IBM Spectrum Storage

IBM Spectrum Scale Strategy Days

Backup of IBM Spectrum Scale file systems

Dominic Müller-Wicke
IBM Spectrum Protect Development

IBM ®

IBM’s statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM’s sole discretion.

Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision.

The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code or functionality.

Information about potential future products may not be incorporated into any contract.

The development, release, and timing of any future features or functionality described for our products remains at our sole discretion.

Agenda

- Large Filesystem Backup
- Performance Considerations

Spectrum Protect / Spectrum Scale Integration Overview

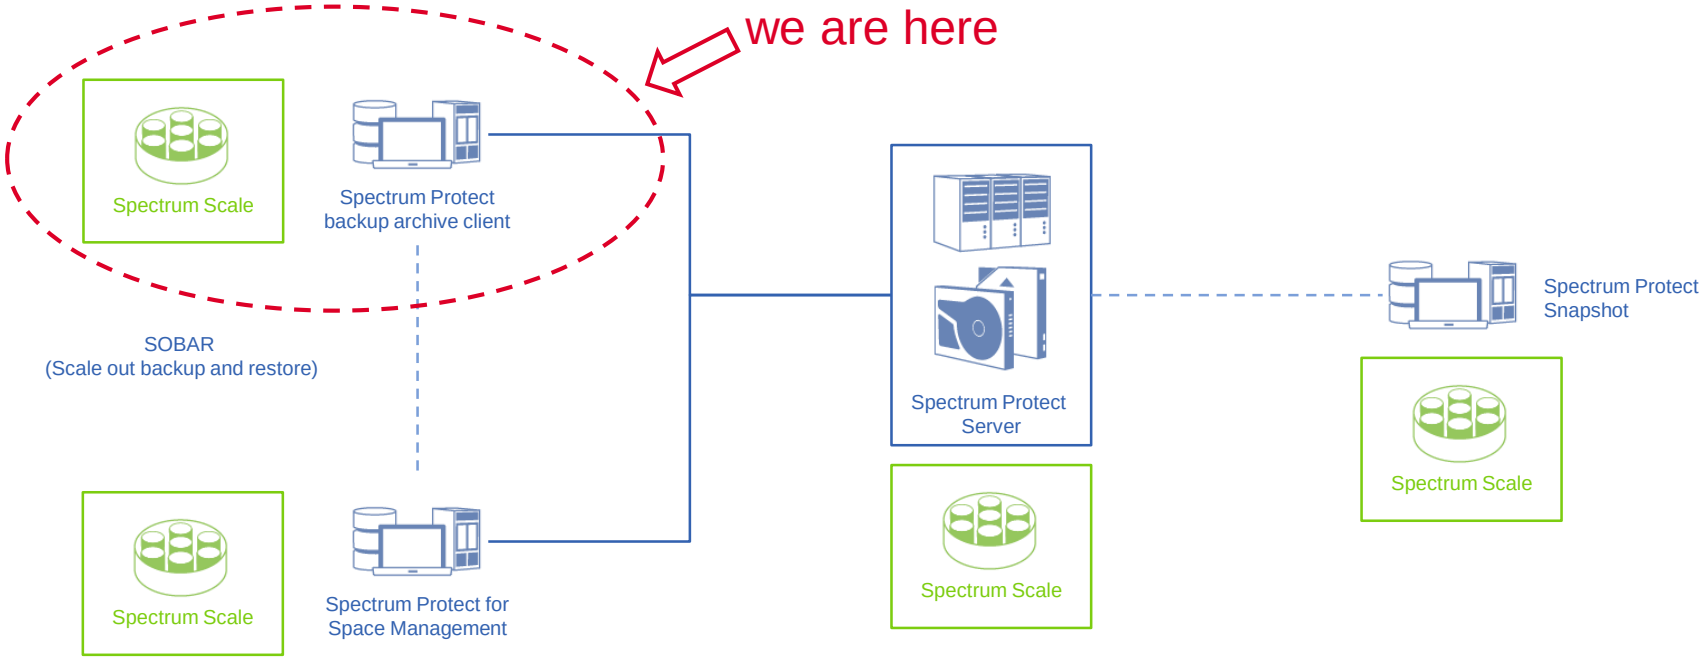

Large Filesystem Backup

IBM Spectrum Protect progressive incremental backup

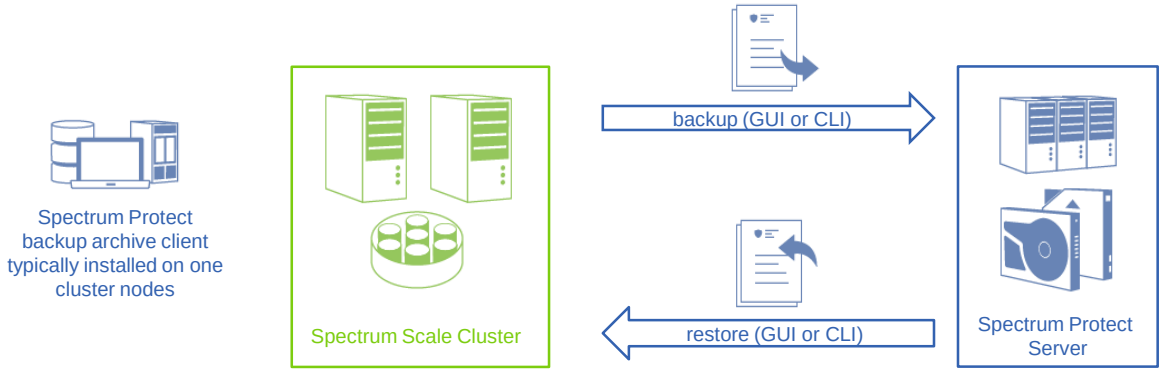

- Environment: Small IBM Spectrum Scale installations with a small number of nodes and file systems. IBM Spectrum Protect backup archive client installed on one or more cluster nodes
- Scalability: Millions of files, Terrabytes of data, up to 25.000.000 Objects (empirical value)
- Processing: Standard IBM Spectrum Protect backup archive client progressive incremental is used to perform file system backup. Potentially a second node for a second file system backup
- Pros: Simple setup and usage
- Cons: Limited performance and scalability

