zadara

OBJECT STORAGE

WHAT IS OBJECT STORAGE?

Object Storage is an alternative way to store, organize and access units of data. It provides a reasonable balance between performance and functionality versus simplicity and scalability. Object Storage enables a minimal set of features: store, retrieve, copy, and delete objects. These basic operations are done via REST APIs that allow programmers to work with the objects. The HTTP interface to Object Storage systems allows fast and easy access to the data for users from anywhere in the world.

THE DIFFERENCES BETWEEN OBJECT BLOCK, AND FILE STORAGE

Object Storage is much more scalable than file storage because it is vastly simpler. Objects are not organized in hierarchical folders, but in a flat organization of containers or buckets. Each object is assigned a unique ID or key. Their keys, regardless of where the objects are stored, retrieve objects. Access is via APIs at the application level, rather than via OS at the file system level. As a result, Object Storage requires less metadata, and less management overhead than file systems. This means Object Storage can be scaled out with almost no limits. Object Storage is easier to use than block storage and overcomes the limitation of fixed size LUNs. It also removes file system limitations such as the folder size or path name length. Unlike block or file, Object Storage does not use RAID for data protection. It simply keeps a number of copies of each object.

EXECUTIVE SUMMARY

In today's data-driven world, enterprises are generating and storing vast amounts of data that must be secured, managed, and accessible at all times. Object storage has emerged as a leading technology for managing unstructured data at scale, providing a highly scalable and flexible solution for storing, managing, and accessing large amounts of data. However, traditional cloud object storage solutions have often been associated with performance and security limitations. This white paper presents Zadara's on-prem object storage solution that addresses these limitations, providing high performance, security, and resilience to meet the needs of today's MSPs and enterprise customers.

Cloud providers using Zadara's technology can now offer S3/Swift compatible object storage. Object Storage is a simple, extremely performant, easy to use, and well protected solution for storing any type of unstructured data, multimedia data, sensors sampling big data, or archiving and backup of any sort. Zadara's Object Storage runs side by side with Zadara's VPSA (Virtual Private Storage Array) on Zadara's zStorage clouds, utilizing the same infrastructure, and providing higher value to our customers.

WHAT IS PRIVATE OBJECT STORAGE?

Private Object Storage is an Object Storage solution built on dedicated compute resources either in a private or public cloud. As such, it guarantees better isolation and security to ensure no performance impact by other users of the object storage. Object Storage in a private cloud enables all of the benefits of the public Object Storage. However, it eliminates the latencies associated with the Internet and the public cloud, and therefore, provides much better performance.

OBJECT STORAGE USE CASES

Object Storage can be used in variety of applications and use-cases, such as:

- **Backup and Disaster Recovery:** Private Object Storage can be used for backup and disaster recovery purposes, where data can be stored locally and replicated to remote locations for redundancy and disaster recovery purposes.
- **Cost Savings:** For organizations with large amounts of data, private object storage can be a more cost-effective solution compared to cloud storage or traditional file and block services, especially for long-term storage.
- **Hybrid Cloud Architectures:** On-prem object storage can be a foundation for a hybrid cloud architecture, allowing organizations to store some data on-premises and some in the cloud, and move data between the two environments as needed.
- **CDN:** CDNs work by caching web content in multiple edge locations (i.e., geographically distributed data centers), so that users can access the content from the location that's closest to them. This reduces latency and improves performance, especially for users who are located far away from the origin server where the content is stored. Object storage is a good fit for CDN deployments, as it provides scalable and durable storage for web content, with low latency and high availability. CDN edge servers can fetch content from the object storage system quickly, without putting too much strain on the origin server or the network.
- **Big Data:** Object storage is also ideal for big data applications that require distributed and parallel access to data, such as Hadoop, Spark, or other data processing frameworks. Object storage allows faster data access and processing. This enables organizations to build big data solutions that can scale to meet their needs, without having to worry about storage limitations or bottlenecks.
- **Archiving:** Object storage is well-suited for long-term archiving of data, such as legal and financial records, medical images, and scientific data. Object storage provides the durability and reliability required for long-term data retention.
- **Cloud Applications:** Object storage is a popular storage solution for cloud-native applications because of its scalability, durability, and flexibility. Cloud applications can store and retrieve data from object storage using simple RESTful APIs.
- **IoT Data Storage:** Object storage can be used to store data generated by IoT devices, such as sensors and smart devices. Object storage provides the scalability and flexibility required to store and process large volumes of data generated by IoT devices.
- **Data Sovereignty and Compliance:** Private Object Storage when designed with local governance in mind can help organizations comply with data sovereignty and compliance regulations, as data is stored within their own data center and under their own control.

