vSphere Replication FAQ

January 08, 2018

Table of Contents

1. vSphere Replication (FAQ)

vSphere Replication is a proprietary, host-based replication engine for VMware virtual machines.

1.1 Introduction and General Information

What is VMware vSphere® Replication™?

vSphere Replication is a proprietary, host-based replication engine for VMware virtual machines. After initial replication of a virtual machine has been completed, changed blocks are tracked and only these deltas are sent to the target location. This approach lowers bandwidth utilization and enables more-aggressive recovery point objectives (RPOs) than manual, full-system virtual machine copies.

What are use cases for VMware vSphere® Replication™?

Replicate one or more virtual machines:

- Within a VMware vSphere cluster or across vSphere clusters at the same site for local data protection, virtual machine migration, and disaster recovery
- Across vSphere clusters at different sites for cross-site data protection, virtual machine migration, and disaster recovery
- From an on-premises data center to a VMware vCloud® Hybrid Service™ provider for disaster recovery
- From one data center to another for use with VMware vCenter™ Site Recovery Manager™

What are the minimum requirements for enabling VMware vSphere® Replication™?

Using vSphere Replication as a standalone solution without VMware vCenter™ Site Recovery Manager™ requires VMware vCenter Server™ 5.1 or higher and VMware vSphere 5.1 or higher. Use of vSphere Replication with vCenter Server 5.0 and vSphere 5.0 requires vCenter Site Recovery Manager 5.0. The edition must be VMware vSphere Essentials Plus Kit or higher. Virtual machines must be configured with virtual machine hardware version 7 or higher. The Guest OS Quiescing feature of vSphere Replication requires VMware Tools™.

What is the terminology commonly used with VMware vSphere® Replication™?

The following provides the component names and a brief, high-level definition:

- **vSphere Replication** : Name of the VMware vSphere feature that enables host-based replication
- **vSphere Replication Management Server** : Virtual appliance containing both the management component for vSphere Replication and the component that receives replicated data (vSphere Replication Server)
- **vSphere Replication Server** : Virtual appliance containing the component that receives replicated data but does not contain the management component
- **vSphere Replication Agent** : Component built into vSphere that tracks the changes in a replicated virtual machine and transmits these changes to a vSphere Replication Management Server or vSphere Replication Server virtual appliance
- **Recovery point objective (RPO)** : Policy that defines the maximum tolerable amount of data loss, measured as a period of time—typically, minutes or hours. For example, a 60-minute RPO means that except for any changes that occurred to the source in the preceding 60 minutes, the replicated copy, or “target,” should contain the same data as the source.

If I configure VMware vSphere® Replication™ with a recovery point objective (RPO) of 60 minutes, does that mean replication will occur at exactly 60-minute intervals? Does that also mean the recovery point is always going to be exactly 60 minutes old?

No. vSphere Replication utilizes an internal scheduling mechanism that takes into account factors such as time required for previous replications to complete, number of concurrent replications, and so on. The schedule is constantly adjusted to avoid violating RPO policies and to help balance the replication workload. As a result, the recovery point of a virtual machine might be less than the configured RPO. For example, a virtual machine configured with an RPO of 60 minutes might have a recovery point that is only 30 minutes old if the most recent replication occurred less than 30 minutes previously and took only a few minutes to complete. In other words, if the RPO is configured at 60 minutes, vSphere Replication will replicate data at various times to maintain a recovery point that is no more than 60 minutes old.

What are the roles of the various VMware vSphere® Replication™ system components?

vSphere Replication Agent and vSCSI Filter Driver : Components built into VMware vSphere that manage the replication process and capture writes issued by the virtual machine to the vSphere host storage subsystem

They provide the following functions:

- **Track changes to the virtual machines and send these changes to a remote vSphere Replication Management Server or vSphere Replication Server virtual appliance**
- **Schedule replications**
- **Coordinate replication of virtual machine configuration and group consistency for virtual machine disks**

vSphere Replication Management Server : Linux-based virtual appliance that contains the management framework for vSphere Replication

One vSphere Replication Management Server is deployed per VMware vCenter Server™ environment. It provides the following functions:

- Authenticate for vSphere Replication
- Map between protected and replica disks
- Allocate storage resources at the recovery site
- Monitor vCenter Server for changes that might impact replication
- Orchestrate the creation of test (with VMware vCenter™ Site Recovery Manager™) and failover images
- Provide vSphere Replication support to vCenter Site Recovery Manager
- Contain the vSphere Replication Server component enabling the vSphere Replication Management Server virtual appliance to serve as a target for replication

vSphere Replication Server: Linux-based virtual appliance that receives replication data from vSphere hosts. vSphere Replication Server sends the data to disk at the target location using the Network File Copy (NFC) protocol. A vSphere Replication Server virtual appliance is deployed in addition to a vSphere Replication Management Server virtual appliance to provide an additional target for replication. As many as nine vSphere Replication Server virtual appliances are supported, plus the vSphere Replication Management Server virtual appliance—which also contains the vSphere Replication Server component—for a maximum of 10 replication targets in a single vCenter Server environment. Deploying vSphere Replication Server virtual appliances enables workload distribution

across hosts and provides some redundancy, although failover of the replication stream from one vSphere Replication Server to another is a manual process.

