Practical Guide to Cloud Computing
Version 3.0

December, 2017

Contents

Acknowledgements

The *Practical Guide to Cloud Computing* is a collaborative effort that brings together diverse customer-focused experiences and perspectives into a single guide for IT and business leaders who are considering cloud adoption. The following participants provided their expertise and time to this Guide and/or its Version updates:

Anita Ali (TD Bank), Claude Baudoin (cébé IT & Knowledge Management), Jeff Boleman (IBM), Asher Bond (Elastic Provisioner), Christian Boudal (IBM), Mike Edwards (IBM), Melvin Greer (Intel), Larry Hofer (Cloud and Security Services), Reddy Karri (Schlumberger), Rajesh Jaluka (IBM), Yves Le Roux (CA Technologies), John McDonald (ClearObject), John Meegan (IBM), Jem Pagan (JNK Securities), Bill Parker (FusionStorm), Sujatha Perepa (IBM), Keith Prabhu (Confidis), Gladwin Rao (IBM), Ram Ravashankar (IBM), Karolyn Schalk (IBM), Prasad Siddabathuni (Independent), Gurpreet Singh (Ekartha), Walter Stochevski (UTC-dX), Joe Talik (Neoris), Joel Thimsen (Perficient), Bill Van Order (Lockheed Martin), Amy Wohl (Wohl Associates), Elizabeth Woodward (IBM), Steven Woodward (Cloud Perspectives).

Revisions

Much has changed in the realm of cloud computing since the *Practical Guide to Cloud Computing, Version 2.0* whitepaper was published in April, 2014. Version 3.0 includes the following updates:

- The Essential Characteristics of Cloud Computing has been revised to reflect the characteristics defined in ISO/IEC 17788.
- The Select Cloud Deployment Model(s) step has been rewritten.
- The Determine Who Will Develop, Test, Deploy and Maintain the Cloud Services step has been expanded to include considerations for maintaining cloud services.
- The Assess and Resolve Security, Privacy and Data Residency Issues step has been expanded to include the treatment of data residency management.
- The Develop a Proof-of-Concept before Moving to Production step has been revised to include defined success criteria.
- All other sections have been updated to reflect the evolution and maturity of both the business and technical aspects of cloud computing.

- References have been added to several new CSCC whitepapers and other supporting industry documentation or standards.

Executive Overview

The aim of this guide is to provide a practical reference to help enterprise information technology (IT) and business decision makers adopt cloud computing to solve business challenges. The *Practical Guide to Cloud Computing* provides comprehensive and actionable information in a single reference.

The cloud computing marketplace has evolved in the three years since we published Version 2.0 of this guide. Cloud computing has moved from an interesting experiment to a proven information technology with multiple vendors, large and small, taking a variety of approaches. Most customers have adopted at least some cloud computing technology; others are using cloud computing as their sole IT infrastructure or are moving in that direction.

There is considerable interest, governance, and support for cloud computing from formal IT departments, but there continues to be a presence of Line of Business (LOB) managers spending their IT budgets on cloud services, often without the knowledge or approval of IT management [1]. Also, large numbers of business customers are still choosing their own cloud technologies and expecting IT to support and manage them after the procurement decisions have been made.

This guide provides a way of evaluating the market from the point of view of your organization's needs and providing information that is helpful in selecting both a cloud architecture and an implementation approach through the use of in-house staff, cloud vendor(s) or both.

From an architectural perspective, we have moved beyond the simple alternatives of private vs. public clouds. Hybrid clouds (including multiple clouds, both private and public, temporarily or permanently interconnected) are commonplace.

Today, it is expected that an organization may have any or all of these models, depending on its needs for speed of execution, available resources, various levels of data protection and security, centralized management and an array of other reasons. A discussion of how to choose the most effective cloud service and deployment model is included in this guide.

Another area of rapid development relates to the building of "cloud native" applications and to the technologies and architectures that have evolved to support them. The use of containers, "serverless" computing, and microservices architecture are becoming commonplace. It is necessary to factor in these aspects of cloud computing into decisions as to which cloud services to use.

The "Roadmap for Cloud Computing" section is the heart of the guide. It details both strategic and tactical activities for decision makers implementing cloud solutions. It also provides specific guidance to decision makers on the selection of cloud service and deployment models. The activities and recommendations in the roadmap take into account the different sizes and IT maturity of customer organizations.

Despite the guidelines provided in this whitepaper, the ultimate selection of cloud solutions and their success depend upon the judgement of IT and business decision makers and their organizational and operational realities.

Rationale for Cloud Computing

Cloud computing offers a value proposition that is different from traditional enterprise IT environments. By providing a way to exploit technologies such as virtualization, application containers, and serverless computing which aggregate and share computing resources, cloud computing can offer economies of scale that would otherwise be unavailable. With minimal upfront investment, cloud computing enables global reach of services and information through an elastic utility computing environment that supports on-demand scalability. Cloud computing can also offer pre-built solutions and services, backed by the skills necessary to run and maintain them, potentially lowering risk and removing the need for the organization to retain a group of scarce highly-skilled staff.

Cloud computing does not exist in a vacuum. Most organizations have a broad variety of applications already running in their data center. For most, cloud computing will extend their existing IT infrastructure. Cloud computing can be dedicated to particular tasks. It can be used mainly for new projects or an organization may use it for surge demand, guaranteeing a certain level of performance for enterprise computing.

Essential Characteristics of Cloud Computing

- *Broad network access*: A feature where the physical and virtual resources are available over a network and accessed through standard mechanisms that promote use by heterogeneous client platforms. The focus of this key characteristic is that cloud computing offers an increased level of convenience in that users can access physical and virtual resources from wherever they need to work, as long as it is network accessible, using a wide variety of clients including devices such as mobile phones, tablets, laptops, and workstations;

- *Measured service*: A feature where the metered delivery of cloud services is such that usage can be monitored, controlled, reported, and billed. This is an important feature needed to optimize and validate the delivered cloud service. The focus of this key characteristic is that the customer may only pay for the resources that they use. From the customer's perspective, cloud computing offers the users value by enabling a switch from a low efficiency and asset utilization business model to a high efficiency one;

- *Multi-tenancy*: A feature where physical or virtual resources are allocated in such a way that multiple tenants and their computations and data are isolated from and inaccessible to one another. Typically, and within the context of multi-tenancy, the group of cloud service users that form a tenant will all belong to the same cloud service customer organization. There might be cases where the group of cloud service users involves users from multiple different cloud service customers, particularly in the case of public cloud and community cloud deployments. However, a given cloud service customer organization might have many different tenancies with a single cloud service provider representing different groups within the organization;

- *On-demand self-service*: A feature where a cloud service customer can provision computing capabilities, as needed, automatically or with minimal interaction with the cloud service provider. The focus of this key characteristic is that cloud computing offers users a relative reduction in costs, time, and effort needed to take an action, since it grants the user the ability to do what they need, when they need it, without requiring additional human user interactions or overhead;

- *Rapid elasticity and scalability*: A feature where physical or virtual resources can be rapidly and elastically adjusted, in some cases automatically, to quickly increase or decrease resources. For the cloud service customer, the physical or virtual resources available for provisioning often appear to be unlimited and can be purchased in any quantity at any time automatically, subject to constraints of service agreements. Therefore, the focus of this key characteristic is that cloud computing means that the customers no longer need to worry about limited resources and might not need to worry about capacity planning;

- *Resource pooling*: A feature where a cloud service provider's physical or virtual resources can be aggregated in order to serve one or more cloud service customers. The focus of this key characteristic is that cloud service providers can support multi-tenancy while at the same time using abstraction to mask the complexity of the process from the customer. From the customer's perspective, all they know is that the service works, while they generally have no control or knowledge over how the resources are being provided or where the resources are located. This offloads some of the customer's original workload, such as maintenance requirements, to the provider. Even with this level of abstraction, it should be pointed out that users might still be able to specify location at a higher level of abstraction (e.g., country, state, or data center).

Refer to the ISO/IEC 17788 standard [2] for more details on cloud computing characteristics, roles, deployment models and service models.

The Benefits of Cloud Computing

In addition to providing access to a shared pool of configurable computing resources (e.g., networks, servers and storage), cloud computing promotes a loosely coupled, composable and highly reusable services environment for agile application development and roll-out. Because virtual instances and supporting services can be provisioned and terminated at any time and the customer organization pays only for the computing resource they are employing, costs can be lower, and they are paid “as you go,” not as an upfront capital expenditure.

Cloud computing enables business agility. Cloud computing provides the ability to make use of computing resources on an immediate basis, rather than a need to first invest time and skilled resources in designing and implementing infrastructure (hardware and middleware) and/or applications, and then deploying and testing it. This leads to faster time to value which may mean greater revenue, larger market share, or other benefits.

In essence, the top six benefits of cloud computing can be summarized as follows:

1. *Achieve economies of scale*. Increase volume output or productivity with fewer resources (computing and human).
2. *Reduce CapEx by moving to OpEx*. The "pay as you go" operational expenditure (OpEx) model, based on demand / utility computing, will help reduce capital expenditure (CapEx) on hardware and software licenses.
3. *Improve access.* Information access can be anytime, anywhere and anyhow through omni-channel access.
4. *Implement agile development at low cost.* Design, development and rollout of new solutions and services using agile methodologies on cloud-based shared development operations.
5. *Leverage the global workforce.* A "follow-the-sun" model can provide 24x7 resources for defining, developing and rolling out new solutions. Cloud computing can be rolled out in multiple data centers around the globe, ensuring that services are close to end users – providing better performance, customer service, and appropriate redundancy.
6. *Gain access to advanced capabilities*. Many of the latest advances in software (such as AI, Blockchain, Data Mining) are available off-the-shelf as cloud services, enabling an organization to gain the benefits of these capabilities with minimal investment.

What is the Importance of Standards-Based Cloud Computing?

