Enabling Replica on Hyper-v Host servers

There are at least four components involved when configuring Hyper-V Replica.

- **Primary Server:** A server running Microsoft Hyper-V version 3.0 (Windows Server 2012) with "Hyper-V Replica" enabled. This Server will participate in the replication with "Replica Server" configured at the other end. The virtual machines running on the "Primary Server" are configured to participate in the replication process. There is no need to configure any other component for Hyper-V Replica on the Primary Server. The Primary Server always runs in the production site or Primary Site. Primary Server monitors the changes on the Virtual Machines' VHD files with the help of modules implemented in the "Replication Engine".
- **Replica Server:** A server running on Windows Server 2012 with Hyper-V role enabled in a disaster recovery site which accepts the replication packets from the "Primary Server". Replica Servers always run in disaster recovery site or "Replica Site" in Hyper-V terminology.
- **Primary Virtual Machine:** A virtual machine configured on the 'Primary Server" which participates in the replication. By default, no virtual machines are configured to replicate. You must enable replication for at least one virtual machine running on the "Primary Server" before "Hyper-V Replica" feature is in action. The Primary Virtual Machine is available to users and will be running all the time unless something goes wrong with it.
- **Replica Virtual Machine:** A virtual machine which is a copy of the "Primary Virtual Machine" and is hosted on the "Replica Server". "Replica Virtual Machine" has the backup copies attached to it which are used during the failover process discussed in this article. The Replica Virtual Machine will always remain OFF and be ready to come online if something issues happen to the Primary Virtual Machine.

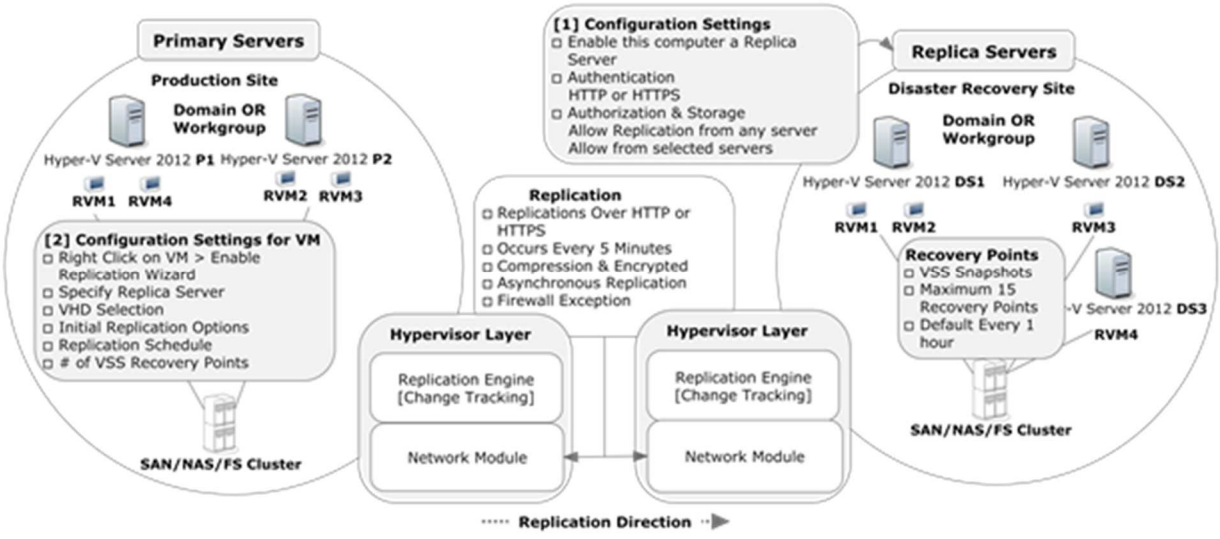

Figure – Hyper-V Replica and its components in Workgroup or Domain Security Model

As shown in **Figure** the Primary Servers (P1 and P2) are running on the production site and the Primary Virtual Machines (RVM1, RVM2, RVM3, and RVM4) are hosted on P1 and P2.

