

Transforming Government

A cloud assurance program guide





Authors

Amanda Craig
Microsoft Trustworthy Computing

Min Hyun
Microsoft Trustworthy Computing

Contributors

Benedikt Abendroth
Gregg Brown
Yen-Ming Chen
Kaja Ciglic
Erin English
Brice Keown
Angela McKay
Paul Nicholas
Jim Pinter
Matt Rathbun
Shawn Veney
John Weigelt



Contents

Executive summary	4	
Introduction	6	
Prerequisite	Leverage existing IT assurance constructs and adapt for the cloud	7
Phase 1	Structure a cloud assurance program	9
Step 1	Identify roles and responsibilities of government, service provider, and third party stakeholders	9
Step 2	Establish governance processes for initial authorization and ongoing authorization	14
Step 3	Establish a continuous improvement loop to maintain program effectiveness and align with broader strategies	16
Phase 2	Develop a cloud assurance program	17
Step 1	Adapt existing data classification scheme for cloud services	17
Step 2	Define requirements baselines by leveraging global standards	19
Step 3	Establish a cloud-adapted risk assessment process	21
Step 4	Establish risk management processes to address unique risk scenarios	22
Phase 3	Implement the program	24
Step 1	Identify and categorize data and systems for cloud migration	24
Step 2	Implement requirements baselines, risk assessments, and approaches to addressing unique risk scenarios	25
Step 3	Select a cloud service delivery model	27
Step 4	Select a cloud deployment model	28
Conclusion	29	



Executive summary

Around the world, public sector cloud adoption is on the rise.¹ Governments are recognizing that cloud services offer enormous value and agility and can unlock vast potential for innovation, security, and resiliency. As a result, governments are moving beyond questions about *whether* to use cloud computing and are now focused on *how* to more efficiently, effectively, and securely integrate cloud services.

Microsoft has designed this cloud security guide to support governments as they develop and implement policies and programs to migrate their data and systems to cloud services. We know that governments have questions, such as: what roles should government and third party organizations have in assuring that a cloud service can be trusted? What is the process through which governments will gain assurance about the cloud service? How will the cloud assurance program leverage existing government information technology (IT) policies, as well as global cloud security best practices? And, once the cloud assurance program is in place, how will governments, ministries, and agencies implement it, enabling them to procure and integrate cloud services?

This document outlines a series of steps that are designed to help governments address these questions and to serve as a starting place as they build their cloud assurance programs. Including best practices gathered from our experience in partnering with governments, this guide will help governments build a robust but agile program, enabling them to stay on pace with the rate of technological change and to take advantage of the newest security and productivity features.

¹ IDC predicts that public sector spending on cloud services will grow to \$128 billion by 2018, more than doubling 2014 spending. *IDC Forecasts Public IT Cloud Services Spending Will Reach \$127bn in 2018 as the Market Enters a Critical Innovation Stage* (2014), <http://www.idc.com/getdoc.jsp?containerId=prUS25219014>.



The guide is organized into phases that will help governments structure, develop, and implement a cloud assurance program. The phases are as follows:

Phase 1 Structure a cloud assurance program	<ol style="list-style-type: none">1. Identify roles and responsibilities of government, service provider, and third party stakeholders2. Establish governance processes for initial authorization and ongoing authorization3. Establish a continuous improvement loop to maintain program effectiveness and align with broader strategies
Phase 2 Develop a cloud assurance program	<ol style="list-style-type: none">1. Adapt existing data classification scheme for cloud services2. Define requirements baselines by leveraging global standards3. Establish a cloud-adapted risk assessment process4. Establish risk management process to address unique risk scenarios
Phase 3 Implement a cloud assurance program	<ol style="list-style-type: none">1. Identify and categorize data and systems for cloud migration2. Implement requirements baselines, risk assessments, and approaches to addressing unique risk scenarios3. Select a cloud service delivery model4. Select a cloud deployment model

Establishing a cloud assurance program is a foundational investment that enables governments to leverage secure cloud solutions to deliver and extend citizen services. Cloud assurance programs pay significant dividends, increasing understanding of the cloud’s robust capabilities and security features and growing technical competence within the government community. Cloud assurance programs position governments to readily empower a mobile workforce, accelerate digital services, and reap economic benefits with the confidence that security requirements are being met as data is transmitted and stored in the cloud.



Introduction

Cloud computing has the potential to transform the government workforce and citizen services, improve productivity and efficiency, and catapult economic growth and innovation.² But, how governments plan to deploy cloud services can either enable or impede such transformation. In particular, establishing processes to determine whether they should trust cloud services and how they can achieve cloud assurance can be especially challenging because of governments' unique set of offerings and variety of data. When discussing cloud computing, a set of key attributes provides users with a level of certainty that their data is sufficiently protected as it traverses across and rests in a cloud system:

- *Cloud assurance* can be defined as ensuring security, compliance, privacy, and trust in cloud services so that the services are functioning as intended. Simply put, customers want cloud service providers to do the right thing—and to prove it.
- *Cloud security* is comprised of the architectural and operational foundation and processes in place to safeguard a physical and virtual system as well as the data and functions that it hosts.
- *Cloud compliance* entails the security measures in place so that a system meets specific requirements defined in standards, regulations, and policies.
- *Cloud privacy* and *control* ensure that customers own their data, which can only be accessed, used, deleted, and shared as determined by the customer.
- *Cloud trust* is having certainty in the security, compliance, transparency, and privacy of the cloud system housing your data.

² *Transforming Government: A cloud policy framework for innovation, security, and resilience* (2015), <http://aka.ms/cloudsecurityprinciples>.



This guide focuses on the phases and steps needed to structure, develop, and implement a cloud assurance program, and it also integrates considerations related to cloud security, compliance, privacy and control, trust, and transparency. Moreover, the steps outlined below have been aggregated based on leading practices adopted by public and private sectors around the world, so governments that utilize this guide can establish a program that builds from risk management principles and ensures real security outcomes. With this in mind, the guide will continue to evolve in tandem with changing cloud requirements and maturing assurance programs. In addition, it incorporates learnings that predate the widespread use of cloud services, described immediately below as a prerequisite to building a cloud assurance program and integrating cloud services.

This guide is designed to complement a white paper, *Transforming Government: Cloud policy framework for innovation, security, and resilience*, which describes six policy principles relevant for transitioning to cloud services. Building on those principles, this guide describes how governments can have assurance in their cloud services and articulates the steps that governments should take to create a cloud assurance program. It also acknowledges that many governments have pre-existing digital security strategies that a cloud security program should build from and integrate.

Prerequisite - Leverage existing IT assurance constructs and adapt for the cloud

Transitioning to cloud services works best when governments:

- Use their existing IT programs, processes, and policies as a starting point
- Identify potential gaps when transferring to the cloud environment
- Update existing guidance and processes or create new ones, adapted for the cloud, as needed

Information assurance has been important for technology users since the traditional boxed product and on-premises systems era. Many governments today have established IT security programs that help them to achieve assurance by assessing risk-based policies, processes (e.g. data classification, risk assessments, and IT lifecycle management), and governance models. For example, [Common Criteria for Information Technology Security Evaluation](#), an authoritative international standard for computer products and systems certification, has been widely adopted by governments around the world.

However, computing has evolved, enabling content delivery to progress from a limited, one-to-one paradigm (i.e. traditional on-premises computing) into a multi-dimensional, many-to-many paradigm (i.e. multi-tenant cloud computing). While this shift in content delivery has already occurred, widely adopted methods to provide assurance that cloud services are operating and secured as intended are still emerging. As they benefit from these shifts and new paradigms that are inherent in cloud services, users of technology will work with technology providers to develop methods through which they can achieve assurance, and they will look to previous information assurance methods as guidance or as a starting place.

