Treasury Board of Canada Secretariat
Secrétariat du Conseil du Trésor du Canada
Canada

Application Modernization Guidance

Building Modern Applications Using Platform Services

Version 3.0

Introduction

Cloud computing has introduced a fundamental shift in the way information system services are delivered. The Government of Canada (GC) has a cloud first policy requirement that is aimed at leveraging this service delivery model. The GC Cloud Adoption Strategy directs Departmental and Agency CIOs to consider the following Cloud service models, in the following order of priority:

- Software as a service (SaaS),
- Platform as a Service (PaaS),
- Infrastructure as a Service (IaaS).

When migrating existing application to the cloud, IaaS is often the target service model. IaaS allows an organization to virtualize many of the same resource available in a data centre and represents the minimum change from status quo. While this may be a pragmatic strategy, in the long-term, it will not lead to full realization of the benefits of public cloud. Departments and Agency are encouraged to leveraging Platform-as-a-Services to achieve the full benefits of lower operational burden and more timely delivery of IT services.

The authors of this document would like to thank the GC Cloud & Computing Network of Expertise for their feedback

Navigating Cloud's Service Models

Three common architectural descriptions used to describe public cloud's three service models are:

Software as a service (SaaS)[1]

For the GC this is public Cloud services that delivers business solutions to end users. Microsoft Office 365 is an example of SaaS solution used by the GC. Other examples include CRM, ERPs and business intelligence. SaaS is typically used by end users such as employees and Canadian cictizens.

Platform as a service (PaaS)[1]

Platform as a Service (PaaS) are public cloud services used to build and deliver business solutions. PaaS is typically used by IT professionals, data scientists, and specialized professionals, less so by end-users. Serverless and back-end-as-a-service are another terms often used to describe categories of platforms as the public cloud service provider manages the servers, storage, and software used to deliver the platform. Examples include analytics, containers, databases and a variety of managed services.

Infrastructure as a service (IaaS)[1]

Infrastructure as a Service are public cloud services used to host software in support of business solutions. IaaS is typically used by IT professionals leveraging networking, storage, and compute resources.

While these service models may be convenient architectural definitions, service providers rarely segment their products along those definitions. Often products will be composed of multiple service models.

SaaS provides a complete business solution, however, for the most part, the SaaS market provides solutions for common business capabilities found across all enterprises such a Customer Relationship Management, Enterprise Resource Planning Systems, collaboration, email, planning tools, amongst others. When the market doesn't offer a complete business solution, PaaS and IaaS allow organizations to build custom business solutions. While IaaS is often attractive as it offers services that mimic services offered in traditional IT models for hosting software, it comes at the price of higher operational burden and increased compliance and engineering activities when compared to other service models. When building new and refactoring existing applications, platform-as-a-service should be an organizations primary choice of service model.

The Technology Stack

When considering the implication of choosing a service model, it is important to be mindful of the technology stack. Users interact with applications or websites that in turn are dependent on middleware

[1] Definitions are standard NIST definitions (https://csrc.nist.gov/glossary/term/)

and runtimes that in turn are dependent on operating systems (collectively referred to as ‘Platform’ in the illustration below) that in turn are hosted on a virtualization and abstraction software, that in turn is dependent on servers and storage, all of which resides in facilities such as data centres. This layering of technology dependencies is known as the technology stack and is illustrated below. Each layer of technology requires patching, upgrades, replacement, and monitoring amongst other activities. Those activities are collectively known as operational activities.

