Veeam Replica

Maria Zubareva
Team Leader EMEA&APAC

Outline

- Replication job and its workflow
- Replication architecture
- Network Mapping and Re-IP
- Mapping/Seeding/Replica from Backup
- Key changes in the .vmx file, replica summary
- Failover and Failback

What is replication?

- **process of copying a VM from its primary location to a destination location (redundant target host)**
- creates an almost identical, functional copy of a VM in a ready-to-start-state (thus providing minimal Recovery time objective)
- allows to have restore points
- onsite and offsite implementations
- provides ability to customize replicated VM settings (name, networking, ip configuration, disks locations)
- allow to use WAN accelerators to minimize network traffic
- controlled failover/failback from Veeam console

Replication architecture

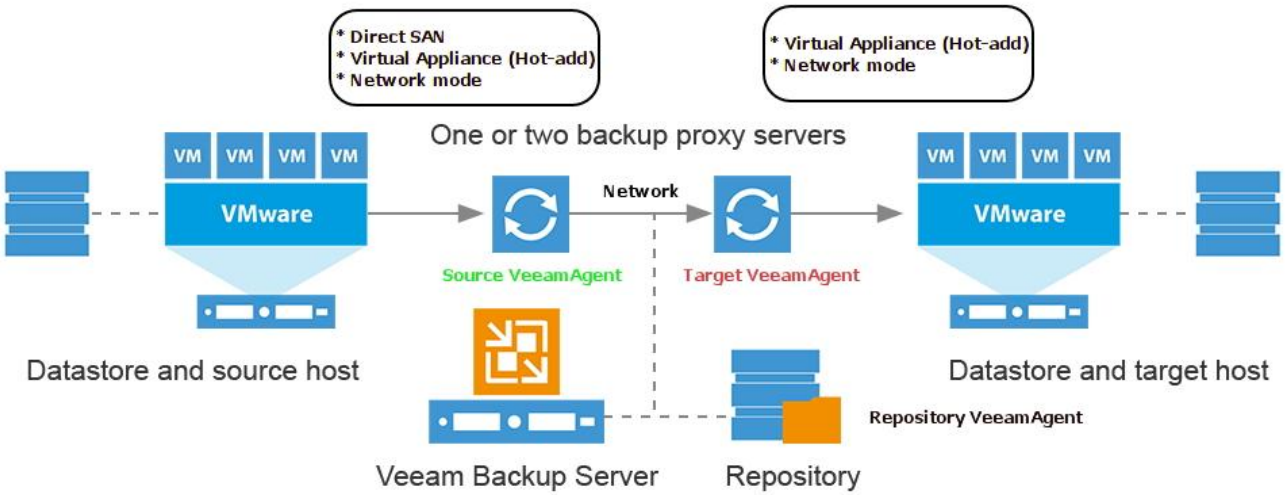

* on target site Direct SAN TM is available during the initial run of replication job

Network mapping

- Changes the VMWare network on the replica VM
- Applied directly to the target .vmx file on replication

Re-IP

- IP address, subnet mask, default gateway and DNS can all be changed for the replica VM
- Applied only on failover
- Windows only

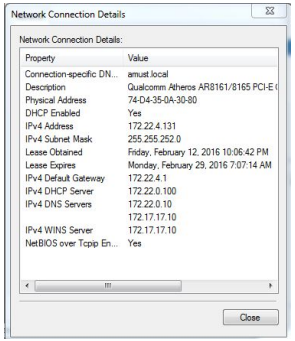

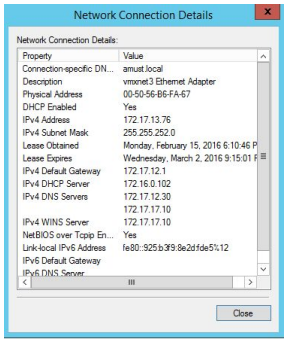

Replica Mapping

- Have an existing VM on the target host that is similar to the production VM
  - either made from the same template or left from a previous replication job
- Point to that VM in the replication job
- Veeam will calculate the digests and transfer the different blocks
  - all existing snapshots on the target VM will be removed first

Replica seeding

- Make a backup locally, on-site
  - Transfer it to the off-site (e.g. on a USB drive) to a Veeam repository
- Rescan the repository so that Veeam registers the backup files
  - The backup will appear under Backups > Imported
- Use this backup as a seed in the replication job
- Once the job starts:
  - Veeam will restore the VM to the target host from that backup
  - Veeam will scan and determine the data blocks that are different between production VM and the restored VM (“calculating digests”)
  - Only the changes are transferred over

Replica from backup

1. Have backup(s) of VM on the target site
2. Choose backup repositories that have backup(s) of the VM
3. During every ran of the replication job Veeam will scan chosen repositories to find the newest data to build a restore point on replica

