IBM Software
April 2010

Effective Storage Management for Cloud Computing

April 2010

Contents:

EFFECTIVE STORAGE MANAGEMENT FOR CLOUD COMPUTING

The U.S. National Institute for Standards and Technology (NIST) has defined Cloud Computing as "*a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction.*"

Cloud-based systems have brought a new, scalable application delivery service model to the market. They can help clients save both capital and operational costs, and further shift capital expenses (CAPEX) to operational expenses (OPEX) while providing increased flexibility for their organizations. However, a key aspect of any type of Cloud deployment is reliable service delivery, on which ready access to the digital information (data) in the Cloud is foundational requirement.

In Cloud Computing environments, applications reside in massively-scalable data centers where compute resources can be dynamically provisioned and shared to achieve significant economies of scale. Storage capacity scales with compute resources, and each needs to be effectively managed to gain the maximum benefits from Cloud.

A 'pay-as-you-go' usage model enables users & companies to better predict and manage expenses, reduce costs, and simplify operations. Also, the cost of adding new services & users should approach $0.

CLOUD CONFIGURATIONS

There are several models of Cloud Computing currently in the market, but they break down into a combination of the services being provided, and the manner in which they are hosted. For example, a Cloud may provide infrastructure as a service

(servers, storage, networking), or software platforms as a service (desktop and server applications, portals, etc.).

The methods of hosting the Cloud can be 'private', where the company owns and manages the Cloud infrastructure, or 'public', where a third-party owns and manages it, or a 'hybrid' model combining both, where the Public Cloud may act as an overflow facility for the Private Cloud or is used to satisfy other application needs.

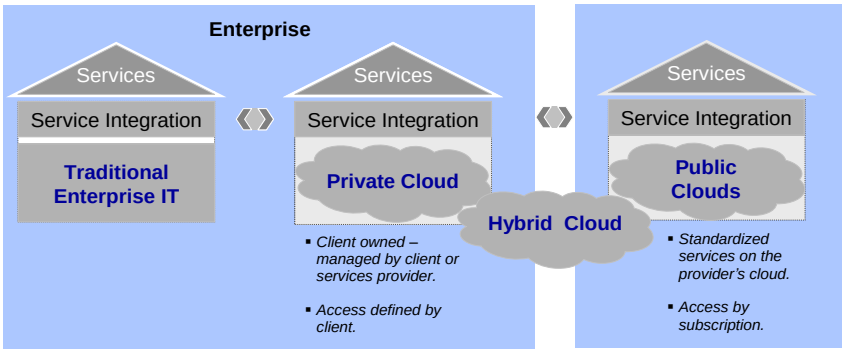

Highlights:

Options for deploying Cloud Computing:
- Build your own Private Cloud
- Subscribe to a Public Cloud
- Do both (Hybrid Cloud)

Data and storage management in the Cloud are critical:
- To provide a reliable, on-demand service experience
- To reduce costs
- To mitigate risks

The keys to effective Cloud storage management include:
- Data protection and recovery
- Data lifecycle management
- Storage utilization and optimization
- Storage resource management

Because all of the data in a Cloud resides in the same shared systems, the effective management of the storage resources and data in the Cloud are of paramount importance in maintaining satisfactory service levels. Cloud services rely heavily on keeping the data and applications they are managing available at all times, and to restore operations quickly following any type of data disaster (database corruption, virus attack, hardware failure, local/regional disaster). Downtime, or failure to provide access to needed storage and data resources, constitutes a failure of the Cloud service, and this should obviously be avoided.

At the same time, effective utilization and optimization of the storage resources in the Cloud infrastructure, and the proper placement of data on different tiers of storage within the Cloud, will help to minimize the overall costs of hardware, software and administration.

IBM Storage Management solutions for Cloud Computing

Data Protection and Recovery

Tivoli Storage Manager provides unified data protection and recovery management across the entire infrastructure, including the Cloud, to provide reliable and predictable access to data. Tivoli Storage Manager (TSM) includes automated data lifecycle management, built-in data reduction technologies, advanced application protection and more. Data managed by Cloud providers must be protected appropriately, according to its value. For business data, that typically means managing multiple copies and multiple versions of every file. Tivoli Storage Manager provides flexible levels of data protection, and can meet a variety of Recovery Point Objectives and Recovery Time Objectives, so multiple levels of service can be offered. More importantly, TSM can be configured to segregate data by client. Cloud providers can ensure their clients that on-site and off-site backup media will not get into the hands of other users, even in a Disaster Recovery scenario.

