storsimple

Cloud-integrated Storage – What & Why

Table of Contents

Overview

This paper is an overview of a new class of storage technology called cloud-integrated storage, or CiS. CiS combines a number of storage technologies, including iSCSI SAN, snapshot, backup, deduplication and compression with storage services offered by cloud service providers. The CiS architecture brings new management capabilities that organizations may want to consider for their own IT implementations. Figure 1 below illustrates the basic topology of CiS.

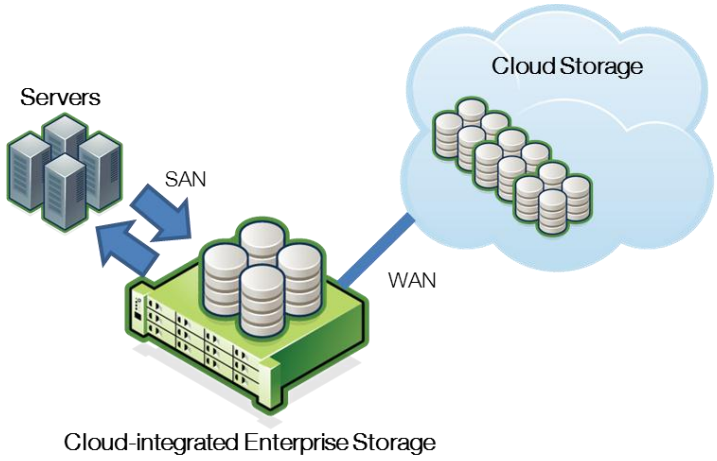

Figure 1: Basic topology of CiS

CiS Architecture

CiS products are built from three fundamental components (see Figure 2)

- SAN storage
- Cloud data management
- Enterprise-class storage platform

Figure 2: Components of CiS

Enterprise-class Storage Platform

Enterprise ready, highly-available storage is protected from a single point of failure with redundant protection for all components, including dual controllers, battery-backed memory, multiple ports and RAID. In addition, hot swappable components and non-disruptive software upgrades mean hardware and software can be upgraded without downtime.

Operating system and platform certifications from Microsoft Windows Server 2008 and VMware are integral parts of what constitutes an enterprise storage platform today. In addition, StorSimple has been extensively tested and deployed with cloud storage providers such as Windows Azure.

Application-consistent snapshots are integrated with CiS systems and use an application's storage APIs to temporarily quiesce operations and get clean point-in-time versions of data. That means data can be quickly and accurately restored without errors or administrators spending additional time resolving data inconsistencies.

Enterprise Tier 2 SAN Storage

CiS uses the industry-standard iSCSI SAN protocol to connect to servers. iSCSI is easily configured for use with both Microsoft and VMware servers and is widely understood by storage administrators.

CiS is intended to run as primary storage for enterprise tier 2 applications, including email, file shares, Microsoft Sharepoint, content management systems, virtual machines/VM sprawl and large unstructured data repositories. It is not built for latency-sensitive applications such as online transaction processing. A good way to visualize tier 2 data is that is follows a lifecycle curve similar to that shown in Figure 3 where data is accessed most heavily in the days immediately after it is created and then less frequently thereafter.

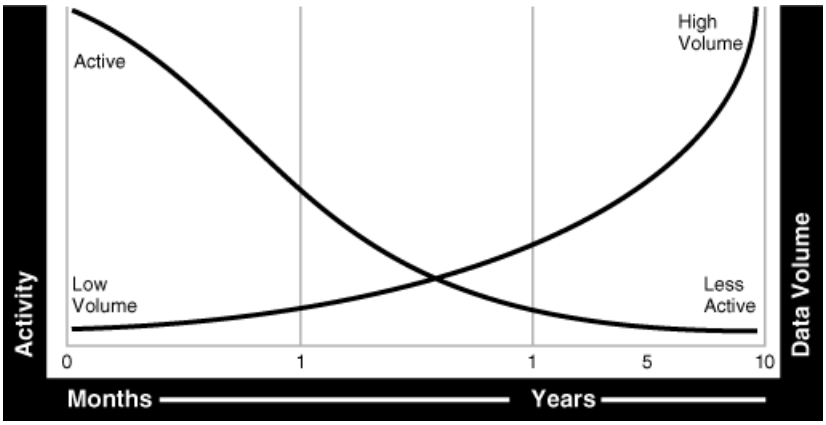

Figure 3 Data Lifecycles. source: Oracle Corporation

Activity-based storage tiering and data ranking

CiS systems use three different types of storage: performance-oriented flash SSDs, capacity-oriented SAS disk drives and cloud storage. Data is moved from one type of storage to another according to its relative activity level and customer-chosen policies. Data that becomes more active is moved to a faster type of storage and data that becomes less active is moved to a higher capacity type of storage.

