



keeperSAFE™

The New Paradigm Shift:
Software-defined Storage

keepertechnology



Table of Contents

Introduction	3
keeperSAFE Overview	3
Software-defined Storage	4
Tightly Integrated with Commodity Hardware	6
Architecture	7
Data Availability	10
Software and Microcode Upgrades	10
Hardware Redundancy and Reliability	10
Data Protection	12
Advanced Data Management with IDMT	14
Connectivity and Integration	15
Reporting and Analytics	17
Conclusion	18

Introduction

Machine-generated data is the fastest-growing segment of the mass storage data market. The digital content explosion combined with our “data-on-demand” society is presenting organizations with unprecedented scaling challenges. The rapid growth of online, near line, and backup data pools has pushed general-purpose storage systems beyond their capabilities and has forced manufacturers and users to rethink their storage architectures across the data storage industry.

This growth is driving storage technology innovation into the marketplace. Keeper Technology (Keeper) has been delivering state-of-the-art integrated technology solutions to commercial and government customers for more than a decade, particularly in industries that need to provide continual online access to data in the petabyte and exabyte scale. As Keeper has grown, we have remained focused on our core mission to provide enterprise class, hyper-converged data storage and data management solutions to our government and commercial customers.

KeeperSAFE is the definition of a true hyper-converged architecture. Building from the ground-up, Keeper has tightly integrated open-source software defined storage into fully vetted commodity hardware. The end result: an affordable, scalable, and elastic appliance package featuring full compute, store, networking, virtualization, and management capabilities in one.

This document describes the evolution of keeperSAFE (kSAFE), a turnkey data storage appliance created by Keeper, as well as its functions and operation. This document also addresses key capabilities and specifications such as the product’s data access, management, security, and other technical requirements as well as its benefit to in-house data center operations. This whitepaper concludes with Keeper’s application as a purpose built data storage solution created for a fully-integrated data management environment.



keeperSAFE Overview

A Next-Generation Platform for Exabyte-Scale Storage

Keeper introduced the first version of kSAFE in January 2011, to address specific use cases within our government customers for managing and protecting the unique combination of large data with billions of files.

The first generation of kSAFE provided capacities up to 100 TB and up to 5 billion files through a NAS presentation layer and leveraged commodity RAID technologies for the data storage layer. Features such as online snapshots, snapshot recovery, local and wide-area replication were utilized in most customer deployments.

The second-generation kSAFE became available in 2013, and allowed for capacities up to 1 PB and file counts up to 15 billion. We also added capabilities for block access over Fiber Channel and iSCSI presentation layers.

Our current, third-generation kSAFE began shipping in 2015, providing virtually limitless scaling for both file counts and capacities due to its ability to store data as scalable objects.

The backend storage infrastructure transitioned from RAID technology to a ‘storage server’ infrastructure where each storage shelf has built-in intelligence, facilitating advanced data management and tiering, all while performing online upgrades, online technology insertion and online technology decommissioning. This generation kSAFE added native integration with Intelligent Data Management Technology (IDMT) and native disk encryption, amongst many other features.

Today’s kSAFE is built upon the open-source Ceph project. Many of the core features for kSAFE follow the Ceph roadmap including support for CephFS and BlueStor support. Keeper makes rolling improvements for overall advancements in management and usability of the GUI interface. Keeper constantly evaluates new storage server platforms for inclusion in the kSAFE ecosystem and anticipates investigating very high-density storage solutions for availability in the near future as well as NVMe solid-state solutions to serve as a tier of storage for certain use cases requiring extremely low latency combined with high scalability.

Software-defined Storage

Infinite Capacity and Elasticity

For more than 20 years, RAID had been the go-to technology for increasing performance while ensuring data redundancy, but disk drive capacity has significantly outpaced access speed. In response to this together with explosive data growth, Keeper transitioned away from RAID as its backend storage infrastructure in 2013.

Traditional scale-up RAID storage architectures face significant limitations when volumes approach the petabyte level. RAID storage solutions distribute data across sub-sets of spindles, which are not fault tolerant by design and limit performance scalability. Additionally, RAID recovery takes too long with today's large disks, and disks are only getting denser.

The Keeper storage ecosystem is an open-source, software-based storage platform that scales across physical, virtual, and cloud resources without limits. It is built upon the open-source Ceph project; Ceph is an object-based storage solution that supports scalable clusters up to the exabyte level. Ceph bypasses the limitations of RAID in many ways – primarily by replicating data in multiple locations and fault domains. This uses less expensive disk controllers and avoids the problems common with RAID and today's large disks.


Manage big, semi-structured, and unstructured data growth while maintaining performance, capacity, and availability.

Open-source technology provides a core capability to power enterprise data management systems, but it takes unique talent to bring these disparate projects together to make them function and perform efficiently. Keeper Technology harnesses the power of open-source within kSAFE, creating a comprehensive, flexible, scalable, and supportable data storage solution.

The kSAFE storage architecture enables algorithms to broadly distribute objects, files and block data homogeneously over disks and nodes to provide maximum utilization of all system resources. And unlike traditional RAID architectures, kSAFE supports this amalgamated technique for system fault tolerance, which is designed to be both self-healing and self-managing.

Figure 1: Algorithmic Methodology of Storing and Retrieving Data

kSAFE delivers extraordinary scalability with thousands of clients accessing petabytes to exabytes of data.



In the kSAFE environment, objects, files, and block data are allocated homogeneously across disks and nodes; this results in a very high fault tolerance, scalability of capacity, and data migration with zero downtime.

Ceph, which is at the core of the Keeper storage ecosystem, uses RADOS (Reliable Autonomic Distributed Object Store) to separate objects from the underlying storage hardware. RADOS ensures flexibility in data storage by allowing application object, block, or file systems to interface simultaneously. The object storage devices are not just for data access, but also allow for serialization, replication, and failure detection.

The CRUSH (Controlled Replication Under Scalable Hashing) algorithm is the underlying technology in the Keeper software architecture. CRUSH deterministically computes where data can be found and should be written. CRUSH provides a better data management mechanism compared to older approaches, and is better suited to hyper-scale storage.

