



Interoperability and Portability for Cloud Computing: A Guide

November, 2014

Contents

Motivation and Considerations	4
Interoperability & Portability Overview.....	5
Basic Definition of Interoperability.....	5
Basic Definition of Portability	5
Interoperability & Portability Challenges.....	6
Elements involved in Interoperability & Portability for Cloud Services.....	7
Interoperability & Portability Scenarios	9
Scenario 1: Customer Switches Providers for a Cloud Service	10
Scenario 2: Customer Uses Cloud Services from Multiple Providers	13
Scenario 3: Customer Links One Cloud Service to Another Cloud Service	17
Scenario 4: Customer Links In-house Capabilities with Cloud Services	19
Scenario 5: Migration of Customer Capabilities into Cloud Services.....	23
Summary of Key Considerations and Recommendations.....	27
Standards for Interoperability and Portability.....	28
Works Cited.....	29
Additional References.....	30
Appendix A: Interoperability Model for Cloud Computing.....	30

© 2014 Cloud Standards Customer Council.

All rights reserved. You may download, store, display on your computer, view, print, and link to the *Interoperability and Portability for Cloud Computing: A Guide* white paper at the Cloud Standards Customer Council Web site subject to the following: (a) the document may be used solely for your personal, informational, non-commercial use; (b) the document may not be modified or altered in any way; (c) the document may not be redistributed; and (d) the trademark, copyright or other notices may not be removed. You may quote portions of the document as permitted by the Fair Use provisions of the United States Copyright Act, provided that you attribute the portions to the Cloud Standards Customer Council *Interoperability and Portability for Cloud Computing: A Guide (2014)*.

Acknowledgements

The major contributors to this whitepaper are: Claude Baudoin (cébé IT & Knowledge Management), Eliezer Dekel (IBM), Mike Edwards (IBM), John Meegan (IBM), Jem Pagan (JNK Securities), Karolyn Schalk (Garden of The Intellect LLC), Gurpreet Singh (Ekartha), Joe Talik (AT&T) and Steven Woodward (Cloud Perspectives).

Executive Overview

Cloud computing is rapidly growing in importance for many organizations, with ongoing take-up of a wide range of cloud services and the transition of both data and applications to cloud computing environments. The topics of *interoperability* and *portability* are significant considerations in relation to the use of cloud services, but there is also confusion and misunderstanding of exactly what this entails. The aim of this Guide is to provide a clear definition of interoperability and of portability and how these relate to various aspects of cloud computing and to cloud services.

Interoperability and Portability for Cloud Computing: A Guide describes interoperability and portability in terms of a set of common cloud computing scenarios. This approach assists in demonstrating that both interoperability and portability have multiple aspects and relate to a number of different components in the architecture of cloud computing, each of which needs to be considered in its own right. The aim is to give both cloud service customers and cloud service providers guidance in the provision and selection of cloud services indicating how interoperability and portability affect the cost, security and risk involved.

Motivation and Considerations

Cloud computing is having an enormous impact on how organizations manage their information technology resources. The abundance of easy to access computing resources enabled by cloud computing provides significant opportunities for organizations, but poses challenges for enterprises in a number of areas. The current cloud computing landscape consists of a diverse set of products and services that range from infrastructure services (IaaS), to specific software services (SaaS) to development and delivery platforms (PaaS), and many more. The variety of cloud services has led to proprietary architectures and technologies being used by vendors, increasing the risk of vendor lock-in for customers. Incidents such as a cloud service providers shutting down operations or the discovery of significant security vulnerabilities in applications have highlighted this risk.

Cloud service customers need to avoid the problem of lock-in, where they run the risk of being tied to a particular cloud service provider due to the difficulty and costs of switching to use equivalent cloud services from other providers. As an example consider an organization using a PaaS (Platform as a Service). A PaaS platform from a particular vendor could support only limited and proprietary web frameworks, languages, libraries, databases, etc. This can lead organizations to develop application architectures dictated by features offered by the PaaS cloud service provider which can lead to their applications being locked to that vendor, essentially non-portable. There are no perfect solutions to completely avoid these problems, but organizations need to consider this issue carefully when selecting cloud services.

As enterprises move to adopt cloud computing in its various manifestations, the issues of portability and interoperability need to be addressed head on by both providers and customers. To mitigate the risk of lock-in organizations should review existing data governance and purchase policies and processes to see if these support a strategy to achieve high levels of interoperability and portability.

In our view, the goals of cloud interoperability and portability – which are to enable cloud users to avoid vendor-lock in, allow for customers to make best use of multiple diverse cloud services that can cooperate and interoperate with each other – are critical to future cloud service adoption and the realization of the benefits of computing as a utility.

Interoperability & Portability Overview

The cloud eco-system is large, with many providers offering a wide variety of cloud services. Understanding the portability and interoperability “of what” is the necessary first step of planning and designing for the use of any cloud service. Clarifying the specific interoperability and portability concerns accelerates identification of the “best fit” options and potential development of solutions. This section provides an overview of the topics of interoperability and portability which is useful in understanding the more detailed descriptions contained in the scenarios and subsequent sections.

Basic Definition of Interoperability

Broadly speaking, *interoperability* can be defined as a measure of the degree to which diverse systems or components can work together successfully. More formally, IEEE and ISO define interoperability as the ability for two or more systems or applications to exchange information and mutually use the information that has been exchanged. To be more concrete, in the context of cloud computing, interoperability should be viewed as the capability of public clouds, private clouds, and other diverse systems within the enterprise to understand each other’s application and service interfaces, configuration, forms of authentication and authorization, data formats etc. in order to cooperate and interoperate with each other.

In cloud computing, the most significant interacting components are those which belong to the cloud service customer which interact with components of the cloud service provider. The nature of the interaction is a network connection using a prescribed interface or API. There are typically multiple separate interfaces, each dealing with a different aspect of the cloud service. For example, there are the functional interfaces of the cloud service itself, authentication and authorization interfaces, interfaces for administration of the cloud services, and business interfaces for billing and invoicing. The ideal of interoperability is that the interfaces are standardized in some way – i.e. they are *interoperable* - so that the customer can switch to another cloud service provider with minimal impact on the customer's components.

It is important to recognise that there are different aspects of interoperability which can be positioned and described as separate layers. Refer to Appendix A for a detailed discussion on an interoperability model for cloud computing which describes each layer and the associated interoperability challenges that need to be addressed.

Basic Definition of Portability

Portability is about the ability to move an entity from one system to another so that it is usable on the target system. The main problem caused by the lack of portability is that it may take considerable effort

to transform the entity from its format on the source system to the format required by the target system. Portability is divided into two separate areas: data portability and application portability:

- *Cloud data portability* is the ability to easily transfer data from one cloud service to another cloud service, without being required to re-enter the data. It is the ease of moving the data that is the essence here. This might be achieved by the source service supplying the data in exactly the format that is accepted by the target service. But even if the formats do not match, the transformation between them may be simple and straightforward to achieve with commonly available tools.

The first aspect of data portability is that there must be a capability to retrieve customer data from the source cloud service and also a capability to import customer data into the target cloud service. This is commonly done through the existence of some API (or web interface) associated with the cloud service – it may be a generic API such as one of the forms of FTP, for example, or it may be a specific API unique to the cloud service. It is important to note that the API used for the source service may not be the same as the API used for the target service and that different tooling may be required in each case.

