



IBM Spectrum Storage

Accelerating and simplifying backup

IBM Spectrum Protect™ with
IBM Spectrum Scale™

Agenda

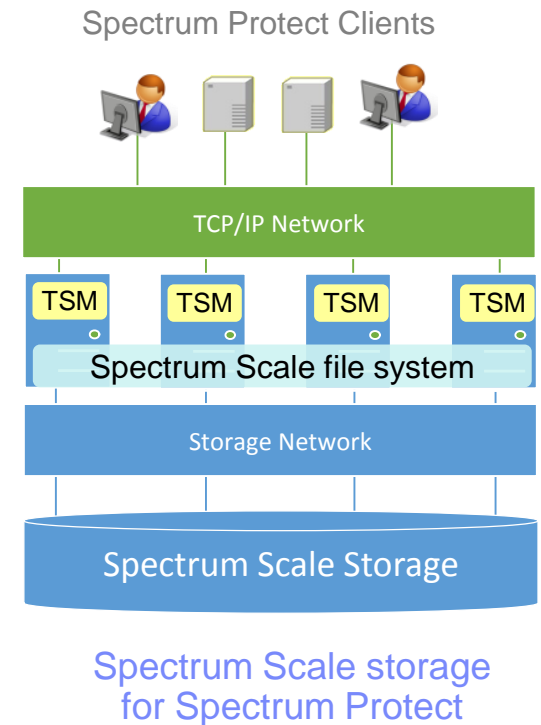
► Overview

Configuration options

Advantages

Spectrum Protect on Spectrum Scale - Overview

- **Multiple** Spectrum Protect (TSM) instances store DB and storage pools in a Spectrum Scale file system (GPFS)
 - Spectrum Scale provides global name space for all Spectrum Protect instances
 - Instances share all file system resources
- Spectrum Protect instances run on cluster nodes accessing the file system and disk directly
- Spectrum Scale file systems balances the workload and capacity for all TSM instances on disk
- Provides standardized, scalable and easy to use storage infrastructure for the multiple instances

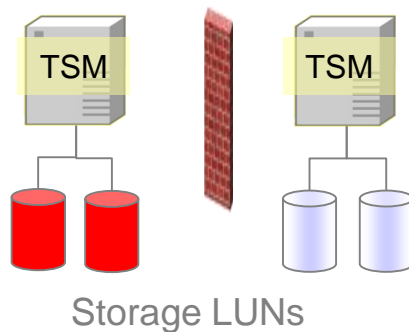


Why ?

Without Spectrum Scale

- Each backup server has its own isolated file system
- Each backup server is tightly coupled to LUNs
- Storage islands appear with underutilized capacity
- Capacity and performance management is challenging
- Scaling and performance may impact apps and users


Spectrum Protect Instance



With Spectrum Scale

- Scale capacity seamlessly and transparently to apps or users under the shared file system global namespace
- File system replication is included
- Build your infrastructure using commodity storage, i.e. no vendor lock in.
- Central administration of all storage

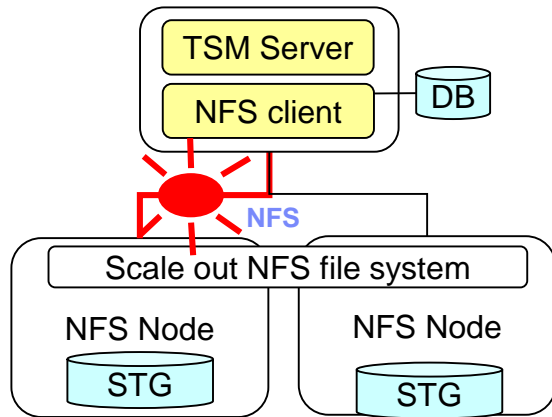
Spectrum Protect Instance



Spectrum Scale does not have the NFS bottleneck


When Spectrum Protect performs I/O to NFS server it is bound to a **single** node

- Because a NFS session is point-to-point
- I/O Performance on single volume is limited by one NFS server



