

Komprise Elastic Data Migration Overview

Businesses today are looking at modernizing storage and moving to a multi-cloud strategy. As they evolve to faster, flash-based Network Attached Storage (NAS) and the cloud, migrating data into these environments can be challenging. The goal is to migrate large production data sets quickly, without errors, and without disruption to user productivity. Komprise Elastic Data Migration delivers a fast, reliable, and cost-efficient migration solution that is easy to use without requiring heavy lifting. This white paper gives a technical overview of the Komprise Elastic Data Migration solution and shares its superior performance results compared with common file migration tools such as rsync and Robocopy.

Why file data migrations are hard

Migrating unstructured file data from one NAS to another and to the cloud is no simple feat. It involves file data that consists of the file content and metadata, including permissions and attributes that aren't always represented identically by all the file servers, which creates a complex process. And since the amount of data to be migrated can be huge, it can mean long migration windows and significant effort. For organizations that require minimal downtime, the disruption of long migration cutover windows poses a major problem. Migrating files fast is easy when the files are large, because the migration overhead is small compared to the time required to transfer the file over the wire. It's a different story when the files are small as the migration overhead can be a multiple of the transfer time.

To address migration to a cloud NAS or to a NAS in another datacenter, modern migration solutions also need to support migrations over WAN connections with acceptable performance. Migrations over a WAN will be much slower than over a LAN due to the higher latency. For these cases, it's important to reduce the back-and-forth protocol "chatter" and automatically partition the task into multiple smaller ones that run in parallel to overcome the performance degradation from the WAN. A strong migration product not only needs to work across different vendors' NAS solutions quickly, but it also needs an intuitive UI that lets you run, monitor, and manage hundreds of migrations simultaneously. And, of course, it must be fast to reduce the downtime when IT cuts over from old file servers to the new ones.

In summary, unstructured data migrations of file and object data are difficult because data sets are large and the migrations can take months, maintaining file fidelity across dissimilar systems is challenging and the network performance can significantly impact migrations, especially to the cloud.

Komprise Elastic Data Migration: Greater Speed and Reliability

Komprise Elastic Data Migration is a high-performing migration solution that boasts the fastest migration speeds for both NFS and SMB data sets. It is a highly parallelized, multi-processing, multi-threaded approach that works at three levels:

- **Multi-level Parallelism:** Maximizes the use of available resources by exploiting parallelism at multiple levels including shares and volumes, directories, files, and threads to maximize performance. Komprise Elastic Data Migration breaks up each migration task into smaller ones that execute across the Komprise Observers. Komprise Observers are a grid of one or more virtual appliances that run the Komprise Intelligent Data Management solution. This parallelism occurs automatically across the grid of Observers. The user simply creates a migration task and can configure the level of parallelism. Komprise does the rest. Read below to learn how Komprise migrations are 27X faster than common tools such as rsync.
- **Protocol-level Optimizations:** Reduces the number of round-trips over the protocol during a migration to eliminate unnecessary chatter. Rather than relying on generic NFS and SMB clients provided by the underlying operating system, Komprise has fine-tuned the client protocols to minimize overhead and unnecessary back-and-forth messaging. This is especially beneficial when moving data over high-latency networks such as WANs.
- **Komprise Hypertransfer:** This technology creates dedicated virtual channels across the WAN, which minimizes the WAN roundtrips. This in turn mitigates SMB protocol chattiness and dramatically improves data transfer rates. **Tests done using a data set dominated by small files show how Komprise accelerates cloud data migration 25x faster than other alternatives.** Additionally, Hypertransfer strengthens security for [data migrations](#) as all file communication passes directly from the on-premises Observers to the cloud Windows Proxies through the private Hypertransfer channel. No part of the migration goes directly from on-premises systems to the cloud filers, and therefore the cloud filers themselves do not need to be exposed to any systems or network outside the cloud if that is not desired by the customer. Read the [white paper](#) to learn more about Hypertransfer and how it compares with other tools.

