



CICompass

Best Practices for Cloud Provider Analysis

APRIL 25, 2022

EWA DEELMAN, JAREK NABRZYSKI, KARAN VAHI, RICK BENSON

Table of Contents

How to choose a cloud provider for a scientific project?	1
Cloud Deployment Models	2
Cloud Computing Infrastructure Requirements	3
Cloud Adoption Risks and Mitigation Strategies	4
<i>Security Risks</i>	4
<i>Performance Risks</i>	4
<i>Availability Risks</i>	5
<i>Operational and Governance/Cost/Control Risks</i>	5
<i>Vendor Lock-in Risks</i>	6
<i>Cloud Migration Strategy Related Risks</i>	6
Summary	8
Acknowledgements	8

This project is supported by the National Science Foundation Office of Advanced Cyberinfrastructure in the Directorate for Computer Information Science under Grant #2127548. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the funding agency.

How to choose a cloud provider for a scientific project?



While public clouds offer the promise of greater efficiency and other benefits to scientific projects, major concerns still arise regarding various risks associated with a full migration to the cloud. There is an entire commerce around managing deployments within cloud environments, suggesting it's not as do-it-yourself as on-premises management.

It is important that scientific projects consider all pros and cons along with risks associated with moving to the cloud. A rushed migration without a clear strategy can end up costing the project more than necessary, with existing legacy applications consuming cloud resources and generating costs at an alarming rate.

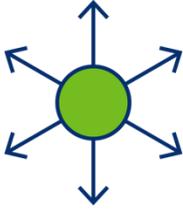
“One of the major issues for our future cloud platform is determining where it can be deployed in a cost-effective and sufficiently capable environment”

Chad Trabant, project manager for Cloud Computing Platform in Seismological Facilities for the Advancement of Geoscience (SAGE) and Geodetic Facility for the Advancement of Geoscience (GAGE).

In this white paper, we present major criteria that should be considered while choosing a cloud provider for any scientific project. We focus on benefits and risks associated with moving to the cloud.



Cloud Deployment Models



There are many cloud deployment models to choose from: IaaS (Infrastructure as a Service), SaaS (Software as a Service), PaaS (Platform as a Service), and CaaS (Container as a Service like managed Kubernetes or other container services).

In an IaaS model, a cloud service provider (CSP) hosts the hardware components originally present in an on-premises data center, including servers, storage, and networking hardware, as well as the virtualization or hypervisor layer. The IaaS provider also provides various services to deliver the infrastructure components.

The SaaS model eliminates the need to install and run applications on the project servers or in their own data centers. This eliminates the expense of hardware procurement, provisioning, and maintenance, as well as software licensing, installation, and support. Generally, a project would pay for this service on a monthly/hourly basis using a pay-as-you-go model.

The third model is a PaaS model, which is suitable for organizations that need support for key services, such as application hosting or Java development. A PaaS provider forms and supplies a strong and optimized environment on which users can install applications and datasets. Organizations using the PaaS model need to focus on deploying and running applications rather than on designing and maintaining the underlying infrastructure. Many PaaS products are geared towards collaborative software development. Such products often include computing and storage infrastructure, as well as versioning systems, compilers, and SDKs.

The CaaS model has evolved as a compromise between IaaS and PaaS. It abstracts the full stack in a very compelling way, without the many challenging surprises existing in both IaaS and PaaS. CaaS already is emerging as the better go-forward model for projects looking into long-term investments in the cloud. CaaS has such advantages as:

- CaaS deploys quickly and lightly on almost any infrastructure;
- CaaS provides commodified, standardized functionality on-premise and/or on public clouds;
- CaaS offers open container technology, which is de-facto standard in the cloud industry;
- CaaS offers unprecedented freedom and flexibility to developers.



Cloud Computing Infrastructure



Any project that migrates to the cloud needs to consider future cloud computing infrastructure requirements. The considerations should include VM infrastructure needs, relational database servers infrastructure, performance, data volumes including backups, network services and inter-VM bandwidth, identity management services, potential container orchestration, monthly/annual volume of egress data, cybersecurity services, and more. Each cloud provider offers different infrastructure services and it is usually hard to navigate through these offerings. The following baseline comparison of the major cloud providers:

Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform (GCP), might help with initial decisions.