Many of the information security program practices already in place can be re-used and adapted for a cloud environment, whereas others (e.g. physical asset management) may need to be re-applied or deprecated. Governments benefit from adopting a holistic approach to IT risk management, which includes cloud computing as part of their overall IT portfolio. The challenge in adapting existing programs and practices to



establish a cloud assurance program is the need for speed. The slow pace of traditional programs—like Common Criteria, for which it could take years to obtain a certification—impede the government’s ability to adopt and benefit from the latest technologies. Because cloud services can be developed, deployed, and improved much faster than traditional IT products can be, cloud assurance programs must be calibrated to match the pace of technology while still meeting the established security bar.

Ensuring both security and speed will be particularly important as instances of hybrid environments (i.e. traditional on-premises systems co-operating with cloud-based systems) continue to emerge. For this common scenario, implementing an approach in which governments start with existing best practices and guidance and adjust where needed provides much-needed continuity, consistency, and efficiency. Moreover, in any scenario in which governments are integrating cloud services, continuity, consistency, and efficiency will be pivotal for programmatic success, and both adapting existing IT assurance constructs for the cloud and utilizing the steps outlined below will help governments work towards those goals.



Phase 1

Structure a cloud assurance program

Structuring a cloud assurance program is a critical first step, and great care should be taken to ensure that there is a clear understanding of program goals before development begins.

For instance, a sound assurance program should not focus solely on *security* outcomes, as achieving a particular security implementation at the expense of poor user experience or performance latency may not be the desired end state. Instead, a sound assurance program should carefully balance security with performance and innovation to support organizational mission and business objectives.

Once there is a clear understanding of the intended outcomes, governments can begin to establish processes in support of them.

To start building a cloud assurance program, as described above, governments should integrate their existing IT security strategies. As part of that process, they may choose to create new authorities or assign existing authorities with cloud-specific roles and responsibilities. Either way, to structure their cloud assurance programs, Microsoft recommends that governments take the following steps:

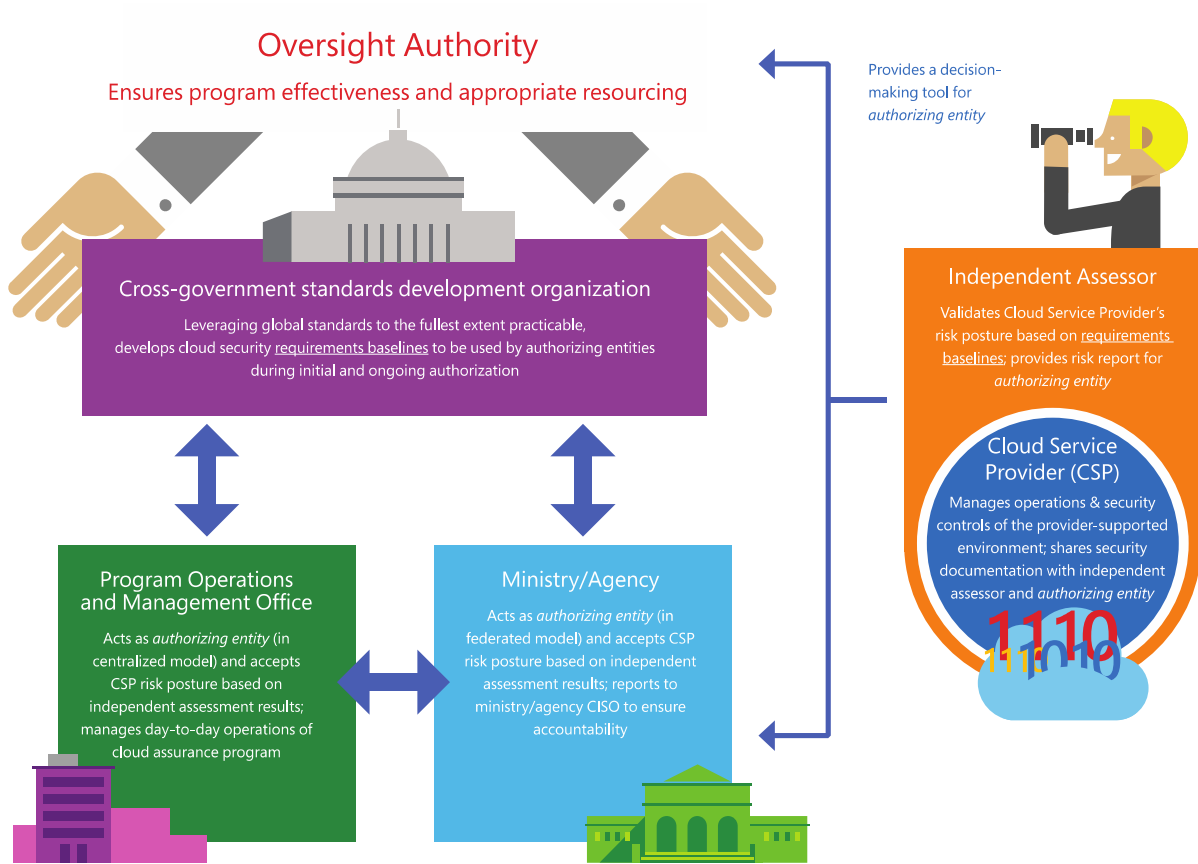
- 1) Identify roles and responsibilities of government, service provider, and third party stakeholders
- 2) Establish governance processes for initial authorization and ongoing authorization
- 3) Establish a continuous improvement loop to maintain program effectiveness and align with broader strategies

Step 1: Identify roles and responsibilities of government, service provider, and third party stakeholders

Governments should begin by identifying and defining government entities' and other stakeholders' roles and responsibilities. This step should also clearly define how different stakeholders interplay and how they will be held accountable in their roles.



Figure 1: Cloud Assurance Program Roles and Responsibilities



As a starting point, Microsoft recommends defining the following stakeholders and capabilities:

1) Oversight authority – ensures assurance program effectiveness and appropriate resourcing.



An oversight authority bears the ultimate responsibility for the success of the cloud assurance program. It is responsible for ensuring the program's delivery against its goals and objectives but is not involved in the daily operations and management of the program. Part of ensuring the program's success is securing sufficient resourcing from both a capability (i.e. expertise) and a capacity (i.e. budget) perspective. To enable continuous improvement, this entity should also be empowered to enforce programmatic and governance changes.

2) Program operations and management office – manages day-to-day operations; may act as Cloud Service Provider (CSP) authorizing entity.



