IBM Hybrid Cloud
White Paper

Hybrid cloud for big data and analytics solutions

Create a hybrid cloud that's seamless, secure, and built for analytics

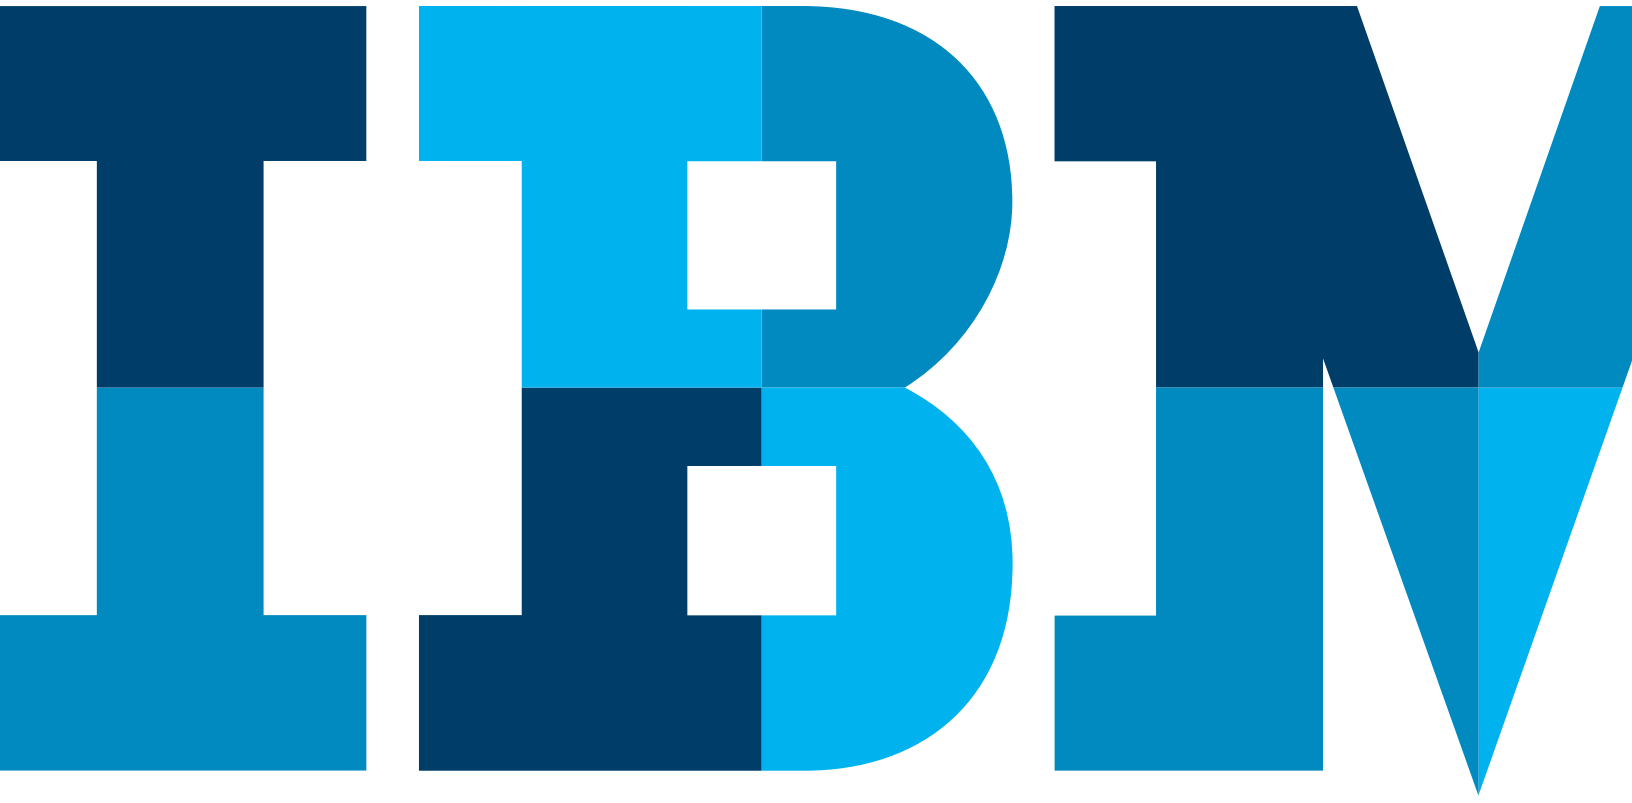

IBM

Contents

Introduction

In our strategic discussions with IT leaders and their C-level business counterparts, we hear time after time that they are focused on — or at least considering — either shifting existing workloads to the cloud, extending existing workloads to the cloud, or building new workloads on the cloud and integrating those with existing workloads. For some, this discussion underscores the need for multispeed IT, a topic that IBM Analytics Group CTO Tim Vincent discusses in our Analytics InsightOut series on Big Data Hub.

Quite often, we see that the need for data security and governance makes some organizations hesitant about migrating to the cloud. This is perfectly understandable given the types of data gathered and used by businesses today, the regulations they must adhere to on both a local and global level, and the cost to maintain data and operational infrastructure. Fortunately, the business model for cloud technology is evolving to enable more businesses to deploy a hybrid cloud, particularly in the areas of big data and analytics.