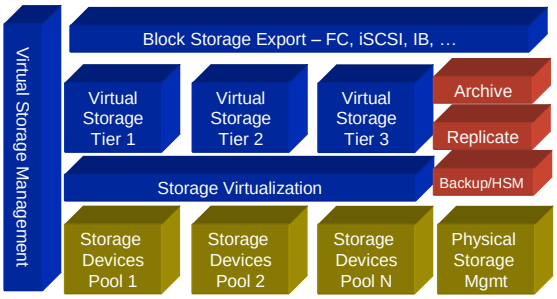

Storage Cluster Model for Cloud

Tivoli Storage Manager FastBack is an advanced continuous data protection and near-instant recovery software solution for business-critical Windows and Linux servers. Tivoli Storage Manager FastBack helps clients reduce the amount of data at risk between backups to almost zero, and reduces the time to recover from almost any data loss to just seconds. Within Cloud environments, FastBack can be deployed to provide enhanced service level protection for the most critical applications, whether on physical or virtual servers.

Tivoli Storage FlashCopy Manager enables fast & frequent backup of critical applications to limit data loss without disruption to operations. It leverages advanced snapshot technologies in IBM System Storage. FlashCopy Manager can work as a standalone solution to provide backup and restore as well as integrate with Tivoli Storage Manager to provide longer-term data management.

DATA LIFECYCLE MANAGEMENT

An important consideration in planning and managing storage in Cloud environments is to make the most efficient use of storage resources by placing data on the most appropriate and cost-effective tier of storage that meets service delivery requirements, and to eliminate data that is no longer needed. The process of automatically moving and deleting data based on business rules can be called data lifecycle management.

IBM Tivoli Storage Productivity Center for Data is a sophisticated data reporting tool that can help set policies and identify data that can be safely removed from the Cloud. With TPC for Data, you can get reports on where your data is, sorted by access or saved dates, by who owns it, by the application that created it, and numerous other filters. From the intelligence you gain from the TPC for Data reports, you can set meaningful policies in your data management software to automatically take the appropriate action on data that shouldn't be clogging up your primary systems. It can also help you to identify and eliminate duplicate data, orphan data, temporary data and non-business data for deletion.

The next step is to automate these data management policies through three distinct processes: migration, archival, and expiration.

The **IBM Tivoli Storage Manager** (TSM) family includes the capabilities that automate the migration of data between multiple tiers of storage based on policies. The net result of these processes is to remove unneeded data from your primary Cloud storage systems, which will reduce or delay your need to acquire more expensive hardware and reduce administrative costs, all without impacting key operational processes.

Tivoli Storage Manager data migration solutions not only help you clean up your primary storage systems to help them run

more efficiently, they can also be used to easily move data to new storage technologies as they are deployed within the Cloud.

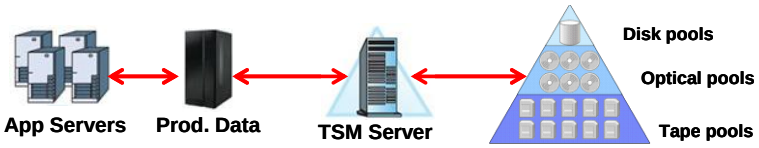

These solutions work transparently in the background, automatically selecting and moving files from primary to secondary tiers of storage, or back again, based on the policy criteria that you set, such as file size or length of time since a file has been opened. And it is all transparent to the users and applications.

Data migration within the Cloud will help you get control of, and efficiently manage, data growth and its associated storage costs by providing automated space management, providing the following key features:

- Storage pool "virtualization" helps maximize utilization of the managed storage resources.
- Restore management is optimized based on the location of the data.
- Migration is transparent to the users and to applications.
- Migrations are scheduled to minimize network traffic during peak hours.
- Automatic migrations occur outside the backup window.
- By setting proper threshold limits, annoying 'out of disk space' messages can be eliminated.

Storage Utilization and Optimization

IBM software solutions to manage storage in Cloud environments are designed for delivering optimized data availability and performance, and improved utilization of storage assets. They can help clients to reduce efforts for managing their storage infrastructure stack in a Cloud by addressing the streamlining of resources to support services, protection & management of data, being able to virtualize the entire storage infrastructure and offer it as a single resource to the Cloud.

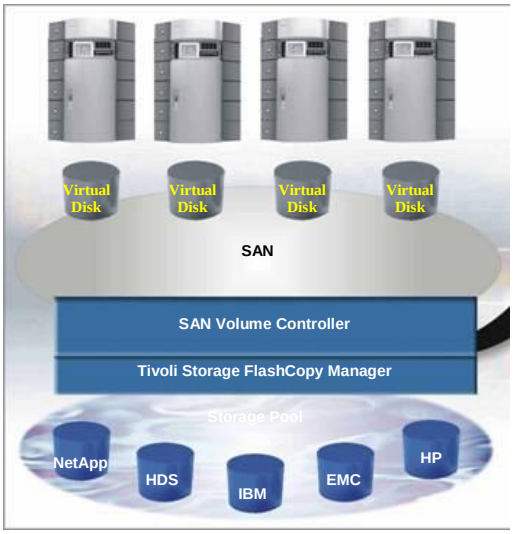

SAN Volume Controller combines heterogeneous storage devices into a single pool of capacity, and manages the placement of data within the virtual storage pool based on business policies

IBM System Storage SAN Volume Controller virtualizes IBM and non-IBM storage to enable resource pooling, thin provisioning and simplified management. SAN Volume Controller (SVC) eliminates 'silos' of storage by enabling the sharing of storage resources among applications. It auto-assigns and releases capacity on demand, addressing a critical requirement of all Clouds – the redeployment of resources as utilization needs change.

