File Storage at Scale

Bjorn Andersson, BlueArc

Who is BlueArc?

Private Company, founded in 1998

Headquarters in San Jose, CA

Highest Performing NAS in the Industry

Proven in the most demanding HPC environments

6th Generation Product

Patented Hardware Accelerated Architecture

Doubled performance and scale with each platform release

BLUEARC®

Select HPC Research Customers

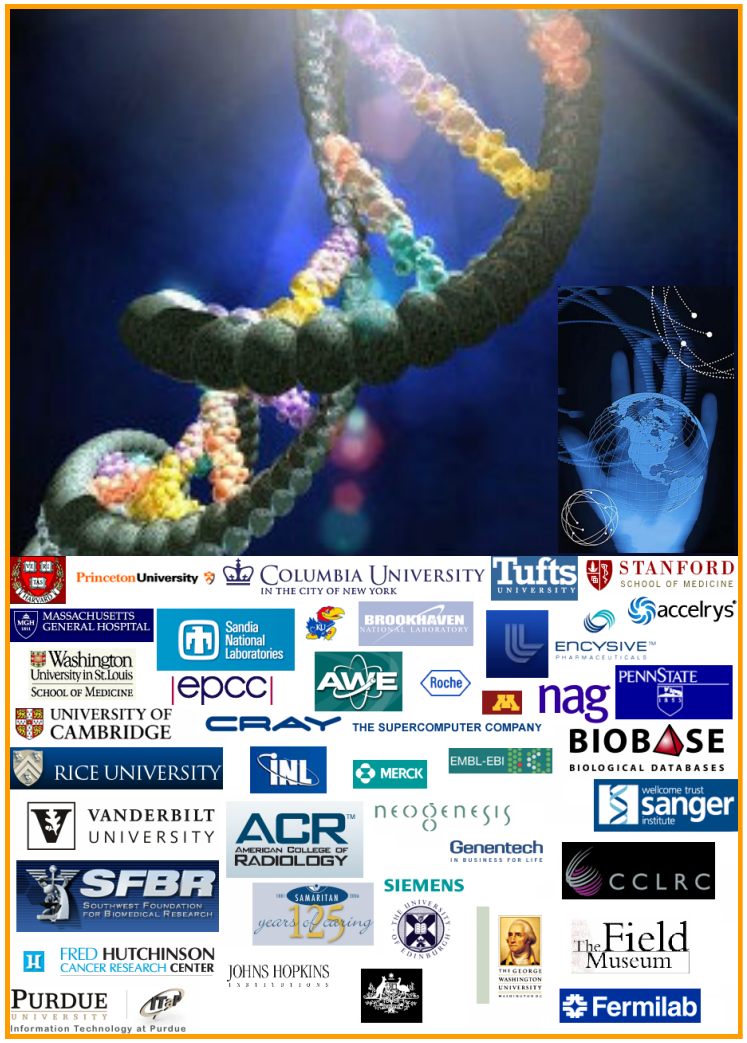

- Atomic Weapons Establishment
- Brookhaven National Laboratory
- Baylor College of Medicine
- Cambridge University
- Chevron
- Columbia University
- Cold Spring Harbor Laboratory
- Commissariat a L'Energie Atomique
- Cray, Inc.
- Duke University
- European Bioinformatics Institute
- Fermi National Accelerator Laboratory
- Fred Hutchinson Cancer Research Center
- Genentech
- Georgia Institute of Technology
- HECToR Partners
- Idaho National Laboratory
- Jet Propulsion Laboratory
- Johns Hopkins University
- Lawrence Berkeley National Laboratory
- Lawrence Livermore National Laboratory
- Merck & Co.
- Massachusetts Institute of Technology
- NASA
- National Cancer Institute
- National Heart, Lung, & Blood Institute
- National Oceanic and Atmospheric Administration
- National Renewable Energy Laboratory
- Oak Ridge National Labs
- Ontario Institute for Cancer Research
- Penn State University
- Princeton University
- Purdue University
- Renaissance Computing Institute
- RWTH Aachen
- Sandia National Laboratories
- Sanger Institute
- Stanford University
- Tokyo Tech
- University of California Los Angeles
- University of Michigan
- University of Minnesota
- Vanderbilt University
- Washington University in St. Louis

The Essence of BlueArc

Build for multi-petabyte **scale** and HPC **performance** requirements while using **standards** and adding **features** more typically found in a general purpose file systems

Getting the Most Out Of a Standard Network File System

User Benefits:
Simplicity and familiarity

Standard clients

NFS Clients

Choice with Multiple, Competing Implementations

Optimized Operation At Scale

Bigger pipes

Robust, Scalable Performance

Decoupled

NFS Servers (Clustered)

High Availability

Hardware Acceleration (bandwidth&iops)
Millions of files

Easy Capacity Scaling

Advanced Tiered Architecture, Open backend

BLUEARC

How is this Architected?

