



VERITAS™



Backup Exec Best Practices Guide 4.0

Here's how to get Backup Exec performing optimally



Backup Exec Best Practice Guide

- The Backup Exec Best Practices Guide includes tips and recommendations to help you plan, install, upgrade, protect and backup your environment more effectively.
- Consult any of the following resources from the Backup Exec User Interface in the **Help and Documentation** menu if you have questions or difficulties :
 - The Administrator's Guide for comprehensive information about Backup Exec.
 - The Backup Exec Help for searchable, topic-based documentation.
 - Backup Exec Support landing page - https://www.veritas.com/content/support/en_US/15047.html
- You should always run a full backup job before and after updating/upgrading Backup Exec, the operating system and/or any applications.
- To achieve optimum performance, not only the Backup Exec Server, but also the underlying infrastructure has to be properly understood, configured and tested as every environment is unique.

Backup Exec Best Practice Guide

DISCLAIMER

- The following Backup Exec Best Practices Guide contains information from a combination of sources. Some parts of the information comes from Veritas and partner field experiences combined with verified/public information from VERITAS and/or other validated publicly available sources.
- Because the guidance given often can change for these rapidly evolving technologies, we strongly encourage you to at all times verify with the available verified/public information for changes. We will regularly update this Backup Exec Best Practices Guide to keep abreast with these changes and maintain the usability of this guide
- For guidelines and best practices on installing and configuring 3rd party applications and platforms, please refer to those 3rd parties for this type of documentations and other resources

Backup Exec Server Tuning

How to analyze your requirements?

- What & Which?
 - Data on which servers do you need to back up?
 - Data/applications: Which are critical and which is less critical?
 - Data/applications: Which need frequent and less frequent backups?
 - Are the recovery SLA's for different types of data and applications?
 - Is/are the connectivity and path(s) from BE server to where the data resides?
- How & by when?
 - To recover meeting the agreed SLA's
 - To recover in a DR scenario
 - Implement DR scenario testing
- Where?
 - Are the backup sets kept for redundancy and for how long to retain source set and its duplicates?
 - Are the documentation of BE configuration to be kept?

INDEX

Topic	p.
Hardware recommendation & tuning	6
Network recommendation & tuning	9
BE Server Hardware recommendation	14
Keep an Eye on your Environment	15
BE in an optimized environment	16
Things to keep in mind	17
Virtualization	21
Optimised Duplication	31
High Availability	38
Tuning Storage	42
Active Directory GRT Capabilities	47

Topic	p.
Instant GRT	49
Instant Recovery	52
Cloud	56
Tuning Backup Exec	61
Tuning VMware Backups	67
Backing up Database Servers	70
Backing up Oracle	73
Backing up Backup Exec	75
Where to find more information	77
End of Support Life Dates	78
Contact	79

Backup Exec Server Hardware Recommendations

Good to know

Servers consist of several parts but we will focus on four main parts:

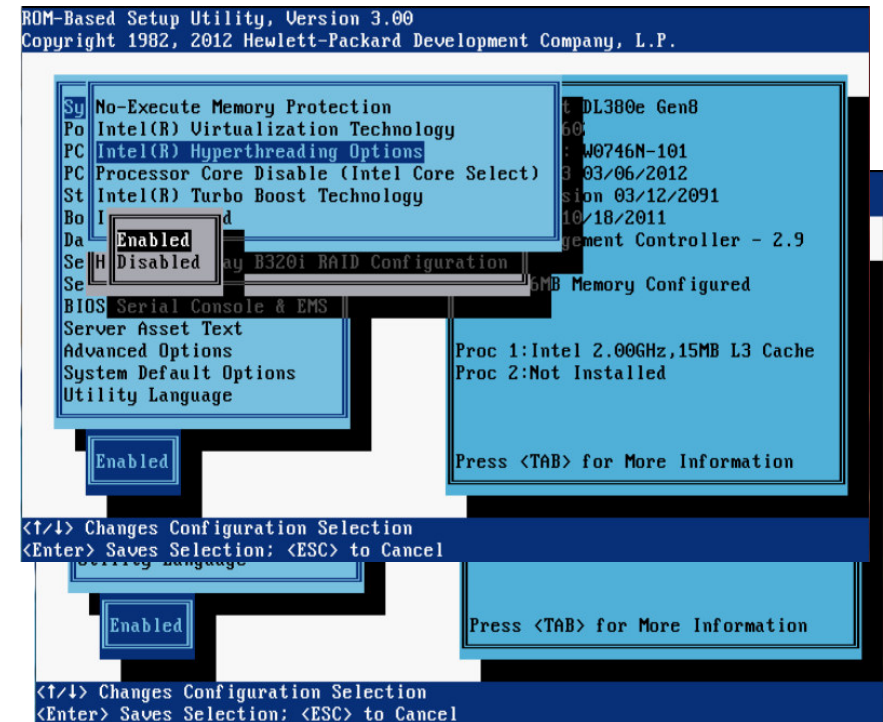
- CPU
- Memory
- Network
- Storage

The following slides will give a more detailed insight to these components and recommendations on hardware usage for Backup Exec servers.

Backup Server CPU Recommendations

Good to know

- As of today, the Backup Exec deduplication process is not capable of using multiple cores.
- Therefore, you should better use CPUs with fewer but **faster cores** than ones with more but slower cores.
- Turn off Hyper Threading in BIOS (to speed up the cores).



Backup Server Memory Recommendations

Good to know

- Memory is critical for Backup Exec.
- For deduplication to work you can calculate the necessary **free** memory (not used by OS or other applications) as follows:
 - 1,5 GB RAM per TB deduplication storage
 - Add another 8 GB for the operating system
- Example: A backup server has 20 TB of deduplication storage. Therefore it will need $20 \times 1,5 + 8 = 38$ GB RAM.
- Since the size of a deduplication storage in Backup Exec is limited to 64 TB, the maximum RAM needed for Backup Exec servers is $64 \times 1,5 + 8 = 104$ GB.

Backup Server Network Recommendations I

Good to know

- Quite often, the backup server's network connection is the bottleneck in the infrastructure.
- Especially, when doing backups directly to tape, keep the following table in mind:

LAN speed \ Tape speed		LTO 4 (190 MB/s)	LTO 5 (230 MB/s)	LTO 6 (320 MB/s)	LTO 7 (600 MB/s)
100 Mbit	(10,5 MB/s)	5,53%	4,57%	3,28%	1,75%
1 Gbit	(117 MB/s)	61,58%	50,87%	36,56%	19,5%
2 Gbit (Team)	(205 MB/s)	107,89%	89,13%	64,06%	34,17%
4 Gbit (Team)	(355 MB/s)	186,84%	154,34%	110,94%	59,17%
10 Gbit	(1000 MB/s)	526,32%	434,78%	312,50%	166,67%

Network Analyses

https://support.symantec.com/en_US/article.TECH87209.html

Troubleshoot Network

<https://support.microsoft.com/en-us/kb/2643970?wa=wsignin1.0>

Backup Server Network Recommendations II

Good to know

Time to copy 1 TB of data (Hours:Minutes)			
LAN speed	Theoretical	Realistic via LAN	Realistic via SAN (iSCSI/FC)
100 Mbit (10,5 MB/s)	27:45	330:00 (~14 days)	./.
1 Gbit (117 MB/s)	02:30	30:00	06:25
2 Gbit (Team) (205 MB/s)	01:25	17:00	03:40
4 Gbit (Team) (355 MB/s)	00:48	09:45	02:10
10 Gbit (1000 MB/s)	00:20	3:30	00:45

Tuning Network

Good to know

- Use a dedicated network for backups, if possible.
- Verify correct name resolution or use hosts file to force Backup Exec to use the “private” network link.
- Use NIC teaming, but don’t overestimate the performance of teamed network links:
 - 1 Gbit-Link will usually transport around 71 MB/Sec.
 - A 2-Gbit team will usually transport around 124 MB/Sec.
 - A 4-Gbit team will usually transport around 215 MB/Sec.
 - (Four NICs teamed will have around 25% loss of performance due to the teaming process.)
- Set fixed network speed and duplex on all servers as well as on the switches.
- For Windows Server 2008/R2: turn off offloading
Chimney Offload and Receive Side Scaling: <http://support.microsoft.com/KB/951037>
 - `netsh int tcp set global chimney=disabled`
 - `netsh int tcp set global rss=disabled`
 - `HKLM\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\EnableTCPA = 0`

Backup Server Disk Recommendations

Good to know

- In order to get significant performance improvements, increase the number of disks in the in the volume. Tests have shown that doubling the amount of disks in a RAID 5 volume (from 6 to 12) **tripled** the overall **performance!**
- So better **use many small disks** than a few larger ones.
 - Note: If number of disks are in the "ten's" you need to think about using RAID 6 instead as the risk of simultaneous double disk failure will increase by the number of disks used.
- Best for performance and redundancy is RAID 1 for the system volume (OS) - (disk type: SSD).
- Best for performance and redundancy is RAID 5 for the storage volume (disk type: SAS).
 - Note: The rotation speed of the drives in the storage volume has nearly no impact on the performance.

Usage	Disk-Type	RAID-Level
System Volume - OS	SSD	1
Storage Volume - Data	SAS	5 or 6

Backup Server Storage Controller Recommendations

Good to know

- Use storage controllers with as much write cache, as possible (≥ 2 GB).

More cache leads to better performance.

HP Smart Array Controller Technology:

<http://www8.hp.com/h20195/v2/GetDocument.aspx?docname=4AA5-4124ENW>

- Don't use RAID controllers without buffered write cache (flash module or battery).

