

Special Paper on Cloud Computing: Portability and Interoperability

“Cloud Switching” and the free flow of data – portability and interoperability of software and data across cloud services

When deciding how to implement and deploy cloud computing technologies, enterprises must consider how easy it will be to access and move software and data components to adapt to current and future needs. This document outlines aspects of interoperability and portability in the context of cloud computing. It also explores attributes of cloud computing that address issues that enterprises face in accessing and moving software and data in cloud computing, as well as provide some background and context for making informed decisions. Particular consideration is given to “hybrid cloud computing”, an increasingly common paradigm that enterprises are adopting for their cloud solutions. Hybrid cloud solutions involve a mix of multiple cloud services or resources that can be deployed to both public and private systems, often in combination with existing enterprise systems.

INTRODUCTION

Previous ECIS papers¹ provide a basis for understanding important aspects of cloud computing, and to highlight potential issues that can restrict the value of cloud solutions. As interest in cloud solutions has grown, technology has evolved to broaden the scope of where cloud services can be used. This evolution has expanded capabilities in many ways. For example, many enterprises seek to maximize their ability to adapt to changes in future needs while simultaneously being able to leverage existing investments in software and data. These different capabilities often have unique dependencies and requirements, and vendors have introduced many enhancements designed to improve the ability of their cloud platforms to meet such requirements. This “multi-dimensional” approach has led to the concept of the “hybrid cloud”.

As described in previous ECIS documents², even relatively homogenous cloud platforms can manifest limitations that can potentially impact competition and customer choice within the Information Technology market merely by the manner in which a cloud service is architected and/or deployed. Also, the importance and value of *data* itself has continued to increase, in contrast to the software which operates on the data; this is one of the main reasons for the interest in adapting cloud solutions to leverage existing infrastructure. With the evolution into an increasingly heterogeneous model, both in terms of software as well as data, hybrid cloud solutions require additional considerations in order to realize the promise of cloud computing.

¹ <http://www.ecis.eu/>

² Ibid.

27 June 2016

THE DEFINITION OF “HYBRID CLOUD COMPUTING”

Definitions are usually of paramount importance to understanding technology. As is often the case with a new idea, especially one that is popular, it can be difficult to develop consensus for the term ‘hybrid cloud’. The word ‘hybrid’ literally means “a thing made by combining two different elements”³. However, such a broad definition leaves significant opportunity for confusion (and abuse). Fortunately, there are efforts to provide common definitions for cloud computing that reduce ambiguity, including ‘hybrid cloud’:

“Hybrid cloud is a deployment model using at least two different cloud deployment models”⁴

A cloud deployment model is a way in which cloud computing can be organized based on the control and sharing of physical or virtual resources, and includes public, private and community clouds. With hybrid cloud, the cloud deployments remain unique entities, but are bound together by standardized or proprietary technology that enables data and application portability. For many uses, hybrid cloud is a combination of private cloud and public cloud deployments, often including integration with existing on-premises enterprise systems.

With this definition as a basis, an appropriate understanding of relevant concepts is easier to develop. Some examples of such related concepts include *portability* and *interoperability*.

PORTABILITY AND INTEROPERABILITY

Two related, but very different, attributes of a component within a computing system are portability and interoperability. Portability is “the ability of software or data to be *transferred* from one machine or system to another”⁵. Interoperability is “the ability of two or more systems or applications to exchange information and to mutually use the information that has been exchanged”⁶. Portability defines the ability to physically *move* software or data from one system to another. Interoperability defines the ability to *interact* between systems via a well-defined interface to obtain predictable results; the software or data continues to reside on the same physical machine after the interaction. Each of these capabilities is important for any given cloud solution, depending upon the specific business requirements.