What are the replication and synchronization types in VMware vSphere® Replication™?

- **Initial full sync:** vSphere Replication creates an empty virtual disk at the target location and replicates all datacontaining blocks in the source virtual disk to the virtual disk at the target location. This operation can be very time consuming, depending on the amount of data that must be replicated and the speed and latency of network connection between the source and target locations.
- **Sync** : Replication of the blocks that have changed since the last replication cycle. This is sometimes referred to as a "delta sync."
- **Full sync** : vSphere Replication compares the source virtual disk to the target copy by using checksums to determine which blocks are "out of sync." The "out of sync" blocks are then replicated from the source to the target. vSphere Replication must read the entire contents of both the source and target virtual disk files. This operation can be very time consuming. This sync method is also used when an offline copy of the source virtual disk file is placed at the target location. These offline copies are typically referred to as "seeds" and are used to reduce the amount of time and network bandwidth required to establish replication.

What databases are required for VMware vSphere® Replication™?

Only the vSphere Replication Management Server virtual appliance requires a database. It can be configured to use an internal database, which is the default configuration. External databases are also supported using Oracle Database and Microsoft SQL Server. vSphere Replication Management Server contains all configuration and monitoring information in its database. It will push some configuration information to vSphere Replication Server virtual appliances, which vSphere Replication Server stores in an embedded MySQL database.

How does VMware vSphere® Replication™ interact with other VMware products?

The VMware compatibility guides are the best sources for verifying compatibility with and support of various combinations of VMware features and products. The following notes also pertain:

- VMware vSphere vMotion® and VMware vSphere Distributed Resource Scheduler™ (vSphere DRS) are fully supported. vSphere Replication will continue without a full sync. VMware vCenter Server™ will not allow migration of a protected virtual machine to a host running a version of VMware vSphere prior to 5.0 because these vSphere versions do not contain the vSphere Replication Agent. It is not possible to manually install the vSphere Replication Agent on a vSphere host prior to version 5.0.
- VMware vSphere® Storage vMotion® and VMware vSphere Storage DRS™ are not supported with replicated virtual machines on vSphere 5.0 and 5.1. These features are supported with vSphere 5.5 and higher at the source location but not at the target. vSphere Storage vMotion is not supported for test virtual machines created by VMware vCenter™ Site Recovery Manager™ and vSphere Replication at the target location. It is not possible to use vSphere Storage vMotion to migrate replication target virtual disks.
- Virtual machines protected by VMware vSphere Fault Tolerance (vSphere FT) cannot be replicated with vSphere Replication.
- Virtual machines protected by VMware vSphere High Availability (vSphere HA) can be replicated. However, when a replicated virtual machine is recovered by vSphere HA, vSphere Replication might require a full sync.

- vSphere Replication can replicate virtual machines with snapshots. Snapshots at the source are not reproduced at the target location. vSphere Replication will recover a virtual machine at the target location, using the latest replicated data regardless of whether this data was originally stored in a regular virtual disk file or a virtual machine snapshot file at the source location.
- Reverting to a virtual machine snapshot that was created after vSphere Replication was configured, at the source location, typically causes vSphere Replication to perform a full sync for a virtual machine. Reverting to a virtual machine snapshot that was created before vSphere Replication was configured, at the source location, pauses replication and requires replication to be manually reconfigured or stopped. This is because replication information is stored in the configuration file (VMX) of a virtual machine. If a virtual machine is reverted to a snapshot taken before replication was enabled, the replication information for the virtual machine will be lost. If reverting to a snapshot changes the configuration of a virtual machine—for example, a virtual disk is deleted—replication will be paused and must be reconfigured.
- A linked-clone virtual machine that is replicated will be recovered as a standard "full" virtual machine by vSphere Replication. Replicating a virtual machine template is not supported. A virtual machine must be powered on for replication to occur. To move a template from its source location to a target location with vSphere Replication, use this work-around: Convert the virtual machine template to a standard virtual machine, power it on, configure replication for the virtual machine, allow the initial full sync to complete, power off the source virtual machine, recover the target virtual machine, and convert it to a template at the target location.
- VMware vSphere vApps™ are not fully supported. Although it is possible to replicate the virtual machines that compose a vApp, vSphere Replication cannot replicate the constructs and policies of the vApp itself. Use this work-around: Replicate the virtual machines contained in the vApp and recover them at the target location. Create a new vApp at the target location with the same configuration and policies as the original vApp. Import the recovered virtual machines into the newly created vApp at the target location.
- VMware vCloud® Networking and Security™ configurations will not be recovered. Therefore, it is recommended that administrators configure vCloud Networking and Security solutions for each location and avoid replicating vCloud Networking and Security virtual appliances. However, standard virtual machines utilizing resources provided by vCloud Networking and Security solutions can be replicated.

Can I use vSphere Replication to replicate a VM between vCenters running different versions of vSphere Replication (eg. VR 5.5 to 6.0)?

