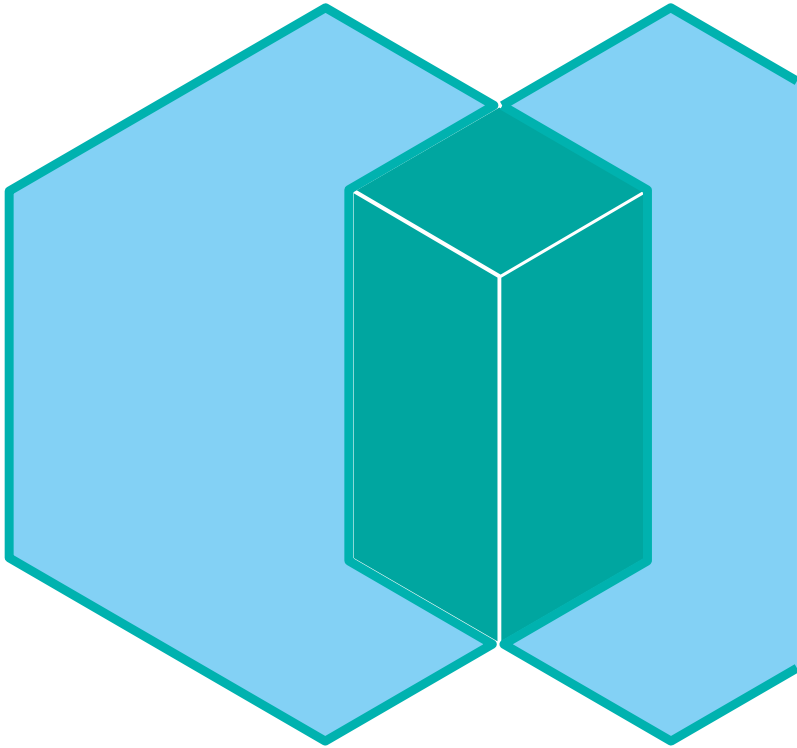

IBM Spectrum Scale

– Recent Updates and Outlook –

Spectrum Scale Expert Workshop 2016 – Ehningen

Mar 9, 2016 – Ulf Troppens

New in Spectrum Scale 4.2
Priorities 2016
Miscellaneous

Outline

New in Spectrum Scale 4.2

The History of Spectrum Scale

This infographic is the genealogy of IBM Spectrum Scale, from it's birth as a digital media server and HPC research project to it's place as a foundational element in the IBM Spectrum Storage family. It highlights key milestones in the product history, usage, and industry to convey that Spectrum Scale may have started as GPFS, but it is so much more now. IBM has invested in the enterprise features that make it easy to use, reliable and suitable for mission critical storage of all types.

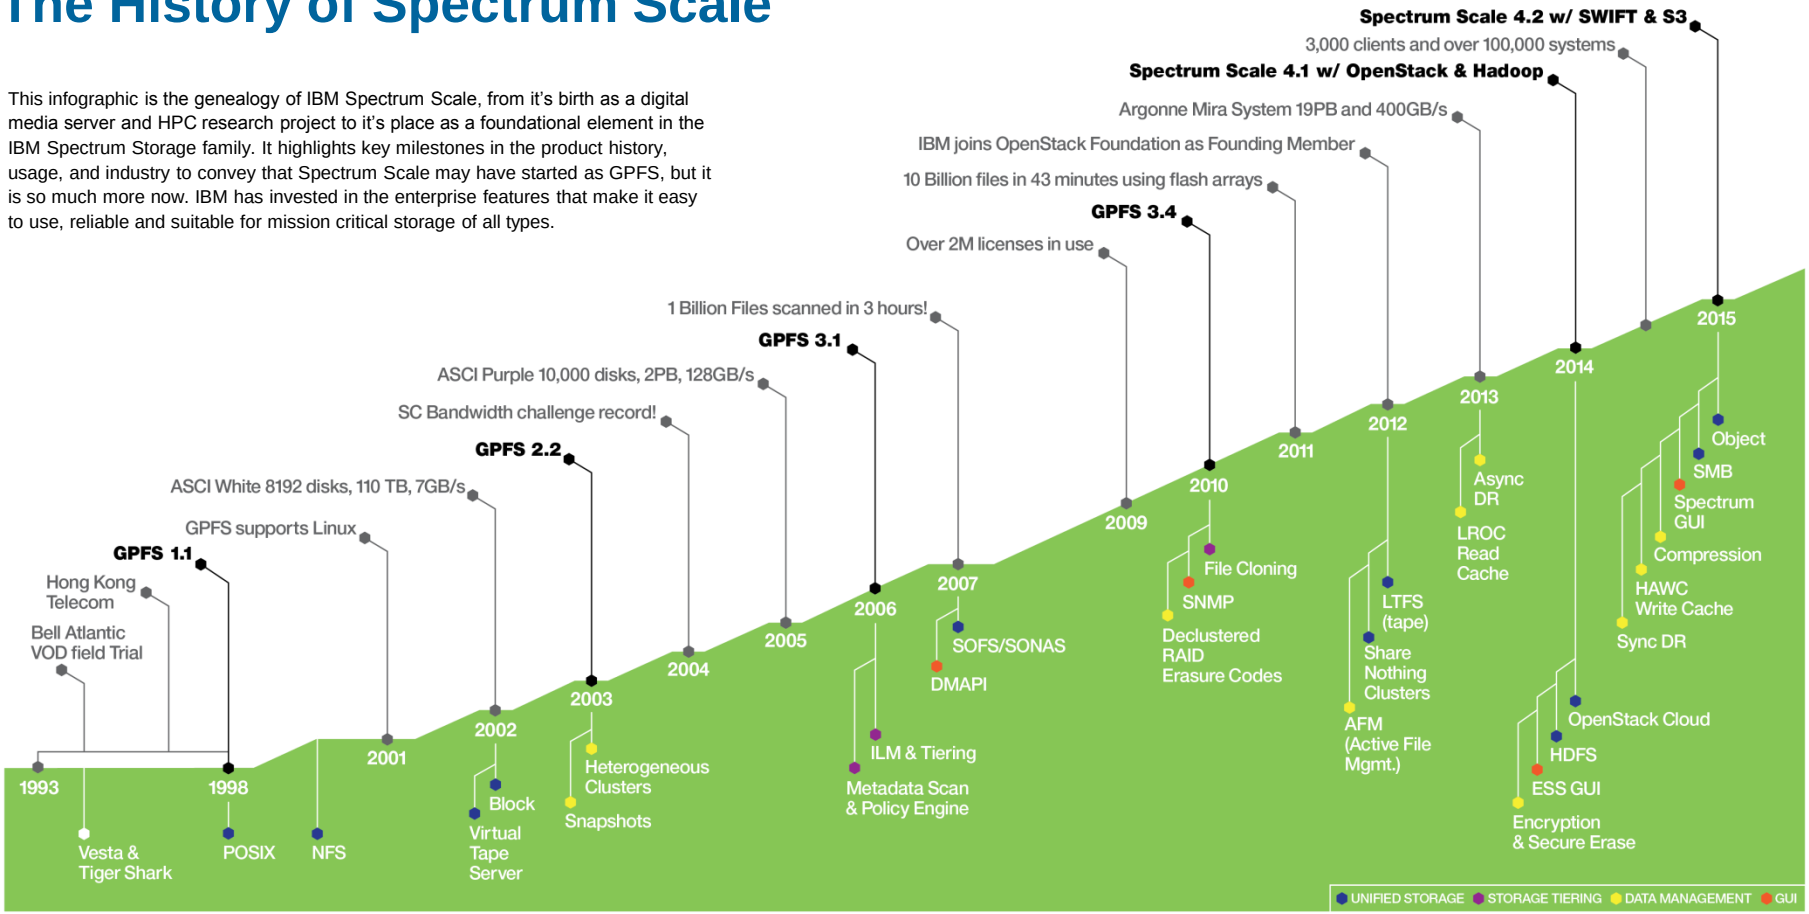

Store everywhere. Run anywhere.