- **Offload engine for high speed data transfer**

- **CPUs for unburdened data management**

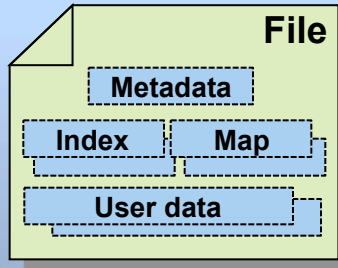

- **Intelligent file objects for high efficiency**

No compromise design:

- **Parallel processing** of core file system (FPGA offload engine)
- Simultaneous access to **CIFS**, **NFS** and **iSCSI** (**pNFS** mid 2011)
- **CPU** runs data management functions not already burdened by file system
- **Object File System Architecture** for Intelligent, flexible data management (SiliconFS)
- **Storage virtualization** and **thin provisioning**

Sustained Performance at Line Rates

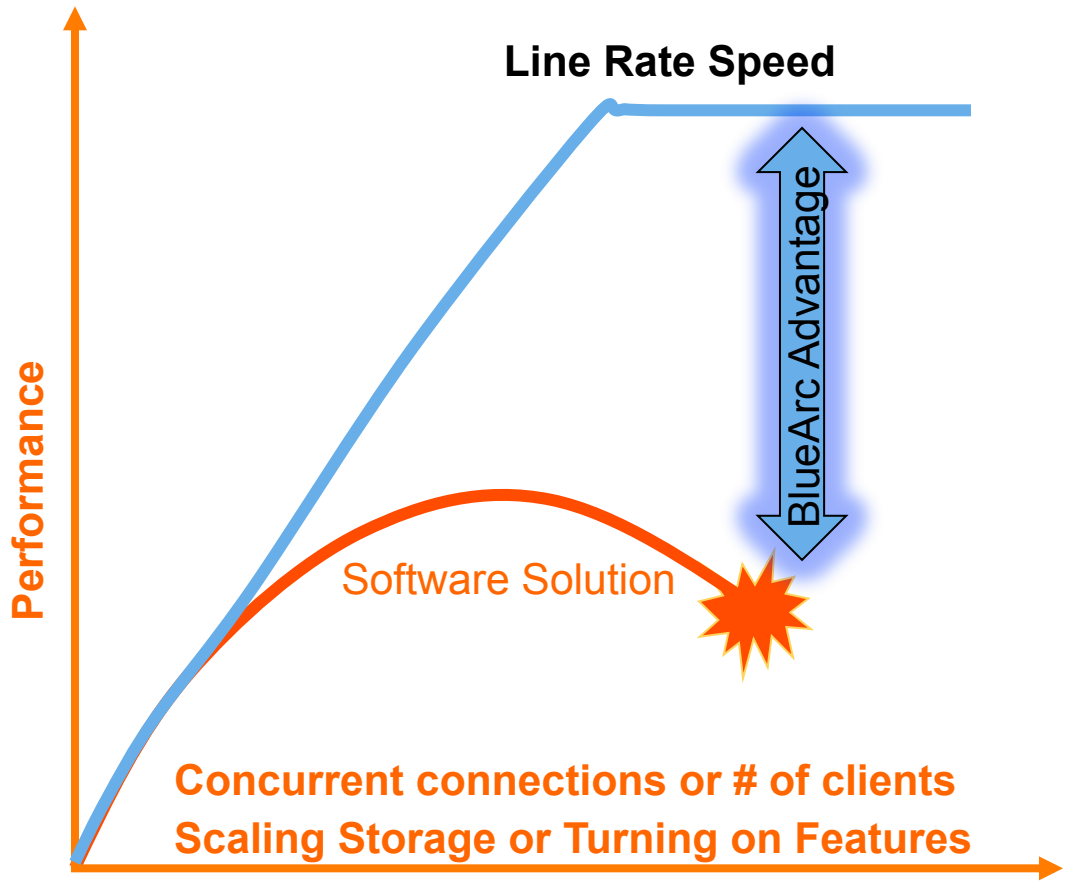

- As concurrent connections or # of clients scale
  - Performance level increases linearly
  - No fall off due to CPU utilization
  - Much higher level of maximum performance
  - If pushed to maximum sustain 100% indefinitely
- Benefits
  - More users per filer
  - More functions per filer
  - Fewer filers & licenses
  - Simplified management