OBJECT STORAGE CHARACTERISTICS

Zadara Edge Cloud's Object Storage service is provided on Zadara's zStorage clouds, side by side with the VPSA (Virtual Private Storage Array) that provides block and file storage services.

These are the key properties of Zadara's Object Storage:

- **Flat Namespace:** In object storage, each object is identified by a unique identifier or key, and the object itself contains the data and metadata associated with the object. Unlike traditional file systems, where files are stored in a hierarchical directory structure, object storage uses a flat namespace, which means that objects can be stored and retrieved independently of any hierarchical organization. This flat namespace allows easier scaling and management of large volumes of unstructured data. Because each object is identified by a unique key, there are no limits to the number of objects that can be stored in the system, and objects can be easily located and retrieved by their key, without the need for navigating a complex directory structure.

- **Unlimited Scalability:** Object storage is known for its unlimited scalability, which is one of its key advantages over traditional storage solutions. Object storage systems can be designed to handle extremely large amounts of data, making it ideal for use cases where data volumes are constantly growing, such as big data analytics, scientific research, and media and entertainment. Unlike traditional storage systems, Zadara's Object Storage can scale almost infinitely by adding additional drives to the cluster. As more drives are added, the capacity of the system increases, and the system can continue to operate without interruption, providing seamless expansion to meet growing data needs.

- **High Availability and Resiliency:** Zadara's Object Storage provides high availability with no single point of failure on both physical and logical layers, and data resiliency, as data is distributed across multiple nodes within the data center.

- **Erasure-Code Data Protection (or 2-Way Mirroring):** Erasure code is particularly good for durability because it can provide a high level of fault tolerance while using less storage capacity compared to traditional data replication methods. For example, a typical replication method would require storing multiple copies of the same data on different storage nodes, which can be costly in terms of storage capacity. However, with erasure code, redundant fragments of the data are generated and distributed across multiple nodes in a way that allows the original data to be reconstructed even if some of the nodes fail. This not only improves durability but also reduces the cost of storage.

- **S3 and Openstack Swift API Support:** S3 & Openstack Swift API are highly popular and widely used object storage API that provides a wide range of capabilities for managing and accessing Object Storage stored data. The API provides a simple and intuitive way to interact with the storage, allowing developers to easily store and retrieve data, as well as perform other common operations like listing objects, deleting objects and managing access. One of the key benefits of the S3 API is its compatibility with a wide range of programming languages and tools, making it easy to integrate with existing applications and workflows.

- **Object-Lock Support:** Zadara Object Storage Immutability ensures data integrity by stopping stored objects from being deleted or overwritten during a specific retention timeframe. With Object Storage Immutability (Object Lock) enabled on an object its retention mode can't be changed, and its retention period can't be shortened.

- **Comprehensive Usage Reporting and Metering:** usage reporting and metering capabilities can be used to track resource usage and allocate costs to different departments or business units, making it easier to manage and optimize storage costs over time. The system exposed comprehensive metering capabilities that allow users to monitor and track resource usage in real-time.

- **Multi-Tenancy at the Account Level:** Multi-tenancy is a critical capability in many storage environments, allowing multiple users or tenants to share a single storage system while maintaining security and data isolation. Zadara's Object storage provides multi-tenancy at both the object storage and account levels, allowing multiple users or applications to share the same underlying storage infrastructure while still maintaining data isolation and security.

- **Scale-Up Performance:** the ability to scale up and down in terms of performance provides a key element in the flexibility and cost effectiveness of the solution, allowing the user to meet changing performance requirements as use-cases are added or changed over time.