Due to safety reasons, most RAID controllers use their cache modules only for read operations, when no battery or flash back module is present, because without such a module, a power loss to the server would almost sure result in an unpredictable data loss.

- Never use RAID controllers to connect to tape drives, always use dedicated HBA's.
- Always ensure that the firmware of the storage controllers is up to date.

Backup Exec Server Hardware suggestions

Small – up to 5 TB

- CPU ≥ 2.6 GHz, ≥ 6 Cores
- System-Volume:
2x SSD ≥ 200 GB (RAID 1)
- Data-Volume:
6x 1TB SATA HDD (RAID 5)
- Memory: 16 GB
- 4x NIC 1 Gbit
- 1x Power Supply
- Optional:
1x SAS for tape device

Medium – up to 7 TB

- CPU ≥ 2.6 GHz, ≥ 6 Cores
- System-Volume:
2x SSD ≥ 200 GB (RAID 1)
- Data-Volume:
8x 1TB SAS HDD (RAID 5)
- Memory: 24 GB
- 4x NIC 1 Gbit
- 1x Power Supply
- 1x SAS/FC for tape device

Large – up to 22 TB

- CPU ≥ 3.0 GHz, ≥ 6 Cores
- System-Volume:
2x SSD ≥ 200 GB (RAID 1)
- Data-Volume:
24x 1TB SAS HDD (RAID 6)
- Memory: 48 GB
- 2x NIC 10GBE
- 2x Power Supply
- 1x SAS/FC for tape device
- 1-2x SAS/FC for SAN Connect

Each type can be scaled out



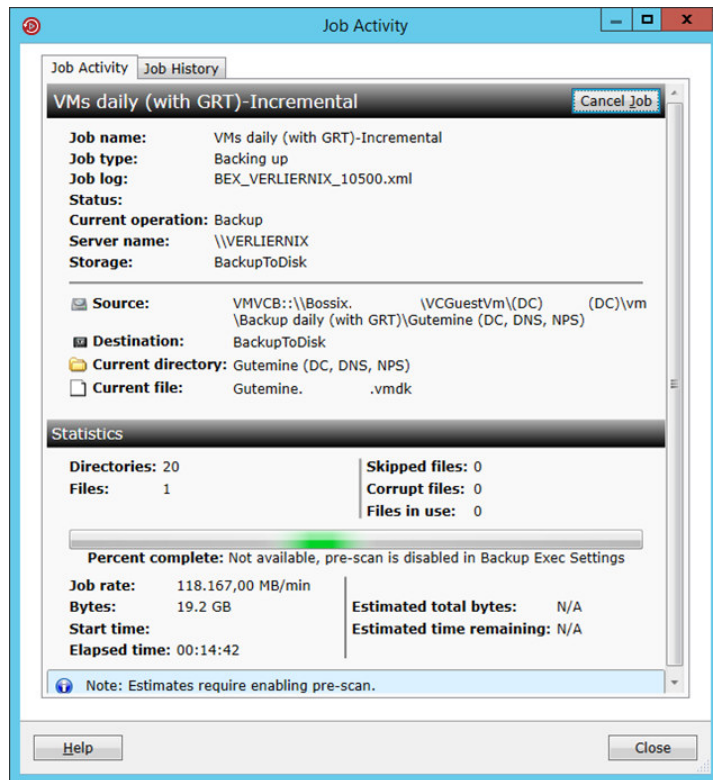
Keep an Eye on the Environment

Network

- Network Performance (Latency) Considerations
 - Ensure less than one percent (1%) packet loss during transmission.
 - Ensure a destination “round trip” latency of 250ms or less.
 - Connection problems will impact job success rates.
- Latency and bandwidth isn’t the same
 - You can have a broad network line to the internet, but its latency will still lead to “slow” backup rates.
- Make sure to test the usability of 3rd party WAN optimizer before using them for backup, restore or client site deduplication, because these aren’t tested configurations and only supported as an alternative configuration.

Backup Exec in an optimized environment

Backup Exec can be pretty fast, if configured correctly



Hardware for BE Server:

- HP DL380e Gen8
- 1x 8-Core Xeon
- 64 GB RAM
- Storage Controller: SmartArray p822 (2GB FBWC)
- Storage Controller for SAN connect: 12 Gb SAS
- 2x 200 GB SSD for OS
- 12x 1 TB HDD for B2D
- 12x 1 TB HDD for Dedup
- 2x 10 GbE NICs (teamed)

Storage for VM:

- 1x HP MSA 2040
- 24x600 GB HDD

Note:

To achieve optimum performance, not only the Backup Exec Server, but also the underlying infrastructure have to be configured, optimized and tested, as every environment is unique.

Things to keep in mind I

Not all 10 GBE ports are the same

- In order to use multiple 10 GbE ports on a switch, the switch must not only offer these ports, but also have the backplane capacity to transport the traffic arriving at the ports.

Example: Cisco Catalyst 6500 can only transport 2x20 Gbit/sec per line card, so cards with 8x10 GBE are 100% overbooked, cards with 16x10 GBE 200%.

- Note:
Especially when talking about iSCSI,
we recommend using deep buffer switches.

Table 1. Cisco Catalyst 6500 Series 10 Gigabit Ethernet Modules Primary Features Comparison

Feature	4-Port 10GbE Fiber Module	8-Port 10GbE Fiber Module	16-Port 10GbE Fiber Module	16-Port 10GbE Copper Module
Ports	4	8	16	16
Optics	XENPAK	X2	X2	<ul style="list-style-type: none">No opticsCopper (RJ-45) connectors
Switch fabric connection	40 Gbps (80 Gbps full duplex)	40 Gbps (80 Gbps full duplex)	40 Gbps (80 Gbps full duplex)	40 Gbps (80 Gbps full duplex)
Oversubscription	1:1	2:1	4:1	4:1

Things to keep in mind II

Verify

- Backup Exec integrated verification does not free you from the need to
 - do random restore tests.
 - check, whether the files you restored are readable/useable.
- Verify on a deduplication volume means that all files from the backup set(s) have to be re-hydrated to get verified.
 - Example: You back up 1 TB of data and achieve a ratio from 5:1, then 200 GB of new data will be stored. During a verification run, the whole 1 TB of data have to be rehydrated and read.
- Verify on Cloud based storage means that all data have to be read which will incur an extra cost = \$ (same as a restore).
- Verify does **NOT** check the **CONTENT** of a file, only its **CHECKSUM**.
- For 3 stage jobs where last stage is tape, we recommend to use verify only on this last stage.

Things to keep in mind III

Deduplication

- Not all data can be deduplicated well
 - i.e. Exchange log files contain >90% of unique data. Therefore, if the overhead of doing the deduplication is a concern, we recommend targeting incremental Exchange backups to backup-to-disk rather than to a deduplication storage.
 - Also the attempt to deduplicate media files (video, pictures and music) will not be very efficient.
- For GRT enabled backups, full and incremental backups need to use the same device type! E.g. You can split between deduplication storage and B2D for Exchange GRT enabled backups Full and incremental backups, but if the backups are not affected (is within the required time window) use the same storage!
- Backup to a deduplication storage is less performant than Backup-To-Disk
 - Changing from Backup-to-Disk-to-Tape to Backup-to-Dedup-to-Tape may in some case extend the backup window beyond what is practical for lesser performant Backup Exec servers
- Especially the attempt to read from a deduplication storage (i.e. for copying data to tape or verify) at the same time, where backup jobs are running will dramatically decrease the performance of the storage device.
- The deduplication storage logon account is not required to be any Windows user account. It can just be a logon account that **only exists** in Backup Exec. This is **recommended** as the account will **not be affected** by a Windows account change of password.

Things to keep in mind IV

Backup Exec account

- Understanding the Backup Exec accounts
 - The Backup Exec System account which will be in most cases also the default logon account.
 - i.e. when you create a job it will use by default this account as default to log on to the selected resources.
 - The Backup Exec Service account for those services that are not using the local system account. This account will in most cases be the same as the Backup Exec System Account.
 - The Backup Exec System account and the account used for the Backup Exec services should be the same.
 - Would recommend not to use the <domain>\Administrator as this account may have restrictions imposed in some applications that can prevent successful backups and restores.
 - Instead copy the Administrator account and name it something meaningful e.g. BEadmin
 - For reference https://www.veritas.com/support/en_US/article.TECH130255

Backup Exec Server – used Ports

Firewall Ports

- By default, Backup Exec uses TCP 10000 for communications between the backup server and the agents installed on source machines.
- If TCP 10000 is already in use by another application, this can be changed per server.
- To do so, following these steps:
 - Stop the Backup Exec Remote Agent service on the source server.
 - Open the file C:\WINDOWS\SYSTEM32\DRIVERS\ETC\SERVICES with administrative permissions.
 - Add the following line at the bottom of the file:
NDMP XXXXXX/TCP #BACKUP EXEC
Where “xxxxx” is the port number you prefer to use, i.e. 10001.
 - Restart the Backup Exec Remote Agent service.
- For verification, whether TCP 10000 is in use, use the following command:
netstat -abn | findstr :10000

Virtualization I

VMware and Hyper-V support different technologies to do backups and restores.

