

Service Provisioning and Lifecycle

Content


- Private
- Hybrid
- Public
- Serverless
- Artifacts

As higher education IT budgets shrink, and we are asked to do more with the same or less, the toolkit universities use needs to evolve. Prior to the mobile web and BYOD, there was, and frankly still is, the client server model. Prior to that, there was mainframe. Universities have seen and survived many technology changes and will see many more.

Technology is an enterprise enabler. It serves to facilitate the pedagogical and research mission of the modern university. The cloud is a framework that can allow us to meet this need, and change with it much more quickly than can be done on premise or manually.

A services inventory is a required first step for this journey. From there, we will travel two parallel paths that then meet in an ideal outcome of a fully automated and scalable service. Responsive to the cyclical, burstable, at times, unpredictable nature of University needs.

Three categories of systems tend to exist on a campus: commercial applications that typically form the foundation of Enterprise Resource Planning systems, custom developed applications that address specific business needs or workflows of the university, and research applications. These types of systems can benefit greatly from cloud methods but your own campus need should be used to determine how far down the path you proceed. For some applications, the cost of refactoring the application to be automatically scalable is too great. For others, retirement and replacement is in the near future so it is not cost effective to invest the resources. It is beneficial to avoid technical debt, but don't do so at the cost of additional capabilities to support the university mission. Avoid the temptation to leverage technology, for technology's sake.



Brad Greer, University of Washington

This graph represents a spectrum. Not all applications will be in public cloud. Not all are in private cloud currently. General movement from left to right (or LL to UR) but not 100%.

<http://staff.washington.edu/brad/cloud1/>

Private

While still in the On-Premise or Private Cloud phase, one of the most important activities you can undertake is to gather an inventory. Document what services are being provided, the applications that are required to run for those services, and then any dependencies that exist. From there, break out each service and assess the critical factors associated with the service. Key elements include:

- number of dependencies of the system and on the system
- complexity of the application itself
- visibility of the service
- size, and compliance or sensitivity of the associated data set

Initially, it is best to start with a visible, but easy, win to grow confidence in the opportunities of an application running in a public cloud.

It is also important to understand the life expectancy of an application, as well as its life cycle. If a legacy application is fragile, or planned to be retired in the near future, it may not be worthwhile to move the application before it is sunset.

Some applications are particularly well suited to living in the cloud. Application that are public facing, mobile device centric, or iterated upon quite frequently thrive in the public cloud.

Hybrid

During the hybrid cloud phase, on-premise private clouds and off-premise public cloud usage at the IaaS and PaaS layers will be in use at the same time. Although some services may be redundant and provided from multiple clouds, IT infrastructure (data centers, networks, storage, and servers) will be converged, virtual, and software-defined. The speed of creation, final scale, and duration of the hybrid cloud phase will depend on an institution's individual strategic plan, current infrastructure deployments, desire to avert risks regarding cloud lock-in and unpredictable cost models.

Designing and deploying operating systems and infrastructure with a hybrid cloud strategy is beneficial in several ways:

- Whether working on the public or private portions of a Hybrid Cloud, the knowledge gained by existing IT staff on 'how clouds work' and associated skill changes will enable more rapid cloud adoption in the future.
- Encourages incremental review and optimization of standardized of existing and available IaaS and PaaS services
- Allows more concurrent projects within IT to adopt and use infrastructure that is automatable and software defined.
- Use of multiple clouds allows deeper understanding of similarities and differences, advantages and disadvantages of cloud offerings.
- Provides an environment to test system migrations between public and private clouds. This will inform future decisions regarding the risks of vendor lock-in, security options, and cost variability.

During the hybrid cloud phase, growing and maturing the use of a DevOps methodology is a significant opportunity. By creating project teams that include both developers and operators, cloud-oriented designs and tools will be learned in parallel throughout the IT organization and will enhance overall strategic organizational alignment. For projects that work on private cloud or public cloud, DevOps will positively influence roles to change from a "break-fix" operations towards roles focused on service delivery and configuration management.

Currently owned IT infrastructure will be replaced, refreshed, or phased out over a variety of planned lifecycles and It may be planned that some IT infrastructure will continue to be on-premise for many years. Given this, it is very important to continuously enable the hybrid cloud ecosystem at every opportunity for infrastructure refresh or update. One strategy to consider would be to require all infrastructure upgrades and replacements be evaluated for suitability in the hybrid cloud model. On-premise storage system replacement is a common example By replacing an existing SAN or NAS based storage system with one that provides hosted block storage, this infrastructure will have a higher utilization and effectiveness for its entire lifecycle.