A hybrid cloud is a combination of on-premises & local cloud integrated with one or more dedicated cloud(s) and one or more public cloud(s). We refer the combination of the on-premises & local cloud with the dedicated cloud(s) as **"private environment"**. The public cloud and private environments are structured so that they operate independently, but communicate with each other via an encrypted connection on a private and/or public network, using technologies that facilitate the portability of applications and data.

A hybrid cloud allows organizations to integrate personal and/or confidential information from the private environment with applications running on the public cloud, while leveraging the public cloud's computational resources and storage. For example, organizations can generate actionable insight by integrating the data from Systems of Record (private environment) with Systems of Engagement in a public cloud or by applying edge-analytics on the devices in the public cloud.

In addition, hybrid cloud increases scalability by allowing organizations to use public cloud resources for situations where the private environment doesn't provide adequate computational power. Furthermore, containers can be used to increase the portability of workloads between private environment and public cloud. Finally, hybrid cloud is ideal for global distribution of applications and data, allowing better management of data sovereignty and compliance.

In this document, we will summarize what hybrid cloud is, why it's important in the context of big data and analytics, and how to implement a hybrid cloud strategy with help from IBM.

What is hybrid cloud?

A hybrid cloud is the connection of the private environment with one or more public cloud(s) as shown in Figure 1. It leverages the best of what each environment has to offer, providing the flexibility to locate data and services based on business need. Data can be located and accessed based on consumption patterns and analytical workload requirements within hybrid cloud environments, providing data & analytics for the different personas where it is needed. Access to all areas of the hybrid cloud environments are managed and controlled to uphold privacy, security and other data governance requirements.

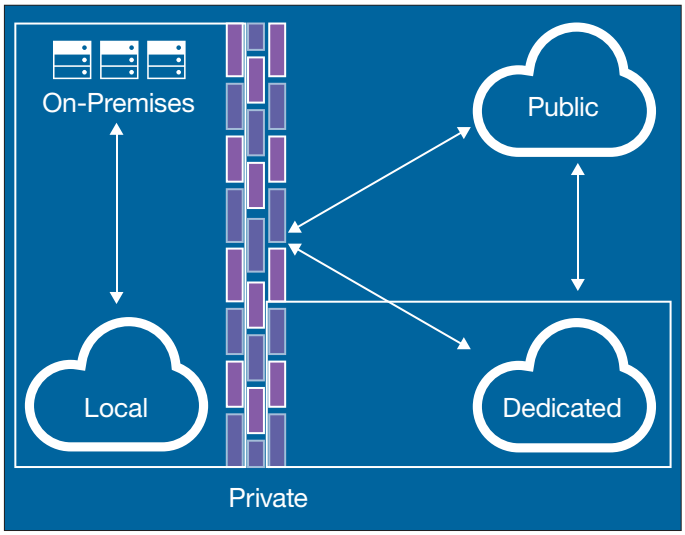

Figure 1: Hybrid cloud

The digital transformation requires a new hybrid cloud—one that's open and flexible by design, and gives clients the freedom to choose and change environments, data and services as needed. This approach allows cloud apps and services to be rapidly composed using the best relevant data and insights available, while maintaining clear visibility, integrated control, governance and security everywhere. As shown in Figure 2, the majority of the Systems of Records usually resides on the private environment, while the Systems of Engagement and Systems of Automation are mostly on the public cloud(s) and the Systems of Insight span across all environments of the Hybrid Cloud[1]. The flexibility and openness of the hybrid cloud allows the data and the associated analytical workload to be placed where it makes the most sense in terms of business needs. The information privacy and security is managed and controlled consistently across all the systems of the hybrid cloud environments.

Our definition of hybrid cloud is consistent with the majority of our clients who want to extend private environment to the public cloud. We believe that private environment is essential to the model because most businesses will always require portions of their data and infrastructure to remain behind the corporate firewall due to industry standards, local regulations, or their own attitudes toward controls. This creates an even more flexible architecture by giving businesses more freedom to choose and change their environments and deploy services and applications more quickly.

We view hybrid cloud strategy as an overall architecture solution, and not just a migration path. The goal is to allow you to extend workloads from a pure private environment model to a Hybrid model that couples private environment and public clouds. The strategy certainly can be used to direct your organization toward the cloud, but it can help you accomplish the integration of the environments in the hybrid cloud.