	VMware	Hyper-V	Agent required
Backup via SAN	Yes	No	No
Restore via SAN	Yes	No	No
Full backup	Yes	Yes	No
Diff./ incr. backup	Yes	Yes	No
GRT for files	Yes	Yes	No
GRT for applications	Yes	Yes	Yes

Virtualization II

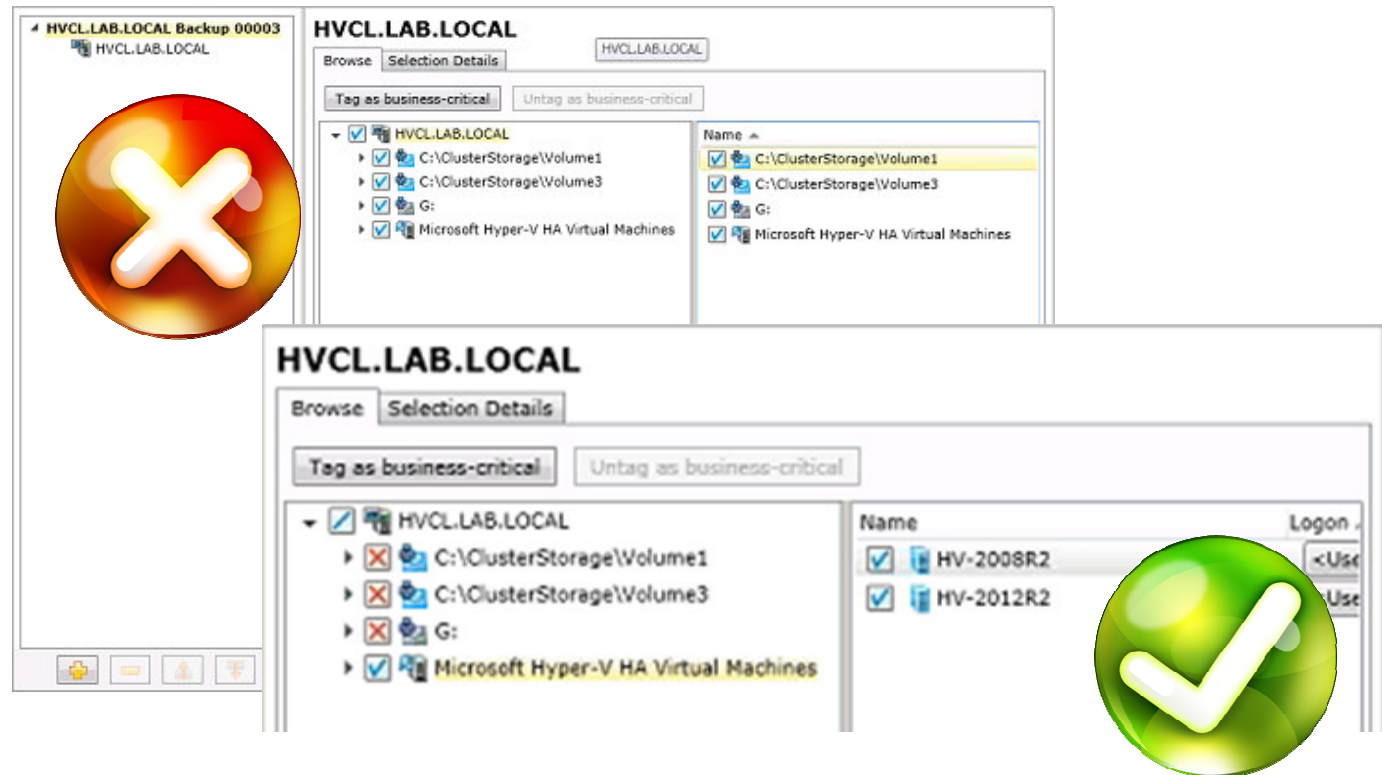
Hyper-V Best Practices

- If Hyper-V Integration Services aren't running, Backup Exec has to put the VM into a saved state to do the backup.
Note: This means, the VM is out of service.
Requirements for using the Agent for Microsoft Hyper-V: <http://www.veritas.com/docs/000058755>
- Install Backup Exec Remote Agent on the VMs you want to backup with GRT for applications.
- Understand that the backup job rate may be slower, when the VM backed up that is on a CSV not owned by the cluster owner.
VERITAS Article 000100805: <http://www.veritas.com/docs/000100805>

Virtualization III

Hyper-V Best – Protecting Cluster Shared Volumes

Do **NOT** select CSVs and VMs in the same selection list, as this will lead to VSS errors.



Virtualization IV

Hyper-V Best Practices – Incremental backups

If you are using the “faster processing method” for the Incremental backups, you will see a better backup performance, but there is more disk space on the Hyper-V storage volumes used, in order to create the snapshot tree required for this method.


Both backup methods have no impact on the full backups.


Hyper-V


Incremental Backup Settings

Choose how you want Backup Exec to process incremental backups:

- ☒ Use the faster processing method
- ☐ Use the standard processing method

 This method requires additional storage disk space on the Hyper-V host as one checkpoint for each virtual machine backed up is always present on the host even after the backup job is complete. This checkpoint is required to track the changes between the current and the next backup job.

 Next backup job will run as a full backup.

 Differential Backups are not supported with the faster processing method. With the faster processing method selected, all differential backups for Hyper-V run as incremental backups.

Virtualization V

Hyper-V Scenario FAQs

- **Q:** Is it possible and supported to install BE on a standalone Hyper-V host and back up the host itself, as well, as the hosted VMs?
A: Yes, this is technically supported and working fine, as long, as there's no additional software running on the host, like Exchange Server.
Please review the licensing guides from Microsoft and Veritas regarding the needed licenses.
- **Q:** Is it possible and supported to install BE on a clustered Hyper-V host and backup the host itself, as well, as the hosted VMs?
A: No, this is not supported.
- **Q:** Can I install the backup server as a virtual machine and backup other VMs?
A: Yes, this is possible.

Virtualization VI

Hyper-V Scenario FAQs

- **Q:** Can a tape storage be attached to a virtualized Backup Exec Server?
A: Yes, connect them by iSCSI, instead of using SCSI pass-through devices attached to the host.
(SCSI pass-through is only supported as an alternative configuration).
- **Q:** Is it possible and supported to back up Hyper-V virtual machines directly from the storage (SAN)?
A: No, this is currently not supported.

Virtualization VII

VMware Best Practices

- Keep VMware Tools up to date, as they are a necessary prerequisite for backing up VMs via the host.
- If you want to install the Backup Exec Remote Agent, do not install the VSS Provider included in the VMware tools, as it will conflict with Backup Exec.
VERITAS Article 00009506: <http://www.veritas.com/docs/000009506>

During the initial install of the Remote Agent, Backup Exec will remove the VMware VSS provider.
VERITAS Article 000042539: <http://www.veritas.com/docs/000042539>

- Install the Backup Exec Remote Agent on all VMs that you want to backup with GRT for applications (File level GRT does **not** require the Remote Agent to be installed).
- Present the SAN LUNs containing the VMs to the backup server.
How to backup ESX(i) using SAN transport: <http://www.veritas.com/docs/000012638>
- SAN transport is not recommended for thin provisioned disk restore, since it is probably slower than NBD.

Virtualization VIII - VMware Hotadd

Transport mode

- If Backup Exec server is virtualized, Hotadd, is the most performant backup method to use.
- Hotadd is not available for restores, have to use NBD/NBDSSL.
- The Backup Exec disk storage shouldn't use the same datastore as the one that hosts the VM's you are backing up.
- If using tape devices with a virtualized Backup Exec Server, connect them by iSCSI instead of using SCSI pass-through devices attached to the host. (SCSI pass-through is only supported as an alternative configuration).

Transport mode priority list:

<input checked="" type="checkbox"/>	SAN - Use the SAN to move virtual disk data
<input checked="" type="checkbox"/>	Hotadd - Use virtual disk files from the Backup Exec server on the virtual machine
<input checked="" type="checkbox"/>	NBD - Do not encrypt the virtual disk data for over-the-network transfers
<input checked="" type="checkbox"/>	NBDSSL - Encrypt virtual disk data for over-the-network transfers

Virtualization IX

Good to know

- Don't backup to the same storage hardware as the VM's are using.
- Don't use virtualized disks as backup targets, use physical disks presented by SAS/FC/iSCSI.
- Every Backup Exec configuration that is supported in a traditional physical environment would also be supported in any virtual environment without qualification.
VERITAS Article 000080758: <http://www.veritas.com/docs/000080758>

Virtualization X - Keep an Eye on the Environment

Storage Location

If possible, backup (VMware) virtual environments via SAN, not LAN.

Storage Location	Backup via NBD	Backup via SAN Transport
Local (non clustered)	Yes	No
NFS Shares	Yes	No
iSCSI LUNs (vmfs)	Yes	Yes
Fiber LUNs (vmfs)	Yes	Yes

- Real World Example: ESX connected to storage via 10 GbE
- Backup of Test-VM (40 GB) from NFS share: **4.600 MB/min**
- Backup of Test-VM (40 GB) from iSCSI LUN: **24.200 MB/min**

Optimised Duplication I

Good to know

- The term “Optimised Duplication” describes the technology Backup Exec uses to replicate data between two deduplication storages connected to two different backup servers in a way that only those blocks are replicated that do not already exist in the target storage.
- To simplify, this means that one backup server does a backup to a locally attached deduplication store and afterwards performs a second stage in the backup policy to copy (duplicate) the data to a second deduplication store that is hosted on a different backup server. The effect that is achieved by only sending over the network the blocks that the receiving deduplication store does not have is comparable to what a client-side duplication backup job does.
- In order for this to work, the environment must have a Enterprise Server Option (“ESO”) licensed and a Central Admin Server (CAS) installed. The CAS may be one of the servers involved in the duplication job, but it doesn’t have to as it can also be from one MBES to another MBES.
- To use a dedicated network just for the Optimised Duplicate job between two Backup Exec servers, create two entries in the HOSTS file of each server using the NetBIOS (only) names of both BE servers and the static IP address’s for the network that is desired to use.


