

# Automated File Data Orchestration in a Parallel Global File System

Presenter: David Flynn, Hammerspace CEO and Founder



# Why Parallel NFS is Relevant Now More Than Ever

## The Current Reality:

- Data orchestration is an absolute requirement across silos, sites, & clouds.
- High-performance requirements have gone mainstream.
- The world is moving to software-defined on commodity infrastructure.
- Linux is ubiquitous → enables a sophisticated, standards-based, open-source client to come built-in (not third-party).

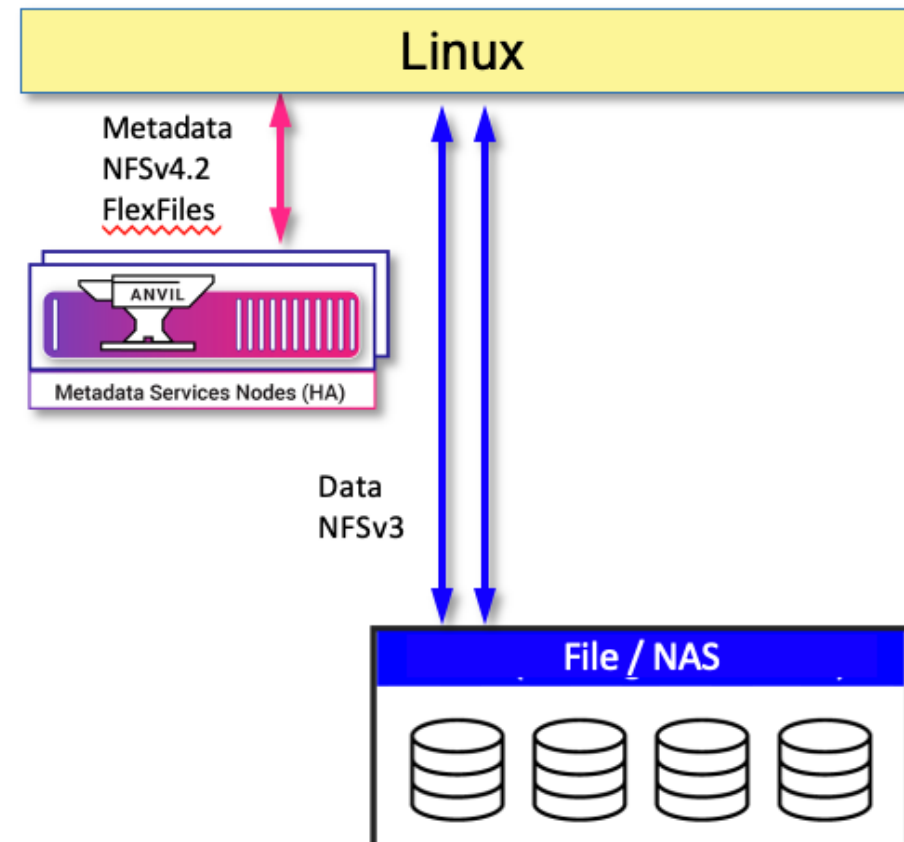
## Therefore:

- NFS 4.2 solves these problems.
  - File access that bridges storage silos, sites & clouds.
  - Parallel file system with no need to install third-party client & management tools.
  - Avoids need to rewrite apps to use object storage.



# Hammerspace Architecture Overview


- **Metadata**
  - Hammerspace “Anvil”
  - Bare-metal, virtual, or container deployment
  - Synchronous replicated cluster for HA
  - Billions of inodes with millions active open
  - Full enterprise NAS data services
  - Instant data-in-place assimilation
- **Client**
  - NFS v4.2 in-box from RHEL 7.6 onward
- **Data**
  - Any NFS v3 NAS
  - Leverages NTAP, Isilon file clone APIs
  - Linear scalable data-path performance





# DSX – Store Function


- Bare-metal, virtual or container deployment
- Parallel, linear scalable performance
- Sources any block storage
  - Direct attached
    - SSD, NVMe, HDD
    - Optional local striping and mirroring
  - Network attached
    - SAN, iSCSI, EBS
- Supports share snapshots and file clones
- Client can mirror writes to multiple DSX nodes
- Or use erasure encoded groups of DSX nodes





# DSX – Mover / Cloud Mover Function


- Bare-metal, virtual or container deployment
- Parallel, linear scalable performance
- Stateless, scale-out
- Fully automatic scheduling
- File to file mobility
  - NFSv3
  - **No interruption to ongoing access**
- File to object mobility
  - S3, Azure Blob, etc. over HTTPs
  - Global dedupe, compression, encryption
  - Transfer & egress optimized