Standards-based cloud computing ensures that cloud services from multiple providers can readily interoperate, based on open standard interfaces. Standards, when appropriately applied, allow workloads to be readily moved from one cloud provider to a different one with minimum effort. Applications created for one cloud computing environment can be employed in another one, eliminating the need to rewrite or duplicate code.

Many different standards have been proposed. Some permit various degrees of interoperability and portability, while others are in fact proprietary environments. Once a proprietary environment is selected, an organization will probably experience some degree of vendor lock-in. This means that integrating applications or services across differing proprietary cloud platforms is likely to require expensive and time-consuming work. The issues associated with lack of interoperability and portability are described in *ISO/IEC 19941 Cloud computing -- Interoperability and Portability* [3].

Some of the proposed standards are based on open source initiatives, for example the *Open Container Initiative* [4]. This has the advantage of making all the code transparent, available for inspection, and more readily suited for an interoperable environment. However, whenever a new technology is attracting a great deal of attention, neither vendors nor customers are likely to wait for mature standards or rich open source environments. They tend to leverage the advantage of early adoption of emerging technology at the price of having to move to a standard (and perhaps an open source) environment at a later date.

Roadmap for Cloud Computing

This section provides a prescriptive series of steps that customer should take to ensure a successful cloud deployment. It takes into account differences in the sizes of organizations and their IT maturity levels. The following steps are discussed in detail:

1. Assemble your team for cloud adoption
2. Develop a business case and an enterprise cloud strategy
3. Select cloud deployment model(s)
4. Select cloud service model(s)
5. Determine who will develop, test, deploy and maintain the cloud services
6. Develop governance policies and service agreements
7. Assess and resolve security, compliance, privacy and data residency issues
8. Integrate with existing enterprise services
9. Develop a proof-of-concept (POC) before moving to production
10. Manage the cloud environment. Depending on the maturity of the organization and the level of adoption of cloud computing, the entry point will change for each new service being evaluated.

Step 1: Assemble Your Team for Cloud Adoption

It is important that the cloud customer[1] establishes a clearly defined team to develop and approve a cloud business strategy and implementation plan for cloud services that will be part of the total IT environment.

Along with cloud computing, an evolution in the partnership between IT and business leaders is driving more robust collaborative decision making. In the past, the recommendations, design, development, deployment and maintenance of the IT environment were primarily driven by the IT department. Cloud computing is creating an evolution wherein business leaders are getting more closely engaged in IT decisions. To support this digitalization of the business, the IT decision maker must expand their thinking in terms of who the key stakeholders are. More and more often you see roles like the Chief Marketing Officer (CMO), Chief Data or Digital Officer (CDO), or Chief Cloud Officer (CCO) influencing IT budgets.

Adoption of cloud services is increasingly viewed as a strategic business decision that allows business not only to improve IT efficiency but also help in achievement of global business goals.

While requirements around cost, performance, longevity, and feature/functionality roadmap are still critical, the modern CIO must also take into consideration input from such key stakeholders as the CEO, CDO, CFO, CISO, CMO, and Development or Product Leadership. The following are some examples of what the C-Suite will be looking for:

[1] Note that this section focuses on the cloud customer. Points of contact between the cloud customers and cloud providers, brokers, and carriers are covered in the later steps of this section.

- CEO: Co-driving innovation and differentiation capabilities within the organization enabling the CEO to focus on growing the business. Cares about finding new ways to differentiate through technology.
- CDO: Mapping digital capabilities to strategic priorities, enabling new functionality or efficiencies, measuring efficiencies and ROI; re-engineering/automation of business processes can improve costs and promote revenue generation. Consolidated monitoring of digital initiatives is critical.
- CFO: Co-driving data privacy, maintenance of financial records, and security of financial reporting systems. Cares about the growth aspect of the finance agenda, while delivering efficiencies through technology adoption.
- CISO: Co-driving information management, risk management, brand protection, third-party relationship management, and other functions beyond their historically technical role. Cares about standardization and security being integrated throughout the technology stack; from physical through logical; IAM, DLP, IDP/IDS.
- CMO: Co-driving Digital Transformation. Cares about the customer experience; new customer touch points, new opportunities for engagement, and more integrated experiences.

This is happening because the business is increasingly seeing cloud computing as a tool to:

- Get closer to their customers
- Improve IT efficiency and Improve product time to market
- Balance the financial portfolio (CapEx vs. OpEx)
- Increase sales/revenue
- Streamline the supply chain
- Extend business processes, making them more accessible by third parties
- Reach new customer segments
- Partner data sharing/integration opportunities
- Standardize the technology stack and related security
- Leverage cloud native security capabilities

Bottom line, the adoption of multiple cloud consumption models for infrastructure needs is being viewed more and more as a strategic business decision. Hence, it is logical that the adoption of cloud computing should be led by senior management with a broader section of the C-Suite playing the role of key advisors.

In essence, the team you build will need to address various aspects of adoption. These resources participate to different degrees in three phases of cloud adoption described in Figure 1. Different skills are required at the different phases of cloud adoption—strategic, tactical and operations planning.

Figure 1: Three Phases of Cloud Adoption

Strategic Planning Phase

During the strategic planning phase of cloud service adoption, CEOs and the senior management team lead the organization to establish the vision, terms of reference, and guidelines.

- *Vision*. The CEO and senior management team should define the overall vision for cloud adoption. It is critical that the business leadership, particularly the executive levels, collaborate and buy in to the vision. The vision should address the future of the business and the acquired differentiation, competitive advantage, and/or value proposition gained through a cloud strategy.
- *Terms of reference.* Terms of reference should be defined early to ensure that adoption stays focused on the target business goals. Effective terms of reference should at minimum address the purpose, goals, guiding principles, roles and responsibilities, and rules of engagement of the teams involved in forming the cloud computing vision and strategy.
- *Guidelines.* Based on the culture of the business, senior management should provide broad guidelines for cloud adoption, including security and privacy posture, data maintenance and location policies, etc. The guidelines will provide a business framework to capture the initial high-level requirements, and support the alignment of the leadership team.

Workload Planning Phase

During the tactical planning phase, typically led by the CIO or CTO, the organization performs both a business and a technical analysis.

- *Business analysis.* This phase requires the oversight of senior business stakeholders such as the CMO, product or application owners, IT (including as necessary the CIO, CTO, CISO, lead architects and business continuity manager), and legal representatives who must review and communicate regulatory compliance and legal requirements. The overall goal is to build a business case and the supporting long-term enterprise strategy for the transition to cloud computing that delivers sufficient return on investment or return on value.

  It is critical that the business drivers for transformation remain the primary focus during this phase. Technology comparisons and early calls on the technology solution could negatively impact a sound understanding of the business requirements if introduced too early in the process.

This phase will provide the foundation for categorizing and identifying suitable applications or workloads to be considered as candidates for implementation in the cloud.

- *Technical analysis.* This phase requires the attention of IT (or digital) stakeholders including the CIO, CTO, CDO, CCO, CISO as well as lead architects, operations personnel, IT security, and senior business managers. The goal of this phase is to develop a Transformation Roadmap which in addition to the Technology Roadmap would include components such as Critical Success Factors, Service Management Capability Maturity Roadmap, Organizational Functional Model, IT Architectures and Operations Models, Processes and Capabilities, to name a few.

On the Technology Roadmap, development of the target consumption models takes the business requirements analysis results from this prior phase, and compares it to the various cloud services and deployment models that are available. It is during this phase that the business case for application or workload migration, as well as application "build versus buy" is developed and presented to the strategic team for their feedback. Business processes and flows will likely also be disrupted and impacted. Business processes and flows will likely also be disrupted and impacted in non-technical areas such as procurement and technical areas such as provisioning of services.

Operations Planning Phase

During the operational implementation phase, leaders from various operations groups work through procurement and implementation details, and establish ongoing operational plans and processes for the cloud deployment.

- *Procurement.* This phase includes negotiations with potential cloud service providers and requires the procurement team, finance, legal, senior business managers, and IT including the CIO, CTO and lead architects to be engaged. Assessment models/instruments should be discussed to provide consistency during the evaluation and contract negotiation phases. For example, a proof-of-concept or weighted scoring matrix may be developed to align business and technology requirements.

- *Implementation.* This phase includes the development, customization, and configuration of solutions which will be deployed in the cloud environment and requires the attention of IT including lead architects, developers and testers as well as operations personnel. This phase should begin the process of identifying potential changes in processes/procedures and provide a high-level gap analysis to identify risk mitigation/management opportunities.

- *Operations.* This phase addresses ongoing operations and management of the cloud services and deployed solutions. Business owners, operations personnel, customer support, and IT including developers and testers are required in this phase. During this phase, the gap analysis moves from a high level to a more detailed examination to determine what (and how) changes to existing operations will be impacted by the cloud computing initiative.

Step 2: Develop a Business Case and an Enterprise Cloud Strategy

To ensure a smooth transition to cloud computing, an organization should develop an overarching cloud strategy which creates the foundation for project-specific adoption. Cloud computing presents interesting business model opportunities to organizations of all sizes. Developing a business case and strategy that clearly articulates how cloud computing will transform key business processes like procurement, marketing, customer acquisition and support, product development, etc. is critical.

Within the context of an enterprise strategy for cloud computing, individual business problems that cloud computing can potentially address need to be identified, and specific business justification must show that cloud computing is the right strategic alternative. High level value propositions for cloud computing, including the shift of capital expenditures (CapEx) to operational expenses (OpEx), cost savings, rapid deployment, elasticity, access to advanced capabilities, etc., are necessary but insufficient unless quantified.

Obtaining executive support for the initiative is critical. Executives from IT, Lines of Business (LOBs), procurement and executive management must review and approve the business plan before proceeding. Getting key executives on-board early in the process will help alleviate potential issues down the line.

When developing an enterprise strategy for cloud computing, the considerations highlighted in the following table should be taken into account.

