OpenIO object storage tiering

March 2018

Table of Contents

Introduction

Managing petabytes of data is challenging, in terms of both cost and efficiency. Many storage solutions try to solve these challenges, but rarely deliver. Scalability is only one aspect of the problem, but performance, ease of use, and flexibility are equally important to make the infrastructure sustainable over time. OpenIO SDS addresses all these aspects and its object storage tiering features ensure that even the largest amount of data can be managed efficiently by any type of organization with a suitable $/GB.

The Challenges of Storing Petabytes of Data

Humans and machines are creating more and more data, and its volume is increasing at an exponential rate. Some of it needs to be stored forever, while, in some cases, the life span of data is just a few seconds. The way this data is accessed is also very different. The number of sources creating data is quickly increasing, and new big data and IoT applications available to a wider set of end users will contribute to this acceleration.

Storing data quickly is only the first step. Most of it becomes cold after a few hours, and only a small percentage of data is frequently accessed. In other cases, it is saved and forgotten, but some data is accessed in a repetitive or seasonal pattern; it's not unusual to see storage systems frequently accessing only 10% of the data they hold, or even less. As storage budgets remain flat, or even shrink, this becomes a problem for all enterprises, cloud providers, and managed service providers.

Addressing Data Storage Issues

The cloud is not the solution. You can surely achieve good $/GB, and TCO is attractive since you don't need to manage your own infrastructure, but there are risks associated with hidden costs, quality, and frequency of access, as well as the bandwidth required to move data around.

The same goes for traditional on-premises storage infrastructures. Building a huge storage repository on site and hiring experienced storage managers to maintain it gives the best performance and control. But this option is expensive, and if the goal is to reduce TCO, you need a highly automated system to move the needle from TB/sysadmin to PBs/sysadmin under management. Unfortunately, traditional storage arrays or NAS solutions are not designed for this approach. Limits and constraints in their architecture lead to complex infrastructures and costly over-provisioning processes to ensure sufficient capacity for seen and unforeseen demand.

OpenIO proposes a different solution to this problem, with all the benefits and none of disadvantages. Instead of a one-size-fits-all approach based on traditional large storage systems or cloud only, we using dynamic tiering, which includes:

- A fast-access tier based on flash memory;
- High-capacity hard disks for near-line data;
- Cloud storage to manage peak demand as well as cold data sets;
- And tape for the best $/GB and long term archiving.

Thanks to OpenIO SDS's rule-based dynamic tiering, a company can easily orchestrate data movements through multiple media and data protection schemes, on premises as well as in hybrid environments, leaving end user- with freedom of choice and flexibility to build a sustainable infrastructure that will fit current and future needs.

The trouble with Traditional Storage Tiering

Tiering is a common solution in storage, though it is rarely executed efficiently. In some cases, the storage administrator decides where data should reside. When the process is automatic, the mechanism is usually based on a simple algorithm, which is not optimized for all possible scenarios, and this can introduce unexpected behavior and performance consistency issues for some workloads.

OpenIO's dynamic tiering solution offers the following benefits:

- Easy to use through a web GUI and APIs;
- Highly scalable and adaptable to different workloads thanks to its policy-based data movement engine;
- Flexible and configurable to quickly adapt to any ongoing scenario;
- Cluster-wide load-balancing to improve performance;
- High application efficiency thanks to transparent data access;
- Simplified creation and management of large and durable data lakes for analytics and other applications.

OpenIO Storage Tiering Architecture

Most tiering software operates only on a schedule. For example, after several days of inactivity, data is moved from the top tier to a lower level. This may be fine for email, where relatively few files are accessed after a day or two, but it can be an issue for other types of data, which can have a high access rate for specific time periods or during unexpected events. Think about movies for example; Christmas classics do not generate any traffic most of the year, but during those few days around the holidays they are watched millions of times. Generalized global tiering policies can easily create performance and latency inconsistencies when a file, such as a video, is retrieved from a slow tier or the cloud before starting the streaming process.

OpenIO SDS's smart backend architecture implements a sophisticated dynamic tiering mechanism that takes into account several parameters including time, metadata, content, data location, events, and more. It is based on indirection tables that include pointers, which address storage spaces located on several tiers (content, container, account, namespace). The administrator can fine-tune the storage policies to send objects onto different tiers based on their content. Data placement can occur either while new data comes in or afterward, through specific data movers.

OpenIO SDS can be configured to have several different performance tiers in a single namespace. It is even possible to mix hot and cold data on a single storage platform without additional software or increasing complexity.

OpenIO SDS's dynamic tiering features include:

- Dynamic policy settings, which can be easily changed and can evolve over time;
- Flexible data placement while ingested (when stored) or later using data movers;
- Indirection tables with pointers enable fast and scalable tiering across storage pools;
- Stretch cluster mode for local or long-distance replication across tiers;
- Granular policy configurations according to metadata (date, pattern, namespace, containers);
- Cold data compression to reduce capacity footprint and improve storage efficiency.

Storage policies define tiering parameters, the level of security to be implemented, the data to be replicated, the kind of data protection mechanism used, and how the data has to be compressed, deduplicated, encrypted, or otherwise processed. All policies are namespace-wide.

OpenIO SDS's dynamic tiering technology is platform-independent and works seamlessly on different CPU and media types (such as Kinetic drives), allowing end users to choose what is best in their current situation, while leaving the door open to future improvements.

Hybrid Cloud for Bulk Content and Big Data

OpenIO SDS's dynamic tiering is ideal for building a hybrid infrastructure, thanks to public cloud providers such as Backblaze B2 or Amazon AWS S3. By offloading some storage capacity to an external cloud provider, end users can seamlessly manage peak demand as well as benefit from the low cost offered by this type of service, especially for cold data that will probably never be accessed again.