Replica Servers (DS1 and DS2) are running in the disaster recovery site and the "Replica Virtual Machines" (RVM1, RVM2, RVM3, and RVM4) are hosted on the DS1 and DS2 Replica Servers.

Tip: Windows Server Hyper-V 3.0, a standalone and free version released with Hyper-V Role enabled, can also be used to participate in the Hyper-V Replication.

There can be more than one "Replica Server" running in the disaster recovery site accepting replication packets from multiple Primary Servers running in the production site as shown in the **figure 1.0** above. Depending on your customer's requirement, you can also host Replica Servers in production site and Primary Servers in disaster recovery site.

Note: One virtual machine can be configured with one "Replica Server". The same virtual machine cannot be used by any other "Replica Servers".

As shown in **figure** , the hypervisor contains a replication engine; this is the same on the primary server and the replica server. On the primary server the replication engine tracks all the write operations of the Virtual Machines on this server. These changes are replicated to the replica server. The replica server in turn has (of course) the same modules but does not track any changes on the Virtual Machines, unless a failover has occurred and the replica server has become the primary server, but this is not normal.

As shown in the **figure** above, the Virtual Machines running on the Primary Server are configured in a file server cluster. The Hyper-V Replica is designed in such a way that it works irrespective of where the virtual machine VHD files resides. The Virtual Machine VHD files can be hosted on a NAS LUN, DAS or an SMB share on a file server or on a CSV (Clustered Shared Volume). There is no need to implement a cluster just for storing the Virtual Machine VHD files but it is recommended to store the Virtual Machines VHD files on a failover cluster to avoid any single point of failure for storage devices.

The Primary Server and Replica Server can authenticate using either HTTP (Kerberos) or HTTPs (Certificate) as indicated in the **figure** above. A firewall exception is mandatory for Hyper-V Replica. Please note that the firewall exception rules for port 80 for HTTP and 443 for HTTPS are already created in the Windows Firewall. The firewall rules just need to be enabled.

A firewall exception is mandatory for Hyper-V Replica Primary and Hyper-V Replica Servers.

By default, replication from Hyper-V Primary Server to Hyper-V Replica Server occurs every 5 minutes and this interval cannot be changed. The packets that travel from Primary to Replica Server are compressed if you have enabled compression for the virtual machine. The encryption is available, or packets get encrypted, only if you have selected "**Certificate-Based HTTPS**" as the authentication mechanism between Primary and Replica Servers.

Tip: You can enable/disable compression per Virtual Machine when configuring a Virtual Machine for Hyper-V Replication.

Hyper-V Replica provides two types of replica copies; "Standard Replica copy" and "Application-Consistent replica copy". These replica copies are called "recovery points" which are explained in the later part of this article.

Hyper-V Replica uses Volume Shadow Copy Service (VSS) to take point-in-time snapshots (Application-Consistent recovery point) of the applications running inside the Virtual Machine and/or Virtual Machine (Standard Replica Recovery Point). A maximum of 15 recovery points can be created, and the default interval for creating these copies is every hour.

If the limit is reached (16th copy), Primary Server merges the oldest snapshot to the base replica VHD.

Replica Backup copies are explained more in detail in a later part of the article.

Primary Servers and Replica Servers can operate in different domains in the same Active Directory Forest as long as they are able to communicate and authenticate.

Tip: Active Directory implements two-way transitive trust relationships between the domains of the same Active Director Forest. So by default every other object in the domain is trusted by the other domains. Hyper-V Replica utilizes the trust-relationship to authenticate Primary and Replica Servers in different domains.

Installing and Configuring Hyper-V Replica

You must have sufficient storage on both the Primary and Replica servers to host the files used by virtualized workloads. There must be network connectivity between the locations that are hosting the Primary and Replica servers. Properly-configured firewall rules that permit the replication between the Primary and Replica sites are the minimum requirements before you can configure and implement the Hyper-V Replica feature.