To replicate between vCenters, vSphere Replication requires the same major version (eg. 5.5.x > 5.5.x) of vCenter and the same version of vSphere Replication (eg. 6.5.1 > 6.5.1) at both source and target. This is the only officially tested and supported configuration. Check the VMware Product

Interoperability Matrix if you have any questions about matching vCenter and vSphere Replication versions.

If a VMware vSphere® host that is performing replication goes offline and VMware vSphere High Availability restarts the protected virtual machine on another vSphere host, will replication be resumed?

Replication will resume, but a full sync will likely be required for virtual machines affected by the vSphere host outage.

If Microsoft Volume Shadow Copy Service (VSS) quiescing is enabled when configuring a virtual machine for replication, will this cause issues with backup solutions that also use VSS?

In most cases when using VMware vSphere® and VMware vCenter Server™ 5.5 or higher, there should not be any issues with protecting virtual machines using both solutions. Occasional issues have been observed—for example, creation of a snapshot fails for a backup job or replication job—but these issues are easily remediated. With VMware vSphere Replication™, quiescing of the virtual machine might fail, but replication continues and notifications will be observed in the VMware vSphere Web Client. Depending on the backup solution in use, an administrator might need to rerun the backup job for virtual machines that failed to back up properly. Additional information specific to vSphere Replication and VMware vSphere Data Protection™ Advanced can be found in this blog article: http://blogs.vmware.com/vsphere/2013/12/improve-availability-with-vdpa-and-vr.html

What replication engine is VMware vSphere® Replication™ using? Is it Changed Block Tracking (CBT)? Are virtual machine snapshots used?

VMware developed the vSphere Replication engine. The method vSphere Replication uses to track changes to a virtual machine is very similar to the CBT mechanism that is part of VMware vSphere Storage APIs – Data Protection. However, it is not CBT, preventing interference with other solutions that utilize CBT—for example, VMware vSphere Data Protection. Virtual machine snapshots are not used at the source unless Microsoft Volume Shadow Copy Service (VSS) quiescing is enabled when configuring replication. In that case, the virtual machine is quiesced, a snapshot is taken to capture the application-consistent state of the virtual machine, the delta (changes) is created for replication, the snapshot is committed, and the delta package of data is replicated. This occurs with every replication cycle.

Can VMware vSphere® Replication™ tasks be scheduled? How does vSphere Replication determine replication schedules?

No. The replication schedule for a virtual machine is determined by the recovery point objective (RPO) set when configuring replication for that virtual machine. The possible value for RPO can be any number of minutes from 15 to 1,440 (24 hours). There is currently no way to schedule replication at specific times. vSphere Replication generates its own replication schedule internally by factoring all replicated virtual machines on each VMware vSphere host.

A 48-hour replication schedule is computed using historic data change rates. The vSphere Replication scheduler calculates and updates this schedule after each delta sync and each time certain events occur, such as a change in virtual machine power state, replication reconfiguration, and so on. vSphere Replication attempts to spread out replication cycles to minimize the number of concurrent instances on a vSphere host. Because replication takes time to complete—especially with large amounts of data and slower network connections—each replica is considered “aged” by the time the current replication cycle completes. To avoid RPO violation, vSphere Replication attempts to complete a replication cycle in less than half of the configured RPO. Estimated transfer time is calculated by averaging the previous 15 delta replications and adding 20 percent to that average transfer time as a buffer.

The examples shown in Figure 1 help illustrate this. In both examples, the RPO is configured at 60 minutes.

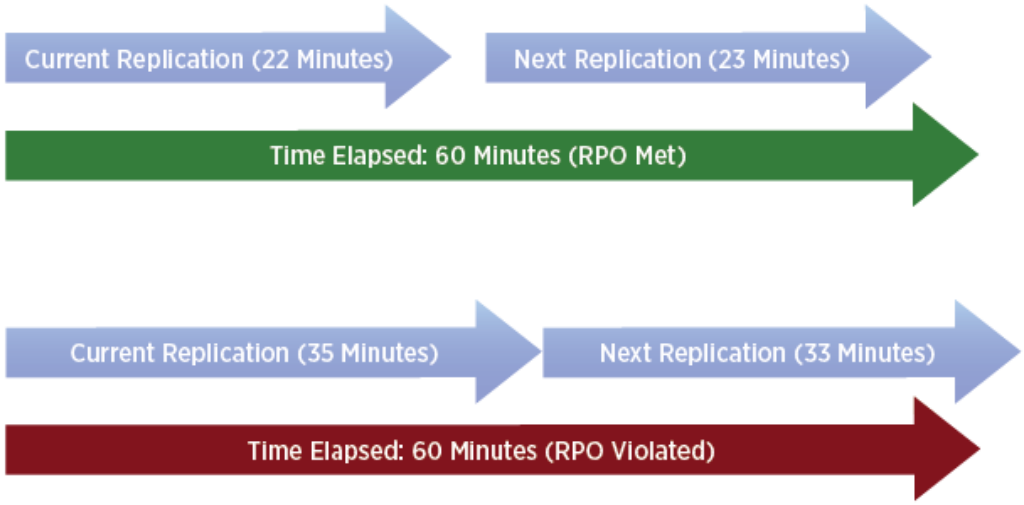

Figure 1. vSphere Replication Schedule Generation

In the top example, the current replication takes 22 minutes to complete. After completion, the data at the target will look the same as the data at the source when replication started 22 minutes earlier. As the next replication begins, the restore point of the target continues to age. The next replication takes 23 minutes to complete. When that cycle is finished, the target has newer data and looks the same as the source did 23 minutes before. Data replication starts again and the cycle continues.