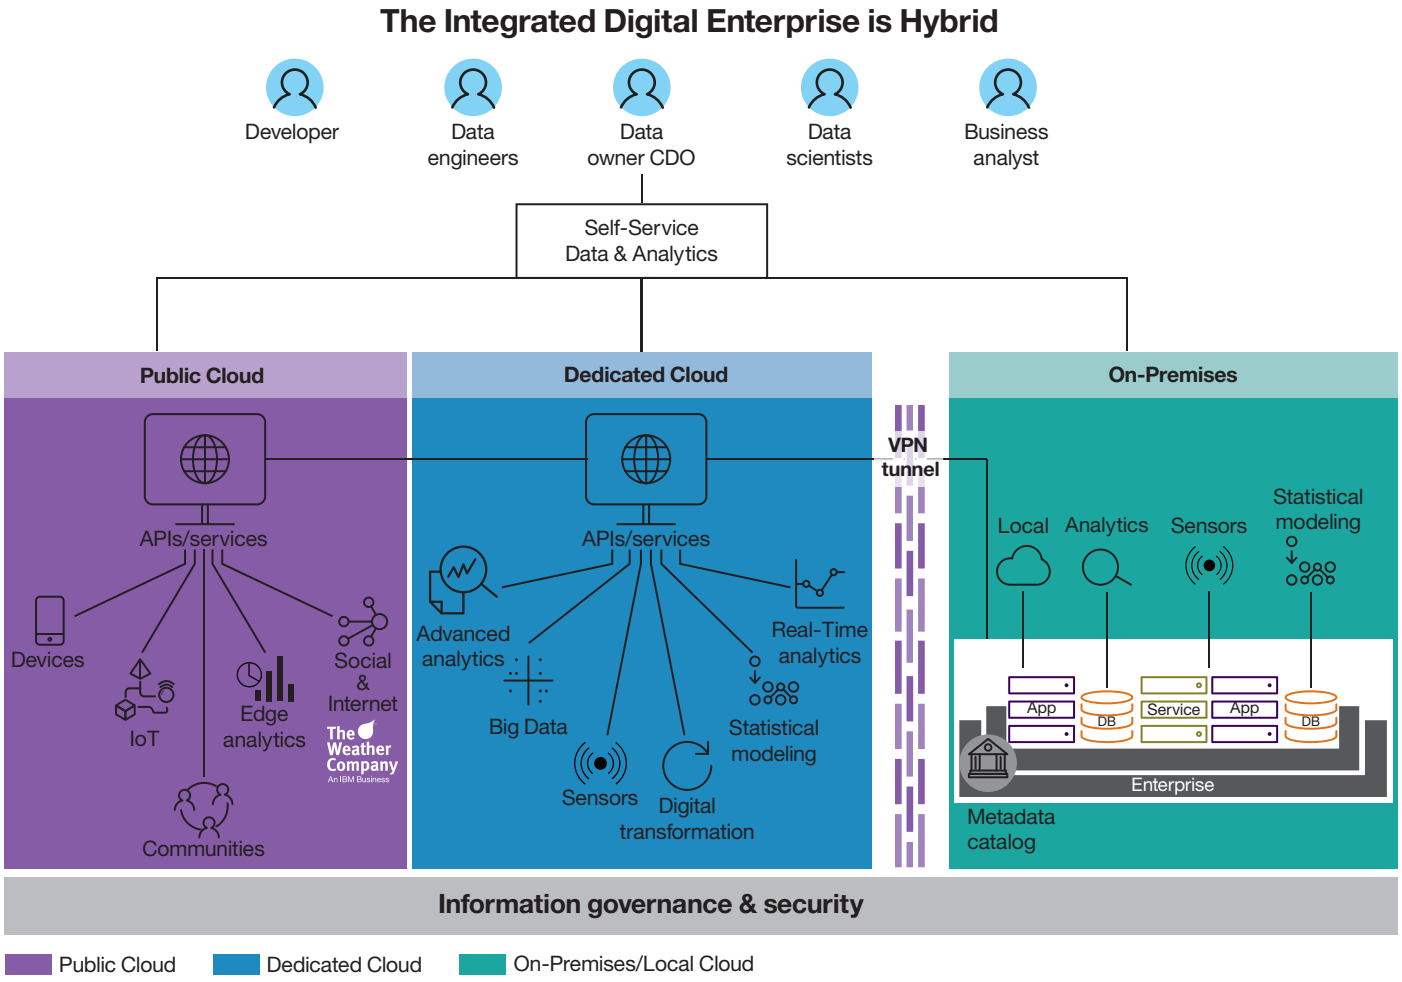

Figure 2: Hybrid cloud for the digital enterprise

Why hybrid cloud for big data and analytics?

A hybrid cloud allows different personas to work with data and analytics capabilities where it makes the most sense for them to do so and this helps to define the requirements where the data and analytics capabilities should be placed/available in the hybrid cloud environments. As a result, analytics workloads can run more efficiently wherever the data is stored.

It's important to have hybrid cloud as an option because location should be one of the first architectural decisions for any analytics project. In particular, organizations need to consider where the data should be stored, and where the analytical processing should be located relative to the data. Meanwhile, legal and regulatory requirements also impact where data can be located, as many countries have data sovereignty laws that prevent data about individuals, finances and intellectual property from moving across country borders.

Systems are going to have multiple centers of gravity which will dictate where processing will occur. For example, if building a data lake as part of a Systems of Insight and the data that feeds the data lake is in the private environment then the center of gravity will be on the private environment and the processing of the data should stay within the private environment. But if the Systems of Insight starts including data born on the public cloud then there could be a second center of gravity.