Analytics with migration for data and environment insights

Unlike most migration solutions, Komprise Elastic Data Migration provides valuable analytics to help you plan your migrations. Komprise data analytics help you understand where data lives, how it is being used and how it is growing. But in addition, Komprise assesses your environment and network so you can identify potential bottlenecks before starting a migration. The data analytics insights make it possible to properly plan and manage your migrations. The analytics show how much data you have on each volume (share), the age of that data (whether by last modified time or last accessed time), and for each volume charts what types of files are present, a histogram of file sizes, and space consumed by files of different sizes.

Insights from Data Analytics

- How much data is on each share
- The age of the data
- The types of files
- How much space it takes
- And more

Ongoing value

Komprise eliminates the sunk costs of traditional migration approaches. After the migration, you can use the full Komprise unstructured data management platform to maximize use of the new NAS file servers or cloud storage. Komprise analyzes data growth and usage across your storage to find cold, inactive data, and projects the ROI of moving cold data to secondary storage such as object storage in the cloud or on-premises. Komprise also moves cold data transparently based on customer-defined policies, so users continue to access the moved data in the same location as before. Komprise helps organizations reduce over 70% of storage costs while managing data growth.

Komprise ACE

ACE (Assessment of Customer Environment) is a program to analyze the expected performance of a customer migration or any movement of data—whether file to file, object to object, file to object, on premises or in the cloud. The ACE tool proactively identifies potential bottlenecks and other issues independent of Komprise **Elastic Data Migration** running in the customer’s environment and takes an hour or less of the customer’s time. Common issues include network and security configurations and file sizes.

Testing with the Android Benchmark

Testing was done to compare the data migration utility rsync for NFS data with Komprise. [Note that SMB migration performance is enhanced with Hypertransfer, per above] Figure 1 shows the architecture of a NAS migration using Komprise. The easy-to-deploy, fault-tolerant, distributed architecture consists of one or more Komprise Observer virtual appliances running at the customer site. The virtual appliances are connected to a Komprise Director, which provides the administrative console UI and runs as a cloud service. Komprise works across on-premises NFS and SMB/CIFS storage and object/cloud storage seamlessly.

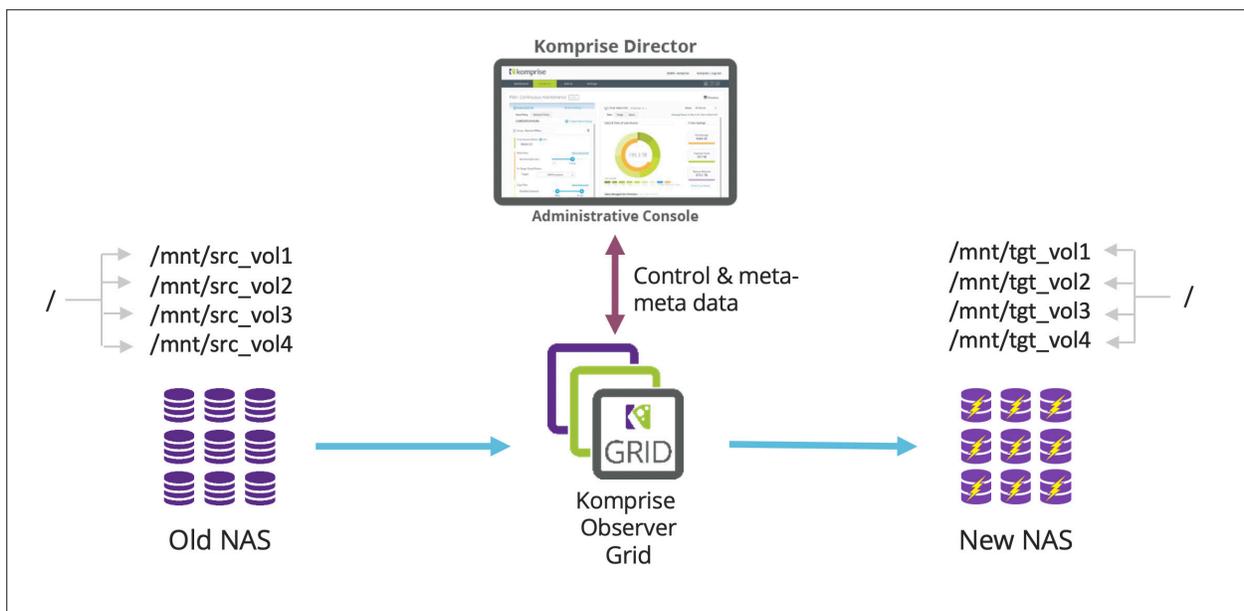


Figure 1.
Komprise Elastic Data Migration Architecture

When migrating data, Komprise performs automatic retries if the network or storage is unavailable. It ensures that all permissions and file data are fully migrated and verifies it using checksums and integrity reporting per file. Performance can be enhanced as needed by increasing the number of threads in an Observer or by adding more Observer virtual machines.