In the bottom example, the current replication takes 35 minutes. After it has completed and the replicated data has been committed to the target, the target looks the same as the source did 35 minutes before. The next replication starts and the restore point of the target continues to age. Again, target data is not updated until the next replication has completed. In this case, 33 minutes later—8 minutes past the RPO of 60 minutes. Assuming that replication continues to take 33–35 minutes to complete, vSphere Replication will get further behind its RPO policy of 60 minutes with each replication cycle. Alerts will be shown in the user interface, but replication will continue. To resolve this RPO violation, the RPO policy should be increased. Adding more available network bandwidth can sometimes—but not always—help resolve the RPO violation, but there are other factors that can affect performance.

Where is VMware vSphere® Replication™ information for the source virtual machine stored?

A virtual machine's configuration is stored in a file with the .vmx extension. The following are examples:

```
hbr_filter.rpo = “240”
hbr_filter.destination = “192.168.2.122”
hbr_filter.port = “44046”
hbr_filter.protocol = “lwd”
hbr_filter.quiesce = “TRUE”
```

For each replicated disk, changed blocks are tracked in memory. In some cases, these changes are then written to a persistent state file (PSF) in the virtual machine home directory—for example, when a replicated virtual machine is powered off. In addition, a PSF contains a demand log for the virtual disk. If a block is changed before vSphere Replication can replicate a previous change for that same block, this data is queued to the demand log for transmission during the next replication cycle.

How does vSphere Replication keep track of changes to a virtual machine? Are the changes buffered?

At the source location, changed blocks are tracked in memory. The actual contents of the blocks are not normally tracked, although there is an exception: If a block is changed (again) while replication is in progress, the contents of the block are copied to the persistent state file (PSF) to ensure the

consistency of the virtual disk at the target location. This can cause the PSF to grow considerably for virtual machines that have very high data change rates. Copies of the blocks in the PSF are kept only until the replication cycle completes.

Does VMware vSphere® Replication™ maintain a virtual machine's snapshot hierarchy?

No. vSphere Replication does not replicate a virtual machine snapshot hierarchy from the source location to the target location. Snapshots are collapsed into a single aggregate virtual machine disk (VMDK) file at the target location. In other words, a virtual machine with snapshots might be configured for replication, but it will be recovered with no snapshots—snapshot data committed—when it is recovered at the target location.

However, vSphere Replication 5.5 introduced Multiple Point In Time (MPIT) recovery. MPIT recovery enables an administrator to recover a virtual machine to the latest replicated copy at the target site and then revert, or "roll back," that virtual machine to a previous point in time. When MPIT recovery is configured, these recovery points appear as virtual machine snapshots at the target location when a virtual machine is recovered using vSphere Replication. There are no dependencies between snapshots at the source location and recovery points at the target location. A virtual machine at the source location with no snapshots can still be configured to utilize MPIT recovery with vSphere Replication.

What is the hbrdisk.RDID file?

This is a redo log found at the target location. This data is applied to the last consistent instance of the base disk after replication has completed successfully. The size of these files depends on the number of changes at the source since the last successful replication cycle. It is possible for a redo log to be the same size as the source if all of the blocks have been changed, but this is unusual. After data has been committed successfully from the redo log to the base disk, a new redo log is started. There are occasions when more than one redo log for a virtual disk is observed—for example, if multiple recovery points are enabled for the virtual machine.

Can powered-off virtual machines be replicated using VMware vSphere® Replication™?

vSphere Replication can be configured for a virtual machine that is powered off. However, data is replicated only when the virtual machine is powered on.

Because replication occurs only when a virtual machine is powered on, is there data lost to the replication process as the virtual machine is shut down?

Any writes that are done during shutdown are not replicated until after the virtual machine starts up again and a replication cycle begins. With VMware vSphere® Replication™ 5.1 and higher, the administrator is given the option to perform a final replication before the target virtual machine is recovered. This enables recovery with no data loss.

Another scenario is during a VMware vCenter™ Site Recovery Manager™ planned migration where vCenter Site Recovery Manager instructs vSphere Replication to replicate outstanding changes from the source location. During a vCenter Site Recovery Manager planned migration, virtual machines that are running are shut down, a final replication is initiated, and the virtual machines are powered on at the target location. This operation ensures that there is no data loss as a result of the vCenter Site Recovery Manager planned migration.

What is the difference between pausing replication and stopping replication?

If an administrator pauses replication, the replication configuration remains in place but data will not be replicated until replication is resumed. If replication is stopped, the replication configuration is removed and replicated data at the target location will be removed. If an administrator wants to keep a copy of the replicated virtual disks at the target location, these steps can be followed:

1. Pause replication.

2. Copy the files that are to be kept at the target location to a different folder or rename the existing folder.
3. Stop replication, to remove the replication configuration.