The second aspect of data portability is the syntax and semantics of the transferred data. The syntax of the data should ideally be the same for the source service and the target service. However, if the syntax does not match (i.e., the source uses JSON syntax but the target uses XML, say), it may be possible to map the data using commonly available tools. If the semantics of the transferred data does not match between the source and target services then data portability is likely to be more difficult or even impossible.

- *Application portability* is the ability to easily transfer an application or application components from one cloud service to a comparable cloud service and run the application in the target cloud service. The ease of moving the application or application components is the key here. The application may require recompiling or relinking for the target cloud service, but it should not be necessary to make significant changes to the application code.

Interoperability & Portability Challenges

For interoperability, there are many challenges associated with cloud computing. In general, the interfaces and APIs of cloud services are not standardized and different providers use different APIs for what are otherwise comparable cloud services.

The greatest level of interoperability is likely to be found for IaaS cloud services, where functionality is often broadly equivalent and there are a number of standard interfaces - some formally standardized such as CDMI, others being de-facto standards in the marketplace. PaaS cloud services have lower levels of interoperability. There are few interface standards for PaaS functionality, although there are some open source platforms that are becoming popular in the marketplace and where different cloud service providers use the same open source platform, their interfaces are either identical or closely equivalent.

It is SaaS applications which present the greatest interoperability challenge today. There are very few standard APIs for SaaS applications - even switching from one SaaS application to another SaaS application with comparable functionality typically involves a change in interface. There is a resultant impact on both end users of the cloud service for any user interfaces and also on any applications or systems belonging to the cloud service customer that use APIs offered by the SaaS application.

There are some practical approaches to handling these interoperability challenges. One possibility is for the cloud service customer to provide an isolation or mapping layer between their own applications and systems and the cloud service interfaces, so that the cloud services are not invoked directly by customer applications. Technologies such as enterprise service buses (ESB) can be used to build these isolation layers. An emerging trend is for customers to write custom code in a PaaS platform in order to tailor access to SaaS applications to suit their organizational needs - these PaaS applications form an isolation layer for other customer systems.¹ Another possibility is to use the services offered by an inter-cloud provider (sometimes called a cloud service broker), who takes on the role of mapping a “standard” interface offered to the customer to a varying set of interfaces offered by a number of different cloud service providers.

For application portability, the biggest challenges are for applications built for PaaS platforms. For IaaS cloud services, there are in practice a number of standards that enable portability of applications, such as OVF and the provision of commonly used operating systems like Linux. PaaS platforms can vary widely between different providers - the app environment can differ substantially, including the way in which platform services are offered to the app code and also in which set of platform services are available. For example, in order to be scalable and elastic a PaaS may enforce a specific way to persist and manage data that may not be supported by other PaaS platforms. The differences between alternative PaaS platforms can require extensive re-engineering of customer code when the code is moved between those platforms. Two trends in the marketplace seem to offer promise for easing the situation - one is the increasing adoption of common open source PaaS platforms such as Cloud Foundry, while another is the emergence of containerization technologies that allow subdivision and independent deployment of parts of an application, such as Docker.

Elements Involved in Interoperability & Portability for Cloud Services

Before examining specific scenarios, it is worth highlighting the various interfaces between the cloud service customer and the cloud service, along with the data that is held within the cloud service.

¹ Some SaaS providers make available a PaaS platform alongside their SaaS services specifically to make it simpler for customers to write and run custom applications based on the API made available by the SaaS service, such as the Heroku PaaS supplied by Salesforce.

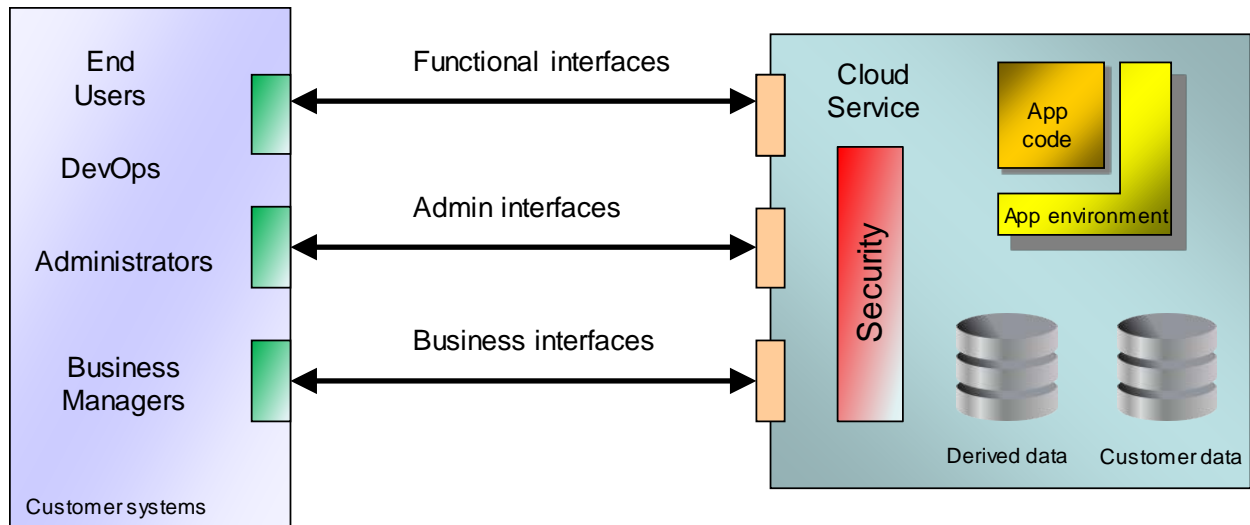


Figure 1: Elements of Interoperability and Portability for Cloud Services

Figure 1 illustrates some of the key customer roles associated with using a cloud service and the components, interfaces and the data associated with that use.

In this diagram, the **App code** within the cloud service can be taken to represent the customer application code in the case of IaaS and PaaS cloud services, but in the case of a SaaS service, the app code would typically belong to the provider and would be managed by the provider.

The **App environment** represents the "API" that the cloud service presents to the app code – and the app code must be able to use this API in order for the app to run. The App environment could be the operating system, or it might be an API offered by some middleware framework, depending on the nature of the cloud service.

The **Security** component represents a set of capabilities which are used to secure the cloud service and includes authentication and authorization of users, encryption of data in motion and at rest, firewalls and technologies for dealing with attacks such as DDoS (distributed denial of service).

Customer data represents the cloud service customer data, which may be held as records in a database or held as data objects in files or in a data store. **Derived data** represents data which is created and stored as a result of the customer use of the service, such as log records or configuration information.

Figure 1 shows three main interfaces between customer roles and the cloud service: the Functional interfaces, the Admin interfaces and the Business interfaces. The **Functional interfaces** are associated with the main functional capabilities offered by the cloud service. The **Admin interfaces** involve capabilities for administering the cloud service and include capabilities such as monitoring the service and managing its behavior including aspects of security such as user identities, authentication tokens and authorizations. The **Business interfaces** involve capabilities relating to the business aspects of the cloud service including subscription information, billing and invoicing.

It is important to understand that each of these interfaces may have multiple forms. 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 app. However, the same capabilities may also be made available as an 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.

Interoperability aspects of a cloud service mainly relate to the three interfaces between the customer and the cloud service – how users and applications in the customer environment interact with the functional, admin and business interfaces offered by the cloud service. It is important to understand that the interoperability of the three interfaces may be independent of each other and that the interoperability of one interface does not guarantee the interoperability of the others.

