

# Know Your Data to Manage it Better with HPE and Komprise

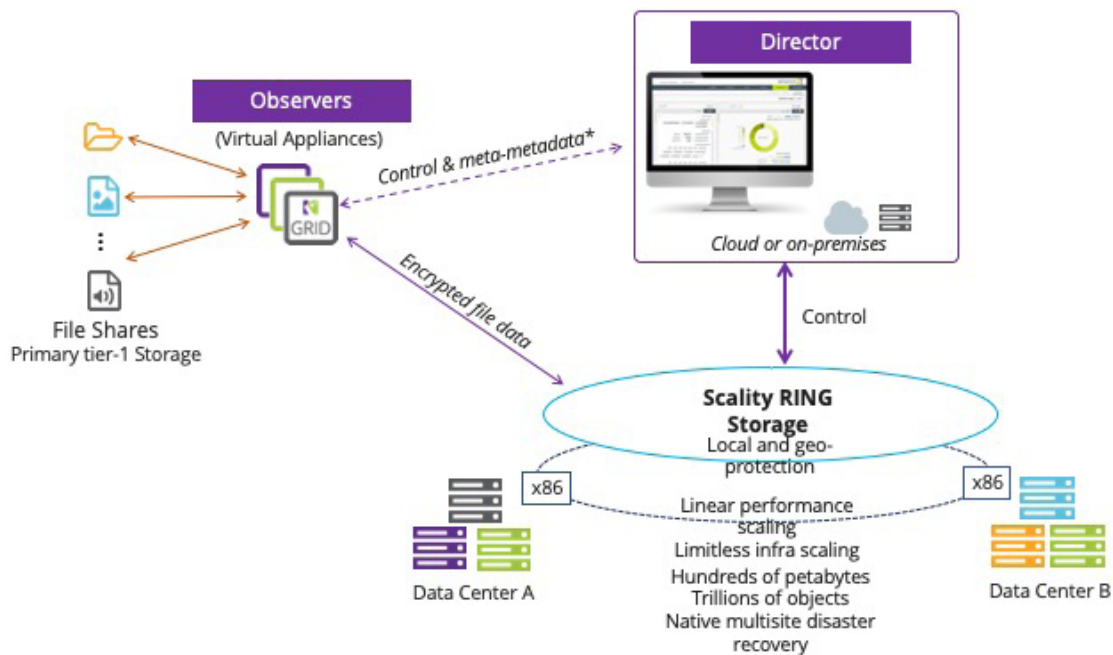
As data footprints quickly expand, managing data more easily becomes critical. Understanding your data is key to managing it in a more strategic and efficient way. Komprise and Hewlett Packard Enterprise (HPE) are working together with customers across multiple industries to help streamline costs, create more efficient, capacity-enhanced storage, and increase the resiliency of their data in active archiving, replication, and disaster recovery.

## Introduction

Research shows that 60% to 80% of business data is infrequently accessed and becomes inactive within months of creation. Yet it's often stored and managed in the same way as active data—at a steep and unnecessary cost. Komprise uses analytics-driven data management to enable businesses to manage their data intelligently by identifying infrequently accessed *cold* data across their NAS storage frames. It also allows them to transparently move targeted data with user-defined policies to more cost-effective options, such as cloud or object storage—without affecting user or application access.

You can ensure that only active *hot* data is kept on your tier-1 primary storage by relegating cold and other targeted data (including explicitly excluded data) to a more cost-effective capacity-enhanced storage system. HPE StoreEasy and HPE Apollo with Qumulo are ideal as high performance, scalable tier-1 NAS storage, while HPE Scalable Object Storage with Scality RING on HPE Apollo is better suited to provide the capacity-enhanced object storage to reduce or eliminate the need to increase the capacity of tier-1 NAS storage, reducing on-going costs and maintenance.

For redundancy and disaster recovery (DR), many businesses need to keep a copy of their data in an object store that protects their data across multiple data centers. Komprise's policy-based tools enable data to be continuously copied, e.g., from an HPE StoreEasy system or an HPE Apollo 4200 with Qumulo, to HPE Scalable Object Storage with Scality RING on HPE Apollo 4000 Storage Servers. This provides protection during a network outage between data centers, and in cases where one data center becomes inaccessible due to error, equipment failure, or disaster. In DR, Komprise helps maintain business continuity by enabling fast access to data in the remote data center.