PROVISIONING

Object Storage is ordered and provisioned just like the VPSA (Virtual Private Storage Array), by request at the Zadara Provisioning portal. The user selects the drive type and the capacity required to be assigned to the newly created Object Storage instance, and the system automatically assigns the needed number of Virtual Controllers (VCs). The user receives access to the object storage that contains a default account. At that point, the administrator can create accounts, assign account administrators, and can begin storing objects.

OBJECT STORAGE USAGE MODEL

Object Storage may contain any number of accounts that are assigned to different organizations. Users of each account can store objects, and organize them in containers (buckets), as illustrated in the diagram below.

- The object storage may contain a high volume of accounts (1,200) as defined by the administrator.
- Accounts are associated with an organizational unit or a department, and contain a list of user-defined containers.
- Containers are uniquely named buckets with flat structures that contain objects.
- Objects are a piece of data of any size. A key that is assigned to it uniquely identifies them.

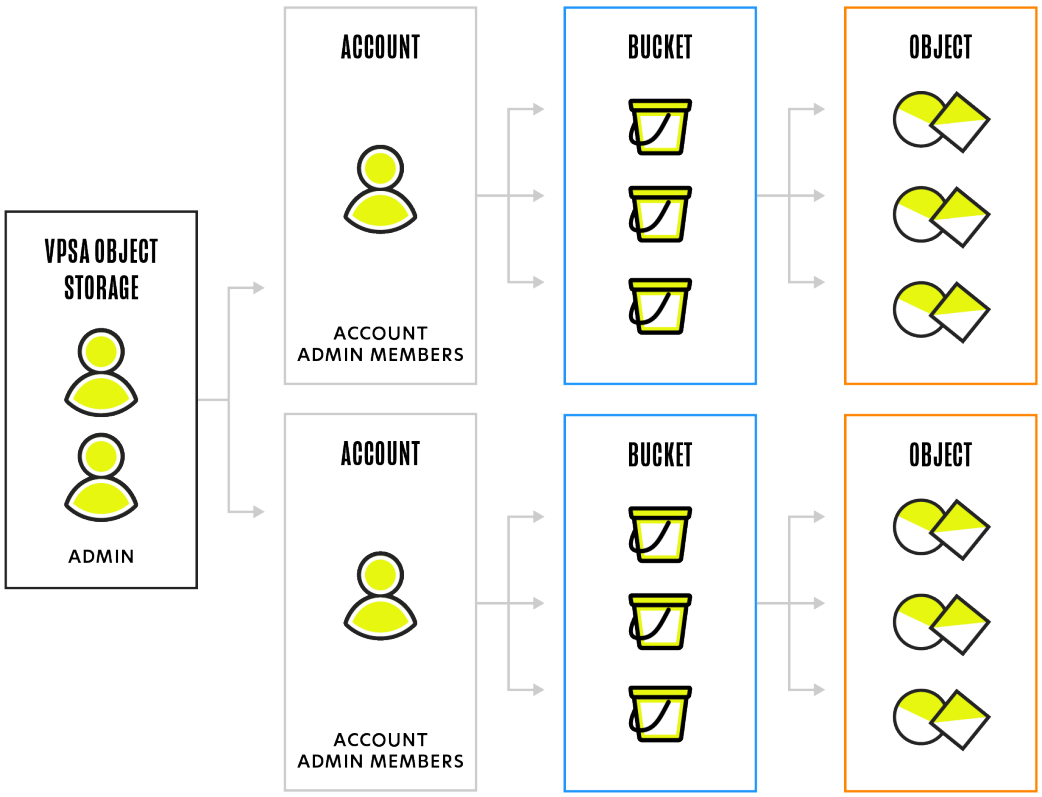

AUTHENTICATION AND PERMISSIONS

There are three types of users assigned to Object Storage:

- **Object Storage Administrator:** responsible for the administration of the object storage. The user (registered in Zadara Provisioning Portal) that orders the object storage becomes the Administrator. By default, the object storage is created with one account (Administrator account) and the Administrator is a member of this account. Administrators can add other users with the same role. Administrator is a super-user with privileges to create accounts and users of any role. Users with the Administrator role can define policies and add/remove drives. Users with an Administrator role can perform containers and objects operations across accounts. The Administrator is also responsible for the Object Storage system settings (like IP addresses, SSL certification, etc.), and has access to the Object Storage metering and usage information.
- **Account administrator:** can create an account and can manage their own account. They can perform any user management operation and containers/objects under their account.
- **Member:** can perform object storage operations according to the permission given by the account administrator, within the limits of that account. These operations include create/delete/list containers and create/delete/list objects.