Remove data-related bottlenecks

Challenge

- Managing data growth
  - Lowering data costs
  - Managing data retrieval & app support
  - Protecting business data

Unified Scale-out Data Lake

- File In/Out, Object In/Out; Analytics on demand.
- High-performance native protocols
- Single Management Plane
- Cluster replication & global namespace
- Enterprise storage features across file, object & HDFS

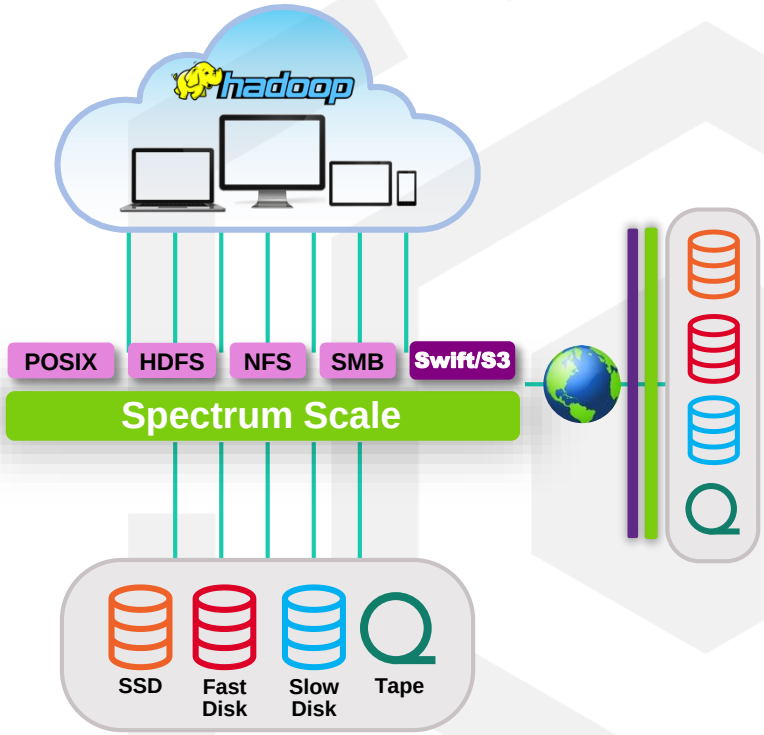

Store everywhere. Run anywhere.

Content Repositories

Challenge

Object storage for static data

- Seamless scaling
- RESTful data access
- Object metadata replaces hierarchy
- Storage efficiency

Spectrum Scale Swift & S3

- High-performance for object
- Native OpenStack Swift support w/ S3
- File or object in; Object or file out
- Enterprise data protection
- Spectrum Scale RAID (ESS)
- Transparent ILM
- Encryption of data at rest and Secure Erase

Store everywhere. Run anywhere.

Analytics without complexity

Challenge

Separate storage systems for ingest, analysis, results

- HDFS requires locality aware storage (namenode)
- Data transfer slows time to results
- Different frameworks & analytics tools use data differently

HDFS Transparency

- Map/Reduce on shared, or shared nothing storage
- No waiting for data transfer between storage systems
- Immediately share results
- Single ‘Data Lake’ for all applications
- Enterprise data management
- Archive and Analysis in-place

➔ Analyze object and file data without copying into HDFS

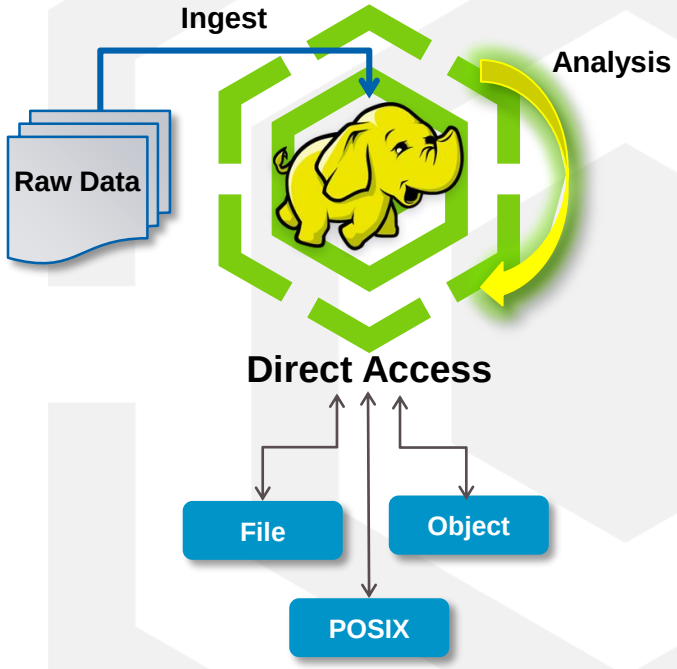

Backup Of Large Spectrum Scale File Systems

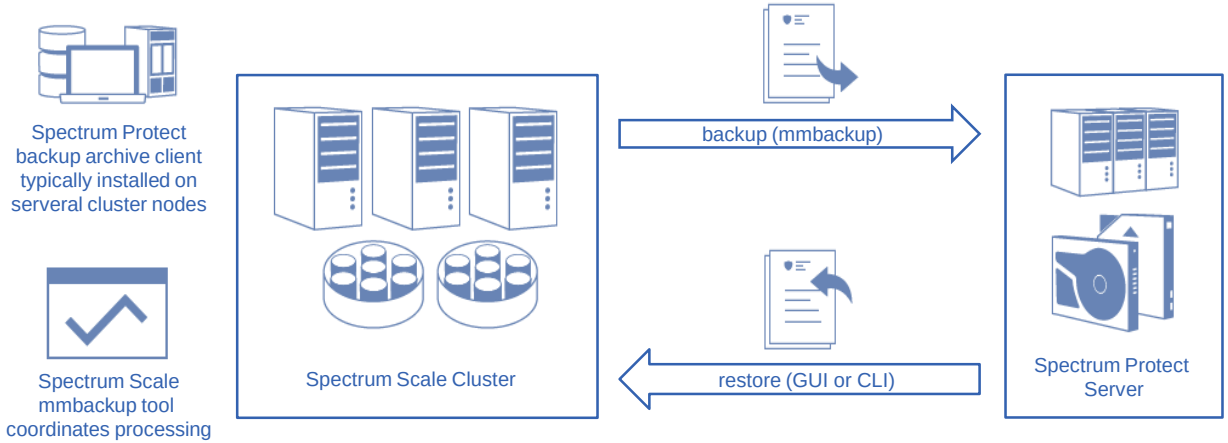

Function

- Massive parallel filesystem backup processing
- Spectrum Scale mmbackup creates local shadow of Spectrum Protect DB and uses policy engine to identify files for backup
- Spectrum Protect backup archive client is used under the hood to backup files to Spectrum Protect Server
- Spectrum Protect restore (CLI or GUI) can be used to restore files

➔ Use any backup program to backup file, object and Hadoop data
➔ Use Spectrum Protect to benefit from mmbackup and SOBAR to backup and restore huge amounts of data