	Pros	Cautions
AWS	Most reliable, most fully-featured, the most number of services. Deep developer knowledge base and rapid pace of new feature releases. Dominates the cloud space with features such as configuration, monitoring, security, auto-scaling. More experienced, enterprise-friendly services. Good open-source tools integration. Global reach. AWS has 31% of the worldwide cloud market.	A number of options and the need for developer learning time can be overwhelming. More rigid pricing model. Weak hybrid cloud solution. In early 2019, a Greenpeace report accused AWS of abandoning its commitment to 100 percent renewable energy. AWS has also made a habit of keeping information about its carbon footprint out of public view.
Azure	Best leverage of Windows technology. Proven track record. Most user-friendly dashboard. Better development and testing tools. Also provides a Hybrid cloud. World's second-largest cloud provider and it grew 48% year-over-year in 2020. Azure pretty efficiently uses green energy as well as provides pretty transparent information about its energy sources.	Problematic if you're not running Windows technology exclusively. Support for other operating systems is rather limited, and support only with selected varieties of Linux. Less reliable – more frequent outages and bugs. Lacking in the DevOps area.
GCP	The best option for containers and machine learning. Rapid release of new features. Often the most competitive on cost. Expertise in DevOps, flexible discounts, and contracts, specifically designed for cloud-based businesses. Google has 7% of the global market and is growing faster than other providers in the database as a service and IaaS. GCP makes it up with stronger support for container and Kubernetes use cases. It is easier to learn compared to AWS and Azure; Google cloud storage options are especially easy to learn and use. Better pricing than AWS and Azure.	Fewer bells and whistles (especially with migration services). Less robust offerings outside containers and machine learning. Not as many features as AWS or Azure, but this might be considered an advantage (lower workforce development costs).

Scientific projects should also check out the offerings from academic providers, such as San Diego Supercomputing Center, Texas Advanced Computing Center, Open Storage Network or JetStream.

It is also worth discussing your project needs with [Internet2](#) and/or [CloudBank](#), both offering some cloud support to academic institutions.

Cloud Adoption Risks and Mitigation Strategies



Despite the many benefits that cloud computing brings, there are also risks associated with migrating to the cloud, and sometimes these risks can outweigh the expected benefits. In every potential cloud migration project, and especially in taxpayer publicly funded projects, we need to carefully analyze whether there is any potential for hidden and migration cost-related undisclosed incidents or other risks related to, e.g., service continuity. It is important to ask whether a cloud is a sustainable or non-sustainable solution for a project.

Cloud adoption-related risks can be grouped in several categories, such as security, performance, availability, operational impacts (including governance and cost controllability), compliance, vendor lock-in, and user experience and workforce development implications. In this section, we discuss the most relevant risks and provide recommendations for risk mitigation. In addition, we can cluster various cloud migration risks into three major categories: mission-oriented risks, cloud service provider risks, and risks associated with different cloud adoption strategies. When moving to the cloud project consortia need to develop a new risk management profile and strategy.

Security Risks



IaaS and CaaS models provide virtualized computing resources over the Internet hosted by a third party. The security concerns are similar to those of running a private data center, but since a project is using virtualized resources that technically belong to somebody else, weaknesses in the owner's security controls can affect a project dramatically, especially when insufficient due diligence was done a priori. There is a huge risk with a lack of understanding of the various security groups and roles provided by the commercial cloud providers, which may lead to serious hacks. Running a secure cloud infrastructure will require strong cybersecurity skills on the project team.

Performance Risks



Most IaaS and CaaS infrastructures are multi-tenant, meaning the physical resources are being shared among potentially many cloud provider users' virtual infrastructures. Many providers also offer multi-tenant VMs. For performance-sensitive applications, such as contracted early warning monitoring, such shared environments might be a source of problems. Bare metal cloud resources might be a better choice for such applications. But if the applications/workloads are highly dynamic, i.e., applications spin up and down rapidly, or if some workloads must be spun up in minutes, run for some hours, and after that can be turned off dynamically, virtualized infrastructures are a great choice. The big question is "what type/size of VMs do we need for particular tasks?" It is important to do some evaluation

and performance requirements analysis at scale to answer this question. Projects need a thorough understanding of their application performance requirements, and we suggest a cloud-based testbed/pilot should be designed using potentially free allocations from various cloud providers to analyze the performance of critical workloads in different cloud infrastructures. The choice will have a great impact on both the annual cost and the performance of the solution. This risk can also be mitigated through proper SLA agreements with cloud providers.

