Use Cases

March 2017

Table of Contents

Introduction

OpenIO has selected a number of use cases to illustrate the value of OpenIO SDS and the positive impact it can deliver when deployed alongside traditional and next-generation business applications.

OpenIO SDS is a pure software-defined storage application. It is open-source object storage software that transforms commodity servers into large storage pools. Its lightweight backend architecture, smart automation functions, easy to use interface, and broad feature set make OpenIO SDS the ideal tool for the most demanding IT environments. And thanks to its growing partner ecosystem, OpenIO is now able to build end-to-end integrated and comprehensive solutions.

Email, video, active archive, big data lake foundations, and generic object storage use cases highlight the benefits of OpenIO SDS.

Email

OpenIO has a proven track record with email services, having worked and sustained large email services for more than a decade for major telecom operators in France. The product currently supports more than 45 million users for web mail applications and associated cloud storage service.

The Challenge

Email continues to be at the heart of many businesses, and it is crucial that email storage solutions prevent data loss and ensure uninterrupted service. The infrastructure supporting the email service must also scale transparently within a data center and across geographies. The workload is essentially made up of small operations, manipulating small files – a few KB each – accessed concurrently, with email objects usually read in the first few hours and then archived on servers.

End users take state-of-the-art email services for granted nowadays; they expect excellent uptime, virtually unlimited quotas, high speed, and bulletproof reliability. But, for service providers, meeting these expectations isn't easy, especially at scale and when the service is provided for a few dollars per year or, worse, for free. Infrastructure sustainability in terms of overall TCO and $/GB is fundamental.

Traditional storage infrastructures aren't flexible enough, and operating costs are too high. Large-scale infrastructures standardize on commodity hardware, and companies like to keep this hardware homogeneous, with single components, or vendors, so it is easy to swap when necessary. For large corporations or telcos who already have strong vendor agreements, this flexibility is clearly an advantage.

As an infrastructure service provider, the storage layer should support multiple messaging engines and protocols, both commercial and open source, based on adoption trends and product capabilities. At the same time, sophisticated data services implemented on the storage layer are now needed to offload demanding tasks such as security threat detection without impacting mail server performance.

The Solution

To make email storage infrastructure flexible and sustainable, a software-defined storage solution, like OpenIO SDS, capable of running on any commodity hardware, is the best solution. OpenIO SDS is hardware agnostic and thus ensures a complete separation between hardware and software. Thanks to its Conscience technology (a broad set of sophisticated load balancing algorithms), OpenIO SDS can run on mixed-node configurations, allowing end users to expand the storage system with nodes of different size, capacity, CPU architecture, power, and hardware generation.

OpenIO SDS has been designed with hyper-scale in mind. It is highly automated, and its capacity, number of nodes, and performance are limited only by the resources available. But no compromise has been made in ensuring data protection and efficiency. OpenIO SDS offers data replication and erasure coding to deliver best-in-class data durability, as well as dynamic automated tiering and compression to improve data footprint efficiency and overall $/GB. System efficiency can be driven by Grid for Apps, an event-driven application framework included in OpenIO SDS, which take advantage of unused cluster computing resources to run specific applications on the storage system, directly where data resides. OpenIO SDS can offload tasks like spam filtering or threat detection from the mail server to the storage platform.

In terms of messaging engines, OpenIO SDS is application-independent and supports key vendors, allowing end users freedom of choice. OpenIO and its partner ecosystem have developed dedicated connectors and integrations for messaging systems with Synacor Zimbra, Cyrus, Dovecot/Open-Xchange, and Synchronoss (Openwave Messaging).

Benefits and advantages

- Proven large deployments with tens of millions of users;
- Standard commercial and open source email engines;
- Pure software-defined storage solution;
- Open source software;
- High data durability with replication and erasure coding techniques;
- Low $/GB thanks to compression, erasure coding, and dynamic tiering;
- Effective TCO.

Digital content

OpenIO SDS's unique architecture can be used to build strong back-end infrastructures for media content services of any size. Content distribution networks, direct streaming, rendering, and near-line archives are currently the most common such applications. With the increase of rich media files in enterprise environments, more and more use cases are surfacing, showing the inadequacy of traditional storage infrastructures.

The Challenge

Digital content is an important part of various business activities, and is increasingly considered a critical corporate asset.

Rich media content (including videos, audio, and graphics) requires a storage platform capable of serving files quickly and through different access methods, depending on available bandwidth and remote device capabilities. At the same time, low $/GB and high resiliency are key factors.