New in Spectrum Scale 4.2

	New Feature	Benefit
Client Experience Focus	• Common interface across Spectrum Portfolio • GUI Phase 1	Easy to learn UI and integration across Spectrum Storage portfolio Simplify common management functions, including • Enabling protocols • Policy driven placement and ILM • Monitoring • Troubleshooting
Object Storage	• Unified File and Object • Extended S3 API support	Single view of data with wither file or object read and write Enable applications originally written for AWS
Big Data & Analytics	• Native Hadoop Support • Ambari Integration	Higher performance and broader integration with HDFS applications to go beyond Hadoop and embrace Map/Reduce ecosystem
Storage efficiency	• Compression of Cold data for File & Object	• Improve Storage utilization & efficiency for Cold data • Efficienciently reduce data size using compression policies
General	• Quality of Service for File • z Linux support • Sudo wrappers	Expanding functionality in Spectrum Scale data aware policy engine: • Performance reservations to meet SLAs – even by time of day • Extending multi-site resiliency features to z-Linux

Speed and simplicity: Performance monitoring highlights

System health
Node performance
Network traffic
Historical trends

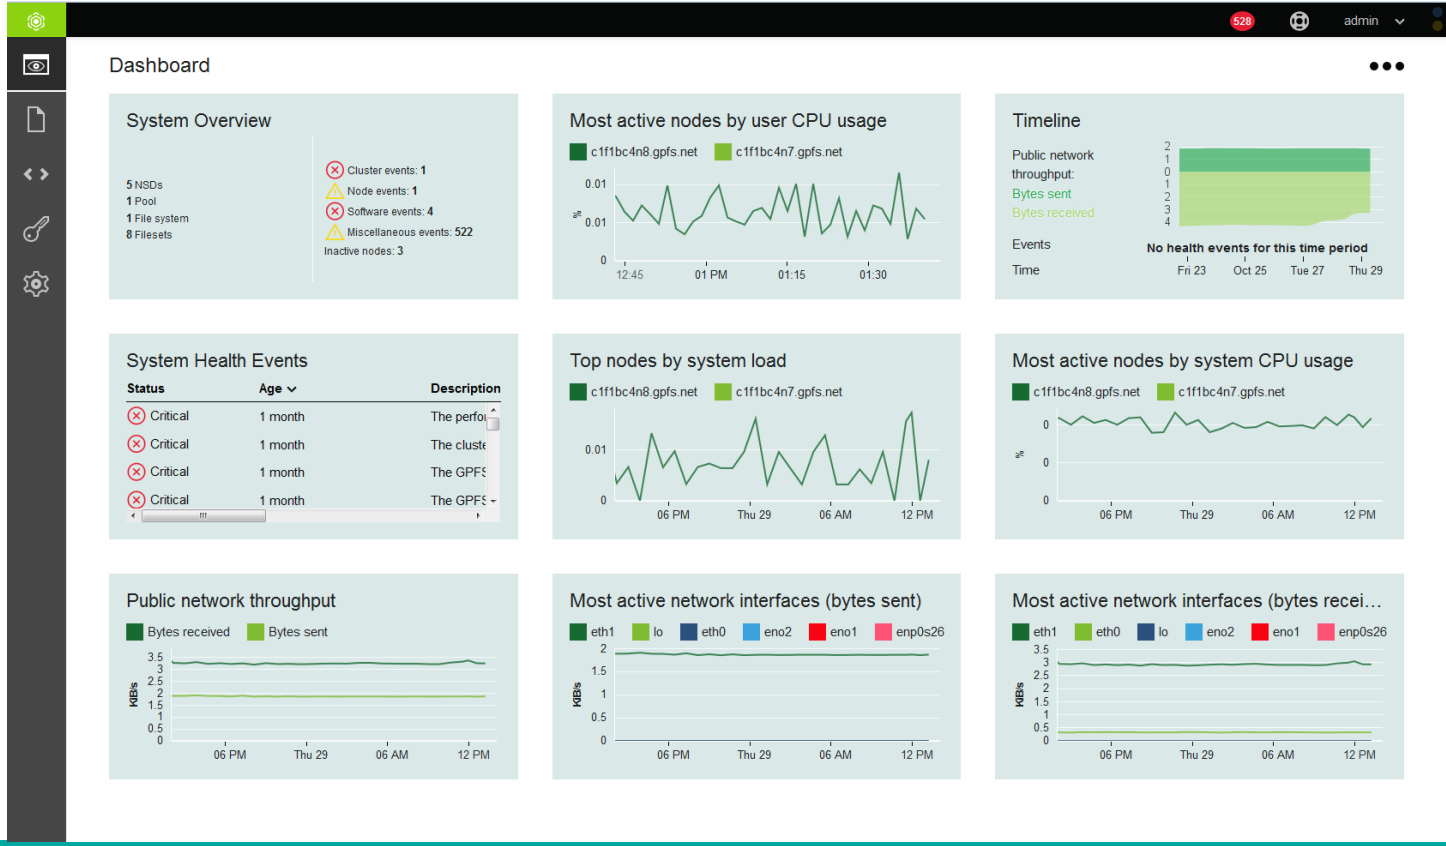

Priorities 2016

Disclaimer

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.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon many factors, including considerations such as the amount of multiprogramming in the user’s job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve results similar to those stated here

2016 Development Priorities

Every year we define a set of goals

- Based mainly on client feedback and market opportunity
- Target is to achieve them within the year

Focus areas

- Problem determination
- Documentation
- Security
- Defect backlog

Functional enhancements

- Improvements for Big Data
- More flexibility for GNR

Hills – Problem Determination

1. An IT administrator who monitors Spectrum Scale can be made aware of the health of his Spectrum Scale components in one cluster, from a single place.

2. An IT Administrator, can perform self-service problem determination by utilizing provided guidance or automated solutions to problems, without contacting IBM Support.

3. An IT Administrator, can pre-check/check Spectrum Scale and its operating environment to avoid potential problems after initial installation or when changes are made, from a single tool.

Simplicity

Subject to change.
Details are under investigation.

- Spectrum Scale provides a plenty of parameters which allow tuning for a broad range of workloads by an expert user
- Simplicity replaces those parameters by a few aggregated parameters which enable an average skilled user to tune Spectrum Scale for the most common workloads

➔ Simplicity is problem prevention

Software Configuration and Tuning - Physical

Spectrum Scale Client Nodes		
Parameter Name	**Value**	**Description**
deadlockDetectionThreshold	0	Disables automatic deadlock detection.
deadlockOverloadThreshold	0	Disables automatic deadlock detection.
flushedDataTarget	1024	Sets the maximum number of open file objects for which data have already been flushed.
flushedInodeTarget	1024	Sets the maximum number of open file objects for which data and metadata have already been flushed.
idleSocketTimeout	0	Disables timeouts for idle sockets.
ignorePrefetchLUNCount	1	Disables automatic determination of maximum prefetch requests based on visible LUN count. The maximum prefetch requests are instead determined by prefetch buffers and prefetch threads.
inodeXWPrefetchThresholdCount	0	Enables prefetching of inode token in exclusive mode.
logBufferCount	50	Sets the number of log buffers.
logBufferSize	1M	Sets the size of each log buffer.
logPingPongSector	0	Disables the use of 'ping-pong' sectors in logging. The feature is unnecessary when using ESS.
logWrapAmountPct	2	Sets the percentage amount by which the log recovery point is advanced during the flushing of modified log entries.
logWrapThreads	128	Sets the number of threads to use for the flushing of modified log entries.
logWrapThreadsPerInvocation	128	will look in docs for concise description
logWrapThresholdPct	20	Sets the log capacity percentage at which the log flushing algorithms are triggered.
maxActiveIallocSegs	8	Sets the maximum number of active inode allocation segments per node.
maxAllocRegionsPerNode	32	Sets the maximum active allocation regions per node for disk allocation.
maxBackgroundDeletionThreads	128	Sets maximum number of threads to use for file deletions.
maxblocksize	16M	Sets the maximum file system block size.
maxBufferCleaners	1024	Sets the maximum number of threads for cleaning data buffers.
maxFileCleaners	1024	Sets the maximum number of threads for flushing data and metadata.
maxFilesToCache	6291456	Sets the maximum number of files to cache.
maxGeneralThreads	2048	Sets the maximum number of non-critical daemon worker threads.
maxInodeDeallocPrefetch	32	Sets the maximum number of threads that prefetch inode tokens of deleted files.

maxBufferCleaners	1024	Sets the maximum number of threads for cleaning data buffers.
maxBufferDescs	2M	Sets the maximum number of buffer descriptors.
maxFileCleaners	1024	Sets the maximum number of threads for flushing data and metadata.
syncWorker1Threads	256	Sets the maximum number of threads to use flush data during explicit sync calls.
worker1Threads	1024	Sets the number of threads used by Spectrum Scale to handle I/O requests.
worker3Threads	64	Sets the number of inode prefetch threads to use.