Is it possible to change the recovery point objective (RPO) of a virtual machine that is currently being replicated?

Yes. VMware vSphere® Replication™ will adjust its schedule to accommodate the new RPO as quickly as possible.

Is there a way to test failover for a virtual machine protected by VMware vSphere® Replication™ without VMware vCenter™ Site Recovery Manager™ or VMware vCloud® Hybrid Service™– Disaster Recovery?

Automated testing that does not interrupt replication requires the use of vCenter Site Recovery Manager or vCloud Hybrid Service – Disaster Recovery. Coordinating a test recovery of multiple virtual machines also requires vCenter Site Recovery Manager or vCloud Hybrid Service – Disaster Recovery. However, there is a basic work-around, which can be performed only one virtual machine at a time: Replication is stopped and must be manually reconfigured. There are two recovery modes in vSphere Replication without vCenter Site Recovery Manager: The first is "recover with recent changes," which requires the source virtual machine to be powered off. After the source virtual machine is in a powered-off state, a final sync is performed to replicate any outstanding changes. The virtual machine is then recovered at the target location. The second option is "recover to the last replica," which can be done irrespective of the power state of the primary virtual machine. The virtual machine can be recovered with its virtual network interface card disconnected to avoid interference with the source virtual machine.

Is there an API or SDK for VMware vSphere® Replication™?

No. There are no supported means of programmatically interacting with vSphere Replication. Articles referencing the vim-cmd hbrsvc command can be found in online searches, but it is not currently supported.

How are upgrades handled in an environment where VMware vSphere® Replication™ is deployed?

vSphere Replication upgrades can be orchestrated with VMware vSphere Update Manager™ or manually through the virtual appliance management interface (VAMI). Using vSphere Update Manager to orchestrate vSphere Replication virtual appliance upgrades is recommended. VMware vCenter Server™ and its supporting components—VMware vSphere Web Client, vCenter Inventory Service, and so on—must be upgraded first to be compatible with new releases of vSphere Replication. Updates to the vSphere Replication components built into VMware vSphere will likely be delivered through a vSphere patch or upgrade. Always check the VMware compatibility guides and matrixes for feature and product interoperability: http://www.vmware.com/guides.html

What happens to a virtual machine's replication state if the host on which the virtual machine was running crashes during replication?

When a replicated virtual machine is powered back on, a full sync will be initiated. After the full sync has completed, regular delta syncs will continue.

What happens to a virtual machine's replication state if the VMware vSphere® Replication™ virtual appliance receiving replication for that virtual machine crashes during replication?

Changes to the virtual machine are tracked in memory and in the persistent state file (PSF), if necessary. Replication will continue after the vSphere Replication virtual appliance restarts. Replication can also be reconfigured to use a different vSphere Replication virtual appliance.

What happens if the host receiving NFC traffic from a VMware vSphere® Replication™ virtual appliance goes offline?

The vSphere Replication Server component will attempt to locate another host that has access to the datastore containing the replica. If an alternate cannot be located, replication is paused until a host with access to the same datastore is located.

What happens if the VMware vSphere® Replication™ Management Server virtual appliance goes offline?

Replication for virtual machines using the vSphere Replication Management Server virtual appliance as the target for replication will stop, but changes to the source will be tracked until the vSphere Replication Management Server virtual appliance is back online. Replication for virtual machines using other vSphere Replication Server virtual appliances will continue as configured, but changes to the replication configuration cannot be made until the vSphere Replication Management Server virtual appliance is back online. Another limitation is the inability to make changes to the vSphere Replication environment such as the addition of a vSphere Replication Server virtual appliance.

1.2 Networking

What are the TCP/IP ports used by VMware vSphere® Replication™?

For details on the network ports used by vSphere Replication with and without VMware vCenter™ Site Recovery Manager™, see VMware Knowledge Base article 1009562: http://kb.vmware.com/kb/1009562

Replicated data is sent directly from each VMware vSphere host that is running one or more virtual machines protected by vSphere Replication to a vSphere Replication Management Server virtual appliance or vSphere Replication Server virtual appliance at the target location. The vSphere Replication Server component passes the replicated data to the Network File Copy (NFC) service of a vSphere host that has access to the target datastore. The vSphere host then writes the data to disk. The management component of the vSphere Replication Management Server does not participate directly in replication but is a management layer that tracks replication configuration and monitors the replication in the environment.

Can a particular VMware vSphere® management interface be configured for replication? Is it possible to configure an additional network adapter in a VMware vSphere Replication™ virtual appliance?

At the source location, replication traffic comes from a management (VMkernel) port. It is possible to configure a static route for replication traffic, but this is currently a manual process performed at the vSphere command line.

Configuring more than one virtual network interface card on a vSphere Replication virtual appliance is supported as of version 6.0 or later. This allows for the segregation of management, replication and NFC traffic. More details on this can be found here.