IBM Scale Out Network Attached Storage (SoNAS) is a multi-petabyte scale out NAS storage offering for information storage and distribution. It is built on IBM's General Parallel File System (GPFS) to store and manage billions of files in a single namespace – something that is critical to larger Compute and Storage Clouds.

It can meet today's storage challenges with quick and cost-effective IT-enabled business enhancements, combining massive scalability, high availability and automated data management to efficiently deliver information around the world.

STORAGE RESOURCE MANAGEMENT

Often, the absence of sophisticated storage management systems in a Cloud results in a lack of visibility into the storage environment. Cloud administrators have difficulty in understanding how much capacity is available, where it is, which applications are accessing it, how secure they are, and whether they are able to meet stringent Recovery Time Objectives. Gaps that may exist between the unpredictable demand for data availability and the ability of business to meet that demand in an efficient way is resulting in unmet service levels, additional downtime, new hardware & operational costs, and lower customer satisfaction.

IBM Tivoli Storage Productivity Center (TPC) is an advanced storage resource management solution that manages virtualized storage environments with tools that can help customers reduce the complexity of storage environments by centralizing, simplifying and automating storage tasks associated with Cloud infrastructures. This integrated solution can help improve storage availability within Cloud environments by combining the assets, capacity, performance, analytics, configuration, operations and data replication from a single management platform. Designed to manage both IBM & non-IBM storage systems, TPC is easy to deploy & manage and its components are tightly integrated with Tivoli service management software to deliver automated management of data and resources.

IBM Cloud Storage Service Offerings

IBM has developed the worldwide capability and capacity to provide integrated Cloud service offerings to meet your storage management needs:

Smart Business Storage Cloud Services
As data volumes grow and the ability to handle various file formats becomes more complex, supporting efficient and cost-effective access to data can be increasingly difficult, with users experiencing reduced performance and outages.

IBM Smart Business Storage Cloud service can help you successfully deploy a high-performance, scalable storage-virtualization solution to facilitate growth and innovation at lower operational costs. Storage Clouds can be deployed at your site, at an IBM delivery center or both.

Onsite Data Protection services provide optimized protection for data located in customer data centers, branch offices or hosting centers through a proven, professionally managed backup and recovery service.

Remote Data Protection services provide automated disk-based backup and recovery for branch and remote sites (PCs and servers) and are for companies that have distributed locations with limited IT services.

Email Management Express provides a common service platform for email continuity, email archiving, emergency notification services and emergency collaboration capabilities.

The new **fastprotect online** continuous data backup service for desktop and laptop computers provides security-rich encryption and helps recover lost data virtually anywhere, anytime, without help desk or IT involvement.

SUMMARY

IBM software and service solutions for smarter storage management in Cloud environments help to ensure that business and IT are fully aligned and supported by integrated service management. They help deliver a workload optimized approach, and offer a choice of implementation options for superior service delivery with agility and speed.

IBM now offers second generation storage management technology for Cloud environments, delivering faster ROI for a smarter planet. The first generation technology was developed several years ago, deployed internally within IBM and used successfully by over 300,000 IBM employees worldwide. The best practices and lessons learned form the base for our second generation software technology, built to support the needs of a smarter planet – more scalable, more reliable, more efficient and more flexible.

Unlike other Cloud Computing solutions in the market using hardware & software combinations that are not tested for scalable results, IBM solutions for storage management in Cloud environments are tested and proven reliable, and also backed by a strong commitment to further research and development of IBM Cloud technology. Across the entire portfolio, IBM storage software products have several things in common – they are fueled by business expertise; built to scale, integrate, and adapt to a customer's changing needs; and have many years of proven reliability in real-world environments.

IBM reduces the complexity of managing Cloud environments by offering a complete portfolio of automated solutions for managing data and storage infrastructure, enabling better efficiency for business resiliency, reducing costs and improving security, while increasing visibility, control and automation of the Cloud storage infrastructure.

FOR MORE INFORMATION

IBM is uniquely positioned to provide a wide range of storage management solutions, even for mixed-vendor environments. We can help you implement a dynamic storage infrastructure that scales to meet your changing business requirements.
To learn more about how IBM Tivoli storage management solutions can help you address data management concerns in your environment, contact your IBM sales representative or IBM Business Partner, or visit ibm.com/software/tivoli/solutions/storage. Additional information can be found at ibm.com/software/tivoli/solutions/reduction

IBM Corporation Software Group
Route 100
Somers, NY 10589
U.S.A.

Produced in the United States of America
April 2010