TRANSFORM

Business Transformation through Cloud Computing

The E-Health Special – Introducing Cloud SOA

Neil McEvoy, CEO, CloudBestPractices.net
neil.mcevoy@L5consulting.net

Welcome to the second issue of TRANSFORM, the e-magazine for the Cloud Best Practices Network that focuses on the powerful business transformation that Cloud computing can enable.

For this issue our goal will be to showcase two key areas: Cloud SOA, the intersection between Cloud designs and the Service Oriented Architecture, and then how this capability might be applied in EHealth.

Cloud in Healthcare

To provide this industry focus on Healthcare we'll be reviewing, and responding to, the recent Cloud strategy published by Canada Health Infoway.

Infoway is the main government body in Canada who sets EHealth technical standards, most notably EHR (Electronic Healthcare Records), and critically the integrations of multiple different systems to achieve Healthcare delivery.

The first white paper from their Emerging Technologies Series, Cloud Computing in Health, sets the scene for the role of these integrative Cloud services:

> *"The original vision for Infoway's EHRS Blueprint is very well suited to the cloud computing model. Like cloud, the Blueprint is a service oriented architecture (SOA) designed to provide interoperability and information sharing services to a broad spectrum of applications, in a highly scalable manner. Much of the e-health infostructure, such as: the Health Information Access Layer (HIAL), clinical domain systems, registry services and dedicated shared services (like consent or clinical decision support), can be considered Software as a Service (SaaS) that is offered by a ministry, health region or health care delivery organization."*

This is a visionary and insightful Cloud strategy document, and through this issue of TRANSFORM we`ll set out to provide an industry response to the key points identified, showcasing:

- **Cloud SOA –** How-To implementation models for the SOA architecture methods described, such as linking the Enterprise Service Bus to Community Cloud services for Client and Provider Registries.
- **Community Cloud Blueprints -** Hybrid cloud implementations that meet the needs of an RHA to provide a community service (such as e-referral) utilizing Federated Identity Management (F/IDM) services.
- **Big Data Analytics -** PaaS that can support the analytical needs of reporting, analysis, dashboards, extraction, transformation and load (ETL) and predictive analytics.
- **Cloud 2.0 –** Mobile and social enablement of legacy applications and processes.
- **Mobile Device Management**

Cloud SOA - Enabling the Extended Enterprise

What is especially intelligent and important about the paper is that it defines the transformational business benefits of Cloud in Healthcare:

- Reduced IT costs, through consolidation/virtualization of IT services.
- Reduced complexity and increased scalability
- Increased Agility, through collaborative care and mobile & social enablement.
- Reduced IT Sales and Acquisition Cycles
- Increased Measurability and Accountability

There is a keen focus on practical delivery, and especially on the possible new application scenarios it can enable, like appointment brokering and scheduling, E-Referrals, Consent Management and Secure E-mail.

This business focus and these scenarios that offer clear benefit to users are key to the success of Cloud adoption, and highlight the primary characteristics of this new trend.

A key transformation that Cloud SOA can enable is "Enabling the Extended Enterprise" – Better integration between collaborating organizations.

For example consider the concept of EHRs, the backbones of eHealth – Electronic patient data, and ask questions like Where is it stored, and Who owns it?

Traditionally IT is based on a core model where this information would be stored in the corporate data-centre, in a single dedicated database, in the hospital.

However this stems from an era when users were hard-wired to the system via green screen monitors, in the same office. Today we access information from multiple Internet applications including mobile devices from anywhere and everywhere all over the world.

Furthermore implementing one healthcare procedure requires multiple applications, as a patients' file is moved between the different specialists involved in the treatment.

Not only are there different applications within the hospital, like Laboratory Information Systems, but also there are others outside of the organizatin too, with "supply chain" partners.

For example in Canada a recent development has been the enablement of pharmacies like Shoppers Drug Mart to dispense flu shots.

What this scenario highlights is that patient care, and the associated information systems, are actually distributed across multiple organizations; our EHR is actually achieved via multiple 'fragments' of data about us, each stored in a different system unique to that part of the treatment.

Therefore the primary goal of technology is to unite these into a single logical patient view, and this is Cloud SOA.

The Personal Data Ecosystem

If we extrapolate these trends further and look at the broader long term view of Cloud, we can see the ultimate evolution towards 'Personal Clouds'.

In short the direction of iPhones et al and consumerization of technology in general will see patients come to 'own' and host their own data.

Each person will choose their preferred Cloud provider for use as a "Digital Vault" for storing all of their personal information from family photos through to legal documents like birth certificates and of course then healthcare records.

Technically as well as socially this architecture is ultimately superior to any other as it eliminates all duplicated storage of the information, locating it instead in the one place that all stakeholders would agree on: Under the users control.

Managing the Canadian Framework for eHealth to Construct a SaaS Cloud Identity

David Chartash, Research Student at Centre for Global eHealth Innovation – dchartas@ieee.org

Software as a Service Within the EHealth Cloud

Based on the Canadian eHealth Architecture Model (see Figure 1.1), the architecture of services delivered to a variety of health professionals is designed as a layered bus based approach. This model, as typified in the associated figure, can be easily constructed in the Software as a Service (SaaS) method, providing each point of service with a software interfacing with data services and databases through multiple layers via the Health Information Access Layer(HIAL) communication bus.

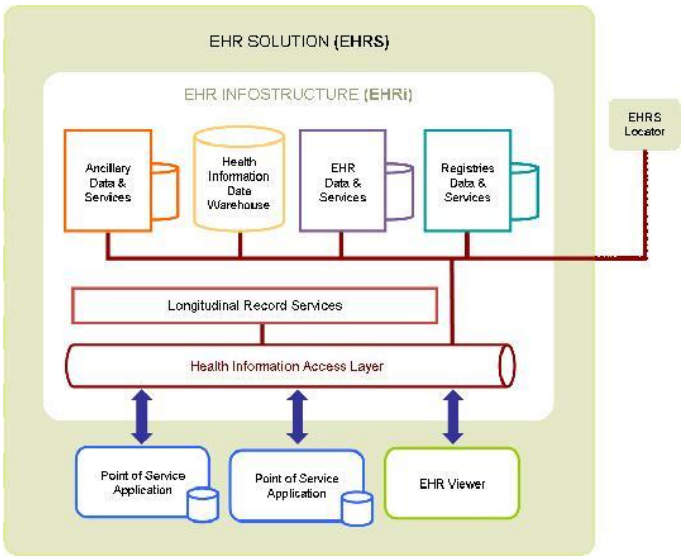

This architecture's design specifies several common services (as illustrated in Figure 1.2), all of which are required to interface through the HIAL to all data storage locations and services. In particular, these services provide for the enhanced features of an Electronic Health Record (EHR) that are often specified by clinical requirements as well as current laws regarding health information in Ontario (such as privacy and interoperability services).

Within the architecture map, these services are not specifically identified or exploded (as such as the documents specifying operation of a virtual medical record by Health Level 7 International [2]). In providing the common services identified by Canada Health Infoway (CHI), it is important to note that these services at point of service operate to improve clinical practice via informatics support. Said informatics support allows for clinical decision support, alarms and alerts, as well as interoperability with medical devices and other components of the service architecture layers (as is illustrated by the HL7 document above [2]).