How does VMware vSphere® Replication™ manage bandwidth consumption for initial and ongoing replication?

vSphere Replication uses separate TCP ports for full-sync and delta replication streams. Bandwidth is not monitored by vSphere Replication, and there is no way to configure bandwidth throttling in vSphere Replication. The data replicated by vSphere Replication is not compressed or encrypted, enabling third-party WAN acceleration solutions to provide some benefit.

How do I estimate the replication traffic bandwidth requirements?

There is a vSphere Replication Calculator available here. This allows for inputting multiple variables like number of VMs, size, change rate, RPO, bandwidth, etc and solving for the unknown.

When thinking about replication bandwidth requirements, there are many factors that influence bandwidth requirements for VMware vSphere® Replication™—for example, the amount of changed data, the frequency of data changes, and the recovery point objective (RPO) configuration for each replicated virtual machine. For instance, if a block changes only once per day, it is replicated only once regardless of RPO configuration. Conversely, if a block changes many times throughout the day, and the RPO is set to a low number such as 30 minutes, the block might be replicated as many as 48 times in one day. As discussed in VMware vSphere Replication Administration, examine how data change rate, traffic rates, and link speed impact the RPO:

1. Identify the average data change rate within the RPO by calculating the average change rate over a longer period and then dividing it by the RPO.
2. Calculate how much traffic this data change rate generates in each RPO period.
3. Measure the traffic against the link speed.

For example, a data change of 100GB requires approximately 200 hours to replicate on a T1 network, 30 hours on a 10Mbps network, and 3 hours on a 100Mbps network. When multiple virtual machines are replicated from a single host, vSphere Replication attempts to balance replication traffic by scheduling transfers to avoid parallel updates where possible.

What happens if there is a network disconnect between the source location and the target location? When connectivity has been restored, will replication continue from where it left off?

Yes. Changes to the source virtual disks will be tracked and later replicated when network connectivity has been restored.

1.3 Storage

Is there a minimum VMware vSphere® VMFS version required for VMware vSphere Replication™?

No. VMFS version support is determined by vSphere. If the vSphere host can access the datastore and the virtual machines residing in it, the virtual machines can be replicated with vSphere Replication.

If using VMware vSphere® VMFS to host VMware vSphere Replication™ protected virtual machines, do the source and destination need the same block size?

No. vSphere Replication captures storage writes in the VMware vSphere host before they are passed to the storage subsystem. This enables vSphere Replication to work with a variety of storage platforms, file system types, and block sizes.

Are there any special virtual machine disk (VMDK) requirements to support VMware vSphere® Replication™?

Thick and thin virtual disk provisioning formats are supported. Virtual raw device mapping (RDM) devices are supported. Physical RDM devices are not supported. VMDKs that are configured with an independent persistent or independent nonpersistent disk mode can be replicated with vSphere

Replication. However, when the virtual machine is recovered at the target location, the recovered VMDK(s) will be in dependent disk mode.

Is it possible to replicate a virtual disk connected to a virtual SCSI controller that is configured for SCSI bus sharing?

Replicating a virtual disk attached to a SCSI controller configured with virtual or physical bus sharing is not supported. SCSI bus sharing must be set to "None" to be compatible with VMware vSphere® Replication™.

How do I prepare the target location for using VMware vSphere® Replication™ "seeded" data?

Copying a virtual disk from the source location to the target location using a method other than vSphere Replication is typically referred to as "seeding." This can be beneficial in cases where there are large amounts ofdata or limited bandwidth available when performing the initial full syncs. vSphere Replication can utilize "seed" copies of virtual disks at the target location to limit the amount of data that must initially be replicated by vSphere Replication. Seeding can be done in a variety of ways. One of the most common methods is copying the source virtual machines to a second storage platform at the source location and shipping that storage to the target location. VMware vSphere hosts at the target location can then mount the storage containing the copies of the virtual machines for use with vSphere Replication. When configuring replication at the source location, the administrator can specify the storage containing the copied virtual machines at the target. vSphere Replication identifies virtual disk files that can be used as "seeds" and queries whether the files should be used for replication.

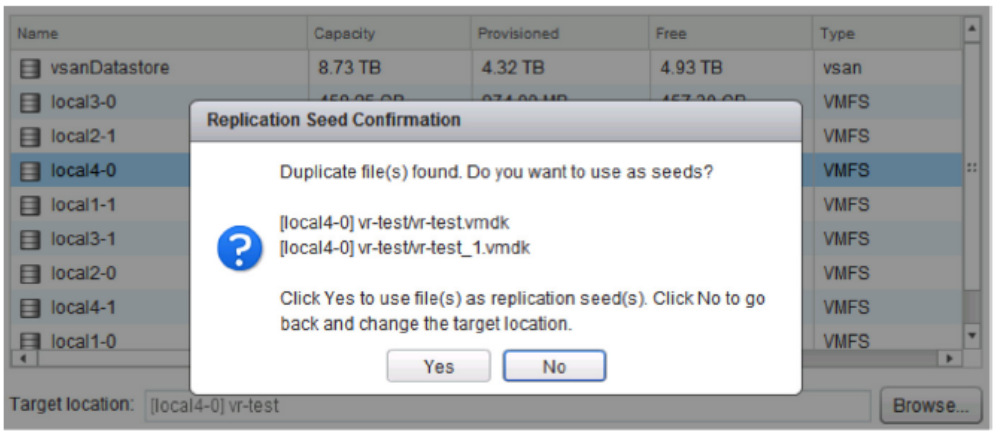

Figure 2. Replication Seed Screenshot

How do I create virtual disk "seeds" for replication targets?

