

Veeam 9.5 VMCE ADO (VMCE-A) Study Guide

SECTION 1 – INTRODUCTION

1.1 – Course Overview

- a. This course provides a design, sizing, & deployment Availability solution to meet RTPO of less than 15 minutes

SECTION 2 – DESIGN AND SIZING

2.1 – DNS & Name Resolution

- a. Ensure all VMware components accessible via forward & reverse DNS lookup
- b. If DNS is not available, a workaround is to add vCenter Server, ESXi Hosts, and Veeam Managed Servers to all Veeam component local hosts file
 1. Recommended to add both the short name (hostname) and FQDN
 2. Recommended to add FQDN of ESXi Hosts when adding them to vCenter
 - a) When using Network (nbd) Mode, the FQDN is returned via VADP, thus the Veeam Proxy Server must be able to resolve the FQDN via DNS

2.2 – Veeam Backup Server

- a. Compute Requirements
 1. **Minimum recommendation: 2 CPU cores and 8GB RAM**
 - a) 1 CPU core / 5GB RAM per 10 concurrent running jobs
 - b) Concurrent jobs:
 - 1) Running backup or replication jobs
 - 2) Jobs with continuous schedule > Backup Copy and Tape jobs
 2. Use 30 VMs per job for better parallelism & resource usage
 - a) Using per-VM, can have appx. 300 VMs
 3. SQL Server Database load is dependent upon concurrent running jobs
- b. Database
 1. Default name = **VeeamBackup**
 2. Items stored in DB:
 - a) Backup infrastructure info
 - b) Jobs settings
 - c) Jobs history
 - d) Sessions
 - e) Files
 3. Default install = **SQL Server 2012 Express**

Veeam 9.5 VMCE ADO (VMCE-A) Study Guide

- a) SQL Express limitations/when to use “full” SQL
 1. Each instance uses only 1GB RAM
 2. Each instance uses only up to 4 cores of the first CPU
 3. Database cannot be larger than 10GB
 4. If unable to accommodate Staging Server req’s due to DBs > 10GB
 5. If environment has > 500 VMs for Veeam or 1000 VMs for VeeamONE
 6. When using advanced DB features such as encryption
4. Backend storage recommendation for config DB = 10K RPM or faster drives

DB Sizing Recommendations

Number of concurrently running jobs	CPU	RAM
Up to 25	2	4 GB
Up to 50	4	8 GB
Up to 100	8	16 GB

Figure 1, CPU & RAM Veeam Database Configuration Guidelines

- c. Operating System
 1. Latest OS recommended; minimum desktop/server OS = Win7 SP1 / Win2K8 SP2
 2. OS support reference: [user guide](#)
- d. Disk Space – Veeam Server Folders and Size
 1. Default install, **40GB** – C:\Program Files\Veeam\Backup and Replication
 2. Logs location, **3GB per 100 VMs** – C:\ProgramData\Veeam\Backup
 - a) Environments with **> 500 VMs**, use separate disk for logs & **no SQL Express**
 3. Backup Catalog, **100MB per 1M files (Linux = 50MB)** – C:\VBRCatalog
 - a) Used if VM guest indexing is enabled; VMs = 10MB per 1M files
 4. vPower NFS, **10GB or 10% per TB** – C:\ProgramData\Veeam\Backup\nfsDatastore
 - a) Used to boot VMs for **Instant VM Recovery**, for config files & redo logs
 - b) vPower NFS Installs
 - 1) Windows Repositories – install on the Repository itself
 - 2) Dedup appliances & CIFS/SMB shares – install on Gateway Server
 - 3) Linux Repository – on managed Windows server close to Repository
- e. Other Software
 1. Strongly advised against high transaction software installed on Veeam Server, such as:
 - a) AD
 - b) Exchange
 - c) Prod SQL DBs
 2. Allowable installs:
 - a) Staging SQL
 - b) Outlook for .pst restores
 - c) Adobe Reader

Veeam 9.5 VMCE ADO (VMCE-A) Study Guide

Chart of Veeam & SQL Requirements

COMPONENT	REQUIREMENT
CPU / RAM	Minimum CPU – 2 Cores Minimum RAM – 8 GB
Database (VeeamBackup)	SQL – Full or Express Editions supported Minimum – MS SQL 2008 Current version – MS SQL 2017
Operating System	Recommended: latest OS version & patch Minimum Server – Windows 2008 SP2 Minimum desktop – Windows 7 SP1
Storage (Folders and minimum space requirement)	Install location – 40GB, C:\Program Files\Veeam\Backup and Replication Logs – 3GB per 100 VMs, C:\ProgramData\Veeam\Backup and Replication Backup Catalog – 100MB per 1M files, C:\VBRCatalog (10MB per 1M on guests) vPower NFS – 10GB (or 10% per 1TB), C:\ProgramData\Veeam\Backup and Replication\NfsDatastore
Other (software)	High I/O transactional applications not supported – AD , Exchange, SQL Supported co-software – Adobe, Outlook (.pst restores), Staging SQL

Figure 2, Veeam Backup Server Requirements Overview

- f. Backup Server Placement > in relation to data flow between ancillary Veeam services for optimal backup performance & restore throughput
 1. Host & Storage discovery – **Broker Service**
 - a) Host scans = **4hrs** ; Storage scans = **1hr**
 - b) View discovery info: Veeam Console > **History** tab > **System view** section
 - c) Stored in configuration DB
 - d) Not recommend to add standalone ESXi Hosts if using shared storage
 - e) Storage scans performed by Proxy servers
 - 1) Veeam Server reads info using storage controller management API
 2. File-level recovery flow
 - a) Two mounts are performed for Windows file-level recovery

Veeam 9.5 VMCE ADO (VMCE-A) Study Guide

- 1) Remote console – displays restore point contents
 - i. 50-400MB may be transferred between Repository & Veeam Console; if this mount is over a latent network, its recommended to install a Console close to the Repository to speed up time it takes to open FLR wizard
- 2) Mount server – performs traffic restore to target VM
3. Veeam Backup Enterprise Manager (EM)
 - a) Self-service portal for admins or Helpdesk to recover files, VMs, emails, DBs
 - b) Remote Console mount can be avoided if file system guest indexing is enabled, as the guest VM content is stored in the Veeam Catalog & presented through EM
4. Veeam Explorers
 - a) Installed by default *with* Veeam Console
 - b) Exchange, SQL, Oracle Explorers use staging servers, so attention should be paid for network connectivity between the mount server & staging servers
 - c) Storage Explorer for Snapshots workflow:
 - 1) Snapshot clone created
 - 2) Block access – mount LUN to ESXi Host & register a temp datastore
 - 3) NFS access – access existing NFS datastore & look for cloned VM
 - 4) Register temp VM within VMware inventory
 - 5) Access VM via VMware API
 - 6) Show content in Veeam Explorer to restore
 - 7) After exiting Veeam Storage Snapshot Explorer, the temp datastore, VM, & LUN Clones will be removed

2.3 – Veeam Proxy Server

- a. Compute Requirements
 1. General info:
 - a) Incorrect proxy sizing is a common cause of performance issues
 - b) A Proxy **task** processes **1 VM disk** at a time; doubling the task count can reduce the backup window by 2x
 - c) Recommended sizing = **1 core / 2GB+500MB RAM per concurrent task**
 2. Proxy sizing calculation methods
 - a) Formula method
 - 1) **Veeam recommended method**
 - 2) **D** – source data in MB (#TB * 104857)
 - 3) **W** – backup window in seconds (#hrs * 3600)
 - 4) **T** – throughput = D / W
 - 5) **CR** – change rate (.)
 - 6) **CF** – cores required for full backup = T / 100
 - 7) **CI** – cores required for incremental backup = ((T*CR) / 25)
 - 8) Example – see Handbook

Veeam 9.5 VMCE ADO (VMCE-A) Study Guide

- b) Rule of thumb (Cores = Total VMs / 30 ; RAM = # of Cores x 2)
 - 1) Fast way to do sizing but not recommended as it tends to oversize
 - 2) **RoT** = 1 CPU & 2 GB RAM for a set of 30 VMs, avg chg rate = 2-3%, 8hr backup window
 - 3) Example: 480 VMs , 48TB storage , 8hr window , CR = 5%
 - i. Using 1 core per task rule & ea. task appx 30 VMs – 480 / 30 = 16 cores
 - ii. Each task (i.e. core) uses about 2GB RAM, so 2 * 16 = 32GB RAM req'd
- 3. Physical Proxies
 - a) Recommended: **dual-socket CPUs with 10 cores each (20 cores total)**
 - b) Increasing cores can lead to network saturation, so recommended to use 2 servers of 20 cores & one 10 GbE nic each instead of one 40 core with two 10GbE nics
- 4. Virtual Proxies
 - a) **No more than 8 CPU cores per Virtual Proxy VM to avoid co-stop scheduling**
 - b) Add more Proxy VMs instead of more cores
- b. Transport Modes – method used by Veeam Proxy to read source VM data from a Host & write the VM data to a target
 - 1. Direct storage (**san**) – VMware Direct SAN / Veeam Direct NFS
 - a) Can be deployed via:
 - 1) Physical Proxy using FC/FCoE/iSCSI/Shared SAS/NFS
 - 2) Virtual Proxy using iSCSI or NFS
 - b) Veeam Direct NFS limitations:
 - 1) NFS cannot be used for VMs with snapshots
 - 2) Direct NFS is only used for first **target side** session of Replication; source will continue using NFS
 - 3) Direct NFS not used for Windows VMs if VM Tools Quiescence enabled
 - 4) If Direct NFS cannot be used, Veeam will use nbd
 - 2. Network (**nbd**)
 - a) **Uses VMware storage stack** to read/write data *through VMkernel interface via NFC protocol port 902*
 - b) Limited to 40% throughput by VMware
 - 3. Virtual appliance (**hotadd**)
 - a) At backup or replication, a VM snapshot is processed, the snap disk is mapped to hotadd proxy to read data (at backup) & write data (at restore), then is unmapped
 - b) Runs using **Veeam Data Mover service** 'below normal priority'
 - c) CBT is disabled for VM Proxies
 - d) **IDE disks not supported** (backup or restore) for hotadd
- c. Guest OS Processing
 - 1. Process
 - a) Veeam performs guest OS inventory to check for VSS-aware apps
 - b) Veeam runs pre-freeze script (if any) for Windows/Linux that utilize other quiescence means

Veeam 9.5 VMCE ADO (VMCE-A) Study Guide

- c) VSS quiescence performed & a **small executable component is deployed** to coordinate VSS & indexing
- d) VM snapshot created
- e) VSS unfreeze/'thaw' performed
- f) Veeam runs post-thaw script (if any)
- g) Backup data transfer & snapshot commit performed
- h) Log file truncation performed with VSS or via native Oracle commands
- 2. Veeam connects to VM for guest processing **over network using RPC**, or **VMware VIX channel** if RPC unsuccessful
- d. Application-Aware Processing (AAIP)
 - 1. **MySQL doesn't support VSS**; must use VMware Tools or HV-IC guest quiescence
 - 2. Supported on Windows 2003 SP2 and later; not supported on non-Windows VMs

2.4 – Veeam Repository Server

- a. Sizing Compute Resources
 - 1. Two calculation methods for CPU
 - a) Per backup job files – **1 CPU core per concurrent job** that processes data
 - b) Per VM backup files – **1 CPU core per concurrent VM** that processes data
 - 1) Formula = *total backup proxy concurrent tasks targeting repo / avg virtual drive count per VM stored on repo*
 - 2. RAM Recommendation: configure **4GB RAM per CPU core** regardless of backup write method used
 - a) **RECOMMENDED: 2 cores/8GB RAM**; *minimum req'd: 2 cores/4GB RAM*
 - 3. **Use same compute resources for Gateway servers**
- b. Choosing Repository Type – it should be highly resilient & scalable
 - 1. Factors to consider:
 - a) Capacity
 - b) Write performance
 - c) Read performance
 - d) Data density
 - e) File security
 - f) Backup file utilization
 - 2. High-performance Repository can be used for recent retention & instant restores
 - 3. A Repository with capacity, slower storage used for long-term retention (BCJs)
- c. Estimating Repository Capacity – Calculator > <https://rps.dewin.me/>
 - 1. Data needed (i.e. sizing)
 - a) Total size of VMs backed up – frequency, retention, fwd/reverse chain
 - b) Make compression & dedup assumptions – generally 2:1 ratio
 - c) Change rate (typically 2-5%)
 - d) Include space for one-off full
 - e) Include space for backup transform operations (ffwd or rev): full * 1.25
 - f) *****Always add space for future growth*****

Veeam 9.5 VMCE ADO (VMCE-A) Study Guide

NOTE: For dedup sizing, consult with the dedup storage vendor

2. Can use Veeam ONE VM Change Rate Estimation report as a pre-deploy assessment of how much space Repositories should use/need
3. Run **Capacity Planning for Backup Repositories Veeam Reports pack report** periodically to analyze Repository free space & estimate future storage

2.5 – Veeam WAN Accelerator

- a. Uses network compression, multi-threading, dynamic TCP window size, variable block dedup, & global caching for offsite Backup Copy & Replication jobs
- b. Determining Required Bandwidth
 1. Compare primary Backup job daily chg stats with WAN-Accelerated BCJ stats
- c. Backup Mode Effect (used on source backup job)
 1. **Sequential I/O reads** – Fwd/FFwd Increm Modes (any .vib's); better performance
 2. **Random I/O reads** – Reverse Increm/Synthetic/Active Full (any fulls)
- d. Source WAN Accelerator
 1. TCP threads: default = 5; max = 100; test stream config scenarios using WANem
 2. Configuring source-side cache irrelevant, VM data is read directly from Repo
 3. Consumes more CPU (due to compression); recommended **4 CPU / 8GB RAM**
 4. Disk size (C:\VeeamWAN\Digests) = (Total source data * 2%) + 10GB working space
 - a) Use SAS or near-line SAS; no benefit using SSD
- e. Target WAN Accelerator
 1. Cache size
 - a) Configuration is applied to the WAN Accelerator pairs
 - b) Recommended 10GB per VM OS *type* (i.e. Win8, Win2016, etc) processed
 - c) Recommended to use cache pre-population
 - d) Recommended to seed initial copy of data to target Repository
 2. Consumes more I/O; recommended **2 CPU / 8GB RAM**
 3. Disk size – 10GB per OS type + 20GB working space
 - a) For 1-to-1 Accelerator ratio, use SAS disks; for Many-to-1 use SSDs
 - b) C:\VeeamWAN\GlobalCache\trg – one folder per paired source with **blob.bin**
 - c) C:\VeeamWAN\GlobalCache\temp – **data.veeamdrf** file transferred to source
 4. Formulas
 - a) Cache = (# of OS types * 10GB) + 20GB working space
 - b) Used disk space = (# of source WAN Accel's) * (Cache size)
- f. WAN Accelerator Deployment Amount
 1. Recommended **4:1** (source to target) WAN Accelerators; 10:1 can be considered
 2. For WAN Links > 100Mbps, advised not to use WAN Accelerators

