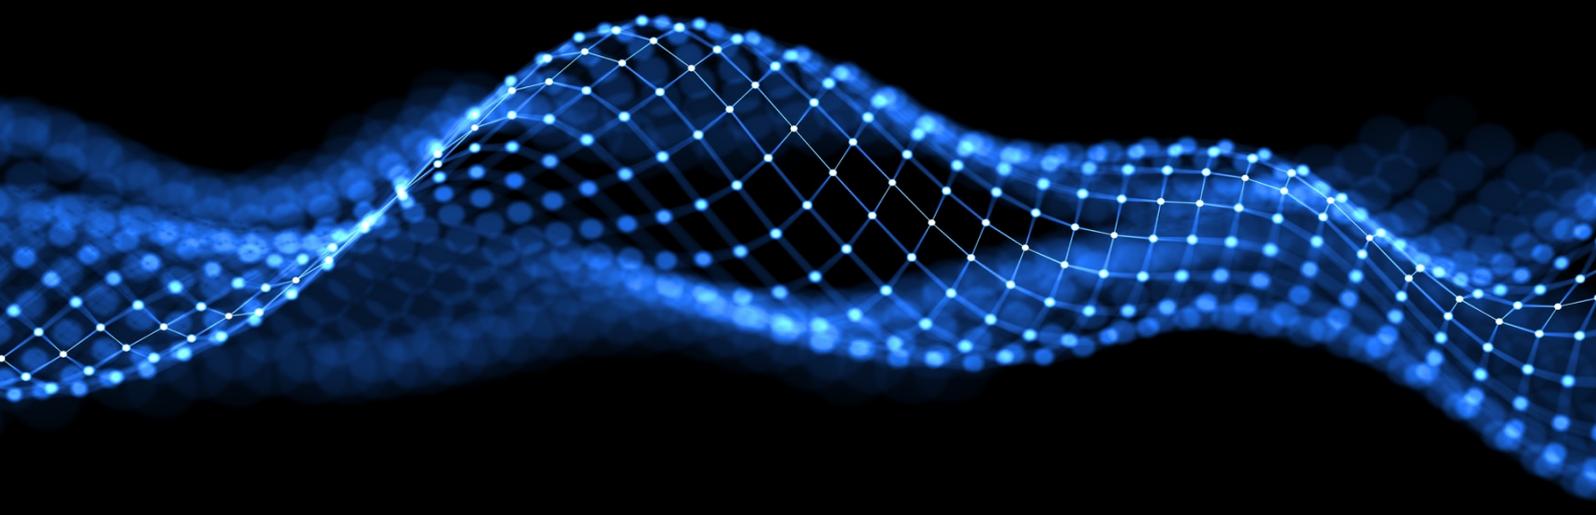


# Scality Zenko



## PRODUCT BRIEF

### The Architect's View



Zenko is an open source multi-cloud data controller, effectively a gateway between multiple object stores. The platform offers more than a gateway though, providing a translation layer between multiple object protocols and allowing users to standardise on the S3 API. With the ability to consolidate search and built-in data management policies, Zenko could be a strong solution for implementing multi-cloud object data mobility.

#### Background

The majority of new data being created today is unstructured, in the form of documents, logs, objects and other binary content. Object storage provides the capability to store vast amounts of data and the public cloud providers have been quick to capitalise on this. Amazon Web Services' (AWS) first public cloud storage offering was S3 (Simple Storage Service) – an object storage platform.

There are now dozens of object storage solutions, either in the public cloud or on-premises, with many supporting the S3 protocol. Although the S3 API is still proprietary and owned by AWS, it has become the de-facto standard for object storage support, even if vendors have their own API solutions.

Some vendors have specifically chosen not to support S3. It's not surprising to find out that Microsoft Azure's Blob (Binary Large Object) storage platform has its own protocol and doesn't offer the S3 API.

Making a choice on an object storage solution depends on a number of factors, including features, cost and reliability. Picking a specific platform is a big commitment because the effort involved in moving data between object stores is significant. Data has inertia and takes time and effort to move. Unstructured data is typically referenced by a web-based object endpoint address, which of course would change during the migration. Having data in two places at once adds complication to both applications and management.

For IT organisations thinking about using public cloud, there's also a cost perspective, as the major hyper-scalers charge for data read from an object store (egress charge). Whilst it's free to add data into platforms like S3, taking data out (or simply accessing it) incurs a charge. This makes it impractical to move data around, simply to reduce storage costs.

