Toronto Cloud Business Coalition

PLANNING FOR THE CLOUD/ CLOUD STRATEGY: ENTERPRISE

Essential guidance on how best to develop plans and strategies to guide cloud adoption and use within large enterprise environments

Lead Analyst: Mary Allen
Contributing community experts: Shawn Rosemarin (VMware), Roy Hart (Seneca College), Joe Belinsky (Moneris), Jeff Cohen (Shoppers Drug Mart), Sangam Manikkayamiyer (Symantec), Stefano Tiranardi (Symantec), Chris Vernon (Symantec), Wil Stassen (Cushman & Wakefield), Matt Starkie (Microsoft)

December 2015

Contents

Planning for the Cloud/Cloud Strategy: Enterprise

Best practice guidance from the Toronto Cloud Business Coalition

Definition and context

There is a sense of urgency around cloud adoption that the TCBC Enterprise Cloud Planning working group believes can lead some to "shoot first, ask questions later." Growing awareness of the competitive advantage that cloud computing has delivered and the opportunities it can help generate has created impatience in many organizations – and within individual business units – to take advantage of cloud without due consideration for adoption criteria that address the needs of the business as a whole. A commonly held view associates cloud with 'agility' and 'rapid time to value' and a common expectation, fueled often by vendor positioning on cloud technologies, is that adoption similarly should be simple' and 'instant'. But cloud deployment is not limited to decisions around technology in isolation, rather it is the working group's contention that a key requirement for successful implementation is for the user organization to "get its house in order before leaping to cloud." This entails planning across multiple cloud-powered activities, whether these involve decision-making related to procurement of SaaS applications, or sourcing the infrastructure needed to deliver an application or service to internal and external customers. Optimal planning focuses on supporting the transitions that will be experienced by people, and in business processes and technology that accompany migration to cloud.

The TCBC Enterprise Cloud Planning working group believes that transformational change management is a foundational principle that frames proper planning for cloud adoption and use. This exercise has multiple components, but in the cloud context consists of four primary requirements:

- Managing the end-user experience, which is likely to change when a cloud based solution is deployed;
- Managing roles and responsibilities, which are likely to change when operations are moved to a cloud provider;
- Business workflow and process change needed to accommodate the addition or subtraction of application features that may accompany migration; and
- IT change and release management, including processes to enable the security of IT environments, and the business of IT (Who pays for what and how?).

Cloud spells transition for IT, but it can also entail loss of control over both the user experience and the technologies that support business and IT processes – dynamics that can and should be considered in cloud adoption planning.

At a technology level, it is also possible to apply a systematic approach to cloud planning which focuses on identifying what problem needs to solved at each logical layer. In the case of

applications, cloud deployment aimed at providing a better way to deliver applications and services operates in the following ways:

- At the device layer, delivering application capability across any device;
- At the control plane layer, optimizing the operations and automation of application workloads and showing/charging back for them appropriately; and
- At the infrastructure abstraction layer, driving maximum utilization across all hardware assets whether service is delivered from within the data centre or via the public cloud.

As a broad approach to cloud strategy, this systematic line of attack has much merit, and is one that the working group would urge end user organizations to consider as a means to comprehensive cloud planning. The goal of the working group, however, is not to address planning considerations at each layer of the stack or in all change management scenarios, but rather to adapt this principle of problem solving in specific areas that have been identified as key issues hindering cloud adoption in the Canadian marketplace.

Aligned with the need to address change requirements in transition planning for people, processes and technology, the working group has identified the following issues as most important with respect to the development of best practices:

1. **Roles and Responsibilities**

    - How does cloud impact roles within the organizational structure and hierarchy and how are these roles evaluated?
    - How can existing IT organizations transition away from traditional infrastructure maintenance to new, higher-level service delivery functions such as enterprise architecture and cloud brokerage?
    - What is the best way to encourage interaction and drive alignment between IT, Finance and the Line-of-Business to ensure cloud decisions are evaluated in the best interest of the company?
    - What is the best way to develop a shared set of values and accountability between development and operations teams to drive greater collaboration and ownership across the application lifecycle?

2. **Skills Training and Upgrade**

    - Cloud adoption and management demands many new skills sets. At a foundational level, enterprise architecture now spans beyond the walls of the data centre and therefore requires diverse experience and training to fill this “knowledge gap.” Specifically, tomorrow’s enterprise architects will need to well versed across on-premise and public cloud architectures in order to determine the best place for a given service to reside.

- The end-user experience will likely change with cloud adoption. How can the organization better predict the potential for staff members to need education and training on the new features, processes, interfaces associated with cloud migration? What is the best mechanism for delivering this essential skills upgrade to ensure minimal productivity hit at launch?

3. **Organizational Controls and Governance**

- Since cloud may entail some loss of control over data management, what must be in place to ensure sensitive data is protected? Is it possible to plan and contract for data governance within service provider cloud environments?
- Which existing processes and policies will translate to a cloud world and which will not? How can an organization optimize workflows and simplify approvals to ensure that the new cloud based service delivers against its agility goals.
- Data stewardship is a critical undertaking in any IT environment; however, cloud has catalyzed this discussion since data is highly mobile in cloud scenarios, and since cloud entails some devolution of control over data to the third-party service provider. What does the user need to understand in terms of data residency requirements from a legal perspective, and what policies need to be in place to provide transparency on where the data has travelled, what it touches (in a multi-tenant environment), and who has access rights?
- What techniques can be developed to re-evaluate standard organizational workflow and approvals as the organization transitions to cloud? How does the organization assess whether or not approvals and process flows need to remain the same, or change within a cloud context? If these change, does the user organization build functions to accommodate transition or hand these off to a third-party and what are the criteria for making this decision?

4. Assessing Cloud Readiness

- What factors should user organizations take into account when deciding if the business and culture are ready for cloud, and if cloud is the best option for the business?
- Which applications are best delivered via SaaS, which are better delivered via public cloud based infrastructure services, and which are best delivered via on-premise private clouds, and what determines this?

- How can advanced infrastructure capabilities such as blueprints, cloud automation and orchestration in internal private cloud or public cloud environments improve application service delivery and agility?

5. Cloud Connections

- How does the organization enable seamless mobility between clouds for test/development and disaster recovery scenarios? What architectural considerations must be enabled to support mobility between clouds?
- How can the organization ensure that cloud services interoperate with on-premise or other third-party cloud services in order to provide seamless application delivery?
- What standards must we set upon our Application Programming Interfaces (API) to ensure simple and efficient communication between services?
- How can a standard set of cloud automation blueprints optimize service delivery across heterogeneous clouds?