This entity operates the program through daily management. In a "centralized" approach, which is described below, it may also issue authorizations (i.e. to a CSP) to demonstrate that a cloud service has an acceptable risk posture and is approved to provide services for government environments. Key responsibilities include:



- Operations
 - Authorize a CSP within a pre-determined time period if performing the risk assessment and within more expedited time frames if leveraging the risk report issued by the independent assessor. This includes establishing processes that support timely CSP authorizations to stay synchronized with the rapid pace of change in cloud technology and the addition of new services;
 - Establish an ongoing authorization approach that focuses on remediating security control failures and highest risks, addresses unique risk scenarios, and certifies new services in sync with the rapid pace of technological innovation;
 - Develop, in collaboration with stakeholders, capabilities for security trend analyses, change management, and incident response; and
 - Accredite independent assessors that will be recognized as the primary certifiers of a CSP's risk assessment and provide performance quality oversight.
- Policy and strategy development
 - Determine the applicability and re-use of other widely recognized industry and public sector-accepted certifications, self-attestations, and associated artifacts;
 - Consult with other ministries or agencies, independent assessors, standards bodies, and CSPs to ensure input into cloud security policies and assurance processes; and
 - Establish an industry working group with commercial cloud providers to advise public sector stakeholders on the governance of the cloud assurance program and to provide accountability for government and industry roles.
- Ministry/agency support
 - Provide assistance to ministries or agencies regarding compliance with requirements, guidelines, and standards; and
 - Support ministry or agency cloud computing risk assessment and authorization counterpart functions by supplying subject matter expertise.
- Metrics reporting to ensure accountability
 - Report annually to the oversight authority on: the status, efficiency, and effectiveness of the program during the preceding year (e.g. length of certification in months); progress made by ministry or agency cloud customers in migrating to cloud services; and progress made during the preceding year in advancing tooling to securely automate processes and reporting functions to increase efficiency and reduce costs.



3) Cross-government standards development organization (CGSDO) – uses global standards for cloud security requirements baselines.



The CGSDO leverages global standards to identify baseline security requirements. Oftentimes, it will tailor the requirements or guidance established by the International Standards Organization (ISO), European Network and Information Security Agency (ENISA), and National Institute of Standards and Technology (NIST), which are internationally recognized standards development bodies or centers of expertise whose work has received broad industry adoption. Key responsibilities of the CGSDO include:

- Issue and maintain requirements baselines, related policies, and guidance that: 1) are consistent with existing laws and globally recognized standards that are tailored for cloud computing services to the fullest extent practicable (see Phase 2 for further detail); 2) vary according to data and system classification level (e.g. high, moderate, low); and 3) map to the distinct cloud services delivery models (i.e. Infrastructure-as-a-Service (IaaS), Platform-as-a-Service (PaaS), Software-as-a-Service (SaaS), helping to streamline the assessment process and leverage dependencies across the service layers);
- Determine an approach that allows ministries or agencies the discretion to add requirements above the baselines, provided that they: 1) clearly articulate the risk analysis that justifies modification of the requirements baseline; 2) show that existing requirements do not provide the same or similar protection; 3) demonstrate that the security value justifies the costs that will be incurred by the CSP and the ministry/agency to meet the new requirement; and 4) whenever possible, accept alternative measures or compensating controls that demonstrate security equivalency;
- Adapt existing data classification scheme for cloud services as necessary and provide guidance for mapping the appropriate baseline considerations to varying classification levels of data; and
- Define a scheme for accrediting independent assessors and recognizing certifications and attestations achieved through other certifications.

4) Ministry/Agency – deploys cloud services; may act as CSP authorizing entity.



These organizations deploy cloud services and are responsible for assessing, implementing, and monitoring security controls that are the ministry's or agency's responsibility. In a "federated" approach, which is described below, ministries and agencies also issue authorizations (i.e. certify CSPs) to demonstrate that a cloud service has an acceptable risk posture and is approved to provide services for their environment. Ultimately, ministries or agencies are the risk owners because they are best positioned to understand and manage their own risks. To successfully fulfill these responsibilities, a ministry/agency should:

- Authorize a CSP within a pre-determined time period if performing the risk assessment and within more expedited time frames if leveraging the risk report issued by the independent assessor;
- Report authorization time frames to the program operations and management office;
- Interface with the program operations and management office, independent assessors, and CSPs to evaluate and provide validation of the assessment;



- Ensure sufficient and ongoing support by establishing a dedicated ministry or agency resource to serve as the primary authorizer of security assessment packages for initial authorization and continuous monitoring; and
- Report metrics to the ministry's or agency's chief information security officer (CISO) or chief information officer (CIO) (depending on organizational structure) to ensure accountability.

5) Independent assessor – *validates cloud service risk posture; provides a risk report for CSP authorizing entity.*



Independent third party assessors serve as the primary certifiers of cloud services. They make a risk-based determination regarding whether CSPs have fulfilled a government's requirements. Independent assessors are responsible for validating and attesting to the quality and compliance of CSP-provided security assessment materials and for developing a summary risk report. These risk reports should reflect an objective representation of the current risks identified by the CSP and the assessor, based on the results of continuous monitoring activities, and should be used as a decision-making tool for the program operations and management office or ministries and agencies in authorizing CSPs. Leveraging these risk reports prevents duplicative technical reviews of the assessment package and focuses attention on the risks and threats of highest concern.

6) Cloud service provider (CSP) – *provisions cloud services; develops security documentation; works with independent assessors as they test and validate controls and with authorizing entity for risk review.*



A CSP provisions cloud-based services (i.e. IaaS, PaaS, SaaS) and manages the operations and security controls of the provider-supported environment, which can vary based on the service model, but at minimum includes underlying infrastructure services. The entity determines the service boundary that will undergo certification and develops the associated documentation, which defines the system itself and the security measures in place. The CSP also works closely with the independent assessor to scope and test applicable controls during initial authorization and ongoing authorization. It also works with the ministry or agency customer consuming the cloud service to ensure understanding of shared controls implementation and responsibility. Lastly, the CSP must also ensure accreditation of any new or modified service or feature affecting the security posture of the service.



Step 2: Establish governance processes for initial authorization and ongoing authorization

Microsoft's experience has demonstrated the value of establishing two distinct processes for cloud assurance. The first process, initial authorization, is a comprehensive evaluation. It is an original assessment of a defined set of cloud services, including a comprehensive evaluation of security practices and controls. The second process, ongoing authorization, is a narrower, targeted review, including regular assessments of a core set of practices, controls, and/or performance indicators to maintain visibility of the CSP's risk posture as ministries/agencies deploy a cloud service. While the scope of assessment and testing for these distinct processes will vary, together, these processes should help government ministries and agencies have assurance in the security and resilience of cloud services as they begin and continue to deploy them.

Initial authorization

Initial authorization is a process for assessing the security posture of a CSP and the risks associated with a cloud system. Often, the process involves the assessment of security controls, which are technical, operational, or managerial measures implemented on a system to address a security risk. While reviewing a CSP's security controls implementation is important and necessary, documentation review of security controls alone does not provide a true indication of a system's or CSP's security posture. Instead, taking a more holistic approach and focusing on security capabilities enables a CSP to demonstrate a range of innovative techniques, which in many instances may *exceed* a government's security objectives.

Initial authorization should be defined by a set of activities that provide a more holistic view of the operational practices in place to secure a cloud system, including:

- Reciprocity of widely recognized security certifications and attestations that have been verified by an independent assessor;
- Validation of new security controls by an independent assessor; and
- Interactive review sessions with the CSP, independent assessor, and government authorizer to address threats and significant weaknesses.

CSPs will demonstrate a wide range of security and certification experience, expertise, and scale. Therefore, cloud assurance programs should also be flexible enough to allow for variable review, monitoring, and reporting based on a CSP's demonstrated security posture. This can be validated by other widely recognized third party certifications or consistently strong performance.

Requirements baselines are most effectively developed centrally, by a cross-government standards development organization (CGSDO) that leverages global standards, and should be implemented based on data classification level (i.e., higher impact data or systems may require higher requirements baselines). Depending on whether a government takes a centralized or federated approach to its cloud assurance program, initial and ongoing authorization may be a function of a program operations and management office or a ministry or agency, respectively. Our experience has shown value in the "do once, use many times approach" reflected in a centralized or federated model. (Additional details on centralized, federated, and decentralized approaches provided in Phase 3.)