Application Portability relates to the capability of moving the App code to or from the cloud service. This typically only applies to IaaS and PaaS services, since in the case of a SaaS service, the App code belongs to the cloud service provider and cannot be ported elsewhere by the customer. One of the most important factors for application portability is represented by the **App environment**. To port code from one cloud service to another one, the target app environment must be usable by the application being ported.

Data Portability relates to the capability of moving data into and out of the cloud service environment. Typically, it is the cloud service customer data which is the concern for data portability. However, some of the cloud service derived data may also be of concern in relation to some cloud services and should not be overlooked. For cloud service customer data, data portability is usually of most concern for SaaS cloud services, since for these services, the content, data schemas and storage format are under the control of the cloud service provider and the customer will need to understand how the data can be imported into the service and exported from the service. For IaaS and PaaS services, it is typically the case that the cloud service customer is in control of the content and schemas for the data, with the service offering basic storage capabilities such as a file system or object store.

Interoperability & Portability Scenarios

This section leverages a set of scenarios to describe interoperability and portability considerations and requirements including recommendations on how to address them.

1. Customer switches between providers for a cloud service
2. Customer uses cloud services from multiple providers
3. Customer links one cloud service to another cloud service
4. Customer links in-house capabilities with cloud services
5. Migration of customer capabilities into cloud services

These scenarios are described in the sections that follow. Note that the scenarios focus primarily on public cloud since this deployment model presents the greatest interoperability and portability challenges to customers.

Scenario 1: Customer Switches Providers for a Cloud Service

This scenario addresses the straightforward case of a cloud service customer currently using a cloud service of cloud service provider A, who wishes to switch to using an equivalent cloud service of cloud service provider B. This use case is a key one from the perspective of the customer since it is essential to enabling customers to take advantage of the marketplace in cloud services and avoid the issue of lock-in to a single cloud service provider.




Figure 2: Customer Switches Providers for a Cloud Service

While this scenario is outwardly simple, the reality is that this scenario touches on many of the issues associated with both interoperability and portability. The exact set of issues will vary depending on the nature of the cloud service.

Interoperability Considerations

For SaaS cloud services, the App code belongs to the cloud service provider.² In this case, moving from provider A to provider B does not involve porting the app code – the app code may be completely different between the two cloud service providers. What is important in the SaaS case is the

² Some SaaS providers make available a PaaS platform alongside their SaaS services specifically to allow customers to write and run code which customizes the use of the SaaS applications, and also to develop and run new applications. In such cases, transitioning to a new SaaS service of a different provider can also include code portability concerns. Refer to the PaaS description in this section for interoperability and portability implications for this type of customer code.

compatibility of the functional interface for the app – in particular, any user interfaces presented to end users and also any APIs made available for use by customer applications. It is probably unrealistic to expect that user interfaces will be identical for cloud service A and cloud service B, however, it is reasonable to expect that similar functionality will be presented in a broadly similar way to reduce the cost and effort of retraining end users.

However, any functional APIs made available by a SaaS service are likely to be used by customer applications - these applications will need to deal with the switch from cloud service A to cloud service B. If the APIs are not interoperable as is quite likely to be the case, then the implication is that any customer applications using the APIs would need to be changed as part of the process of switching from cloud service A to cloud service B (see the [Interoperability & Portability Challenges section](#) on page 6).

For IaaS and PaaS services, interoperability is not an issue for any functional interfaces offered by the app code when moving from provider A to provider B since the customer owns the application and its functional interfaces. Similarly, the user interfaces presented to end users are likely to be the responsibility of the app code and will not be directly affected by the cloud service itself – as long as the app code can be ported, then the user interface will also port and be available when using cloud service B. However, the cloud service APIs used to upload, deploy and control the app code in the cloud service are an interoperability concern, since tooling used by the customer operations staff uses these APIs and this tooling needs to connect to cloud service A and then to cloud service B as part of the migration.

Portability Considerations

For SaaS services, it is typical that the format and the content of the cloud service customer data is in the hands of the cloud service provider. Thus **data portability** is a major consideration in moving from cloud service A to cloud service B if the cloud service is a SaaS service. Ideally the data format should be the same for both service A and service B. In addition to the format, the data content (extent and semantics) should be the same for service A and service B. Data portability can still be achieved if the formats are different between service A and service B, since there are straightforward standard tools that can be used to perform some data transformations – and where there are not standard tools, it may be possible to build a custom tool. Differences in the extent or in the semantics of the data are much more serious and could be a major barrier to achieving data portability.

For IaaS services, the **data format** for cloud service customer data is usually in the hands of the customer, since the facilities provided by the cloud service are typically relatively low level, such as providing volumes for binary file or object storage (i.e. the cloud service does not know or care about the detailed format of the customer data). As a result, data portability is not likely to be a major concern for IaaS services.

Similar considerations can apply to data portability for PaaS services, but the situation is often more complex. For customer data, the PaaS service may provide instances of databases ready-to-use – in which case, the actual database(s) provided may be sensitive to the data format of the customer data, although there are generalized formats (CSV, XML, etc.) which are supported by many types of database. How data is loaded into a PaaS cloud service and how it is retrieved needs to be examined by

the customer and migration from cloud service A to cloud service B involves data portability questions that need to be answered.

For IaaS and PaaS cloud services, the App code belongs to the customer and the question of **application portability** is of utmost importance – what does it take to move the App code from cloud service A to cloud service B? The first question is what format each cloud service accepts for the app code – for example, do they accept the same format of VM image? Another question that may arise is whether the machine architectures for the two services are the same (this is one aspect of the app environment). Finally, there is the question of the compatibility of the app environment, which can involve the version of the operating system being used, the libraries available to the application and any service APIs used by the application.

In the ideal case, application portability implies that the app code which runs on cloud service A will run on cloud service B without any changes. There may be cases where application portability cannot be achieved without some changes – in this case, it is the amount of change and the nature of the change that must be considered. Rebuilding the application code, possibly against a different version of the operating system or a different version of the libraries used by the code, may be simple and low cost. Redesigning the application code to adapt to changed interfaces is likely to be more costly and is less desirable.

For PaaS services, the considerations for App code portability can be much more complex. The app environment can consist of a substantial stack of software, with many APIs which are used by the app code. In addition, many capabilities can be presented to the app code as services of various kinds (databases, messaging, rules engines, etc.). The app code may have dependencies on a particular set of services, via their APIs, and it is vital to know that the set of services available within cloud service A is matched by the set of services available in cloud service B.

Admin Interface Considerations

Cloud service admin interfaces, which are used to monitor and manage applications, are significant for all forms of cloud services. These admin interfaces may involve web applications or other visual interfaces and may also involve APIs.

Moving from cloud service A to cloud service B requires that the admin interfaces are compatible (particularly in the case of visual interfaces) and also interoperable (particularly in the case of APIs).

It may be the case that the admin interface is divided into separate sections dealing with particular capabilities – for example, monitoring and reporting capabilities may be delivered by one interface, while management and administration capabilities may be delivered by a different interface.

Business Interface Considerations

The business interfaces apply to all forms of cloud services and include the capabilities relating to subscription management, billing and invoicing.

Moving from cloud service A to cloud service B requires compatible and/or interoperable business interfaces to ensure that the tools or program components used by the cloud service customer for business capabilities can be used successfully following the move.

Security Considerations

Security aspects of the cloud service include authentication and authorization of users and administrators of the cloud service, the configuration and operation of encryption both for data stored within the cloud service and also for data transmitted to and from the cloud service, firewalls and the configuration of other security capabilities.