Table 1: Key Elements of Strategic Planning

Element of Strategic Planning	**Strategic Planning Activities**
Educate the team	• All team members (IT, business, operations, legal and executives) must be educated on what cloud computing is and what it is not. • Establish a common definition of cloud computing (including terminology) for the entire organization so everyone is speaking the same language. • Using cloud computing is an iterative process in which new services build on previously implemented services add value to existing IT environments.
Establish both short and long term plans	• Create an organization-wide master blueprint and roadmap for adoption. • Map cloud computing benefits against existing business problems to identify potential solution areas. • Anticipate the variety of disruptions that may occur both inside and especially outside IT (service levels, security, legal, vendor management, etc.). • Leverage long-term planning to reduce risk of vendor lock-in by considering interoperability, portability and ease of integration up front.

Element of Strategic Planning	**Strategic Planning Activities**
Understand required services and functionality	● Determine business case and potential ROI and/or potential new revenue opportunities. ● Leverage enterprise architectures, standards and industry frameworks to help accelerate the collection of service information and improve consistency. ● Customer facing services require separate categorization and analysis from internal services.
Execute a thorough cost analysis [5]	The overall cost of application migration to cloud computing must include the following elements: ● On-going cloud service costs ● Service management ● License management ● Application re-designs ● Application deployment and testing ● Application maintenance and administration ● Application integration ● Cost of developing cloud computing skills ● Human resources and talent management implications ● Additional tools/ services/ processes
Assess the impact to service levels [5]	For each application being migrated to cloud computing, consider the impact on the following application characteristics: • Application availability • Application performance • Application security • Privacy • Regulatory compliance • Data residency
Identify clear success goals and metrics to measure progress	● The team sponsoring the project must include success factors in their proposal. ● Metrics need to be agreed to by executives making the final decision to proceed with the project. ● Define baselines for the existing service before launching the new service in order to determine its impact. ● Clearly identify trigger points to be measured. ● Develop a cloud adoption roadmap.

Element of Strategic Planning	Strategic Planning Activities
Consider the existing IT environment	• Develop a complementary cloud adoption strategy with a focus on integrating and leveraging existing technologies and standards. • Develop a strategy to ensure that any existing services to be migrated to cloud computing will continue to comply with standards. • Leverage reusable internal services to improve delivery efficiency of customer facing services.
Assess the current operational support model	• Validate the impact on existing operational support models fused with cloud enabled operational support model. • Develop a strategy to Integrate support models. • Develop a strategy to coexist based on the scale of cloud adoption. • Develop a strategy to interoperate and the information required to do along with tools that will be required.

Understand legal/regulatory requirements	• Customers of cloud services must understand the responsibilities associated with their respective national and supranational obligations for compliance with regulatory frameworks and ensure that any such obligations are appropriately complied with. Some examples of legal/regulatory constraints upon electronically stored information are as follows:
	  ◦ Physical location of the data
	  ◦ Data breach
	  ◦ Personal data privacy
	  ◦ Data destruction when the corporation no longer wants the relevant data available or transfers it to a different host
	  ◦ Intellectual property, information ownership
	  ◦ Law enforcement access
	  ◦ Service availability
	• With over 150 countries having ratified the United Nations Convention on the Rights of Persons with Disabilities and an increasing focus on accessibility regulations, it is important to establish a plan for ensuring accessibility compliance.
	• Understand cloud service specific deployment standards and compliances required by various industries, for example, FISMA and FedRAMP for U.S. Federal government agencies.

Identify required skills	• Map required skills against available skills. • Develop a plan to enhance internal skills to address potential gaps. • Consider external skills as an option for addressing gaps.

Element of Strategic Planning	**Strategic Planning Activities**
Track results for an extended time	● Reinforce that the objective of implementing the new cloud service has been achieved. ● Identify any trends that may need to be addressed to improve the existing service or contract for a new service to take advantage of the trend.
Understand the exit process [6] [7]	● An exit clause should be part of every cloud service agreement. ● Understand the details of the exit process including the responsibilities of the cloud service provider and cloud service customer. ● The exit process should include detailed procedures for ensuring business continuity. It should specify measurable metrics to ensure the cloud provider is effectively implementing these procedures. ● The most important aspect of any exit plan is the retrieval and preservation of cloud service customer data.

Step 3: Select Cloud Deployment Model(s)

In order to determine the cloud deployment model(s) that best suits your company's business requirements you must take into consideration the factors in Table 2[2].

It is important to note that "private cloud", where cloud resources are dedicated to one customer, can take one of two forms: 1) on-premises private cloud, where the cloud environment is implemented on the premises of the customer, and 2) private hosted cloud (sometimes called dedicated cloud) where the cloud environment is on the premises of a cloud service provider, but where the resources involved are not shared with any other customers.

Table 2: Considerations for Selecting a Cloud Deployment Model

Consideration	Private	Hybrid	Public
Criticality of cloud services	Private clouds are appropriate for mission critical applications and compliance sensitive services necessary for business continuity. Critical data availability is key to deciding whether to keep the workload on-premise.	Hybrid deployments help take advantage of public cloud features for certain non-critical workloads, while retaining business critical data and applications on-premise.	Public clouds are more appropriate for services that are not mission critical and do not require access to sensitive information. They can also be more appropriate for organizations that do not have the resources to ensure high availability of on-premises systems.
Type of workload	Private on-premises may be preferable for applications that have very stringent latency requirements.	Cloud bursting of on-premises core business capabilities during seasonal surge is an example; replicating selected customer information to a lightweight cloud database for quicker access by mobile apps is another.	Public cloud is suitable for workloads that require access to high volume data (e.g., real-time analytics running against very large data stores), for workloads with highly variable load patterns, and for access to advanced services that may be difficult to implement on-premises.

[2] The information in this section is based on information from the NIST Cloud Computing Synopsis and Recommendations document, Special Publication 800-146 [8]. The Community deployment models are not called out explicitly in this section since they are similar to the Private deployment options.

Consideration	Private	Hybrid	Public
Migration costs	With a Private deployment model, installing and managing cloud software may incur significant cloud software costs even if non-allocated hardware exists within a consumer organization. Expenses may be mitigated if the organization has adopted a service oriented architecture environment and moves to an expense formula for internal departments.	Hybrid deployments, by definition, are transitory between private and public, and hence a point in migration to/from public cloud. Costs can be controlled based on need and maturity.	Public clouds have low upfront costs for the use of cloud services. The implications are similar to the outsourced private cloud scenario except that additional security precautions need to be taken into account.
Elasticity	With Private (On-premise), finite resources are available since computing and storage capacity is fixed and has been sized to correspond to anticipated workloads and cost restrictions. If an organization is large enough, it may be able to provide enough elasticity to clients within the consumer organization.	Hybrid deployments typically include a component of public cloud services; availability of resources is only limited by acceptable security limitations.	Public clouds can generally be considered unrestricted in their size. Additionally, they can generally use multi-tenancy without being limited by static security perimeters, which allows a potentially high degree of flexibility in the movement of customer workloads to available resources.
Security threats	With Private (On-premise), consumers have the option of implementing appropriately strong security to protect resources against external threats to the same level of security as can be achieved for non-cloud resources.	Organizations utilizing Hybrid deployments can choose to limit the kind of data/services that are exposed to the public, thus helping mitigate threats.	With a Public model, customers have limited visibility and control over information regarding security. The details of provider system operation are usually considered proprietary and not available for examination by customers. Certification of cloud services may provide a level of assurance to customers.

Consideration	Private	Hybrid	Public
Multi-tenancy	With Private, risks are mitigated by restricting the number of possible attackers: all of the clients would typically be members of the customer organization or authorized guests or partners.	Multi-tenancy on the public part of a Hybrid cloud can be limited to either periodic use (bursting) or by exposing limited information/services to the public cloud. Sensitive and mission critical areas can still reside on-premises to minimize risks.	With a typical Public model, a single physical machine may be shared by the workloads of any combination of customers. In practice, this means that a customer's workload may be co-resident with the workloads of competitors. This introduces potential reliability and security risks, though evolution of technology and practices have helped reduce both.
Compliance	Organizations using an on-premise model typically have more control but also more responsibility in ensuring compliance with various industry and government standards (HIPAA, GDPR, etc.) since they control the entirety of the infrastructure.	Hybrid deployments help organizations stay compliant by providing the customer with the ability to choose the environment most suitable for compliance requirements – on-premises or public cloud. Care needs to be taken however, to ensure compliance where solutions transition between the two.	Many public cloud services are explicitly certified for compliance with one or more standards and/or regulations. It is necessary to select cloud services with appropriate certifications - or cloud services which the customer can get certified appropriately.
Environment portability	For on-premise deployments, portability is unlikely to pose a significant challenge.	Portability is typically an issue with the public cloud part of a Hybrid deployment. In addition to avoiding potential vendor lock-in, organizations need to ensure access via API and a streamlined integration approach to all parts of the applications/infrastructure.	Organizations need to investigate portability and minimize risk of vendor lock-in before they proceed with deployment on public clouds.

Consideration	Private	Hybrid	Public
Disaster Recovery/ Failover	Depending on the size and maturity of the organization and the nature of the application, DR and failover provisioning might involve significant costs and workload. It might require multiple physically separated data centers, for example.	Organizations might choose to leverage public cloud to act as a failover option for their internal applications, or for disaster recovery, but they need to ensure regular synchronization between on-premises and public cloud systems.	Many large cloud service providers have multiple data centers across the globe and provide DR and failover options with standardized and documented procedures.

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

Today, hybrid cloud deployment is commonplace. An organization leverages a combination of private and public cloud deployments depending on its needs for speed of execution, available resources, need for data protection and security, and an array of other reasons.