IBM Spectrum Scale mmbackup on file system level

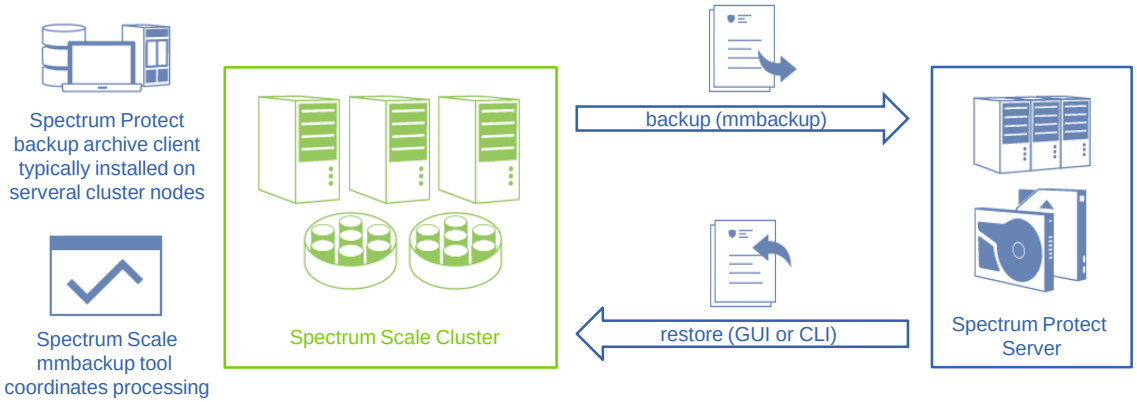

- Environment: Medium IBM Spectrum Scale installations with a single digit number of nodes and file systems. IBM Spectrum Protect backup archive client installed on several cluster nodes
- Scalability: Tens of millions of files, Tens of terrabytes of data, up to 1.000.000.000 Objects (empirical value)
- Processing: IBM Spectrum Scale mmbackup scans file system and IBM Spectrum Protect data base and generates list of backup candidates. IBM Spectrum Protect backup archive client used from mmbackup to perform file system backup.
- Pros: Simple setup and usage, Good performance and scalability
- Cons: All data goes to one IBM Spectrum Protect server

M: IBM Spectrum Scale mmbackup on file system level

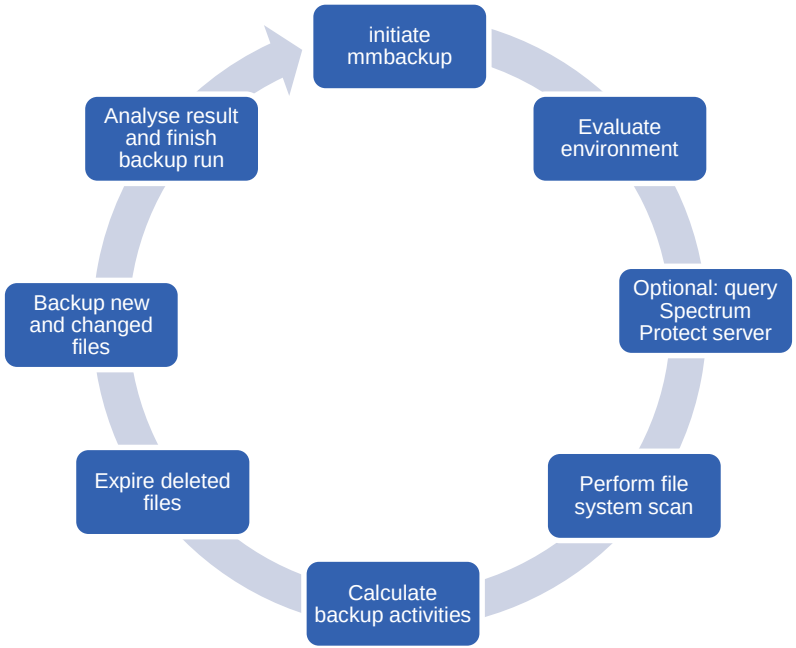

Backup cycle:

- After start mmbackup evaluates the cluster environment and verifies product versions and settings
- Optional the Spectrum Protect server is queried for existing backup information. In other cases existing shadow DB is used for processing
- The policy engine is used to generate a list files currently eligible for backup activities
- Compare existing shadow DB and scan result to calculate file lists for required backup activities
- Expire all files deleted in the file system since last backup run
- Incremental backup all files with changed metadata in the file system since last backup run
- Selective backup all files with changed data in the file system since last backup run
- While backup activities ongoing update shadow DB inline
- Analyse backup results from all used cluster nodes and finish backup cycle by selective backup the current shadow DB

