IBM Spectrum Scale Strategy Days 2021

Jochen Zeller IT Architect





Backup Consolidation with Veeam and Spectrum Scale



jochen.zeller@sva.de

/ THE "WHAT" AND "WHY"

- Veeam® Backup & Replication[™] is a solution for simple, reliable and flexible protection of cloud, virtual and physical workloads. In brief: a backup and recovery software.
- In almost all installations Veeam is using disk storage for backup repositories.
- The customer tendered a 2 x 3.2 Pbyte storage solution for Veeam including operations.
- They already use IBM Spectrum Scale in the IBM Spectrum Protect environment and as HPC file system.
- We offered a high available and scalable repository based on one IBM Spectrum Scale file system.



/ VEEAM INFRASTRUCTURE CONFIGURATION OVERVIEW – IN GENERAL



- 1. The backup server triggers a backup job
- The backup proxy e.g. a VM in the same ESX cluster - copies VM data from the source ESX server to the backup repository server
- 3. The repository server stores the backup data to the file system and runs merge or copy processes on the backup data. The repository is placed on a IBM Spectrum Scale file system.

/ WHY IBM SPECTRUM SCALE AS STORAGE FOR VEEAM?

With IBM Spectrum Scale, we

- share all backup repositories on all nodes.
- use CES to make the backup repository IPs high available, and therefore the repositories themselves.
- do not have silos, all capacity and performance is available to all nodes and repositories.
- can add / remove / replace storage capacity and nodes online, without Veeam knowing anything about it.
- have a perfect IO behavior for large capacity drives.
- do not have support for the BlockClone technology like ReFS has. (With ReFS and the BlockClone technology, only the pointers of the modified blocks are updated. The contents of the blocks are not written to disk a second time. The updated Veeam backup file is generated from the new pointers to these blocks. This feature saves space and reduces IOs for copying data dramatically).



/ OUR VEEAM REPOSITORY SOLUTION



/ THE FIRST TIME CHALLENGE ...

Veeam Scale-Out Backup Repository



- From an application perspective, it seemed to make sense to use the scale-out repository. It allows
 repositories distributed across multiple nodes and thus offers more flexibility. And this together with IBM
 Spectrum Scale that sounds promising!
- But! The scale-out repository distributes the work for a job across two repo nodes. One node reads, another node writes. And the data transfer goes over the network, here 10GBE. This led in total to a load > 200, performance problems, very long running jobs this combination is not promising.



/ THE FIRST TIME CHALLENGE ...

IBM Spectrum Scale file system block size

- As backend storage we use NetApp E-Series with NL-SAS drives for data and SSD drives for meta data.
- We configured Veeam to organize the backups into a single file for each node and backup job. This leads to less than 100.000 large files.
- Most Veeam IOs are 1MB, some 512KB. Most of the IOs are sequential.
- Because of "large_files+large_seq_IOs" we started with 16MB file system block size to get max streaming performance out of the storage.
- BUT! We must create a daily snapshot as a protection against ransom ware. What happens during an FS Block update when using snapshots?







/ WHY IS THE SNAPSHOT BEHAVIOR IMPORTANT FOR THE SOLUTION?

- Veeam is working with backup chains and creates "synthetic full backups" out of a full and incr backups



- The merges on IBM Spectrum Scale are very IO intensive because all data to merge must be read and written (e.g. instead to Microsofts ReFS – this allows to repoint the blocks to a new file).



/ WHY IS THE SNAPSHOT BEHAVIOR IMPORTANT FOR THE SOLUTION?

- To give you an idea how much data a merge changes:

# mmlssnapshot dp -dblock-size G							
Snapshots in file system	dp: [data	and meta	adata in G	B]			
Directory	SnapId	Status	Created			Data	Metadata
@GMT-2021.02.16-05.00.53	353	Valid	Tue Feb 1	6 06:00:55	2021 1	07212	6
@GMT-2021.02.17-05.00.53	354	Valid	Wed Feb 1	7 06:00:55	2021	21723	2
@GMT-2021.02.18-05.00.54	355	Valid	Thu Feb 1	8 06:00:55	2021	41142	3
@GMT-2021.02.19-05.00.54	356	Valid	Fri Feb 1	9 06:00:55	2021	67147	4
@GMT-2021.02.20-05.00.54	357	Valid	Sat Feb 2	0 06:00:56	2021	11392	1
@GMT-2021.02.21-05.00.54	358	Valid	Sun Feb 2	1 06:00:55	2021	3406	1
@GMT-2021.02.22-05.00.54	359	Valid	Mon Feb 2	2 06:00:55	2021	63106	4
@GMT-2021.02.23-05.00.54	360	Valid	Tue Feb 2	3 06:00:56	2021	59135	3
@GMT-2021.02.24-05.00.55	361	Valid	Wed Feb 2	4 06:00:56	2021	1091	1

- With 16MB fs block size we saw a multiple of reads and writes on the disks than on Spectrum Scale.



/ OUR CURRENT SOLUTION

- We use Veaam "classic" backup repositories.
- Each CES IP in combination with a directory in a fileset represents a backup repository.
- If a node fails or goes in maintenance, the CES IP moves to another node and the repository is online again. Only active tasks for the current VMs will fail, the job continues to work with the following VMs after the IP move is finished.
- We recreated the file system with 512KB block size because of the snapshots behaviour. This increased the performance of most jobs and reduced the backend disk IO (instead of the 16MB fs block size).
- We needed to make Veeam a little more HPC-like to better leverage the power of Spectrum Scale.
 We increased the number of parallel processing tasks (1 task processes one .vmdk) per job to 60.
 The recommendation and default value is 4.





/ THE OVERALL BIG PICTURE





/ PERFORMANCE CAPABILITY OF THE STORAGE

mmdsh -N all '/usr/lpp/mmfs/samples/perf/gpfsperf mixrw seq /dp/zuf2/tmp/perf.test.\$(hostname) -r 512K -n 200G -th 8 -fsync'

8 nodes, 8 threads per node, 50% seq read & 50% seq write @ 512K record size







/ A TYPICAL WEEK ...





/ SUMMARY

- IBM Spectrum Scale is a very good solution for large Veeam Backup Repositories
- Veaam Scale Out Repositories are not aware of cluster file systems, so you have to use classic repositories
- When using Spectrum Scale snapshots, be aware of copy-on-write
- Increase the number of parallel tasks for the Veeam jobs
- Our customer is very satisfied with this solution, we are currently installing on the next site





/ CONTACT



JOCHEN ZELLER

IT Architect

jochen.zeller@sva.de



