ISG™ imagine your future®

GOING PUBLIC:

Three Key Considerations for Your Public Cloud Infrastructure Strategy

Steve Follin

The Market Reality

The race to the public cloud is now at full throttle among large enterprises. With few exceptions, enterprises are using public cloud services, including infrastructure as a service (IaaS), software as a service (SaaS) and platform as a service (PaaS). For most large enterprises, the near- and medium-term cloud model is "hybrid" or "multi-cloud," a combination of data center infrastructure services and services from one or more cloud providers.

Extending infrastructure to include cloud-based services is ultimately about the applications.

As enterprises develop their plans to integrate public cloud infrastructure services into their organizations, it's helpful to think of cloud data center services as an extension of current data center services – including identity management, security, service management, compute, storage and communications – as opposed to entirely new environments that need to be established from scratch. This paradigm helps leaders think about how best to shift workloads across an enterprise's data centers and cloud providers.

Below are three key considerations to keep in mind as you craft and update your cloud migration strategy.

1. Build your cloud strategy to accommodate large variances to planned cloud resource volumes.

Extending infrastructure to include cloud-based services is ultimately about the applications. Enterprise cloud strategy and migration timing should be driven by the benefits to the business and the application owners who often care more about improving agility and scalability than lowering costs.

An enterprise's cloud strategy, scope and timeline should be driven by an application migration and transformation roadmap developed collaboratively with the IT infrastructure team. It's important to recognize that, despite application owners' best intentions, those roadmaps can be wildly inaccurate due to many factors both within and beyond the application owners' control. Two examples illustrate this point:

- One enterprise developed an aggressive cloud migration roadmap, only to see that roadmap discarded when the CIO was promoted and replaced with someone whose priorities were refactoring applications for cloud, rather than rapid lift and shift.
- In a second example, an enterprise developed a three-year cloud migration plan across its divisions that had mixed results: one division hit its migration targets while another encountered a multi-year delay due to technical challenges in developing its new, "cloud native" applications.

Cloud roadmap delays often happen when application owners are only minimally involved in planning. A crucial component of developing a cloud migration roadmap is identifying "what's in it" from the application owners' points of view. This answer can vary widely. In some cases, there's little benefit perceived by application owners. In other cases, the application owners are anxious to move quickly to cloud hosting. Their reasons range from internal benefits, such as improved DevOps capabilities or availability of microservices, to direct customer benefits, such as better performance and scalability.

Enterprises need a "Plan B" for the possibility that cloud resource volumes do not ramp up as planned, as well as a decision point for when Plan B should be seriously considered.

It's not unusual to see cloud migration timelines slip 12 – 18 months or for migrations of legacy applications to cloud services to be cancelled altogether. In the latter case, abandoning migration may be due to a decision – based on migration cost estimates – to retire those legacy applications and replace them with cloud-based solutions. Replacements of that magnitude often take three years or more, thus reducing planned cloud server and storage volumes for a significant period of time.

Recognizing the potential variances in projected cloud resource volumes is critical when developing financial models and making operational plans for existing data centers. Enterprises need a "Plan B" for the possibility that cloud resource volumes do not ramp up as planned, as well as a decision point for when Plan B should be seriously considered. Senior management needs to be aware of the impacts of application migration delays or changes in expected volume of new applications planned for cloud hosting so it can avoid surprises and make proper resources available, including existing data center upgrades, if necessary.

2. Make stability and performance the #1 priority of cloud application migration.

The major public cloud providers offer a tantalizing array of services with tools that purport to simplify, automate, measure and report on cloud-related tasks and services. Integrating new application capabilities, re-platforming apps and refactoring applications for the cloud are the domain of the applications teams. The IT infrastructure and applications teams must work collaboratively to develop plans and timelines for implementing the cloud infrastructure capabilities needed to support the applications, as well as the operational processes required to "keep the lights on." Implementing these operational processes – such as capacity management, alerting, backup and disaster recovery – forms the platform upon which DevOps capabilities are built.

For applications designed and developed for cloud environments or "cloud-native" applications, application owners naturally collaborate with the infrastructure team as part of the development process, to identify defects and eliminate surprises in deployment. This same degree of collaboration is typical for refactoring data center applications for

cloud hosting, a process that bakes infrastructure considerations into the new application architecture and design. For example, it's not unusual for an application's database services hosted in an enterprise's data center to be replaced with a public cloud provider's PaaS database services when the application is migrated to the public cloud.

