

The Critical Role of Storage in Hybrid Cloud Environments

Hybrid Cloud Solutions from IBM

By John Webster

September, 2016



Evaluator Group

Enabling you to make the best technology decisions



Evaluator Group

Today's Cloud Experience

Enterprise use of cloud computing has evolved from the side lines of shadow IT to the meteoric growth in enterprise usage of private, public, hybrid, and multi-cloud environments for a wide range of applications. Rapid growth in cloud usage now leads to complex private, hybrid and multi-cloud environments where a lack of administrative skills in cloud computing coupled with still maturing cloud management applications are now the biggest challenges to ongoing growth in cloud usage within the enterprise.

While it is clear that public clouds continue to be the most popular cloud platforms, private clouds have emerged from within the same environments and now dominate the cloud landscape as well. Therefore, it is no surprise that most enterprises see bridging the two—the hybrid cloud—as a strategic endpoint. In fact, survey data as shown in Figure 1 below indicates that a large majority of all enterprise cloud users (70%) are currently engaged in hybrid cloud implementations or pilot projects.

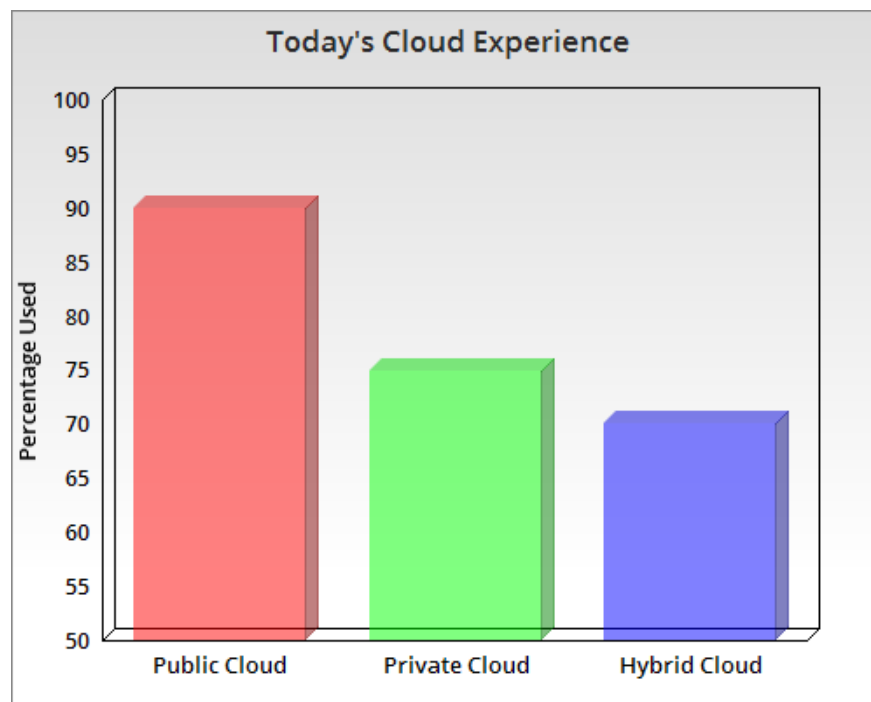


Figure 1. Percentages of US enterprises using public, private and hybrid clouds. Note that 70% of enterprises now use hybrid cloud architectures in some form. (Source: Evaluator Group Cloud Trends Analysis based on publicly available survey data.)

The ability to deliver on-premises performance and manageability with data stored in the cloud is a key requirement of the enterprise hybrid cloud. As the use of platform-as-a-service (PaaS) and cloud applications grows, there is increased pressure to create an agile cloud experience with data stored on-premises. Business application users want the speed, agility and responsiveness of the public cloud. Enterprise IT executives have gotten the message and want to deliver that same experience, but with an ability to control cost, assure quality of service (QoS) and comply with data management requirements for security, accessibility, and governance. A strategic approach built upon active data management across on-premises, private and public clouds is necessary. Active data management enables centralized IT to solve cost, performance, and corporate governance issues as the enterprise moves forward with hybrid cloud computing.