With Spectrum Scale, Spectrum Protect performs I/O to **all** disk in parallel

- I/O performance on a single volume is done to all disk in parallel



Key Advantages for Spectrum Protect with Spectrum Scale

Better storage utilization – multiple TSM server share the same storage


Better operational efficiency with one storage for all TSM server

Seamless scalability of storage capacity and performance

Flexible disaster protection – asynchronous, synchronous or Native RAID

High availability in clustered file system

Cost efficiency by utilizing standard infrastructure components



Agenda

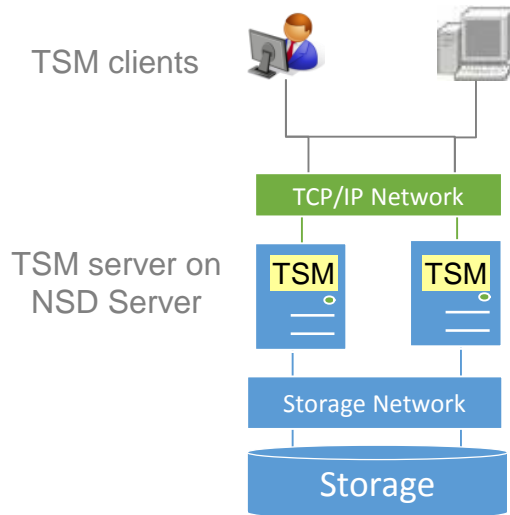
Overview

▶ **Configuration options**

Advantages

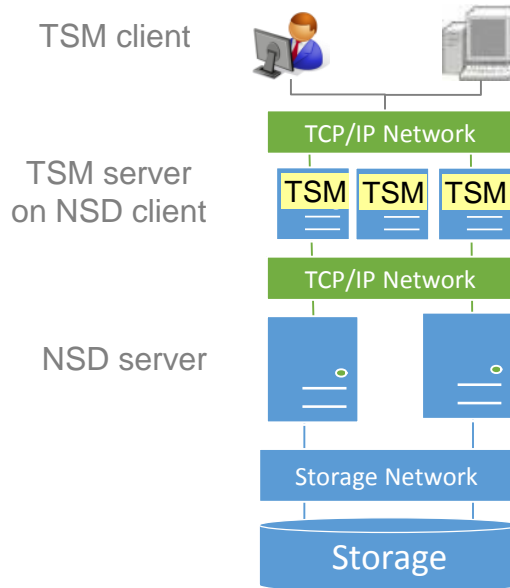
Deployment options

On NSD server



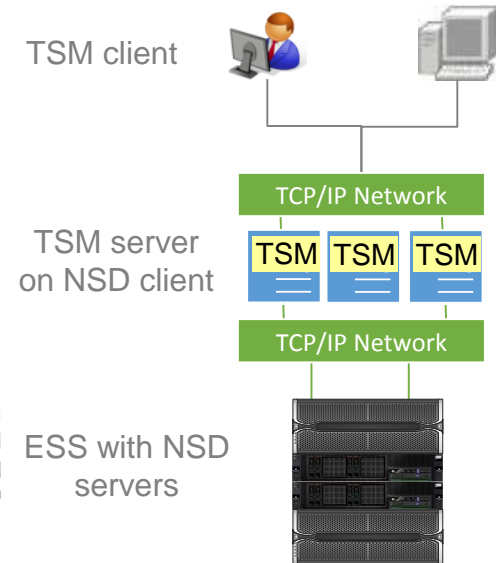
- TSM server on NSD servers
- Direct SAN storage access
- Requires less infrastructure
- If Spectrum Scale is used just by TSM only