For example, "cloud bursting" is a concept in which a consumer uses on-premises IT resources for routine workloads but optionally accesses one or more external private or public clouds during periods of high demand. Different cloud deployment variants may also be appropriate for particular organizational functions or roles. For example, an organization may elect to process sensitive data such as payroll information in a private hosted cloud but use a public cloud for new software development and testing activities.

The IT maturity of an organization along with its size will have a significant impact on the service deployment decisions that are made:

- Larger organizations with mature IT environments may lean initially towards private cloud deployments and may transition some workloads to hybrid and public deployments over time.
- SMBs and new companies without existing infrastructure may transition more rapidly to public cloud deployments. SMBs have much to gain in terms of cost savings, IT capacity, lower skill requirements and improved application functionality that was not available to them previously. Security issues with Public deployment must be taken into consideration. As a result, SMBs are advised to initially consider Hybrid deployments, putting new or modified applications to public deployment in the early transition phases. New companies without existing infrastructure have the ability to quickly grow their workloads without the startup costs of a data center.

For a more detailed overview of Hybrid Cloud Computing, please refer to CSCC whitepaper, *Practical Guide to Hybrid Cloud Computing* [9].

Step 4: Select Cloud Service Model(s)

While the business value of cloud computing is compelling, many organizations face the challenge of staging a gradual adoption of cloud service capabilities, incrementally advancing their IT environment.

There are a variety of ways that organizations today are leveraging the benefits of cloud computing. Many patterns of implementation start with an infrastructure virtualization or application containerization project to establish a foundation that enables future cloud service adoption. Conversely, some companies are simply consuming business or IT solutions from a public cloud outside their organization.

As depicted in the figure from the *NIST Cloud Computing Reference Architecture* [10], the three most common cloud service models are Software as a Service (SaaS), Platform as a Service (PaaS), and Infrastructure as a Service (IaaS). To determine the service models that best suit your company's business requirements, the potential benefits and issues of each model must be given careful consideration. Additionally, the IT maturity of an organization along with its size will significantly impact the service model decisions that are made.

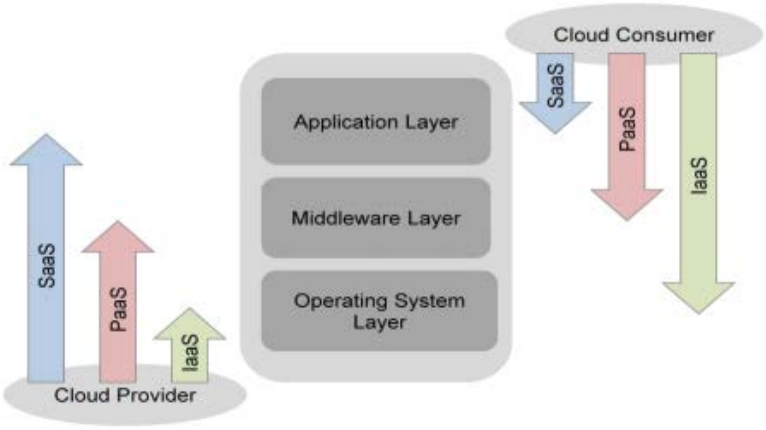

Figure 2: NIST Cloud Computing Reference Architecture - service models

Software as a Service (SaaS)

SaaS involves the acquisition of a complete application or business service running as a cloud service. It is the cloud computing equivalent of buying a packaged application; one that typically requires minimal configuration before it is ready for use. SaaS allows businesses to benefit from the "pay-as-you go" concept in addition to being highly scalable, offering flexibility to companies to provision and de-

provision based on business needs. This consumption based model for software eliminates many of the high start-up costs for initial licensing and installation of software delivered in a traditional model.

SaaS gives businesses complete freedom from managing IT infrastructure and the entire software stack which enables them to concentrate on using the features of the service to achieve their business objectives. Business solutions implemented as cloud services provide customers the flexibility to choose the approach that is best for their company by making it possible to consume and execute business processes, analytics, and applications in the cloud.

SaaS can be categorized under two broad headings:

- *Horizontal SaaS offerings.* These are SaaS offerings that are typically applicable to organizations across a range of business sectors. Some of the common SaaS applications are in the areas of email, customer relationship management (CRM), productivity, collaboration, human resources (HR), analytics, etc.
- *Sector-specific offerings.* With the proven success and maturity of the horizontal SaaS offerings, sector-specific SaaS offerings are emerging. These include applications in the areas of logistics and supply chain management (SCM) and healthcare, for example.

SaaS has the following key features:

- SaaS offerings are accessible over the public Internet which makes it very easy to roll them out to a large audience within a short period of time.
- SaaS works on a usage-based pricing model which enables businesses to subscribe to only those services that it needs and for the required number of users.
- SaaS typically offers a standard feature set which allows some level of configuration for individual customers but typically no customization.
- Organizations can reduce their capital expenditures (CapEx) on the procurement of software licenses by adopting SaaS offerings on a subscription basis.
- Deployment of SaaS offerings is typically much shorter than deployment of traditional packaged solutions. This enables businesses to make use of any short “window of opportunity” that may present itself.
- SaaS upgrades are typically instantaneous and service providers are responsible for deployment. They get tested prior to deployment and the process is transparent to the users.
- SaaS offerings are typically scalable as vendors plan for scalability in their cloud solutions. This enables businesses to scale up rapidly if the business needs dictate.

- Security and privacy issues are the responsibility of the cloud service provider. Cloud service providers have a strong business requirement to ensure that their solutions deal with these issues properly.

- Availability of the solution, including the backup of customer data, is usually handled by the cloud service provider, thereby eliminating the need for users to maintain their own disaster recovery procedure for these solutions.

Approaches for Adoption of SaaS

The approach for adopting SaaS offerings will differ based on the IT maturity of the organization. For simplicity, two approaches are described below: one for large organizations and one for SMBs. However, given that each organization is unique with its own challenges, it is recommended that organizations evaluate both options and determine a strategy that addresses their unique requirements.

Organizations with mature IT systems have already implemented in-house packaged applications. Having spent years with these systems and making significant investments in hardware, software and management, they are reluctant to drop these systems which have stood the test of time. These applications often have the highest costs per unit of functionality to enhance, support, and operate, therefore significant cost reductions can be realized. Unfortunately, migrating this class of application (legacy or non-virtualized) will incur higher project costs.

Organizations with nascent IT systems, especially SMBs, may not have made significant investments in IT. The reasons for this vary from cost concerns to the complexity of managing such deployments in-house. However, with the emergence of subscription-based SaaS offerings, such organizations now have an option to adopt SaaS solutions for business needs which was not possible previously.

Table 3: SaaS Adoption Approaches

SaaS Adoption Approach for Large Organizations	SaaS Adoption Approach for SMBs
Large organizations can take the following approach to SaaS adoption: 1. Analyze SaaS offerings in terms of Total Cost of Ownership (TCO) & Return on Investment (ROI) and risks such as vendor lock-in, interoperability, and existing IT infrastructure — especially network and data center infrastructure. 2. Define a clear SaaS strategy for both private and public implementations before adopting specific SaaS offerings. 3. Consider SaaS for global business functions that would deliver improved ROI in a cloud environment. Consider SaaS for rapidly evolving business environments where new requirements are likely to emerge, such as social business and Web campaigns. 4. Evaluate SaaS offerings when packaged applications need to be renewed due to a software or hardware refresh which involves additional purchases. 5. Adopt new disruptive SaaS solutions (perhaps sector-specific) to maintain or extend competitiveness.	SMBs can take the following approach to SaaS adoption: 1. Analyze SaaS offerings in terms of TCO/ROI and risks, such as vendor lock-in, interoperability, and existing IT infrastructure — especially network and data center infrastructure. 2. Define a SaaS strategy for both private and public implementations before adopting specific SaaS offerings. 3. Re-evaluate business processes and identify those that can be enhanced through use of applications that can help improve competitiveness with larger organizations. 4. Identify availability of SaaS offerings for these specific processes. 5. Evaluate the various SaaS offerings from a business and technical perspective.

Platform as a Service (PaaS)

PaaS provides an integrated development and runtime platform for creating, deploying, and managing custom applications in a cloud service. Based on the standardization and automation of a common set of topologies and software components, the platform provides elasticity, efficiency, and automated workload management. A PaaS environment dynamically adjusts workload and infrastructure characteristics to meet existing business priorities and SLAs. The big advantage of a PaaS offering is that it provides a ready-deployed software stack that caters to the development and deployment of custom applications in a cloud computing environment, sharply reducing the effort required by developers and operations staff.

PaaS helps eliminate the need for developers to work at the image-level, enabling developers to completely focus on application development. It also helps reduce software design steps and enables faster time-to-market using predefined workload patterns.

For a complete understanding of PaaS environments and how they are used, see the CSCC whitepaper, *Practical Guide to Platform-as-a-Service* [11].

The incentives for an organization to transition to a PaaS environment differ based on the size and IT maturity of the organization. For large organizations, a key motivation for considering PaaS is the ability to quickly and inexpensively develop and deploy new applications. Large organizations have additional incentives for considering a move to PaaS. PaaS provides:

- Highly standardized and automated provisioning of predefined workloads
- An integrated development and runtime platform for specific workloads
- Consistent pattern-based deployments for the most common workloads
- Strong support for cloud native applications, including technologies such as containers, "serverless" computing, and microservices
- DevOps capability that facilitates communication, collaboration and integration between software developers and IT professionals. Refer to Step 6 of the CSCC whitepaper, *Convergence of Social, Mobile and Cloud: 7 Steps to Ensure Success* [12]
- Integrated workload management for SLA enforcement, dynamic resource management, high availability and business priorities
- Awareness and optimization of workloads based on business priorities and SLAs
- Consolidation of workloads under a simplified management system

Smaller organizations can benefit from the ready provision of running instances of major portions of the software stack required by a custom application, such as databases and messaging infrastructure, removing the need to have skilled staff to set up, run and maintain what can be complex software.