Why Hybrid Cloud

The hybrid cloud that bridges private and public clouds to enable data and application portability is now a common destination for enterprise cloud initiatives across all industries. The reason becomes clear when one adds up all of the potential of the benefits of hybrid cloud vs other cloud computing models. These include:

A Total Cloud World View – Line of business (LoB) users are concluding that, while use of cloud-based applications has become business-critical, they have neither the time nor the staff expertise to connect and support these applications. As a result, LoB users are looking increasingly to enterprise IT to set cloud usage policies as well as select public clouds and private cloud technologies. Cloud selection and management has become a collaborative effort between LoBs and enterprise IT. A hybrid cloud consolidates assets and gives enterprise IT administrators a complete cloud world view. Due to the number of vendors and protocols that may be involved, primary storage lags compute and application deployment today.

Agile Applications Management – The need for agility is common across enterprise cloud initiatives. Public clouds are the model of agility that enterprise IT wants to replicate in private clouds. A hybrid cloud architecture enhances agility by enabling applications to be hosted and re-hosted quickly on the most appropriate platform. To be fully effective, the data must follow the applications.

Cost control – Many public cloud users have experienced escalating and complex monthly bills from public cloud service providers. Building a bridge between public and private clouds allows applications and data to be re-hosted on premises where the cost of application support and data storage can often be reduced and cost control can be more effectively applied. Public cloud can still provide a reliable, lower-cost alternative to building a second site or for temporary deployments.

Operational control – Enterprise IT administrators commonly complain that they have little management control over public cloud computing instances. A hybrid cloud architecture lets them choose which management environment is the most efficient. Applications that need more IT administrative attention can be hosted on-premises or moved as needed.

Optimized Application Delivery and On-going Support – The common use of DevOps processes by cloud users denotes an application instantiation processes that streamlines application development and testing while more tightly integrating deployment and updating with IT operations. With a hybrid architecture, developers can build cloud native applications on the most appropriate “side” of the cloud and host them on premises or off at will.

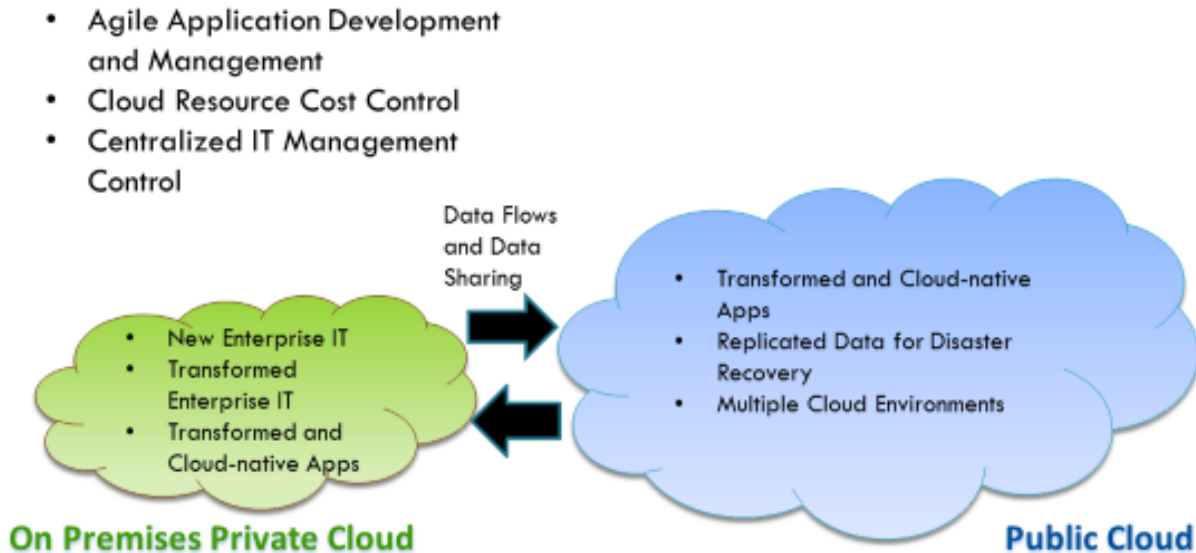
Data Sharing Across Cloud-based Applications – Without a hybrid cloud infrastructure, public cloud data remains locked up there as is the case on the private cloud side. Hybrid clouds promote data sharing for a wide range of traditional and emerging use cases including mobile application delivery, Big Data analytics and IoT (discussed in more detail below)

Improved security – public cloud security has improved greatly since the early days of public clouds and may now exceed that of many enterprise data centers. In fact, survey data indicates that security is no longer the number one inhibitor for enterprise cloud usage.