Before I start to explain how to configure Hyper-V Replica, I ought to explain that Hyper-V Replica works with two security models:

- Workgroup Security Model
- Domain Security Model

In "Workgroup Security Model", you configure Hyper-V Replica in a Workgroup environment but this model requires a certificate based authentication model which is outside the scope of this article. In this article we're going to use the domain security model, based on the regular Active Directory domain. As a high level process, four steps are performed to implement Hyper-V Replica. We'll assume that you have

already installed two or more Windows Server 2012 servers and enabled those Hyper-V Roles on them which are going to participate in the replication. Please also make sure to join these servers to the domain if you have not done so already.

- Step 1 - Enabling Hyper-V Replica in the disaster recovery site
- Step 2 - Enabling Virtual Machines for replication.
- Step 3 - Enabling the Firewall Rules
- Step 4 - Finalizing the configuration

Step 1 - Enabling Hyper-V Replica in the disaster recovery site

There are three configuration options available when you click on the Hyper-V Settings > "**Replication Configuration**" setting in the Left Pane as shown in the **figure 1.1** below:

- Enable this computer as a Replica Server
- Authentication and Ports
- Authorization and Storage

Enable this computer as a Replica Server

To enable the Hyper-V Replica, check "**Enable this computer as a Replica Server**" on the Hyper-V Server Settings page as shown in the figure 1.1 below:

