SARDINA

FishOS Ceph by Sardina Systems

Advanced Open-Source Software-Defined Storage Solution



sales@sardinasystems.com
European-headquartered | Globally Operating

Hyper-Efficient, Scalable, Growth-Focused:
Technology, Product, Enterprise, AI Data Centers

Sardina Systems Proprietary. Copyright (C) 2014 - 2026. All rights reserved.
FishOS® by Sardina Systems. Fully automated cloud solutions for enterprises based on OpenStack, Kubernetes, and Ceph.

1 How Storage Architectures Have Evolved

2 Key Features and Capabilities of Distributed Storage

3 User Interface (UI) Capabilities and Examples

4 Building Storage at Scale: Speed, Stability, Simplicity, and Cost Efficiency

5 Ensuring Business Continuity with Robust Disaster Recovery

6 Proven Results: Customer Success Highlights



Sardina Systems Proprietary. Copyright (C) 2014 - 2026. All rights reserved.

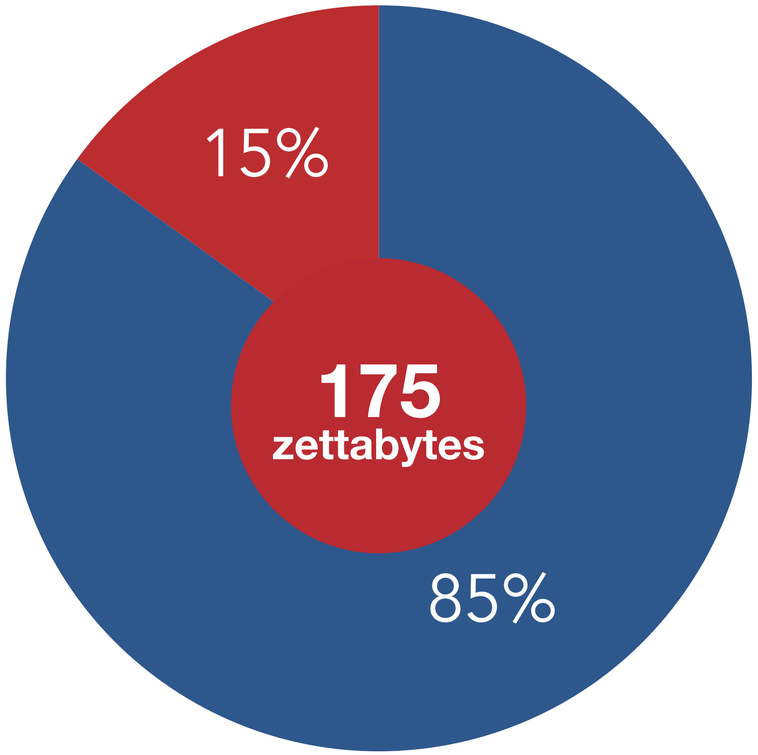
RAPID EXPANSION OF DATA VOLUMES

Data and storage are inseparable — and both are expanding relentlessly

Autonomous vehicles (AVs):
Each vehicle produces around 30 TB of data per day, adding up to hundreds of petabytes over the course of the R&D cycle.

Low-temperature electron microscopy:
A single experiment generates 50–100 TB of data, adding up to tens of petabytes annually.

Next-generation 8K video:
Each audio-video stream reaches 5 Gb, with every channel producing up to 20 PB of data annually.



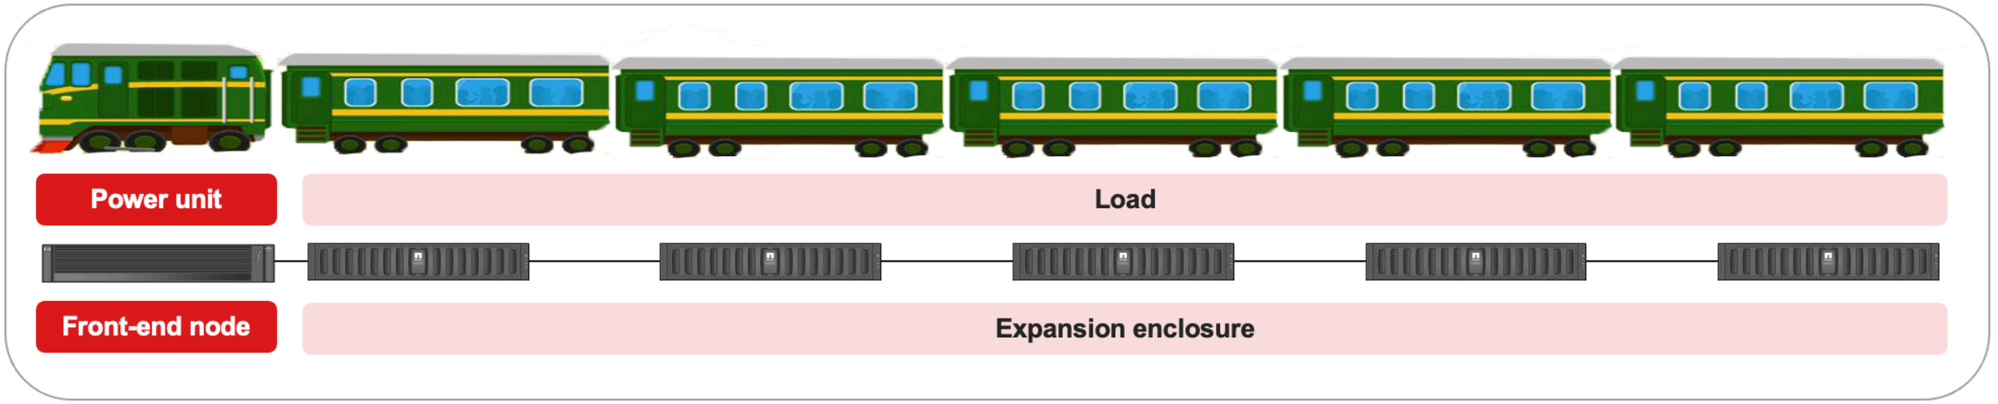
Estimated global **data volume in 2025 is 175 zettabytes**, with **unstructured data making up 85%** of the total. The number is projected to more than **double by 2028, reaching 394 zettabytes.**

*According IDC report

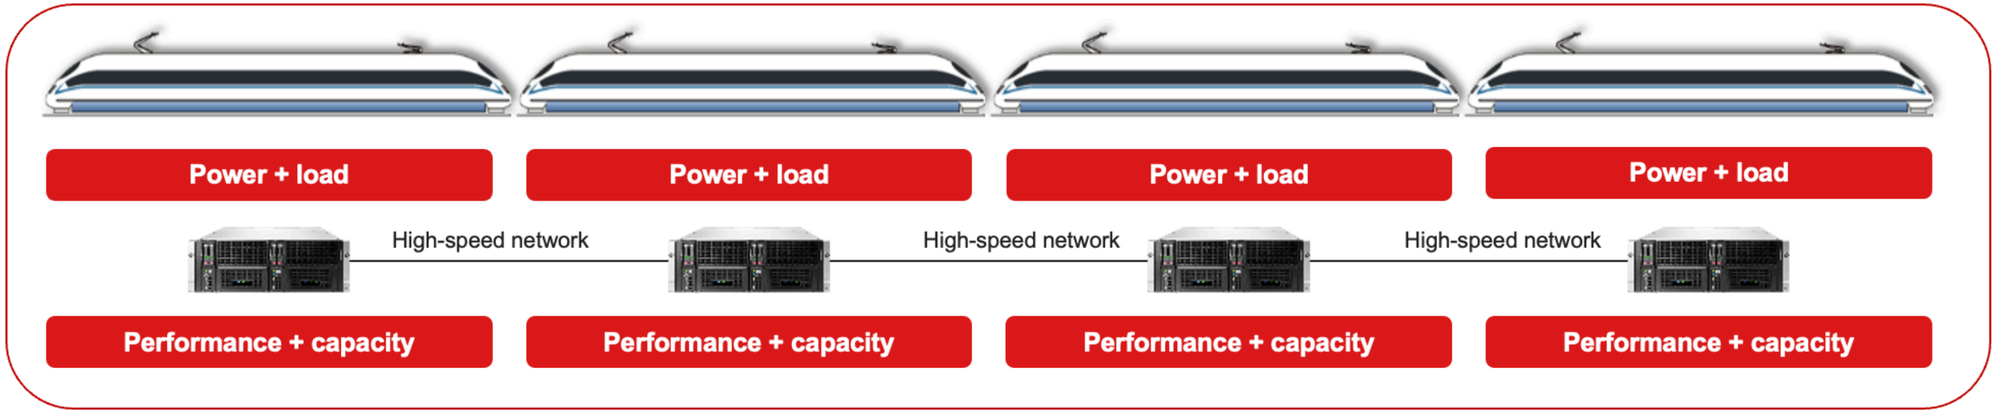
Sardina Systems Proprietary. Copyright (C) 2014 - 2026. All rights reserved.
FishOS® by Sardina Systems. Fully automated cloud solutions for enterprises based on OpenStack, Kubernetes, and Ceph.

RAPID EXPANSION OF DATA VOLUMES

Traditional architecture



Distributed architecture

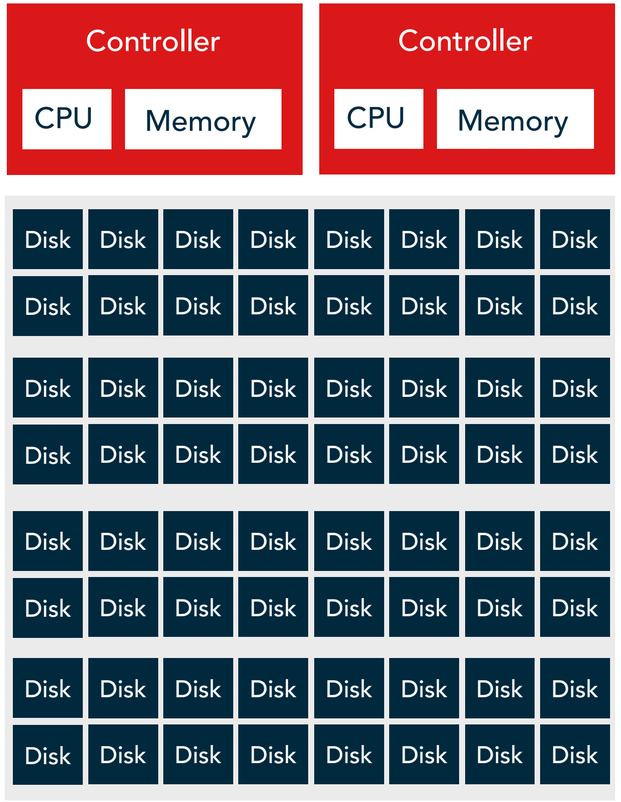


Sardina Systems Proprietary. Copyright (C) 2014 - 2026. All rights reserved.
FishOS® by Sardina Systems. Fully automated cloud solutions for enterprises based on OpenStack, Kubernetes, and Ceph.

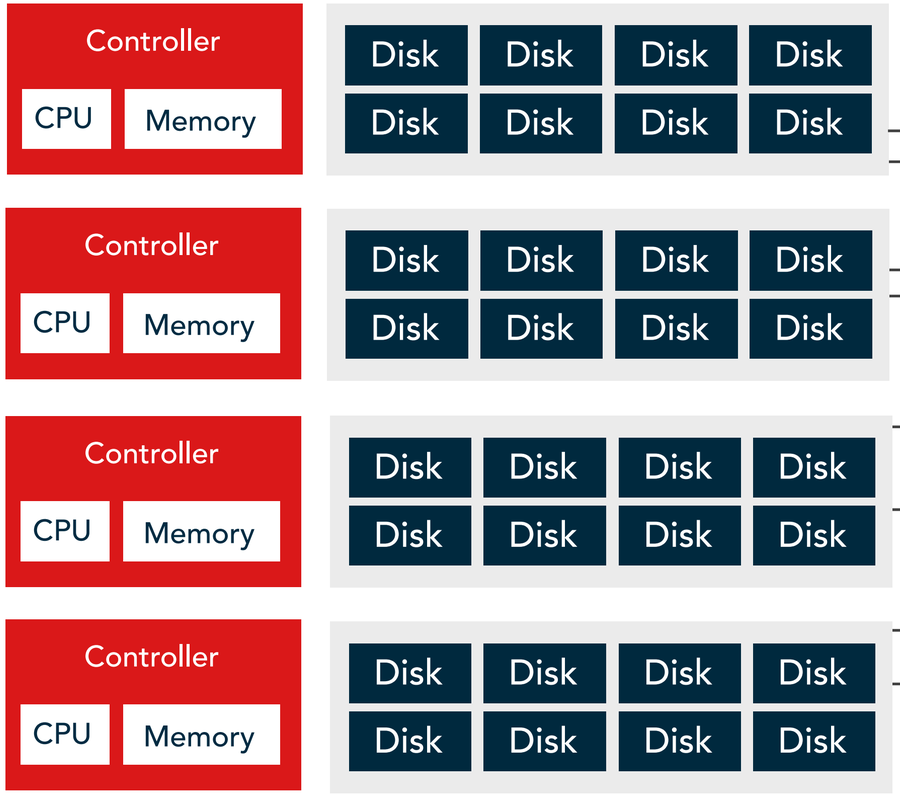
ARCHITECTURAL DIFFERENCES:

DISTRIBUTED VS. TRADITIONAL STORAGE

Traditional Centralized Storage with Scale-Up Design (TESS)



Scale-Out Approach in Distributed Storage



Static performance controller	Performance increases in proportion to nodes
Petabyte capacity with hundreds of disks	EB-Level storage with thousands of disks
Redundant controllers with RAID protection	Redundancy through erasure-coded storage
Locked-in hardware — expensive, limited agility	Standard hardware — affordable, fast to evolve

Sardina Systems Proprietary. Copyright (C) 2014 - 2026. All rights reserved.

FishOS® by Sardina Systems. Fully automated cloud solutions for enterprises based on OpenStack, Kubernetes, and Ceph.

3 KEY BENEFITS OF DISTRIBUTED STORAGE

1 Flexible Scaling and High Scalability

2 Concurrent Read/Write and High Performance

3 Standard Hardware and reduced CapEx and OpEx

Sardina Systems Proprietary. Copyright (C) 2014 - 2026. All rights reserved.

FishOS® by Sardina Systems. Fully automated cloud solutions for enterprises based on OpenStack, Kubernetes, and Ceph.

1 How Storage Architectures Have Evolved

2 Key Features and Capabilities of Distributed Storage

3 User Interface (UI) Capabilities and Examples

4 Building Storage at Scale: Speed, Stability, Simplicity, and Cost Efficiency

5 Ensuring Business Continuity with Robust Disaster Recovery

6 Proven Results: Customer Success Highlights



Sardina Systems Proprietary. Copyright (C) 2014 - 2026. All rights reserved.

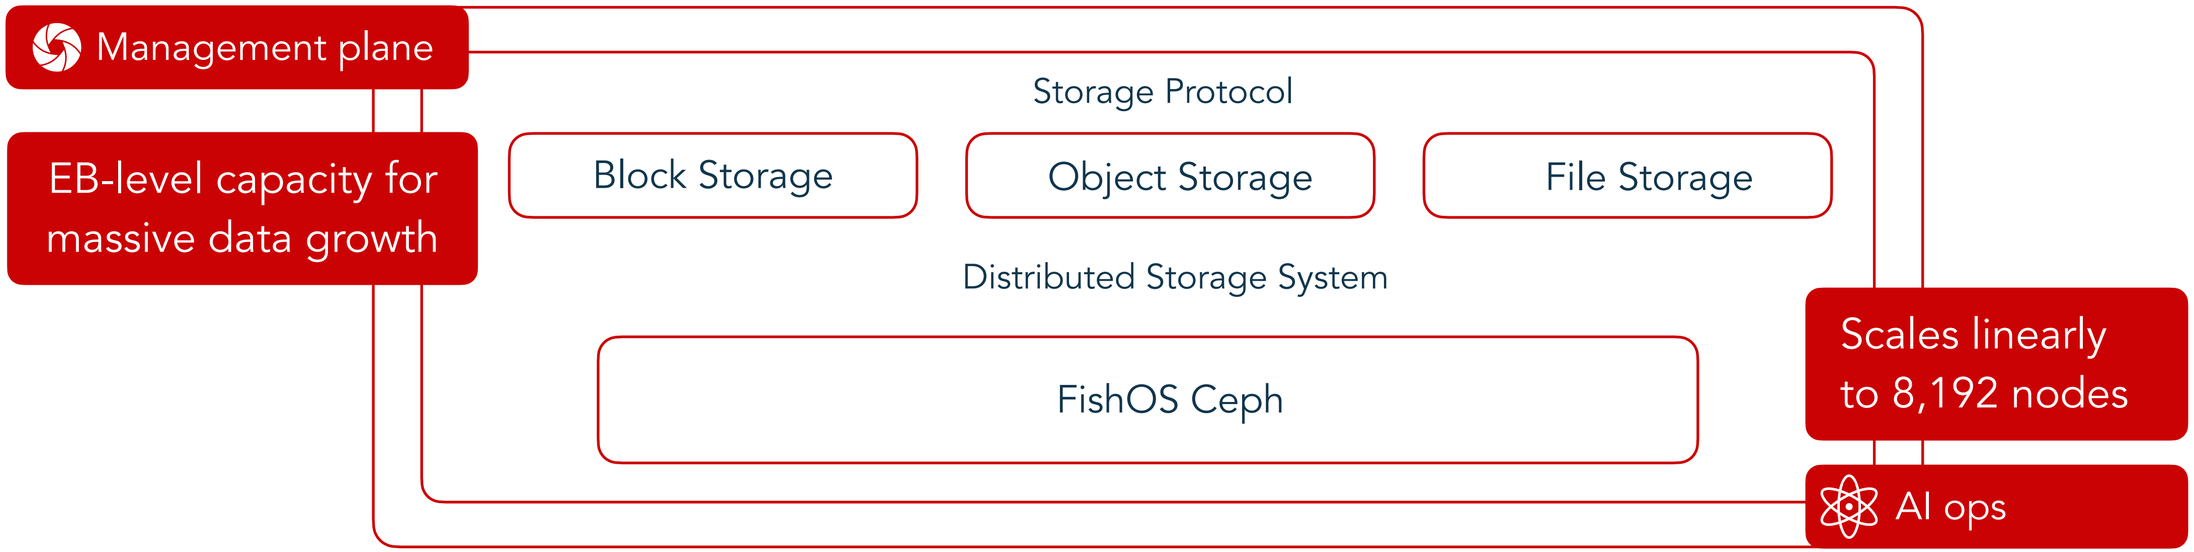
FishOS® by Sardina Systems. Fully automated cloud solutions for enterprises based on OpenStack, Kubernetes, and Ceph.

SARDINA SYSTEMS SDS: NEXT-GENERATION DISTRIBUTED STORAGE

One platform, two protocols, serving all industry scenarios with unified management, simplified operations, and seamless linear scalability.