BLUE ARC

Transparent Data Mobility That Really Works!
Tiered Storage for Persistent Data

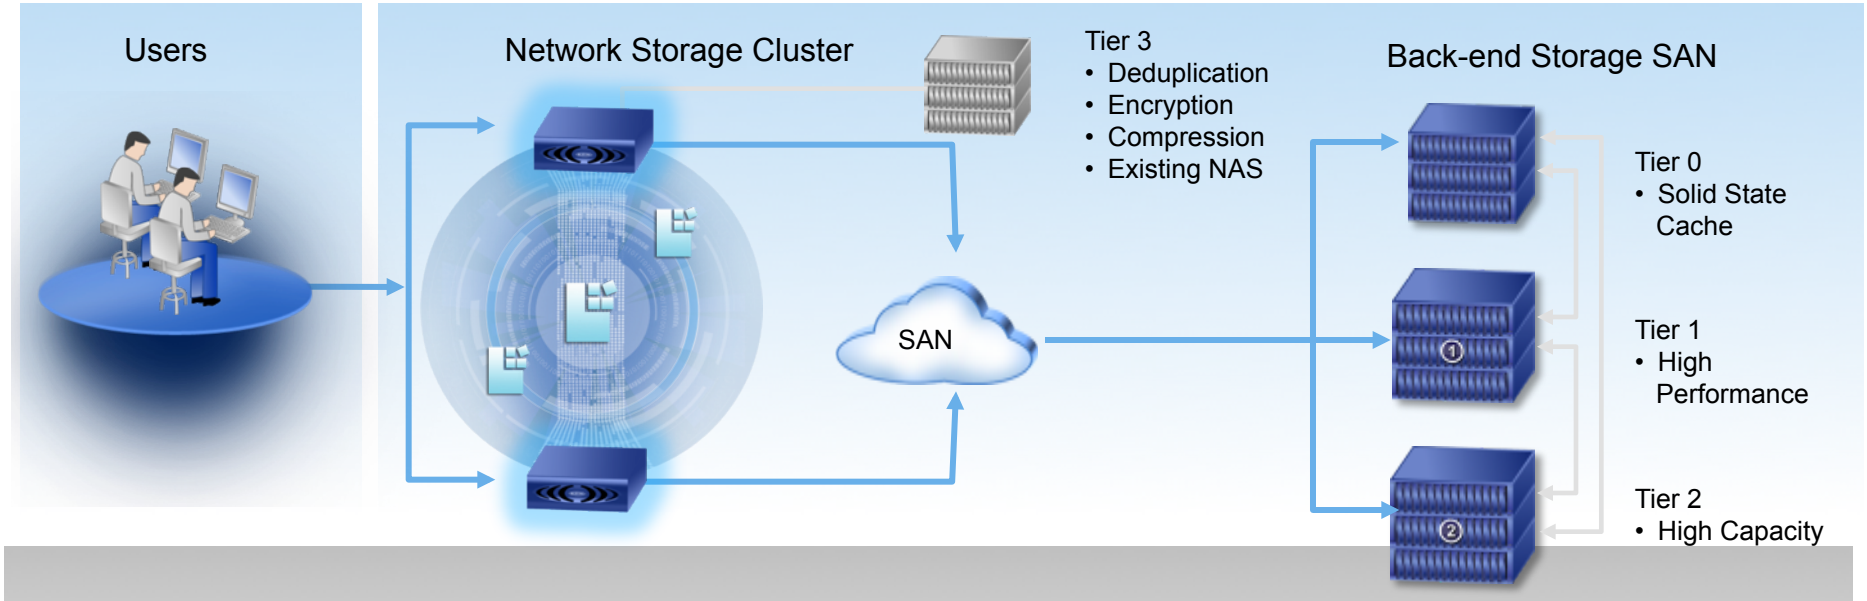

The seamless migration of data across storage tiers within a single namespace

Ease data management and reduce costs

- Automatic and transparent data migration between tiers
- Rules-based policy engine reduces manual intervention
- Third-party or external storage devices as an integrated tier
- Reduced dependence on high performance tier for peak demands

BLUEARC

Robust NFS Performance

HNAS 3200, AMS 2500	300GB x 144 HDDs, 8D+1P	36 Disks	72 Disks
Sequential Read %	I/O Block Size	Throughput in MB/s	Throughput in MB/s
100	32KB	529	934
100	256KB	557	919
100	1024KB	574	923
75	32KB	572	806
75	256KB	552	764
75	1024KB	545	763
50	32KB	583	819
50	256KB	523	718
50	1024KB	530	722
25	32KB	567	857
25	256KB	528	752
25	1024KB	520	781
0	32KB	611	894
0	256KB	593	818
0	1024KB	566	796