Optimised Duplication II

Optimised Duplication overview

1. The CAS performs a backup job to “it’s” deduplication storage.
2. The CAS connects to the MBES to request access to the deduplication storage connected to the CAS.
3. The CAS copies the chunks to the second deduplication storage that do not exist there and are needed for restores.

Note:

In this example the role of the CAS can either be installed on any of the MBES's or on a 3rd machine.



Deduplication with OST Appliances I

Good to know

- „Open Storage Technology (OST)“ is a software interface that allows Backup Exec to manage external storages in terms of tracking what's happening inside.
- In other words, the deduplication process itself is outsourced to the hardware appliance, while Backup Exec still gets all information and catalogs needed to do restores.
- This outsourcing has many advantages but also disadvantages (more next slides).

Deduplication with OST Appliances II

Good to know

The bright side:

- OST provides
 - a hardware-based deduplication engine that removes the load from the backup server.
 - application-aware replication.
- This leverages the hardware replication capabilities to copy data between sites, but manages the process through policies set in Backup Exec.
- Because of the tight integration, Backup Exec's access to the replicated data is transparent and allows direct restores from all replicates, wherever they are located.
- Since OST appliances have their own storage management, they are not limited to the maximum size of 64 TB, as Backup Exec's internal deduplication storages are as of today.
- Because of their specialized hardware and firmware, OST appliances are often better performing than the (Windows-based) internal deduplication storages.

Deduplication with OST Appliances III

Good to know

The dark side:

- Restores from GRT enabled backups need to be staged during the restore process, similar to restoring from tape.
- Therefore an appropriate amount of local storage in the backup server needs to be available during restores (approx. 1.2 - 1.5 times the size of the largest container backed up [vmdx, edb etc.])
- Replication between OST appliances is only supported, if both appliances are from the same vendor and use the same OST plug-in.
- The availability and compatibility of client-side deduplication is limited to the capabilities of the OST plug-in provided by the storage vendor.
- The cost of the OST Appliances itself.

Deduplication with OST Appliances IV

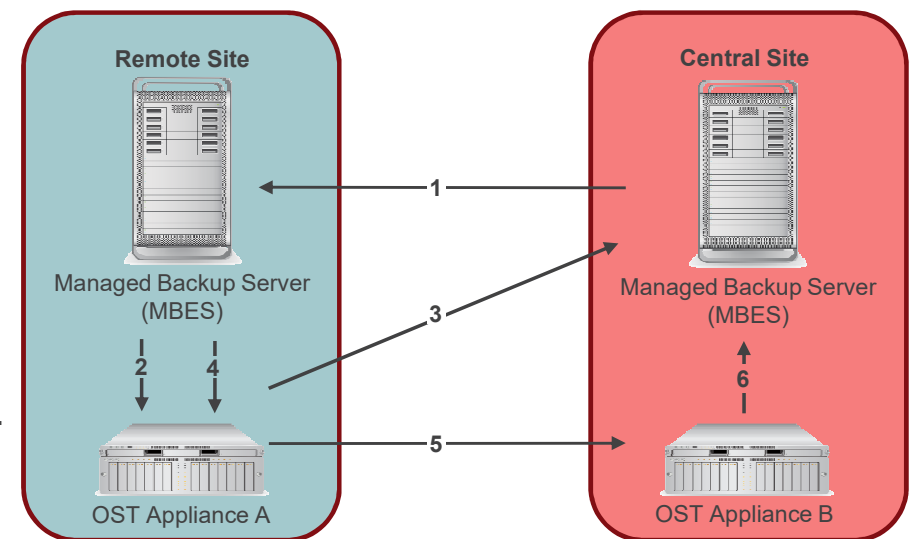
Good to know

- Replication, or “Optimised Duplication” as Veritas calls it, is **MANAGED** by Backup Exec but **PERFORMED** at the storage level.
- Therefore, the OST Appliances must be able to connect to each other by their replication protocol (For details see next slide).
- Please refer to the vendors’ documentation on which firewall port(s) to open.

Deduplication with OST Appliances V

Optimised Duplication in Detail:

1. The CAS instructs the MBES to perform a backup job to the OST appliance in the remote site and to duplicate the data to the OST appliance in the central site afterwards.
2. The MBES writes the data to “his” OST appliance.
3. The OST appliance sends the information needed for cataloging to the CAS.
4. After the backup job is done, the MBES instructs the OST appliance in the remote site to...
5. ...replicate the job’s data to the appliance in the central site.
6. The OST appliance in the central site sends the data information needed for cataloging to the CAS so that restores can be performed from both OST appliances.



High Availability I

Good to know

- Many applications and databases can be deployed either as a standalone installation or as a highly available instance that is spread over multiple physical or virtual servers.
- When backing up databases and applications on standalone servers, you can decide whether you want to back up the resources via the virtualization host or via an agent installed inside the virtual application server.
- Deploying applications highly available means that the application runs as a sort of virtual instance that can be hosted by any one of the members of the clustered system.
- Since Backup Exec has to bring the application into a consistent state in order to be backed up, it must “talk” to the virtual instance instead of one of the cluster nodes.
- Those highly available instances will be displayed in Backup Exec as a dedicated resource entry in the BACKUP AND RESTORE pane.


High Availability II

Good to know


- To fully protect a database or application cluster, it is necessary to back up
 - The local drives of all cluster nodes to restore them in case of a local failure (i.e. an operating system error).
 - The application in a separate job using its own resource record.
- Example: To protect a two-node Exchange-DAG, you will have to back up three different objects:
 - Node A (local hard drive(s) and system state).
 - Node B (local hard drive(s) and system state).
 - The clustered resource itself (DAG) (which will include the information store).
- Therefore, even if the highly available application instance is running on virtual servers, it cannot be backed up using the Agent for VMware and Hyper-V, as this would process the VMs one by one, without protecting the application at all.

High Availability III

Exchange DAG

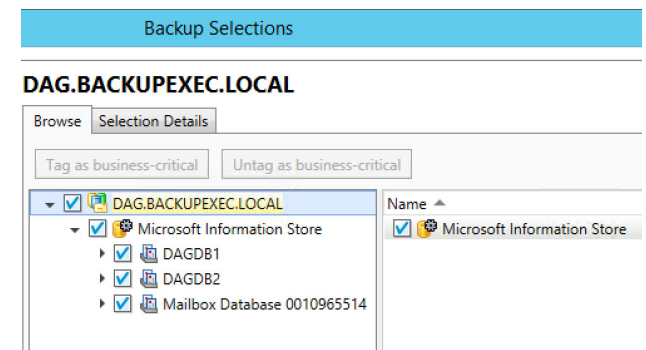


Installing RA on **two** DAG nodes automatically displays **three** entries in Backup Exec GUI.



Only backups of the cluster nodes are usable for **SDR**.

Database files are **automatically** excluded.



Only the selection list of the **DAG** allows to select the **Exchange** resources
The DAG node can also be added manually, if needed.

High Availability IV


SQL Cluster

SQL Cluster


Server	Active Alerts	Status	Last 7 Days of Backup Jobs
SQL-DESAL-000. SQL Cluster DTC		Backed up	Sa Su Mo Tu We Th Fr
SQL-DESAL-001. SQL Cluster		Backed up	Sa Su Mo Tu We Th Fr
SQL-DESAL-002. SQL Cluster Instance2		Backed up	Sa Su Mo Tu We Th Fr
SQL-DESAL-003. SQL Cluster Instance3		Backed up	Sa Su Mo Tu We Th Fr
SQL-DESAL-004. SQL Cluster Instance4		Backed up	Sa Su Mo Tu We Th Fr
SQL-DESAL-005. SQL Cluster Instance5		Backed up	Sa Su Mo Tu We Th Fr
SQL-DESAL-006. SQL Cluster Instance6		Backed up	Sa Su Mo Tu We Th Fr
SQL-DESAL-007. SQL Cluster Instance7		Backed up	Sa Su Mo Tu We Th Fr
SQL-DESAL-CM1. SQL Cluster Node		Backed up	Sa Su Mo Tu We Th Fr
SQL-DESAL-CM2. SQL Cluster Node		Backed up	Sa Su Mo Tu We Th Fr



Installing RA on **two** Cluster nodes automatically displays **multiple** (here:9) entries in Backup Exec GUI.



Only backups of the cluster **nodes** are usable for **SDR**.



Only the selection list of the **clustered instance** allows to select the **SQL** resources.

Tuning Backup-To-Disk Storages

Good to know

- Create all arrays with a stripe size of 64 kB.
- Format all volumes in Windows with a block size of 64 kB.
- If possible, attach multiple arrays to different controllers or buses on the backup server.
- Don't use software RAIDs to combine multiple arrays into one larger one; rather use storage pools inside Backup Exec.
- Use **local volumes** for **B2D** over UNC paths (Windows fileshares or NAS CIFS shares), especially when using **GRT** enabled backups.

Tuning Deduplication Storage

Good to know

- Create all arrays with a stripe size of 64kB.
- Format all volumes in Windows with a block size of 64kB.
- Don't underestimate the size of the deduplication database (PostgreSQL)
May reach up to 8% of the size of the deduplication storage.
- Assure that your deduplication storage is not running out of disk space, as this can result in a damaged deduplication storage.
We recommend creating a dummy file on the volume that can be deleted when disk space is running low.
 - Use the following command to create a 100 GB file:
`#fsutil file createnew <drive letter>:\DeleteOnlyIfDedupSpaceLow.txt 104857600000`
 - When volume for deduplication storage was extended, the services need to be restarted in order to recognize the change.