Veeam 9.5 VMCE ADO (VMCE-A) Study Guide

VEEAM COMPONENT SIZING RECOMMENDATION TABLE

COMPONENT	SIZE
Veeam Server	2 cores / 8 GB RAM (1 core/5GB per 10 concurrent)
Database	2 cores / 4 GB RAM (per 25 concurrent jobs)
Proxy Server	Cores = (# of VMs / 30) ; RAM = Cores x 2
Repository Server/Gateway Server	2 cores / 8 GB RAM
Source WAN Accelerator	4 cores / 8 GB RAM
Target WAN Accelerator	2 cores / 8 GB RAM

Figure 3, Veeam Component Sizing Overview Table

2.6 – Interactions With Hypervisors

- a. Interaction With vSphere
 1. Special characters (issue with **VMware only**)
 - a) VMs with special characters (i.e. %) will not get processed by Veeam
 2. vCenter Server
 - a) Connectivity with vCenter is one of top reasons for failed Veeam jobs
 - b) vCenter must be able to handle increased transactional workload for environments with many concurrent/near-CDP Veeam jobs or random command timeouts occur
 - c) Connection between Veeam & vCenter must be available, consistent, reliable
 - d) Recommend Veeam jobs & Backup Service stopped during vCenter maint
 - e) Active-Active Metro Cluster recommendations
 - 1) For vCenter 6.5, implement vCenter HA
 - 2) Pre-6.5 – replicate vCenter to other datacenter, & use DRS affinity to pin vCenter VM to a specific ESXi Host
- b. Impact of Snapshot Operations
 1. Snapshotting is vital to Veeam operation, so be aware of snapshot limitations:
 - a) VMs with bus-sharing disks (MSCS), do not support snapshots
 - b) VMs with physical RDMs do not support snapshots
 - c) VMs that use direct-path I/O devices do not support snapshots
 - d) VMs with (persistent) independent disks have snapshot limitations
 - e) Since Veeam **v7 R2**, there is a limit of 4 snapshots per datastore; this limit can be modified using the **MaxSnapshotsPerDatastore** registry change on the Veeam server (HKLM\Software\Microsoft\Veeam\Veeam Backup and Replication)

Veeam 9.5 VMCE ADO (VMCE-A) Study Guide

2.7 – Veeam ONE

- a. Deployment Scenarios
 - 1. **Typical** – preferable for SMBs, installs all components on 1 physical/virtual server; for environments with < 100 Hosts/1500 VMs
 - a) Local or remote SQL Server instance is req'd; comes with SQL 2012 Express
 - 2. **Advanced** – installing components on separate machines, separating client/server roles; for > 100 Hosts/1500 VMs
 - a) Veeam ONE Server and Veeam ONE Web U/I installed on separate machines
 - b) Veeam ONE Monitor Client installed on client machines for user access
 - c) SQL Server (see 'Typical'); **NOTE:** for 1000+ VMs, remote SQL recommended
- b. Data Collection Modes
 - 1. Install type during install
 - a) Typical – for environments with less than 100 Hosts/1500 VMs
 - b) Advanced – for environments with greater than 100 Hosts/1500 VMs
 - 2. Switching between modes after install
 - a) Go to:
%ProgramFiles%\Common Files\Veeam\Veeam ONE Settings\
VeeamOneSettings.exe, then **Deployment Configuration > Save > Restart Veeam ONE Monitor and Veeam ONE Reporting Services**
 - b) Scalability Mode Change Effects (**Typical vs Advanced**)

1) VMware real-time performance query timeout:	300sec vs 1200sec
2) VMware historic performance query timeout:	900sec vs 1800sec
3) MinTreeCheckInterval:	10sec vs 1min
4) PageUpdateDelay:	.5sec vs 1.5sec
5) GetVmsnapshotTimeout:	900sec vs 3600sec
6) CollectthreadTimeout:	6hrs vs 12hrs
7) DashboardCaptureTimeout:	3600sec vs 7200sec
8) Hyper-V cache retention settings:	-1(infinite) vs 120sec
9) Hyper-V collection job interval:	60K ms vs 120K ms
10) Ignore disconnected Hyper-V hosts:	false vs true

SECTION 3 – INFRASTRUCTURE ASSESSMENT

3.1 – Introduction

- a. Covoiturage
 - 1.
- b. Infrastructure
 - 1.

Veeam 9.5 VMCE ADO (VMCE-A) Study Guide

3.2 – Proof of Concept

- a. Small-scale POC typically consists of:
 1. vCenter, ESXi Hosts, **10-20 VMs**
 2. Veeam Server – single VM; **4 CPU cores; 6-8GB RAM**
 3. Usually customer-driven – Veeam Evaluator Guide used as reference
- b. Enhanced Evaluation
 1. Designs often change after initial assessment due to network, storage, security requirements, as well as new customer demands
- c. Workshop Example
 1. Infrastructure discovery
 - a) First datacenter
 - 1) Virtualization platform & version used
 - 2) Storage system, type, & connection used
 - 3) Storage virtualization?
 - b) Second datacenter
 - 1) Storage replication used?
 - 2) Active/active cluster used? (i.e. vCenter HA?)
 - c) Network connection info:
 - 1) 10GbE LAN?
 - 2) WAN between 2 DCs?
 - 3) VMkernel interface link speed?
 - 4) vCenter Server – physical or virtual; location
 - d) Production data info:
 - 1) Number of VMs
 - 2) Used data
 - 3) Number ESXi Hosts & sockets (licensing)
 - 4) Number of Clusters
 - e) Create first implementation draft
 - 1) Start with Repository, using customer demands (see Handbook)
 - 2) Next is Proxy – physical vs virtual
 - 3) Veeam server – VM, and replicate to 2nd DC
 - 4) Storage – list arrays/types (e.g. IBM, EMC)
 - 5) For proper Proxy & Backup Mode selection, it is vital to know where data to be backed up is located
 - f) Define OS/applications
 - 1) Create list of OSs
 - 2) Create list of apps – most critical first; SQL/SP used?
 - g) Define business critical apps – influences RPO discussions & DR planning by:
 - 1) Defining number or biz crit VMs
 - 2) Determining slower performance ok at DR (i.e. use Inst VM Recovery)
 - 3) Use all Veeam possibilities – Full Restore for small VMs; Instant Recovery for large VMs, & Replica Failover for Databases or high I/O VMs

Veeam 9.5 VMCE ADO (VMCE-A) Study Guide

- 4) Determine failure (SPOF) areas – storage (1 SP, datastore, or array); network (single link/latency); does Veeam fit / align with these areas
2. Network & firewall
 - a) Can be used between Veeam components, except RPC inspection function
 - b) It's best to place Veeam components in same network segment as virtual infrastructure to allow efficient/fast usage of bandwidth

VEEAM COMPONENT	PORT
Veeam Backup/Enterprise Service	TCP 9392
Veeam Installer Service	TCP 6160
Veeam Transport Service	TCP 6162
Veeam Data Mover Service	TCP 2500-5000
Veeam Backup Manager > vCenter	TCP 443

Figure 4, Veeam Services Ports

3. Veeam Component Location (Proxy, Repository, Veeam, Veeam ONE, Ent Mgr):

VEEAM COMPONENT	LOCATION
Backup Proxy & Backup Repository	In the VMkernel Network
Veeam Server & Veeam ONE	Next to vCenter Server
Veeam Enterprise Manager	Internal DMZ

Figure 5, Veeam Component Infrastructure Placement

4. Restore points
 - a) Customer requirements
 - 1) Read sample requirements in Handbook
 - 2) Retention policy planning is most time-consuming as it could lead to expensive storage requirements
- d. Preparation
 1. Enhanced POC conditions:
 - a) **Close to prod environment** with firewalls used
 - b) Storage used similar to prod
 - c) Veeam storage integration used when possible
 - d) **100-200 VMs** for good working load
 - e) **All major apps used** to test restore
 2. Prep steps
 - a) Use Veeam User Guide, ADO Handbook, & Veeam Best Pract Guide
 - b) Check firewall ports
 - c) Create vCenter acct for Veeam ONE – **read-only+datastore browsing+CIM**
 - d) Setup storage integration

Veeam 9.5 VMCE ADO (VMCE-A) Study Guide

- e) Verify virtualized domain controller present if using SureBackup
- f) Customer – prepare anti-virus, storage (avoid low-end NAS), & OS/DB installs
- g) Setup Veeam Server & components with customer

3.3 – Assessment Using Veeam ONE

- a. Storage Latency Report (**Total Host Latency / Datastore Average Latency** metrics)
 - 1. Identifies storage systems under heavy pressure or max load; run at least 24hrs
 - 2. Useful when jobs show **source** as ‘bottleneck’
- b. VM Change Rate Estimation Report (**Top 5 VMs with Largest and Least Change Rate**)
 - 1. Helps identify VMs with high change rate at block level
 - 2. Assists with backup target storage planning
 - 3. **Not able to detect in-guest iSCSI or NFS**
 - 4. Config VMs to run at start of Backup/Replic jobs to address long job runtimes
- c. VM Configuration Assessment Report (**VMware**)
 - 1. Helps assess VM readiness for Veeam backup (**VM Overview , Potential Issues**)
 - 2. **Potential Issues** pie chart:
 - a) VMware CBT
 - 1) To utilize CBT, snapshots *should* be removed
 - 2) CBT is supported by VMware Hardware 7 and above
 - b) Virtual disks
 - 1) Detects VMs with Independent disks, which are skipped by Veeam
 - 2) Detects disk 2TB+ on vSphere 5.5 and earlier (disks must be < 1.9TB)
 - 3) Detects disks running SCSI Bus Sharing (MSCS); not supported by Veeam
 - c) Application-aware image processing (AAIP)
 - 1) VMware Tools must be installed/running to use AAIP
 - d) Datastore free space
 - 1) Datastores must have more than 10% free space
 - e) Other
 - 1) Detects VMs with names having: @ / \ < >
- d. Configuration Assessment Report (**Hyper-V**)
 - 1. Cluster optimization
 - a) Hosts use same CPU vendor/model
 - b) Hosts use same storage controller vendor
 - c) Hosts use same Hyper-V version
 - d) Hosts use same NIC vendor
 - 2. Host optimization
 - a) Hosts not running Balanced power plan (**High Performance** recommended)
 - b) Hosts do not have RDP printers mapped
 - c) Host vCPU / Logical CPU ratio is <= 8
 - 3. VM optimization
 - a) Hosts do not have Checkpoint (snapshot) Redirection verification
 - b) Verification of VM Integration Services running; required for AAIP

Veeam 9.5 VMCE ADO (VMCE-A) Study Guide

- c) VM VSS Requestor state on VMs – must be Started/Auto or Auto(Delay)
- d) Verification no Checkpoints (snapshots) so CBT can be utilized
- e) Datastores need > 10% free space
- e. Infrastructure Overview Report
 - 1. Shows configuration properties of VMs, Clusters, Datastores, Hosts, & Networks
- f. Active Snapshot Report
 - 1. Shows **top 5** snapshot age (oldest) & size for select VMs across a time range
 - 2. Best practice suggests **deleting snapshots older than 3 days**
 - 3. **NOTE:** select to exclude Replication snapshots; **snaps are Restore Points since v6**
- g. Orphaned VM Snapshots Report (shows **Datastore Free, Wasted, Used Space**)
 - 1. Snapshots on Datastores but not shown in VMware Snapshot Manager
 - 2. Snapshot Hunter process – scheduled by **Veeam Backup Service**
 - a) VMware Snapshot Consolidation method
 - b) Hard Consolidation without Quiescing – performed **ONLY** if no user snaps
 - c) Hard Consolidation with Quiescing – performed **ONLY** if no user snaps
 - d) Performed up to 4 times within a 4hr interval; email sent if all attempts fail
 - e) Backup windows are considered for the process; i.e. won't be performed if interval not within window; email will instead be sent & job has **Warning**
 - f) Process viewable: **History > System view**

SECTION 4 – SECURITY

4.1 – Permissions

- a. Hyper-V
 - 1. **Local admin** for product install
 - 2. **Local admin** for Source Hyper-V server
 - 3. **Root** for Linux Repository (if used)
 - 4. **Write** permission on target/share folder
 - 5. Off-Host Proxy processing VMs on SMB3 Share:
 - a) Off-Host Proxy **LOCAL SYSTEM** account needs **Full Access** on SMB3 Share
 - b) Off-Host Proxy must be located in **same domain as SMB3 share**
 - c) SMB3 server domain must be trusted by Off-Host Proxy domain
- b. vSphere
 - 1. vCenter Server – granular permissions with limited scope; full admin recomm
- c. Application-Aware Processing
 - 1. **Local admin** account for VSS processing on backed up VM using AAIP
 - a) If UAC must be enabled, **SID-500 'administrator' account** needs to be used
 - b) If no sid-500 acct can be used, UAC must be disabled:
 - 1) Win 2008/R2 – change UAC setting to 'Never Notify'
 - 2) Win 2012/R2 – reg change: **DWORD: EnableLUA = 0**
HKLM\Software\Microsoft\Windows\CurrentVersion\policies\system