Independent of Read/Write mix

Independent of block sizes

Scales

BLUEARC

SPECsfs®2008 Performance

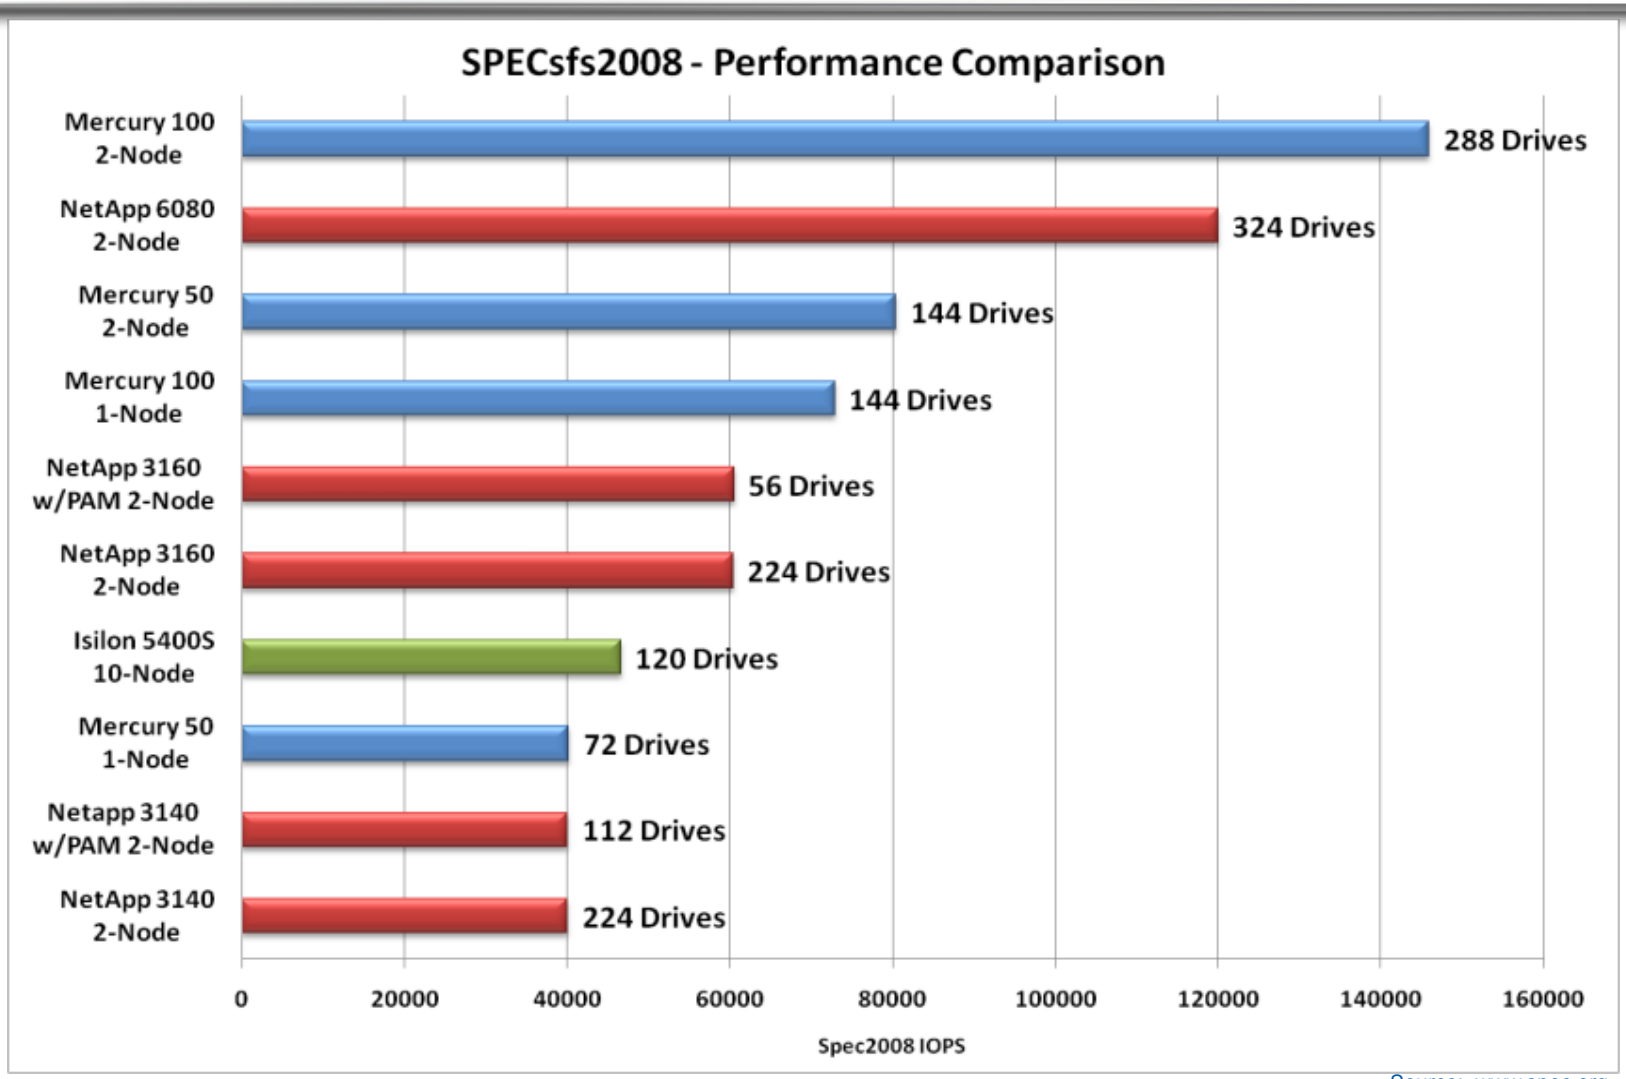

Source: www.spec.org

BLUEARC®

Strong Nodes
BlueArc File Server Platforms

	Mercury 55	Mercury 110	Titan 3100	Titan 3200
				
Product Class	Lower Mid-range	Mid-range	Mid-range	High End
Cluster Nodes	2	Up to 4	Up to 8	Up to 8
Max Storage Capacity	4 PB	8 PB	8 PB	16 PB
Performance (specSFS IOPS)	60,000	100,000	100,000	200,000
NFS Throughput	700 MB/s	1100 MB/s	1200 MB/s	1900 MB/s
Storage Options	All BlueArc storage array options are available with each platform			
Software / File Services	All software and filesystem options (NFS, CIFS, iSCSI) available			

