Getting Departments to Care About Storage Savings

How showback and simplicity get the archiving buy-in you need

IT has had to plead with users to archive files—ANY files—to alleviate rapidly filling NAS. But with massive unstructured data growth and rising storage and backup costs, this bristly dynamic has to change.

komprise

Many companies are choosing a Storage-as-a-Service (STaaS) approach to centralize IT's efforts for each department. But convincing department heads to care about storage savings is a tough task without the right tools. Data management and archiving is viewed by users as an extraneous hassle and potential disruption that fails to answer "What's in it for me?"

Massive
STaaS Growth
expected by 2027[1]

This white paper explains how to make STaaS successful by telling a compelling data story department heads can't ignore. This coupled with transparent data archiving techniques that do not change the user experience are critical to successful systematic archiving and significant savings.

Learn how using analytics-driven showback can help secure the buy-in needed to archive more data more often. Once they understand their data—how much is cold and how much they could be saving—the conversation quickly changes.

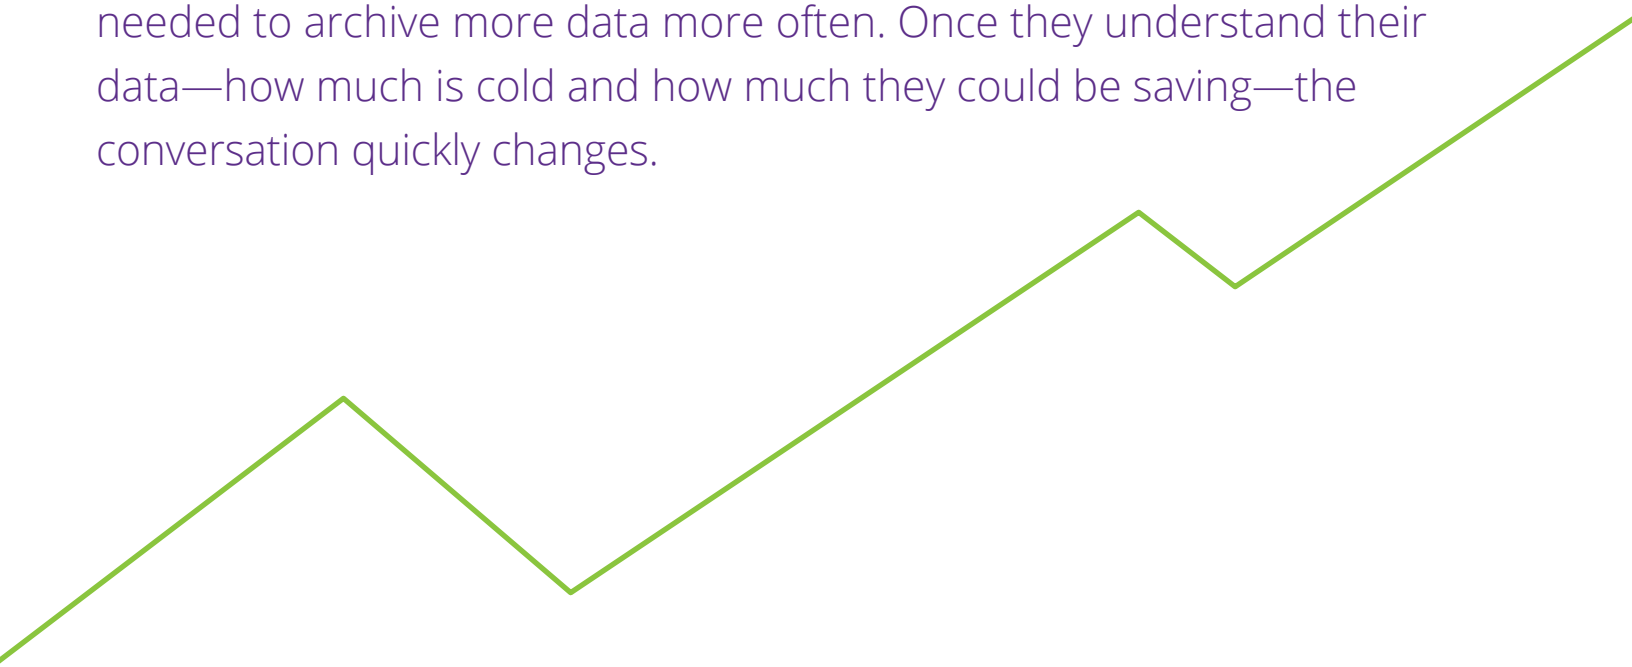

[1]"Global Storage as a Service (STaaS) Market Analysis to 2027" ReportsWeb.com, September 2020

Why Departmental Users Just Aren't Into Storage

Let's first look at the reasons why users are averse to The Storage Discussion.

- **They're data hoarders**—It's far easier to save a file than discover later you've deleted something necessary. Saving and keeping files is a tried-and-true safety net that creates a surprising amount of growing data that gets increasingly expensive to manage over time.
- **They're busy**—Focused on day-to-day business productivity, users don't have time to spend managing data and figuring out what should be deleted or archived.
- **They've been burnt before**—Chances are they've had a bad past experience with now-you-see-it-now-you-don't data archiving. Frustrating access to moved files has left a sour taste, and they're not keen on inviting that disruption into their department again.
- **They're not buying it**—They're tired of heavy-handed corporate mandates to save costs. They ring hollow, and IT's push can be seen as an empty threat.