Ongoing authorization

After granting initial authorization, governments must be confident that CSPs that achieve initial authorization are sufficiently maintaining their security status as they stay on pace with rapidly evolving cloud technology. In today's fast-paced digital era, information exchange occurs on-demand and among many interconnected devices, and securing hyper-scale volumes of data requires governments to engage in near-real time risk management of significant threats and vulnerabilities through ongoing monitoring of system changes and security measures.

Ongoing authorization enables this much-needed, agile operational visibility and improves upon the static security practices of traditional computing, which has often been developed, deployed, and improved more slowly and so resulted in a less dynamic approach to risk management. In establishing ongoing authorization for cloud services, governments should take into account the importance of speed and agility to mitigate new and evolving threats and realize the security benefits of new technologies and features. They must also be able to assess new services on pace with upgrade cycles, which can occur as frequently as monthly or quarterly.

Skewed risk rating from over-emphasis on low impact vulnerabilities

Risk management strategies require organizations to prioritize risks and apply more resources toward mitigating higher risks. Generally, higher risks are reflected as "high" findings, while lower risks are reflected as "low" findings. However, in a commercial cloud environment, "low" findings can reach hundreds or thousands when accounting for vulnerabilities by host (due to the large number of hosts managed by CSPs). The effort of reviewing a low finding on multiple hosts to determine validity is a highly manual process that is incompatible with how CSPs internally track risks. In addition, government customers must spend significant resources reviewing reports fraught with low findings, which provide minimal indication of a CSP's overall security posture. Instead, to establish an agile process, governments should focus on impactful vulnerabilities, including moderate or high findings, and control failures, deferring the tracking of low findings to CSPs' internal processes.

For example, government users should not have to lag behind commercial organizations that leverage Microsoft Azure's Operations Management Suite, a recently released information technology management solution that, among other capabilities, allows users to identify missing system updates and perform log analytics in their cloud environment. To the fullest extent possible, initial and ongoing authorization latencies should not limit governments' access to modern and upgraded cloud security solutions.

An effective ongoing authorization process focuses on assessing CSP practices through a subset of metrics that capture information related to the risks and threats of greatest concern, including but not limited to:

- Status of CSP's vulnerability and configuration management and threat assessment practices;
- Any significant breach in the certified environment;
- Significant levels of service interruption in a certified environment (i.e. breach of availability);
- Any significant and systemic weaknesses identified from recent continuous monitoring reports (e.g. high risk findings that have not been remediated);
- Control implementations that have changed since the last assessment; and
- Newly added controls or identified risks that were not considered in previous assessments.

Governments can then develop templates to capture these reporting requirements for CSPs to supply on a periodic basis as deemed appropriate for sufficient operational visibility (e.g. monthly, quarterly). CSP reports should be validated by an independent assessor to ensure sufficient technical review and verification prior to a



government's risk review. Governments can use these third party validated reports to verify that CSPs are implementing effective practices to secure their environments and maintain their authorizations.

Step 3: Establish a continuous improvement loop to maintain program effectiveness and align with broader strategies

As any cloud assurance program matures, it faces inevitable shifts in priorities, areas in need of greater efficiencies, and evolution of broader cybersecurity strategies. Therefore, having a process of continuous improvement is vital to ensuring that program operations and management are periodically re-evaluated and improved and that coordination between policies is maintained. As a starting point, governments and ministries should leverage the outcome-focused metrics reporting provided to the oversight authorities, CISOs, or CIOs to help to verify the effectiveness of the cloud assurance program and identify areas in need of adjustment.

Establishing an effective continuous improvement program also requires incorporating the expertise and experience of industry stakeholders, including CSPs that architect and operate cloud systems and independent assessors that regularly analyze and validate the security of those systems. Governments can determine whether to solicit industry feedback or to regularly engage with a formally recognized industry body. Regardless of the model, continuing to adapt will ensure that a government's cloud security program stays in step with ever-evolving technology and the cybersecurity landscape.

Ultimately, continuous improvement is essential to achieving a risk-based approach,³ which is foundational for any information and communications technology (ICT) system but especially important for cloud computing. In order for governments to access rapid advances in service quality and security, they must develop an agile cloud assurance program.

³ NIST Risk Management Framework (RMF), <http://csrc.nist.gov/publications/nistpubs/800-37-rev1/sp800-37-rev1-final.pdf>, 5-7. For continuous improvement, at the organization level, approaches to risk identification or mitigation need to be reviewed; at the mission level, decisions about data sensitivity need to be reviewed; and at the information system level, selection and deployment of security controls need to be reviewed. Similarly, the NIST Cybersecurity Framework explains that risk should be assessed at the executive level, business/process level, and implementation/operations levels, and that learnings from each level should feed back into the others, <http://www.nist.gov/cyberframework/upload/cybersecurity-framework-021214.pdf>, 12.



Phase 2

Develop a cloud assurance program

With the structure now in place, governments should next look to the relevant authorities to develop the cloud assurance program. In developing the program, the authorities will be creating the structures, such as requirements baselines, that will help to substantiate trust between CSPs and government cloud users. The relevant authorities should take the following operational steps to ensure the successful development of their cloud assurance program.⁴

- 1) Adapt existing data classification scheme for cloud services
- 2) Define requirements baselines by leveraging global standards
- 3) Establish a cloud-adapted risk assessment process
- 4) Establish risk management process to address unique risk scenarios

Step 1: Adapt existing data classification scheme for cloud services

Governments have been classifying their systems and data for decades. Data classification enables an organization to categorize its systems and data based on the level of sensitivity and criticality. This qualification allows for CSPs and data owners to manage data according to its relative value or use rather than treating all data equally,⁵ resulting in more effective data protection. These processes have become increasingly important in the shared responsibility model of cloud computing since customer-generated data, while in the custody of the CSP, must be classified by the data owner.

⁴ Steps 1 and 2 are best undertaken by the CGSDO and steps 3 and 4 are best managed by the program operations and management office.

⁵ *Data classification for cloud readiness*, Microsoft Trustworthy Computing (2014).



Implementing a data classification scheme is usually a two-part process:

1) Select a government-wide terminology model. Although many governments may already have a government-wide terminology model in place for on-premises technology, as they transition to cloud services, governments may consider whether this terminology model is sufficient or should be adapted for cloud systems. The CGSDO is best positioned to make this determination.

2) Identify and categorize data and systems according to that model. As described in Phase 3, the ministry/agency data owner is best positioned to identify and classify its data and systems.

When adapting a terminology model, CGSDOs should consider doing so in a way that continues to reflect the relative *sensitivity* and *criticality* of their data and systems.

For instance, supervisory control and data acquisition (SCADA) systems, which control the distribution of electric power, contain both real-time sensor data and routine administrative data; real-time sensor data is not impacted by a loss of confidentiality but is highly impacted by a loss of integrity or availability, whereas routine administrative data is minimally impacted by a loss of confidentiality, integrity, or availability.⁶

Sensitivity indicates the likely impact resulting from compromise, alteration, or misuse of data; the more sensitive the data or system, the more essential its confidentiality and integrity. *Criticality* indicates the likely impact resulting from loss of availability of the data or system; the more critical the data or system, the more essential its availability. Sensitivity and criticality are both essential in making data or system categorization decisions.

