



**WHITE PAPER: Cloud Computing**

# Nine Storage-Related Attributes of an Enterprise Cloud

Todd Loeppke  
Vice President, Storage Architecture, Office of the CTO, Savvis

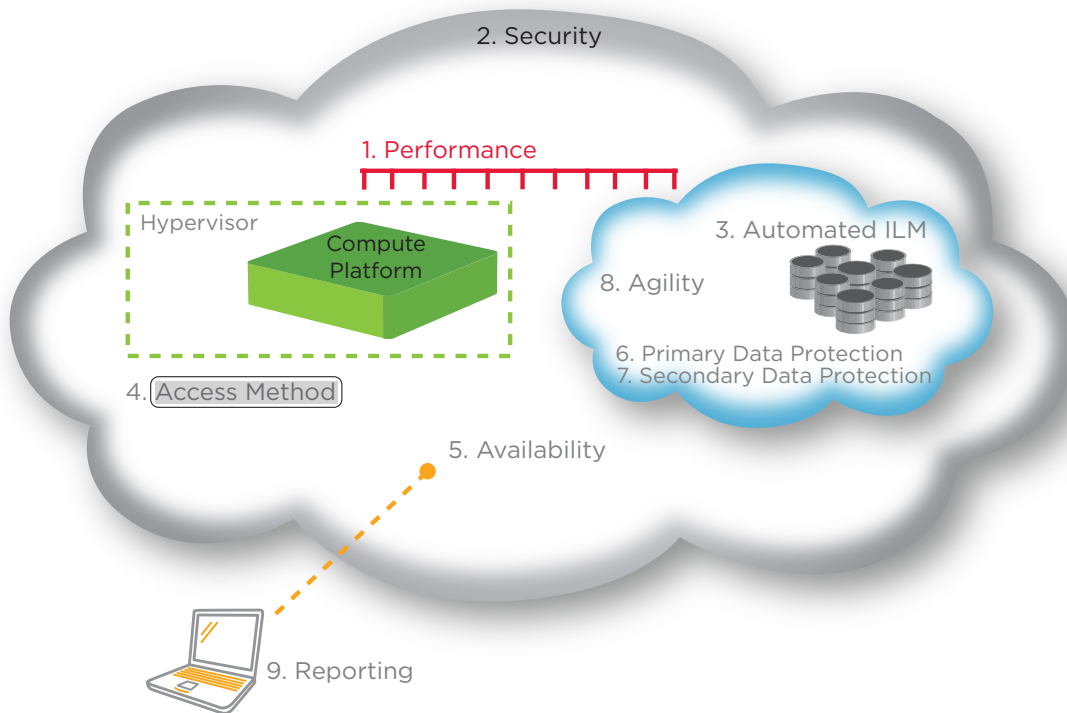
## **Table of Contents**

2	Introduction	9	Secondary Data Protection
4	Performance	9	Storage Agility
5	Security	10	Storage Reporting
6	Automated ILM Storage	11	Conclusion
7	Storage Access Method	11	About the Author
8	Availability		
8	Primary Data Protection		

## Introduction

Amid all the recent attention being paid to cloud computing, storage has largely been relegated to the background. Most cloud offerings today are narrowly viewed as a collection of CPU cores, a fixed memory allocation, low-tier spinning storage, and possibly a rudimentary network with a block of Internet-facing IPs. Still, there have been some interesting technology advancements related to cloud computing and storage recently — specifically access methods using Web services — so that accessing storage is no longer confined to device files and Network File System (NFS) mount points.

Typical “enterprise features” of data storage and management evolved over generations of IT architecture innovation. Storage architects recognize these features as essential to business-critical and production applications, but they have largely been absent from the cloud. The purpose of this white paper is to describe nine of these critical storage-related attributes as they relate to an Enterprise Cloud.