Showback coupled with transparent data archiving is critical.

It is for these reasons that data management needs to be simple and frictionless for users, and also be able to easily show them the value of what IT is trying to accomplish. This is where showback coupled with transparent data archiving is critical.

Showback: Making STaaS Work for Everyone

Using the STaaS model, each department is now on the hook for their own storage. But to get managers to care, IT essentially needs to create a kind of "storage utility bill" for each department based on their use. To create an effective, easy-to-read showback report takes the right reporting tools to tell a compelling story.

The Right Tools

While there are tools to identify the capacity used on the storage devices, they don't work at the data level. More importantly, the tools don't show you how to save costs and develop and execute policies that manage the storage cost effectively. Companies find that simply enforcing limits or threatening departments with quotas and budgets are not effective in controlling data sprawl.

The Right Story

Below are the steps to tell a department's data storage story in a way that makes them understand and get motivated to save up 70% on storage costs.

1 Departmental Data Discovery: Crafting Their Storage Story

To compile their data, you need to set up groups of shares that pertain to those departments and analyze them to show each department their usage and how much it's costing them to store it. This unique report is generated for each department, and it contains the following information:

- How much data they have
- How fast it's growing
- Who's using this data
- How much is cold vs. hot
- How much it's costing them

These are the cold, hard, data facts that they will most likely never have seen.

2 Meeting of the Minds: Having The Storage Discussion

Meet with each department head to review their report showing usage and costs and the ROI of tiering off cold data with the resulting savings.

Oftentimes the cost savings alone seals the deal, but there's a host of other reasons they need to be aware of.

Ongoing cost savings

Most are impressed with the deep savings that comes from archiving their cold data. But remind them that all those savings continue to accrue: quarter after quarter, making a serious improvement in their budget.