Translation between mmbackup and dsmc commands

mmbackup command	dsmc commands used during execution	comments
all mmbackup commands	dsmc query sessions	
	dsmc query options	
initial mmbackup run for file system	dsmc restore /<fsname>/.mmbackupShadow*	try to get previous shadow db.
	dsmc query backup -detail	if no shadow db can be found or restored full filespace query, single-threaded. if no filespace can be found dsmc will create new filespace during backup processing.
mmbackup --rebuild OR mmbackup -q	dsmc query backup -detail	if no shadow db can be found or restored full filespace query, single-threaded.
mmbackup -t full	dsmc query inclexcl	
	dsmc selective backup -filelist	for all file system objects.
mmbackup -t incremental	dsmc query backup /<fsname>/.mmbackupShadow*	if no shadow DB is present.
	dsmc query inclexcl	
	dsmc expire backup -filelist	for file system objects deleted in file system to be expired in IBM Spectrum Protect DB.
	dsmc incremental backup -filelist	for file system objects that need POSIX attributes and migration state updates in IBM Spectrum Protect DB. Note: if the file has ACL or EA metadata the file data will be sent to the IBM Spectrum Protect server.
	dsmc selective backup -filelist	for file system objects that need file data update.
mmbackup --expire-threads		defines number of dsmc expire commands started per node.
mmbackup --backup-threads		defines number of dsmc selective or incremental commands started per node.

http://www-01.ibm.com/support/docview.wss?uid=swg21999651

IBM Spectrum Scale mmbackup on root directories level

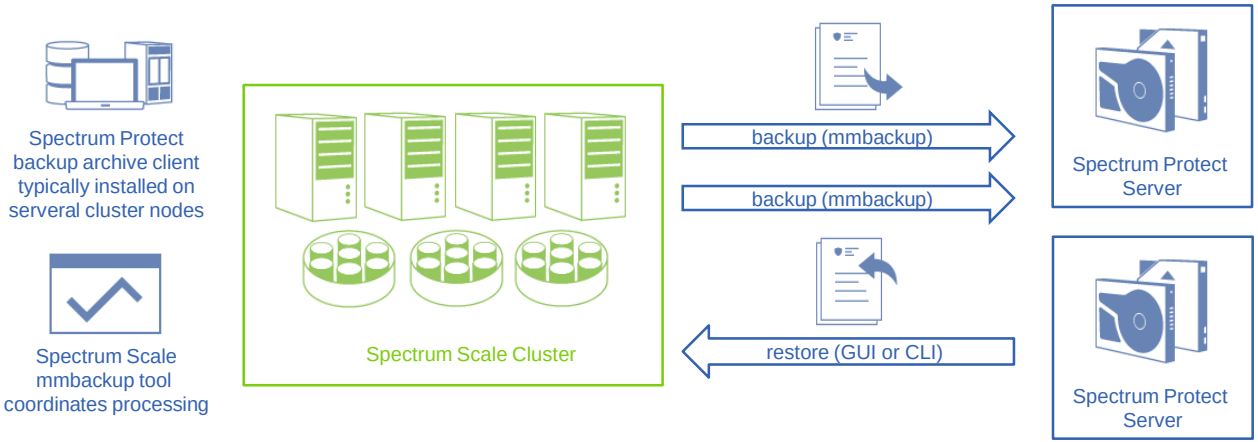

- Environment: Large IBM Spectrum Scale installations with a double digit number of nodes and file systems. IBM Spectrum Protect backup archive client installed on several cluster nodes
- Scalability: Hundreds of millions of files, Hundreds of terrabytes of data

IBM Spectrum Scale mmbackup on root directories level

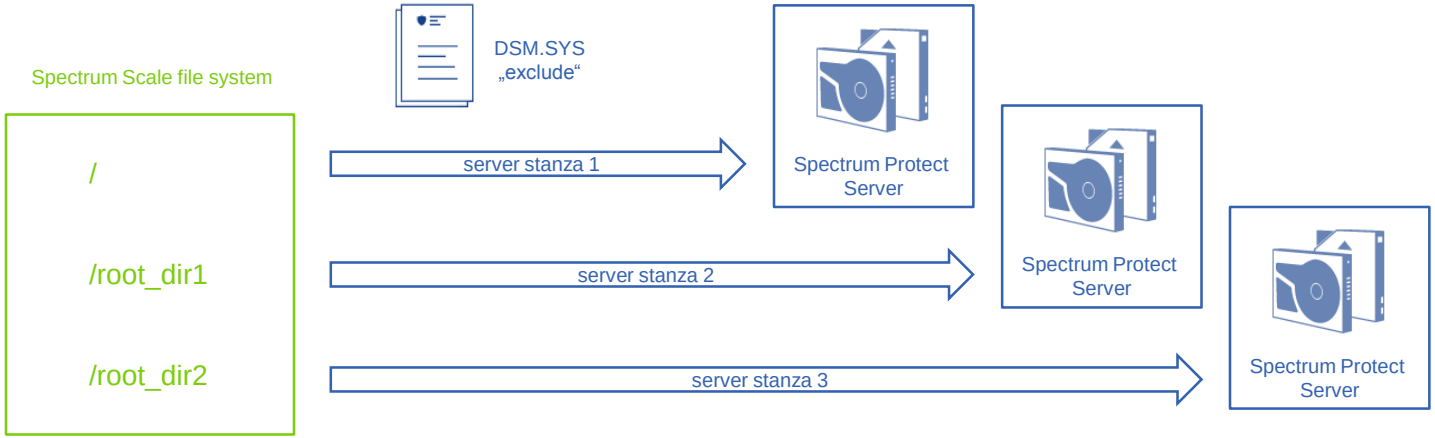

- Processing: IBM Spectrum Protect exclude processing used to divide file system into backup parts on root directory level. One part goes to one server. IBM Spectrum Scale mmbackup is used to backup parts to different servers using –tsm-servers option.
- Pros: usable on existing data w/o full backup, scalable, IBM Spectrum Protect server housekeeping can be parallelized
- Cons: complex planning and setup, IBM Spectrum Scale mmbackup sequential processing

Petascale Data Protection

The singificant grow of data faces storage providers with new challanges. Beside the administration and maintenance of disk pools for large file systems the data protection and data archiving of big data clusters causes serious demands. The following slides describe a solution for data protection for large scaling environments with IBM Spectrum Protect and IBM Spectrum Scale.