## Defining Terms

Upon reviewing cloud computing blogs, white papers, and press releases, it becomes clear that the term cloud computing is used to describe a variety of architectural concepts and service offerings. This paper focuses on infrastructure cloud services such as servers, storage, and security components with the goal of differentiating cloud services that are suitable for business-critical applications from those that aren't. To that end, the following terms will be used:

- **Enterprise:** An enterprise is a company with hundreds of millions of dollars in revenue that deploys many types of computing environments. These include everything from the most basic build-and-teardown test environments to processing-intensive, customer-facing, and revenue-generating production environments. The more mission-critical an enterprise's application are, the greater the need for increased levels of security, performance, and availability. This paper focuses specifically on the data storage and management requirements of such mission-critical enterprise applications.
- **Mass market cloud:** In the context of this paper, this term refers to the storage features that usually accompany the cloud computing solutions on the market today.
- **Enterprise cloud:** A pool of abstracted, highly scalable and managed compute, network, storage, and security infrastructure packaged to deliver flexible service level agreements (SLAs) and consumption-based billing, and capable of hosting the spectrum of applications required by an enterprise.

The following discussion of nine cloud storage attributes includes general comments about each attribute, how each relates to mass market clouds, and how they pertain to an enterprise cloud.

## Attribute No. 1: Performance

Performance costs money. In a well-architected application, performance and cost are balanced. The key to achieving this is to match an enterprise's business performance requirements to the right technologies, which in turn requires that the enterprise translate its requirements from business language to IT metrics. Since this translation is difficult, enterprises often end up with static IT architectures that cannot meet the changing performance requirements of the business. Enterprise cloud computing provides a platform that is better suited to react to changes in performance requirements.

Storage I/O in early cloud platforms typically possessed relatively high latency. That's because vendors have focused more on making the data in the cloud readily accessible rather than improving SLAs related to performance, bandwidth guarantees, or I/Os per second (IOPs). And there are two main reasons that latency remains relatively high: the type of access method and the type and configuration of the storage media being deployed.

The access method consists of a combination of multiple layers of protocols (e.g., SOAP, NFS, TCP, IP, and FCP) over a physical layer of the OSI Model. Data access that includes a shared physical layer (like Ethernet) and several layers of protocols (like SOAP or NFS) generally introduces more latency than a dedicated physical layer (like Fibre Channel) running FCP. Most mass market clouds also include the Internet in the data access, which contributes to data access latency.

For storage media, most mass market clouds use SATA drives in a RAID or JBOD (just a bunch of disks) configuration. Since SATA drive (sometimes referred to as near-line drive) performance is generally inferior to that of an enterprise drive (usually a solid-state, SAS 2.0 or Fibre Channel drive), the resulting storage device performance falls below application requirements.

When you combine a relatively low-bandwidth/high-latency access method with low-performance storage media, enterprises get an overall storage subsystem that fails to support many of their critical business application requirements. The resulting solution is often suitable only for test and development purposes.

In contrast, enterprise clouds need to provide many different performance tier options. As performance requirements change — for example, as an application is moved from development into production — the platform should be able to accommodate that change. Ideally, storage within an enterprise cloud should have multiple **performance knobs** that can be adjusted so the business performance requirements can be matched to the appropriate level of I/O performance.

Ultimately, to satisfy enterprises' high-end storage performance requirements, cloud solutions must use platform technologies that perform at or above the level of today's "enterprise technologies". Fibre channel Storage Area Networks (SANs) are generally used for this purpose. In addition, how the technology is implemented is just as important as the technology itself. In a hypervisor environment, virtual machines with enterprise requirements must be configured to perform at consistently high levels.

## Attribute No. 2: Security

Security and virtualization are often viewed as opposing forces. After all, virtualization frees applications from physical hardware and network boundaries. Security, on the other hand, is all about establishing boundaries. Enterprises need to consider security during the initial architecture design of a virtualized environment.

Data security in the mass market cloud, whether multi-tenant or private, is often based on trust. That trust is usually in the hypervisor. As multiple virtual machines share physical logical unit numbers (LUNs), CPUs, and memory, it is up to the hypervisor to ensure data is not corrupted or accessed by the wrong virtual machine. This is the same fundamental challenge that clustered server environments have faced for years. Any physical server that might need to take over processing needs to have access to the data/application/operating system. This type of configuration can be further complicated because of recent advances in backup technologies and processes. For example, LUNs might also need to be presented to a common backup server for off-host backups.

Businesses need to secure data in the enterprise cloud in three ways. The first involves securing the hypervisor. The primary goal: To minimize the possibility of the hypervisor being exploited, and to prevent any one virtual machine from negatively impacting any other virtual machine.<sup>1</sup> Enterprises also need to secure any other server that may have access to LUNs, like an off-host backup server. The other area that needs to be addressed is the data path. Enterprises need to pay attention to providing access paths to only the physical servers that must have access to maintain the desired functionality. This can be accomplished through the use of zoning via SAN N-port ID virtualization (NPIV), LUN masking, access lists, and permission configurations. Last, there should be options for data encryption in-flight and at-rest. These options might be dependent on the data access methods utilized. For data under the strictest compliance requirements, the consumer must be the sole owner of the encryption keys which usually means the data is encrypted before it leaves the operating system.