Figure 2: Classification levels

Below, we introduce a terminology model based on the U.K. government’s approach. In 2013, the U.K. government overhauled the way that it classified data, drastically reducing its number of classification levels to three.⁷ While governments may have differing compliance needs that drive them toward more classification levels, the advantages of the U.K. government’s approach are its simplicity and transferability to the majority of workloads considered for cloud migration. After its overhaul, the U.K. government realized that up to 90 percent of its public sector business could be classified as “official,” resulting in that data being “cloud ready” immediately.⁸

Sensitivity and Criticality	UK Model
High	Top Secret
Medium	Secret
Low	Official

⁶ Federal Information Processing Standards Publication (FIPS PUB 199), *Standards for Security Categorization of Federal Information and Information Systems*, NIST (2004), <http://csrc.nist.gov/publications/fips/fips199/FIPS-PUB-199-final.pdf>.

⁷ Richard Kemp, *Seeding the Global Public Sector Cloud: Part II – The UK’s Approach as Pathfinder for Other Countries* (2015), <http://www.kempitlaw.com/wp-content/uploads/2015/10/Part-II-Seeding-the-Global-Public-Sector-Cloud.pdf>.

⁸ *ibid*



Step 2: Define requirements baselines by leveraging global standards

Cloud security requirements should reflect standardized sets of security controls corresponding to each classification level (i.e. a higher baseline should correspond with higher sensitivity and criticality). These should be comprehensive enough that they minimize the need for ministries or agencies to add their own particular controls but not so broad that they encompass one-off controls that are not widely used. The baselines should also be adopted government-wide, with ministries or agencies maintaining discretion to add specialized requirements to address unique risk scenarios.

To maximize efficiency in developing the baselines, the CGSDO can leverage a number of widely reputed, tested, and adopted international standards and national frameworks. Alternatively, governments can develop their own requirements, map them to global certifications, and then validate the sufficiency of applicable certifications through an independent assessor. Both options provide a government-tested approach that delivers security assurance and economies of scale. In addition, governments should scope requirements differently for IaaS, PaaS, and SaaS. The activities supporting this step are:

- 1) Leverage global standards.** We recommend governments look to globally-recognized frameworks to structure their evaluation of a CSP's risk posture and develop requirements baselines. This will ensure the practicability and efficiency of developing, implementing, and demonstrating compliance against a core set of risks that are common across sectors, ministries/agencies, and governments. Adopting this approach helps to streamline the certification process, achieve scalability, and stay synchronized with cloud innovation cycles.

Focusing on the desired security objective behind a specific requirement can help governments rightly focus on *what* security outcomes they want to achieve rather than *how* to achieve them. While controls provide a valuable way to consistently measure security practices, our experience has shown that organizations that focus primarily on controls implementation limit their access to the benefits of best-in-class security innovation.

Instead, as governments evaluate whether CSPs meet their security bar, they should clearly articulate their security outcomes and allow the CSP to develop the optimal techniques for meeting (if not exceeding) those outcomes. Most governments will likely begin with existing prescriptive security controls frameworks with a process in place to accept alternative implementations. However, as cloud assurance programs mature and CSPs continue to rapidly innovate improved security features, controls implementation details will become increasingly irrelevant and, ultimately, defunct.

While the desired end-state is a framework based on customer-defined security outcomes and CSP-determined security techniques to meet those outcomes, it will likely result from a progressive dialogue that requires collaboration across the cloud assurance stakeholder community. As governments concurrently work to continuously improve their cloud assurance programs to this desired end-state, we offer interim steps that governments can implement today. These interim steps enable governments to define requirements baselines, which can also provide a starting point for dialogue with cloud assurance stakeholders.



One helpful reference is the Cybersecurity Framework, developed by NIST, which maps risk management steps to existing standards, guidelines, and practices for reducing cybersecurity risk.⁹ The Cybersecurity Framework has five overarching functions that provide a high-level, strategic view of risk management: identify, protect, detect, respond, and recover.

Each function contains multiple categories, which disaggregate the functions from discrete areas into concise statements of desired outcomes.¹⁰ The categories are further broken down into subcategories, which provide even more specific statements of desired outcomes and, where applicable, controls and practices.¹¹ Importantly, the Framework's subcategories are formally mapped to a number of widely-adopted international and national standards, practices, and security controls, including ISO 27001 and NIST Special Publication (SP) 800-53 Revision 4 (Security and Privacy Controls for Federal Information Systems and Organizations).

Examples from Australia, New Zealand, Singapore, UK

Many governments around the world are using ISO 27001 as their starting place to manage common risks. For instance, 70 percent of the controls included in Australia's 2015 Government Information Security Manual, New Zealand's Cloud Computing Information Security and Privacy Considerations, Singapore's Specification for Multi-Tiered Cloud Computing Security, and the UK's G-Cloud map to the controls included in ISO 27002.

Other resources commonly used by governments include ISO 27018, Service Organization Controls (SOC) 1 and SOC 2, and NIST SP 800-53 Revision 4, and the wide use and overlap among these existing best practices helps to demonstrate their sufficiency as a starting place for any government. In fact, mapping ISO 27001 and SOC 1 and 2 to NIST SP 800-53 Revision 4 demonstrates roughly an 80 percent overlap in security practices.

Global certifications are available for ISO 27001, making it highly compatible with existing best practices like the Cybersecurity Framework.¹² For governments looking for more in-depth guidance to leverage, NIST SP 800-53 Revision 4 can supplement ISO 27001.¹³ In addition, ISO 27018, the first global cloud privacy standard, establishes guidelines for protecting personally identifiable information (PII).¹⁴ Governments can obtain third party certification assurances to verify that CSPs are adhering to ISO 27001 for security and ISO 27018 for privacy.¹⁵

⁹ NIST Cybersecurity Framework, <http://www.nist.gov/cyberframework/upload/cybersecurity-framework-021214.pdf>.

¹⁰ For example, within the "protect" function, a category is access control, which requires that "access to assets and association facilities is limited to authorized users, processes, or devices, and to authorized activities and transactions."

¹¹ For example, within the "access control" category, a subcategory is "remote access is managed."

¹² To achieve the ISO 27001 certification, CSPs must: evaluate their risks, assessing threats, vulnerabilities, and impacts; design and implement a comprehensive set of information security controls and other risk management tools to mitigate risks; and design and implement an overarching risk management process, resulting in continuous monitoring of and improvements in risk mitigation. A CSP's compliance with ISO 27001 thus demonstrates to government and other customers a continuous and long-term investment in information security best practices and resilience. Likewise, to achieve the ISO 27018 certification, CSPs must process PII in conformity with customer requirements; define processes for the transmission and destruction of PII; and document and communicate data security infringements to customers, among other requirements.

¹³ NIST SP 800-53 Revision 4, <http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-53r4.pdf>.

¹⁴ http://www.iso.org/iso/catalogue_detail.htm?csnumber=61498.

¹⁵ *Building trust in the cloud: Creating confidence in your cloud ecosystem*, EY Insights (2014), <http://ey.com/GL/en/Services/Advisory/Building-trust-in-the-cloud>; *ISO 27018—Protection for Personally Identifiable Information*, CIS, <http://www.cis-cert.com/Pages/com/System-Zertifizierung/Cloud-Computing/Data-protection/ISO-27018/ISO-27018.aspx>.



2) Recognize and re-use applicable certifications and artifacts. Once frameworks have been reviewed, governments should move forward by basing their cloud requirements on or mapping existing requirements to them. Then, they can have an independent assessor validate the mapping and controls implementation against the initial certification.¹⁶ The outcome of this assessment should identify:

- Controls satisfied under the initial certification;
- Controls not addressed that require testing; and
- Controls that require re-testing.