6. Exit Strategies

- How difficult, costly or complex is the process of "unrolling" cloud deployments if business circumstances change?
- How seamlessly can an organization move its application workloads between private and public cloud platforms?
- How can organizations identify and avoid vendor lock in with cloud providers, and what contingency plans should be in place around portability of data, services and "logic"?
- How important a consideration is an exit strategy in the decision to move to the cloud, and what is an appropriate planning horizon? Is the time frame in planning for cloud exit equivalent to planning for ERP – i.e. a twenty-to-thirty-year decision?

Figure 1. Planning for the enterprise cloud: managing layers of change and control

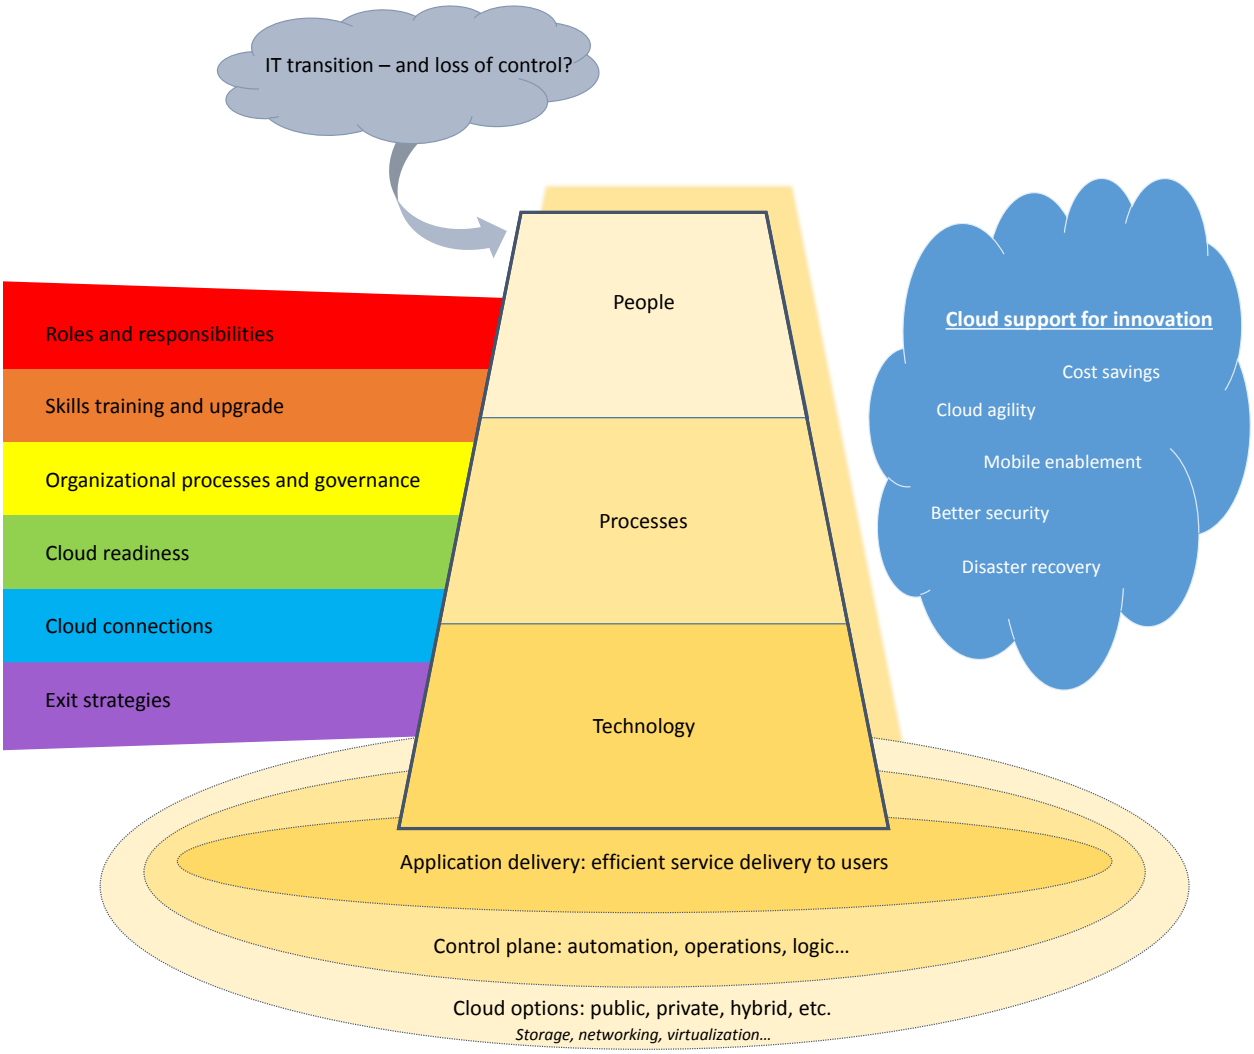

Source: TCBC/InsightaaS, 2015

Reference sources

For cloud standards definitions and implementation, guidance see:

- National Institute of Standards and Technology. NIST Cloud Computing Standards Roadmap. US Department of Commerce. Special Publication 500-291.
- OpenStack Documentation. http://docs.openstack.org/

Business objectives

At its inception, cloud was viewed primarily as a means to cutting costs in IT service delivery. However, with greater experience of cloud, organizations are beginning to sketch out a much larger value proposition for the technology. While some cost savings may be available, the ability of cloud computing to support innovation within the organization is now viewed as the key business benefit. Without this support for innovation, the cost curve in and of itself is not nearly compelling enough to be the single or largest driver of cloud adoption. In fact, the Enterprise Cloud Planning working group has identified five primary adoption drivers that are aligned with an organization's business objectives.

Cost savings

The sharing of resources has enabled higher utilization rates for infrastructure in multitenant cloud. Combined with other technology efficiencies, this sharing produces cost savings that may be passed on from the provider to the consumer. Other sources of savings for the user organization include outsourcing of IT support and maintenance. New payment models may also appeal to the end user organization: a shift to OPEX for on-demand resources from CAPEX investment in IT equipment may reduce overprovisioning and large up-front costs, and SaaS subscription-based spending may better accommodate some user budgets.

Cloud agility

Cloud technologies enable business agility on a number of levels. While SaaS applications may deliver a better user experience and greater reliability to enhance worker productivity, 'try and buy' options for cloud solutions increase the likelihood that organization procures the most appropriate technologies. Similarly, use of cloud infrastructure can improve development cycles by reducing the time required for stand-up of minimally viable products from months or weeks to hours, and by enabling fast testing of new products for commercial viability. The end result is shortened time to market and more rapid capture of demand, as well as the ability to test the limits of demand for existing products in new geographic or other markets. Through instant access to infrastructure resources and on-demand consumption, IaaS offers opportunity for easier product/market experimentation, as well as potential for creating new, agile operating models for cloud native businesses.