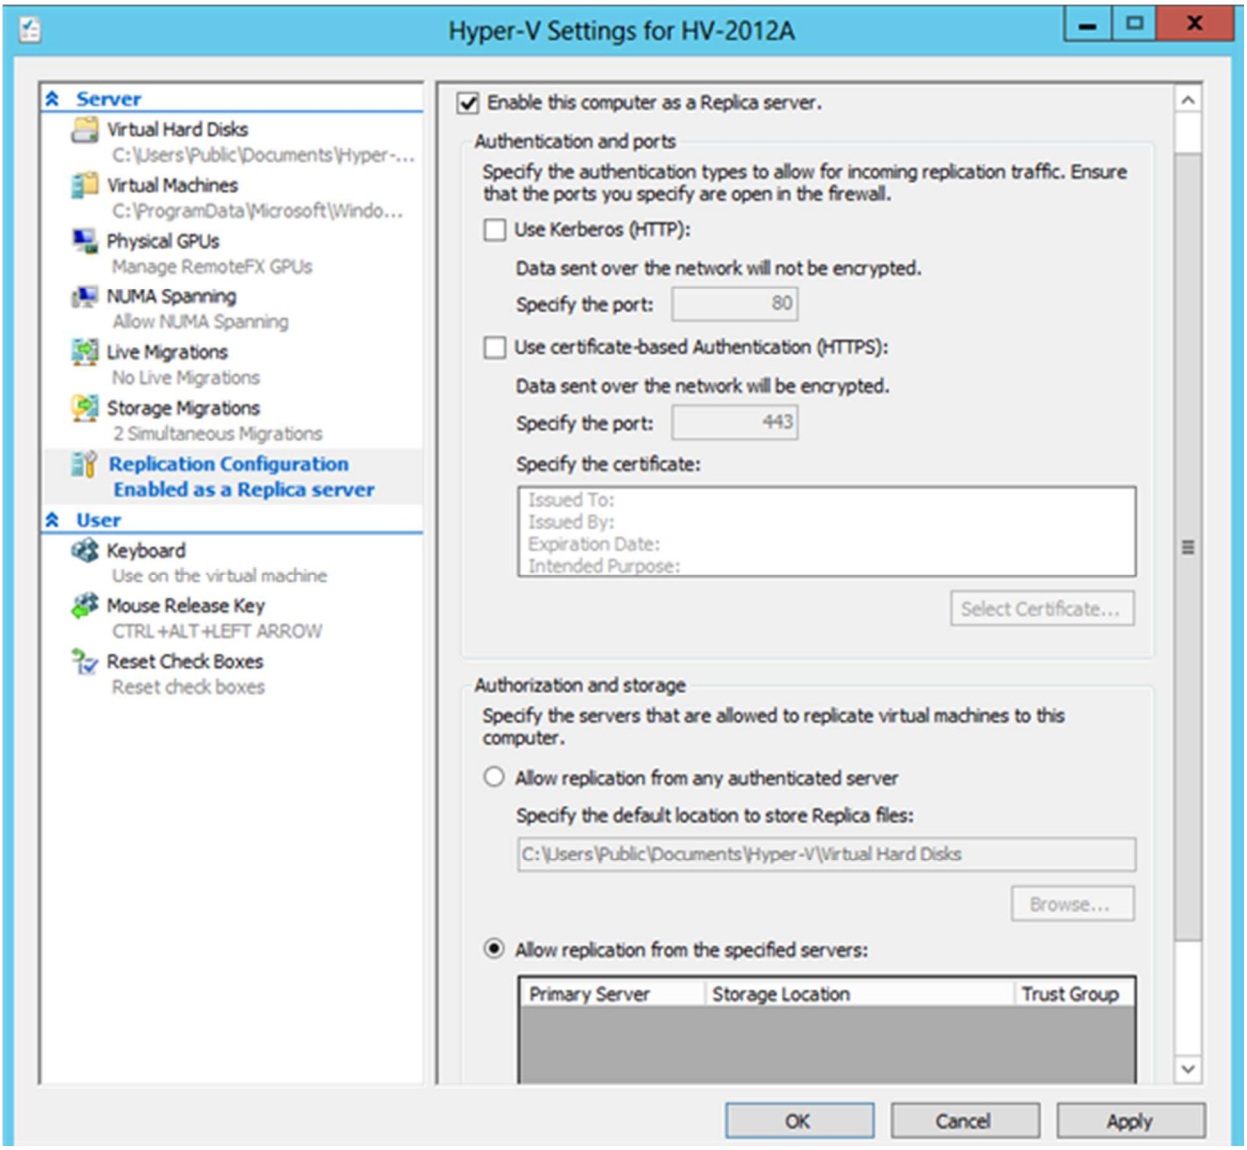

Figure 1.1 – Hyper-V Replica Settings Page

Note: By default, the Hyper-V Replica is not enabled with the first install.

Tip: This is the only configuration setting provided in Hyper-V Manager UI to enable Hyper-V Replica on the Replica Server other than configuring a Primary Virtual Machine to participate in the replication which you will perform on Primary Server.

Authentication and Ports configuration

Hyper-V Replica and Primary Servers will authenticate over one of two protocols; either "Kerberos HTTP" or "Certificate-based Authentication using HTTPS". You can select the authentication mechanism using the "**Authentication and Ports**" form on this settings page.

Tip: HTTP port 80 and HTTPS port 443 are, by default, enabled in the firewall exception on the Replica Server. If you change the port number here, you must be sure to modify the existing Hyper-V Replica firewall rule and then enable the firewall exception so the Replica Server can receive the packets from Primary Server.

Note: You can use Certificate-based Authentication (Using HTTPs) if you need the data to be sent over the network be encrypted. The data will not be encrypted if you select Kerberos Authentication (HTTP).

Authorization and Storage configuration

For authorization, you have a choice:

- "**Allow Replication from any authenticated Server**" option allows this Hyper-V Replica Server to accept virtual machine replication packets from any Primary Server which has successfully authenticated using one of the authentication mechanisms discussed above.
- "**Allow Replication from the specified Servers**" option allows this Hyper-V Replica Server to accept virtual machine replication packets only from the servers which are part of the list you have configured.

Both options; "**Allow Replication from any authenticated Server**" and "**Allow Replication from the Specified Servers**" allow you to specify the storage location for the replicated Virtual Machines. This Is useful for configuring different storage locations for Virtual Machines being replicated from different Hyper-V Primary Server. You can also create the Trust Groups so that Replica Servers accept replication packets only from the trusted servers.

Click "**Ok**" or "**Apply**" to configure this server as a Hyper-V Replica Server.

Please note:

- By default, no authentication type is selected and if you try to "**Apply**" the configuration or try to configure this server as a Hyper-V Replica Server, the following error message is shown.