- This slide deck corresponds to the whitepaper „Peta Scale Data Protection“
- Link to the paper:
  https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/Tivoli%20Storage%20Manager/page/Petascale%20Data%20Protection
- The paper describes a data protection approach scaling up to hundreds of petabytes for an IBM Spectrum Scale file systems using IBM Spectrum Protect backup-archive client and IBM Spectrum Protect for Space Management. The focus of this paper is to provide configuration guidance for the setup and operation of the data protection processes in such an environment.
- This paper also introduces the concept of different service levels for data protection on file system and fileset level.

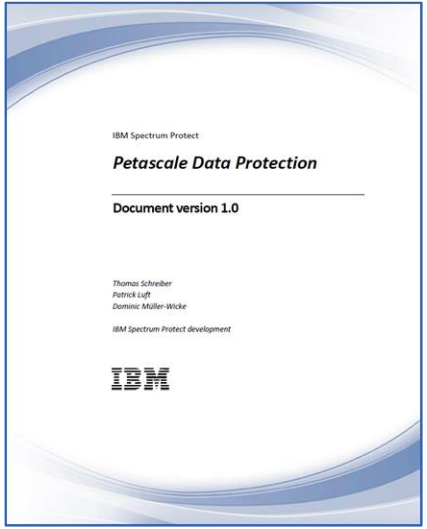

Peta Scale Data Protection – Architecture

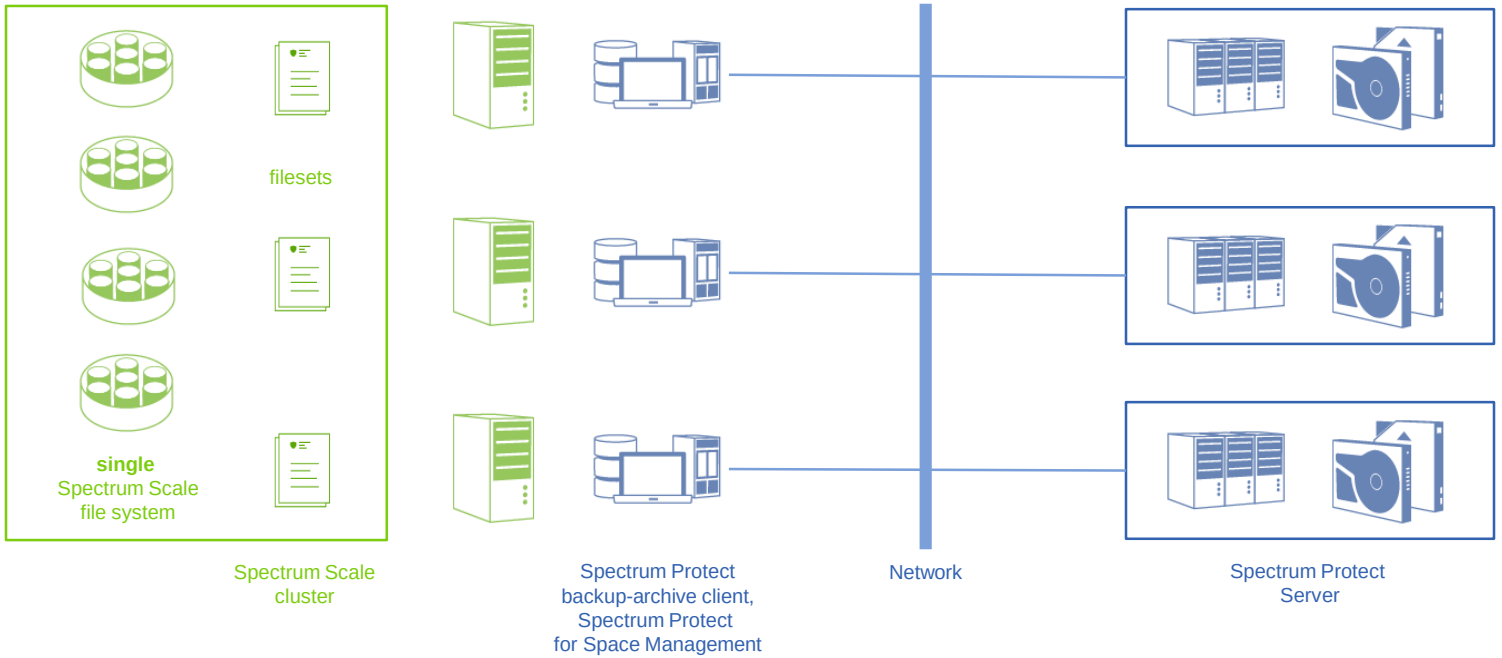

Peta Scale Data Protection – Key features

- Extreme scalability due to multiple Spectrum Protect servers to protect a single Spectrum Scale file system
- High backend storage media flexibility due to multiple supported storage technologies (disk, tape, cloud) for a single file system
- High QoS flexibility due to fine grain data protection approach (fileset level)
- Integration between Spectrum Protect client products warranted (inline copy)
- Ultra fast disaster recovery with SOBAR supported

Peta Scale Data Protection – Technology

- Key technology behind the solution is Spectrum Protect „active server binding“ that is implemented from Spectrum Protect for Space Management and used from Specturm Protect backup-archive client.
- Usage of Spectrum Protect for Space Management (HSM) for file migration is optional, but file system management is required for active server binding. HSM is mandatory if fast disaster recovery with SOBAR is planned.
- The first time a file will be send from file system to the Spectrum Protect server (backup or HSM) it will be bound to the specified server.
- Granularity of backup and HSM processing is Spectrum Scale fileset level. The backup and HSM processing for each fileset is independent from others.
- Active server binding is visible for Spectrum Scale policy engine scans