On NSD client



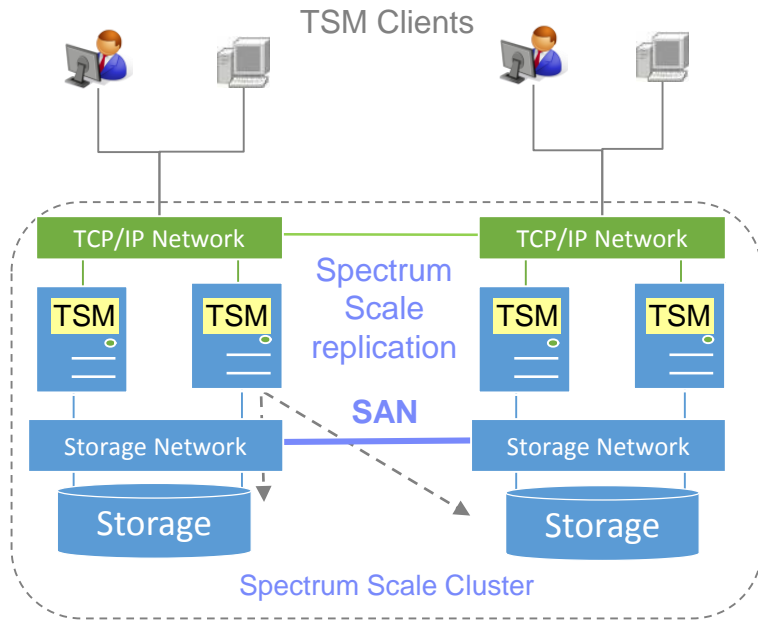
- TSM runs on client connected to Scale / ESS NSD server
- SAN or LAN access from TSM server to Scale/ ESS server
- Requires more infrastructure
- If Spectrum Scale is used by multiple applications

With Elastic Storage Server



Replication options for Spectrum Scale and Spectrum Protect

Spectrum Scale synchronous replication

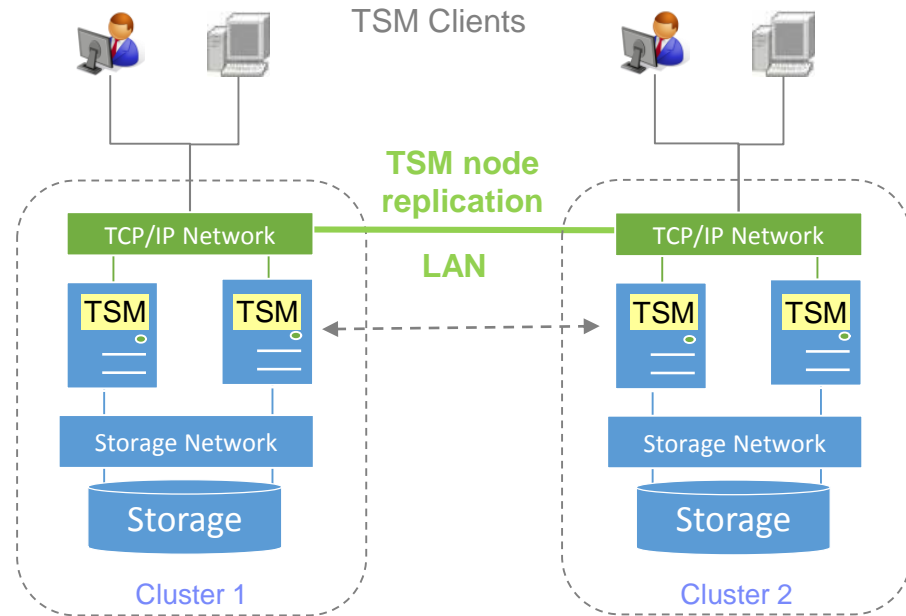


Synchronous replication via LAN or SAN

One cluster with active instances in both sites

Made for high availability and workload balancing

Spectrum Protect node replication



Asynchronous replication by TSM via LAN

Two independent clusters with TSM instances

Made for disaster recovery

Configuration guidance

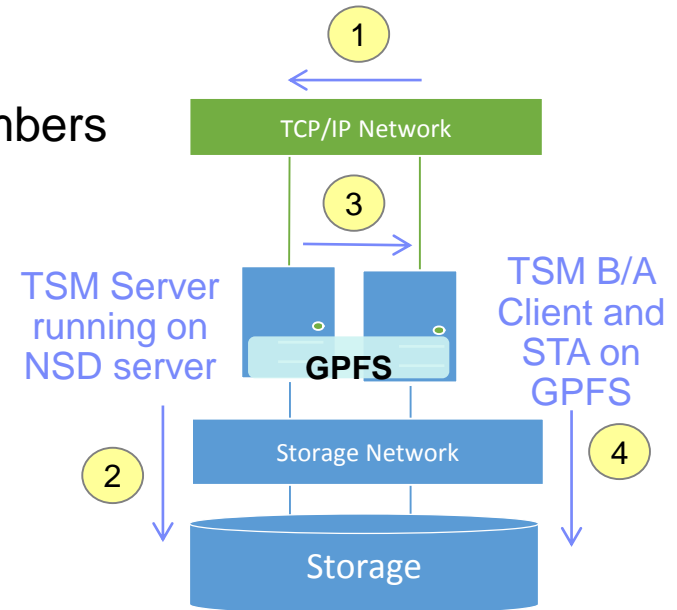
- Separate network adapters for incoming client workload and outgoing TSM server workload
- Separate storage for TSM DB, Logs and storage pools
 - Different file systems for DB, Logs and storage pools
 - Block size for DB and active Log: 64K
 - Block size for storage pools: 2 – 8 MB
- Consider creating filesets for TSM instances within file systems
 - Each instance has a subdirectory (file set) in each file system
- With deduplication use Flash or SSD storage for TSM DB, Logs
- With FILE device class configure data format = native when possible
 - Fixed block I/O provides better storage pool performance
- Set parameter DIRECTIO=NO and DIOENABLED=NO in TSM servers

