

# Cloud Myths and Realities

**Version 1.1**

**A Discussion Paper from the OMG Cloud Working Group**

**September 2023**

**Document mars/2023-09-05**

This discussion paper presents guidance that has been carefully crafted by industry and technology experts from OMG's Cloud Working Group, and approved for publication by OMG's Middleware and Related Services (MARS) Platform Task Force. While the paper has not been endorsed by OMG's Board of Directors and is not a standard, readers are encouraged to follow this guidance, participate in wider discussion on this topic, and help OMG evolve it by participating in future work of the Cloud Working Group.

# Table of Contents

- Copyright Notice ..... 3
- Acknowledgements..... 3
- Overview, Background, and Rationale..... 4
- The Myths ..... 5
  - Myths About Scope and Applicability ..... 6
  - Myths about Costs ..... 8
  - The Infinite Resource Myth..... 9
  - Myths About Exiting the Cloud ..... 10
  - Myths About Security ..... 11
  - Myth About Availability ..... 13
  - Myths About Management..... 14
- Conclusion and Recommendations ..... 15
- Appendix A: ChatGPT’s Answer to the Myths Question..... 16
- Appendix B: Stories from the Trenches ..... 19
- Appendix C: References ..... 20

## Copyright Notice

Copyright © 2023 Object Management Group. All rights reserved. You may download, store, display on your computer, view, print, and link to the *Cloud Myths and Realities* discussion paper at the OMG Cloud Working Group Web site subject to the following: (a) the Guidance may be used solely for your personal, informational, non-commercial use; (b) the Guidance may not be modified or altered in any way; (c) the Guidance may not be redistributed; and (d) the trademark, copyright or other notices may not be removed. You may quote portions of the Guidance as permitted by the Fair Use provisions of the United States Copyright Act, provided that you attribute the portions to the *OMG Cloud Myths and Realities Discussion Paper v1.1 (2023)*.

## Acknowledgements

Development of the *Cloud Myths and Realities* paper was a collaborative effort that brought together diverse customer-focused experiences and perspectives into a single guide for cloud customers.

The paper's main authors were:

- Troy Anderson – independent consultant
- Claude Baudoin – cébé IT & Knowledge Management (editor)

The following participants also provided their knowledge and time to this effort through suggestions and comments:

- Jeff Boleman – IBM
- Conrad Kimball – Boeing
- Marios Panagakis
- Lisa Schenkewitz – IBM
- Prasad Siddabathuni – EDIFecs

## Revision History

Version 1.0	Approved June 2023	Original version
Version 1.1	Approved September 2023	Includes an additional myth suggested by Bruce Speyer (Texas Department of Information Resources)

## Overview, Background, and Rationale

Since the cloud computing model became a mainstream phenomenon, about fifteen years ago, users have been regularly submitted to a barrage of confusing claims and counterclaims. On the one hand, the providers of cloud services have naturally been keen to convince prospective customers that everything would be simple and cheap if they adopted their solutions; their motivation was clear, but their message lacked any subtlety and conveniently omitted many qualifications and caveats. On the other hand, many people have been spreading FUD (fear, uncertainty, and doubt): suppliers of traditional hardware and software solutions, including hosting service providers, wanted to preserve their market; analysts and consultants make more money if the customers feel that the decision to move to the cloud is risky; some executives are worried – often justifiably – about regulatory and security issues; and IT managers want to find a reason to maintain the size of their departments and of their shiny computer rooms, which are the visible signs of their importance.

Ever since the Cloud Standards Customer Council (CSCC), the predecessor to the OMG Cloud Working Group, was formed in 2011, our mission has been to be the voice of the cloud customer, and we have carried it out by publishing many guides and papers that explain to cloud customers (actual or potential) how to evaluate, select, implement, and govern cloud solutions deliberately and safely. Our flagship publication, the *Practical Guide to Cloud Computing* [1], now in its third edition, should in theory allow its readers to pursue a rational path to success without falling prey to the exuberant promises of the vendors or the denigrations of the naysayers.

Many websites and publications address the myths of cloud computing. If you ask Google for “cloud myths and realities,” it tells you there are about 18 million results, which include titles such as “the top N myths of cloud computing” where N is just about any number between 5 and 20. Some of these sources are useful (see references [2], [3], [4], [5]), but many are derivative, repetitive, simplistic, or sometimes poorly reasoned (and even less well written). In addition, some authors have a hidden agenda: to avoid this issue, buy *our* solution, or hire *us* to advise you.

As a result, it seemed useful to collect, analyze, and in many cases debunk some of the myths – positive or negative – using the collective experience of the members of OMG’s Cloud Working Group and the critical eye of members of OMG’s Middleware and Related Services (MARS) Task Force, which vets the working group’s papers.