**Figure 1:** Solution architecture for Komprise with HPE Scalable Object Storage with Scality RING on HPE 4000 Apollo Storage Servers

## Solution Overview

The Komprise data analytics software requires no dedicated hardware and runs as a scale-out grid of virtual machines. Its distributed architecture consists of one or more Komprise Observer virtual appliances running close to the primary storage on the customer’s premises, and connected to a Komprise Director virtual appliance that can run in the cloud or on-premises. Instead of using traditional SQL databases that limit scalability and can create bottlenecks, Komprise Observers analyze and aggregate metadata. Komprise works seamlessly across many data sources using NFS, SMB/CIFS, or S3 object/cloud storage—including file servers from Dell/EMC, HPE, NetApp, and Windows.

## Understand Your Data, Wherever it is

Komprise profiles the data across all storage devices with visibility and easy-to-use analytics that help you answer questions, such as:

- How much data is being managed across all my storage devices?
- What are the types of files and how much is my IT currently managing?
- What’s the distribution of our file sizes?
- Who is accessing which files?
- How fast is our file storage growing?
- What is the associated hosting cost of data at my current locations?

## Creating “What-If” Scenarios

After storage has been analyzed, Komprise software lets the customer create different “what-if” scenarios. By specifying various data management objectives, they can quickly see the projected impact on storage capacity and data-hosting costs. When the plan has been configured for the data to move or be copied, the Komprise Observer grid will handle transferring data to the target object stores based on the customer-defined policies. The data is moved easily and transparently. The process is completely

invisible to the user because the original file is replaced with a Komprise dynamic link that looks and operates like the original file. When a user accesses the link, Komprise provides access to the file as if it were still on the original tier-1 NAS storage.

Figure 1 shows how Komprise software moves or copies data from a primary, tier-1 NAS storage and writes it to HPE Scalable Object Storage with Scality RING on HPE 4000 Apollo Storage Servers, optimizing the use of primary storage and capacity-enhanced object stores.

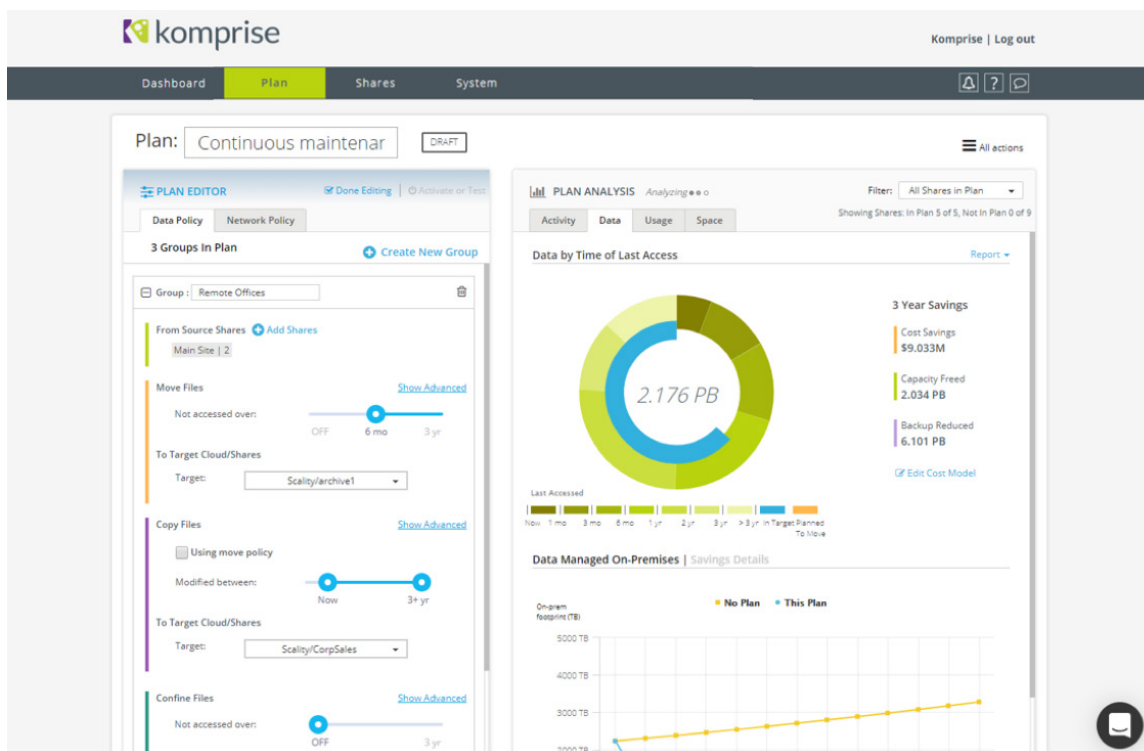


Figure 2: Komprise Director management console showing analysis of data usage

## Solution Components

### Komprise Intelligent Data Management

Komprise Intelligent Data Management consists of intuitively designed features that give IT administrators the data visibility and insight and the power to make smarter decisions.