LAN free backup to Spectrum Scale

- TSM server and TSM client are GPFS cluster members
 - Both can access GPFS file systems

■ How it works

- 1 TSM client request volume from TSM server in GPFS file system
- 2,3 TSM server mounts volume in file system and gives handle to TSM client
- 4 TSM client writes directly to volume in GPFS file system



- LAN free backup not possible with ESS – but server free
- LAN free backup not possible with container pools (inline dedup)
 - Deduplication only possible with legacy dedup on file volumes

Agenda

Overview

Configuration options


▶ **Advantages**

Better operational efficiency

- Single point of management
 - Storage resources for all Spectrum Protect instances are centrally managed

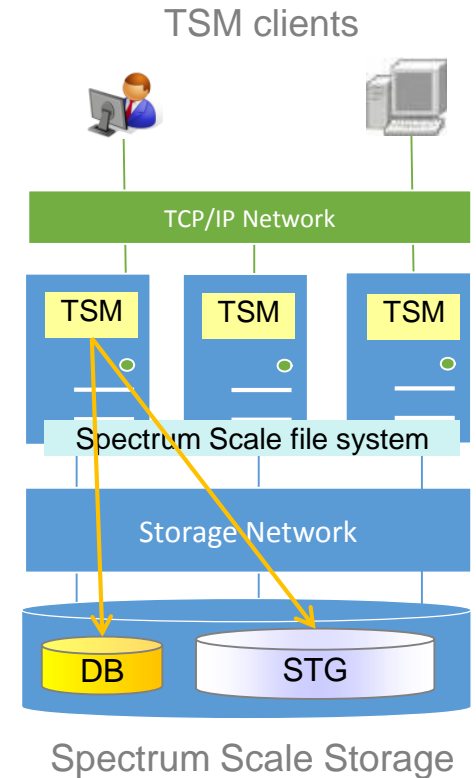
- Always-on
 - Spectrum Scale supports highly resilient architectures
 - Adding storage capacity is transparent to Spectrum Protect instances

- Transparent resource balancing
 - Instances requiring more resources can benefit from instances requiring less



Lower cost

- No extra storage resources for DB required
 - All Spectrum Protect instances store DB and storage pools in one Spectrum Scale cluster
- Better storage utilization
 - Storage capacity is provided to all Spectrum Protect instances in global namespace
- No extra infrastructure required
 - Multiple Spectrum Protect instances can run on one Spectrum Scale node
- Use and reuse standard Infrastructure




Scaling parallel performance

- **Direct access to Spectrum Scale storage**
 - Spectrum Protect stores directly in Spectrum Scale storage, no overhead for conversion of files to blocks (like with NFS)

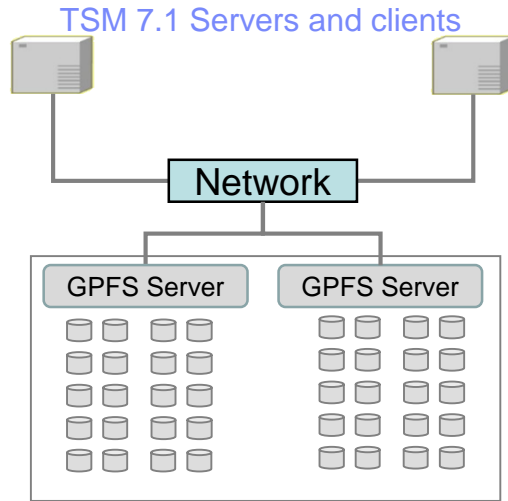
- **Parallel usage of all available disk**
 - Spectrum Protect writes in parallel to all disk
 - No dedicated connection or session between TSM server and GPFS storage (no NFS bottleneck)

- **Fine tuning for DB and storage pools file systems**
 - Variable file system block size allow to adopt specific workloads






How fast can a backup with Spectrum Protect on Spectrum Scale be?