"Lifting and shifting" existing data center applications may invite a less collaborative approach since timeframes often are compressed and teams may employ automated code conversions or apply standard conversion processes. In these cases, performance and stability surprises may result. For example, communication latency issues may cause processes to time-out and fail, degrading application functionality. Defects of this type can be challenging to resolve, as they may not be easy to replicate in production or test environments.

Only after an enterprise has substantial experience moving an application or an application bundle workload to the cloud will it have the insights it needs to optimize performance across clouds.

A successful migration should quickly establish the required application functionality, stability and performance in the cloud. Implementing portability across cloud providers, for example, should not be a priority (although cloud portability should be an application design consideration for new or refactored applications). There's no way to avoid all surprises when moving data center applications to the cloud, but teams can recognize this and plan for a stabilization period. If performance falls short of meet expectations, teams should be prepared to spend the time needed to troubleshoot and remediate when stability and performance issues occur.

3. Plan to implement cloud portability in critical applications within three years of "cloud go-live."

Until recently, "cloud lock-in" had not been a significant risk for most enterprises because they did not have substantial services, particularly production services, in the cloud. But early cloud adopters are now seeing the effects. Some of these enterprises have strategic relationship agreements that were established three to five years ago and are now expiring. Their negotiation leverage has diminished since their switching costs and risks are high. Major cloud providers now have solid books of business and no longer need to offer exceptional discounts to their posted rates to build their volumes. Of course, customers continue to benefit from the decreasing resource rates characteristic of the public cloud services market. Additionally, as customers refine their cloud operations, asset management and development capabilities, they have the potential to wring out cost savings by reducing cloud sprawl, for example, or optimizing CPU and storage utilization.

These lock-in considerations have two implications:

- Cloud portability needs to be part of application and infrastructure architecture and design
- Enterprises cannot stop at stability and performance; applications and infrastructure must evolve to support application portability

While developers and infrastructure staff will be keen to implement cloud portability from the outset of migration to cloud-based services, this degree of cloud performance management is premature. Only after an enterprise has substantial experience moving an application or an application bundle workload to the cloud will it have the insights it needs to optimize performance across clouds. As long as cloud portability was envisioned at the design phase, implementation at a later date will be tractable.

Conclusion

Almost all large U.S. enterprises have committed to moving a substantial number of production applications to a public cloud. Transforming the infrastructure to support this migration can be viewed as extending the enterprise data centers to include cloud-based infrastructure resources, including servers, storage, network, firewalls and load balancers. Decision makers should keep these three points in mind:

- Infrastructure planners should assume that the actual application migration timeline may vary considerably from their plan; delays of 12 months are not unusual.
- When migrating applications, focus first on stabilizing and meeting performance requirements.
- After application stabilization, implement and operationalize portability across clouds to avoid cloud provider lock-in, and plan for portability during application design or refactoring.

ISG's Cloud Advisory team helps enterprises assess their applications for "cloud-readiness," create future-state visions and the application portfolio transformation roadmaps for achieving those visions. ISG uses proprietary and commercial tools coupled with proven methodologies and experience to assess the current state and create the future-state vision, roadmap and detailed business case/budget. Contact us to find out how we can help.

ABOUT THE AUTHORS

GOING PUBLIC:
Three Key Considerations for Your Public Cloud Infrastructure Strategy

STEVE FOLLIN

Steve has more than 25 years of experience in IT sourcing, solution design and strategy. His recent work includes helping a media company consolidate its global infrastructure and renegotiating existing infrastructure services agreements for two manufacturing clients. With Steve's guidance, both manufacturing clients were able to cut monthly costs by more than 20 percent and add cloud-based services. Time and again, Steve has demonstrated his ability to lead sourcing initiatives for infrastructure, cloud and data center services, and data center consolidations. He has experience across a wide range of industries, including the public sector.

ISG™

imagine your future®

ABOUT ISG

ISG (Information Services Group) (NASDAQ: III) is a leading global technology research and advisory firm. A trusted business partner to more than 700 clients, including 75 of the top 100 enterprises in the world, ISG is committed to helping corporations, public sector organizations, and service and technology providers achieve operational excellence and faster growth. The firm specializes in digital transformation services, including automation, cloud and data analytics; sourcing advisory; managed governance and risk services; network carrier services; technology strategy and operations design; change management; market intelligence and technology research and analysis. Founded in 2006, and based in Stamford, Conn., ISG employs more than 1,300 professionals operating in more than 20 countries—a global team known for its innovative thinking, market influence, deep industry and technology expertise, and world-class research and analytical capabilities based on the industry's most comprehensive marketplace data. For more information, visit www.isg-one.com.

Let's connect NOW...

ISG™

imagine your future®

