Practical Guide to Hybrid Cloud Computing

February, 2016

Contents

Acknowledgements

The *Practical Guide to Hybrid Cloud Computing* is a collaborative effort that brings together diverse customer-focused experiences and perspectives into a single guide for IT and business leaders who are considering cloud adoption. The following participants have provided their expertise and time to this effort: Martin Borrett (IBM), Pierluigi Buratti (IBM), Candice Campbell (Ernst & Young), Mike Edwards (IBM), Maryann Hondo (State Street), Pietro Iannuci (IBM), Gopal Indurkhya (IBM), Heather Kreger (IBM), Petra Kopp (IBM), John McDonald (CloudOne), John Meegan (IBM), Jem Pagan (JNK Securities), Raffaele Pullo (IBM), Gopinath Rajagopal (IBM), John Sanders (CIO Management Group), Karolyn Schalk (IBM), Karl Scott (W. Capra Consulting Group), Alessio Tarenzio (IBM), Pamela Wise-Martinez (Pension Benefit Guaranty Corporation) and Gary Zeien (IBM).

Executive Overview

Hybrid cloud is attractive because it enables cloud service customers to address their business needs by leveraging the wide ranging capabilities of public cloud service providers – in particular, the low cost and leading-edge functionality available – and at the same time using private cloud deployment for more sensitive applications and data. Interlinking cloud-deployed applications and data with traditional non-cloud enterprise applications and data is also an important part of hybrid cloud deployments.

Today, hybrid cloud deployment is commonplace. An organization leverages a combination of private and public cloud deployments depending on its needs for speed of execution, available resources, need for data protection and security, and an array of other reasons. A discussion of how to choose the most effective cloud service and deployment model is included in this guide.

The aim of this guide is to provide a practical reference to help enterprise information technology (IT) managers, business decision makers, application architects and application developers understand the hybrid cloud computing deployment model and how it can be used to solve business challenges rapidly and cost effectively.

As such, this guide focuses primarily on considerations from the perspective of cloud service customers and not cloud service providers. In addition, hybrid deployments involving on-premises private cloud services interoperating with public cloud services are the main focus of the paper. Other hybrid scenarios, such as deployments that include public cloud offerings from different providers, are acknowledged and positioned but they are not discussed in detail.

The first part of the guide defines hybrid cloud computing, explains why this deployment model is essential for addressing business requirements and outlines the key considerations that customers must take into account as they start their transition. The perspective of different customer roles (business, design, developer and operations) are also taken into account since their requirements and corresponding considerations vary significantly.

The "Prescriptive guidance for successful implementation of Hybrid Cloud" section is the heart of the guide. It details both strategic and tactical activities for decision makers implementing hybrid cloud solutions. It covers all the essential technical considerations for hybrid cloud deployment including integration, connectivity, governance, management, security and privacy. It provides specific guidance and best practices for decision makers in each of these areas.

What is Hybrid Cloud Computing?

Cloud computing in general is well described in the CSCC *Practical Guide to Cloud Computing*, *V2* [1]. The alternative deployment models of public cloud and private cloud are discussed, with the considerations that apply to using each of them. Hybrid cloud computing is a deployment model which involves combining the use of multiple cloud services across different deployment models – in particular, combining the use of public cloud services with private cloud services.

The ISO 17788 *Cloud Computing Overview and Vocabulary* standard [2] defines hybrid cloud as "a cloud deployment model using at least two different cloud deployment models" – where the potential deployment models are public cloud, private cloud and community cloud.

Based on this definition, there are many combinations of cloud resources that can be leveraged in a hybrid cloud deployment. These combinations can also involve a mix of different cloud service models, Infrastructure as a Service (IaaS), Platform as a Service (PaaS) and Software as a Service (SaaS). For example, a hybrid cloud deployment could combine an on-premises private IaaS cloud service with a publicly hosted SaaS application. It is also important to recognize that private cloud services could be hosted on-premises or alternatively hosted off-premises in a dedicated part of a cloud service provider's data center.

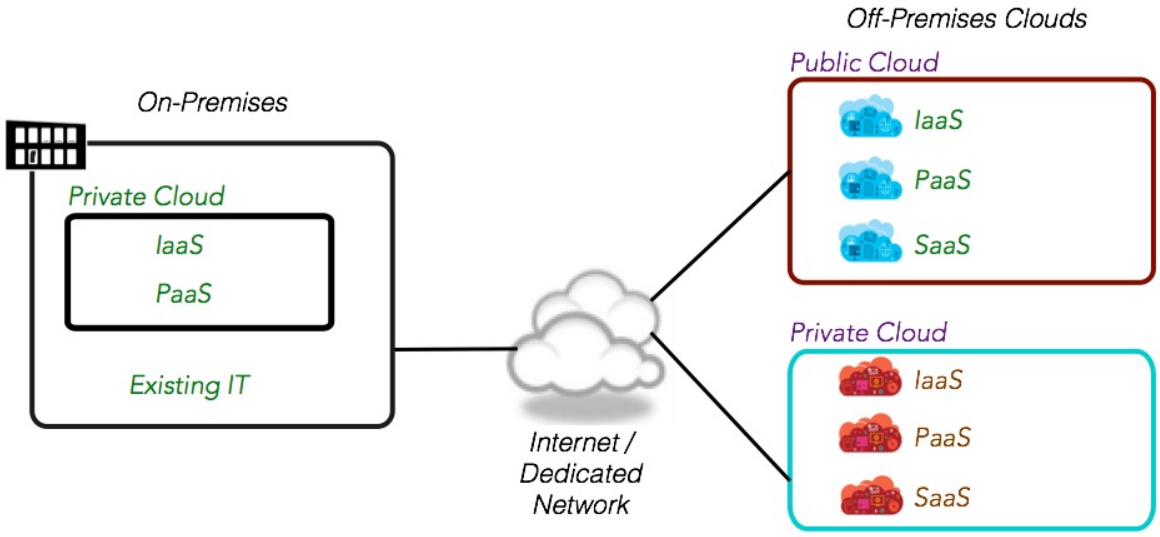

Figure 1: Hybrid Cloud provides the enabling capabilities to bridge environments, layers, and resources such that it is seamlessly automated

Most enterprises are not "born on the cloud," therefore most cloud resources typically need to be connected to significant on-premises IT systems. For these enterprises, the most common hybrid cloud system architecture is where on-premises resources (private cloud and non-cloud) are combined with off-premises cloud resources, providing the business with a variety of new and innovative capabilities. This paper focuses on this combination.

Here are key considerations that need to be taken into account:

- *Integration*: Hybrid cloud computing is about aggregation and integration of capabilities and services from cloud service providers with on-premises resources, leveraging the best-of-breed. The emergence and evolution of core integration capabilities, such as Service Oriented

Architecture (SOA), Representational State Transfer (REST) Application Programing Interfaces (APIs), and cloud management and orchestration frameworks have opened up new options for integrating cloud services. In addition, new capabilities from a wide variety of services (on-premises and off-premises) can now be aggregated to provide the business with a broader set of capabilities on which to act.

- *Composition*: Hybrid cloud is about combining services and capabilities in a way that supports desired business objectives with agility and within budget, while at the same time dealing with risks and regulations in a satisfactory way. A hybrid cloud also enables businesses to utilize these services in varied durations and usage models. A set of services (either IT focused or business focused) could be used over a long time based on predetermined business needs, or they could be used for a limited period based on specific business events.

- *Organizational Impact*: People in many different roles are impacted by hybrid cloud in a variety of ways. The value of hybrid cloud is highly dependent on what is important to your organization.

Hybrid cloud means different things to different people

Different roles in an organization are typically concerned with aspects of enterprise IT, which at high level are either a focus on "using it", "building it" or "running it."

