



Treasury Board of Canada
Secretariat

Secrétariat du Conseil du Trésor
du Canada

Canada

Guidance on Aligning Procurement Strategy with Cloud Strategy

Scaling with the SSC Cloud Framework Agreement

The Shared Services Canada Cloud Framework Agreement has been foundational for the growth of cloud adoption in the Government of Canada. It has allowed dozens of departments and agencies to gain access to cloud services that have been qualified to host Protected B Medium Integrity and Medium Availability services and data. It allows departments and agencies to access cloud while minimizing lengthy and burdensome procurement processes. Departments turned to the framework agreement during COVID to gain easy access to cloud services to deploy services to Canadians.

The framework agreement has been structured in such a way that to allow minimum burden of competition when purchases against the agreement are small, but the level of competition required grows as the size of the purchase grows.

The framework agreement's simplicity at the onset can become a source of consternation as the department or agency scales its usage. As more and more organizations scale their cloud usage, departments have sought advice and guidance on how best to leverage the framework agreement while not getting surprised later.

A typical example to illustrate the type of situations departments have become concerned about we will use the following example:

1. A department will make a purchase against the framework agreement for \$100,000 of IaaS services. Typically from a provider they want to 'try' or someone on the teams may already have used somewhere else. Since the dollar-value of the purchase is low, the level of justification for selecting the provider is also low.
2. The department deploys a landing zone and configures the cloud platform for their requirements. The department invests in a security assessment and authority to operate (ATO). The team grows its skills and experience with the platform. The foundation for cloud is poured.
3. The department continues to deploy more workloads as their abilities and requirements grow.
4. The department adds an additional \$500,000 to their original purchase to scale their usage of the cloud environment.
5. When the department wants to add another \$1M to their original purchase, they become aware that they are approaching a dollar-value where they will be required to compete their requirement among the providers on the framework agreement.
6. The department feels anxious as they have not invested heavily in building out a platform with their current IaaS provider. A competitive process may require them to change providers. Similarly, they may keep their current environment, but may need to start deploying workloads among other IaaS providers requiring additional landing zones, operations, and skill sets.

The solution to this type of scenario is early planning and aligning the organization's procurement strategy with their cloud strategy. This guidance document will use Gregor Hohpe's multi-cloud taxonomy to discuss the types of multi-cloud architecture an organization may want to achieve and how to use the SSC framework agreement to achieve that architecture. Hohpe's taxonomy helps to illustrate that multi-cloud does not have a single definition or architecture.

Even if your organization's strategy is not multi-cloud, but instead to invest in a single provider, this guidance document will still apply. Although multi-cloud may not be your strategy, it will likely be a pragmatic end-point. Often organizations invest in one primary cloud provider, but through restructuring, sacred-cow projects, or ungoverned usage, departments find themselves having to support a second or third provider, even if to a minor degree.

Competative Thresholds on the Framework Agreement

The SSC cloud services framework agreement is designed to escalate the level of competition required as the level of expenditure rises. [The Government of Canada Cloud Brokering](#) site can provide more details, however, at the time of writing this document, the agreement included this term:

1.4.4.3 Purchases greater than \$4,500,000:

When the total estimated cost of a requirement is greater than \$4,500,000 (before applicable tax(es)), the CA will conduct a Requirements Confirmation Request ("RCR"). The RCR documents will be sent to all FA Contractors by email. Upon receipt, FA Contractors shall indicate their interest in participating in the RCR within the timeframe specified in the email, according to their ability to fulfill the mandatory technical requirements defined within the RCR documents. Interested FA Contractors will be given a minimum of 10 FGWDs to respond with the information requested in the RCR documents such as demonstration of technical capabilities, pricing commitments, committed service levels, etc. The FA Contractor will not be paid for responding to an RCR and a response is at the Contractor's discretion.

What Is Multi-Cloud?

Every organization uses more than one cloud provider, but this is not necessarily multi-cloud. Multi-cloud is when multiple cloud providers are used to deliver the same capabilities.