Application availability – in addition to security, application availability has also improved and for some applications may exceed that of on premises.

New features for application development and support -The public cloud service provider can offer an array of application services that could be harder to duplicate within a private cloud. Hybrid cloud allows enterprise IT to leverage these features without having to build them on the private cloud side.

Hybrid Cloud in the Enterprise



Integrating Private and Public Clouds—the Hybrid Cloud Destination

Storage for Hybrid Cloud—a CIO Focal Point

In spite of this new and rapidly evolving IT environment, data and the ability to preserve, protect and manage it is still as fundamental to information delivery as it was at the dawn of mainframe computing. A siloed and complex data storage environment can slow delivery of cloud-based applications and negatively impact agility. It can also increase exposure to data loss as the enterprise move forward with cloud computing initiatives. Bring clarity to the enterprise cloud strategy going forward by laying a scalable foundation for cloud data management.

But just because cloud is a new computing environment, it doesn't mean that all of the time-honored IT practices data no longer apply. It has usually been the case that operational challenges such as the assurance of application performance and data integrity can be met by implementing the right storage technologies. And that is still the case with cloud computing.

The good news is that a foundation for enterprise production quality storage may already exist in many enterprises. Upwards of 80% of hybrid cloud implementations are for business continuance and disaster recovery meaning that storage based links between on premises private clouds and public clouds have already been established via data copying between the on premises data center and the remote public cloud site. These disaster recovery capabilities can serve as a data layer foundation for additional data

services that support bi-directional application movement for cloud bursting and application development. Exposing copied data to cloud applications—cloud native as well as cloud-transformed applications—is a key enabler for sharing persisted data and creating a foundational data layer to support hybrid cloud applications.

However, hybrid clouds bring a new set of storage requirements that include:

Storage agility to enable overall IT agility. Enterprise IT can't create an agile services delivery engine while the underlying data layer is static, fragmented and unresponsive to the immediacy of new business requirements.

Automated Quality of Service (QoS) to assure that business users always see a consistent level of application responsiveness in spite of differing and often-changing workload conditions.

Open, secure data access enabled by a diverse set of protocols so that data can be accessed by a range of applications present in hybrid cloud environments. Multi-protocol support also breaks down data silos and allows data to be shared and moved more freely and at reduce cost.

Global location-awareness of where data is placed, the latency associated with placement, and the cost of moving it around.

Responsiveness to existing IT management and data governance and policies, access control and security measures, data integrity and protection best practices, and methods for complying with governmental and industry-specific mandates.

Transparent workload migration to the most appropriate platform at any given time. Often this means that the cloud storage environment will be called upon to manage metadata separately for responsiveness and predictability as well as offer a machine learning capability to predict when and to where a workload resides.

Centralized, automated storage infrastructure management that functions across private and public cloud storage environments and can be integrated with cloud management platforms.

Integration with the company's virtualization platform preferences as a foundation for hybrid cloud. Many IT organizations are in the process of transforming their existing IT infrastructures into private clouds. This is particularly true of VMware users who are instantiating vCloud and vRealize as well as Microsoft Hyper-V users planning on making AzureStack the central to their transformational strategies. Integration also includes management interfaces so that storage can be managed from the perspective of a cloud administrator.

The IoT Use Case

As mentioned, disaster recovery is a common premise for instantiating a hybrid cloud from the perspective of data replication between on-premises and public cloud storage. But we know that IT executives want to build-out a far more robust hybrid cloud infrastructure for application mobility coupled with data mobility. An Internet of Things (IoT) use case for hybrid cloud provides an example of what the foundational data layer could look like. Here, the ability to share data becomes an important consideration.