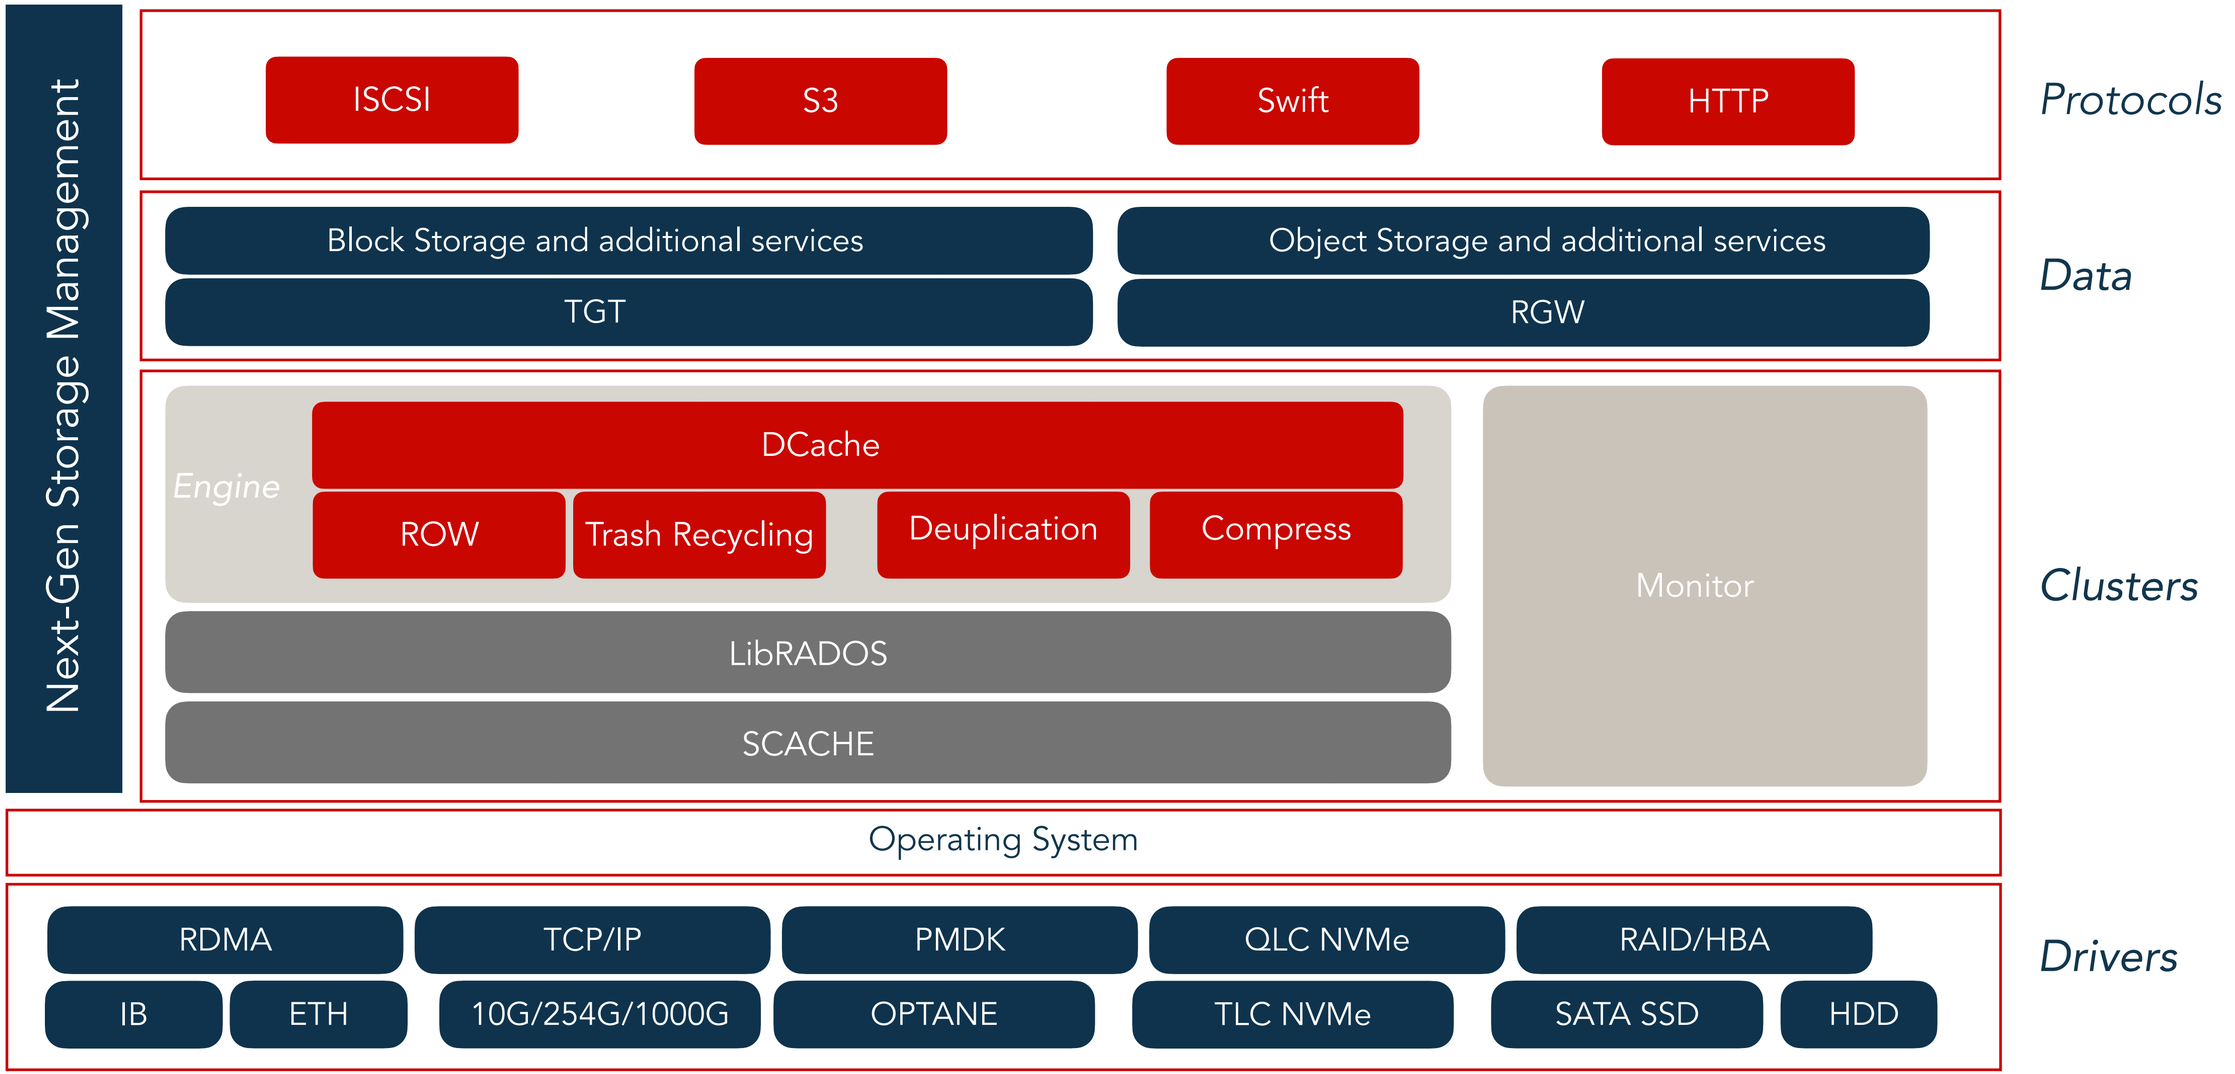
Virtualization & Containerization for:
Smart Healthcare, Optical Recognition, Digital Machine Platforms, Archive Storage, Video Surveillance, Smart School and more.

Massive-scale, high-speed, stable, easy-to-manage, and cost-effective — delivering reliable data storage for every industry.



Sardina Systems Proprietary. Copyright (C) 2014 - 2026. All rights reserved.
FishOS® by Sardina Systems. Fully automated cloud solutions for enterprises based on OpenStack, Kubernetes, and Ceph.

FISHOS DISTRIBUTED STORAGE ARCHITECTURE



Sardina Systems Proprietary. Copyright (C) 2014 - 2026. All rights reserved.

FishOS® by Sardina Systems. Fully automated cloud solutions for enterprises based on OpenStack, Kubernetes, and Ceph.

NETWORK OVERVIEW

Service Network: Handles client data write operations.

Storage Network: Facilitates data exchange between storage nodes.

Management Network: Transmits data related to system management, maintenance, and monitoring.

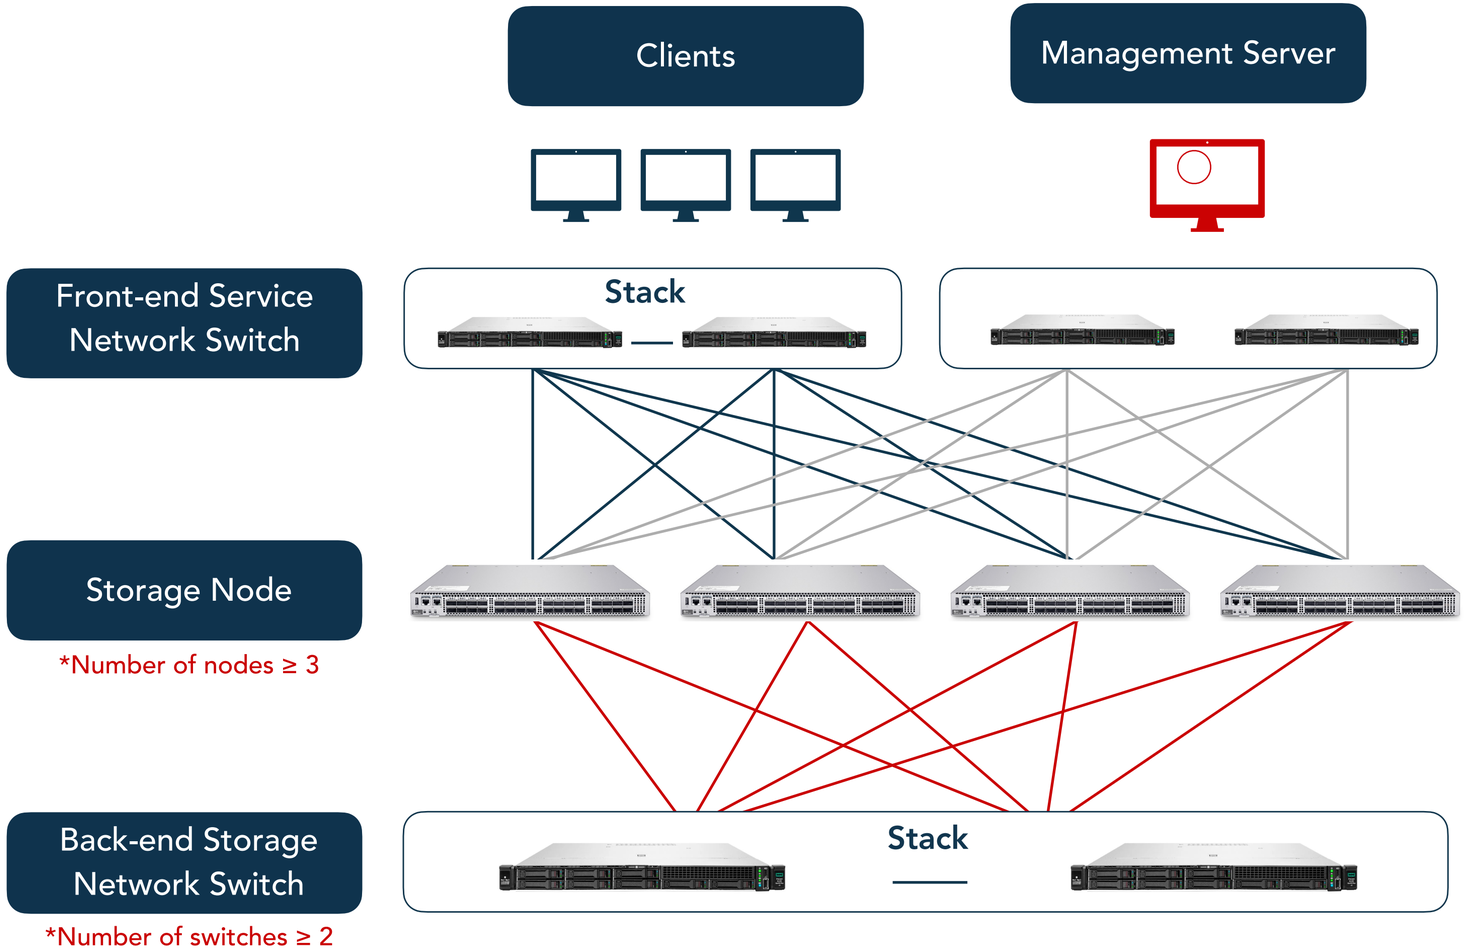
Key Capabilities:

- Compatible with 10GE, 25GE, and 100GE Ethernet adapters.
- Supports RoCE v2 (RDMA over Converged Ethernet version 2) for efficient, low-latency communication.

Recommended Configuration:

- For both the service and storage networks, it is best practice to deploy a dual-port Bond4 (802.3ad LACP) setup to ensure network redundancy and improve fault tolerance.

FISHOS DISTRIBUTED STORAGE NETWORK ARCHITECTURE



Sardina Systems Proprietary. Copyright (C) 2014 - 2026. All rights reserved.

1 How Storage Architectures Have Evolved

2 Key Features and Capabilities of Distributed Storage

3 User Interface (UI) Capabilities and Examples

4 Building Storage at Scale: Speed, Stability, Simplicity, and Cost Efficiency

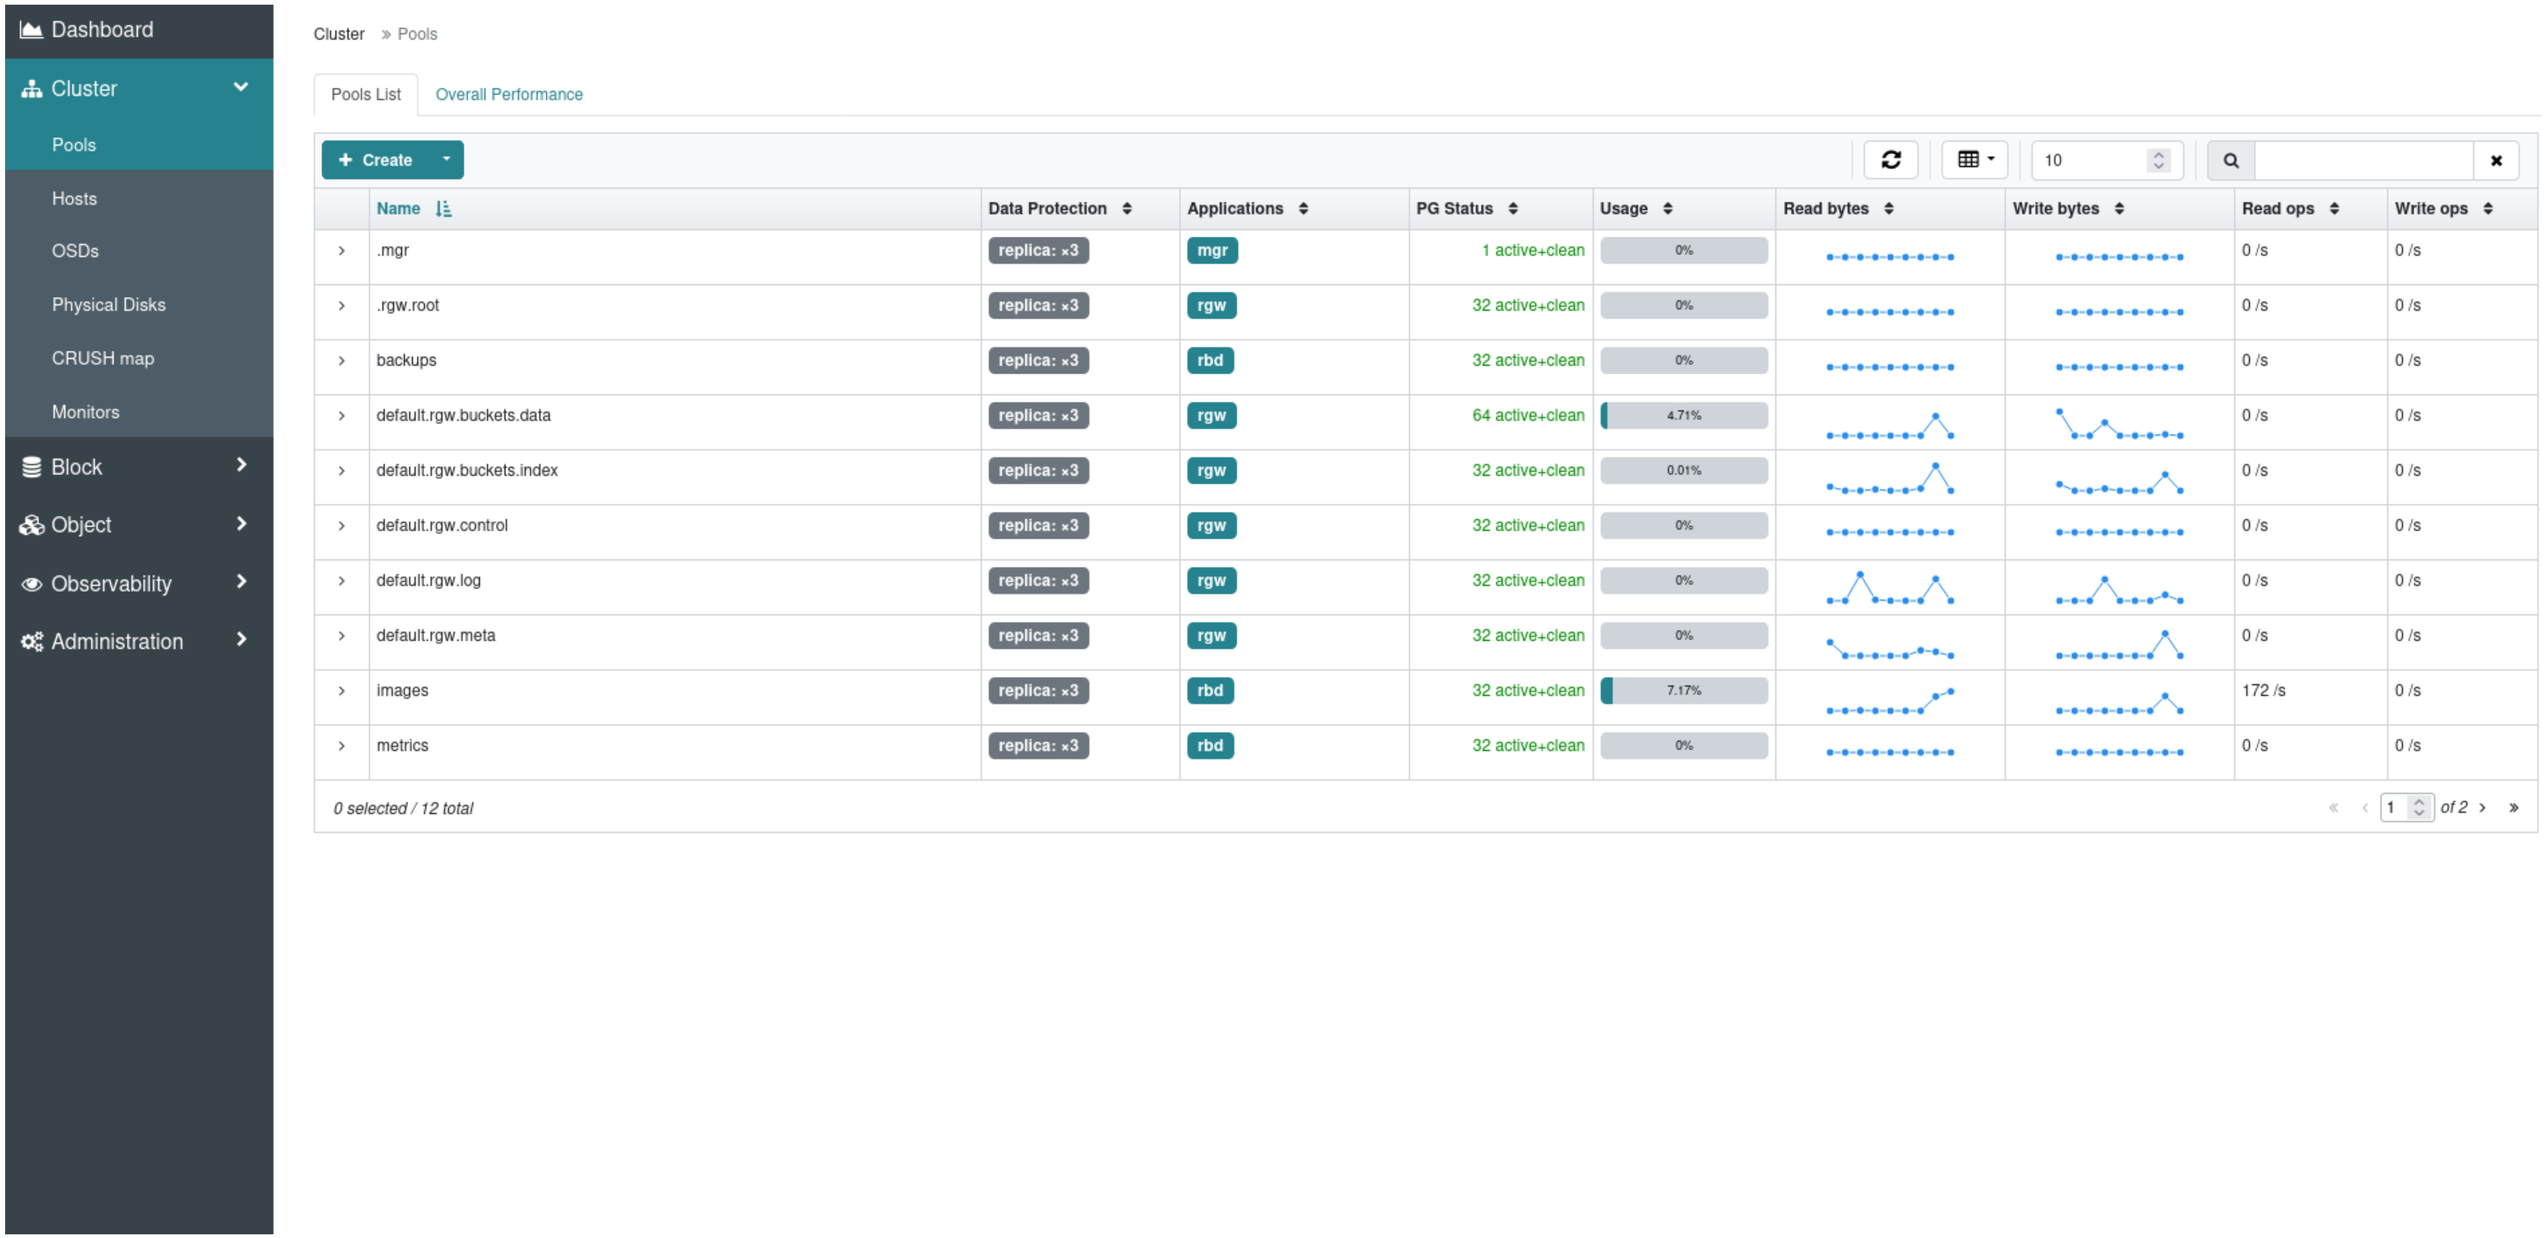
5 Ensuring Business Continuity with Robust Disaster Recovery