Data set: Android Open Source Project

The speed of Komprise Elastic Migration was tested using the Android Open Source Project data set¹. The Android data set was chosen because it is representative of the files at typical EDA companies: a high file count of very small files with some large files. The Android data set was 74 GB, with the specifics shown below:

File count	Data set size	Directory structure	File sizes
990K	74 GB	6 levels deep; files in every leaf directory Includes 139,469 directories	Mixed, with high count of small files and some medium/large files: Small: 500 B – 100 KB Medium/Large: 10 MB – 1.5 GB

Komprise Elastic Data Migration provides analytics about the data that is to be migrated, which is very useful for planning large-scale migrations. Below are some of the charts from the analysis to show the characteristics of the Android data set. The chart below shows the common file extensions in the Android data set. “Pack” and “archive” are the most common, followed by binaries and source code.

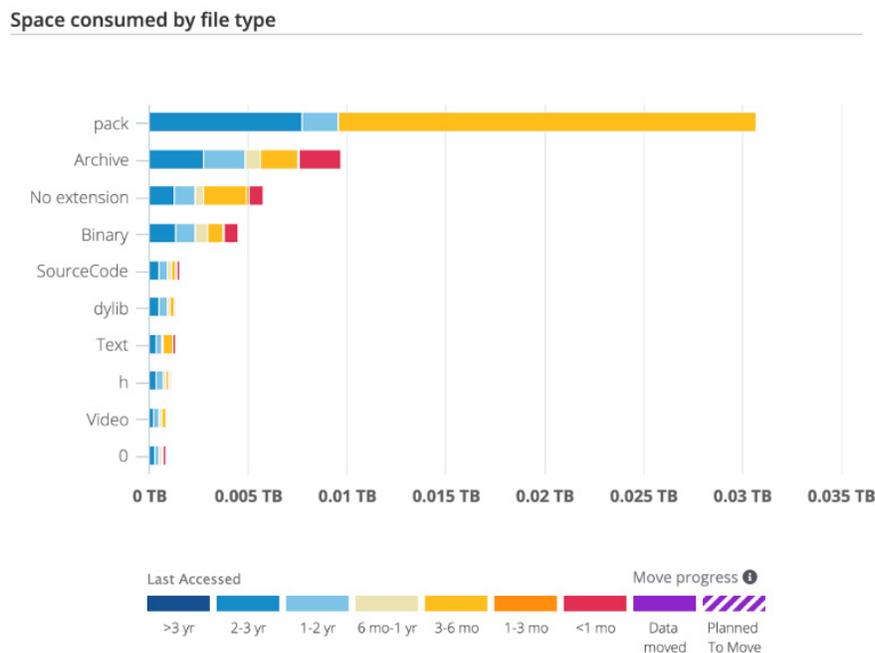


Figure 2.
Types of files reported by Komprise for the Android source code

¹ Android Open Source Project data set was downloaded from <https://source.android.com/setup/build/> to an NFS volume mounted on a source file server.

As shown in Figure 3, most of the space is consumed by files larger than 100 MB but most of the files are small (1 KB–10 KB) and almost all are less than 100 KB (Figure 4). Note that the Komprise analytics reports not only show the types of files, their sizes and the number of files, the reports also break down each category by “age” or the time they were last accessed. This information is used by Komprise to tier cold files based on user-defined policies.