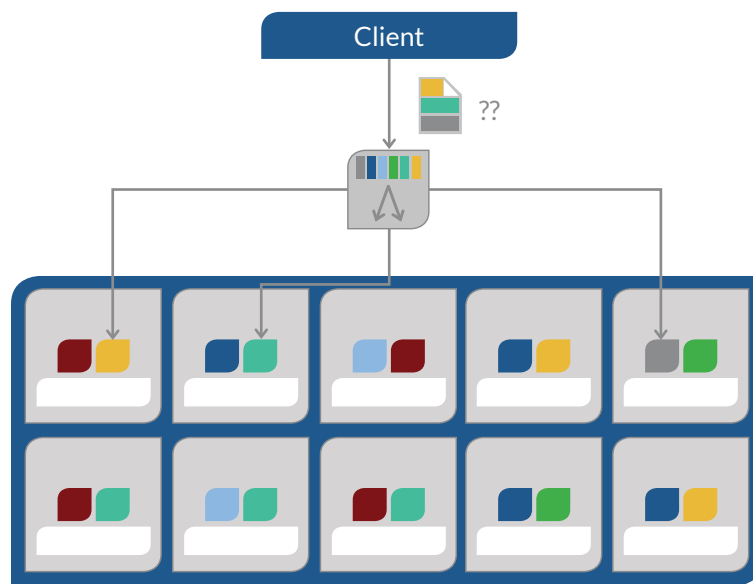
CRUSH uses a weighted distribution methodology to determine how data is stored, retrieved, and algorithmically allocated between all available disks and nodes; it cleanly distributes the work to all the clients in the cluster.

CRUSH Uses Intelligent Data Replication to Ensure Resiliency and Elasticity

CRUSH is aware of the infrastructure and relationships between the various components: server, rack, switch, power circuit, data center, etc. These components are treated as fault zones, and CRUSH places data and replica locations to ensure that the contents are safe and accessible even if some of the components fail. This eliminates single points of failure, and enables use of commodity hardware to build a highly resilient storage system.

When components fail or new disks or nodes are added to the cluster, kSAFE autonomously adjusts to the new layout and moves data in order to re-establish even distribution and fault tolerance. As the cluster grows by adding more disks or nodes, the power to manage that data grows at the same pace. As more

Figure 2: CRUSH Data Mapping



kSAFE creates a map of the storage node cluster to homogeneously store and retrieve data across the complete system.

clients access the cluster, the power to find data also grows. There are literally thousands of open source projects that could assist in the creation of the keeperSAFE appliance. Selecting the best available and then ensuring that they operate seamlessly together with hand-selected, best in-class hardware, requires significant data storage industry experience as well as decades of involvement in open source software development and integration. As an enterprise data storage and data management company, Keeper Technology is uniquely positioned to qualify and integrate the best of breed open source projects into keeperSAFE.

Infinite Capacity

kSAFE is built to scale to the exabyte level and back.

Unlimited Flexibility and Performance

RADOS enables object, block, or file systems to interface simultaneously.

Intelligent Storage Nodes

Nodes actively collaborate, replicate data, consistently apply updates, detect node failures, and migrate data.

Strategic Data Placement

The CRUSH algorithm ensures efficient data placement as data is added, changed, and removed.

Minimal System Administration

Clustering increases performance, avoids data bottlenecks, is self-healing and self-managing.

No Single Point of Failure

Customers have no awareness of system failures because their data remains available.

Tightly Integrated with Commodity Hardware

Because kSAFE is based upon open source software technologies, the natural question is: Why can't I use the same or similar open source technologies to create my own kSAFE 'lookalike' product? The short answer is, 'It is feasible, but you have to realize what you are getting into when creating a one-off, open-source solution in an enterprise-class environment.'

Participating in the open-source community is a two-way street. There is an expectation that companies utilizing open source will dig into the source code and contribute enhancements back to improve the quality of the open-source code. It takes time and commitment to develop credibility within the open-source community and Keeper Technology is already there, and fully engaged through our team's monitoring of previous and ongoing bug-fixes, enhancements, and assistance.

Bridging the Gap Between Hardware and Software

Keeper Technology has spent tens of thousands of hours and millions of dollars in facilities and capital equipment investment to create a sustainable, supportable data storage product line. It is one thing to install software downloaded from the Internet and make it work to some level of functionality; it is entirely another matter to create an enterprise-class data storage solution that is tightly integrated with hardware.

Even for those customers capable of 'rolling their own' open source solutions, additional software development is necessary to fully integrate the solution, and the long-term maintenance of the environment can be challenging. There is always help available from the open source community, but the response is typically related to level of engagement with, and commitment to, the open-source community.

Hardware Qualification

The Keeper Technology engineers continually run commodity hardware solutions through an extensive evaluation process to ensure that we present a reliable, sustainable solution to appeal to a broad range of customers. In our initial search for a viable platform, we investigated physical hardware reliability and maintainability specifications as well as the availability of features for software system information needed to detect and correct problems and failures during normal operations. Keeper Technology acquired several hardware platforms to evaluate for long-term operations, and only the best platform survived. Hardware source selection is an ongoing process to bring new form factors and capabilities to market.

Software Development Qualification

The Keeper Technology software development team has worked tirelessly to fill in the gaps when open source software is not available or viable and ensure that the hardware and software work well together. As with most products, error management and reporting are key requirements for understanding if a system is running properly. Keeper collects tens of thousands of data points per hour, analyzes them, and rolls that information up into our own intuitive management GUI that presents both software and hardware status, performance information, and a host of configuration and operational data.

System Management Qualification

One of the unique things about combining open-source software with data storage hardware is that the open source software has been written very generically so that it does not depend upon any one vendor's implementation of hardware.

Keeper Technology provides comprehensive enclosure management services to detect failing and failed components from disks, as well as power supplies and system interconnects. kSAFE then promptly and accurately displays that information on our Keeper GUI so a systems administrator can provide the appropriate response. Having these features is a critical differentiator when buying a solution versus attempting to develop your own open source solution. Without these advanced management features that Keeper has developed, the level of effort and risk to managing a system is immeasurable.

Keeper allows all of our customers to benefit from the agility, elasticity, and transformative nature of today's open-source software combined with qualified data storage hardware in a repeatable, supportable data storage appliance.

