

This document provides guidelines to use when setting up your disaster recovery environment using Zerto Virtual Replication to ensure the optimum performance for your environment.

The following guidelines are provided:

- [General Sizing Considerations](#)
- [WAN Sizing Requirements with Zerto Virtual Replication](#)
- [Repository Sizing Required for Offsite Backups](#)

General Sizing Considerations

Zerto Virtual Replication is designed to be scalable, supporting large environments as effectively as it supports small environments. When installing Zerto Virtual Replication it helps to know in advance what the normal usage is going to be in terms of the following:

- [Number of Virtual Machines Requiring Protection](#)
- [VRA Considerations](#)
- [Database Considerations](#)
- [Zerto Cloud Connector Considerations](#)

Number of Virtual Machines Requiring Protection

A Zerto Virtual Manager is installed on a Windows Server machine (2003 SP2 or higher, 2008, 2008R2, 2012 or 2012R2). This machine must be configured with enough CPUs and RAM to handle the number of virtual machines that can be protected or the number of peer sites paired with this Zerto Virtual Manager, as follows:

NUMBER OF VIRTUAL MACHINES OR PEER SITES		NUMBER OF CPUS	RAM SIZE
VIRTUAL MACHINES	PEER SITES		
Up to 750 virtual machines	Up to 5 peer sites	2 CPUs	4GB
Between 751-2000 virtual machines	Up to 15 peer sites	4 CPUs or 2 Dual Core CPUs	4GB
Over 2000 virtual machines	Over 15 peer sites	8 CPUs or 4 Dual Core CPUs	8GB

Each Zerto Virtual Manager can manage up to 5000 virtual machines, either being protected or recovered to that site.

You can only protect a virtual machine in a VPG when the virtual machine has no more than 60 disks. In VMware vSphere environments this requires 4 SCSI controllers each with a maximum of 15 disks. In Microsoft Hyper-V environments, this can be a combination of IDE and SCSI disks, where each virtual machine can have up to 2 IDE controllers each with a maximum of 4 IDE disks and up to 4 SCSI controllers each with a maximum of 15 disks, so that the total of IDE and SCSI disks does not exceed 60 disks.

The following limitations apply when protecting to AWS:

- You can only protect virtual machines in the protected site that are supported by AWS in the recovery site.
- You cannot protect machines that have a disk larger than 16TB.
- Amazon Web Services supports virtual machines running a Windows operating system with up to 26 volumes, including the boot disk.
- Amazon Web Services supports virtual machines running a Linux operating system with up to 40 volumes, including the boot disk.
- Virtual machines running a Linux operating system with more than 1 volume cannot be failed back from AWS after a failover or move operation.
- A Zerto Cloud Appliance can manage up to 200 virtual machines being recovered.

VRA Considerations

Each VRA can manage a maximum of 1500 volumes, whether these are volumes being protected or recovered.

The amount of memory allocated to a VRA when it is installed determines the maximum buffer size for the VRA for buffering IOs written by the protected virtual machines before the writes are sent over the network to the recovery VRA. The recovery VRA also buffers the incoming IOs until they are written to the journal. If a buffer becomes full, a `Bitmap Sync` is performed after space is freed up in the buffer.

AMOUNT OF VRA RAM	VRA BUFFER POOL SIZE
1GB	450MB
2GB	1450MB
3GB	2300MB
4GB	3,300MB
5GB	4,300MB
6GB	5,300MB
7GB	6,300MB
8GB	7,300MB
9GB	8,300MB
10GB	9,300MB
11GB	10,300MB
12GB	11,300MB
13GB	12,300MB
14GB	13,300MB
15GB	14,300MB
16GB	15,300MB

The protecting VRA can use 90% of the buffer for IOs to send over the network and the recovery VRA can use 75% of the buffer. That is, for example, a protecting VRA defined with 2GB of RAM can buffer approximately 1305MB before the buffer is full and a `Bitmap Sync` is required.

Zerto recommends using the default amount of RAM.

Note: The number of virtual machines that a VRA can support is not dependent on the amount of VRA RAM.

Database Considerations

Zerto Virtual Replication requires a database to manage information for disaster recovery. SQL Server Compact (SQL CE) is included as an internal database with Zerto Virtual Manager.