Beyond understanding that the layered architecture specifications of the pan-Canadian EHR architecture, and translating this architecture to a SaaS architecture, a clear understanding of the point of service interfaces to the data services is important. For this reason, the operating process is detailed in Figure 1.3 to a limited extent, in which some of the associated actors are detailed at point of activity. Complementing this figure is Figure 1.4, which details the final point of interaction between one actor's service (built as a software tool), the EHR viewer and the data and other registries.

In the aforementioned paragraphs, the architecture and services are illustrated as identified by CHI, the framework and structural guidance non-profit for the pan-Canadian EHR network. These services, built as software components of an abstraction layer, are manageable as cloud SaaS operationalized components for integration into actor practice in each point of communication with the HIAL communication bus.

Therefore, describing the entire model for EHR architecture with an understanding that the entire architecture is intended to be provided by software as services to integrate with all data systems and registries at the hospital or provider level. From the aforementioned backbone, building a SaaS model is trivial, additionally, providing these services by a distributed cloud model allows for a simplification of the concept across multiple providers and multiple operators.

Virtualized Desktops as a Cloud-based Hospital Informatics Platform

Following the architecture identification mentioned above, building SaaS cloud based system requires a link to the cloud services on each computer accessing the system. At the physician or provider level, this would require extensive support for a complex series of

software and hardware linkages. Building virtualized desktops to provide for a software based solution (and accompanying support) to access the cloud services allows for configurability and modularity as for the device ecosystem at point of care, particularly given the wellspring of mobile computing and hybrid general and fixed purpose computing devices.

In order to develop virtualized desktops, particularly those that offer a SaaS deployment, careful consideration of the services provided and deployment culture will ensure adoption and in healthcare, safe and meaningful use of the equipment. For this reason, adherence to the meaningful use guidelines [3, 4] developed in the United States, as well as advice from clinical partners to develop informatics solutions that capture not only EHR efficiency, but clinical efficiency to improve hospital understanding of the usefulness of Health Information Technology (HIT) services.

With the above in mind, eschewing traditional business intelligence, a more important component of building a cloud based framework in ehealth is the introduction of more complex business intelligence. Said complexity is proposed as a method to integrate cloud-based health IT infrastructure into the framework of health quality and evidence produced by the Ministry of Health and Long Term Care, as a method of improving adoption and ensuring that provincial scientists and policymakers deem HIT as a major benefit to their practice (subsequent institutionalization of the services provided will ensure continued adoption and maintenance). The standard of evidence has yet to be developed, although briefly the points of consideration that should be built while watching the entry of new technologies into the market (and Health Quality Ontario's push for a greater amount of evidence based medicine contributing the deployment assessments) will follow. In the device market, this evidence primarily consists of utility measures to qualify cost/benefit ratios for incremental deployment adjustments.

For clinical informatics, this appears to be initially constructed by building an indexed searchable database of all elements of care, from text based notes and discharge summaries to vital signs, orders and diagnoses. Once a formal structure for health quality evaluation is established for informatics solutions in the "post-EHR" market, particularly via the expansion of technology assessment units providing cost-effectiveness and cost-utility analyses at both the market and hospital level, we will see a coherent strategy for specific points of business intelligence to be realized. In the interim, focusing in the market criteria displayed in the American Meaningful Use Criteria will ensure that some level of business intelligence functionality is realized prior to the market changes expected in the future.

Data Abstraction for Health Research and Cost-Effectiveness

In adopting multiple forms of data interchange protocols, particularly with respect to the production of business intelligence, the need for complete data abstraction through mediated software services/interfaces is a highlight of forward thinking towards effective quality evaluation.

If the blueprint for SaaS based services continues to construct linked abstraction layers connected to services, all querying database layers, this seems like a trivial and constructive method of linking private clouds with the academic sector, through permission optimized restrictions. That said, databases such as the patient administration database at the Institute for Clinical and Evaluative Sciences remains a major contributor to the healthcare quality evaluation from major sources of research and economic modelling. With this in mind, building into the framework of cloud services an access layer capable of pulling anonymous, decision limited data for quality research and economic modelling with regards to coding decisions and complex care parameters will only improve decision making opportunities brought about and performed through the cloud.

[1] A. C. Inc and Sextant, "A 'conceptual' privacy impact assessment (pia) on canada's electronic health record solution (ehrs) blueprint version 2," tech. rep., Canada Health Infoway, 2008.

[2] K. K. Guilherme Del Fiol, Robert Jenders and H. Strasberg, "Hl7 version 3 domain analysis model: Virtual medical record for clinical decision support," tech. rep., Clinical Decision Support Work Group, 2012.

[3] C. for Medicare and M. Services, "Medicare and medicaid programs; electronic health record incentive program stage 2," tech. rep., Department of Health and Human Services, 2012.

[4] O. of the Secretary, "Health information technology: Standards, implementation specifications, and certification criteria for electronic health record

technology, 2014 edition; revisions to permanent certification program for health information technology," tech. rep., Department of Health and Human Services, 2012.

[5] S. Ratajczak, "Electronic health record infostructure (ehri) privacy and security conceptual architecture," tech. rep., Canada Health Infoway, 2005.

[6] C. H. Infoway, "Ehrs blueprint: An interoperable ehr framework executive overview," tech. rep., Canada Health Infoway, 2006.

Guarding Against Cloud Silos With C-SOA

By Rafat Shaheen, Enterprise Solution Architect, Rackspace – raft.shaheen@rackspace.com

Cloud Service Oriented Architecture (C-SOA): Loosely-Coupled Highly-Orchestrated Service Delivery

Cloud Service Oriented Architecture (C-SOA) is an architectural approach to leverage cloud computing resources while utilizing Service Oriented Architecture (SOA) disciplines to drive substantial business value. In this convergence, SOA provides the underlying enterprise platform to consume cloud services.

Historically, as business requirements evolved, enterprises continued to deploy new systems for almost every new application suite. These systems in turn were deployed with respective servers, storage, networks and processes.

The constant additions of these systems — sometimes compounded by mergers and acquisitions — resulted in silos or islands of systems. Consequently, they are more difficult to leverage or integrate resulting in more complex tightly-coupled enterprise architecture. This complexity affected the ability to maintain these systems and was costly to change as business needs continued to evolve.

The need of this level of agility is a key aspect of survival to many companies. C-SOA has the potential to prevent building silos again – this time in the cloud, which will be far more complex and difficult to reconcile. Service reuse, extensibility, abstraction, business value and agility are, not surprisingly, common shared attributes of both SOA and cloud computing.

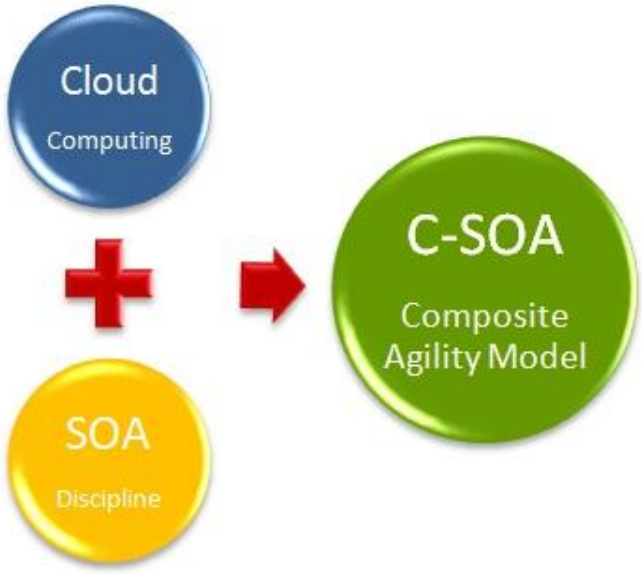

SOA is simply an architectural pattern to guide the structure and assembly of systems to deliver business value. This level of service abstraction provides a genuine focus on structure for a valid design. SOA, as a framework, allows systems to be represented and further abstracted as services. Combined with a process orchestration layer and a monitoring capability, it is possible for these services to be reused and consumed without major integration or significant development efforts.