6 Proven Results: Customer Success Highlights



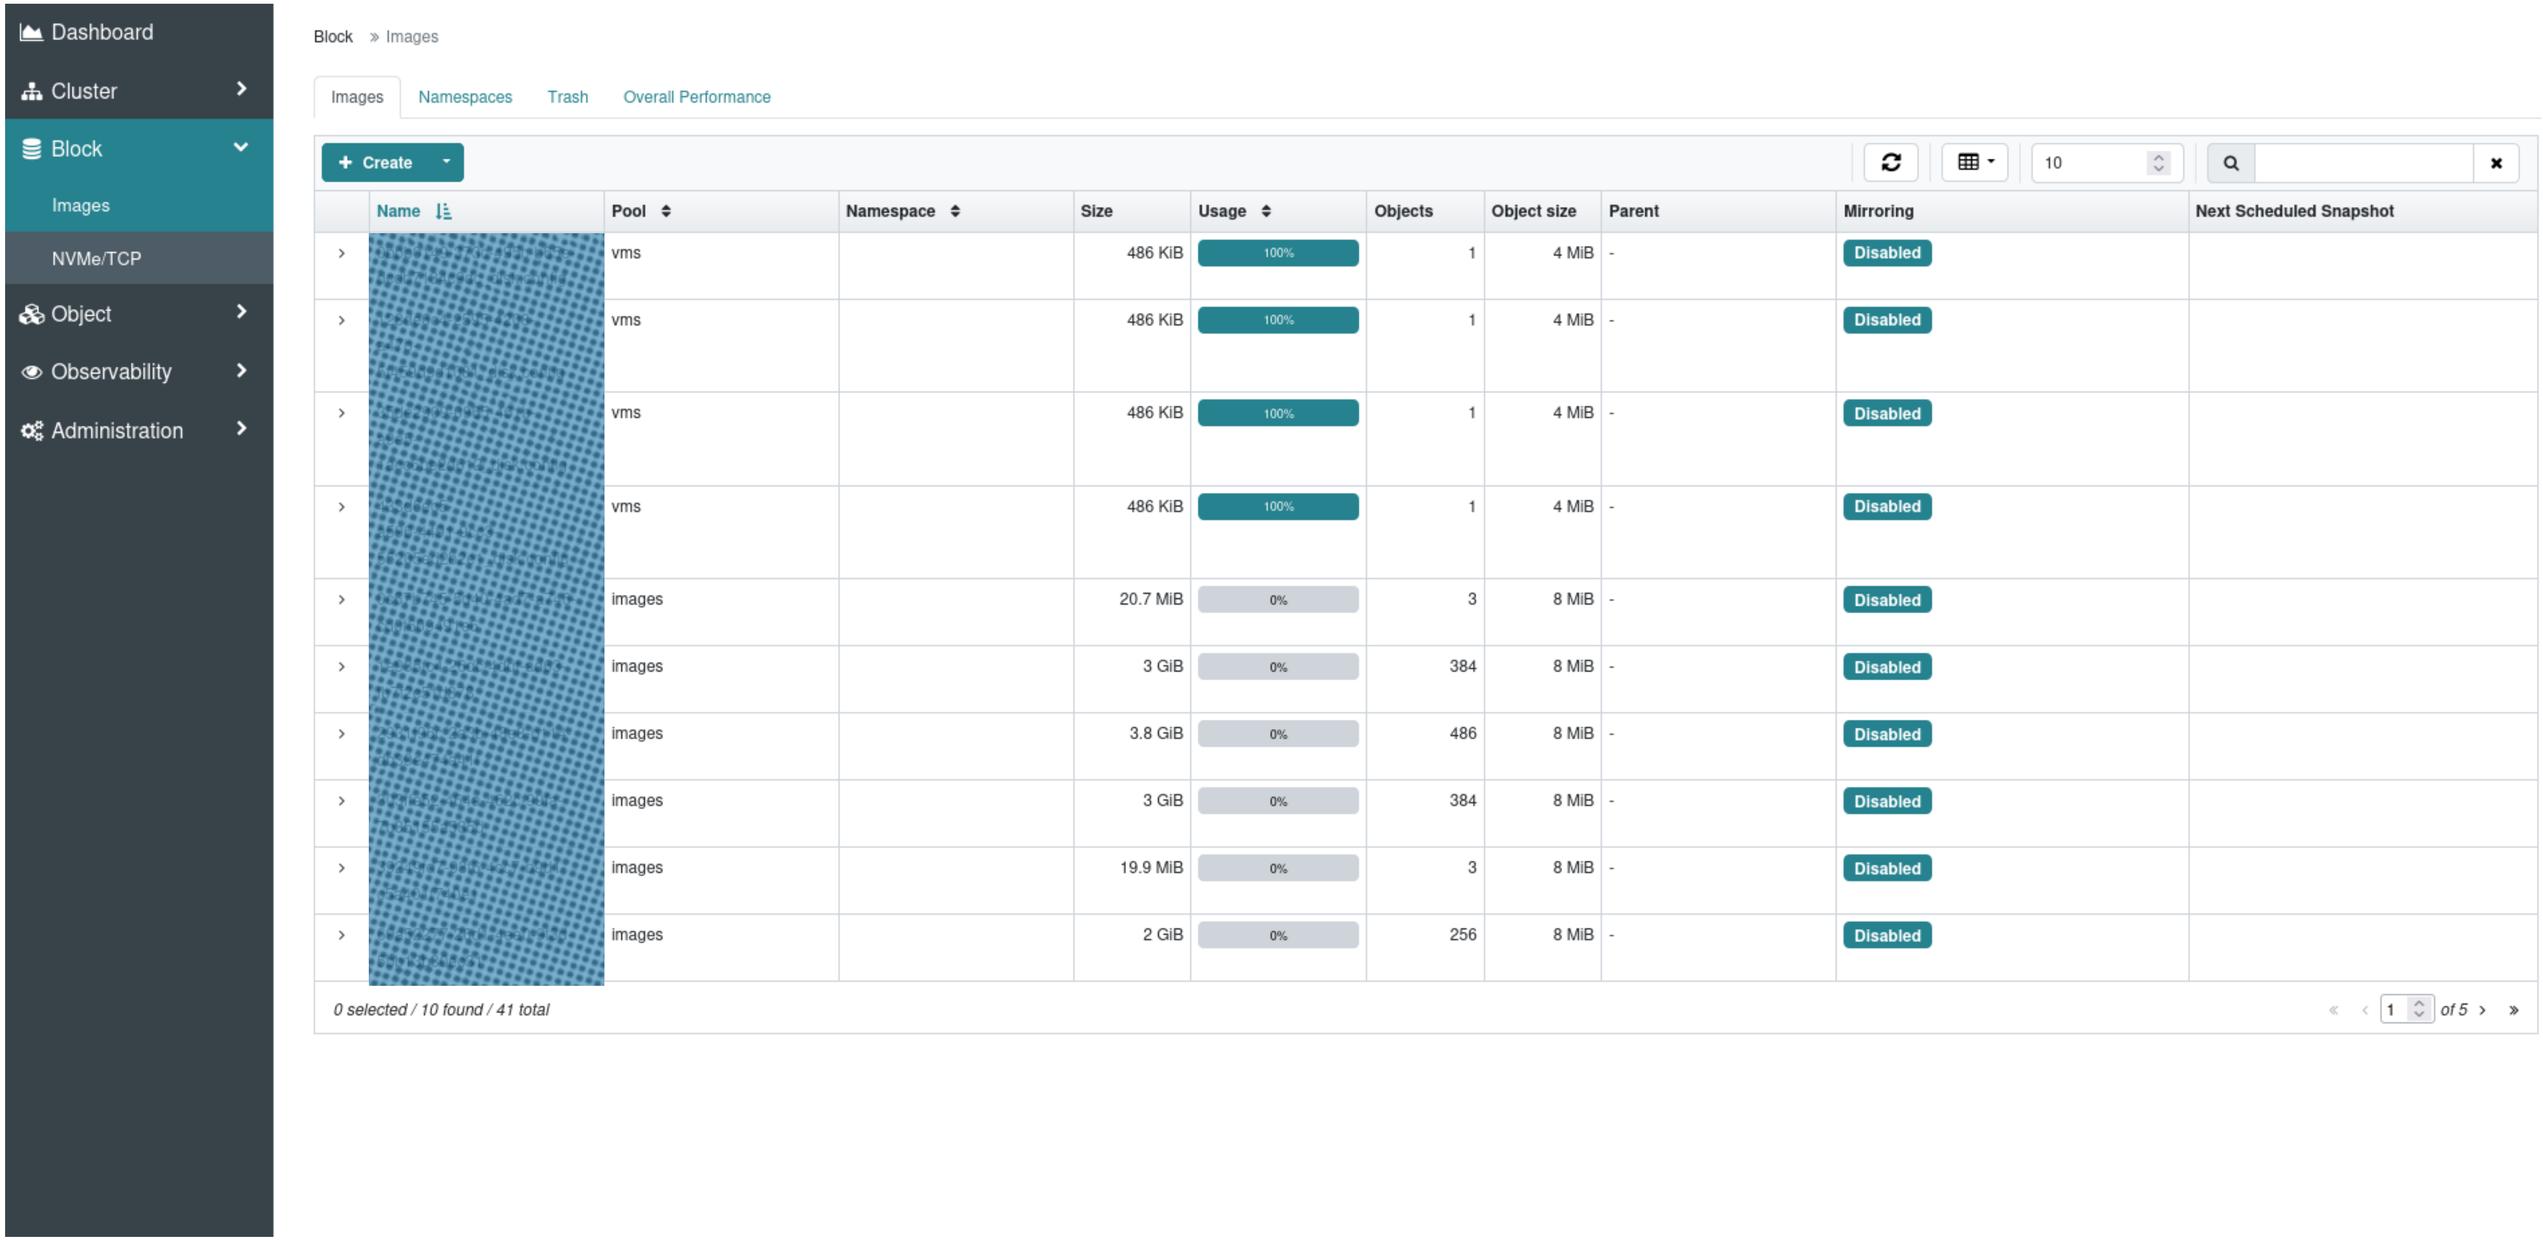
Sardina Systems Proprietary. Copyright (C) 2014 - 2026. All rights reserved.

FISHOS CEPH UI



Sardina Systems Proprietary. Copyright (C) 2014 - 2026. All rights reserved.
FishOS® by Sardina Systems. Fully automated cloud solutions for enterprises based on OpenStack, Kubernetes, and Ceph.