SQL CE has a maximum database size of 4GB. The larger the environment, the larger the database size that is required by Zerto Virtual Replication. When the database size reaches around 60% capacity, performance starts to degrade as database lookups take longer. Zerto recommends using Microsoft SQL Server for medium size and larger environments to prevent performance degradation. During the Zerto Virtual Replication installation you can specify that you want to use an SQL Server database for the Zerto Virtual Replication database. Zerto supports SQL Server Enterprise, SQL Server Standard, and SQL Server Express editions.

Note: If SQL Server is used, it is your responsibility to make sure that database downtime is planned in coordination with your disaster recovery and business continuity requirements. During database downtime, there will be inconsistencies between Zerto Virtual Managers, such as the management of checkpoints, resulting in problems if a recovery is required.

Use the following tables to determine the deployment size and thus, the database to use for Zerto Virtual Replication. The criteria used to determine the deployment size are: number of hosts, sites, virtual machines, and VPGs. When the deployment sizes are different across the criteria, always choose the database based on the larger deployment size. For example, an

environment with only 2 sites but 100 hosts is considered for Zerto Virtual Replication purposes as a medium-very large sized site and should use SQL Server.

Number of Hosts

NUMBER OF PROTECTING HOSTS	DEPLOYMENT SIZE	DATABASE TO USE
Up to 40 hosts	Very small-Small	Zerto Virtual Replication internal database
Over 40 hosts	Medium-Very large	SQL Server

Number of Sites

Usually, more hosts means more sites, especially when referring to cloud service providers (CSPs). The number of sites affects your decision to use the Zerto Virtual Replication internal database or SQL Server.

NUMBER OF SITES	DEPLOYMENT SIZE	DATABASE TO USE
Up to 4 sites	Very small-Small	Zerto Virtual Replication internal database
Over 4 sites	Medium-Very large	SQL Server

Number of Protected Virtual Machines

The number of protected virtual machines affects your decision to use the Zerto Virtual Replication internal database or SQL Server.

NUMBER OF VIRTUAL MACHINES	DEPLOYMENT SIZE	DATABASE TO USE
Less than 400	Very small-Small	Zerto Virtual Replication internal database
Greater than 400	Medium-Very large	SQL Server

Number of VPGs

The number of VPGs affects your decision to use the Zerto Virtual Replication internal database or SQL Server.

NUMBER OF VPGS	DEPLOYMENT SIZE	DATABASE TO USE
Less than 200	Very small-Small	Zerto Virtual Replication internal database
Greater than 200	Medium-Very large	SQL Server

Zerto Cloud Connector Considerations

A Zerto Cloud Connector requires the following, regardless of the size of the site:

- 4GB disk space
- At least 1GB of reserved memory
- 1 vCPU

Zerto recommends using a 10Gbps NIC for each Zerto Cloud Connector, enabling it to handle 10Gbps of traffic.

WAN Sizing Requirements with Zerto Virtual Replication

When preparing your deployment, you need to verify that the connectivity between any two paired sites has bandwidth capacity that can handle the data to be replicated between the sites. This section provides information about WAN requirements when using Zerto Virtual Replication. It will help you to estimate the WAN requirements for disaster recovery,

and describes how to collect and analyze the performance statistics for the virtual machines to protect using Zerto Virtual Replication. These actions should be performed to ensure that connectivity between the protected and recovery sites has sufficient bandwidth capacity to handle the amount of data to be replicated between the sites.

Zerto Virtual Manager employs sophisticated data compression algorithms to reduce the bandwidth required between the protected and recovery sites. Although this compression can be very effective in reducing bandwidth requirements, its efficiency is highly dependent on the specific characteristics of the data being replicated.

The following files are provided in the WAN Sizing zip to help with WAN sizing:

- Zerto Virtual Replication WAN Sizing Estimator
- Example scripts for WAN Sizing

Note: Zerto Virtual Replication requires a minimum dedicated bandwidth of at least 5Mbps. You can also work with third-party WAN optimization and acceleration technologies, such as those supplied by Riverbed Technologies and Silver Peak.

Supported WAN Links

Zerto virtual replication requires a minimum bandwidth of 5Mbps both upstream and downstream from each protected site and from each recovery site. Zerto replicates using TCP/IP based connectivity and supports the following types of WAN links:

- LAN
- MPLS
- IP Link over Fiber
- Site-to-site VPN

Zerto Virtual Replication does not support replicating through NAT translated devices via port forwarding. For replication to occur through NAT translated devices, all replication traffic must be within a VPN.