Availability Risks



Cloud disasters happen. We have witnessed big cloud providers facing disastrous moments. It is hard to predict when a disaster in the cloud will occur and how serious its impact will be. No matter whether on-premise or in the cloud, all we can do is have a precisely designed plan of how we will respond to and recover from such disasters. The [disaster recovery template](#) may help in the planning process. Big questions to answer are: “Can we afford minutes/hours/days/weeks of data unavailability?”; “What risks are we able to take?”; “What would be the consequences for the project if the data is not available for a number of hours/days/etc.”; “If we move to the cloud, where are our disaster recovery services located?” These are the business decisions that will affect both the cost of the projects and the operational practices. One way to provide disaster recovery services is to have a separate infrastructure on-premises (in a separate data center) or in a different cloud. Managing and monitoring a separate disaster recovery data center might be very difficult and expensive. Cloud-based disaster recovery, on the other hand, can help projects deal with the most common problems of running a separate disaster recovery data center. Potential benefits include but are not limited to: close to zero capital cost; scalability; and minimal time of standing up and provisioning a backup infrastructure and data (potentially measured in minutes).

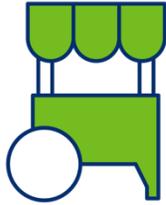
Operational and Governance/Cost/Control Risks



Governance and the ability to control the cost might become significant challenges and risks related to cloud computing. Understanding the exact cost of outsourcing infrastructure or service and being able to perform very accurate cost analysis in order to decide whether cloud migration would be economically worthy for a project is very important, but at the same time, due to the nature of the cloud, it is an almost impossible task to accomplish. In our opinion, the ability to predict and control the cost of cloud deployment and operations is the highest risk for projects.

As in every project, IT governance should safeguard the proper implementation and use of all cloud assets. It should be of the highest priority for projects to set appropriate policies and procedures for staff managing the infrastructure as well as for its users. One of the key issues in today’s cloud infrastructures is to have full control over the provisioning, de-provisioning, and operations of the infrastructure to control both the cost of the cloud infrastructure (computer, storage, and data transfer) and data quality. IT teams must adapt their traditional IT governance and control processes to include the cloud. There are also several other ways to keep cloud costs under control, such as monitoring utilization and right-size instances, automating shutdown of temporary workload instances, shutdown workloads during certain hours, or move workloads to cheaper cloud/region.

Vendor Lock-in Risks



Cloud vendor lock-in refers to a situation when the cost of switching to a different provider is so high that the customer is essentially stuck with the original vendor. Such factors as workforce skills, an architecture tailored to a particular vendor's low-level services, or the need to avoid interruptions might contribute to the potential cost of moving to another vendor. Projects need to analyze all the risks associated with their architecture and the vendor, especially past vendor performance, quality of services, and technology roadmap as well as vendors' past technology changes, whether there are any incompatibilities with previous versions of services, vendors' financial standing, and any dramatic price increase for their services in the past. Proof-of-concept deployment may help mitigate risks associated with provider lock-in. Also, it is important to have strategies for moving the project data and services to another cloud provider or to on-premise infrastructure. Hybrid and private cloud deployments may also dramatically reduce the vendor lock-in risk.

Cloud Mitigation Strategy Related Risk



Cloud migration is the process of moving some or all organization's digital operations to a cloud. There are three main types of cloud migration an organization can perform — on-premises to cloud, cloud to cloud, or cloud to on-premises. When performing any of these three migration types, there are five methods and strategies an organization can use:

1. **Lift and shift** – moving applications to the cloud as-is. This is also sometimes referred to as rehosting. The “lift and shift” approach involves moving the application into the cloud as-is and hoping it works. This is actually possible for many applications, and cloud vendors make it as easy as possible by providing various flavors of operating systems.

The advantages of this approach include:

- No code or architecture changes.
- Easy migration of services.
- Easy compliance and security management.

The disadvantages include:

- Legacy applications are usually not scalable and do not allow for distributed workloads as cloud-native applications do.
- On-premise applications might suffer from latency or performance issues after migration since they were not optimized for the cloud.
- No application is bug-free. After moving an app to the cloud, more risks and problems may occur.

