

IBM Spectrum Scale Protocols – Overview –

Tomer Perry, Ulf Troppens

With contributions from Ingo Meents, Michael Diederich, Sandeep Patil, Monica Lemay

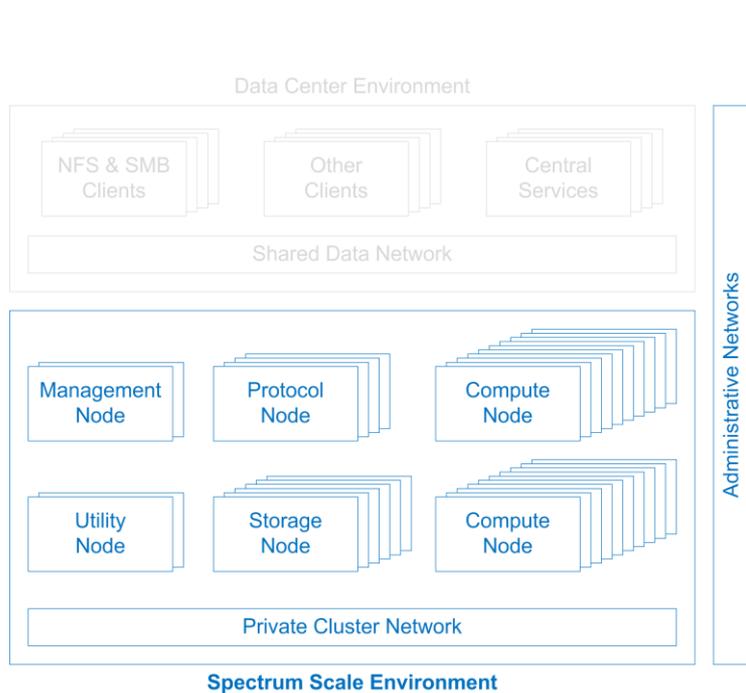


Prelude

Question to audience:

- Who uses Spectrum Scale CES?
- Who uses Spectrum Scale CES without Compute Nodes?
- Who uses CNFS?
- Who uses homegrown Samba/CTDB?
- Who uses something different on top of Spectrum Scale for remote access?

Spectrum Scale Environment



- The Shared Data Network provides remote access to the Spectrum Scale environment.
- The Private Cluster Network connects all components of the Spectrum Scale environment.

Compute Nodes (NSD Clients)

- Run applications to access and analyze data stored in one or more Spectrum Scale filesystems
- Most nodes of a Spectrum Scale environment are Compute Nodes.

Storage Nodes (NSD Server)

- Provide the storage capacity for the Spectrum Scale filesystems

Protocol Nodes (Cluster Export Services)

- Provide remote access to Spectrum Scale filesystems using NFS, SMB and Object

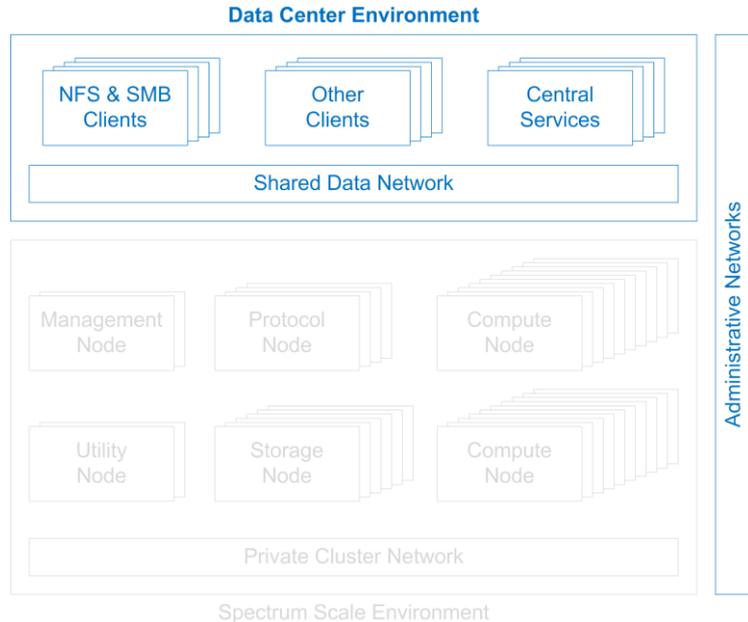
Utility Nodes (Data Management Nodes)

- Dedicated nodes for heavy-weight data management tasks such as backup, tiering, hybrid cloud workflows.