WAN Optimization Considerations

This section briefly discusses some factors that can affect system optimization and that should be considered in addition to those discussed in [“Estimating the Bandwidth Requirements Between Sites”, on page 6.](#)

Zerto Virtual Replication uses the maximum bandwidth available as well as optimization and internal compression to maintain low RPOs. This removes the need to use third-party WAN acceleration devices. However, Zerto Virtual Replication does work with third-party devices, including those supplied by Riverbed Technologies, Silver Peak, and others.

Note: Zerto Virtual Replication WAN compression can be disabled for any VPG as required. In addition, bandwidth can be throttled as necessary and QoS priorities can be specified for each VPG.

If the WAN link becomes saturated, disconnected, or the data change rate is too high to maintain continuous replication, Zerto Virtual Replication enters a Bitmap sync mode. In this mode, VRA memory stores an index of which blocks have changed. When the condition causing the Bitmap sync is resolved, the bitmap changes are read from the source disk and replicated, and continuous replication resumes.

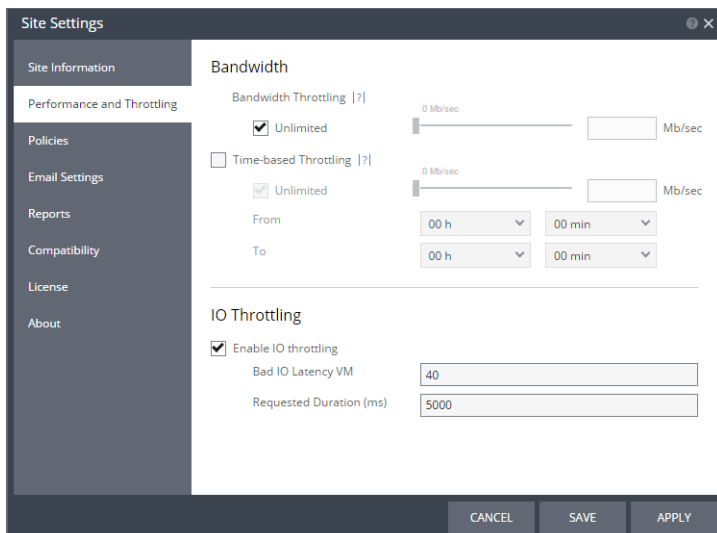
A Delta sync is used when a WAN link is down for too long or when a VRA is rebooted with protected VMs running. In this case, the source and target virtual disks are compared, changes in the protected site are replicated to the recovery site, and continuous protection resumes.

Bandwidth Throttling to Temporarily Free Bandwidth

Zerto Virtual Replication can be configured to throttle the bandwidth it uses during a specific time period. During this period, replication does not use more than the bandwidth selected. At the end of the period, bandwidth throttling ends.

To configure time-based throttling:

1. In the Zerto user interface, select *Site Settings > Performance and Throttling*.
The *Performance and Throttling* dialog is displayed.



2. To specify time-based throttling, check the `Time-based Throttling` checkbox.
3. Specify the maximum bandwidth for the period.
If the `Unlimited` checkbox is checked, the bandwidth is always unlimited.
A value of 0 Mb/sec means that the bandwidth used is determined automatically by Zerto Virtual Replication.
You can use the slider to set the Mb/sec.
 - a) In the `From` fields, select the start time for throttling.
 - b) In the `To` fields, select the end time for throttling.
4. Click `APPLY` or `SAVE`.

Using Swap Disks to Lessen the Bandwidth Load

A swap disk is a virtual disk that Zerto Virtual Replication replicates and then isolates from further changes. If temporary files on a protected virtual machine are moved to a swap disk, although the disk itself is replicated to the recovery site, changes to it are not replicated. This reduces replication traffic, journal size, and storage use.

Swap disks are configured for the individual disk in the virtual machine in the VPG definition.

Preseeding to Initially Lessen the Bandwidth Load

In preseeding, the entire virtual machine is block-copied and moved to the recovery site. Only changes made after the preseeding are replicated.