Primary drivers for big data and analytics in the hybrid cloud

- **Integration:** Organizations need to integrate data that is stored and managed in a hybrid environment across the private environment and public cloud(s). Typically, these organizations need to integrate Systems of Engagement and/or Systems of Automation (IOT) applications, such as social media, customer management systems, and devices, with Systems of Insight, such as predictive and real-time analytics hosted on public clouds, and mission-critical applications and data stored on servers in the private environment (Systems of Record).
- **Brokerage/management for workload and resource optimization:** Different workloads have different requirements for security, speed, resources and storage. Many organizations are driven to hybrid cloud because they want the option to place the data and the analytical workload where it makes the most sense based on the business requirements. These organizations want the ability to optimize cost, performance and agility, while also enjoying the flexibility to move data and analytical workloads between private environment and public cloud.
- **Portability:** Another major case for hybrid cloud is the need to ensure portability of analytical workloads and data. In order to manage costs and effectiveness, IT management needs to be able to move workloads and data to whatever platform best meets changing customer demands. This capability requires IT to consider the feasibility of the new analytical workload and data on a specific hybrid cloud environment based on the overall hybrid cloud architecture.
- **Compliance:** A hybrid cloud allows for distributing global applications, data and workloads across geographically dispersed private environment and public cloud(s) where the requirements for data sovereignty, compliance, privacy, identity management, and data protection could imply for the data and consequently the workload to be placed on a specific environment in a specific country. An organization can choose to deploy cloud environments that are already compliant with regulatory requirements (such as HIPAA, PCI, and SOX), and are located in a specific country to comply with local privacy and data sovereignty laws.

Considerations for implementing a hybrid cloud strategy for big data and analytics

The key considerations for implementing a hybrid cloud strategy include:

- **Cultural shift:** One of the biggest challenges in moving to a hybrid cloud is establishing and promoting a collaborative, service-oriented approach for provisioning data & analytics and self-service capabilities to be able to extend private environment to the public cloud.
- **Varying levels of hybrid sophistication:** A hybrid cloud strategy can have different levels of sophistication: deep integration between private environment and public cloud(s), or more simplistic, static, point-to-point connections using a virtual private network (VPN), a secured gateway, and an API manager designed to expose systems of record data (private environment) to systems of engagement (public cloud).

Figure 3 summarizes different hybrid cloud scenarios.

