

Slash Your Storage Costs for Good

Komprise Data Management with Western Digital Cloud
Object Storage

Executive Summary

Your data is among your organization's most precious assets. Too often though, it also represents a huge waste of resources. How can both be true? Because the value of your data changes as it ages. What's critical, active and valuable today may or may not be down the road. And yet, organizations hold onto more of their data, for longer, than ever before.

In fact, up to 85% of a typical organization's data may be "cold," inactive data that hasn't been accessed in months.¹ But in most cases, that cold data is stored and managed the same way as "hot" or active data. There's a good chance that, right now, you're spending hundreds of thousands of dollars annually—maybe millions—on high-performance storage for data you rarely touch. And, with enterprise data doubling every two years, the cost is only going to go up.

If you're going to keep pace with data growth without blowing your budget, you'll need a way to store cold data long term, without using the same storage, at the same cost, as you use for active data. The problem: most legacy storage solutions have no way to identify inactive data and automatically move it to a less expensive long-term storage tier. Or at least, they can't without breaking file hierarchies and disrupting processes for users or applications that need to access older data.

It's time for a fresher, smarter approach to data management. And it starts with a simple two-part solution:

- **Use Western Digital cloud object storage for archived and unstructured data.**

Cloud object storage gives you great performance, excellent durability and high reliability for inactive and unstructured data—for a fraction of the cost of high-performance NAS storage.

- **Use Komprise Data Management to lower cold data costs by up to 70%².**

Komprise software analyzes your data growth and usage across current storage to show how much you can save by moving cold data to cloud object storage. Then, it continuously, automatically moves inactive data to your cloud object storage tier, based on policies you define—without disrupting your users or applications.

This paper provides a technical overview of the combined Western Digital and Komprise solution, detailing how it can help you slash storage costs and manage exploding data growth. Read on to learn more.

¹ Based on Komprise customer demographics. Details at: <https://www.komprise.com/wp-content/uploads/2018/03/Biz-Value-Report-final-web.pdf>

² Savings calculated using Komprise customer-reported averages for current storage costs.

The Growing Data Management Challenge

Data is growing faster than ever before. Nearly 90% of all the world's data was created in just the last two years³. Analysts project the world's data generation will reach 163 Zettabytes—that is, more than 16 trillion Gigabytes—by 2025⁴. Much of this data is, or at least could be valuable. The question is how to retain it all while keeping within flat budgets.

Most organizations recognize that much of their data—potentially 50 to 85%—is infrequently accessed within months of creation. And yet, they continue storing and managing cold data the same way as active data, with the same expensive high-performance NAS storage. IT leaders would like to explore more affordable scale-out storage options, such as cloud object storage. But to date, there have been no simple solutions to identify cold data and automatically move it to cloud object storage tiers, without disrupting users and applications.

Legacy storage solutions just weren't designed for the kinds of data management challenges organizations now face. Typically, they are:

- **Too costly:** They require expensive enterprise licenses and upfront infrastructure investments.
- **Too complex:** Most have multiple moving parts to manage, such as storage agents, hardware, software, databases and more.
- **Too brittle:** Legacy solutions rely on static stubs that can be corrupted or orphaned, agents that need to be updated as the storage evolves, and in-depth rules that specialists need to define and maintain.
- **Too disruptive:** Many legacy systems introduce a significant performance hit due to the management overhead they generate. And in almost all cases, they can't move data to a different storage tier in a way that's transparent to users and applications.

It's time to shift to a modern storage architecture that can handle today's massive and growing data scales, and that can store data more intelligently and cost-effectively.

³ IBM Marketing Cloud. "10 Key Marketing Trends for 2017." 2017. <https://www-01.ibm.com/common/ssi/cgi-bin/ssialias?htmlfid=WRL12345USEN>.
⁴ IDC. "Data Age 2025: The Evolution of Data to Life-Critical." April 2017. <https://solutionsreview.com/data-management/idc-data-creation-to-reach-163-zettabytes-by-2025/>

The Smart Storage Solution: Cloud Object Storage with Intelligent Data Management

Now, you can slash storage costs by continuously, automatically moving inactive data to a cloud object storage tier—without disrupting your users or applications. By combining Western Digital's ActiveScale™ cloud object storage system with Komprise Data Management software, you can take advantage of the scalability, performance, and economics of private cloud object storage for more of your data. Komprise transparently moves your data from high-cost primary storage to low-cost cloud object storage, while preserving the file-based access on which your file hierarchies, applications and users rely.