The security aspects have a number of parts – some parts apply to the running of the cloud service itself, other parts consist of the administration of the security components, including setting up or modifying user identities and the capabilities they are authorized to use. It is important to check that equivalent security capabilities available for cloud service A are available for cloud service B.

The support of some technologies by the cloud services makes the transition from one cloud service to another one potentially simpler – an example is the support of third party Id and Access Management function, where it is possible that the IdAM system could be one installed and operated by the cloud service customer and where that system is used by both cloud service A and cloud service B with few changes – not requiring the porting of a large set of data about users and not requiring the need to change the interface used to administer the user data. Other standard technologies that can help include the increasingly common support of OAuth 2.0 and OpenID that permit use of commonly available ID and access management services across a range of cloud services.

Recommendations:

- For SaaS, ensure user interfaces, APIs, protocols and data formats are well defined for cloud services. Whenever possible, insist on standard APIs, protocols and data formats.
- For PaaS, ensure that the application environment is based on open technologies to increase the number of viable alternative cloud service providers which can facilitate migration if a change in provider is warranted.
- For IaaS, ensure that the cloud service accepts standard or widely accepted application packaging formats such as OVF and that any interfaces and APIs are open and/or standard.
- Insist that your cloud service provider supports key open technologies (open standards and/or open source) for admin and business interfaces.
- Leverage the support of third party ID and Access Management functionality to authenticate and authorize access to cloud services.

Scenario 2: Customer Uses Cloud Services from Multiple Providers

This scenario concerns the case where a cloud service customer uses one (or more) cloud services from cloud service provider A and one (or more) cloud services from cloud service provider B. The cloud service(s) from provider A might be equivalent in functionality to the service(s) from provider B, or they

may have different functionality, depending on the business needs that the customer is aiming to satisfy.




Figure 3: Customer Uses Cloud Services from Multiple Providers

An example where the cloud services would have equivalent functionality is the case where the customer uses two providers to get resilience (i.e., continue to access necessary computing resources in the case where one of the providers has a service outage). A different example is where the best SaaS service for one capability is offered by provider A while the best SaaS service for a second capability is offered by provider B – and the customer wants to use both capabilities to satisfy business needs.

This scenario is becoming more mainstream - enterprises have multiple choices when it comes to cloud services because cloud vendors are ready to provide unique capabilities. A best of breed cloud implementation enables organizations to choose the best cloud service for each particular capability. From an interoperability and portability perspective, this scenario touches on many of the issues discussed in Scenario 1 above.

Interoperability Considerations

Whether the functional interface of cloud service 1 needs to be interoperable with the functional interface of cloud service 2 depends on whether the two cloud services are dealing with equivalent functionality or different functionality.

In the case where the cloud services deal with equivalent functionality then it is likely that the same customer components will interact with both services –as a result it is best if the two services use the same or interoperable interface(s). If services with equivalent functionality do not have interoperable interfaces then customer components will need to be updated to support the two interfaces – not an ideal situation.

Where the two cloud services deal with different functionality, then the need to use the same or interoperable interfaces is lower – indeed, for very different functionality it is probably not reasonable to expect interoperable interfaces. However, there may be aspects of the two functional interfaces that

should be based on the same technologies. One example is the technology used for Identity and Access Management, since this is a common feature of most cloud services.

Portability Considerations

Data portability can be a requirement between cloud service 1 and cloud service 2. It is usually not a significant issue for IaaS services, since the cloud service customer is typically in control of the data formats used for these services. For PaaS services, the cloud service customer usually has a lot of control of the data formats, but this may be limited where the PaaS services make use of particular database technologies – for example, in the case where cloud service 1 uses one database technology while cloud service 2 uses a different database technology.

For SaaS services, data portability can be a significant issue where cloud service 1 is equivalent to cloud service 2 (as in the case of scenario 1). Even in cases where the two SaaS services are not equivalent in functionality, there may be a need to use some data extracted from cloud service 1 in the operation of cloud service 2 – in this case, it is best if the data extract has the same format, extent and semantics for cloud service 1 and cloud service 2. If not, then some form of data transformation may be required in order for the customer to use both cloud services successfully.

For IaaS and PaaS services which involve the deployment of application code into the cloud service, it can be important for the app environment to be equivalent between the cloud services. This is necessary where application portability is required when the same app code gets deployed to cloud service 1 and code service 2. However, it can also be important in cases where different app code is deployed to the different cloud services, since the developers may well want to use the same knowledge, technologies and tooling to build both sets of app code, which may be difficult if the app environment offered by the cloud services differ substantially. For IaaS services, machine image formats are an important component of application portability. A portable machine image format improves portability across different service providers. An example of a standard machine image format is Open Virtualization Format (OVF).

Admin Interface Considerations

Administering cloud services from multiple cloud service providers means interacting with two sets of admin interfaces used to monitor and manage each of the cloud services. If the two interfaces are not interoperable, the implication is duplication of effort. There might be multiple user accounts to manage, multiple sets of access controls to maintain, and multiple administration portals to learn and use. Even a relatively simple task, like checking to see which machine instances are running, can become a multi-step process.

In the ideal case, the customer staff should use one set of tools and applications to monitor and manage all cloud services, irrespective of which cloud service provider is being used. This requires that any admin APIs are interoperable and that any admin visual interfaces are composable (for example, based on standard web technologies that can be integrated on a single browser).

Where the admin interfaces of two cloud services are not equivalent, it is necessary for the customer to consider techniques and technologies which can reduce the impact of the differences – for example, by using admin tooling that has adapter layers that can connect to multiple different admin APIs.

Business Interface Considerations

Using cloud services from provider A alongside cloud services from provider B implies a requirement on the cloud service customer to integrate the capabilities relating to subscription management, billing and invoicing. This is necessary in order to keep a good grip on expenditures and is also a requirement for allocating what are likely to be dynamic costs to the right internal budgets and projects.

In the ideal situation, the business interfaces offered by provider A are compatible and interoperable with those offered by provider B. This enables the customer to use a single set of business tools to manage the usage of all the cloud services. Where the business interfaces are not all interoperable, then the customer should look for business tools that can perform mapping or transformation of each of the different business interfaces offered by the different providers. An example would be a tool that can cope with a range of invoice formats.

Security Considerations

Using multiple cloud services from different providers requires that different aspects of security be carefully assessed – one aspect pertains to the running of the cloud service itself, other aspects consist of the administration of the security components, including setting up or modifying user identities and the capabilities they are authorized to use. The ideal situation is that these aspects are interoperable between the different providers, enabling a single set of tooling and procedures to be used by customer staff.

In some cases, this can be achieved by delegation of capabilities from the cloud service to the customer, such as the support of third party Id and Access Management, which may be a system owned and run by the customer. In other cases, the provision of capabilities using a standard interface and standard technology is a useful approach – for example with respect to data encryption.

Recommendations:

- Refer to recommendations for Scenario 1.
- For SaaS, ensure user interfaces, APIs, protocols and data formats are identical (or have a clear mapping) for equivalent functionality running on different cloud service providers.
- Consider implementing an Enterprise Service Bus (ESB) to perform interface, protocol and data transformations to address differences between cloud services from different providers.
- For PaaS, ensure that the application environments (web server, database server, etc.) supported by different cloud service providers are compatible.
- Consider the use of an intermediary (an “inter-cloud provider”) to help address and solve the issues of integration, portability and interoperability of multiple cloud services.

