

The Problem

There is a long-held belief that complexity drives cost in IT. This is true of storage infrastructure in particular, especially as all industries are struggling with how to reconcile the cost of storing, protecting, and managing increasing volumes of unstructured data. The problems of data management today are a direct result of the complexity that has been architected into storage systems since the Y2K frenzy in 2000.

At issue is the fact that there are many types of files - each of which have different business value over their lifecycle. Data structures today are inherently complex. Instead of being simple files in a folder, there is also rich metadata, and other external information about the files that is relevant to workflows and data value. Multiple copies of the files may need to be managed across different storage platforms for data migration, archiving, collaboration, and data protection. IT environments are also more complex. Customers have heterogeneous (mixed) storage platforms and different file systems. They also have data that lives in the cloud, on flash, tape, and on multiple storage tiers, each of which is optimized to only a portion of the data requirements at that moment. Files that are seldom used, for example, don't need to be stored in the same way that active data are. Some subsets of data may need very high-performance all-flash arrays for a period of time, while others may work well with object storage, tape, or cloud. While administrators can manually move data across different storage to meet evolving requirements, this process can interrupt users and adds IT management overhead.

The problem is that there is not a "one-size-fits-all" storage solution that addresses all of these data requirements. Even with years of advancements in storage technologies, the management of storage infrastructures across multiple vendors platforms has only gotten more complex, especially as IT organizations have to do more with static or shrinking operational budgets. It is against this environment that users and administrators are crying out for simplification. Hierarchical Storage Management (HSM)

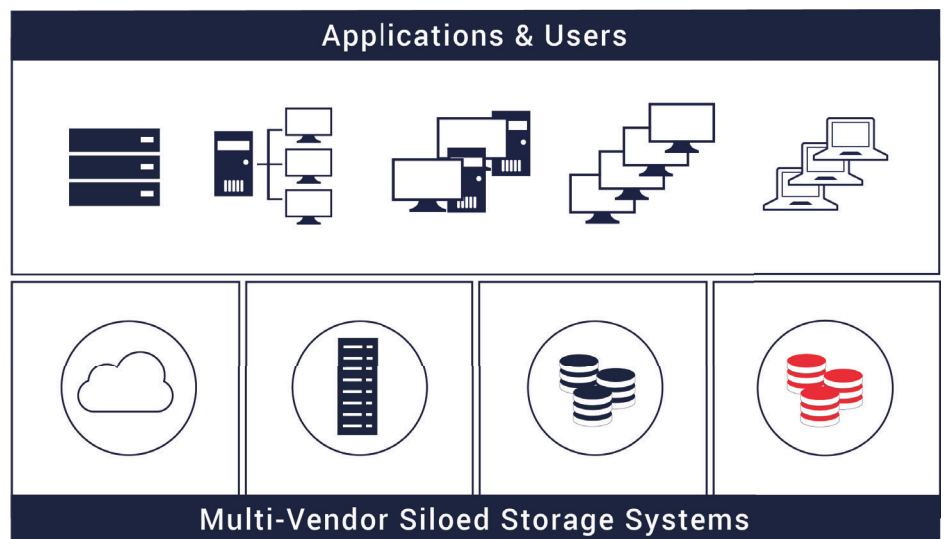
This white paper discusses how autonomous data management removes complexity from heterogeneous storage environments while reducing OPEX by 80%

solutions were one way to help this: Software would automate migrating data from primary online storage to offline or nearline disk or tape. But classic HSM solutions are no longer sufficient to address modern data workflows.

Storage vendors try to solve this complexity by offering end-to-end storage-centric “ecosystems.” This only adds to the problem, creating a vendor “silo” that locks users into a particular file system or vendor’s solution stack. This type of solution might work fine for users that can commit to a single, one-vendor solution for all their storage needs. But such vendor-locked solutions typically come with a premium price tag, and significantly limit customers’ freedom to take advantage of other storage technologies that may emerge. Most users have several different file systems, and multiple vendor solutions, which are incompatible, and don’t work well together. Too often IT organizations are left with no choice but integrate their own mix of storage solutions, which adds risk, complexity, and significantly adds to operational expense.

Today data is managed by the storage silo in which data is trapped, resulting in:

- ✓ **Complexity**
- ✓ **Decreased IT Workload**
- ✓ **Lack of Data Visibility and Value**
- ✓ **Reduced Productivity**



Cost of Ownership and ROI

These days, IT departments are finding their data storage strategies subjected to the same rigorous accounting scrutiny as are the non-IT initiatives of other business units. Concerns about total cost of ownership (TCO) and return on investment (ROI) are pressuring IT planners to contextualize their funding requests with utilization and other empirical metrics.