With adaptive, analytics-driven intelligence from Komprise, you can automatically move cold data to more efficient ActiveScale cloud object storage. And you can redirect your IT resources to initiatives that drive more value for your business.

How Much Can You Save with Smarter Data Management?

A typical business* using Komprise Data Management and Western Digital ActiveScale cloud object storage System can realize up to:

- **88%** cost savings on cold data storage
- **57%** three-year overall storage savings
- **\$10.4 million** three-year savings
- **Immediate ROI** on the updated storage investment

* Based on an average 4-petabyte NAS environment with 60% cold data.¹

Inside Komprise Data Management

Komprise is an intelligent, analytics-driven data management solution that analyzes data growth and usage across your current storage to project the ROI of moving inactive/cold data to secondary cloud object storage. Komprise then automatically, continuously moves data based on policies you define. Your data is moved transparently, so users and applications can continue to access the moved data the same way they did before. By automatically moving inactive data to less expensive cloud object storage, Komprise is helping organizations around the world manage their data growth and slash cold storage costs by more than 70%.

Unlike legacy storage management solutions, Komprise is:

- **Simple:** Komprise is simple to deploy and operate, requiring no proprietary interfaces or complex infrastructure setup.
- **Open:** Komprise works using open standards — NFS, SMB/CIFS and REST/S3 — without requiring you to employ proprietary stub files or agents.
- **Vendor-agnostic:** Komprise is built on open standards and works with any storage supporting those standards. You can keep using your preferred vendors, even as you shift to lower-cost cloud object storage for inactive data.
- **Analytics-driven:** Komprise captures analytics from deep insights into your organization's data usage and growth. It uses these analytics to deliver an ROI-driven approach to data management, and to optimize your storage based on your unique data needs.

- **Transparent:** Even after Komprise moves data to cloud object storage, the data remains fully accessible from the sources as files, exactly as before, and as files or objects on the target.
- **Scale-out:** Komprise scales out on-demand. There are no central bottlenecks, databases or servers to limit scalability.
- **No lock-in:** Your data is always accessible from your source storage and your target storage, even if Komprise is taken offline.
- **Software-based:** Komprise does not require dedicated hardware or upfront infrastructure investments.
- **Non-intrusive:** Komprise analyzes and manages data in the background, outside the hot data and metadata paths, with no impact to storage and network performance.
- **Adaptive:** Komprise throttles back as needed when your storage or network are in active use, so you never have to monitor or schedule when Komprise runs.

Komprise Architecture Design Principles

Komprise's advantages derive from its unique architecture, which is designed around four basic principles:

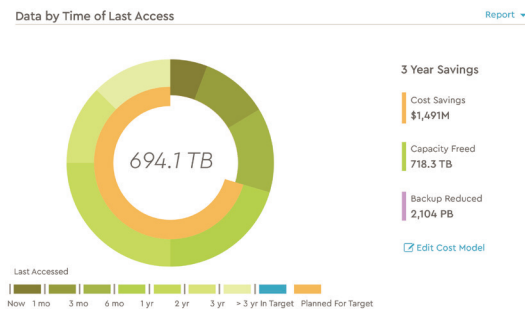
- **Distributed scale-out architecture:** A distributed scale-out architecture allows Komprise to scale with today's massive data growth. Start with one Komprise virtual machine and simply add more to scale, balance load and create fault-tolerance. There are no central databases or bottlenecks to limit scalability. Resiliency and high-availability are built in without requiring any dedicated infrastructure.
- **Storage-agnostic and non-intrusive:** Komprise is storage-agnostic and non-intrusive, so it can work with any environment and adapt as those environments change. Komprise is built on open standards and moves data without putting agents on your storage or requiring you to use proprietary, error-prone static stubs. Instead, Komprise uses dynamic links, which are highly resilient and native to filers. Komprise also adapts to your environment, so it runs invisibly in the background —without interfering with active usage of your primary storage.
- **Redundant, transparent hierarchical file system:** Komprise uses the data sources and targets themselves to create a redundant, overarching namespace. By using open standards and creating a distributed design, Komprise provides unified visibility and management. All permissions and attributes are fully preserved, and no agents or changes to the source NAS are needed. Moved data is accessible as files from the source and as files/objects on the target, so there is zero lock-in.
- **Distributed search and analytics:** Komprise provides deep insight through a distributed search and analytics framework on all metadata across storage. This empowers you to gain a deeper insights and greater control over your data. At the same time, the system itself uses those insights to optimize and drive how Komprise moves, archives, replicates and manages data to suit your specific environment and your unique data needs.