Security Work 2016

Subject to change.
Details are under investigation.

Sudo wrapper / no root ssh

- Make GUI functional

File encryption (on rest)

- Consumability improvements in the configuration of SKLM
- Support for the Vormetric key server
- File encryption performance (whitepaper)

Authentication

- GUI admin user can authenticate via external AD or LDAP server (delivered with 4.2.0-1)
- External Keystone SSL support for object

Miscellaneous

- Spectrum Scale security best practices (whitepaper)
- Multi-region object deployment with a highly available keystone service (whitepaper)

Miscellaneous

IBM Spectrum Scale Experten Workshop: Agenda

Tag 1: 9. März 2016

10:00 – 10:15 Uhr	Welcome	Goetz Mensel
10:15 – 10:45 Uhr	IBM Storage Technical Strategy Summary	Robert Haas
10:45 – 11:15 Uhr	IBM Spectrum Scale Update & Directions	Ulf Troppens
11:15 – 11:45 Uhr	Agile collaboration with customers	Alexander Wolf-Reber
11:45 – 12:15 Uhr	Cloudy jigsaw puzzles	Harald Seipp
12:15 – 13:15 Uhr	Lunch	
13:15 – 13:45 Uhr	IBM Elastic Storage Server Update & Directions	Falk Steinbrück
13:45 – 14:15 Uhr	Large Scale Video Server - ein ESS und CES Erfahrungsbericht	Jochen Zeller
14:15 – 14:30 Uhr	Short Break	
14:30 – 15:00 Uhr	Accelerating and simplifying backup with Spectrum Protect on Spectrum Scale	Andre Gaschler / Nils Haustein
15:00 – 15:30 Uhr	Data protection with Spectrum Protect	Dominic Müller-Wicke
15:30 – 16:00 Uhr	Break	
16:00 – 17:00 Uhr	Problem Determination	Mathias Dietz
17:00 – 17:30 Uhr	GUI	Markus Rohwedder
17:30 – 18:00 Uhr	IBM Spectrum Scale File Protocols - NFS and SMB on CES nodes	Ingo Meents
18:00 – 18:15 Uhr	Meet the developers	Goetz Mensel
18:15 – 20:30 Uhr	Get-together	All

Outlook
Solutions
Access

IBM Spectrum Scale Experten Workshop: Agenda

Tag 2: 10. März 2016		
09:00 – 09:30 Uhr	Licensing	Heiko Lehmann
09:30 – 10:00 Uhr	Workflows with AREMA	Ulrich Voigt
10:00 – 10:30 Uhr	Spectrum Scale for Hadoop	Olaf Weiser
10:30 – 11:00 Uhr	OpenStack integration	Harald Seipp
11:00 – 11:30 Uhr	Break	
11:30 – 13:00 Uhr	Spectrum Scale Updates	Olaf Weiser
13:00 – 13:45 Uhr	Lunch	
13:45 – 15:30 Uhr	News from Almaden	Sven Oehme
15:30 – 16:00 Uhr	Closing	Goetz Mensel

Outlook

Solutions

Access

Trival VM and Open Betas

IBM Spectrum Scale Trial VM

This Trial VM offers fully pre-configured IBM Spectrum Scale instance in a virtual machine based on IBM Spectrum Scale 4.2 GA version. The download bundle includes the virtual image and the requisite guides (Quick Start guide, Explore guide and Advanced guide) allowing you to try the key features in minutes. Use the Quick Start guide for installation instructions. The Explore guide provides step-by-step instructions to try our unified file & Object as well as GUI functionality.

Use IBM Spectrum Scale Forum or mail to scale@us.ibm.com to ask questions and to give your feedback.

Date	Type	Description	Download
14 Jan 2016	Evaluation	VM with pre-configured IBM Spectrum Scale	Download

IBM Spectrum Scale GUI Open Beta

The IBM Spectrum Scale GUI (Graphical User Interface) is the graphical interface for IBM Spectrum Scale. This can be used in conjunction with our existing command line interface. You are invited to test the latest pre-GA beta version of our GUI here. You can try this GUI with your existing 4.1+ installs of Spectrum Scale but to use all features we suggest the latest GA version which is 4.2.0-1.

Try the beta build below. Use IBM Spectrum Scale forum or mail to GUI Feedback to ask questions and to give your feedback.

Date	Type	Description	Download
18 Feb 2016	Open Beta	IBM Spectrum Scale GUI	Download

IBM Spectrum Scale transparent cloud tiering

The IBM Spectrum Scale transparent cloud tiering is a new hybrid cloud storage capability for IBM Spectrum Scale. This feature allows cloud storage to be used as a storage tier in the same manner as other IBM Spectrum Scale storage tiers. This open beta of IBM Spectrum Scale will support on-premise object storage, including Cleversafe, as well as off-premise cloud object storage. You are invited to test the beta version of this technology that is planned to be generally available with an upcoming release of IBM Spectrum Scale. We encourage you to watch the demonstration video of this capability **here**.

IBM Spectrum Scale transparent cloud tiering Open Beta Drop 2 code is now available [1 March 2016]. This update includes:

- Namespace backup with periodic checkpoint copies in the cloud
- Ability to manage deletion of older file versions on the cloud
- Updated documentation
- Various defect fixes

Use IBM Spectrum Scale Forum or mail to ibmmcstr@us.ibm.com to ask questions and to give your feedback.

Date	Type	Description	Download
1 Mar 2016	Open Beta	IBM Spectrum Scale transparent cloud tiering	Download

https://www.ibm.com/developerworks/servicemanagement/tc/gpfs/evaluate.html

IBM Elastic Storage Server (ESS)

Integrated scale out data management for file and object data

Optimal building block for high-performance, scalable, reliable enterprise storage

- Faster data access with choice to scale-up or out
- Easy to deploy clusters with unified system GUI
- Simplified storage administration with IBM Spectrum Control integration

One solution for all your data needs

- Single repository of data with unified file and object support
- Anywhere access with multi-protocol support: NFS 4.0, SMB, OpenStack Swift, Cinder, and Manila
- Ideal for Big Data Analytics with full Hadoop transparency with 4.2

Ready for business critical data

- Disaster recovery with synchronous or asynchronous replication
- Ensure reliability and fast rebuild times using Spectrum Scale RAID's dispersed data and erasure code

Advantages of Spectrum Scale RAID

Use of standard and inexpensive disk drives

- Erasure Code software implemented in Spectrum Scale

Faster rebuild times

- More disks are involved during rebuild
- Approx. 3.5 times faster than RAID-5

Minimal impact of rebuild on system performance

- Rebuild is done by many disks
- Rebuilds can be deferred with sufficient protection