Veeam 9.5 VMCE ADO (VMCE-A) Study Guide

- d. Veeam Explorer for Microsoft Active Directory – target (restore) DC:
 - 1. **Administrator** for target AD
 - 2. **Exchange Management Organization member** for auto mailbox reconnect for recovered user or group
 - 3. AD object recovery logic
 - a) Tombstone object exists – it's used to restore security attributes, objectSIDs, & objectGUIDs; **DISABLE AD RECYCLE BIN FEATURE**
 - b) No tombstone object exists – Veeam creates new object during recovery & sets all attributes to same value as in backup, but access rights may be lost
 - 4. AD object recovery considerations:
 - a) Use AD Explorer for proper object restore
 - b) Use backups newer than tombstone lifetime interval; use ADSIEdit/DSQUERY
 - c) When moving an object to other domain within Forest, no tombstone will remain, and thus object cannot be recovered to the original domain
 - d) Group Policy restores > existing policies are deleted & restored ones are written
- e. Veeam Explorer for Microsoft SQL Server
 - 1. **sysadmin fixed** role on SQL Server VM backing up or on granular roles
 - 2. **sysadmin fixed** role on target SQL Server restoring DB to
 - 3. **Read/Write** for connecting to Target/Staging SQL server for log replay
 - 4. Acct running Veeam SQL Explorer have **Read/Write** on folder exporting DB to
 - 5. For log truncation Veeam will use NT AUTHORITY\SYSTEM for SQL 2012/2014 or LOCAL SYSTEM for SQL 2005/2008/R2 if the specified account is unsuccessful
- f. Veeam Explorer for Microsoft Exchange
 - 1. **Full Access** to Exchange DB & logs
 - 2. **Read/Write** to files in the folder with the DB
 - 3. Access rights for item recovery can be granted via impersonation or providing the user with Full Access to the mailbox
 - 4. Public Folder restore:
 - a) Account used should **own a mailbox** on the target Exchange Server
 - b) Account should have **an org management role** on target Exchange Server
 - 5. Mailbox restore:
 - a) Mailbox owner acct on target Exchange server has **Full Access** for mailbox
 - 1) Full access can be granted through impersonation or rights assignment
 - 2) If acct is a non-mailbox owner, rights must be granted via impersonation
- g. Veeam Explorer for Microsoft Sharepoint
 - 1. **SP farm admin / sysadmin fixed role** accts for source VM guest processing
 - 2. **sysadmin fixed role** for Sharepoint Explorer acct on staging SQL server
 - a) If using ADFS to connect to staging SQL – can use current or specified account
 - b) If using Forms Auth to connect to staging SQL – can only use specified acct
 - 3. Account for target Sharepoint server where item are being restored to:
 - a) If permissions for restored item are inherited from parent: **Full Control** for item

Veeam 9.5 VMCE ADO (VMCE-A) Study Guide

- b) If permissions not inherited & item replaces existing item: **Contribute** for item & **Full Control** for its parent
- 4. **Full Control** on target SP where SP site is restored to; if user is not a site/system admin, the restore user's permissions will not be restored
- h. Veeam Explorer for Oracle
 - 1. Connection to source server (guest OS & Oracle DB)
 - a) Single account in Guest Processing step of Backup job setup
 - 1) Windows
 - i. **Local admin** to connect to VM, access file system & admin share
 - ii. **Read/Write** for DB file hierarchy access (Full Control recommended)
 - iii. **ora_dba** group member
 - 2) Linux
 - i. **root** or acct elevated to root via sudoers
 - ii. **oinstall** & **asmadmin** group member
 - b) Two accounts – in Guest Processing step & on Oracle tab of Backup job
 - 1) Windows – **Local Admin / Read/Write / SYSDBA** via **ora_dba**
 - 2) Linux – **root / SYSDBA** via local **dba** group
 - c) **NOTE: identical accounts are needed for all DBs on source server**
 - 2. Connection to target server
 - a) Windows – **Local Admin & Read/Write** for DB hierarchy
 - b) Linux – **root**; authentication by user/pwd or by private key
 - 3. Connection to staging server
 - a) Windows – **Local Admin & SYSDBA**; Oracle 12c = pwd & *Logon As* permission for Oracle Home user
 - b) Linux – **root & dba & asmadmin** member

4.2 – Veeam Backup Enterprise Manager

- a. Roles
 - 1. Portal admin – full access
 - 2. Restore operator – access VMs within restore scope & perform restores
 - 3. Portal user - access VMs & reports within restore scope; perform restores
- b. **NOTE:** for self-service discovery, reverse DNS must be functional
- c. Restorable Items:
 - 1. Entire VMs
 - 2. Guest files
 - 3. MS Exchange items
 - 4. MS SQL DBs
 - 5. Oracle DBs

Veeam 9.5 VMCE ADO (VMCE-A) Study Guide

SECTION 5 – DESIGN SCENARIO: PART 1

5.1 – Architecture Meeting

- a. First Datacenter
 1. See Sect 3-2, c.
- b. Based on assessment info gathering (**Sect. 3**) & customer requirements, a solution can be designed on paper
 1. Changes may be needed after mtg with other depts or customer changes demands; both can lead to higher costs & delay implementation
- c. Architecture meeting
 1. Done to limit unexpected design changes discussed above
 2. Answer any questions by all depts
- d. Second Datacenter
 1. See Sect 3-2, c
- e. Networking
 1. See Sect 3-2, c
- f. Amount of Production Data
 1. See Sect 3-2, c
- g. Applications
 - a) See Sect 3-2, c
- h. Implementation Draft
 1. Start with Repository – if data needed in 2 DCs, discuss 2 Repos & use of BCJ
 2. Proxy – physical vs virtual; determine mode (san, hotadd, nbd); disk type (thick?)
 3. Backup Server – physical vs VM; replicate if VM, but not its underlying Datastore
 4. Other components – e.g. tape?
 5. Use **all** Veeam possibilities to come up with best RTO/RPO solution design
 - a) Recovery options:
 - 1) Classic VM Restore – small VMs
 - 2) Instant Recovery – large VMs
 - 3) Replication – high I/O DBs
 - b) Determine SPOF areas

5.2 – Scenario

- a. Datacenters
 1. DC1 – main office & DC at same location (different floors)
 - a) Read mock customer DC1 infrastructure data in Handbook
 2. DC2
 - a) Read mock customer DC2 infrastructure data in Handbook
- b. Current Solution
 1. See Handbook
 2. Pain points:

Veeam 9.5 VMCE ADO (VMCE-A) Study Guide

- a) See Handbook
- c. Customer Data Protection Requirements / Other Information
 - 1. See Handbook

SECTION 6 – OPTIMIZATION

6.1 – Protecting Veeam Backup Server Configuration

- a. Recommended to replicate Veeam & remote DB VM to reduce overall RTO
- b. Recommended to place a config backup copy at DC2 via File Copy job
 - 1. If config backup lost/damaged, you can import backups to Veeam
 - 2. Can also use standalone Extract utilities to restore VMs, files, & app data

6.2 – Veeam Backup Server Database

- a. MS SQL Server Edition – DB instance named **VeeamBackup**
 - 1. MS SQL 2012 Express – included with Veeam
 - a) Recommended to upgrade DB to SQL 2014 Express
 - b) SQL Express limitations:
 - 1) Each instance uses up to only 1GB RAM
 - 2) Each instance uses up to 4 cores of the *first* CPU socket
 - 3) DB size maximum = 10GB
 - 2. When to use “full” SQL Server version:
 - a) For **> 500 VMs**, mainly due to CPU/RAM constraints, not DB size
 - b) If unable to configure Staging SQL for SQL/ Sharepoint due to **DB > 10GB**
 - c) When needing **to match license level and/or feature-base of recovered DB**

SQL Server Edition	Express	Standard	Enterprise
Supported Processors	1	4	8
Max Memory	1GB	64GB	OS max
Max DB size	10GB*	524PB	524PB

Figure 6, SQL Server Maximums

***NOTE:** SQL Express 2008 and earlier versions support only 4GB DB maximum

- b. Database Placement
 - 1. Since Veeam v9, Explorer Staging SQL Servers can be remote, so no longer a requirement for Veeam to have Standard/Enterprise SQL locally
 - 2. Why Local (on Veeam Server)?
 - a) Recommended for lowest latency & highest performance
 - 3. Why Remote?

Veeam 9.5 VMCE ADO (VMCE-A) Study Guide

- a) SQL Clustering used with AlwaysOn Availability Groups
- b) Simplified Veeam recovery
- c) Enterprises have dedicated virtual SQL Clusters
- c. SQL Server Configuration Tips
 1. Consider SQL Std/Ent for Staging of Explorer for SQL/Sharepoint DBs & configure all features prod DBs use
 2. Install same license level, latest version, and latest cumulative updates
 3. For encryption DB restores, install valid cert on Staging SQL Server
 4. Verify the SQL **sysadmin** role is enabled for proper log file processing SQL 2012+
 5. Place SQL **tempdb** on fastest disks
- d. Modifying Database Connection Settings
 1. Utilize the DBConfig utility (`Veeam.Backup.DBConfig.exe`)
 2. DB connection types:
 - a) sa auth – verify Veeam UI & PoSH changes use this
 - b) Windows auth – verify Veeam UI or PoSH user has proper DB rights (**KB1853**)

6.3 – Proxy Server (vSphere)

- a. Direct Storage Access – VMware Direct SAN or Veeam Direct NFS
 1. Pros
 - a) Very fast, reliable, predictable backup perf using 8Gb FC or 10Gb iSCSI/NFS
 - b) Zero impact on ESXi Hosts or VM prod networks for data transport
 - c) Ability to do Full VM Restore since Veeam **v8+** if VMs use thick disks
 - d) NFS – fastest restore of NFS since Veeam v9 due to using multiple threads
 - 1) Provides optimized STUN times due to multithreaded reads
 - e) Both Direct SAN & NFS **can be used for all Replica source read data**, and *initial Replica target data* (thick-disk only for SAN)
 2. Cons
 - a) Requires physical server; can use virtual for iSCSI Direct SAN
 - b) For FC/iSCSI Direct SAN, VM restores can only be done for VM thick disks
 - 1) Datastream needs to coordinate with vCenter in background which can lead to slow restore speed, as such consider using hotadd for restore
 - c) Difficult to configure – phys servers, Zoning/iSCSI config, driver updates
 3. Recommendations
 - a) Use storage vendor for multipath driver & for optimal settings
 - b) Attaching numerous volumes to Proxy can increase logging & VM lookup job processing time; disable process: `HKLM\Software\Veeam\Veeam Backup and Replication\VDDKLogLevel, DWORD = 0`
 - c) Best performance vs cost -> use fewer Proxies with more cores; 2 sockets/10-12 cores each deemed optimal
 4. Security considerations
 - a) Veeam changes Windows SAN policy to 'Offline(shared)' to prevent Windows bringing volumes online or writing to volumes

Veeam 9.5 VMCE ADO (VMCE-A) Study Guide

- 1) To do restores via Direct SAN, disable automount & enable the volume:
`diskpart > automount disable`
 - b) Disable Disk Mgmt via Group Policy: User Policy > Admin Templates > Windows Components > Microsoft Management Console > Restricted/Permitted Snap-Ins > Disk Management (enable)
 - c) Veeam deploys VMware VDDK into Proxy to coordinate Reads (& Writes for restores) with vSphere, so **VMware controls R/W streams**
 - d) Avoid excessive admin access to Proxies
 - e) Present LUNs as read-only LUN Masking/Zoning to Proxies on FC switches
 - f) If volume was brought online by mistake & has Windows signature, review **VMware KB1002168** & contact support
 - g) To access NFS, Proxies need in NFS export policy member for **R/W access**
5. Summary
- a) Use Direct Storage for fastest backup & reduced ESXi Host load
 - b) Consider adding VM Proxies for hotadd mode restores
 - c) Direct NFS is recommended for best NFS backup speed; no negative effects, and can be used on physical or virtual proxies
- b. Backup From Storage Snapshots
1. Minimizes VMware snapshot load
 2. Can help reduce RTPO
 3. Does not need access to source datastore
 4. Because CBT blocks need mapped to storage snap, processing can take longer
 5. For big environments, may be best to do a mix:
 - a) VMs with small change rates – don't use Backup from Storage Snap
 - b) VMs high change rates or big disks – use Backup from Storage Snap
- c. Virtual Appliance Mode
1. Recommended for high dynamic environments difficult to maintain access to new datastores, to reduce administrative overhead
 2. **Preferred mode for VMware vSAN**
 3. Deploy at least 1 Proxy VM per Cluster since Datastores are in Clusters typically
 4. For local Datastores, use 1 VM Proxy per ESXi Host; otherwise **nbd** is used
 5. NFS datastore considerations
 - a) Relies on LCK files in VM folders to determine VM ownership
 - b) During hotadd operations, the Host the Proxy is on takes temp ownership of VM being backed up & makes change to its LCK file, which can cause high STUN and make VM unresponsive
 - c) Ensure VM Proxy is on same Host VMs wanting to backup are on
 - 1) Give preference to VM Proxy: **HKLM\Software\Veeam\Veeam Backup and Replication\EnableSameHostHotAddMode, DWORD = 1or2**
 - 2) Value = **1**: if Proxy available, Veeam uses it; if busy, Veeam will use another Proxy
 - 3) Value = **2**: if Proxy available, Veeam uses it; if busy, Veeam will use **nbd**

Veeam 9.5 VMCE ADO (VMCE-A) Study Guide

- d) Hotadd with NFS typically requires many VM Proxies which some environments do not prefer; so if Direct NFS can't be used, just use **nbd**
- d. Network Mode
 - 1. Considerations
 - a) Maintain good network connection between ESXi VMkernel ntwk & Veeam
 - b) Veeam selects Proxies based on network subnet
 - 1) If Automatic selection is configured, Veeam selects Proxies in same subnet as ESXi VMkernel interfaces
 - 2) If Proxies are in different subnet than ESXi Hosts, manually select them so Proxies from other sites don't process data; select ones in the same site
 - 3) For multi-branch orgs, manually select Proxies per site for streamlining
 - c) Veeam reduces NBD NFC transport connections to 7 to protect VMware
 - 1) To increase connections: **HKLM\Software\Veeam\Veeam Backup and Replication\ViHostConcurrentNfcConnections, DWORD = # (higher than 7)**
 - 2) Increasing NFCs can cause failed jobs due to more ESXi RAM buffer usage; mitigate this by increasing ESXi buffer from 16,384 to 32,768 MB (**VMW KB2052302**)
 - 3) NBD only uses 40% of VMkernel bandwidth due to VMware throttling mechanisms on the VMkernel interfaces
 - 4) Slow backup & restore on 1GbE (only 20-30MB/s)
 - e. Backup I/O Control
 - 1. 'Enable Storage Latency Control' – **monitors** storage **read latency** using VMware ESXi Host **metrics taken every 20secs**
 - 2. Veeam starts processing disks residing on same datastore with a 40-60sec offset, which can degrade job performance