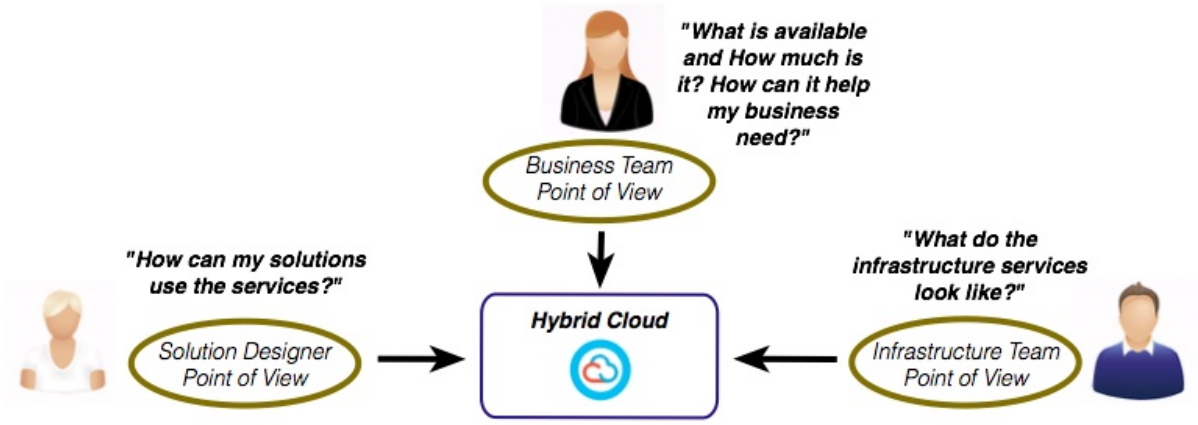

Figure 2: Cloud Service Customer Points of View for hybrid cloud computing

Solution Designer Perspective: "Give me more options"

A solution designer typically looks to hybrid cloud for options and capabilities that are not available in an on-premises or off-premises only solution.

For designers and developers some of the critical success factors related to hybrid cloud include:

- *Flexibility*. Repartition/deploy elements of a solution based on changing technology services or improved capabilities.
- *Speed*. Quickly leverage new capabilities in their solutions no matter where they are available in the hybrid cloud.
- *Capacity*. Quickly leverage the capacity of off-premises clouds to provide peak support for on-

premises systems.
- *Consistency*. Support continuous delivery of applications across the hybrid cloud leveraging common tools and processes.
- *Agility*. Design and develop solutions in such a manner that where they are deployed across a hybrid cloud can be adjusted in a seamless manner.

Infrastructure Perspective: "More things to manage"

Another perspective is that of the infrastructure or service provider team. This team typically looks at hybrid cloud based on SLAs, processes, and resources (compute, storage and networking).

For the infrastructure team some of the critical success factors related to hybrid cloud include the ability to:

- Have visibility into all resources and services being utilized and managed across the hybrid cloud.
- Monitor all infrastructure resources in a consolidated and federated manner following current processes.
- Access all resources on the hybrid cloud to perform activities as required (e.g., console login).
- Manage the deployed infrastructure to perform patching, auditing, and security management as required by the organization's SLAs.
- Enforce consolidated security across the hybrid cloud, including consolidated audit logs.
- Control provisioning of workloads based on organizational policies.

Business Team Perspective: "What is available? How much is it? What value does it provide?"

For business teams, some of the critical success factors related to hybrid cloud include the ability to:

- Have a consolidated view of what is running where and what the financial implications are by organization/department.
- Understand the SLAs of a solution as a whole based on different providers.
- Have consolidated invoices that can integrate with the organization's financial systems.
- Access the systems running on the hybrid cloud without regard for where they are running. The business users should not know or care.
- Perform financial "what if" modeling leveraging the services of different cloud providers.
- Reduce costs and increase business flexibility of leveraging hybrid cloud.
- Understand software licensing exposure across the cloud providers being utilized.

As you can see, these roles can look at hybrid cloud very differently. This will come into focus more when we look at the value, the scenarios, the considerations and the challenges related to hybrid cloud.

Why Hybrid Cloud is important to the business

In terms of IT, businesses today operate at multiple speeds. To meet the challenges of competition and the rapidly evolving marketplace, businesses need to be agile and innovative, particularly for mobile and web applications used by customers. At the same time, the stable processes and systems that keep the business running cannot be replaced easily and instead evolve at a slower pace. This is *two speed IT*. One part, the *steady speed*, delivers enterprise-strength IT services, and the other part, *fast speed*, enables exploitation of new digital business opportunities. However, success comes through optimizing support for both types of IT. Hybrid cloud can help provide a solution to this need.

Three key value points of hybrid cloud computing are:

1. ***Enables a highly cost-effective, rapidly responsive and elastic IT, better aligned with the business needs in order to support two speed IT.***

   A hybrid cloud can be designed for dynamic consumption, interconnection, orchestration and control of all types of cloud services. Hybrid cloud enables a more responsive and elastic IT that is able to quickly respond to the demands and needs of both steady speed *systems of record* (SoR) and new fast speed *systems of engagement* (SoE). Workloads are enabled through automation and programmable IT through open infrastructure APIs and composable infrastructure services.

2. ***Provides a portfolio of business and IT services that leverage the best capabilities of cloud service providers, enabling flexibility in what can be built and where it can be deployed.***

   Using cloud services provides new tools and data for innovation. The business is no longer constrained by what they have available on-premises. In essence, the "catalog of services" that they have at their disposal increases dramatically.

3. ***Enables the business to innovate faster while leveraging existing systems and capabilities.***

   Simply put it comes down to time and money. The ability to compose services from both on-premises and off-premises (i.e., hybrid) is a key enabler to increase speed to market and to reduce costs.

Key Considerations for Hybrid Cloud Computing

Given an understanding of what a hybrid cloud is, the business value it provides, and use cases for how it can be leveraged is the first step on the journey to hybrid cloud computing. Some of the benefits of hybrid cloud (e.g., capacity, flexibility, elasticity, service portfolio and resource cost) have to be weighed against a core set of considerations. The next step is to develop an understanding of the key considerations faced in implementing a hybrid cloud in your organization. This section provides a summary of what these are and why they are important. Here are the key considerations:

- *How to determine the placement of solution components.*
  What should go where and how should it be designed? Should a solution only include the usage of private dedicated or local cloud resources? How do you leverage available public cloud services?

- *How to integrate with existing enterprise systems.*
  How can existing business applications along with existing management and monitoring systems be leveraged? How will the internetworking be accomplished?

- *How to handle an increase of management complexit*y.
  How do I manage resources running in different cloud services, particularly public cloud services? The lifecycle of resources involved in supporting key business operations can be

transient (short lived and/or move automatically). How is that handled by existing management and reporting systems?

- *How to ensure that security is considered in all aspects of the hybrid cloud.*
  How do I ensure that both on-premises AND off-premises cloud environments are secure? How about the data stored off-premises? How about the data in transit?

- *How to deal with rapidly evolving and partially mature technologies.*
  The speed of capabilities being deployed and or changed in a hybrid cloud changes at a different speed for *fast speed* resources vs. *steady speed*. How does the organization support this?

- *How to implement common operational services such as backup and disaster recovery in a hybrid cloud.*
  Given that there are multiple providers in play, how can the different backup/recovery solution and networking options enable a seamless environment that meets an organization's SLAs?

- *How to ensure adherence to regulatory and compliance requirements.*
  How will you ensure data placement, data encryption, personal information protection, contractual management (e.g., software licensing) adhere to the appropriate regulations?

Prescriptive guidance for successful implementation of Hybrid Cloud

This section provides a prescriptive series of steps that should be taken to ensure successful hybrid cloud computing deployment from the perspective of a cloud service customer. It takes into account differences that result based on the size of the organization and its IT maturity level. The following steps are discussed in detail:

- Determine cloud deployment model for applications and data
- Integrate with existing enterprise services
- Address connectivity requirements
- Develop governance policies and service agreements
- Assess and resolve security and privacy challenges
- Manage the hybrid cloud environment
- Consider a backup, archive and disaster recovery plan

Depending on the maturity of the organization and the level of adoption of cloud computing, the entry point will change for each new service being evaluated.

For the following sections, it is useful to keep in mind a picture of the components and their interconnections that are involved in hybrid cloud computing, as shown in Figure 3. This divides the components of concern to the cloud service customer into three groups – the customer's own **in-house systems**, components running in **private cloud services** (whether on-premises or off-premises) and components running in **public cloud services**. Figure 3 also highlights the three types of interfaces that exist between these groups – namely, the **functional interfaces** of the applications and services, the **administration interfaces** used to manage and control the applications and services (this includes

security administration), and the **business interfaces** used to manage subscriptions, billing and payments.

The customer's in-house systems are divided into in-house applications, management applications and security systems, plus in-house datasets and databases. For both private cloud services and for public cloud services, the major components described include the app code (for applications running within the cloud service) with its application environment, cloud service customer data and cloud service derived data (which includes logs) and various security components.

The major concern of hybrid cloud deployment is to ensure effective and efficient integration of all the components in the three groups in Figure 3, taking into account the interfaces between them.