To encourage a competitive environment, fairness amongst providers, and choice for IT organizations, the GC is a multi-cloud organization. The degree to which each individual department and agency is multi-cloud may vary. The GC is not a monolith organization and no one-size-fits-all approach works.

Examples of Multi-Cloud:

Collaboration. Office 365 from Microsoft and Workspaces from Google are competitive cloud services. They allow users to interact, collaborate, co-edit, email, conference, etc etc... Having both services available in your organization can be considered a multi-cloud scenario.

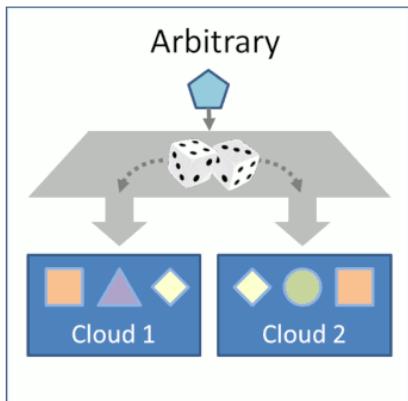
Case Management. Salesforce and Dynamics 365 are both low code platforms for delivering customer experiences, case management, and workflows. While differentiated, they are competitive cloud services. Having both of these cloud services in an organization would be considered a multi-cloud scenario.

Application Hosting. AWS, Google Cloud, and Azure offer competing services for hosting workloads amongst other other capabilities. Using more than one of those providers work be a multi-cloud scenario.

Application Hosting and Office 365. While using AWS and Office 365 is two different cloud providers, because they offer different capabilties, this is not considered a multi-cloud scenario.

Aligning Multi-Cloud Architecture to Procurement

Gregor Hohpe, in *Cloud Strategy; A Decision-based Approach to Successful Cloud Migration* points out that multi-cloud means many things to different people. To try and illustrate the different multi-cloud approaches that may exist, he has established a multi-cloud taxonomy. This guidance document will not repeat the fullness of his work, but instead the reader is highly encouraged to read *Cloud Strategy; A Decision-based Approach to Successful Cloud Migration* or see a brief overview on his [The Architect Elevator](#) blog.

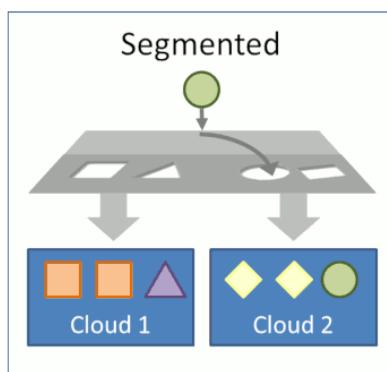


Arbitrary

In an arbitrary multi-cloud architecture, the use of multiple providers is largely ungoverned and left to local decision making without a single strategy. This is likely not an organization's multi-cloud strategy, but instead where they simply are with the intention to move out of this state. In this architecture one branch of the organization may be using one provider, because it's a best fit for them, but another project may be using a different provider because that is what they inherited.

To achieve this architecture with the framework agreement, not a lot of foresight is required. Each new business requirement can translate

to a new buy against the framework agreement allowing competition to decide the best provider for the workload.



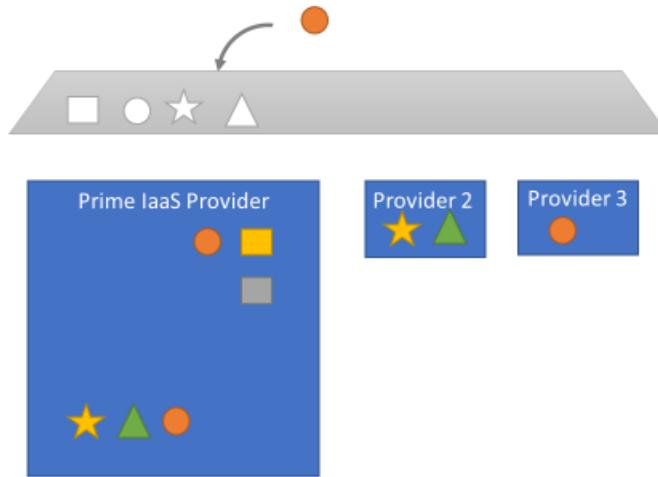
Segmented

In a segmented architecture, workload characteristics are used to determine which provider will be used to host a workload. An organization may wish to exploit the strengths of one provider over the other. For example, machine learning workloads are allocated to one provider with enterprise windows/linux workloads to another. The allocation across providers may not be even.