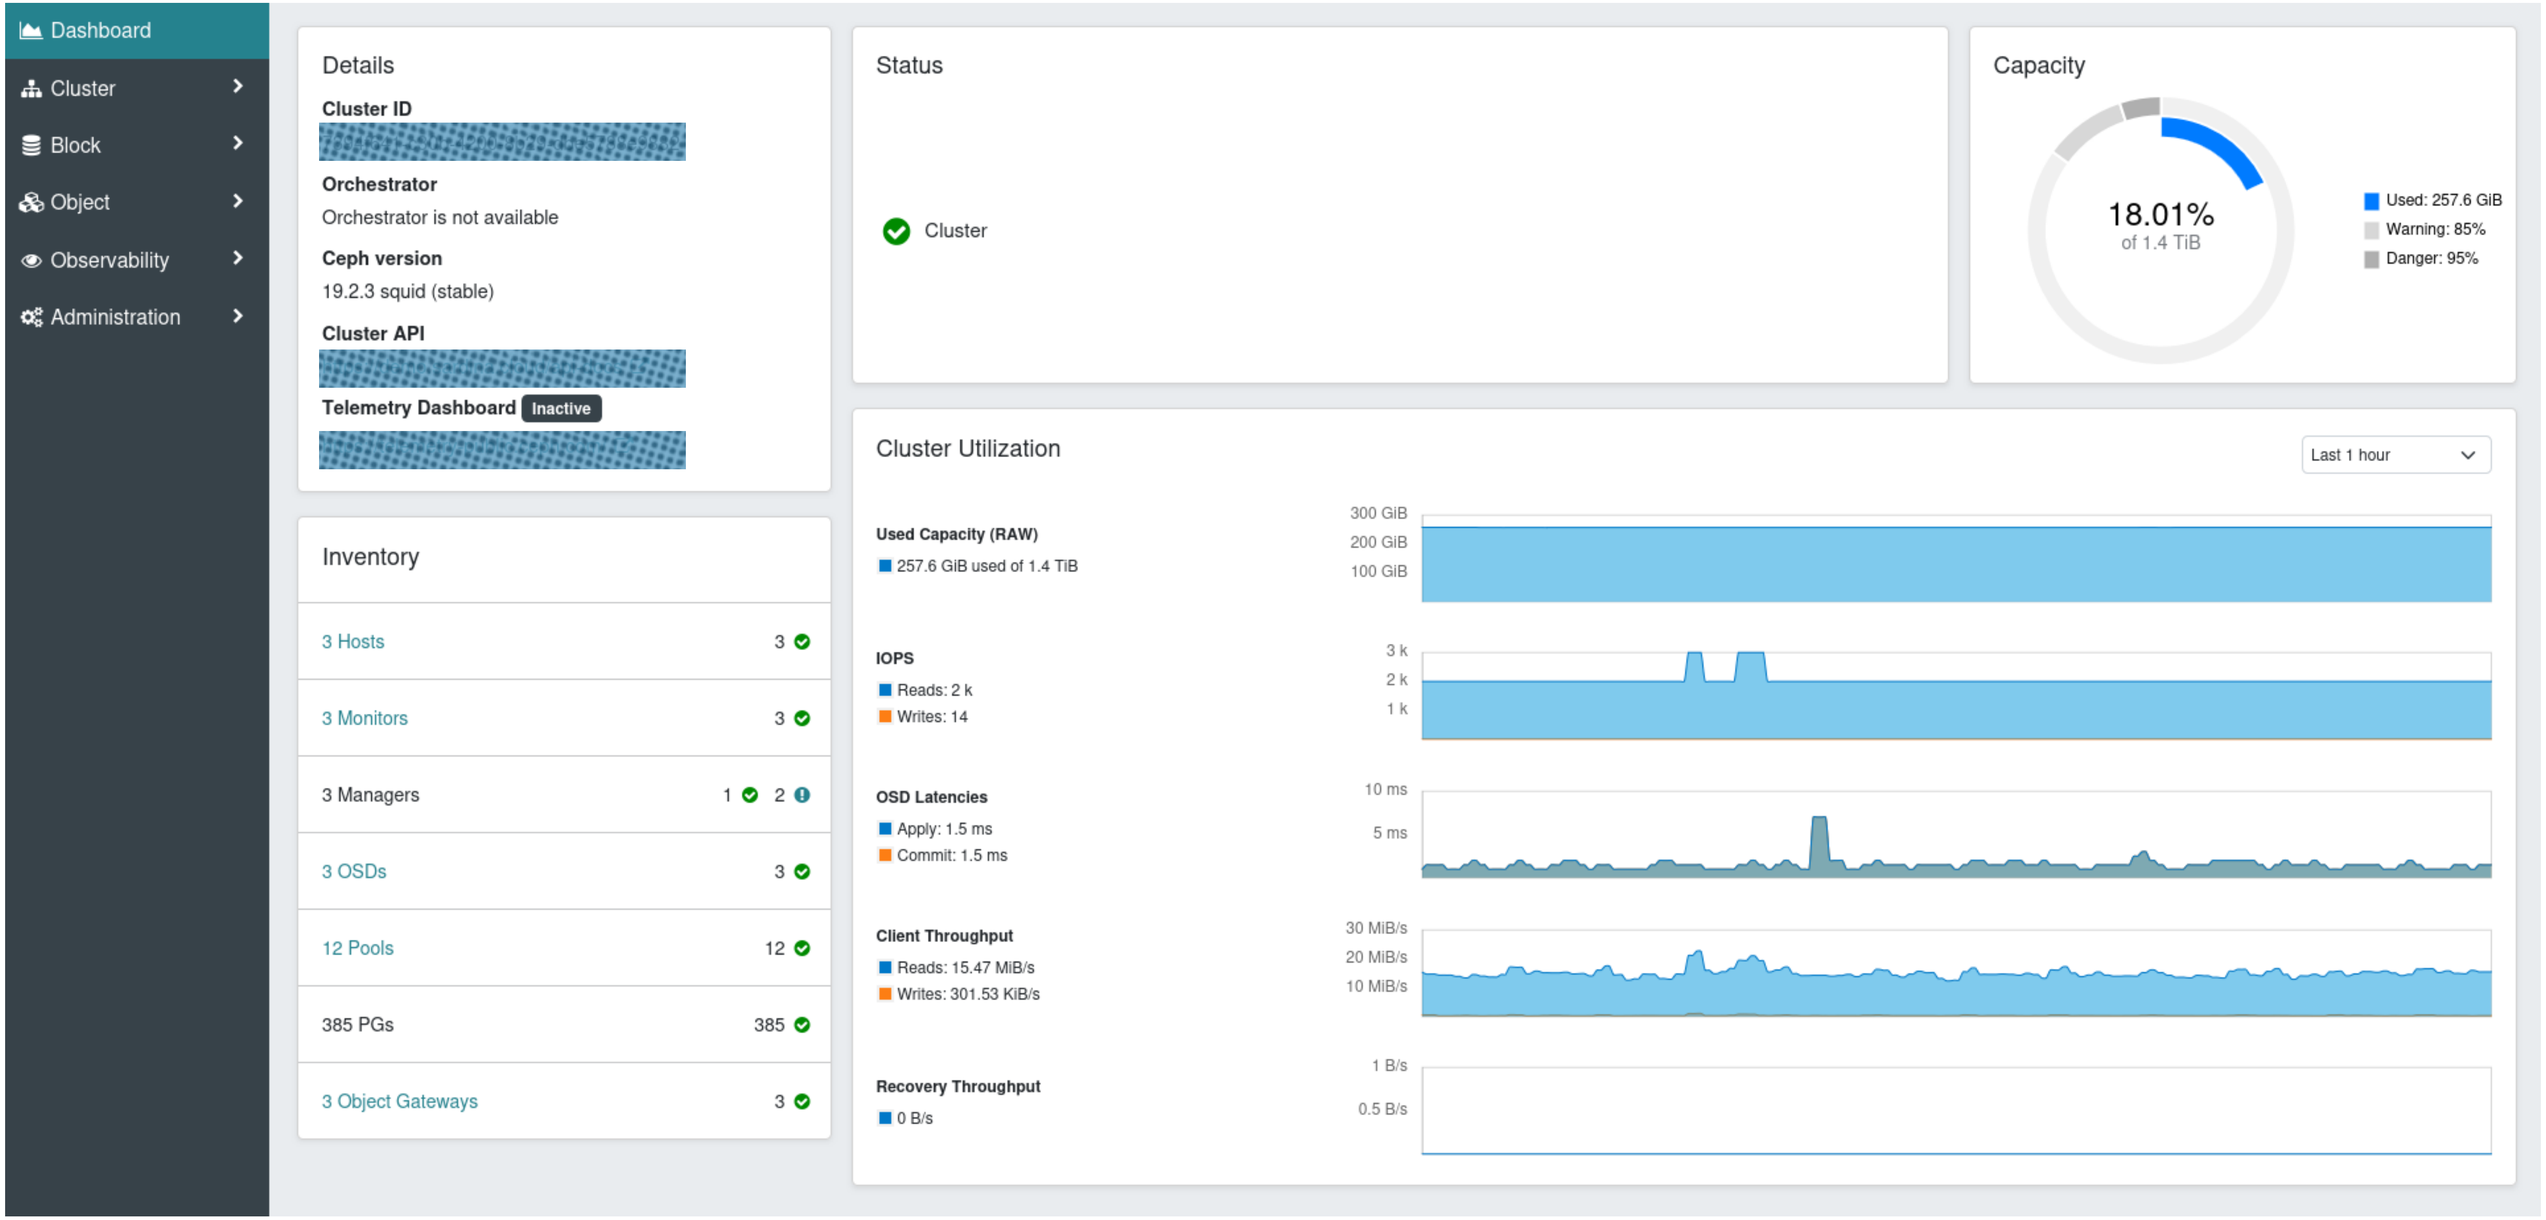
FISHOS CEPH UI



Sardina Systems Proprietary. Copyright (C) 2014 - 2026. All rights reserved.

FishOS® by Sardina Systems. Fully automated cloud solutions for enterprises based on OpenStack, Kubernetes, and Ceph.

FISHOS CEPH UI



Sardina Systems Proprietary. Copyright (C) 2014 - 2026. All rights reserved.
FishOS® by Sardina Systems. Fully automated cloud solutions for enterprises based on OpenStack, Kubernetes, and Ceph.

1 How Storage Architectures Have Evolved

2 Key Features and Capabilities of Distributed Storage

3 User Interface (UI) Capabilities and Examples

4 Building Storage at Scale: Speed, Stability, Simplicity, and Cost Efficiency

5 Ensuring Business Continuity with Robust Disaster Recovery

6 Proven Results: Customer Success Highlights



Sardina Systems Proprietary. Copyright (C) 2014 - 2026. All rights reserved.

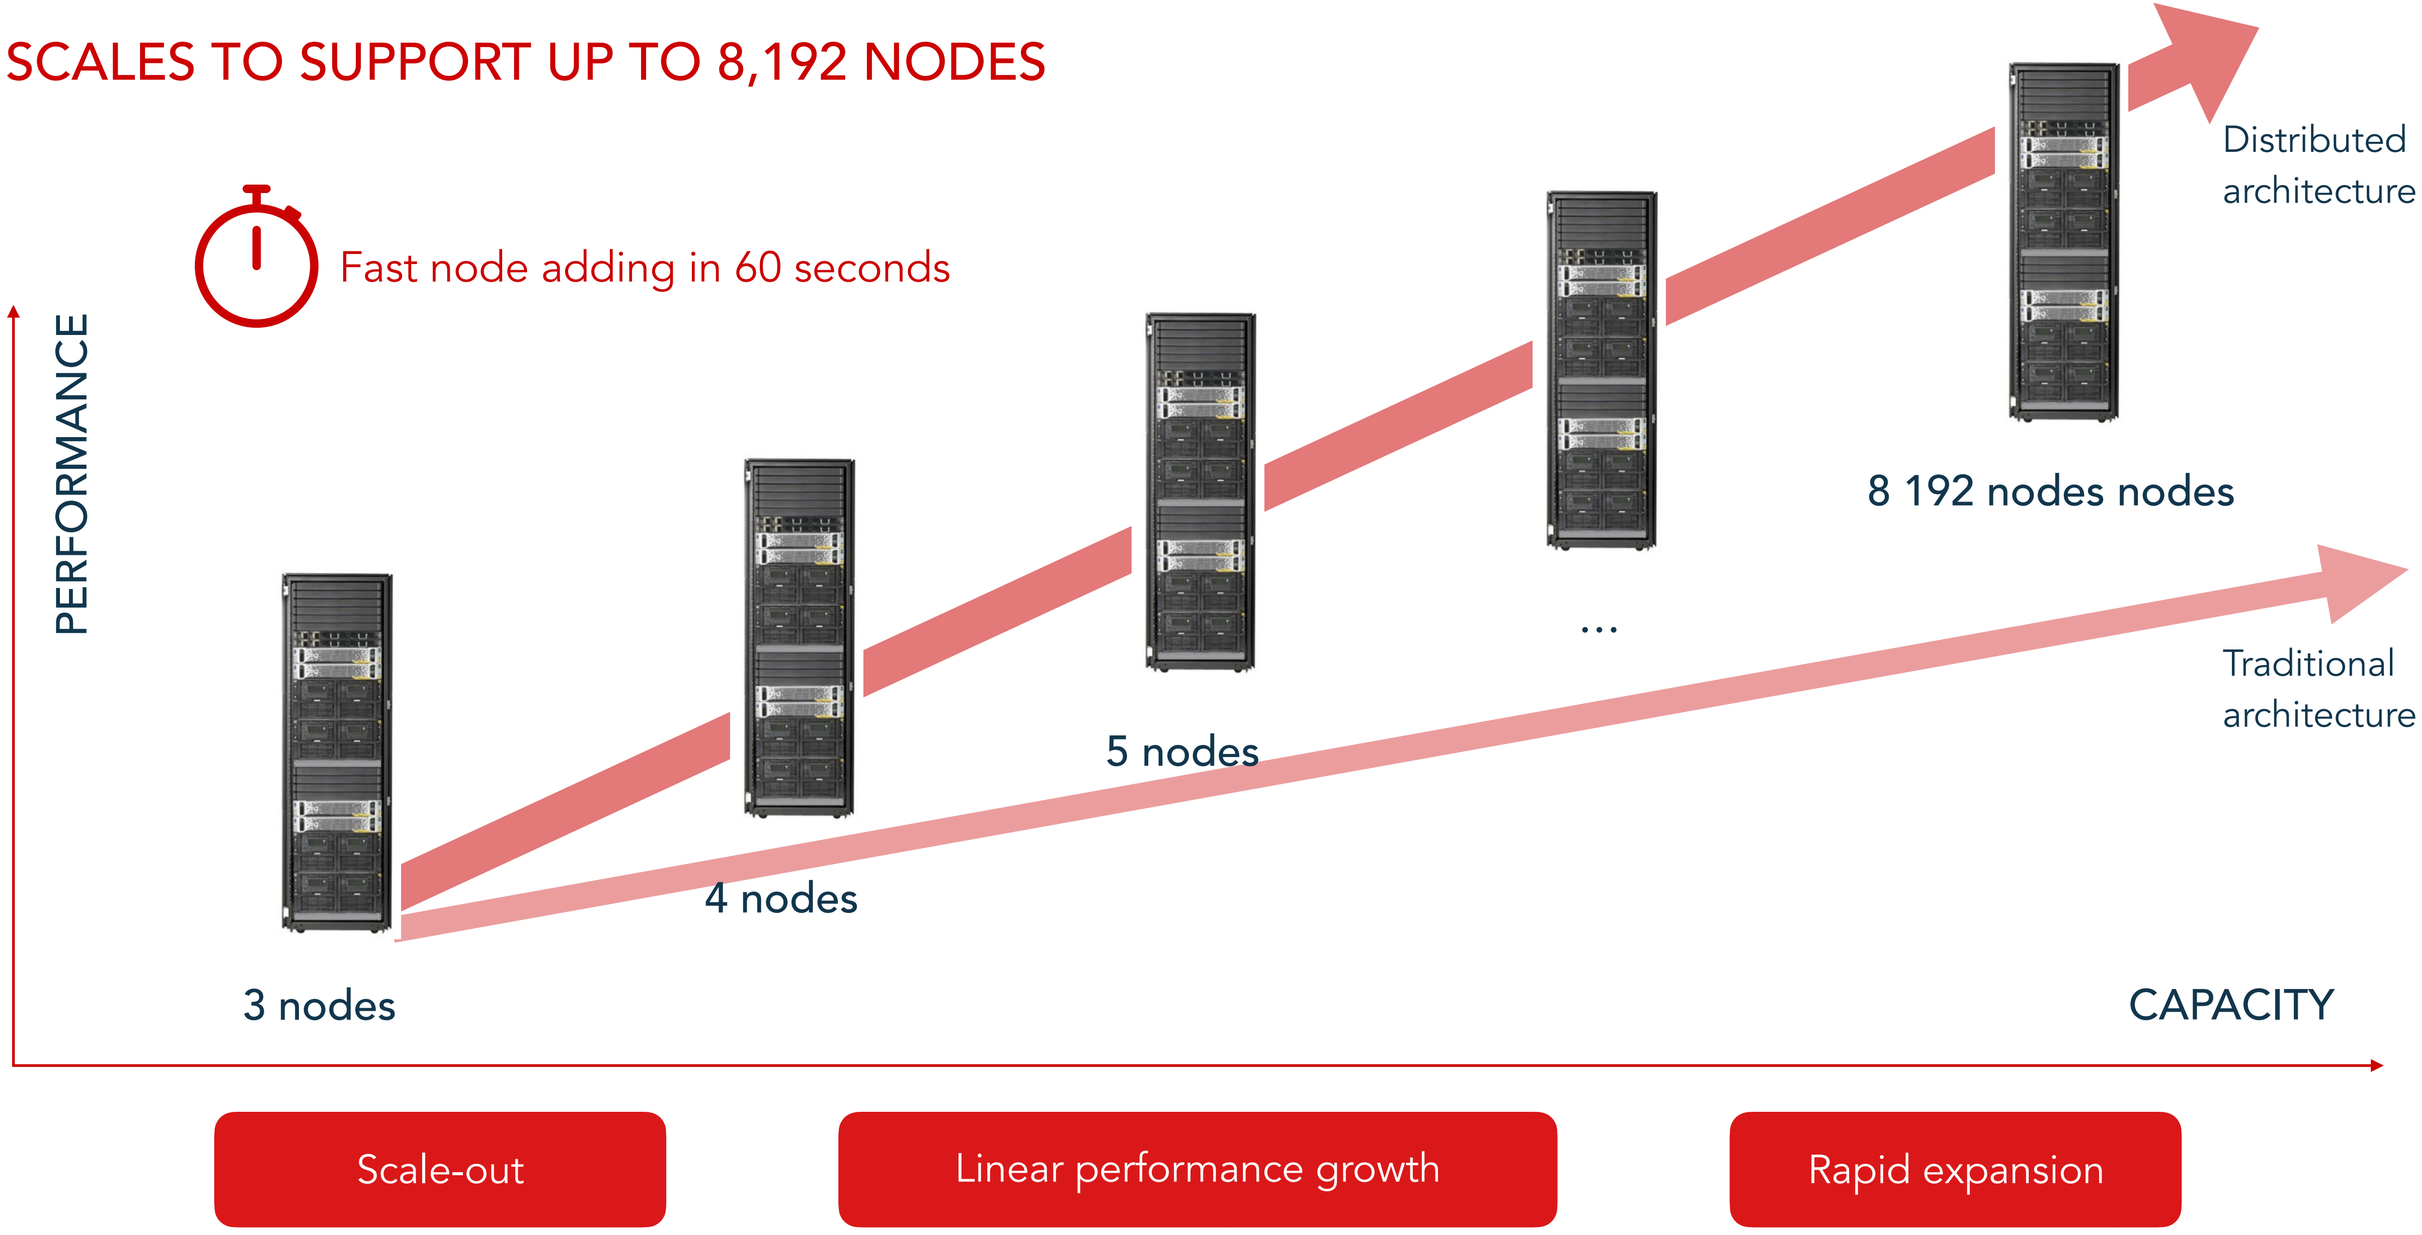
FishOS® by Sardina Systems. Fully automated cloud solutions for enterprises based on OpenStack, Kubernetes, and Ceph.



1. Massive Scale

Sardina Systems Proprietary. Copyright (C) 2014 - 2026. All rights reserved.

SCALES TO SUPPORT UP TO 8,192 NODES



Scale-out

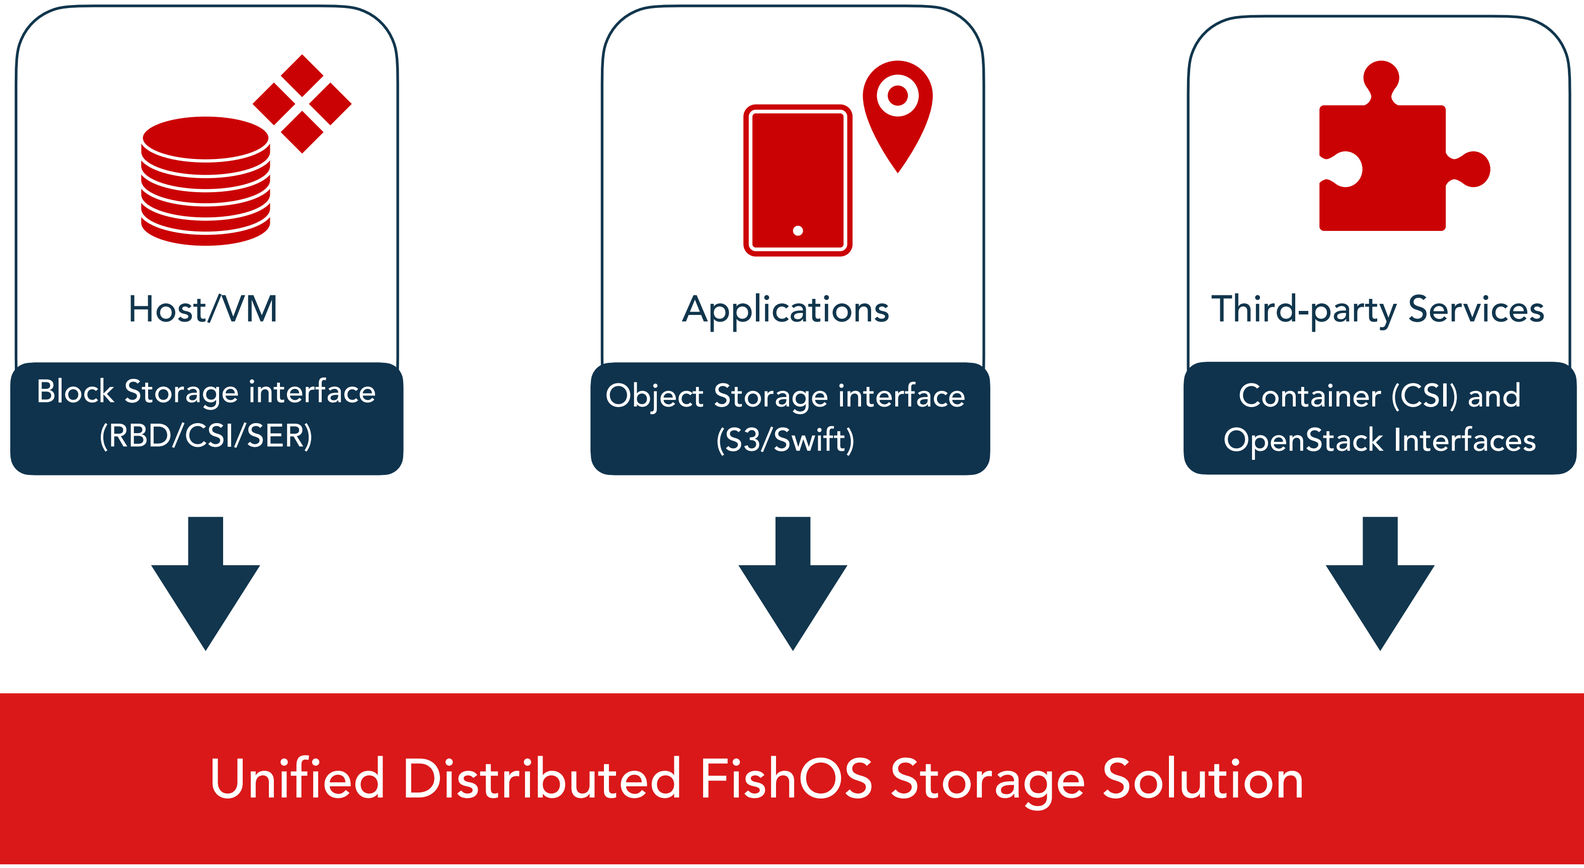
Linear performance growth

Rapid expansion

Sardina Systems Proprietary. Copyright (C) 2014 - 2026. All rights reserved.