Scenario 3: Customer Links One Cloud Service to Another Cloud Service

In this scenario, the customer uses two cloud services, but one of the cloud services is used directly by the other one. The ability to link cloud services together in support of a single application or an integrated set of applications is emerging as a useful approach where different cloud services can each provide specific capabilities which can be even more effective when linked together. Business technology leaders are quickly coming to a point of understanding that a single cloud solution provider may be challenged to meet the needs of their entire organization.




Figure 4: Customer Links One Cloud Service to Another Cloud Service

A cloud service linking to another cloud service can be illustrated using a SaaS-to-PaaS example. In this example, the SaaS application delivers the functionality/feature sets that an organization needs as a business solution; however, the SaaS application may not have the advanced functionality to deliver analytics and business intelligence. The organization can leverage additional capabilities by using a PaaS service from another cloud service provider and develop a custom analytics application that consumes data from the SaaS solution via an API, combining it with other data sets, which could drive additional revenues or further market differentiation.

Interoperability Considerations

Linking one cloud service with another cloud service requires that the second cloud service has a well-defined API that the first cloud service can utilize remotely. It is assumed that the cloud services leverage SOA techniques (such as REST interfaces, stateless interactions) to facilitate invocation. The impact of performance variations on the use of the API due to network constraints need to be considered.

The category of the first cloud service and the category of the second cloud service influence the nature of the connection between the two services. Where the first cloud service is either an IaaS or a PaaS

service, then the application code running in that service belongs to the cloud service customer and the main concern is that this code can successfully utilize the API of the second cloud service. It is likely that the IaaS or PaaS platform will be capable of supporting invocations of remote service APIs, although there may be issues to consider relating to security capabilities such as authentication credentials and encryption technologies.

If the first cloud service is a SaaS offering, the situation is more complex, since in this case, the application code belongs to the cloud service provider. In this case, the application code of the first cloud service must be structured to enable the use of the API of the second cloud service. For this to be possible, there is probably a requirement for the API to be standardized in some way (possibly a “well known” de-facto standard rather than a formally standardized one). In addition, there will be a need for the customer to be able to configure the first cloud service to use the second cloud service.

Regarding the second cloud service - if it is an IaaS or PaaS service, then the code running there belongs to the customer and the customer is in control of the API which it offers to the first service. The more interesting case is where the second cloud service is a SaaS service, where any API is dictated by the cloud service provider. In this case, it is important that the API is offered using standard technologies - ideally the whole interface is defined by a standard, but in the common case where it is not standardized, then at least the basic protocols should be standard (e.g. use of REST/JSON or REST/XML protocols and data formats). In the common case where the API is not standardized, there is a danger that the customer will get locked-in to the particular cloud service and find it difficult and expensive to move to a different provider - as a result the customer should consider mitigation techniques (see the [Interoperability & Portability Challenges section](#) on page 6).

Portability Considerations

This specific scenario does not involve moving or transferring an entity from one system to another, therefore, portability is not an issue. Refer to scenarios 1 and 2 above for portability considerations relating specifically to cloud services.

Admin Interface Considerations

For the customer to administer cloud services from different providers means that multiple admin interfaces need to be taken into account. Administering multiple providers can mean duplication of effort and the need to adapt to different interfaces from each provider. Refer to scenario 2 above for a thorough discussion on admin interface considerations.

Business Interface Considerations

For the customer to utilize cloud services from different providers requires that the customer deal with the business interfaces offered by each cloud service provider. Refer to scenario 2 above for a thorough discussion on business interface considerations.

Security Considerations

From a security perspective, linking cloud services from different providers requires not only appropriate security measures for each cloud service individually, but also requires security measures to be applied to the connection from the first cloud service to the second cloud service. The connection

between the two services will require appropriately strong ID and Access Management capabilities to be applied. This will require that the technology for this is supported at both ends of the connection (i.e. cloud service 1 must be able to send appropriate credentials when it invokes cloud service 2). It may also be required to encrypt the data sent between the two cloud services - this requires that they mutually support encryption technology of the right strength.

The support of third party ID and Access Management functionality, where it is possible that the IdAM system could be one installed and operated by the cloud service customer and where that system is used by the different cloud service providers, would significantly reduce the need to port/duplicate user security information.

Recommendations:

- Refer to recommendations for Scenario 2.
- Ensure that services provided by the cloud service providers leverage SOA design principles and can utilize and expose APIs to enable interoperability.
- Consider the use of an intermediary (an “inter-cloud provider”) to help address and solve the issues of integration, portability and interoperability of multiple cloud services.
- Ensure that the security technologies supported by the second cloud service are usable by the first cloud service when it uses the capabilities of the second cloud service.

Scenario 4: Customer Links In-house Capabilities with Cloud Services

As more enterprises consider their 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. In addition, enterprises will also have to assess how other in-house capabilities will be leveraged in their cloud strategy. These capabilities include people, processes and of course technology. Developing an ‘in-house’ view of cloud adoption based on these critical criteria is the challenge that enterprises face during the early cloud computing adoption phase and continues as more workloads and projects are committed to cloud services.

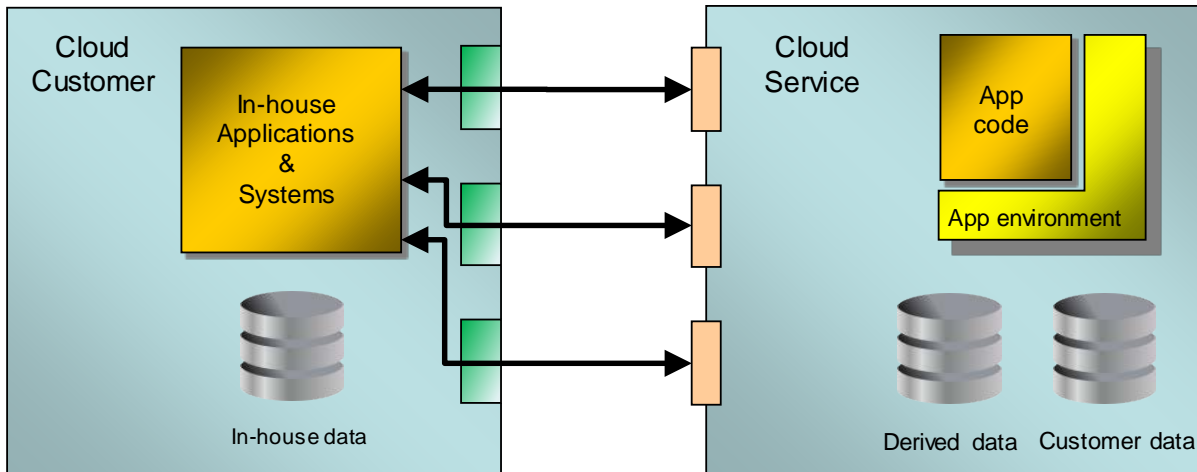


Figure 5: Linking In-house Capabilities to Cloud Services

As new cloud services are deployed, the need to connect them with various on-premises applications and systems becomes important. Cloud service owners 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 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 continue to be monitored and managed by an existing suite of on-premises IT tools.

Refer to the CSCC whitepaper “Migrating Applications to Public Cloud Services: Roadmap for Success” for a detailed discussion on integration considerations for connecting cloud services with on-premises services.