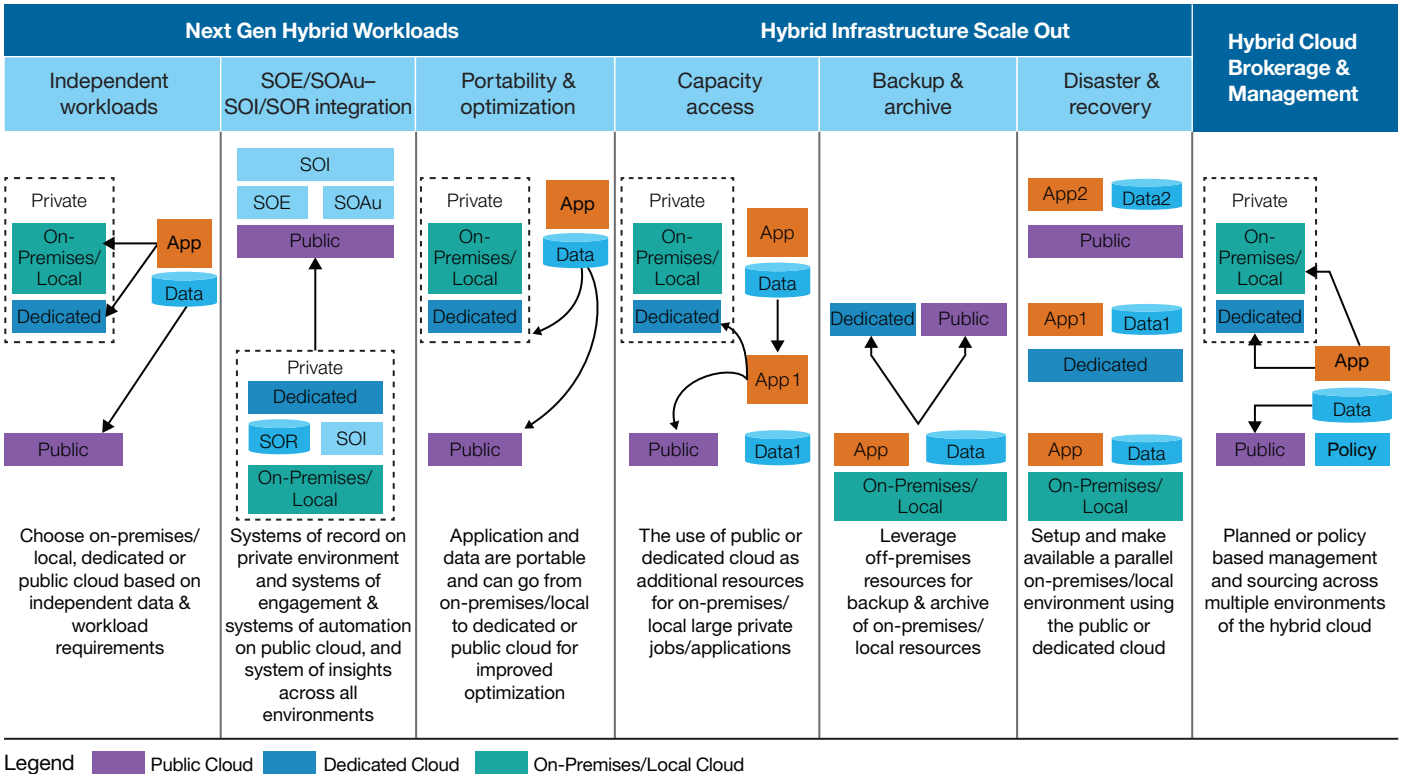

Figure 3: Hybrid cloud scenarios

The use case that is one of the primary drivers of hybrid cloud is the integration of Systems of Engagement (SOE) and Systems of Automation (SOAu) with Systems of Insights (SOI) and Systems of Records (SOR). This is the primary case for extending private data to the public cloud. Organizations want the full agility and flexibility to generate personalized customer offers and respond to market dynamics. This use case requires a new type of architecture, such as an enterprise data lake; that is, a group of fit-for-purpose repositories that are well managed, governed, protected, connected by metadata and provide self-service access[2]. How a typical enterprise data lake may fit in the hybrid cloud strategy landscape, including different IT domains such as Systems of Record, Systems of Engagement, Systems of Automation, and Systems of Insight? It depends on the data gravity of the organization. If the data gravity stays inside of the private environment, then the data lake will be defined inside of the private environment. If the data gravity is shared between the private environment and public cloud, then the data lake will be defined across both private environment and public cloud. If the data gravity stays inside of the public cloud, then the data lake will be defined on the public cloud.

Some of the key topics that need to be considered when planning a hybrid cloud strategy are:

- **Data and analytics**—where to store the data based on the analytical workload and the type of access?
- **Data movement and replication**—where to store the data to avoid data movement and replication?
- **Data preparation and integration**—what is the best environment to do data preparation/integration?
- **Data sovereignty and compliance**—where to store the data based on data sovereignty requirements and which compliance regulations are required?
- **Data governance and security**—how to secure the data across all environments and which access control needs to be in place?
- **High availability and disaster recovery**—do we need high availability and disaster recovery for this application/data in the cloud?
- **Network configuration and latency**—do we need to change the network configuration to satisfy the latency requirements of an application in the cloud?
- **Portability**—do the application and associated data need to be portable to different environments?
- **Scalability**—does this application require the capability to scale up and/or down?
- **Resource orchestration**—how to define the resource orchestration required for all workloads?

These aspects are discussed in more detail below.

Data and analytics (descriptive, predictive, and prescriptive)

How does an organization decide where to put data on a hybrid cloud and how to use it? What's the best strategy to balance sharing and mobility with the need for privacy and security on a hybrid cloud? Data is such a valuable asset or a form of intellectual property that organizations are naturally concerned about moving it. The future of the hybrid cloud strategy for most traditional organizations is still unclear. Across industries, traditional organizations are still trying to figure out the best use cases to move to the cloud, and most have not even started a hybrid cloud strategy yet.

One of the most important hybrid cloud considerations for all organizations is data gravity. To limit data movement, production workloads should be processed where the data is stored. This means that an organization should consider the analytical workloads that they will be processing when deciding where to locate their data. Another important consideration is where the data discovery & exploration and advanced analytics will happen when deciding where to locate the data.

A hybrid cloud data lake (based on IBM's data lake strategy) is important to the hybrid cloud strategy because it helps define the data topology of the hybrid cloud, based on the following characteristics:

- **Data consumption patterns**—How will different personas work with the data across the multiple environments of the hybrid cloud? This will define where the data should be available.
- **Analytical workloads**—What is the best environment for each analytical workload? This will help to define where the data should be stored.
- **Integration with other data repositories**—What is the best way to do data integration across the multiple environments of the hybrid cloud? This will define where the data integration process should run based on data gravity.
- **Data movement requirements**—How to minimize data movement across the multiple environments of the hybrid cloud? This will help to define where the data should be stored.
- **Data sovereignty and compliance requirements**—What is the best environment for the data sovereignty and compliance requirements? This will help to define where the data should be stored.
- **Data governance and security requirements**—What are the access rules and security constraints required for each kind of data across the multiple environments of the hybrid cloud? This will define the level of security and who can have access to the data.

A hybrid cloud data lake helps organizations overcome data silos by providing the data and metadata with information governance based on business needs, and allowing users to access the data they need from where it makes sense. This type of on-demand data engagement produces insights when they are needed most. Organizations should strive to get as much data into the data lake as possible, while implementing information governance policies, standards and tools to control user access and data provenance. The data lake also allows self-service capabilities and multispeed IT, enabling users such as data scientists and business analysts to locate data by using the information catalog to easily conduct analytics without IT assistance, and to generate their own insights. The data lake can be on the private environment or can also on the private environment and public cloud together. The best way to define where the data lake should be is to consider the key attributes when doing a data topology exercise as described above.

Organizations are using the cloud for a range of business use cases, including reporting, sandboxes, production, IoT analytics, and much more. Cloud analytics services offerings are evolving and becoming more popular, especially with business customers. Having a hybrid cloud architecture that can provide access to new technologies as they emerge without requiring IT departments to learn, install or support them can be a significant accelerator for business analytics solutions. Today's new SaaS offerings are increasingly targeting specific business areas such as churn detection as a service, fraud detection as a service, and marketing campaign as a service. These offerings can improve business outcomes much faster than in-house efforts would. They also require the integration of private environment with the public cloud - another driver for a hybrid cloud strategy.

The goal of any analytics solution is to provide an organization with actionable insights for smarter decisions and better business outcomes. Different types of analytics, however, provide different types of insights. There are three principal types of analytics: descriptive (What has happened?), predictive (What could happen?), and prescriptive (What should we do?).