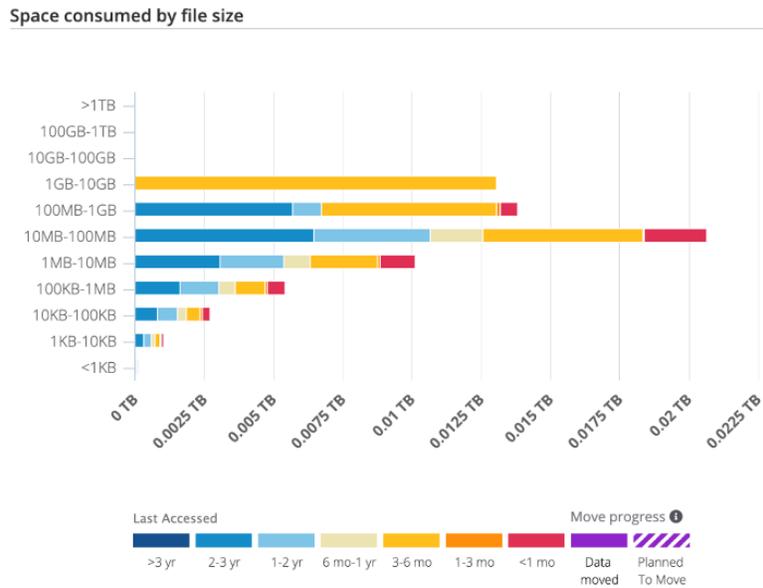


Figure 3. Komprise analytics report showing amount of space consumed by files of different sizes

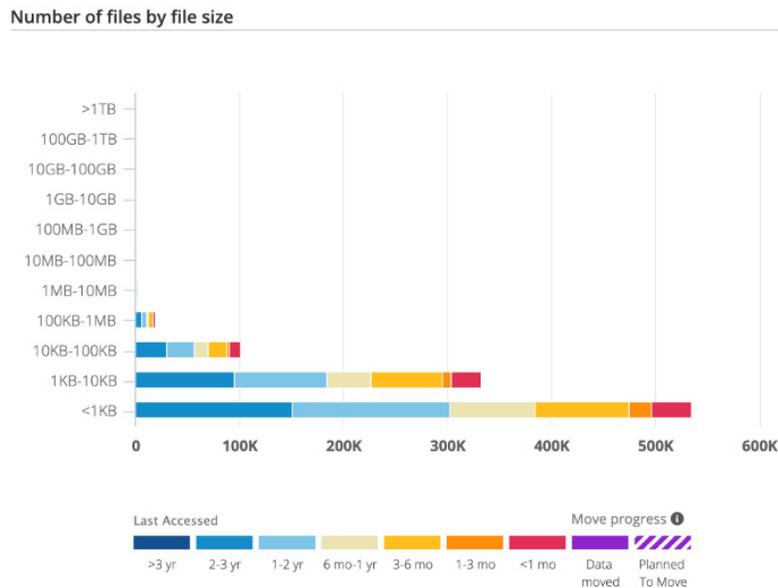


Figure 4. Komprise analytics report showing the number of files for files of each size

Migration console Figure 5 is a sample screenshot of the UI of the migration tab. This figure shows a summary of all the ongoing migration tasks and their status. It also shows details of one particular migration task. Note that for this task, it shows details of the multiple migration iterations and details of any errors or issues found during these iterations. A CSV export providing details of a single migration task or all the tasks can be downloaded. This UI makes it easy to manage and report on hundreds of simultaneous migrations. All migration functions, including creating a migration task, obtaining its status and error reports, and stopping or pausing it, can be done via an API. This makes it possible to integrate Komprise Elastic Data Migration into an existing management platform.