Veeam 9.5 VMCE ADO (VMCE-A) Study Guide

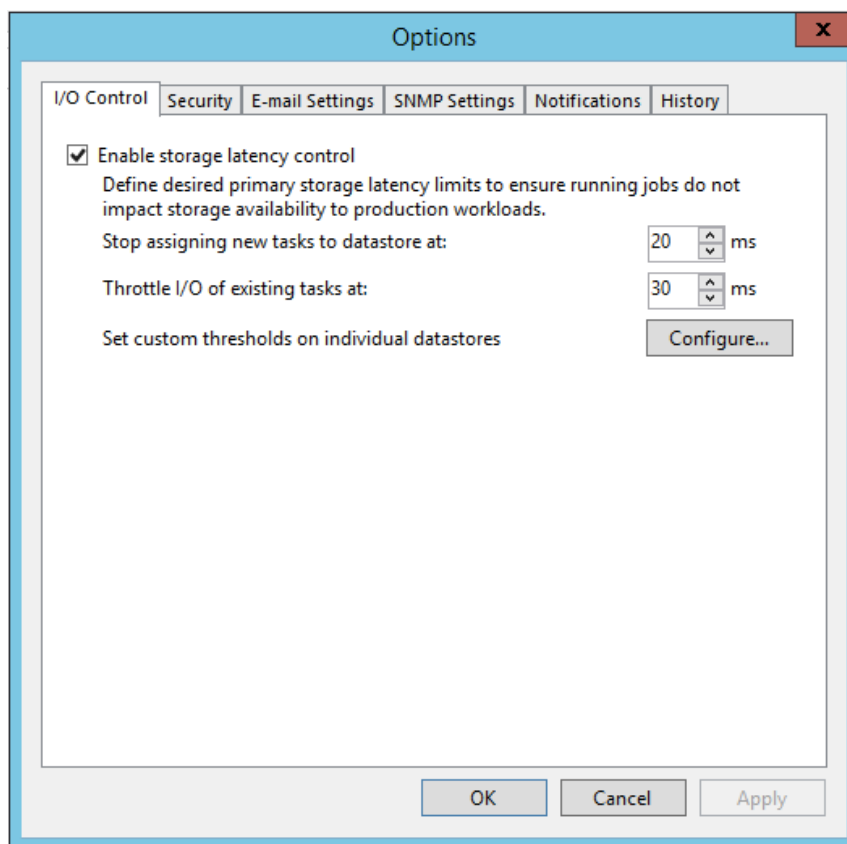


Figure 7, Backup I/O Control Configuration Window

6.4 – Proxy Server (Hyper-V)

a. Backup Methods

1. Online backup - recommended method for Hyper-V, with no VM downtime
 - a) Process for **MS Hyper-V 2008R2 / 2012** – native Microsoft Hyper-V approach
 - 1) Veeam interacts with MS Hyper-V Host VSS Writer & requests VM backup
 - 2) MS Hyper-V Host VSS Writer passes backup request to VM guest OS HV-IC
 - 3) HV-IC (VSS Requestor) requests backup of VSS-aware apps
 - 4) VM VSS Writers are instructed to quiesce app data
 - 5) VM VSS provider takes **internal VM snapshot**
 - 6) VM returns from read state to read/write state & internal snapshot passed to the VM HV-IC
 - 7) HV-IC notifies the MS Hyper-V Host VM is ready for backup & the Host VSS provider takes a Volume snapshot (**external snapshot**) of the Volume the VM disks resides on, & triggers auto-recovery process
 - 8) The Volume snapshot is presented to Veeam, which reads the data in either On-Host or Off-Host backup mode
 - 9) When the backup is complete, the Volume snapshot is deleted

Veeam 9.5 VMCE ADO (VMCE-A) Study Guide

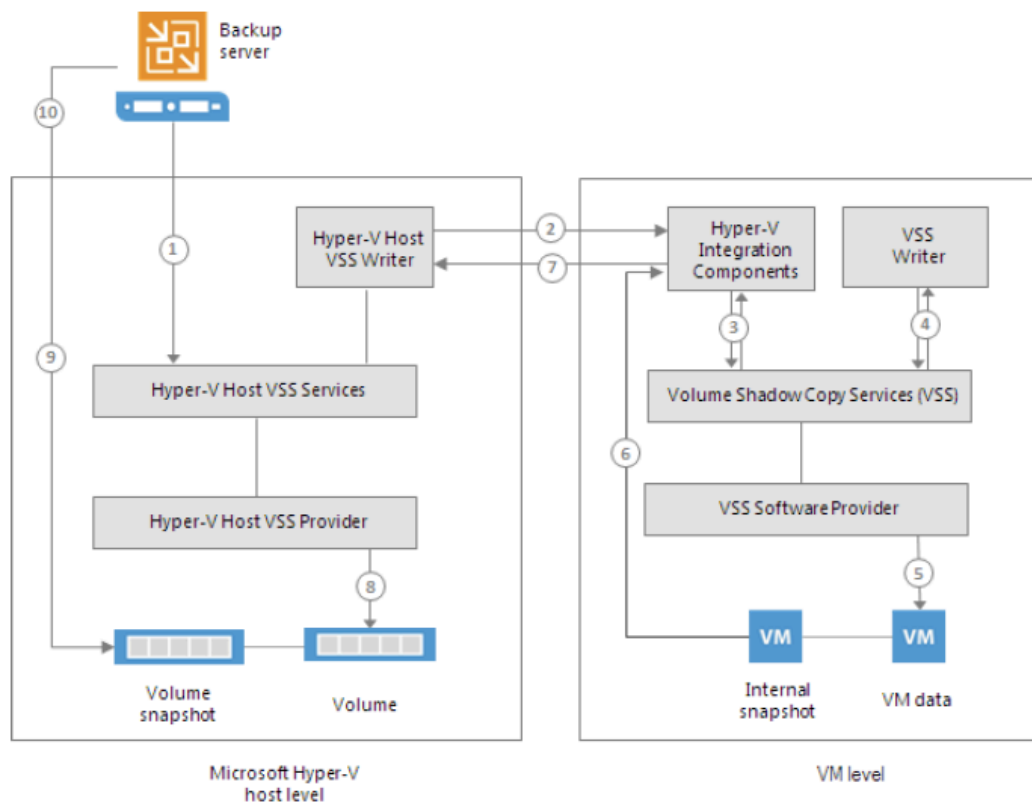


Figure 8, Hyper-V Server 2008R2 / 2012 Backup Process

- b) Process for **MS Hyper-V 2012R2** – uses backup & auto-recovery checkpoints
- 1) Veeam interacts with Hyper-V Host VSS Writer & requests backup of VM
 - 2) MS Hyper-V Host VSS Writer passes backup request to VM guest OS HV-IC
 - 3) HV-IC (VSS Requestor) requests backup of VSS-aware apps
 - 4) VM VSS Writers are instructed to quiesce app data
 - 5) After VM apps are quiesced, a backup checkpoint is created for each virtual disk & writes are directed to a **GUID.avhdx** file
 - 6) Then each virtual disk has an auto-recovery checkpoint created, named **GUID-autorecovery.avhdx**; (not dependent on backup checkpoint)
 - 7) The VM guest OS Internal Services Shadow Copy VSS provider takes an **internal snapshot** & passes it to the VM guest OS HV-IC
 - 8) HV-IC notifies the Hyper-V Host the VM is ready for backup
 - 9) The Host performs auto-recovery, mounting the checkpoint to the VM
 - 10) The MS Hyper-V Host VSS provider takes a Volume snapshot (**external snapshot**) of the Volume the VM disks resides on, & the auto-recovery checkpoint is merged with the VM
 - 11) Volume snapshot presented to Veeam & data read via On/Off-Host mode
 - 12) When the backup is complete, the Volume snapshot is deleted, & the backup checkpoint is merged with the VM

Veeam 9.5 VMCE ADO (VMCE-A) Study Guide

- c) Process for **MS Hyper-V 2016** – uses production checkpoints (pt-in-time image)
- 1) Veeam requests a production checkpoint of a VM via VSS or FS freeze
 - 2) Hyper-V Virtual Machine Management (VMM) Service (Windows VMs) or System Freeze (Linux VMs) brings VM guest OS data to a consistent state
 - 3) MS Hyper-V VMM Service creates production checkpoint (**avhdx**) file for each VM virtual disk & writes to this file
 - 4) On-Host Backup
 - i. Veeam reads VM disk data in read-only state, and the production checkpoint is merged with the VM after backup is complete
 - 5) Off-Host Backup
 - i. The Host VSS provider takes a Volume snapshot (**external snapshot**) of the Volume the VM disks resides on, and the production checkpoint is merged with the VM
 - ii. Veeam reads the data from the presented Volume snapshot
 - iii. When the backup is complete, the Volume snapshot is deleted

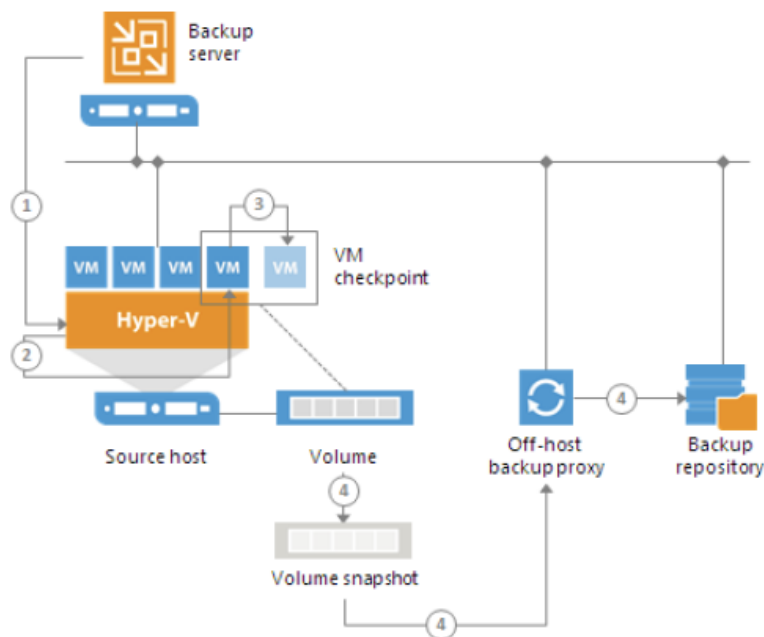


Figure 9, Hyper-V Server 2016 Backup Process

2. Offline backup (saved-state backup)
 - a) Native MS HV approach to quiesce VMs before taking Volume snap
 - b) Requires downtime! VM forced to a 'saved state' (hibernation) to create a stable image; i.e. not transactionally consistent
 - c) **This method not applicable to VMs on Hyper-V Server 2016**
 - d) Process:
 - 1) Veeam interacts with MS Hyper-V Host VSS & requests backup of a VM
 - 2) Host VSS Writer forces VM to saved state
 - 3) The VM guest OS hibernates & CPU/RAM content is written to dump file

Veeam 9.5 VMCE ADO (VMCE-A) Study Guide

- 4) Host VSS Provider takes Volume snapshot & VM returns to normal state
- 5) Volume snap is presented to Veeam & reads data
- 6) When the backup is completed, the volume snap is deleted
- e) If online method can't be used, it is recommended to use Veeam Crash-Consistent method instead of this (offline) method

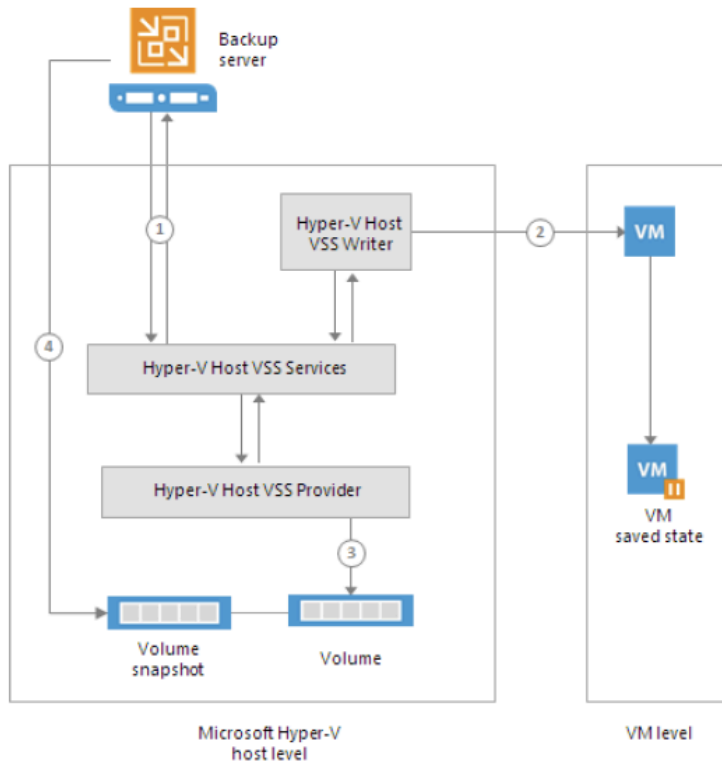


Figure 10, Offline Host Backup (notice pause icon in 'Saved State' VM)

b. Backup Modes

1. On-Host backup

- a) VM data processed on source Host VM resides, which may produce high CPU/network overhead to the source HV Host
- b) Rules for Assigning On-Host backup role in CSVs
 - 1) **MS Hyper-V 2012/ 2012 R2** – when MS CSV Software Shadow Copy Provider is used to create snapshots, the On-Host Proxy role is assigned to the Host owning the CSV
 - 2) **MS Hyper-V 2008 R2** – when the VSS software or hardware provider is used to create snapshots, the role is assigned to the Host the VM is registered

2. Off-Host backup

- a) Dedicated Host used as a Proxy & acts as a data mover, retrieving data from source VM volume, processes it, then transfers it to the destination
- b) Benefits:
 - 1) Processing load not on source Hyper-V Host

Veeam 9.5 VMCE ADO (VMCE-A) Study Guide

- 2) Control of Backup/Replication job traffic
- c) Uses transportable shadow copies to create a Volume snapshot the VM resides on & mounts the snap on another Host in the same subsystem (SAN)
- d) Requirements:
 - 1) **Must be physical**; VM with Hyper-V role enabled can be used for testing
 - 2) Source Hyper-V Host & Off-Host Proxy must be same Hyper-V version
 - 3) In backup job properties: Off-Host method & Host used must be selected
 - 4) Source Hyper-V & Off-Host Hosts must connect to shared storage via SAN
 - 5) To manage Volume Shadow Copies, a VSS Hardware Provider (supplied by storage vendor) must be installed on Source Host & Off-Host Proxy; Provider must support transportable shadow copies
 - 6) If VM disks reside on a CSV with dedup enabled, the Off-Host Proxy must be 2012R2 or 2016 with dedup enabled
 - 7) The Off-Host Proxy can't be part of a Hyper-V Cluster using the same CSV
- e) Process:
 - 1) Veeam triggers a Volume snapshot from source Hyper-V Host
 - 2) The snapshot is detached from source Host & mounted to Off-Host Proxy
 - 3) The Off-Host data mover retrieves VM data from the mounted Volume snap, processes it, then transfers it to the destination
 - 4) After processing, the Vol snap is unmounted from the Proxy & deleted
- c. Backup I/O Control
 - 1. Same as with VMware, but **metrics taken every 10secs**
 - 2. Be aware of the same offset issue as noted in VMware section