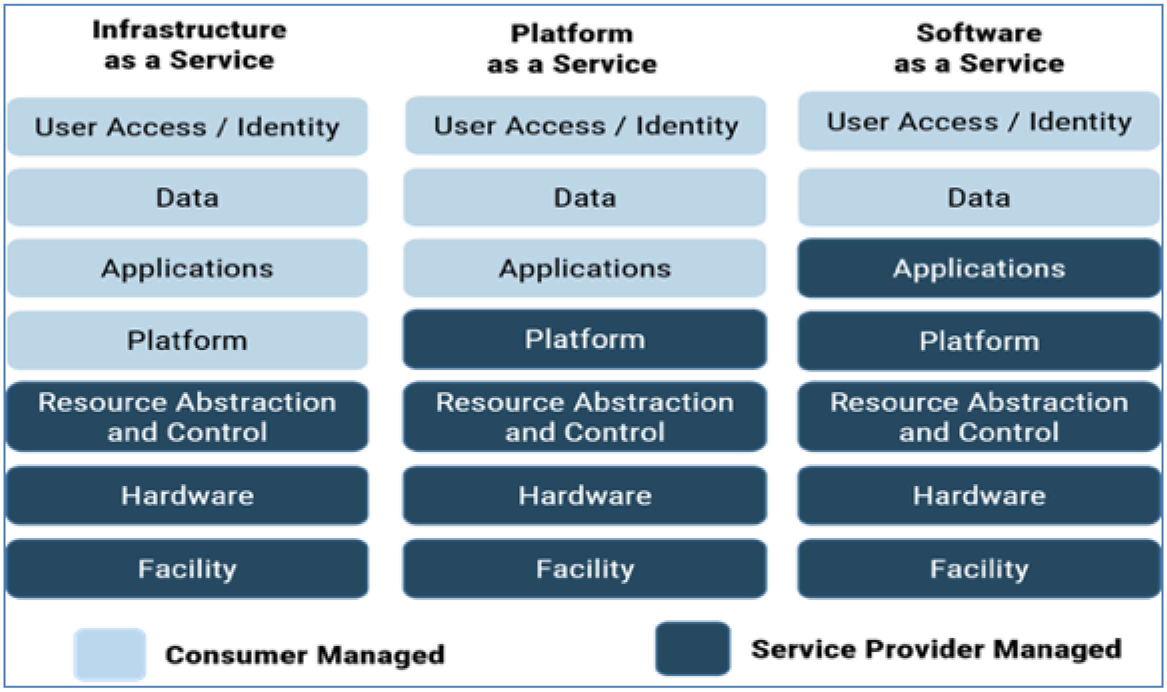

Infrastructure as a Service	Platform as a Service	Software as a Service
User Access / Identity	User Access / Identity	User Access / Identity
Data	Data	Data
Applications	Applications	Applications
Platform	Platform	Platform
Resource Abstraction and Control	Resource Abstraction and Control	Resource Abstraction and Control
Hardware	Hardware	Hardware
Facility	Facility	Facility

(Diagram: https://www.canada.ca/en/government/system/digital-government/modern-emerging-technologies/cloud-services/cloud-security-risk-management-approach-procedures.html)

For the most part, in traditional IT models, the Government of Canada manages the entire technology stack and in turn is burdened with all the operational activities required by each layer of the technology stack. Each layer is on a different lifecycle management schedule. Facilities may have long maintenance and replacement cycles, while platforms may have far shorter lifecycles. The management of these technology layers, while burdensome, are all focused on delivering one outcome; applications that provide business value in the form of automation and positive user experiences for Canadians. With public cloud services, the management of technology layers is divided between the cloud service provider and their consumer, in this case the GC. This is known as the shared responsibility model for cloud services. The GC remains focused on the upper layers of the technology stack and are related closest to delivering business value.

Faster Delivery of IT Services with Less Effort

When compared to Infrastructure-as-a-Service, Platform-as-a-Service provides developers, data scientists, and business users with the ability to deliver business solutions with less effort and faster delivery times. As the majority of the technology stack is managed by a public cloud service provider, the Government of Canada can maintain its focus on delivering business value.

By reducing the need to engineer software and infrastructure into solutions, users benefit from:

1. No patching of software. Patching of operating systems and applications is number two on the Canadian Centre for Cyber Security (CCCS) Top 10 IT security actions. While keeping software patched may seem simple enough when thinking about deploying a single application, across the Government of Canada, this creates tens of thousands of software assets to be patched. At the time of writing this guidance the GC has over 20,000 instances of Windows 2008 that require an upgrade to Windows 2016, an undertaking expected to require more than three years of effort. The cycle of funding, managing, and performing software upgrades never stops. Platforms, given their serverless architectures, greatly reduce or remove the need to patch software.

2. Reliability is inherited from the underlying cloud provider managed resources, however the user often must make configuration choices to achieve the level of reliability required for a particular application or service.

3. Ease of scaling. Small pilots require low effort, cost, and time to start. Scaling can be achieved through configuration.

