

# SCALA STORAGE

## A Market Proven Parallel File System For Data Intensive Storage

### ABOUT SCALA FILE SYSTEM

#### **What is SCALAf/s?**

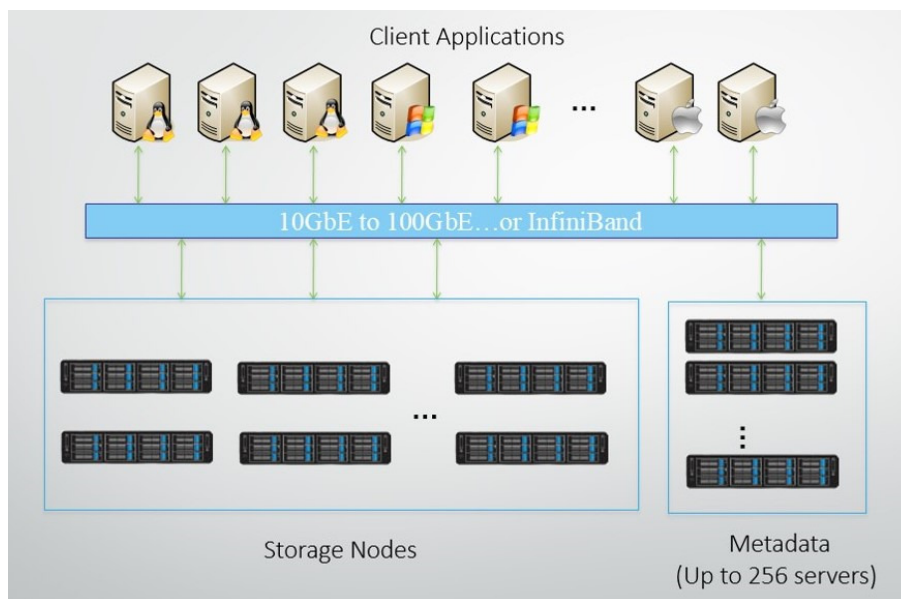
With more than **1 Exabyte** in **customer deployment**, our underlying technology is a **market proven** parallel file system developed to **efficiently handle** multiple variations of **I/O intensive** workloads.

With great **success worldwide**, this proprietary system is **the solution** for over **700 customers**.

The current **largest cluster** installed has more than **300 storage nodes** of **95PB** and over **200GB/s I/O** throughput.

**Worry-free** solution with **24/7 customer support** and cluster management.

#### **System Architecture**



#### **Why is SCALAf/s special?**

Data files are **transparently distributed** over **multiple nodes**. By simply increasing the number of servers and disks in the cluster, you can **seamlessly scale** file system **throughput and capacity** to the needed level.

**SCALAf/s** is **the right choice** for a storage system with the best **combined values** of **high performance, scalability, reliable data protection** and **affordable cost**.

**All-in-One Storage** for **object, block, and file-based data**.

### SUPPORT

With **direct contact** to the **system developers**, SCALAf/s offers **easy management** with designated consultants. Our **highly-trained** and experienced **engineering team** is **available 24/7**.

**No matter how large** or small the system our staff always makes **customer care** a **top priority**.

#### Uncompromised Reliability

After hardware setup, software **installation and updates** can be **done in one hour**. Adding servers for **increased performance** and **capacity** requires **no downtime**. With web based **graphical monitoring and administration**, most cluster **issues can be fixed** remotely **without operation interruption**.

**Support Options** include **Next Business Day, Remote or On-Site, Advanced Hardware Replacement**.

#### Solution Guarantee

Whether your goal is to increase productivity or have a better ROI, **we guarantee usage satisfaction on all SCALA systems**.

# SCALA STORAGE

## BEST COMBINED VALUES

Without controllers or gateways, our file system allows **concurrent access** between **all application clients** and **storage servers** in the cluster.

### Best I/O Throughput

Always **saturate** what **the hardware** can offer. Using **quality drives** and **server configurations** with **dual 10GbE ports**, the **average throughput** of a single SATA or SAS drive is **45 to 100MB/s**. This allows you to **fulfill** your throughput requirements **without using higher cost hardware** such as SSD drives..

### Reliable Data Protection

With **N+M File-level Erasure Coding**, SCALAf<sub>s</sub> **distributes data** content at the **file level** across **multiple storage servers**. When **N+2 erasure coding** is applied the cluster can **sustain operation** with up to **two simultaneous failures**, whether it is an **individual drive** or a **whole node**.

While **traditional** hardware or software **RAID** needs to **rebuild** an **entire drive**, SCALA **rebuilds only** the **files** that are **affected** while using the **entire cluster to rebuild**. Thus it delivers much **faster data recovery** in a **fraction of the time** that **traditional RAID** architectures require. **No downtime or reboot**. **Recovery of 1 terabyte in less than 20 minutes**.

By directory base, the system offers optimum data protection levels for different files and **better capacity utilization, up to 90% with 16+1**.

### Fully POSIX-Compliant


SCALAf<sub>s</sub> is **compatible with all** software applications, x86 based servers and IP networks. No need to make application changes. POSIX enables the **addition of clients** and servers **without downtime**.

SCALAf<sub>s</sub> supports Linux kernels up to **the latest version** and Linux distributions including Debian/Ubuntu, SLES/OpenSuse, or RHEL/Fedora.

### Computational Storage

Besides **metadata management**, SCALAf<sub>s</sub> storage servers can also run **client applications** or **computing tasks** on the **same physical nodes**.

Such a converged structure is **cost-effective** as it provides **computational storage** so the cluster of servers become **shared storage plus data processing units**.



### Linux / Windows / macOS

Native clients, all kernel modules that do not require any patches.