# DSX – Portal Function – Legacy Client Support


- Bare-metal, virtual or container deployment
- Parallel, linear scalable performance
- Stateless, scale-out
- Virtual IPs with fail-over
- NFS v3, SMB 2.x/3 and S3
- Global file locking
- Extensive Caching
  - Metadata
  - Read data
  - Write-back and write-through caching as appropriate





# DSX – Containerized Microservices

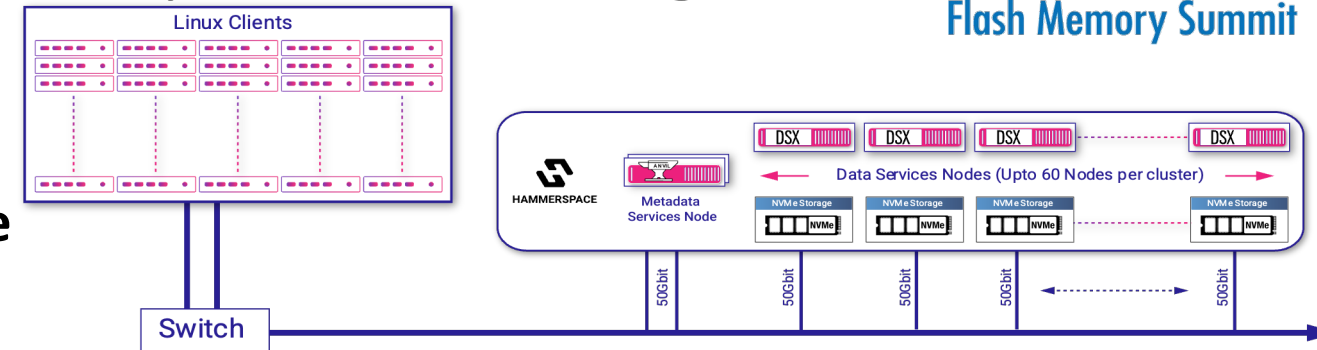
- Deployment flexibility
  - Co-resident on client nodes (hyper-converged)
  - Dedicated storage-only nodes
- Eliminates networking hops
  - Port, cost and latency reduction
- Bypasses serialization over NFS
  - IO short-circuits in the kernel
- Achieves full NVMe performance
  - Tens of Gbytes per second
  - Millions of IOPS
  - Microsecond latency





# Example: Linear Scalability Saturating Infrastructure

- Performance testing showed **linearly scale** from small to large:
  - Saturating the network for throughput-dependent workloads.
  - And saturating the backend storage for IOPS-dependent workloads.
- Testing showed **16 DSX nodes hit 1.17 Tbits/s** with 32kb file sizes with low CPU overhead.
- In testing for raw IOPS with this configuration, the same test using small 4k files achieved 6.17m IOPS.
- Testing showed linear scalability to limits of network and storage, by adding more nodes.



- **Test Suite:**
  - 192 clients
  - 16 DSX Nodes
    - Can scale to 500 DSX nodes per cluster x 16 clusters.
    - DSX Nodes can be mixed instance types.
- I/O Pattern Randomized – 90/10 R/W mix
- NFS Exports were mounted with NFS 4.2