Architecture

Because of its Lego-like architecture, kSAFE allows the building of data storage that scales incrementally without worrying about capacity overload or without having to undergo disruptive forklift upgrades. As components fail or new devices are added, kSAFE autonomously adjusts to the new layout and moves data to reestablish even distribution and fault tolerance.

kSAFE natively supports OpenStack and other Linux physical and virtualized storage environments. It fully leverages a unique combination of mature open-source data storage projects and Keeper Technology developed code, all tightly integrated with commodity hardware, providing an open solution that is free from vendor lock-in.

Based on customer-required configuration, all hardware and software components are delivered as a complete kSAFE product.

Key Characteristics of a kSAFE

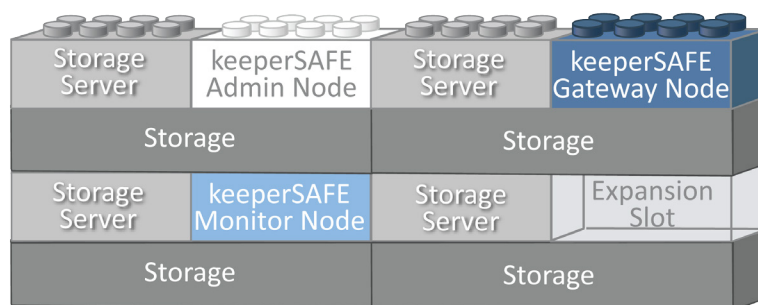
The kSAFE storage platform provides common data storage features and capabilities of cloud architectures, including Storage as a Service, by providing a well-integrated and closely aligned combination of hardware and software capabilities.

While there are no maximum limits in capacity, amount of data, or number of files/objects kSAFE can grow to, there is a minimum base configuration from where kSAFE can begin. An entry-level kSAFE configuration should include:


- (1) Admin Node
- (3) Monitor Nodes
- (3) Storage Nodes
- (1) 10/40 Gb Ethernet Data Switch
- (1) 1 Gb Ethernet Management Switch

System capacity growth is accomplished through a straightforward, on-line addition of new storage shelves, thus increasing on-line and near-line capacities independently.

Figure 3



Each storage shelf has built-in intelligence, facilitating advanced data management tiering while mastering online upgrades, online technology insertion and online technology decommissioning.



kSAFE Growth: Advanced Enterprise Configuration

kSAFE is capable of scaling in multiple dimensions. As storage nodes are added to the system, they increase the overall capacity as well as the total number of spindles. These nodes add a balanced amount of bandwidth to the backend and frontend networks. If configuration parameters are chosen to optimize for performance, each storage node added to the system can add as much as 1 GB/s of performance.

kSAFE gateways also support a scale-out model. For most protocols, gateways can be clustered together to achieve higher aggregate performance. The more gateways that are added to the system, the more throughput the clients will have available and the more data that can be pushed to the storage nodes on the back end.

All gateways, regardless of protocol (IDMT, Block, NAS, Object) are configured with a total of 4 x 10 GbE ports – connections to both the customer network for client access and also to the kSAFE frontend network for access to the storage nodes. The CPU, memory, and local storage are tailored to the functions of the gateway.

This model results in all gateways having approximately 1 GB/s of available performance in the system. Given that a kSAFE could scale to many hundreds, if not a thousand or more storage nodes, and a hundred or more gateways, the performance of a full up kSAFE could be an aggregate of well over 100 GB/s.

Keeper constantly evaluates new storage server platforms for inclusion into the kSAFE ecosystem and anticipates exploring other very high-density storage solutions for availability in the future, with NVMe solid-state solutions to serve as a tier of storage for certain use cases requiring extremely low latency combined with high scalability.

Figure 4: Advanced Enterprise Architecture





System Expandability

Scalability and ease of expansion are critical to any data-centric organization's overall effort to increase revenues: allowing for organic growth in capacity in order to take on more projects and preventing loss of business due to the system capacity overload. The kSAFE storage platform augments the speed, ease of use, and ultimately the time to product completion on an exponential scale.

kSAFE is capable of scaling in multiple dimensions. As storage nodes are added to the system, they increase the overall capacity as well as the total number of spindles. They also add a balanced amount of bandwidth to the backend and frontend networks.

From a hardware perspective, because kSAFE is an object-based storage system with a modular design, the system scales in both performance and capacity by simply adding additional storage shelves.

Intelligent Storage Nodes

The lifecycle of the overall systems can be extended and migration efforts can be significantly shortened because storage nodes can easily be replaced with nodes of the next generation – non-disruptively while the system is online and providing uninterrupted data services to users and applications.

What might take traditional storage systems days or weeks to add 2 PBs to the enterprise environment, with kSAFE customers can be ready to start writing data in under an hour. Nodes actively collaborate, replicate data, consistently apply updates, detect node failures, and migrate data.

Other scalability and expandability features include:

- Incremental Upgrades: Upgrade individual components as needed. Imagine initiating a project that requires a PB of storage and where management no longer needs to buy the entire PB. Start with 200 TBs and add increments of 200 TBs as storage is needed, when it is needed.
- Future-proof Scalability: Easily, transparently, and incrementally insert and remove technologies and capabilities.
- Component-based Scalability: If and when a three-year-old 4TB drive fails, simply swap it out with a newer 8TB drive. kSAFE gives you full use of the entire drive capacity allowing double the capacity within the same form factor.
- Performance Scaling: Being object-based storage at its core, kSAFE performance scales by aggregating the capabilities of all the drives in the system.

The kSAFE software scalability advantage is closely aligned with its hardware capabilities. Scaling of the CRUSH map is infinite. There are virtually no maximum limits in capacity, amount of data, or number of files/objects kSAFE can grow to.

Data Availability

kSAFE leverages open source software-defined storage, commodity hardware, and a Keeper-developed management layer to deliver the highest levels of availability possible. The distributed nature of kSAFE is fully redundant with all critical components that are replicated, mirrored, or otherwise protected against failures.

Software and Microcode Upgrades

Software and firmware upgrades are handled with no downtime and no disruption. If a running service is affected by a particular software upgrade, it is performed as part of a failover or similar fashion such that it is non-disruptive to clients of the system.

