

This document provides **sizing** guidelines when setting up your disaster recovery environment using Zerto Virtual Replication, in order to ensure the **optimum performance for your environment**.

The following guidelines are provided:

- [“General Sizing Considerations”, on page 1](#)
- [“WAN Sizing Requirements with Zerto Virtual Replication”, on page 4](#)
- [“Supported WAN Links”, on page 5](#)
- [“WAN Optimization Considerations”, on page 6](#)
- [“Estimating the Bandwidth Requirements Between Sites”, on page 9](#)
- [“Repository Sizing Required for Offsite Backups”, on page 15](#)

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”, on page 1](#)
- [“VRA Considerations”, on page 2](#)
- [“Database Considerations”, on page 3](#)
- [“Recommendations and Considerations When Using an External SQL Server”, on page 4](#)
- [“Zerto Cloud Connector Considerations”, on page 4](#)

Number of Virtual Machines Requiring Protection

Your disaster recovery environment using Zerto Virtual Replication 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.

For specific environment details, click to open the relevant requirements document.

- [Requirements for vSphere Environments](#)
- [Requirements for Microsoft Hyper-V Environments](#)
- [Requirements for Microsoft Azure Environments](#)
- [Requirements for Amazon Web Services \(AWS\) Environments](#)

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 the **total** of IDE and SCSI disks must **not exceed 60** disks. Each virtual machine can have up to:
 - **2 IDE** controllers each with a maximum of **4 IDE** disks
 - **4 SCSI** controllers each with a maximum of **15** disks

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 data for scalable business continuity software solutions.

- During the Zerto Virtual Manager installation, the user is able to select whether to install and use an **embedded** SQL Server (**localdb**) as the database.
- Alternatively, and also during the installation, the user is able to choose whether to instead select and use an **external** SQL Server instance.
- When updating from Zerto Virtual Replication version 4.5.x, SQL CE will continue to be used as the embedded database.
- From Zerto Virtual Replication version 5.5, as part of scaling up the Zerto Virtual Manager, sizing requirements were increased.
- The larger the environment protected by Zerto Virtual Manager, the larger the database size required to support it.
- When using an embedded SQL Server, the size of the environment that can be protected by Zerto Virtual Manager, is limited.
- **Failure to follow the sizing guideline can result in software errors, and not just performance degradation.**
- After installing Zerto Virtual Manager, you are able to migrate from using an embedded SQL Server, to using an external SQL Server. For guidelines, considerations and procedures, see [Migrating the Zerto Virtual Replication Database to Microsoft SQL Server](#).
- When using an external SQL Server, Zerto supports SQL Server Enterprise, SQL Server Standard, and SQL Server Express editions. Zerto recommends using SQL Server Standard or Enterprise edition, and also recommends placing them on a separate VM from where the Zerto Virtual Manager is running.

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 the Zerto Virtual Managers, such as the management of checkpoints.

Use the following table 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
- VPGs

When the deployment sizes are different across the criteria, always choose the database based on the **larger** deployment size.

DATABASE TYPE	MAX NUMBER OF VPGS WHICH CAN REPLICATE TO THE CURRENT/LOCAL SITE	MAX NUMBER OF PEER SITES
SQL CE (upgrade from 4.5.x)	200	10
External SQL Server	10,000	Up to 100

Recommendations and Considerations When Using an External SQL Server

Before migrating your database, you must review this section.

Note:

To perform a **database migration**, Zerto Virtual Replication version must be **4.5Ux and later**.

- Zerto recommends running the external database on a separate VM.
- Zerto recommends using faster storage for the database. For a Zerto site protecting more than 250 incoming VPG, we highly recommend using fast storage for the database (SSD).
- The SQL Server must use at least 4 CPUs. See the following table:

NUMBER OF VPG PROTECTING TO A SITE	SQL SERVER MEMORY	DISK SPACE
Up to 1000 incoming VPG	8 GB	10 GB
1000 - 2000 incoming VPG	16 GB	25 GB
2000 - 5000 incoming VPG	16 GB	50 GB
5000 - 10000 incoming VPG	16 GB	100 GB

- Although Zerto Virtual Replication supports SQL Server Enterprise/Standard/Express editions, when migrating to an external SQL Server service, Zerto recommends using Standard or Enterprise edition.

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 9](#).

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.

See the following sections:

- [“Bandwidth Throttling to Temporarily Free Bandwidth”, on page 7](#)
- [“Using Swap Disks to Lessen the Bandwidth Load”, on page 7](#)
- [“Preseeding to Initially Lessen the Bandwidth Load”, on page 7](#)
- [“Setting the Priority for Using the WAN by Multiple VPGs”, on page 8](#)
- [“Pausing Replication”, on page 8](#)

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.