To the extent that governments re-use existing certifications and artifacts, they continue to gain efficiencies and reduce costs that ultimately cascade to the government customer.

3) Scope requirements based on cloud service delivery model. In cloud computing, a fundamental construct is the layered and distributed nature of systems and applications. Given that software applications leverage an underlying infrastructure, and thereby inherit their underlying security controls, having distinct requirements defined for IaaS, PaaS, and SaaS can streamline the assessment process. More specifically, a SaaS application undergoing certification can reference the accredited security assessments of the IaaS/PaaS that it leverages and be assessed for only those new requirements unique to the SaaS layer. Governments can thus expedite ministry or agency consumption of the “latest and greatest” cloud services, features, and security techniques, which are released at a frequency that could not otherwise be supported. In the absence of existing requirements baselines for IaaS, PaaS, and SaaS, governments should consult CSPs and independent assessors to develop them.

Step 3: Establish a cloud-adapted risk assessment process

After the CGSDO defines a government’s cloud security requirement baselines, the program operations and management office needs to establish a process through which the CSP-authorizing entity (i.e., the program operations and management office or a ministry/agency) assesses risks by utilizing the independent assessor’s CSP risk posture report. Within their evaluations, independent assessors should develop a summary risk report: a concise and objective representation of the risks identified by the CSP and independent assessor. The risk report should also demonstrate whether the cloud service’s security features and controls result in a security level consistent with the relevant requirements baseline/s.

As the authorizing entity evaluates the risk report, it should follow a process consistent with ISO 31000, a global risk management standard. ISO 31000 recommends the following:

1) Identify important risk events (i.e. threats and vulnerabilities), including those of a strategic, operational, compliance-related, and technical nature.¹⁷ The first activity in performing a holistic risk assessment is to identify potential threats, vulnerabilities, and consequences of concern. As ISO 31000 and *Assuring the Security of Cloud Services*¹⁸ explain, governments are encouraged to bring together a variety of

¹⁶ Independent assessors should perform this validation since their roles require proficiency in cloud architecture, operations, and security controls implementation.

¹⁷ *Assessing compliance & risk for cloud computing deployments*, <http://az370354.vo.msecnd.net/publicsector/government/MIC0675%20Cloud%20Field%20Booklet%20A4%20S6R1.pdf>.

¹⁸ James Kavanagh, *Assuring the Security of Cloud Services*, <http://aka.ms/safehandbook>; *Assessing compliance & risk for cloud computing deployments*, <http://az370354.vo.msecnd.net/publicsector/government/MIC0675%20Cloud%20Field%20Booklet%20A4%20S6R1.pdf>.



stakeholders, including public and private sector entities, to catalogue relevant risk scenarios as part of the process. To streamline what could otherwise become a cumbersome process, *Assuring the Security of Cloud Services* recommends that governments identify only their most important risk events. To that end, it also provides a starting place, a catalogue of 50 risk events commonly assessed by governments. The catalogue can also be leveraged to help governments undertake threat modeling or to create conceptual hierarchies that depict the various ways in which an asset can be attacked (i.e. depicting risk scenarios).

2) Determine risk likelihood, impact, and tolerance for each of those events. For each risk event, governments should also assess potential consequences by considering:

- Risk *likelihood* (probability on a scale from rare to almost certain that a risk event will happen);
- Risk *impact* (consequence of a risk event on a scale from minimal to catastrophic); and
- Risk *tolerance* (the level of risk an organization is willing to accept for a risk event, measured on a scale from very low to extreme).

These will vary according to the risk event and the high, medium, or low sensitivity or criticality of the data or system affected.

For example, a breach of high sensitivity data may have a catastrophic impact and fall into the possible category on the likelihood scale; such an event would likely then be coupled with a very low risk tolerance. Alternatively, a breach of low sensitivity data may be rare on the likelihood scale and have minimal impact, resulting in an extremely high risk tolerance.

Step 4: Establish risk management processes to address unique risk scenarios

In both on-premises and cloud environments, governments can mitigate, avoid, transfer, or accept risk. Unique risk scenarios may be **mitigated** through outcome-based requirements, **avoided** through cloud architecture or delivery decisions, **transferred** through service level agreements or insurance, or **accepted** when tolerable to achieve efficiencies.

For instance, governments might avoid risks for which they have an extremely low tolerance and accept risks for which they have an extremely high tolerance. Most common risks can be reduced or mitigated through the leveraging of global standards like ISO 27001. However, under a very narrow set of circumstances, an authorizing entity may identify risk scenarios that require additional security measures, especially for highly sensitive or critical data or systems.

1) Mitigate. Governments that face a unique risk scenario that no existing standard or certification sufficiently reduces may want to develop a new requirement to address that particular risk. If the risk is government-wide, then the CGSDO is the appropriate entity to set the requirement. If the risk scenario is unique to a ministry/agency cloud consumer, then customized requirements may be developed by the authorizing entity. While the authorizing entity must be given the flexibility and authority to add controls, this discretion should be subject to the CGSDO's oversight so that the level of additional security gained is balanced against the cost of implementing additional controls or operating at higher parameters.

To empower the authorizing entity to make this assessment, any new requirement should:



- clearly articulate the security value proposition of the added control or heightened parameter;
- show that existing requirements do not provide the same or similar protection;
- demonstrate that the security value justifies the costs that will be incurred by the CSP, and ultimately the ministry or agency, to mitigate the unique risk scenario; and
- Accept alternative measures or compensation controls that demonstrate security equivalency.

Given the pace of cloud innovation and the inherent variability in cloud architectures, customized requirements should be adaptable enough to allow for alternate implementations. Such adaptability is achieved through requirements that are outcome oriented, meaning that they are focused on *what* outcome is needed, not on *how* best to achieve that outcome. By focusing on outcomes rather than methods, governments permit CSPs to find the most innovative and practical solutions, which may be native to a CSP's service offering or require an innovative solution to be developed by the CSP.

Ascertain the security value-add of new requirements

Governments should consider the additional complexity involved when requiring customized security measures for cloud environments. Cloud-based systems are architected to be automated, on-demand, scalable, and continuously updated with enriched feature sets. As a result, excessive one-off, customer-specific requirements can substantially challenge the provisioning of seamless, agile, and low-cost cloud services.

Entirely new requirements and certifications should only be developed when governments require that CSPs take on activities to achieve security outcomes beyond what those cloud providers already achieve through existing standards and audit regimes. If new requirements do not remove the need for existing certifications or address substantial net-new domains, then they will create marketplace confusion and result in inefficiencies rather than increased security.

2) Avoid. Governments may also identify unique risk scenarios and decide to avoid those risks by choosing a different cloud architecture or delivery mechanism. Cloud architecture refers to cloud service and deployment models, including IaaS, PaaS, and SaaS solutions, as well as public, private, or hybrid solutions. Cloud delivery mechanisms include direct service provisioning by a CSP, as well as service provisioning partnerships, often between global CSPs and local CSPs or technology companies. Responsibilities for and implementation of security controls vary according to cloud service and deployment model choices. Therefore, risk scenarios related to customer control requirements may respond most noticeably to architecture decisions that alter these responsibilities and implementations. This activity is further described in Phase 3.

3) Transfer. Unique risk scenarios may also be transferred through service level agreements (SLAs) or insurance. SLAs are a particular form of contract relevant to cloud service offerings. CSPs may offer SLA commitments related to various aspects of service quality, including uptime or service availability; for instance, Microsoft's Azure offers, at minimum, 99.9 or 99.95 percent availability for many services.¹⁹ In addition, government customers can leverage insurance to transfer unique risk scenarios, particularly those resulting in a high financial exposure.