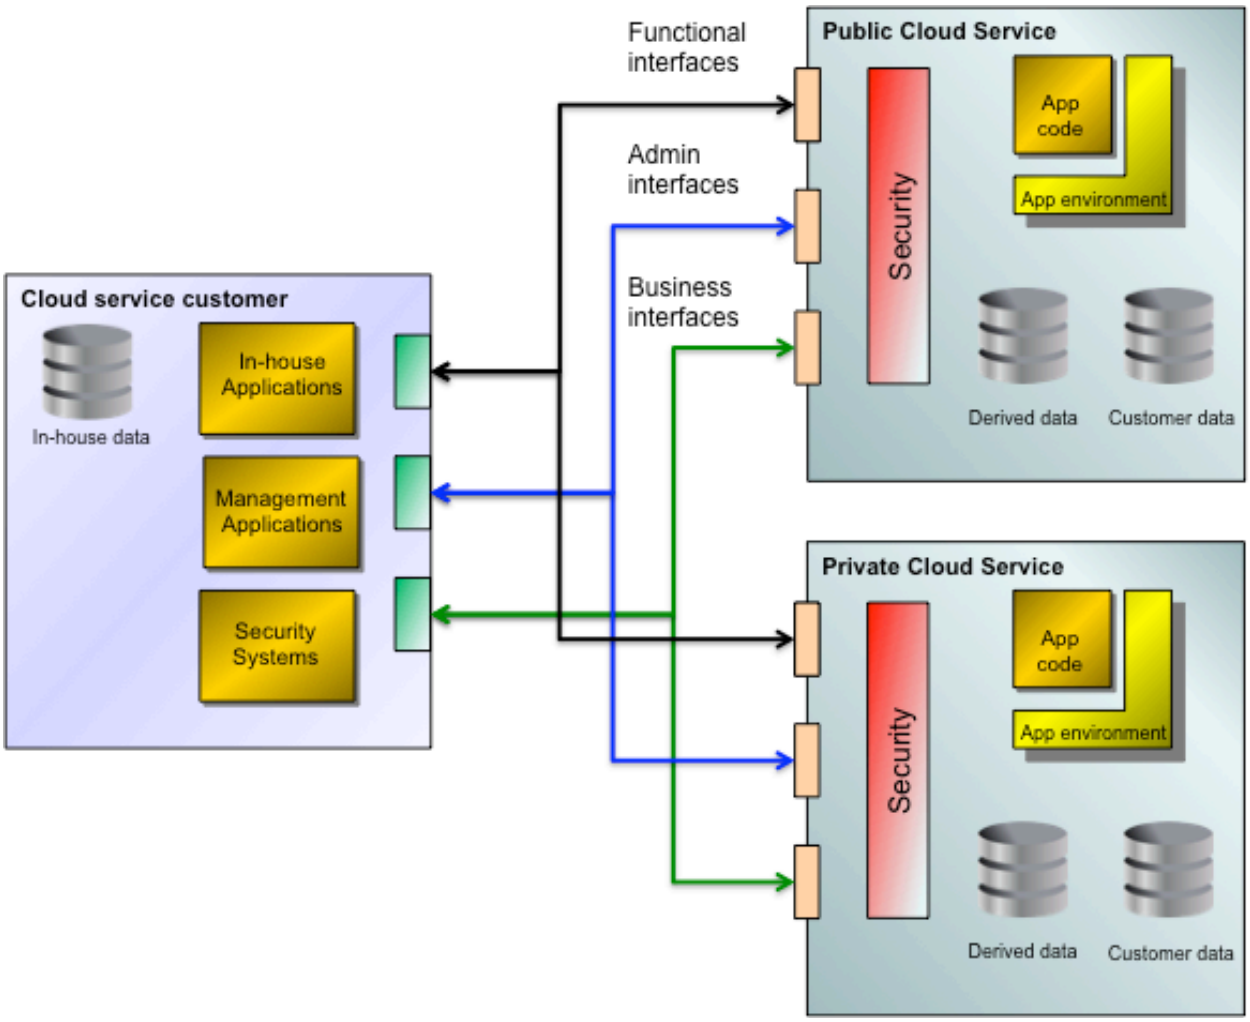

Figure 3: Components and connections for Hybrid Cloud Computing

Determine Cloud Deployment Model for Applications and Data

Defining the right application and data deployment model is a critical step in the journey to use hybrid cloud. The right deployment model positions IT to align with business needs and enable achievement of objectives. The cloud deployment model must address the following:

- What cloud resources should be deployed
- What applications, services and data should go where

The two speeds of IT is a major consideration when defining which cloud resources to deploy. In the *steady speed,* IT must continue to support the traditional enterprise applications to keep the business running. This implies that current on-premises services must be maintained. Additionally, off-premises services, such as private dedicated IaaS, can be used to increase infrastructure flexibility for enterprise services.

The *fast speed* requires IT to deliver applications and services more quickly whether it's scaling existing services to satisfy spikes in demand or providing new applications quickly to meet an immediate business need. The use of public cloud services may be required to satisfy these requirements.

After sorting out the cloud resources to use, the next step is determining where applications and data should go. Here are the four basic options for application placement in a hybrid cloud environment:

1. Place the application and its datasets into a public cloud environment connecting to existing enterprise applications and datasets on-premises as required.
2. Place the application and its datasets into a private cloud environment connecting to existing enterprise applications and datasets on-premises as required.
3. Place the application on-premises (or into a private cloud environment) and link to public cloud services as required to obtain new or specialized capabilities.
4. Place some components of the application and its datasets into a public cloud environment while placing other components into a private cloud service and/or an on-premises non-cloud environment, linking them together as required.

Deployment decisions must consider these four options. IT architects should assess the right application architecture to achieve maximum benefit. This includes understanding application workload characteristics and determining the deployment model for multi-tier applications. For example, understanding where heavy processing is performed and the interaction between presentation, business logic, integration, and data layers is key to making the right application deployment decisions. This has to be understood to ensure acceptable application response times are maintained in the hybrid model.

Cloud service brokerage is emerging as an important component for delivering consistent and accurate placement and prioritization planning for applications in hybrid cloud environments. Cloud service brokers function as intermediaries between cloud providers and enterprises and help strengthen the relationships providers have with customers by offering planning, integration and management services. They provide tools to govern and control environments across private and public clouds covering cloud spend, security and resiliency aspects from the application infrastructure design through the final cloud managed services that are actually provisioned and utilized.

Regarding the related question of which *cloud service model* to use for the application:

- *IaaS*: well suited to cases where the application and its required software stack already exist and are well understood by the enterprise - move the whole set of software into one or more virtual machines with associated storage resources.

- *PaaS*: well suited to applications built for runtimes and services supported by a complete middleware platform, especially where the enterprise wants to be relieved of the burden of maintaining and operating complex software stacks if required by the application. Also particularly suited to creating new custom applications with minimum effort and risk.

- *SaaS*: well suited to cases where the enterprise wants to avoid the costs and risks of developing custom applications, where standard off-the-shelf cloud services provide the required business capabilities.

The summary below represents the high-level steps that need to be taken to rationalize application and data deployment. These steps should be incorporated in the IT planning process. This is recommended to ensure the model is connected to where the business and IT is going.

1. Identify the business objectives and strategic imperatives

2. Align with the IT strategy and future state architecture to achieve it

3. Develop application and data inventory

4. Map business processes to applications

5. Identify disposition of applications – retire, consolidate, maintain, upgrade/enhance

6. Define deployment decision criteria and apply to the retained (consolidate, maintain, upgrade/enhance) applications

7. Document deployment model

A key step in the process is applying decision criteria to determine the deployment of applications (step 6 above). Flexibility, security, speed, cost, control, locality, and service levels represent the criteria used when deciding between public and private (on and off-premises) cloud deployment. The team performing the analysis must have a strong understanding of application business and technology requirements in order to make informed deployment decisions.

- *Flexibility* – Elasticity requirements must be considered when assessing public and private cloud usage. Seasonal demand for services, such as Black Friday ecommerce activity, drives a significant increase in capacity requirements for a short period of time. Public cloud services are preferable when elasticity is a major concern as purchasing on-premises capacity to support this kind of demand is not cost effective. Mature organizations provide “bursting” to public cloud services when privately hosted services exceed defined thresholds.

- *Security* – Security and privacy are major considerations when evaluating application and data deployment and the data being used. IT should understand compliance requirements and establish cloud deployment policies. Traditionally, less sensitive data such as publicly available information have been candidates for the public cloud while highly sensitive information may be better hosted on-premises or in a private cloud service. Security in public cloud services continues to mature and is improving, potentially enabling the hosting of highly sensitive information in such services to be considered. The organization has to assess the available controls in the public cloud service and architect security services to deal with the risks associated including those related to multi-tenancy.