This results in a core agility of the underlying architecture to respond to changing business requirements. With C-SOA, the focus is to bring value to business and achieving a true partnership between IT and business lines to advance the enterprise in the face of competitive pressures and changing business needs.

Some enterprises take undue risks by creating quick connections to cloud resources from enterprise systems in an ad-hoc fashion. These point solutions will further impact the realization of a true and flexible composite service layer. C-SOA advocates using a disciplined architectural approach to enterprise SOA and expanding its resource pools to include distributed cloud resources.

Metrics like reliability, performance, recoverability, availability, security and dependency affect services

delivery and should be taken well into architectural considerations.

Cloud computing with its elasticity features addresses a core SOA design issue related to ubiquitous access and service scalability. Likewise, cloud-computing will benefit from SOA's governance and sound architectural principals particularly for shared services. The results of a survey by Gartner research of CEOs and other senior business executives are depicted in the graphic below. The survey results show that 60 percent of the respondents indicated that they felt constrained by IT as they managed their enterprises out of the recent economic downturn.

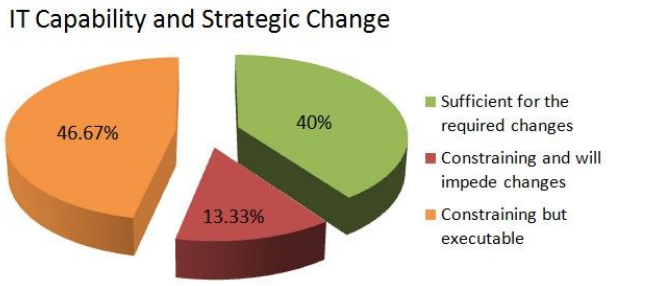

Gartner Report: Understanding and Measuring Business Value of SOA (March 2011)

The benefits of C-SOA include:

- Increased collaboration and federation
- Improved interoperability and agility
- Aligned business and IT goals and efforts
- Diversified choice of service providers
- Increased ROI while reducing IT operating cost and overhead

The characteristics of SOA services (loosely coupled, abstracted, autonomous, reusable, stateless, compose-able, discoverable and standardized) can be fully supported by cloud computing mechanisms. The core of our C-SOA discussion is rightfully the spectrum of services delivered irrespective of source or location. Services are assembled into processes which then define behavior of business applications. Processes can be further orchestrated to support business workflows for an end-to-end business operation.

Clearly, Cloud IT transformation will involve people in addition to process, and technology. To successfully implement C-SOA in the enterprise, it is imperative that a dedicated team is established to:

- Define, build and manage the creation of C-SOA shared services
- Create standards, procedures and a governance framework
- Develop a strategic plan for services to meet the business' strategic goals

This is the team that will be charged with creating an agile IT foundation to enable the business to compete, adapt and capture market opportunities. But how do you measure the agility core of an enterprise? Agility can be measured by the speed at which your changes can be implemented, the extent of the changes that can be tolerated at a given point in time and the ease of implementing those changes.

Technology Spotlight – Kantara UMA (User Managed Access)

The Canada Health Infoway Cloud strategy document calls for F/IDM (Federated Identity Management) and Consent Management capability.

An example of a relevant technology innovation is the Kantara Initiative and their UMA program. Read more in this white paper Controlling Data Usage with User-Managed Access, by Eve Maler.

A Use Case for Kantara UMA in Healthcare

By Adrian Gropper, Medical Device Developer and Privacy Advocate
agropper@healthurl.com

Introduction

State health information exchanges (HIE), having adopted standardized secure (Direct) email for their providers are now drafting RFPs for patient-authorized aggregation, discovery and transfer of health records.

They will seek identity management, record location and access management services that are simple, cost-effective and likely to be supported by EHR vendors either voluntarily or as a result of federal mandates.

Problem Scenario

From the HIE perspective, they need technology, policy and process to manage at least three data flows: (1) the matching of patient ID across independent EHRs; (2) the ability to locate the EHRs that have data for a particular patient; and (3) the list of patient authorizations that control transfer of information from one EHR to another.

The policy and process for these three functions need to follow applicable state and federal regulations and to provide meaningful privacy protections lest the patient simply "opt-out" of participating in the HIE in favor of point-to-point or "sneakernet" alternatives.

The HIE must deal with the fact that they do not have an in-person relationship with the patient (or the physician) by federating with the data holders that do, with independent identity assurance services or by engaging the patient in-person.

Current Situation

Matching of patient ID across independent EHRs is not a simple matter. Errors can result in "duplicate" patients that inadvertently hide critical information or in "merging" of two separate patients into one that can have major safety and privacy consequences. The current state-of-the-art in patient matching involves complex probabilistic algorithms and human review of a significant percentage of cases.

This need for trusted staff intervention adds cost, and hinders the ability to engage patients as stewards of their own data and authorizations. Fortunately, recent initiatives and regulations in support of secure email for both patients and providers known as the Direct Project, are making it possible to give each patient in the HIE a voluntary global identifier in the form of a Direct email address.

Direct email addresses are associated with a standard PKI certificate accessible via DNS and LDAP. Most HIEs already plan to issue Direct addresses to physicians and, with the addition of appropriate process, this capability can be extended to patients as well. There remains the issue of convincing data holders to request and store the patient's Direct email in the EHR so that they can include it in HIE transactions.

Problems

A state Health Information Exchange (HIE) faces two typical problems: participation and sustainability. Participation is an issue because the data holders (primarily hospitals and physician practices) have limited incentives to share patient information outside of their own private network and the Electronic Health Record (EHR) software vendors that serve the providers charge for each interface and therefore prefer a point-to-point architecture.

Sustainability is an issue because hospitals and providers see the cost of an HIE as an added expense, one that often does not offset any specific existing cost. In the current fee-for-service healthcare payment model, the primary beneficiaries of HIE have been seen as the patients and the payers but patients have not been in a position to purchase HIE services and payers have limited leverage over the data holders.

The situation, however, is shifting in favor of the state or regional HIE for two reasons: global payments (that put the provider at risk for patient expenses outside of their private network) are slowly replacing fee-for-service and federal incentives are driving EHR vendors to install EHRs that include information exchange capability at no additional cost.

Gaps in Other Technologies

Healthcare involves many practice groups, labs, imaging centers, pharmacies, nursing homes, schools, state and federal institutions. Some of these are as large as Social Security (disability) or a multi-$billion hospital chain but others could be a solo practice or local pharmacy.

The larger places employ tens of thousands of people in various roles and have to safely manage role-based access restrictions. The small ones are too small to enter into separate reciprocal federation agreements with the large ones. Because of this diversity, OAuth, OpenID Connect and SAML are difficult to scale.

Proposed Improvements

The ability to locate EHR data holders for a particular patient and to manage the authorizations that control transfer between the data holders and requesters is where UMA comes in.