6.5 – Repository Server

- a. General Info
 - 1. Server that runs Transport Service & provides a destination folder for backups
 - 2. Stores: Backup files, VM copies, & Replica metadata
- b. Requirements
 - 1. Consideration factors:
 - a) Capacity
 - b) Write performance
 - c) Read performance
 - d) Data density
 - e) File security
 - f) Backup file utilization
 - 2. Characteristics a Repository should be:
 - a) Should be resilient as it'll be hosting essential data
 - b) Should be scalable, allowing the backup to grow
 - 3. Org's retention requirements may differ; thus, different Repositories can be used
 - a) High-performance Repository for recent retention & instant restores
 - b) Slower, more capacity for long-term retention (i.e. Backup Copy job use)

Veeam 9.5 VMCE ADO (VMCE-A) Study Guide

c. Impact On Destination Storage (I/O):

METHOD	IMPACT ON STORAGE I/O
Active full	1 write I/O
Forward incremental	1 write I/O
Forever fwd incremental	3 I/O – 1 read/2 writes
Reverse incremental	3 I/O – 1 read/2 writes
Synthetic full	2 I/O – 1 read/1 write
Synthetic full with transform	4 I/O – 2 reads/2 writes

Figure 11, Repository Storage I/O Impact

d. Considerations

1. Windows or Linux Repository?
 - a) Main difference is in how each handles NAS: SMB vs NFS
 - b) Linux can generally handle more throughput
 - c) For smaller environments, best to keep all same OS to avoid complexity
 - d) Other considerations involve cost and administration burdens
2. Physical vs virtual
 - a) Can use either, but keep in mind storage & transport media will be heavily occupied; as such, best to use physical
3. SMB/CIFS Gateway server
 - a) Auto selection – may be good when there are multiple Proxies
 - b) Manual selection
 - 1) Helpful when SMB share is remote to avoid synthetic tasks over the WAN
 - 2) Always recommended to use a server close to the SMB share
 - 3) Helps keep data flow under control
 - 4) Can configure multiple Proxies pointing to the SMB share to have multiple streams, potentially increasing performance
4. Load balancing (automatic selection)
 - a) Only 1 Gateway Server per Backup chain is used to write data
 - b) For Per-VM, multiple Gateway Servers may be started & used concurrently
 - c) 1 Proxy should have enough resources to run multiple Gateway instances
 - d) Best practice is to use a mid-size/Enterprise-class NAS with good I/O perf
 - 1) If low-end NAS is used, it is recommended to use a block-based storage protocol (i.e. iSCSI)
5. Repository cache
 - a) Target data mover caches metadata from backup files instead of requesting from those files multiple times over the course of a job run, improving

Veeam 9.5 VMCE ADO (VMCE-A) Study Guide

- performance of storage with poor random I/O performance & reducing overall storage load
- b) **ONLY ENABLED WHEN USING 64BIT OS**; & increases RAM by 2GB per job
- c) Using Auto Selection for Gateway Server for Data Domain & HPE StoreOnce disables the cache
- d) Disabling built-in dedup disables cache
- 6. Cloud Repository
 - a) Limitations:
 - 1) Backup Copy jobs cannot use a Cloud Repository as a Backup job source
 - 2) Transaction log backup not supported on targeted Cloud Repositories
 - 3) Backups in Cloud Repository cannot be used for Instant VM Recovery, Multi-OS File-Level Recovery, or Restore to Azure
 - 4) Parallel processing limit is set by SP & applies to **direct transport mode** only; built-in WAN Accelerators still process disks **sequentially**
 - e. Configuration Guidelines
 - 1. Block alignment – job’s Advanced button > Storage tab, ‘Storage Optimization’
 - a) Default setting = “**Local (1MB)**” for Backups ; “**LAN (512KB)**” for Replicas
 - 1) With 2x Veeam compression, the write would be 512KB (of 1MB) or less
 - b) Configure storage stripe size close to block size to mitigate RAID write penalty; always check with storage vendors for recommended configs
 - 1) In the ‘field’, **256KB** & **512KB** have shown to be optimal stripe size values
 - 2. File system formats
 - a) **Set Cluster/Allocation Unit Size to 64KB** to mitigate fragmentation
 - 3. Large file switch for NTFS
 - a) On NTFS vols with dedup, especially on Win2012R2: **format fs:ntfs /L**
 - 4. Keep Backup file size under control – backup chains (Full+Incr) < **10TB**
 - f. Scale-Out Backup Repository (SOBR) – available since Veeam v9
 - 1. File-placement policies
 - a) Performance – separates .vbk files from .vib files
 - 1) Provides good performance for merging operations (synthetic, Reverse, FFwd), generating reads/writes IOs on differentiated storage
 - b) Data locality – job tries to keep all backup files on same extent
 - 1) Provides safest data placement; great performance if used properly
 - c) Policies are not strict; if extent is inaccessible, Veeam uses another extent with enough space to house backup file(s)
 - d) Veeam reserves 1% storage on each extent for metadata (.vbm) & merge operations
 - 2. Scale-Out Backup Repository (SOBR) and LAN
 - a) “Master” & “slave” (auxilliary) data mover
 - b) SOBR imposes heavy LAN traffic for Proxy load balancing & master/slave data mover communication

Veeam 9.5 VMCE ADO (VMCE-A) Study Guide

3. Extent load-control settings
 - a) If extent reaches max concurrent task limit, it won't be used as target, possibly leading to unwanted file distribution across extents
 - b) **It's safe to uncheck this setting as long as extent h/w limits aren't reached**, & let limits be set by number of parallel tasks assigned to the Proxy
4. Veeam Extent selection order (*also used for Backup Copy jobs*):
 - a) Extent availability (accessibility)
 - 1) If extent & backup files not accessible, Veeam triggers a full
 - b) SOBR placement policy
 - c) Load control settings (extent max concurrent tasks)
 - d) Extent with most free space
 - 1) Veeam estimates space for full to be **50%** of source VM data
 - 2) Veeam estimates space for increments to be **10%** of source VM data
 - 3) Reverse: Veeam estimates **10% for increm & additional 10% for full**
 - e) Availability of file in current backup chain – extents housing increment files have higher priority
5. Performance policy extent selection
 - a) Veeam stores .vbk's & .vib's *in same backup chain* on separate extents
 - b) The above selection conditions are also applied here

6.6 – Per-VM Backup Files

- a. Default setting for SOBR; needs manually enabled for standalone Repositories
- b. Using this may reduce dedup ratio to nil, while compression ratio still intact
 1. Leverage BitLocker for space reclamation on NTFS as dedup compensation

6.7 – Deduplicaton Appliances

- a. Considerations
 1. Suggested to use as secondary storage (long retention) not restore performance
 2. Since Veeam transform processes (synthetic full, reverse increm, fwd increm) are random I/O-based, it is recommended to use dedup appliances with **Active Fulls** with a window long enough to perform those fulls
 - a) **NOTE**: EMC Data Domain DDBoost & HPE StoreOnce Catalyst technologies are designed to handle synthetic (full)/transform operations
- b. Primary Backup Repository
 1. Unless using Exagrid dedup appliance that has Landing Zone feature, restore operations will not be as fast (Exagrid restore is only fast for latest restore point)
 2. Since Veeam v9, use per-VM with dedup appliances
- c. Backup Copy Job (BCJ) Repository
 1. BCJs apply transform operations to backup chains by design, so additional job processing will occur; yet, if GFS is configured, it will save storage space
 - a) If using Data Domain or StoreOnce, additional time is minimal

Veeam 9.5 VMCE ADO (VMCE-A) Study Guide

2. Archive Backups (GFS)
 - a) **Active full** is recommended to reduce load on appliance & improve BCJ performance
 - 1) BCJs use Synthetic by default, which uses random I/O, & Dedup Appliances are optimized for sequential I/O
 - 2) Rehydration/Dehydration, then writing synthetic tasks via random I/O severely degrades performance, thus why Active Full is recommended
 - b) Select 'Read The Entire Restore Point From Source Instead of Synthesizing From Increments' option to enable Active Fulls for GFS Archives
 - c) Network load will be higher, but dedup appliance performance will increase
 - 1) If BCJ data transfer does not fit the interval due to the network load, the BCJ interval is extended
- d. Dedup As a Virtual Tape Library (VTL) Container
 1. Recommended job compression config to **None or Dedup-Friendly**
 - a) Otherwise, job would write compressed data making dedup less efficient
 2. Tasks not supported with tape jobs:
 - a) SureBackup
 - b) Instant VM Recovery
 - c) Direct VM Restore
 - d) Direct Guest Object Restore
- e. Optimizing for Non-Integrated Dedup Storage
 1. Set the 'Align Blocks' option for dedup systems with fixed block size (4KB block)
 - a) Setting is good for dedup, but could lead to fragmentation & unused space
 2. Compression can be configured for low-bandwidth links, but the Gateway server must be configured to decompress data to enhance dedup rates
- f. Achieve Parallel Write Operations by:
 1. Increase number of jobs to the dedup appliance to generate more write streams
 2. Leverage multiple CIFS Gateway Servers for more TCP-IP connections
 3. Leverage per-VM backups
 4. **NOTE:** be weary of the max concurrent streams each dedup appliance supports
- g. Restores with Dedup appliances
 1. Veeam Explorers perform intense read operations, thus making long FLR wait time; as such, a full VM restore may be quicker as it generates sequential reads
 2. Recommended to use Veeam file-indexing so restore can be done via Enterprise Manager (EM); restore performance won't be enhanced, but user experience will as wait time will be less

Veeam 9.5 VMCE ADO (VMCE-A) Study Guide

h. Job Tuning for Dedup Appliances

SETTING	RECOMMENDATION
Compression	Set to Optimal in job, & configure 'decompress on target' in Repository Advanced properties if low-bandwidth link
Veeam inline dedup	Performance can be increased by disabling this. But, if CBT fails, Veeam will have to stream all data; thus, <u>if current performance is satisfactory, leave enabled</u>
Block size	To limit metadata & optimize sequential writes, may want to set to local target (4MB), for faster backups but produce slower restores. Optionally, smaller block produces slower backup but faster restores
Encryption	Because of heavy effect on dedup rates, this is not recommended

Figure 12, Deduplication Appliance Job Tuning Settings

i. Vendor-Specific Recommendations

1. EMC Data Domain DDBoost capabilities:
 - a) High dedup between Gateway & DD appliance, since dedup is at source
 - b) Better LAN parallelization
 - c) Seamless Veeam transforms
 - d) Supports only **59 restore points** per backup; **90 write streams**
2. ExaGrid
 - a) Since Veeam **v8**, Exagrid can be its own Data Mover, so no Gateway needed
 - b) Jobs are pre-configured with recommended Exagrid & Veeam settings
 - c) **10 max concurrent tasks** allowed for a whole appliance
 - 1) Per-VM – config all jobs to total 10 concurrent tasks to the appliance
 - 2) Non per-VM – limit concurrent tasks to appliance to be **1**
 - d) For Veeam v8 and older, create Exagrid Repository per Backup job to avoid jeopardizing dedup ratio & to utilize parallel processing
 - e) Landing Zone should have **125% of a full** backup for restores (IR, restore, etc)
3. HPE StoreOnce
 - a) Jobs are pre-configured with recommended StoreOnce & Veeam settings
 - b) If using Catalyst, set a couple Veeam registry keys:
 - 1) If Share is not set as 'Distant': **HKLM\Software\Veeam\Veeam Backup Transport\UseLowBandwidthMode = 1**
 - 2) If Share set as 'Distant':
HKLM\Software\Wow6432Node\Veeam\VeeamBackupTransport\PayloadChecksumsDisabled | BodyPayloadCompressDisabled ; both = 1

Veeam 9.5 VMCE ADO (VMCE-A) Study Guide

6.8 – WAN Accelerator

- a. Source
 - 1. Default TCP threads = **5** ; max = **100**
 - a) Setting TCP threads is 1st step & only set for the source WAN Accelerator ensuring different source Accelerators can have different TCP thread settings
 - 2. Formula for setting threads: **link speed x 1.5**; test using WANem
- b. Target
 - 1. Cache size set on target & applies to **pair** of Accelerators
 - 2. Recommended to configure 10GB cache per guest OS processed
 - a) Multiple Linux OS types are considered as 1
 - 3. Recommended to use cache population upon config completion
 - 4. Recommended to seed the cache to further reduce WAN saturation

6.9 – Tape

- a. Tape Server/Drives/Drivers
 - 1. Tape server – any Windows server managed by Veeam
 - a) For small environments can be the Veeam server; but If backup to tape will be used a lot, use dedicated server as tape traffic loads the tape server
 - b) Use the latest Veeam version & Server OS for tape server as they will have the latest tape speed optimizations & invoke the best performance
 - c) Supported tape drives – **LTO-3 and higher**
 - d) Driver settings:
 - 1) IBM – use non-exclusive setup & start .exe with admin rights
 - 2) HPE – not installable with downloadable install.exe; instead run install.exe & choose extract. Then use Device Manager > Update driver
 - 3) If unable to test speed, assume jobs with LTO5/6 = 40MB/s
- b. Tape & Backup Mode – depending on the backup mode, backup to tape either copies files to tape or synthesizes a full backup
 - 1. Reverse incremental:
 - a) Standard media pool – job will backup the full & ignore the increments (.vrb)
 - b) GFS media pool – job will synthesize full with .vrb on specified days
 - 2. Fwd increm with active/synth fulls – backup chain on tape will be a copy of the backup chain (full+increm) on disk; virtual full option is ignored
 - 3. Fwd increm withOUT active/synth fulls – full files synthesized from existing RPs on disk; virtual full option can be configured
- c. Encryption
 - 1. Hardware encryption used by Veeam if provided by tape device & enabled
 - a) If encryption used, & backups already encrypted, backups are encrypted twice
 - b) Restoring double-encrypted files on same Veeam server decrypts automatically; from another Veeam server, pwds need entered twice

Veeam 9.5 VMCE ADO (VMCE-A) Study Guide

2. If h/w is not supported, **256 AES software encryption** used
 - a) Veeam generates encryption key stored with data on tape
- d. Media Pool / Set / Vault
 1. Media pool – group of tapes managed by Veeam
 - a) Service media pool – created & managed automatically; cannot be created manually
 - 1) Empty media – can be used with other Pools
 - 2) Unknown media – placed in Unrecognized Pool to not overwrite
 - 3) Some content media – placed in Imported Pool & overwriteable
 - 4) Exhausted/broken media – placed in Retired Pool & no longer used
 - b) Custom tape media pool – created manually, & to which data can be written
 - 1) Can create as many as needed
 - 2) Media can be assigned manually or automatically
 - 3) Configure each pool according to its purpose (i.e. overwrite protection)
 - 4) Can be spanned over multiple tape libraries (Veeam v9 and above)
 - c) Custom GFS media pool – manual pool used for wkly, mntly, qrtly, yrly data
 - 1) Same options as Custom Pool
 2. Media set
 - a) Media pool subset containing backups; if 1 full is on set, it is ‘self-sufficient’
 3. Media vault
 - a) Used to organize offline media
 - b) Export tapes to vault manually or automatically via tape job or media pool properties