- *Speed & Automation* – The ability to quickly deploy applications and satisfy business needs is a major consideration when making deployment decisions. The immediate availability of resources and capabilities in public cloud services makes this model desirable when fast deployment is required. Additionally, existing public cloud services can be leveraged to reduce the time to build new capabilities for the business.

- *Cost* – The cost associated with hosting the application over the term of usage is a key factor when considering application and data deployment. Thorough understanding of private cloud services costs (e.g., monthly server cost) and understanding public cloud variable costs, such as outbound internet bandwidth usage, is critical to establishing an apples-to-apples comparison. The costs of developing and maintaining application capabilities in-house also need to be compared with the alternative of using off-the-shelf cloud services with an ongoing subscription model.

- *Locality* – The location of application and data must be considered when deploying services to public and private cloud. Locality considers latency of services to end users (includes application and data integration) to ensure a great user experience. For example, cross border latency can have a significant impact on response times. Hosting services in country for the user community will improve the experience. Additionally, data sovereignty concerns may encourage hosting of services in a public cloud where a data center presence does not exist.

- *Service Levels* – Service levels associated with mission critical applications are a major consideration when assessing application deployment - including availability, response times, throughput, etc. Mapping service levels associated with public cloud services and private cloud services to the requirements of applications helps to identify any mismatches. For example, availability of public and private cloud services must be understood, along with the options to architect around known limitations.

- *System Interdependencies* – How dependent is a system on other systems in your organization? This is key since you may have a workload that meets all the criteria above, however, if it has multiple integration points that cross the hybrid divide, it may not be suitable for a hybrid cloud.

Thoughtful planning is required to develop a hybrid cloud deployment model that enables business to innovate. Identifying the right resource model and working through deployment decisions are key steps

in aligning the model to suit your business. It's important to get this right as course correction can be disruptive.

Integrate with Existing Enterprise Systems

As enterprises consider their hybrid cloud computing strategy, they will inevitably face the challenge of addressing how they will leverage their existing in-house IT investment with their newly adopted cloud services. As new cloud services are deployed, the need to connect them or the applications running on them with various on-premises applications and systems becomes important (as depicted in Figure 4).

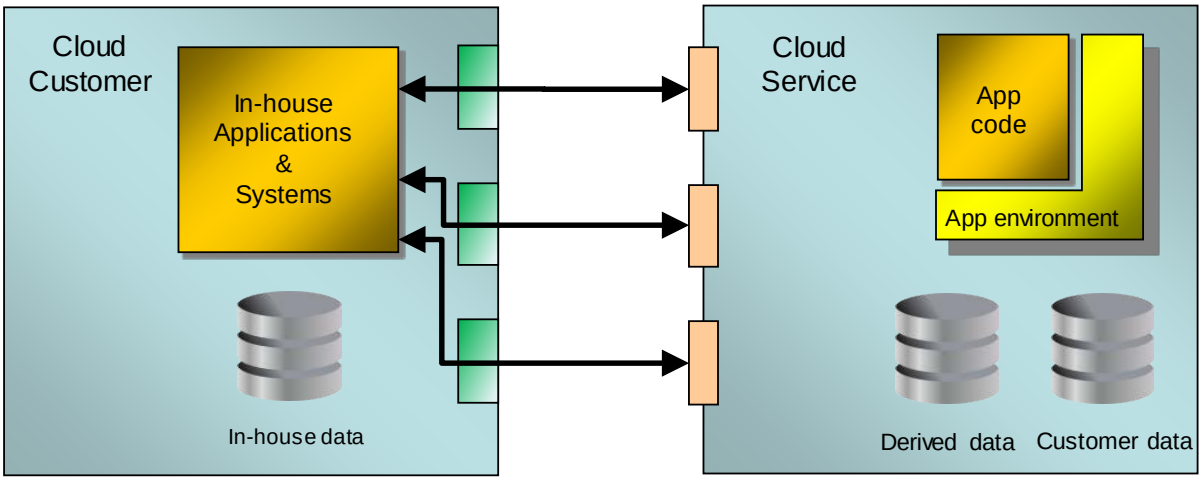

Figure 4: Linking In-house capabilities to Public Cloud Services

Cloud service customers need to understand the impact of these connections and address it. Integration between applications is typically classified into three types:

- Process (or control) integration, where an application invokes another one in order to execute a certain workflow
- Data integration, where applications share common data, or one application's output becomes another application's input
- Presentation integration, where multiple applications present their results simultaneously to a user through a dashboard or mashup

The interoperability aspects of a cloud service mainly relate to these three connections or interfaces between the customer and the cloud service. It is also important to understand that the interoperability of the three types of integrations may be independent of each other. The purpose of these integrations may be to perform an end-to-end workflow that crosses the boundaries between multiple business capabilities or systems (for example, entering a transaction in an accounts receivable system when a customer places an order in an e-commerce application). Another form of integration is when the cloud service must be monitored and managed by an existing suite of on-premises IT tools.

To ensure that integration for hybrid cloud environments goes smoothly, customers should take the following actions:

- Insist that the cloud service provider supports key open technologies (open standards and/or open source) for admin and business interfaces.
- Ensure that on-premises applications are leveraging SOA design principles and can utilize and expose APIs to enable interoperability with private or public cloud services.
- Examine whether existing in-house systems are available to deal with the business aspects of using cloud services (subscription information, billing and invoicing). If they are not available, consider installing new systems to cover these aspects; if they are available, consider how those systems can connect to the business capabilities of the cloud service(s).
- Consider implementing an Enterprise Service Bus (ESB) to perform interface, protocol and data transformations to address differences between on-premises systems and cloud services.
- Leverage the support of third party ID and Access Management functionality to authenticate and authorize access to cloud services.
- If cloud service(s) need access to on-premises APIs or data, address the security issues raised by enabling access to these capabilities from the cloud environment - for example, put in place suitable API Management capabilities to prevent unauthorized access.

In many cases, the challenge is not so much "integration" as it is "re-integration" or "maintaining the integration" between pieces of the entire system that are coupled in a certain way. Customers must identify integration points of existing systems and applications to ensure functionality is not lost once their hybrid cloud solution is deployed.

For example, the main capabilities of the application may be presented in the functional interfaces to end users as a web browser application or as a mobile application. However, the same capabilities may also be made available as an application program interface (API) for consumption by custom applications written or purchased by the customer and running on the customer's systems. In the cloud service environment, APIs are typically defined by a programmatic interface based on a common protocol such as REST/JSON or SOAP.

The use of application program interfaces (APIs) in hybrid cloud environments is a good mechanism for managing internal and external business services and delivering connectivity between existing software systems and those hosted in cloud services. APIs make capabilities of an application or service available for use by other programs through a defined and controlled interface and protocol (e.g., REST over HTTP is commonly used for APIs). API management technologies are emerging as a necessary set of capabilities required to support the use of well defined, controlled and secure APIs.

Refer to the CSCC whitepaper, *Interoperability and Portability for Cloud Computing: A Guide* [3], for a more detailed discussion on integration considerations for connecting cloud services with on-premises systems.

Address Connectivity Requirements

Connectivity in the hybrid cloud is a key domain for successful deployments. Not only must it provide communication between components, it must also support the hybrid cloud solution by:

- Meeting requested service levels (e.g., user experience);

- Adhering to the security policies and
- Enabling the overall IT management strategy

There are several opportunities that arise when dealing with network connectivity and hybrid cloud solutions. There are considerations and decisions to make that are deeply tied to solution deployment that also depend on the features and flexibility that the cloud service providers offer. The interconnections shown in Figure 3 can help customers understand the various elements that need to be connected in a hybrid cloud deployment.

Successful deployments rely on, and often start with, a sound knowledge of the workload. When using a hybrid cloud deployment, workloads may involve interactions with in-house systems, remote users, and a variety of cloud services. Therefore, it is essential to know the characteristics of all of the different communication paths that comprise the solution, including maximum latency, bandwidth or packet loss, which is critical for expected service levels.

This section highlights examples of where hybrid cloud solutions can simplify the network design and operational phases and under what circumstances connectivity requires special focus to meet expected results.

- *Network Links* - Depending on aspects like bandwidth, latency, service levels and costs, the physical connectivity over the wide area network (WAN) can vary in terms of technology. Hybrid cloud solutions can rely on point-to-point links as well as the Internet to connect data centers (on-premises and cloud provider). The selection of the connectivity type depends on the analysis of aspects like performance and availability. Some hybrid cloud workloads may not tolerate the variations in throughput of the Internet. For these workloads, a point-to-point connection between the cloud provider and the enterprise data center can be deployed.