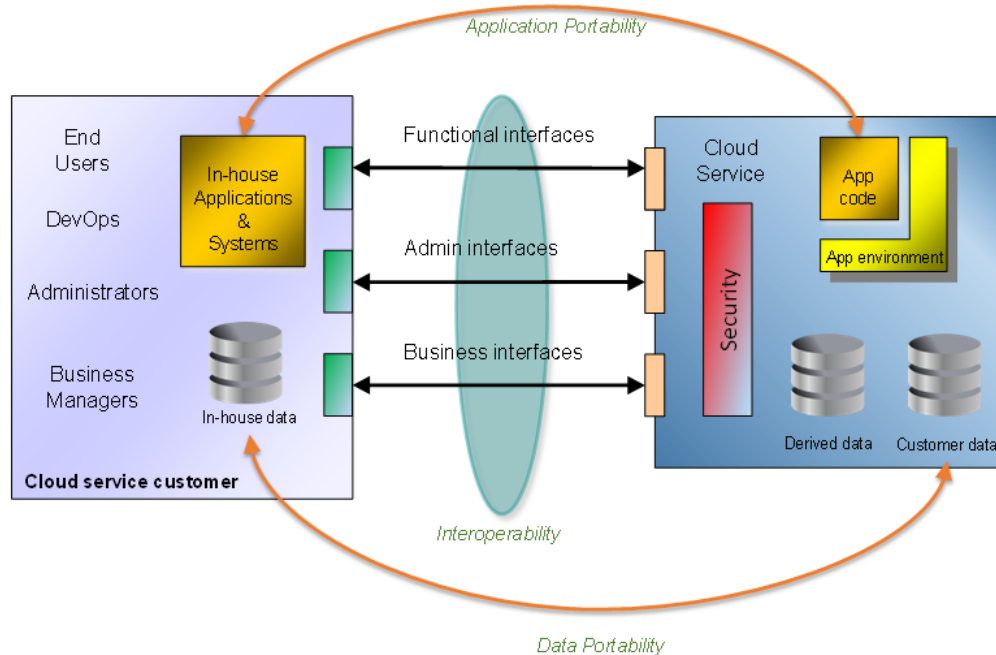
³ <http://www.oxforddictionaries.com/definition/english/hybrid>

⁴ ISO/IEC 17788:2014: http://standards.iso.org/ittf/PubliclyAvailableStandards/c060544_ISO_IEC_17788_2014.zip

⁵ <http://www.oxforddictionaries.com/definition/english/portability>

⁶ [ISO/IEC 17788:2014, 3.9](#)

27 June 2016



Elements of Interoperability and Portability for Cloud Services

For example, an enterprise may have a large database of information within their own enterprise that they wish to use with a new cloud solution, but they want to continue to maintain control of their valuable data within their own datacentre. For such a solution, interoperability between the cloud service and the database is of prime importance. The same enterprise may also be building another cloud solution for a brand new market they are considering expanding into. This new market is nascent, so the enterprise wants to minimize the amount of funds they risk for their exploration. For this solution, a public cloud service can offer an attractive value because only the resources actually used must be paid for and the amount of resources used can be scaled to match the exact number of customers using the solution. However, if the new business proves successful, the enterprise may wish to be able to move the software from the public cloud into their own datacentre at some point in the future. For this solution, portability of the software and/or the data between the public cloud platform and their private cloud infrastructure is of prime importance.

SOFTWARE PORTABILITY VS. DATA PORTABILITY

Portability is often discussed for software (or “applications”), but the portability of data can also be of significant importance. Changes to business strategy, market successes or failures, and new regulatory requirements may require relocating data to different geographic locations. Another consideration is the sheer volume of data that is created; some datasets may grow so large that it makes more sense to move the *software functionality* that uses the data instead of the traditional model of moving the data to the

27 June 2016

software. Thus the ability to move software *and* data individually has inherent value.

Software portability is usually affected by the dependencies the software has on other computing components such as “middleware” software runtimes, operating systems, databases or hardware architectures. Software portability is made more difficult in cases where these dependencies change in some way when moving from the source system to the target system, for example a change in the hardware architecture or a change of operating system.

In a similar manner, data portability is affected by the format and content expected and supported by the target system. This is sometimes influenced by the nature of the data storage system, but more commonly it is influenced by the software that interacts with the data. How intimately data is related to the software that uses the data depends upon a number of factors, and is primarily determined by architectural design decisions made during software development. For cloud computing, these decisions may be constrained by the type of cloud service model⁷ being used. IaaS and PaaS solutions can utilize technologies such as file systems and DataBase Management System (DBMS) technologies, which must also be available in the target system in order for data to be portable between one cloud service and another. SaaS solutions by definition contain the software that operates on data, and this software usually dictates the content and format of the data used by the cloud service.