Management Nodes

- Provides administration services (e.g., Spectrum Scale GUI, Zimon Collector, Compute Cluster Login Node, Compute Cluster Management Node).

Data Center Environment



- The Shared Data Network provides remote access to the Spectrum Scale environment.
- The Private Cluster Network connects all components of the Spectrum Scale environment.

NFS&SMB Clients

- Users and applications accessing data stored on a Spectrum Scale filesystem using NFS and/or SMB

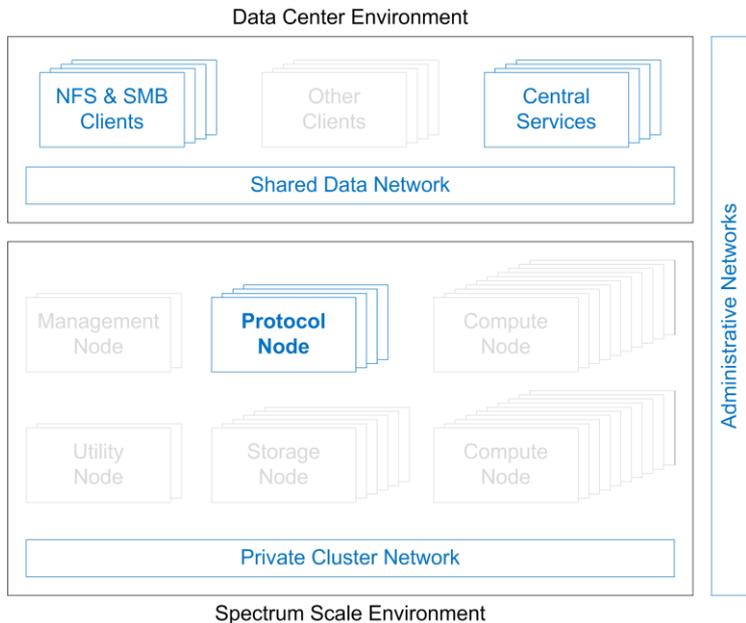
Other Clients

- User and applications accessing data stored on a Spectrum Scale filesystem using other services that are outside the scope of this presentation (e.g., Swift/S3, HDFS, Aspera, rsync, scp, etc.)
- Administrative workstations (e.g. GUI client, REST API client, SSH client, etc.)

Central Services

- External infrastructure services required for the whole solution such as Authentication and ID mapping (e.g. AD, LDAP), Time synchronization (e.g., NTP), Name resolution (e.g., DNS), etc.

CES Protocol Nodes



CES Protocol Nodes for NFS and SMB interact with

- NFS & SMB Clients
- Central Services (e.g. authentication, DNS, NTP)
- Shared Data Center Network
- High-Speed Private Network

This charts deck provides guidance on how to use Spectrum Scale Protocol Nodes in context of a larger Spectrum Scale environment

- When to choose Protocol Nodes?
- Software components
- External Dependencies
- Network Requirements
- NFS and SMB Considerations
- Miscellaneous