- *Network Virtualization* - The virtualization strategy that the enterprise may have adopted on-premises can be an enabler for the connectivity of services in the hybrid cloud. Virtual switching may or may not be available from the cloud service provider and this is a key factor to consider when designing the solution. Software defined networks can address challenges like managing the IP address plan across sites together with the virtualization of network functions like firewalls and load balancers that may be deployed in their virtual forms in the cloud using the preferred vendor.

- *Security* - Security in the connectivity domain needs to be evaluated and understood; the cloud provider network security standards and approach may need to match the overall network security policies, guidelines and compliance. For example, the choice of encrypting and authenticating traffic on the WAN can be evaluated against encryption at the application level. The decision goes beyond networking connectivity and spans different domains like systems (for the computing resources) and applications. Technologies such as VPNs can be employed to provide secure connections between components running in different environments.

- *Internet Infrastructure* - When the Internet is selected to link the different data centers (on-premises or between cloud service providers), cloud service providers may leverage multiple carriers with diversified ingress/egress paths. This leads to an increased level of availability without the complication of having to manage multiple contracts with different suppliers, which may be based on time commitments that do not match the temporary needs of the user.

- *IP Address Plan* - Cloud service providers may provide a lower degree of flexibility in the choice for ranges of IP addresses to assign to resources in cloud services. With this constraint, overlaps with the enterprise network in the customer data center may occur and must be addressed accordingly. A typical approach is to use Network Address Translation (NAT) techniques.

- *End to End Management* - The management system that is deployed on-premises may need to be adapted to work with cloud services, not only for the monitoring, provisioning and de-provisioning of cloud resources, but also to facilitate problem identification activities that may span multiple systems that have different governance boundaries.

- *Availability of Distributed Denial of Service (DDoS) Protection Services* - The availability of cloud services to protect from DDoS attacks or provide services like Domain Name Services (DNS) enables the customer to concentrate on the business applications rather than dedicate resources to acquire, set-up and maintain such services. This is especially applicable to the hybrid cloud for workloads that have components deployed into a cloud service that will be exposed to the Internet.

- *Service Continuity* - Where a multi-site high availability and disaster recovery strategy is used, connectivity plays a pivotal role in meeting both the recovery time objective and the recovery point objective.

In summary, connectivity supports the deployment of solutions in the hybrid cloud with techniques, tools and processes that are known, and can also leverage new paradigms like software defined network to provide a broader set of resources that address newer challenges and different priorities in the requirements.

Develop Governance Policies and Service Agreements

The cloud services combined to create a hybrid cloud computing environment are likely already operating under a governance framework and service agreements. The governance planning process for the hybrid cloud environment involves a gap analysis and the harmonizing of existing policies to assure that the use of the cloud services is governed effectively.

Governance for cloud computing refers to the system by which the provision and use of cloud services are directed and controlled. [2]

Internal cloud governance is used for the application of design-time and run-time policies to ensure that cloud services are designed, implemented and delivered according to specified expectations.

External cloud governance is used to form an agreement between the cloud service customer and the cloud service provider concerning the use of cloud services by the cloud service customer.

It is important to ensure that cloud governance practices mesh with SLAs and other contractual elements of the cloud service agreement. The individual governance practices used by cloud service customers and cloud service providers exist on a continuum from simple to sophisticated and are encapsulated within their role. It is the responsibility of each role to implement governance according to their needs. [4]

Governance of hybrid cloud computing is an extension of governing cloud services. Cloud service customers must identify and work with the right stakeholders from both business and IT to iteratively develop and maintain appropriate policies, service agreements, and service management.

Who to include when developing policies and agreements

The ISO/IEC document 17789 *Information technology — Cloud computing — Reference architecture [4]* provides a detailed description of roles and responsibilities that can assist in identifying domain owners. The functional role descriptions help customers map the correct personnel in their organization to the roles itemized in 17789 that are typically part of the governance planning process:

- Internal operations management
- Cloud services development
- Cloud broker role (business relationship)
- CSP operations representation
- Product Management or Line of Business representation
- Financial analyst
- Risk/compliance
- Legal/contract

Hybrid cloud environments require a governance model that encompasses all of the environments used in the deployment. The extent to which a cloud service provider participates in governance related activities, such as change communication, can be used as a differentiator when choosing between CSPs for particular workloads.

Assess existing compliance and governance frameworks

Existing compliance and governance frameworks may not take into account the added complexity of hybrid cloud environments including new touch points between internal and external systems and their geographical location. When cloud computing is used to extend an on-premises system, it is likely that the governance in current use does not cover the cloud components of the hybrid architecture.

In general, special consideration should be given to what is done at or above the Service Responsibility Line (SRL) illustrated in Figure 4. Typically, the cloud service provider offers management services at or below the line as part of the standard cloud service offering, which varies between different cloud service models. Meanwhile, the cloud service customer is usually responsible for the elements that are above the SRL. Governance and service levels for the hybrid cloud computing environment must define the hand-offs, overlaps and gaps between systems, service providers, locations, etc.

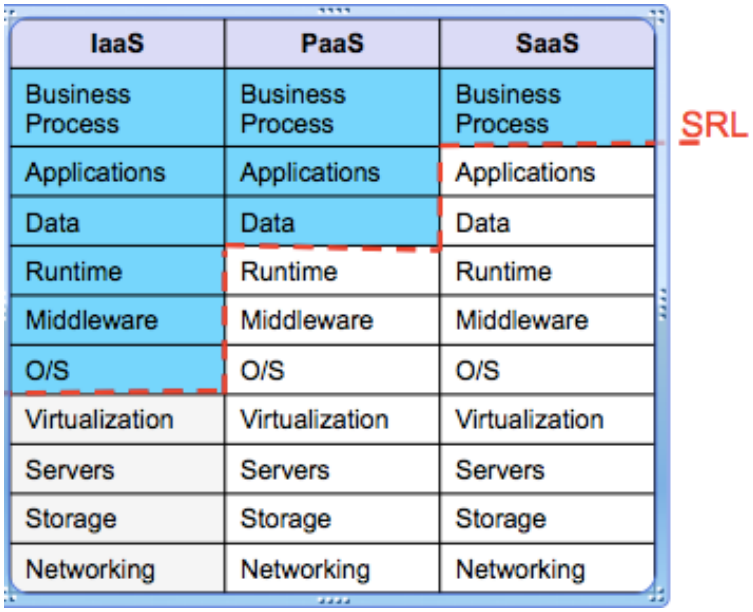

Figure 4: Service Responsibility Line for Different Cloud Service Models

Change Management and Change Communications

Governance of hybrid cloud computing environments requires the coordination of internal cloud governance and one or many external cloud governance agreements and relationships. Given the nature of hybrid cloud computing, the governance surrounding traditional on-premises systems and applications needs to be reviewed to ensure that there are no gaps in governance or service expectation.

Special attention needs to be paid to change management and communication. The hybrid cloud computing architecture could include multiple cloud service providers. The scope of what needs to be managed and who needs to be engaged is larger compared to environments based solely on the use of a single cloud service. For change management there are some areas that need more oversight, typically those that are associated with automation, self-service or backup and disaster recovery. These include:

1) Service catalogs - Organizations that create front-ends for the consumption of external cloud services specific to their users need to assure that content describing offerings, linkage to policy and other information and guidance are maintained consistently.
2) Access controls - The transformation and connectivity between cloud services, whether private or public, rely on the secure passing of credentials and consistent availability. Changes here could shut down communication between application components.
3) Data storage - The expansion of physical and logical data placement within the hybrid environment requires special attention to geographies, datacenter compliance, and systems access when planning for backup and disaster recovery.

Service Level Agreements

Governance and service agreements for hybrid cloud computing environments need to take into account multiple communication touch points, change management cycles, responsibility hand-offs, and

geographies. The overarching governance for the specific hybrid environment is impacted by each included service and system, which may have its own architecture and governance model that needs to be considered when planning the whole. While it is largely automation that unifies the elements of the hybrid environment, governance and service agreements need to be defined for the automation.

Hybrid cloud environments also create new challenges in resolving disputes or responding to a security breach. The multiple integration and touch points between hybrid components mean it may not be clear what or who is responsible for an outage or incident. Contracts may limit resolution of customer disputes to a pre-selected arbitrator. Decisions on what workload may be run on what portion of a hybrid cloud needs to consider the quality of service, the remedies for outages and incidents, as well as the more usual security and performance concerns. Without clear agreements across each area of what source of metrics is authoritative, cloud service providers might get into a defensive posture when there are problems with a cloud service where a cloud service customer uses outputs from their own service monitoring tools to identify the problems.