### Know your data with Dynamic Data Analytics

#### *Insight for smarter decisions*

With the *Komprise Director* management console, shown in Figure 2, IT administrators can analyze data usage and growth to plan and to create policies for moving and managing data, automating the ongoing management of data across storage sources and targets. It generates several charts, including the time of last access and space consumed by file type, file size, top users, and top groups, as shown in Figure 2 above.

## ***Make an informed storage plan***

After storage has been analyzed, a *Komprise Plan* is created for minimizing storage capacity and data-hosting costs. The Komprise management console will project the estimated capacity that will be freed up on the primary storage and the anticipated cost savings. The cost model in the management console can be modified to use your own data-hosting costs, providing a customized ROI calculation. Komprise software projects the three-year savings based on historical growth rates of data. This provides no-risk planning for the most effective data management—*before* actually moving any data.

## ***Create virtual data lakes***

When you need to dig deeper into your data beyond trends, *Komprise Deep Analytics* can help. We provide an intuitive way to search and find specific files that fit your exact criteria across all your storage. Simply build your specific queries to get both summary information and detailed reports on the files that fit your criteria. Tag the data you find, files being created, and ones you're learning about. This dynamic approach allows you to run tag-based queries and build real-time virtual data lakes on the fly, without having to first move the data. You can continually leverage these data lakes for applications like Big Data, AI, and ML. This is available both via a user-interface and an API.

## **Move your data without user impact with Transparent Move Technology™**

Once the Komprise Plan is finalized, it's activated and *the Komprise Observer grid* moves the data transparently. Using *Transparent Move Technology™* data is moved from the tier-1 NAS (e.g., HPE StoreEasy or HPE Apollo with Qumulo) to the target (e.g., HPE Scalable Object Storage with Scality RING), based on the user-defined policies. The transferred data still appears to exist on the source storage system just as before because Komprise delivers file-based access to all the moved data. Users and applications experience no disruption. Komprise moves files as objects and allows object-level access to the moved data from the store, as well as file-level access from the original source storage device using dynamic links.

## ***Flexible and scalable to meet your growing data needs***

Komprise Intelligent Data Management deploys in minutes and yet easily scales-out to handle massive data growth. As the load on the Komprise Observer grid grows with your data, additional Observers can be added on-demand. All Komprise Observers are managed as a single, fault-tolerant grid, simplifying management while delivering a robust, efficient, and scalable, high-performance solution.

## **HPE Storage Solutions**

### **HPE StoreEasy**

HPE StoreEasy is a leading, easy-to-manage NAS product family that provides a centralized space for securely storing documents, images, audio, and video files.

### **HPE Apollo 4000 Storage Servers**

HPE Apollo 4000 Storage Servers are density-optimized platforms, purpose-built to service object storage solutions that provide robust object storage solutions. HPE Apollo 4000 systems are configured to serve as storage nodes in a Scality RING cluster, providing both storage and connector functionality.

### **HPE Scalable Object Storage with Scality RING**

The HPE Scalable Object Storage with Scality RING running on HPE Apollo 4000 storage servers provides a software-defined storage (SDS), petabyte-scale data storage solution that is designed to interoperate in the software-defined data centre. HPE Scalable Object Storage with Scality RING utilizes the HPE Apollo 4000 storage servers for high-density storage capacity, enabling enterprises to deploy Scality RING with efficiency.

## HPE Apollo 4200 with Qumulo File Fabric

Apollo 4200 Gen9 servers running the Qumulo File Fabric (QF2) software is a highly scalable file storage solution that runs in the data center. The highly economical Apollo 4200 with QF2 file storage solution provides real-time analytics to let administrators easily manage data no matter how large the footprint or where it's located globally.

## Summary

With data growing rapidly and IT budgets tightening, businesses need a simpler way to efficiently manage their data sprawl while cutting costs. Combining Komprise Intelligent Data Management with HPE scale-out storage solutions relieves primary storage from the pressure of rapidly growing data by making it easy to move older, inactive data to more cost-effective object storage. This extends the life of primary storage by freeing up the space used by inactive data and achieves cloud economics with on-premises deployments.

## Learn More

Learn more about how Komprise and HPE can help you cut costs, free up primary storage capacity, and strengthen data protection. Contact [sales@komprise.com](mailto: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](mailto:info@komprise.com)  
Visit: [komprise.com](http://komprise.com)

For media requests email  
[marketing@komprise.com](mailto:marketing@komprise.com)  
©2019 Komprise, Inc. All rights reserved.