keeperSAFE: Technical Overview

keeperSAFE™

Keeper Technology data storage, management and protection

Integrated, future-proof and elastic solutions

keepertechnology

In This Issue:

Access | Manage | Protect

Keeper Technology delivers enterprise-class storage technology solutions to government and private sector clients. We provide a broad range of technology products for government program support, data management services, IT department support, and solution delivery.

Like many of our customers, we faced the daily challenges of maintaining an ever-growing volume of unstructured data and struggled with maintaining traditional data storage solutions. So we created something better – keeperSAFE, a private cloud, object-based, scale-out storage platform that has redefined data access, management and protection.

keeperSAFE leverages a unique combination of open source projects and Keeper developed intellectual property, all tightly integrated with commodity hardware to provide an open solution that is free from vendor lock-in. Our single platform provides storage elasticity and scalability from 100's of terabytes to 100's of petabytes, while simultaneously supporting cloud architecture and NAS and Block interfaces.

Keeper solutions provide a secure, available, and protected ecosystem to connect users and applications to data in order to quickly translate that data into information.

Taking on the Challenges of Data Storage

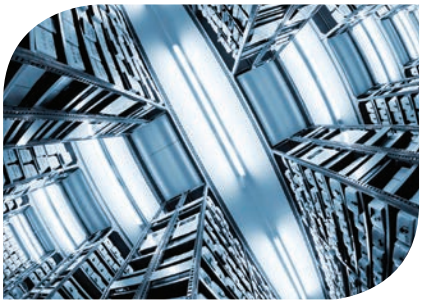

Your data is a critical organizational asset with tremendous market value and strategic implications. Whether in the private or public sector, it can no longer be a support afterthought. Data demands high reliability, fast performance, and low maintenance, all within budget. The backup paradigms developed in the past aren't sufficient for today's big data systems. Traditional add-on solutions have significant limitations or challenges.

For more than 20 years, RAID has been the go-to technology for protecting against disk failures, but disk drive capacity has quickly outpaced access speed. RAID recovery takes too long with today's large disks, and disks are only getting denser. With technologies like RAID6, it may take a 4TB drive days to complete a rebuild. During that time, users are exposed to simultaneous disk failures, undetected bit errors, and extended periods of degraded performance.

Legacy RAID Vs. keeperSAFE

RAID VENDOR	keeperSAFE
Proprietary Software	Based on Open Source
Proprietary Hardware	Commodity Hardware
Based on Logical Blocks	Based on Scalable Objects
Long RAID Initialization	Quick, Parallel Initialization
Long RAID Set Rebuilds	Short 'SWARM' Disk Rebuilds
Access Constraints via Single Controller Pair	Access Expands as Storage Shelves are Added
Number of Drives Supported Limited by Controller Pair	Virtually Unlimited Number of Drives via Storage Shelf Expansion
Comprehensive Management GUI	Comprehensive Management GUI

The keeperSAFE Difference

keeperSAFE presents an entirely innovative solution to today's data storage challenges by utilizing cloud technologies fused with open source breakthroughs. The result is a cost effective, flexible, scalable and efficient storage solution to address the enterprise-class data storage and management requirements of today – and the future.

Purpose-Built, Object-Based Storage Platform

- Limitless Scaling & Elasticity (incremental growth – add or remove components)
- Automated, Real-Time Data Protection – eliminates bottlenecks, aggregate resources
- Extreme Resiliency, Easy to Operate and Maintain

True 'Scale-Out'

- Resource Aggregation Without Bottlenecks
- Objects Provide the Backbone
- Capacity Optomization

Cost Effective

- Open Source + Commodity hardware (leveraging what you have) + Keeper technology
- Features a revolutionary System Wide Auto Rebuild Mechanism (SWARM)
- Reduces Future O&M Costs

ACCESS	**No** performance bottlenecks	**Unlimited** scalability of performance and capacity	**Zero** downtime	
MANAGE	**Central** management of huge data stores	**Unified** multi-protocol access	**Seamless** introduction of new storage	
PROTECT	**Fully** redundant	**Self-Healing**	**Geographic** dispersion	**Automatic** rebuild of failed disks to nodes

keeperSAFE Architecture

KeeperSAFE's foundation is based on intelligent objects. Object storage provides a scalable capability for today's mix of unstructured data. Fixed disk sectors and inodes are inflexible and have significant scaling issues. Object based storage provides the traditional benefits of network storage while adding new benefits:

- Object solutions are platform independent.
- As abstract containers of data, objects can be shared across heterogeneous platforms locally and across networks.
- Each object has a unique identifier not tied to a physical path name, enabling location independence.
- The address space is flat (instead of a hierarchical tree structure), enabling linear scaling by adding nodes.