Test setup



GPFS Native RAID system with NL-SAS disk



Positioning Spectrum Protect with Spectrum Scale

Multiple Spectrum Protect servers on Spectrum Scale can benefit from

- Global name space provided within high available cluster
- High availability with synchronous replication across sites
- Better cost efficiency
- Better performance
- Ease of use



Thank You

Additional References

Storageneers



Scale out backup with TSM and GSS: Performance test results

Elastic Storage with GPFS Native RAID performance test results with Tivoli Storage Manager over 40 GBit Ethernet

Is a scale out NAS system the same as a scale out file system?

Software Defined Data Protection – The next generation of backup

The Register[®]

Mirror, mirror on the wall, who has the best TSM backend of all?

Big Blue stuffs data into backup at GIGABYTES/sec

Links

- **Spectrum Scale Home Page**
<http://www-03.ibm.com/systems/storage/spectrum/scale/overview.html>
- **Spectrum Scale Knowledge Center:**
http://www-01.ibm.com/support/knowledgecenter/STXKQY/ibmspectrumscale_welcome.html
- **IBM Spectrum Scale Wiki**
[http://www.ibm.com/developerworks/wikis/display/hpccentral/General+Parallel+File+System+\(GPFS\)](http://www.ibm.com/developerworks/wikis/display/hpccentral/General+Parallel+File+System+(GPFS))
- **Spectrum Scale Redbook:**
<http://www.redbooks.ibm.com/abstracts/sg248254.html?Open>
- **Presentation: Spectrum Protect with Spectrum Scale – Overview and Positioning (IBM Internal)**
<https://w3-connections.ibm.com/files/app#/file/13c05ddc-b5f8-454b-ad25-7c20635c5de4>
- **Presentation: Spectrum Protect with Spectrum Scale – Guidance for solution design and sizing (IBM Internal)**
<https://w3-connections.ibm.com/files/app#/file/63eca9a4-2469-4d36-8537-e41ff6278717>
- **Spectrum Scale wiki: configuring Spectrum Scale for reliability:**
<https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/General%20Parallel%20File%20System%20%28GPFS%29/page/Configuring%20GPFS%20for%20Reliability>
- **ESS Sizing Guide (IBM Internal)**
<https://w3-connections.ibm.com/wikis/home?lang=en-us#!/wiki/GPFS%20Sales%20Wiki/page/ESS%20Capacity%20and%20Performance>
- **TSM Blueprints**
<https://ibm.biz/TivoliStorageManagerBlueprints>
- **TXNGroupMax and TXNBytelimit**
 - http://www-01.ibm.com/support/knowledgecenter/SSGSG7_7.1.1/com.ibm.itsm.srv.ref.doc/r_opt_server_txngroupmax.html
 - http://www-01.ibm.com/support/knowledgecenter/SSGSG7_7.1.1/com.ibm.itsm.client.doc/r_opt_txnbytelimit.html
- **The Modular Data Retention concepts**
https://ibm.biz/IBM_MDRC
- **TSM Architecture and Design Service**
https://escc.mainz.de.ibm.com/e_dir/e4uweb.nsf/site.xsp?off=EUSTO120
- **TSM Server option DIRECTIO**
<http://www-01.ibm.com/support/docview.wss?uid=swg1IC73238>
- **TSM Server performance tuning guide:**
ftp://public.dhe.ibm.com/software/products/TSM/current/b_perf_tuning_guide.pdf
- **Edision Group whitepaper:**
<http://www.theedison.com/index.php/library-edison/221-ibm-spectrum-scale-outperforms-emc-isilon-as-a-backup-target>