Some momentum in favor of UMA comes from the significant likelihood that OAuth will be part of EHR incentive regulations in future years. Pilot projects such as RHEx have already shown the way to using OAuth and OpenID Connect as part of a data exchange scheme and have UMA on the roadmap.

Other projects, such as Kairon, have begun to tackle the algebra of state, federal and patient mandates that govern authorization. The most active public-private effort in this area today is the Automate Blue Button Initiative, or ABBI, that seeks to pilot and validate simple and scalable OAuth-based access to EHRs.

Solution Scenario

UMA would build on top of the Direct project PKI certificate infrastructure to bootstrap OAuth-secured interactions between all of the participants including the institutions, the staff and the patients. OpenID Connect would be enlisted to manage individual attributes and support role-based access management.

Policy statements will be designed to facilitate dynamic discovery of OAuth-protected resources in a manner consistent with the wide disparity in the size of participating institutions.

Patients will be first-class citizens in the UMA-based solution with convenient access to their identity providers and authorization managers. Patients will benefit from single-sign-on when it's implemented in a way that's transparent and consistent with whatever level of privacy they seek in their personal situation. Patients will also benefit by centrally managing access to their highly distributed healthcare data.

UMA Healthcare Use Case - The State Health Information Exchange

Now, it's up to HIEs to make the transfer of information among data holders practical and that means managing patient authorizations and adding value by aggregating patient information across the multiple independent data holders that serve a typical patient.

To summarize, the federal push for patient-controlled View, Download and Transmit and the mandate for secure Direct email accessible to patients and providers is in full swing but currently serves only point-to-point transfers.

Key Management Strategies for the Hybrid Cloud

Robert Griffin, Chief Security Architect, EMC RSA. Robert.griffin@rsa.com

Abstract

The rapid increase in the use of hybrid cloud services also increases the risk of exposure or loss of information and identities. Data encryption can play an important role in reducing this risk, but only if the keys used can be effectively managed. This article explores strategies for effective key management for data encryption for cloud deployments and the trade-offs that need to be considered.

Introduction

There is increasing recognition of the critical role that security, privacy and compliance have to play in decisions about what workloads to deploy into private, public and hybrid clouds.

Recent guidance by EMC, for example, calls out the importance of evaluating the trust requirements of workloads and the trust capabilities of the target cloud environment in determining the feasibility of deploying a particular workload to a particular cloud. (1) It also provides a valuable discussion of the dimensions of trust to be considered in cloud deployments, captured in this diagram.

Figure 1: The Dimensions of Trust

The Security for Business Innovation Council report calls out critical issues such as isolation of tenant data in multi-tenant cloud environments, visibility into the cloud service provider infrastructure in order to perform security assessments, and availability of information from cloud service providers for compliance purposes. (2)

To address these and other issues, many private, public and hybrid cloud infrastructures already use cryptography in many ways as part of their security strategy. For example, public/private key pair support is available from a number of vendors for securing data in use. Certificate-based protection for data in motion is broadly deployed in the pervasive use of SSL/TLS. Both cloud service providers and independent technology vendors provide symmetric key encryption for data at rest. In all these areas, cryptography technology already plays an essential role in security, privacy and compliance for the cloud.

Data encryption using symmetric keys is widely recognized as a core technology for securing data within any cloud environment. Early guidance from the Cloud Security Alliance, now in its third version, discussed why encryption is needed in cloud deployments, emphasizing that "strong encryption with key management is one of the core mechanisms that Cloud Computing systems should use to protect data." (3) Data encryption is also called out in the CSA control matrix. (4)

Key Management for Data Encryption in the Hybrid Cloud

The most recent CSA guidance expands this discussion of encryption significantly, while continuing to emphasize the critical role that data encryption plays.

But as that guidance also calls out: "Encrypting data has little value if both providers as well as users of cloud services do not vigorously enforce the processes around key management." (5) What are those key management processes that a cloud-enabled enterprise should enforce? CSA provides a number of recommendations, including "where possible, keys should not be stored in the cloud and must be maintained by the enterprise or a trusted key management service provider." (6) That is, certain deployment models for key management are preferable. But what are the deployment models for key management that an enterprise could consider? And what are the trade-offs among these various models in terms of security, cost and other considerations?

In developing use cases for the most recent version of the OASIS Key Management Interoperability Protocol (KMIP), we identified three primary key management deployment models, all of which are represented in currently-available commercial offerings. (7) These deployment models differ across four major dimensions:

- *Where are keys created?*
- *Where are keys used?*
- *Where are keys stored?*
- *Where are key policies managed?*

In the first of the deployment models, key management and data encryption remain in the enterprise.

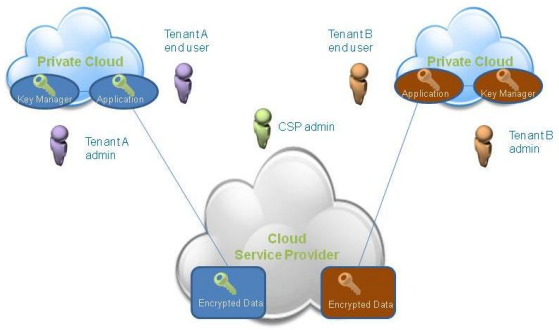

Figure 2: Key Management Deployment Model 1

Encrypted data is sent to the cloud service provider, but never the encryption keys, which are created, used and stored in the enterprise; key policies are also managed in the enterprise. The cloud service provider receives only the encrypted information, such as for backup or archive purposes.

A key management administrator within the enterprise is responsible for such activities as authorizing applications for access to keys, defining key policy, monitoring the key management system and so on. A system administrator within the cloud service provider is responsible for managing shared or isolated resources for storing the encrypted data, monitoring the infrastructure and so on. End users access applications within the private cloud, including data backup applications that manage the movement of encrypted information to and from the cloud.

As noted by NIST in Special Publication 800-144, "Guidelines on Security and Privacy in Public Cloud Computing", this is a common model for key

management for the cloud. (8) It is supported by such commercial off-the-shelf products as TrendMicro SecureCloud™. (9)

It is also a model supported at least to some degree by many SaaS, PaaS and IaaS vendors, although sometimes at an additional charge to the tenant; that additional charge reflects the impact of this model on the cloud service provider's ability to perform de-duplication on the stored data and thus increasing storage cost, or to perform data mining or other processing on the stored data, reducing the value that the provider can derive from the data. In addition, there are a number of important variants on this model. For example, Porticor takes advantage of partially homomorphic encryption and split keys to allow some processing on the encrypted data by the cloud service provider. (10)

This approach has significant advantages for the enterprise. It provides the simplest model for the security, privacy and governance of keys and unencrypted data, insofar as both remain solely within the purview of the enterprise. It also helps the enterprise to prevent uncontrolled dissemination of unencrypted information.

However, even though retaining control of keys, the enterprise should ensure that the encrypted data is restricted in access and dissemination within the provider's environment, and that well-governed processes for deletion of encrypted data in the provider's environment are supported; otherwise the risk of various attacks on the encrypted data, such as brute force attacks, increases.

This first deployment model has the major disadvantage that it restricts the enterprise's ability to take advantage of cloud capabilities beyond data storage. The second deployment model addresses this disadvantage by moving the application to the cloud, while retaining key management within the enterprise.