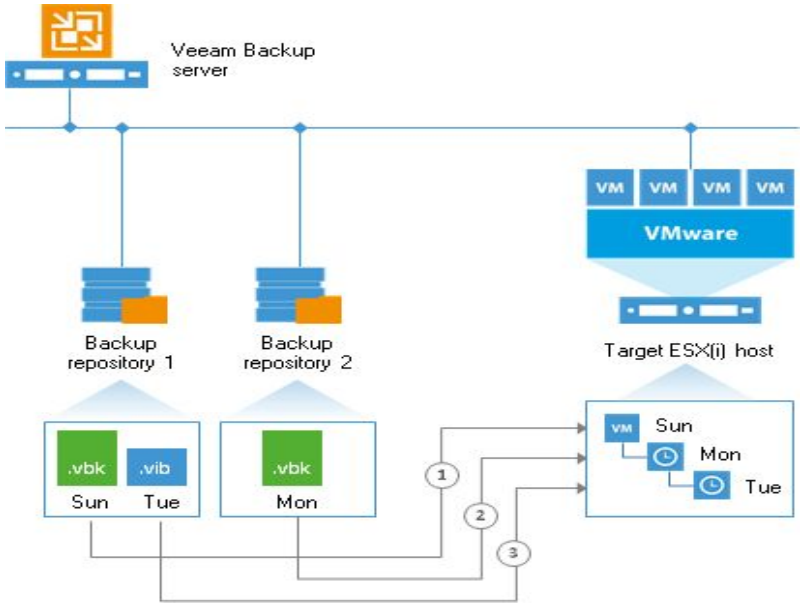

Snapshot-based replica

- Native VMWare snapshots are used as restore points on the replica VM
- Initial run: a copy of the VM is made and a snapshot is created
- Subsequent runs: only the changes are transferred and are put into new snapshots
- Retention: removes oldest snapshots

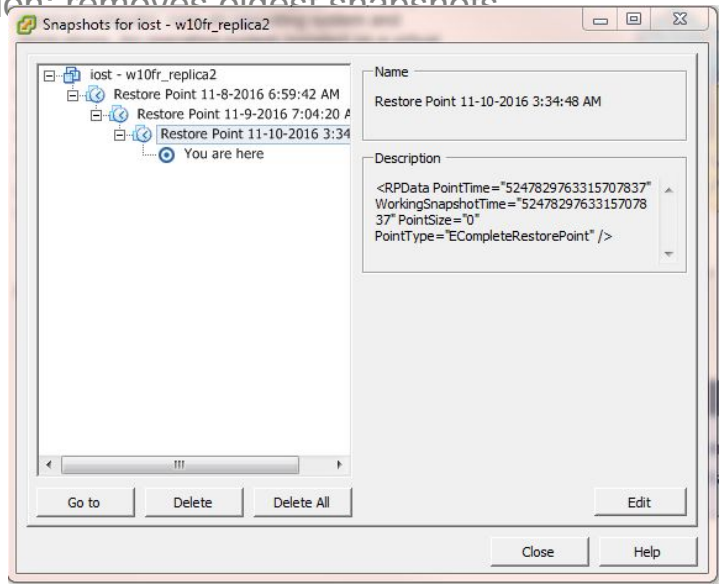

Replication job workflow

1. Detect best source and target proxies for each VM in the job
2. Start Veeam agents: Source, Target and Digest.
3. Read and parse source .vmx.
4. Snapshot source VM (with optional VSS guest-processing), “VEEAM BACKUP TEMPORARY SNAPSHOT”.
5. Determine objects (files) for replication: .vmx, .vmxf, .nvram, .vmdk.
6. Revert replica VM to the latest restore point (for subsequent runs)
7. Compile target .vmx, upload it to the target datastore, upload small files .vmxf, .nvram. (1st run)
8. Register replica VM on the target host (1’st run)
9. Add VeeamReplicaSummary parameter to the replica VM (1’st run)
10. Create empty disk(s), attach it to replica VM (1’st run).
11. Configure network adapters as per network mapping rules.
12. Snapshot replica VM, “Veeam Replica Working Snapshot”.
13. Transfer actual source disk data (from source: NBD/hot-add/SAN; to target: NBD/hot-add/SAN for initial run only) to the target disk file that was snapshotted.
14. Save disk blocks checksums (“signatures”, or “digests”) data to a .vbk file on a repository.
15. Delete “VEEAM BACKUP TEMPORARY SNAPSHOT” on the source VM.
16. Revert and delete “Veeam Replica Working Snapshot” snapshot on replica VM.
17. Create snapshot on replica, “Restore Point dd.mm.yyyy hh:mm:ss”.