Interoperability Considerations

Linking in-house capabilities with cloud services requires that on-premises functionality and data needed by the cloud service are clearly identified and vice versa. For each of these functions and data sources, there must be a well-defined API in place that can be utilized remotely. If the on-premises applications and cloud services leverage SOA techniques (REST interfaces, stateless interactions, etc.) then the integration effort should be reduced. If not, the impact of redesigning the applications and services to provide suitable interfaces could be significant. In addition, the impact of performance changes due to network constraints need to be considered as part of the integration.

For SaaS services, the cloud service provider must make available the necessary information (such as API descriptions and security requirements) to address functional integration requirements since the provider controls the cloud service. For PaaS services, the customer should be able to address most of the functional integration requirements, since the app code running in the cloud service is controlled by the customer, although provider assistance may be needed to address middleware specific integration requirements. For IaaS cloud services, the cloud customer should be able to fully address all functional integration requirements since it is expected that the interfaces offered and used by the app code running in the cloud service are controlled by the cloud service customer.

Portability Considerations

This specific scenario does not involve moving or transferring an entity from one system to another, therefore, portability is not an issue. Refer to scenarios 1 and 2 above for portability considerations relating specifically to cloud services.

Admin Interface Considerations

Linking cloud services with on-premises capabilities requires a thorough examination of the monitoring, management and support requirements for both environments as an integrated or holistic solution. Processes that are equivalent in both environments should have monitoring, management and support models that support post-production deployment, and should link in-house technology with cloud-based technology to formulate a unified solution.

Existing support models and processes must be revisited to determine if there is a need to modify the current processes to address potential disparities between the on-premises components and the cloud services. Critical functions such as backup/recovery, disaster recovery, fail-over and high availability requirements must be re-examined with cloud services as a part of the scope and planning process. Maintenance plans and support processes must also be re-visited to determine how the integrated solution will be monitored, managed and maintained on ongoing basis.

Ideally, a consolidated monitoring and management platform is possible as production environments require constant maintenance and support to ensure performance, security and availability to meet the expectations of the business units and their user base. Can the management tools used in-house still be used, or is it necessary to adapt to new monitoring and management facilities supplied by the cloud service? The answer to these questions will depend on the admin interface offered by the cloud service(s) and in particular whether the interface conforms to existing interoperable standards, which can be used by the in-house tools and systems. It may be necessary to consider adapter code or the use of mapping capabilities such as an ESB to facilitate the integration in cases where the cloud service admin interface is not directly interoperable with the in-house systems.

Business Interface Considerations

Business interfaces involve capabilities relating to the business aspects of the cloud service including subscription information, billing and invoicing. In many cases, these business interfaces do not exist for on-premises applications. As a result, the introduction of cloud services may require customers to adopt new business related user interfaces and/or APIs defined by the cloud service provider possibly involving

the purchase or development of in-house systems for the management of business capabilities. In cases where on-premises business management systems do exist, cloud customers may need to adopt new tools and systems, or acquire adapters to match existing systems to the interfaces offered by the cloud service provider.

Security Considerations

From a security perspective, linking in-house capabilities with cloud services requires strong authentication and authorization services to ensure proper access is being granted to potentially sensitive services and data. The support of third party ID and Access Management, where it is possible that the IdAM system used to control access to the application running in the cloud service could be an existing one installed and operated by the cloud service customer, would significantly reduce the need to port or duplicate user security information.

Given that the connection to the cloud service is likely to traverse the public internet, consideration should be given to what encryption should be applied to that connection – and should any customer data stored in the cloud service also be encrypted. Not only must the cloud service support the necessary encryption capabilities, but items such as certificates and encryption keys must be managed in a way compatible with the customer’s security policies. Technologies such as Virtual Private Network can also be considered for securing access to the cloud service.

In the case where the cloud service accesses APIs or data that are supplied by in-house systems, there is a further, very significant set of security considerations that must be addressed. The APIs and/or data access implies the creation of new, publicly exposed interfaces, which need careful control since they offer a potential new attack surface. Access control, firewall configuration, denial-of-service countermeasures, encryption techniques must all be considered in relation to these new APIs. The deployment of suitable API Management capabilities may be part of the response to these considerations.


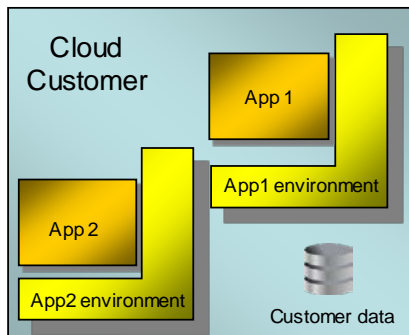
Recommendations:

- Refer to recommendations for Scenario 1.
- Ensure that on-premises applications are leveraging SOA design principles and can utilize and expose APIs to enable interoperability with remote cloud services.
- Examine whether existing in-house systems are available to deal with the business aspects of using cloud services - 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.
- 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.

Scenario 5: Migration of Customer Capabilities into Cloud Services

This scenario addresses the case of a customer currently running an application or service on-premises who moves that capability to a public cloud environment. This use case is a key one from the perspective of the customer since cloud computing offers several benefits, including cost savings and business transformational improvements that can have a positive impact.

before...



after...

Figure 6: Migration of Customer Capabilities into Cloud Services

Assessing applications and workloads for readiness for migration to a cloud service is required for organizations to determine which applications and data can – and which cannot – be readily moved to a public cloud environment and what service models (IaaS, PaaS, or SaaS) can be supported. It often makes sense to start with the lowest-risk applications—those with minimal customer data and other sensitive information—or applications that can take advantage of the elasticity of cloud computing.

The table below highlights suitable and less suitable types of applications for migration to cloud computing. Refer to the CSCC whitepaper “Migrating Applications to Public Cloud Services: Roadmap for Success” for a detailed discussion on the steps customers should take to ensure successful migration of existing applications to cloud computing.

Suitable Candidates for Cloud	Less Suitable Candidates for Cloud
<ul style="list-style-type: none"> • Applications that are used by a group of mobile workers to manage their time and activity, and that contribute only limited info to the company's broad management information databases. • Applications that are run infrequently but require significant computing resources when they run. • Applications that are run in a time zone different from that where your company's IT personnel are located. • Development, testing and prototyping of application changes, even if the final applications will be run on your own infrastructure. • Service Oriented Architecture (SOA) applications. 	<ul style="list-style-type: none"> • Applications that involve extremely sensitive data, particularly where there is a regulatory or legal risk involved in any disclosure. • Applications now being run on the company's private network and that are very performance-sensitive. • Applications that require frequent and/or voluminous transactions against an on-premises database that cannot be migrated to a cloud service. • Applications that run on legacy platforms that are typically not supported by cloud providers.

Interoperability Considerations

For SaaS cloud services, the application code belongs to the cloud service provider. In this case, migrating an on-premises application or service to a public cloud service provider does not involve porting the application code (i.e., the on-premises application is being replaced by the application offered as a cloud service). What is important in the SaaS case is the compatibility of the functional interface for the application – and, in particular, any user interfaces presented to end users and also any APIs made available to other customer applications. It is probably unrealistic to expect that user interfaces provided by the SaaS application will be identical to the on-premises application; however, it is reasonable to expect that similar functionality will be presented in a broadly similar way to reduce the cost and effort of retraining end users.