- How much would be saved by archiving cold data

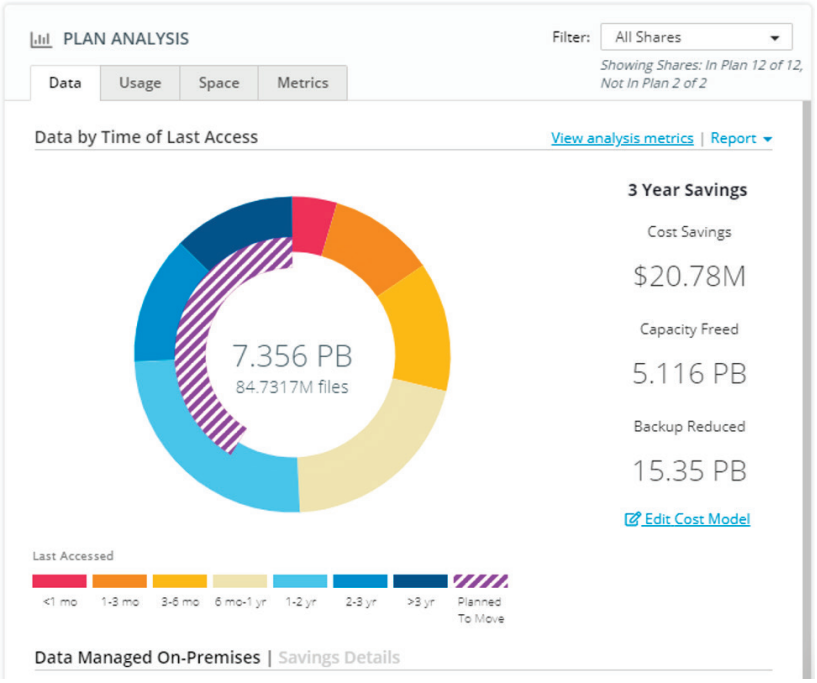

Projected cost savings over the next 3 years, based on archiving policies set by user.

No effort, no hassle

Remind users that they won't need to lift a finger to make any of this happen. Once you set agreed-upon archiving policies, everything is automated. It's great if they want to be involved in setting the initial policies, such as how cold data must be before tiering it off, but after that, Komprise does the rest. They'll only need to look at their future reports to see the wisdom of their choice partnering with IT to better manage their department's data and save.

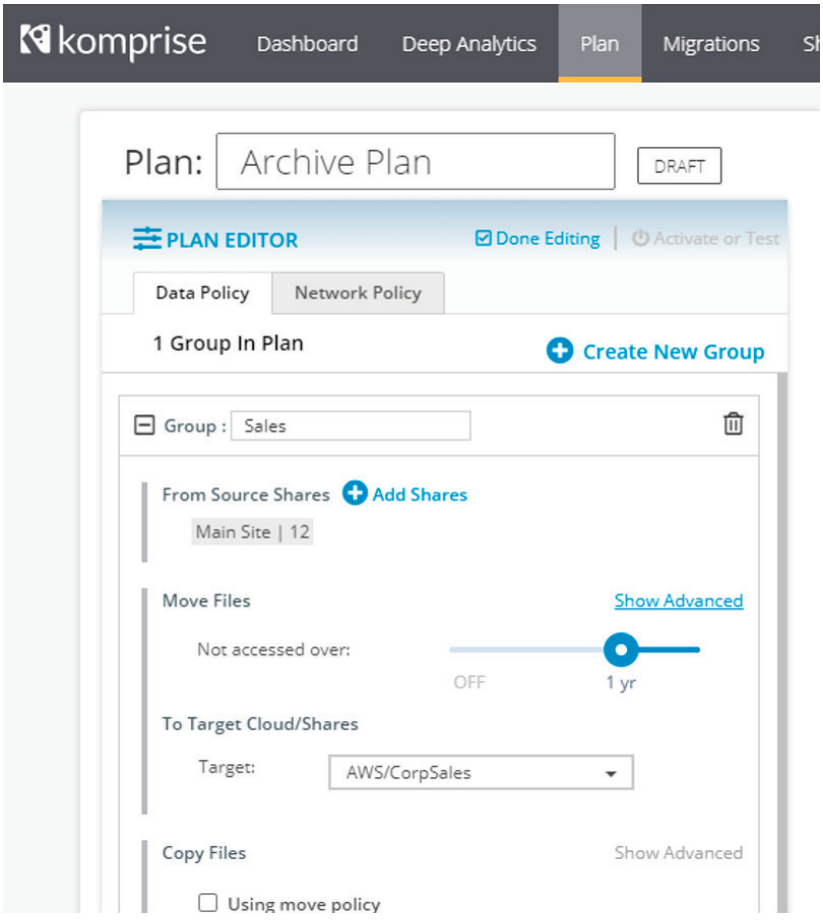

Users set agreed-upon archiving policies and Komprise does the rest.