Key changes in vmx file

Source VM	Replica VM
displayName = "al-nodisk"	**displayName** = "al-nodisk_replica"
Suffix can be customized in the job	
ethernet0.generatedAddress = "00:50:56:86:3c:4f"	**ethernet0.generatedAddress** = "00:50:56:86:6b:a0"
Mac address is assigned automatically by Vsphere for all new and cloned VMs	
	uuid.action = "keep"
UUID is preserved to avoid issues with software licences in the Guest OS	
uuid.location = "56 4d 64 ad 9f 9f 2c 70-31 e9 b7 e2 40 bd cf 82"	**uuid.location** = "56 4d 58 9d 47 0b 0c 5a-26 ff 86 b3 6b 46 c2 fd"
UUID is generated based on VM files location	
ctkEnabled = "TRUE"	**ctkEnabled** = "False"
CBT is disabled because the VM is turned off and updated in this state by Replication Job	
scsi0:0.fileName = "al-nodisk_1.vmdk"	scsi0:0.fileName = "al-nodisk_1-000003.vmdk"
Replica always has a snapshot	
doesn't exist	**VeeamReplicaSummary** = "<Summary>..
Replica VM is tracked and identified by summary field	

Replica failover

- Switches over from the original VM on the source host to its VM replica on the target host
- Applies Re-IP rules to have the replica VM's IP settings match the DR site network

Replica failback

- Switches back to the production VM
- Synchronizes restored VM with it's replica state by transferring only differential blocks

Replica failover workflow

1) Revert to the selected restore point
2) Apply Re-IP rules to replica machine (involves mounting replica's system disk, modifying Windows registry)
3) Update replica summary field -> add failover status
4) Power on replica machine
5) Update replica status in Veeam database (replica will be locked while in failover status) and GUI

Undo Failover

1. Revert to the latest restore point
2. Update replica summary field -> remove failover status
3. Power off replica machine
4. Update replica status in Veeam database (unlock replica storage)

Permanent Failover

1. Add the original VM to the excluded objects list in the replication job
2. Delete all snapshots on replica VM
3. Remove digests from the repository
4. Remove replica summary field
5. Delete restore points from Veeam database

Replica failback workflow

1) Check replica and source vm configurations
2) Detect source and target proxy modes
3) Power-off original VM
4) Create working snapshot on the original VM: “Veeam Replication Failback Snapshot”
5) Calculate digests for the 1st disk (both for source and replica VMs)
6) Replicate disk content (only blocks that differ) – update source vm to pre-failover replica state
7) Repeat steps 5 and 6 for each disk
8) Power-off replica vm in order to transfer changes that have been made during failover (replicate snapshot state) and transfer the data
9) Create snapshot on replica VM: “Restore point dd-tt” (failback protective snapshot)
10) Revert and delete “Veeam Replication Failback Snapshot” on the source VM
11) Undo Re-IP rules if any
12) Power on source VM.

Undo Failback

1) Revert replica to the pre-failback state
2) Update replica summary field -> return failover status
3) Delete failback protective snapshot
4) Power on replica

Commit Failback

1) Edit source VM object in replication job (restore to different location)
2) Update replica summary field -> remove failback status
3) Revert replica back to the latest restore point

Failover Plan

- Offers a possibility to failover several VMs, one after another with some time delay between (optional)

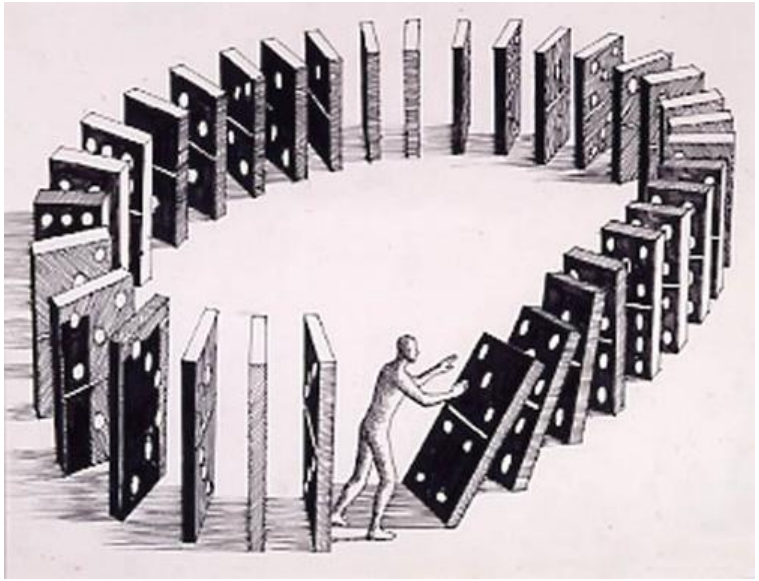

Planned Failover

- Can be used if you're planning to turn off your production VMs or to migrate it to a new location with Veeam Replica

Planned Failover workflow

- Planned failover triggers incremental run of the replication job for the selected VM
- The source VM is powered off
- Failover from original VM to it's replica
- Replica VM is turned on

Thank you!

VEEAM
IT JUST WORKS!™