One other area that enterprise clouds should address is how data is handled on break/fix drives and reused infrastructure. There should be well defined break/fix procedures so that data is not accidentally compromised. When a customer vacates a service, a data erase certificate should be an option to show that the data has been deleted using an industry standard data erase algorithm.

---

<sup>1</sup>For more information on this topic, refer to the Savvis white paper "Securing Virtual Compute Infrastructure in the Cloud" at <http://www.savvis.com/cloudcompute>

### **Attribute No. 3: Automated ILM Storage**

Information lifecycle management (ILM) has been at the heart of a very effective marketing campaign by vendors who sell multiple tiers of storage. Although the value proposition behind the ILM concept is simple — align the cost of storing data to the business value of the data — the real challenge comes in the actual execution of such an objective because most so-called ILM solutions are not granular enough to achieve this goal.

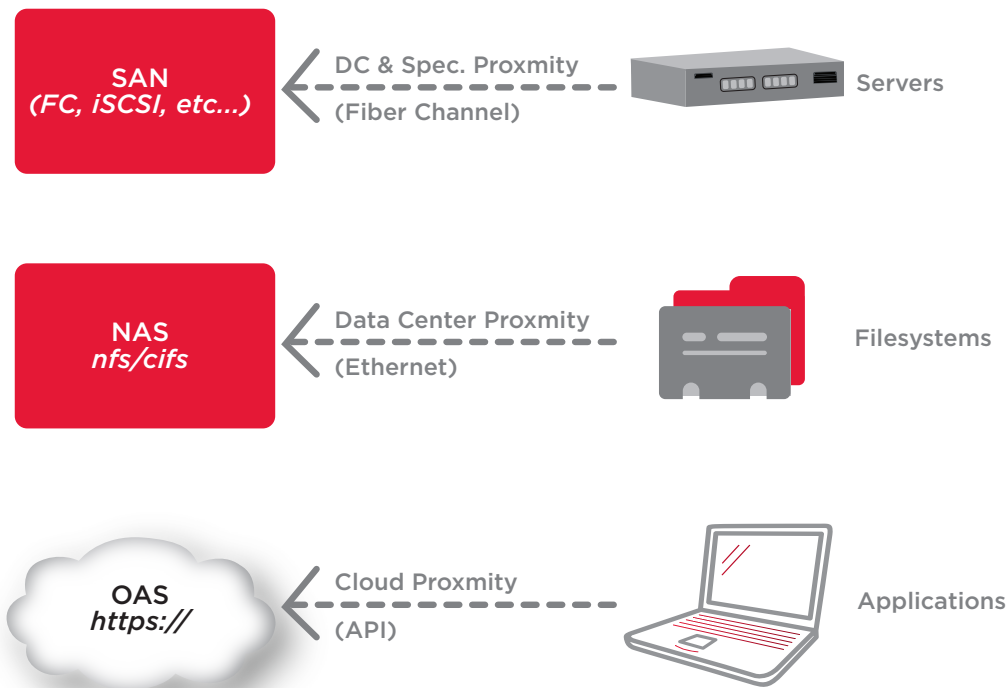
To date, ILM has not been implemented in mass market clouds. The reason is twofold. First, the spinning media used in most clouds is usually found in the bottom tier of a typical ILM solution. With no lower tier to move data to, ILM can't be deployed. Second, the complexity and cost of implementing an ILM strategy that is granular enough to be effective has been incompatible with cloud economics.

According to some industry reports, 70 percent of data is static. By storing the right data on the right media, enterprises can cut costs. Add the savings they can realize from deploying a cloud computing platform, and the financial benefits of implementing ILM in the cloud are significant. This should be possible without breaking applications or adding unnecessary complexity to operations.

