Extending NAS to Google Cloud Storage with Komprise

Classic NAS has become an expensive tier of storage for seldom-accessed data. Learn how to use the Google Cloud Platform (GCP) service Cloud Storage and Komprise to actively archive and replicate data to the Google Cloud without disrupting users and applications.

Storage IT Challenges

Enterprise NAS devices are typically refreshed at three- to five-year intervals. And while IT admins may be aware that a large proportion of data on this expensive storage is cold, migrating it can be complex and risky. Archived data that's needed by users and apps often causes operational disruption. Permission to archive cold user data is rarely given, and when it is, identifying and migrating the correct data to the cloud is an extensive, cumbersome manual process involving spreadsheets, reporting tools, and various software applications. As your data footprint grows, so too do these challenges. Now there's a smarter solution.

Streamline migrating data to the cloud

Komprise coupled with Google Cloud Storage automatically identifies and moves cold data by policy from any NAS to Google Cloud Storage. This is accomplished without disruption because the moved data still looks like it's stored on the primary NAS. When a user or an application accesses this data, Komprise automatically recalls it with a transparent bridge of object data to files. IT can manage their storage farms efficiently, automatically and seamlessly—without having to ask for user permission.

Another IT issue is the rate unstructured data is being generated. Because backing it all up is simply too costly, in most cases, such data is never backed up. With Komprise, you can replicate unstructured data to durable Cloud Storage, providing an automated, simple way to help protect your data and facilitate DR.

Capabilities

Analyze Data Usage and Growth Across Storage

Komprise provides Dynamic Data Analytics across storage silos to identify how much data is hot and how much is cold and to help answer the following:

- What types of files are they?
- What's the distribution of file sizes?
- Who is accessing which files?
- How fast is file storage growing?
- How much data is inactive?

Charts provide a quick visual representation of the data profile.