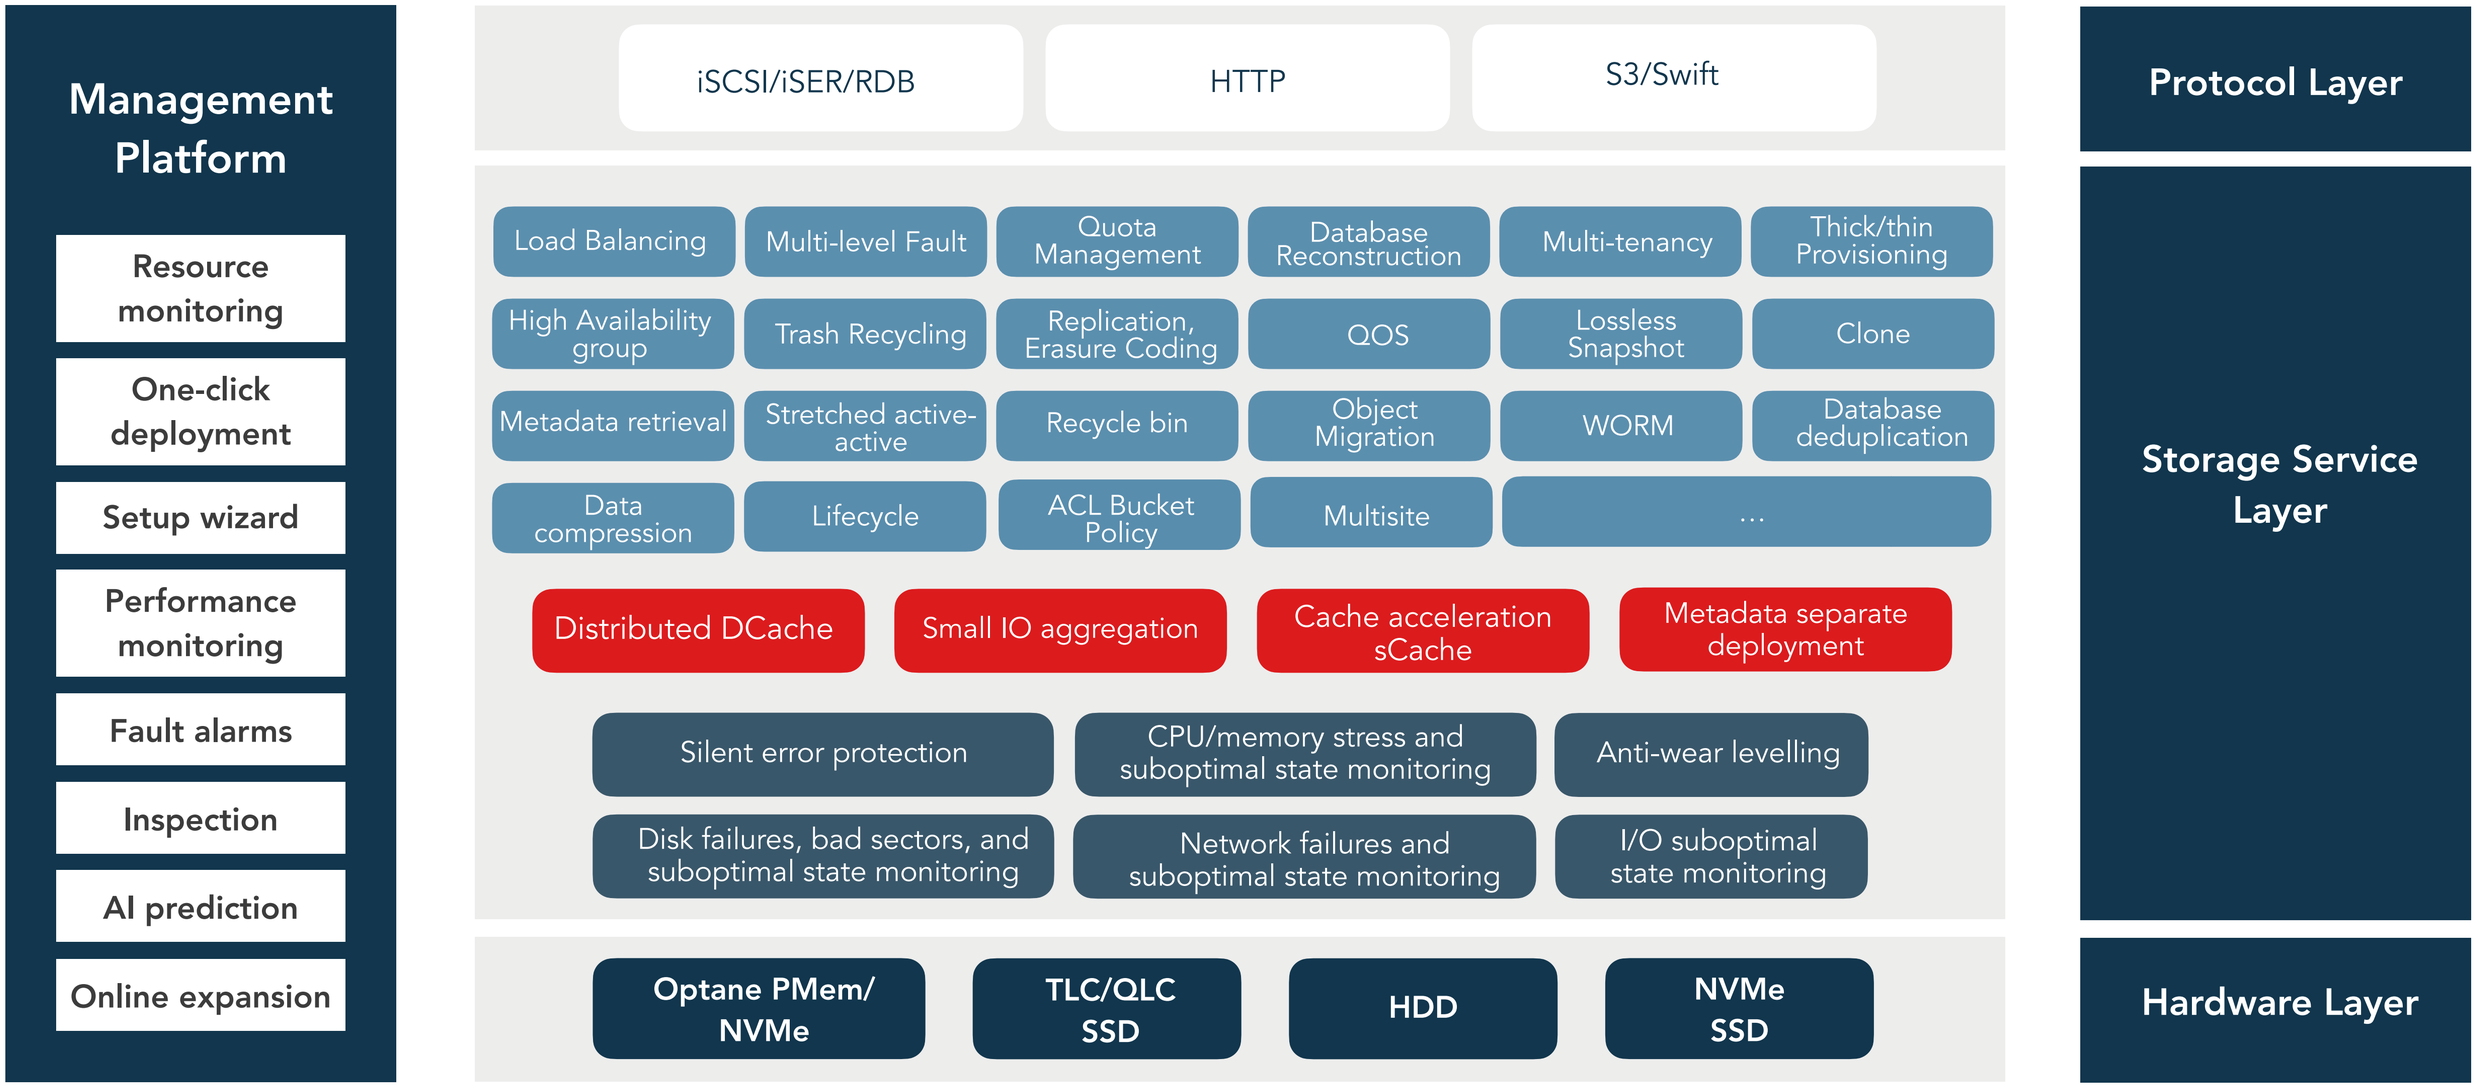
FishOS® by Sardina Systems. Fully automated cloud solutions for enterprises based on OpenStack, Kubernetes, and Ceph.

WIDE RANGE OF SUPPORTED SERVICES



Sardina Systems Proprietary. Copyright (C) 2014 - 2026. All rights reserved.
FishOS® by Sardina Systems. Fully automated cloud solutions for enterprises based on OpenStack, Kubernetes, and Ceph.

RICH ADVANCED SOFTWARE FEATURES



Sardina Systems Proprietary. Copyright (C) 2014 - 2026. All rights reserved.

FishOS® by Sardina Systems. Fully automated cloud solutions for enterprises based on OpenStack, Kubernetes, and Ceph.

BROAD COMPATIBILITY WITH MAJOR SOFTWARE AND HARDWARE PLATFORMS

Operations Platform

Via universal interface support



Virtualization

Via universal interface support



Database

Full compatibility with relational and non-relational databases



Cloud platform

Cinder and CSI integration for seamless cloud storage access



Optimized for diverse workloads with support for CPUs, GPUs, and FPGAs

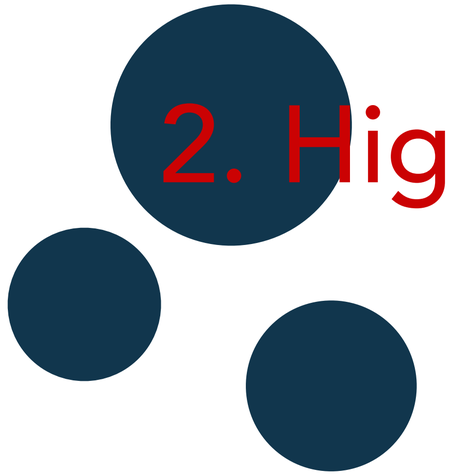


Supports a wide range of storage controllers, network adapters, disks, and QAT acceleration.



Sardina Systems Proprietary. Copyright (C) 2014 - 2026. All rights reserved.

FishOS® by Sardina Systems. Fully automated cloud solutions for enterprises based on OpenStack, Kubernetes, and Ceph.

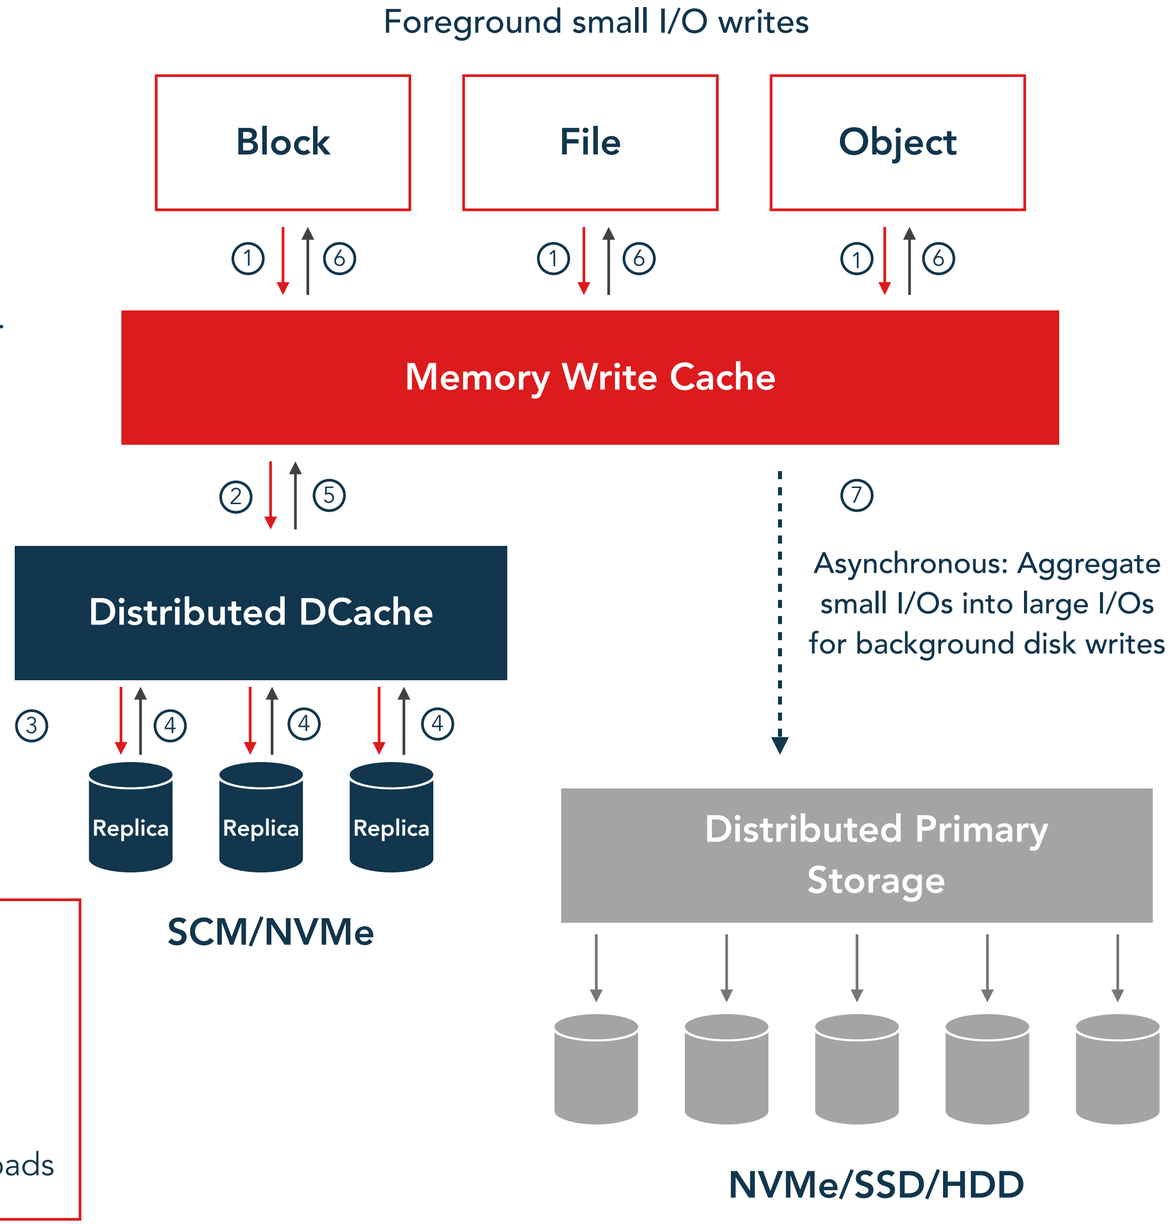


2. High Performance

Sardina Systems Proprietary. Copyright (C) 2014 - 2026. All rights reserved.

ARCHITECTURAL-LEVEL DISTRIBUTED CACHE SUPPORT

- **Immediate Dcache Acknowledgment**: Small I/O writes return instantly after being written to distributed cache, ensuring the **shortest I/O path** and **faster response times**.
- **Small I/O Aggregation (ROW)**: Utilizes **ROW (Read-Optimized Write)** small I/O aggregation technology to merge scattered writes into system-aligned large I/Os. This minimizes the overhead from erasure coding (e.g., zero-padding due to partial stripe writes) and boosts write efficiency.
- **Large I/O Direct Passthrough**: High-volume I/Os bypass the cache and are written directly to disk, which enhances throughput and extends disk lifespan by reducing write amplification.



Customer Value:

- Accelerated application response with **immediate small I/O commit**
- Improved efficiency through **smart I/O path optimization**
- Longer disk life and **better performance consistency** for both small and large I/O workloads

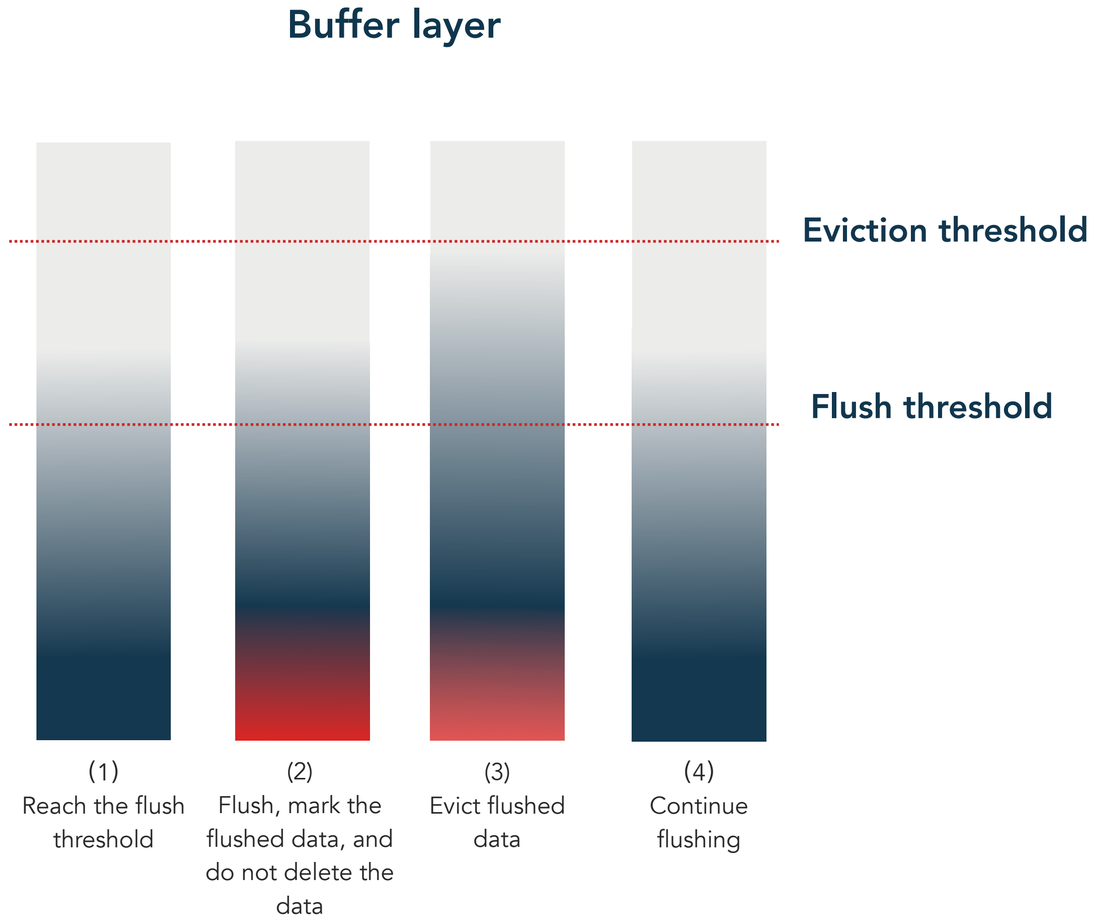
Sardina Systems Proprietary. Copyright (C) 2014 - 2026. All rights reserved.
FishOS® by Sardina Systems. Fully automated cloud solutions for enterprises based on OpenStack, Kubernetes, and Ceph.