4) Accept. Because mitigating, transferring, and avoiding unique risk scenarios will likely add costs or bear other undesirable outcomes (e.g. performance latency), a government may ultimately decide to accept some risks when the decision to do so is tolerable. Governments that leverage global standards and best practices are in the best position to undertake risk acceptance because such standards and practices provide a reasonable set of security domain coverage for cloud services.

¹⁹ SLA summary for Azure services, <https://azure.microsoft.com/en-us/support/legal/sla/summary/>.

Phase 3

Implement the program



With the cloud assurance program structured and developed, ministries and agencies must next work with the authorizing entity to certify and integrate cloud services. As the data owner, the ministry/agency also has specific roles, responsibilities, and obligations in the context of cloud computing risk management.

For instance, ministries/agencies are responsible for appropriately classifying data or enforcing policies for managing user account access to specific data and systems. Microsoft’s experience in operating a hyper-scale, commercial cloud environment has shown that customer controls implementation is a significant source of security vulnerability.

Ministries and agencies will need to select a cloud service delivery and deployment model that addresses their business and security requirements, using the following steps:²⁰

- 1) Identify and categorize data and systems for cloud migration
- 2) Implement requirements baselines, risk assessments, and approaches to addressing unique risk scenarios
- 3) Select a cloud service delivery model (IaaS, PaaS, SaaS)
- 4) Select a cloud deployment model (public, private, hybrid, community)

While CSPs manage the physical infrastructure and host operating system and virtualization layers, customers are responsible for ongoing risk assessments and management of the environments that they operate—including the guest operating systems and other associated application software.

Step 1: Identify and categorize data and systems for cloud migration

This step covers the processes for identifying and categorizing data and systems according to the terminology model adapted by the CGSDO. Some questions that governments should be considering as they undertake this step include: what are the roles necessary to ensure that systems and data are treated according to their

²⁰ Note that while these steps are described sequentially, they can just as effectively be undertaken concurrently.



relative sensitivity and criticality? What functions and workloads does my organization want to move to the cloud? And how sensitive or critical are the systems and data?

To address these questions, ministries and agencies need to leverage their government-wide terminology models and have in place a data classification process that includes the following activities:

1) Clarify data handling roles. Ministries and agencies should first clarify the various data ownership and handling roles—such as data owners, custodians, administrators, and users²¹—so that appropriate levels of access and accountability are in place. Typically, *data owners*, the original creators of the data, will first assign a data classification based on the government-wide terminology model. This is consistent with what is envisioned in the NIST Risk Management Framework.²² At any point, the data owner may delegate ownership, assigning it to a *data custodian*; in the cloud context, that custodian will often be a CSP. *Data administrators* are typically in charge of various functions for the relevant data set, including archiving or storing data or monitoring logs. In addition, there may be *data users*—people or organizations that require access to the relevant data to do their jobs or fulfill their responsibilities. Data owners, administrators, or custodians may grant access to data users.

2) Identify functions, systems, and data that will be migrated to cloud services. Next, ministry or agency data owners should identify the functions, data, and systems they want to transition to cloud services. This decision will likely begin with a discussion around the functions that are most critical to move to the cloud (e.g. e-mail, collaboration workloads) and then transition into a more granular discussion on the specific systems and data supporting that particular function.²³ Once data and systems are identified, they must then be classified according to the sensitivity and criticality terminology model, typically a three-tiered structure. As stated earlier and exemplified in Figure 2, this classification tier will determine the applicable security controls or baselines needed for each of the designated levels.

3) Categorize data and system sensitivity and criticality. This step applies the government-wide data classification terminology model to data and systems considered for cloud migration. Ministries and agencies should consider *how* data owners within their organizations will categorize both existing data and new data that is continuously created.²⁴

Step 2: Implement requirements baselines, risk assessments, and approaches to addressing unique risk scenarios

In this step, authorizing entities, independent assessors, and cloud-consuming ministries/agencies implement the requirements baselines and risk assessment processes. A significant question that governments must

²¹ *Data classification for cloud readiness*, Microsoft Trustworthy Computing (2014)

²² NIST RMF, <http://csrc.nist.gov/publications/nistpubs/800-37-rev1/sp800-37-rev1-final.pdf>, (Step 1 – Categorize Information System, Task 1-1); *Data classification for cloud readiness*, Microsoft Trustworthy Computing (2014).

²³ Data exists at rest, in process, or in transit and may be either structured or unstructured. Structured data will often be found in databases and spreadsheets, whereas unstructured data exists in places such as documents, source code, and email. Moreover, data does not exist in isolation but as part of a system, and the function of that system should be identified and taken into consideration.

²⁴ Several types of processes exist, including manual processes, location-based processes that classify data based on a user's or system's location, application-based processes such as database-specific classification, and automated processes that use various other technologies.

²⁴ Extreme data sensitivity may drive an organization to use manual data classification, a relatively resource-intensive option, but application- or location-based processes are useful for quickly categorizing data like health or financial records or email.



answer for themselves is whether the ministries/agencies or the program operations and management office should act as the CSP-authorizing entity. Ministries and agencies act as authorizing entities in a federated model, and the program operations and management office acts as an authorizing entity in a centralized model. Governments may also evaluate decentralized approaches to cloud certification functions.

As indicated in Figure 3 below, there are three approaches to cloud certifications functions, each one having benefits and downsides.

In a **centralized model**, ministries and agencies leverage centrally-developed and -evaluated requirements baselines without adding new requirements, and the program operations and management office acts as the CSP-authorizing entity.

In a **federated model**, ministries and agencies leverage centrally-developed requirements baselines but may add requirements to address unique risk scenarios, and ministries and agencies also act as the CSP-authorizing entity. The federated model results in slightly less efficiency than the centralized model, but it empowers ministries or agencies as the primary authorizers of a cloud service. As ministries and agencies migrate workloads to cloud services, they own the risk for those workloads, meaning that they are responsible for protecting their own data and complying with requirements.

In a **decentralized model**, ministries and agencies develop their own requirements baselines and act as CSP-authorizing entities.

Figure 3: Approaches to Cloud Certification Functions

Key Questions	Centralized	Federated	Decentralized
Who develops the requirements baselines? →	CGSDO	CGSDO – and ministries may add requirements for unique risk scenarios	Ministries
Who acts as the authorizing entity? →	Program operations and management office in partnership with independent assessor(s)	Ministries in partnership with independent assessor(s)	Ministries in partnership with independent assessor(s)
What are the benefits? →	Maximum efficiency and cross-government learning	Efficiency and cross-government learning; ministry risk ownership	More ministry risk ownership
What are the drawbacks? →	Less ministry risk ownership	Less efficient than centralized model	Significant inefficiency and higher costs, less cross-government learning

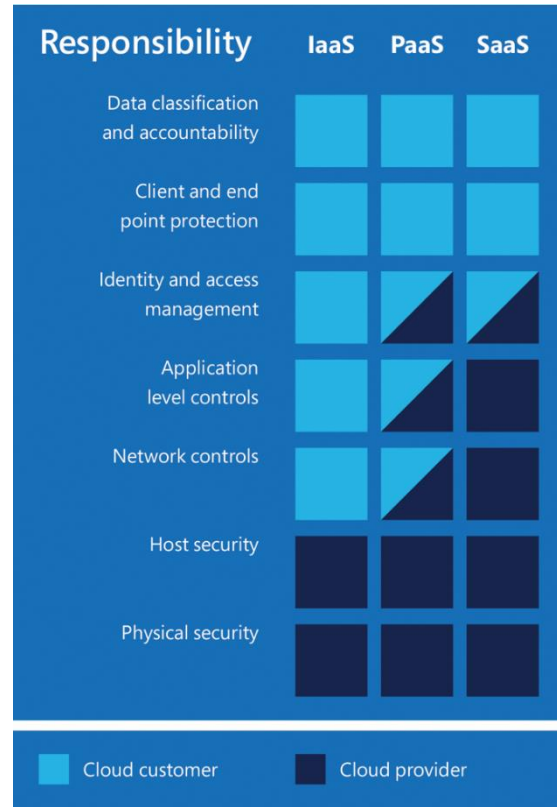


Step 3: Select a cloud service delivery model

As ministries and agencies determine which data, systems, and services they want to migrate to the cloud and how they will manage risks, they should also consider which cloud service and deployment models are most fitting for their needs. As indicated in Figure 3, there are three major types of cloud services models: IaaS, PaaS, and SaaS.