Setting the Priority for Using the WAN by Multiple VPGs

Zerto Virtual Replication enables prioritizing replication for each VPG. The priorities are:

- High
- Medium
- Low

Prioritization is used when there is bandwidth contention between sites or when the load on the VRA is more than can be accommodated. It enables maintaining a low RPO on the most critical VPGs. As bandwidth becomes available, lower priority VPGs are synchronized and returned to a continuous protection state.

Pausing Replication

Zerto Virtual Replication enables VPG replication to be paused, for example, during maintenance.

Estimating the Bandwidth Requirements Between Sites

Estimating the bandwidth requirements between the protected and recovery sites involves the following:

1. Collect data characteristics for protected VMs.
2. Calculate the estimated bandwidth requirements.

Note: When the recovery site is Amazon Web Services (AWS), you estimate the required bandwidth for the protected machines, either in vSphere or Hyper-V as described below.

- [“vSphere Environments: Collecting Data Characteristics for VMs”](#), below
- [“Hyper-V Environments: Collecting Data Characteristics for VMs”](#), on page 9

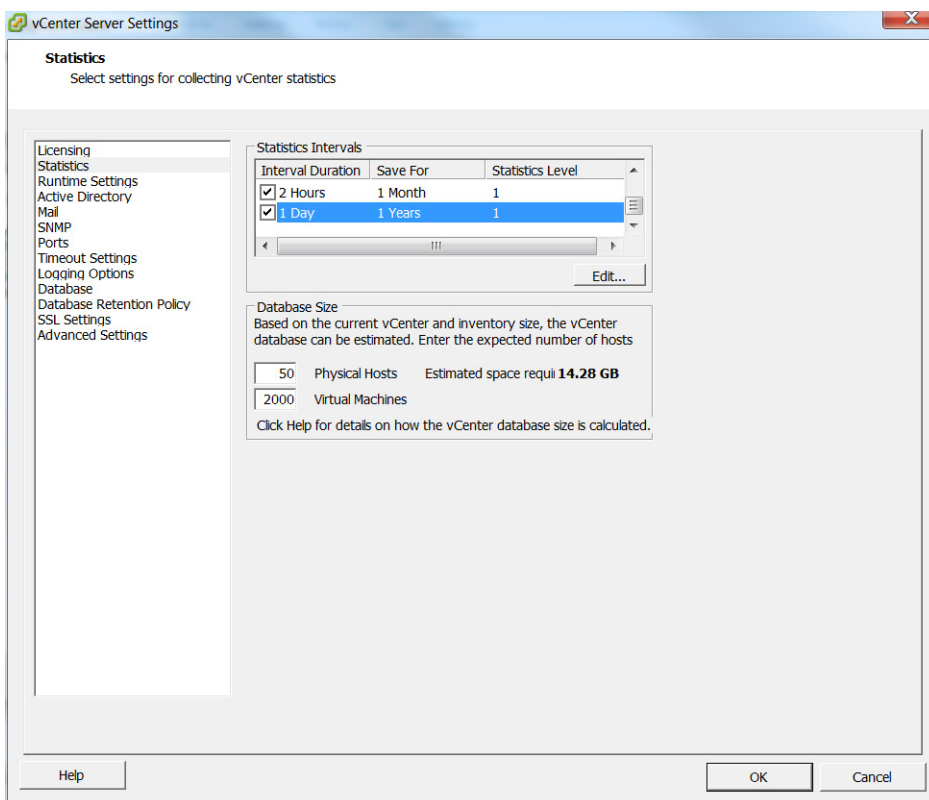
vSphere Environments: Collecting Data Characteristics for VMs

Before you can collect the required data, you must first enable data collection in vCenter Server.

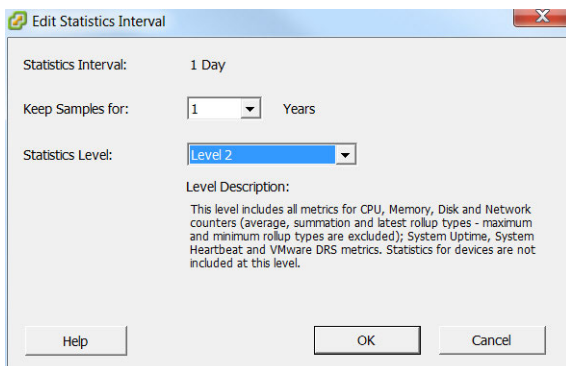
Enabling vCenter Server Data Collection

To enable vCenter Server data collection:

1. Connect to the vCenter Server.
2. In the Administration menu item, select `vCenter Server Settings`.
The `vCenter Server Settings` dialog is displayed.
3. Select `Statistics`.