Disclaimer

- This information is provided on an "AS IS" basis without warranty of any kind, express or implied, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Some jurisdictions do not allow disclaimers of express or implied warranties in certain transactions; therefore, this statement may not apply to you.
- This information is provided for information purposes only as a high level overview of possible future products. PRODUCT SPECIFICATIONS, ANNOUNCE DATES, AND OTHER INFORMATION CONTAINED HEREIN ARE SUBJECT TO CHANGE AND WITHDRAWAL WITHOUT NOTICE.
- USE OF THIS DOCUMENT IS LIMITED TO SELECT IBM PERSONNEL THIS DOCUMENT SHOULD NOT BE GIVEN TO A CUSTOMER EITHER IN HARDCOPY OR ELECTRONIC FORMAT.

Important notes:

- IBM reserves the right to change product specifications and offerings at any time without notice. This publication could include technical inaccuracies or typographical errors. References herein to IBM products and services do not imply that IBM intends to make them available in all countries.
- IBM makes no warranties, express or implied, regarding non-IBM products and services, including but not limited to Year 2000 readiness and any implied warranties of merchantability and fitness for a particular purpose. IBM makes no representations or warranties with respect to non-IBM products. Warranty, service and support for non-IBM products is provided directly to you by the third party, not IBM.
- All part numbers referenced in this publication are product part numbers and not service part numbers. Other part numbers in addition to those listed in this document may be required to support a specific device or function.
- MHz / GHz only measures microprocessor internal clock speed; many factors may affect application performance. When referring to storage capacity, GB stands for one billion bytes; accessible capacity may be less. Maximum internal hard disk drive capacities assume the replacement of any standard hard disk drives and the population of all hard disk drive bays with the largest currently supported drives available from IBM.

IBM Information and Trademarks

- The following terms are trademarks or registered trademarks of the IBM Corporation in the United States and / or other countries: IBM, IBM Spectrum Storage, IBM Spectrum Protect, IBM Spectrum Scale, IBM Spectrum Accelerate, IBM Spectrum Virtualize, IBM Spectrum Control, Tivoli, IBM Elastic Storage
- Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
- UNIX is a registered trademark of The Open Group in the United States and other countries.
- Microsoft Windows is a trademark or registered trademark of Microsoft Corporation.
- Isilon is a registered trade mark of EMC Corporation in the United States and other countries
- Other company, product, and service names may be trademarks or service marks of others.

Disclaimer – Performance and Capacity

Performance is based on measurements and projections using standard 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.

There are many ways to set up a filesystem, the figures shown are for a filesystem implemented with parameters for "maximum performance" (typically 16MB blocksize). If more efficient space utilisation is required for small files (e.g. 512KB blocksize) then the achievable performance will be affected.

Throughput figures (MB/sec) are based on benchmarks using the IOR benchmark and the ESS set up with 8+2P data protection and a combined Data+Metadata storage pool. Where actual results for a model/storage combination are not available, figures are estimates which have been scaled from real benchmarks. IOR tests were performed using Infiniband interconnect and RDMA, figures for other network interconnects may vary significantly.

IOPS figures are "uncached, raw" IOPS tested using the gfsperf benchmark.

The Separate capacity columns assume 7% to 10% of the filesystem capacity is set aside and dedicated for Metadata (MD), this is a typical figure but can vary depending on filesystem requirements (e.g. a very large number of files may require >7%). The figure shown is approximately 2%-3% usable MD with 3-way replication assumed. The combined Data+MD figure provides an indication of maximum filesystem capacity, a combined MD+Data is the default way in which most filesystems work. Separated MD and Data pools provide better performance, total capacity is seen as reduced because the whole of the MD space is assigned explicitly and is seen as separate- this is the recommended way to set up a GPFS filesystem. The ESS default is to use separate MD and Data pools- 7% (GUI) or 10% (ESS scripts) is a typical MD capacity (approx 2% actual, 3-way or 4-way replicated).

These calculations take into account the equivalent of 2 spares worth of disk space per Declustered Array (DA) which is the default and can be changed.

These calculations also take into account GNR overheads (checksums) etc.

These calculations assume all vdisks are of the same RAID type, although mixing is allowed in practice.

This does not take into account any GPFS file system level replication which will reduce the net usable space appropriately.