To do this, enterprises should use a policy-based block-level ILM approach regardless of the access method or application type. By tracking data attributes at the block level, there is no need to perform data profiling or data movement at the operating system level. This approach is also independent of the operation system type and access method used to store the data. Not only does it optimize capital spent on storage while maintaining performance (all writes are done on the highest tier), it also reduces power consumption by “waterfalling” unused blocks of data to the lowest tier. This makes sense because near-line storage space uses about 20 percent of the power as enterprise storage space.

#### Attribute No. 4: Storage Access Method

As shown by the diagram below, there are three mainstream ways to access computing storage space. They are **block-based** (SAN or iSCSI), **file-based** (CIFS/NFS), and through **Web services**. Block and file-based access are most commonly found in enterprise application designs, and enable greater control of performance, availability, and security. At this point, most mass market clouds leverage Web services interfaces like SOAP and representational state transfer (REST) to access data. Although this is the most flexible method, it has performance implications. Ideally, an enterprise cloud provides all three access methods to storage to support different application architecture.



### **Attribute No. 5: Availability**

IT infrastructure maintenance windows have largely been eliminated due to the necessity for enterprises to support users in multiple time zones with around-the-clock availability. Although SLAs are typically tied to availability, they can be difficult to measure, from a business perspective, due to the cascading effect of multiple infrastructure component SLAs.

As mentioned earlier, I/O performance in the mass market cloud is often only best effort. If a cloud platform is dependent on parts of an infrastructure that are not managed by an internal IT group, then putting redundant infrastructure components and paths in place is the best way to mitigate the risk of downtime. Although cloud storage service providers continue to increase availability while watching costs, the SLAs in the current market do not meet the needs of enterprises' business-critical applications, as they often have caveats that exclude situations outside of the cloud provider's control.

At the high end of the enterprise cloud, the storage system resembles an in-house enterprise storage solution, with multiple paths, controllers, separate fabrics, RAID technology, end-to-end infrastructure control/monitoring, and a mature change-management process. At the low end, however, storage availability is equivalent to today's mass market cloud SLAs. To provide the range of SLAs required by enterprises, enterprise cloud providers must leverage robust infrastructure designs and innovative technologies that have been well-tested.

### **Attribute No. 6: Primary Data Protection**

Primary data is data that supports online processing. Primary data can be protected using a single technology, or by combining multiple technologies. Some common methods include the levels of RAID, multiple copies, replication, snap copies, and continuous data protection (CDP).

Primary data protection within the mass market cloud is usually left up to the user. It is rare to find the methods listed above in mass market clouds today because of the complexity and cost of these technologies. A few cloud storage solutions protect primary data by maintaining multiple copies of the data within the cloud on non-RAID-protected storage in order to keep costs down.

Primary data protection in the enterprise cloud should resemble an in-house enterprise solution. Robust technologies like snap copies and replication should be available when a business impact analysis (BIA) of the solution requires it. APIs for manipulating the environment are critical in this area so that the data protection method can be tightly coupled with the application.

The main difference between in-house enterprise solutions and storage in an enterprise cloud is how the solution is bundled. To maintain the cloud experience of deployment on demand, options must be packaged together so the service can be provisioned automatically. The result is a pick list of bundled options that typically meet a wide variety of requirements. There may not be an exact match in the frequency of snap shots, replication, and the like, for a customer's requirements. Nonetheless, most users will usually sacrifice some flexibility to realize the other benefits of operating within an enterprise cloud.

### **Attribute No. 7: Secondary Data Protection**

Secondary data consists of historical copies of primary data in the form of backups. This type of data protection is meant to mitigate data corruption, recover deleted or overwritten data, and retain data over the long-term for business or regulation requirements. Typical solutions usually include backup software and several types of storage media. Data de-duplication might also be used, but this can raise issues in a multi-tenant environment regarding the segregation of data.

There are solutions (commercial and public-domain) that can be added to mass market cloud storage offerings to accomplish secondary data protection, but it is rare for the mass market cloud providers to package this together with the online storage. Although the reasons vary, in some instances SLAs related to restore times and retention periods can be difficult to manage.

Whether the solution is a private or a multi-tenant cloud platform, control, visibility, and restore SLAs are critical for secondary data protection. Initiating a restore should be straightforward and should happen automatically once the request is submitted. Users should be able to count on some predictable level of restore performance (GBs restored / amount of time) and should be able to select the length of retention from a short pick list of options. Finally, users should also be able to check on the status of their backups online. Since frequency and retention determine the resources required for storing backups — and thus the cost — online status of usage and billing should be viewable by the consumer to avoid surprises at the end of the billing period.