IoT is commonly based on the use of sensors or other electronic devices that are given the ability to transmit data via wired or wireless network. These devices generate continual streams of data and, depending on the number of devices for a given application as well as the types of data transmitted, an IoT network can generate very large volumes of data. It is now common practice to “land” this data initially in a large scale storage environment often called a “data lake.” And while the data lake can be located within a public cloud or instantiated on premises, the public cloud has initially become the destination of choice because of the unpredictable nature of IoT data streams. At this point in time, enterprise IT is still climbing up the IoT learning curve. Because of unexpected data volumes, public cloud storage capacity can be activated quickly and on an as needed basis.

At these data aggregation points, a number of processes are typically applied including data “cleansing,” data correlation from the convergence of multiple streams as well as with contextual data, and time-based trends analysis. The resulting structured or unstructured data can then be fed into a wide array of different analytics processing platforms including Hadoop, Cassandra, MongoDB, as well as traditional data warehouses.

Hybrid cloud storage comes into play in the following ways:

1. Raw IoT data is directed to a public cloud. Examples are Databricks for Spark Streaming and IBM SoftLayer. Here, a number of processes such as querying and the convergence of stream data with enterprise data from the private cloud side can be performed. The results can then be stored for longer term retention at the public cloud site as historical data for further analysis and replicated to at private cloud site to make it available to other applications running on premises.
2. Alternatively, raw IoT data is directed to an on premises private cloud where analytics processes are performed. The results can then be stored for longer term retention and further analysis on premises and replicated to public cloud storage to make it available to applications running there and/or for disaster recovery purposes.

IBM's Hybrid Cloud Storage Portfolio

IBM's Spectrum Storage portfolio offers a range of storage platforms and software suites that span from high performance, low latency all flash arrays to highly scalable, distributed content repositories. They also address all phases of data management including data protection, secure archive and automated infrastructure management.

All storage platforms are software-driven and feature a common set of enterprise data center grade storage attributes. However, in that hybrid cloud is the stated architectural direction for many enterprise IT organizations going forward, we believe that it is important to position these attributes within the context of a hybrid cloud deployment:

Hybrid Cloud Storage Requirements	IBM Spectrum Storage Attributes
Storage agility	<ul style="list-style-type: none"> • Modular software driven to deploy data and storage services quickly on or off premises • Scale performance and/or capacity at any time • Availability as pre-integrated systems or as software only
Open, secure data access	<ul style="list-style-type: none"> • Support for multiple, open APIs • Secure multi tenancy • Built-in key management
Transparent workload migration	<ul style="list-style-type: none"> • Multiple data copy and migration facilities that function without interruption to production systems residing on either side of the hybrid cloud • Data virtualization

Responsiveness to policy management	<ul style="list-style-type: none"> Centralized, automated policy enforcement and management via software for data governance and compliance
Global location-awareness	<ul style="list-style-type: none"> Automated copy data management across on and off premises application deployments for data protection, test/dev and big data analytics applications
Virtualization platform preferences	<ul style="list-style-type: none"> VASA, VADP support VMware vCloud and vRealize support Future support for Microsoft AzureStack
Automated QoS	<ul style="list-style-type: none"> Automated storage tiering for performance Workload isolation in a multi-tenant environment Data availability across software upgrades, hardware refreshes, and in the face of disk, node, and site failures
Automated Infrastructure Management	<ul style="list-style-type: none"> Control management and security layer that spans multiple storage environments and multiple platforms

Storage platforms within the IBM Spectrum Storage portfolio are available as software only, as cloud services and as integrated storage solutions. They include:

Spectrum Accelerate – Scale-out block for all flash and hybrid flash/disk arrays. IBM FlashSystem A9000 and A9000R together with XIV systems use this software. Spectrum Accelerate is also available as a cloud service.

Spectrum Virtualize – Virtualized block storage for Flash and disk arrays that virtualize internally and externally attached storage arrays including third party arrays. Systems built with this software include IBM Storwize family, FlashSystem V9000, SAN Volume Controller, and VersaStack converged systems.

Spectrum Scale – Unstructured data storage scalable to PBs. IBM’s Elastic Storage Server is based on Spectrum Scale software.

Spectrum Archive – Automated archiving software that moves data from disk to tape

Spectrum Protect – Data protection and resiliency software including deduplication. Spectrum Protect forms the core of many cloud backup and recovery offerings from IBM and others.