Because of its building-block architecture, the keeperSAFE solution allows for data storage builds that scale infinitely, and are flexible to run and administer. This flexibility is the key to a "future-proof" system that can change to meet your operational needs. It also allows keeperSAFE to utilize legacy systems and transition and migrate when it makes financial and business sense.

Modular Storage Servers

- Up to 432 TB Raw Storage per 4u
- 10 GbE Back-End Data Distribution
- 10/40 GbE Front-End Data Access

Data Access

- Native Block, Object, and Shared FS
- NFS-3 & 4, SMB-1, 2, 3
- iSCSI, FC Block

Available PCIe SSD Cache

The core open source storage elements of keeperSAFE include Reliable, Autonomous, Distributed Object Store (RADOS) and Controlled Replication Under Scalable Hashing (CRUSH.) RADOS is the software core that manages the mapping of the infrastructure and calculates the location of data objects.

CRUSH distributes and manages the homogeneous allocation of the data between all available disks and nodes. The key function of CRUSH is that every participating component, from clients to admin nodes to storage server, understands the overall architecture and the roles and capabilities of other components. With this shared understanding, all of the components become highly independent. CRUSH allows every component to calculate the location of a data object simply by applying the hashing algorithm to the object ID and the system map. With this algorithmically determined method of storing and retrieving data, the system ensures highly parallelized operations that result in the ability to replicate data quickly and efficiently, protecting against single point failure – or even multiple failures - with little or no impact to customer applications.

CRUSH is in the LINUX kernel, having matured over several years. A standard driver uses the algorithm so any LINUX user can calculate where data is stored, eliminating bottlenecks of traditional solutions. This provides infinite scaling of object location lookups, while NFS, pNFS, StorNext, GPFS, Lustre, Panasas all stop scaling at some point because of contention with metadata access.

High performance is a natural result of keeperSAFE architecture.

keeperSAFE Performance

High performance is a natural result of keeperSAFE architecture. With an object- based core, the system aggregates the performance of all of the components using a network mesh to leverage aggregate network bandwidth. RAID is not scalable; its architecture is linear and reaches saturation levels due to its central data distribution design and reliance on clustering additional hardware and software components for scale-out.

Next Generation Data Management

- Limitless Scalability
  - Distributed, Coordinated Storage Intelligence
  - Combine the Capabilities of Many Drives
- Leverage Point Technologies
  - SSD Cache
- Leverage High-Speed Networking
  - Back End Network for Automated Data Placement
  - Front-End Network for User Access

Capitalizing Legacy Systems

You have invested countless resources in your current data management solutions. Keep those systems, and make the cloud upgrade necessary to position your infrastructure for inevitable changes. The issue isn't if you can, but rather how to do it best. keeperSAFE seamlessly integrates with your existing legacy data management system with no downtime. It assimilates into your existing environment by directly supporting existing protocols, including:

- NAS: SMB 1, 2, 3; NFS 3, 4
- Block data: Fibre channel, iSCSI

The breakthrough technology behind keeperSAFE not only enhances your current system, solving today's data storage and access issues, but also builds a system capable of evolving for future needs. Add new technologies – keeperSAFE delivers a scalable cloud solution capable of integrating with RESTful protocols such as S3 – or remove technology, or even deploy new components, with no loss of access.

keeperSAFE Management

Managing and maintaining a data storage system in a mission-critical environment is filled with challenges. So we designed keeperSAFE to be a more elegant solution, seamless and trouble-free. It features centralized software deployment, network management, log file management, cluster management, configuration, administration and maintenance, and seamless integration of all modes and network components.