There are four logical tiers in the system, two at the SSD level and one each in the SAS and cloud storage levels. The four tiers are shown below in Table 1.

Tier name	Storage type	Data activity	Reduction applied
Native	**SSD**	**New, most active**	**None**
Hot	**SSD**	**Existing, most active**	**Deduplication**
Warm	**SAS**	**Between hot and cool**	**Full**
Cool	**Cloud**	**Least active**	**Full**

Table 1: Logical tiers in a StorSimple CiS system

The third column in Table 1, data activity, indicates the relative activity level for the various tiers. The native and hot tiers have the most active data, with the distinction being that data in the native tier has been recently added to the CiS system and data in the hot tier was has been in the system longer. The cool tier has the least active data and the warm tier has data that is neither the most or least active.

Data ranking is an automated process that combines the activity level of data with its age and applicable policies to determine what tier data is placed on.

Data reduction and thin provisioning

The fourth column in Table 1 indicates the type of data reduction technology used in the various tiers. The native tier has none, the hot tier uses *deduplication* (or dedupe) and the warm and cool tiers use full reduction, which means data is compressed as well as deduped. Notice that the progression from native to warm tiers implies that data is first deduped before it is compressed.

Dedupe reduces the amount of data stored in the system by identifying data duplicates and removing excess copies. Dedupe is particularly effective in virtual server environments. Compression reduces the amount of data stored in the system by identifying strings of repeated data values and replacing them with encoded shorthand

Another capacity-conserving technology in CiS systems is *thin provisioning*, which allocates storage capacity as it is needed, as opposed to reserving capacity in advance. All storage in a CiS system is thinly provisioned.

Cloud Data Management

CiS provides a broad set of data management tools that enable customers to use cloud storage in ways that are familiar to them, including archive and backup storage. Many of the data management tools in CiS systems are automated to reduce the workload on administrators.

Cloud snapshots are point-in-time copies of data that are stored on cool tiers in the cloud. All cloud snapshots are fully reduced (deduped and compressed) to minimize the amount of storage consumed.

Cloud data reduction and cloud storage WAN optimization refers to the fact that data transferred and stored in the cloud by a CiS system has already been fully reduced. This minimizes the cost of cloud storage as well as the transaction costs and WAN bandwidth associated with storing data in the cloud.

Cloud as a Tier (CaaT) refers to the automated use of cloud storage as the cool tier in a CiS system. Data that is ranked lowest is sent to a cool tier in the cloud where it remains until it is accessed again and promoted back to the warm tier.

Cloud thin provisioning is thin provisioning applied to the cloud storage used by a CiS system. Cloud thin provisioning allocates cloud storage automatically as it is needed instead of reserving cloud capacity in advance, ensuring CiS customers do not pay more for storage than necessary.

CiS systems provide **volume-level cloud mapping** between storage volumes on CiS systems and the location in the cloud – or clouds. Different volumes can have cool tiers on the same or different cloud services. Every CiS system keeps a **metadata map** that describes the state of the system and provides an image of the volume's contents at the time a snapshot is taken. This map is typically .1% the size of the stored data.

AES-256 encryption is applied to all data transmitted and stored in the cloud by CiS systems to ensure its security.

SHA-256 hashing is applied to all data transmitted and stored in the cloud as a means to guarantee data integrity.

Cloud clones are the equivalent of a synthetic full backup that have all the current data for a volume at the time of the last snapshot. They are stored in the cool tier for use in disaster recovery scenarios, but they occupy separate repositories from cloud snapshots and can reside within the same or a different cloud service as the volume's cloud snapshots. Figure 4 shows that cloud clones are located in different repositories from cloud snapshots and that they can use the same or a different cloud service.