Cleaning up Deduplication Storage

Good to know

- Backup Exec automatically deletes expired media/backup sets and makes the space available to the Deduplication Folder.
- This automated process takes place every twelve hours (12:20 & 00:20). The log file SCHED_QUEUEPROCESS.LOG (located in DeduplicationFolderPath\Log\Spad) records when queue processing gets initiated.
- The process can also, when needed, be initiated manually.

Please find detailed instructions here:

- Manual Space Reclamation for Deduplication Storages in Backup Exec
VERITAS Article 000017049: <http://www.veritas.com/docs/000017049>
- Please note: The reclamation process needs some temporary space to run. Therefore, if the deduplication storage ran out of disk space, the cleanup process cannot run.

Disk performance

Good to know

- To find out what the base disk performance is on volumes used for backup disk based devices, you can use the standalone tool NBPerfchck.exe.
It can also be used to find out, what the network throughput is to a disk volume from a remote disk volume.
- Keep the test results as part of your deployment documentation and use them as a baseline in case of issues.
- Veritas recommends a minimum disk performance level of 130 MB/sec for B2D and Deduplication read and write operations. Ideally everything above 200 MB/sec is good.
- VERITAS Article 000095782: <http://www.veritas.com/docs/000095782>

Windows OS versions and AD GRT Capabilities

Good to know

- To perform a GRT-enabled backup of a Windows Server Active Directory Server, you **must** use a Backup Exec server that uses the same or a newer version of Windows Server OS as then the Active Directory Domain Controller(s).
- After restoring Active Directory user objects, you must assign them new passwords and then re-enable the user accounts.
- Active Directory GRT restores do not require the Domain Controller to run in Directory Services Restore Mode (DSRM), instead they can be run online.

Traditional Active Directory recovery process

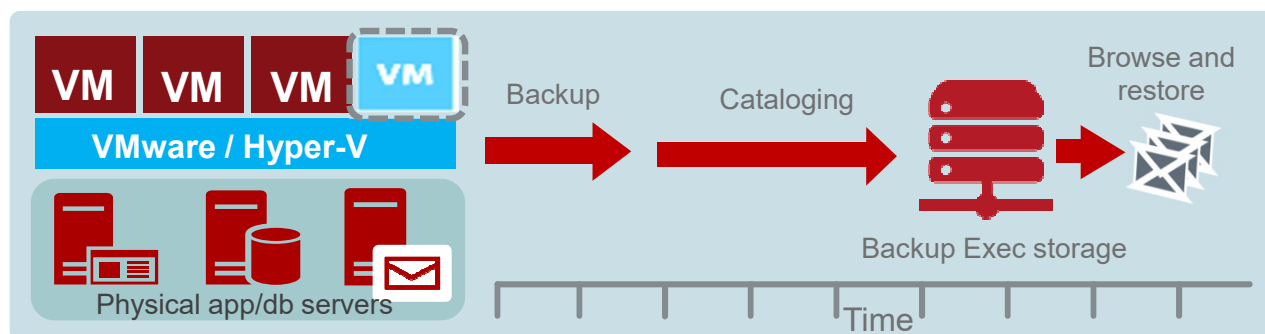
Good to know

Non-authoritative restore (Repair defective Domain Controller)	Authoritative restore (Reset Active Directory to a previous state)
Restart the server in Directory Services Restore Mode (DSRM)	Restart the server in Directory Services Restore Mode (DSRM)
Restore System State from Backup	Restore System State from Backup
Reboot the server	Remove the network cable
Inbound replication occurs	Reboot the server
	Use NTDSUTIL to elevate objects to recover
	Attach the server to the network
	Outbound and inbound replications occur

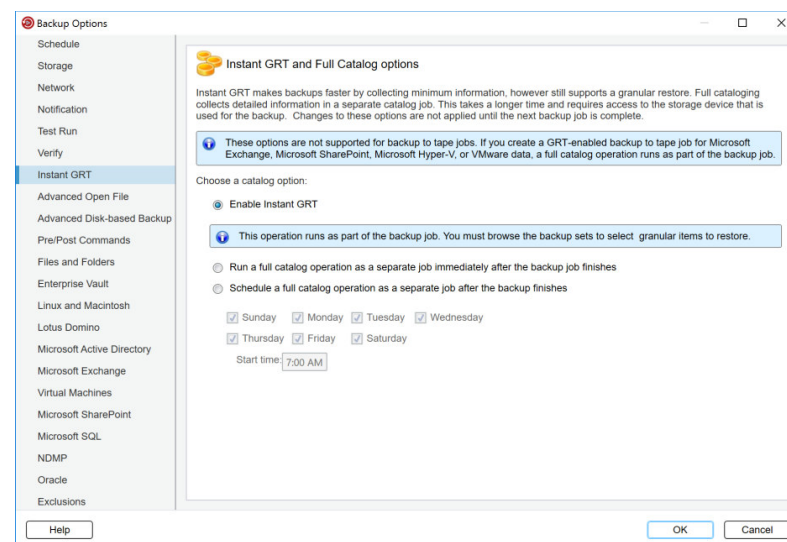

Instant GRT

Instant GRT reduces the overall time the backup job takes to complete by collecting minimum catalog information, however still supports a granular item restore.

Traditional Backup



Instant GRT Backup (physical & virtual)



Instant GRT

Good to know

- Upgrades do not change existing default and settings.
- If a new device type (for example deduplication) is created, the default will be Instant GRT.
- The Instant GRT operation is applicable for GRT-enabled backups for Active Directory, Exchange, SharePoint, Hyper-V, or VMware data (physical and virtual).
- Instant GRT was always used for agent based Active Directory backups.
- Instant GRT isn't applicable for backup to tape jobs.
If you create a GRT-enabled to tape job for Exchange, SharePoint, Hyper-V or VMware data, a full catalog operation runs as part of the backup job.
- All Instant GRT backup data sets have to be on the same disk storage type.
If you use a RDX device, ensure that the complete backup set (full and incremental) is on "one" cartridge.
If not practicable then backup to a B2D or deduplication storage device.
- Full cataloging option is still available (immediately after the job or scheduled).

Instant GRT vs. Full Catalog

	Instant GRT (Quick Catalog)	Full Catalog (Immediate or Scheduled)
Search wizard support for GRT restores	Not applicable	Can search catalogs from GRT backups for GRT restores
Backup set browse	Instant GRT by dynamically browsing backup sets from GRT backups for GRT restores	Browse catalogs from GRT backups for GRT restores
Catalog job	Cataloging is completed as part of the backup job with no delay	Separate catalog job can be configured to run immediately after backup job or scheduled
Catalog information collected	Only the minimum required catalog data is collected (faster) *	Detailed catalog data is collected (slower)
Catalog data	Smaller; contains minimum required catalog data	Larger; contains detailed catalog data
Time to browse for GRT restore	GRT data in the backup set is read when backup sets are expanded (slower)	GRT data is already available from detailed catalog data (faster)

Instant Recovery

Key Features

- Powers on the virtual machine (regardless of it's size) directly from the backup storage.
- Makes the VM visible in vCenter or Hyper-V manager as soon as it's recovered.
- Recovered VM can be migrated online to production by either VMware vMotion or Hyper-V Live Migration.

Instant Recovery

Requirements

First of all, the requirements depend on whether you are using it with Hyper-V or VMware:

Microsoft Hyper-V:

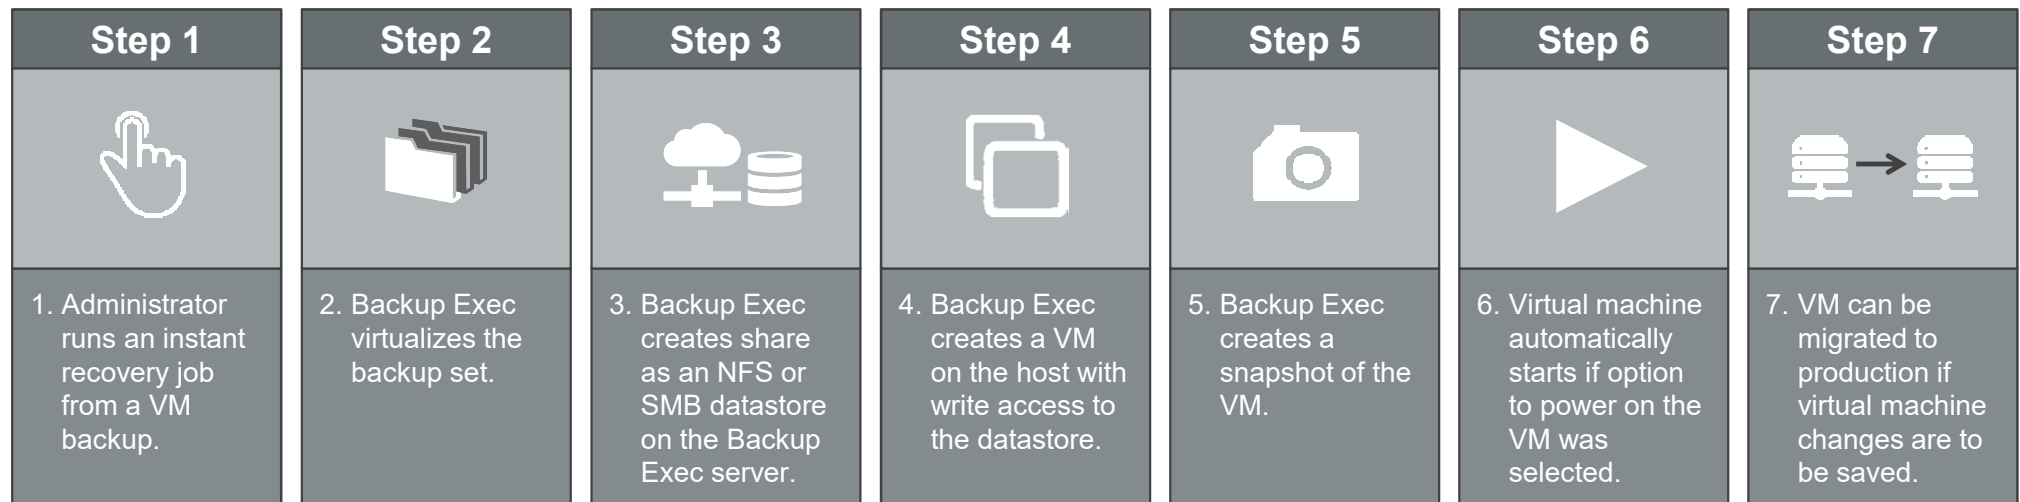
- Backup Exec server on Windows Server 2008 SP2 or later.