Better fault tolerance

- End to end checksum
- Much higher mean-time-to-data-loss (MTTDL)
  - 8+2P: ~ 200 Years
  - 8+3P: ~ 200 Million Years

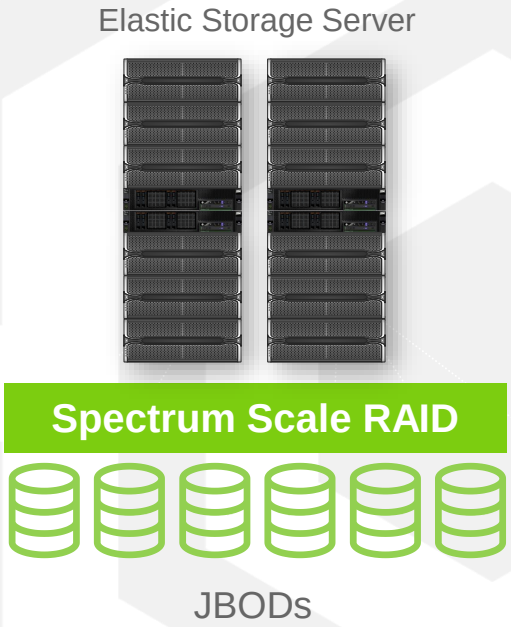

IBM Spectrum Control

Overview of all File Storage Systems

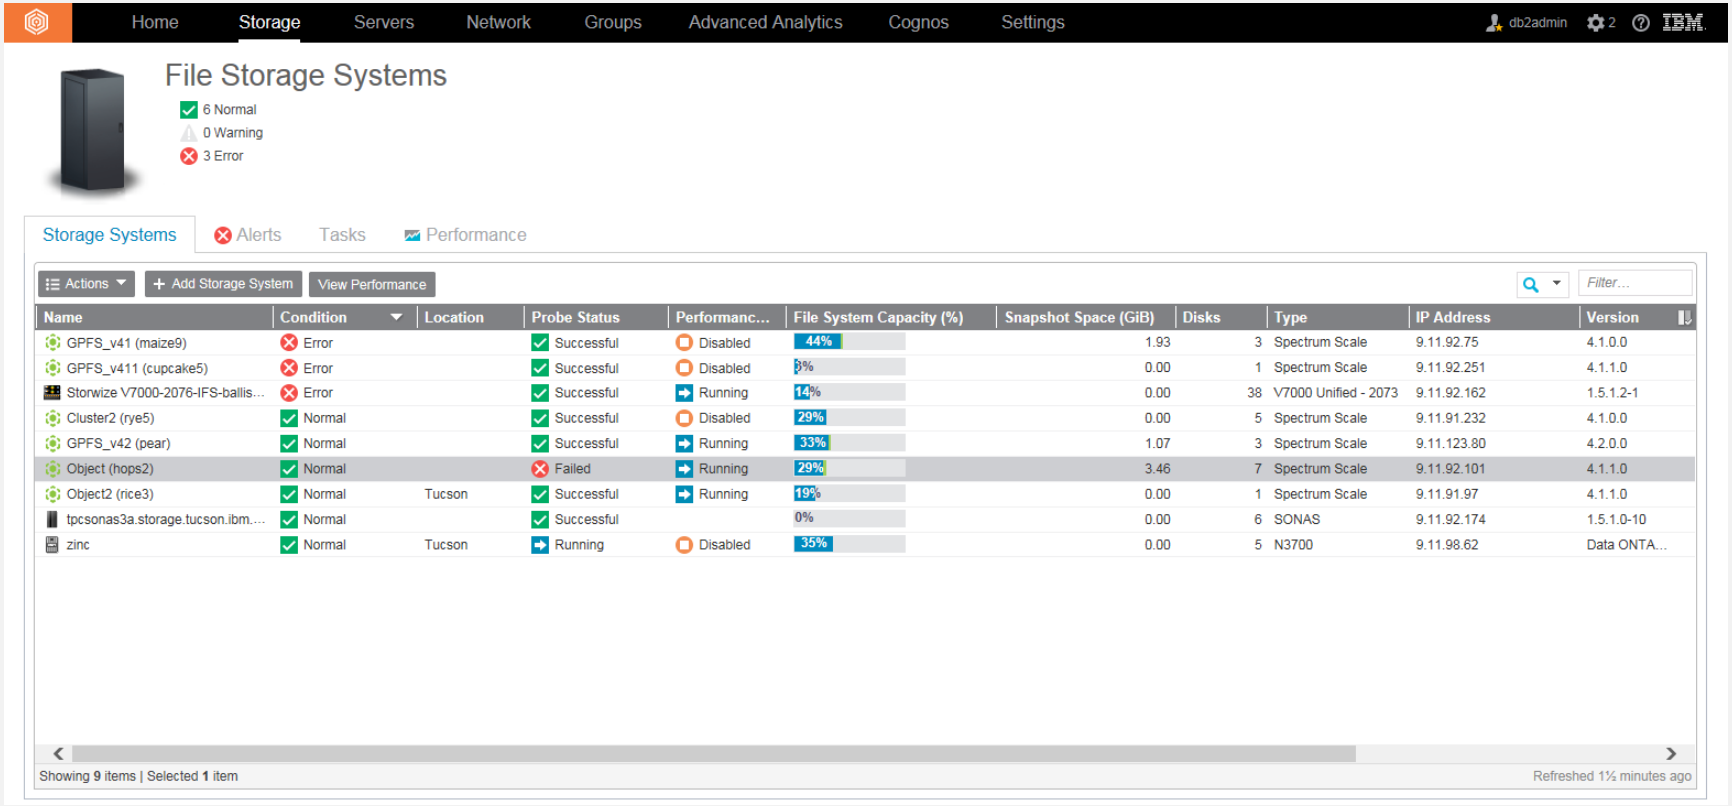

SAN-attached storage troubleshooting

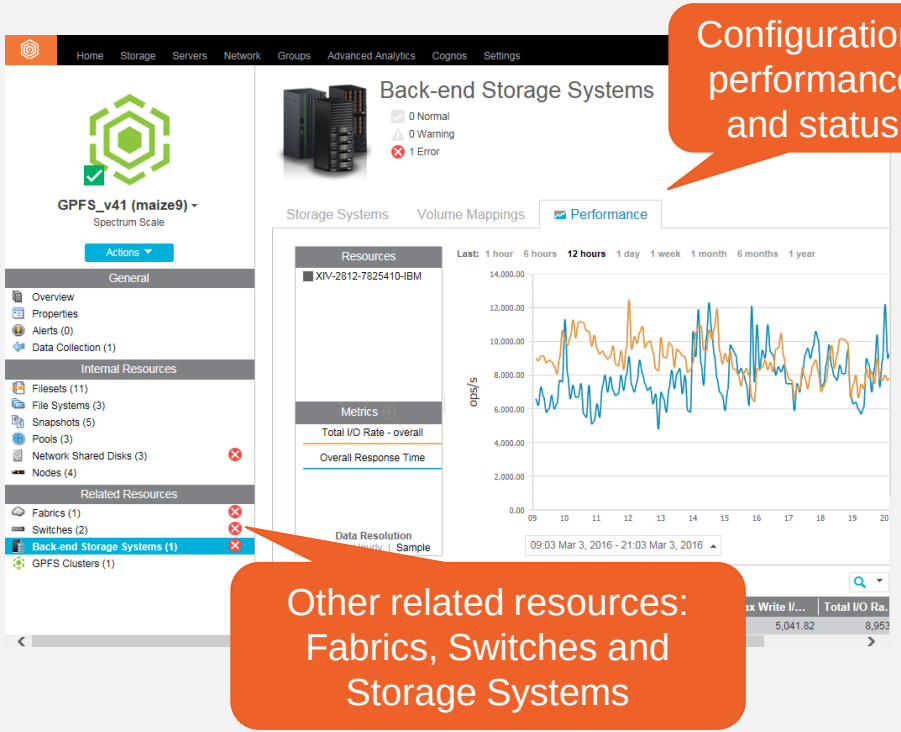

Clusters may use NSDs built off of SAN attached storage such as FlashSystems or other block storage systems. Latency within the Spectrum Scale file system may be due to issues within the storage or the fabric connections.