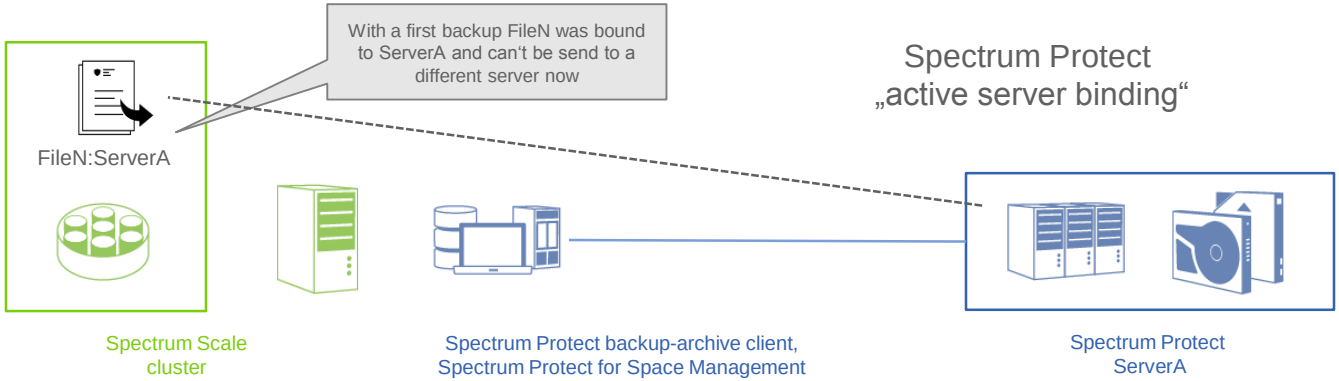

Peta Scale Data Protection – Usage Scenario

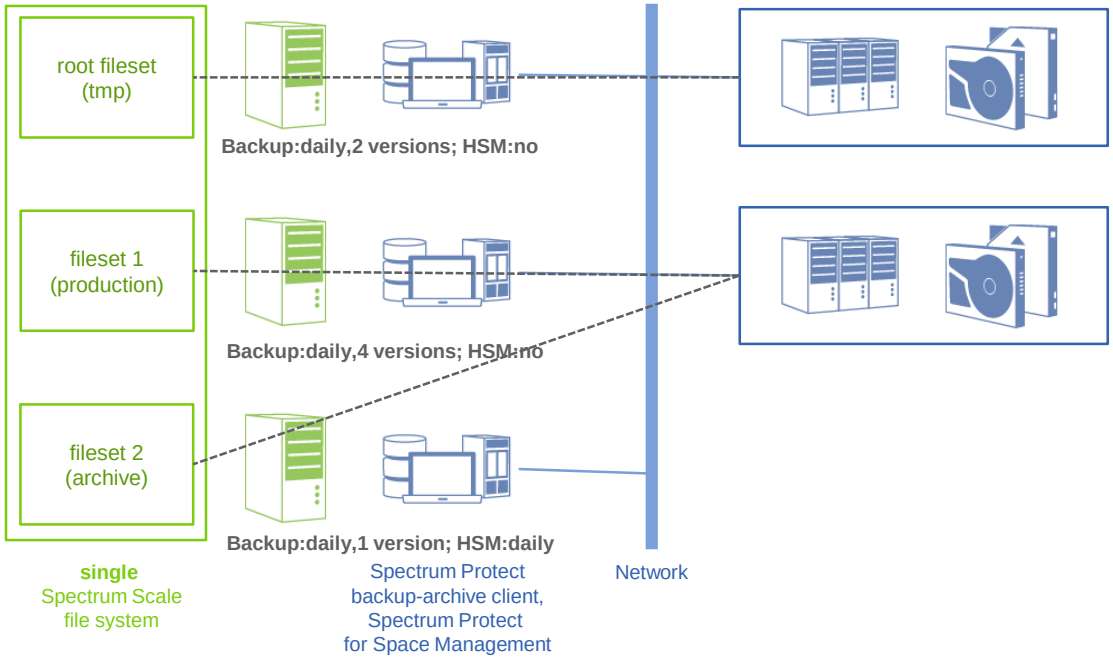

- Spectrum Scale file system contains three filesets (root, 1, 2)
- root fileset has binding to server 1
- fileset 1 and 2 have binding to server 2
- fileset 1 contains production data that is frequently changed and needs more backup versions. Due to frequent changes HSM is not required here
- fileset 2 contains archive data that is typically unchanged after creation. Due to this the data will be archived to high latency and low cost media with HSM

Performance Considerations

IBM Spectrum Scale mmbackup – performance considerations

File system scan and preprocessing	Include / Exclude processing	Filelist processing	Parallelism in backup and expire processing
Considerations to optimize general options and workload balancing	Considerations to optimize canadiate selection	Considerations to optimize filelist size	Considerations to optimize session and transaction handling

Workload on cluster nodes

1. mmbackup and the policy engine use a significant amount of memory and CPU cycles on all used nodes in the cluster
2. The amount of data and metadata I/O is high for both policy engine (inode scan) and mmbackup (shadowDB)