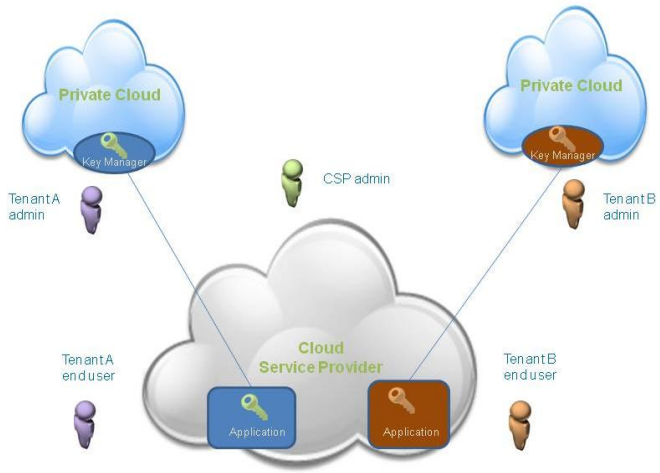

Figure 3: Key Management Deployment Model 2

In this second deployment model, the encryption keys are created and stored in the enterprise, where key policies are also managed. But the keys are used in the cloud. The cloud service provider receives transient copies of the keys, which are used for encrypting and decrypting information in the application processing performed within the provider's infrastructure. A key management administrator within the enterprise is again responsible for authorizing applications for access to keys, defining key policy, monitoring the key management system and so on. But end users now access the applications in the cloud. The system administrator within the cloud service provider, in addition to managing shared or isolated resources for application processing, monitoring the infrastructure and so on, may now also be responsible for monitoring and attesting to key usage in the provider infrastructure, particularly in terms of guaranteeing that keys are not persisted or disseminated.

This deployment model is supported by several commercial off-the-shelf products, including AFORE CloudLink® (11) and Gazzang zNcrypt™ (12), both of which support the transient and transparent movement of keys to the cloud environment. It is also supported by service-based models for key management, such as the JSON profile for the KMIP protocol. (13)

As mentioned earlier, this approach has the significant advantage of enabling the enterprise to take greater advantage of cloud resources, particularly when moving its own applications to a PaaS or IaaS provider.

It provides a relatively simple model for the security, privacy and governance of keys, insofar as both remain largely within the purview of the enterprise; however,

agreements that stipulate the non-persistence of keys and that provide for attestation and audit of this agreement are now required. In addition, unencrypted data is now present to a greater or lesser extent within the cloud service provider infrastructure, significantly increasing the requirements both for specification of technology and process securing that information and for attestation and audit that the agreed-to technology and process are in fact being used and are effective.

The third deployment model moves key management into the cloud.

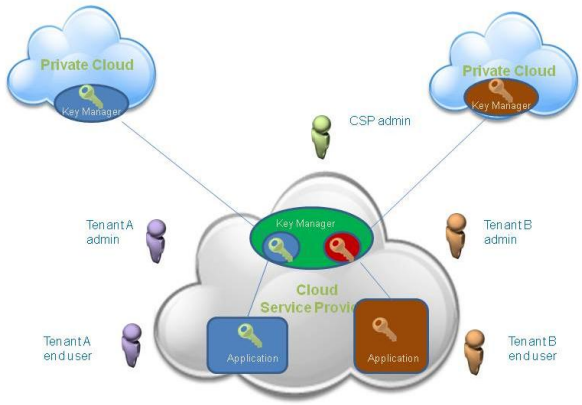

Figure 4: Key Management Deployment Model 3

In this third deployment model, the encryption keys are created, stored, used and managed in the cloud service provider, either within a shared infrastructure (as shown in figure 4) or in an isolated, tenant-specific part of the provider infrastructure. In this model, keys that persist in the provider infrastructure are used for encrypting and decrypting information in the application processing, also performed within the provider's infrastructure. End users access the applications in the cloud. Responsibility for administering those keys may be owned by a tenant-specific key management administrator who also operates in the cloud, authorizing tenant-specific applications for access to tenant-specific keys, defining tenant-specific key policy, monitoring tenant-specific key management system and so on. Or all responsibility for keys for a given tenant or all tenants may now be the responsibility of the system administrator within the cloud service provider, who also manages shared or isolated resources for application processing, monitors the infrastructure and monitors and attests to key usage in the provider infrastructure.

The approach of having key management performed within the cloud service provider is very common, although often with limited tenant-specific capabilities for key administration. (14). In a variant on this model, shown above in figure 4, the enterprise may retain a key management infrastructure responsible for key creation and key policy management. In his approach, indicated in information currently available regarding Microsoft Azure™ Trust Services, keys are transferred as needed or at specific intervals to the cloud service provider. (15)

This model has the advantage of enabling the enterprise to take even greater advantage of cloud resources, potentially reducing the operational cost of key management. It is important to note, however, that the liability for data loss or exposure typically remains with the enterprise, so that the requirements for specification and audit of data security and privacy remain critical. The definition of technology, process and structure for effective security and privacy of data entrusted to a cloud service provider is difficult, even with the help of reference architectures and control matrices. The negotiation with a cloud service provider the instrumentation, attestation and audit of those capabilities, again including process and structure as well as technology, can be even more difficult. And the effective audit of the agreement can be most difficult of all, complicated by the provider's responsibility for protecting each tenant, making it difficult for them to share details on resource allocation and processing within a shared infrastructure.

The Role of Standards in Key Management for the Hybrid Cloud

There are a number of fundamentally important standards already being used in support of key management for the hybrid cloud. Transport-level Security (TLS) is virtually universal in securing communication channels. Public Key Infrastructure standards define the certificates and keys used in identity identification and signing. Cryptographic standards establish the core algorithms and mechanisms used in data encryption.

Especially for the second and third deployment models, we see the OASIS Key Management Protocol (KMIP) mentioned earlier as taking an increasingly important role.

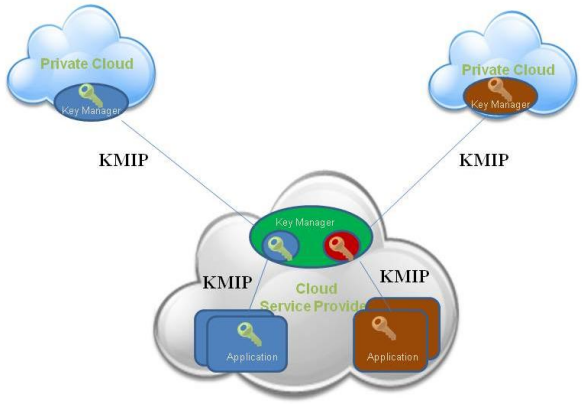

Figure 5: KMIP in Cloud Key Management

KMIP is intended to make it easier for any key manager to talk to any cryptographic environment and to other key managers. When key management is in the cloud service provider, for example, KMIP enables the key manager as a shared resource to communicate effectively with tenant-specific applications and enterprise-resident key managers. An interoperable protocol is an essential enabler for key management across the hybrid cloud, enabling organizations to take advantage of all three models as they are appropriate within the complex world of today's enterprise.

Conclusion

There are certainly many issues that still need to be addressed in effective key management for the cloud. For example, as with the cloud in general, one of the most difficult questions in key management for the cloud is how for an enterprise can ensure that the cloud service provider has a secure key management environment.