DATA PROTECTION AND HIGH AVAILABILITY

Object Storage data protection and reliability is achieved by replicating objects across multiple storage nodes and disk drives. The user can configure the number of replicas to suit their needs. The minimal Zadara object storage configuration utilizes two storage nodes (for 2-Way Mirroring) or three storage nodes (for erasure code). The object storage platform (and the entire zStorage cloud platform) has no single-point-of-failure throughout its architecture).

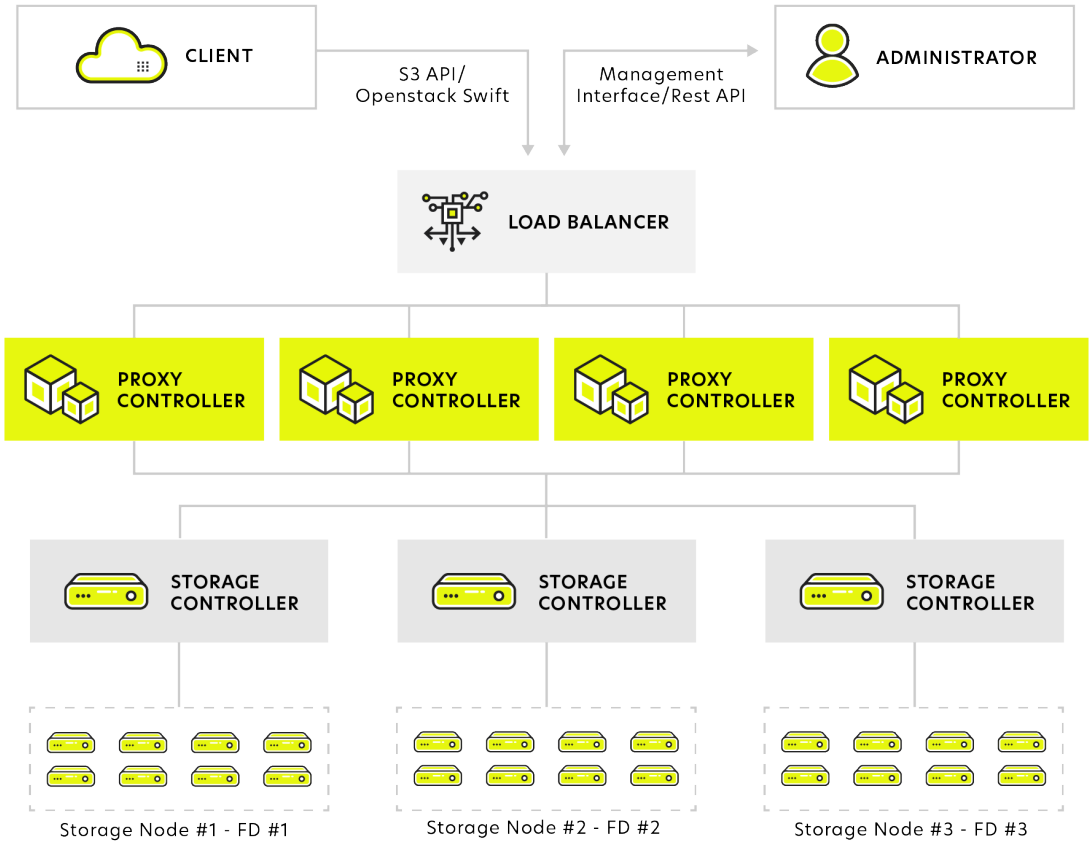

ARCHITECTURE

Zadara's Object Storage is built from two main building blocks: Proxy (Controller) and Storage (Controller) . The proxy is the interface to the users or the application using the data objects. The Proxy node receives the API request, authenticates the requester, decides what storage node should handle the coming request, and passes it over. Once the appropriate storage node handles the request, the proxy returns the answer to the caller. The storage node is responsible for storing the objects on the drives, and updating its metadata in the databases. The load balancer controls the traffic from all clients to all Proxy Nodes. The Object Storage is provisioned with an internal load balancer ready to use out-of-the-box, but users can configure the external load balancer outside of the Zadara environment, to handle the most intense performance requirements.