Applications of this kind generate a specific type of workload: mostly sequential large file reads, but, depending on the size and type of content, IOs can be smaller and scattered. Concurrent access to many files from several locations and devices at the same time is common, and, in order to ensure that all devices can access the files, the same content must be accessible through as many protocols as possible (NFS, SMB, S3 APIs, etc.).

The Solution

Deployed on commodity x86 and ARM servers, and using standard Linux distributions, OpenIO SDS is a complete object storage solution that includes an event driven application framework. This core feature, named Grid for Apps, allows data storage and applications to run on the same hardware to boost efficiency and reduce TCO. For example, video files can be encoded in real time when streamed, or multiple versions can be created for different purposes while ingested.

OpenIO SDS provides a broad set of industry-standard file sharing protocols including SMB, NFS, FTP, and AFP in addition to native file access based on FUSE for local access. Remote internet access and application integration is enabled through native APIs and libraries, as well as native and REST APIs, including S3 and Swift.

Scalability, reliability, and the high parallelism made possible by Conscience technology are other key characteristics of OpenIO SDS. It allows businesses to build media content repositories of any size, and its high flexibility allows the platform to adapt to several use cases.

Benefits and advantages

- OpenIO Grid for Apps consolidates applications and data storage for unmatched efficiency;
- Conscience technology provides continuous, automated load balancing for the highest parallelism, throughput, and consistent performance;
- This mature solution is deployed with several video services that require fast file transfers and media asset management;
- Pure software-defined storage solution leverages commodity hardware;
- Open source software;
- High data durability thanks to replication and erasure coding techniques;
- Transparent dynamic tiering so end users have the best combination of $/GB and performance;
- Low TCO.

Active archive

With organizations of all kinds storing more and more data for longer periods of time, it is essential to implement a storage system that can handle large amounts of data, scale over time, and perform retrievals as transparently as possible. In this scenario, a low $/GB is fundamental for infrastructure sustainability, while file and API-based access methods can help end users access data from legacy as well as modern web-based and mobile applications.

The Challenge

An efficient active archive offers a low $/GB combined with the ease of use of standard file and API interfaces. It must provide a transparent tiering mechanism to take advantage of faster media for frequently accessed data, as well as slower, high capacity drives, cloud, and tapes for files that won't be accessed for a long time.

Workload is characterized by a majority of writes and relatively few reads, but data durability and the lowest possible TCO remain essential. In highly regulated environments, WORM and compliance are also important.

The Solution

OpenIO SDS is the platform of choice for secondary storage workloads, especially for data archiving. OpenIO SDS runs on commodity servers leveraging standard hardware components such as low-cost, high-capacity hard drives. OpenIO SDS can build resilient clusters of independent servers, and deliver multiple-nine data durability with replication and erasure coding. WORM is also offered as part of the data storage platform, and compliance for specific regulations can be added easily with external third-party partners or directly by the platform.

SMB, NFS, AFP, and FTP are standard access methods alongside proprietary and standard APIs like Swift or S3.

OpenIO has launched a partner ecosystem to validate, certify, and accelerate market adoption. Among them, secondary data software players are represented by multiple occurrences.

Benefits and advantages

- OpenIO uses industry-standard file sharing protocols, such SMB, NFS, AFP, FTP, and FUSE for seamless data access;
- Long term retention with WORM capabilities;

- Automated dynamic tiering for the best $/GB and high performance for frequently accessed files;
- Software-defined storage solution leverages commodity servers;
- Open source software;
- Data availability and durability, thanks to sophisticated replication techniques and erasure coding;
- Low TCO.

IIoT and Big Data

We are entering in a new world full of small, isolated devices that can produce huge amounts of data, but, while connectivity is evolving at a similar pace, there are many reasons to store, preserve, and process this data where it is created. At the same time, consolidating data in large repositories (data lakes) enables users to run big data applications and get the most out of their data.

The Challenge

Building an Industrial IoT strategy is a top priority for many large organizations, but there are several challenges to consider. IoT devices are generally small, and local resources are limited. This means that data resiliency, security, and high availability are hard to achieve. A seamless, robust, and scalable data plane is the key to building infrastructures capable of supporting next-generation IoT applications as well as big data analytics in large data centers.