Today - Troubleshooting an issue may involve hunting through the Spectrum Scale GUI, Brocade Network Advisor (or Cisco Fabric Manager), and storage system element managers.

With Spectrum Control - A storage team can start from a node or file system and trace performance through the fabric to the SAN attached storage.

Multi-cluster environments

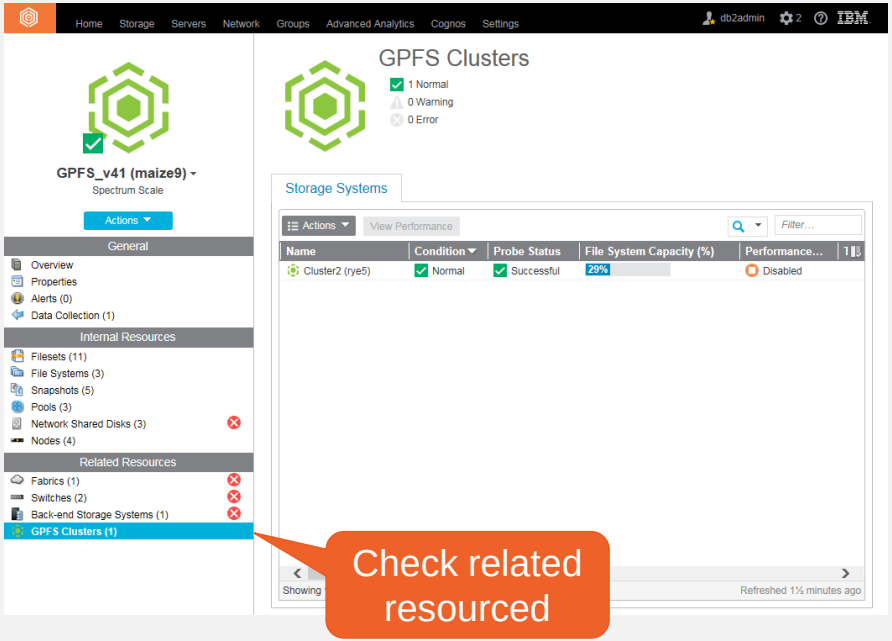

Many Spectrum Scale customers that we talked to have more than a sincle cluster, typically some of which are client only and storage only. You would have a better idea of how normal this is.

Today - If a storage team wants a complete view of their Spectrum Scale environment, they have a few choices:

- Jump between multiple Spectrum Scale GUIs
- Write their own home grown tools
- Purchase a product that can monitor multiple clusters.

With Spectrum Control - Storage teams can see their entire Spectrum Scale environment at a glance, easily comparing capacity and workloads across multiple clusters.