Any functional APIs made available by a SaaS service should in the ideal case be the same as the interface provided by the on-premises application or service that is being replaced (i.e. the APIs are **interoperable**). In the common case where the APIs are not interoperable, then the implication is that any customer applications using the APIs will need to be changed as part of the process of replacing the on-premises application with a SaaS equivalent. Depending on the extent of customer applications affected, to minimize changes, it may be advisable to create a common mapping layer that converts old API calls to the format required by the new SaaS offering or deploy an Enterprise Service Bus (ESB).

In the case of PaaS and IaaS cloud services, the user interfaces presented to end users and the APIs presented to external applications are likely to be the responsibility of the application code - and this code is the responsibility of the cloud service customer. However, these interfaces may be impacted when the on-premises application is migrated to a public cloud service due to the nature of the network connection from the on-premises location to the cloud service. If the on-premises application already leverages SOA techniques (REST interfaces, stateless interactions, etc.) then the impact of migration should be minimal. If not, the impact to other applications invoking the migrated application or service could be significant. In addition, the impact of performance changes due to network constraints need to be considered as part of the migration.

Portability Considerations

For SaaS services, it is typical that the format and the content of the cloud service customer data is in the hands of the cloud service provider. Thus, data portability is a major consideration when migrating an on-premises application to a public cloud SaaS service. Ideally the data format and data content (extent and semantics) should remain the same after SaaS migration. Data portability can still be achieved even if the data formats are different between the on-premises application and the SaaS application since there are straightforward standard tools that can be used to perform data transformations – and where there are not standard tools, it may be straightforward to build a custom tool. Differences in the extent or in the semantics of the data are much more serious and could be a major barrier to achieving data portability, particularly if the SaaS application requires more data than is available from the on-premises application, or if the SaaS application has a different meaning for some of the data.

The impact of migrating an on-premises application to a PaaS cloud service should have minimal impact on *data portability* unless data format and data content need to be altered as part of the migration. There may be exceptional cases where a database provided as part of the PaaS environment may be sensitive to the data format of the customer data, although there are generalized formats (CSV, XML, etc.) which are supported by many types of database. How data is loaded into a PaaS cloud service and how it is retrieved needs to be examined by the customer. Migration from an on-premises data store to a PaaS data store may involve data portability questions that need to be answered.

For PaaS migration, since the application code belongs to the customer the question of *application portability* becomes important – what does it take to move the on-premises application to the PaaS cloud service? One of the most important factors for application portability is represented by the **App environment** shown in Figure 2 above. In effect this represents the "API" that the cloud service presents to the application code – and the application code must be able to use this API in order for the application to run. In the ideal case, application portability implies that the application code which runs on-premises will run on the PaaS service without any changes (i.e., the App environment is compatible). If the App environments differ, the application code will need to be changed to account for the differences.

The requirements for migrating an on-premises application to an IaaS service tend to be lower since the entire software stack is migrated: the application code itself, plus any supporting code it requires – potentially including the underlying operating system. To achieve this, it must be possible to package the complete software stack as one or more virtual machine (VM) images, which can then be imported into the cloud service and executed there.

Whether the software stack involved will work in a virtual machine environment may depend on whether there is use of specialized device drivers or hardware devices that are unlikely to be supported by an IaaS cloud service provider; an application depending on these capabilities is not a good candidate for migration.

Assuming the software stack will work in a virtual machine environment, the question of *application portability* becomes less of an issue for migration to an IaaS service. That is, the on-premises application should run on the IaaS service with few if any changes. For IaaS services, the data format for cloud service customer data is usually in the hands of the customer, since the facilities provided by the cloud service are typically relatively low level, such as providing volumes for binary file or object storage (i.e. the cloud service does not know or care about the detailed format of the customer data). As a result, the impact of migrating an on-premises application to an IaaS cloud service should have minimal impact on *data portability*.

One consideration which applies to the migration of customer application code to either a PaaS or IaaS cloud service is whether the code is capable of taking advantage of the capabilities of the cloud service. It may be possible to port the code to the cloud service and get it to run successfully, but to enable capabilities such as elasticity and scaling may require recoding or redesign. However, this may be handled as part of a phased migration process, where full exploitation of the capabilities of the cloud service is part of the later phases, following migration of the basic functionality of the customer application.

Admin Interface Considerations

The management and monitoring of the migrated application running in the cloud service must be considered including the capabilities to deploy and configure the application, modify the resources assigned to the application (CPU, storage, etc.) and monitor the application's usage and status. Can the tools used in-house still be used, or is it necessary to adapt to new monitoring and management facilities supplied by the cloud service? These considerations could be some of the most significant when migrating an application to a cloud service, since it is likely that the in-house facilities for monitoring and control will not match the equivalent facilities available with the cloud service.

These admin interfaces may involve web applications or other visual interfaces and they may involve APIs. Migrating an on-premises application to a cloud service requires that the admin interfaces are compatible (particularly in the case of visual interfaces) and also interoperable (particularly in the case of APIs). If not, significant and potentially costly adaptation will be required. In the most extreme cases, it may be necessary to change the customer systems used for the various administration tasks, introducing new systems capable of interacting with the facilities made available with the cloud service.

Business Interface Considerations

Business interfaces involve capabilities relating to the business aspects of the cloud service including subscription information, billing and invoicing. It may well be the case that such capabilities do not exist at all for an on-premises application – and even where they do exist, it is likely that the interface(s) provided to support them will not match the equivalent interface(s) made available by the cloud service provider.

Moving an on-premises application to a cloud service provider requires compatible and/or interoperable business interfaces to ensure that the tools or program components used by the cloud service customer for business related capabilities can be used successfully following the move. In the most extreme case,

this may involve the customer adopting new tools and systems, or acquiring adapters to match existing tools to the interfaces offered by the cloud service provider.

Security Considerations

Refer to scenario 4 above for a thorough discussion on security considerations.

Recommendations:

- Refer to recommendations for Scenario 4.
- For SaaS, consider compatibility with on-premises applications and the migrated cloud service. Ensure user interfaces, APIs, protocols and data formats are well defined for migrated cloud services. Whenever possible, insist on standard APIs, protocols and data formats.
- For PaaS, ensure that the application environment (web server, database server, etc.) supported by the cloud service provider is compatible with your on-premises application environment.
- Examine the cloud service provider interfaces for administration and business capabilities and ensure that they can be used directly or integrated with existing or new in-house systems.

Summary of Key Considerations and Recommendations

Here is a summary of the key considerations and recommendations for cloud service customers as they assess the interoperability and portability strengths and weaknesses of potential cloud service providers:

- Ensure that on-premises applications are leveraging SOA design principles and can utilize and expose APIs to enable interoperability with remote cloud services.
- Consider implementing an Enterprise Service Bus (ESB) to perform interface, protocol and data transformations to address differences between on-premises applications and cloud services as well as differences between cloud services from different providers.
- For SaaS, ensure user interfaces, APIs, protocols and data formats are well defined for cloud services:
 - Whenever possible, insist on standard APIs, protocols and data formats
 - Consider compatibility of on-premises applications and the cloud services
 - Consider compatibility of cloud services provided by different cloud service providers
- Address the following challenges for PaaS cloud services:
 - For application migration purposes, ensure that the application environment (web server, database server, etc.) supported by the cloud service provider is compatible with your on-premises application environment.
 - For portability between cloud services, ensure that the application environment is based on open technologies to increase the number of viable alternative cloud service providers which can facilitate migration if a change in provider is warranted.

