WHITE PAPER

QUISTOR BACKUP AS A SERVICE

QUISTOR

CONTENTS

INTRODUCTION

Performing a backup of your data should be one of the key activities of your IT department in order to be prepared for any disaster or equipment failures that might occur.

The public cloud, in our case – being an Oracle Partner – the Oracle Cloud, can provide you with a secondary (off-site) location to host a copy of your databases, files, and more. Since the first introduction of the public cloud, the cloud has matured significantly and reached such a level that placing a backup in the cloud feels like having a backup on your own premises.

As many of you may know there are endless possibilities on how you can perform a backup. In this paper, we are going to guide you through the Oracle Cloud technologies used for creating a highly secure backup solution for our customer.

WHY BACKUP TO THE CLOUD?

Choosing a cloud-based solution for your backups has major advantages compared to the traditional way of creating off-site backups by writing, shipping, and storing tapes at an off-site location. Below you can see the most important benefits:

Fast Provisioning & Scalable

Offsite cloud storage provisioned on-demand with on-demand capacity expansion capability.

Low Cost

No capital expenditure, low operating expenses, includes ASO/ACO licensing, simple pay per use model.

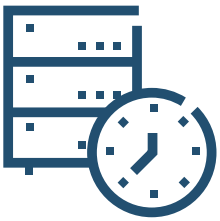

24x7 Accessibility to Offsite Storage

Data are accessible securely via the internet from anywhere.

Secure and Reliable

Data is secured with client-side encryption and transmitted securely to the cloud.

CUSTOMER'S PREVIOUS BACKUP STRATEGY

In the customer's previous setup, they had two backup procedures in place: a regular one and a non-regular one. The regular backup procedure contained a daily on-premise backup in location 1 and a monthly full backup via tape stored in their satellite location 2.

The non-regular backup procedure was applicable for those files which were of the essence for the customer's production environment. These processes were running via separate applications or scripts. This meant that maintaining the backup environment required a lot of management.

Another disadvantage was that the processes were not registered, and the various processes were only known to the responsible people.

In short, they had a backup infrastructure based on tapes and several separate scripts and their objective was to replace it for one uniform backup procedure.

CUSTOMER'S OBJECTIVES

The customer's objectives regarding the implementation of a cloud-based backup solution were the following:

1. Modernizing the current tape system to a future-proof solution;
2. Periodic backups that are stored outside the main office (Location 1) and can be started from (any) other location for business continuity;
3. Shorten the restoration time to guarantee business continuity in case of an emergency;
4. To be able to independently restore files from the backup to the production environment;
5. One central backup environment and procedure;
6. Supporting or replacing the existing archive environment for files matching the retention periods;
7. Setting up integrity checks on made backup files and being able to periodically perform disaster recovery tests.

Above objectives were mapped against a set of very strict rules and requirements. To give you an idea, we have highlighted the most relevant technical requirements. The requirements for the Managed Service Provider we disregarded, as we will solely focus on the technical requirements related to the solution implemented.

BACKUP SOLUTION REQUIREMENTS

The new backup solution must be able to go back up to thirty days to restore specific files, servers or entire environments. Furthermore, it should be able to perform weekly full backups with (minimum) daily incremental backups. The estimated data growth is 3% per year, whereas the storage capacity for backup and archive purposes can grow with at least 15% over 5 years without the need for hardware expansion. However, the storage capacity of the backup must be scalable and should provide the flexibility to add additional capacity (licenses or hardware).

Data must be stored and sent securely and according to the most recent encryption protocols. It should also support the iSCSI network protocol for connecting to storage. Further, the solution provided must contain a bandwidth of at least 1GB/s. In terms of information security, they had to meet the ISO 27001 guidelines.

The solution should be able to independently repair damaged data, even if this happened after the backup has already been written to the storage location.

The archive environment should be able to provide the flexibility to store data for as long this as preferred by the customer. This means that if the (maximum) retention period has been reached the solution should provide the flexibility, in case requested by the customer, to extend the archive retention period.

RECOVERY REQUIREMENTS

The solution offered must facilitate backup to an external location with at least a daily interval (RPO). Files must be retrievable up to thirty days due to the requirements of the business. The restoration times apply for a maximum of 4 hours for the production environments and most critical processes for continuation (RTO). All in all, this is about 30Tb. Based on the past, 2 restore activities per year were considered with 200GB in total over 4 years.