Visibility into external cloud services can be a challenge in the hybrid cloud environment. The hybrid cloud model can involve multiple components, provided by different vendors that form a single service for consumption by the end user. Each element of automation and orchestration for the hybrid cloud environment needs to operate consistently and according to pre-defined expectations. The self-service functionality that is so valued in the cloud computing model often creates higher expectations of availability and communication. This is due to the increased visibility that the Lines of Business might have to what was previously IT only operations.

The *Practical Guide to Cloud Service Agreements V2.0* [5] lays out a series of steps for evaluating cloud service agreements and provides guidance for each step.

Identify gaps in measurement and management visibility

Each cloud service provider has some way of monitoring and reporting their cloud service offering. These capabilities are often proprietary to the specific provider. Cloud service customers establishing a hybrid cloud solution need a way to look across providers and components. The ideal method is a management tool or suite that offers “a single pane” view and aggregates logs and other operational data for metrics and reporting. Transaction monitoring becomes more complicated and for highly regulated workloads, such as a securities trading application, the ability to get the requisite level of detail across geographies and providers becomes a key to decisions on where components reside in the hybrid environment.

For specific compliance requirements, customers should include specific reporting needs in the overall governance scheme so that audit response is simplified.

Assess and Resolve Security and Privacy Challenges

Security and privacy are key concerns for any customer using cloud computing, whatever the deployment model. The CSCC white paper *Security for Cloud Computing: 10 Steps to Ensure Success* [6] describes the general concerns relating to security and privacy when using cloud computing. The risks and approaches described in that white paper apply generally to hybrid cloud. Additional considerations

that arise from the multiple cloud computing environments involved in hybrid cloud are discussed in this section.

Hybrid cloud computing involves one or more private cloud services and one or more public cloud services, plus some customer in-house applications and systems, as shown in Figure 3. Applications and services may exist in each of these environments, along with data, for example, in the form of files or databases. Connections exist between the different environments, which can be functional, administrative and business related.

The additional challenges for the cloud service customer in handling security and privacy for hybrid cloud deployment, over and above those which apply to other forms of cloud service deployment, mainly concern the interfaces between the different environments, the movement of applications and data between the environments and the organized control of assets across these environments. Security needs to be applied consistently across this whole *system.*

A discussion of the interfaces and the movement of applications and data involved with cloud computing environments is covered in the CSCC white paper *Interoperability and Portability for Cloud Computing: A Guide* [3]. Migrating applications to cloud environments is discussed in the CSCC white paper *Migrating Applications to Public Cloud Services: Roadmap for Success* [7].

For hybrid cloud deployments, the interfaces between components running in the different environments and the spread of data between those environments are the sources of additional risk that must be addressed with suitable controls. This can include some loss of control over assets and data placed into a cloud provider's systems. Despite this inherent loss of control, cloud service customers still need to take responsibility for their use of cloud computing services in order to maintain situational awareness, weigh alternatives, set priorities, and effect changes in security and privacy that are in the best interest of the organization.

Hybrid cloud deployment requires the cloud service customer to ensure close and regular contact with each cloud service provider and to be aware of the cloud service agreement that applies to each cloud service to ensure the alignment of the relevant provisions for security and privacy. In hybrid cloud, controls and provisions need to be harmonized across all the systems used by the cloud service customer. In particular, the protection of personal data stored and processed on the provider's systems must be ensured. Cloud service customers must also ensure appropriate integration of cloud services with their own systems for managing security and privacy.

Taking the list of security risks associated with cloud computing presented in the CSCC security white paper, here are specific risks that relate to hybrid cloud deployment:

- *Governance*. In a hybrid cloud deployment, customers must ensure that appropriate governance processes apply to the public cloud services, as well as private cloud services and in-house systems.

- *Responsibility ambiguity*. Made more complex where applications and data move between

private and public cloud services and where different services are aggregated.

- *Multiple interfaces.* Dealing with cloud services from multiple providers compounds the risks since it is likely that each will have its own set of security and privacy characteristics.

- *Authentication and Authorization*. A hybrid environment could mean that gaining access to the public cloud environment could lead to access to the on-premises cloud environment.

- *Isolation failure*. Multi-tenancy and shared resources are defining characteristics of public cloud computing. Public cloud services could provide a back-door to the customer's other computing systems when they are integrated in a hybrid deployment.

- *Compliance and legal risks*. The cloud service customer may have compliance for in-house and private cloud systems, but this may be compromised by inappropriate integration with public cloud services – each public cloud service must be checked for its compliance status.

- *Handling of security incidents*. The detection, reporting and subsequent management of security breaches needs to be shared between the customer and the cloud provider. Notification rules need to be in the cloud service agreement so that customers are not caught unaware or are informed with an unacceptable delay.

- *Management interface vulnerability*. Interfaces to manage public cloud resources (such as self-provisioning) are usually accessible through the Internet. Since they allow access to larger sets of resources than traditional hosting providers, they pose an increased risk, especially when combined with remote access and web browser vulnerabilities.

- *Application protection*. Protecting applications becomes more complex if they rely on cloud services in a mix of different environments, or where the applications themselves are moved between environments (e.g. "bursting" from an on-premises environment to a public cloud environment).

- *Service unavailability*. Hardware, software or communication network failures could occur with cloud services, or with on-premises systems, or the data links between them. It is wise to design hybrid systems to cope with these failures gracefully.

- *Insecure or incomplete data deletion*. Where data and applications are spread across different environments, it is likely to be harder to ensure correct data deletion across all of them.

- *Visibility and Audit*. It is important for the cloud service customer to be able to get visibility into activities in all the different environments used in the hybrid cloud solution. Interoperability and data portability are key issues in this, to enable the customer to use a common set of tools and systems.

As customers transition their applications and data to use hybrid cloud, it is critical for them to maintain, or preferably surpass, the level of security they had in their traditional IT environment.

Hybrid cloud computing may create new security risks but it also provides opportunities to provision improved security services that are better than those many organizations implement on their own.

Cloud service providers may offer advanced security and privacy facilities that leverage their scale and their skills at automating infrastructure management tasks. This is potentially a boon to customers who have few skilled security personnel.

Following the CSCC security white paper, it is appropriate to use a series of steps for hybrid cloud customers to evaluate and manage the security and privacy aspects of their use of cloud services, with the goal of mitigating risk and delivering an appropriate level of support. At each step, hybrid cloud requires consideration of the different environments being used and the connections between them.

Specific considerations:

- ID and Access Management (SSO & Federation) – ideally there should be a single IdAM system (potentially one already in use in-house by the customer's organization)
- Secure networking – consider use of facilities such as VPN between the on-premises environment and the cloud services
- Encryption – needs to be in place for all sensitive data, wherever it is located – this requires a planned approach to key management
- Perimeter security (firewalls, DDoS attack handling, etc.) – needs to be coordinated across all environments with external interfaces

Manage the Cloud Environment

Implementing service management in hybrid cloud environments may be challenging due to different factors. For example, cloud providers may or may not provide their own management processes and tools, or public cloud services may or may not be manageable from the traditional service management tools running in an on-premises environment.

However, implementing a single point of management is a must for any enterprise IT department in order to be able to fulfil the requested SLAs and QoS. This becomes even more important when part of the infrastructure or the services are not directly in the control of the IT team.

The major business driver for implementing a hybrid cloud service management system or to extend the on-premises one is to get a single point of control for end-to-end resources and services. However, leveraging the services and the built-in management tools from a cloud provider sometimes might reduce the management costs in terms of licences and effort.

Here are the critical steps cloud customers must take to plan a management solution for hybrid cloud:

1 – Analyze the management processes and use cases you need to implement

Start from the major service management ITIL processes and depending on your needs decide which ones you need to implement. The following list describes the subset that makes sense in a hybrid environment:

- *Service strategy* – financial management (metering/billing), etc.
- *Service design* – availability management, capacity management, etc.

- *Service operation* - event management, request fulfilment, incident & problem management, common service operation activities (for example, backup/archive), etc.
- *Service transition* - configuration management, software asset/license management, change management, patch management, etc.