SCACHE INTELLIGENT TIERED FLUSHING

Implements a smart data flushing mechanism across storage tiers to optimize performance and resource usage.

- When the **flush threshold** is reached, data is offloaded to a **lower storage tier**.
- Upon hitting the **eviction threshold**, previously flushed data is **evicted** to free up space.
- Based on **access frequency (heat)**, frequently used data can be **promoted** to a **higher-speed tier**, ensuring fast access to hot data.

This dynamic approach balances performance and capacity, while maintaining responsiveness across mixed workloads.



Sardina Systems Proprietary. Copyright (C) 2014 - 2026. All rights reserved.

FAST READS WITH INTELLIGENT PREFETCHING

Boosts sequential read performance by proactively loading data into memory based on intelligent access pattern detection.

Sequential Stream Detection: Identifies sequential I/O patterns and prefetches data from primary storage into the read cache.

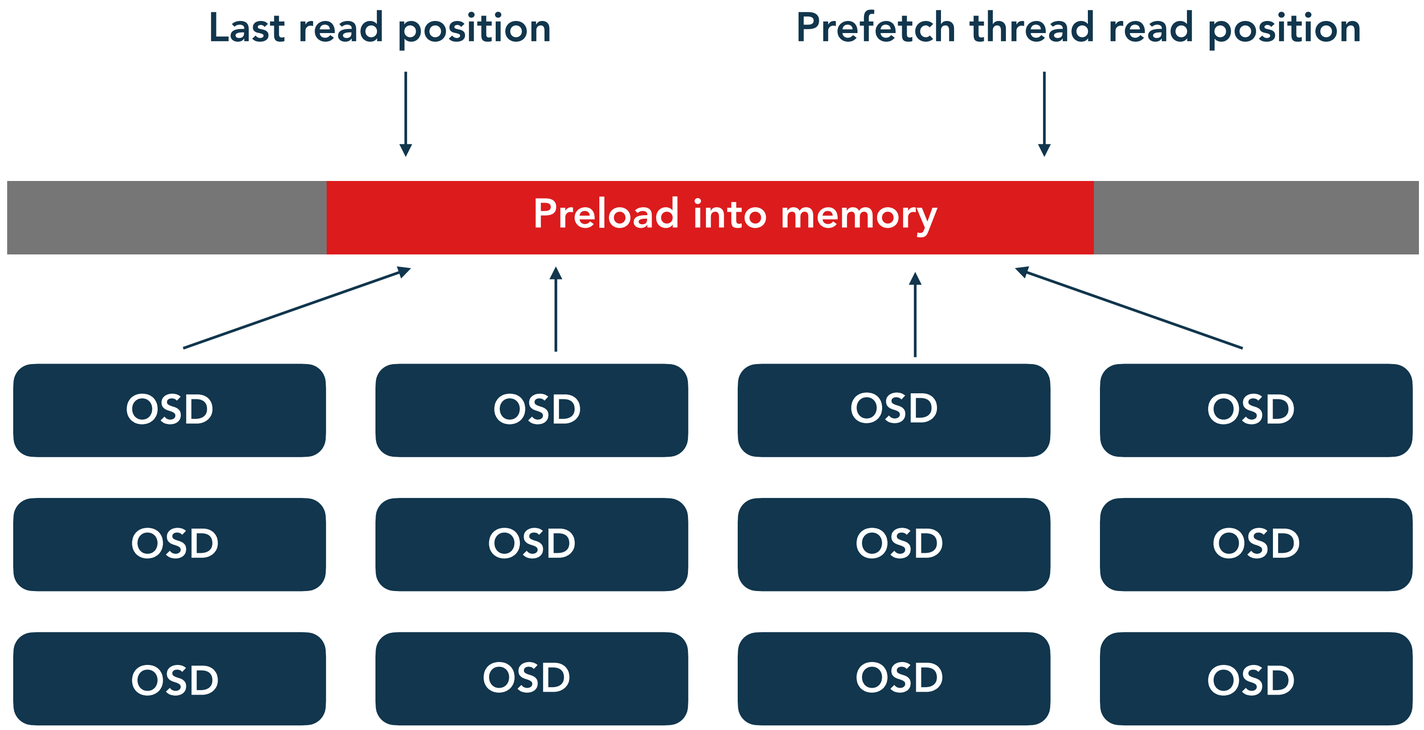
Dynamic Identification Strategy: Prefetch behaviour adapts to real-time service characteristics, optimizing detection accuracy.

Prefetch Scope Optimization:

- **Block size** is adjusted based on data access heat.
- **File prefetch size** scales as a power-of-two multiple of the I/O size, maximizing cache efficiency.

Key Benefit:

Significantly accelerates read operations for sequential workloads while maintaining intelligent resource utilization.



Sardina Systems Proprietary. Copyright (C) 2014 - 2026. All rights reserved.

PROTOCOL ACCELERATION: RDMA-ACCELERATED BLOCK STORAGE

Reduced CPU Usage:

Leverages **kernel bypass** to enable direct data transfer in user space, eliminating costly context switches between user and kernel modes.

Zero-Copy Transmission:

Enables **direct memory access** from the local system to remote storage without intermediate data copies, significantly reducing CPU cycles.

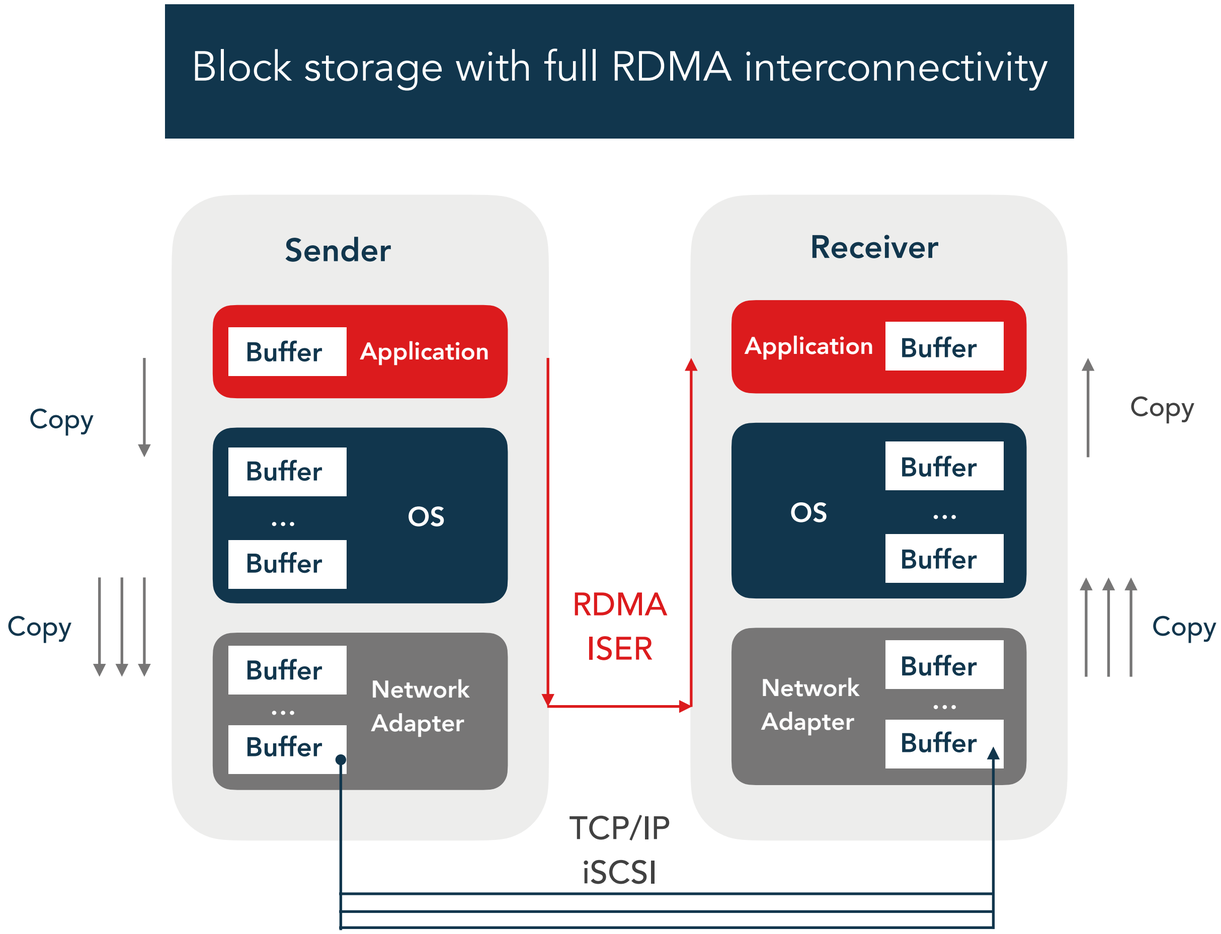
High Bandwidth & Low Latency:

- Supports **multi-buffer aggregation** and **stream splitting** for optimized throughput.
- Enhances iSCSI performance via **iSER (iSCSI Extensions for RDMA)**, accelerating both control and data paths.

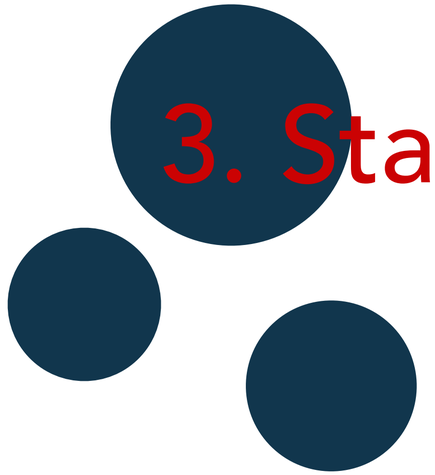
Customer Value:

Lower CPU load, freeing resources for application workloads

Improved data transfer efficiency and **faster response times** in high-performance environments



Sardina Systems Proprietary. Copyright (C) 2014 - 2026. All rights reserved.



3. Stability

Sardina Systems Proprietary. Copyright (C) 2014 - 2026. All rights reserved.

REDUNDANCY MECHANISM

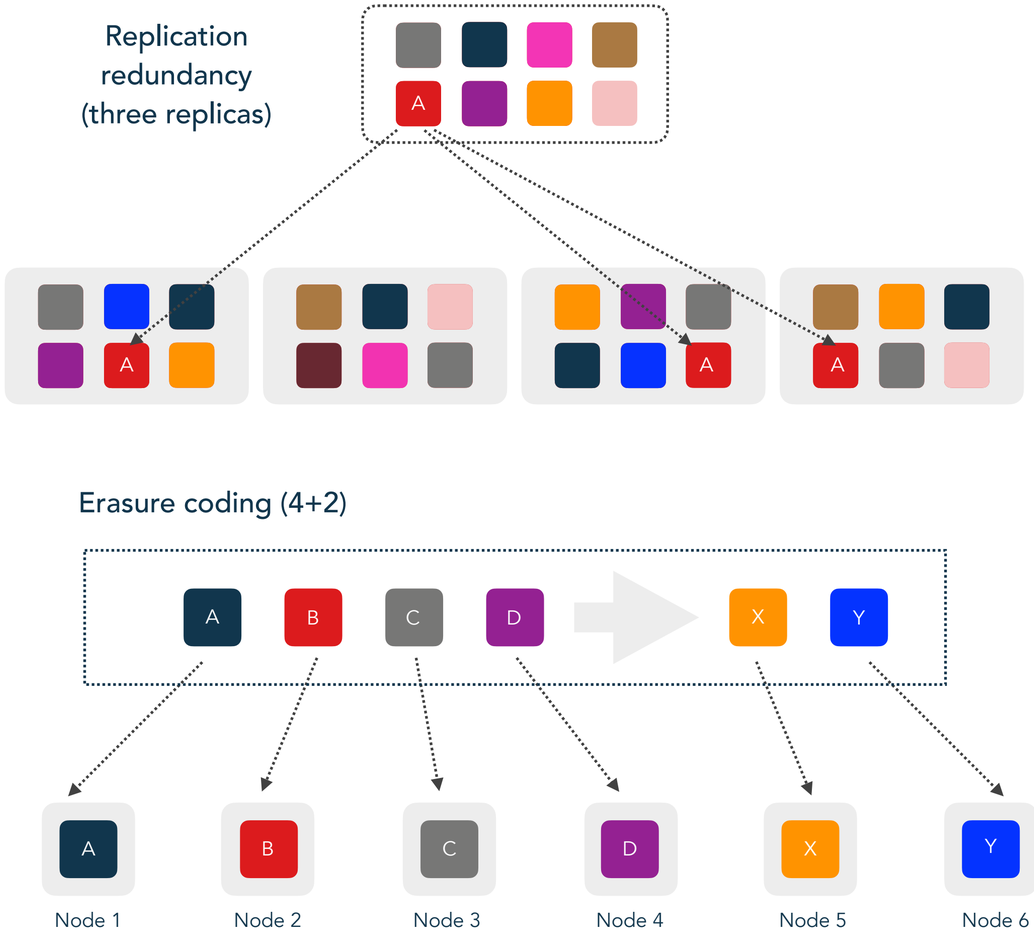
Implemented through replication and erasure coding for robust data protection.

Replication Support:

- Supports 2 to 8 replicas for configurable data redundancy levels
- Achieves up to 11 nines (99.999999999%) data reliability
- Enables rapid recovery and high availability in failure scenarios

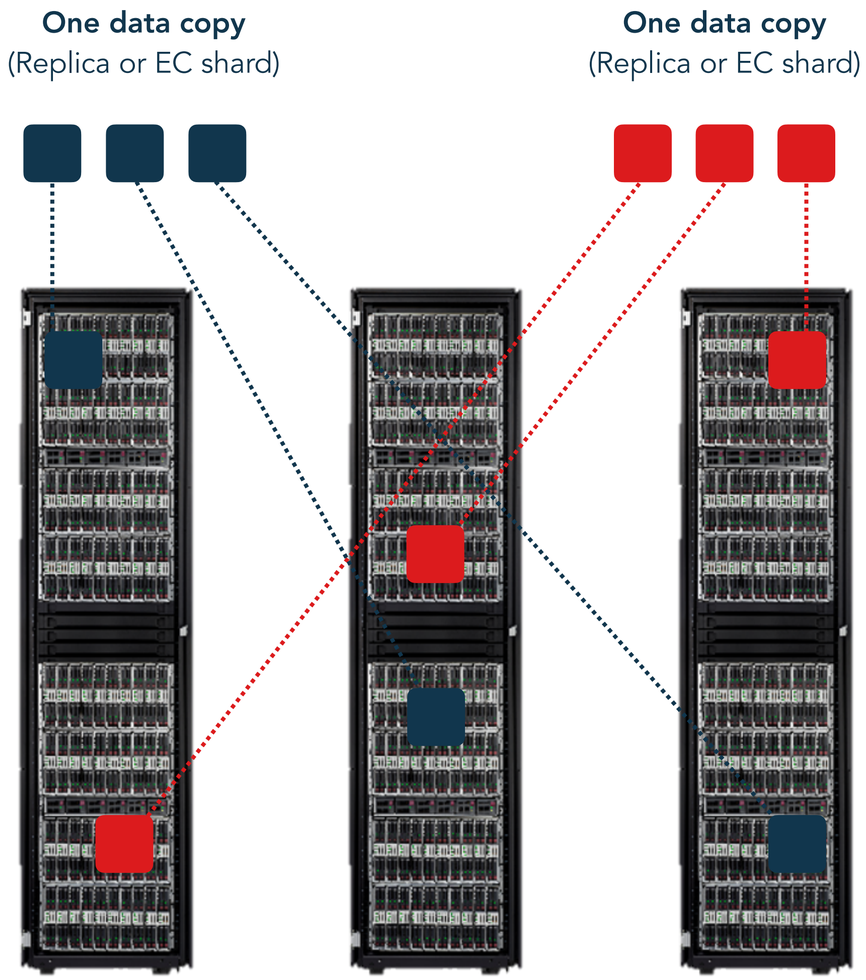
Erasure Coding (K+M):

- Fully supports K+M erasure coding with flexible parameters:
$2 \leq K \leq 22$, $1 \leq M \leq 4$
- Tolerates failure of any M nodes or disks without data loss
- Minimum required nodes: K + M for data protection; M for fault tolerance
- Supports up to 22+2 erasure coding, balancing storage efficiency with resiliency



Sardina Systems Proprietary. Copyright (C) 2014 - 2026. All rights reserved.

FAULT DOMAIN ISOLATION BY RACK AND POWER



- **Replication and erasure coding shards** are distributed across **distinct racks and power zones**, ensuring fault domain isolation.
- Guarantees that a **single rack or power failure** does **not impact data availability or service continuity**.
- Enables **high availability** and **service resilience** even in the event of localized infrastructure failures.

Sardina Systems Proprietary. Copyright (C) 2014 - 2026. All rights reserved.
FishOS® by Sardina Systems. Fully automated cloud solutions for enterprises based on OpenStack, Kubernetes, and Ceph.

COMPREHENSIVE DATA RECONSTRUCTION CAPABILITIES

Ensures fast, flexible, and resilient recovery across a distributed storage architecture.