SOLUTION IMPLEMENTED

So, what exactly is the solution that was implemented?
Quistor chose a modern managed cloud-based backup solution that uses the Oracle Cloud Infrastructure (OCI). The Oracle Cloud is a quality platform that guarantees more than just availability. Oracle offers end-to-end SLAs related to performance, availability and manageability of the services.

Quistor was convinced that the combination of OCI with Quistor Managed Services fits in with the objectives and technical requirements of this customer. Running the backup servers in a scalable cloud environment makes the solution future-proof and allows it to easily scale to the expected (3% data) growth and unexpected (data) growth of this customer, without the need for physical hardware.

In addition, the solution should provide one uniform backup procedure where the data are stored outside the headquarters. In case of emergencies, it makes it easy to back up from any entity to ensure business continuity. The below architectural visual shows their newly Oracle Cloud-based Backup Solution.

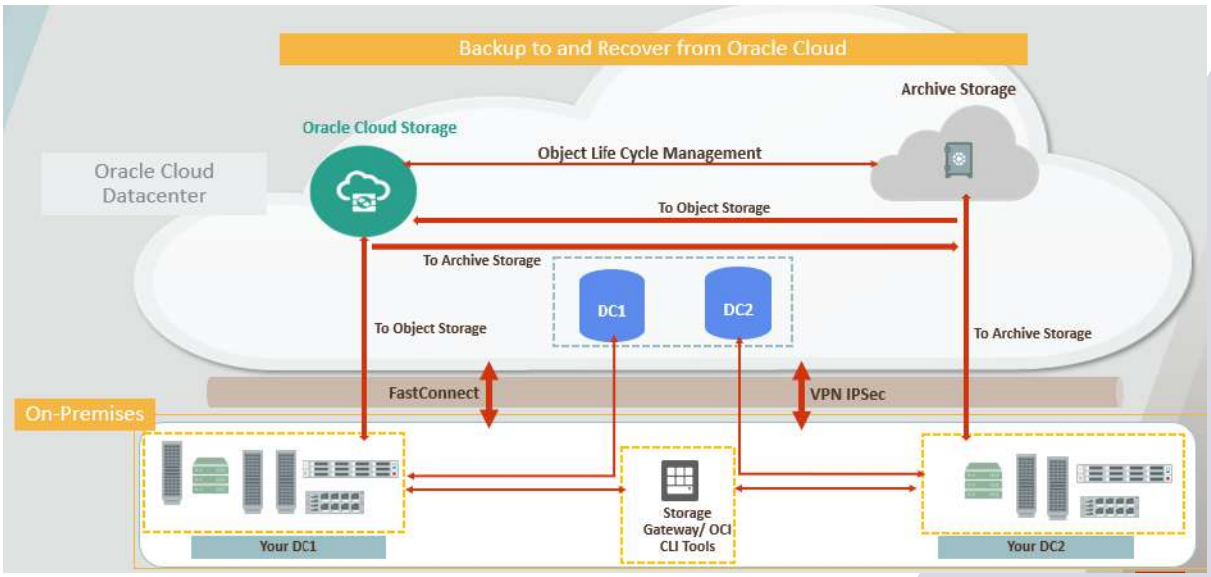

INTRODUCTION TO ORACLE CLOUD TECHNOLOGIES

What exactly are the Oracle Cloud technologies used to create this customer's backup solution?

Oracle Cloud Datacenter

Oracle deploys its cloud in data center regions. For this customer, we selected the Amsterdam region as the new external location for hosting the backup solution. Oracle Cloud Datacenters have availability domains and each availability domain contains at least 3 fault-tolerant domains. Fault-tolerant domains are clustered and mutually synchronized. Through the technical setup of the Oracle Cloud and by placing the storage of both locations in the Oracle Cloud we could provide a highly available solution that ensures business continuity in case of an emergency or failures.

OCI VM Shapes

A hardware shape is a standard configuration of CPUs, internal memory and other hardware components. The standard shapes provide the right balance between cores, memory and networks. All shapes come with an X7-based standard compute with Intel Xeon Platinum 8167M Processor, Base frequency 2.0 GHz and a maximum turbo frequency of 2.4 GHz. This backup solution requires three VM Standard 2.2 shapes, containing 2 OCPUs and 30Gb of internal memory. One Oracle OCPU is representing 2 Virtual Cores.

VM Standard 2.2 shapes → DC1
VM Standard 2.2 shapes → DC2
Oracle Storage Gateway

Block Storage

Oracle Cloud Block Storage can provide this customer with consistent high performance. This means 60 IOPS per GB, up to a maximum of 25,000 IOPS per volume, backed by Oracle's first in the industry performance SLA. The Oracle Cloud Block Storage service will be used for storing the backups of the customer that are accessed more frequently. Block storage is added to the above-described shapes with a total of 100Tb per month.