Across industries, hybrid cloud is helping organizations more efficiently deploy predictive and prescriptive analytics use cases, including:

- **Retail:** Create personalized marketing campaigns using real-time actionable insights into customer shopping behavior across all channels, including sentiment analysis.
- **Healthcare:** Analyze clinical data, identify trends, detect patterns and predict outcomes. Help healthcare organizations anticipate change, so they can plan and carry out strategies to improve patient outcomes.
- **Financial services:** Allow faster and more efficient fraud detection at the point of sale.
- **Telecommunications:** Allow more efficient network optimization and failure detection.

Figure 4 shows the data analytics lifecycle on hybrid cloud and information flows on three distinct steps. The first step includes searching and gathering information using the information catalog on the private environment and/or public cloud and using data from all environments of the hybrid cloud. The second step includes developing business insights (analytics) on the private environment and/or public cloud and using data from all environments of the hybrid cloud. The third step includes deploying those insights into the business process on any environment of the hybrid cloud where it makes most sense before repeating the cycle. The hybrid cloud allows different personas to work on different environments and data topology zones, collaborating with each other to generate new insights.

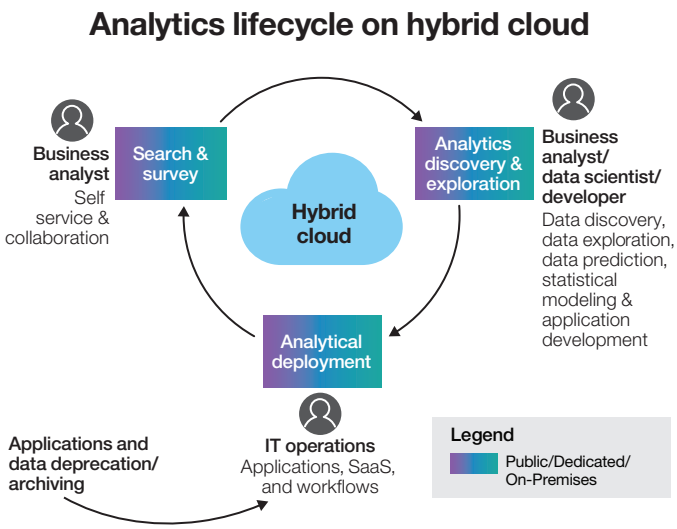

Figure 4: Analytics lifecycle on hybrid cloud

Data movement

One of the biggest challenges of hybrid cloud is the movement of data within the hybrid environments. By its very nature, data can be cumbersome and difficult to move. In addition, data needs to be carefully protected and controlled as it's being moved. Copying large volumes of data to the public cloud currently takes too much time and bandwidth, but doing it incrementally with micro-batches, streaming or asynchronous replication can be an option to reduce time and latency. Today, all cloud providers require copying data to a disk appliance to move the data to the cloud if the volume is greater than 5 TB.

Data synchronization is another challenge. When changes are made on the prime copy (private environment), how do you replicate the changes to the public cloud instances to prevent out-of-sync information?

The current options for data movement among the hybrid environments includes the use of a disk appliance for large data volumes, the use of traditional file transfer protocols for small data volumes, and high-speed file transfer tools for medium/ large data volumes.

Examples are:

- Import/Export using a disk appliance/USB drive for large data volumes; data can be compressed and encrypted
- FTP/SFTP using TCP protocol for small data volumes; can be susceptible to latency
- API using a REST API call or messaging API via Kafka for small data volumes; can be secured with HTTPS, and can be susceptible to latency
- Tsunami UDP using a file transfer protocol including TCP for control and UDP for data to transfer small/medium data volumes over very high speed long distance networks
- Parallel Remote File Transfers using bbSCP, bbFTP, or gridFTP for small/medium data; data can be compressed but it cannot be encrypted

- Data Replication using WANDisco for medium/large data volumes; can be scaled out to maximize bandwidth, and can be used for DR replication
- High-speed File Transfer using Aspera for medium/large data volumes; supports compression, encryption, and auto scaling

The current options for moving data within hybrid cloud suffer from the low speeds and high latency provided by the typical wide area network (WAN). Latency in these environments can reach into the milliseconds, compared with microseconds expected in the private environment. This means public cloud instances may be inefficient and slow. Therefore, latency it is a very important factor when considering different scenarios on a hybrid cloud such as:

- **Operational applications on the cloud**—the data should be collocated with the application to avoid latency
- **Large volume data movement**—should avoid large scale data movement even with VPN and compression. All cloud providers require copying data to a disk appliance to move the data to the cloud if the volume is greater than 5 TB.
- **Incremental data movement**—micro batches, streaming or asynchronous replication are the best options to move data to the cloud.

There are a few possible ways to reduce/improve latency in hybrid cloud environments.