Further, it is critical to understand that the portability of software or data is not a simple “true or false” metric. The level of ‘portability’ (as a measurement) can be defined as the “amount of effort required to move from one system to another”. For software, if the source and destination systems are running the same operating system and hardware the “portability” is likely to be easy, perhaps as simple as copying some files. However, if the source and destination systems run a different operating system on different hardware, it may take considerable effort for a large team of people to make the changes necessary for the software to work on the destination system. Both examples are technically “portable”, but obviously one is more “portable”, and therefore will cost much less to move. The more extreme scenario is typically referred to as ‘software migration’.

Data portability can be measured across a similar spectrum, depending on the capabilities of the file system or data management software and also depending on the requirements of the software interacting with the data on the target system. If the source and destination cloud services store the data in a common format (“common syntax”) the data is likely to be easily transferred between them, perhaps as simple as copying some files between the source and destination cloud services. Other cloud services may require an “export/import” process, where data is converted to a format suitable for the destination cloud service- which may be a standardized interchange format. Such a process can take a long time, especially for large data sets. In addition to the data format, there is the question of whether the data has the same *meaning* between the source and target cloud services (“common semantics”). Portability is likely to be more complex in the case where the destination cloud service requires different content than the source is capable of providing. Finally, a target cloud service may require data to be loaded via a custom interface; in this case special *software* may have to be created to move data from the source. As is the case for software, there is a spectrum for data portability that can require a variable amount of effort, cost and risk to enable.

⁷ See previous ECIS papers.

27 June 2016

PORTABILITY AND INTEROPERABILITY DIRECTLY AFFECT THE ABILITY TO SWITCH CLOUD VENDORS

Many vendors offer cloud services. However, even cloud services which offer similar capabilities don't always have identical interfaces, even though many technologies exist to enable interoperability between cloud services. Many cloud services use common interoperable protocols; for example, many cloud services have interfaces that use ReST over HTTP (or HTTPS), which is a well-known and popular standard. However, the interfaces offered by the cloud services using these protocols are not typically standardized- and there are often detailed differences in the way in which individual service operations behave.

It is differences in the interfaces offered by different cloud services that make interoperability between one cloud service and an alternative cloud service difficult. If a cloud service customer adapts their systems to work with one particular service, they may not be able to choose an equivalent cloud service from a different provider without having to adapt their systems again for the new provider; this limitation is due to the lack of interoperability.

Portability is also a critical attribute of a cloud service that directly dictates a customers' ability to choose alternative cloud service providers without making extensive investments to change or adapt the software and/or data component(s) so they can be used on a different cloud service. The cost, time and risk of adapting existing software and data must be taken into account when considering switching cloud service providers.

In order to interoperate with a service or porting data and applications between providers, many considerations need to be taken into account. Issues such as transport protocols, encoding syntaxes for communication messages or data, semantics, and organisational and legal policy issues need to be addressed. Of these, semantics is key as without a mutual understanding of the data being exchanged or ported, or without a mutual understanding of an expected outcome when two systems interoperate, interoperability and data portability will be ineffective. As such, when switching providers or porting data, it is essential to establish a mutual understanding of (i) the exchanged data and, (ii) behavioural semantics.

The ability to switch to alternative cloud service providers is not restricted solely to customers that seek lower prices or better value. Cloud service providers are also enterprises themselves and they may change their offerings based on their own business requirements, which may not align with those of their customers. Before cloud computing, enterprises owned and managed their own infrastructure, so once a system was deployed the enterprise could be relatively confident they would enjoy its functionality for the planned lifetime of that system. One of the values of cloud computing is that customers no longer have to maintain their own infrastructure, and they can pay for only the actual resources they consume. However, a significant downside to this model is that the customer also becomes fully dependent upon the capabilities offered by the cloud service provider. If the cloud service provider modifies or stops offering a feature or interface that a customer relies upon, the customer may be forced to incur significant costs to simply keep the solution they already had.

ENABLING THE FREE-FLOW OF DATA, WITHOUT RESTRICTION

In addition to being able to place and move data to the most appropriate location for existing business requirements, another aspect that should be considered is the potential *unknown* future value of data. The evolution of most technologies has at some point involved finding a completely new use for existing resources, and many of these new uses could not have been predicted. Historically it has taken considerable effort to adapt software and data to be useful in ways beyond that originally envisioned.

