

Komprise Native-Mode File Transfer

April 2019

Komprise Solution Overview

Komprise Intelligent Data Management enables businesses to manage their data intelligently by analyzing and transparently archiving, replicating, and managing data. To reduce the amount of high-performance storage required (since high performance generally comes at a high cost), Komprise identifies stale data across the storage landscape, transparently tiers this data to cheaper secondary storage of the customer's choice, and conducts ongoing lifecycle management to move the data to further cheaper storage classes (e.g., cheaper object or cloud or NAS) – regardless of vendor.

Komprise's patented solution performs this data management without any changes to the hot data and metadata paths, without any scaling limits, without any storage agents, without any costly hardware or infrastructure, and without any vendor lock-in. Komprise can also put a copy of the data per customer policy to another storage target of the customer's choice and incrementally update the copy on a nightly basis. This feature provides data replication at a fraction of the cost and without the lock-in of an identical mirror.

Komprise is designed for massive scale, using a distributed architecture that consists of one or more Komprise Observer virtual machines running at the customer site connected to a Komprise Director that can run either as a cloud service or on-premises. The Observer virtual appliances run in the background without interfering with active work on the storage and the network, and are organized in a lightweight grid that enables an adaptive, scale-out architecture that can be throttled back as needed.

Komprise Data Transfer

When Komprise moves stale data from a primary storage source to a secondary storage target, the target can be on-premises NAS, on-premises object storage, cloud-based object storage, or tape. When the secondary storage target is object storage (whether on-premises or cloud-based), Komprise can transfer the file in one of the following data formats:

- Chunked (with or without encryption)
- Chunked and compressed (with or without encryption)
- Native



Chunked data format

In chunked format, when files are transferred to secondary target object storage, they are stored as a set of chunks, each typically 5 MB in size. (The chunk size is configurable.)

Chunked and compressed data format

In chunked and compressed format, when files are transferred to secondary target object storage, they are stored as compressed chunks to reduce file size and thus the required storage space on the secondary storage target. Handling the compression and decompression of compressed chunks, however, does consume compute resources and can affect file access performance.

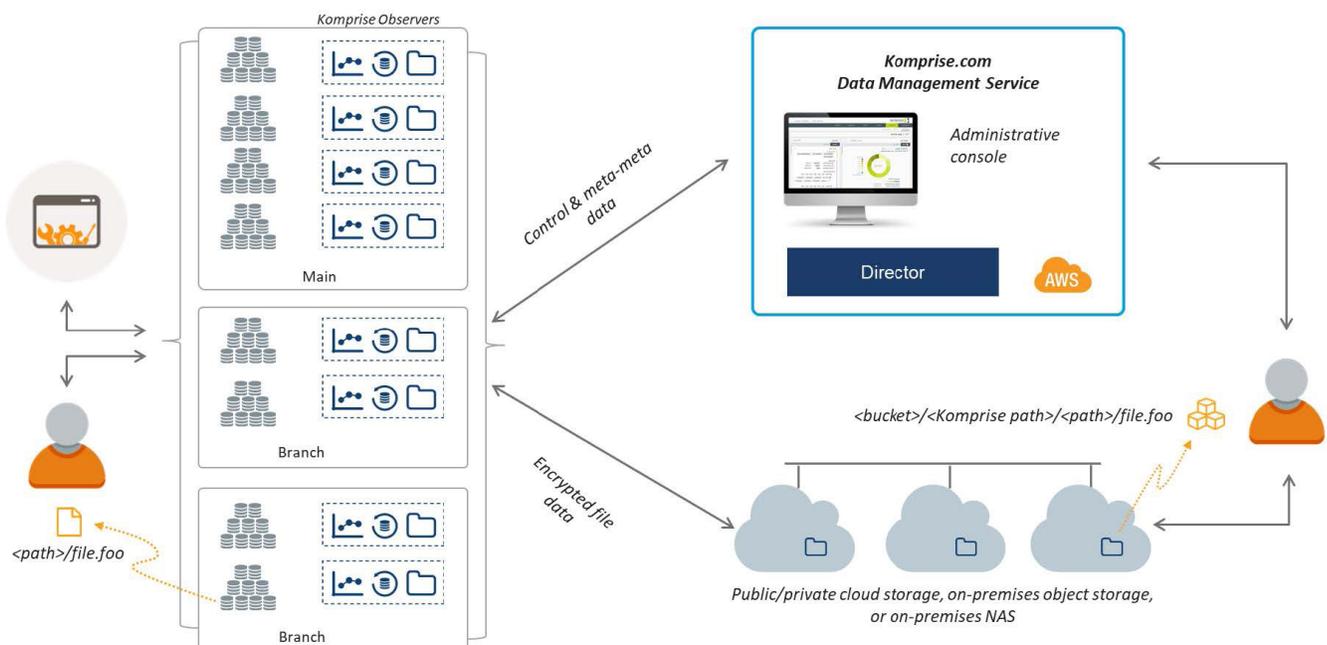
Encrypted format

All Komprise data transfers are sent using secure data transmission (SSL over HTTPS). In addition, for the chunked and chunked-and-compressed data formats, chunks can be encrypted (uses AES 256-bit symmetric key) prior to data transmission for customer-controlled, end-to-end file transfer security.

Native data format

In native format, when files are transferred to secondary target object storage, they are stored as individual objects. Native mode enables users to directly access the files (as objects) from the target storage. One current limitation of the Komprise solution is that once files are stored natively on the secondary target object storage, they become read only. Direct editing of files (objects) on the target storage must be prevented, else Komprise will no longer be able to provide access to those moved files.

The diagram below shows the Komprise solution architecture, as well as how users can access files stored in secondary target object storage either through Komprise or natively using the secondary storage's protocols and APIs.





Advantages of native data format

There are several advantages to using native format over chunked, chunked-and-compressed, and encrypted formats:

- Use of native format prevents lock-in by enabling users and applications to read files directly on targets without going through the source or Komprise share.
- Native format optimizes PUT and GET costs when using cloud storage by storing files on targets as single objects instead of multiple file chunk objects.
- Native format transfers are 20% - 60% faster than those of other data formats since these leverage multi-part uploads and certain optimizations.

When a user wants to access a file in native mode, she would navigate to the file on the object store using a slightly different path than on the source.

Example:

Assume the original file on the source lived in <path>/file.foo. After the file is archived to the designated object store, if the user wanted to access the file natively from the object store, she would navigate to <bucket>/<Komprise path>/<path>/file.foo to access the file natively as a complete object. The <Komprise path> includes information that identifies the user's Komprise deployment, the file's source, and the archive target – information that the user can find from the Komprise UI console.

Summary

Komprise is a low impact, software solution that is designed to manage storage of hot and cold data differently and enable customers to save an estimated 70% on primary storage costs, while avoiding vendor lock-in and taking advantage of whatever storage tier makes the most business sense for them, today and in the future. With data accessible in its native format no matter in which tier it resides, Komprise customers achieve maximum flexibility for their users while driving down costs.



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](https://www.komprise.com)

For media requests email:
marketing@komprise.com