2. **Refactor** – modifying applications to better support the cloud environment. The applications are re-architected to fit the cloud better. This strategy includes modifying the existing software, usually a large chunk of it. The process requires lots of software engineering effort and thus is much more complex and time-consuming than other migration strategies.

Advantages of this strategy include:

- Long-term costs can be potentially reduced as refactoring allows you to achieve a better elasticity of applications, e.g., scaling as needed to reduce resource consumption.
- Cloud-native and microservices architectures allow apps to rapidly adapt to new requirements by adding or modifying functionality in a fairly inexpensive way.
- Native, decoupled application components provide high availability of services.

The disadvantages include:

- Vendor lock-in: The more cloud-native the application is, the more coupled to low-level cloud services and high-level APIs the application is. Moving to another cloud will require another refactoring.
- The high upfront cost of re-factoring, involving very advanced, cloud-savvy software engineers.
- In software refactoring, the more we change the application, the more risks of introducing new bugs into the code. New bugs can delay the migration and can cause outages once the applications are deployed in the cloud.

3. **Re-platform** – moving applications to the cloud without major changes but taking advantage of the benefits of the cloud environment. This strategy lies somewhere between the aforementioned strategies. Common modifications include, e.g., the ways apps interact with databases to benefit from automation and elastic database infrastructure, or enabling better scaling and leveraging reserved resources in the cloud environment with minimal code changes. It seems this strategy is very relevant to many projects.

Advantages of re-platforming include:

- Migration can be done in steps. You can move some workloads to the cloud-first, experiment with the environment, learn lessons, and move on to other workloads.
- It does not require too much resource/time/cost compared to refactoring.
- Many cloud features can easily be used, e.g., auto-scaling, managed storage, managed data processing services, among others.

The disadvantages of re-platforming include:

- To see potential benefits from re-platforming, a high level of workload automation is required, i.e., time from skilled cloud engineers is required to develop basic automation of workloads.
- It is important to design the process in an optimal way and manage the scope of the work to be done to re-platform your applications. Failing to do so may turn the re-platforming project into a full refactoring project generating unexpected costs. This often means that

only the most common cloud components are being utilized while re-platforming. Specialized components often require more dramatic changes to your application. Organizations need to be driven by the business value of all re-platforming decisions.

4. **Rebuild** – rewrite the application from scratch. This strategy is not relevant/recommended for projects that have many large scale legacy applications.
5. **Replace** – retire the application and replace it with a new cloud-native application. This strategy is not relevant/recommended for projects that have many large scale legacy applications.

Summary

In this white paper, we focused on both benefits and risks of cloud migration for any scientific collaboration project. It is important that such projects consider all pros and cons along with risks associated with moving to the cloud. A rushed migration without a clear strategy can end up costing the project more than necessary, with existing legacy applications consuming cloud resources and generating costs at an alarming rate. There is no one-size-fits-all approach. It is important to start with defining the business value that cloud technology transformation can enable. The project needs to identify where the impact will be derived quickly and safely, whether through cloud capabilities, cost efficiencies, or risk mitigation. **Big questions are: in this new environment, how fundamentally will operations and workflows change; what skills will be needed to manage project services on the cloud; how the project's organizational structure will have to be adapted; and what impact on both the project teams' and users' behaviors will the cloud migration have?**

Glossary

API	Application programming interface
AWS	Amazon Web Services
CaaS	Container as a Service
CSP	Cloud Service Provider
GCP	Google Cloud Platform
IaaS	Infrastructure as a Service
IT	Information technology
PaaS	Platform as a Service
SaaS	Software as a Service
SDK	Software development kit
SLA	Service-level agreement
VM	Virtual machine

Acknowledgements

This white paper is delivered by the NSF Cyberinfrastructure Center of Excellence CI Compass. CI Compass provides expertise and active support to cyberinfrastructure practitioners at NSF Major Facilities in order to accelerate the data lifecycle and ensure the integrity and effectiveness of the cyberinfrastructure upon which research and discovery depend. To contact us please send an email to: contact@ci-compass.org.

CI-Compass would like to acknowledge several contributors from the [National Science Foundation](#) funded [GAGE](#) and [SAGE](#) facilities. Our special appreciation goes to Rick Benson¹, Chad Trabant, David Mencin, Jerry Carter, and other IRIS and UNAVCO facility staff.

¹ Rick Benson passed away to a tragic accident on February 6, 2022.