- ❖ **Linux:** all versions from kernel 2.6 and up
- ❖ **Windows:** XP, Windows Server 2003 and up
- ❖ **macOS:** 10.5 and up

### High Level Data Availability

The **automated self-monitoring mechanism** can **single out and report** failed hardware at both the **disk or server level**. Once **inactive hardware is detected**, it will be **isolated for read-only operation** or **taken out**. The system will then start a **self-healing process without any operation interruption**.

Cluster or single server product, SCALA Storage is ideal for HPC, AI, and Big Data Analytics.

# SCALA

STORAGE

## ENTERPRISE FEATURES

Offering SAN and Object storage options our system supports Hadoop, Oracle/SQL, VMware, KVM, and Xen. SCALAs works with iSCSI, CIFS, NFS, HTTP, S3, Swift, HDFS, and Cinder.

### Single Cluster Threshold

	Theoretical	Actual Deployment
Storage nodes	4,096	333
Metadata servers	256	32
System capacity	EB	95PB
Number of files	Unlimited	50 Billion

### Professional Support

- ❖ **Cluster Monitoring:** Free support access via emails, phone and live chat. Our consultants can remotely access the system and run diagnostics to ensure cluster condition. On-site support is also available to keep customer business running smoothly.
- ❖ **Software Maintenance and Update:** Once installed, enjoy free software upgrades and access to a vast suite of enterprise features, such as clone snapshot and replication.
- ❖ **High Quality Hardware:** All hardware including replacements must go through pre-configuration testing so installation and repair will be done in a time sensitive manner.

### Selected Features

- Load balance switch, hardware evenly shared system workload
- Runs on platforms such as x86, OpenPOWER, ARM, and Xeon Phi
- Re-export through Samba, NFS, FTP, HTTP, iSCSI
- Support for group/user ACLs and quota
- Fully active network with automatic failure detection
- Supports InfiniBand, GigE, multiple subnet and bonding
- Cold data sanity check, automatic repair, no downtime
- WORM directory avoids modification of saved data

### Customer Industries

- ❖ Oil and Gas
- ❖ Scientific Computing
  - Genomics
  - Cryo-electron Microscopy
  - Satellite Imaginary and Observatory
  - Geographical Data and Mapping
  - Meteorology/Climate
- ❖ Higher Education
- ❖ Media and Entertainment
- ❖ Telecom and Internet
- ❖ AI and Big Data
- ❖ Video Surveillance



Internet



Oil & Gas



Life & Sciences



Media & Entertainment



Video Surveillance



Public Govt

# SCALA STORAGE

## STATISTICS

### Metadata Cluster

Standalone or embedded with storage servers. **Designed to fulfill any performance requirements, extremely high IOPS or file open rate.**

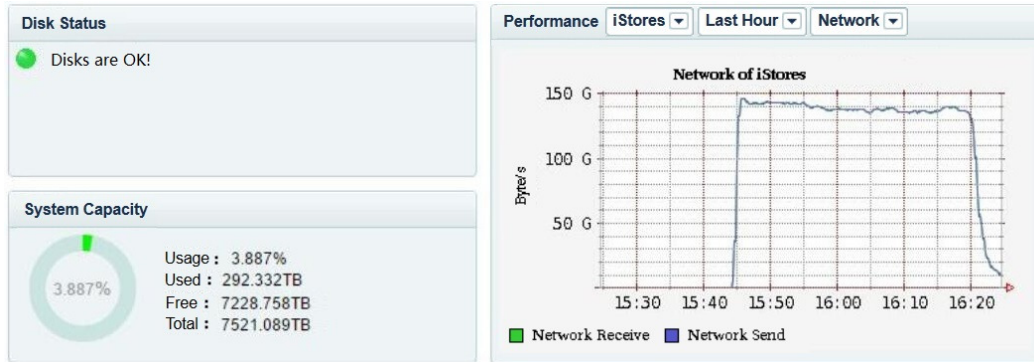
Scalable up to 128 pairs (or 256 nodes). In a testbed with 8 storage servers (each with 24x 2TB 7200r SATA), 160 client processes and dual 10GbE network, a metadata pair can provide a sustained file creation rate of more than 20,000 creates per second

### Linear Throughput Increase


Unlike the competition who must use high-end hardware to boost performance SCALAfS provides predictable and sustainable throughput without deterioration overtime- **with only commodity servers and inexpensive hard drives!**

With highest level of scalability, system throughput will always increase when adding servers and drives.

*Customer on-site 7.5PB, near 150GB/s throughput (102 4U 24-bay storage nodes, dual 10GbE network, 4TB SATA)*



*Block device SSD IOPS vs. Ceph (3 2U 12-bay storage nodes, dual 10GbE network, 240GB SSD, 3 client)*



# SCALALA

STORAGE

## REFERENCES

### Comparison Chart

Features	Isilon Nitro	IBM Spectrum Scale	Lustre	SCALAfS
Snapshots	Yes	Yes-Complex	No	Yes
Independent capacity/performance scaling	No	No	No	Yes
Scale to thousands of nodes	No	Yes	Yes	Yes
QoS	Yes	No	No	Yes
N+M Data Protection	No	No	No	Yes
Encryption	Yes	Yes	No	Coming
S/W only, H/W independent	No	Yes	Yes	Yes
IB & GbE Support	No	Yes	Yes	Yes

### Happy Customers

We take pride in the fact that most of our customers who are now at the Petabyte level started with only a few hundred terabytes.

- First Oil & Gas customer in 2009, I/O throughput 2x greater than StorNEXT FS
- Higher Education: University of Florida, Georgia Southern University
- Scientific Computing: Direct Electron (San Diego) on electron detection for biological molecules
- Video Surveillance: Dante Security