6.10 – Veeam Backup Enterprise Manager

- a. When to Deploy Enterprise Manager (EM)
 1. If wanting to use encryption for Backup/BCJ
 2. For org’s that have ROBOs
 3. For org’s wanting first-line support to perform in-place restores (delegation)
 4. For deployments spanning multiple stand-alone Veeam instances (licensing)
 5. Automation; when access to Veeam RESTful API is needed
 6. If EM is not deployed, password loss protection won’t be available

6.11 – Indexing

- a. Use-case (when to implement)
 1. Only when wanting to do advanced searches or use 1-click file level restore
- b. Sizing Veeam Catalog
 1. **2MB per 1M files** per VM on Enterprise Manager server
 - a) **10MB per 1M files** required **on Windows** guest VM; **5MB for Linux**
 2. **Veeam Catalog Service** responsible for maintaining index data for all jobs on each Veeam server, then replicated to EM

Veeam 9.5 VMCE ADO (VMCE-A) Study Guide

3. Standard licensing – index data kept for only RPs still on the Repositories
 4. Ent/Ent+ licensing – index data a default 3mos, or longer
- c. Example:
1. 2 VMs, 10M files each, backed up 2x/day for 30 days for 3mos
 2. Windows guest VM = 100 MB space needed
 3. Linux VM = 50 MB space needed; Linux needs ½ due to **mlocate** not indexing metadata
 4. Enterprise Manager = 20 MB x 60 restore points = 1.2GB x 3mos = 3.6GB total

6.12 – Anti-Virus

- a. Folder exclusions for A/V solutions:

TYPE	EXCLUSIONS
Folders	C:\Program Files\Veeam C:\Program Files (x86)\Veeam C:\Program Files\Common Files\Veeam C:\Program Files (x86)\Common Files\Veeam C:\ProgramData\Veeam C:\VeeamFLR* All Veeam Repository \ WAN Accelerator folders VBRCatalog (HKLM\Software\Veeam\Veeam Backup Catalog) NFS (HKLM\Software\WOW6432Node\Veeam\Veeam NFS)
Folder on VM guest OS (VSS only)	C:\ProgramData\Veeam\Backup\VeeamVssSupport
Files	VeeamAgent.exe VeeamAgent64.exe \.vmdk*

Figure 13, Anti-Virus Exclusions

6.13 – Failover Plans

- a. Start a group of VMs automatically, with a time delay between each VM in the list
- b. Limitations:
 1. Can only start 10 VMs simultaneously; more than 10 VMs can be added to a group, but the 11th VM won't start until at least 1 of the 1st 10 VMs has started

Veeam 9.5 VMCE ADO (VMCE-A) Study Guide

- c. Finalizing
 1. Undo failover – reverting to orig VMs & discarding all changes; can be done on whole group **in groups of 5** with a 10sec interval between each group of 5
 2. Permanent failover
 3. Failback – VMs must be done individually; 3 failback options (orig, restored, etc)
- d. Options
 1. Scripts supported: .bat , .cmd , .ps1 , .exe
 2. VMs are started in order in the list, so rearrange VMs if needed (i.e. DNS , DC)
 3. Time delay – time Veeam waits before starting the next VM in a list; otherwise Veeam will start all VMs simultaneously (in groups of 10)

SECTION 7 – DESIGN SCENARIO: PART 2

7.1 – Scenario

- a. Design adaptation to customer changing environment
- b. Datacenters
 1. DC1
 - a) See Handbook
 2. DC2 – colocation facility
 - a) See Handbook
- c. Updated Data Protection Requirements
 1. See Handbook for requirements

SECTION 8 – AUTOMATION

8.1 – Veeam Backup Server

- a. Unattended Install
 1. Components must be installed in strict order depending on server install type
 2. Veeam Server:
 - a) Prerequisites:
 - 1) MS SQL Server 2005 or greater
 - 2) MS Visual C++ 2010 SP1 Redist Pkg
 - 3) MS SQL Server 2012 System CLR Types
 - 4) MS SQL Server 2012 Mgmt Objects
 - b) Install order:
 - 1) Veeam Backup Catalog
 - 2) Veeam Server

Veeam 9.5 VMCE ADO (VMCE-A) Study Guide

- 3) Veeam Explorers:
 - i. Veeam Explorer for AD
 - ii. Veeam Explorer for Exchange
 - iii. Veeam Explorer for SQL
 - iv. Veeam Explorer for Oracle
 - v. Veeam Explorer for Sharepoint
3. Veeam Console – just install the Console
4. Veeam Enterprise Manager:
 - a) Prerequisites:
 - 1) MS SQL Server 2005 or greater
 - 2) MS Visual C++ 2010 SP1 Redist Pkg
 - 3) MS SQL Server 2012 System CLR Types
 - 4) MS SQL Server 2012 Mgmt Objects
 - 5) IIS Components:
 - i. Default Document
 - ii. Directory Browsing
 - iii. HTTP Errors
 - iv. Static Content
 - v. Windows Auth
 - 6) Update 4.0.3 for .NET 4.0 (MS KB2600211)
 - b) Install order:
 - 1) Veeam Backup Catatlog
 - 2) Veeam Backup Enterprise Manager
- b. **Unattended Upgrade**
 1. **Not supported**; just uninstall current version then install new version in unattended mode
 2. Win 2003 Servers are not supported in Veeam v9; if they are connected to Veeam, disconnect prior to upgrade or install will fail
- c. **Unattended Update** Install
 1. **Supported**, via script running in elevated privileged mode
 2. Script: %patch% [/silent] [/noreboot] [/log] [VBR_AUTO_UPGRADE=1]
 - a) “patch”(.exe path) & “silent” are required; see below for remaining options
 3. Script example: C:\Temp\VBR_update.exe /silent /q /log
C:\Support\Veeam.log /c /t <path> VBR_AUTO_UPGRADE=1

Veeam 9.5 VMCE ADO (VMCE-A) Study Guide

OPTION	PARAMETER	REQUIRED	DESCRIPTION
%patch%	path	yes	Path to update executable (on VBR or share)
/silent	-	yes	Sets user-level interface to no
/q	-	yes	User interaction = 'no'
/c	-	yes	Auto-extract update file; if not added, UI mode is used
/t	path	yes	Used with /c; Path to temp folder of extracted update files
/noreboot	-	no	Suppresses reboot
log	path	no	Specified path to update log file location
VBR_AUTO_UPGRADE=1	boolean	no	Veeam components update automatically after VBR server Backup Service is started

Figure 14, Unattended Update Install Script Options

4. Install results

- a) To verify if update completed successfully, check the last exit code via CMD PROMPT or PoSH: `cmd.exe = %ERRORLEVEL%` ; `PoSH = $LastExitCode`
- b) To verify component update success, use the Veeam Console

8.2 – VMware vSphere Tags

- a. Policy-Driving Availability
 1. Tag advantages:
 - a) An object can have multiple tags
 - b) Tags can be applied by any user with sufficient privileges
 - c) Searches/filters can be done with tags in single mode or boolean operators
 - d) Tags are immediate
- b. Apply Tags Manually
 1. Time-consuming
 2. Error-prone
 3. Environment size = administrative effort

Veeam 9.5 VMCE ADO (VMCE-A) Study Guide

- c. Apply Tags with Automation
 - 1. Accurate
 - 2. Quick
 - 3. Frees up admin time
 - 4. Can use multiple tools – PoCLI, vRA, vRO, Puppet, Chef, Veeam ONE
- d. Apply Tags with Veeam ONE – sync vSphere tags via VONE Business View
 - 1. Synchronize groups with tags
 - a) Configure VONE Biz View by **enabling ‘update of tags’ option** to translate & write its categorization data to vCenter Server
 - 1) A **manual sync can be done via VONE Reporter** through Collection job > Object properties
 - 2. How objects are tagged
 - a) Category – assigning Biz View category value maps to a vCenter tag category
 - b) Group – assign a Biz View tag value which maps to a vCenter tag
 - c) Upon data collection, Biz View detects virtual infrastructure groups & tags them on vCenter side
 - 1) If there are no matching tags/tag category on vCenter, Biz View creates the tag/category
 - 2) If a Biz View group has no objects, no tag/category is written to vCenter
 - 3. Enabling updates of tags
 - a) **vSphere 5.5 or higher is required**
 - b) Remove default/predefined Biz View groups/categories BEFORE enabling to prevent needless data being transferred to vCenter
 - c) Biz View > Configuration tab > Options (left menu) > Operation Mode section > **select ‘Use vSphere Tags’**
- e. “No Backup” Tag
 - 1. Ensures VMs with this tag do not require protection (backup)
 - 2. Other VMs without ‘no backup’ tag then need evaluated why not backed up
 - a) VONE alarms can be configured to notify of VMs with no restore points (RPs)

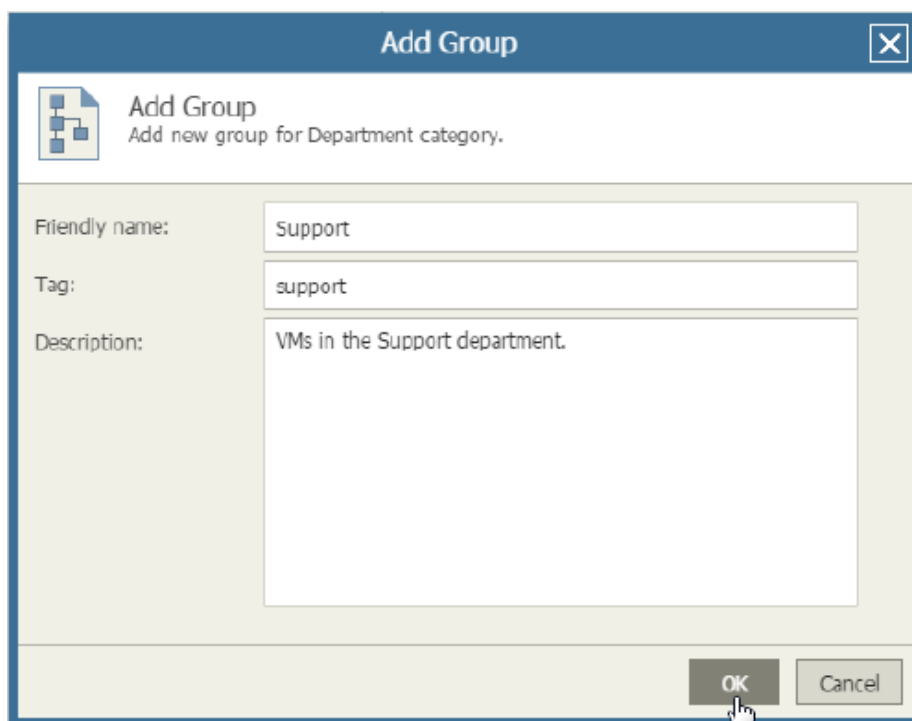
SECTION 9 – AUDIT AND COMPLIANCE

9.1 – Categorization

- a. Creating & Managing Groups – Veeam ONE Business View allows categorizing virtual environment from a business standpoint rather than technical; presents data in business terms; every category created must be assigned at least 1 group
 - 1. Static groups – objects to administer explicitly; objects may/may not share common attributes

Veeam 9.5 VMCE ADO (VMCE-A) Study Guide

2. Dynamic groups – objects administered/updated automatically during data collection by Business View; based on search expressions that share common attributes; **can only contain Hosts, Storage, & VM objects**, not Clusters
- b. Creating Static Groups
1. Veeam ONE Business View > Configuration tab > Groups, via menu on left
 2. Choose a category (shown as tabs)
 3. Click 'Add' at top of group list
 - a) Specify Group Name
 - b) Specify vCenter Server tag to be matched to the group (required, even if not mapping vCenter tags to Biz View groups)
 - c) Enter Description if desired, then OK; repeat for additional Groups



The screenshot shows a dialog box titled "Add Group" with a close button in the top right corner. Below the title bar, there is a header section with a tree icon and the text "Add Group" and "Add new group for Department category.". The main area of the dialog contains three input fields: "Friendly name:" with the value "Support", "Tag:" with the value "support", and "Description:" with the value "VMs in the Support department.". At the bottom right of the dialog are "OK" and "Cancel" buttons, with a mouse cursor clicking the "OK" button.

Figure 15, Static Group Creation, Veeam ONE Business View

- c. Creating Dynamic Groups
1. Veeam ONE Business View > Configuration tab > Groups, via menu on left
 2. Choose a category (shown as tabs)
 3. Click 'Grouping Expression' at top of group list
 4. Enter an expression Biz View must use to create groups & distribute virtual infrastructure objects into the group
 5. Repeat for additional dynamic groups

Veeam 9.5 VMCE ADO (VMCE-A) Study Guide

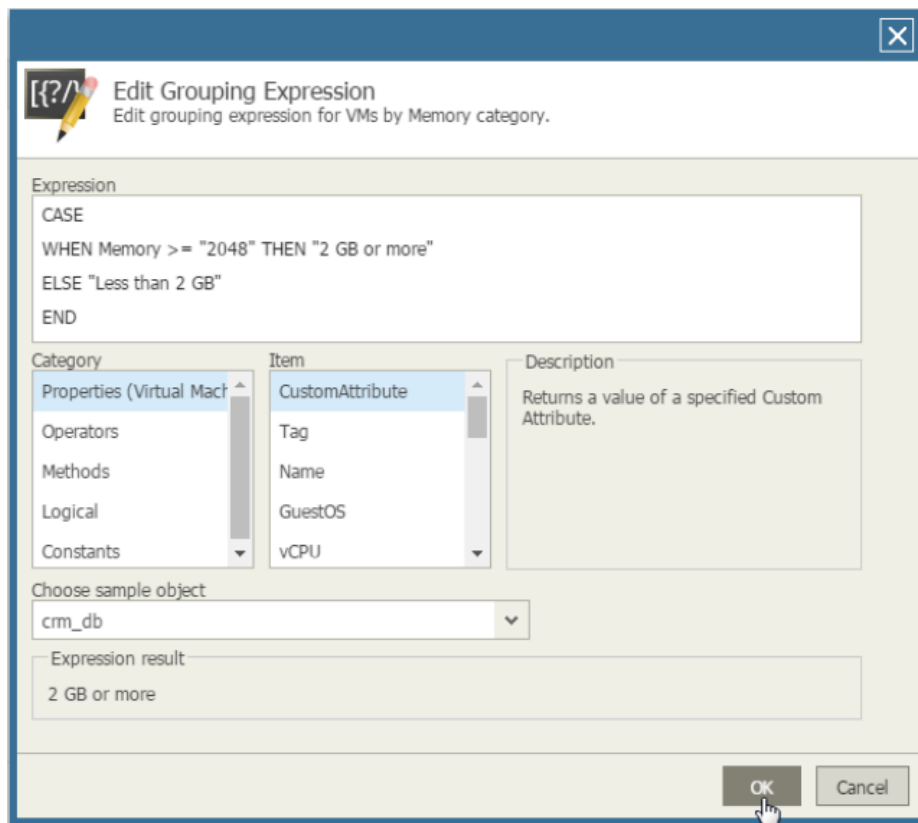


Figure 16, Dynamic Group Creation, Veeam ONE Business View

9.2 – Audit

- a. Provides IT Depts assurance of adequate protection; helps reduce risk of data loss through intentional/unintentional tampering or reconfiguration
- b. Infrastructure Changes Audit Report
 1. Helps identify recent **changes made in the virtual infrastructure**
 - a) **Top 10 Host Changes, Top 10 VM Changes, Top 10 Datastore Changes**
 - b) **Number of Changes Made (day)**
 - c) **Modifications Per User**
- c. Job Configuration Change Tracking Report
 1. Helps audit **changes made in the backup environment**
 - a) **Job Modifications by User**
 - b) **Job Modifications by Day**