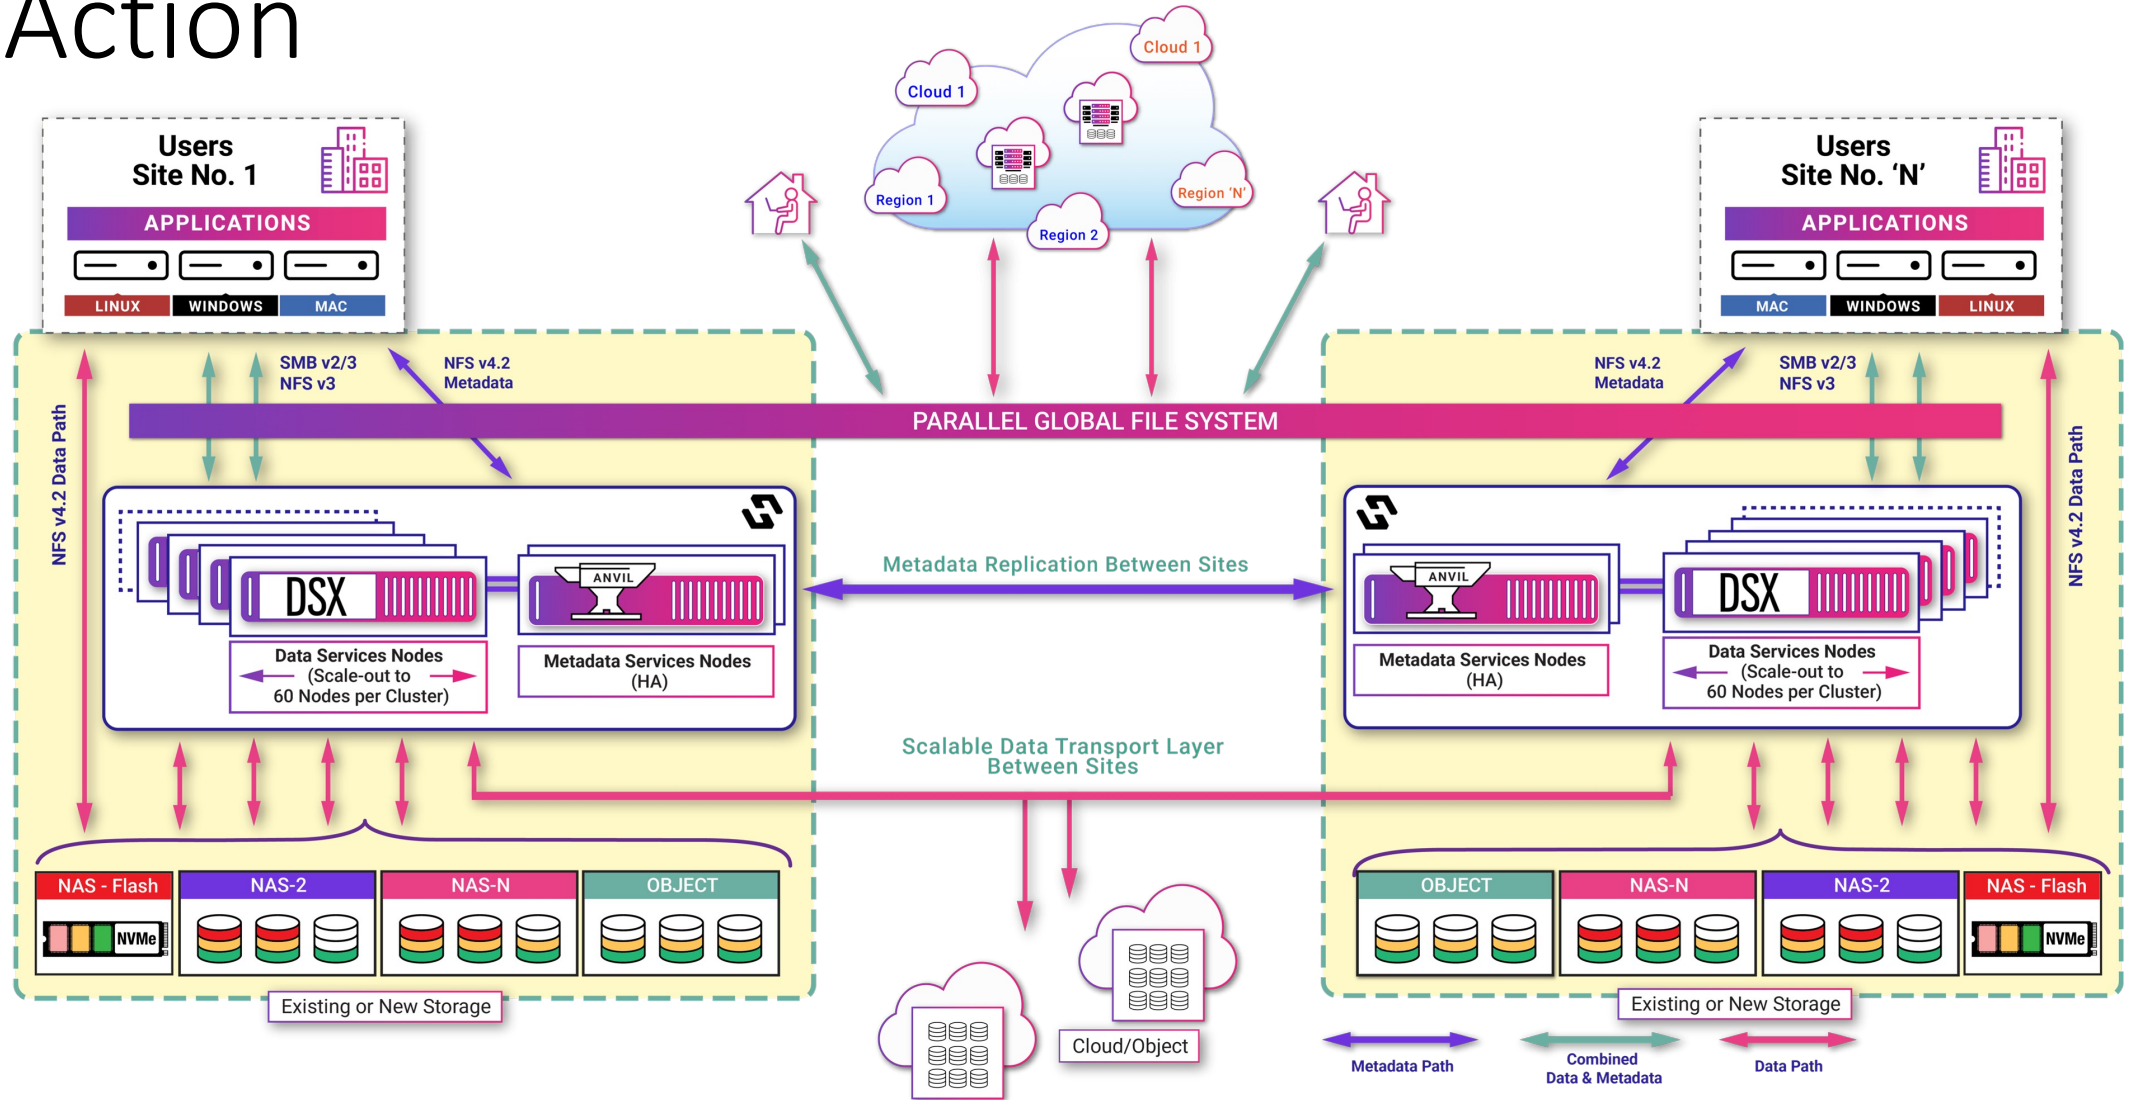


# NFS4.2 – Target Advantage Areas

- **NFS 4.2 - Sweet Spots:**
  - Scale-out distributed high-performance file-based workloads.
  - Stateful file access at scale globally across block, file, & object.
  - No client software required – included in standard Linux distributions.
  - Runs on commodity hardware.
  - Supports any on-prem or cloud storage of all types from any vendor.
- **With Hammerspace:**
  - Supports decentralized environments:
    - Global file system spanning silos & sites
  - Actionable metadata, including custom metadata driving objective-based policies across any storage type and location.