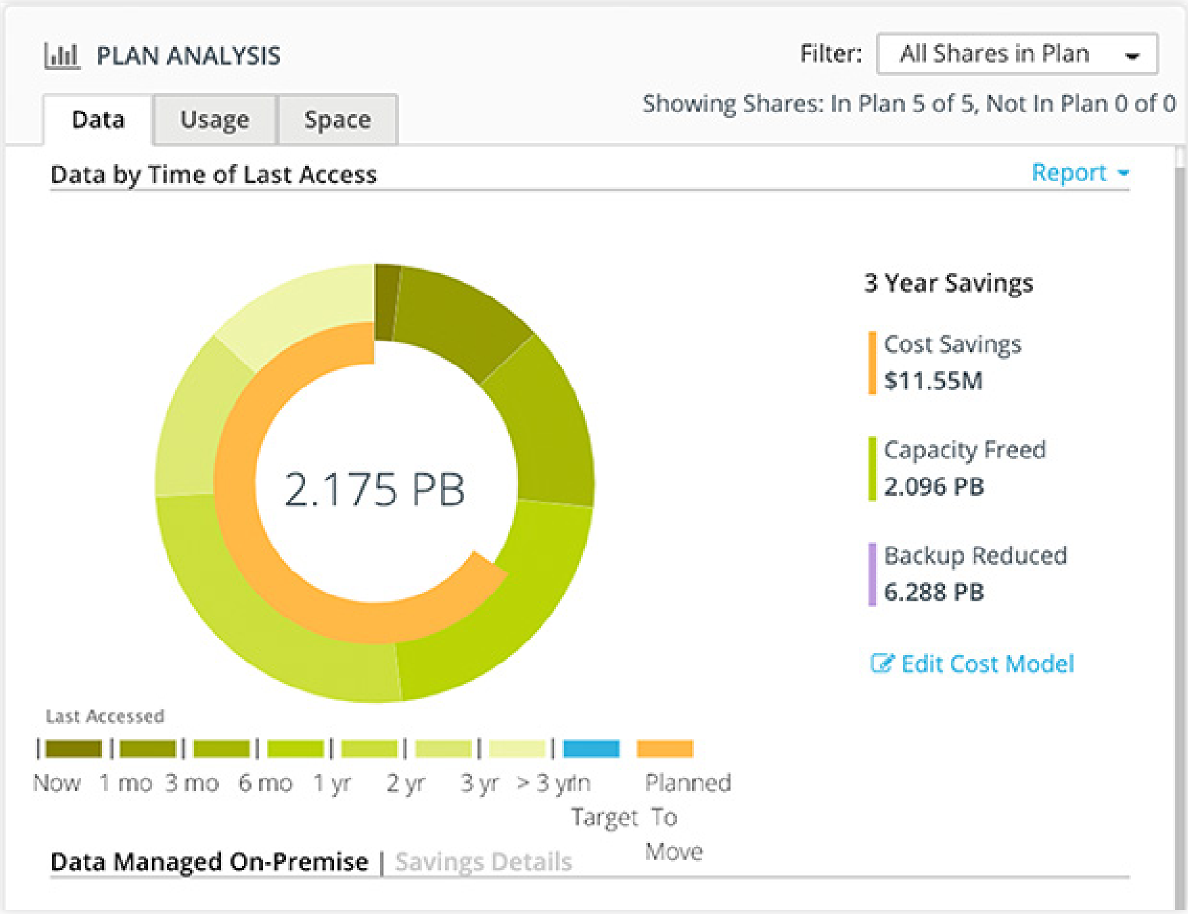

PLAN ANALYSIS

Filter: All Shares in Plan

Showing Shares: In Plan 5 of 5, Not In Plan 0 of 0

Data by Time of Last Access

3 Year Savings

Cost Savings
$11.55M

Capacity Freed
2.096 PB

Backup Reduced
6.288 PB

Figure 1: This donut chart shows that Komprise has analyzed 2 PB of data. The colored buckets show when and how much data was last accessed. The orange ring shows the administrator's move policy: all data that has not been accessed in over five years is slated to be moved.

For more granular decision-making, Komprise also provides access and aging information that's based on file type, size, owner, group, and directory.

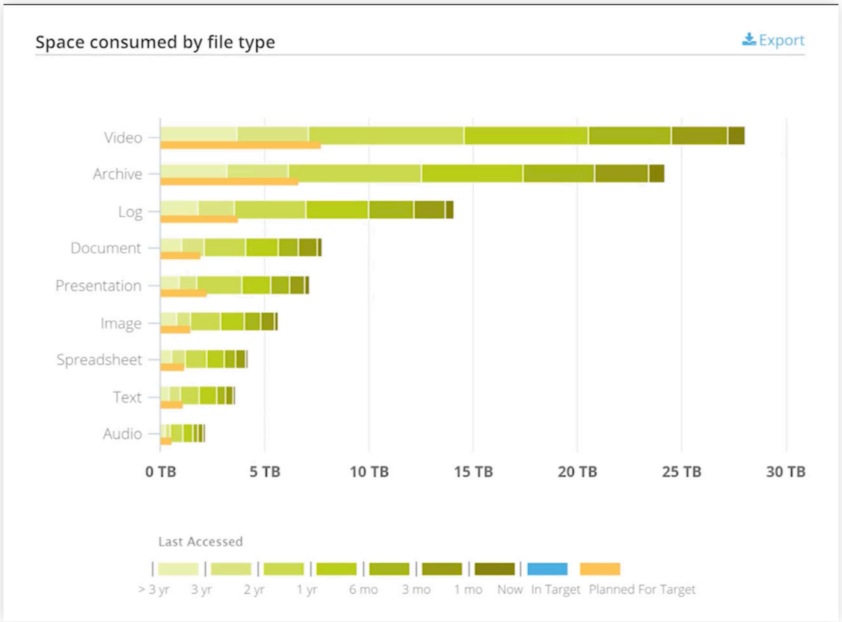

Figure 2: See what kind of data is being used how often.