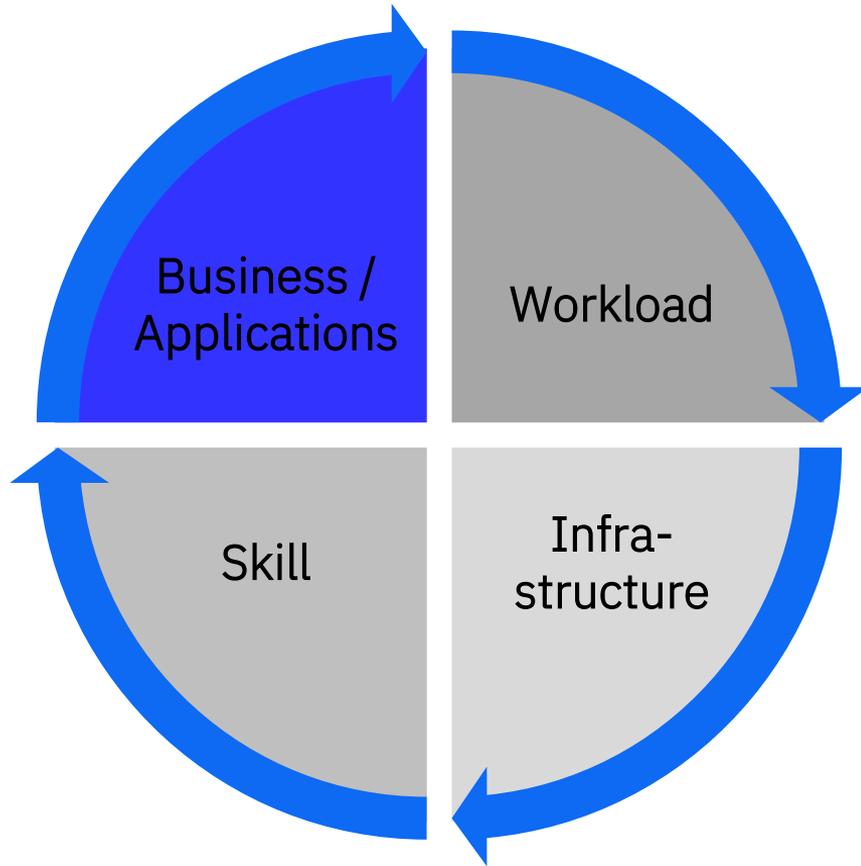
- The Shared Data Network provides remote access to the Spectrum Scale environment.
- The Private Cluster Network connects all components of the Spectrum Scale environment.

IBM Software-Defined Storage solutions for unstructured data workloads

IBM Offering	IBM Spectrum Scale & ESS			IBM Spectrum NAS			IBM Cloud Object Storage	
Workloads	Compute Clusters	Big Data analytics, ML/DL	High speed backup & restore	File serving for Virtual Machines	Home Directories / General purpose file serving	NAS for Microsoft Apps	Archive, Content Repository & distribution	Cloud & Native Object Applications
Key capabilities	High Performance			Simple with Native File Protocols			Multi-site and Native Object Storage	

- Spectrum Scale is targeting on HPC, technical computing, analytics and machine learning/deep learning.
- IBM added Cluster Export Services (CES) to Spectrum Scale to support workflows in those fields that have steps that require SMB, NFS or object access.
- Spectrum Scale works very well for selected SMB and NFS workloads.
- Choose Spectrum Scale with Compute Cluster for High Performance workloads (e.g. Genomics Blueprint) <http://www.redbooks.ibm.com/abstracts/redp5479.html?Open>
- Spectrum Scale Protocol Nodes provide NFS and SMB, but Spectrum Scale it is not an Enterprise NAS solution.

Choosing the right solution



- The **business requirements** determine the required **applications**.
- The **applications determine** the generated **workload**.
- The **workload determines** the required **infrastructure**.
- The **infrastructure determines** the required **skills**.
- The available **infrastructure and skills determine** the capability to support the **business**.

Contrasting file-based workloads

		Parallel File System	Network Attaches Storage (NAS)
Workload	Applications	Broad range of scientific applications, big data and analytics, ML/DL, parallel applications	Broad range of office applications, roaming profiles, etc.
	Scalable Performance	High (large data sets, fast metadata operations, high throughput, low latency)	Medium/Low (average performance and scaling needs)
	Consistency	Strict (Node see updates from remote nodes immediately)	Eventual (Node may see updates from remote nodes after a delay)
Infrastructure / Features	Access to clients	Controlled (Limited number of privileged users)	Wild west (End user have root access to laptops, etc.)
	Client OS Interoperability	Limited (number of operating systems, number of versions, number of architectures)	Flexible (Broad range of different OS versions including very old OS versions and architectures)
	Predominant Client OS	Linux	Linux, Windows, macOS
	Protocol	Proprietary (e.g., Spectrum Scale NSD)	Standard (NFS, SMB)
	Number of clients	Thousands (<16k)	Tens of thousands
	Network	Private Cluster Network	Shared Data Center Network
Skills	Deployment Model	Software Defined Infrastructure	Hardware Appliance
	Client Software	Additional software package for access to parallel filesystem	NFS and or SMB are included in the operating system
	Admin Skills	System administrators (Deep skills in Linux, networking, system software, etc.)	Storage administrators (Mostly management of storage appliances)