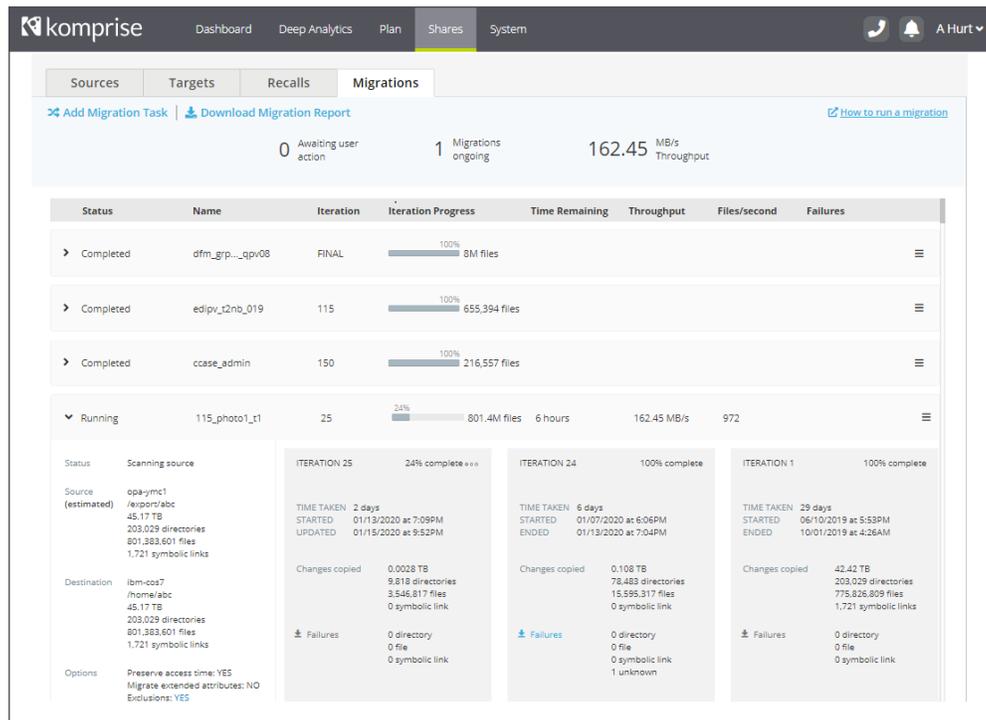


Figure 5. A sample screenshot showing Komprise Migration tab UI

The migration environment

We ran multiple iterations of migrating the Android data set using the open source *rsync* migration utility and using the Komprise migration solution. The environment in which the tests were conducted consisted of a source file server with 10K SAS drives and a flash-based destination file server. They were connected via a 10 Gbps Ethernet link. A WAN was simulated using a WAN emulator between the migration software and the destination file server. A latency of 30 ms was used to simulate WAN conditions.

The results

When tested on the LAN with the Android data set, Komprise was consistently **27 times faster** than *rsync*². While *rsync* took several hours to complete over a LAN, Komprise completed the migration task in a few minutes.

When we ran the migration tests over the simulated WAN connection, *rsync* could not complete the test even after 48 hours. In comparison, Komprise was still able to complete the task in minutes. As the number of Komprise Observers scaled from 1 to 4, Komprise was able to **cut the degradation** due to higher latencies on WAN **by 250%**.

Komprise beat *rsync*

- 27 times faster
- Cut degradation from WAN by 250%

The table below shows the test results, both with and without using Komprise Elastic Data Migration as WAN degradation relates to the LAN performance.

# of Komprise Observers	Elastic Data Migration	Degradation due to WAN
1	Off	5X slower
2	On	3X slower
4	On	2X slower

². Results were obtained in an isolated, high performance lab environment. Customer experiences may vary depending on a variety of factors, including but not limited to: the specific source and destination filer servers, the load upon those servers, the network environment and performance, the specific virtual machines on which the Komprise virtual appliances are run, and the specific data set migrated.

Summary

Komprise Elastic Data Migration provides high-performance data migration at scale, solving critical issues that IT professionals face with these migrations. Komprise makes it possible to easily run, monitor, and manage hundreds of migrations simultaneously. Unlike most other migration utilities, Komprise also provides analytics along with migration to provide insight into the data being migrated, which allows for better migration planning. While Elastic Data Migration is available as a superior, standalone migration solution, it is also a part of the full Komprise Intelligent Data Management platform.

Known as a Smart Data Migration, Komprise can be used to identify and transparently archive cold data with no change in how users and applications access the cold data. This dramatically reduces storage costs without disrupting users or existing data protection workflows. Komprise Intelligent Data Management also provides Deep Analytics to answer ad hoc questions about your data and Smart Data Workflows to simplify and automate the process of finding the right data sets and enriching them for improved search and use in AI.

Learn more at komprise.com/migration



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
©Komprise, Inc. All rights reserved.