2 – Analyze on-premises management tooling

One of the simplest approaches is to extend the management processes and tools already being used and leverage them to manage the resources/workloads in the cloud. The main points to consider:

- How well do existing management tools support the new cloud infrastructure?
- At which level (application, middleware, operating system and infrastructure) do on-premises tools manage the cloud resources?
- How much will it cost (effort and licenses) to install, integrate and maintain current management agents on the cloud infrastructure?
- Do the on-premises management tools provide APIs to facilitate eventual integrations?

3 – Analyze cloud service provider management functions and the service responsibility line

Another option is to leverage the management capabilities provided by the cloud platform to manage both the cloud resources/workloads and, if possible, on-premises resources/workloads.

A fundamental aspect to understand when analyzing the management capabilities from a cloud provider is the so called s*ervice responsibility line* and the specific management capabilities offered by the service provider (see Figure 4). Some CSPs supply additional management services for components above the SRL; for example, performance monitoring tools, backup tools, patch management for operating systems, disaster recovery services, etc.

To exploit the cloud service provider management functions, the following points need to be considered:

- Does the cloud service provider support the necessary management processes?
- How much would it cost to manage the cloud resources/workloads?
- Can the cloud services manage on-premises resources?
- Do the cloud management tools provide APIs to facilitate integration?

Architectural patterns for hybrid cloud management solutions

There is no “one-size-fits-all” solution for hybrid cloud management. Experience indicates that there are three design patterns that are used to build a hybrid cloud management solution. These patterns can be used alone or in combination.

The three patterns are represented in Figure 6.

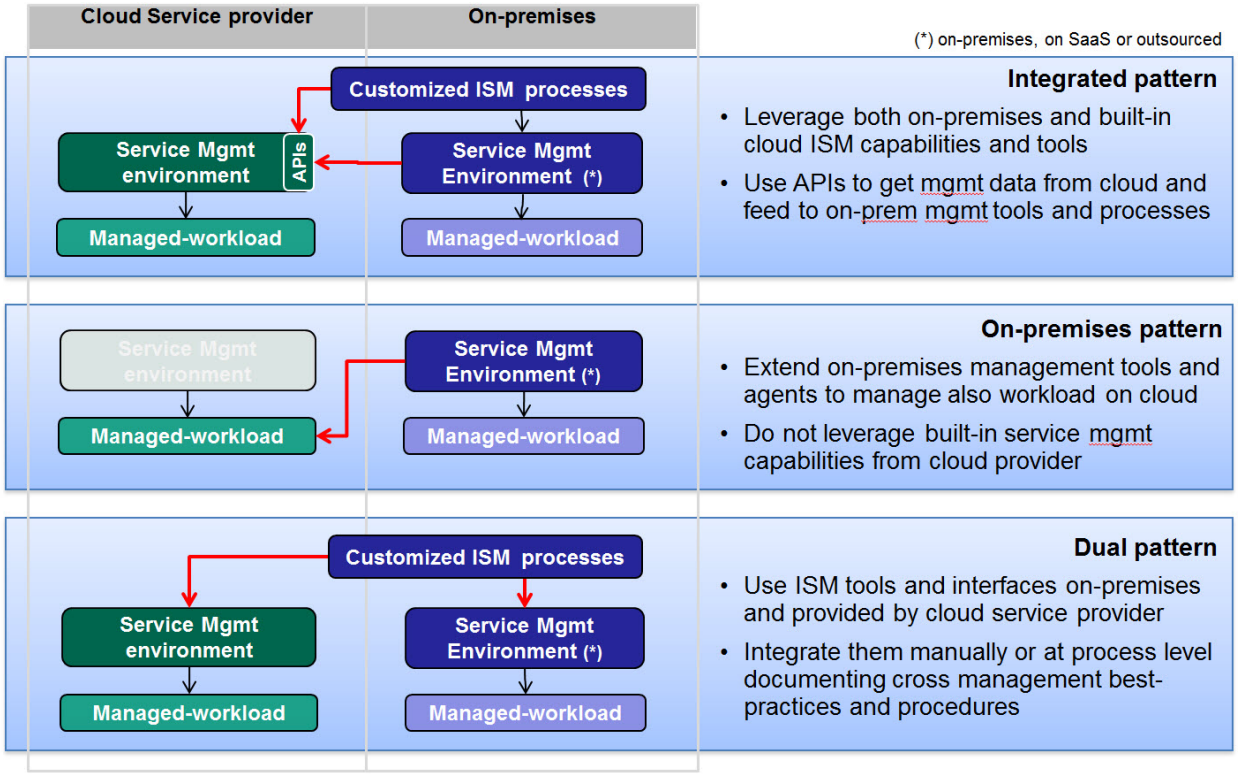

Figure 6: Patterns for hybrid cloud management solutions

The **"Dual Pattern"** is the starting point and it is based on the assumption that cloud service customers use two different management stacks/tools: those used to manage on-premises resources/workloads and those provided by the CSP. Both need to be integrated at the management process level.

An advantage of the dual pattern is that it can be a simple starting point. If cloud customers are able to gather, consolidate and integrate management data manually, then there is a good opportunity for automation.

The downside of this approach is that it requires customizing management processes with manual steps or investing in developing automation.

To summarize, you can use dual pattern if you want to integrate a few management capabilities across on-premises and cloud workloads and for scenarios that cannot be implemented with the other methods (for example, operating system patch management when the OS is completely managed by the service provider).

The **"On-premises pattern**" extends the scope of the on-premises management tools to manage the cloud resources and workloads, usually by installing management agents on the cloud resources/workloads. The management of cloud resources/workloads is typically done above the service responsibility line since commercial management tools do not have visibility of what happens below the service responsibility line.

The advantage of this pattern is that it's relatively simple to implement and does not require big changes to cloud service customers' management processes.

The downside is that the installation of additional management agents increases management costs (the agents themselves need management) and there may be additional license costs. Another drawback is that on-premises management tools might not integrate well with cloud service provider capabilities (e.g., agents might not intercept infrastructure or middleware events) and some scenarios might not be covered (e.g., alarms/tickets automatically created in an IaaS cloud service cannot be intercepted).

To summarize, the on-premises pattern is recommended if cloud customers rely on the cloud service provider to manage the cloud resources/services below the service responsibility line and they want to manage a limited amount of workloads (VMs or services). Consider that in many cases cloud customers may need to implement a manual process (see dual pattern) to realize those management scenarios that cannot be implemented because of limitations of the on-premises management tools in managing the cloud environment.

The **"Integrated pattern"** is an evolution of the "dual pattern"; it leverages both the on-premises and cloud-provider management services using existing APIs to automate the gathering, transfer and consolidation of management data from both environments.

The advantage of the integrated pattern is that it provides a single point of management from the on-premises tools while it does not require installation of management agents on the cloud services. In some cases APIs allow deeper integration.

The downside is that it needs a new bridge component between on-premises tools APIs and cloud service provider APIs.

To summarize, the integrated pattern can be more cost-effective to manage a certain number of cloud services; the number depends on the costs of creating and maintaining the integration code. There could be cases where this approach cannot be used because the cloud service provider or the on-premises management tools do not provide appropriate APIs.

Consider a Backup, Archive and Disaster Recovery Plan

A hybrid cloud computing environment requires careful planning of backup, data archive and disaster recovery mechanisms - various hybrid cloud deployment options should be considered.

Backup and Restore

Backup and restore in hybrid cloud computing has the same high-level requirements as it does inside a traditional data center. However, traditional backup and restore mechanisms lack the accessibility (internet), flexibility (on-demand scalability) and protection (off-premises) that a remote, cloud-hosted backup solution provides. Customers still require a mechanism to handle data deletion (accidental or malicious) or data corruption. Some customers may also choose to leverage backup data for disaster recovery (DR).

The backup solution must be reliable, highly available, scalable and secure. Some of the key challenges for backup up in a hybrid cloud environment:

- If the data is split between on-premises data center and public cloud, some form of backup synchronization may be required to ensure that the restored data is consistent.
- A single backup tool must support the management of the backup of data that is split between different sites (on-premises or public cloud) and different geographic locations.
- Data must be secured by using a highly available encryption mechanism.
- If primary data is migrated to the cloud, the backups that exist for this data must be managed such that data can be restored whenever and wherever necessary - the frequency of these data restores must be coordinated with the CSP.
- Solution delivery teams have to be trained to manage backups in the hybrid cloud.

Data in the hybrid cloud can be backed up in one of three ways:

- Architect a solution from scratch using backup products
- Use a backup service
- Use the native cloud backup services

The method chosen for backup will depend on the service level agreements (SLAs), cost and support requirements. For example, when the primary data is in the public cloud, the backup service provided by the CSP may be sufficient to meet the customer requirements. If the primary data is on-premises in a non-cloud environment, a backup solution with scheduled backup operations produced by a backup manager application may be necessary.