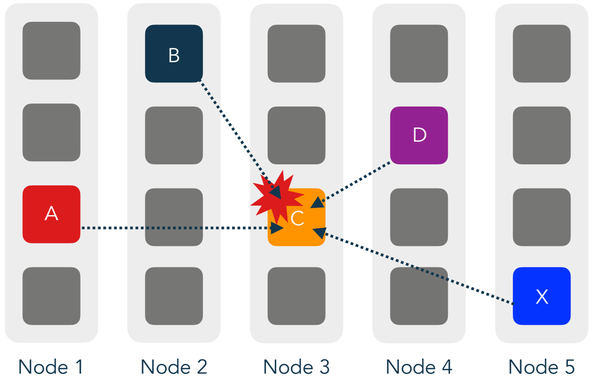
Reconstruction Scenarios:

- Triggered automatically when new nodes or disks are added to the cluster
- Initiated after sustained node/disk failures or manual disk removal

Flexible Reconstruction Policies:

Three-tier policy control:

- Adaptive — balances performance and recovery dynamically
- Service First — prioritizes foreground service continuity
- Reconstruction First — accelerates data repair when needed



Multi-node/multi-disk participation in reconstruction

High-Speed Parallel Recovery:

- Distributed parallel reconstruction across the full resource pool
- No need for dedicated hot spare disks
- Achieves up to 15 minutes per TB reconstruction speed under optimal conditions

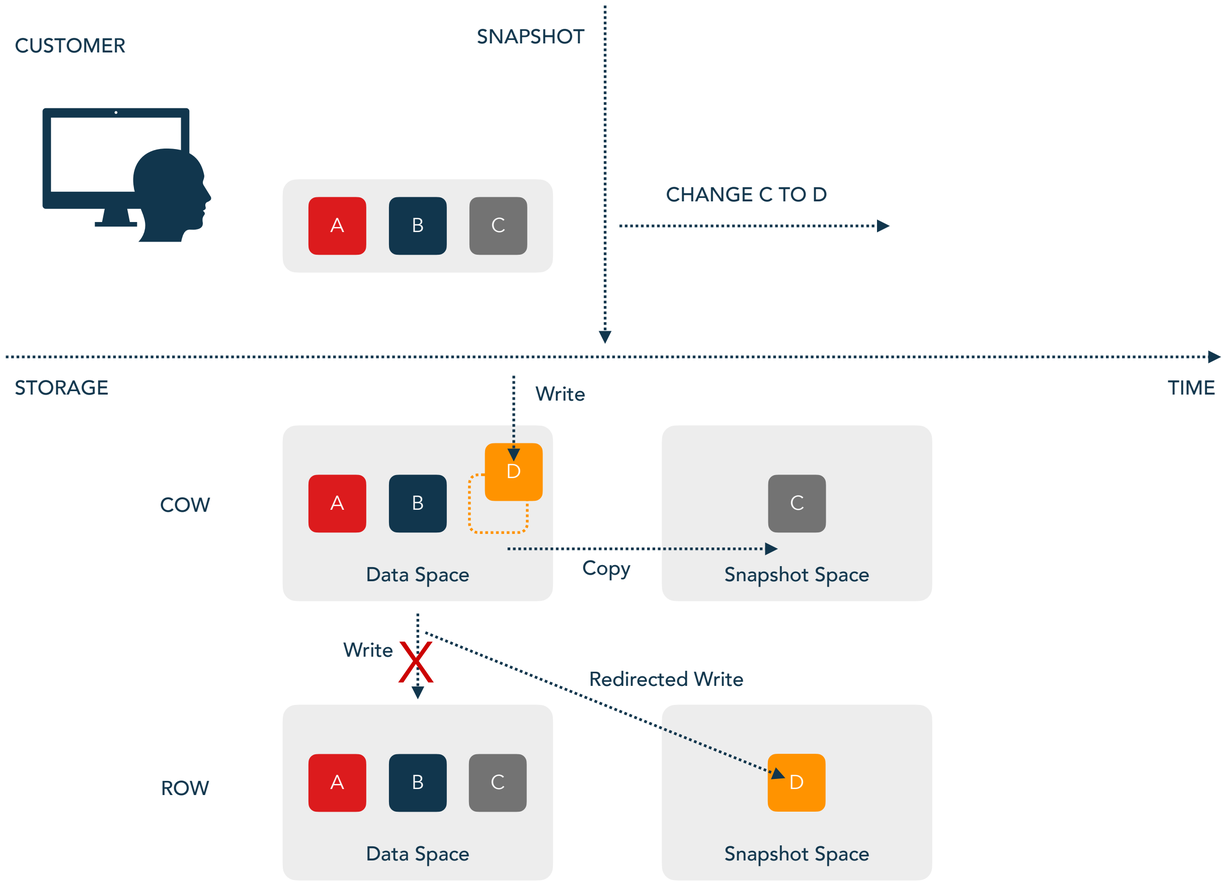
Customer Value:

- Decentralized data layout removes single-node read/write bottlenecks
- Automatic and rapid recovery with no disruption to active workloads
- Maintains high service availability even during repair operations

Sardina Systems Proprietary. Copyright (C) 2014 - 2026. All rights reserved.

FINANCIAL FORECAST, CASH FLOW & USE OF FUNDS: SALES, MARKETING, R&D

Directory or volume-based snapshots



Technical Capabilities:

COW (Copy-on-Write):

- Operates on continuous physical blocks in the original volume
- Frequent metadata (MDS) operations
- Delivers high data read efficiency

ROW (Redirect-on-Write):

- Avoids redundant data copies
- Ideal for OSD write-intensive environments
- Performance impact kept under 5%

Implementation Strategies:

- Supports manual, scheduled, and second-level snapshot creation
- Automatic overwrite of excess snapshots to conserve storage space

Application Scenarios:

- Fast data recovery from accidental deletions
- Safe environments for business testing and validation

Customer Value

Enables recovery of user-deleted data via snapshots with minimal impact on original volume performance, with read/write overhead below 5%

Sardina Systems Proprietary. Copyright (C) 2014 - 2026. All rights reserved.

DATA CONSISTENCY VERIFICATION: PREVENTING SILENT ERRORS

Normal Verification

- Runs during **idle system periods**
- Compares **metadata consistency** across all replicas
- If inconsistencies are detected:
  - Triggers **interface alarms**
  - Automatically initiates **data recovery**
- **Performance impact ≤ 5%**
- Supports **priority adjustment** to balance with production workloads

Deep Verification

- **Default schedule**: every **two weeks**
- Performs **full data content validation** (not just metadata)
- Detects silent errors in data stored over long periods
- Upon detection:
  - Triggers **alerts**
  - Initiates **automated recovery**
- **Performance impact ≤ 22%**
- Fully configurable: **adjust priority and verification frequency**

Sardina Systems Proprietary. Copyright (C) 2014 - 2026. All rights reserved.

NETWORK ISOLATION FOR ENHANCED SYSTEM RELIABILITY

Improves overall cluster stability by proactively identifying and isolating suboptimal network conditions:

- Prevents **repeated MON elections** and **OSD flapping** caused by unstable network conditions
- Enhances system **availability** and **performance consistency**

Network Monitoring Metrics

Continuously evaluates network health through:

- **Packet loss and error rate monitoring**
- **Latency monitoring**
- **Network reachability checks**
- **Stability assessment** across all communication paths

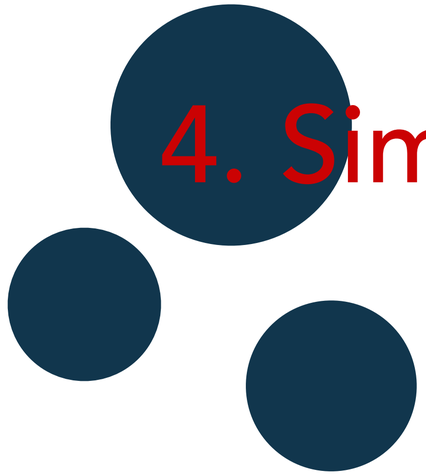
Link Management & Recovery

- **Threshold-based detection** of abnormal links
- **Policy-driven execution** for isolation or recovery actions
- **Automated recovery decisions** based on real-time network conditions

Default Thresholds:

- **Packet loss rate** ≥ 10%
- **Packet error rate** ≥ 10%
- **Latency** ≥ 200 ms
- **Observation period**: 60 seconds

Sardina Systems Proprietary. Copyright (C) 2014 - 2026. All rights reserved.



4. Simplified operations

Sardina Systems Proprietary. Copyright (C) 2014 - 2026. All rights reserved.

EASY DEPLOYMENT & MAINTENANCE

Rapid Deployment

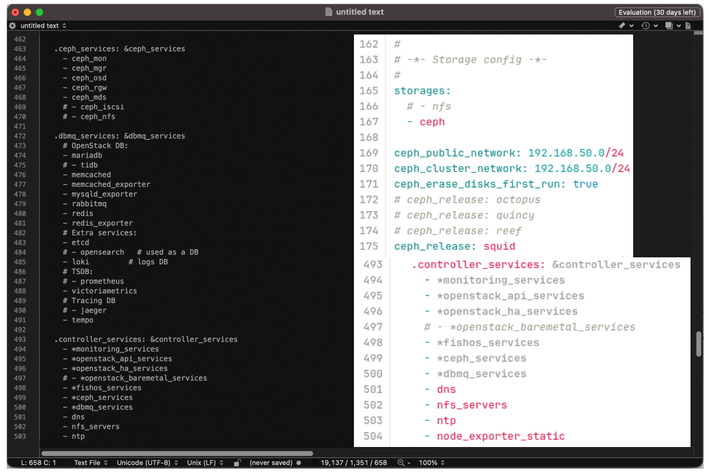
- 60-second deployment with simplified workflows
- One-click scanning and batch node addition
- Supports IPv4/IPv6 dual-stack environments

Comprehensive Monitoring

- Real-time cluster monitoring with alerting
- Disk failure detection and predictive analytics
- Detailed management reports for capacity, performance, and fault trends

Streamlined Operations (Ops)

- Failure localization for troubleshooting
- Unified policy management
- Online, non-disruptive upgrades to avoid service interruptions







Sardina Systems Proprietary. Copyright (C) 2014 - 2026. All rights reserved.

FishOS® by Sardina Systems. Fully automated cloud solutions for enterprises based on OpenStack, Kubernetes, and Ceph.

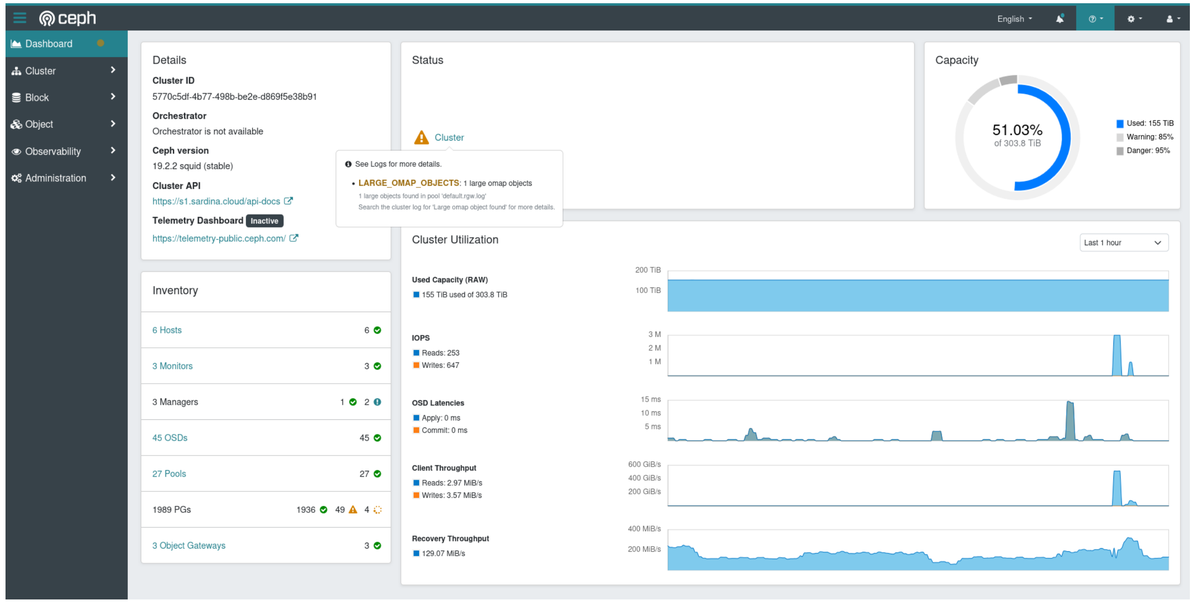
INTELLIGENT O&M TOOLKIT

Cluster Inspection

- Supports **file, block, and object storage** architectures
- Includes routine health checks, hardware status inspections, pre-deployment validation, pre-upgrade checks (online & offline), pre-expansion readiness assessments

Visual Operations Toolbox

- Includes Network topology diagrams, Bandwidth and IOPS visualization, Latency statistics dashboards
- Supports **.log and cluster file visualization** for intuitive monitoring and analysis



Scenario-Based Parameter Checks

- Validates system parameters for specific cases such as: Sobey, media asset management, centralized procurement, and PACS
- Flags abnormalities when current parameters deviate from expected values
- Displays detailed **actual and expected values** for corrective actions

Performance Testing

- Integrated **Vdbench** tool for I/O performance validation
- Supports tests across different **protocols** and **data sizes**
- Outputs include IOPS, latency, and bandwidth metrics
- Helps assess storage system behavior under various load profiles

Troubleshooting Tools

- One-click **log and system info collection**
- Built-in **issue diagnosis and analysis** modules
- Collects module-specific logs to aid root cause analysis

Sardina Systems Proprietary. Copyright (C) 2014 - 2026. All rights reserved.
FishOS® by Sardina Systems. Fully automated cloud solutions for enterprises based on OpenStack, Kubernetes, and Ceph.

INTELLIGENT FAULT & TREND PREDICTION

Disk Failure Prediction

- Real-time **health status** displayed for every disk: *Optimal*, *Faulty*, or *High-Risk*
- **High-risk disks** flagged up to **2 weeks in advance**, with:
  - Predictive insights
  - Detailed logs and alerts
  - Actionable maintenance recommendations
- **Prediction accuracy ≥ 80%, false positive rate ≤ 0.1%**

Performance Prediction

- Uses **AIOps** to model performance trends from **days to months ahead**
- Detects **cyclical workload patterns** (e.g., peak vs. off-peak hours)
- Assists with operational planning:
  - **Data backup scheduling**
  - **Maintenance window alignment**
  - **Upgrade timing optimization**

Storage Capacity Forecasting

- Predicts **capacity growth trends** for up to **12 months**, at both **disk pool** and **cluster** levels
- Generates **expansion recommendations** to prevent resource exhaustion and ensure smooth scaling

Sardina Systems Proprietary. Copyright (C) 2014 - 2026. All rights reserved.



5. Maximized value

Sardina Systems Proprietary. Copyright (C) 2014 - 2026. All rights reserved.

HIGH-RATIO ERASURE CODING: MAXIMIZE DISK SPACE EFFICIENCY

K+M Erasure coding	Minimal Number of Nodes	Space Usage
2+1	3	66%
4+2:1 (ALL)	3	66%
4+2 (ALL)	6	66%
8+2:1 (file, object)	5	80%
8+2 (file, object)	10	80%
8+3	11	72%
8+4	12	66%
16+2	18	88%
22+2	24	92%

Flexible Erasure Coding Support

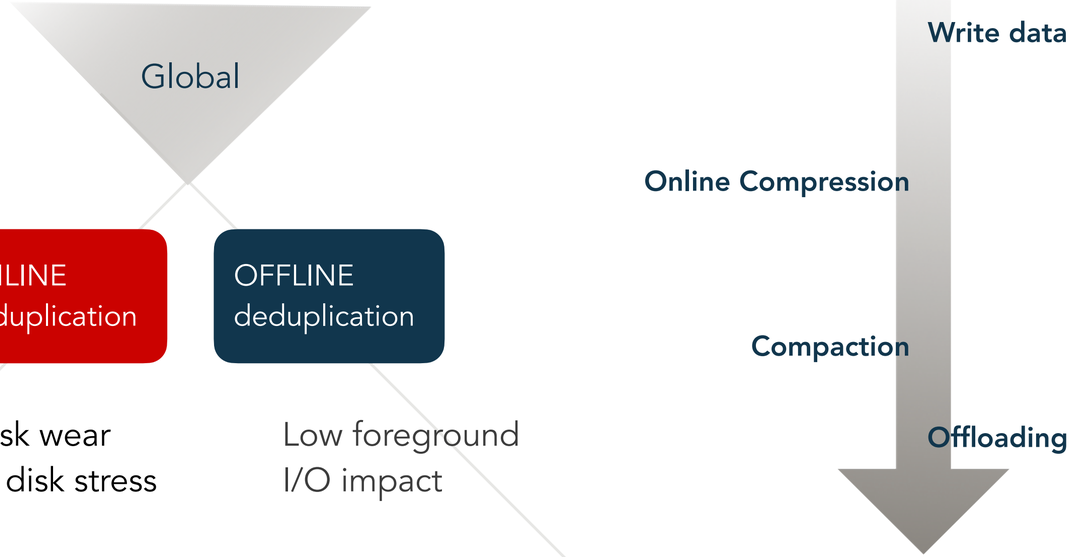
- Supports up to **22+2 (K+M)** configurations
- Flexible coding range: **2 ≤ K ≤ 22**, **1 ≤ M ≤ 4**
- With **K+M** erasure coding, data remains intact even if **any M nodes or disks fail**
- **Minimum node requirement** for K+M:1 scheme: **(K + M) / M**



Achieves up to **175% increase in effective disk capacity** using the same hardware configuration. **Reduces redundancy overhead compared to traditional replication.**

Sardina Systems Proprietary. Copyright (C) 2014 - 2026. All rights reserved.

GLOBAL DATA DEDUPLICATION & COMPRESSION: UNLOCK MORE USABLE SPACE



Efficient Fixed-Length Deduplication:

- **8KB Fixed-Length Deduplication**: Identifies and eliminates redundant data blocks with fine granularity
- **High Availability**: Dual-hash verification using **SHA-1 + CRC** ensures data integrity
- **Low Performance Overhead**: Less than **40% impact** under full-load scenarios
- **Memory Efficient**: Optimized **fingerprint indexing** and **cold fingerprint tables** minimize memory usage

Software compression:

- Compression algorithms: LZ4, GZIP, Deflate
- **Support online compression and decompression**

Hardware compression:

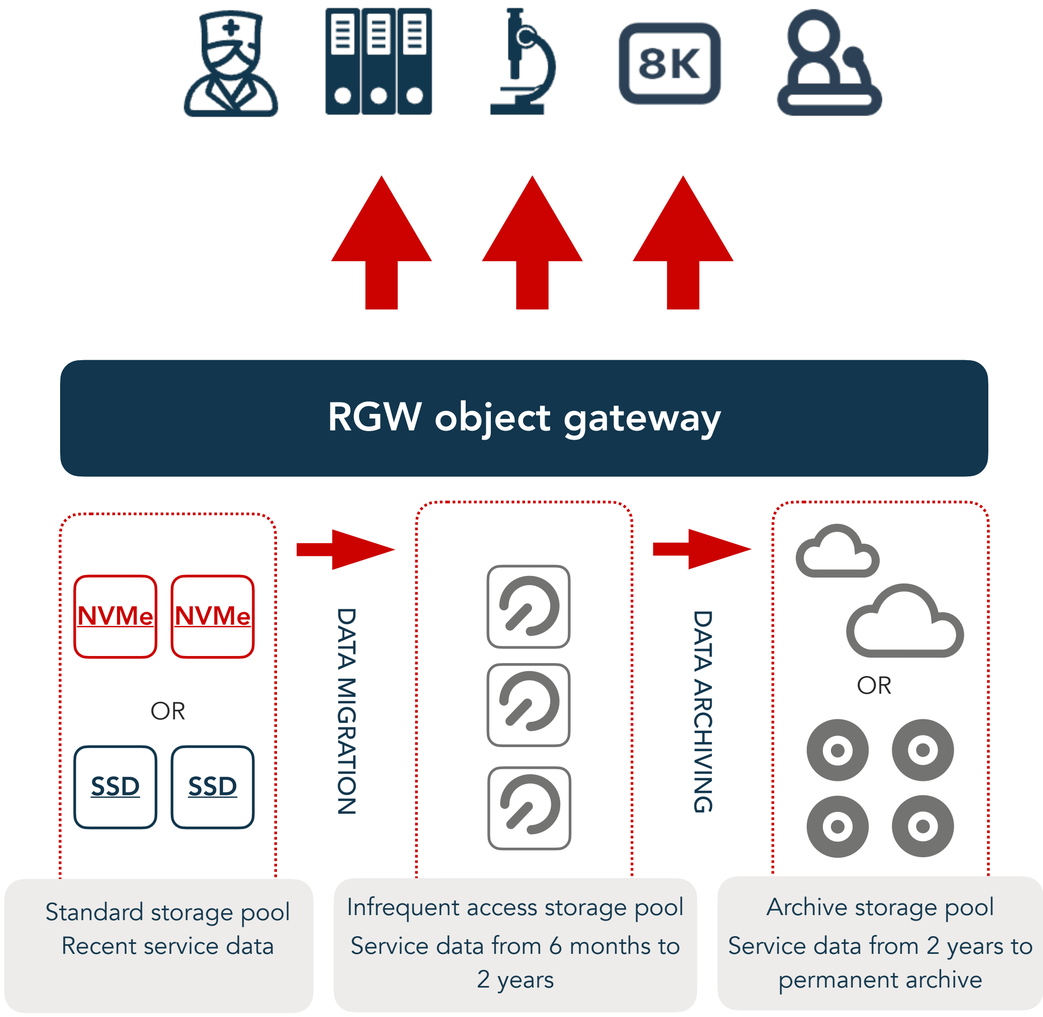
- **Accelerator card**: Support offloading algorithms with Intel QuickAssist Adapter 8960/8970 to **reduce CPU load**
- Compress algorithm: GZIP
- **Support asynchronous compression**

Dynamic Traffic Control:

- Supports **traffic control for compaction and compression** processes
- Dynamically adjusts resource usage to **minimize impact on user workloads**, especially under **high-load scenarios**
- Ensures smooth business operations while maintaining storage efficiency tasks in the background

Sardina Systems Proprietary. Copyright (C) 2014 - 2026. All rights reserved.