9.3 – Compliance

- a. Protected VMs Report – helps validate your organizational backup protection state
 1. Report charts
 - a) **Protected VMs**

Veeam 9.5 VMCE ADO (VMCE-A) Study Guide

- b) **VM Last Backup State** – failed, success, warning
- c) **VM Last Backup Age** – no restore points, outside RPO, within RPO
- 2. Parameters
 - a) Scopes – Folders (VMware only;) virtual infrastructure
 - b) vCD objects
 - c) Business View objects
 - d) RPO – max amount of data acceptable to lose; age of latest Backup/Replica
 - e) Exclusion mask –VM exclusions; explicit entry or via wildcard (*), separated by semicolon (e.g. *_R&D)
 - f) Job type – Backup, Replica, BCJ
 - g) Include (or not) VM templates in report
- b. Alarms – to notify about missing or out-dated Backups, according to RPO req's
 - 1. Process:
 - a) VeeamONE Monitor > bottom of left Inventory Pane > Alarm Management
 - b) In Alarm tree on left > under VMware or Hyper-V > Virtual Machine
 - c) In Alarm Management pane, double-click **VM No Backups** pre-defined alarm
 - d) General tab > click to 'Enable' alarm
 - e) Rules tab > click txt box to configure rules:
 - 1) Rule Type: i.e. RPO
 - 2) Interval: i.e. 24hrs (or mins/days)
 - 3) Job Type: i.e. Backup (or Replication)
 - 4) Severity: Warning (or Resolve, Information, Error)

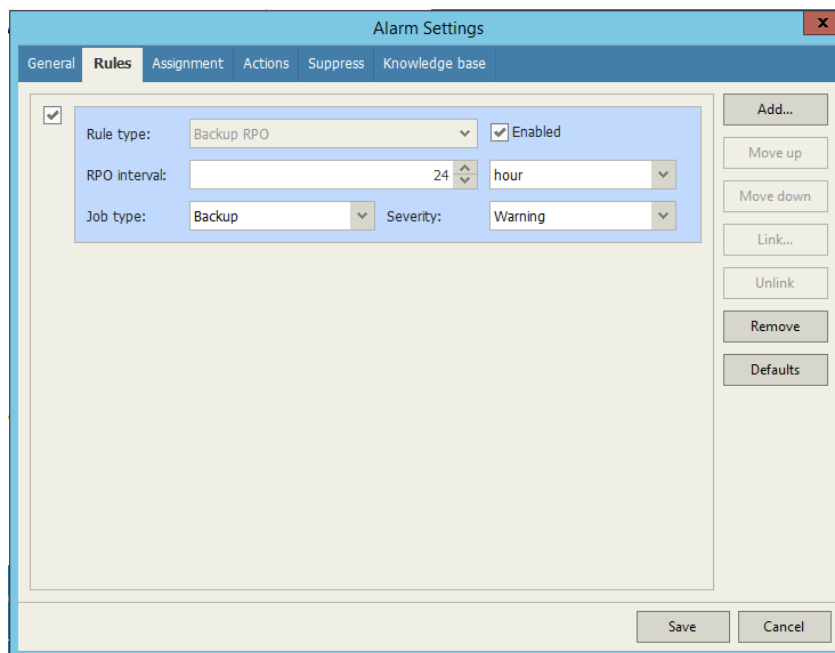


Figure 17, Alarm Configuration Settings

- f) Assignment tab > select listed item & 'Remove'
 - 1) Click 'Add' and browse VI to add VMs

Veeam 9.5 VMCE ADO (VMCE-A) Study Guide

2. Result
 - a) The first alarms will be triggered by Alarm interval setting (i.e. 24hrs)
 - b) **If no backups created, the VM's creation time will be the latest backup time**
- c. VM Backup Compliance Overview Report – returns list of VMs not having minimum number of backup copies; as well as mission-critical VMs meet 3-2-1 Rule
 1. **VM Compliance Status** – Compliant VMs, Non-Compliant VMs
 2. Same parameters as Protected VMs Report, but main param = **Req'd # of Copies**

9.4 – SureBackup & SureReplica

- a. Custom Roles
 1. Server roles – pre-defined test sets, located in .xml files at:
%ProgramFiles%\Veeam\Backup and Replication\SbRoles
 2. VM tests – heartbeat signal via VMware Tools, ping via VM Netw, or app script
 3. To create custom role:
 - a) Copy a .xml file from SbRoles folder
 - b) Modify tag using an UUID generator
 - c) Required parameters = **SbRoleOptions , Role , SbRole , Id , Name , Options , SbVerificationOptions , HeartbeatEnabled , PingEnabled**

```
<SbRoleOptions>
  <Role>
    <SbRole>
      <Id>d307ca96-4ca6-42d1-8fe7-16b03700e222</Id>
      <Name>Domain Controller (Authoritative Restore)</Name>
      <ApplicationName>Active Directory</ApplicationName>
    </SbRole>
  </Role>
  <Options>
    <SbVerificationOptions>
      <ActualMemoryPercent>100</ActualMemoryPercent>
      <MaxBootTimeoutSec>1800</MaxBootTimeoutSec>
      <AppInitDelaySec>120</AppInitDelaySec>
      <TestScripts>
        <TestScripts>
          <TestScript>
            <Name>Domain Controller</Name>
            <Type>Predefined</Type>
            <TestScriptFilePath>Veeam.Backup.ConnectionTester.exe
          </TestScriptFilePath>
            <Arguments>%vm_ip% 389</Arguments>
          </TestScript>
        </TestScripts>
      </TestScripts>
      <HeartbeatEnabled>True</HeartbeatEnabled>
      <PingEnabled>True</PingEnabled>
    </SbVerificationOptions>
  </Options>
</SbRoleOptions>
```

Figure 18, Sample SbRole

4. For Linux, may need to run code within the VM using **Putty\plink.exe**, located in Veeam install directory
 - a) plink.exe exit codes are passed to SureBackup job for reporting

Veeam 9.5 VMCE ADO (VMCE-A) Study Guide

- b) If using Putty with SSH private key, connect manually initially to accept the SSH thumbprint or the SureBackup job will timeout
- c) **puttygen.exe** can be used to create a private key
- b. Custom Test Scripts
 1. Most pre-defined roles use a simplified telnet called ***ConnectionTester.exe**
 - a) This just verifies network services are available
 2. Veeam considers an **exit code=0** as passed; all other codes are failures
 3. Same account to run Veeam services is needed to run scripts
 4. To test services that don't use TCP, run **query-service.ps1**
 - a) Script takes IP of VM & a given service; if service doesn't exist or not running, the script terminates & produce **error code 1** & auto powers-off the data-lab
 - b) To run the script, PoSH Execution must be 'RemoteSigned' on Veeam server
 - c) Account to run script must be a local admin on testing VM
 - d) Windows firewall on VM must allow remote connections or be turned off
- c. Virtual Lab in Complex Environments
 1. VMware vSS
 - a) Virtual Lab proxy appliance & isolated netwk must run on the same ESXi Host
 - b) Use Basic or Advanced Single-Host lab configurations
 2. VMware vDS – generally used for SureReplica or DRS Clusters
 - a) Verify Veeam & proxy appliance are in same subnet to prevent routing
 - b) Use Advanced Multi-Host lab configuration
 3. DR
 - a) Recommended placing Veeam server at DR site for 1-click failover/failback

SECTION 10 – TROUBLESHOOTING

10.1 – Logs

- a. Logs Location
 1. Most all logs are located in: **C:\ProgramData\Veeam\Backup**, or sub-directory under **..\Backup**
 - a) Agent logs
 - 1) NFC & VDDK Agent logs: **..\Backup** (e.g. **Agent.NfcXXXX.log**)
 - 2) "Source" & "Target" Agent logs: **..\Backup\Job-Name-Folder**
 - b) Job logs
 - 1) Backup/Replication: **..\Backup\Job-Name-Folder** (e.g. **DC1-VMBackup.log**)
 - c) Service logs: **..\Backup\Svc.Veeam'SvcName'.log** (e.g. **Svc.VeeamTransport.log**)
 - d) Console logs: **..\Backup\Console** (e.g. **CONSOLE_'username'_localhost.log**)

Veeam 9.5 VMCE ADO (VMCE-A) Study Guide

- e) Configuration Backup logs: `..\Backup_Configuration_Job*.log`
 - f) File-level recovery session logs:
 - `..\Backup\FLRSessions\OSrestoreType\`
 - 1) 2 logs – `date-stamp.log` of the FLR ; sub-directory of VM used for FLR
 - g) Instant VM Recovery: `..\Backup\IRvmName\`
 - h) Task logs: `..\Backup\Job-Name-Folder\` (e.g. `Task.VMName.log`)
- b. Most Useful Logs
- 1. Backup/Replication
 - a) Task log – start with this log; log for each VM in job-named folder
 - b) Job log – deals with before & after VM processing; log for each job
 - 2. Guest Processing
 - a) Guest VSS (**VeeamVssSupport** / **VeeamVixProxy**) logs & Windows events
 - b) Hyper-V issues:
 - 1) Verify HV-IC is up to date
 - 2) Restart Hyper-V Volume Shadow Copy Requestor service
 - 3) 'Guest Processing Skipped' error can't be troubleshooted with logs
 - 3. Hyper-V
 - a) Snapshot creation issues – logs from Proxy, source Host; source Agent & SnapshotCreator logs; Windows Event logs
 - 4. VMware ESXi
 - a) Snapshot issues are all in **Task logs**, except Storage Snapshot: Job* logs
 - b) For config files, look at **Agent.<job>.Source** or **.Target** logs
 - c) For disk files, look at **Agent.Vddk/VixDiskLib** logs
 - d) For Hotadd, look at **Agent.<job>.Source** or **.Target.Hotadd.Disk** log
 - e) For non-Hotadd, look at **Agent.<job>.Source** or **.Target** logs
 - 5. Backup Repository
 - a) Always check **Agent.Target** logs
 - b) Use Windows Events for vague hardware-related errors
 - 6. Backup Copy Job
 - a) For direct transfer, look at same logs as Backup jobs (Task*.log)
 - b) Use Service logs if WAN Accelerators used (**Svc.VeeamTransport*.log**)
 - 7. Full VM/VM Files/HDD Restore
 - a) VM folder under **Vm.<vm-name>.***
 - 8. Veeam Explorers
 - a) Start with Windows FLR
 - b) Then, look in Explorer type folders: **ExchangeExplorer\Logs*.log**
 - 9. Catalog/Indexing-Related
 - a) Occurs often with backup or tape jobs so start with **Job.** or **Task.** logs
 - b) Then use **Svc.VeeamCatalog** log
 - c) For indexes not being removed: **Util.CatCleanup.log**
- c. Reading/Searching Logs
- 1. Search terms
 - a) **> Error** or **> Warning**

Veeam 9.5 VMCE ADO (VMCE-A) Study Guide

- b) **Err |** or **WARN |** have same function but in agent logs
- c) **Has been completed** – useful when no HTML report, or get lost correlating logs with a report
- 2. Other search terms
 - a) Backup/Replication job/task logs
 - 1) Job options – xml file showing job properties
 - 2) [JobSession] Load – **shows bottleneck** that's shown in job statistics
 - 3) Preparing point – tells if backup is incremental, active full, reverse increment
 - 4) Starting agent – useful for locating target agent log
 - 5) Starting SSH – same as 'agent' but for Linux
 - 6) [ProxyDetector] – shows how Proxy & transport mode are selected
 - 7) [VimApi] Create / [VimApi] Remove – creation & removal of snapshots
 - 8) Disk – shows state of VM or replica disks, snap info, disk capacity, etc
 - b) Agent log
 - 1) Err | & WARN |
 - c) Backup service log (**Svc.VeeamBackup* .log**)
 - 1) **== Name** (2 spaces between '==' and 'Name'): shows all configured jobs, whether succeeded or failed, and their state

Result: [Success], State: [Stopped]

Figure 19, Backup Service Log Job Results

- d. How Logs Created
 - 1. Every log is created by a process (.exe) running on a Windows server; e.g. service logs are created by its associated Veeam service
 - a) Veeam Backup Proxy service writes **Service.VeeamTransport** service
 - 2. Services deploy additional processes as needed
 - a) **Backup service** deploys job manager processes; writes manager, job, & task logs
 - b) **Transport service** deploys the agent process, which writes agent logs
 - c) Some errors may be in multiple logs
 - d) If a process crashes, hangs, or fails to start, it won't write to a log
 - e) Redirection – a single process writing to many logs
 - 1) Older logs have a higher number, e.g. target.log , target1.log (older)
 - f) Rotation – log archival sent to storage; controlled by registry keys (**Veeam KB1825**)
- e. Job/Task Logs (VMware/Hyper-V)
 - 1. Agent IDs
 - a) Agent error – designated in logs by **[AP]**
 - b) To find what Agent log reported errors in Task logs, search for the 'id' after **[AP]** in Task log (e.g. **[AP] (ae84)**)
 - 2. Storage statistics
 - a) Found in **Job* logs** (passed from Agent logs)

Veeam 9.5 VMCE ADO (VMCE-A) Study Guide

- b) Compressed / Dedup file size listed first (Fig. 15 below); then file size before compr/dedup 2nd; Dedup ratio (58%), then Compression ratio (44%)

```
66984534016\n
[21.02.2014 21:36:45] <14> Info [AP] (d329) output:
257704365692\n
[21.02.2014 21:36:45] <14> Info [AP] (d329) output: 58\n
[21.02.2014 21:36:45] <14> Info [AP] (d329) output: 44\n
```