Public

The third phase of service evolution focuses on fully utilizing the public cloud where the majority of an institution's workload is supported in off-premise environments. The institution is transitioning from purchasing technology infrastructure for anticipated load to a utility model for consuming a technology function. The technology functions are being served by highly scalable, external providers leveraging shared infrastructure to serve multiple clients across numerous industries. Infrastructure is now being uniformly treated as code and is provisioned by developers and testers on demand rather than system administrators. Services are being automatically deployed, managed, and scaled. As an institution migrates to a public-cloud model, it needs to account for its people, process, and technology strategy going forward.

Adopting a public cloud model introduces a distinct shift in IT staffing strategies for Universities. On-premise service models focus on hiring roles (e.g., system administrators, data center operators, developers) to support the full technology stack. The hybrid cloud model starts to emphasize new roles (e.g., vendor manager, systems integrator, architect, business analyst) in the organization and provide opportunities for interested staff to develop new knowledge, skills, and abilities. The public cloud model is dominated by the new IT roles and now requires existing staff to change within the organization. This required change can naturally be met with apprehension. During this time period, it is critical to (1) communicate the vision and cloud strategy, (2) focus on career planning and development for IT staff, and (3) minimize the impact to existing operations. Ultimately, a public cloud migration is about refocusing your staff from commoditized operations to value added activities.

The public cloud model thrives on process standardization. The more standardized process an institution can adopt; the greater benefit a public cloud will yield for an institution. A variety of service models are offered to a University for achieving its administrative needs. Four primary models -- each with an increasing order of standardization magnitude -- can be consumed by an institution:

1. Infrastructure as a Service (IaaS): customizable application on standardized system infrastructure;
2. Platform as a Service (PaaS): customizable application on standardized application and system infrastructure;
3. Software as a Service (SaaS): configurable application on standardized application and system infrastructure; and
4. Business Process as a Service (BPaaS): automated business processes on standardized SaaS, PaaS, and IaaS

Most commonly, a University will deploy a combination of the above service models to accomplish its business processes. The greatest challenge with pushing toward a fully standardized model is the change threshold among an institution's students, faculty, and staff.

Finally, an institution needs to address its technology choices and migration strategy to the public cloud. The previous hybrid model is a natural first step, so most institutions will already have some capacity of workloads running in a public cloud. Like most industries, an organization will start with deploying its new and non-critical workloads, move onto migrating its larger workloads, and finally, migrate its core workloads. This typical migration path generally results in the applications with the most sensitive data and workload being left for last. Before getting to this point, an institution needs to develop a holistic vendor, migration, and security strategy. For example, selecting single or multiple vendors; refactoring, replacing, or rehosting applications; adopting new business processes; determining the exit strategy and cost for vendors; and addressing privacy and security concerns with security reviews, appropriate terms and conditions, and Business Associate Agreements (BAAs). For Infrastructure as a Service, one should consider a design strategy that uses “containers” instead of “servers”. Infrastructure built-on containers is simpler to manage, allows you to manage infrastructure more like an application, and can be more easily moved between clouds (public and private).

Serverless

As we move to our last foreseeable phase of service evolution, we cycle back to a strongly-provisioned client-server model. At this point, endpoints are robust enough to run any required application logic or user interface element. The traditional, three-tier web application (database, application, web) has been dissolved into a data store hosted database, and a highly functional mobile application that interacts with the data store hosted database using a set of structured APIs.

It is important to identify applications that are a good fit for this model. In the near term, student-centric applications and ones used by highly mobile staff, custodial and maintenance professionals for instance, are a good fit. Over time, we will likely see a greater emphasis placed on mobile capabilities, with mobile applications moving from a nice to have supplementary capability to the primary way information is accessed. This is likely to include access from campus WiFi networks, as well as common LTE networks.

Campuses effectively operate as a small city would, with many of the prerequisites required for an Internet of Things ecosystem to exist. Connected sensors - one of these key components - currently generate data points on a wide array of elements on campus such as building environmental controls, landscaping systems, and laundry machines. These are just some examples of data points that the community could collect and make accessible using a database and APIs. Aggregating access to these elements makes it easier to drive insights for these data point as well as using them to inform how we can better serve our students and communities.

The expectations of our students, faculty, and staff are continuously evolving. We should be providing capabilities to enable them to meet their own needs and empower the campus community to serve their individual constituencies. Our incoming freshmen buy the latest devices available a mere number of weeks before they arrive on campus.

Beyond mobile first and backends-as-a-service lays an unknown, but quickly approaching future state. As the pace of IT capabilities quicken, we will continue to see new IT paradigms be developed and arrive to disrupt existing systems. This change is not new, from mainframe to client-server, then to virtualization, and now to cloud, IT has always been up to the challenge.

Artifacts

- [Indiana University - Sharing Institutional Data with Third Parties](#)
- [CMU Considerations for Implementing Software Solutions](#)
- [Harvard IT DevOps Tools](#)