VMware vSphere:

- Backup Exec server on Windows Server 2012 or later.
- Server for NFS (Windows Role) installed on Backup Exec server.
- Ports TCP/UDP 111 and 2049 must be open between ESX and Backup Exec for NFS traffic.

Instant Recovery

Instant Recovery Workflow













- Note: Veritas recommends that you remove an instantly “Recovered VM” that is no longer needed in Backup Exec after the VM has been successfully migrated.
The backup set of the virtual machine will remain untouched on the Backup Exec storage and can be reused.
- The additional load on the Backup Exec storage may impact the overall performance of the Backup Exec server.

Instant Recovery

Standard virtual machine recovery vs Instant Recovery

Standard vs. Instant Recovery Comparison

Different data and application restore challenges require different solutions

Standard Recovery Process	 Restore time is dependent on the size of the VM and the network and storage speed.	Instant Recovery Process	 Instant recovery takes less time and jobs are not dependent on virtual machine size.
	 Transfers all data from the backup set to the host.		 Does not transfer virtual machine data to the instantly recovered VM.
	 Data is moved to the virtual (VMware/Hyper-V) host.		 Uses backup set image for all read operations and a snapshot on the virtual host for all write operations.
	 Uses virtual (VMware/Hyper-V) host storage.		 Uses Backup Exec server storage until you migrate the instantly recovered VM.
	 The Backup Exec server and the virtual host can be restarted.		 The Backup Exec server and the virtual host cannot be restarted (BE Services can be restarted). Instantly recovered VM can be restarted.

Cloud Connectors

The Backup Exec Cloud Connector feature provides seamless and secure integration with 3rd-party cloud storage services, which enables direct-to-cloud backups and disk-to-cloud backups.



Common Cloud Terminology

Good to know

- **Access key ID** - An alphanumeric code that allows access to the cloud storage.
- **Bucket/Container/etc.** - A logical unit of storage that stores objects, such as data and metadata.
- **Secret key** - An alphanumeric code that allows access to the cloud storage.
- VERITAS Article 000108140: <http://www.veritas.com/docs/000108140>

Pros and Cons of Common Cloud-based Backup Scenarios

Good to know

Backup Scenario	Pros	Cons
Back up directly to the cloud	This is the simplest operation. It does not require any additional space on the local backup storage device.	<p>Since cloud backups and restores may be slow depending on available bandwidth, they may not fit in your backup window.</p> <p>A menu option to back up directly to the cloud is not available. To back up directly to the cloud, you must create a backup-to-disk job, and then edit the storage properties to select the cloud storage device.</p>
Back up to disk, and then duplicate to the cloud	<p>This operation provides a quick restore from a local copy.</p> <p>This is usually faster than backing up directly to the cloud, so your backup window may not be impacted.</p> <p>A menu option is available for this operation, so configuring this type of operation is easier than backing up directly to the cloud.</p>	This option requires additional disk space on the local backup storage device.
Back up to a deduplication storage device, and then duplicate to the cloud	<p>This operation provides a quick restore from a local copy.</p> <p>In addition, it reduces the amount of disk space required for a local copy.</p>	The backup will no longer be deduplicated when it is copied to the cloud.

Cloud

Good to know

- GRT restore have to be staged to a local path on the Backup Exec Server (default: C:\Temp).
- Verify on Cloud based storage means that all data have to be read which will incur a cost \$ (same as a restore). Restores will cost money (check your service provider's latest pricing information).
- Backups and restores to and from the cloud will be slower then jobs using on premise devices.
- Backups and restores require an internet connection.
- Don't leave the only copy of your business critical data on a Cloud device.

Cloud

Good to know

- To back up to disk and then duplicate to the cloud, you must configure two types of storage in Backup Exec: a local disk storage device for staging purposes and the cloud storage.
- Create specific buckets to use exclusively with Backup Exec.
- Use a different bucket for each cloud storage device. Do not use the same bucket for multiple cloud storage devices, even if these devices are configured on different Backup Exec servers.
- Ensure that bucket names contain only lowercase letters, numbers, and dashes or hyphens. Also, ensure that bucket names do not begin with a dash.
Buckets are not available for use in Backup Exec, if the bucket name does not comply with the bucket naming conventions.

Tuning Backup Exec I

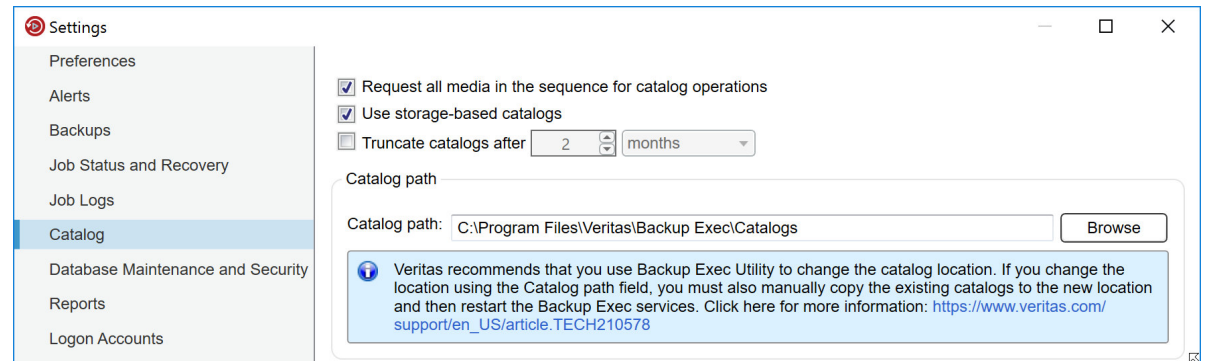
Move Catalogs

- Move the “Catalogs” folder to a performant volume (i.e. the Backup-To-Disk storage or at best another dedicated fast IO storage with adequate capacity).

Note: If you change the catalog location this way, you’ll have to move the files from the old location to the new one afterwards manually, as this is not done automatically.

Notes:

- 1TB front-end backups generate approx. 50 GB of catalogs.
- Moving the catalogs mitigates the following risks:
 - Running out of disk space on the system volume
 - In case of reinstall/upgrade the OS the catalog data is not affected.



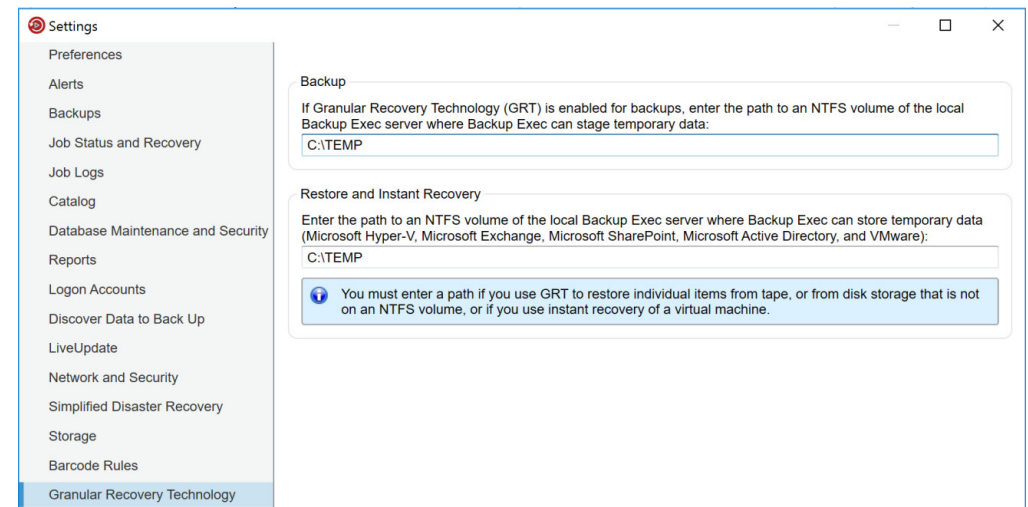
Tuning Backup Exec II

Move GRT-Temp

- GRT restores from a tape/OST/VTI/Cloud storage device need to be staged to a local disk path. This is the **Restore and Instant Recovery** path! The **Backup** path can in most cases stay as default but make sure it is excluded from anti virus and malware scans.
- Default staging location path is C:\Temp.
- Change the Restore and Instant Recovery path to a volume which has sufficient space to stage.

Notes:

- Multiply the total size of the backup sets you are restoring from by at least 1.25.
- e.g. recovering an Exchange DB of 500 GB from tape requires at least approx 625 GB.
- Moving the path mitigates the following risks:
 - Running out of disk space on the system volume.