Use cases and SLAs determine where to locate the backup repository. If the use case for backup is to handle data deletion / data corruption, the backup data is usually co-located with the primary data to enable quick restores. In some cases, the backup data is in offsite location for reasons of either cost or disaster recovery.

When data is backed up to a public cloud service, the customer has to pick from the choices offered by the CSP for the location and type of the backup repository.

Archive

Data is archived in the hybrid cloud for the same reasons as in the traditional on-premises data centers:

- Move inactive / infrequently used data to less expensive storage (e.g., old email)
- Long term retention of data such as customer records, media files (images, video) etc.

The challenges for archive in relation to hybrid cloud are the same as those for backup. If the data has to meet regulatory requirements, such as HIPAA or Sarbanes-Oxley, it has to be stored in a secure, tamper-proof, read-only environment for as long as needed by law.

The method of archiving data depends on the requirements. If data is on-premises, the solution could use well-known products such as NetApp SnapLock or IBM FileNet. If data is in the public cloud, data is archived to locations supported by the CSP.

Create a Disaster Recovery Plan

Creating a Disaster Recovery Plan (DRP) for a hybrid cloud solution varies only slightly from creating a DRP in the traditional world, with the need to consider how cloud resources are handled.

Using a public cloud service provider for DR is possible, but comes with a few challenges:

- **Resource readiness not predictable and may have wide variance**: DR is associated with a defined RTO (Recovery Time Objective), which indicates the time required to restart operations in emergency conditions. A component of that RTO is the time required to provision the emergency infrastructure (DR servers). Most CSPs do not provide any written commitment on resource provisioning time.

- **Large scale resource request may take time to fulfill**: DR may lead to a massive concurrent request for resources, which typically happens for large scale disasters. Can the CSPs fulfill such a burst of requests?

- **Cloud service providers are not in the DRaaS business**: Cloud customers are solely responsible for developing, debugging, and executing recovery procedures. Most CSPs do not offer managed DRaaS, but only sets of components that customers can select, combine and run to build their own DRaaS.

A deep analysis of current IT patterns and dependencies is required to determine what components need to be available in case of a disaster and when they need to be back available to protect your business.

When using a hybrid cloud solution, evaluate the feasibility of a split among different sites (and/or providers), and the feasibility to recover using cloud services, by verifying whether the cloud provider has all the IT platforms needed for your system, both in terms of hardware (IT infrastructure) and software.

Choose the DR solution model

At a high level the DR solution can be classified in three different models, each model being associated to different performance in terms of time to restart (RTO: recovery time objective) and maximum data loss allowed (RPO: recovery point objective).

- **Active-Active**: this type of solution implies the use of active components on multiple sites, involving both running the production workload and active data replication required to maintain a copy of what is running in the other site. The primary and secondary data centers must have low network latency to avoid performance impacts. These solutions usually require the IT resources in the second data center to always be in use. This has cost implications.

- **Active-Passive**: in this type of solution the only active components in the DR site are those required to sustain the data replication process. All the other components are provisioned on-demand.
- **Passive**: the required components in the DR site, including data, are provisioned on-demand - data are re-loaded from a back-up copy

In case there is a disaster, the DR solution must meet the following requirements:

- Ability to access DR configuration information within the timeframe that has been agreed upon with the cloud provider.
- Access to a data center site with availability requirements and metrics consistent with the source site.
- Ability to restart operations in full compliance with expectations (RTO and RPO).

Evaluate if the Hybrid-DR cloud-based solution satisfies all your requirements

Having defined the evaluation and benchmarking criteria, customers need to define the benefits and impact that a cloud-based hybrid DR entails.

For both cloud-to-cloud DR (same provider) and for on-premises to cloud DR, consider the following options:

- Evaluate if the data replication available from the provider provides you control over replication and/or data consistency
- Adopt software defined storage (SDS) that contains the required data replication functions
- Adopt software based data replication
- Leverage application/database based data replication

Customers should compare new cloud-based solutions with the ones offered by DR providers:

- The solutions could cover the platforms needed.
- Some providers might tailor their services to meet your requirements.
- Some experienced and knowledgeable DR providers could guarantee the solution throughout its lifecycle. The provider could address all the above requirements and maintain co-responsibility for the SLAs, regardless of the data center that is used (customer, DR provider, public cloud provider) or the architecture chosen.

When all these points have been addressed, customers can compare costs and performance of viable options and choose the solution with the best price/value that fits their DR needs.

Summary of Keys to Success

Key Factor	Description
1. Determine cloud deployment model for applications and data	• Determine the right resource model – on-premises private cloud, hosted private cloud, or public cloud • Rationalize application and data environment • Apply decision criteria to define the right deployment model – flexibility, security, speed & automation, cost, locality, service levels, and system interdependencies • IT architects consider options for application placement in the hybrid cloud (see four options)
2. Integrate with existing enterprise systems	• Put in place controlled interfaces by which components in cloud services can access applications and/or data in on-premises systems – consider technologies such as API Management • Consider the administration and business aspects of the integration as well as the functional integration of the systems • Demand that the cloud service provider supports standards for the interfaces to their cloud services
3. Address connectivity requirements	• Consider the requirements of each link between components that spans two or more cloud services or on-premises system and ensure that appropriate connectivity is available to support those requirements • Consider the use of network virtualization if available • Ensure that the connectivity capabilities can support resilience and disaster recovery requirements
4. Develop governance policies and service agreements	• Assess existing compliance and governance frameworks, identify gaps and harmonize processes • The need for thorough and efficient change management and communications increases with the addition of multiple cloud service providers • Allow adequate time to educate and habituate changes across the organization • Identify gaps in measurement and management visibility
5. Assess and resolve security and privacy issues	• Understand the interfaces between components running in private cloud services, in public cloud services and on-premises and apply appropriate and consistent security controls to each of them • Evaluate the location of all datasets in the hybrid cloud deployment and ensure the application of consistent access controls and encryption • When migrating application components between environments, be careful to check that the security controls in place for the new environment meet or exceed those in place for the old environment • Apply technologies across all the environments that are part of the hybrid cloud deployment – for example, single IdAM system and Single-Sign-On
6. Manage the cloud environment	• Enable management of the complete hybrid cloud system, spanning all the environments used • Either adapt and integrate existing on-premises management tools or consider use of new cloud based management services, based on cost and functionality • Look for APIs and integration points for management capabilities rather than fixed-function management applications

7. Consider backup, archive and disaster recovery plan	• Monitor the frequency of backup and archiving as this will drive cloud service provider costs • For public cloud workloads and components, make certain legal agreements are in place, as necessary • The location of the backup and archive data will determine the Recovery Time Objective (RTO) for restore and retrieval. The RTO will be lower if they are collocated with the primary data. Restoring over the WAN will introduce latencies that will increase the RTO. The architectural decision in this regard must be made to meet the SLAs. • Determine what resiliency and backup capabilities are provided out-of-the-box for the cloud services portion of the hybrid cloud deployment • For offsite backup and archiving of sensitive, proprietary or financial data, make certain the cloud service providers' physical location is permissible and acceptable given legal and regulatory constraints

Works Cited

[1] CSCC: *Practical Guide to Cloud Computing.* http://cloud-council.org/resource-hub.htm#practical-guide-cloud-computing-v2

[2] ISO/IEC 17788: *Cloud Computing Overview and Vocabulary.* http://standards.iso.org/ittf/PubliclyAvailableStandards/c060544_ISO_IEC_17788_2014.zip

[3] CSCC: *Interoperability and Portability for Cloud Computing: A Guide.* http://www.cloud-council.org/resource-hub.htm#interoperability-and-portability-for-cloud-computing-a-guide

[4] ISO/IEC 17789: *Cloud Computing Reference Architecture.* http://standards.iso.org/ittf/PubliclyAvailableStandards/c060545_ISO_IEC_17789_2014.zip

[5] CSCC: *Practical Guide to Cloud Service Agreements V2.0.* http://www.cloud-council.org/resource-hub.htm#practical-guide-to-cloud-service-agreements-version-2

[6] CSCC: *Security for Cloud Computing: 10 Steps to Ensure Success.* http://cloud-council.org/resource-hub.htm#security-for-cloud-computing-10-steps-to-ensure-success

[7] CSCC: *Migrating Applications to Public Cloud Services: Roadmap for Success.* http://cloud-council.org/resource-hub.htm#migrating-applications-to-public-cloud-services