Approaches for Adoption of PaaS

Organizations with mature IT systems have made significant investments in their development and runtime platforms along with significant investments in human resources associated with solution development and testing. As a result, they will initially look to refactor these assets as they transition to cloud computing. Custom applications which may benefit from the use of a PaaS offering could include new applications written to support mobile devices or applications which support social computing.

In many cases, SMBs do not possess the resources to invest significantly in development and runtime platforms and they lack the in-house human resources to develop and test home-grown applications. Many SMBs are reliant on Independent Software Vendors (ISVs) to deliver their application functionality. As a result, they might be dependent on a cloud service provider to support a PaaS environment that is consistent with their ISV's applications.

Table 4: PaaS Adoption Approaches

PaaS Adoption Approach for Large Organizations	PaaS Adoption Approach for SMBs
The following steps provide a recommended approach for PaaS adoption by large organizations: 1. Analyze PaaS offerings in terms of total cost of ownership (TCO) / return on investment (ROI) and risks such as vendor lock-in and interoperability with existing IT infrastructure. 2. Define a clear PaaS strategy for both private and public implementations before adopting specific PaaS offerings. 3. Identify early offering candidates based on specific criteria (for example, low risk to the business). 4. Consider starting with either the Private (On-premise) or Private (Outsourced) deployment model which provides a good initial transition to PaaS for both mission critical and non-mission critical workloads with relatively low risk. 5. Consider a platform that leverages existing expertise – i.e., a development team experienced in Java will likely gravitate to a Java-based platform.	The following steps provide a recommended approach for PaaS adoption by SMBs: 1. Analyze PaaS offerings in terms of TCO/ROI and risks such as vendor lock-in, interoperability and existing IT infrastructure. 2. Define a PaaS strategy for both private and public implementations before adopting specific PaaS offerings. 3. Determine if there's sufficient in-house development resources to justify the use of a PaaS environment – if not, SaaS may be the best alternative.

If sufficient in-house development resources exist, both the Private (Outsourced) and Public deployment models are viable options. Selection will be dependent upon the nature of the applications being developed and deployed.

Infrastructure as a Service (IaaS)

The incentives for an organization to transition to an IaaS environment differ based on the size and IT maturity of the organization. For SMBs, the primary motivation for considering IaaS is capital expense reduction and access to IT capacity that would otherwise not be available. For large organizations with potentially several data centers and departmental silos in different geographical locations, there are additional incentives for considering a move to IaaS. Incentives include addressing low server utilization, high administrator-to-server ratios, data center sprawl, proliferation of ad-hoc IT solutions, and desire for improved, more centralized control of IT assets.

While many organizations today are using virtualization to consolidate their IT infrastructures, hardware consolidation is only one benefit of virtualization. Organizations that move beyond virtualization with

IaaS capabilities such as integrated service management, automation and rapid provisioning can realize significant benefits:

- Reduction in IT operating expenses and capital expenses by improving resource utilization and administrator-to-server ratios
- Faster time to market through increased efficiency and automation of standardized solutions
- Simplified, integrated management, including real-time monitoring and high-scale low-touch provisioning
- Greater visibility into business processes and system performance to identify redundancies and bottlenecks
- Scaled operations that can meet market dynamics and business strategy

Approaches for Adoption of IaaS

Organizations with mature IT systems already have significant investments in both infrastructure hardware and in-house IT management skills. As a result, they will initially look to refactor these assets as they transition to cloud computing.

Organizations with nascent IT systems, especially SMBs, may not have made significant investments in their IT systems. As a result, they will be incented to transition more rapidly to infrastructure services that are delivered and managed by an external cloud service provider.

Table 5: IaaS Adoption Approaches

IaaS Adoption Approach for Large Organizations	**IaaS Adoption Approach for SMBs**
The following steps provide a recommended approach for IaaS adoption by large organizations: 1. Analyze IaaS offering in terms of total cost of ownership (TCO)/return on investment (ROI) and risks such as vendor lock-in, interoperability and existing IT infrastructure. 2. Define a clear IaaS strategy for both private and public implementations before adopting specific IaaS offerings. 3. Start with an infrastructure virtualization project to establish a foundation that enables future cloud adoption. 4. Consider moving to a Private (On-premise) deployment model which provides a good initial transition to IaaS with relatively low risk. 5. Consider Private (Outsourced) and Public deployment models which can potentially deliver added business value. Closely consider security and reliability issues as well as integration with existing enterprise services. 6. Consider Disaster Recovery capabilities if needed 7. Consider managed services to re-focus IT teams on more added value activities	The following steps provide a recommended approach for IaaS adoption by SMBs: 1. Analyze IaaS offerings in terms of TCO/ROI, risk (vendor lock-in/interoperability/existing IT infrastructure). 2. Define an IaaS strategy for both private and public implementations before adopting the IaaS offerings. 3. In many cases, the Private (On-premise) deployment model will not be feasible given insufficient ROI associated with consolidating a relatively small number of existing IT assets. 4. Consider the Public deployment model which provides access to computing and storage capacity at the lowest cost. 5. Application migration and administration costs must be considered for Public and Private (Outsourced) options. 6. Consider Disaster Recovery capabilities 7. Consider managed services to re-focus IT teams on more added value activities

Large organizations should consider the following types of projects as good candidates to begin transitioning to a Private (On-premise) cloud-enabled data center:

1. *Consolidate and virtualize your infrastructure*. Realizing the benefits of cloud computing begins with the foundation, including the efficient and effective consolidation and virtualization across server and storage platforms, to begin building a cloud infrastructure.

2. *Standardization through image management.* Reducing the sprawl of versions through image management addresses the visibility and control of virtualized images to reduce operational costs associated with virtualization proliferation in the data center. Image management allows clients to better utilize virtualization as an enabler of standardized high-quality service delivery. When clients implement effective image management, they are better suited to progress into a

cloud computing model.

3. *Management of the virtual environment.* Organizations can expand beyond infrastructure virtualization with integrated service management, automation, provisioning and self-service to more quickly deploy IT services, increase resource utilization, and better manage their cloud environments and ultimately reduce operational cost.

As illustrated in this section, there are numerous considerations that need to be evaluated when selecting a service model that best meets your company's business requirements. An effective initial approach is to identify a contained business area where cloud could be impactful, identify one or more cloud service models that could be effective in addressing the requirement, and initiate a proof of concept to assess the feasibility and ROI of the alternatives.

Step 5: Determine Who Will Develop, Test, Deploy and Maintain the Cloud Services

Determining the most effective method to design, develop, deploy and maintain new cloud applications can be a struggle. In many cases, there is no right answer. The direction will be based on the needs and capabilities of the organization. There are essentially four options for the organization to consider[3]:

- In-house development and deployment
- Cloud provider development and deployment (if available)
- Independent cloud service development provider
- Off the shelf purchase of a cloud application service

Table 6 examines the pros and cons of the various options for acquiring a new service.

[3] The design of the cloud application service is omitted since that should originate from the enterprise and will require the efforts of the IT, business and administrative teams. The new service must have functional capabilities which meet the requirements of the target users and will have also a positive ROI. Designing a cloud service is an extended discussion which will not be covered here, other than to state that ensuring that the design process is followed will be critical to the development and deployment activities.

Table 6: Options for Acquiring a New Cloud Application Service

Options for acquiring a new service	Skills	Startup considerations	Updates to services	Testing, deployment and support
In-house development and deployment	Dependent on internal skills, acquired training, and availability of in-house resources to develop and support new applications.	Should reduce the learning curve on how to link to legacy services.	The enterprise owns the cloud application and can incorporate future updates and maintenance based on their internal processes and schedule.	Offers potentially tighter controls and governance during the testing, deployment and support process. In-house test and operations managers can work closely with IT and business leaders to ensure thorough support is provided.
Cloud provider development, deployment and support	The cloud provider's area of expertise is cloud computing which should translate into a shorter development and deployment timeline especially with the first cloud service.	A cloud provider will have to be educated on the legacy services which will be linked to the cloud service (APIs, data formats, security, etc.).	If the cloud provider does the maintenance for new features, the enterprise needs to understand costs and the expected responsiveness to complete requested updates.	Requires coordination between the enterprise development and operations teams with the cloud provider development and test teams.
Independent cloud service development provider	Should have proven experience and expertise on the specific cloud application service under consideration, thereby reducing development, testing and deployment costs.	Will require education and production knowledge of the legacy services and infrastructure which will be linked to the cloud service.	Will require coordination and a structured engagement with the enterprise implementation team and also the cloud provider implementation team in order to test and deploy the cloud service.	Ongoing updates, testing, and governance could be more complex and costly as well as take longer given the need to coordinate three parties as opposed to two.

Options for acquiring a new service	Skills	Startup considerations	Updates to services	Testing, deployment and support
Off the shelf purchase of a cloud application Software as a Service (SaaS)	Ensure that the application meets all the business requirements for the enterprise and all the open standards and API requirements of the enterprise.	Validate the level of effort required to adjust business processes accordingly and map the off the shelf data formats to the enterprise's data formats.	Determine who will be responsible for the modification, testing deployment, and maintenance activities.	Ensure the total cost of ownership of the off the shelf service offsets the costs for modification. If the time to production-ready deployment is significantly shorter, then the off the shelf option should be considered.

Selecting a methodology for implementing a cloud application can vary depending on whether the customer is a large organization or a SMB. Typically, the skills available within a SMB are targeted towards supporting existing applications and it may make more sense to consider contracting for resources from a cloud service provider or skilled independent development firm. Large organizations may have the flexibility to re-assign internal skills to a cloud project and accommodate the transition to cloud internally.

As evident from the above analysis, there are tradeoffs for each of the options listed. Organizations will have to avoid the risk of assuming that experienced virtualization skills will automatically transfer into cloud computing skills without additional staff development. An investment in cloud training and best practices should be decided during the early phases of the effort. Ultimately, the organization needs to take its own unique requirements into account when deciding what best meets their business needs. This effort can translate into leveraging several of the options in parallel, based on the needs of a particular cloud service.