Mobile enablement

Serving as the central repository for corporate data and applications, cloud technologies are powering the emergence of a new mobile workforce that can access content anywhere, anytime. Productivity benefits derived from workplace flexibility are augmented by cloud's ability to support collaboration between workers in remote locations, and communication between businesses and their customers, supply chain partners and other stakeholders.

Better security

IT security is an ongoing challenge for many organizations, and the issue continues to become more pressing with evolution of the threat landscape, and business practices such as BYOD that

introduce additional risk. A successful defence requires the knowledge and expertise of the cloud service provider, which is often in a better position to engage and maintain needed security expertise. In addition, the cloud provider is likely to have invested in the acquisition of security certifications that require skills and ongoing investment that may be out of reach of the end user organization.

Disaster recovery

Business continuity is an increasingly prevalent use case for cloud technologies. The cost and complexity associated with maintaining mirror data centre facilities in locations that are far enough apart to avoid simultaneous outage from physical disaster, yet close enough to support the low latency requirements of modern applications are driving increasing numbers of organizations to outsource DR to lower cost, third-party cloud storage specialists.

The working group's view is that a good starting point for cloud planning is to identify the pain point(s) that cloud is expected to address, and ensure that deployment will address a sufficient number (two or more) of these adoption drivers. Combined, these cloud-based capabilities work to support business innovation, helping organizations to sustain their momentum in the market. Successful implementation is less about "cool technology" than about alignment of cloud with business objectives: an ability that is increasingly critical for the organization in ultra-competitive markets, and understanding of which IT leaders are increasingly expected to demonstrate. The role of IT has changed. It no longer functions simply as a buyer of IT equipment; increased scrutiny and tighter budgets mean IT must now must explain to the CFO what business problem a project addresses before gaining approval. IT is a value organization in any business, but must be able to outline the business value in technology implementations in order to remain relevant.

Planning can be instrumental in this process. Inadequate planning, on the other hand, results in ad hoc adoption that fails to capitalize on the synergies available through cloud orchestration, as well as growth of 'shadow IT' and its associated security and compliance risks. While the IT organization has traditionally been responsible for governing and operating all IT services, cloud adoption ease – of SaaS applications, for example – has produced a shadow IT phenomenon in which cloud resources are sourced by various individuals or business departments outside the purview of IT. Staff who may not have proper authorization or training in contractual obligations sign on to cloud services, and through misuse of corporate identity may bind the employer organization to terms in an agreement which are not well understood or that incur unacceptable risk/liability, as failure to properly plan around the legal liability for data can leave organizations open to considerable exposure.

The challenge is to ensure the IT organization governs all workloads or services, whether or not it operates them, a goal that may be realized through planning, processes and policies that holds IT accountable for governance of services, regardless of where these are operating. Many organizations are now maturing processes that can identify areas where shadow IT might be occurring and working to enable a transition plan that will introduce greater

accountability; however, most likely remain unaware of just how much cloud they are actually consuming. The working group agrees that "shadow IT is not necessarily a bad thing but it needs to be managed."

Figure 2. Supporting enterprise business objectives with cloud planning and strategy

Source: TCBC/InsightaaS, 2015

Best practice and processes in enterprise cloud planning

Roles and Responsibilities

The discourse over cloud's impact on staffing within the adopting organization has been ambiguous. While the IT community's first response to cloud automation was fear for the displacement of traditional IT workers, this concern has been mitigated through the articulation of new roles and responsibilities that would allow IT staff to grow and develop expertise beyond "keeping the lights on" – beyond the performance of routine maintenance tasks in the data centre that could be automated through cloud. A good deal of attention was also devoted to the "cloud broker," a service function that would be responsible for the dynamic procurement and management of services potentially spanning private-public and hybrid cloud.

But how has this shift played out in real life, and what needs to be in place from a people perspective to achieve effective cloud transformation?

At a basic level, the impact of cloud varies from organization to organization, and depends on the impetus for change within. In 'top down' scenarios where the driver of cloud adoption is the CXO who creates cloud by principle, the IT department is likely to embrace the operational efficiencies that can be achieved in certain areas – service desk, for example. In cases where the Line-of-Business manager drives cloud adoption, focus is on agility for the rapid delivery of business value, with little regard for the cloud provider's underlying functionality, or the IT department's capability to support it in the longer term. The optimal scenario is for these two worlds to meet: for the business unit manager to source cloud within the security and governance parameters established by IT, and for IT departments to aim not only for the efficiency gains that virtualization of IT infrastructure can bring, but also for the adoption of solutions with the characteristics that are unique to cloud – on-demand, self-service and pay-per-use.

To effect this culture shift, it is the working group's belief that cloud implementation, in both private and public contexts, demands first a new way of thinking about service delivery on the part of IT. The delivery of efficient IT services is no longer about "cost savings," it's about saying yes to more innovation and driving "revenue gains." "Embrace change; it is coming" as one group member put it, through more open collaboration and interaction at top levels between IT and the business units to better align operational planning with longer term business needs, and through better coordination of operational and development activities. To accelerate the pace of change, the working group recommends adoption of the agile "DevOps" approach in which development and operations teams break down walls and deliver seamless services with a single shared goal – to help foster innovation.

If operations in traditional legacy environments can be characterized by what the working group describes as an "assembly line, batch mentality," adoption of public cloud services calls for thinking in real time about integration points between corporate applications and data and external systems, also referred to as a "service based architecture." While integrations are required in traditional, monolithic IT environments, they tend to be hard written into the applications within millions of lines of code and dependant on their original architecture. Service based architectures, in which micro services deliver on the premise of interoperability without dependency, are different in that each individual service is designed to be flexible in its integration with other services. Traditional IT organizations need to move "up the stack" to focus more on the way that services are architected, built and maintained across the enterprise and less on the underlying infrastructure or platforms that they sit upon. But managing multiple integration points involves different skill sets, different management practices and more timely response: in twenty years' time the working group expects the IT organization look different than it does today – and feature an expanded role for integration specialists, and

cloud brokers tasked with adapting the organizations unique business and IT needs to available services.

At an organizational level, another area that cloud will shape is data sovereignty and data lifecycle management. Understanding where data resides and what permissions exist, can be more complicated questions in a cloud environment, a consideration that is likely to entail an expanded role for the CISO in ensuring appropriate data governance is in place, as well as new requirements in the development of information architectures. Today, for example, in the vast majority of SaaS solutions the SaaS provider retains all data. In future, it may be that SaaS will allow for multi-site architectures where data sits "on-premise," Or perhaps cloud organizations will begin to expand into the business of data lifecycle management, contracting for the delivery of data retention, compliance and security requirements.