The top four ways to reduce latency and maintain performance

1. **Location, location, location:** It takes more time to cover longer distances, even when moving at the speed of light. So, to reduce hybrid cloud latency, organizations should connect to a public cloud facility that is geographically close to its private environment.
2. **Dedicated connections:** The Internet is a single network that is shared between billions of users. This means public Internet connections may experience traffic congestion, causing bottlenecks and increased latency between private environment and public cloud(s). To avoid this and reduce hybrid cloud latency, organizations can establish dedicated private connections between their private environment and public cloud(s).
3. **Optimize traffic:** Even with caching, there are instances when data has to be moved within a hybrid cloud, and this can cause an increase in network traffic. Latency will occur each time a packet moves across the Internet, but applications can reduce the cumulative latency of those packets by moving fewer of them. Data compression allows an organization to pack more data into each packet, which lowers the number of packets needed to move a data set.
4. **Optimize workloads:** Developers shouldn't overlook the importance of workload design and its influence on network latency. Workloads designed to use massive data sets, or deal with critical, time-sensitive data, can be extremely sensitive to network latency. Architecting these workloads to better accommodate hybrid cloud environments can alleviate some latency issues.

Data preparation and integration

Data prep isn't just for data pros. What separates data preparation from more traditional data management activities is really nothing more than the level of technical knowledge that is required. While data management is appropriate for database administrators, IT managers, and others who are close to the data in their everyday jobs, data preparation is more approachable for a less technical audience. More marketing managers, sales directors, and financial decision makers are taking an active role in improving the quality of their data, and thus taking more control over the ultimate quality of their decisions using a hybrid cloud environment.

Effective analytics relies heavily on data prep. The old expression "garbage in, garbage out" still applies, but today's organizations are going beyond the basics to improve data quality by empowering business users to perform self-service data prep and by providing new tools on the cloud to perform those tasks. Even the highest quality data, in massive volumes, can lead to tremendous amounts of time wasted in searching, gathering, integrating and preparing data for analysis. Best-in-class organizations are intensely focused on using technologies and activities aimed at enhancing both the quality and the timeliness of their data in order to improve their analytics. To accomplish that requires an environment that allows self-service & collaboration with access to all data from the organization. A hybrid cloud strategy helps organizations to define the environment for their business to improve the quality and the timeliness of the data.

The three phases of achieving hybrid cloud data integration

1. **Exposing private data to SaaS apps:** The first stage in developing a hybrid integration platform is to expose Systems of Record and Systems of Insight data to SaaS applications on the cloud via API management and a secure gateway. For example: a financial institution wants to do data monetization by exposing data and insights to customers and third party vendors from the systems of records and the systems of insights (on private environment) to the systems of engagement (in the public cloud).
2. **Hybrid cloud data lake:** As the volume and variety of data increases, enterprises need to have a data topology strategy based on data tiers, consumption patterns, analytical workloads and data gravity to define the different information zones of the hybrid cloud data lake including landing zone, provisioning zone, shared operational zone, data warehouse & data marts zone, self-service & collaboration zone, and analytics zone. The data topology strategy will define where the data should be stored and processed on the hybrid cloud data lake taking in consideration the type of consumptions, the type of the analytical workloads, the type of data integration among data repositories, the type of data movement required, the type of data sovereignty & compliance required, and the type of data governance and data security required.
3. **Real-time analytics with streaming data:** Businesses today need insight at their fingertips in real time. In order to prosper from the benefits of real-time analytics, they need an infrastructure that can support it. As big data analytics and IoT data processing moves to the cloud, companies require fast, scalable, elastic and secure platforms to transform that data into real-time insights. This is being accomplished with edge-analytics running in the cloud consuming the information from the IoT devices. The information is also sent to a data repository in the hybrid data lake for further discovery & exploration.

Security

Security continues to be the primary concern surrounding public cloud adoption, and these challenges also apply to hybrid cloud. As organizations consider which application can run where, they need to consider the compliance, identity management and data protection needs of those application workloads. Also, organizations need to consider the following security aspects:

- Is a virtual private network required?
- Do data in motion and at rest need to be encrypted?
- Are secure gateways required?
- Is a secure communication protocol (HTTPS) required?
- Is a secure encrypted tunnel (SSH Tunnel) required?
- Is LDAP authentication required?
- Where do the application and data need to be located?

Privacy regulations may limit certain workloads from crossing geographical boundaries. In addition, regulatory requirements such as HIPAA, PCI and SOX require the infrastructure where data resides for a specific application to be compliant.

On the identity management and credential side, operations teams need to ensure user permissions and unique credentials propagate from a private environment to a public cloud. Lastly, assuring the public cloud provider has the basic data protection and cryptographic mechanisms in place and is diligent about updates and patches with minimal disruptions is paramount.

Mind the gap: The biggest potential point of failure for hybrid cloud deployment is where the public cloud and private environment offerings meet. At this point, a gap often exists between in-house security protocols and third-party security standards. If this gap is not closed, malicious actors or malware could slip through it. Meeting this challenge requires a new breed of IT professional, one who is familiar with both the rigors of in-house penetration testing and the more flexible nature of public cloud environments. With the right amount of oversight, it is possible to close this gap and improve hybrid cloud security at its weakest point.

Deal with data: What's the most valuable resource a company owns? Its data. Yet in public and hybrid cloud deployments, the security of data is often overlooked in favor of ease of access and usability considerations. This, of course, leads to an increased security risk. To meet this challenge, companies must define a specific data handling and encryption strategy including encryption key management, encryption for data in motion and at rest, secure gateways, virtual private network, where the data should be located, before deploying cloud services. This eliminates the problem of ad hoc data security —which, by nature, is reactive rather than proactive —and replaces it with reliable, repeatable security protocols that can be applied both cloud-wide and companywide. Furthermore, organizations need an efficient data governance and classification strategy to ensure that data has the correct classification anywhere on the hybrid cloud, and the data access granted to different users is based on the classification of the data.