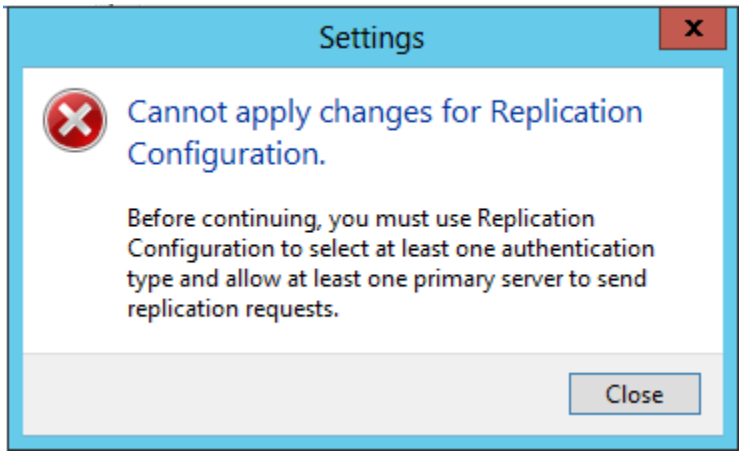

Figure 1.2 – Configuration Page Error Message for Authentication

- Similarly, you must enable the "Firewall Exception Rule" for Hyper-V Replica to operate successfully. The message for enabling the firewall rule is shown when you hit the "**Apply**" button on the "Enable Replication Configuration" Page as shown in the **figure 1.3** below.

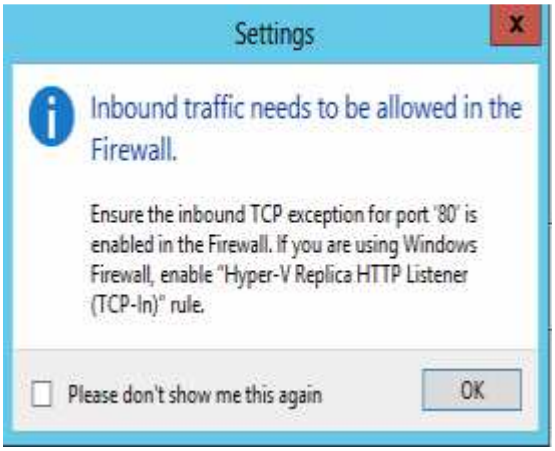

Figure 1.3 – Firewall Rule Enabling Message

There are two types of backup copies provided by Hyper-V Replica for Virtual Machines:

- **Standard Replica Backup copy**
- **Application-Consistent Backup copy**

The "**Standard Replica Backup Copy**" can be further divided into two types:

- **Latest Recovery Backup Copy**
- **Point-In-Time Backup Copy**".

The "**Latest Recovery Backup Copy**" is a backup copy that is created by the Primary Server for Primary Virtual Machine and only one backup copy is created and managed. This backup copy is created whenever any changes occur to the base VHD file of virtual machine. The changes are merged into the base VHD and are replicated every 5 minutes. The Primary Server maintains "single" log/ Standard Replica backup copy in which it keeps the changes to be replicated till replication interval occurs. Selecting "**Only the latest recovery point**" in the wizard enables you to have this type of backup copy as shown in the **figure 1.11** above.

The "**Point-In-Time Backup Copy**" allows a Primary Server to keep multiple backup copies for a Primary Virtual Machine. The "**Point-In-Time Backup Copy**" is created every hour by the VSS component and a maximum of fifteen backup copies can be created as highlighted in the wizard shown in **figure 1.11** above. The one-hour interval for creating these backup copies cannot be changed. Selecting "**Additional recovery points**" in the wizard enables you to have this type of backup copy. Selecting "**Additional recovery points**" enables other options under this category which is to specify the additional recovery points in the "**Number of additional recovery points to be stored**" text box.

Both "**Latest Recovery Backup Copy**" and "**Point-In-Time Backup Copy**" are sent to Replica Server. Enabling "**Additional Recovery points**" will put more overhead on the "Primary Server" and requires more system resources for processing and enough storage to store the "**Point-In-Time Backup Copy**".