From the perspective of TCO, there is good reason to put storage in the crosshairs. Depending on the analyst you want to reference, storage accounts for between 33 and 70 cents of every dollar spent on IT hardware annually. Moreover, analysts estimate that the demand for storage capacity of unstructured data is growing at an unprecedented rate -- depending on the user over 200% per year, rather than the 40% annual rate that was estimated in 2010.

Not surprisingly, the cost trajectory of storage acquisitions has attracted significant attention. Tactical stop-gap technologies like compression and deduplication have been enlisted to squeeze the maximum amount of data into the minimum amount of storage space. However, these techniques do not address the real source of storage expense, which is only partially explained by equipment and software costs.

In fact, the cost to acquire and deploy storage is only a fraction of the cost to own and operate complex storage infrastructure. By some estimates, annualized storage TCO has a 6:1 cost ratio – with the cost to acquire, license, and deploy representing only 1/6th of the total cost to own and operate storage infrastructure (calculated on an annual basis). The other 5/6ths of the cost to own the infrastructure consists of management and administration (labor costs), backup and archive, facility costs including floor space, power and cooling, and other day-to-day operational expenses.



The real storage TCO impact is usually operational expenses (OPEX) -- administration and management, facility costs, data management costs -- more than the capital expenses (CAPEX) – acquisition, licensing and deployment costs. This fact is inescapable and makes it imperative for companies to do whatever they can to drive down OPEX in storage.

The increasing number of contemporary storage choices, whether monolithic or software-defined, actually makes administration and management more challenging, not less. Greater complexity in storage requires administrators to possess greater skills and experience which costs more money for

IT resources. Integrating disparate storage solutions into unified topologies and moving data between them requires manual intervention, scripting, and other workarounds. And every maintenance upgrade or technology refresh, even minor ones, carries with it a risk of downtime, disruption to users, and more hands-on effort by IT staff.

At issue is the complexity of the solution. A more complex environment requires more attention by storage admins or greater dependency on the vendor. Complex value-added functionality can actually obfuscate the global management of infrastructure and data, driving more manual processes. Complex, manual data management processes generally require more operators, driving OPEX higher. And, complex infrastructure, administered by more operators, is generally more prone to expensive downtime, which requires more complexity in the form of data protection and preservation.

Object-based storage and cloud solutions are often mentioned at this point as a savior of storage TCO, but for companies with diverse use cases and workflows, cloud/object storage alone is not the answer. Users have opted to leverage clouds for specific applications or as extensions to their infrastructure. Adding a cloud silo to an already complex infrastructure often increases, rather than decreases, complexity and TCO.

“The reality is that most IT organizations simply have no way to measure how efficiently their most expensive storage resources are being utilized”

-Marc Staimer | Dragon Slayer Consulting

Achieving a Measurable Improvement in ROI

Storage infrastructure and data management are cost centers to every organization, regardless of the data type or use case. Even -- or perhaps especially -- when data is the company's primary asset, the ROI for how those assets are stored and protected directly impacts the bottom line.

This goes beyond the way ROI is often calculated on storage vendor web sites, which may only measure how an individual storage element compares with an alternative solution. What is typically omitted from such ROI calculations is the impact on complexity for IT staff to manage another storage type, or disruption to users when their data moves from one platform to another. Looking at such measurements in isolation of the whole picture is a storage-centric attempt to answer data-centric problems.

The oft-quoted business maxim that “you can’t improve what you can’t measure” applies directly to this situation. The reality is that most IT organizations simply have no way to measure how efficiently their most expensive storage resources are being utilized across the enterprise. It is easy enough to know how fast the data is growing in each storage tier, and when more storage will need to be added. But with most storage tools today, it is very difficult if not impossible to provide detailed real-time metrics of storage utilization across all storage types, active vs passive data under management, project or departmental storage cost allocations, and other information necessary to optimize data placement across storage resources, and thus to maximize storage utilization.