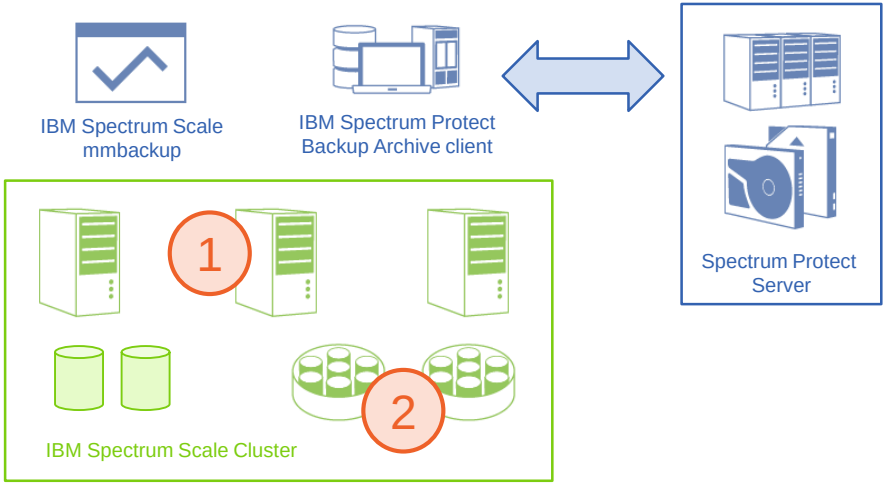

- Ensure high I/O performance of your storage system and storage network
- Serialize backup of different file systems
- Share workload of backup between nodes

Example: Cluster has nodes N1 to N4 and files systems FS1 to FS4

- Run mmbackup on nodes N1,N2 for FS1 and on nodes N3,N4 for FS2, ... in parallel
- After FS1 has finished run mmbackup on nodes N1,N2 for FS3, ...

Workload on cluster nodes

3 Both tools create temporary files in global and local working directories what requires free space

4 mmbackup‘s server queries and multiple backup sessions load network and server

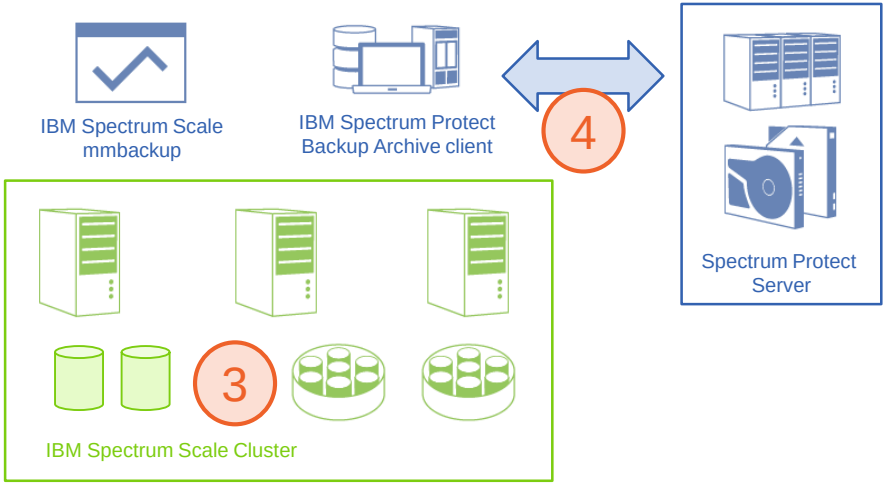

- Choose your global and local working directories wisely. Prevent out of storage conditions
- Check and shrink your log files on a regular base. Note that the Spectrum Protect Backup-Archive client 7.1.6 has instrumentation enabled per default now.
- Ensure high network bandwith of all used cluster nodes to the IBM Spectrum Protect server. Use different networks for client and server workloads.

General recommendations

- Use latest versions: Spectrum Scale 4.x.x and Spectrum Protect 7.1.x have good improvements for mmbackup
- Do not mix OS types. Run mmbackup on either AIX, xLinux, pLinux or zLinux nodes
- Rebuild of the shadowDB takes time. Use option –q or –rebuild as seldom as possible.
- Consider Spectrum Protect character limitations (especially for environments including Windows or machine generated file names)
  - Files with control-X, control-Y, carriage return and the new line character in their name can't be backed up to Spectrum Protect.
  - Use `QUOTESARELITERAL` (if mmbackup is used with --noquotes), if file names contain " or ' .
  - Use `WILDCARDSARELITERAL`, if file names contain * or ?.
- Do not use the below Spectrum Protect processing options:
  - `SUBDIR YES` (performance killer, costs inactive backup versions)
  - `QUIET`
  - `SCROLLPROMPT`
  - `SCROLLLINES`

  } (show stopper)

General recommendations

Recommendations on mmbackup options

- `-S:` Use snapshots to reduce transaction failures due to re/moved files
- `-m:` DO NOT USE! Use fine grain options instead
  - `--expire-threads`
  - `--backup-threads`
- `-B:` DO NOT USE! Use fine grain options instead
  - `--max-backup-size`
  - `--max-expire-size`
  - `--max-backup-count`

Include / Exclude Processing

- Spectrum Protect offers a rich set of include and exclude options to control which files and directories are backed up.
- mmbackup is building these options into its policy for backup

- Include and Exclude options may have significant impact on scan performance

- Some rules to consider:
  - Use as few EXCLUDE statements as possible
  - Aviod using INCLUDE. Use EXCLUDE instead
  - Do not use „EXCLUDE /dir/.../*“ . Try EXCLUDE.DIR instead.
  - Avoid EXCLUDE and INCLUDE for the same subtree, like

    ```
    exclude /home/dominic*
    include /home/dominic/important*
    ```

  - If INCLUDE is only used to assign right management class in Spectrum Protect

    „INCLUDE <pattern> MGMT“

    use mmbackup service flag is used MMBACKUP_IGNORE_INCLUDE

    ```
    export MMBACKUP_IGNORE_INCLUDE=1
    ```

- Technote on this theme: http://www-01.ibm.com/support/docview.wss?uid=swg21699569