OBJECT DATA LIFECYCLE MANAGEMENT



Storage Policies

Supports three storage classes tailored for different access patterns:

- Standard – for frequently accessed data
- Infrequent Access – for data accessed less often
- Archive – for long-term retention of cold data

Policy Configuration

- Define rules by object name prefix to target specific subsets of data
- Apply rules at the bucket level to cover all objects within a storage bucket

Time-Based Actions

- Configure lifecycle actions based on days since last object update
- Transition objects to Infrequent or Archive storage
- Delete expired data to reclaim space

Full Lifecycle Support

- Manages data from initial write to migration and deletion/archival
- Supports application-transparent access:
- Automatically scans based on configured timing and migration speed
- Enables direct access to archived data without rollback

Sardina Systems Proprietary. Copyright (C) 2014 - 2026. All rights reserved.

SARDINA SYSTEMS DISTRIBUTED STORAGE

Massive-scale, high-speed, stable, easy-to-manage, and cost-effective —
built to power modern data-intensive workloads.

Massive-Scale Architecture

- Up to 8,192 nodes for extreme scalability
- Supports multiple storage protocols (block, file, object)
- Compatible with diverse hardware models
- Rich suite of advanced software features
- Fully integrated with mainstream ecosystems for seamless deployment

Simplified Management

- Rapid deployment & easy maintenance
- Integrated diagnostics and intelligent maintenance tools
- AI-powered fault prediction and trend forecasting

Cost-Effective Efficiency

- Data utilization up to 92%
- Frees valuable storage capacity
- Optimizes high-performance storage use

High-Speed Performance

- Accelerated writes: data tiering enhances write throughput
- Prefetching boosts sequential read performance
- Built-in RDMA for ultra-low-latency communication

Enterprise-Grade Stability

- Replication, erasure coding, and snapshot capabilities
- Withstands failure of up to 4 racks
- Distributed reconstruction ensures fast recovery
- Network anomaly detection and data consistency checks

Sardina Systems Proprietary. Copyright (C) 2014 - 2026. All rights reserved.

FishOS® by Sardina Systems. Fully automated cloud solutions for enterprises based on OpenStack, Kubernetes, and Ceph.

1 How Storage Architectures Have Evolved

2 Key Features and Capabilities of Distributed Storage

3 User Interface (UI) Capabilities and Examples

4 Building Storage at Scale: Speed, Stability, Simplicity, and Cost Efficiency

5 Ensuring Business Continuity with Robust Disaster Recovery

6 Proven Results: Customer Success Highlights

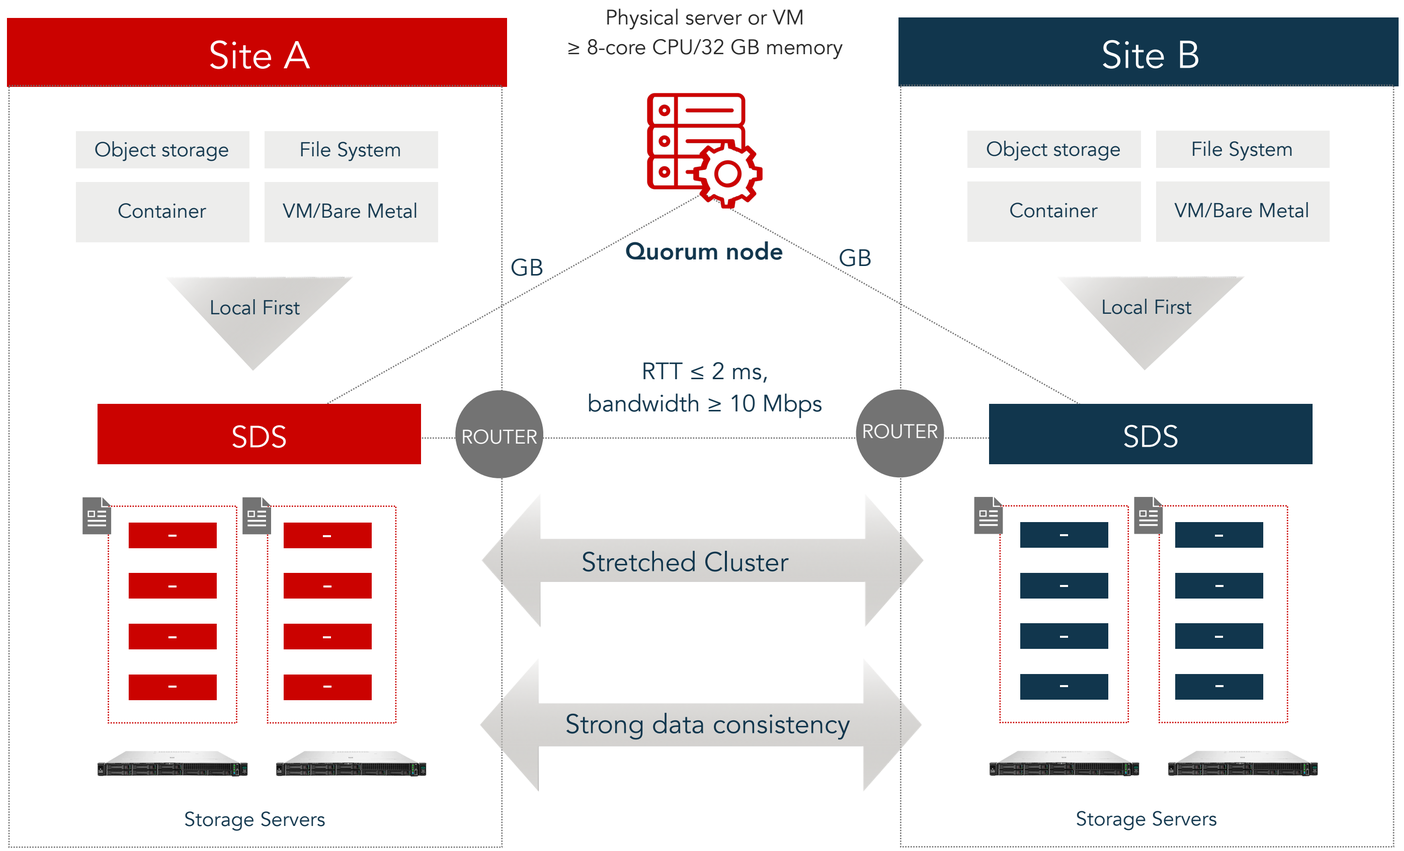


Sardina Systems Proprietary. Copyright (C) 2014 - 2026. All rights reserved.

FishOS® by Sardina Systems. Fully automated cloud solutions for enterprises based on OpenStack, Kubernetes, and Ceph.

METRO DUAL-ACTIVE DISASTER RECOVERY (BLOCK & OBJECT STORAGE)

Built on a stretched cluster architecture, this solution enables dual-active data centers within the same city for zero data loss and continuous service availability.



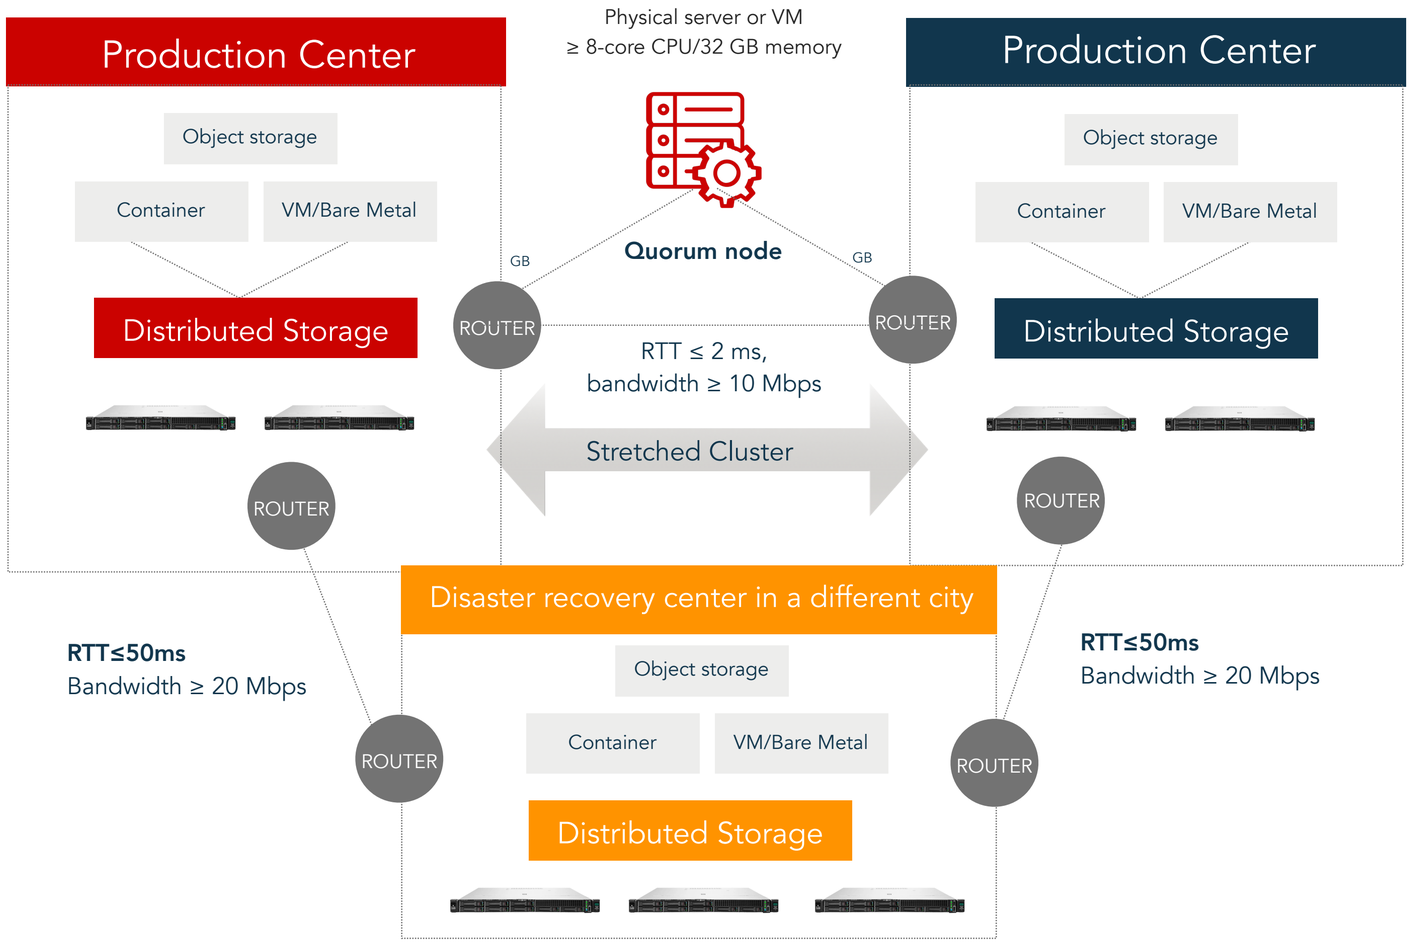
Key Features

- **RPO = 0, RTO ≈ 0** – real-time protection, instant failover
- **Dual-active sites** with strong consistency and 24/7 availability
- **Layer 3 network support** reduces interconnect costs
- **Quorum node** prevents split-brain scenarios
- **4 replicas across 2 sites**, saving **33% capacity** vs. 3+3 setups

Sardina Systems Proprietary. Copyright (C) 2014 - 2026. All rights reserved.
FishOS® by Sardina Systems. Fully automated cloud solutions for enterprises based on OpenStack, Kubernetes, and Ceph.

TWO-SITE THREE-CENTER DISASTER RECOVERY (BLOCK & OBJECT STORAGE)

A **stretched cluster-based** solution enabling business continuity across three data centers — two active sites and one remote DR center.



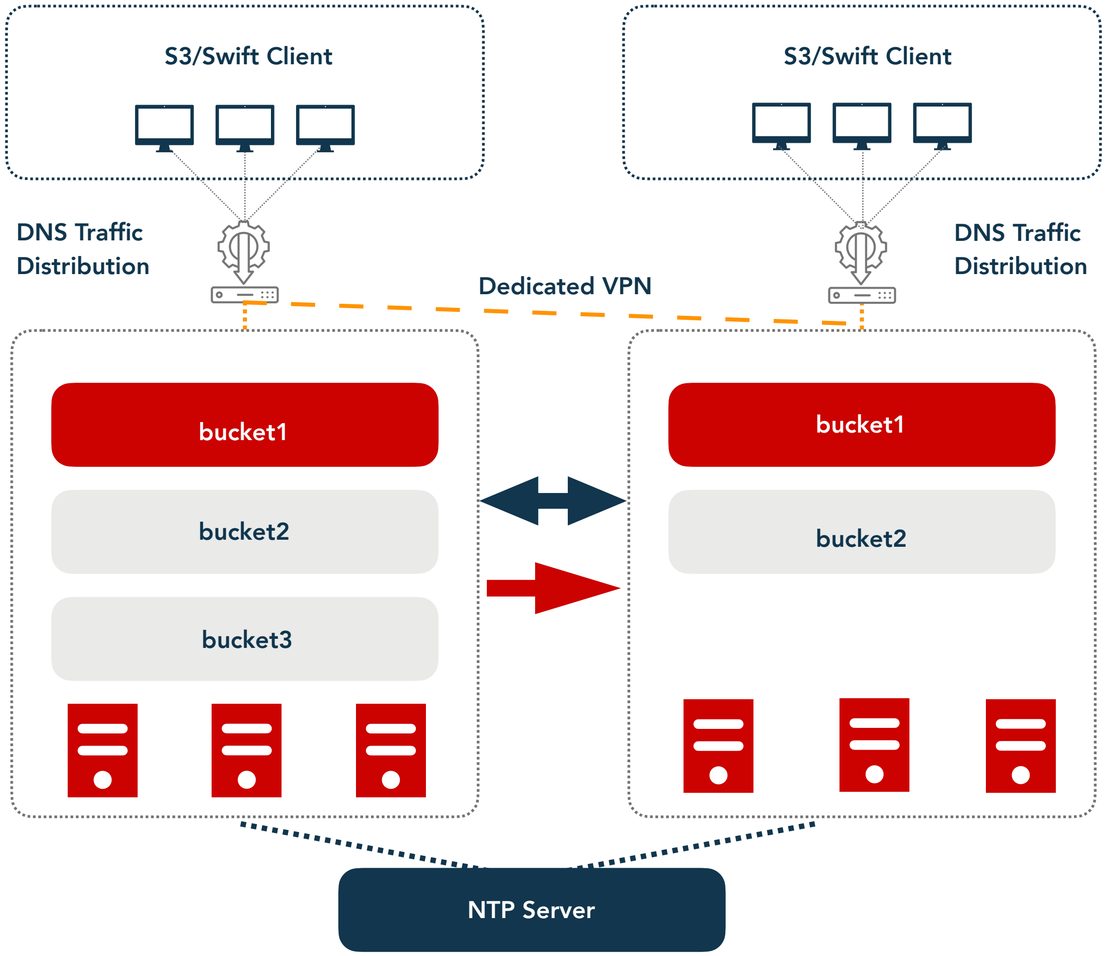
Key Features

- **Dual-active centers** (block/object) with **simultaneous dual writes**
- **4 replicas across 2 sites**, saving **~33% capacity** vs. 3+3 setups
- **RPO = 0, RTO ≈ 0** for active sites (**Level 6 DR**)
- **Remote replication center** with **RPO ≥ 5 minutes**
- **No geographic restrictions** for the remote DR site

Sardina Systems Proprietary. Copyright (C) 2014 - 2026. All rights reserved.
FishOS® by Sardina Systems. Fully automated cloud solutions for enterprises based on OpenStack, Kubernetes, and Ceph.

MULTI-SITE OBJECT STORAGE

A flexible object storage solution supporting **disaster recovery across multiple data centers** with configurable synchronization strategies.



Deployment Models

- **Short reach (<20 km)**: Stretched cluster
- **Long haul (≥20 km)**: Dual-site or three-site architecture
- **Disaster Recovery Level 5, RPO ≥ 3 minutes**

Flexible Synchronization

- Initial **full sync**, then **incremental sync** at defined intervals
- Supports **read-only** or **read-write** modes for secondary sites
- **Bucket-level granularity**: One-way or two-way sync per user configuration

Minute-level RPO

- 3 min, disaster recovery level 5
- Support configuring multiple clusters as multi-sites, each cluster can independently handle read/write services

Sardina Systems Proprietary. Copyright (C) 2014 - 2026. All rights reserved.

1 How Storage Architectures Have Evolved

2 Key Features and Capabilities of Distributed Storage

3 User Interface (UI) Capabilities and Examples

4 Building Storage at Scale: Speed, Stability, Simplicity, and Cost Efficiency

5 Ensuring Business Continuity with Robust Disaster Recovery

6 Proven Results: Customer Success Highlights



Sardina Systems Proprietary. Copyright (C) 2014 - 2026. All rights reserved.

FishOS® by Sardina Systems. Fully automated cloud solutions for enterprises based on OpenStack, Kubernetes, and Ceph.

FishOS by Sardina Systems

Validated by Leaders, Backed by Associations



Silver Member
of the OpenInfra Foundation



Silver Member
of the Linux Foundation



Winner of the IDC HPC
Innovation Award, 2015



Cloud Project of the Year
at DCS Awards, 2024



Winner of the DCD
Open Data Center Project Award, 2017

Sardina Systems Proprietary. Copyright (C) 2014 - 2026. All rights reserved.
FishOS® by Sardina Systems. Fully automated cloud solutions for enterprises based on OpenStack, Kubernetes, and Ceph.

SARDINA

Hyper-Efficient, Scalable, Growth-Focused:
Technology, Product, Enterprise, AI Data Centers

FishOS Ceph by Sardina Systems

Advanced Open-Source Software-Defined Storage Solution



sales@sardinasystems.com
European-headquartered | Operating Globally

ABC SYSTEMS AG
APPLIED BRAINWARE & COMPUTER SYSTEMS

Schlieren, Bern, Zug +41 43 433 6 433 www.abcsystems.ch

Sardina Systems Proprietary. Copyright (C) 2014 - 2026. All rights reserved.
FishOS® by Sardina Systems. Fully automated cloud solutions for enterprises based on OpenStack, Kubernetes, and Ceph.