make use of the Hyper-V VSS Writer. This VSS writer is unable to access pass-through disks, so those disks are typically excluded from backups.

Another popular mechanism for storage connectivity is iSCSI. Microsoft fully supports iSCSI connectivity for virtualized Exchange Servers. If you are planning to use iSCSI, however, then it is best to run the iSCSI initiator in the host OS. Although it is supported to run the iSCSI initiator in the guest OS, doing so can lead to diminished performance in some situations.

One last storage best practice I want to pass along is that you should always consider the impact your storage architecture will have on live migrations. For example, suppose you decide to connect a virtual Exchange Server to Fibre Channel storage by way of virtual Fibre Channel connectivity.

In this type of situation, you would need to make sure that any host server that the VM could potentially be live migrated to (or failed over to) has the required host bus adapter and the appropriate virtual Fibre Channel configuration.

Considerations for disk I/O

So far the discussion of storage for virtual Exchange Servers has focused largely on supportability. However, the issue of performance is nearly as important.

Exchange Server has always been a very disk I/O-intensive application. In the last few Exchange Server releases, Microsoft has made changes to the way the mailbox database works in an effort to drive down the disk I/O requirements. Even so, a big part of ensuring that Exchange Server performs well is making sure that the storage subsystem can deliver sufficient I/O.

Microsoft's recommendation for virtualized Exchange Server environments has always been to use a dedicated set of physical disks for each mailbox database. That way, the disks do not have to share the I/O load with other processes.

While I certainly don't dispute this recommendation, the use of dedicated disks drives up storage costs and could be problematic for smaller organizations. In these situations, an alternative solution could be to use the Quality of Service Management feature found in Windows Server 2012 R2 with Hyper-V. This feature allows you to throttle disk I/O in an effort to ensure that no VM monopolizes the disk I/O for your shared storage.

The perils of state manipulation

At the beginning of this white paper, I mentioned that one of the fundamental rules for virtualized Exchange Servers is to avoid manipulating state data. This refers to any situation in which a virtualized Exchange Server's memory contents are dumped to disk for use at a later time. There are two main ways in which you must avoid state manipulation.

First, virtualized Exchange Servers must never be put into a saved state. There are a number of misconfigurations (such as failure to enable the Integration Services) that can result in a virtual Exchange Server being automatically placed into a saved state during a live migration or a backup operation. The use of saved states is not supported. If it becomes necessary to momentarily take an Exchange Server offline, you must reboot the server rather than placing it into a saved state.

Another common situation in which saved states can be an issue is with hypervisor snapshots. Hypervisor snapshots (or checkpoints as they are called in Windows Server 2012 R2 with Hyper-V) are a protective mechanism that allows a VM to be rolled back to an earlier point in time. Hypervisor snapshots are not supported for use with Exchange Server.

To give you a quick example of why snapshots are not supported, let's look at what might happen if a server containing message queues were to be snapshotted. The snapshot would contain an image of the server as it existed at a specific point in time. If such a snapshot were to be rolled back, the message queue's contents would also be reverted to an earlier point in time. Some of the possible side effects might include:

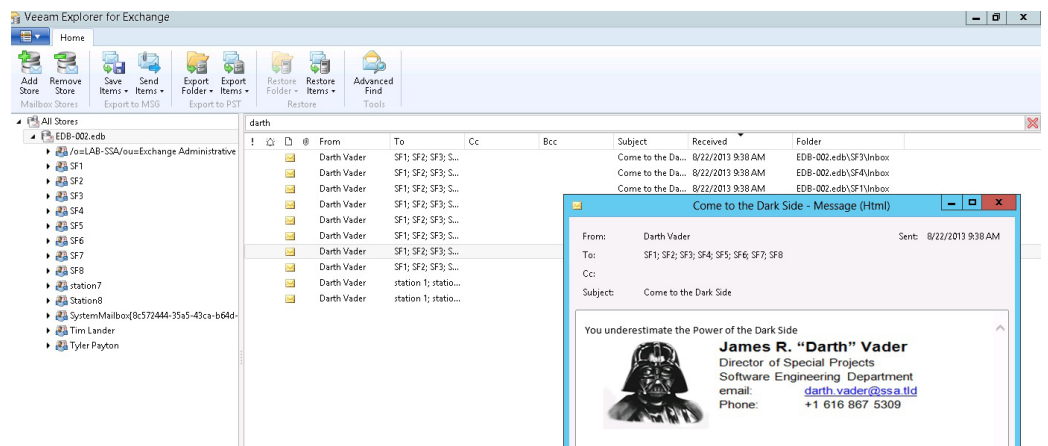
- Old messages being resent
- Current messages that were in the queue at the time of the rollback would likely be lost
- Database corruption might occur as a result of duplicate transactions