Directory trees and files

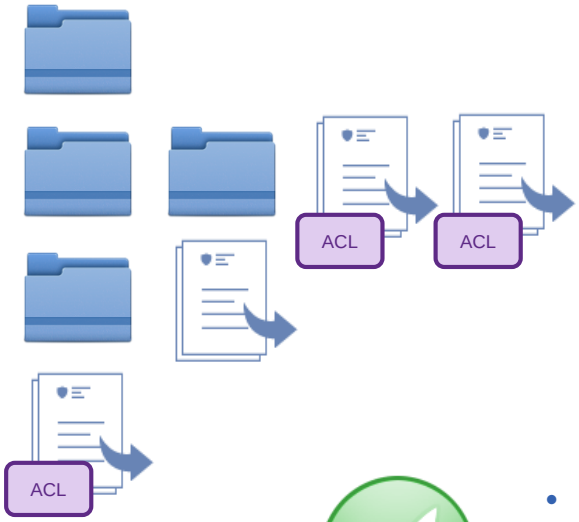

1. The combination of filsystem mount point, directory path and file name is unique identifier for a backup object. Changes lead to new backup.
2. The IBM Spectrum Protect server stores ACL / EA metadata in combination with file data. Changes lead to new backup.

- Prevent move or rename of files or directories. These changes lead to a new backup of all affected files
- If ACL or EA metadata is used prevent changes of POSIX attributes. These changes lead to a new backup of all affected files

Filelist processing

- IBM Spectrum Scale mmbackup processes filelists
- Three global filelists contain all files and file system objects that must be expired, updated, send to the IBM Spectrum Protect server. The backup processing happens on small chuncks of the global lists

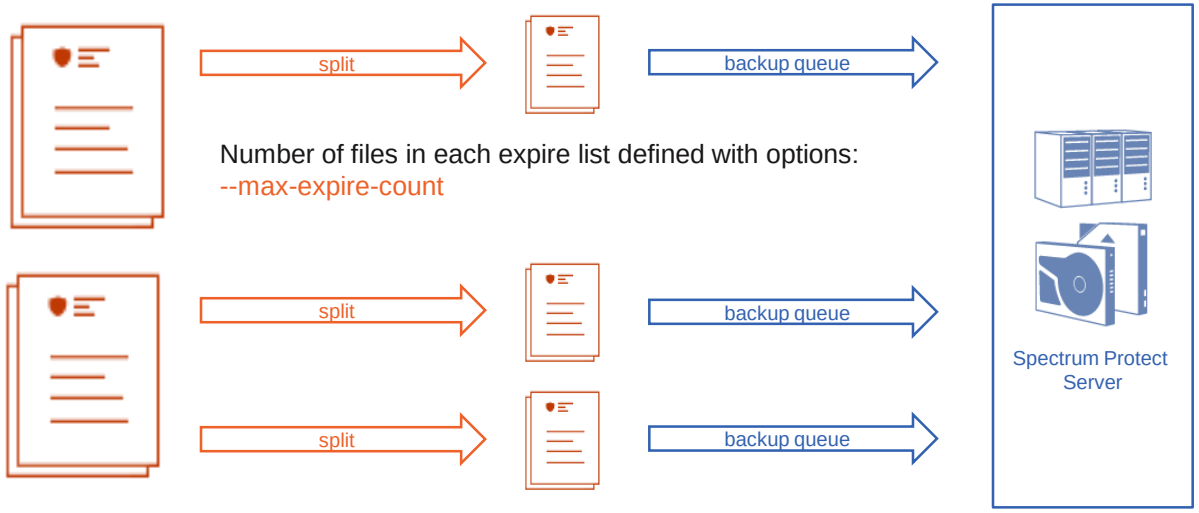

Filelist processing

- Each file list processed from mmbackup starts one IBM Spectrum Protect Backup-Archive client dsmc command via CLI
- Depending on the server sessions related settings one or more sessions will be opened from each process
- Server logon and sessions creation is more expensive compared to backup or expiration transactions

- Use higher values for –max-backup-count/size if you don‘t observe transaction issues
- Use the multiple of the value you use for the mmbackup option –max-backup-count for the IBM Spectrum Protect option TXNGROUPMAX.
- Use the multiple of the value you use for the mmbackup option –max-backup-size for the IBM Spectrum Protect option TXNBYTELIMIT.
- Both settings ensure multiple (perfectly alligned) transactions inside a single server session.

- Use the multiple of the value you use for the mmbackup option –max-expire-count for the IBM Spectrum Protect option TXNGROUPMAX.

Option RESOURCEUTILIZATION

- RESOURCEUTILIZATION option defines number of consumer and producer threads in IBM Spectrum Protect Backup-Archive client dsmc command for backup
- Table shows values and session numbers for backup
- Expiration processing uses only one sessions

Value	#Sessions (send+query)
1	1
0 (default), 2	2 (1+1)
3, 4	3 (2+1)
5, 6	4 (3+1)
7	5 (4+1)
8	6 (5+1)
9	7 (6+1)
10	8 (7+1)

Calculate the number of server sessions for backup

Number of nodes specified with mmbackup paramter -N

Number of threads specified with mmbackup paramter --backup-threads

Number of threads specified with BA client option RESOURCEUTILIZATION

Number of parallel mmbackup runs (file system or fileset)

- The number of parallel backup sessions must be below your setting for MAXSESSIONS
- The available mount points defined with MAXNUMMP must be higher than this calculation result.
- The maximum values for a given MAXNUMMP can be caluclated as follows:

```
#backup-threads = #mount-points / (#nodes * (RESOURCEUTILIZATION[VALUE] - 1))
```