It is important to use vmkfstools to create the "seed" copies of virtual disks, especially for object stores such as VMware vSAN™. See VMware Knowledge Base article 1028042 for more information: http://kb.vmware.com/kb/1028042

How much disk space should be allocated for replication at the target location?

A safe and conservative estimate is two to three times the size of each virtual machine disk (VMDK) being replicated. For example, if a source virtual disk is 40GB, the following amount of capacity at the target might be allocated:

1. 40GB for the "base," or consistent, replica
2. 40GB for the redo log—that is, the current replication cycle

3. 40GB for an extra replica such as a VMware vCenter™ Site Recovery Manager™ recovery plan test

It is possible to allocate less capacity. In a few scenarios, more than three times the storage might be needed, but this is uncommon. In all cases, VMware vCenter Server™ alarms should be set to monitor free space on the datastores where the target replicas are placed.

What is the worst-case target location storage usage by VMware vSphere® Replication™ for a virtual machine?

As much as four times the capacity of the source virtual disk might be needed:

- Capacity for the target "base" virtual disk
- Space for the most recently completed replication cycle (yet to be committed to the "base" disk)
- Additional space for the replication cycle in progress
- When VMware vCenter™ Site Recovery Manager™ is in use, another copy for an active vCenter Site Recovery Manager recovery plan test

If MPIT recovery is enabled, additional capacity will be required for these recovery points. This additional capacity can be significant if the source virtual machine is very active—that is, with many blocks changing frequently—or if many recovery points are retained. vSphere Replication supports retention of as many as 24 recovery points.

What is the worst-case source location storage usage by VMware vSphere® Replication™ for a virtual machine?

The persistent state file (PSF) for a virtual disk can be slightly larger than the source virtual machine disk (VMDK) file. PSFs are stored in the virtual machine's home directory regardless of where the VMDKs are stored. An overly large PSF is very uncommon but can occur when an extremely large number of blocks in the source VMDK are updated rapidly during replication.

What happens if the storage at the target location goes offline for a prolonged period—a few hours, for example?

Virtual machines at the source location will not be adversely affected. Changes that occur in the source virtual machine will be tracked. If the failure occurs during a replication cycle, the changes are moved to the demand log, which is part of the persistent state file (PSF). When the target datastore is available, changes will be replicated.

If a new virtual disk is added to a virtual machine that is currently being replicated, what is the process to ensure that the new disk is replicated?

Replication will pause when the new disk is added, a VMware vCenter Server™ alarm will be triggered, and replication must be reconfigured to include the new virtual disk.

If a virtual disk that is being replicated is removed from a virtual machine, what is the process to ensure that this disk is no longer replicated and is removed from the target location?

When a replicated virtual disk is removed from the source, the virtual disk is no longer replicated. The target virtual disk files can be observed for some time after the virtual disk is removed at the source. This is especially true if MPIT recovery is enabled. After the replicated copy of the virtual disk is no longer needed, VMware vSphere® Replication™ deletes the unnecessary files at the target location. The exception to this is if the original virtual disk at the target site was a "seeded" copy. In that case, an administrator must manually remove the virtual disk files at the target location when they are no longer needed.

Is it possible to change the size of a virtual machine disk (VMDK) while it is being replicated with VMware vSphere® Replication™?

If an administrator attempts to resize a VMDK that is protected by vSphere Replication, the resize operation will fail regardless of the virtual machine power state. Use the following process to resize a VMDK protected by vSphere Replication:

1. Perform a vSphere Replication recovery of the virtual machine. Select the "Use latest available data" option if the source virtual machine is to remain powered on. Uncheck "Power on the virtual machine after recovery" in the last step of the recovery wizard.
2. Stop replication of the source virtual machine.
3. Resize the disk(s) of the source virtual machine, as needed.
4. Resize the disk(s) of the target virtual machine in the same manner as the source.
5. Note the location of the target virtual machine's files—that is, datastore and folder.
6. Remove the recovered virtual machine from VMware vCenter Server™ inventory at the target location.
7. Delete all files except VMDK files.
8. Configure replication for the source virtual machine. When selecting the target location, specify the datastore and folder of the recovered virtual machine at the target location.
9. Click "Yes" in the "Replication Seed Confirmation" window.
10. Complete the remaining steps in the configure replication wizard. vSphere Replication will perform a full sync using the replication seed target.

If VMware vSphere® Replication™ is committing changed data to the target virtual disk (s) and there is an unexpected outage, will vSphere Replication retry the update after the outage has been resolved?

If the VMware vSphere host receiving the NFC traffic goes offline, the vSphere Replication virtual appliance will attempt to reopen the virtual disk through another vSphere host and commit the changed data. A full sync is typically not required in this scenario. If the vSphere Replication virtual appliance receiving replication traffic goes offline, operations "in flight" from the source location will be retried until the vSphere Replication virtual appliance is back online. Similar activity takes place when a WAN outage occurs or storage at the target location becomes unavailable.

