



# ExtremeStor™

## IBM Spectrum Scale V5 With Expert Architectural Design, Implementation and Support

### Highlights

#### IBM Spectrum Scale V5: New Levels of Storage Performance and Efficiency

- Substantial improvements in I/O performance
- Significantly reduced inter-node software path latency to support the newest low-latency, high-bandwidth NVMe technology
- Improved performance for many small and large block size workloads simultaneously from new 4 MB default block size with variable sub-block size
- Improved metadata operation performance to a single directory from multiple nodes

#### High Performance Parallel File System

IBM Spectrum Scale stripes data in parallel to multiple servers with multiple disks, for extreme, scalable performance

#### Expert Architecture, Solution Design, Tuning, and Implementation

Extensive analysis of existing and future needs, comprehensive solution architecture and validated hardware and software build that ships fully integrated

#### Intelligent Automated Tiering

Data placement policy engine intelligently tiers data from high-performance flash to low cost storage across a single namespace; places data based on a file's "heat" – easier and more economical than tiering to more expensive, high latency external object or cloud storage

#### Built-in Data Protection and Services

Snapshots, Fail-over, Replication, RAID 6, Quotas, Tiering, Compression, Storage pools, ILM data placement policies, and Active File Management (AFM)

#### Exclusively Focused on Technical Computing and Data Management

Applied Data Systems is 100% focused on technical computing, specializing in CPU and GPU computing, low latency networking and IBM Spectrum Scale high-performance storage

#### IBM Spectrum Scale V5 with Expert Architectural Design and Support

Applied Data Systems ExtremeStor is based on the latest and most innovative IBM Spectrum Scale release v5, with expert architectural design and performance tuning by globally acknowledged IBM Spectrum Scale experts. Applied Data Systems delivers IBM Spectrum Scale on fully integrated, top quality, industry standard hardware for maximum performance, reliability and data protection. Expert support is provided by Applied Data Systems engineers, backed by IBM providing second level support.

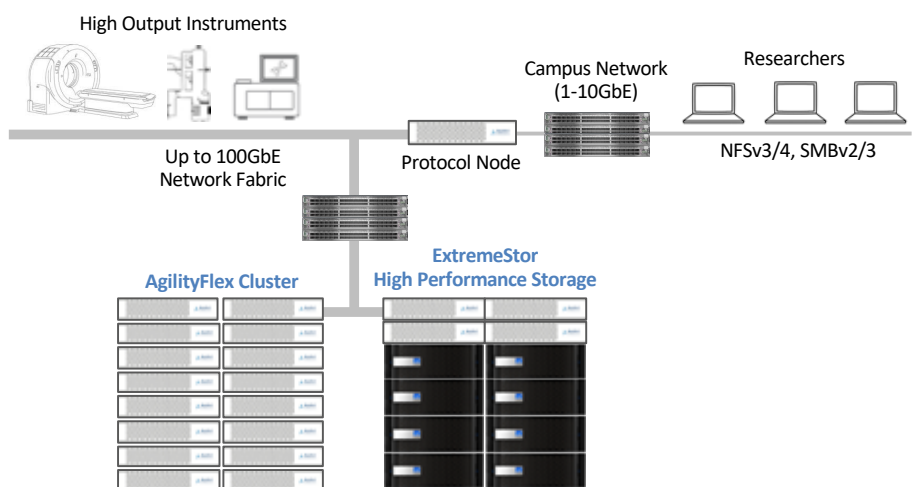
#### Extreme Performance and Limitless Scale

IBM Spectrum Scale is a parallel file system, trusted by 1000's of organizations, that stripes data in parallel to multiple disks across multiple servers, delivering both extreme performance and high capacity at low cost. A layer of ultra-fast, low-latency, high-bandwidth NVMe flash technology can also be integrated with data automatically tiered to lower cost storage pools as access becomes less frequent.

#### Accelerate Life Science, Machine Learning and Research Workflows

An example scientific computing workflow is shown below, where data is ingested from high output instruments, sequencers, cameras, and microscopes into ExtremeStor high-performance, multi-tier, single namespace storage. Data is then processed and analyzed by the AgilityFlex Compute Cluster, also provided by Applied Data Systems. ExtremeStor storage provides the cluster the necessary high throughput and low latency to accelerate data hungry compute node processing, speeding completion.