Tip: "Additional recovery points" will appear as "Standard Replicas" when you use any of the failover types discussed in the next part of this article. "Additional Recovery Points" are sometimes called "Crash-Consistent" backup copies.

The "**Replicate incremental VSS copy every**" option (as shown in the **figure 1.11** above) is different from "**Only the latest recovery point**" and "**Additional Recovery points**". This is sometimes referred as "Application-Consistent" backup copy. An "Application-Consistent" backup copy is created by the Hyper-V VSS Writer for the applications which are running in the Primary Virtual Machine. The Hyper-V VSS Writer communicates with Hyper-V VSS Requestor Service in Primary Virtual Machine, running as part of the Hyper-V Integration Service Components, to create an "Application-Consistent" backup copy.

Tip: "Hyper-V VSS Requestor Service must be running in the Primary Virtual Machine before "Application-Consistent" backup copies can be created. The service must also be running in the Replica Virtual Machine before the "Application-Consistent" backup copies can be applied.

By default, interval for creating Application-consistent backup copies is every 1 hour but you can change the interval by moving the slider to the hour of your choice at a maximum of 12 hours as shown in the **figure 1.11** above.

Tip: Hyper-V Replication happens every 5 minutes but the "Application-Consistent Backup Copy" is created every 1 hour. That means the applications running inside the Primary Virtual Machine have a good backup copy only after 1 hour!

Tip: Creation of backup copies (discussed above) for Primary Virtual Machine takes place at Primary Server. Replica Server is just notified of a new copy. Replica Server then keeps these copies as the historical data for recovery purpose if you have selected "Additional Recovery Points" option in figure 1.11 above.

Because it has an impact on the performance of Hyper-V Server if you enable the "**Replicate incremental VSS copy every**" option, you must enable this option only if you think that an application running inside the Virtual Machine will require this functionality. For example, for any application hosted in Primary Virtual Machine which has its own VSS Writer to recover its data then you must enable it.

Coming back to the configuration page; the configuration page also shows the space required for storing the default (4 copies) recovery points for the Primary Virtual Machine which is 2.7 GB. This is shown in the above figure 1.11. If you divide this figure by four, you will get size for storing one recovery point. It completely depends on the size of the Primary Virtual Machine.

The backup copies, as discussed above, are created by the Volume Shadow Copy Service.

If you need to restore a Primary Virtual Machine or bring the Primary Virtual Machine online at the Replica Server in case if something goes wrong on the Primary Server, then there will be following recovery points provided for your selection:

- Latest Recovery Point - This one is always available and is the last replication packet sent by the Primary Server to Replica Server
- Standard Replica_<TimeStamp> - Up to 15 copies
- Application-Consistent_<TimeStamp> - Up to 15 copies

Good practice to check following items to make sure that the newly implemented and configured Hyper-V Replica environment is healthy:

a. Hyper-V identifies the Virtual Machines on both the Primary and Replica Servers by the GUID of Virtual Machine. So you can rename the Primary Virtual Machine and Replica Virtual Machine so that it is easy for you to identify the virtual machine type running on both the servers. For example, I would rename Primary Virtual Machine to "Production VM" and Replica Virtual Machine to "DR VM".
b. Make sure to check that the Primary Virtual Machine is always running or ON and Replica Virtual Machine running on the Replica Server is OFF.
c. Perform replication health check for Primary Virtual Machine to ensure that replication is happening.

d. Perform a "Test Failover" on the Replica Virtual Machine to ensure it comes online in case of any disaster at Primary Server. "Test Failover" is explained in detail in the latter part of this article.
e. Please make sure that the Hyper-V VSS Requestor Service is running on both Primary and Replica Virtual Machines; otherwise Application-Consistent backup copies will "not" be created.
f. The **EnableWriteOrderPreservationAcrossDisks** option must be set for Virtual Machine hosting applications that save data across VHD files. The option determines whether all Virtual Machine VHD files are replicated to the Replica Server to the same point in time