Is it possible to replicate the same virtual machine with a storage array replication solution and VMware vSphere® Replication™?

Yes. It is technically possible, but it can lead to unwanted results. For example, if a virtual machine is recovered by vSphere Replication, the original (source) virtual machine is still being replicated by array replication, resulting in two different versions of the same virtual machine at the target location. This also requires twice the amount of bandwidth because the data is being replicated two times.

If replication is disabled for a virtual disk, is it possible to reenable it?

Replication for a virtual disk can be reenabled.

Are there risks or adverse effects when disabling replication for a virtual disk that is part of a virtual machine configured with multiple virtual disks?

It might take a considerable amount of time for the remnants of the disk for which replication was disabled to be removed from the target location, especially if MPIT recovery has been enabled. This is

because VMware vSphere® Replication™ retains data at the target location until it is no longer needed for recovery.

1.4 Performance and Monitoring

What performance information is available about VMware vSphere® Replication™?

Impact on virtual CPU performance for a virtual machine that is protected by vSphere Replication is approximately 2 to 6 percent. In nearly all cases, this is not an issue because the vast majority of virtual machines
are not CPU constrained.

Are there logs that show VMware vSphere® Replication™ virtual appliance information?

vSphere Replication virtual appliance logs are stored in /opt/vmware/hms/logs, and they can be gathered by logging in to the vSphere Replication virtual appliance management interface (VAMI) or VMware vCenter™
Site Recovery Manager™ UI. vSphere Replication Server logs are stored in /var/log/vmware and they are collected using the vSphere Replication VAMI or vCenter Site Recovery Manager UI.

What performance overhead exists on the source hosts?

The replication scheduler built into VMware vSphere® uses an insignificant amount of CPU and memory to compute the replication schedule of all the virtual machines on the host.

What performance overhead exists on the target hosts?

VMware vSphere® Replication™ virtual appliances receive the replication traffic. They utilize compute and networking resources in a manner similar to any other virtual machine. Additional storage load is placed upon
the VMware vSphere hosts as vSphere Replication transfers replicated data to storage using the Network File Copy (NFC) protocol. During a full-sync operation, checksum calculation operations are distributed across
multiple hosts to minimize CPU impact on any one host. Several factors influence where contention might occur. For example, if there is much bandwidth—1Gbps, for example—the amount of NFC traffic might tax the host or
even the underlying storage system.

When a recovery point objective (RPO) violation occurs, will a VMware vCenter Server™ alarm be triggered?

vCenter Server alarms can be created to notify administrators when an RPO violation occurs and when it is resolved.

How far behind can replication get before a full sync is initiated?

A full sync is triggered if there is a disk content ID mismatch between the source and target or if there is an error condition—for example, a corrupted demand log file at the primary—which should be rare. The system currently
doesn't trigger full sync if replication falls behind the configured recovery point objective (RPO). A content ID mismatch usually occurs if there are out-of-band changes made to virtual disks—for example, if a source host or
VMware vSphere® Replication™ virtual appliance at the target location crashes and comes online again or when replication is paused and later resumed.

Are there VMware vCenter Server™ alarms that provide an indication of when an additional VMware vSphere® Replication™ Server virtual appliance should be deployed?

Typically, administrators monitor vSphere Replication Server network traffic and CPU utilization. If these alarms are triggered frequently, deploying another vSphere Replication Server virtual appliance and moving some of the replication streams from the original appliance to the new appliance can help distribute the load.

If Guest OS Quiescing is enabled, what happens when errors are received within the guest OS for Microsoft Volume Shadow Copy Service (VSS)? Are there any alarms or alerts tied to the Guest OS Quiescing feature?

Replication will continue if VSS quiescing fails. The VSS provider in VMware Tools™ will attempt to quiesce down to the application level. If that fails, Guest OS Quiescing is still attempted. If it fails, the files at the target site will
be a crash-consistent copy of the virtual machine. Any failures to properly quiesce a virtual machine will generate a warning event. A VMware vCenter Server™ alarm can also be created to raise an alert when there is a failure.

What VMware vCenter Server™ alarms are available for use with VMware vSphere® Replication™?

The complete list of vSphere Replication alarms can be found in vSphere Replication documentation. Popular alarms that can be configured include “RPO violated,” “Remote vSphere Replication site is disconnected,”
“VR Server disconnected,” “vSphere Replication paused,” and “No connection to VR Server.”

1.5 About the Authors

Jeff Hunter is a Senior Technical Marketing Architect at VMware with a focus on vSAN, business continuity and disaster recovery solutions. He has been with VMware for more than six years, prior to which he spent several years implementing and administering VMware virtual infrastructures at two Fortune 500 companies. Follow Jeff on Twitter: @jhuntervmware

GS Khalsa is a Senior Technical Marketing Manager at VMware focusing on business continuity, disaster recovery and data protection solutions. GS has been with VMware since 2013. Prior to VMware, GS spent time as both a customer and VMware partner. Follow GS on Twitter: @gurusimran