27 June 2016

Many technologies have been developed to attempt to improve the flexibility and agility of adapting resources to new scenarios, and this goal has been a driving force in the evolution of cloud computing since its inception. Most traditional enterprise systems – before cloud computing – have focused on adapting the *interoperability* of systems as opposed to the *portability* of data. With the advent of cloud computing the pace of technological innovation has dramatically increased. This places an ever-increasing importance on the ability to “repurpose” software and data for new usage. Restrictions in the free-flow of data via the lack of interoperability – and increasingly via the lack of portability – impede innovation and limit value.

CONSIDERATIONS FOR EVALUATING INTEROPERABILITY OF CLOUD SERVICES: FUNCTIONAL, ADMINISTRATIVE, MANAGEMENT AND BUSINESS CAPABILITIES

When it comes to interoperability, there are multiple, different interfaces between a customer and a cloud service provider, as shown in the Figure on page 2. The **functional interfaces** (which may include human user interfaces as well as programmatic interfaces) are made available to access the service being offered, and they allow the customer to perform its business activities that necessitated a cloud solution. To ensure the smooth operation of the customer's use of cloud services, and to ensure that those cloud services are running well with the customer's existing ICT systems and applications, **administration interfaces** are also provided. Finally, **business interfaces** are typically used to allow a customer to perform functions such as discovering and selecting appropriate cloud services, invoicing and paying for usage, and other financial and legal activities.

While the underlying technologies for each type of interface may be similar, different considerations may need to be taken into account. The business interfaces must work with existing in-house enterprise systems, and the administrative interfaces must work with in-house IT management systems – these are typically “off the shelf” IT management solutions. Each of these interfaces must be compatible with the in-house systems, or must be adaptable to them via the use of integration technologies. There are two challenges to interoperability here:

1. What happens when you want to change cloud service providers?
2. What happens when you want to replace an in-house system?

The more a customer can rely on standard interfaces, the more flexibility they will have in switching cloud service providers as well as switching between different in-house systems.

EXPLORING ASPECTS OF PORTABILITY AND INTEROPERABILITY RELATIVE TO IAAS, PAAS AND SAAS

Interoperability concerns, challenges and issues are largely similar regardless of the type of service being used. While functional and user interfaces are likely to change because of the nature of the service, the business functions are likely to be very similar. Administration interfaces will differ in the types of entities that can be managed, but the style of these interfaces is likely to remain the same, based on commonly used existing management systems.

When it comes to portability, the type of service being used determines the degree of software and data portability available. There is a spectrum in terms of ownership and management of hardware resources, software and data. In a pure in-house system, the customer owns and manages all of these and has total responsibility. At the other extreme, SaaS customers have no ownership and perform little management of hardware resources or software, and only have access defined by the interfaces that the provider chooses to make available. Progressing between these two points, IaaS offers the management of virtual

27 June 2016

machines, storage and networks, where the customer usually has substantial control over their data and the software components they load into the cloud service. For PaaS, there are fewer resources to manage as they are abstracted by middleware. Here the data is still largely in the customer's control, but it is likely that the customer will choose the target cloud service on the basis of the middleware features provided (e.g. a particular DBMS, or a particular runtime) in order to ensure compatibility with existing software components and datasets.

Software portability is not a concern for SaaS by its very definition (the software within the cloud service belongs to the cloud service provider) and data portability depends on the interfaces offered by the provider for loading and unloading data, including the data formats offered as well as their semantics. IaaS services usually provide straightforward software portability, since the resources provided are lower-level and support most common software frameworks. Hardware architectures can be a concern, since differences between hardware may require adapting software and/or data to the destination cloud service. For PaaS services, there is often a focus on software portability – usually the destination cloud service provider is chosen specifically for their support of the necessary software frameworks and middleware; without such support, the customer would face the substantial task of adapting their software to use different frameworks and middleware. Depending on the architecture of the software that is running on a PaaS service it may be possible to port data independently from porting the software, especially if the data is held in standardized file systems or databases.

HYBRID CLOUD SOLUTIONS PLACE INCREASING IMPORTANCE ON PORTABILITY AND INTEROPERABILITY

Throughout history, a common trend in technological evolution has been to “decompose” larger systems into smaller parts. There are many reasons for this trend, but a common one is to be able to leverage the functionality offered by the smaller part in more ways. For example, rather than having every single cloud service provide its own mechanism to authenticate users it makes more sense to have one service that authenticates users for *all* cloud services. Over time this decomposition happens continually- the parts get smaller and smaller, and are used across more and more other services. The end result is improved efficiency and operation but with an ever-increasing inter-dependency across numerous services. A failure of any one of the small parts can impact a large number of dependent services.