OBJECT STORAGE AND STORAGE ARRAYS, SIDE-BY-SIDE

Zadara's Object Storage resides side by side with Zadara's VPSA (Virtual Private Storage Array). The object storage infrastructure is a collection of Virtual Controllers (VC) and drives across the cloud storage nodes. VC's are dedicated Virtual Machines (VM) running the Proxy and Storage Nodes as described above. Just like the VPSA, Object Storage has dedicated drives assigned to it to store the data. Object Storage automatically distributes the load across the storage nodes in a way that guarantees the system durability, even in case of Storage Node failures. Drives assigned to object storage are organized in policies according to the redundancy level required (2-way or Erasure Coding). When drives are added to the configuration, Object Storage automatically redistributes the data across all drives.

Object Storage uses large SATA drives (per customer order) to store object contents, and SSD drives (automatically assigned) to store the users and container databases. By default, each Zadara Object Storage is built from two Proxy Virtual Controllers and at least two Storage Controllers. Users with high demanding performance requirements can scale-up their object storage performance by adding additional Proxy Controllers. This can be done on a self-service basis from the provisioning portal.

API INTERFACES & MANAGEMENT

Object Storage fully supports both AWS S3 API and OpenStack Swift API. For API details refer to the API documentation. Both are REST API over HTTPS.

Zadara's object storage solution supports the industry-standard APIs, including AWS S3 API and OpenStack Swift, enabling customers to seamlessly integrate the solution into their existing cloud environments. This ensures that customers can easily migrate their data, without the need for complex and time-consuming data migrations. Zadara's Object Storage solution also includes a proprietary REST API for system management, enabling customers to manage and monitor their storage infrastructure using a simple and intuitive interface.

By supporting a variety of APIs, Zadara's object storage solution provides customers with maximum flexibility and interoperability, allowing them to integrate the solution with a wide range of applications and services. This ensures that customers can easily leverage the benefits of object storage, while also minimizing the risk of vendor lock-in and ensuring that they can easily move their data between different storage systems as their needs evolve over time.

PERFORMANCE

Performance is a key consideration for any storage solution, and object storage is no exception. Zadara's Object Storage solution is designed with performance in mind, leveraging the latest hardware and software technologies to deliver extremely high performance. Our solution uses a distributed architecture that scales up and out linearly, enabling customers to achieve high levels of throughput, operations per second and low latency.

SECURITY

Data security is a top priority for any storage consumer, and Zadara's Object Storage solution is designed to meet the most stringent security requirements. Zadara's solution includes a range of security features, including encryption at rest and in transit, and access controls. Zadara's solution uses industry-standard encryption algorithms to protect data both in transit and at rest, ensuring that data is always secure and protected from unauthorized access.

To further enhance security, Zadara's solution includes access controls that enable customers to define fine-grained access policies based on user roles and permissions. This ensures that only authorized users have access to sensitive data, and that data is protected against unauthorized access or tampering.

BILLING AND CHARGEBACK

Zadara's Object Storage provides usage and billing reports both at the storage and the account level. The reports contain details regarding the capacity consumed and upload/download traffic. Usage reports can be exported as CSV files for further processing in external billing tools.

CONCLUSION

Object Storage complements Zadara's offerings for public and private clouds.

In addition to block and file storage services, Zadara's Object Storage solution provides extremely high performance, security, and resilience, making it an ideal choice for any user who needs to manage and store large amounts of unstructured data. Our solution leverages the latest hardware and software technologies to deliver industry-leading performance, while also providing robust security features and high-availability mechanisms to ensure data availability and resilience. Whether you need to store large amounts of unstructured data for archival purposes or support mission-critical applications that require high levels of performance and availability, in a location of your choice — Zadara's Object Storage solution is the perfect choice.