Many environments still rely on magnetic tape or optical storage for data protection (backup, BC/DR, archiving) as well as bulk storage for big content or other big data needs. By adopting OpenIO SDS, it is possible to replace legacy data silos, such as tape libraries, and use the cloud instead; or use them as a tier alongside the cloud.

In addition to leveraging performance acceleration through cache buffers, data is secured through encryption. The data footprint impact is reduced by the use of compression and parity-based erasure codes; more data can be stored with less overhead. The key to this approach is automated tiering with configurable policies that define how, when, and where data and storage are managed.

OpenIO SDS Dynamic Tiering for Complex Email Environments

One of the most interesting applications for OpenIO SDS's dynamic tiering is for large mailbox storage. It can reduce storage cost by up to 90% while simplifying operations, dramatically lowering TCO.

Fast SSDs can be used for recent and hot email conversations while larger and cheaper SATA disks can store the rest. System administrators can easily enable automated tiering and define policies to balance resource and application requirements. For example, 10% of data can be active and placed in the fastest media pool (first tier). It can then be moved automatically to a second, high-capacity tier – slower and cheaper - after a few days without changing how it will be accessed by the application.

To achieve the best performance, the SSD pool can be configured with object replication (creating three data copies). To maximize capacity utilization, data saved on hard disks in the secondary pool can be protected using 8+3 or 14+4 Erasure Codes (EC), along with compression. This can reduce cost while improving the data footprint efficiency and enhancing availability.

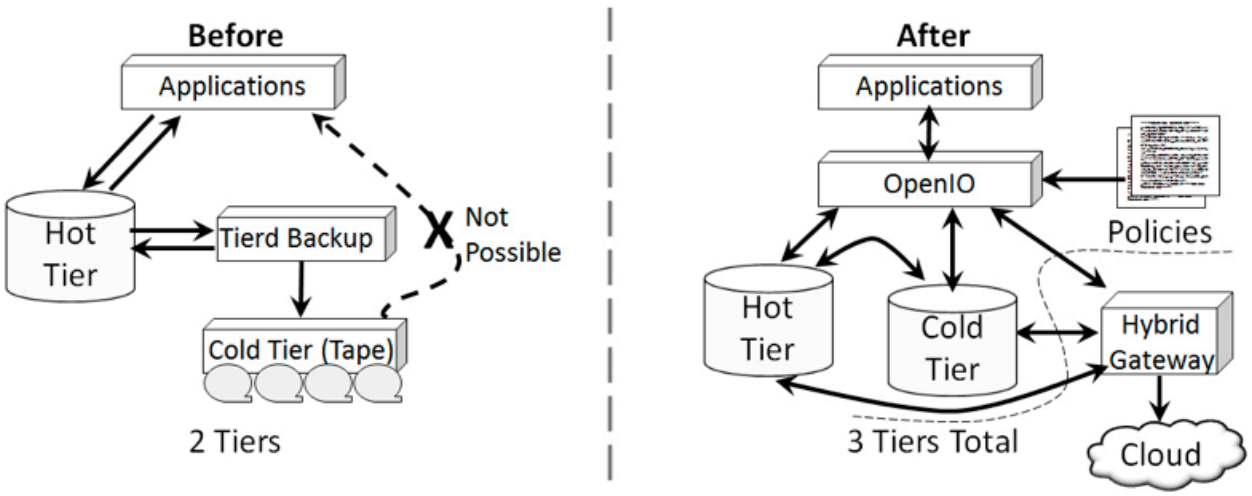

Figure 1 : Before and after OpenIO object storage tiering is enabled

The benefits of OpenIO SDS dynamic tiering for email infrastructures include:

- Improved overall storage performance compared to large HDD-based solutions;
- Lower storage costs thanks to the balanced use of SSDs and Disks;
- Increased storage capacity thanks to erasure codes and compression of the capacity tier;
- Improved control over data placement, durability, and performance;
- Transparent to applications and end users.

OpenIO SDS uses faster flash SSDs more intelligently to boost effectiveness while leveraging the capacity of lower cost SATA disks for efficiency. An SSD storage tier for hot data is used to store eight days of email. Leveraging user-defined storage QoS policy rules, older emails are transparently re-tiered to the slower, lower cost SATA pool. The fast tier is protected with mirror replication. The slower pool is protected with erasure coding, and compression further reduces costs.

Summary

OpenIO SDS's robust, scalable object storage leverages dynamic tiering to improve storage efficiency and reduce costs. Automatic tiering policies, defined by system administrators, take into account multiple parameters to ensure optimal data efficiency, performance, availability, and durability. Data is automatically moved from fast, expensive media to cheaper, large-capacity disks and cloud storage. This is fully transparent to applications, and improves overall infrastructure TCO. OpenIO SDS's scalability and flexibility allow end users to increase the size of their media pools without impacting performance or incurring painful rebalancing operations.

OpenIO has been working with object storage since 2006. Since the first production version of SDS, in 2008, this solution has been deployed in environments of all sizes, from 3 nodes up to 10+ PB infrastructures. Its use has spread from email to the consumer cloud, video archiving and streaming, video surveillance, healthcare applications, voice call recordings, and more. OpenIO's solution has been open source since 2012, and it has engendered a steadily growing community.

The future of storage is here. Find out more at www.openio.io, and follow @openio on Twitter.

Next-Generation Object Storage and Serverless Computing

openio.io

FR
2 Bis Avenue Antoine Pinay
59510 Hem

FR
37 Rue de Liège
75008 Paris

JP
1-35-2 Grains Bldg. #61
Nihonbashi-Kakigara-cho,
Chuo-ku, Tokyo , Japan 103-0014