4. Make sure that the *Statistics Level* value for all interval durations up to and including the one day duration is at least 2. If any of the durations have a value less than 2, do the following, starting with the smallest interval:
 - a) Select the interval and click *Edit*.
 - b) Change *Statistics Level* to *Level 2*.
 - c) Click *OK*.



5. Repeat step 4 for all the values up to and including the 1 day interval duration.
6. Click *OK* and wait for at least a day before using the aggregate usage data.

Collecting Data Characteristics for VMs

You can collect data characteristics for the virtual machines in a VPG in one of the following ways:

- Via vSphere Client console performance statistics.
- By running a script to collect the data characteristics.

Note: The script samples supplied with the download, require vSphere PowerCLI and permissions to access the vCenter Server using the script.

- By using operating system performance monitors, such as the Microsoft Performance Monitor utility for Windows operating systems or the `iostat` command for Linux operating systems.

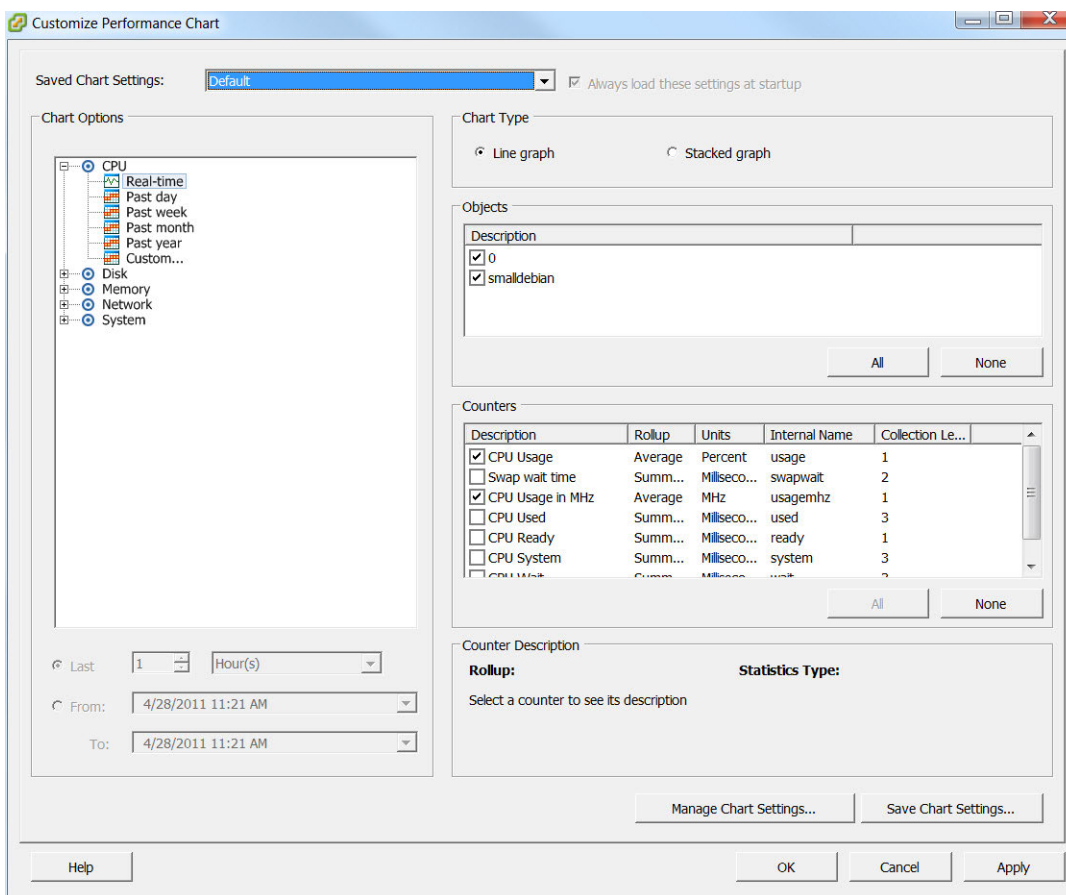
Collect data for a minimum of one day. Collecting this information impacts on performance and therefore the collection period should be long enough to gather a true representation of usage but not too long. The first procedure described below, to collect data characteristics for the VMs via the vSphere Client console performance statistics, uses a timeframe of one day and the second procedure, to collect data characteristics for the VMs by running a script to collect the data characteristics uses a timeframe of seven days.