How it Works

Komprise uses an easy-to-deploy, easy-to-scale distributed architecture. The architecture consists of one or more Komprise Observer virtual machines running at the customer site connected to a Komprise Director, which can run either as a cloud service or on premises. Komprise works across any on-premises NFS, SMB/CIFS and cloud object storage system seamlessly.

Get Up and Running in 15 Minutes

Deploying Komprise is simple. Just download the Komprise Observer and point it at existing storage via NFS or SMB/CIFS. Within minutes, Komprise provides analytics on how much data you have across your storage, how it's being used, who is using it and how fast it is growing. Even on petabytes of data, Komprise provides nearly instantaneous results using a patent-pending dynamic sampling technique.

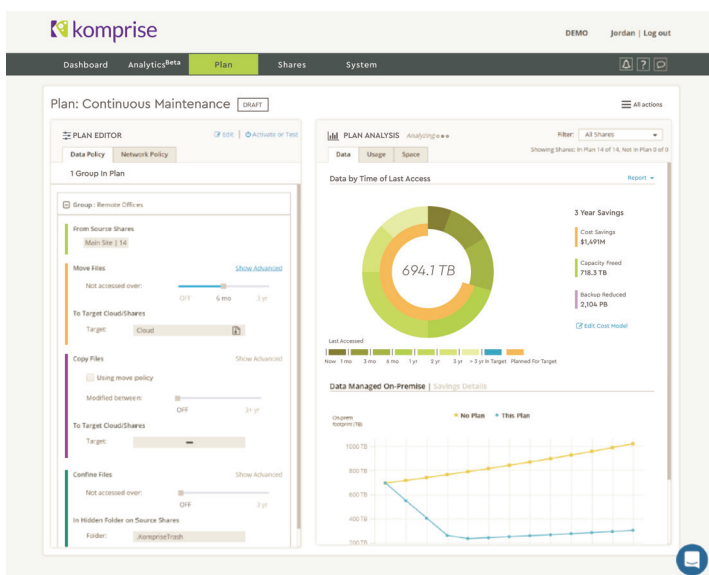


Understand Your Storage Environment with Data Analysis and Interactive ROI Visualization

Set your data management objectives and play with “what-if” scenarios to understand the projected impact on your footprint and costs. Komprise interactively projects:

- How the data footprint at the source will change based on the chosen objectives
- The cost impact of the policy
- The projected ROI of the policy

This provides a no-risk way to plan capacity and determine the most effective data management approach before moving data.



Move Data with No Changes to User and Application Access

Once you're satisfied, simply activate the plan and Komprise moves the data transparently. It doesn't use any proprietary agents or static links on the storage system—mechanisms that can break file hierarchies and traditionally cause problems. Users and applications experience no change at all. They continue to see and access the data the same way they did before.

Handle Data Migrations with Ease

Komprise Intelligent Data Management now includes NAS migrations as a data management capability, so you can eliminate the guesswork and sunk costs associated with migration. Simply pick your source and destination, and Komprise handles the rest.

Komprise also uses a fault-tolerant, highly available architecture. It resiliently moves your data and handles failures automatically, so your data migrations run reliably. It also preserves file hierarchy, NTFS permissions and all the metadata. With Komprise ensuring the full integrity of your data and metadata, migrations become much simpler.

Inside Western Digital ActiveScale Cloud Object Storage System

The more innovative your organization, the more unstructured data you're likely to store. Video, audio, sensor data, emails, IoT and social media—organizations are keeping everything. Which is why most find their unstructured data growing much faster than their transactional data. Savvy companies can transform and analyze vast amounts of inactive and unstructured data to glean hugely valuable insights. But to do it, they need to be able to access it at scale, with good performance, and store it for a very long time. In this context, more organizations are looking to cloud object storage. Increasingly, they're turning to Western Digital to help them create a “Data Forever” architecture.