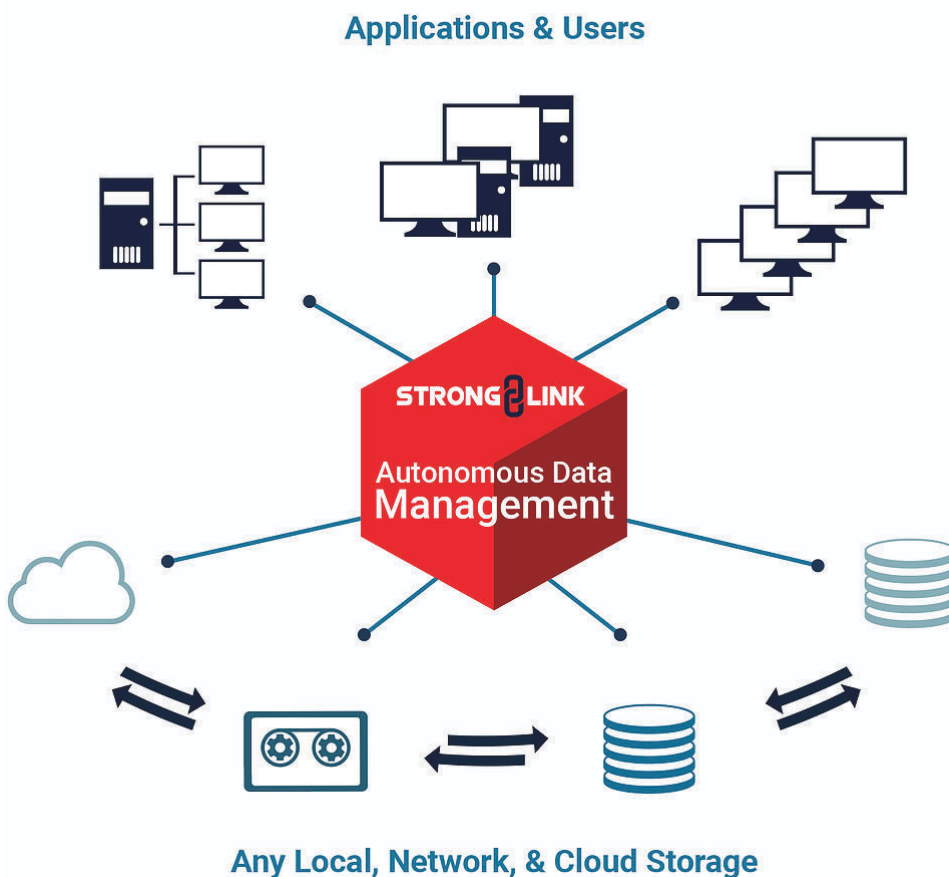


Figure 3. StrongLink intelligently manages data across file, network, cloud, and tape regardless of vendor.

Autonomous Data & Storage Management is the Solution

The StrongLink software platform solves these problems, removing complexity in heterogeneous storage infrastructures. StrongLink eliminates the complexity that consumes IT teams’ resources, drive down costs while enforcing data protection. Bridging otherwise incompatible file systems and storage types into a global namespace, StrongLink maximizes storage utilization, automates data migration, tiering, and protection. Most important, StrongLink provides visibility into the files and their metadata providing real-time,

dynamic reporting and data insights. This visibility is key so users can understand their current storage environment for utilization and efficiency, predictive behaviour to maximize storage services and resources, and extract data value to increase productivity.

StrongLink Data Insights enables:

- ✔ Visualization and statistical analysis
- ✔ Automate pattern discovery
- ✔ Pattern-based prediction
- ✔ Root cause analysis
- ✔ Increases user productivity
- ✔ Maximize storage utilization & efficiency

Files from all storage are connected and are presented in the StrongLink global namespace with a single file system view. Users maintain a persistent view of their data at all time and only see data based on permissions mapped from their AD credentials. This means users can simply focus on using their data, and stop wasting cycles trying to find or retrieve it.



Figure 4. StrongLink delivers visibility and data insights across all storage types, from any vendor.

StrongLink puts customers in control, combining metadata-driven insight and powerful policy engines, to solve the problems related to managing multiple storage protocols and use cases. StrongLink is Data Centric, it captures, normalizes and aggregates metadata from all data sources and uses it to drive the StrongLink Autonomous Engines. Data services are driven by the Autonomous Engines with user defined policies. Workflows can be automated to meet data life cycle management, global visibility enables collaboration, and data protection and provenance is always enforced with automated data verification and audit trails.

StrongLink is Storage Aware, it maximizes storage utilization efficiency by automatically placing data on the ideal storage according to IT-defined policy. Data life-cycle, tiering & migration is dynamically

managed by the Autonomous Engines that evaluates the state of the data, the status and availability of infrastructure, and the storage services all within the context of the data's policy. StrongLink dramatically reduces the number of operators required to manage the infrastructure while ensuring that you get the most from the capacity that you have. It allows IT administrators to automate data policies across storage tiers and locations. It gives IT the ability to perform seamless migrations and tech refreshes without interrupting user access.