Contrasting NFS and SMB workloads

		Spectrum Scale Protocols	Enterprise NAS Appliance
Workload	Use Case	Enable workflows that require occasional access via NFS and or SMB to data stored in Spectrum Scale.	Enable a broad range of NFS and SMB intensive workloads for scientific, business and office applications.
	Applications	Typically qualified and certified by admin team. Limited 3rd-party application support.	Broad ISV ecosystem with many certified and supported applications (e.g., VDI, Microsoft Exchange)
	NFS and SMB Features	Mandatory NFS and SMB protocol features are sufficient to support scientific applications.	Support for many optional NFS and SMB features required to support office and business applications.
	Number of clients	Limited (see NFS and SMB considerations)	Tens of thousands of NFS and or SMB clients
	Actual Workload	To some degree predictable to admin team.	Mostly unpredictable to the admin team.
	Concurrent update	Maintenance window required	Mandatory
Infrastructure / Features	Optimization Goal	High performance (dedicated networks, etc.)	Operational efficiency (thin provisioning, etc.)
	Authentication	Typically a single authentication source (AD or LDAP) is sufficient (managed by same team)	Integration in existing enterprise authentication infrastructure (managed by different team)
	Fast access	High-performance access via Compute Nodes.	No direct access to internal local filesystem.
	Network	Data Center Network + Private Cluster Network	Data Center Network
	Backup	Proprietary (e.g. Spectrum Protect, mmbackup)	Standardized (Built-in NDMP support)
	Quota / Snapshots	Reasonable quota and snapshot policies	Sophisticated quota and snapshot policies
	Antivirus	Limited support	Many certified antivirus solutions
Skills	Integration	Designed, build and integrated by customer	Designed, build and integrated by appliance vendor
	Customizability	Software with many configuration options	Appliance with limited configuration options

When to use Spectrum Scale Protocol Nodes?

Reason that Protocol Nodes are feasible

- Clients with low performance requirements. Access via the Shared Data Network meets the needs.
- Clients that only require occasional access to the data.
- Clients that are connected to a remote network (e.g., client cannot be connected to the Private Cluster Network).
- Eventual consistency is good enough (e.g., transient inconsistencies due to NFS metadata caching are acceptable).

Reasons to avoid NSD clients

- Clients that are not stable (e.g., frequently rebooted, disconnected from Cluster Network)
- Clients that run an operating system version and architecture that is not supported by Spectrum Scale NSD clients.
- Clients where the vendor does not allow to install the Spectrum Scale NSD clients (e.g., genome sequencer, HIL station).
- Clients that have very constrained resources (e.g., sufficient CPU, sufficient memory).
- Clients with limited or no access control (e.g., untrusted people have root access).

When to use NSD Clients?

Reason that NSD Clients are feasible

- Clients that need frequent and high-performance access to the data.
- Clients that can be connected to the Private Cluster Network.
- Clients that run an operating system version and architecture that is supported by Spectrum Scale NSD clients.
- Clients where the vendor allows to install the Spectrum Scale NSD clients (e.g., Compute Node).
- Clients with access control (e.g., only trusted people have root access).

Reasons to avoid Protocol Nodes

- Clients with high-performance requirements. Access via the Shared Data Network is not sufficient.
- Eventual consistency is not acceptable (e.g., transient inconsistencies due to NFS metadata caching are acceptable).