Step 6: Develop Governance Policies and Service Agreements

Cloud Service Agreements (CSAs) should be evaluated in conjunction with specific needs, expectations, governance processes, and other cultural considerations. Data residency requirements, such as the EU's Global Data Protection Requirements (GDPR) also contribute to the need for clarity in CSAs. Service agreements vary greatly based on the deployment models involved (Public, Private, Hybrid), the service models they support (SaaS, PaaS, IaaS), and the specific cloud service. The CSCC has published two guides [6] [7] that help cloud customers evaluate cloud service agreements, leveraging a prescriptive ten step roadmap.

The ISO/IEC 19086-1:2016 standard [13] published in 2016 *"Seeks to establish a set of common cloud SLA building blocks (concepts, terms, definitions, contexts) that can be used to create cloud Service Level Agreements (SLAs)"*. 19086-1 also clarifies the differences between a CSA and an SLA, *"A cloud service level agreement (SLA) is a part of the cloud service agreement and details can vary for different cloud*

service categories, cloud services and different cloud deployment models." Cloud service agreements are generally intended to protect cloud providers from litigation, rather than assure a high level of service for customers. Public cloud service agreements are usually non-negotiable, making it even more critical to read and understand them in detail. The CSCC whitepaper, *Public Cloud Service Agreements: What to Expect and What to Negotiate Version 2.0* [7] offers guidance on options and rationale.

In addition, cloud service agreements are often cascading, leading to more challenges regarding the accountability and governance of the end-to-end cloud solution. The relationship between multiple CSAs is increasingly important as more organizations adopt hybrid cloud computing. The line of responsibility can vary across components of a complex hybrid architecture. The CSCC whitepaper, *Practical Guide to Hybrid Cloud Computing* [9] offers guidance on both governance and CSA content. A cloud service provider may use peer service providers, where the service agreements offered by those peer providers may be of significant interest to an overall understanding of responsibility for escalation or auditability.

The cloud service customer has final responsibility for performing due diligence, understanding their cloud service agreements and potential impact to their business. It is also essential that the cloud service customer has the necessary internal strategy, governance, and processes in place to use these services wisely and well. For example, to embark on adoption of cloud services without a strategy for the adoption of open standards or a policy on data residency could put an organization at risk. Particular attention should be given to organizational development and operations culture as cloud services provide for flexibility and need to be governed and managed without compromising the agility expected from cloud services.

In general, cloud service agreements can be decomposed into three major artifacts: "Customer Agreement," "Acceptable Use Policy" (or Terms & Conditions), and "Service Level Agreement," although the boundaries between these artifacts are vague and can vary from provider to provider. Bear in mind that these three artifacts may change at different times, independently from each other. Evaluation of cloud service agreements should include the following considerations.

- *Policies*. What policies and processes does your organization have that constrain cloud service decisions? These might include who signs off on subscriptions, how subscription payments are approved, enterprise architecture guidance, security or data access. Consider also contractual obligations the cloud services customer has with their customers or suppliers that might be affected.
- *Culture*. Are there cultural considerations, where Service Agreements can potentially mitigate concerns?
- *Governance*. Good governance requires transparency and accountability that leads to appropriate decisions that foster trust and assurance. What you need governed is, of course, a key consideration.
- *Objectives*. While developing and/or evaluating cloud service agreements, overall objectives and expectations will drive many of the discussions and approaches.
- *Metrics/Measures*. Cloud service agreements, especially aspects of Service Level Agreements, require consistent measurement. Measures and metrics will be used to validate service levels

and determine when remediation needs to be applied or resources dynamically allocated to assure service level objectives are met.

- *Terms and Conditions/Acceptable Use Policies*. Cloud service agreements may have specific terms, conditions, and use policies that need to be considered. This includes, but is not limited to: exclusions, limitations, usage and disclaimers.
- *Service Level Agreements*. A document stating the technical performance promises made by the cloud service provider, remedies for performance failures, and how disputes are to be discovered and handled.
- *Remediation and Compensation*. When fault and failures occur, they define what compensation is offered and the responsibilities of the parties involved. They also should clarify what disaster recovery/business continuity services are available in the event of service interruption.

A cloud service agreement does not absolve the cloud service customer of all responsibilities. Ongoing vigilance is required to ensure that service users continue to receive the expected level of service and compensation when service levels are not maintained. Customers should maintain a continuous level of responsibility by receiving direct feedback on the service level objectives of the service and be aware of any additional features (such as service monitoring and metrics collection) which may be needed.

Step 7: Assess and Resolve Security, Privacy and Data Residency Issues

Security and privacy are two of the issues that concern would-be cloud adopters the most. Depending on the domain in which they work (various industries, government, education, etc.), these concerns may rank just above or below those about availability and performance. More recently, data residency has been added to the mix. Practical guidance in this area should be used to manage risks while avoiding overreaction and paralysis, which often result from the concerns.

Understanding the Concerns

Security and privacy concerns were raised as soon as the cloud computing model appeared. In the early 2000s, organizations were already struggling to maintain adequate security in the presence of increasingly effective malware and other security threats. The general model of security viewed the enterprise as a fortress, with ramparts of firewalls and virus scanners isolating the inside from the outside. The cloud computing model means that some enterprise resources are outside the "fortress," leading many CISOs to believe that a cloud service could not be secure. At the same time, privacy rules were being tightened worldwide, and it seemed improbable that Personal Identification Information (PII) could be kept under appropriate control outside of the enterprise.

In more recent years, the risk of violating laws and regulations when moving data (including, but not limited to personal data) across borders has been exposed through well-publicized cases, adding a third related but distinct concern. In the CSCC whitepaper, *Data Residency Challenges [14]*, we define data residency as *"the set of issues and practices related to the location of data and metadata, the movement of (meta)data across geographies and jurisdictions, and the protection of that (meta)data against unintended access and other location-related risks."*

Assessing the Risks

Security, privacy and data residency are risk management issues, and should be treated using the same formal approaches: evaluate the probability and the impact of the potential threats, prioritize the risks accordingly, design and implement mitigation measures, test them, and keep monitoring the situation.

In *Security for Cloud Computing: 10 Steps to Ensure Success* [15], the CSCC takes a very broad view of the risks involved (see sidebar).

Adopting a cloud solution does not imply that the provider is solely responsible for security, privacy, and data residency issues. The customer and provider jointly share responsibilities. For example, in the case of an Infrastructure-as-a-Service (IaaS), the provider is responsible for the compute, storage, and network while the customer is responsible for hardening the operating system and middleware layers, implementing appropriate identity and access management, managing privileged access, encrypting data at rest and in motion, configuring the network and firewall to reduce the possibility of a breach, and so on.

Cloud Security Risks

- Loss of governance
- Compliance and legal risk
- Responsibility ambiguity
- Isolation failure
- Data protection
- *Insecure or incomplete data* deletion
- Handling of security incidents
- Service unavailability
- Management interface vulnerability
- Vendor lock-in

The following five considerations allow the risks to be discussed rationally:

- Many of the security, privacy and data residency concerns raised by cloud computing have existed since the first forms of IT outsourcing were introduced. These challenges should be seen as variants on previously existing issues, not totally new ones.
- The levels of security and privacy that are achieved in-house are often no higher than are achieved by cloud services. Cloud service providers typically have many more resources to assign to security design and monitoring than a single customer does, and providers have a strong business case for good security since a breach could undermine their entire business. Data residency is different, as this is clearly an issue that is exacerbated using cloud solutions.
- Inside threats are much higher than many would guess. Insiders were responsible for 39% of all data breaches in 2015 [16]. These breaches can be a result of malicious intent, accidental, or both. Insider breaches are often under reported and much harder to detect. Enterprises need to mitigate such risks with targeted programs to address deficiencies in training, communication, monitoring, policies, process, etc.
- Once a customer's information is in a cloud service, an attacker may have more difficulty finding it than if it is held on-premise. Therefore, a cloud solution can be *more secure* than an in-house system.

- Cloud customers must take responsibility for their use of cloud services, not abandon the responsibility to the providers. This includes understanding which data resides in the cloud service, where it may be residing, what its level of confidentiality is, how sensitive it is, whether it is encrypted, who has access to it, and so on.

Ten Sub-Steps for Security

The CSCC whitepaper, *Security for Cloud Computing: 10 Steps to Ensure Success*, provides 10 specific steps (within this guide, they become sub-steps of Step 7) to manage cloud computing security.