Zero impact accessing moved data: Transparent Archiving

A common pushback users have is the frustration of accessing moved data. Some have suffered through archiving that used stubs or agents, which created disruption and a productivity hit to their team when files needed to be accessed again. Those concerns are eliminated with Intelligent Data Management.

Transparent Archiving: Flawless File Access

Files archived by Komprise are replaced with a standard file system construct called a symbolic link. That link will have the same name, attributes and permissions as the original file. This allows users to still "see" the file in their original directories. Directory listings still show the file as if it is still on the original source file system.

In the rare event when a user clicks on the symbolic link of a cold, archived file, a file system request is sent to Komprise automatically. Komprise accesses the file from the archive storage (typically a cloud storage) and streams the content back as a response to the file system request. The user is none the wiser. There is an added latency. Instead of say 3 to 5ms it could be 50ms. The actual latency will depend on the archive storage and the speed of the interconnect. In general, cloud storage with a reasonable interconnect works well. For those requiring very fast response, users can opt to use a local object store or lower cost NAS for archiving purposes.

Komprise caches files so that latency from on-going access is reduced significantly. This also eliminates repeated egress costs for files archived in public clouds such as AWS and Azure.

Komprise allows you to transparently archive data, so file access simply isn't an issue. No stubs, agents, or proprietary software. Moved data is found and opened exactly like before—whenever and wherever data has been moved. A tiny arrow in the file icon is the only evidence a file has been moved. That's Transparent Archiving, and it doesn't get much easier than that.

Making an Archiving Plan

Some department heads want to be involved in deciding what data stays and what gets tiered off to cheaper storage. The Komprise console makes it easy to work together creating policies so there's no surprises.

With the Komprise Plan, IT can play out "what if?" scenarios using their own costs to see how savings vary with different archive policies. IT simply chooses the department's shares, the target to archive to, and the policy to govern which data to archive. Department heads can specify how "cold" data should be before it gets archived, and IT can readily see the potential cost savings from that department.

Once the department and IT mutually agree upon the policy, the Plan can be activated and Komprise will automatically archive existing cold data, and continuously archive data as it ages. Reports available in the UI allow IT to provide constant feedback to both their departmental customers and to executive management.

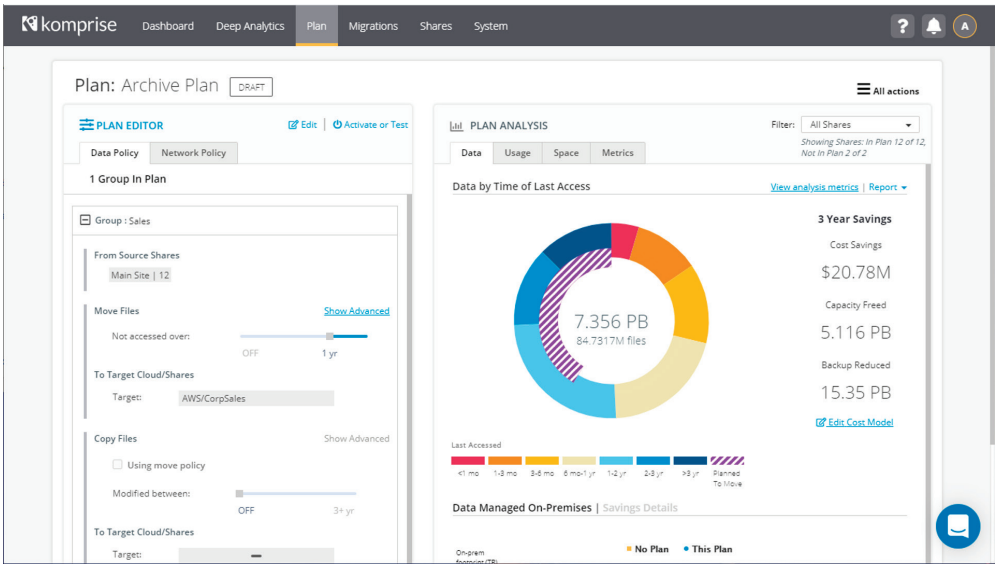

The Komprise console makes it easy for users to play "what if" scenarios and view the projected cost savings based on those policies.