Boot Volumes

The images for the customer are stored in the Oracle Cloud Boot Volume in order to easily create a compute instance. This allows you to act fast in case of a disaster because a similar hardware configuration can be realized in a very short time.

Object and Archive Storage

The Oracle Cloud Infrastructure offers two distinct storage class tiers to address the need for both performants, frequently accessed "hot" storage, and less frequently accessed "cold" storage. By using the Storage tiers, we were able to maximize performance where appropriate and minimize costs for this customer.

We implemented the Object Storage tier for the data to which they needed fast, immediate, and frequent access. In this case, the Archive Storage tier was ideal for storing data that is accessed infrequently and requires a longer retention period. One of the customers' requirements was to have the flexibility of an unlimited retention period. For this reason, in the technical setup, we used Archive Storage to store the backups. It is also more cost-effective than Object Storage for preserving cold data.

The Data stored in Oracle Archive Storage can be accessed through the same interfaces as the Oracle Object Storage (API, SDK, CLI), but unlike Oracle Object Storage the retrieval of data is not immediately. This means that you have to restore the data from the archive storage to an Object Storage and start the backup restore process from there. The Oracle Cloud Block Storage service will be used for storing the backups of the customer that are accessed more frequently and instantaneously.

To download an object from Archive Storage, you must first restore the object. The restoration takes at most an hour from the time an Archive Storage restore request is made, to the time the first byte of data is retrieved. The retrieval time metric is measured as Time To First Byte (TTFB). How long the full restoration takes, depends of course on the size of the object. However, the customer's requirement to bring back their specific files, servers or even full environment within a 30 days timeframe could be easily be met with this backup solution.

Oracle Storage Gateway

Quistor installed the Oracle Storage Gateway on an Oracle Cloud Compute instance. The Storage Gateway is the bridge between the customer's on-premise data center and the Oracle Cloud and is an effective way to migrate data to the cloud. Also is the data easily accessible via the Gateway, enabling the customer to restore specific files, servers or entire environments when needed.

In principle, the gateway can be compared to a traditional Network Attached Storage (NAS) system when data is written to it. However, the Storage Gateway also moves the data in the back to OCI object storage and puts it in a bucket of the object storage cloud service. Once in an object storage bucket, the data lives as objects that can be viewed, managed or used from the OCI or through the gateway interface.

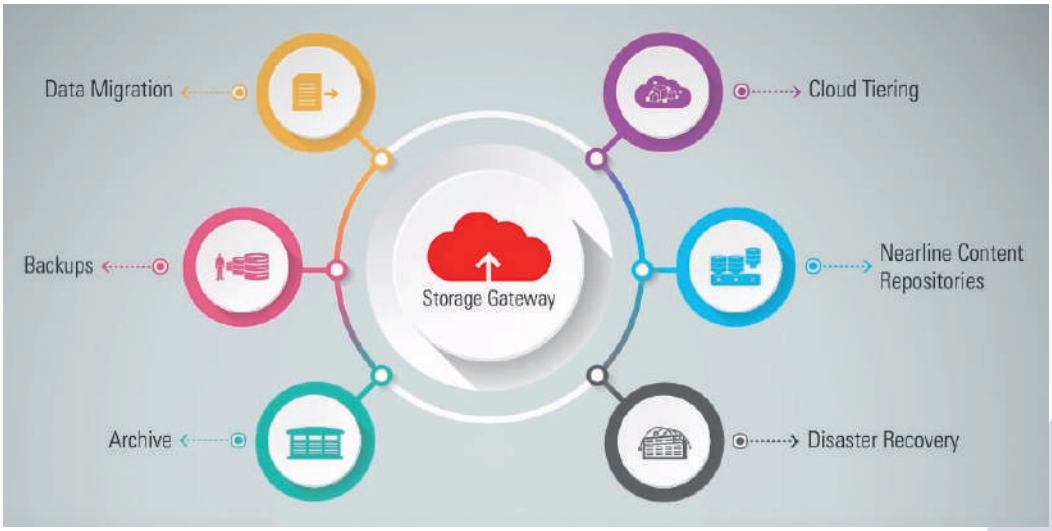

Applications store and retrieve objects from OCI Object Storage through file systems that you create in Storage Gateway. Storage Gateway exposes an NFS mount point that can be mounted to any host that supports an NFSv4 client. The Storage Gateway mount point maps to an Object Storage bucket.