The entire software upgrade process, whether it is firmware, microcode, or SAFEos (the kSAFE operating system), seamlessly is managed by the kSAFE Admin Node. If the upgrade being applied on a node causes that node to briefly be taken offline, then the upgrade is performed in a rolling manner, thus having no impact to clients of the system. This is one of the benefits of a software-defined clustered system. The redundancy of the architecture allows for one or more nodes to be taken offline either briefly or for extended periods with no disruption in service to the clients.

Hardware upgrades are straightforward with new components simply added to the system. New hardware is recognized by the system, configured, brought online, and made available to the software layers as additional resources. All of this is done with no disruption to ongoing activities, no interruption of service, and no downtime.

Hardware Redundancy and Reliability

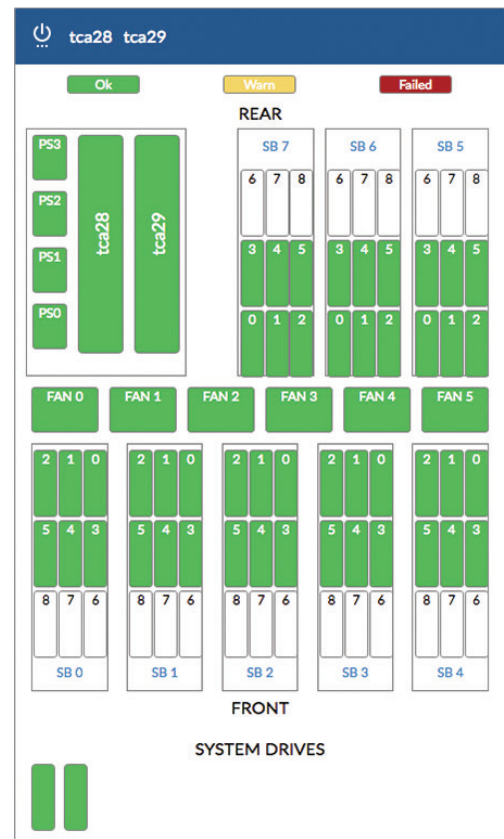
The unprecedented reliability of a kSAFE environment begins with extremely reliable hardware. Keeper continually tests and assesses hardware offerings in the development lab and integration center at Keeper Technology headquarters in Ashburn, Virginia.

Based on our staff's more than 150 combined years of experience operating and maintaining mission-level storage environments, we qualify hardware based on a number of factors:

- Quality of manufacturing process
- Degree of redundancy internal to the node
- Ease of management
- Customer and field replaceable units
- Reliability during our own physical and runtime testing
- Third-party reviews and experiences

Based on this qualification process, only the most reliable and easily maintainable systems are chosen as kSAFE hardware options. Once hardware is approved, then the management modules are written. Using low-level protocols (SNMP, IPMI, etc.) we communicate with, monitor, and manage all levels of the node from the baseboard management controller to fans, power supplies, disks, boards, etc. All of this information is captured in a database and made available for display through the GUI and for analysis.

Figure 5: kSAFE GUI



In a kSAFE environment, a single storage shelf can generate more than 15,000 data points per hour. In the GUI, this information is coalesced, refined, and presented through an intelligent use of design and color to convey the most meaningful information to the user as quickly and easily as possible. The data from all nodes of a kSAFE also feed into the analytics module that can present a variety of user-defined graphs as well as perform predictive failure analysis.



kSAFE is Architected with No Single-Points-of-Failure

In the event of hardware failure (disks or nodes) the system is able to seamlessly proceed without downtime. kSAFE automatically recreates new data copies (replicas) for those that have been lost during the failure of components and data copies are available on other locations within the system.

Additionally, this replication provides the ability to maintain and track multiple copies of the same data in different locations for data protection, collaboration, and/or disaster recovery.

Every critical component that would cause a disruption of service to clients in case of failure is engineered as redundant at the hardware node level. These include:

- Data Network Switches
- Admin Nodes
- Monitor Nodes
- IDMT Gateways
- Storage Servers
- Near Line Storage Nodes

Many of these components share additional protection through internal hardware redundancy (power supplies, fans, etc.) and also at the software level (erasure coding, checksums, etc.). Ultimately, the varying degrees of kSAFE design, testing, and system architecture exceed the stated requirement of “no single-points-of-failure.”

The No-Downtime Benefit to the Analyst

If the soul of an enterprise is in its data, then its heart is in the user. Analysts and their tools—both deriving answers from the continual onslaught of information—are critical to the success of any enterprise. KeeperSAFE provides a stable and reliable information backbone, ensuring that data is accessible and dynamically protected.

As new algorithms and approaches are proven, more compute and storage resources are dynamically allocated in real time, all without interruption, providing the quick reaction capability that leads to competitive advantages, informed decisions, and solutions. The heart and soul are satisfied, leading to tangible advantages in efficiency and mission capabilities without breaking the budget.

Data Protection

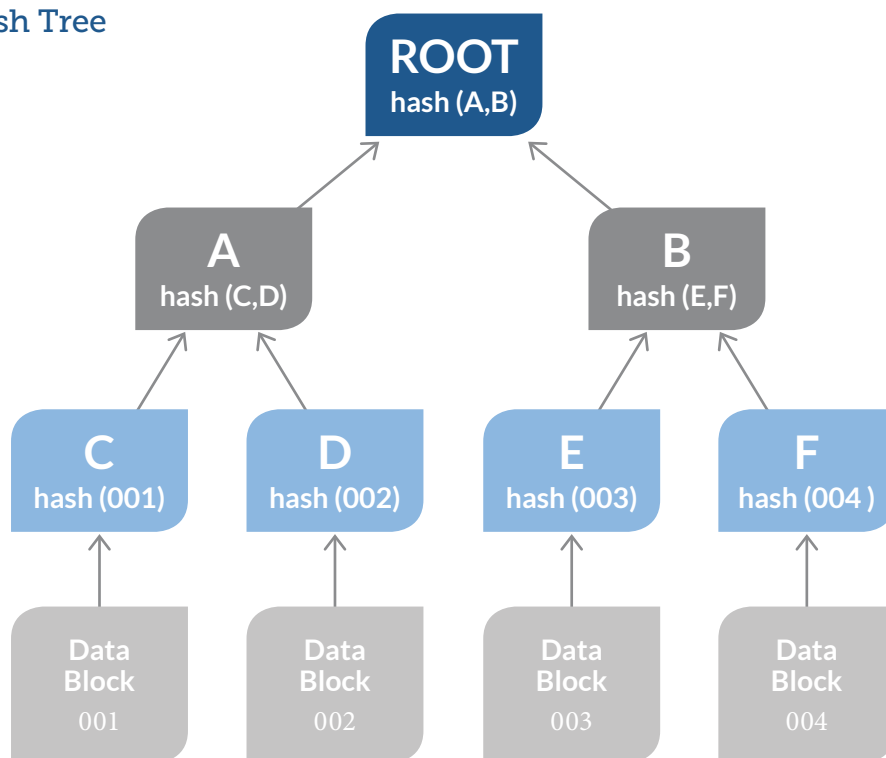
At Keeper Technology, data protection is taken very seriously. We employ multiple types of protection at various levels with many data checks along the way. This is best illustrated by looking at data protection layer by layer.

Near-line tier

In the kSAFE near-line tier, data is protected through a combination of 256-bit checksums at multiple levels along with triple-redundant local erasure coding across the disks, which, in itself has an error rate of 1 bit in 2 million years.

The checksums are stored in a hash tree with each node holding the checksums of all the blocks beneath it. This protection can be visualized in two dimensions. Since this is archive storage and not meant to be dynamic, the erasure coding and checksum parameters are optimized for long-term storage and are not user configurable.

Figure 6: Merkle Hash Tree



The primary tier of storage, on the kSAFE storage nodes, is designed to be dynamic and thus, has more configurable data protection options. At the disk level, disks can be encrypted or left unencrypted. Encryption is performed in software and does not rely on the more expensive self-encryption drives. Rather, encryption can be applied to any data drives in the kSAFE, including SAS, SATA, or SSD. The encryption is KMIP v1.1 compliant and can store encryption keys on an external, KMIP compliant key manager.

Within kSAFE, storage pools are created on top of the disks. Each storage pool has an administrator defined protection methodology. Currently, the two supported protection methods are replication and erasure coding. For replication, objects are protected by copying them to other disks, other storage nodes, or other racks,

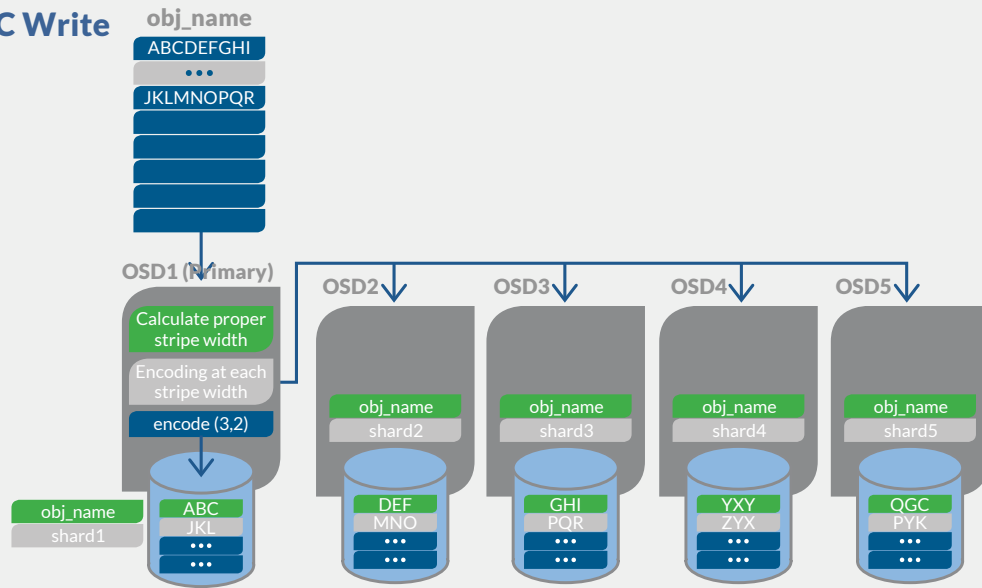
depending on the size of the system. The number of copies is configurable, with the default being three copies.

For erasure-coded storage pools, objects are encoded with redundancy information and spread across a number of disks, nodes, or racks, depending on the size of the system. The erasure coding parameters can be configured, but are already optimized by the system based on configuration.

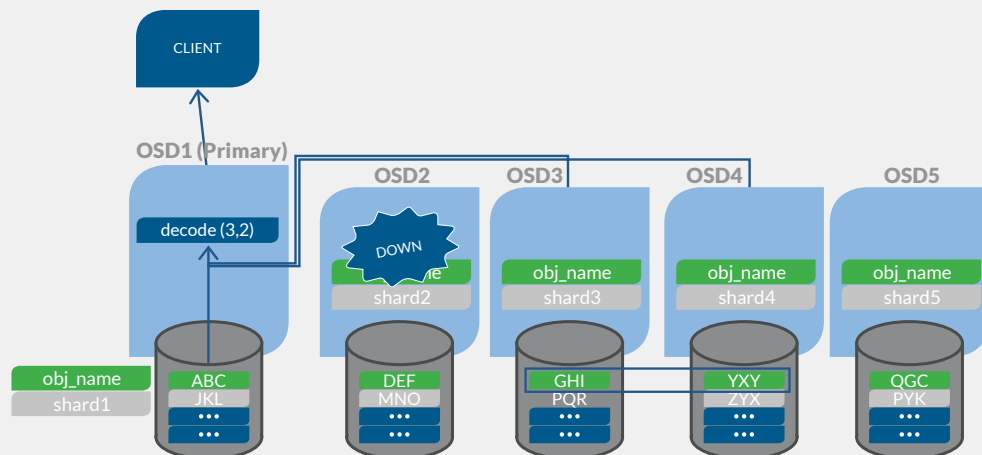
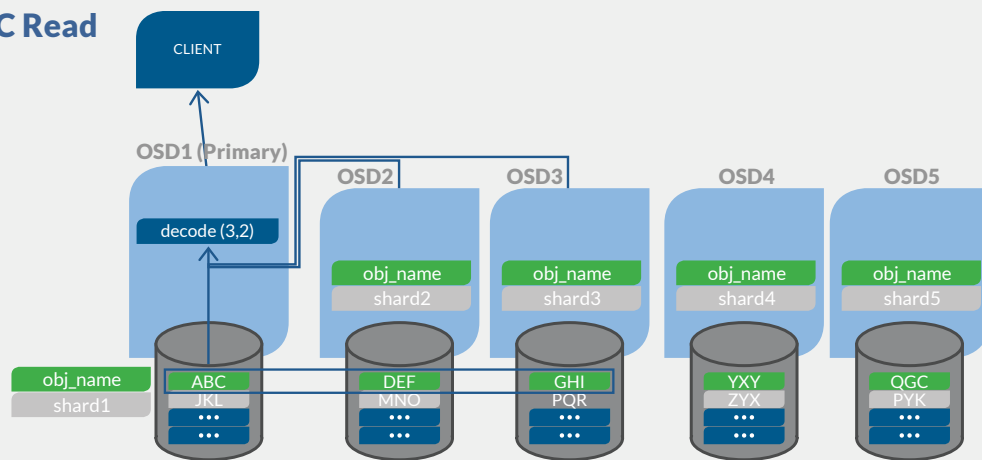
The trade-offs of changing these protection methods or parameters are both useable space and performance. Increasing the protection level decreases usable space, and the replication method of protection has better performance than the erasure-coded protection method.

Figure 7: Erasure Coding Parameters

EC Write



EC Read



Storage at the petabyte- and exabyte-level requires more than data replication. By utilizing erasure coding, data will be encoded in the primary OSD and then spread to the corresponding OSDs. This reduces storage requirements while maintaining the same availability, which will dramatically save on total cost of ownership.

If there is any data missing, Ceph automatically reads from the parity and then does the decoding.

Advanced Data Management with IDMT

Most organizations need a dynamic and flexible data management capability that can accommodate for policy and business need adjustment as both technology and economic requirements change. Keeper Technology provides one of the broadest solutions on the market for true-tiered storage. kSAFE with IDMT (Intelligent Data Management Technology) has the internal tiering and integration capabilities to provide seamless movement of data.

The policy-based tiering and data replication features of IDMT add data intelligence, management, and governance to digital repositories. Its powerful analyzer engine can index and manage existing data stores, while improving data efficiency, access, and workflow.

Comprehensive Tiering Strategy

Leveraging multiple storage tiers is accomplished by placing data on the tier that is most appropriate to meet business specific objectives. Storage tiers are typically defined by the cost and performance characteristics of the specific tier. These cost and performance parameters then dictate what data should reside on the target tier and for what time period.

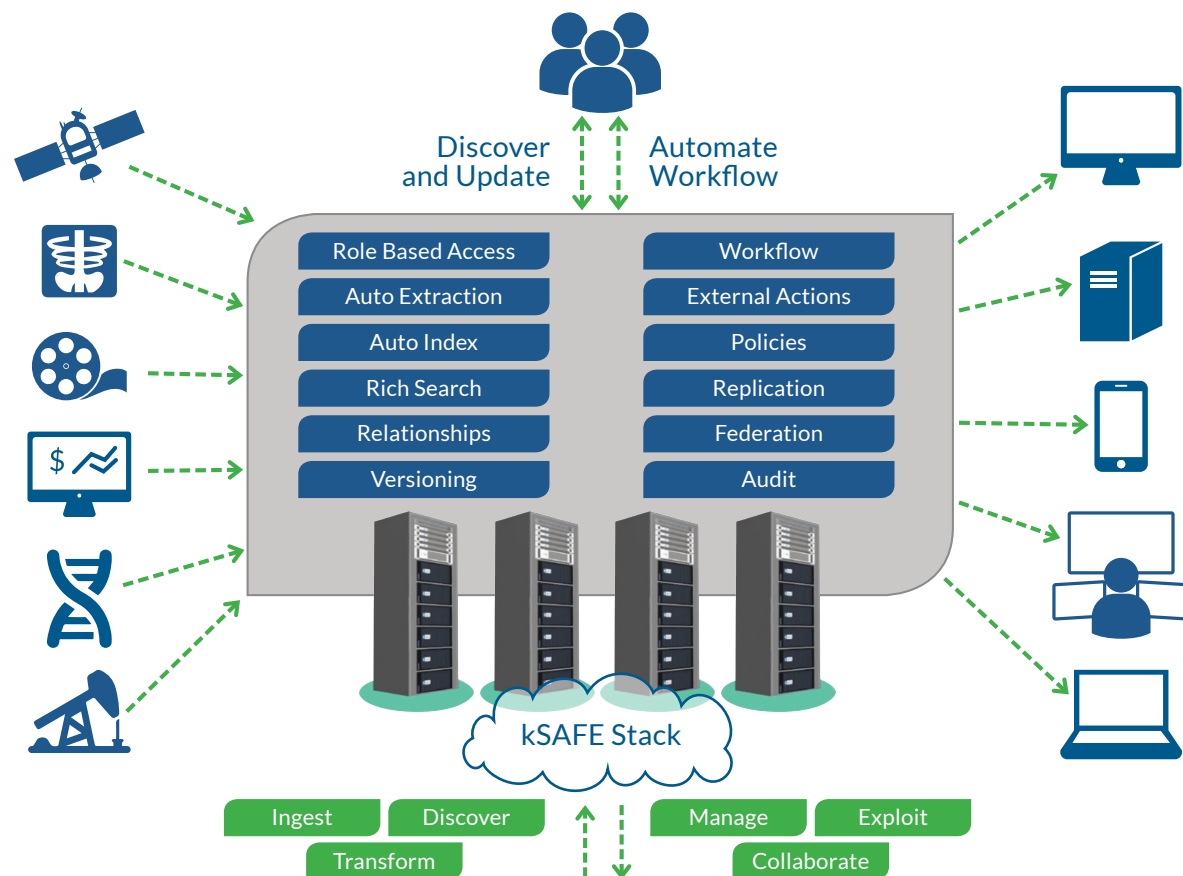
Many solutions are severely restricted on the data parameters that may be utilized for tier selection by making tiering determinations solely based upon file size, age, last access time, file name, or file extension. kSAFE with IDMT is unique in that any and all metadata characteristics such as: technical, business, user-defined, extracted, etc. can be used for determining where data should reside and why.

Replication and Federation

IDMT has a flexible, policy-based, replication capability. Since policies are driven by metadata, and all metadata is treated equally in IDMT, replication policies can be defined with nearly limitless capabilities. Replication can be used for data protection within a site, with practical critical alerts and real-time reporting, so that self-healing can occur when a bad object is found. Replication between sites and geographic regions provides data protection against regional incidents. Unalterable audit tracking and versioning is ideal for collaboration, security, and incident response.

Additionally kSAFE integrates with third-party appliances that can also perform policy-based data tiering, including replicating data to multiple disk storage repositories and replicating and/or relocating data to tape storage systems.

Figure 8: Advanced Data Management with kSAFE with IDMT™



Connectivity and Integration

Most large organizations have a variety of systems, applications, and solutions that require varying degrees of integration and access to a common set of data. They also have significant history and systems use, and require a wide range of protocols and interfaces to support that connectivity. Fortunately, very few organizations can match the variety of systems, legacy of applications, and variations of connectivity requirements of the Federal Government.

Since Keeper Technology has over a decade of experience designing, installing, and supporting storage systems in some of the most demanding agencies of the Federal Government, we offer an unparalleled understanding of such requirements.

The primary storage platform, kSAFE, supports a variety of access methods and protocols:

- **Block:** RBD, Cinder, Glance
- **NAS:** NFS (v3/v4), SMB (v3)
- **Object:** S3, http/ReST, Swift
- **File:** Clustered, Shared File System

The NAS protocols can simultaneously be used to access the same data. Likewise, the Object protocols can be simultaneously used to access the same objects. SAFEos, the Shared File System, and NAS protocols can simultaneously access the same data.

File Growth with Scale-out NAS

In traditional scale-up NAS storage scenarios, an enterprise would commit a significant investment to a large storage array, then gradually fill up that space over time. The kSAFE scale-out NAS approach breaks from this trend by enabling companies to expand capacity gradually. Instead of purchasing excess capacity up front, companies increase their investment in storage volume incrementally, which is both efficient and cost effective. There is no limit to storage capacity with scale-out network-attached storage.

Integration with Intelligent Data Management Technology (IDMT)

IDMT is a powerful data management platform, that when integrated with kSAFE, leverages the metadata locked within a digital asset. This metadata may be dynamically augmented through the addition of user-defined metadata and attributes. IDMT then enables customers to quickly access information for a given task, actively managing and protecting data in ways that also reduce infrastructure costs, and facilitating the broadest framework for collaboration, whether across incompatible systems, data types, or between remote locations.

IDMT has more than 20 different defined types that can be used for metadata. From the simple types (string, integer, Boolean) to more specific types (date, email address, URL) to even using other assets as a metadata type types.

The IDMT policy engine allows for advanced data management and tiering capabilities that manage and track data as it is relocated or replicated from storage tier to storage tier and from site to site. Complex searches can be constructed and performed, smart data policies can be built and implemented, and simple to complex workflows can be created and executed.

Seamless inter-operability with the Cloud

KeeperSAFE with IDMT functions as a gateway in a true hybrid environment. Keeper delivers a scalable cloud solution capable of integrating with RESTful protocols such as S3. Managing storage in house improves performance, access, scale and efficiency. KeeperSAFE efficiently leverages IDMT's policy engine to ensure data is easy to locate and stored in the best location for its value.



Management and Monitoring

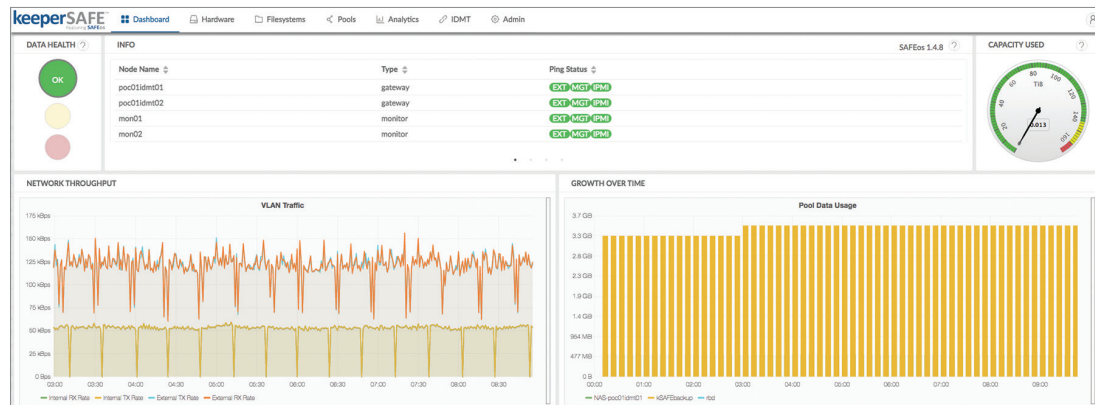
One of the unique things about combining open-source software with data storage hardware is that the open-source software has been written very generically so that it does not depend upon any one vendor's implementation of hardware.

The Keeper Technology software development team has worked tirelessly to not only fill in the gaps when open-source software is not available or viable, but to ensure that the hardware and software work well together. As with most products, error management and reporting is a key area that is required to know if a system is running properly.

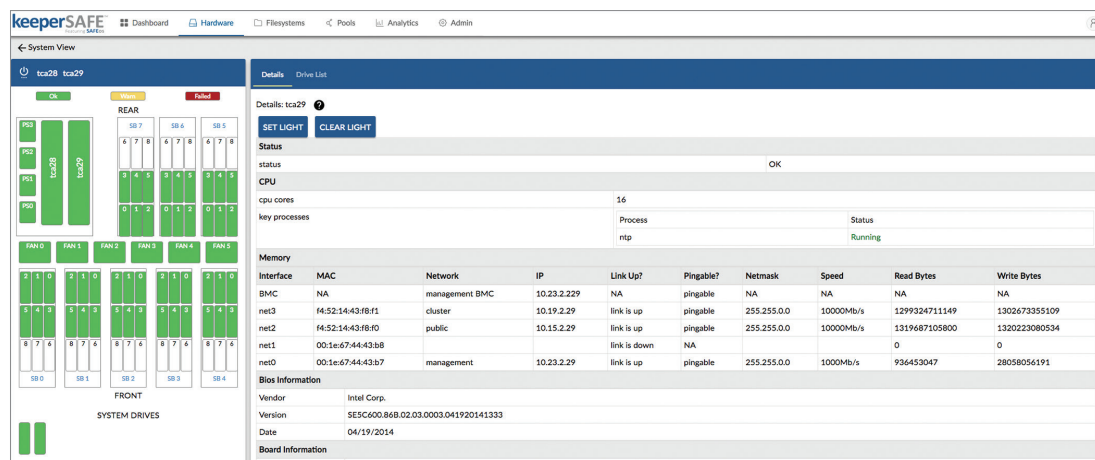
Keeper Technology develops and provides comprehensive enclosure management services to detect failing and failed components from disks, as well as power supplies, and system interconnects. That information is then promptly and accurately displayed on the kSAFE management GUI so a systems administrator can provide the appropriate response.

kSAFE collects tens of thousands of data points per hour, analyzes them, and rolls that information up into our intuitive management GUI that presents software status, hardware status, performance information, and a host of configuration and operational data.

Figure 9: kSAFE Dashboard GUI: Customizable, Unified Status of Hardware and Software Health



kSAFE's GUI dashboard provides a single screen overview of the health of the system from both the hardware and software perspective, including a glimpse into performance and capacity used.



The integrated management capability of the Hardware Tab allows for drilling down to a detailed view of the underlying component infrastructure, simplifying diagnostics, system expansion, and component replacement.

Management and Monitoring (continued)

With appropriate management modules for all node types in the system, the entire kSAFE can be installed, configured, managed, upgraded, and replaced from the active kSAFE Admin Node. The GUI is used for the vast majority of this, supplemented by the CLI for a few customer-specific configurations.

Software is installed and upgraded from the Admin Node. SAFEos is delivered as a single download and contains all microcode, firmware, and SAFEos packages needed to operate the system.

In addition to management modules, monitoring modules are developed for each node type to ensure as much data about the system as possible is available. This goes all the way down to parameters like power supply voltages, fan speeds, CPU temperatures, disk bit error rates and temperatures, etc. All of this data adds up quickly; as much as 15,000 data points or more are generated per hour for a single storage shelf in a kSAFE system.

Overall system health is monitored through data fusion of the information output from the OSS and the monitoring modules of the various node types. All of this information is distilled, prioritized, and presented in an easy-to-digest form in the kSAFE GUI. Furthermore, an SNMP monitoring station can be configured to enable receiving of SNMP traps or to support SNMP queries of the system.

Reporting and Analytics


The kSAFE GUI has the ability to generate basic reports about hardware in the kSAFE system. The Analytics tab within the GUI provides access to a rich graphing package that supports the generation of user-defined graphs from any of hundreds of parameters within the system.

Many of the thousands of telemetry data points collected per hour by the system can be added to graphs in support of many types of analytics.

Another form of reporting is based on audit trails. Audit trails of all types can be enabled on the system. All audit information is stored as immutable objects in the database and, like queries, can be accessed from the Desktop, the CLI, or through the API with audit-specific functions. The audit system enables complete, unalterable tracking of all actions.

Figure 10: GUI Analytics with Rich Graphing Capabilities





Conclusion

The digital content explosion is presenting organizations with unprecedented scaling challenges. Experts predict a 4300% increase in annual data generation by 2020. To manage this big, semi-structured, and unstructured data growth, while maintaining performance, capacity, and availability, requires ingenuity and long-term vision.

Keeper Technology has been delivering integrated technology solutions to commercial and government customers for more than a decade, particularly in industries that need to provide continual online access to data in the petabyte and exabyte scale. As we have grown, we have remained focused on our core mission to provide enterprise class, highly differentiated data storage and data management solutions to our government and commercial customers.

The keeperSAFE storage architecture is a robust, highly scalable block and object storage platform that augments the speed, ease of use, and ultimately the time to product completion on an exponential scale.

Unlike proprietary, hardware-based storage, which limits your ability to keep up with today's enormous data stores and requires significant investment in over-sized storage arrays, Keeper Technology gives customers a software-defined storage platform that scales across physical, virtual, and cloud resources.

kSAFE leverages open source, software-based storage, commodity hardware, and a Keeper developed management layer that delivers the highest levels of availability possible. The distributed nature of kSAFE is fully redundant with all critical components that are replicated, mirrored, or otherwise protected against failures.

What might take traditional storage systems days or weeks to add 2 PB's to the enterprise environment, with kSAFE customers can be ready to start writing data in under an hour. Our Ceph-based solution actively collaborates, replicates data, consistently applies updates, detects failures, and migrates data.

The bottom line is Keeper allows all of our customers to benefit from the agility, elasticity, and transformative nature of today's open-source software combined with qualified data storage hardware in a repeatable, supportable data storage appliance.

We have the technical talent, business experience, and long-term vision to help implement, optimize, and manage storage infrastructures as needs evolve. We are committed to helping our customers get the most value out of their investment in data storage, management and security by leveraging our expertise, dedication, and enterprise-grade services and support.

Notice: This whitepaper contains proprietary information protected by copyright. Information in this whitepaper is subject to change without notice and does not represent a commitment on the part of Keeper Technology. Keeper Technology assumes no liability for any inaccuracies that may be contained in this whitepaper.

Keeper Technology makes no commitment to update or keep current the information in this whitepaper, and reserves the right to make changes to or discontinue this whitepaper and/or products without notice.

No part of this document may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the personal use, without the express written permission of Keeper Technology.

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

keeperSAFE and keeperCARE are trademarks of Keeper Technology, LLC in the U.S. and other countries. Other names herein may be trademarks of their respective owners.