Note: When running vCenter Server versions before version 5.x, if any of the virtual machines use NFS storage, metrics for the NFS storage are not generated by the vCenter Server.

To collect data characteristics for the VMs via the vSphere Client console performance statistics:

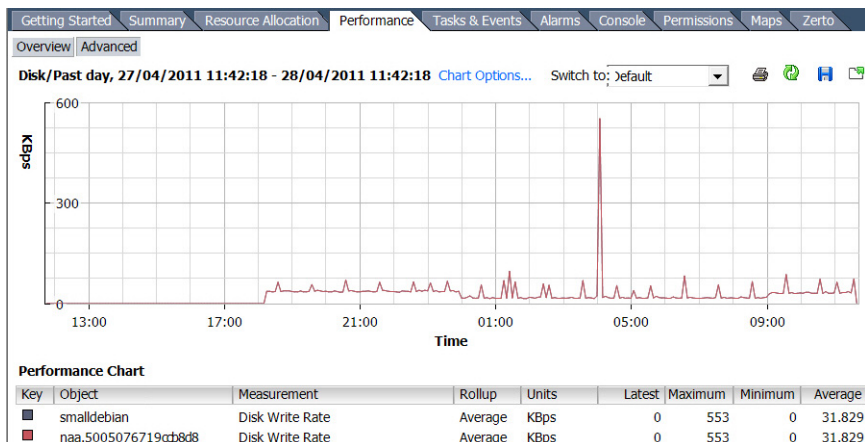
1. In the vSphere Client console select the VM and open the *Performance* tab.
2. Click *Advanced*.
3. Click the *Charts Options* link.

The *Customize Performance Chart* dialog is displayed.



4. In *Chart Options*, drill-down in *Disk* and select *Past day*.
5. In *Counters*, click *None* to clear all the selections and then select *Disk Write Rate* or *Write Rate*.
6. Click *OK*.

A chart similar to the following is generated:



Use the chart for the average write rate of the VM.

To collect data characteristics for the VMs via a script:

Note: The following script and the samples supplied with the download, require vSphere PowerCLI and permissions to access the vCenter Server using the script.

- Run a script similar to the following:

```
$report = @()
Get-VM | %{
    $stats = Get-Stat -Entity $ -Stat disk.write.average -Start (Get-Date).adddays(-7) -
    ErrorAction SilentlyContinue
    if($stats){
        $statsGrouped = $stats | Group-Object -Property MetricId
        $row = "" | Select Name, WriteAvgKBps, WriteAvgMBps
        $row.Name = $_.Name
        $row.WriteAvgKBps = ($statsGrouped |
            where {$_.Name -eq "disk.write.average"} |
            %{$_.Group | Measure-Object -Property Value -Average}).Average
        $row.WriteAvgMBps = $row.WriteAvgKBps/1024
        $row.WriteAvgKBps = "{0:N2}" -f $row.WriteAvgKbps
        $row.WriteAvgMBps = "{0:N2}" -f $row.WriteAvgMBps
        $report += $row
    }
}
$report | Export-Csv "C:\ZertoOutput.csv"
```

Note: If you want a value other than seven days, change the value of the `adddays()` function. For example to collect data for three days, use `adddays(-3)`.

Use the resulting file, `C:\ZertoOutput.csv`, for the average write rate of the VM.

Note: Versions of this script are included in the download with this document.

Hyper-V Environments: Collecting Data Characteristics for VMs

You can collect data characteristics for the virtual machines in a VPG in one of the following ways:

- By using operating system performance monitors, such as the Microsoft Performance Monitor utility for Windows operating systems or the `iostat` command for Linux operating systems.
- By using Windows PowerShell in Windows Server 2012 to collect network utilization (and other information). When using metering ACLs, you can measure the total network traffic sent and received by a virtual machine. To collect performance characteristics for the virtual machines in a VPG, using PowerShell, do the following:
 - Turn on resource metering for the relevant virtual machines, if it is not already enabled.