- **IaaS** pools hardware resources for compute, storage, and connectivity capabilities, over which a customer can deploy and run operating systems and applications (i.e., PaaS and SaaS).
- **PaaS** delivers application execution services and often an operating system, enabling customers to create and deploy their own applications (i.e., SaaS) with greater agility.
- **SaaS**, also referred to as “on-demand software,” delivers ready-to-use applications, such as e-mail, customer relations and management systems, or Microsoft Office, on scalable cloud infrastructure.

Figure 4: CSP v. Customer Responsibility



As data custodians, CSPs add a security dependency between data and their owner, requiring CSPs and ministry or agency customers to coordinate as they implement controls. As depicted in Figure 4, in each of the cloud service models, the responsibility for various security functions is divided between the CSP and the ministry or agency customer. In any service model, the CSP manages the underlying cloud infrastructure—the datacenters that power the cloud service. However, responsibility for various security controls otherwise varies. As a result, risk scenarios related to customer control requirements may respond most noticeably to architecture decisions that alter these responsibilities and corresponding levels of control. Systems and data sets over which governments want to retain greater structural control, for instance, may be more suitable for IaaS or PaaS solutions, within which governments have more flexibility regarding security implementations. Alternatively, for SaaS solutions, CSPs take on a great degree of responsibility for the implementation of security controls, reducing the breadth of customer responsibility compared to IaaS or PaaS solutions.

In any service model, coordination between CSPs and ministry or agency customers is key. Therefore, in addition to assessing CSPs, governments should also carefully assess ministry or agency implementation of security controls; cloud environments result in shared security responsibilities between CSPs and customers. In each service model, government customers and CSPs may have full or shared responsibility for certain security controls. For instance, SaaS providers are responsible for managing service-level capabilities, which include employing security best practices such as penetration testing and defense-in-depth to protect against cyber threats. SaaS providers are also responsible for physical and data security in the form of employee access controls, encryption of data in transit, and enabling strong authentication. However, customer responsibilities



include user identity and access controls, device management, and data management (e.g. rights management services, data loss protection), which are unique activities that the customer must implement. These security activities, which are under the customer's purview, empower the customer to control, access, and protect its own data.

Step 4: Select a cloud deployment model

Similar to the process of selecting a cloud service model, selecting a cloud deployment model can impact customer risk scenarios and should be driven by a ministry's or agency's mission needs and business requirements. Customers have four choices when evaluating cloud deployment models:

- *Public clouds* are multi-tenant environments in which CSPs make available to the general public their infrastructure, including storage and applications.
- *Private clouds* are environments operated solely for a single organization; they may be managed by that organization or by a CSP.
- *Community clouds* are environments operated for a group of organizations; they may be managed by those organizations or by a CSP.
- *Hybrid clouds* are when public, private, or community clouds remain distinct but are bound together with on-premises ICT by common technology that enables data and application portability.²⁵

Customer decisions about whether to deploy public, private, community, or hybrid cloud platforms are often driven by perceptions of the level of security and customer control. Public cloud models provide distributed resources, resulting in unprecedented efficiencies, cost savings, and resiliency, and the newest features and security techniques are applied to the multi-tenant environment first because of the expansive, world-wide user base that it supports. According to Forrester research firm, there is increasing evidence that more enterprises are adopting public cloud platforms as "best, not only for customer-engagement apps but for analytics and core-business apps as well."²⁶

Alternatively, private cloud models enable greater customization. Nevertheless, meeting customers' security objectives may not directly correlate to the need for a private, dedicated infrastructure. Large CSPs, like Microsoft, have robust capabilities for managing a shared infrastructure while still providing significant and auditable assurances of the security of customer data, including through logical isolation. In other words, while dedicated private cloud solutions can be more specialized, a multi-tenant public cloud is still subject to the same security controls. In addition, due to the large customer base and demand, multi-tenant public and community clouds are prioritized for certification.

Data hosted in the cloud often moves between different services and devices, and given the global nature of commerce and of cloud services, data may also need to move across borders. Some CSPs offer customers choice in where their data resides, mitigating concerns about data sovereignty.²⁷ In other contexts, governments may opt for dedicated, private cloud solutions. Requiring all public sector data to be subject to

²⁵ More information about these deployment models is available in *Transforming Governments. Transforming Government: A cloud policy framework for innovation, security, and resilience* (2015), <http://aka.ms/cloudsecurityprinciples>.

²⁶ <http://www.itwire.com/business-it-news/cloud/72765-%E2%80%98disruption%E2%80%99-ahead-in-maturing-public-cloud-market-forrester.html>

²⁷ Microsoft Azure and O365 enable their customers to choose where they data resides, <https://azure.microsoft.com/en-us/support/trust-center/>; <https://products.office.com/en-us/business/office-365-trust-center-welcome>.



data sovereignty concerns is not consistent with fostering an open, global Internet or with cloud-first principles, and in most circumstances, with effective data classification, governments can ensure that relevant data stays within the confines of a regional selection and travels only between countries with data transfer agreements in place. However, under a very narrow set of circumstances (i.e. top secret data), a data residency requirement may be appropriate. Where local cloud service provisioning is preferable to avoid unique risk scenarios related to extremely sensitive data, service provisioning partnerships between global CSPs and local CSPs or technology companies may be considered.

Conclusion

The goal of a government cloud assurance program is to manage information security risks while still enabling that government to take advantage of the many benefits and opportunities of cloud services. Achieving that goal requires risk-based decision making at every step of a government's process of developing and implementing a cloud assurance program.

Each of the steps outlined in this paper contributes to this risk-based approach. Having an effective governance model in place establishes the roles and responsibilities necessary to consider risk and efficiency and to determine whether new technologies are able to be consumed. In addition, determining data and system sensitivity and criticality requires a government to weigh the relative risks related to the confidentiality, integrity, and availability of different data sets and systems. Leveraging global standards enables governments to achieve a high level of security with maximum agility and efficiency, and assessing and managing unique risk scenarios not mitigated by global standards solidifies a risk-based approach. Governments that establish ongoing authorization processes also ensure that highest priority risks are regularly evaluated.

Ultimately, a risk-based approach must also be instilled through continuous improvement, a process during which governments evaluate how effectively risks are being managed and how risk priorities might be shifting. With a risk-based, agile cloud assurance program deployed, ministries/agencies can migrate workloads to the cloud of choice and assume control of the customer-facing environment—effectively achieving cloud assurance and trust.