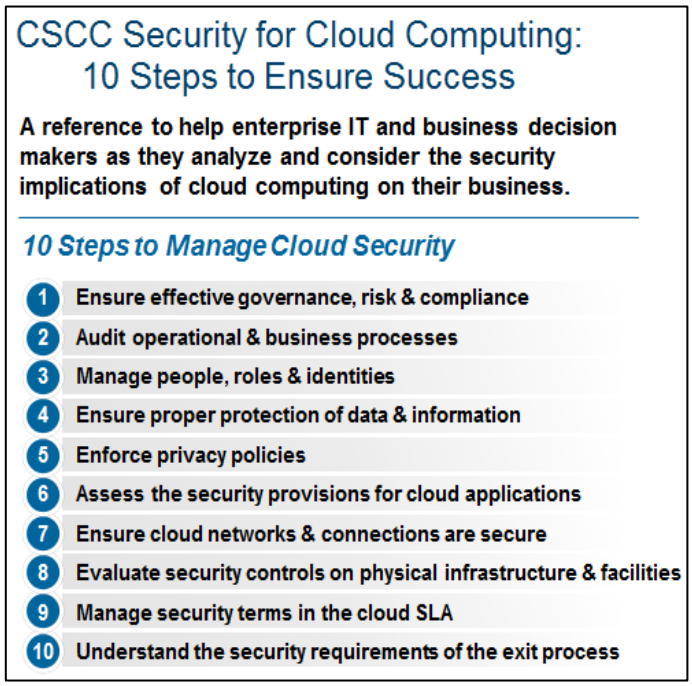

- Step 1, customers must understand the specific laws and regulations (data retention, privacy, disclosure requirements, residency, etc.) that apply to their business.
- Step 2 provides guidance to obtain professional security audits of the cloud service, and to monitor usage for suspicious activity. Audits should be consistent with general security standards such as ISO 27001/27002, and with cloud-specific standards such as ISO 27017 and 27018.
- Step 3 ensures proper user identification, strong authentication, and role-based access control to resources, possibly using federated identity management and single sign-on.
- Step 4 is about assigning a security classification to all data (without forgetting proprietary application code and system images, which should also be protected against theft and tampering).
- Step 5 relates specifically to the acquisition, storage and use of PII, including limiting access to it, storing it securely, specifying who (the cloud service provider or the customer) is responsible for what, and for monitoring compliance.
- Step 6 consists of understanding what security responsibilities the customer has, which differs according to the choice of deployment model (IaaS, PaaS, or SaaS).
- Step 7 ensures that the provider's internal network, as well as the connections between the customer and the cloud services are protected and monitored against external threats.
- Step 8 concerns the physical security of the computer center and building, protection against accidents and the environments (fires, earthquakes, flooding, etc.), screening of provider personnel, disposition of removable media, and so on.
- Step 9 is about Cloud Service Agreements, and is in fact so crucial that, based on work that the CSCC did on this specific topic, it deserves its own subsection below.

- Step 10 is about what happens to customer data during and after the termination of the use of a cloud service, including the complete removal of customer data from all tiers of storage, including any cached or backup copies, by the provider.

Any of these steps can lead to the decision that the project is not feasible. Just like with any good safety policy, the people who assess the situation should have the right to declare that they have found a showstopper and that the project must be halted until the problem is fixed.

Five Sub-Steps for Data Residency Management

In the CSCC whitepaper, *Data Residency Challenges* [14], we offer five steps to assess and manage data residency risks:

1. Establish a governance structure – typically a team with representation from business lines, IT security, legal or compliance, and representatives from the organization's geographic areas.
2. Ensure proper metadata management by having a complete enterprise "data landscape" or information model in one place.
3. Define all the policies and rules on sensitive data elements – what can be located where, what needs to be anonymized or encrypted, etc.
4. Establish reports to monitor the application of policies; measure how much data resides in which country or jurisdiction; and identify deviations.
5. If possible, implement tools to track the provenance and pedigree of information – and how it moves across boundaries as it gets processed and transformed.

Security in the Cloud Service Agreements

Since a cloud service customer always transfers *some* responsibility to the provider, it is important to understand what the service agreements say about the relative roles and responsibilities of the parties.

The CSCC *Practical Guide to Cloud Service Agreements* [6] calls out several issues with the current state of service agreements:

- Privacy and security considerations appear in different documents, with inconsistent titles and language.
- Most agreements impose stringent security obligations on the customer to protect the cloud provider – who decides unilaterally that a security violation occurred – but rarely any similar obligations or penalties regarding the harm that the provider might inflict on the customer.
- Privacy terms usually protect about the customer representatives' contact information, but not the customer's own users, who may be millions of end users.
- Escalation mechanisms are not specified or do not include response time commitments.

In the CSCC whitepaper, *Cloud Service Agreements: What to Expect and What to Negotiate* [7], customers are advised to examine these documents, request clarification, and negotiate what can be negotiated, possibly accepting to pay more for a higher tier of service with more acceptable terms.

Data residency may be completely absent from CSAs. In other cases, the vendor may assure the client that their servers are in-country, therefore there is no issue, but does not commit to maintaining this situation. Further, many support and operations staff may not be in-country who may access / download logs which may contain sensitive or private data. Therefore, make sure to look beyond the infrastructure and software for data residency requirements. Finally, the largest cloud vendors may offer the possibility to specify where the customer's data may (or may not) be stored, sometimes at an extra cost. These terms, or the absence thereof, must be studied to ensure compliance with business policies as well as applicable laws.

Step 8: Integrate with Existing Enterprise Systems

Hybrid cloud architectures are a common implementation pattern today, particularly for large organizations where there is a significant investment in existing applications and systems. The need for compliance with government legislation and/or industry standards is another common scenario supporting hybrid architectures. The term "two-speed IT" is widely used to describe the difference in governance and operational policy between the legacy and cloud native components of the hybrid cloud implementation. For hybrid cloud the adoption of cloud services typically involves integration of the cloud services with existing applications and systems. Integration may be bidirectional and may involve configuration changes or technical changes to the existing applications and systems and/or the creation of new integration components. In some cases, the cloud integration is specifically for implementation of DevOps tooling to reduce the cost and increase the speed of modifying or enhancing on-premises applications.

Integration involves a number of different components, both within the organization and within the cloud service provider. The components include:

- *Data*, where applications and services share common data, or synchronization of some kind is required between data in-house and data in a cloud service.
- *Process integration* between applications/services, where one application or service invokes operations provided by another as part of some workflow.
- *Management capabilities*, which include the monitoring of cloud services and the control of cloud services. These include security capabilities such as Identity and Access Management.
- *Business capabilities* including usage reporting, invoicing and payments.

Further, organizations may adopt cloud services to extend their business processes so that they become more accessible to others in their business ecosystem. In this case, some adaptation of existing applications and services may be necessary. When extending business processes organizations need to spend sufficient time reviewing and adapting their internal policies and processes for managing applications so these can work at the speed of cloud.

There are several ways of establishing links between cloud services and existing applications and systems. If the organization has already established a direction of adopting open standards for data formats or for communication protocols and APIs, then the integration of cloud services should build on

what has been already implemented. This increases the opportunity for achieving interoperability between cloud services and the enterprise applications and systems.

If the organization has not implemented a discipline of adopting open standards, then the new cloud services can be used to set the baseline for the necessary integration components. A clear plan for adopting open standards will help enable interoperability and portability for cloud services and simplify the process of integrating new cloud services, independent of where or how the new cloud service is acquired.

The use of open standards for data formats and for APIs and protocols can assist in the process of integration of cloud services. Ideally, the cloud service itself should utilize open standards – but the existing in-house applications and systems may need some adaptation and updating to conform to those standards. Any work done to update and adapt the existing applications and systems to use open standards should pay for itself in the long run, especially where there is an ongoing commitment to migrate more and more functionality to cloud services over time. One approach to adapting existing applications and systems is to create an adapter component which can translate between the existing applications and systems and the standard data formats, APIs and protocols used to communicate to and from cloud services.

The costliest method of integrating new cloud services into the organization will be to initiate a project to develop custom code for each new cloud service as it is implemented. If this process is followed there are many downsides:

- Increased development costs and time required to integrate the use of the new cloud service
- Increased maintenance costs to add new capabilities
- Reduced flexibility to integrate new services using the same legacy service
- Increased costs and time to move a cloud service to a new cloud provider
- Higher costs to establish a disaster recovery plan

Security integration is usually a key element of the use of cloud services. One of the common requirements is for the integration of the organization's Identity and Access Management (IdAM) system with the cloud service – it is undesirable for the organization to have to administer a separate IdAM system for each cloud service – the best arrangement is for the cloud service to delegate authentication capabilities to the organization's IdAM system. This will require both the cloud service and the IdAM system to support one of the common standards for this capability such as OAuth 2.0 or SAML 2.0.

Step 9: Develop a Proof-of-Concept before Moving to Production

Once the business case for cloud computing is complete and both business drivers and projected ROI are established, it is important to obtain a final ‘go’ or ‘no-go’ decision from senior management. The senior management buy-in should include at a minimum, a review of the proposal, projected costs, timeline, risks, resulting benefits and a rollback plan. If there is an agreement, the next step is to assemble a proof-of-concept (POC) team, comprised of the following resources:

- *Information Technology*. This team should be composed of architects, systems administrators, database administrators, senior developers, and customer support (help desk) resources.
- *Functional Representation*. This team includes at a minimum designated individual(s) within the organization that will manage the continued alignment of the cloud computing solution with business user and key stakeholder expectations during the POC.

Assuming that the POC is successful and meets or exceeds expectations, design, development and implementation activities for the production instance of the cloud service can be fully engaged. Make sure new service meet following success criteria:

- *Costs.* Hosting, maintenance and sustaining costs are lower than current costs.
- *KPI (Key performance indicators).* System uptime, data processing and delivering results, etc. are equal or better than the current environment.
- *User satisfaction.* This is a very important success criterion - If users are satisfied, new implementations will be successful.
- *Infrastructure prep and delivery.* This KPI measurement includes the time to prepare the infrastructure for development, QA, staging, and production.
- *Problem resolution.* Average time between reporting a bug and resolving the issue.
- Security. Systems are secure and required access control measures are in place and tested. Administrators have required access when needed to do their job.
- Monitoring and reporting. New environments have required monitoring and reporting in-place.
- *DR (Disaster Recovery).* Provisions are in-place and verified to start a DR environment in the event the primary hosting environment fails.

Implementing a new cloud service requires the same discipline as implementing a non-cloud service. The implementation team needs to ensure the following success criteria activities are completed:

- Verify the cloud service delivers required functionality in a test environment.
- Verify that all processes work as intended.
- Verify data recovery activities, formatting, migration, and ETL (extract, transform and load) capabilities function.
- Verify integration with management and monitoring systems.
- Ensure that the help desk can address questions and problems quickly.
- Develop a back out plan should there be an unexpected problem in the early stages of production so as not to impact users.
- Verify identity and access management in the new cloud environment.
- Verify administrative staff have required access to perform admin activities within the new cloud service.
- Compare the time it takes for data and system recovery activities to complete in the new cloud service vs. its current state.