Figure 20, Storage Statistics Logging

3. Agent logs

a) **VeeamAgent.exe**

- 1) Core component deployed by Transport service, acts as a data mover
- 2) Generates Agent*.logs
- 3) Main **Agent.Source log** only records read operations for VM config files
- 4) VM disk backup logging is not in Agent.Source logs

4. NFC troubleshooting

- a) NFC protocol used to download/upload non-disk VM files (.vmx, .log, etc)
- b) An API call is made to vCenter or directly to ESXi Host, which then determines which ESXi Host to connect to via NFC to get requested files
- c) Two common NFC problems:
 - 1) “NFC storage connection unavailable” (always 1 of 3 causes); use **Agent.<JOBNAME>.Source.log**, or **Agent.Target** if Replication
 - i. Failure of getting IP Address = update DNS
 - ii. Failure of connecting to Host via port (902) = firewall (probably)
 - iii. File download fails = account connecting to vCenter permissions issue
 - 2) “NFC server (x.x.x.x) is busy” – usually a problem with ESXi VPXA service
 - i. The VM’s .vmx file will show as ‘unassociated’ (icon looks ‘plain’ in vSphere Client)
 - ii. Attempt a VMotion to resolve
 - iii. VMware may have to reinstall the vpxa service

5. Transport connections (client/server agents)

- a) There are always 2 agents, even when source/target are on same server
- b) Server agent is started on the Repository by default, for backups
- c) “No such host is known” – verify IPs are routable/correct, & server agent is listening (port 2502)

6. VeeamGuestHelper/VeeamVssSupport (guest VM):

C:\ProgramData\Veeam\Backup\VeeamVssSupport.log)

- a) **First item to verify is 1st line of log has DNS name of VM backing up**
- b) SQL VSS writer needs to have its data set as Drives, not UNC paths
- c) Data is encrypted before transport to Veeam server
- d) Veeam VSS Writer metadata can also be retrieved via `list writers` cmd
 - 1) Veeam Freeze Writer managed by **VeeamVssSupport** can be ignored
 - 2) ‘freeze’ & ‘thaw’ timestamps cannot be > 60secs ; 20secs for Exchange

Veeam 9.5 VMCE ADO (VMCE-A) Study Guide

- 3) 'Wait timeout 900 seconds' error typically is app miscfg (i.e. PSCONFIG)
- 4) 'Freeze writer is unfrozen' can be safely ignored
- f. Job/Task Logs (VMware only)
 1. Always only 1 Job log & 1 Task log for each VM in the Job folder
 2. **Task logs** will show VM (h/w, disk mode), ESXi Host, vCenter, & Datastore info
 - a) "Mode" (disk) will give clue on backup support (e.g. independent disks aren't supported); as well as VM h/w compatibility with ESXi Host version
 3. CBT (VMware)
 - a) CBT info shown in "source disk agent logs"
 - b) CBT is also used for Active Full backups, using "*" as ChangeID
 - c) Limitations of "*":
 - 1) Not supported on NFS (agent log warnings)
 - 2) Not supported on thick-eager zero disks (no log warnings)
 4. Snapshot creation/removal
 - a) Search for using terms: "Create Snapshot", "RemoveSnapshot", "VimApi"
 - b) Snap removal is usually a longer running process than creation
 - c) Snapshot consolidation info: **VMware KB 2053758**
 - d) "removeChildren False" denotes Veeam only removes snaps it created
 5. Transport modes
 - a) For **san/hotadd** job failover to **nbd**, look in **Agent.Source.Hotadd.disk***
 - b) **hotadd** failures tend to be shown in logs via "Mounting VM vim.xxx"
 - 1) **Veeam KB1054** covers hotadd failures
 - c) SAN transport
 - 1) VDDK will attempt to connect to a Volume via its LVID; if it cannot it'll failover to **nbd** or fail altogether
 - 2) If a Volumes `naa.` or `eui.` number isn't listed, verify it's in the Client; a Proxy rescan may be needed
 - d) Network transport
 - 1) Simplest mode, solely requiring network connection from Veeam source Proxy to the ESXi Host via NFC over ports 902 & 443
 - 2) Most common issues are due to NFC, so the task will fail before VMDKs are processed (See 'NFC Troubleshooting' above, & Veeam **KB1198**)

10.2 – Common Issues

- a. Veeam Backup & Replication
 1. **Transform operations of large file on Win2012/R2 dedup volume failing**
 - a) **Message:** "Error:| Client error. The requested operation could not be completed due to file system limitation. Failed to flush file buffers. File: filename.vbk"
 - b) **Causes:**
 - 1) Veeam not able to write data because of NTFS limitation
 - 2) Manual un-dedup (Expand-DedupFile) fails with same error

Veeam 9.5 VMCE ADO (VMCE-A) Study Guide

- 3) Repo volume must be formatted for large files – `format /fs:ntfs /L`
- c) Resolution:
 - 1) Disable Windows dedup
 - 2) Undedup backup files – copy them to a temp folder, delete orig files, move copied files back to original location/directory
 - 3) Disable NTFS compression on Repository
 - 4) Issue is resolved in Win2012 R2 but /L is still needed:
cmd prompt > `format h: /fs:ntfs /L`
 - 5) See **MS KB967351** for Win2008
2. **File-level restore from VM with ReFS or Windows dedup fails**
 - a) Windows FLR may fail:

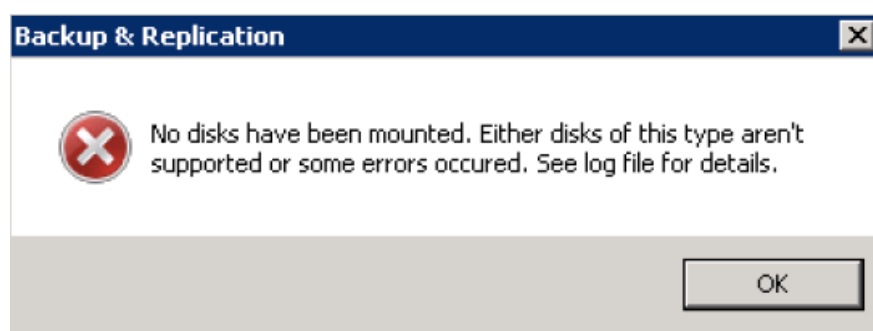


Figure 21, Window FLR Failure Message

- b) Verify the source VM the files are on is using: FAT, FAT32, NTFS, ReFS
- c) Verify Veeam Server on 2012/R2 with data dedup enabled
- d) Troubleshooting steps
 - 1) Clear temp files in `C:\Users\AppData\Local\Temp` or `C:\Windows\Temp` if Veeam services run under LOCAL SYSTEM
 - 2) Open elevated CMD PROMPT to run **vdk.exe**
NOTE: for installs upgraded from v6.5, the path is:
`C:\Program Files\Veeam\Backup and Replication\Backup;`
for v9, the path is:
`C:\Program Files\Common Files\Veeam\Backup and Replication\Mount Service`
 - i. `cd "path listed above"`
 - ii. `vdk remove`
 - iii. `vdk install`
 - 3) Confirm account used for Veeam Backup service has 'Perform volume maintenance tasks' rights
`GPO:Computer Config > Windows Settings > Security Settings > Local Policies > User Rights Assignment`
 - 4) Reboot the Veeam Server
- b. Hyper-V
 1. **VSS not engaged with AAIP enabled** ('Guest processing skipped')
 - a) Integration services

Veeam 9.5 VMCE ADO (VMCE-A) Study Guide

- 1) Incorrect version – check Hyper-V Host Event Viewer log
 - i. Event references “Hyper-V Shadow Copy Requestor... version” error
- b) Guest OS
 - 1) 2012 R2 VMs are not supported for AAIP on Hyper-V 2008 R2 Hosts
 - i. If Windows App Event ID 13 is present in guest VM, create registry key: HKLM\Software\Microsoft\WindowsNT\CurrentVersion\SystemRestore DWORD: **ScopeSnapshots** = 0
 - 2) Verify Windows is on latest SP
 - 3) In Disk Management, verify:
 - i. Volume & Storage Area match in Vol properties > Shadow Copies tab
 - ii. Verify no dynamic disks in the guest VM
 - iii. Verify no guest VM Volumes are FAT or FAT32; must be NTFS or ReFS
 - iv. Verify guest VM System Reserved partition has > 42MB; if not, create Shadow Copy
 - 4) Start or Restart Hyper-V Volume Shadow Copy Requestor service
 - 5) If all has been tried, restart the VM, then restart HV Shadow Copy Req svc
- c) VM configuration (Host level)
 - 1) Verify VM has SCSI Controller, even if not used
 - i. Hyper-V VMMS > Admin Event log: “virtual machine cannot be backed up since it has no SCSI controller...”
 - 2) Verify VM Settings > Backup (Volume snapshot) enabled
 - 3) Verify VM VHD files are not sitting in root of CSV
- d) Verification (Host level) – to test Hyper-V VSS writer
 - 1) Open CMD Prompt elevated
 - 2) `diskshadow /L C:\writers.txt`
 - 3) `list writers`

10.3 – Troubleshooting Mode

- a. SureBackup / SureReplica
 1. To start Troubleshooting Mode
 - a) Open SureBackup stats > rt-click any VM > select Start
 2. Troubleshooting Mode means errors won't cause the Lab to shut down
 - a) Helpful in measuring app boot times via vPower NFS
 - b) Helpful when implementing custom app test scripts
 3. When done, stop the session manually
- b. Windows servers may change their Firewall profile to Public, leading to application test scripts not responding to TCP socket connections

Veeam 9.5 VMCE ADO (VMCE-A) Study Guide

APPENDIX A

Summation of Registry Key Modifications & Their Purpose

PURPOSE	REGISTRY KEY CHANGE
To change the number of max allowable snapshots per VMware Datastore	HKLM\Software\Veeam\Veeam Backup and Replication DWORD: MaxSnapshotsPerDatastore = #
To disable UAC when using AAIP & configuring local admin account for VSS, when local admin account is not named administrator (Win2012+)	HKLM\Software\Microsoft\Windows\CurrentVersion\policies\system DWORD: EnableLUA = 0
Disable Veeam Volume Logging	HKLM\Software\Veeam\Veeam Backup and Replication DWORD: VDDKLogLevel = 0
To give preference to a Hotadd Proxy to an ESXi Host on which a protected VM resides for Direct NFS; to prevent high stuns, since it takes ownership of VM & changes LCK contents of VM	HKLM\Software\Veeam\Veeam Backup and Replication DWORD: EnableSameHostHotAddMode = # Value: 1 – Uses Proxy; if unavailable, will use another Hotadd Proxy Value: 2 – Uses Proxy; if unavailable, will uses nbd mode
To modify Proxy Network Mode data transport NFC connections higher than the default of 7	HKLM\Software\Veeam\Veeam Backup and Replication DWORD: ViHostConcurrentNfcConnections = #
Settings when using HPE StoreOnce Catalyst dedup appliance; a. if Share isn't configured as 'distant' (low bandwidth key) or is 'distant' (payload keys)	HKLM\Software\Veeam\Veeam Backup Transport\ DWORD: UseLowBandwidthMode = 1 HKLM\Software\Wow6432Node\Veeam\Veeam Backup Transport\ DWORD: PayloadChecksumsDisabled BodyPayloadCompressionDisabled, both = 1
Win2012/R2 VM guest running on Hyper-V 2008 R2 , and Event ID: 13 with AAIP/VSS issues	HKLM\Software\Microsoft\WindowsNT\CurrentVersion\SystemRestore\ DWORD: ScopeSnapshots = 0
Default folder paths – Catalog / NFS	HKLM\Software\Veeam\Veeam Backup Catalog\ DWORD: VBRCatalog = [path value] HKLM\Software\Wow6432\Veeam\Veeam NFS\ DWORD: NFS = [path value]

Veeam 9.5 VMCE ADO (VMCE-A) Study Guide

APPENDIX B

Summation of Important Veeam Folders for Logs, Tools, Install Files (NOTE: This is not an exhaustive list, but will contain most log folder descriptions)

ITEM DESCRIPTION	LOCATION
Veeam Backup & Replication (default) install folder	C:\Program Files\Veeam\Backup and Replication\Backup
Server Role (.xml) files to run tests against VMs for SureBackup	C:\Program Files\Veeam\Veeam Backup and Replication\Backup\SbRoles
Script file to test VMs for SureBackup	C:\Program Files\Veeam\Veeam Backup and Replication\Backup\Veeam.Backup.ConnectionTester.exe
Main Log folder location; the following Logs will be designated as under this main folder using “\Backup*.log”. If a Log is in a different location, the full path to that location will be shown	C:\ProgramData\Veeam\Backup
Backup/Replication Job Log folder location	\Backup\JobNameFolder\Job.JobName.*.log
Backup/Replication Task Log folder location	\Backup\JobNameFolder\Task.VmName.*.log
Guest Processing (AAIP) Log folder location (on VM guest); Use Veeam KB: KB1855	\Backup\VeeamVixProxy.log
Repository Log folder location	\Backup\JobName\Agent.JobName.Target.*.log
Veeam Explorer Log folder location	\Backup\ExplorerNameFolder\Logs\
Catalog or Indexing Log folder location	\Backup\Svc.VeeamCatalog.log
VM / Files restore Log folder location	\Backup\VMnameFolder\Vm*.log

Veeam 9.5 VMCE ADO (VMCE-A) Study Guide

APPENDIX C

Summation of Veeam ONE Reports and Their Purpose

REPORT	PURPOSE
Storage Latency Report	Helps identify storage systems under heavy pressure, load, & latency; also used to help view storage when source is bottleneck
VM Change Rate Estimation Report	Helps identify VMs with high change rates to help plan backup target storage
VM Configuration Assessment Report (VMware Only)	Helps assess the readiness of VMs for Veeam Backup
Configuration Assessment (Hyper-V Only)	Analyzes Hyper-V configuration against recommended settings & best practices
Infrastructure Assessment Report	Shows configuration properties of servers, clusters, Hosts, datastores, & networks
Active Snapshot Report	Aggregates historical data showing VM snapshot age & size
Orphaned VM Snapshots Report	Detects snapshots that reside on datastores but don't show in Snapshot Manager
Infrastructure Changes Report	Shows changes to virtual infrastructure
Job Configuration Change Tracking Report	Shows changes made to backup environment
Protected VMs Report	Displays a list of VMs protected by up-to-date backups and replicas, as well as unprotected VMs of out-dated/missing backups or replicas
VM Backup Compliance Overview Report	Shows a list of VMs that do not meet the minimum requirement of backup copies; 3-2-1 Rule analysis

Shane Williford, Sr. Systems Engineer

For Public Use, but please give credit to the author if posted on social media or in blogs