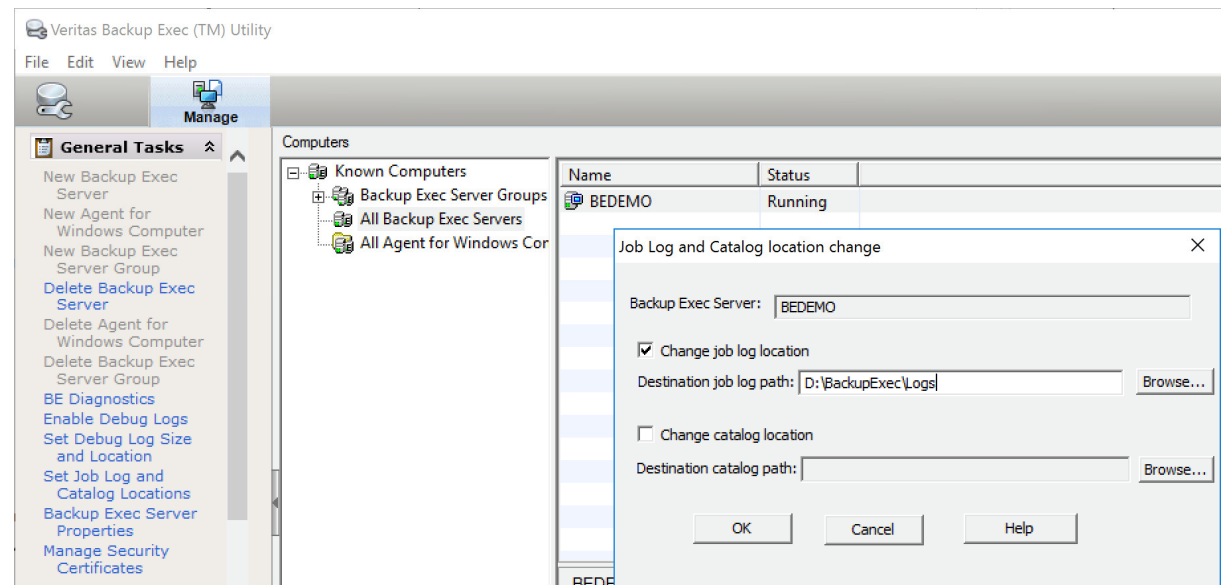
Tuning Backup Exec III

Move Job Log Files (using BE Utility)

Consider moving the “Job Logs” folder to a performant volume.

Notes:

- The target folder must exist, else the move will fail.
- Moving the job logs mitigates the following risks:
 - Running out of disk space on the system volume.
 - In case of reinstall/upgrade the OS the job logs data is not effected.



Tuning Backup Exec IV

Database Maintenance Settings

- Change the DB Maintenance time to a time before the backups starts in order to backup the configuration changes.
- Review how long you really require job logs (default setting in most case is sufficient).
- We recommend to enable “Perform database consistency check” to maintain health.
- Optimize database size:
To prevent “bloated” database tables keeping redundant information.

☒ Enable Backup Exec database maintenance

Last time maintenance was performed: 08.03.2016 19:00:00

Perform database maintenance daily at: 19:00:00

Delete aged data

☒ Delete aged data

Job History:

☒ Keep job history for data on media that have current overwrite protection periods

☐ Keep job history for specified number of days: 90

Number of days to keep data before deleting it from the Backup Exec database:

Job logs: 30

Alert history: 7

Reports: 14

Audit logs: 30

☒ Perform database consistency check

☒ Save contents of database to the Backup Exec data directory

☒ Optimize database size

Backup Exec Database Encryption Key Management

The database encryption key is required to migrate or recover the Backup Exec server.

Path:

☐ Remember the export path. By clicking this check box, you consent to let Backup Exec retain and display the export path during import operations.

Tuning Backup Exec V

Backup Exec Database Encryption Key

- Export the database encryption key.
- Define a path to where you want to export the key.
- You can tick the checkmark that Backup Exec keeps record of the export path, if you want.


Note: You must provide the key, whenever you want to import the database into a Backup Exec environment.

Backup Exec Database Encryption Key Management

The database encryption key is required to migrate or recover the Backup Exec server.

path:

☐ Remember the export path. By clicking this check box, you consent to let Backup Exec retain and display the export path during import operations.

 You must export the database encryption key to ensure that you can migrate or recover the Backup Exec server later. Export the key to:

- Either a physical volume that is assigned to a drive letter or a network share that is specified as a UNC path
- A drive other than the drive on which Backup Exec is installed
- A secure location so that an unauthorized user cannot use the key to access the database

Tuning Backup Exec VI

Good to know

- Turn off Antivirus/Malware software (known to affect GRT enabled backups).
- We recommend excluding the following folders from any type of virus protection (real time and scheduled scans):
 - Catalogs folder
 - Logs folder
 - GRT temp paths (both)
 - Data\Database folder
 - All backup-to-disk folders
 - The deduplication storage

Note:

Keep in mind that tape based backups (stored offline) cannot get infected by malware like Cryptolockers (Locky Ransomware).

Tuning VMware Backups I

Buffer Tuning

HKLM/Software/Symantec/Backup Exec for Windows/Backup Exec/Engine/VMware Agent/

Registry Key Name	Default Value	Rec. Value for 1GB NBD	Rec. Value for SAN/10GB NBD
Enable Buffered Reads	1	1	1
Numbers of Read Buffers	4	10	16
Size of Read Buffers	1024	4096	8192
Enable Buffered Writes	1	1	1
Numbers of Write Buffers	4	10	16
Size of Write Buffers	1024	4096	8192
Write Thread Priority	0	1	1
Read Thread Priority	0	1	1

VERITAS Article 000085311: <http://www.veritas.com/docs/000085311>

Tuning VMware Backups II

VMware Tuning

- Configuration tool (including readme) available at

PingUs: <https://www.pingus.de/wp-content/plugins/download-attachments/includes/download.php?id=141>



VMware Snapshot Management

Good to know

- Improve Snapshot Handling (Retry) - Hybrid Type Snapshot =
VERITAS Article 000086956 - <http://www.veritas.com/docs/000086956>
 - HKEY_LOCAL_MACHINE\SOFTWARE\Symantec\Backup Exec for Windows\Backup Exec\Engine\VMWare Agent\
 - “Use hybrid snapshot method” value data = 1
 - “Snapshot retry attempts” value data = 5
 - “Snapshot retry interval” value data = 60
- Increase Snapshot Cleanup attempts =
VERITAS Article 000087296 - <http://www.veritas.com/docs/000087296>
 - HKEY_LOCAL_MACHINE\SOFTWARE\Symantec\Backup Exec for Windows\Backup Exec\Engine\VMWare Agent\
 - “AutoOrphanSnapshotRemoval” value data = 1
 - “RemoveSnapshotRetryCount” value data = 2

Backing up Database Servers I

Good to know

When backing up (agent-based) database servers like Exchange, SharePoint, SQL, Oracle, Enterprise Vault or Active Directory, keep the following topics in mind:

- For each server create (at least) two job policies:
 - One for the operating system and system state for SDR purposes (weekly or less).
 - One for the database(s) and its related resources (Daily or more).
- Schedule the policies in a way that they don't run in parallel, as this will lead to VSS issues and failed jobs.

Backing up Database Servers II

Good to know

When backing up (host based) SQL servers and you have applications on other servers that use databases on the (remote) SQL, consider the following:

When doing host-based backups, the smallest instance you can select, is the VM.
In other words, you cannot unselect any of the items inside a VM during host-based backups.

Since SQL is a VSS aware application, all databases are flagged as backed up during the host-based, too.

This may lead to inconsistent backups of the application running on the other VM, because “someone else” backed up “its” SQL database (and maybe even truncated the log files) without informing the application.

If you are running such applications, you should exclude the SQL server from all host-based backups and protect it via an agent-based backup policy only.

Backing up Database Servers III

Exchange DAG – Preferred Node for Backups

- Choose one DAG node to do all backups of the databases from, independent, if they are active or passive.
- Can drastically reduce the time it takes to do a backup of a widely distributed DAG, when the node nearest and/or most performant for the Backup Exec server is set as preferred.
- Takes precedence over the job settings regarding backup from active or passive database.

Backing up Oracle (on Windows) I

Good to know

When backing up Oracle database servers keep the following topics in mind:

- Install the Backup Exec Remote Agent on the Oracle server.
- Open the Backup Exec Agent Utility on the Backup Exec Server.
- Click on “Change Settings”.
- On the “Oracle” tab click “new”.
- Select the instance(s) you want to protect and enter a username and password with SYSDBA rights on that instance.
- On the “Database Access” tab enable the “Enable the Backup Exec Server...” checkbox.
- Specify a username and password for a user that has administrative permissions on this server.
Note: You can use the Backup Exec System Account (BESA) here.

Backing up Oracle (on Windows) II

Good to know

- If necessary, change the IP port used to access the Oracle installation.
- On the backup server click the Backup Exec button, select “Configuration and Settings” and then click “Backup Exec Settings”.
- In the left pane, click “Oracle”.
- Enter the name of the Oracle server on which the instance is installed.
Note: specify the name exactly, as it is shown in the “Backup and Restore” tab. Else Backup Exec might not be able to match the two entries.
- Click “Add”.
- Enter the username and password you entered in the “Database Access” tab on the Oracle server before.
- Create a job policy for backing up the Oracle instances.