This includes specifying to the cloud service provider what security controls are required and establishing mechanisms for verify that the security controls are in place and effective. We also need new mechanisms to give enterprises visibility into their keys in the cloud, such as checking that keys have not been lost. Above all, we need more comprehensive and detailed understanding of the issues that need to be considered when looking at key management for the cloud.

But key management for the cloud is already a reality. Our hope is that this article has provided at least some insight that will help you in assessing the issues and opportunities for key management in the cloud.

1) *"Optimizing the Journey to the Cloud: Balancing Trust, Economics and Functionality".* http://www.emc.com/collateral/emc-perspective/h8539-cloud-optimiz-ep.pdf
2) *"Charting the Path: Enabling the Hyper-Extended Enterprise in the Face of Unprecedented Risk: Recommendations from Global 1000 Executives." Page 6. Security for Business Innovation Council. 2009.* http://www.rsa.com/go/innovation/index.html
3) *"Security Guidance for Critical Areas of Focus in Cloud Computing v2.1." Cloud Security Alliance. Page 60.* https://cloudsecurityalliance.org/guidance/csaguide.2.1.pdf
4) *"Cloud Control Matrix v1.3." Cloud Security Alliance.* https://cloudsecurityalliance.org/research/ccm/
5) *"Security Guidance for Critical Areas of Focus in Cloud Computing v3.0." Cloud Security Alliance. Page 133.* https://cloudsecurityalliance.org/guidance/csaguide.v3.0.pdf
6) *"Security Guidance for Critical Areas of Focus in Cloud Computing v3.0." Cloud Security Alliance. Page 132.* https://cloudsecurityalliance.org/guidance/csaguide.v3.0.pdf
7) *"Use Cases for Key Management Interoperability Protocol (KMIP) version 1.2."* http://docs.oasis-open.org/kmip/usecases/v1.0/kmip-usecases-1.0.pdf. *Several other organizations have defined use cases related to the cloud that are useful in assessing key management deployment models, even though the key management use cases may not be directly addressed, such as (*http://docs.oasis-open.org/id-cloud/IDCloud-usecases/v1.0/cn01/IDCloud-usecases-v1.0-cn01.html*) and the Cloud Computing Use Cases group (*http://cloudusecases.org/*).*
8) *Johnson, Wayne and Grance, Tim. Guidelines on Security and Privacy in Public Cloud Computing. Page 30. NIST Special Publication 800-144. 2011.*
9) http://www.trendmicro.com/us/enterprise/cloud-solutions/secure-cloud/index.html
10) http://www.porticor.com/
11) http://www.aforesolutions.com/products/cloud-security-management/cloudlink-overview/
12) http://www.gazzang.com/products/zncrypt
13) https://www.oasis-open.org/committees/download.php/47348/kmip-json-profile-v1.0-wd01.pdf
14) http://www.microsoft.com/en-us/sqlazurelabs/labs/trust-services.aspx
15) *For example, Salesforce supports assignment of a tenant-specific master key data encryption, but limits tenant-specific key administration to key archive/renewal, export/import and deletion.* https://login.salesforce.com/help/doc/en/security_keys_using_master.htm

Migrating Data to the Cloud Shouldn't Mean Sacrificing Control -- or Ownership

by Elad Yoran, CEO of Vaultive, Inc. – Jordan.bouclin@svmpr.com

Introduction

Cloud computing holds the promise of dramatically lower costs, improved scalability, enhanced user experience, greater ease-of-use and increased operational flexibility. However, for many risk-conscious organizations, there are still significant hurdles to overcome before they can fully embrace the cloud.

The primary issue that these organizations face is a structural one: how to control, secure and protect data that is processed by a third-party service. Whether the need is driven by concerns about safeguarding intellectual property, meeting compliance requirements for encrypting data-at-rest, maintaining data residency or increasingly navigating the ambiguity of legal protections for data in the cloud, these organizations need the ability to independently retain ownership and control of their data.

As a result, concerns such as data security and risk mitigation, regulatory compliance, unauthorized data disclosure and access, and international privacy/data residency regulations remain significant barriers to widespread enterprise cloud adoption.

These issues need to be resolved to address the requirements of the legal team, as well security or compliance officers, before moving an organization's data to the cloud. A summary of considerations for each appears below.

Data Security and Risk Mitigation

- ✓ Insulate data from the cloud service provider, prevent unauthorized access
- ✓ Separate ownership and control of data processed by third-party services

Regulatory Compliance

- ✓ Satisfy requirements for encryption of data-at-rest
- ✓ HIPAA, HITECH, GLBS, PCI-DSS, Basel II require or recommend encryption of data-at-rest

Unauthorized Data Disclosure and Access

- ✓ Provide a seat at the table for data owners in regards to law enforcement requests for data from cloud service providers

International Privacy/ Data Residency Regulations

- ✓ Encryption key residency and server-side processing against encrypted data sets

Addressing these requirements that are frequently barriers to organizations effectively adopting cloud-based services requires an approach that strikes a balance between security and service functionality. To ensure maximum data security and control in the cloud computing model, each business should ultimately encrypt data before it goes to the cloud -- and the encryption keys should be controlled by the organization that owns that data. By maintaining control of encryption keys, organizations effectively separate their data from the cloud provider's applications.

Smart encryption technology incorporating encryption of data-in-use is critical in this equation. It provides the separation of data from application by enabling processing of cipher text, which allows encrypted data to be searched, indexed and sorted within a cloud-based application without first having to be decrypted.

As a result, the cloud provider never has access to customer data in an unencrypted form and organizations are able to maintain complete control over their data -- while ensuring maximum security and regulatory compliance. Maintaining server-side processing and operations such as search, sort and index are critical to maintaining service functionality.

The Need for Organizational Cloud Data Controls

Concerns about data control and ownership are coming up more frequently because cloud computing has become increasingly mainstream. Cloud service providers, in response, have made strategic investments to directly address these concerns in order to encourage broader adoption of cloud-based services.

By implementing controls and processes for the cloud environment (such as those provided by the Cloud Security Alliance's Cloud Control Matrix), cloud service providers can credibly argue that they deliver more safeguards than an individual customer could within their own on-premise environments. In many instances, users gain the same economies of scale in security that they would in other IT operational areas, since the cloud service provider can make more extensive investments in security technology, monitoring and processes.

However, the security of the cloud service provider environment and ownership and control of data are related but discrete issues that end users must consider and evaluate in tandem. Ultimately, the customer -- not the cloud service provider -- should be responsible for the security and encryption protection controls necessary to retain data ownership and control.

Roles and Responsibilities	Cloud Service Provider	Service Customer
Availability	✓	
Physical Security	✓	
Resiliency (Disaster Recovery)	✓	
Patch Management	✓	
Breach Monitoring	✓	
Data/ownership		✓
Compliance	✓	✓

Figure 1: The basic responsibilities of the cloud service provider should focus on the security, resiliency, scalability and manageability of their service. The responsibilities of organization remain the same regardless of where corporate data resides or is processed: maintaining ownership and direct control of that data.

From a structural perspective, third party cloud-based services pose a challenge to traditional methods of securing data. Traditionally, encryption has been used to secure data resident on internal systems, or to protect data moving from one point to another. Ensuring that data remains encrypted in place within a third-party provider's environment and throughout the data lifecycle — but is seamlessly available to authorized users — presents a new set of technical challenges.