keeperSAFE features a dashboard, for easy, unified management of the system. Customize your dashboard for monitoring, analyzing and trouble-shooting. Even a small, three-shelf system, for example, can collect more than 45,000 data points per hour. (continued on next page)

keeperSAFE Management (continued)

keeperSAFE installs automatically, and offers one-step re-initialization, and online firmware upgrades. The modular architecture allows for upgrades as needed. This includes components and technology, including:

- Admin nodes
- Monitor nodes
- Gateway
- Shelf
- Drive
- Storage
- Monitor nodes
- Increased shelf density
- New access through Gateway upgrades

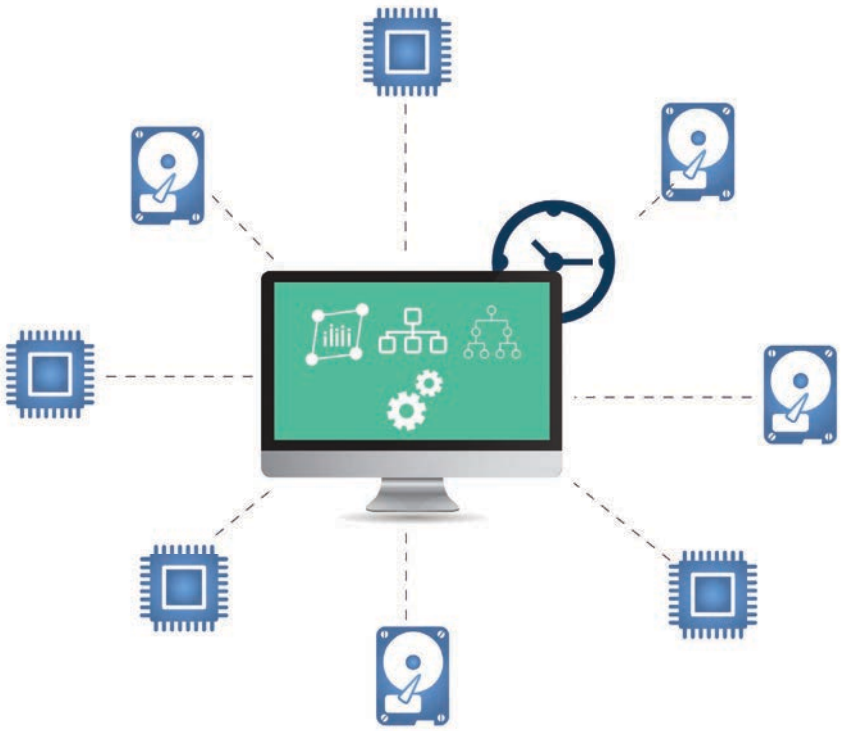

Access | Manage | Protect

keeperSAFE is a purpose built, scalable storage ecosystem. It's an evolutionary technology designed to ease the switch from a traditional, dedicated bare-metal application environment to cloud-based architectures without having to invest in a new storage infrastructure. Key aspects of the technology protect your investment, allowing your data system to evolve and grow while eliminating the most common pitfalls of traditional systems.

FEATURE	BENEFITS
Self-Healing	• Reduces response time by removing the human-involvement element
Future-proof/elastic	• Granular nature of the system means it will continue to evolve • Adapt to new needs and new technologies • End the cycle of disruptive migrations
Reducing operational costs	• Utilize legacy systems • Scale as needed • Reduce capital "forklift" investments • Eliminate downtime and data access delays

Keeper Technology protects your greatest asset – your data – while also considering the bottom line. The cost of managing and maintaining a useful and accessible system are held in check by design. From keeperSAFE's foundation of open source technology, to the utilization of legacy hardware, to future upgrades and component change outs, to online data migration, we've created a solution matched to the long lifespan of your organizational data.

keeperSAFE™

Scalable. Flexible. Secure.

Affordable data storage the way the cloud intended.

keepertechnology | 21740 Beaumeade Circle | Suite 150 | Ashburn, VA 20147
P [571] 333 2725 | F [703] 738 7231 | solutions@keepertech.com | www.keepertech.com