Backing up the Backup Exec Server

Good to know

In case of a disaster restore, the most important machine to recover will be your backup server. In order to prepare for this scenario, you should do the following:

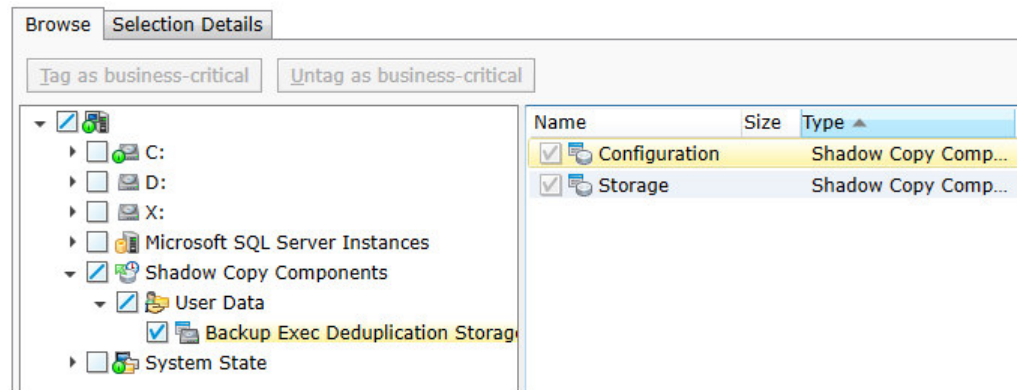
- Create a SDR backup of all components of the backup server and exclude only the volumes that contain B2D- and deduplication storages. (How to backup your deduplication storage folder will be explained later)
- Configure the alternate storage path for SDR files in the Backup Exec settings to point to a network share.
- Let the backup run either to a B2D device located outside the backup server (USB drive or network share) or (better) to a tape device (Locky Ransomware).
- Create a SDR boot medium including all necessary device drivers for your backup server's hardware.
- Note: SDR Restores of the Backup Server from deduplication storages are not possible, as the deduplication database is not available in the SDR preboot environment.

Backing up the Deduplication Storage

Good to know

In order to back up your deduplication storage folder to tape, do the following:

- Create a new backup policy.
- Give it a strikingly name (i.e. “<server name> Dedup”).
- Unselect all items in the selection list by removing the checkmark on the server name level.
- Include the “Shadow Copy Components” resource by clicking on its checkmark.
- Select a tape device (recommnded) as the target for the job.



Where to find more information

- Backup Exec Administrator's Guide
<http://www.veritas.com/docs/000116155>
- Backup Exec Software Compatibility List (pdf)
<http://www.veritas.com/docs/000115689>
- Backup Exec Hardware Compatibility List (pdf)
<http://www.veritas.com/docs/000115688>
- Backup Exec License Guide:
<http://www.veritas.com/docs/000024885>
- Backup Exec Product Use Rights Document (EULA):
<https://www.veritas.com/about/legal/license-agreements.html>
- Download trialware:
<https://www.veritas.com/content/trial/en/us/backup-exec-16.html>
- Backup Exec Support
https://www.veritas.com/content/support/en_US/15047.html
- HCL, SCL, Admin Guide all Versions
<http://www.veritas.com/docs/000017788>
- Backup Exec Software Compatibility List (HTML)
https://download.veritas.com/resources/content/live/SFDC/116000/000115689/en_US/be_16_scl.html?_gda_=1478691974_6d7aaa946e3d03c1ded90004f997872e
- Backup Exec Hardware Compatibility List (HTML)
https://download.veritas.com/resources/content/live/SFDC/116000/000115688/en_US/be_16_hcl.html?_gda_=1478693436_ae17ca77970eec73889eca399ccf67bb

End of Support Life Dates

Version	End of Life Date	End of Standard and Start of Partial	End of Support Life Date
BE 2010	5 March 2012	5 September 2014	1 February 2017
BE 2012	2 June 2014	2 December 2016	5 May 2017
BE 2014	5 April 2015	6 October 2017	7 November 2018
BE 15	7 November 2016	7 November 2017	7 November 2018
BE 16	7 November 2017	7 November 2018	TBA

VERITAS™

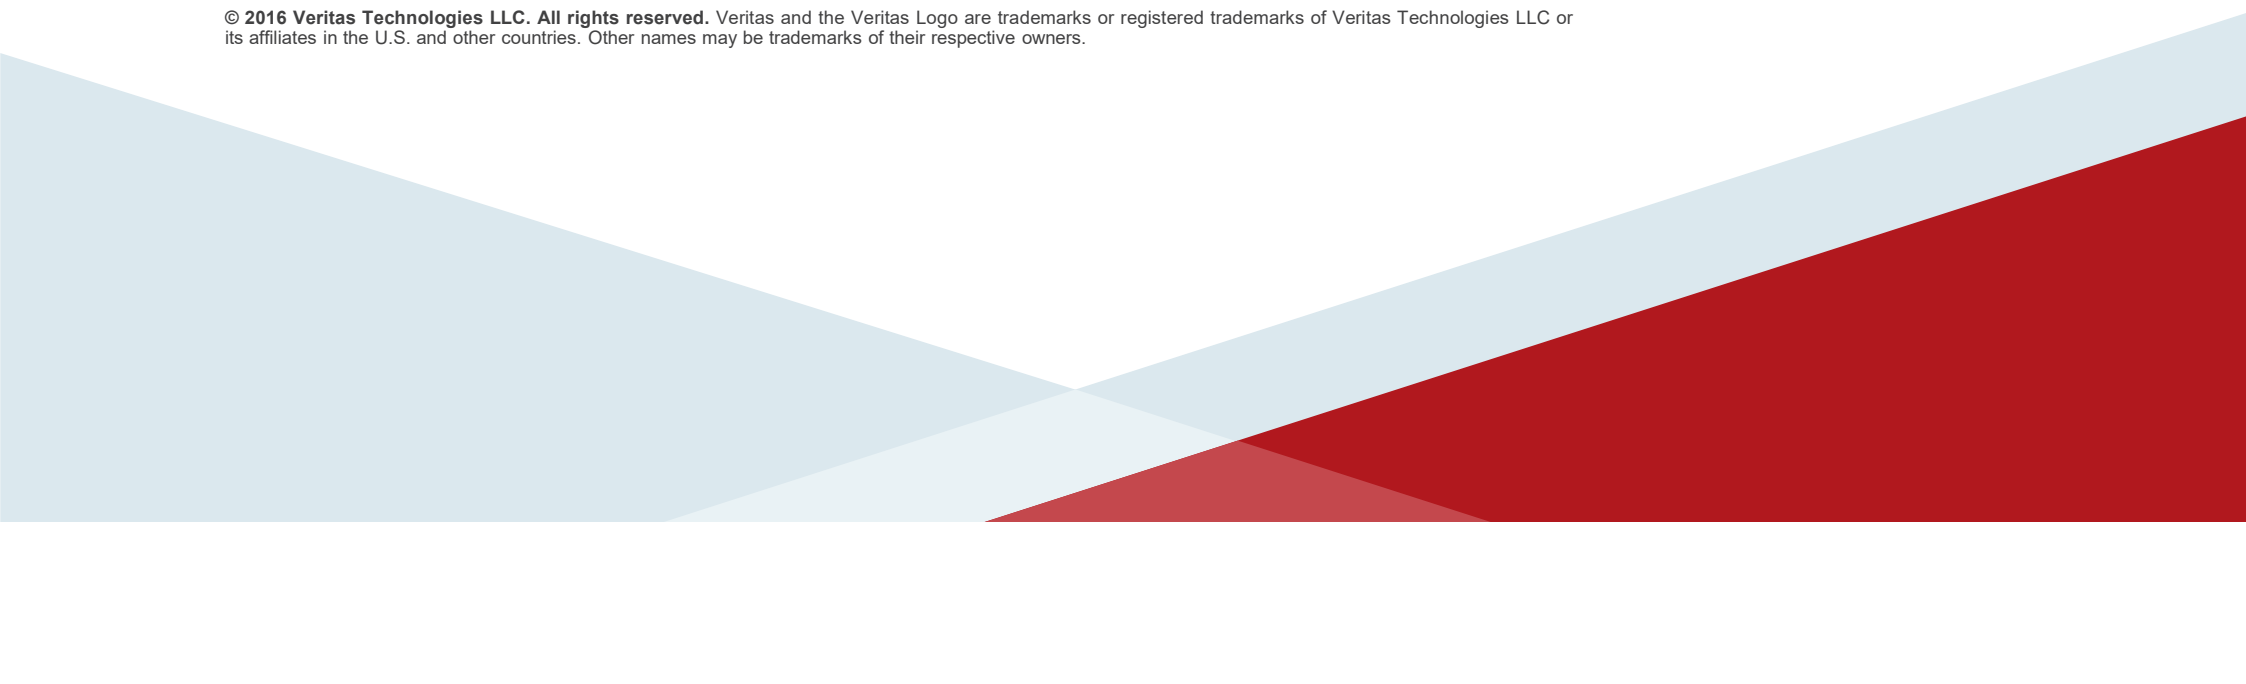
Thank you!

Klaus Kresnik

klaus.kresnik@veritas.com

@KlausKresnik

© 2016 Veritas Technologies LLC. All rights reserved. Veritas and the Veritas Logo are trademarks or registered trademarks of Veritas Technologies LLC or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

The bottom of the slide features a decorative graphic consisting of several overlapping triangles. On the left, there is a large light blue triangle. To its right, a smaller, lighter blue triangle overlaps it. Further right, a large red triangle overlaps the light blue ones, extending towards the bottom right corner of the slide.

VERITAS™

IT'S TIME FOR VERITAS.



#BackupExec Best Practices Guide 4.0