This is the architecture most departments probably have or will have; however, the segmented will be far more asymmetrical than depicted above.

Example 1: Multi-cloud-ish

In the GC, it is observed that most organizations are gravitating towards a single IaaS and single CRM provider. Much like most departments gravitate towards a single Relational Database Provider (Microsoft, Oracle, IBM), departments are selecting a single IaaS and single CRM provider. The logic behind this is simple; the investment required to skill, configure, operate, monitor, and assess each provider is high. Each new provider added requires those investments to multiply, simply to gain similar/same capabilities. For this reason, as organizations scale, they often seek a single IaaS provider. This does not mean, however, with the passage of time, they will not have to operate a second or third provider. Much like databases, organizations select a single vendor, but because of organizational restructuring, inherited technology, and sacred-cow projects, they will also have to support other vendors, but to a lesser degree. Your intention may be to later consolidate workloads from provider 2 and 3 onto your prime later, but the technology may be incompatible or requires a high degree of investment in refactoring.



In this scenario your usage of the framework agreement should align to your strategy. The organization desires a single IaaS cloud provider for most of its workloads. It is your intention to commit your organization's resources and data to this one provider. Once you are ready to commit to deploying, operating, and security assessing your landing zone within a provider, you should begin the selection of your prime IaaS provider

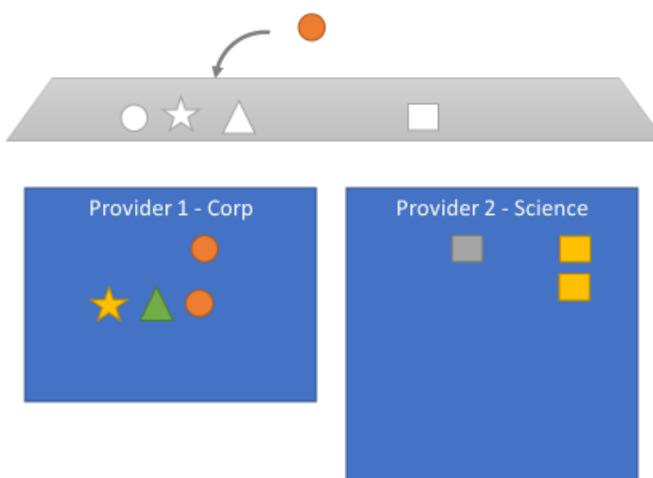
- As an organization you should use the framework agreement to solidify your selection of your prime provider early in your adoption activities.

- Your competitive process should reflect your intention is to select your single, prime, cloud provider and other cloud providers will only be used by exception.
- Your strategy, barring other influences, is to stay with this provider for the foreseeable future: decades perhaps.
- Signal that it may be your intention to shift towards an enterprise agreement with the provider later once your usage scales to the point of requiring such an agreement.

Much like selecting an organization's prime database provider, your requirements such reflect that your organization is selecting its prime IaaS provider for the long-term.

Examples 2: Corporate vs Science

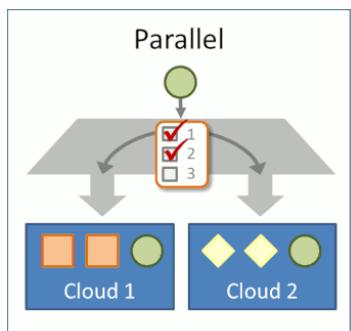
A commonly observed segmentation in the GC is science vs corporate workloads. It is not unusual to see corporate CIO shops having an affinity to one provider while scientists to another. This is logical as the nature of the work being conducted is very different. Even if both were in the same provider, often science desires a very different landing zone than corporate workloads. Scientists collaborate with colleagues outside of the GC. Their workloads are often much more ephemeral as they continually model new scenarios. Science is often much more elastic and requires more exotic forms of compute. Even when one provider is used, the configuration, skills, and monitoring activities are so different from corporate workloads, the work is the same as having two providers. Allowing scientists to use a different cloud provider from corporate workloads is a pragmatic approach for many departments and agencies.