Calculate the number of server sessions for expiration

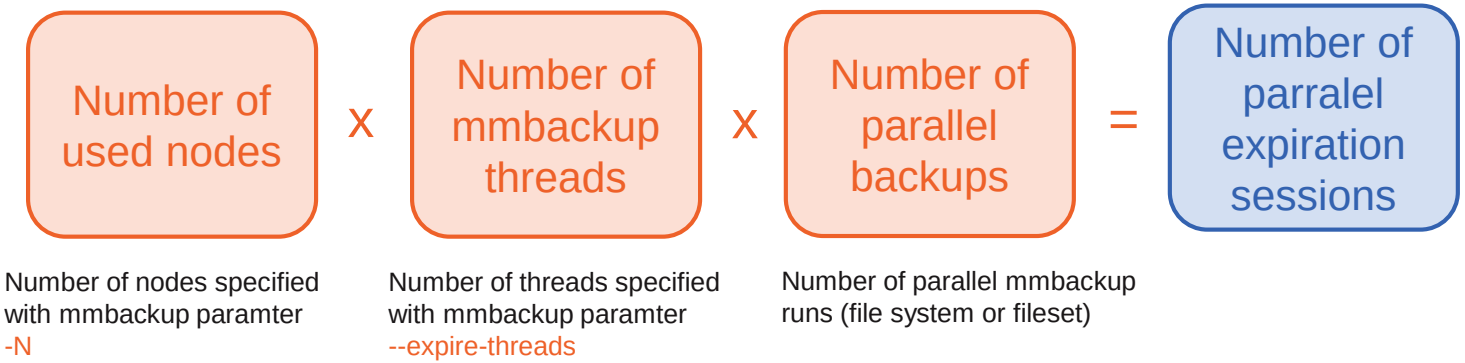

- The number of parallel expire sessions must be below your setting for MAXSESSIONS
- Play with the values for –backup-threads and –expire-threads if your business process allows this
- Keep in mind that parallel processed file system backup might have different schedules for expiration and backup processing
- Keep in mind that parrallel restore processing could happen

Result evaluation

IBM Spectrum Protect Backup-Archive client generates four different return codes:

- 0 : All operations completed successfully.
- 4 : The operation completed successfully, but some files were not processed. There were no other errors or warnings.
  - The file satisfies an entry in an exclude list.
  - The file was in use by another application and could not be accessed by the client.
  - The file changed during the operation to an extent prohibited by the copy serialization attribute.
- 8 : The operation completed with at least one warning message.
- 12 : The operation completed with at least one error message (except for error messages for skipped files).

Result evaluation

IBM Spectrum Scale mmbackup reduces the number of return codes to 0, 1 and 2:

- 0 : 100% perfect execution with all the intended files are now backed up.
- 1 : IBM Spectrum Protect experienced a 4, 8, or 12 return code and some file or files were not processed.
- 2 : means a more severe problem happened on IBM Spectrum Scale mmbackup side of the processing

Note: You should no longer trust the shadowDB in this case. Repair techniques:

- Find .mmbackupShadow.#.<tsm-server>.{filesys,fileset}.old and rename {.old,} to make mmbackup use last known good shadow DB
- OR: If old Shadow DB is not found in file system for some reason, try
  - `dsmc q backup "/<gpfs mount>/.mmbackupShadow.*"`
  - And locate a backed up, recent Shadow DB file to restore to resume backups
- OR: Rebuild the shadow DB file from current inventory (very expensive in time) by using "--rebuild" or "-q" option
- OR: Run mmbackup -t full (very very expensive) and back up everything in the file system

IBM Spectrum Protect for Space Management and container pools

- IBM Spectrum Protect for Space Management was successfully tested with container pools for both options disk and cloud
- The IBM Spectrum Protect server function In-Line copy of files is not supported for container pools. Therefore backup of files that are migrated may fail with an error.
- If you use IBM Spectrum Protect for Space Management AND IBM Spectrum Protect Backup-Archive client on the same file system data ensre that the option **MIGREQIRESBACKUP YES** is enabled.
- Work with your users that renaming of files and move operations are prevented as good as possible.
- Search the web for IT16799 (will show up in the next days)

- Beginning with Spectrum Scale 4.2.1 the command mmchfileset can be used to prevent POSIX changes on files and directories (see: manpage mmchfileset „ `--allow-permission-change`“)

References

IBM Knowledge Center
IBM Spectrum Scale: http://www.ibm.com/support/knowledgecenter/STXKQY/ibmspectrumscale_welcome.html
IBM Spectrum Protect: http://www.ibm.com/support/knowledgecenter/SSGSG7/landing/welcome_ssgsg7.html

IBM Spectrum Protect resources landing page
http://www.ibm.com/support/docview.wss?uid=swg21684850

Petascale Data Protection
https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/Tivoli%20Storage%20Manager/page/Petascale%20Data%20Protection

Overview on Spectrum Protect – Spectrum Scale Integration
https://www.ibm.com/developerworks/community/wikis/home/wiki/Tivoli%20Storage%20Manager/page/Integrating%20IBM%20Tivoli%20Storage%20Manager%20with%20IBM%20Elastic%20Storage

Configuration of Spectrum Protect for Spectrum Scale AFM
https://www.ibm.com/developerworks/community/wikis/home/wiki/Tivoli%20Storage%20Manager/page/Configuring%20IBM%20Spectrum%20Scale%20Active%20File%20Management

Spectrum Protect for Space Manangement whitepaper
Setup policy driven threshold migration: http://www.ibm.com/support/docview.wss?uid=swg27018848
Setup cross platform cluster: http://www.ibm.com/support/docview.wss?uid=swg27028178

YouTube
IBM Spectrum Protect - mmbackup general functions https://youtu.be/3PMO4Sdegs0
IBM Spectrum Protect - mmbackup tweaks for max performance https://youtu.be/sg4FrZHi99Y
IBM Spectrum Protect using Scale for db, logs & storage pools https://youtu.be/vIobC2MDIlE

Thank you

Presenter vCard: