

What's Inside

- Simplify Storage Management 2
- Leveraging the Power of Metadata 3
- Simplifying Global Data and Storage Management 4
- Policy-based Data Management and Data Protection 5
- StrongLink Interfaces 6
- StrongLink Architecture 7
- StrongLink Scalability 8
- Technical Specifications 9

StrongLink: Data and Storage Management Simplified

StrongLink® is a powerful solution that combines AI (artificial intelligence) with metadata-based workflow engines to enable end-to-end management of both your data and your storage resources. This is a new and unique approach to managing your entire data environment that automates data classification and life cycle management for any data, on any storage, anywhere.

StrongLink leverages the power of metadata to create a single global namespace across any storage device from any vendor. This enables administrators to optimize storage resources without impacting user workflows, and enables users an easy way to globally search for, access, and manage data across any storage type, including cloud.

StrongLink eliminates silos and vendor lock-in

The problem with data and storage management is there is no one-size-fits-all storage solution that can address the entirety of the organization's data problems.

As data volumes soar, this often results in customers choosing to lock into a single storage vendor in the hope of minimizing incompatibilities. The problem is this also increases risk, limits choices, and usually means spending more than is necessary on storage infrastructure.

StrongLink was specifically designed to eliminate this problem, empowering you to globally manage your data using your existing storage, or bridging otherwise incompatible storage silos. StrongLink creates a logical abstraction of the physical storage within the enterprise and can bring together multiple file systems into a global namespace. This gives organizations global visibility into all storage resources.

How StrongLink Helps You



SIMPLIFY DATA AND STORAGE MANAGEMENT

Powerful storage provisioning tools and policy-based workflow automation give administrators complete control of all their data, globally managing it across any local, remote, or cloud storage type regardless of vendor.



REDUCE STORAGE COSTS

Optimize infrastructure costs automating data placement to right storage type for each class of data or use case, without adding complexity.



DATA PROTECTION

All files are double check-summed for data integrity, and all user or system activity is logged in an immutable audit system. Incremental file versioning enables file and file system rollback. Data provenance is assured to meet compliance and governance requirements.

StrongLink Core Capabilities



GLOBAL SEARCH

Search for any file or object based upon metadata, in an intuitive query interface average users can easily use. Searches can be local, or across a multi-site global namespace.



DATA MOBILITY

Reduce backups and increase availability by automating data migration between on-site and off-site storage, cloud and archive. Simplifies management of redundant copies for disaster recovery, or multi-site access.



ORCHESTRATION

StrongLink integrates the functions of storage resource management, storage service management, and data management into a unified platform for orchestrating data access across the entire enterprise.

StrongLink eliminates data silos so users can access any storage type via any protocol.

Access any data on any storage in a global namespace

By aggregating all data and metadata from multiple storage systems into a global namespace, StrongLink presents a virtual file system that spans across the entire enterprise.

- Users and applications see a file system hierarchy based upon their use case, permissions, and workflows regardless of which storage silo the data lives.
- The Cognitive Data Management policy engine automates secure data placement, to ensure data is available in the correct storage and location, at the right time, and without disrupting the user's workflow.
- Utilization patterns can be monitored to ensure that stale idle data is not stranded in expensive primary storage, and is automatically migrated to the right storage to accommodate workflows, data life-cycle, etc. Multiple copies can be centrally managed by policy for data redundancy, or archive.

All of this is done transparently to users and applications. Workflows are preserved with StrongLink's cognitive engines that ensure the right data is always available and accessible.

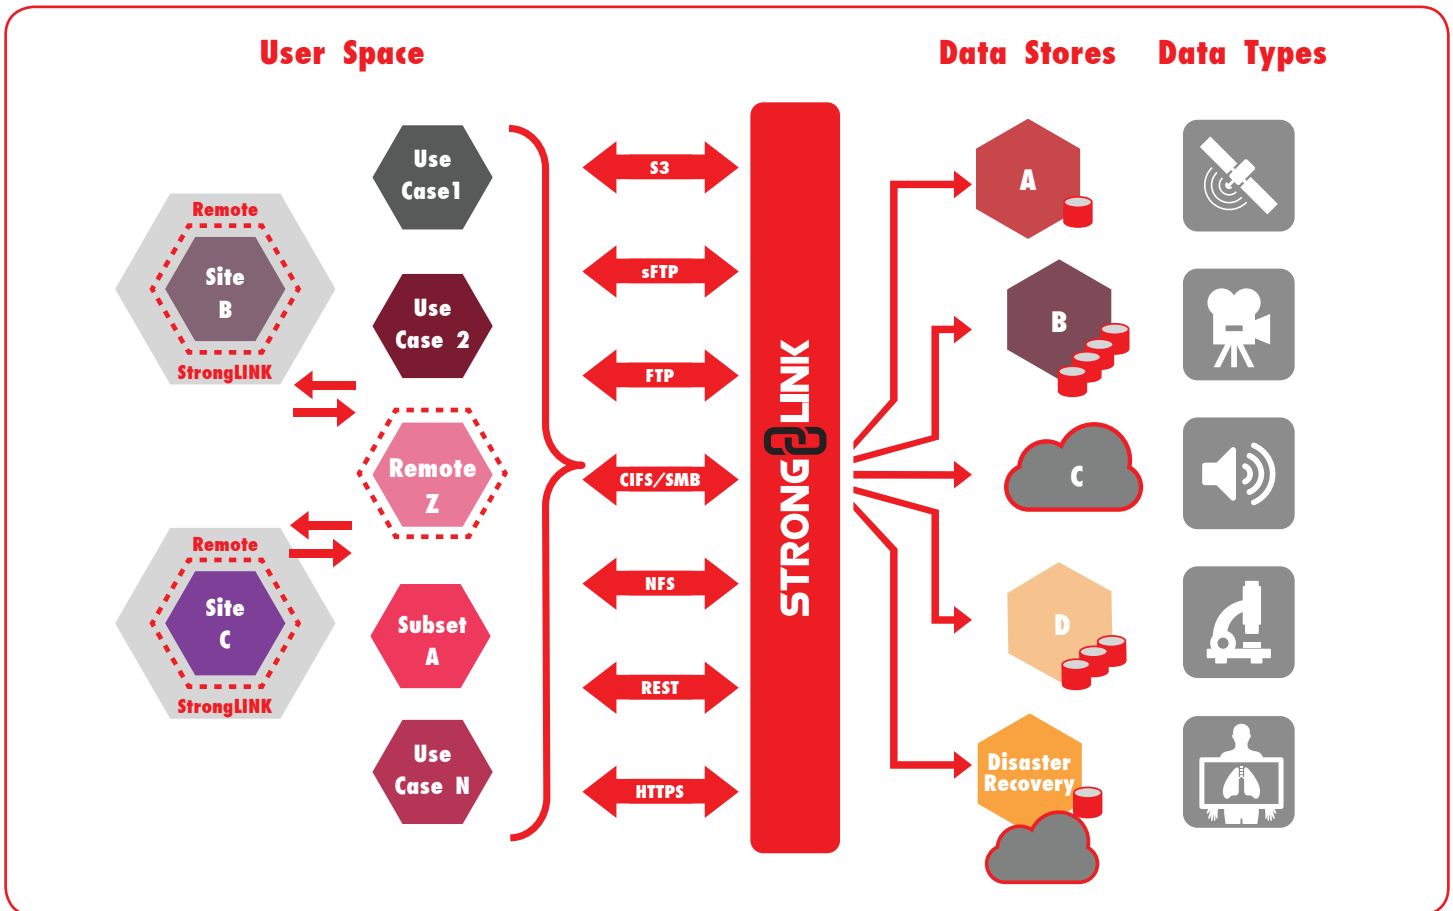


Fig. 1 - StrongLink enables any storage type to be accessed via any protocol, virtualizing all storage resources. Administrators can automate data placement and storage management without impacting users. Users simply see their files.

Rather than trying to physically normalize all data at the infrastructure level, StrongLink does this with metadata to enable global management across all storage silos.

StrongLink is powered by Metadata

StrongLink takes a data-centric approach, coalescing multiple metadata types into an aggregated management framework. This enables StrongLink to provide global control of all data across any storage resource, without needing to alter the underlying infrastructure.

StrongLink leverages metadata to automate and combine storage resource management, storage service management, and data management into a single unified platform for orchestrating data. This is what we call Cognitive Data Management; powerful automation tools for both users and system administrators to actively manage all their data, and the underlying storage infrastructure.

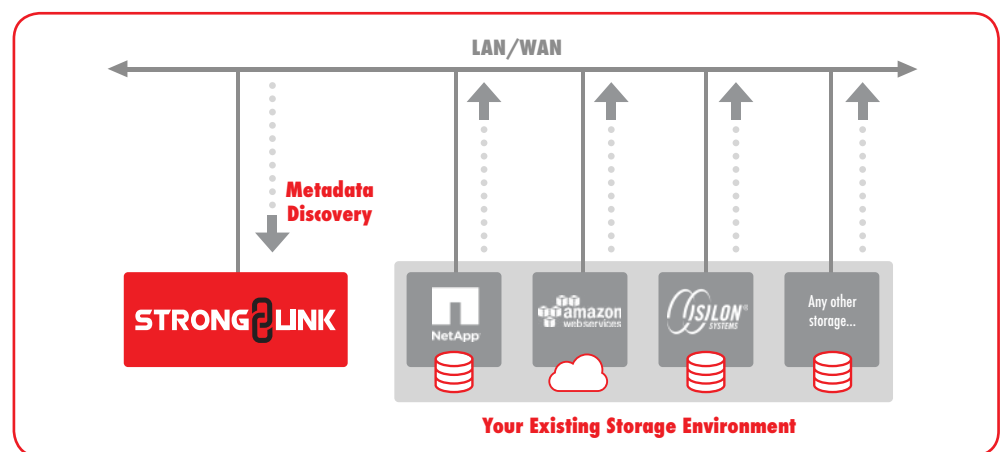


Fig. 2 - StrongLink harvests multiple types of metadata from files on any storage type, which can be used for search, policy-based data migrations, and much more, with a global view across all storage silos.

User-defined metadata

In addition to file system metadata and other rich metadata that can be harvested from file headers or contents, StrongLink also enables system administrators and authorized users to add custom metadata tags to files. When combined with the other metadata elements within a file or which may have been aggregated into the StrongLink learning engine from external databases or MAM systems, there is no limit to the ways users can use such tags to automate complex workflows. In this way, users can easily search for and manage content from different systems that might be related by project tag, use case, or any other variable required. Any metadata fragment can be part of a query, or used to trigger a policy or action.

- StrongLink automatically harvests multiple types of metadata from all the files on any existing file system or object storage. This data classification step creates a detailed picture of all existing data in every silo and storage location. In this way, StrongLink can virtualize namespaces of all or selected parts of the datasets, in addition to having rich contextual information about the data for use in building policies, resource allocation, and accommodating specific use case requirements.
- Users and applications can access the virtual namespaces with the same protocols they use today according to their permissions (AD, LDAP, etc.) or use cases. This provides a global view beyond an individual silo and no disruption to current workflows.
- Automate tasks such as storage tiering, dynamic capacity balancing, and non-disruptive data migration, completely transparent to users and applications.
 - Users simply see all data for which they have permissions all the time, even as it is optimized across different storage types.

StrongLink Simplifies Global Data and Storage Management

StrongLink is data-aware, enabling global data and storage management across all existing storage types.

Reduce Storage Costs and Sprawl

StrongLink optimizes storage resources and increases storage utilization. StrongLink's AI engine provides administrators with the metrics to monitor utilization, assist in capacity planning, and to conduct predictive analytics to enable proactive capacity planning.

- Predictive analytics to assess utilization across all devices, and enable proactive capacity planning.
- Global control of data placement policies to move data to the appropriate tier, or shunt it away from over-subscribed storage, transparently to users and applications.

The StrongLink process:

- Automated Data Classification
- Storage Discovery & Classification
- User and Project-based Quotas
- Storage Analytics
- Storage Pooling = StrongLink SmartPools:

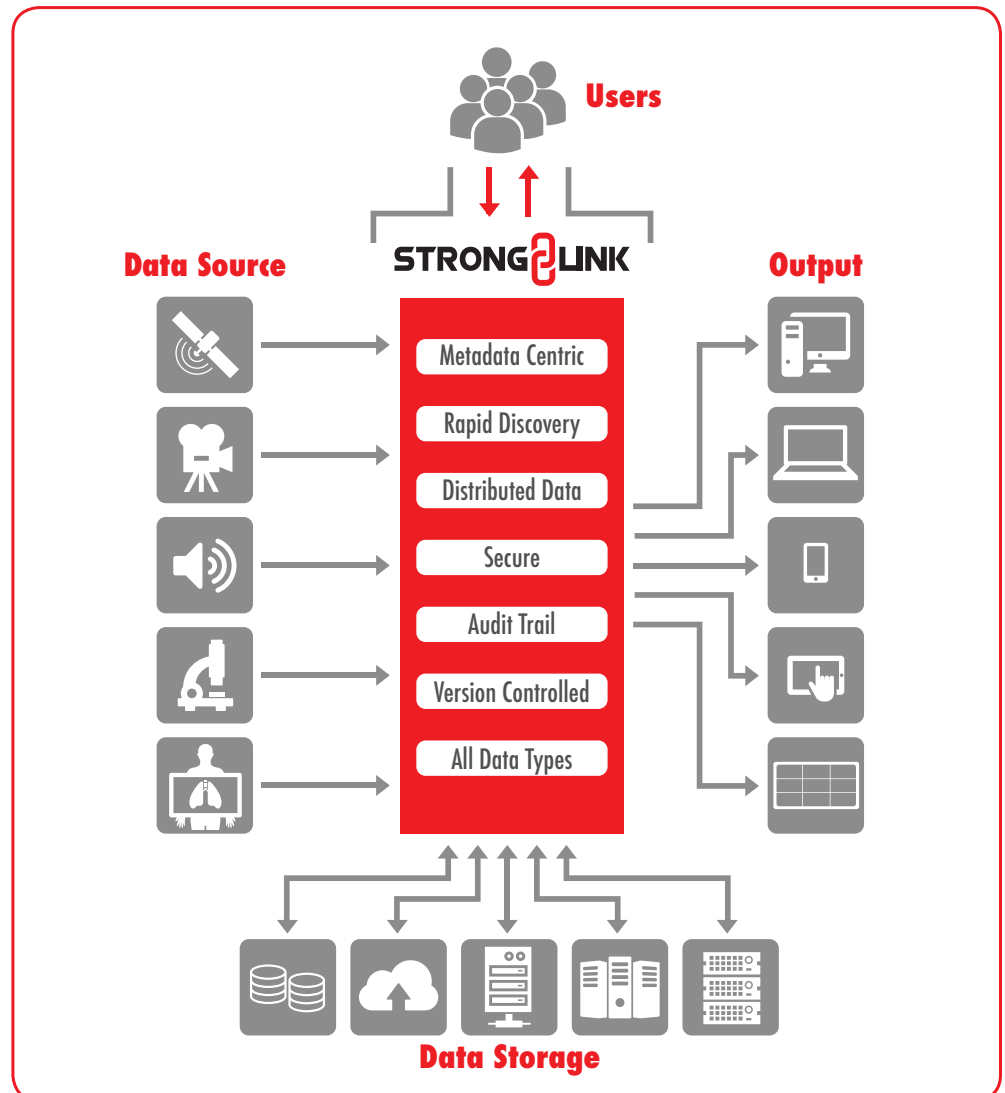


Fig. 3 - By aggregating all metadata types, StrongLink becomes data-aware. Data management and storage management may now be globally controlled at both the user and administrative levels.

With StrongLink, administrators can easily create granular policies based on storage, data, and use case, which would otherwise be impossible to do.

Policy-based Data Management

StrongLink's data-centric approach enables administrators to implement sophisticated data lifecycle policies to precisely manage data placement and data protection strategies at a granular level, even down to the individual file. Policies may be triggered by one or more metadata types.

- **Automated tiering, or data placement between any storage type or location**
 - Policies can be driven by any metadata type.
 - Enable active archives from any source to any target on/off-premise or cloud.
- **Automated workflow policies**
 - Pre-fetch data from nearline or offline storage to production storage based upon workflow profiles, project tags or other variables.
 - Copy or migrate data from one type of storage to another automatically, as part of a production pipeline or other workflow.
 - Predictive analytics to enable global resource management and pro-active planning.
- **Non-disruptive data migration**
 - Global namespace is presented as one or more virtual file systems, so infrastructure changes are transparent to users and their applications.
 - Add new storage types, or remove old storage without interrupting users. They simply see their files as before.
- **File Deduplication**
 - Each file and object has a unique id (UUID), and multiple hash codes to ensure no hash collisions could create a false positive. In this way, StrongLink can ensure a single instance policy within primary storage to reduce redundant copies of files, while also managing replica copies for DR or other protection strategies.

Data Protection and Provenance

Data protection is job number one for StrongLink, and as such the system includes multiple data integrity checks and data protection capabilities to assist administrators to globally protect all data under management.

In addition to automated data placement/migration policies, StrongLink enables administrators to manage multiple instances of a file across any local or remote storage types for redundancy, disaster recover, or to enable remote collaboration and accessibility.

- **Audit System**
 - The StrongLink audit system maintains by default an immutable record of all user or system actions.
 - Reports can be used to validate file provenance and integrity, and enable forensic analysis.
 - The audit system also supports compliance requirements, such as the EU GDPR regulations that will apply in 2018.
- **File Versioning**
 - StrongLink includes a powerful incremental file versioning capability.
 - File and file system versioning is optimized to reduce storage consumption. This enables users to roll back the clock to view files and file systems to a previous point in time.
- **Eliminate Traditional Backup with Automate DR & Active Archiving**
 - Data protection capabilities are built into the StrongLink for resiliency and appropriate redundancy.
 - Redundant copies of files may be automatically migrated to a disaster recovery site, cloud, or active archive platform.
 - Since the StrongLink system is both data-aware and storage-aware, such policies can be implemented differently for different classes of data or use cases.

Single Pane of Glass Storage and Policy Management

StrongLink includes a browser-based control panel and dashboard for storage resource management and policy creation for administrators, and an Explore view for end-users, according to their role. The StrongLink Control Panel may be accessed from any browser on any OS, without the need to install any local client.

- Real-time status of each StrongLink node, and storage resource.
 - Health status of each StrongLink server node, including I/O monitoring, and memory/CPU utilization.
 - Constellation status, including aggregate I/O and job queues.
- Storage management to provision new stores for both in-band or out-of-band configuration.
- Create storage pools or SmartPools, with simple drag and drop interface.
- Create granular data management policies based upon global or user-created metadata query filters. The system ships with a library of standard policies and filters. Administrators can customize these or create their own policies to support workflow based on any metadata fragment.
- Reporting templates are included in StrongLink to enable administrators to generate utilization reports for storage, user activity, and other system metrics. Administrators will be able to customize these reports, or create their own.



Fig. 4 - The StrongLink control panel enables global management of storage and data policies. A rich reporting engine can automate utilization and audit reports across all storage types, to provide the intelligence needed to enable pro-active planning.

Global Search

Users may access data stores via standard network file share protocols, exactly as they do today. In addition, the system also includes an integrated interface, which provides users with additional functionality to find, and manage their data.

- Browse all storage, whether external stores or managed stores, with unified access across different protocols.
- Drag-and-drop data between different stores.
- Users may browse all aggregated metadata attached to files, folders and namespaces.
- Metadata forms editor, so users may build their own custom metadata.
 - User-created metadata can be used manually or part of automated policies.
- Query editor, so users can easily build multi-variable queries based upon any metadata fragment.
 - No need for users to learn complex query semantics or call the IT department to find their data.

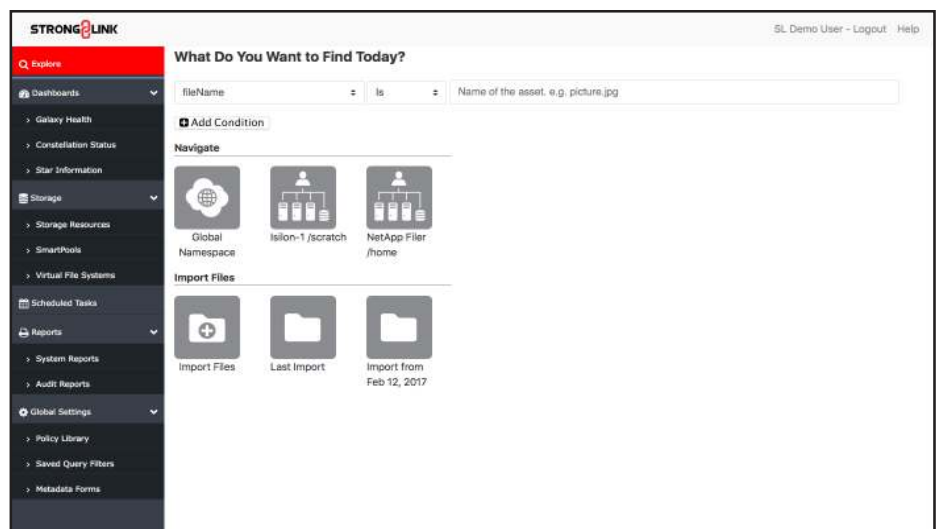


Fig. 5 - StrongLink enables users to create metadata, build queries, and take action across the entire global namespace.

StrongLink is designed to be flexible, scalable and fully resilient, with no single point of failure for always-on operation.

StrongLink Architecture

StrongLink is a no-master server application that is designed to fully leverage today's multi-core CPUs, virtual environments, and cloud infrastructure.


- Application-level self-healing capability that is far superior to OS-based HA solutions.
- A fully protected system that can operate 24x7, without the need for backup windows .
- Parallel, distributed architecture for maximum efficiency of multi-core processors.

Scalability Options

The StrongLink architecture is modular to enable systems to start small, and scale at any time to accommodate increased system load, changing I/O requirements, or other needs. Although StrongLink may be deployed as a single node, in most StrongLink deployments, nodes will be part of a Constellation, consisting of at least three nodes in a fully redundant and self-healing configuration, with no single point of failure.


- **Modular Configuration:**
 - No single point of failure
 - Can deploy in single node (Star) or multi-node (Constellation) or globally in a multi-site hyper-cluster (Galaxy).
 - Nodes can be dynamically added at any time to accommodate changing infrastructure or use cases.
- **Scale out and Scale Up:**
 - No capacity limits
 - No capacity fees
- **Performance Driven:**
 - Support 100's of billions of objects
 - I/O tunable
- **Self Healing:**
 - Resilient
 - Redundant

Star
(A single node)



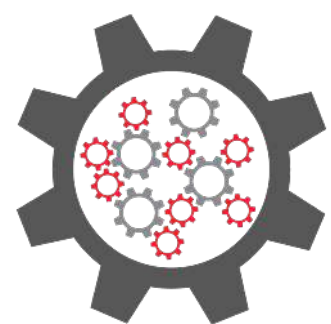
- StrongLink nodes are called Stars, and a single-node configuration is intended for small stand-alone use cases where high availability is not required, or where storage is connected out-of-band.
- Use cases for a single node might be as an automated active archive engine, search, cloud gateway, or other use cases that are out of the environment's primary data path.

Constellation
(Cluster of 3 or more nodes)



- A group of StrongLink nodes, or Stars, in a self-healing configuration, with no-single-point of failure.
- Constellations begin with at least three Stars, and may be expanded at any time to accommodate increased I/O or user requirements.
- In-band architectures, or 24x7 uptime requirements, will need a Constellation.

Galaxy
(Hyper-cluster)



- A StrongLink Galaxy is a loosely coupled group of Constellations in a global hyper-cluster.
- The Galaxy couples multiple sites together into a single global namespace.
- Galaxies are usually deployed to implement Disaster Recovery strategies, to tie together remote offices, or to improve quality of service when accessing the same data at multiple locations.

StrongLink can manage storage directly, for real-time versioning and data controls, or reference data from external stores, or both.

Connecting to Storage:

StrongLink will work with any storage type, whether file, object or cloud.

Storage may be connected and of the following ways:

- **In-Band, fully managed storage**

- Users and applications access data stores via any of the standard access protocols, including CIFS/SMB, NFS, S3, sFTP, HTTP.
- StrongLink will present all data as virtual namespaces, which can present any or all of the underlying storage silos or file systems via any protocol.
- In the in-band deployment, all I/O is managed through the StrongLink constellation, and as such all changes are updated in real time, and are protected by the resilient self-healing architecture.

- **Out-of-Band, storage managed by reference**

- StrongLink is deployed alongside existing storage.
 - Client systems continue to mount and access your storage directly, with your storage seeing StrongLink as another client system.
- This architecture is appropriate when StrongLink is used as a smart archival system, for data protection, or as a storage gateway to cloud, object or other types.

- **Hybrid, In-Band/Out-of-Band**

- A hybrid deployment includes both in-band and out-of-band storage connectivity to get the best of both.

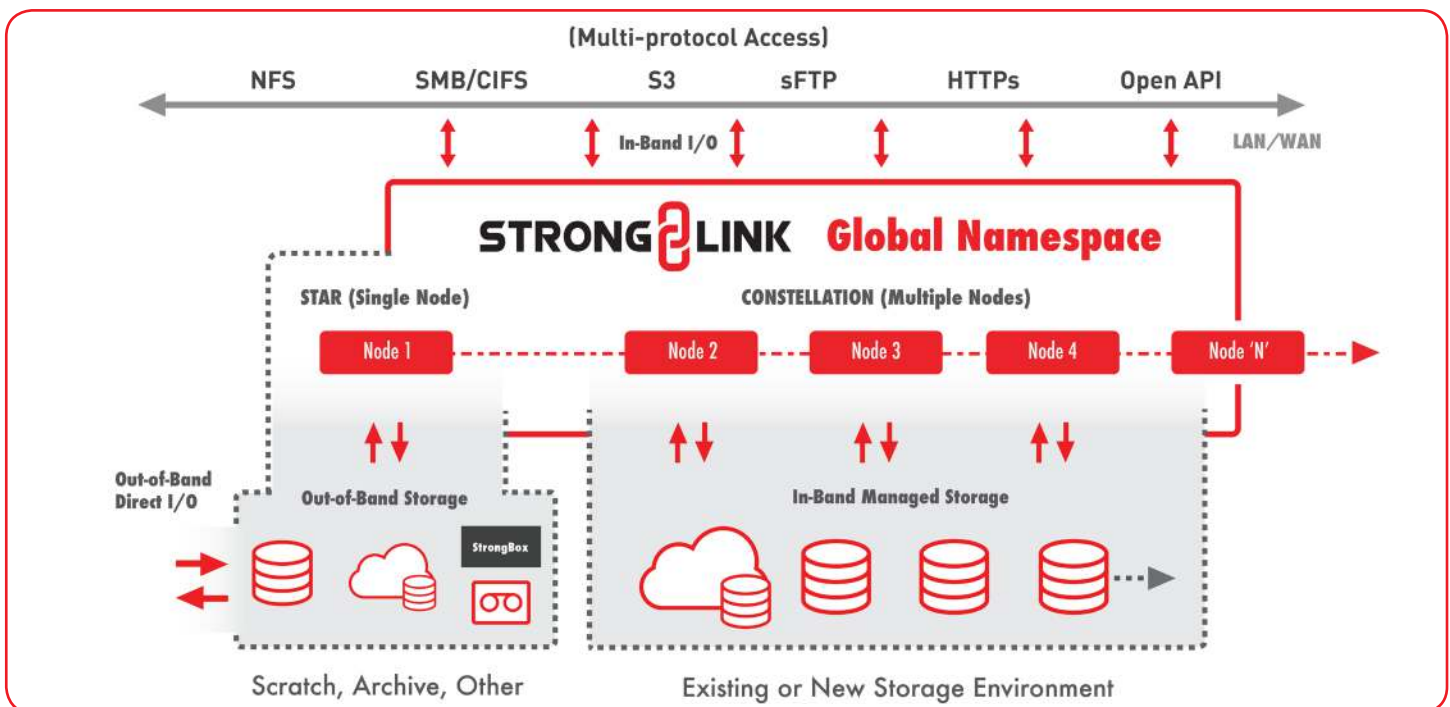


Fig. 6 - StrongLink may be configured either in-band, out-of-band, or as a hybrid architecture. The configuration can start with a single node for smaller use cases and expand as needed into a fully self-healing configuration.

StrongLink Technical Specifications

StrongLink is typically delivered fully installed on physical servers supplied by StrongBox Data Solutions. While StrongLink may be deployed at the customer site on a hypervisor, such deployments are not recommended for high-throughput use cases or for extremely large file sizes, due to the system overhead inherent in VM environments.



StrongLink Physical Specifications	
Form Factor	1U server
Dual Power Supply	550W
AC Power	100-240 V AC, Auto Ranging 50 Hz/60 Hz
Width	19 in.(434 mm, 482,4 mm w/ bezel)
Depth	23.9 in. (607 mm)
Height	1.685 in. (42.8 mm)
Weight	19.9 kg (43.87 lbs)
Internal Storage:	2 x 200GB SSD - for OS 4 x 400GB SSD - for Data
Remote Management Ethernet Port	1GE RJ45 – IDRAC
Dual 10GE Ethernet Port SFP+	(Front End Network)
Dual 10GE Ethernet Port SFP+	(Back End Network)
Quad 1GE Ethernet Port	RJ45 (one will be used for the StrongLink Web Interface)

StrongLink Software Minimum OS Requirements	
StrongLink Control Panel	Web Browser Runs on: <ul style="list-style-type: none"> - Linux - Mac OS - Windows Browser Versions: <ul style="list-style-type: none"> - Chrome - Chromium - Safari - Firefox

StrongLink Environmental Specifications	
Operating Temperature	-10°C to 35°C
Operating Humidity Range	10% to 80% relative humidity with 26°C maximum dew point (maximum wet bulb temperature)
Heat	2133 btu/hr Note: maximum inrush current is 55A for 10ms or less

STRONG LINK

Any Storage. Any Workflow. Anywhere.

For More Information:

strongboxdata.com/stronglink