In this type of architecture, usage of the framework agreement would require two RFPs to be issued. The first RFP's scope would focus on corporate workloads. Much like example 1. The second RFP would focus on competing for the prime IaaS provider for science-based workloads.

Like the previous example, once an organization is at the point of committing to deploying a landing zone, operations, and security assessments, they should proceed with a competitive process to establish their prime provider. In this example, it would be two competitive processes: one for corporate the other for science. The two competitive processes need not proceed at the same time.

- As an organization you should use the framework agreement to solidify your selection of your prime provider early in your adoption activities. Once for corporate, once for science.
- Your request for proposal should reflect your intention is to select your single cloud provider for corporate workloads and another for science workloads. Other cloud providers will only be used by exception.
- Your strategy, barring other influences, is to stay with these providers for the foreseeable future: decades perhaps.
- Signal that it may be your intention to shift towards an enterprise agreement with the provider at a later date once your usage scales to the point of requiring such an agreement.

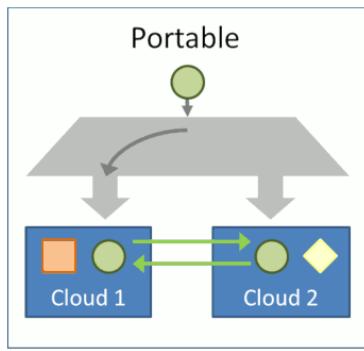


Parallel

This is currently a rare scenario in the GC. As Hohpe points out, deploying a single workload across two providers is used when a service requires an availability level that is higher than can be met with a single provider's multiple availability zones and regions.

It is unlikely a department or agencies full application portfolio would require this level of availability. It is worth pointing out that most cloud providers can provide an availability level of 99.99%. Instead, it is likely one or two highly mission critical services that may require this.

Using the SSC framework agreement, an organization would likely want to compete the services of two providers for those few services requiring higher than 99.99% availability.



Portable

This architecture is often desired when an organization seeks to prevent the selection of a single provider. Instead, the organization seeks to lock themselves in with an abstraction layer that decouples the provision, management, configuration, and operations of workloads from the cloud provider's native services. The level of lock-in is moved up from the cloud provider to the abstraction plane. In this architecture, workloads are able to move between cloud providers with no refactoring or time delay.

From a procurement perspective selection of cloud providers would follow after the selection of the abstraction plane because the cloud providers selected must be compatible with the abstraction plane selected.

- An organization should seek to select the provider of the abstraction plane.
- The organization should use the framework agreement to solidify the selection of multiple IaaS providers who are compatible with the abstraction plane selected.
- Your request for proposal should reflect your intention is to select multiple cloud providers for workloads.

Guidance

- 1) Recognize there are many multi-cloud architectures
- 2) Recognize than the GC is a multi-cloud organization, but departments may have varying degrees of multi-cloud models
- 3) While your organization's strategy may not be multi-cloud, assume multi-cloud is evitable
- 4) Align your usage of the SSC Cloud Framework Agreement to your multi-cloud strategy and desired architecture
 - a) To ensure fairness in the competitive process, signal your strategic intentions as part of the requirements.
 - b) Allow for future flexibility by focusing describing your strategic outcomes for cloud usage
 - c) Commit to long-term decisions regarding selection of providers to ensure amortization in foundational configuration, skills, and compliance over years of usage
- 5) Compete for your prime cloud provider(s) long before reaching competitive thresholds.
 - a) It is recommended that when an organization is ready to invest in deploying a landing zone, operations, monitoring, and security assessment, the organization compete its requirement and solidify selection for the long-term
- 6) While the GC is not currently contemplating enterprise agreements with cloud IaaS providers, as usage grows, this may be a logical next step. Signal your requirement to migrate to such agreements in the future.