Example use cases and workloads for Spectrum Scale Protocol Nodes

Use Case	Workload
Data ingest (e.g., sensor, camera, genome sequencer, ETL)	Mostly write, file size depends on application
Data export (e.g., move processed data to external location)	Mostly read, file size depends on application
Interactive access to data (e.g., read results, visualization)	Mostly read, file size depends on application
Low performance data analysis (e.g., batch processing, grid computing)	Read and write, depends on application

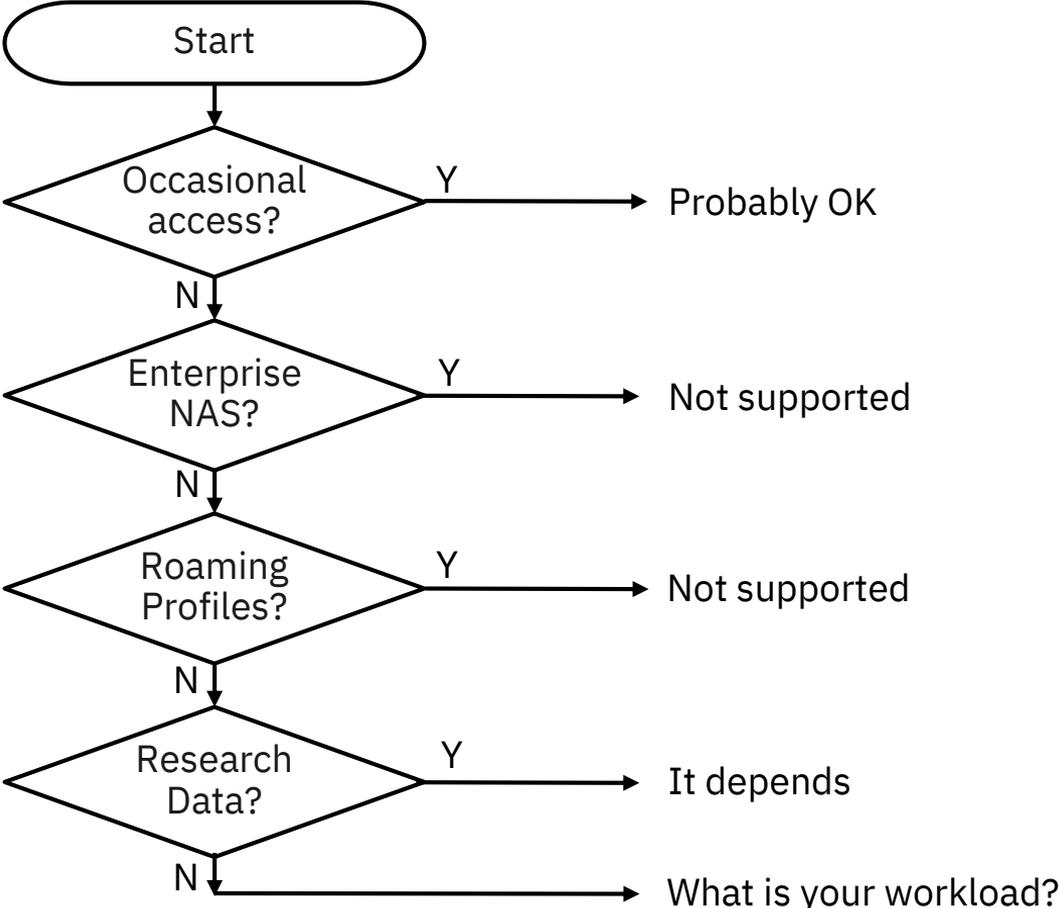
Success Factors

Successful deployments of Spectrum Scale and Spectrum Scale Protocols depend on

- System administrators (Deep skills in Linux, networking, system software, etc.)
- End-to-end skills to architect, implement, operate and troubleshoot the whole Spectrum Scale Environment including software, servers, storage and networks as well as additional functions such as backup, workload scheduling and monitoring
- Capabilities of open source protocol implementations (e.g. Samba, Ganesha) meet the customer requirements
- The trade-offs and limitations of integrating a parallel file system (Spectrum Scale) and open source implementations of NFS (Ganesha) and SMB (Samba) are acceptable for the administrators, the end users and the business owner
- Availability of low latency and high throughput Private Cluster Network (see later section)
- The right workload

- Start with a small environment and use elementary features only.
- Acquires skill in a stable production environment.
- Incrementally grow environment and adoption of advanced features.

Postlude



**Thank
you!**