This kind of proposition can be very complex, as cloud adoption in the case of highly regulated industries in particular demonstrates. In banking, a common practice involves the deployment of public cloud for a certain service, followed by the creation of a backup environment for this in on-premise infrastructure. This data replication and the local run of applications is designed to satisfy auditors of compliance with data retention and archiving policies in the (frequent) event full risk and accountability is not guaranteed by the public cloud service provider. "In the haste to embrace the cloud, you need to measure twice and cut once," one working group member noted, applying creative architectural solutions that consider the entire lifecycle of the data to ensure that data processes and protection remain intact, without the layering on of additional overhead to accommodate company policies.

In the cloud era where architecture is out of the sand box and nebulous, a key role that is evolving is the enterprise architect who has world experience, who can visualize service delivery across multiple clouds and business functions, and who can apply simplicity of design to information management. In the working group's view, the creation and support for what the analyst firm Gartner calls the "Vanguard Architect" is of critical importance in cloud adoption as the level of insight that these have will be a huge factor in determining the appropriateness of the architecture that will be built to support IT service delivery. At the same time, the working group believes that the "Vanguard Architect" should not be a role singled out from others but rather a philosophy for all architects to embrace.

Skills Training and Upgrade

Cloud demands new skills at many levels of the organization. While the leadership team must learn to envision cloud potential and develop strategies for driving the transformational change that cloud can bring, business unit managers must familiarize themselves with cloud capabilities to adjust product/market cycles to the pace of IT service delivery that cloud delivers. For their part in ensuring successful implementation, information managers must develop understanding of cloud's potential impact on data location and governance. Though infrastructure staff may believe that information resides within the boundaries of a specific region, it is the information/privacy officer who has responsibility for understanding the risks associated

with cloud deployment and for initiating information management processes that mitigates this risk. He/she must learn how to ensure that data within a service provider environment is not replicated to another country, or if it is, what privacy/access legislation would apply, how to classify data in terms of sensitivity and potential risk and whether other cloud storage arrangements, including data retention cycles and disclosure provisions, might compromise data privacy. In cases where this knowledge base is not available internally, the information or decision support specialist must know where to access it – and whether this information is best accessed in contractual discussions with the cloud vendor or through discussion with a third party consultant. In addition, the business of cloud demands knowledge of the legal ramifications of data migration to an external provider: who is responsible (the cloud service provider or the customer) for the impact of data breech, and how might cloud outsourcing contracts affect insurance policies/premiums that the adopting organization has in place to manage data risk.

In many ways, the enterprise architect straddles the worlds of business and IT operations, acting as a bridge that links awareness of the business and information needs of the organization with their execution through IT services. As such, the enterprise architect must develop new levels of understanding about the flow of information and business logic beyond the walls of the enterprise to cloud repositories for storage and processing, creating appropriate access points for business along the way.

At the technical level, the working group points to a number of new areas that will demand skills upgrade to support migration to cloud: cloud ops, for example, which is designed to enable cloud automation, operations and financial transparency across clouds. In addition, network architects and administrators must learn how to build out a “virtual network” that allows workloads to easily span Layer 2/3 between clouds.

At a practical level, the notion that one professional would possess all required networking skills is unrealistic. The working group believes that teams with business analysis skills, with technology skills from a networking perspective, and technology skills from a security perspective will all be needed to perform the cloud appropriate audits, to implement security controls at the right exit point and the right entry point into the network, and in conjunction with business owners, to manage access to applications and data that reside in the cloud. To ensure coordination of tasks and that the right skill sets are in place, the group notes that new levels of collaboration between technologists, network and security administrators and the business units are critical to successful cloud implementation.

Skills upgrade may also be required at the end user level. While application training in the case of SaaS adoption vs. deployment of a new software on-premise in traditional IT environments may be similar regardless of the infrastructure it sits on, this is not the case in all instances. For example, proprietary applications by their very nature are designed to address the unique needs of the organization from a business process or workflow perspective: with popular SaaS applications, on the other hand, application capabilities are typically developed to meet

the need for standard delivery of functionality across the broadest possible base of customers. With SaaS, clients are encouraged to minimize customization to accelerate deployment and minimize expensive implementation services and post-sales support. Successful cloud adoption – including extensive uptake of the cloud-based service – depend on educating business unit managers of the limits of customization with popular CRM or sales management apps, and on adjusting workflow and providing end user training on the new technology to build familiarity with the new interface and on new processes.

Organizational Controls and Governance

Broadly defined, 'governance' is the combination of processes and workflow implemented by corporate boards to inform, direct, manage, secure and monitor the activities of the organization as is works toward the achievement of its goals. Often associated with fiscal transparency and fairness, in the cloud context, governance may be understood as guidance for building processes in support of the organization's migration objectives, while addressing potential risk issues associated with transition to the new IT paradigm that cloud represents. Its companion – 'control' – is defined as the actions taken by management, the board or other parties to manage risk and to increase the likelihood that the established goals will be achieved. With cloud implementation, these objectives may include cost savings, greater business agility, faster time-to-market for products and services, or simply access to new IT functionality, outcomes that are more likely to be realized with effective planning, sufficient action and control.

The enterprise planning working group suggests that the technical approaches to control governance in cloud deployment vary with the organization, and depend on company goals, on how much governance can be put in place to manage the risk associated with cloud migration. In addition, control statements can be implemented using different technologies. In the case of highly confidential financial information, such as the enterprise's profit and loss statements or balance sheets, for example, a control statement would say that this information should not move out of the enterprise before earnings are announced, and several different technologies could be deployed to ensure adherence. These might include web or email filtering mechanisms or data loss filtering when information is transferred to the cloud, or the encryption of financial data and its storage in a secure repository that only the auditors can access. A high-level control statement would say that anything which has been labelled as confidential needs to be handled in a certain way, and IT would work with security personnel to identify the right technology to secure the data when it needs to leave the organization – as when data moves to a cloud repository.

While transferring data governance matters to the cloud service provider would be appealing due to simplicity of the solution, ultimately, responsibility as well as legal liability for data resides with the customer organization, which also has corporate interest in ensuring levels of process – for auditability, for example, which some standardized provider infrastructure offerings may not be ready to deliver. In assessing governance and control matters that are

important to cloud migration, the working group suggests enterprise planners consider the following issues:

Data stewardship

With the dissolution of central records bodies back in the 1980s, discipline around information management and retention schedules was lost in most organizations. The last thirty years have been characterized by undisciplined information collection across most enterprises, and reliance on search capabilities, rather than organization of data, for information retrieval. Issues with data retention have been compounded by the proliferation of places that data can now be stored – including traditional, on-premise databases, branch or central office locations, hosted and cloud repositories – making the timely and effective retrieval and disposal of information for legal or regulatory purposes highly problematic.