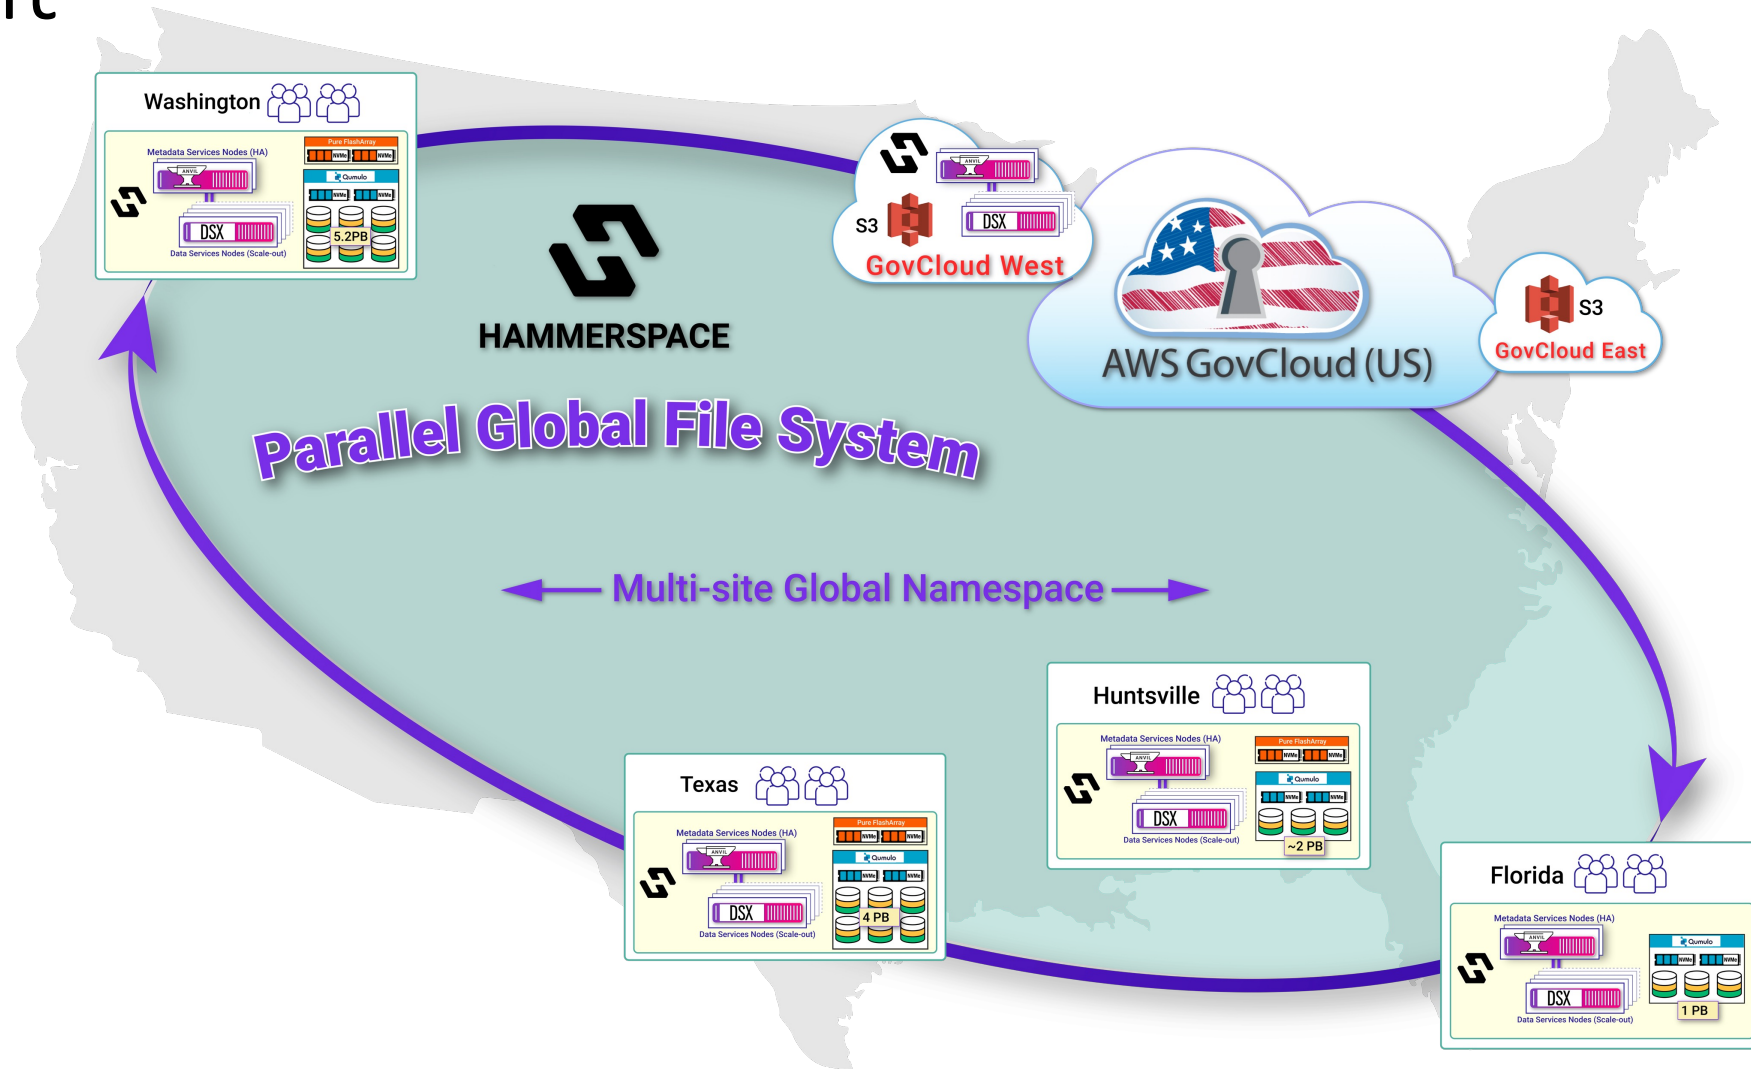


# Unstructured Data Orchestration System in Action





# Data Orchestration Powering Space Flight





# Summary

## File Access and Orchestration for Any Data, on Any Storage, Anywhere.

- NFS 4.2 solves global high-performance file access.
- Flexfiles Layouts provide flexibility to bridge block, file & object at scale, globally.
- Enables transparent live data mobility.
- Supports software-defined commodity model.
- Leverages existing ubiquitous NFS client  
→ No third-party client required.
- Supports extreme scale-out & high-performance file workflows across silos, sites, & clouds.



# Thank You!

David.Flynn@Hammerspace.com