The screenshot shows the 'Site Settings' dialog box with the 'Performance and Throttling' tab selected. On the left is a sidebar with links to Site Information, Performance and Throttling, Policies, Email Settings, Reports, Compatibility, License, and About. The main area is divided into two sections: 'Bandwidth' and 'IO Throttling'. In the 'Bandwidth' section, 'Bandwidth Throttling' has a link to help. Below it, 'Unlimited' is checked, and 'Time-based Throttling' is unchecked. There are two identical sets of controls for 'Unlimited' and 'Time-based Throttling', each with a slider and a 'Mb/sec' input field. The 'Time-based Throttling' section also includes 'From' and 'To' time selectors with hour and minute dropdowns. The 'IO Throttling' section has 'Enable IO throttling' checked, and two input fields: 'Bad IO Latency VM' with the value 40, and 'Requested Duration (ms)' with the value 5000. At the bottom are 'CANCEL', 'SAVE', and 'APPLY' buttons.

2. **Time-based Throttling** is **Unlimited** by default. To change this, deselect Time-based Throttling, and enter a value for the bandwidth for a selected period of time:
 - **From:** The hour and the minute to start the throttling, using a 24-hour clock.
 - **To:** The hour and the minute to end the throttling, using a 24-hour clock.

For Example: If you know that the bandwidth needs specific throttling during a certain period, during the daily peak transaction period you can override the throttling of the bandwidth for these specific times.
3. 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. **Collecting data characteristics for protected VMs.** See the following:
 - “vSphere Environments: Collecting Data Characteristics for VMs”, on page 9
 - “Hyper-V Environments: Collecting Data Characteristics for VMs”, on page 13
2. **Calculating the estimated bandwidth requirements.** See the following:
 - “Calculating the Estimated Bandwidth Requirement”, on page 14

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

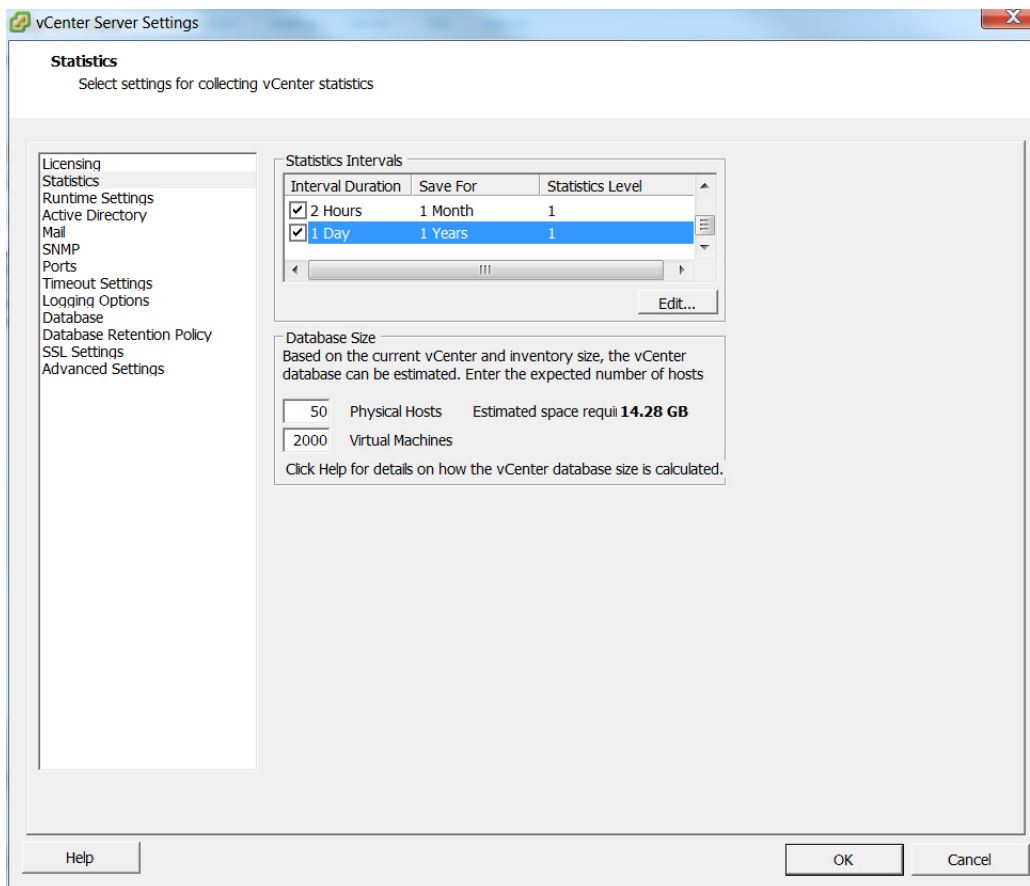
Before you can collect the required data, you must first enable data collection in vCenter Server.