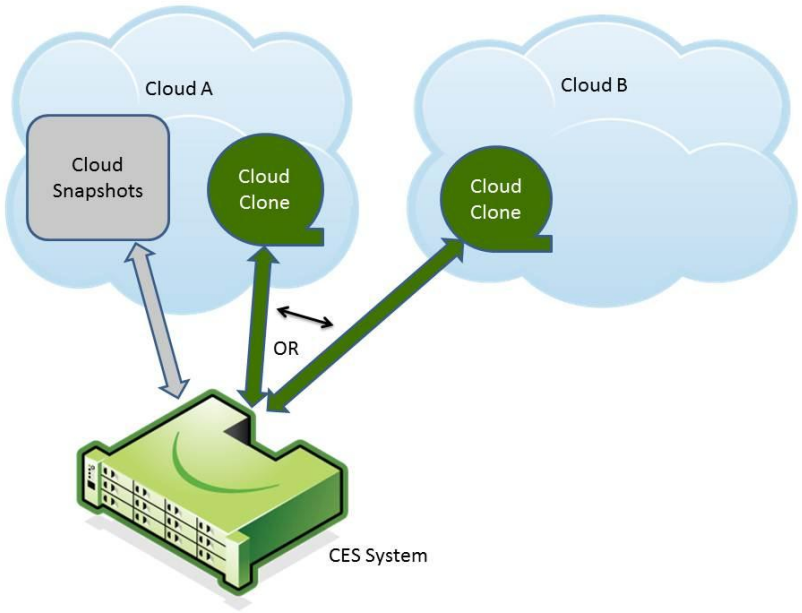

Figure 4: Location of cloud snapshots and cloud clones

A **thin restore** is a disaster recovery process where a CiS system downloads data from the cloud. The first thing downloaded is the metadata map and then users and applications can start accessing their working sets and download them. As data is downloaded it is ranked and placed in either the appropriate tier. The thin restore process is shown in Figure 5.

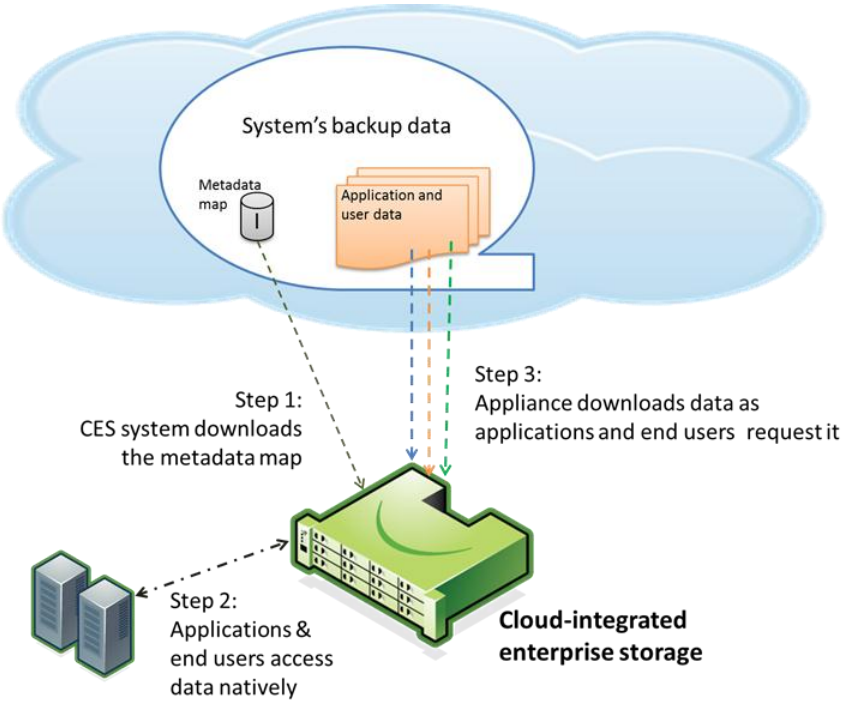

Figure 5: The thin restore process

Thin restores tend to have extremely short RTOs (Recovery Time Objectives); because systems can begin accessing data after the metadata map is downloaded. Thin restores do not restore cool data that does not belong to any working sets.

Figure 6 below contrasts RTOs for three different backup scenarios. The top line in red represents cloud storage used for virtual tape. The problem with this approach is that all the data has to be restored back through the WAN where network performance is virtually prohibitive. The middle line in black is traditional tape stored in an off-site facility where the gating factor is the time it takes to retrieve all the tapes needed from an off-site vault. The green line on the bottom is CiS cloud snapshots or clones and where applications and users can be accessing after downloading a relatively small amount of data.