The Solution

OpenIO SDS is a unique and flexible object storage platform that can be installed on the smallest IoT devices as well on powerful data center servers. This makes it possible to build remote, distributed storage infrastructures for small devices and connect them to central, high-capacity repositories for computation-intensive applications. Standard protocols such as native HTTP interfaces, S3, and Swift ease data access, while Grid for Apps makes it possible to run applications triggered by events directly where data resides.
Thanks to the unique Conscience technology and its unmatched flexibility, OpenIO SDS and Grid for Apps are suitable for the most advanced IoT strategies.

Benefits and advantages

- Pure software-defined software available for small devices and large data center servers;
- The most flexible object storage infrastructure on the market;
- Conscience technology for the best resource utilization;
- Grid for Apps framework to process data where it is generated, consolidated, or where it resides;
- Open source solution;
- Same data plane designed to connect the smallest of IoT devices directly with the largest of data lakes.

Backup repository

Traditional backup repositories, such as VTLs, have limits and constraints when it comes to scalability and TCO. At the same time, their data footprint reduction capabilities are reduced because compression and deduplication is now included in many backup platforms, saving bandwidth and improving backup windows. In addition, new data formats leverage compression and encryption and cannot be compressed any more. Object storage can be the right solution to mitigate these issues and manage short- and long-term backups through a combination of on-premises, low-cost storage, together with tape and cloud servers.

The Challenge

Organizations of all sizes are looking for new backup solutions that offer scalability, ease of use, and a low $/GB. They want to manage short-retention backups on disk for fast recovery, but they also want a second tier in the cloud or on tape for long-term backup or archiving purposes. This kind of infrastructure should be easy to manage, and allow remote replication as well as electronic vaulting.

The Solution

Instead of stacking several VTL appliances on file-based repositories, OpenIO SDS can be used to create enterprise backup solutions compatible with S3 and Swift APIs to store from a few hundred TB up to several PB in a single domain . Dynamic tiering can be used to move older backups to tape or to the cloud transparently. Thanks to this approach, and the nature of the protocols involved, OpenIO SDS can also be configured as a central repository for remote backup sites allowing end users to simplify disaster recovery procedures.

Benefits and advantages

- Erasure coding for the best $/GB;
- Several media types supported for on-premises tiering;
- Cloud gateway for hybrid cloud configurations and dynamic tiering;
- Lower TCO compared to traditional backup solutions;
- Open source and software-defined solution.

Private Cloud

OpenIO SDS is the perfect solution for building a private cloud storage infrastructure and enabling the delivery of next-generation storage and collaboration services to organizations of all sizes. Object storage is quickly becoming the persistent storage layer of choice for next-generation applications, and, thanks to its characteristics of scalability, performance, efficiency, and ease of use, OpenIO SDS can be quickly adopted for a single service (sync & share, for example) and grow from there through consolidation of legacy and new services.

The Challenge

Distributed enterprises need new forms of storage to cope with high-demand applications for the web and for mobile users. With users and applications accessing data from any location at any time, and from any device, data mobility is essential to improve productivity, and therefore competitiveness. To control data while maintaining costs at a reasonable level, it is necessary to change the way storage is accessed and consumed. Traditional storage systems aren't designed to work with internet protocols (like S3 APIs, for example), and aren't able to distribute data over geographies to make it quickly and safely accessible from anywhere.

The Solution

OpenIO SDS is a distributed storage system that can be accessed via traditional protocols as well as APIs like S3 and Swift. Its architecture is designed to allow end users to deploy it on local, stretched, and geo-distributed configurations, so data can be always accessible and available from any kind of application at anytime, and from anywhere.
External remote gateways, as well as specific collaboration software (like SME enterprise Sync&Share), consolidate large amounts of data in a single logical domain, protected through multiple replicas or erasure coding, while making it accessible from remote locations (such as and mobile devices).
Multitenancy is another key characteristic of OpenIO SDS. The solution can run applications and workloads enabling the consolidation of data from several traditional storage systems, lowering overall infrastructure TCO.

Benefits and advantages

- Flexible configuration for local, stretched, and geo-distributed deployments;
- Multitenancy for consolidating data and applications;
- S3 and Swift compatibility for broad third-party application and gateway support;
- Certified enterprise Sync&Share solution;
- Back-end object storage for web and mobile applications;
- Lower TCO compared to traditional storage.

object storage grid for apps

openio.io

FR
3 avenue Antoine Pinay
Parc d'Activité des 4 Vents
59510 Hem

US
180 Sansome Street, Fl4
San Francisco, 94104 CA

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