- “Enabling vCenter Server Data Collection”, on page 9
- “Collecting Data Characteristics for VMs”, on page 10

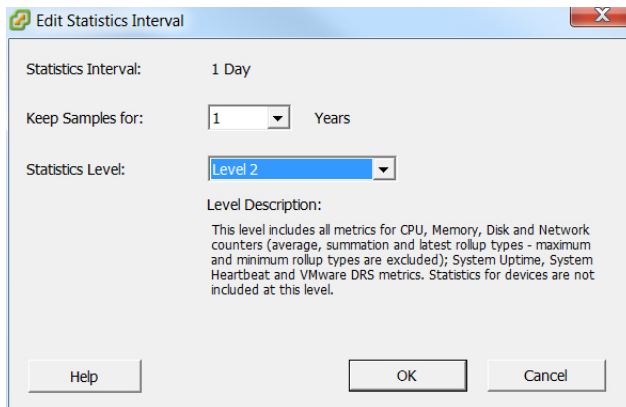
Enabling vCenter Server Data Collection

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. See [Collecting Data Characteristics for VMs via vSphere Client Console Performance Statistics](#).
- Or -
- By running a **script** to collect the data characteristics. See [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.
- Or -
- 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.

Use one of the following procedures:

- [Collecting Data Characteristics for VMs via vSphere Client Console Performance Statistics](#) uses a timeframe of **one** day.
- [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.

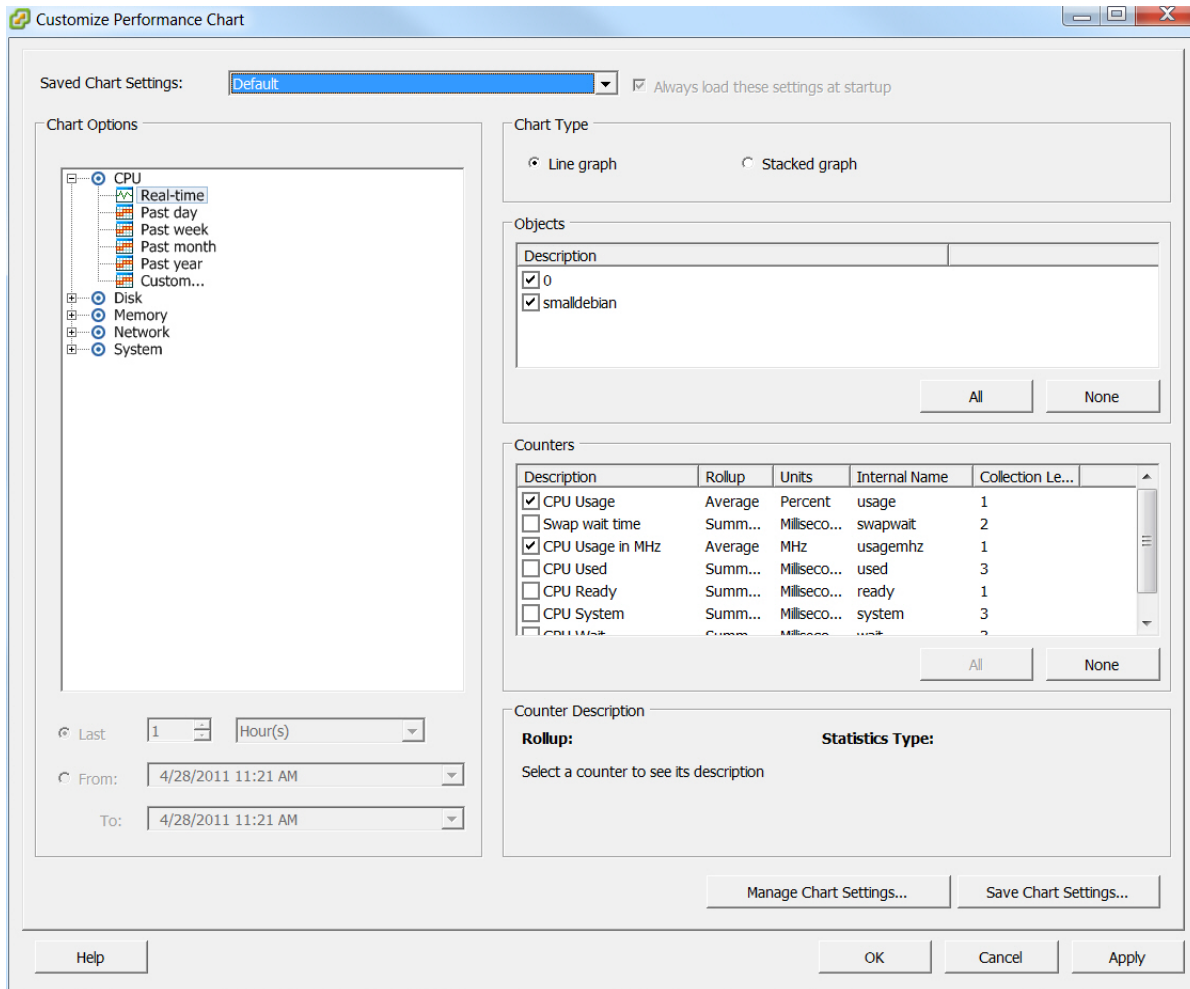
Collecting Data Characteristics for VMs via vSphere Client Console Performance Statistics