Get compliant: Compliance is a high-profile buzzword across the tech industry. That's because if cloud deployments are not compliant with industry or government regulations, companies could face monetary fines and sanctions. Meeting this challenge requires a shift in local IT focus away from pure technology management, toward a larger-scale view that's focused on ensuring compliance across the hybrid cloud environments. Companies need to have the required compliance level that an application requires on a specific environment of the hybrid cloud. The compliance level can vary from privacy regulations associated with data sovereignty rules not allowing data from crossing geographical boundaries to regulatory requirements such as HIPAA, PCI and SOX. The different compliance levels are easily enforced with a private environment but it is more difficult to enforce on the public cloud where you don't have control.

Hybrid cloud for big data and analytics

Hybrid cloud requires a new approach for both IT and the business. The goal for an organization as a whole is to extend the private environment investments to the cloud, to modernize some of the private environment applications to the cloud and to provide a **seamless hybrid experience** taking in consideration the different aspects already presented before—what we at IBM refer to as the North Star. The North Star delivers a consistent experience across private environment and public cloud. To help our customers achieve this, IBM delivers a comprehensive strategy for all fundamental areas of the hybrid cloud infrastructure.

Data and analytics: Provide new tools that help users generate new insights with minimal programming with a plug-and-play approach and common services for all kinds of databases (hybrid cloud data lake), thus enabling self-service analytics, multispeed IT, and collaboration among different personas.

Data movement and replication: Provide new options for data movement, replication and sharing that improve network bandwidth and minimize latency, such as dedicated connections.

Data preparation and integration: Deliver new tools with distributed computing that allows push-down of data transformation, integration and analytics processes to the data, considering where the data resides, at rest or in motion, for digital, mobile and IOT. Automation is needed to ingest and persist data and generate metadata at the speed the business needs. Also, an easy integration (plug-and-play) is critical among different cloud components, including storage, computing engines, application runtimes, frameworks, services, applications and APIs.

Data sovereignty and compliance: Implement improved cloud environments that are compliant with all industry and government regulations. Add more cloud environments across the globe to allow better management of data sovereignty constraints.

Data governance and security: Improve security with full managed access, full data protection, full visibility of security risks, and optimized security operations across private environment and public cloud.

High availability and disaster recovery: Enable full HA/DR, backup and archive capabilities for every component that requires them, and on every environment of the hybrid cloud.

Network configuration and latency: Implement an easy network configuration based on software defined networking (SDN) and improved network latency based on location optimization, dedicated connections, cloud cache, network traffic optimization, and workload optimization.

Workload portability: Use containers to make workloads more portable among private environment and different public cloud providers, thus providing a cloud-agnostic application development capability that allows quick and easy deployment.

Conclusion

It's clear that cloud computing is disrupting and transforming the way organizations work in the digital world. However, confusion remains about the true potential of hybrid cloud. Most companies have not yet developed a hybrid cloud strategy, and many are simply using cloud infrastructure to run their existing applications. Also, there are huge gaps when comparing the solutions available on the public cloud with the solutions available on the private environment.

The only way for an organization to succeed with hybrid cloud is to be able to extend the private environment to the cloud by exposing the private environment data to the cloud, by migrating/modernizing some of the private environment applications to the cloud, by integrating some of the workloads & data across the private environment and public cloud and by having a seamless & consistent experience for all workloads across all the different environments of hybrid cloud. To accomplish this, they need to work with a cloud provider that can deliver the comprehensive cloud capabilities included in IBM's North Star. These capabilities make IBM uniquely positioned to be the partner that helps organizations like yours finally make the most of what hybrid cloud has to offer.

For more information

To learn about related IBM offerings, visit **ibm.biz**/BdsPW2.

To learn more about IBM perspectives on IBM technology, visit **ibm.biz**/BdsjbJ.

About the authors

Marcio Moura

Mr. Marcio Moura is Executive Architect in the IBM Analytics Group CTO Office. Marcio builds long-term strategic partnerships with IBM clients in the areas of Big Data & Analytics, Data Lake Strategy, Cognitive Computing and Hybrid Cloud. He has been leading, designing and implementing solutions and transformation programs across several industries including Energy & Utilities, Financial Services, Retail, Telecommunications, and Transportation.

Dr. Qiqing (Christine) Ouyang

Dr. Qiqing (Christine) Ouyang is Distinguished Engineer and Master Inventor in the IBM Analytics Group CTO Office. She manages a global team to build long-term strategic partnership with clients in the areas of Big Data & Analytics, Cognitive Computing and Cloud Computing. She has worked on strategy development for IBM enterprise data lake and hybrid cloud, and has 70 US patents.

References

1 For more information on the integrated systems, see http://blogs.forrester.com/brian_hopkins/15-04-27-systems_of_insight_will_power_digital_business

2 http://www.redbooks.ibm.com/Redbooks.nsf/RedpieceAbstracts/sg248274.html?Open

IBM Corporation
IBM Global Technology Services
Route 100
Somers, NY 10589

Produced in the United States of America
March 2017

IMW14912-USEN-00