Encryption of Data-in-Use is a Critical Cloud Computing Best Practice

Encryption of data-in-transit and data-at-rest has long been recognized as best practices to enforce the security and privacy of data, regardless of where it resides. Moving to the cloud introduces the need to add another dimension to the existing two states of encryption: encryption of data-in-use.

According to the Cloud Security Alliance's Encryption Implementation Guidance, organizations should implement encryption of data-in-use to ensure that data is secured for the entire duration of its lifecycle (at-rest, in-transit and in-use).

The guidance notes that it is critical that the customer, and not the cloud service provider, is responsible for the security and encryption protection controls necessary to meet their organization's individual requirements.

The not-for-profit Alliance recommends that organizations encrypt their data before it leaves the trusted network, and retain the encryption keys (rather than the cloud service provider). The Encryption Implementation Guidance further states that once data is safely transmitted to a cloud service provider, it should be stored, transmitted and processed in a secure way.

Figure 2: The not-for-profit industry association, the Cloud Security Alliance, recommends that organizations implement encryption of data-in-use to ensure that data is secured for the entire duration of its lifecycle (at-rest, in-transit and in-use).

Data-in-use refers to the actual processing of data in computer memory, in addition to on-screen display and presentation of the data. Existing encryption solutions render encrypted information useless to most

applications, with their value lying primarily in storing that information in an encrypted format.

Employing encryption-in-use for data resident on a third-party cloud server to address control requirements poses a technical challenge to these traditional approaches. Enforcing cloud data controls raises two related issues:

How to insulate corporate data resident on cloud-based servers from third-party access and maintain the encryption when the data is processed in the cloud.

How to ensure that functionality is preserved and the full business benefits of a migration to the cloud are realized.

A generic approach to cloud-based services will inevitably fall short in addressing these two core requirements of encryption of data-in-use: addressing control requirements and preserving comprehensive service functionality.

Traditionally, if cloud-hosted data was encrypted, basic server-side operations such as indexing, searching and sorting records became impossible. Once cipher text was put into a SaaS application, some of the features of the program no longer worked, and the user experience suffered as a result. Encryption of data-in-use is designed to deal with the specific cloud services use case, and is intended to be complementary to existing investments in encryption, rather than a replacement.

Encryption of data-in-use is distinct from point-to-point encryption (encryption of data-at-rest) in that it remains encrypted when in use. In contrast to encryption designed for servers or virtual machines within an infrastructure as a service environment -- which must be decrypted to run operations -- data-in-use encryption allows organizations to feed encrypted information into cloud applications without requiring modifications to the applications, and still perform operations such as search, sort and index of against the encrypted data in the cloud.

While authorized users at the organization holding the encryption keys can access their data, the data remains inaccessible even to the cloud service provider administrators with full access to the customer instance.

How Encryption of Data-in-Use Differs from Tokenization

Tokenization is gaining acceptance among customers who are concerned about the compliance implications of moving data to a third-party service and satisfying data residency requirements related to sensitive data.

However, tokenization and intelligent encryption of data-in-use should not be confused. Tokenization involves replacing a value that is not required for an operation with an arbitrary value, typically in order to reduce the scope of compliance mandates related to data such as Social Security numbers, credit card details or patient records. One of the principal drawbacks of tokenization is that implementation requires a large database of the original data which still resides on premise. This database needs to be backed up and synchronized across datacenters.

While this approach does address data residency requirements and serves to limit the scope of compliance mandates from the cloud-based service, much of the value from migrating to the cloud is offset by additional investment in maintaining and synchronizing the tokenization database.

Encryption of data-in-use, in contrast, allows for server-side processing of all encrypted data on a third-party service (supporting operations like search, sort and index), while meeting compliance requirements.

When intelligent encryption of data-in-use is applied, and the organization's IT department retains ownership of the encryption keys, the state of encryption is maintained even when that data is processed in a third-party environment, while ensuring preservation of application functionality and a seamless user experience. Additionally, decryption is policy-driven and automated.

Conclusion

The addition of encryption of data-in-use empowers the organization to retain full ownership and control of data during the entire process, including when the data is out of its trusted network and in the cloud, while ensuring maximum security and regulatory compliance.

Mitigating the risks associated with control and ownership of proprietary data residing on third-party cloud-based servers is emerging as a critical consideration.

Ensuring that cloud data remains under an organization's control at all times — even when processed in a third-party environment—addresses material concerns related to data security, regulatory compliance, unauthorized

data disclosure and access, and international privacy/ data residency regulations. This allows organizations to reap the benefits of cloud computing, as well as the investments made by cloud service providers in resiliency, performance and security.

Elad Yoran is a recognized expert on information security market and technology trends. Yoran has 20 years of experience in the cyber security industry as an executive, consultant, investor, investment banker and several-time successful entrepreneur. He is also a member of a number of technology, security and community Boards, including FBI Information Technology Advisory Council (ITAC); Department of Homeland Security Advisory Board for Command, Control and Interoperability for Advanced Data Analysis (CCICADA); and Cloud Security Alliance New York Metro Chapter.

Cloud Best Practices for Open Data

Nuccio Piscopo, Solution Director, BIReady - nuccio.piscopo@gmail.com

MaaS and UMA implementation

As different tenants including insurance, clinics, pharmaceutical companies, ambulatories access healthcare systems, sensitive information should be provided only to Authorizing and Authorized users. MaaS (Model as a Service) might allow building and controlling shared healthcare Cloud-ready data, affording agile data design, economies of scale and maintaining a trusted environment and scaling security.

Still, MaaS can play a crucial role in defining the requirements of access control for healthcare multitenant cloud systems: constraints such as delegation, privileges assignment, separation of duty and temporal access are designed "on-premise" along the DaaS lifecycle. With MaaS, hData and hService conceptual functions might be designed "early" and classified in terms of tenant, data security, coherence, outage, availability, geo-location and to secure an assisted h-Service deployment.

MaaS provides "on-premise" h-Data, h-Service agile design

Applying MaaS to design and deploy healthcare services means realize faster and positive impacts on the go-live preparation with Cloud services.

Modeling is the key guideline in the DaaS architecture [1]. In fact, models are "a priori" on-premise databases because they collect behaviours, documents and information concerning structures, access rights, security and database scaling, partitioning and evolution [4]. With MaaS, access rights to data sources can be defined into the data model whereas DaaS defines proper workflows for healthcare service accesses and operations [10] [12].

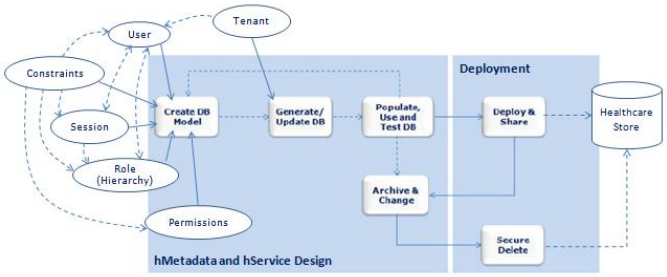

Still, the data model contains information and properties about the deployed h-Service as well as the databases and partitioning architecture:

- Model partitioning contains location constraints specifying where data is stored in case of data retrieving/breaches. The database architecture and partitioning are defined at the model level;
- When creating data models, MaaS allows to define properties for confidentiality, availability, authenticity, authorization, authentication and integrity;
- Models contain configuration and administration properties. Changes to security attributes and configurations to system participants should be only available to healthcare system administrators not to the cloud system administrators (Provider); partitions define DBMS deployment, tenants, hardening, machines, scripts and stored procedures.