Komprise allows you to run "what if" scenarios and get subsequent capacity needs and cost savings in seconds. Want to know what would happen if you moved all data untouched in over a year to Cloud Storage? Get an instant analysis based on your data, historical data growth patterns and which Cloud Storage class you want to use: Regional, Multi-Regional, Nearline, or Coldline.

Control Data Moves with Policies

You can easily set policies to automatically move and copy your data based on your organization's needs using Transparent Move Technology (TMT™).

- The move policy continuously moves inactive and cold data to Cloud Storage as the data ages. Identifying and moving cold data eliminates the ongoing need to increase the capacity of on-premises NAS storage.
- The copy policy allows you to select different conditions for copying data. E.g., only replicate to the cloud data that's been modified in the last year.
- Specific with a policy that obsolete data be removed (moved to a NAS trash folder) rather than moved or copied to a new storage platform.
- If certain data should not be moved or copied, create specific exclusions using file types, size, and folders.
- Create custom policies for data that has unique needs; Komprise dynamically calculates the estimated capacity that will be freed up and your projected cost savings.

Information Lifecycle Management

Komprise uses tiered Cloud Storage to further reduce costs. Through policies that you set in Komprise, you can tier data from Nearline storage to the less expensive Coldline storage based on the age of and lack of access to the data after you have moved it to Nearline storage. Both provide similar access times, so you can reduce costs further by using Coldline storage without affecting your ability to access the data when you need it.

Architecture

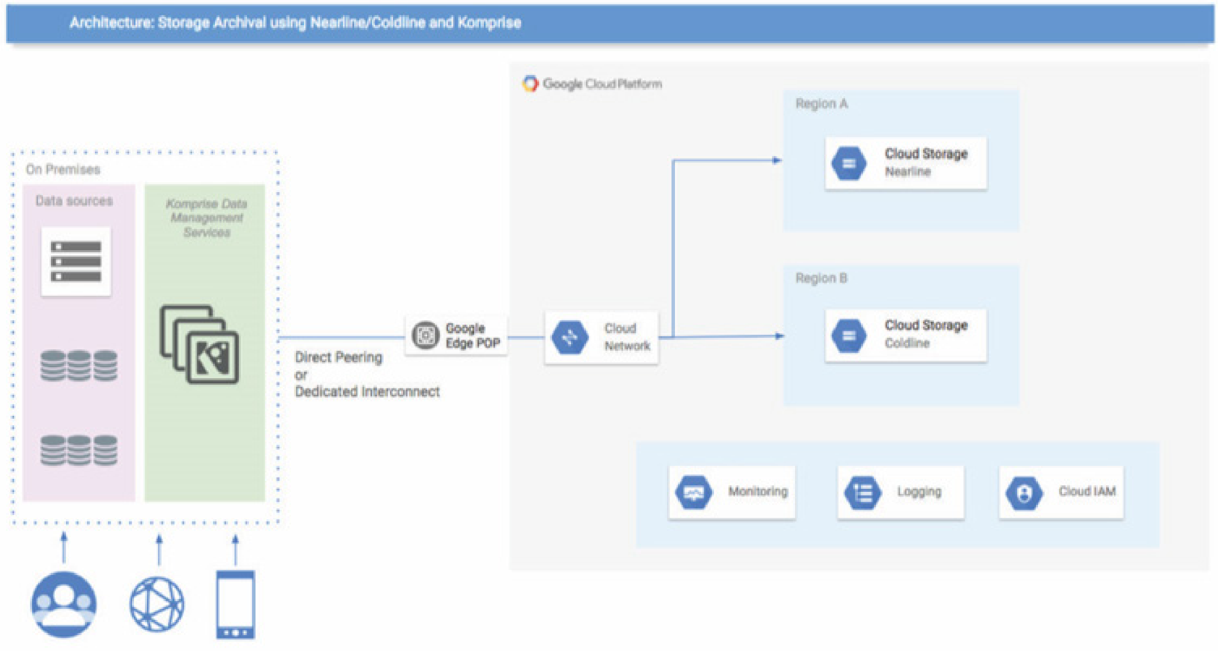

Figure 2: Depicts the architecture of a typical solution.