Results files are stored back on ExtremeStor shared storage where researchers access data from their workstations over standard NFS/SMB protocols to perform local workstation-scale data analysis or kick off subsequent HPC-scale analysis on the AgilityFlex Cluster – where it can access data quickly again. Cold data can be migrated by policy to low-cost, high-capacity drives – within the same namespace – without the added complexity or expense of tiering to external object or cloud storage. External tiering (tape, object, cloud) is also available.



ExtremeStor High Performance, Multi-tier, Single Namespace Architecture

### Performance of Flash with the Capacity and Economics of Disk

Economic high capacity disks can be accessed in parallel for high throughput and capacity, combined with a high performance flash tier along with server-side flash based Local Read Only Cache (LROC) and Highly Available Write Cache (HAWC) for maximum performance

### Consistently Fast Performance

IBM Spectrum Scale scatters data across spinning disk, removing the impact of fragmentation, so there is no degradation in performance as the system fills

### High Speed Metadata Scanning

Empowers ILM policy data movement and data management at scale

### Variable Block Size

Supports a wide range of file system block sizes to match I/O requirements

### Simplified Data Management

Always-on operation, even during capacity and performance expansion, consolidated management console

### Enhanced GUI

Monitors performance, capacity, network, and enhanced maintenance and support, including interaction with IBM support

### Reduces Storage Costs

Built on open, industry standard hardware, roll in new technology over time from multiple vendors

### Genuine IBM Support

2nd and 3rd tier support provided by the IBM Spectrum Scale experts at IBM directly

### Enterprise Reliability and Availability

Rock solid file system and high scale data management trusted by 1000's of organizations combined with best of breed components for extreme performance and data protection

### Flexible Building Blocks

ExtremeStor is delivered as a modular, repeatable, and highly supportable solution consisting of best of breed industry standard components

### Simplified Data Management at Scale With Extensive Data Services Included

In addition to delivering affordable high performance at scale, ExtremeStor includes integrated management tools and an advanced GUI to easily administer billions of files and petabytes of data – all under a single global namespace. These include Quotas, Snapshots, Fail-over, Compression, Quality of Service, Filesets, Multi-protocol access (NFSv3/4, SMBv2/3, FTP, HTTP, S3, HDFS and POSIX parallel file system client), Storage pools, RAID 6, ILM data placement policies, and Active File Management (AFM) WAN caching – all included at no extra charge.

### Intelligent Tiering and Comprehensive Information Life Cycle Management

IBM Spectrum Scale includes a powerful information lifecycle management toolset that includes a rule-based policy engine and ultra-fast metadata scanning to tier data internally across flash and disk storage pools. Frequently accessed files are stored in higher-performing (NVMe Flash SSD) storage pools, while files that are not accessed frequently are stored in more economical, high capacity, slower disk-based storage pools. Migration policies transparently move data from one storage pool to another without changing the file's location in the directory structure. Automated analysis of data usage patterns pull data back up to higher performance tiers as needed.

### Flexible ExtremeStor Modular Building Blocks

Applied Data Systems delivers ExtremeStor as a modular, repeatable, and highly supportable solution consisting of best of breed industry standard components. Fully custom systems are also available.



Model	ES17X2	ES27X2	ES67X4
Throughput	Up to 15 GB/s	Up to 30 GB/s	Up to 60 GB/s
NSD's	2	2	4
Included Arrays	1 x E5760 Hybrid Disk	2 x E5760 Hybrid Disk	4 x EF570 All Flash
Included Expansion	0	2	4
Maximum Expansion	Up to 6	Up to 12	Up to 12
Maximum Disks	Up 480	Up to 960	Up to 288
Disk Capacity	4, 8, 10, 12TB		
Data Protection	RAID levels 0, 1, 5, 6, and 10		
Protocols	Dual active controller with automated I/O path failover NFSv3/4, SMBv2/3, FTP, HTTP, S3, HDFS and Spectrum Scale POSIX parallel file system client for extreme performance		
Data Services	Snapshots, Replication, Quotas, Tiering, Compression, Encryption		
Interconnects	10GbE, 40GbE, 100GbE Ethernet EDR/FDR InfiniBand		
Scalability	Building blocks can be scaled to increase capacity and performance while retaining single global namespace		
Integration	Validated hardware and software solution ships fully integrated		
Warranty	3 years base, extendable to 5 years		



For more information from Applied Data Systems:  
844.371.4949 | info@applieddatasystems.com | www.applieddatasystems.com