DaaS assists h-Service deployment and provisioning

Since MaaS meets h-Service requisites and data properties to set rules and prepare h-Service provisioning and deployment, reliability and security features are verified and updated together with data infrastructure and h-Data protection constraints along the DaaS lifecycle. Models provide preconfigured DaaS properties [2]: when MaaS is applied, data models help meet Cloud h-Service directives. In detail, we might define two operational areas [5] [11] [12]:

1. The left scenario, the on-premise design zone, is the healthcare technical system administrators zone. Here, data models simplify the designing, tracing and updating of metadata. Ordering on-premise models and database architectures should be the primary goal for all healthcare bodies in the Cloud [6] [8]. Through the data model, properties such as service security range, DB partitioning and scaling, multi-tenancy, geo-location and all requested assets are defined "early" in the DaaS lifecycle.
2. The area on the right side is the provisioning and deployment zone. Multiple on-premise databases can be scaled up or down based on deployment needs. Models provide continuity with the databases' structure to extend to the Cloud preconfigured levels of security, compliance and what has been registered inside the data models. The right zone should be managed by a database service hosted platform fed by data models [5]. Models should always be the main reference throughout the lifecycle and provide operational input to the DaaS states.

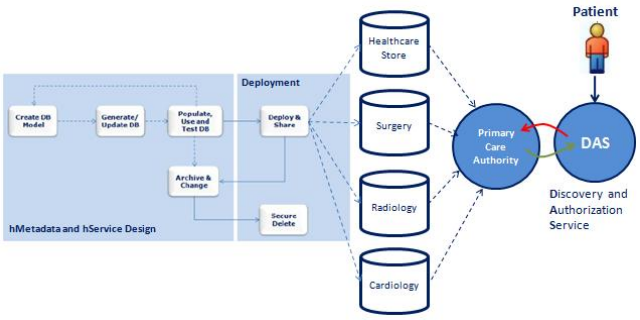

Healthcare users can be assigned to multiple tasks having defined permission and constraints. These access rights are defined at model level by MaaS. Changing the way to assign constraints, role, and permission should be set by MaaS that updates data deployment and provisioning through the DaaS lifecycle.

Conclusion

MaaS might provide the real opportunity to offer a unique utility-style model life cycle to accelerate cloud data optimization and performance in the healthcare network.

MaaS applied to healthcare services is the right way to transform the medical service delivery in the Cloud. MaaS defines "on-premise" data access security [10] [12], coherence, outage, availability, geo-location and an assisted service deployment. Models are adaptable to various departmental needs and organizational sizes, simplify and align domain-specific knowledge combining the data model approach and the on-demand nature of cloud computing. MaaS can support UMA scenarios requisites such as sharing models, nature of policies and claims, cardinality, collocation and host-AM relationships [11] [12].

References

[1] **N. Piscopo** - ERwin® in the Cloud: How Data Modeling Supports Database as a Service (DaaS) Implementations
[2] **N. Piscopo** - CA ERwin® Data Modeler's Role in the Relational Cloud
[3] **D. Burbank, S. Hoberman** - Data Modeling Made Simple with CA ERwin® Data Modeler r8
[4] **N. Piscopo** - Best Practices for Moving to the Cloud using Data Models in theDaaS Life Cycle
[5] **N. Piscopo** - Using CA ERwin® Data Modeler and Microsoft SQL Azure to Move Data to the Cloud within the DaaS Life Cycle
[6] **N. Piscopo** – MaaS (Model as a Service) is the emerging solution to design, map, integrate and publish Open Data http://cloudbestpractices.net/2012/10/21/maas/
[7] **N.Piscopo** – MaaS, DaaS Workshop, Awareness, Courses Syllabus;
[8] **N. Piscopo** - Applying MaaS to DaaS (Database as a Service) Contracts. An introduction to the Practice http://cloudbestpractices.net/2012/11/04/applying-maas-to-daas/
[9] **N. M. Josuttis** – SOA in Practice
[10] **H. A. J. Narayanan, M. H. Güneş -** Ensuring Access Control in Cloud Provisioned Healthcare Systems
[11] **N. Piscopo** - MaaS applied to Healthcare - Use Case Practice http://cloudbestpractices.net/2012/12/10/maas-applied-to-healthcare/
[12] **Kantara Initiatives** - http://kantarainitiative.org/confluence/display/uma/UMA+Scenarios+and+Use+Cases

Introducing OASIS PACR – Compliance Framework for Cloud Computing

Like many organizations the Healthcare sector faces strict information security compliance requirements such as HIPAA, and the new technologies and open standards like KMIP profiled in this magazine are the component parts needed to meet these requirements.

These will enable a secure model where the data is encrypted and the keys controlled by their users and organizations, and suppliers will increasingly offer tailored Cloud services that enable and enhance this approach.

Furthermore regulatory programs will emerge to audit and classify services in line with specific legislative requirements.

Trusted Cloud infrastructure (TCi) for E-Health

Ultimately these evolutions will enable "Trusted Cloud infrastructure" (TCi) – Cloud services that operate and offer a recognized security level in these terms (e.g. "HIPAA Compliance") and are proven to do so via some form of independent audit.

The key is that this rating should be auditable to a given framework requiring some form of a relevant supplier accreditation system, such as the UK's G-Cloud program. This caters for a standardized service catalogue and the mechanisms for registering and approving suppliers.

For heavily regulated sectors like Government, Healthcare and Financial services, these frameworks will be the key to unlocking their large-scale adoption of Cloud services, via providing a supplier marketplace system tailored to their industry needs.

This creates a means by which buyers can distinguish different 'grades' of service relevant to their business goals and legislative mandates, in particular mapping information security levels (like IL0-4 in the UK) to a corresponding Cloud service.

Globally the American NIST organization is recognized for their baseline definitions of Cloud Computing (IaaS, PaaS, SaaS, ..etc), and these can be mapped into eHealth scenarios the same way, and by extension also cater for the same supplier accreditation framework too.

In Healthcare these could be applied in scenarios such as the 'RM&R' eReferal process in Canadian healthcare, which could be implemented on a Community Cloud, making this process available to many organizations via a shared service model.

This could be outsourced to a third-party hosting firms, where the associated security model configuration can be verified against this global Community Cloud baseline, and also where the latest innovations can be monitored and adopted.

For example the latest development from NIST includes the remit to identify the 'Geolocation' of Cloud-hosted data, meaning where exactly is it located geographically, as this determines the laws that are applied.

Healthcare is a pertinent example of this relationship – For example in some provinces in Canada it is legislated that personal healthcare data cannot leave the local region, let alone the country.

In general most if not all Cloud hosting regulations all favour a local in-country hosting model, and so these are simple but critical progress steps, essential to an audit.

These developments set the scene for a new OASIS open standards working group, called 'PACR' standing for Public Administration Cloud Requirements.

OASIS is an international standards body responsible for many XML Web services open formats, now also expanding into the Cloud ecosystem through working groups like TOSCA (Orchestration) and CAMP (Application portability), with PACR intended to focus on how to define and implement Compliance frameworks for regulatory needs.

The group has been submitted for member feedback until 31st December 2012, with a view to launching in 2013.

Read more here.

The Cloud Best Practices Network offers workshop services to explore and plan these roadmap journeys. Contact@cloudbestpractices.org