StrongLink Autonomous Engines

- Query engine** - Global queries across all storage
- Analytics engine** - Life-cycle management, QoS, & capacity optimization
- Visualization engine** - Data Insights & reporting
- Dynamic Data Mover** - Migrate, I/O balancing, tiering
- Metadata engine** - Extract, auto-classification, tag, organize & manage
- Workflow engine** - Automated policies based on schedule and triggers

StrongLink makes storage scaling significantly simpler. Whether you wish to expand the capacity or add a different type of storage to your infrastructure, StrongLink enables you to seamlessly incorporate the new resources into its managed pools, adjusting the distribution of data dynamically and in accordance with data management policies. StrongLink also significantly reduces OPEX, by automating tasks that were once performed manually, such as re-distributing data across new capacity.

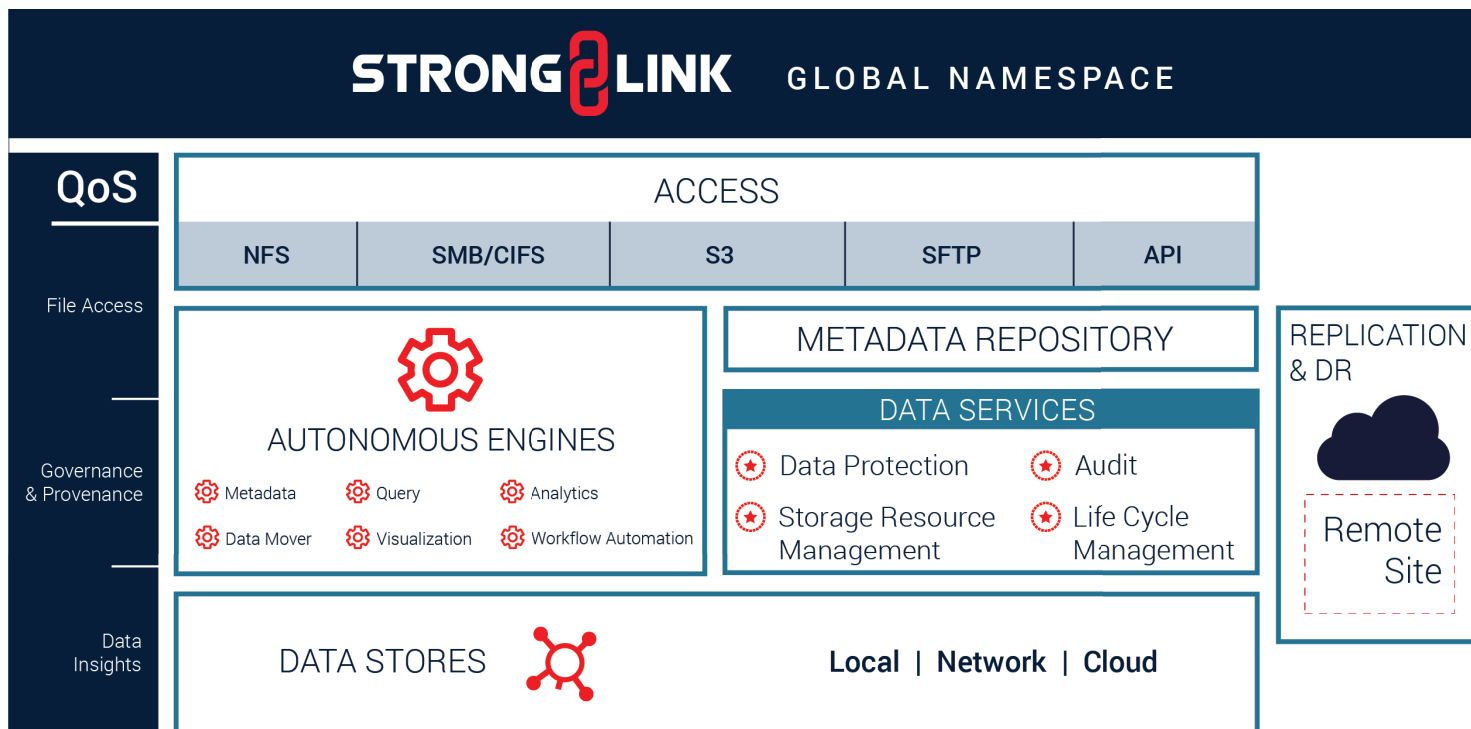
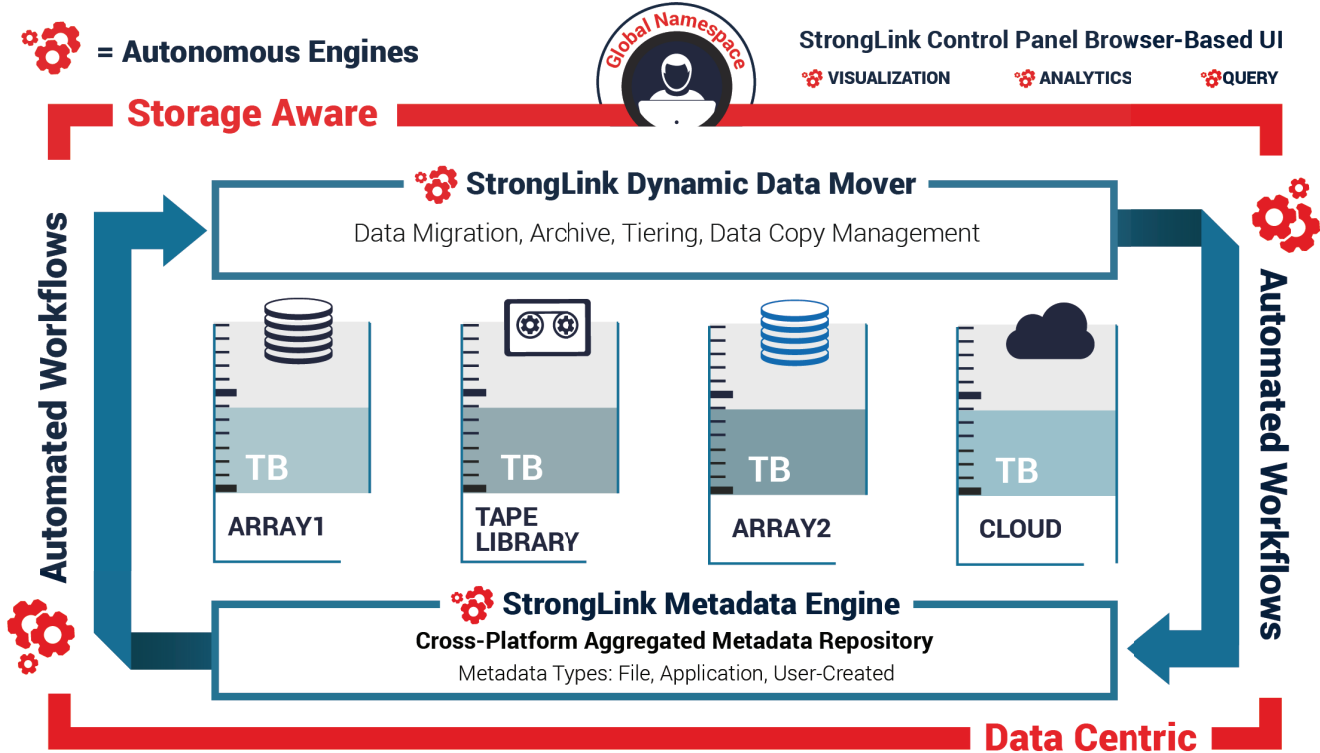


Figure 5. StrongLink architecture



StrongLink provides data protection, preservation, and privacy services via its own stack, but it can also work with third party backup, archive, or security software that you may currently be using. Ultimately, StrongLink associates data with the “best fit” protection, preservation and privacy services available in your infrastructure automatically, according to policies and priorities you decide. And, StrongLink’s intelligent, metadata-driven data management capabilities ensure that no single source of services becomes overloaded by too much demand.

Conclusion

Removing complexity and maximizing the ROI of storage and data management requires relentless (and continuous!) simplification, as well as intelligent automation. Visualization and analytics for data insight provide users control of their data and environment. Users requires a solution that simplifies two huge data problems: Complexity of multi-vendor data environments and, IT burden. StrongLink spans all storage types and vendor platforms, thus facilitating simple administration, optimizing capacity allocation and efficient storage utilization, and enabling rapid, agile provisioning with metrics in response to business needs. StrongLink is the solution to meet these requirements. StrongLink removes complexity, reduces cost and enforces data protection.

Contact us at: +1 512.349.0300 | info@StrongLink.com