### **Attribute No. 8: Storage Agility**

Storage agility simply means being able to adjust storage needs as the business requires. Ultimately, this depends on the ability of the operating system to see storage as it changes, and the access method being used. Managed Operating System (OS images provided by the cloud provider) usually have the greatest agility when it comes to increasing disk space since the drive, mount point or logical volume manager naming standard are managed by the cloud provider. Custom images (OS images supplied by the customer) can still add space but the final configure items will be up to the customer since the exact configuration of the disk space is not known by the cloud provider.

Mass market cloud offerings probably address this area the best of all nine areas discussed here. Most solutions have the ability to add incremental storage in some predefined amount. Removing space is also an option, but is usually done at the volume or mount-point level. As mentioned above, the ability of the operating system to react to these changes is usually the limiting factor.

The storage supporting the enterprise cloud needs to be scalable and billed for in a way the customer can understand. Although adding and removing storage space is important, users prefer to pay for just the space they are using. They also want the ability to make adjustments to their storage like moving data from one storage type or access method to another. Reports on their usage from a Web-based portal are also critical. This type of functionality helps them control cost and provide intelligence for business planning purposes.

### **Attribute No. 9: Storage Reporting**

When a company considers outsourcing all or part of its IT infrastructure, one frequent concern is the loss of visibility of that particular technology. Customers need the ability to understand the state of their environment from both a capacity and performance perspective. To address this concern, robust storage reporting through a customer portal is needed to instill confidence that storage is operating effectively.

Storage-related reporting in the mass market cloud has been fairly basic up to this point. Most vendors provide standard reports on usage; in some cases they also provide some basic performance stats are available — either from the provider, or via shareware or third-party tools.

The enterprise cloud has some advantages over traditional enterprise storage in that the infrastructure usually lends itself to a single storage vendor solution. This makes reporting more straightforward, since data from multiple vendor platforms doesn't have to be translated to produce a report with a single look and feel. Detailed information on historical and real-time usage, along with a few key performance indicators — both historical and real-time — should be visible 24/7 via the customer portal. Ultimately, to allay enterprises' fear of loss of control, the cloud provider should enable the most comprehensive and accurate reporting capabilities possible and make visibility into the storage system utterly transparent.

## **Conclusion**

A robust enterprise cloud should not be focused myopically on CPUs, memory, disks, and a range of IP addressees. Businesses instead should include the nine attributes presented in this paper when formulating an enterprise cloud platform strategy. By doing this, enterprises will finally be able to embrace cloud computing as an improved platform for running their businesses.

## **About the Author**

Todd Loeppke is the Technical Vice President of Storage Architecture in the Office of the CTO at Savvis. Todd provides storage technology leadership for infrastructure and product development, M&A support, and next-generation storage-related platform evaluation. Before joining Savvis, Todd held storage, UNIX and high availability architecture infrastructure/implementation positions at Maryville Technologies (a Midwest IT consulting firm) and SBC (now AT&T). Todd holds a B.S. in Electrical Engineering from New Mexico State University at Las Cruces.



To find out more about Savvis visit  
[www.savvis.com](http://www.savvis.com) or call **1.800.SAVVIS.1.**

---

**Global Headquarters**

1 Savvis Parkway  
St. Louis, MO 63017  
Tel 1.800.SAVVIS.1  
(1.800.728.8471)  
[www.savvis.com](http://www.savvis.com)

**Canada**

6800 Millcreek Drive  
Mississauga, ON  
L5N 4J9  
Tel 1.877.387.3764  
[www.savvis.ca](http://www.savvis.ca)

**EMEA**

Eskdale Road  
Wokingham  
Berkshire RG41 5TS  
United Kingdom  
Tel +44 (0)118 322 6000  
[www.savvis.co.uk](http://www.savvis.co.uk)

**Asia Pacific**

50 Raffles Place  
Singapore Land Tower  
#13-01  
Singapore 048623  
Tel +65 6768 8000  
[www.savvis.sg](http://www.savvis.sg)

**Japan**

7th Floor  
Kyodo Building  
(Jinbocho 3cho-me)  
3-29 Kanda Jinbocho  
Chiyoda-ku  
Tokyo 101-0051  
Japan  
Tel +81.3.5214.0151  
[www.savvis.jp](http://www.savvis.jp)