BLUE ARC

Example:
Genome Sequencing Aggregated Workload

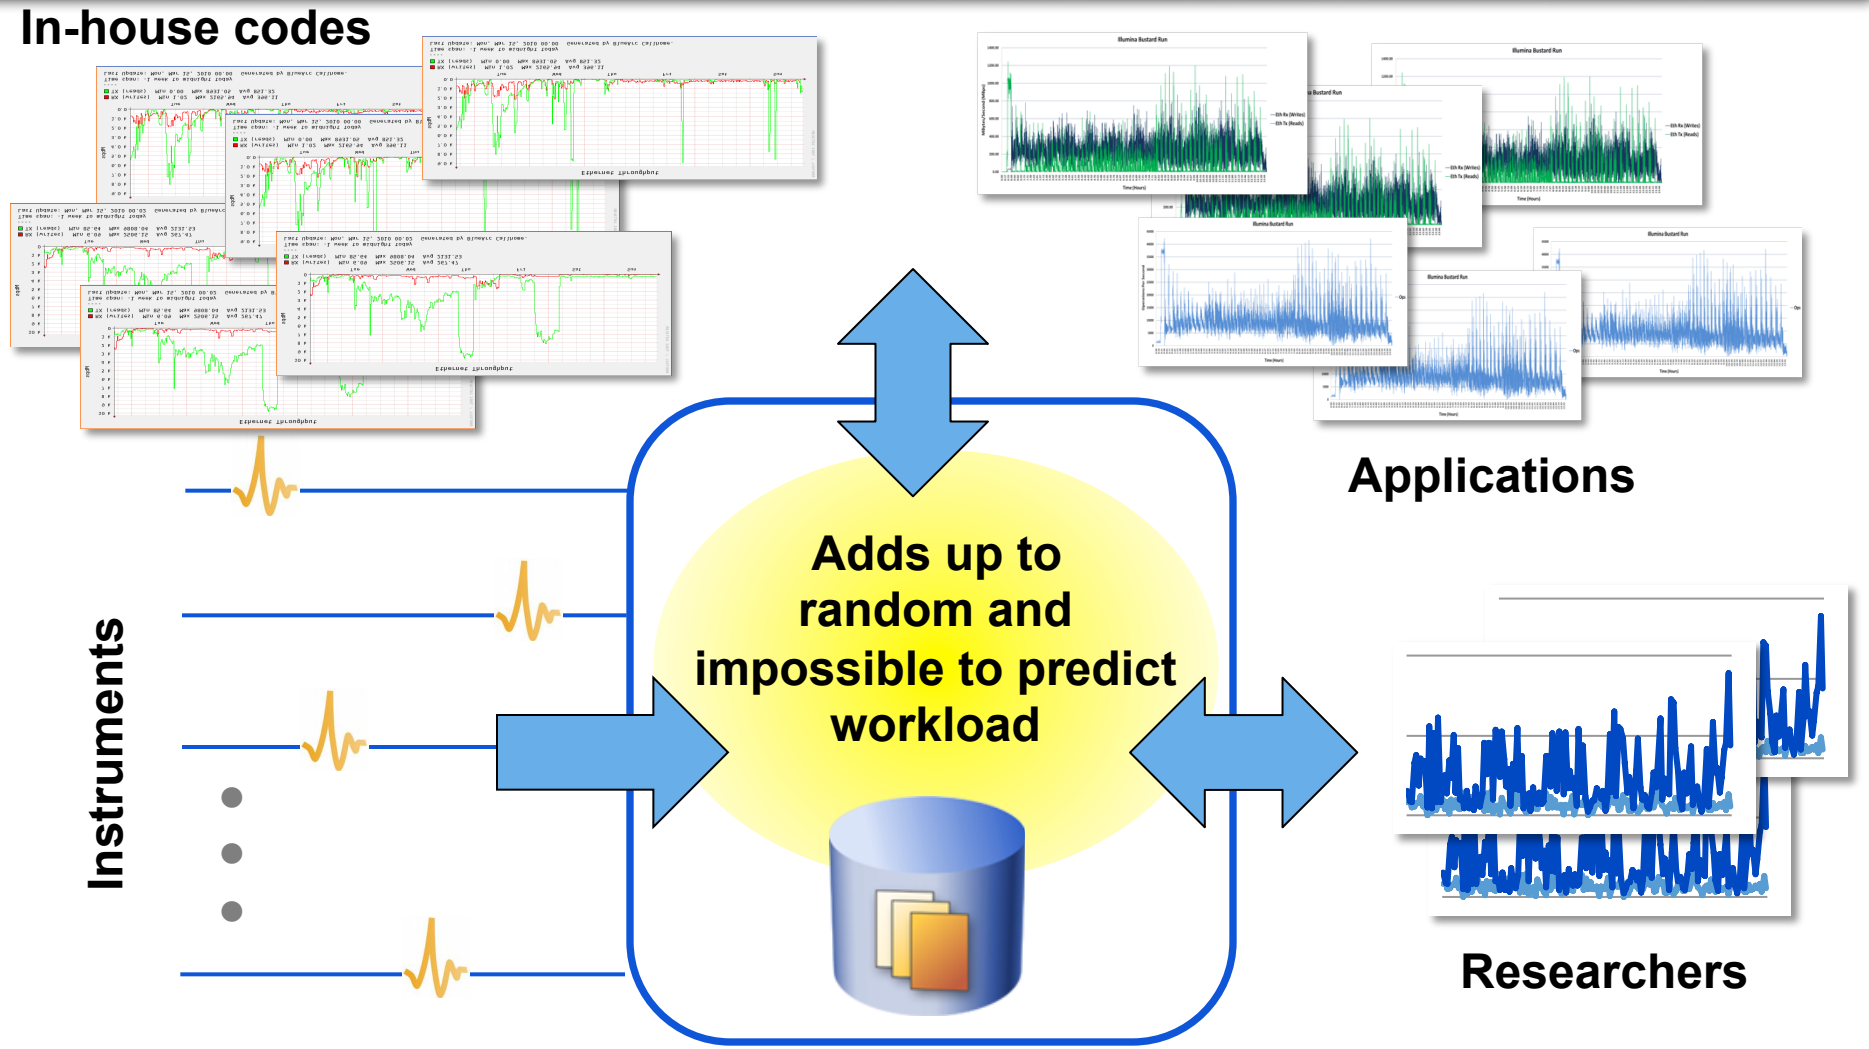

BLUE ARC

Back to the Traditional Network File System

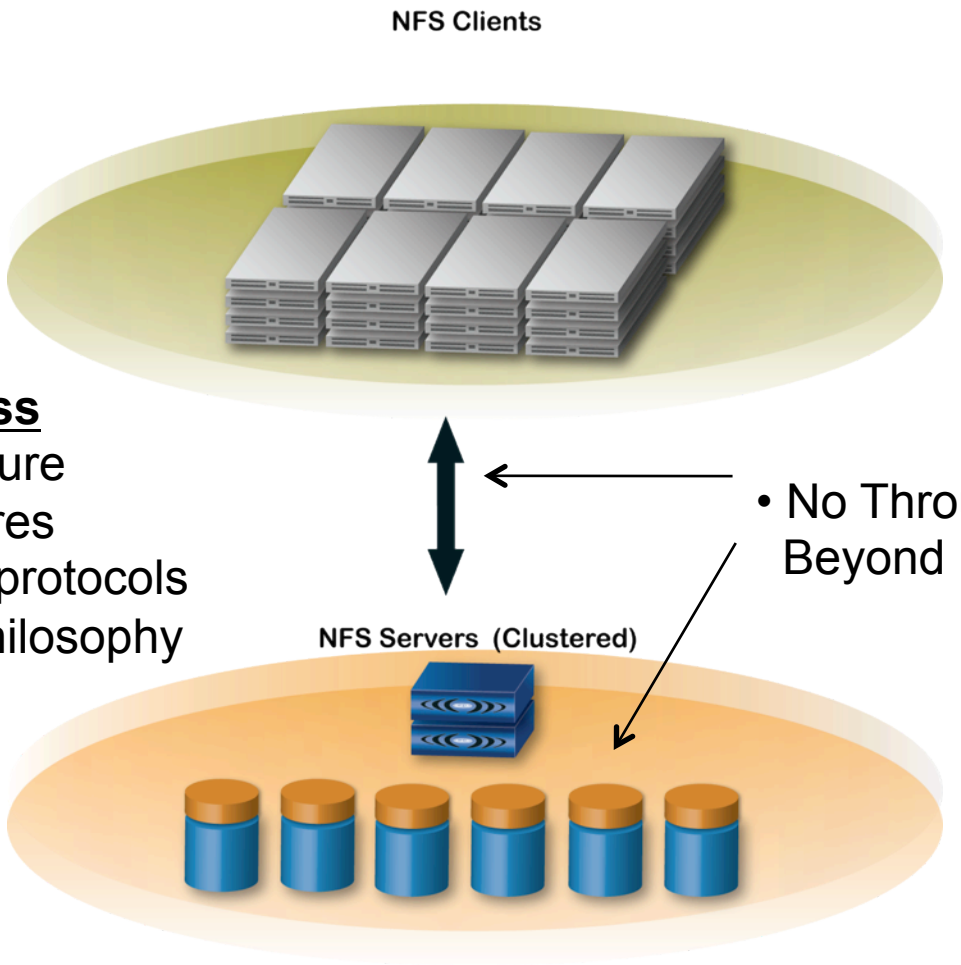

Lots of Goodness

- Proven architecture
- Enterprise features
- Open, standard protocols
- Open storage philosophy

BUT

- No Throughput Aggregation Beyond Line Rate

BLUE ARC

How Does pNFS Change This?