Multi-cluster environments II

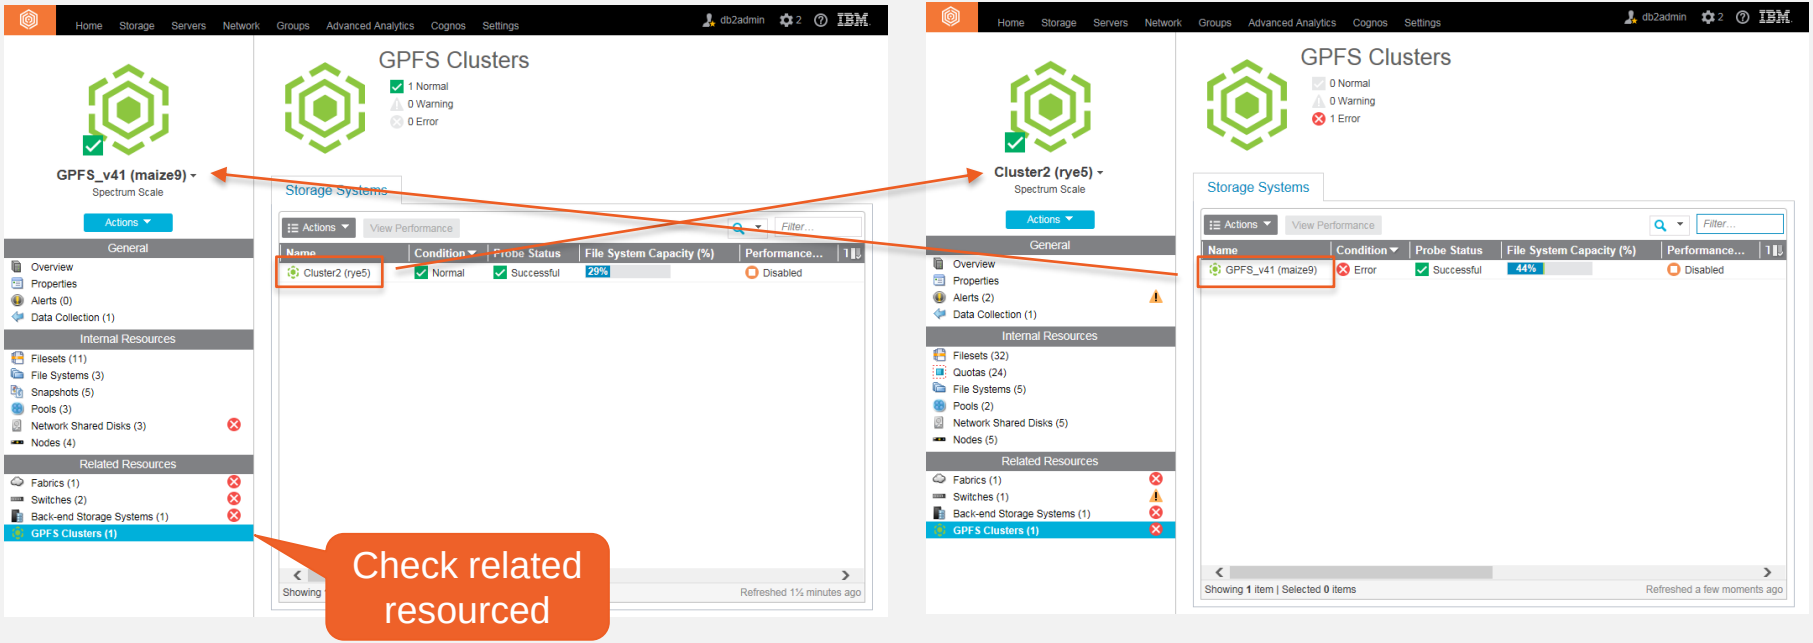

Multi-cluster environments: Cross-Cluster mounts

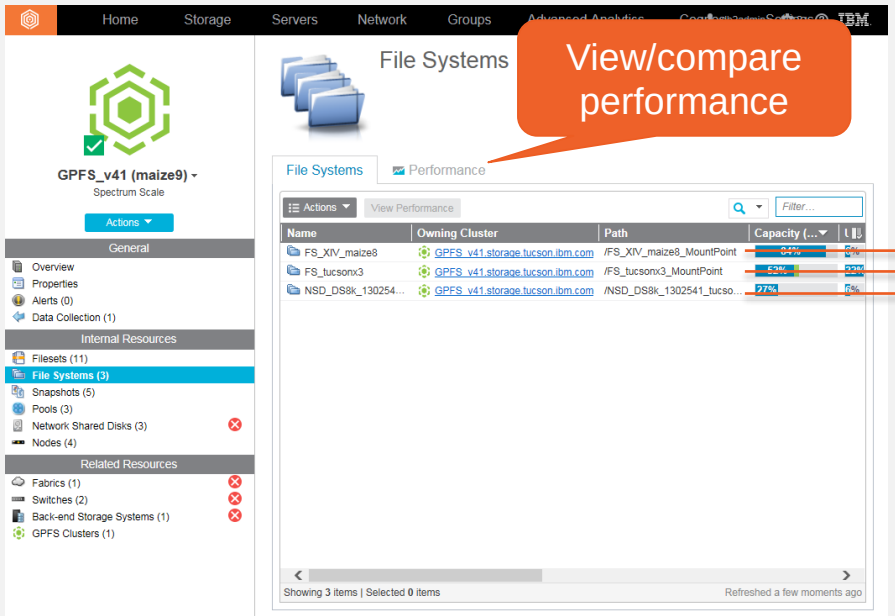

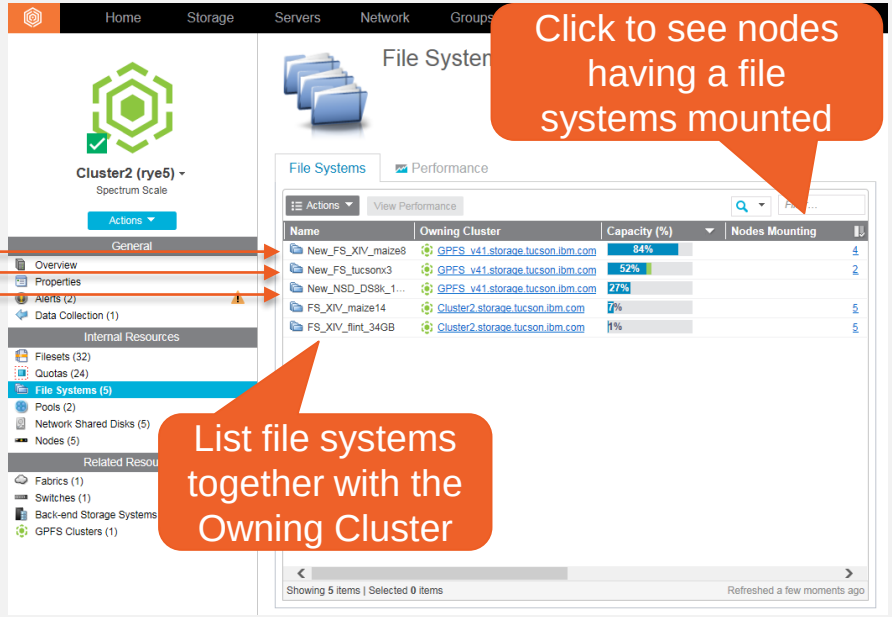

IBM Spectrum Control

Application oriented monitoring

An even in a storage environment an application consists of many components, in this context these are filesets.

Today - Troubleshooting and reporting is difficult because the components like file sets, shares, network and backend storage resources are not available in a single dashboard

With Spectrum Control - A Spectrum Scale admin defines which resources belong to an application. From a list of applications (or departments) the admin can open a panel that shows all the information in a single place.

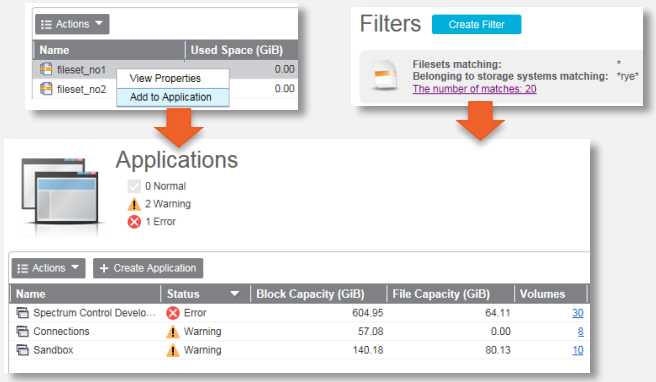

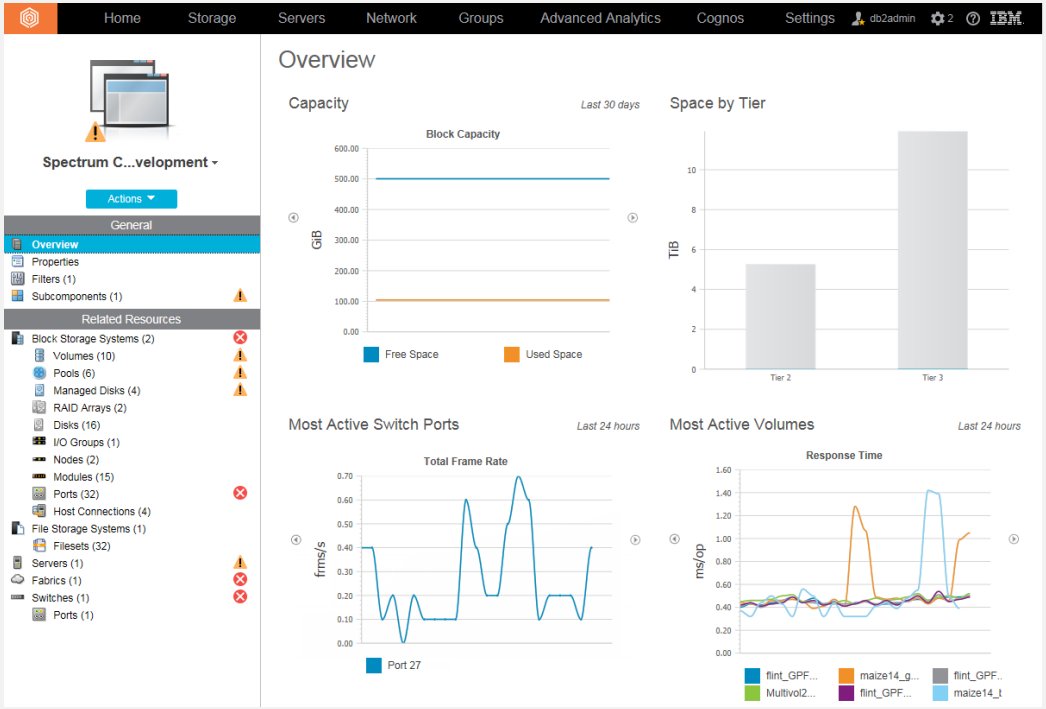

Snapshot backup of Applications

Spectrum Control Advanced edition includes Spectrum Protect Snapshot (aka FlashCopy Manager)

Today – Automation of snapshot bases backups with offload to tape has to be individually developed and maintained

With Spectrum Control - A Spectrum Protect Snapshot can be used to integrate application consistent backups, offloading the backup to tape, and maintain a backup history that's available in Spectrum Control.