- Adjust the collection frequency, if necessary.
- Collect the relevant statistics.

Zerto recommends that you collect data for a minimum of one day. When you have enough statistics, you may want to turn off resource metering since data collection can impact performance.

Turning on Resource Metering

By default, resource metering is not enabled. To turn on resource metering for one virtual machine, enter the following PowerShell command:

```
Get-VM <VM-name> | Enable-VMResourceMetering
```

To turn on monitoring for all virtual machines on a server at one time, enter the following PowerShell command:

```
Get-VM | Enable-VMResourceMetering
```

Once you enable resource metering, Hyper-V begins to collect data. You can reset metering at any time, which discards the data that has been collected up to that point.

If resource metering is enabled but no NetworkAdapterAcls are configured, Hyper-V configures them to measure total network traffic. To measure network traffic through an IP range, configure the NetworkAdapterAcls for the IP range before running Enable-VMResourceMetering.

Adjusting the Collection Frequency

By default, the collection frequency is once every hour. You can change the collection frequency, but understand that data collection can impact performance. To change the collection frequency, enter the following command:

```
Set-VMHost -ComputerName <host-server-name> -ResourceMeteringSaveInterval <HH:MM:SS>
```

The collection frequency is always set at the host server level. You cannot adjust the collection frequency per virtual machine.

For example, if you enter 01:30:00, resource consumption will be collected every hour and a half.

Collecting and Viewing the Relevant Statistics

To view resource usage for one virtual machine, enter the following command:

```
Get-VM <VM-name> | Measure-VM
```

Resource metering data can be displayed for all of the virtual machines that are running on a host. To see data for all of the virtual machines on a host, enter the following command:

```
Get-VM | Measure-VM
```

You can configure PowerShell to display only certain statistics. To do this, you must know the object names that PowerShell assigns to each statistic. You can see the object names by entering the following command:

```
Get-VM | Measure-VM | Select-Object *
```

For example, when working with Zerto Virtual Replication, you are interested in network traffic. To list the network traffic for each virtual machine, enter the following command:

```
Get-VM | Measure-VM | Select-Object VMName, NetworkMeteredTrafficReport
```

You can use VM Network Adapter ACLs to measure network activity to and from a specific network. For example, to meter network traffic for a special subnet or IP address:

```
Add-VMNetworkAdapterAcl -VMName <VM-name> -Action Meter -RemoteIPAddress 10.10.0.0/16 -
Direction Outbound
```

Turning off Resource Metering

To disable the collection of performance statistics, enter the following PowerShell command:

```
Disable-VMResourceMetering -VMName <VM-name>
```

Calculating the Estimated Bandwidth Requirement

Note: The *Zerto WAN Sizing Estimator* is included in the download with this document.

Use the average write rate for the virtual machines in a VPG in the *Zerto WAN Sizing Estimator* to estimate the minimum bandwidth required.

For each VM you also must decide whether compression will be enabled for the VM, based on the data characteristics.

To estimate sizing using the Zerto WAN Sizing Estimator:

1. Open the *Zerto WAN Sizing Estimator*.
2. Enter the following information in the *VM data* sheet:
 - The VM name.
 - The Write KB/s data, based on the statistics gathered in the previous task. Use a period for the decimal mark.
 - Define whether compression is enabled for this VM: Select Yes or No.
 - The application data characteristics: Select Compressed or Compressible.

Note: The *Zerto WAN Sizing Estimator* colors the cell red if you decide to employ compression on compressible data and orange if you decide to avoid compression for compressible data.

The *Zerto WAN Sizing Estimator* estimates the total bandwidth needed for your deployment, using a minimum value of 5 Mb/sec. The estimation is displayed on the top of each page of the *Zerto WAN Sizing Estimator*.

Note: The *Zerto WAN Sizing Estimator* also includes an *Initial Sync Calculator* sheet which you can use to estimate the initial synchronization time given the disk sizes or the virtual machines in a VPG, the available bandwidth and whether or not pre seeding is used. The result is an estimate and could be increased by extensive write I/O activity and/or storage latency. The *Initial Sync Calculator* takes into account the actual ongoing change rate data from the *VM data* sheet as well as automatically using the compression parameter from the *VM data* sheet.