Another parallel trend is the creation of wholly new capabilities, which can be used to extend existing applications and solutions to cover new areas previously not possible. Recent examples include the creation of “cognitive computing” services, which are able to bring artificial intelligence techniques to bear on unstructured or semi-structured data. These sophisticated capabilities are now offered as sets of cloud services, which can be “plugged into” existing applications to give them the ability to service businesses in new and important ways, without requiring extensive new development activities.

Hybrid cloud solutions offer great agility, flexibility and value, but they also increase the inter-dependencies across applications and the many cloud services they depend upon. It is clear that as these inter-dependencies increase, the ability to move and connect all the various cloud “parts” becomes more important.

STANDARDS HOLD GREAT PROMISE TO ENHANCE INTEROPERABILITY AND PORTABILITY FOR CLOUD COMPUTING

It is clear that interoperability and portability are both very important to realize the full value of cloud computing. Having a common, standardized definition for the mechanisms that are used for different systems to exchange information has enabled countless technologies to thrive. The Internet itself would

27 June 2016

never have been possible without the TCP/IP networking standards, which allow any and all computers to connect to each other- much like any telephone in the world is capable of connecting to any other telephone. With the multitude of vendors offering cloud services, interoperability and portability will continue to grow in importance- especially as customers become increasingly dependent upon those cloud services.

Although there are already many standards defined for specific technologies used in cloud solutions, there are also efforts underway from leading standards organizations to clarify definitions, create common formats for improved interoperability and define specific contexts that vendors can adopt to improve compatibility and portability across a wide variety of cloud solutions⁸. For example, security, management and status monitoring are general concepts that are applicable- and important- across all cloud services. Differences between vendors and their cloud services in any of these areas can have significant effects on the overall interoperability and portability of *all* the cloud solutions using those services. By having common, well-defined standards across multiple cloud services enterprises will have better options and more choices in meeting their needs as technology and markets evolve.

THE DIGITAL SINGLE MARKET AND THE HYBRID CLOUD

The various components of the Digital Single Market (DSM) Strategy, which was adopted in May of 2015, include efforts to remove “constraints on the ability of individuals and businesses to move from one (on line) platform to another” and a “Free flow of data’ initiative that tackles restrictions on the free movement of data for reasons other than the protection of personal data within the EU and unjustified restrictions on the location of data for storage or processing purposes.” In addition, the DSM seeks to “Boost competitiveness through interoperability and standardization.” In a similar timeframe, the General Data Protection Regulation (GDPR) has become EU law – and while one of its major goals is to ensure appropriate treatment of personal data, its second major goal is to ensure the “free flow of data” within the EU. The GDPR contains an explicit requirement for a data subject to obtain and transfer their personal data in a “structured, commonly used, machine-readable and interoperable format”.

As this paper demonstrates, the concepts of interoperability and portability in the context of cloud computing are nuanced and require care when applying them in public policy. For example, determining whether software or data is “portable” cannot be answered with a simple yes or no. It is rather a question of determining the degree of interoperability and portability on a sliding scale, with the answer ranging between “very difficult” to “relatively easy”

Regarding interoperability, while the Commission is to be commended for exploring avenues to promote it, it is necessary to understand that this is a complex topic involving numerous types of interfaces. It is important to note that the growth of hybrid cloud deployments will increase pressure on vendors to provide better levels of interoperability. Understanding and channeling that pressure through the global, industry-led standardization process will be critical.

As a final comment, although this is a complex topic it is important to keep policy considerations regarding interoperability and portability in the cloud focused at a broad EU level. The Commission is to be commended for continuing to keep this topic on its agenda. As was apparent at the recent DG

⁸ For example, the ISO/IEC 19941 standard for Cloud Computing Interoperability and Portability is currently under active development

27 June 2016

Connect-hosted Consultation Workshop on the Free Flow of Data⁹ of May 18, 2016, a large number of cloud vendors will happily avoid this topic and its relevance, and instead seek to avoid engaging with policy makers in terms of specific measures that could be adopted.

There is still much work needing to be completed to ensure an open, thriving market. ECIS shall continue to engage on this topic.

⁹ <https://ec.europa.eu/digital-single-market/en/news/consultation-workshop-free-flow-data>