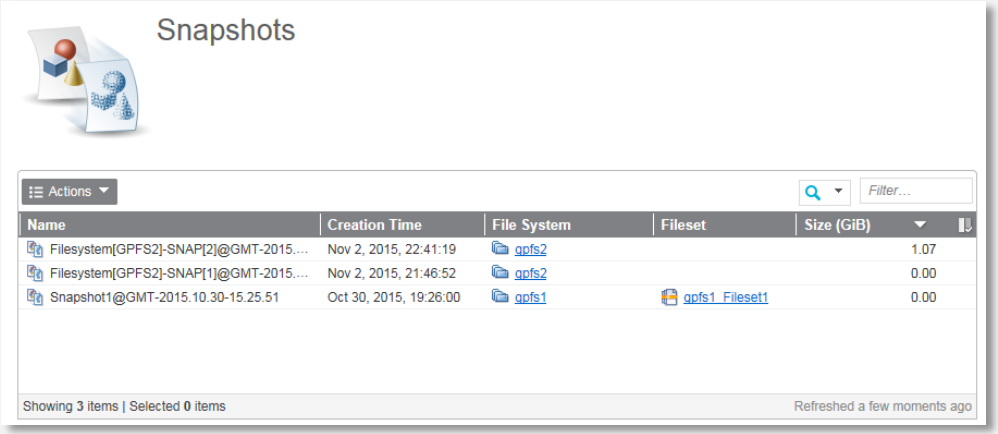

Notes:

- Minimum Spectrum Protect Version 4.1.1.2 (1Q15)
- Minimum Spectrum Protect Version 4.1.4 (1Q16) with offload backup to Spectrum Protect (aka TSM) link

- Miniumum Spectrum Scale Version: 4.1.0.5

Thank you.

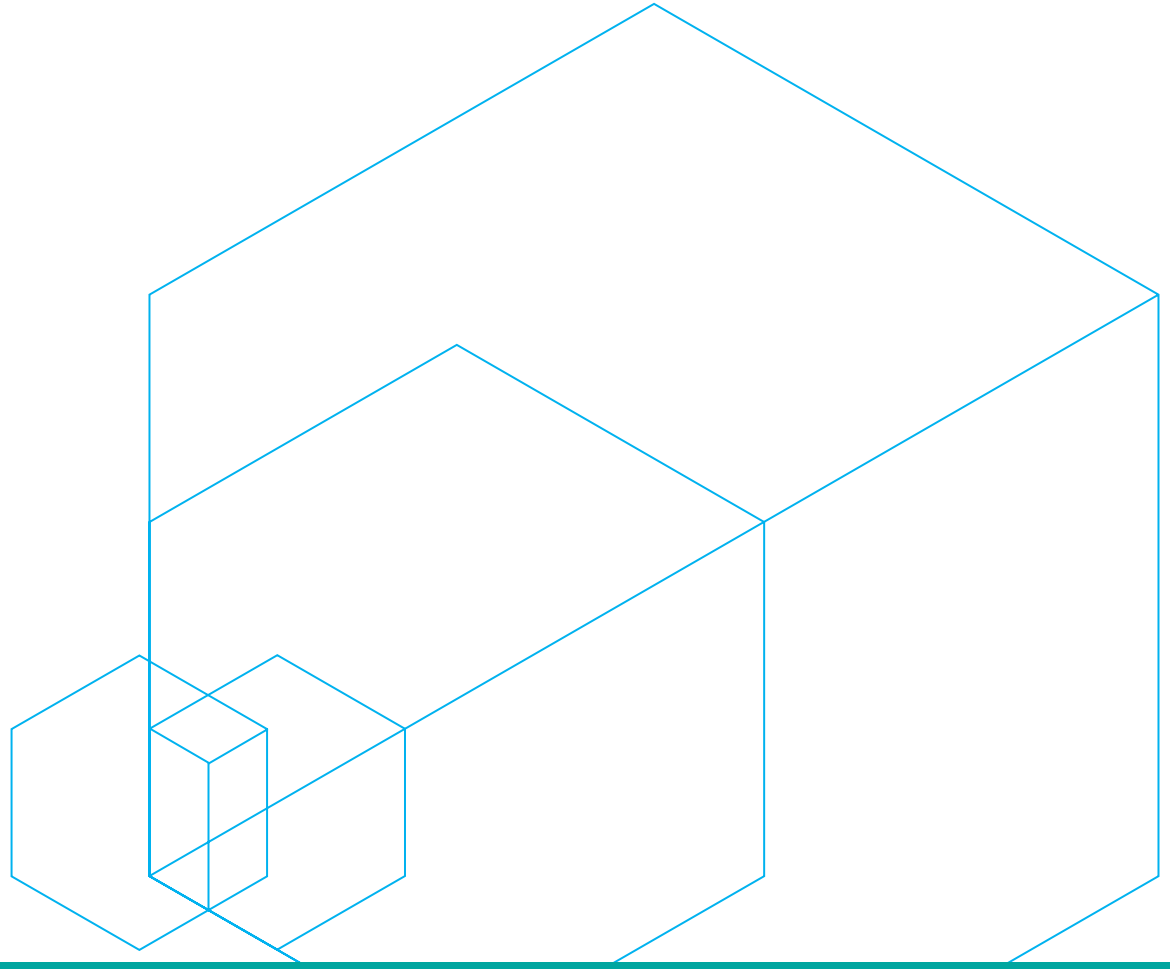

ibm.com/systems

Legal notices

Product data has been reviewed for accuracy as of the date of initial publication. Product data is subject to change without notice. This document could include technical inaccuracies or typographical errors. IBM may make improvements and/or changes in the product(s) and/or program(s) described herein at any time without notice. Any statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business. Any reference to an IBM Program Product in this document is not intended to state or imply that only that program product may be used. Any functionally equivalent program, that does not infringe IBM's intellectually property rights, may be used instead.

THE INFORMATION PROVIDED IN THIS DOCUMENT IS DISTRIBUTED "AS IS" WITHOUT ANY WARRANTY, EITHER OR IMPLIED. IBM LY DISCLAIMS ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NONINFRINGEMENT. IBM shall have no responsibility to update this information. IBM products are warranted, if at all, according to the terms and conditions of the agreements (e.g., IBM Customer Agreement, Statement of Limited Warranty, International Program License Agreement, etc.) under which they are provided. Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products in connection with this publication and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. IBM makes no representations or warranties, ed or implied, regarding non-IBM products and services.

The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents or copyrights. Inquiries regarding patent or copyright licenses should be made, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 1 0504- 785
U.S.A.

Information and trademarks

Notes:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

This presentation and the claims outlined in it were reviewed for compliance with US law. Adaptations of these claims for use in other geographies must be reviewed by the local country counsel for compliance with local laws.

Special notices

This document was developed for IBM offerings in the United States as of the date of publication. IBM may not make these offerings available in other countries, and the information is subject to change without notice. Consult your local IBM business contact for information on the IBM offerings available in your area.

Information in this document concerning non-IBM products was obtained from the suppliers of these products or other public sources. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. Send license inquires, in writing, to IBM Director of Licensing, IBM Corporation, New Castle Drive, Armonk, NY 10504-1785 USA.

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

The information contained in this document has not been submitted to any formal IBM test and is provided "AS IS" with no warranties or guarantees either expressed or implied.

All examples cited or described in this document are presented as illustrations of the manner in which some IBM products can be used and the results that may be achieved. Actual environmental costs and performance characteristics will vary depending on individual client configurations and conditions.

IBM Global Financing offerings are provided through IBM Credit Corporation in the United States and other IBM subsidiaries and divisions worldwide to qualified commercial and government clients. Rates are based on a client's credit rating, financing terms, offering type, equipment type and options, and may vary by country. Other restrictions may apply. Rates and offerings are subject to change, extension or withdrawal without notice.

IBM is not responsible for printing errors in this document that result in pricing or information inaccuracies.

All prices shown are IBM's United States suggested list prices and are subject to change without notice; reseller prices may vary.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

Any performance data contained in this document was determined in a controlled environment. Actual results may vary significantly and are dependent on many factors including system hardware configuration and software design and configuration. Some measurements quoted in this document may have been made on development-level systems. There is no guarantee these measurements will be the same on generally-available systems. Some measurements quoted in this document may have been estimated through extrapolation. Users of this document should verify the applicable data for their specific environment.

DCP03213USEN-06