- pNFS adds parallel I/O to the NFS protocol
  - Eliminates the file server bottleneck
  - Provides parallel data paths, even for a single file
- pNFS is part of the NFSv4.1 standard
  - Approved by IETF Dec, 2008
  - RFCs completed editorial review Oct, 2009
  - RFC numbers issued Jan, 2010
- Multiple implementations are in development
  - Client software is expected to be embedded in leading OS distributions
- The Only Industry Standard Parallel File System

BlueArc pNFS Architecture
Leveraging our Technology Portfolio

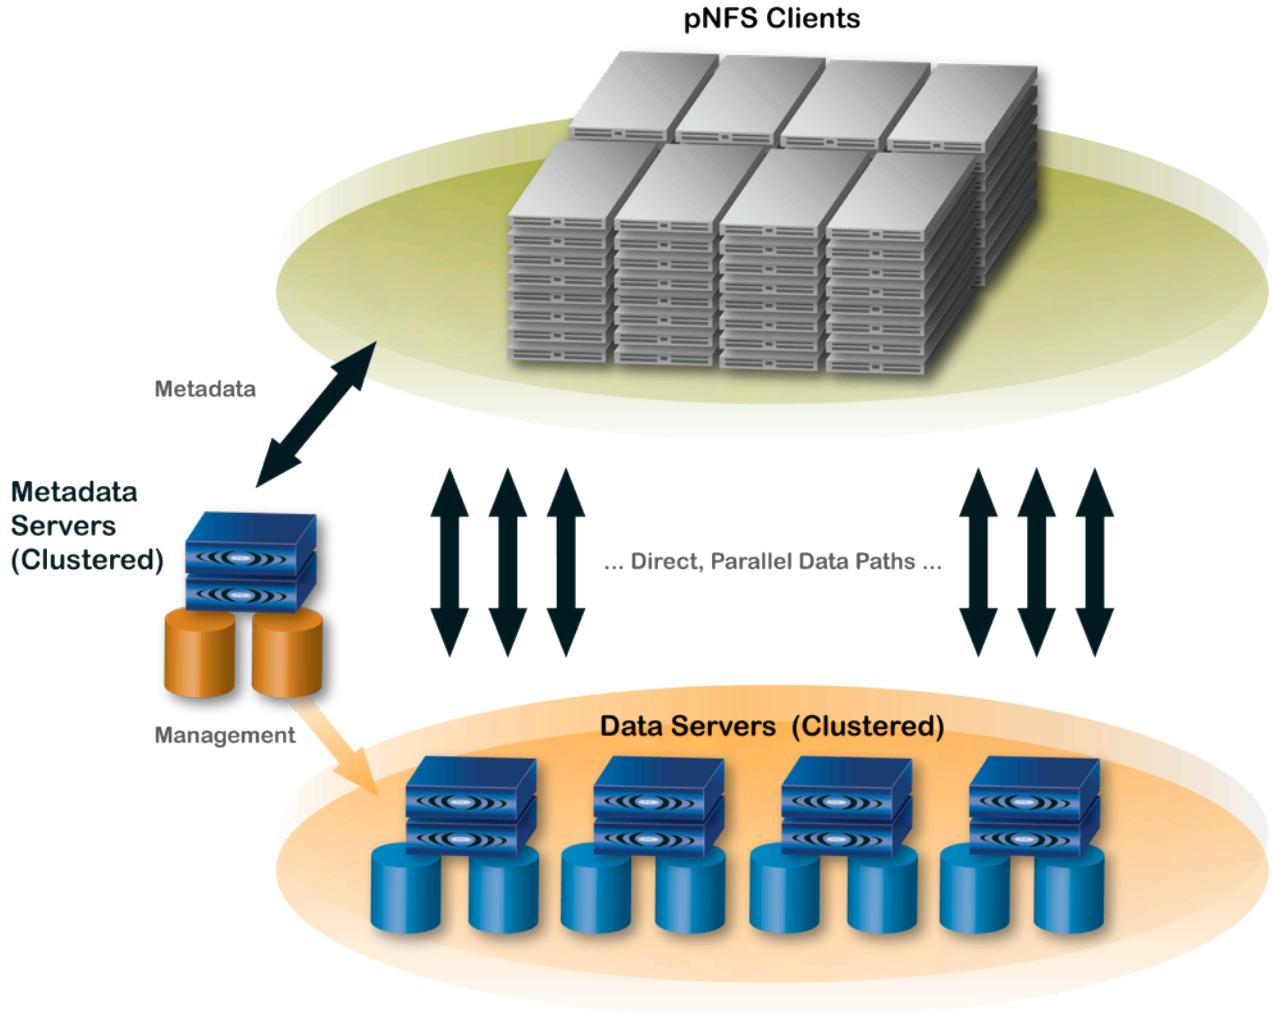

BLUE ARC

Scale-Right Storage for Mixed Environments

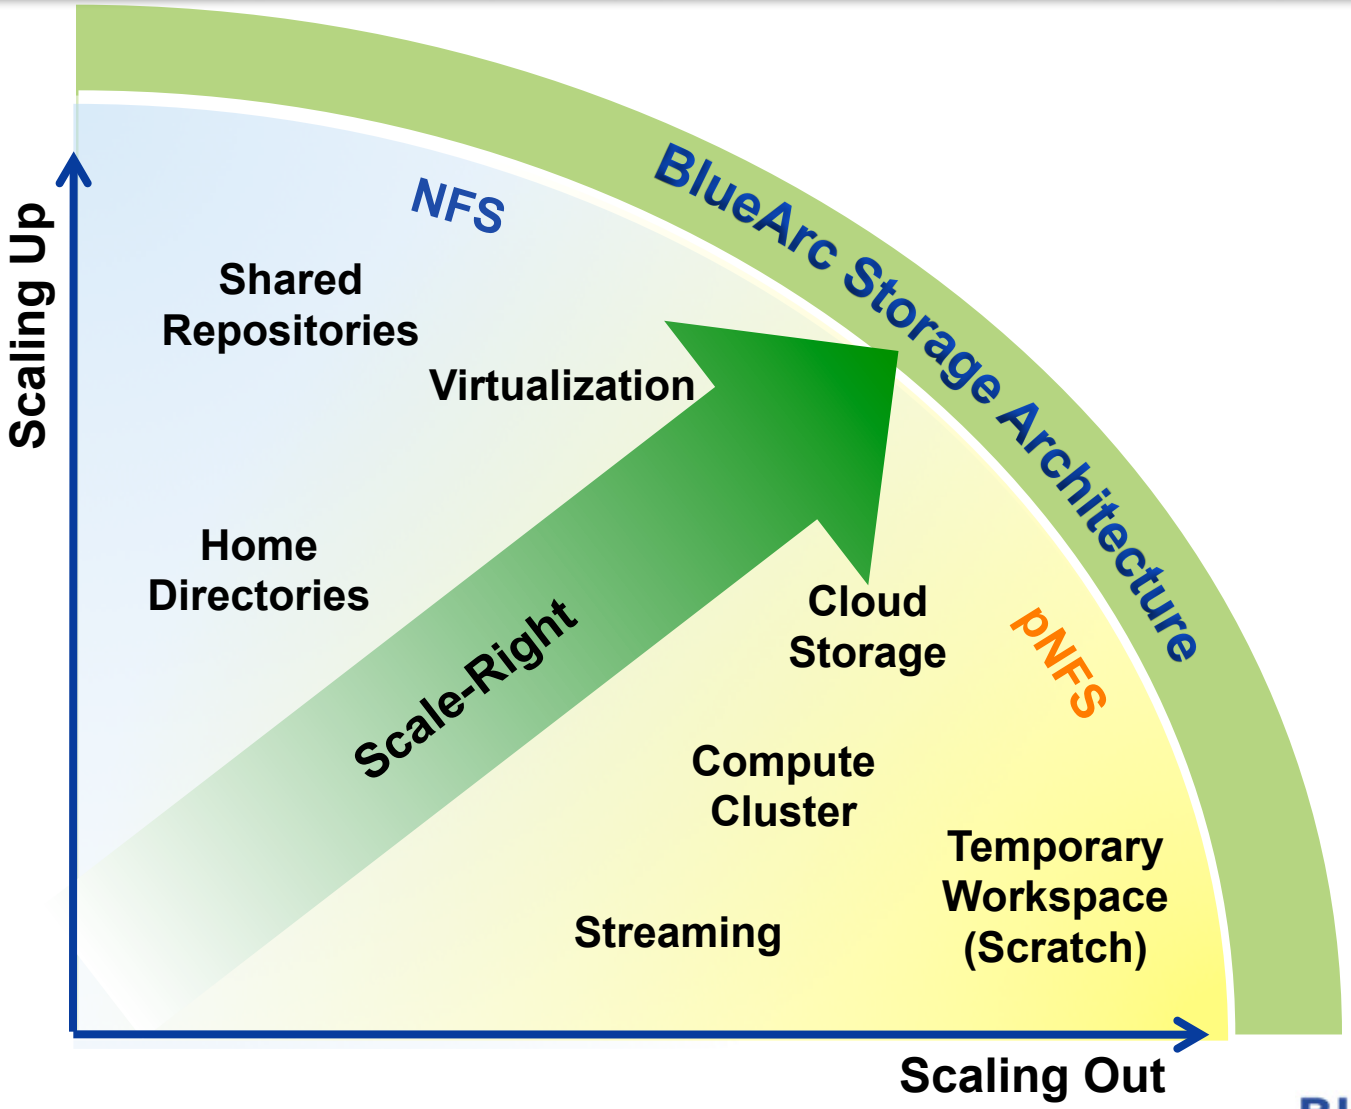

BLUE ARC


BLUE ARC®
Thank You!