#### What is Zenko?

Zenko is an open source and commercial enterprise software platform developed by object storage vendor Scality. The software acts as a data controller across multiple object stores, both on-premises and in the public cloud. The aim of Zenko is to provide a location-neutral interface that allows IT organisations to gain the best use of object storage, without having to be in a constant state of data migration between one platform and another.

Zenko acts as a single point of truth for metadata and object access. The physical storage of data is abstracted away, providing only a single logical endpoint used to access content. The Zenko controller handles the writing of data to back-end storage, based on policies determined by the IT team managing the platform.

There are now three components to Zenko. CloudServer is a lightweight and free object store, compatible with Scality S3 RING and S3 API, aimed at helping developers test their code without incurring public cloud charges. Zenko Orbit is a cloud-based management portal for Zenko resources and is part of the enterprise edition. Zenko controller is the software component that exposes an S3 API, using internal components including Clueso for

metadata search and Backbeat for data workflow. Typically, the controller is deployed as a virtual instance in public cloud.

## Benefits

Rather than provide multiple endpoints for developers to use for storing content, Zenko exposes a single logical address for storing data. If another cloud offers cheaper storage, content can be written to the new location and optionally migrated in the background. This means a single logical name space can span multiple vendors, either for cost efficiency or to provide for resiliency against cloud failure.

Although cloud failure happens only rarely, the impact can be significant. End users and IT departments that want to protect against total cloud failure can use Zenko as a way to write content to multiple locations and continue to have access, even if one cloud provider fails.

One significant benefit of Zenko is in centralising search requests. When multiple object stores exist, search has to be executed against each one and the results aggregated together. With a single namespace and point of truth, search becomes much easier to implement. Search functions can be moved away from the storage platform and scaled/optimised for performance.

One final benefit to consider is the standardisation of APIs. Although many platforms support S3 as a standard, some, like Azure don't. Zenko provides a translation layer between non-standard protocols and S3, allowing developers to use a consistent interface, while gaining the benefit of using non-standard object storage platforms.

## Market Positioning

Scality positions Zenko as an open source platform that helps to drive adoption of object storage. With the release of version 1.0 in September 2018, Scality added paid support and professional services for enterprise customers. Version 1.0 supports Scality RING, AWS S3, Azure Blob Storage, Google Cloud Storage, DigitalOcean Spaces and Wasabi Cloud.

The use of Zenko is very much targeted at developers looking to have a single standard interface to test their code, irrespective of backend platform. At date of publication of this paper, Scality claimed around 2.5 million Docker image pulls of the software, however that doesn't translate to usage or storage capacity managed.

## Caveats

Having a single interface is a great idea, however the issue with this kind of technology is lock-in to the platform. Fortunately, Scality has chosen to write data at the back-end in native format, so Zenko can simply be removed from the equation without affecting any data. The loss is simply the stored accumulation of searchable metadata.

It's worth remembering that Zenko will have performance limitations as data is funnelled through virtual instances running the software. This could be a limiting factor on deployment in the future and will certainly be an area of increased development focus.

Now that Zenko has reached 1.0 and has paid support options, we can expect to see development and evolution of the product to continue.

## Reference Information

- Zenko website - <https://www.zenko.io/>
- Zenko on Github - <https://github.com/scality/Zenko>
- Presentation on Zenko at Cloud Field Day 2 - <https://vimeo.com/227505124>

Further details on Scality and Zenko can be found in the following Architecting IT blog posts:

- Scality Introduces Zenko – A Multi-Cloud Data Controller (12 July 2017) - <https://tsa.io/brk-post-8c44>
- Storage Field Day 14 Preview: Scality (2 November 2017) - <https://tsa.io/brk-post-c765>
- Seeding the Object Storage Market with Free Licences (4 September 2017) - <https://tsa.io/brk-post-84fd>
- Object Storage Capabilities Series - <https://tsa.io/brk-series-object>

Information in this document is provided for guidance only and is offered for use without warranty or other guarantees. Brookend Ltd accepts no liability whatsoever for decisions made from the basis of this information. All statements made are expressed in good faith, based on briefings and other publicly available information. Where available, links are provided to sources used for the production of this document. Scality has been or is an existing client of Brookend Ltd.

Copyright © 2018 Brookend Ltd, all rights reserved. The exploding box logo is a trademark of Brookend Ltd, all other logos and trademarks acknowledged to their respective owners. DocID: BRKBD0103. Last revised 5 October 2018. Revision 1.1.