Resurrecting Projects? Bulk Recall

Often departments know an upcoming project will require massive amounts of data to become hot soon. Say they're ready to start working on Shrek the (Fourth?) or they're going to chip out V2 of the CPU. IT can help their team ramp to peak productivity by bringing that cold data back onto a higher-performance tier that makes more sense for a hot project. Simply setup a bulk recall in Komprise, pick the directories and files you want to bring back, and Komprise automatically manages the rest.

Chargeback: The stick becomes the carrot

Chargeback reports have typically been used as "the stick," but many have found that when you follow these steps, they actually turn into a "competitive carrot."

When IT starts department discussions with the usage report, they get archiving buy-in and see the savings add up. Future chargeback reports become viewed as a competitive tool that encourages them to archive even more aggressively to create even more savings. Once they see how their budget is being gobbled by storage, they're much more likely to try to create savings that allow for other more strategic expenditures.

Case Study

How an Ivy League university saw the storage savings light

> "
>
> ## University Boosts Adoption of Secondary Storage and Saves Costs
>
> In our 300 years, no one had ever deleted a file...nobody knew what data existed. Now we have the visibility to show departments when and how their data's being used—or not. Moving it to the best location saves them and us money.
>
> **Steve DeGroat**
> Manager Enterprise Storage, Top Ivy League Research University

This Ivy League university provides Storage-as-a-Service (STaaS) to its various departments. While Central IT offers different classes of storage at different price points, getting departments to adopt a mix of primary and secondary storage was difficult. Because Central IT lacked visibility into the data on the department's storage, they couldn't illustrate how much of data was hot or cold and how much they could save by archiving cold data.

Departments were also concerned that moving would disrupt student, faculty, and application access, so data tended to stay wherever users put it. Data was not only consuming expensive storage space, but the cold data was being protected and replicated the same way as hot data, which is highly inefficient and costly.

The Solution

Using Komprise, Central IT was able to analyze how each department's data was growing and being used. They were able to provide each department with the analysis of how the data was being used, by who, and what data was cold. They also gave a cost analysis showing how much the department would save by transparently archiving the cold data with Komprise.
Each department was then able to set their own specified policies and use them to move and transparently archive their data, without any changes to users or applications.

The Results

By using Komprise, Central IT was able to:

Provide data insights to departments so the university could plan intelligently with analytics of data usage and growth.

Save storage costs by offloading cold data, which they'd estimated by running "what-if" scenarios for each department.

Provide uninterrupted data access for students, faculty or their apps whether in the cloud, as objects, or on secondary storage.

Set data archival and replication policies by department without affecting users, and without any changes to network or storage infrastructure.

Enhance collaboration between Central IT and departments by sharing insights into departments' data for greater understanding.

Getting Buy-In: What's Not for Users to Love?

With the automated, analytics-driven approach to data management that Komprise Intelligent Data Management provides, you can easily address users' concerns and turn the tables on the storage discussion. It takes simplifying the task of data management and removing the obstacles that data archiving has historically presented.

End Users Are:	Komprise Understands:
Data Hoarders	There's no need to part with their cold data, it just doesn't need to sit on expensive storage tiers.
Busy	This is why data management should run automatically in the background, and data be archived, so users don't notice any difference.
Cautious	With Transparent Archiving there's no hunting for moved data—it's right where it's always been and opens the same way.
Dubious	Change the conversation. Present storage savings as a way to free up funds to put toward more exciting things.

Don't blame users for being skeptical

They just need convincing that storage savings can be that radically easy. Lucky for you Komprise makes that possible because it's radically different.

Learn More

Go to **Komprise.com/product** to learn more.

Komprise, Inc.
1901 S. Bascom Ave. Suite 500
Campbell, CA 95008
United States

For more information:
Call: 1-888-995-0290
Email: **info@komprise.com**
Visit: **komprise.com**

For media requests email
marketing@komprise.com