If you virtualize Exchange, then what?

Today we shouldn't have to convince people to virtualize, in general. However, Exchange is one of the most critical applications and it's understandable that people may be hesitant at first. It's worth noting that there are a number of new options that come exclusively to a virtualized Exchange environment. One noteworthy option is Veeam® Explorer™ for Microsoft Exchange. It is part of Veeam Backup & Replication™ and is a very powerful tool that provides granular recovery and e-discovery from an agentless Exchange backup as a VM.



This tool effectively opens an .EDB database right from a backup to provide granular recovery and e-discovery capabilities. Let's consider the e-discovery situation. It can be a very challenging situation if you, as an Exchange administrator, get the question: "Who sent this email to this person?" and you don't have a tool to provide that answer. The figure below shows a Veeam Explorer for Microsoft Exchange search where any messages from Darth Vader are shown:



The results, as well as a browsing view of each mailbox, can be restored, forwarded as an attachment and more. Veeam Explorer for Microsoft Exchange is also included in Veeam Backup™ Free Edition, though some of its restore capabilities are limited.

Conclusion

As you can see, there are a number of considerations that must be taken into account when virtualizing Exchange Server. As a general rule, you will get the best experience if you follow Microsoft's requirements, avoid using dynamic hardware allocations or overcommitments, and avoid manipulating state data.

About the Author



Brien Posey is a freelance technical writer who has received Microsoft's MVP award eleven times for his work with Exchange Server, Windows Server, IIS, and File Systems Storage.

Brien has written or contributed to about three dozen books, and has written well over 4,000 technical articles and white papers for a variety of printed publications and Web sites.

In addition to his writing, Brien routinely speaks at IT conferences and is involved in a wide variety of other technology related projects.

About Veeam Software

Veeam[®] is Protection for the Modern Data Center[™] - providing powerful, easy-to-use and affordable solutions that are Built for Virtualization[™] and the Cloud. **Veeam Backup & Replication[™]** delivers **VMware vSphere backup**, **Hyper-V backup**, recovery and replication. This #1 VM Backup[™] solution helps organizations meet RPOs and RTOs, save time, eliminate risks and dramatically reduce capital and operational costs. **Veeam Backup Management Suite[™]** provides all the benefits and features of Veeam Backup & Replication along with advanced monitoring, reporting and capacity planning for the backup infrastructure. **Veeam Management Pack[™]** (MP) extends enterprise monitoring to vSphere through Microsoft System Center and also offers monitoring and reporting for the Veeam Backup & Replication infrastructure. The **Veeam Cloud Provider Program** (VCP) offers flexible monthly and perpetual licensing to meet the needs of hosting, managed service and cloud service providers. VCP currently has over 4,000 service provider participants worldwide. Monthly rental is available in more than 70 countries from more than 50 Veeam aggregators.

Founded in 2006, Veeam currently has 23,000 ProPartners and more than 91,500 customers worldwide. Veeam's global headquarters are located in Baar, Switzerland and the company has offices throughout the world. To learn more, visit <http://www.veeam.com>.



Protection for the
Modern Data Center



To learn more, visit <http://www.veeam.com/backup>