You can estimate the WAN sizing required without using the *Zerto WAN Sizing Estimator* using the following procedure.

To estimate sizing without using the Zerto WAN Sizing Estimator:

1. For each virtual machine in the VPG multiply the KB/sec, based on the statistics gathered, by 8 and divide the result by 1024 to provide an answer in Mb/sec. Divide this result by 2 if compression is enabled for the VM and the data is compressible.
2. Sum the results of step 1.

```
WAN Mb/sec = SUM(KB/sec * (8/1024/(1 or 2 if compressible data that will be compressed)))
```

The result is an estimate of the required Mb/sec for the WAN.

Note: If the result is less than 5 Mb/sec, you must use a minimum dedicated bandwidth of at least 5 Mb/sec.

Repository Sizing Required for Offsite Backups

Zerto Virtual Replication provides an offsite backup option that enables saving the protected virtual machines offsite for up to one year in a state where they can be easily deployed.

Offsite backups are managed by the Virtual Backup Appliance (VBA). The VBA is installed as part of the Zerto Virtual Replication installation as a Windows service. During an offsite backup, the VBA communicates with the VRAs on the recovery site to create the virtual machine files, such as the configuration and virtual disk files and these are saved in a repository. The offsite backups are fixed points saved either daily or weekly in the repository. Before you can create an offsite backup for virtual machines, you must first create one or more repositories for the offsite backup jobs. The size of the repository is determined according to the following:

- The sum of all virtual machine sizes that are being backed up.
- The frequency backups are saved.
- The length of time the backups are saved.
- Compression considerations.

Note: When an Amazon Web Services (AWS) bucket is defined as the repository for the offsite backups, the size is unlimited.

The equation used to calculate the repository size is the size before any storage deduplication reduction. This deduplication reduction depends on the actual data as well as the change rate and storage capabilities and must be determined per environment, using the deduplication data provided at the site. For information about Windows Server 2012 deduplication and offsite backups, refer to [Zerto Virtual Replication Offsite Backup and Deduplication document](#).

After initializing the VPG, Zerto Virtual Replication periodically checks that the schedule to run an offsite backup has not been passed. At the scheduled backup time, the offsite backup is run and the offsite backup file stored in the specified repository.

Offsite backups are kept on the recovery site for the retention period specified in the VPG. However, over time the number of stored offsite backups is reduced to save space. For details about the reduction of backups over time, refer to *Offsite Backups* in the *Introduction to Protecting Virtual Machines* section in the administration guide.

The maximum number of stored offsite backups is as follows:

	MAXIMUM NUMBER OF BACKUPS WHEN RUN:	
	DAILY	WEEKLY
1 week	7	1
1 month	11	5
3 months	13	7
6 months	16	10
9 months	19	13
12 months	22	16

To calculate the minimum repository size required when compression is used, use the following equation:

$$\text{Size} = ((\text{sum of all virtual machine sizes in VPG to be backed up}) * (\text{Maximum Number of Backups})) / 2$$

To calculate the minimum repository size required when compression is not used, use the following equation:

$$\text{Size} = (\text{sum of all virtual machine sizes in VPG to be backed up}) * (\text{Maximum Number of Backups})$$

The virtual machine sizes are the used sizes and not the provisioned sizes.

For example, if a VPG contains 3 virtual machines with the following used sizes: 220GB, 435GB, 745GB and backups are required on a weekly basis to be held for 9 months, the storage required when compression is used = $(220+435+745) * 13 / 2 = 9100\text{GB}$ and double this amount if compression is not used.

Note: The above equation assumes compression halves the required storage. If you see that this is not the case, based on the data at your site, use the compression figure from your site.

ABOUT ZERTO

Zerto is committed to keeping enterprise and cloud IT running 24/7 by providing scalable business continuity software solutions. Through the Zerto Cloud Continuity Platform, organizations seamlessly move and protect virtualized workloads between public, private and hybrid clouds. The company's flagship product, Zerto Virtual Replication, is the standard for protection of applications in cloud and virtualized datacenters.

www.zerto.com

For further assistance using Zerto Virtual Replication, contact Zerto support at support@zerto.com.