Western Digital has been a groundbreaking innovator and leader in data storage for decades. With the newest generation of Western Digital's ActiveScale cloud object storage systems, organizations can economically store petabytes of data over multiple generations of storage with confidence. ActiveScale provides a framework to reduce operational costs, improve customer service and deliver enhanced data availability at petabyte scale.

ActiveScale P100 and X100 systems represent the next generation of cloud object storage systems for large-scale growth, simplified management and improved storage economics. Take advantage of ActiveScale cloud object storage systems to economically manage the avalanche of unstructured data inundating your data center and squeezing your IT storage budget.

The ActiveScale Advantage

ActiveScale provides a level of data durability that tape and even conventional NAS systems can't achieve: 19 nines of data durability with advanced erasure coding. At the same time, it delivers excellent performance, especially for the massive object payloads that typically characterize unstructured data. ActiveScale delivers 8 Gigabyte-per-second mixed read/write throughput, even at cloud scale. It lowers operational costs with helium-filled drives that require less power and cooling. And it makes it easy to protect your data by storing it across multiple regions.

Unlike other cloud object storage solutions, ActiveScale delivers:

- **Extreme data durability:** Ensure long-term durability with advanced erasure coding. And protect your data integrity with BitDynamics™ continuous data scrubbing, proactive data integrity checks and self-healing.
- **Flexibility:** Choose the optimal encoding option depending on your capacity and durability goals. And optimize performance with configurable small object policy.
- **Affordability:** Western Digital designs, develops, and manufactures the full breadth of storage solutions from component hardware, firmware, and software to systems. Since we understand every aspect of workloads and device characteristics, we can make informed decisions on the best selection of components to get the job done. We also understand those device characteristics when we design the enclosures to ensure that the vibration, cooling and mechanical design will make the most of the underlying storage technology. As a result, we can deliver superior price/performance and innovation.
- **Availability:** With Western Digital BitSpread® technology, you can maximize high availability—without the high costs of replication. Protect your data using low-cost, easy-to-manage geo-spread dispersion.
- **Simplicity:** Manage your entire storage namespace with a single pane of glass. And draw on ongoing cloud-based storage analytics for a deeper understanding of your data.

The Komprise + ActiveScale Modern Scalable Architecture

The combination of Komprise and ActiveScale delivers a scale-out architecture that can grow on demand, so you can handle today's massive scale of data through intelligent automation. With built-in high availability, Komprise handles scale-out by allowing you to simply add more virtual appliances as the amount of data you're managing grows.

Komprise also employs a distributed fault-tolerant scale-out design. As the load on the Komprise Observer grows, you can add additional Observers on-demand. Komprise manages all Observers as a single, fault-tolerant grid, making management simple while delivering a robust, efficient, scalable, high-performance solution.

Meanwhile, ActiveScale can consolidate multiple tiers of storage, including tape, with highly-scalable cloud object storage that delivers up to 8GB/sec read/write throughput. It reduces your operational costs with helium-filled drives that require less power and cooling. It delivers up to 19 nines durability with advanced erasure coding, along with proactive data integrity checks and self-healing. And it protects your data even in the event of a full data center outage with the ability to cost-effectively deploy it across multiple regions.

Start Slashing Your Storage Costs Today

The data you have to manage is growing much faster than your budget. If you're going to keep pace, you need a simpler, more efficient way to manage data sprawl. High-performance NAS will continue to play an important role in your data environment. But for the huge and growing amount of inactive and unstructured data your business generates, there's a better, less expensive solution.

Western Digital ActiveScale gives you industry-leading cloud object storage performance for managing inactive and unstructured data, with superior reliability and durability compared to NAS. Which means you can store more data, for much longer, at a fraction of the cost of your primary hot storage. Next, add Komprise analytics-driven data management software. In just minutes, you can identify inactive data and calculate how much you can save by moving old data to your cloud object storage tier. Komprise then transparently, automatically and continuously moves cold data to the storage targets of your choice. No complexity, no user and application disruptions, no big hits to the performance of your system. Just seamless, scalable storage for massive amounts of data, at a much more sensible price.

Ready to start managing your data more effectively and slashing your cold storage costs by 70% or more? Contact your Western Digital account representative or visit: www.westerndigital.com

Komprise + ActiveScale Modern Scalable Architecture