In addressing this issue, it is incumbent on the enterprise to know where data resides, how to find information within the data, and what information will be deleted by when. Retention schedules are highly complicated, and vary depending on the nature of the data. In healthcare, for example, data associated with a patient must be kept for the life of that individual, but as a transitory record, an email between hospital administrators may not need to be saved longer than two years, even though it may contain information that is of critical importance to the institution. Engineering and other information on a building must be retained as long as the structure stands – potentially hundreds of years.

In most organizations, responsibility for the protection of key information has been delegated to the Chief Information and Security Officer (CISO), who works with the CIO to ensure that private information is protected and governed appropriately. CISO and the Chief Privacy Officer are relatively new roles within the corporation that typically hail from administrative offices that used to have responsibility for policy on physical records; however, these often assume a conservative stance in regard to information management policy due to data loss challenges. For example, in the case of a data breach, organizations are typically compelled to report that all data has been compromised – a practice that can overestimate risk but satisfy requirements for transparency. The working group suggests that data stewardship that includes a schedule around retention/deletion of data will help ensure that data is not stored longer than it need be, or longer than the organization is required to produce it from a legal perspective, is an important step to data governance in the cloud. Securing disparate storage repositories with policy-based security will also ensure that damage from a potential breach is minimized.

Identity management and encryption

Enforcement of data protection policy can only be accomplished with a focus on policy-based identity management and encryption. Data may be secured, for example, by identifying an individual or organizational role and granting permissions to access specific data sets based on that identity, and by limiting access to those individuals only who have appropriate credentials. Industry best practice is now to follow the path of “least privilege” giving employees only access to what they absolutely need and re-qualifying access to sensitive data on an ongoing basis.

Encryption of data at rest and in transit to and from the cloud repository that does not affect application performance is another key technology designed to fulfill data protection requirements. These technologies, when combined with general security precautions, provide an approach that is different in kind from an earlier "Castle and Moat" model where firewalls would be built around servers to deliver protection. In a cloud world, the organization does not own the servers, the servers move, as do the buildings that house them, and VPNs to manage data traffic over the Internet do not scale directly with the multipoint transmission of corporate data.

Successful security management depends in many ways on balancing IT interest in security with the business end users' need for an experience of security technology that does not unduly interfere with work activities. Single sign-on identity management solutions implemented with single or dual factor authentication schemas are a recent approach to the challenge of reducing complexity for the user while maintaining protection for corporate data that is touted to work well in multi-cloud or hybrid environments. The working group advises caution in use of this technology, though, as it relies on the use of one master key to provide user access to many different processes. While the user need not remember more than one set of credentials, if the key or credentials are compromised, hackers gain access to all cloud applications. In order to address these concerns, the addition of a second factor of authentication is recommended to mitigate such risk. A Single Sign-On (SSO) platform with integrated strong authentication, access control and user management will provide a better experience for users, control for administrators and security for the organization.

A more balanced strategy involves decision-making around which applications may be single sign-on enabled, and which may need an additional layer of inspection. As example, the working group points to the approach taken by the banks' anti-fraud and anti-money laundering groups, which dictates the placement of additional fraud prevention engines to run on top of the transaction data. Even though bank may offer single sign-on, personal online banking trackers can flag more than 10 account transfers within a specified time using the fraud engine, and outlier behaviours will trigger a call to action – a follow up automated text/phone call that improves security and with it, the customer experience.

Workflow transitions

As organizations transition to cloud, change in workflow processes and approvals operates on three levels: at a strategic, organizational level, and at the application and infrastructure layers. Any evaluation of techniques should start from the perspective of business and regulatory risk, considering how and where use of cloud impacts relevant workflows. Where cloud is introduced as a component of a process that carries business or regulatory risk, the enterprise will need to take three steps. The first, which is important to compliance, is to identify where and how external processing and storage of data affects how core systems of record may be accessed or altered. The second step is to evaluate supplier SLAs, identify where and how they do or do not correspond to business and/or regulatory requirements, and work to aligning the

SLA with corporate practice and regulatory guidelines. This process can take several forms: some cloud vendors will adjust SLAs to accommodate a client and others will not. If a supplier will not, then the organization must decide whether they will alter their business practice and expectations or seek another supplier. The third step involves creating a "plan B" to address business risk in the event the cloud service becomes unavailable.

A good starting point for identifying and managing process change resulting from cloud application deployment is to first consider which existing internal processes will be disrupted by the move to cloud. This decision is based on a variety of factors, such as tolerance to security risk and other corporate drivers, or even corporate cultural norms that admit the use of public cloud services or that insist that data be kept under lock and key on-premise. In still other cases, business drivers, such as data residency or uptime requirements, will have primary influence over the identification of applications that can be adapted to cloud delivery models.

For example, unlike the deployment of packaged applications in traditional IT environments, use of SaaS applications may not allow for the levels of customization that the organization is used to, or that is possible with enterprise development of proprietary software designed to serve specific business needs. In these situations, transition to SaaS may entail alignment or change of business process to match application functionality or capabilities, a process that is best addressed in advance of deployment through "business blueprinting" – mapping out a business function from end-to-end to see all elements of workflow. In procurement of SaaS-based applications, the user organization should take into account the costs business change management might entail, as well as the SaaS application's ability to withstand industrial usage, balancing these factors against the cost and convenience of using cloud-based applications. Ultimately the rule of thumb here is to avoid any customization unless absolutely necessary. Although organizational change is hard in the long-term, the upfront effort of customizing a solution to suit your business coupled with the ongoing costs of supporting it can have a dramatic impact on project success and delivering a positive ROI.

In the IT department, the key shift is from task-oriented processes to vendor management. For example, removing responsibility for maintenance activities such as patch management, which may be automated in a cloud context, will simplify process for IT operators; however, responsibilities must be parsed out, delegated either to in house staff or to the cloud provider. In many ways, the biggest challenge, and one that spans infrastructure and applications, is the alignment of technology system change with process change. Areas that can be overlooked in migration to cloud infrastructure, but which require detailed process documentation include: management of CMDB (change management database), change control, ticketing systems, approval processes (for access to cloud apps and data), data classification to determine sensitivity and storage in different cloud modes, and in implementations featuring cloud self-service, deciding who has access and what are their permissions. To ensure successful cloud migration, enterprises are advised to support implementation with plans that does not neglect process change.

Assessing Cloud Readiness

An organization's readiness to adopt cloud varies with the application and with the state of existing capabilities. In cases where the on-premise solution is not reliable or not meeting the needs of the business user, and a cloud solution is available that has greater potential to meet current and future needs at an equivalent or lower cost, the organization's decision for cloud is a straightforward proposition. Email is a good example: many organizations have opted for SaaS delivery of a reliable, relatively low cost productivity solution that offers an additional adoption incentive, namely the replacement of constraints on storage in existing infrastructure with flexible storage resources for the end user and a lower TCO for the business.

In other cases, cloud deployment is driven by the end user, who may be looking for additional features that are not currently available via on-premise systems. Rogue use of popular file sharing applications, for example, would suggest that the user culture is ready for cloud adoption, even if the IT department continues to have concerns over loss of control of upgrade schedules, feature sets or the look and feel of the SaaS application. As one member of the working group put it, in assessing cloud readiness from a cultural perspective, it is important to consider the "low hanging fruit in terms of the services that you operate that are not up to current standards," and address these first to build buy-in from the end user community and the business as a whole. Another good indicator of readiness for cloud is the presence of 'workarounds' that users create, such as replication to ensure remote access to applications and data – emailing GBs of data for work outside the office – which may result in additional risk associated with loss of control over where data is stored. As a means of managing this activity, moving to more accessible cloud-based platforms and applications would offer significant benefit as it could simultaneously meet users' productivity requirements and reduce data-related security and audit threats for the organization as a whole.

Cloud readiness may also be evaluated from a risk perspective, where an understanding of business requirements is critical. Key questions the organization may ask to assess risk tolerance include: is the application targeted at a B2B or B2C relationship? What kind of adherence to security practices, policies and standards are required? And what risk tolerance can be entertained for migration of that specific application to the cloud? Answers to these questions may indicate whether on-premise, cloud or hybrid approaches are better options for particular applications.

Beyond this focus on data security needs for specific applications, policy may also play a role. In some organizations, for example, a 'cloud first' policy established at executive levels and based on the balance of generalized risk evaluation and cloud benefit may serve to encourage greater momentum towards cloud delivery.

Key issues in enterprise cloud planning best practices

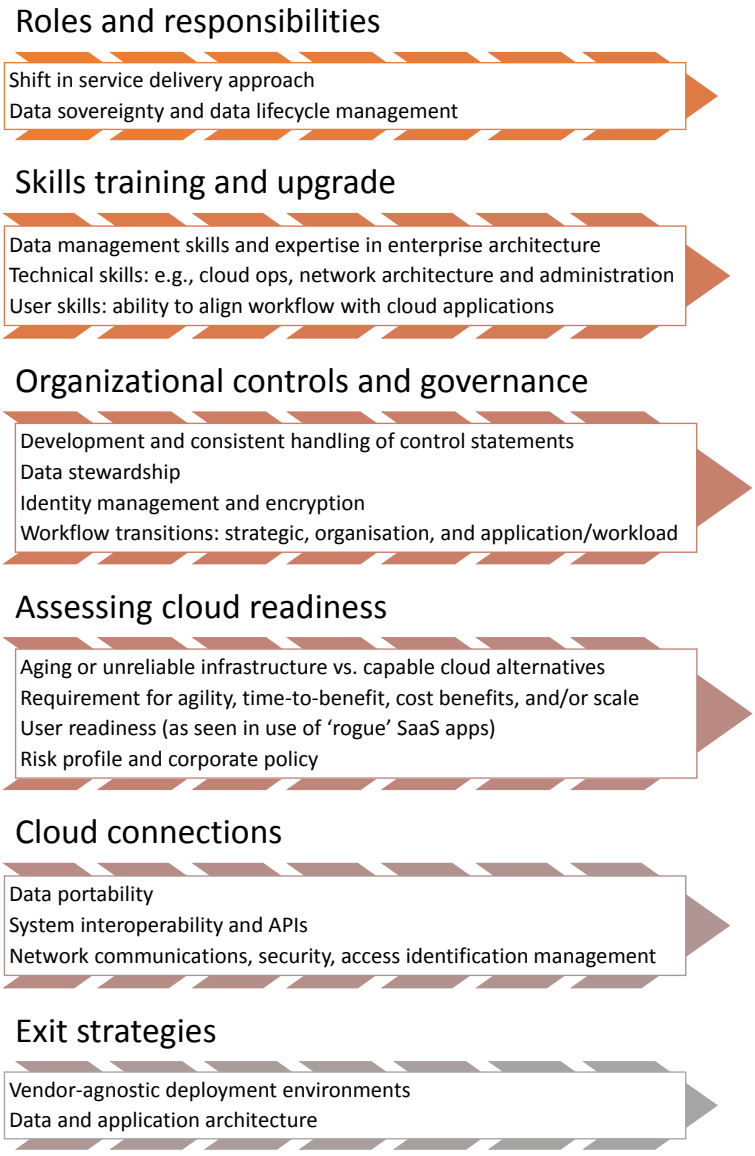

Source: TCBC/InsightaaS, 2015

Moving down the stack, the application's ability to run on cloud infrastructure is another critical factor in assessing cloud readiness. Many legacy applications that were developed using older coding languages are not capable of running without interruption in distributed operating environments that (likely) include virtualized storage and virtualized networking, or of withstanding load balancing in dynamic environments characterized by continuous virtual machine migration. In these situations, the adopting organization may consider a couple of options – rebuild the application using more modern programming languages (Python, Java, etc.), continue to fix problems as they occur, or replace the application. When proprietary applications are differentiated for vertical markets or specific use cases and otherwise unavailable, rebuilding with new software may be indicated. If the application is at end of

lifecycle, the replacement option may be the most appropriate; however, in this situation, the working group advises taking time before migration to properly prepare and cleanse data in advance of transition to cloud.

This focus on the application is also a critical input to an organization's readiness to deploy in public provider resources. Ideally, the enterprise should be able to point to existing business process workflow, and have the same process mapped onto the cloud service as well. Cloud readiness for application migration in the public infrastructure context entails ensuring that the business process flow can also be extended to the cloud environment.

The cloud readiness equation also involves weighing the cost and agility benefits that may substantiate the business case for cloud adoption. Evaluating the relative cost of cloud vs. on-premise deployment is a complex equation that includes estimates for capital expense, staff time and licensing; and the terms fluctuate depending on the number of servers involved, on the server utilization in existing, traditional infrastructure, and by the type of workload or application. The challenge in building IT-centric cost models has led many organizations to focus on specific areas of cost that may be optimized through cloud deployment. As example, the working group cites the case of one large Canadian organization that was finding it difficult to attract new staff talent due to older applications, older infrastructure and the lack of user access to applications. In addition to an IT infrastructure assessment, ROI for this organization included the potential to address operational cost issues associated with talent acquisition.

Much of the discourse around cost control associated with advanced cloud capabilities has focused on the potential to reduce staff engaged in infrastructure maintenance task – "keeping the lights on" – and their reassignment to higher-level activities. This proposition assumes the easy transition of workers. The training and experience of IT professionals would indicate otherwise: a network administrator, for example, will have very different skills than a developer. That said, automation of tasks in everyday operations, such as provisioning servers or terminating unused instances through advanced configuration tools, load sharing across compute, storage and network resources, and dynamic management of the rollout of complex systems at the software layer, can reduce the potential for human error in manual management of these tasks.

While automation of processes such as the provisioning of resources for new employees may be available in more traditional infrastructure, automation is a core competence in cloud-based, software-defined infrastructure that can rapidly decrease the time needed for resource provisioning and consequently increase organizational agility. At the same time, advanced capabilities that are now built into private cloud or service provider public infrastructure may improve an organization's cloud readiness by addressing many of the skill gaps identified as a barrier to an organization's ability to realize cloud benefits. Blueprinting or intelligent architecture solutions can enable a rapid and massive scale-up of cloud resources that cannot be accomplished manually. Cloud orchestration provides the workflow required to optimize this scale of infrastructure, and at the application layer, orchestration works to integrate siloed

resources (including multiple clouds in hybrid scenarios) needed for the organization to take full advantage of the insight offered through the interaction of cloud data and apps. Armed with advanced management tools, which may operate across on-premise, private, hybrid and public cloud environments, organizations may improve cloud readiness and speed time to deployment.

Cloud Connections

The need to manage across silos of information assets is a familiar challenge in organizations with traditional IT infrastructure. Evolving out of enterprise acquisition of new corporate entities with their own systems, operation of individual business units or branch locations with unique IT resource needs and procurement budgets, or the pre-cloud tendency to assign applications and data to specific server resources, information silos are an impediment to business collaboration. In the cloud era, this tendency towards information silos can be exacerbated by the relatively greater ease of adopting cloud resources, and the ad hoc deployment of SaaS apps and cloud infrastructure by business unit managers who are increasingly influential in ever more distributed IT purchase decision making.

But there are techniques that can help address issues with cloud silos. Application Program Interfaces (APIs) are a set of routines, protocols and tools for building applications that provide building blocks for the programmer to use, and specify how software should communicate with operating systems or other control programs. In a cloud context, APIs can support the orchestration of disparate clouds or cloud applications to optimize business operation integration. APIs may also function at a process level, facilitating the user organization's communication with external clouds ease functions such as cloud onboarding. For example, migration to cloud may entail redesign of access to the cloud application, which takes into account risk components such as how access is created, revoked and integrated into user resource provisioning and application lifecycle management. Unless the IT organization has already carried out integration into processes defined by proprietary cloud platforms, an API architecture must be developed that can handle security, sign on and communication between clouds.

In theory, when APIs are leveraged across platforms and applications, adopting organizations can more quickly integrate with cloud provider systems and create integration between clouds. Open APIs accelerate this process, as long as proper documentation on use of the API has been generated from within the open community.

There can be other obstacles to API-based linkages as well. Vendors sometimes deliver connectors that meet the needs of only a small subset of the user community, or that do not address the need for extraction of data – which may compromise data feeds and customized reporting. As a member of the working group explained, "When you first start peeling back the layers to think about how you will integrate, even for organizations that have undertaken automation initiatives, what you can send to the cloud and what you can get back may not be at the level of detail or the quality that would make integration with orchestration

platforms useful." In the access management example noted above, though this may be difficult from a cost or development perspective, the group advises enterprises to "aim high" in their use of APIs for cloud connection: a secured, authenticated link into the API should be established, and the business should maintain the same level of control that is used in manual provisioning of end user access, including the ability to create direct access for many accounts very quickly as well as bidirectional flow of configuration data at a granular level. Though many early stage clouds were built for the unidirectional flow of information – assuming customers would move to vendor platforms but never away – bidirectional data flow enhances usability.

Wherever possible, the working group advises enterprises adopting cloud to look to leverage open, well-documented APIs in their application architectures and to be cognizant of potential issues in leveraging proprietary APIs. A well-integrated API can ease development tasks for the organization, and will also help create a good and transparent user experience for customers of the adopting organization. This is especially important in areas such as epayments, where a bounce back to different sites in transaction processing may compromise the trust an organization is trying to build in ecommerce. The working group notes, however, that transparency to the user does not mean that security and risk aspects of API integration on the back end are neglected. There should be keen focus on identity and access management components, a task that may be further complicated by remote or mobile access to cloud applications and business need for a simplified user experience.

A strategy around single sign on identity management that addresses use of federated IDs or other tools is critical, especially in the situations where solutions are comprised of multiple, integrated cloud systems and applications. Another critical input to ensuring seamless application delivery in a multi-cloud environment (or indeed in use of single clouds) is "BYOK" – "Bring Your Own Key" approaches – in which the user organization maintains the encryption key to its own data, protecting it against the eventuality that one or more provider clouds will be breached. Offered by third party providers as a cloud service, BYOK capabilities may offer the user of integrated cloud applications a simpler approach to securing multiple applications and associated data.

A technique developed by public cloud providers to ease deployment of infrastructure is cloud 'blueprints', otherwise known as 'recipes' or 'templates'. Essentially, the blueprint is an automated script that enables automated, repeatable creation of a set of infrastructure components, including an operating system, application and configuration. Designed to support rapid scale provisioning, blueprints may introduce complexity when the goal is service delivery across heterogeneous clouds, as each blueprint is aligned with a particular vendor platform, and requires familiarity with different native tool sets. Today, a blueprint that is created for a specific cloud is tightly coupled to that cloud: providers use proprietary hypervisors, different distributions of operating systems and applications that are unique to their cloud. For the customer looking to deploy multiple clouds, this conveys a requirement for adaptation to each platform, including learning each native blueprint and the use of separate tool sets. A generic

blueprint system that would allow the user to create a common tool set that could be used across multiple clouds would represent a logical middle step for management of multiple cloud blueprints.

Cloud connections operate at three layers:

- At the base is networking, where moving applications in traditional environments has required reassignment of IP addresses that were tied to physical infrastructure in the networking fabric. By extending switch fabric and routing through software defined networking across multiple environments and clouds, ubiquity of communications is achieved in the IP schema.
- At the next level up is security, where firewall rules and access control lists operate differently from cloud site to cloud site. At this level, a key question is, "is it possible to extend layer 4-7 feature sets, virtualizing capabilities like firewalls and load balancing (as in NFV) so they are not proprietary to a particular cloud service provider premise, but rather ubiquitous across sites?"
- A final layer involves access identification management – BYOK encryption, for example.

To ensure cloud interoperability and the seamless movement of workloads, communication and cabling links must be established across these three areas.

Exit Strategies: understanding how to check out of a cloud environment

Clearly, moving *to* the cloud requires a great deal of planning. However, this planning exercise needs also to consider how to move *from* a specific cloud environment if necessary.

At a high level, the working group advises that enterprises should first consider the architectural decisions that need to be made in order to deploy an application as a service. While there are relatively few difficulties in moving infrastructure components such as the virtual machines from cloud to cloud, the situation is different with applications.

The fastest way to develop an app on public cloud infrastructure is to leverage the provider's platform-dependent micro services; however, in return for accelerated time to value, the user is unable to move the application to another provider's cloud service without breaking and rebuilding the application. While it may be possible to develop an application on what one member of the working group called a "cloud independent platform" to enable shift of the entire application infrastructure to multiple clouds, development may well take longer without the rich deployment feature sets available with proprietary micro services within cloud dependent platforms.

The working group advises enterprises planning deployment of applications on cloud to consider, "how much time, effort and money is it going to cost me to move this application in the event that I have to change clouds?" and what accelerated time-to-market advantage the organization would achieve by tightly coupling the application with the service provider cloud. Netflix, for example, has been three to four years ahead of competitors in terms of bringing

steaming content services to market as it has taken full advantage of AWS services, but the company is now unable to port its service to another provider without significant expense.

What is the best way to the balance of the efficiency benefits gained in writing natively to the cloud provider platform – architecting the solution to take full advantage of unique platform features – against the flexibility won by writing to an agnostic platform? Ultimately, cloud mobility will entail assessment of three decision criteria: what architecture 'gets me there the quickest' (likely to be proprietary to a specific cloud); what architecture provides maximum flexibility to move between clouds; and is more rapid time-to-market worth the cost of vendor lock-in? When an application or service has been built for the cloud and is heavily dependent on provider functionality, rollback of the application may prohibited by complexity or cost. It may be necessary, for example, for organizations looking to re-host an application to rebuild functionality that may only be available as an out-of-the-box feature with a specific provider platform.

In decisions around cloud application mobility, login visibility into cloud provider systems is another important consideration. In the case of a provider breach, for example, it is critical that the user organization have granular login information on the consumption of cloud services in order to develop the required replication, and also to respond to stakeholder questions about levels of organizational risk arising from the provider breach. A provider's ability to deliver this transparency through features built into the application and infrastructure, and confirmed in contractual agreements, should be a key consideration in the evaluation different provider services. This need for transparency may also impact the user's ability to migrate applications from one provider to another: even if the enterprise has granular information from its primary provider, it must have access to comparable data in comparable data formats to allow porting of the application to a next, or additional provider.

Good exit strategies must be grounded in a good backup plan, where either the application or physical environment is in place to roll back cloud services to in-house resources or another provider environment. As the source of migration, the internal user organization may be challenged to establish appropriate staging, testing and migration processes; however, as each component of the application is built, the user should consider if it is possible to roll back the application and data from a specific cloud environment, and how this might be accomplished. This is especially difficult with born-in-the-cloud services – such as the one highlighted in the Netflix example – where rollback represents migration to a net-new environment – essentially, starting over from scratch.

The complexity of the service to be moved in many ways determines the complexity of the migration. Simple applications may be easier to move than more intricate applications, and modern applications may be easier to rebuild than their traditional, legacy counterparts. Going forward, the working group looks forward to an increasing community of third party integration service companies who will build "babble fish" capabilities to enable the rapid and automated re-architecting of modern applications. This approach would offer greater tie-in to new

infrastructure than does today's app migration, and would contribute to application optimization through ongoing introduction of new value-added features. Managed cloud services like these would allow customers to take better advantage of continuous innovation occurring in the provider community.

Ultimately, planning for cloud exit – insurance in the eventuality that cloud applications will need to migrate – returns to considerations of data architecture. Many applications have been built organically and have their own data; however, very few organizations have established common understandings around where data is, who pushes it, who is the steward, where is the data pushed from, where does it reside, and what is being done to secure it. In the working group's view, data architecture considerations are prerequisite to cloud strategy as a whole. A lack of data architecture planning may impact the potential for cloud integration/portability and may also imperil an enterprise's regulatory or governance imperatives. While data governance concerns are common to traditional IT environments, cloud may exacerbate these issues as the application of data governance or policies may be overlooked, neglected or applied after the fact in cloud migration scenarios.

To ensure that data governance is encompassed in cloud planning, the working group encourages enterprises to "think about the exit before you think about the entrance." It is likely that an organization will eventually encounter a situation requiring migration; it is incumbent on the cloud planner to ensure that this does not impose insurmountable challenges to data protection and operational viability.

In the final analysis, exit strategies should involve input from multiple C-level executives. The development of exit strategies should be owned by the CTO, and overseen by the CFO, who addresses risk management by assessing contractual agreements, whether the application can be run without the cloud service, and what might be the impact on the business in the event that the enterprise needs to migrate away for a specific supplier. This evaluation can and should be informed by the CIO and CRO (risk officer), who can offer input on data governance and migration in and out of cloud.

Reference sources

Reference sources helpful in illustrating cloud planning best practices include:

- Governance, Risk & Control. The Institute of Internal Auditors. https://na.theiia.org/standards-guidance/topics/Pages/Governance-Risk-and-Control.aspx
- Michael O'Neil. The Death of Core Competency: A management guide to cloud computing and the zero-friction future. InsightaaS Press, 2014.
  The book is structured in four sections that combine to provide management insight into where and how cloud is important to business strategy as well as guidance on how to align cloud investments and activities to obtain the maximum benefit from cloud

strategies. The book is available as a paperback from Amazon, for Kindle in the US, in Canada, in India, and in all other ebook formats.

About this document/for further information

This document has been prepared by the Planning for the Clou/Cloud Strategy: Enterprise working group of the 2015/2016 Toronto Cloud Business Coalition. Key members of (and contributors to) this working group included:

- Shawn Rosemarin, VMware
- Roy Hart, Seneca College
- Joe Belinsky, Moneris
- Jeff Cohen, Shoppers Drug Mart
- Sangam Manikkayamiyer, Symantec
- Stefano Tiranardi, Symantec
- Chris Vernon, Symantec
- Wil Stassen, Cushman & Wakefield
- Matt Starkie, Microsoft

The Toronto Cloud Business Coalition (TCBC) is a partnership focused on accelerating cloud adoption and use in the GTA and across Canada. It includes individual and corporate members from many different cloud stakeholder communities: IT management from both enterprises and SMBs, global IT and cloud vendors, Canadian 'Born in the Cloud' (BITC) suppliers, ecosystem/channel firms, academics, corporate finance experts, training providers, associations, executives at large with deep experience in the cloud industry, and other experts interested in developing best practices in key areas.

TCBC's activities are underwritten by our corporate members, including:

In addition to this document, TCBC and its members have developed guidance and frameworks on many essential cloud practice areas; we also regularly engage in events ranging from informal discussions to formal, large scale panels and presentations. For more information, please visit our website, or contact us at inquiries@businesscloud.to.

TCBC is operated by InsightaaS.