There is a file to object transparency between Storage Gateway and Object Storage:

- A Storage Gateway file system directory on a local host of the customer maps to a bucket with an identical name in Oracle Cloud Infrastructure Object Storage for this customer.
- Any file written to a Storage Gateway file system is written as an object with the same name in the associated Object Storage bucket. The system of this customer will then store associated file attributes as object metadata.
- The customer can access Object Storage objects directly using native APIs, SDKs, third-party tools, the HDFS connector, and the Oracle Cloud Infrastructure CLI and Console. You use the Refresh operation in Storage Gateway to ingest any data that was added or modified directly in Object Storage.

File Transparency from Oracle Storage Gateway is one of the parts that guarantees business continuity on one site when one of the sites is impacted by a disaster or the other way around.

STORAGE GATEWAY AND OCI CONCEPTS

The following concepts were essential to working with OCI Storage Gateway.

File system

Quistor will create a Storage Gateway file system on localhost to map its files and directories to objects with the same names in a corresponding OCI Object Storage bucket.

File system cache

Storage Gateway's configurable file system cache enables the asynchronous and optimized movement of data to the cloud. The file system cache serves as both a write buffer and a read cache for data storage and retrieval. The write buffer contains data that was copied to the disk cache and queued for upload to OCI. The read cache contains frequently retrieved data that are accessible locally for reading operations. To meet the RTO/RPO requirements we have pinned files in the cache for quick access. You can pin files to the cache for file systems connected to either the Object Storage Standard or Archive tier.

Metadata

The metadata associated with a Storage Gateway file is stored as custom metadata for the corresponding object in OCI Object Storage. Examples of file metadata include object id, creation date, modification date, size, and permissions. Storage Gateway caches all metadata for the file system locally.

NFSV4

NFS is an established and widely adopted distributed file system protocol for handling network storage. NFS lets client computers mount file systems on remote servers and access those remote file systems over the network as though they were local file systems. Storage Gateway performs the NFS to REST API translation needed to interact with OCI Object Storage.

Data Integrity

The built-in data integrity checks ensure that data is validated as it moves through the data path from Storage Gateway to Oracle Cloud Infrastructure Object Storage. Checksum verification helps ensure data integrity. Metadata integrity checks ensure that the metadata is in a consistent state. The checksum for each file can be read using a custom interface.

Connectivity

In this case, the connection between the local infrastructure and the cloud was created via IPSec because their current connectivity provider is not partnered up with Oracle. IPSec VPN establishes an encrypted network connection over the internet between your network or data center and your Oracle Cloud Infrastructure Virtual Cloud Network (VCN). It's a suitable solution if you have low or modest bandwidth requirements and can tolerate the inherent variability in internet-based connections. Oracle FastConnect bypasses the internet. Instead, it uses dedicated, private network connections between your network or data center and your VCN.

Encryption

By default, all data uploaded to the Storage Gateway is encrypted at rest, using the AES 256 encryption algorithm. Data is uploaded and downloaded from the cloud over secure SSL endpoints using the HTTPS protocol.

Within the Oracle Storage Gateway, a filesystem was created in which the backup files are stored and where they can be retrieved from if necessary.

CONCLUSION

In the past decade, cloud providers have overwhelmed you with all the benefits you can obtain when moving workloads to the cloud. As mentioned at the beginning of this white paper, there are endless possibilities to create a backup solution. However, we strongly believe that the public cloud can offer you the most sustainable and most affordable backup possibility.

Today's public cloud offerings have reached such a high level when it comes to performance, security measures, availability, flexibility, competing costs, and more that in whatever upcoming choice you must make as an IT department you should consider the cloud as your next step.
In this case, the customer chose a modern managed cloud-based backup solution that uses the Oracle Cloud Infrastructure. As a result, they have an extremely secure, high-performance backup solution that is scalable, future proof and customized to their requirements and objectives.

HOW QUISTOR CAN HELP

If you find this solution interesting and you want to find out more, don't hesitate to get in touch with Quistor cloud experts via info@quistor.com.

WHAT IS NEXT?

- Initial meeting
- Understand your situation and preferences
- Technical investigation
- Analysis of your current Backup Infrastructure setup
- Development of a Detailed plan including:
- Strategic advice on Cloud solutions
- Detailed plan for Implementation
- A Cloud Architecture (Visual)
- Investment overview / TCO

QUISTOR