- For IaaS cloud services, ensure that the cloud service accepts standard or widely accepted application packaging formats such as OVF and that any interfaces and APIs are open and/or standard.
- Insist that your cloud service provider supports key open technologies (open standards and/or open source) for admin and business interfaces:
 - Common, open interfaces makes it easier to support multiple providers simultaneously
 - Common, open interfaces makes it easier to switch from one provider to a different provider
- Address the following security challenges:
 - Leverage the support of third party ID and Access Management functionality to authenticate and authorize access to cloud services - ideally either a system owned and run by the customer or equivalent capability provided by a supplier, which could itself be a cloud service.
 - 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.
 - When using multiple cloud service providers, ensure that the security technologies supported by the second cloud service are usable by the first cloud service when it uses the capabilities of the second cloud service.
- Examine whether existing in-house systems are available to deal with the business aspects of using cloud services - 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).
- Examine the cloud service provider interfaces for administration and business capabilities and ensure that they can be used directly or integrated with existing or new in-house systems.
- Ensure key interoperability and portability requirements are included in your cloud service agreement.
- Consider the use of an intermediary (an “inter-cloud provider”) to help address and solve the issues of integration, portability and interoperability of multiple cloud services.

Standards for Interoperability and Portability

To date, most of the focus for cloud interoperability and portability standards has been at the IaaS layer although activity at the PaaS level is starting to accelerate. In addition, there are several security standards that enable and facilitate cloud computing interoperability even though they are not exclusive to cloud computing. Cloud computing customers should determine the level of support for the following standards by prospective cloud service providers. Lack of support for these standards is likely to result in interoperability and portability challenges down the road.

- [Open Virtualization Format \(OVF\)](#). A packaging standard developed by the Distributed Management Task Force (DMTF) that is designed to address the portability and deployment of virtual machines.
- [Cloud Data Management Interface \(CDMI\)](#). A standard defined by the Storage Networking Industry Association (SNIA) that defines the functional interface that applications will use to create, retrieve, update and delete data elements from the cloud.
- [Open Cloud Computing Interface \(OCCI\)](#). A set of open specifications delivered through the Open Grid Forum that defines a protocol and API for all kinds of cloud computing management tasks.
- [Topology and Orchestration Specification for Cloud Applications \(TOSCA\)](#). A standard developed by OASIS that enables the interoperable description of application and infrastructure cloud services, the relationships between parts of the service, and the operational behavior of these services (e.g., deploy, patch, shutdown).
- [Cloud Application Management for Platforms \(CAMP\)](#). A standard developed by OASIS that defines an interoperable protocol that cloud implementers can use to package and deploy their applications.
- [Cloud Auditing Data Federation \(CADF\)](#). A standards developed by DMTF that defines open standards for cloud auditing.
- [LDAP](#), [OAuth](#), [OpenID Connect](#) and [SAML](#). Standards that enable third party ID and Access Management functionality.
- [US FIPS 140-2](#). Standard that specifies the security requirements to be satisfied by a cryptographic module utilized within a security system protecting sensitive information.

In addition to standards, there are a number of open source projects that are having a positive impact on cloud computing interoperability and portability. Open source projects that have an open governance model (i.e. not controlled by a single company) and attract a broad supporting ecosystem are the best candidates for creating de facto standards. In the IaaS space, [OpenStack](#) is an example of an open source project that is building significant industry momentum. Open source projects for PaaS are emerging. Examples of PaaS open source projects that are starting to build significant industry support include [Cloud Foundry](#), [Heroku](#), [OpenShift](#) and [Docker](#).

Works Cited

- [1] European Journal of ePractice. *Three Dimensions of Organizational Interoperability*.
<http://www.epractice.eu/files/6.1.pdf>
- [2] Cloud Standards Customer Council (2011). *Practical Guide to Cloud Computing*.
www.cloud-council.org/2011_Practical_Guide_to_Cloud%20Computing.pdf
- [3] Cloud Standards Customer Council (2012). *Practical Guide to Cloud SLAs*.
www.cloud-council.org/2012_Practical_Guide_to_Cloud_SLAs.pdf
- [4] Cloud Standards Customer Council (2013). *Public Cloud Service Agreements: What to Expect and What to Negotiate*
<http://www.cloud-council.org/PublicCloudServiceAgreements2.pdf>

Additional References

Cloud Technology Partners. Driven by Multi-Cloud, the Public IaaS Market Begins to Equalize
<http://www.cloudtp.com/2014/06/27/driven-multi-cloud-public-iaas-cloud-market-begins-equalize>

Managing Elasticity Across Multiple Cloud Providers.
http://hal.inria.fr/docs/00/79/04/55/PDF/Managing_Elasticity_Across_Multiple_Cloud_Providers.pdf

US Government Cloud Computing Technology Roadmap Volume II Release 1.0 (Draft) – Useful Information for Cloud Adopters
http://www.nist.gov/itl/cloud/upload/SP_500_293_volumell.pdf

Appendix A: Interoperability Model for Cloud Computing

When considering real-life interoperability between two systems, it is useful to have a model for the information exchanged, since there are in practice a number of aspects to the exchange that are relevant when connecting systems.

There are a number of interoperability models –one that is most relevant to the customer use of cloud services is the 4 level model described in work conducted for the European Union [1].

#	Layer	Aim	Objects	Solutions
1	Technical	Technically secure data transfer	Signals	Protocols of data transfer
2	Syntactic	Processing of received data	Data	Standardized data exchange formats, e.g. XML
3	Semantic	Processing and interpretation of received data	Information	Common directories, data keys, ontologies
4	Organizational	Automatic linkage of processes among different systems	Processes (workflow)	Architectural models, standardized process elements

In the EU 4-level model, exchange of information between two systems can be considered as taking place on 4 levels which are layered so that the higher layers utilize the lower layers:

- *Technical interoperability* is associated with the protocol used for information exchange – for example the use of the REST HTTP protocol over TCP/IP. The concern is the basic exchange of data between some endpoints.
- *Syntactic interoperability* concerns the format of the data that is exchanged – examples here include XML data structures or JSON data streams. The concern is the data being exchanged.

- *Semantic interoperability* concerns the meaning and structure of the data exchanged – examples here can include XML schemas (for the structuring) alongside directories or ontologies that describe the meaning of elements in the structures.
- *Organizational interoperability* concerns the context in which the data is exchanged – that is, the sending system has an expectation that the receiving system will use the exchanged data in a specific way, typically as part of some larger overall process.

Ideally, two interacting systems will be interoperable on all 4 layers. However, practically speaking, two systems can interact successfully even if interoperation is not possible on all 4 layers. For example, at layer 1, system A may communicate using the REST HTTP protocol while system B may communicate using the MQTT protocol. It may be possible to get them talking to each other by using a protocol adapter (such as an ESB).

Similarly, if the two systems differ at the Syntactic level, it may be possible to enable them to interoperate using a syntax translator – an example may be a mapping between data encoded in XML versus data encoded in JSON.

Two systems which don't interoperate at the Semantic level are likely to be much more problematic. The implication is that the two systems have different data schemas – and it may well be the case that data from one system has no meaning or is unusable on the other system. It may not be possible to create an adapter that can be used to enable the two systems to communicate successfully.

For interoperability at the organizational level, the requirement is that the processes or activities on the customer systems are matched with the processes or activities of the cloud service. Lack of interoperability at this level is likely to be highly problematic – the implication is, for example, that the provider's systems don't have capabilities expected by the customer's systems. This can be reflected in a mismatch of APIs – for example, the set of operations exposed by the provider API may be different from the set of operations expected by the customer systems.