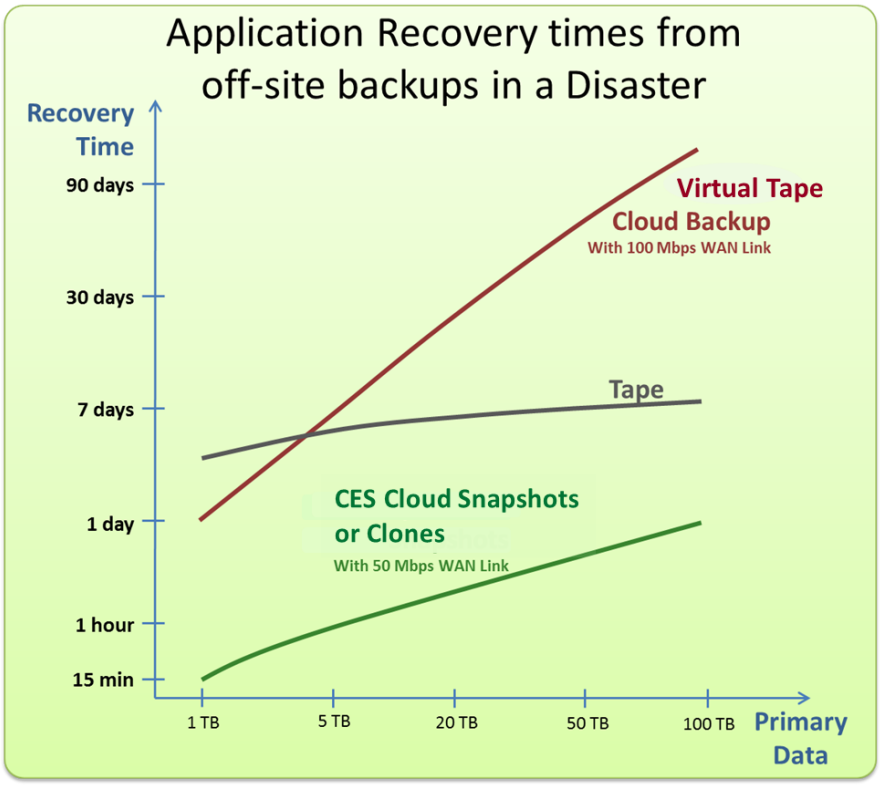

Figure 6: Comparison of RTOs for cloud-based virtual tape, traditional tape and CiS

Location-independent recovery refers to the ability to perform thin restores from any location with a suitable Internet connection. This differs from legacy disaster recovery operations that are restricted to running at specific recovery sites. Location independence adds an additional level of redundancy to the recovery process and does not require the capital investment that traditional replication solutions do. A customer with multiple data center locations can use CiS systems running in any of those locations to recover from disasters in any of the other sites. Similarly, a single CiS system can act as a spare for any of the others, providing an extremely cost-effective DR implementation.

The Big Picture

Figure 8 below shows a composite of all the various elements in CiS systems and its three major pillars: enterprise SAN storage, cloud data management and a highly-available enterprise-class platform.

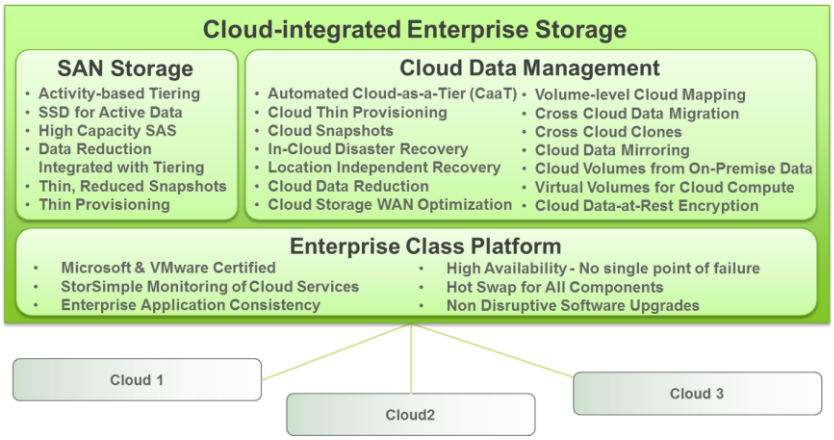

Figure 8: The big picture of cloud-integrated enterprise storage

Summary

CiS is a new type of storage system that integrates SAN storage and cloud data management on a highly-available platform. CiS was developed to meet the requirements of enterprise tier 2 applications where there is a relatively small working set requiring SAN performance and a relatively large amount of data that is dormant, or cool, that is not accessed often, if ever.

A rich set of storage management features automate some of the most time-consuming manual processes in traditional storage environments. CiS systems streamline data protection processes such as backup and DR, significantly reducing the cost and time involved with both.