The POC can be implemented either in-house or directly on a public cloud service. While a public cloud service provides benefits like quick provisioning and scalability, it is important that organizations perform testing using a representative data set rather than production data to ensure data security. It is also important to recognize that there may be differences between the POC and target cloud environments that will have to be addressed upon migration to the production environment.

Once all the testing has been completed and all of the stakeholders have signed off that their area is working properly, the new cloud service can be put into full production when the following activities are completed:

- Business contracts agreed to and in place
- SLAs agreed to and in place
- Customer support (help desk) educated/trained and in place. Help desk can either be within the organization or with the cloud service provider
- Post implementation management or operations plan completed

Step 10: Manage the Cloud Environment

The responsibility within the customer organization for the successful operation of cloud services is shared by the CIO, who has overall responsibility, and the manager of customer support who manages the day to day operational challenges. Any problems which cannot be resolved must be escalated to the CIO to ensure that all avenues to resolve the problems have been executed. If the problems cannot be resolved, then the options written into the SLA can be invoked.

The technical and customer support requirements vary based on the service model, deployment model, and hosting option selected:

- For a Private (On-premise) cloud, the management of the cloud will be consistent with the management of the existing services within the organization.
- For Private (Outsourced) and Public clouds, the responsibility for management of the cloud service(s) will be laid out in the Cloud Service Agreement. The Cloud Service Agreement will establish processes for identifying a problem, indicate who is responsible and depending on the impact of the problem, what resources are brought to bear to resolve the problem.

A disaster recovery process must be defined and implemented to protect the organization and its digital assets. Responsibility for this process can vary depending on the nature of the service – for example, for a SaaS service it will often be the responsibility of the cloud service provider, while for IaaS services, it may be the responsibility of the cloud service customer. The disaster recovery process must be verified prior to deploying and enabling production service. When required, the customer support manager within the organization is responsible for initiating the disaster recovery process. There must be trained individuals in both the cloud service provider and the cloud customer side who can ensure that the recovery process is completed properly and can verify no data loss has occurred.

There must be a documented service agreement and service delivery process between the cloud customer and the cloud service provider which must cover the process for incident reporting and responses to the individual reported the incident. Each problem should have a severity assigned to it to reflect the impact and the resulting urgency for resolution. If an individual within the organization cannot get an incident resolved through the cloud service provider, the issue should be escalated to the customer support manager. The customer support manager will assess the severity of the incident and take the appropriate action.

In addition to technical and customer support, management of the cloud environment entails handling of change requests made by the business to meet its changing business requirements. An effective change management process needs to be implemented to ensure that business requirements are gathered, validated, developed, tested and deployed. Further, in the likely scenario of having multiple cloud vendors, the customer must ensure that vendor management processes are clearly defined to obtain optimum results.[4]

Summary of Keys to Success

Table 7 summarizes a few of the critical keys to success for any organization embarking on a cloud computing journey.

Table 7: Summary of Keys to Success

Key to Success	Summary
Establish executive support	• Senior management team must understand and take responsibility for the successful adoption of cloud services. • Pressures will come from various key players in any cloud decision: lines of business (LOB), IT, finance, legal, procurement, security and the user community. • The IT community is most concerned about global access and impact on networks, security, user performance, etc. The key to their support is a globally-aware architectural plan for cloud implementation. • LOBs are likely to be concerned about the business value offered and the time to market for the solution. • Finance and procurement are most concerned about saving money. The key to executive support is a well-thought ROI rationale and calculation. • Users are often most concerned with time to market and scaling the environment in lock-step with changes to the business. The key to executive support from this group is to demonstrate increased agility and higher elasticity from the cloud.

[4] For services offered in a public cloud environment there may not be the ability to request customization of a service as the offering is used by more than one consumer and only the data is segregated.

Key to Success	Summary
Address organizational change management	● Management must understand and address the pressures and cultural transformation introduced by cloud computing on the organization. ● Cloud computing will introduce change to the normal IT development and deployment processes, breaking down many organizational barriers and norms. ● At the heart of change is fear of loss - primarily, loss of control. The change must have a well-managed, well-planned process for mitigating fear of loss and understanding the significance of a shared responsibility model with an external cloud service provider (if present). Embracing change is critical to success.
Establish commitment	● The organization must be fully committed to developing and executing a strategic plan for cloud computing within the enterprise. ● Adoption of cloud computing should be led by senior management including the CEO and CFO with the CIO, CTO and CISO playing a role of key enablers.
Carefully evaluate cloud service agreements to ensure critical business needs are adequately addressed	● Buy service, not servers. Look for complete managed services where you rely on the cloud provider to integrate all parts into a complete solution. ● A properly negotiated service agreement will ensure there is a partnership between the customer and provider for the overall success of the service.
Address federated governance	● Cloud services are by nature distributed and federated, but most command-and-control systems for managing IT are hierarchical. ● To succeed, some degree of distributed control and federated governance is necessary to match the model of cloud service delivery. ● Before making a decision on a cloud service provider, it is important to understand how the cloud service will be managed and what processes need to be integrated into the existing IT environment.
Handle security and privacy	● At the heart of security is trust. Often cloud providers have a deeper awareness of what is required to provide good security than the customers they serve. However, the customer and cloud service provider must work together to establish a trust relationship and shared responsibility model to establish an optimal security and privacy environment. ● Document the level of security required to properly protect the service and data and let the provider confirm how the requirements will be met. Objectively measure the provider's true security capabilities. ● It is critical that sensitive information does not find its way into the wrong hands. The customer and provider are jointly responsible for ensuring that the data has appropriate protection, consistent with the statements in the CSA.

Key to Success	Summary
Comply with legal and regulatory requirements	• An organization must be aware of and plan for adherence to legal and regulatory requirements, including those related to security, privacy and accessibility. Failure to comply can derail the cloud computing effort and result in costly lawsuits.
Define metrics and a process for measuring impact	• There is truth in the adage that "People do what you inspect, not what you expect." • Track your cloud adoption rate (e.g., ratio of cloud-enabled vs legacy applications) to your strategic goals to ensure you are achieving your expected benefits from cloud adoption. • Create operational metrics and monitoring processes which define steady state success. Define how the metrics will be measured. • Use metrics to assess cost savings and revenue enhancement, and to validate SLA compliance, including elasticity, availability, performance globalization, etc. • By measuring results, there will be a baseline from which to make better decisions for future cloud services with the goal of continual ROI improvement.

Cloud computing offers a value proposition that is different from traditional enterprise IT environments. With proper focus on the key success factors, the promise of cloud computing can be realized.

Works Cited

[1] RightScale 2017 State of the Cloud Report (2017). https://assets.rightscale.com/uploads/pdfs/RightScale-2017-State-of-the-Cloud-Report.pdf

[2] ISO/IEC 17788:2014 Information technology — Cloud computing — Overview and vocabulary. http://standards.iso.org/ittf/PubliclyAvailableStandards/c060544_ISO_IEC_17788_2014.zip

[3] ISO/IEC 19941 Information technology — Cloud computing — Interoperability and Portability. https://www.iso.org/standard/66639.html

[4] Open Container Initiative. https://www.opencontainers.org/

[5] Cloud Standards Customer Council (2013). *Migrating Applications to Public Cloud Services: Roadmap to Success.* http://www.cloud-council.org/deliverables/migrating-applications-to-public-cloud-services-roadmap-for-success.htm

[6] Cloud Standards Customer Council (2015). *Practical Guide to Cloud Service Agreements*. http://www.cloud-council.org/deliverables/practical-guide-to-cloud-service-agreements.htm

[7] Cloud Standards Customer Council (2016). *Public Cloud Service Agreements: What to Expect & What to Negotiate.* http://www.cloud-council.org/deliverables/public-cloud-service-agreements-what-to-expect-and-what-to-negotiate.htm

[8] NIST Cloud Computing Synopsis and Recommendations document, Special Publication 800-146. http://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-146.pdf

[9] Cloud Standards Customer Council (2016). *Practical Guide to Hybrid Cloud Computing*. http://www.cloud-council.org/deliverables/practical-guide-to-hybrid-cloud-computing.htm

[10] National Institute for Standards and Technology (2011): Special Publication 500-292 *NIST Cloud Computing Reference Architecture.* http://ws680.nist.gov/publication/get_pdf.cfm?pub_id=909505

[11] Cloud Standards Customer Council (2015). *Practical Guide to Platform as a Service*. http://www.cloud-council.org/deliverables/practical-guide-to-platform-as-a-service.htm

[12] Cloud Standards Customer Council (2013). *Convergence of Social, Mobile and Cloud: 7 Steps to Ensure Success.* http://www.cloud-council.org/deliverables/convergence-of-social-mobile-and-cloud-7-steps-to-ensure-success.htm

[13] ISO/IEC 19086-1:2016 Information technology — Cloud computing — Service level agreement (SLA) framework — Part 1: Overview and concepts. https://www.iso.org/standard/67545.html

[14] Cloud Standards Customer Council (2017). *Data Residency Challenges*. http://www.cloud-council.org/deliverables/data-residency-challenges.htm

[15] Cloud Standards Customer Council (2015). Security for Cloud Computing: 10 Steps to Ensure Success. http://www.cloud-council.org/deliverables/security-for-cloud-computing-10-steps-to-ensure-success.htm

[16] Blankenship, Joseph (2016). Hunting Insider Threats. Forrester Research report number 134865.

Additional References

National Institute for Standards and Technology (2011): *NIST Cloud Computing Standards Roadmap.* http://www.nist.gov/itl/cloud/upload/NIST_SP-500-291_Version-2_2013_June18_FINAL.pdf

National Institute for Standards and Technology (2014): *NIST Cloud Computing Related Publications.* http://www.nist.gov/itl/cloud/publications.cfm

National Institute for Standards and Technology (2014): *NIST Cloud Computing Program.* http://www.nist.gov/itl/cloud/