Spectrum Control – Automated storage infrastructure management software for provisioning capacity management, availability monitoring, and reporting. Spectrum Control is also available as a cloud service, called Storage Insights.

Spectrum Copy Data Management – copy creation and cataloging for backup, development or deployment.

Acquiring and Financing IBM Spectrum Storage

IBM Spectrum Storage can be acquired in the form of a pre-integrated hardware/software systems, as software only that can be loaded onto industry standard servers, or as a service from IBM. However, infrastructure leasing is also an alternative that should be considered for hybrid cloud deployments because it matches operational agility with financial agility.

Infrastructure Leasing via IGF

Leasing compliments cloud infrastructure agility with financial agility. Leasing is an off-balance sheet acquisition model that is highly conducive to the cloud pay-as-you-go utility model. IGF can provide an operating lease that adds pay-as-you-go and financial agility features that are a good fit with cloud and IT agility¹. The essential point of an IGF operating lease is to transfer ownership of the storage infrastructure to IGF while enterprise IT installs, uses, and afterward returns the infrastructure, thereby accruing all of the benefit from its usage without all of the capitol acquisition strings attached to it.

IGF offers three distinct alternatives to capital acquisition:

¹ IGF lease offerings are contingent upon availability in the country where the customer is located and credit approval.

Fair Market Value (FMV) Lease: IGF typically retains ownership while the customer uses the infrastructure.

Full Pay-out (FPO) Lease: Ownership of the infrastructure can be transferred from IGF to the customer at the end of the lease for a nominal fee—usually \$1.00.

Loan to Purchase: A customer takes out a loan through IGF and owns the infrastructure when all of the payments have been made. This option is attractive in countries where a lease cannot be offered since it still provides cash flow benefits with payments over time.

IGF can also allow a customer to take agility-related advantages of the modular characteristics of Spectrum Storage. During the lease term, customers can add capacity and additional racks as well as upgrade modules without upgrade surcharges or increasing monthly payment amounts by extending to original lease term.

Variable Monthly Payment and Deferred Payment Plans

An IGF lease can be structured such that the payment required at the beginning of the lease term is low, and progressively increases toward the end of the lease when. This is particularly beneficial in situations where the current business climate is difficult or there is a need to start a new cloud project with limited funding at the inception of the project. Furthermore, if there is little to no money left for additional infrastructure in the current year's budget, IGF will offer the customer an ability to install new infrastructure and defer the start of the lease and monthly payments to the following year.

Conclusion

Cloud computing now dominates enterprise IT decision making. However, among the possible cloud computing architectures now practiced, hybrid cloud is the most commonly sought-after end point in the progression from traditional IT to cloud-enabled IT. This we believe is true because the hybrid cloud offers an ability integrate all enterprise cloud resources under a single umbrella that can be holistically managed by enterprise IT. From an information perspective, hybrid clouds also offer an ability to integrate data stores and share data across cloud-based applications.

As such, hybrid clouds create a new data environment for persistent storage that comes with new requirements for performance, protection, security, and manageability. Multiple strategies and technologies will be required. Understanding this, IBM is now positioning its Spectrum Storage product portfolio within the hybrid-cloud as well as optimizing each offering for hybrid cloud deployment. The result is a storage product lineup that offers the broadest coverage of hybrid cloud storage needs than any other currently available.

About Evaluator Group

Evaluator Group Inc. is a technology research and advisory company covering Information Management, Storage and Systems. Executives and IT Managers use us daily to make informed decisions to architect and purchase systems supporting their digital data. We get beyond the technology landscape by defining requirements and knowing the products in-depth along with the intricacies that dictate long-term successful strategies. www.evaluatorgroup.com @evaluator_group

Copyright 2016 Evaluator Group, Inc. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or stored in a database or retrieval system for any purpose without the express written consent of Evaluator Group Inc. The information contained in this document is subject to change without notice. Evaluator Group assumes no responsibility for errors or omissions. Evaluator Group makes no expressed or implied warranties in this document relating to the use or operation of the products described herein. In no event shall Evaluator Group be liable for any indirect, special, inconsequential or incidental damages arising out of or associated with any aspect of this publication, even if advised of the possibility of such damages. The Evaluator Series is a trademark of Evaluator Group, Inc. All other trademarks are the property of their respective companies.