Veeam Repository Best Practices

VEEAM USER GROUP

Background

Why talk about this?

What are the benefits of applying best practices?

Topics At A Glance

Identify your design requirements

What do basic low end and high end repositories look like?

Repositories to avoid

Looking at disk and RAID setup

Veeam Job Optimizations

Protecting your backups

Identify Your Requirements

What is your Recovery Point Objective (RPO) and Recovery Time Objective (RTO)?

How long do you need to retain backups?

Are you currently meeting these objectives?

Considerations

Reliability

Speed (backup and restore)

Time for complete recovery (DR event)

$$$$

Low End Repository

Physical Windows or Linux machine

Local storage

VM guest iSCSI

High End Repository

Dedicated dedupe storage

e.x. EMC Data Domain, ExaGrid, HP StoreOnce

What shouldn't you use

Low-end NAS (reliability / availability)

CIFS

VMDK on VMFS

Windows Server 2012 R2 Deduplication *

Disks and RAID

Disk Type	Approx IOPS	RAID 10 Reverse Incremental	RAID 5/50 Reverse Incremental	RAID 6/60 Reverse Incremental	RAID 10 Synthetic Full	RAID 5/50 Synthetic Full	RAID 6/60 Synthetic Full
7.2K SATA	50	~5 MB/s	~2.8 MB/s	~2 MB/s	~8.5 MB/s	~5 MB/s	~3.5 MB/s
10K SAS	75	~7.5 MB/s	~4.2 MB/s	~3 MB/s	~12.8 MB/s	~7.5 MB/s	~5.2 MB/s
15K SAS	100	~10 MB/s	~5.6 MB/s	~4 MB/s	~17 MB/s	~10 MB/s	~7 MB/s

Maximum performance per spindle.

RAID 10 when possible (2x write penalty; less capacity)
RAID 5 next best choice (4x write penalty; greater risk due to rebuild times)
RAID 6 most severe write penalty (6x)

Consider per-VM based file backups.

RAID Volumes

Stripe size

- Typical IO for Veeam is 256KB to 512KB
- Windows 2012 defaults to 64KB (`fsutil fsinfo ntfsinfo c:`)
- Arrays may be as low as 32KB
- Lots of wasted I/O (512 / 32 = 16 IOPS)

At least 128KB stripe size is recommended.

Arrays

Fill all the drive bays

- Performance degradation when expanding
- More spindles = More performance

File Systems

NTFS

Larger block size does not affect performance; 64KB block size is recommended.

Use the /L parameter when formatting drives

Server 2012 – Max file size 256 TB

Backup Jobs

Space vs. Performance

Reverse Incremental backups are 3 x I/O per block

Forever Forward Incremental backups are 2 x I/O per block

Be mindful of the repository load – can be throttled

Primary Storage

Dedupe

- Avoid due to speed / performance
- If you have to:
  - Look for vendor-specific integrations
  - Avoid jobs w/ full backup transformation
  - Use active Full vs Synthetic Full

Secondary Storage

Dedupe

- Better use case than primary
- Test Backup Copy retention performance
  - Consider Active Full option (vs. transformation)

Replication

- Array level not recommended – replicates corruption
- Use Backup Copy Job – performs its own health check

Backup Job Settings

Dedupe

Keep on

Compression

Keep set to Optimal
Enable “Decompress before storing” repository option
Minimal CPU impact

Other Considerations

Protect Your Repository

- Malicious users, Cryptolocker, bad firmware, etc.
- Rotating media

Consider Scale Out Backup Repositories

- Spread the load across multiple repositories
- Not all backups / VMs sitting in the same repository

References / Sources

Backup Repository Best Practices: 2015 Edition

https://www.veeam.com/veeamon/free-sessions

Veeam Forums

https://forums.veeam.com

The Expert Guide to VMware Disaster Recovery and Data Protection

https://www.veeam.com/wp-vmware-data-protection-disaster-recovery-expert-guide.html