4. The scope of performance and security monitoring is greatly reduced.

5. Cloud service providers deliver new features making it easy to maintain pace with the rate of technology change and improvements.

Given comparable capabilities, using a platform service is recommended over hosting of software using IaaS as it reduces operational burden and reduces time required to deliver business solutions

Challenges of Adopting Platforms

Given the benefits of PaaS, it would seem to be a natural choice, however organizations often experience challenges adopting PaaS for the following reasons:

1. Existing applications must be refactored to take advantage of platform services. Custom built applications are often tightly coupled. For example, applications may include commands that only work for a given database product or version. A ‘move and improve’ strategy can be used where existing applications are migrated to the cloud with the fewest changes possible and refactored over time. Consider using the principles found in 12-factor applications to avoid tightly coupled application architectures.

2. As portions of the technology stack are under the management of a cloud service provider, organizations may fear being locked-in as their span of control is reduced. The risk of lock-in is often less or the same as being locked into software hosted on IaaS, however the opportunities

gained of using platform services is greater than hosted software. For more information about how to evaluate technical lock-in, refer to the document *Application Modernization Guidance: Evaluating Technology Lock-in & Exit.*

3. Migrating existing applications to platform services may mean abandoning software licensing investments. Placing importance on past investments as justification for continuing those investments is known as a sunk cost fallacy. It can be difficult for an organization to overcome the emotional bias that software assets should be abandoned in favour of new solutions. To help ease this transition, the decision to abandon existing software licenses can be made over time. As the adoption of platform services grows, the case for continued investment or abandoning software licenses can be periodically re-evaluated.

4. For many IT professionals, the transfer of operational activities such as patching and upgrades to platform services constitutes a significant portion of their current jobs. The prospect of having these activities being replaced by a cloud service provider can invoke an emotional resistance to using platforms. The purpose of platforms is not to reduce jobs, but instead to refocus effort towards providing a better user experience. The need to provide Canadians with better digital services outstrips the GC's capacity to deliver them. The use of platform services allows the GC's valuable IT workforce to remain focused on digital, citizen-facing services.

5. Platform services are often designed around a zero-trust security model. In a zero-trust model, all interactions between systems and users are authenticated. This is a shift from past focus on perimeter-based security models where a perimeter is established and a higher degree of trust exists between all users and systems within that perimeter. To address this challenge, many platform services offer a hybrid model where zero-trust security is still the primary security model, however, the platform services can also be brought within a logical perimeter.

Example: Database-as-a-Service

In the research paper The Future of the DBMS Market Is Cloud, Gartner Inc provides data that demonstrates that database growth in the cloud now out-paces traditional data centre database growth. Public cloud database platforms, or Database-as-a-Service (DBaaS) offer the benefits of lower operational effort, ease of scaling, and new features that are continuously being delivered. Despite these benefits, adopting DBaaS can be a struggle for organizations. While lower operational burden is of benefit to achieve digital goals, for the database administers whose daily job has been occupied by those activities, the fear that comes with job evolution can cause resistance. Additionally, many organizations have large contracts with incumbent database product suppliers for which hundreds of millions of dollars have been invested in software purchase and maintenance. Organizations experience difficultly deciding how to sunset those investments. The simplest way to start using DBaaS is for new applications. This will provide an organization with the ability to explore DBaaS' benefits and provide the workforce with the ability to learn a new operating model for databases. No immediate decisions regarding existing database contracts are required. Eventually, however, organizations will start refactoring their applications for DBaaS.

Guidance

When evaluating platform services, the following guidance should be considered:

- When a technology capability can be delivered as software hosted on IaaS or PaaS, choose PaaS first.
- When evaluating the risk of technology lock-in, consider the guidance found in the document *Application Modernization Guidance: Evaluating Technology Lock-in & Exit*. Consider the benefits of lower operations and faster delivery of technology, not only the risk of exiting the technology in the future.
- When architecting with platform services, follow the principles of 12-factor applications.
- When evaluating the value of past software investments, avoid the sunk cost fallacy. Continuously evaluate the value of software licensing agreements.
- Provide a learning path and explain the evolution of operational roles to those who fear the impacts of using platform services.
- Refer to the technology stack to help define platform services, their value, and the responsibilities of the public cloud provider versus those of the consumer.