Komprise runs as a hybrid cloud service with a grid of one or more Komprise virtual appliances, called Observers and Proxies, deployed on premises. The grid has a highly parallelized, scale-out architecture. Observers analyze data across on-premises NAS storage, move and replicate data by policy, and provide transparent file access to data that is stored in the cloud. Komprise Proxies encapsulate the extended Server Message Block (SMB) or Common Internet File System (CIFS) metadata and permission structure for compatibility with modern object architectures and accelerate file transfer to Cloud Storage. Finally, a Komprise Director virtual machine (VM) runs in the cloud and provides the management console.

Figure 3: Before and after of files moved by Komprise TMT.

Features

Komprise Intelligent Data Management

SCALE OUT: Komprise does not require any dedicated hardware and runs as a scale-out grid of VMs that are managed as one logical unit. There are no centralized databases, which allows Komprise to grow on-demand to handle data at massive scale. The grid is highly available and so long as at least one Observer is healthy, access to all moved data remains intact. Komprise does not store data, and simply moves data through SSL to Cloud Storage, which is HIPAA compliant.

NON-DISRUPTIVE: A typical challenge with traditional storage services is that they can disrupt end-user access. Komprise preserves the directory structure as well as file attributes on the target, unlike cloud migration tools that strip data off file attributes and move blocks to the cloud that can only be accessed and understood using the application going forward. With Komprise, end users can continue accessing files just as they always did, because the location of data is transparent to them.

HIGH PERFORMANCE: Many migration solution providers significantly reduce the performance of storage during data moves. Komprise, however, is invisible to the hot data path and does not get inline. It adaptively throttles back when the storage systems are actively in use so that Komprise analytics runs non-disruptively in the background. This means that the performance of the active data is unchanged and may even improve as the primary storage becomes less overloaded.

NO STATIC STUBS: A stub, which contains the location to which a file has been moved, can be easily deleted or corrupted, orphaning the files that were moved to the target storage. As file systems grow to hyperscale, multi-petabyte size, managing these stubs becomes increasingly challenging, requiring large and complex database management to protect them to maintain data accessibility. Komprise delivers transparent access by using standard protocol constructs when moving data. When you move a file, a symbolic link containing all the properties of the original file is left behind as a pointer. Users and apps continue to see and can open the file from the original location keeping all the permissions and access control intact. No invasive agents or stubs means no disruption to users, applications, or the data protection workflows.

Security

Komprise ensures that data is protected and encrypted by default providing the following security options for moving data to GCP:

Encryption in Transit and at Rest

In this default mode, data is transmitted between Komprise observer and GCP using SSL, and Google encrypts the data using AES 256-bit symmetric key encryption using Google keys before storing the data. The keys are managed by Google, and Komprise never receives the encryption keys. During access, Google decrypts the data and sends it securely over HTTPS using SSL to the Komprise observers. Data is then transferred to end users accessing the data.

End-to-End Encryption

In this mode, data is encrypted on Komprise Observers using AES 256-bit symmetric key encryption before transferring to GCP. During access, the Komprise Observer retrieves encrypted data from Google that is transmitted in encrypted format. The Komprise Observer then decrypts the data using the Data Encryption Key and then sends it to the user. This is a heightened security mode where data is only available through the Komprise grid and not directly in Google.

Extend Your NAS to Google Cloud Storage with Komprise

With the Google Cloud Platform service Cloud Storage and Komprise, you can realize the benefits of actively archiving and replicating data to the Google Cloud without disrupting users and applications.

Learn More

Learn more about how Komprise and Google can help you cut costs, free up primary storage capacity, and strengthen data protection. Contact **sales@komprise.com**

Komprise, Inc.
1901 S. Bascom Ave. Suite 400
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