Use the following procedure to collect data characteristics for the VMs via the vSphere Client console performance statistics.

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

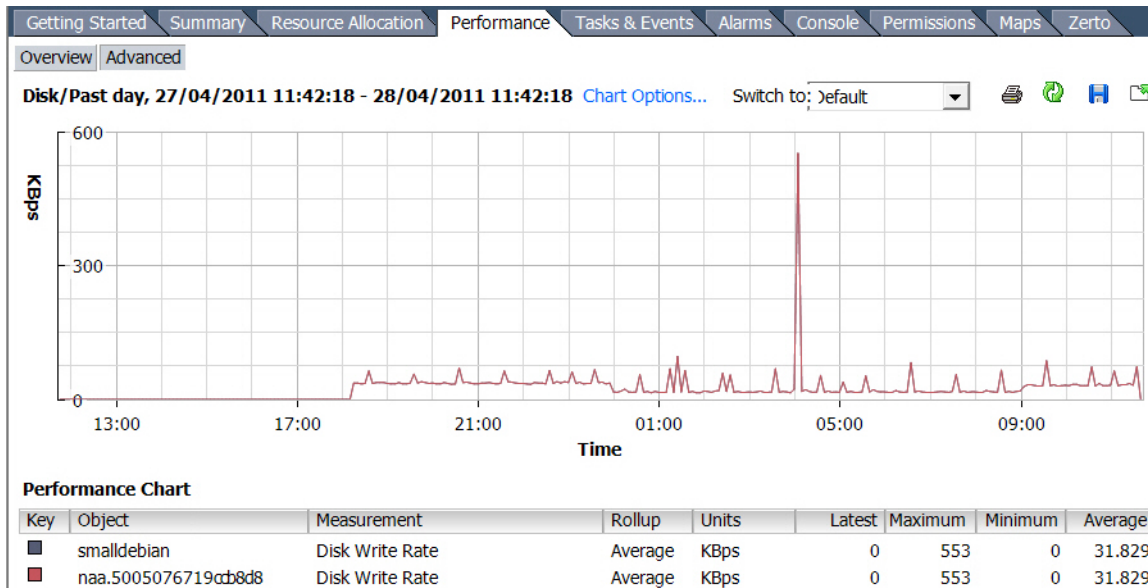
1. In the vSphere Client console select the VM and open the **Performance** tab.
2. Click **Advanced**.
3. Click the **Charts Options** link.

The Customize Performance Chart dialog is displayed.



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

A chart similar to the following is generated:



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

Running a Script to Collect the Data Characteristics

Use the following procedure to run a script to collect the data characteristics.

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

To collect data characteristics for the VMs via a script:

1. 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"
```

2. 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)**.
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.
4. Continue with ["Calculating the Estimated Bandwidth Requirement"](#), on page 14.

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”, on page 13](#)
- [“Adjusting the Collection Frequency”, on page 13](#)
- [“Collecting and Viewing the Relevant Statistics”, on page 13](#)
- [“Turning off Resource Metering”, on page 14](#)

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>

Continue with [“Calculating the Estimated Bandwidth Requirement”, on page 14.](#)

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.

Use one of the following procedures:

- [Estimating sizing using the Zerto WAN Sizing Estimator](#)
- [Estimating sizing without using the Zerto WAN Sizing Estimator](#)

Estimating sizing using the Zerto WAN Sizing Estimator

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**.

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.

3. 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.

TIP:

- 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 preseeding 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.

To estimate the **WAN sizing** required **without** using the Zerto WAN Sizing Estimator, using the following procedure.

Estimating sizing without using the Zerto WAN Sizing Estimator

To estimate sizing without using the Zerto WAN Sizing Estimator:

1. For **each** virtual machine in the VPG multiply the KB/sec (which is based on the statistics gathered) by 8, and divide the result by 1024 to provide an answer in Mb/sec.
 - If **compression is enabled** for the VM and the data is compressible, divide this result by 2.
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, see the section *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.

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.

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.

Repository Sizing Required for Offsite Backups

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**](#).