The paper covers such hotly debated topics such as cost, security, management, and more. For each of the myths we have identified, we present three elements:

- The myth itself – i.e., a statement that is often made by proponents or detractors of cloud computing, and that many customers may believe, but that we assert are either false or greatly exaggerated.
- The reality – what we think is true about the subject of the myth, which may not be the exact opposite of the myth but is at least significantly different. The reality is also often more complex than the myth, which is usually based on oversimplification of the issues.
- A discussion of why we assert that the first statement is a myth and the second is correct, sometimes with suggestions for further reading.

The authors hope that this paper will be a useful companion to our *Practical Guide to Cloud Computing* and other papers, and will help cloud customers see through all the misinformation that unfortunately abounds in this domain.

## The Myths

We describe fourteen of the most common myths in this discussion paper. This is not an exhaustive list.

### **Myths About Scope and Applicability**

- Everything should/will move to the cloud and there is no reason to move back
- Cloud solutions are only for startups and small companies that do not have an IT department
- Cloud computing is the same as Software as a Service
- Cloud services and hosted services are the same thing

### **Myths About Costs**

- Cloud solutions are (always) cheaper than on-premises solutions, and this is the main reason to move to the cloud

### **The Infinite Resources Myth**

- The cloud will automatically scale up and deploy whatever resources we need, essentially offering infinite resources to the customer

### **Myths About Exiting the Cloud**

- Migrating out of the cloud is as straightforward as it was to migrate into the cloud

### **Myths About Security**

- The cloud is less secure than your own data center
- It does not matter where a cloud service is hosted
- Managing encryption is easy thanks to the providers' BYOK (Bring Your Own Key) offering

### **Myth About Availability**

- There are more outages in the cloud than in your data center

### **Myths About Management**

- Once you move to the cloud, the Cloud Service Provider does everything for you
- You can apply your existing policies and procedures to manage the cloud environment consistently with your on-premises environment
- Disaster recovery is built in to cloud services

## Myths About Scope and Applicability

<b>Myth 1</b>	<b>Everything should/will move to the cloud and there is no reason to move back</b>
<b>Reality</b>	There are valid reasons to retain certain systems and solutions in-house, or to insource solutions after they have been in the cloud for a while.
<b>Discussion</b>	<p>A cloud solution may have been initially adopted in order to deploy a needed capability very quickly or on a small scale. Once the solution is stable and its usage grows, the relative costs of an on-premises vs. a cloud solution may change, and the organization has a better sense of what they need.</p> <p>Another situation in which a customer may not move to the cloud is when it has developed a proprietary system that gives it a competitive advantage, therefore there is no equivalent Software as a Service (SaaS) offering, and there is little technical or financial incentive to move this code onto an Infrastructure as Service (IaaS) or Platform as a Service (PaaS) solution.</p> <p>In some cases, the size of the data makes it virtually impossible to use a cloud service. This would be the case of large scientific datasets, such particle accelerator data, seismic studies in the oil and gas industry, meteorological data, etc.</p> <p>Finally, while there are myths around security (to be discussed later), there may be regulatory constraints that prevent the customer from moving to the cloud. This can be the case for financial or medical systems in smaller countries where there is not a suitable domestic cloud provider, and the law prohibits storing the data outside of the country. See OMG's paper on <i>Data Residency Challenges</i> [6].</p>

But now that we have debunked the myth that everyone and everything should move to the cloud, beware of an opposite myth:

<b>Myth 2</b>	<b>Cloud solutions are only for startups and small companies that do not have an IT department</b>
<b>Reality</b>	Cloud solutions may be the obvious choice for an organization to avoid building an in-house infrastructure but they can be (and are) used by companies of any size.
<b>Discussion</b>	<p>For a start-up company, cloud solutions are an obvious way to create a small but complete IT environment quickly. Typically, this involves collaboration tools (e-mail, calendar, messaging, teleconferencing, project management), document storage, and several key applications such as ERP (or at least accounting) and CRM. The costs can be low at first if the company has very few employees, and usage can grow progressively as the business expands instead of incurring a high initial investment.</p> <p>A medium-sized business that has already licensed and deployed its own internal applications may choose to migrate to the cloud if it outgrows its on-premises resources. A company that expands to multiple locations, and therefore would need to set up a private cloud, may find it more logical to move to the public cloud at that point.</p>

Another scenario is when an organization deploys new online services to be used by consumers. Since those consumers will access the services over the Internet, there is no benefit to having the servers and the software in-house, and a cloud provider's infrastructure (including its large typically large Internet bandwidth) is better able to cope with the escalating or irregular consumer traffic.

Even a large business with existing technical infrastructure may want to leverage the cloud to dynamically provision and relinquish development environments (using a platform-as-a-service model).

The recent popularity of remote working arrangements means that employees regularly access their own company's systems over the Internet. This negates the advantage of hosting many applications internally.

Modernization is another reason for large organizations to move to the cloud. If an older application needs to be replaced, cloud solutions should be evaluated against the option of migrating to a new in-house solution. One of the advantages of the cloud is that in the future, software and hardware upgrades will be performed by the cloud provider, typically much faster than an end user would do.

Finally, cloud-only applications are starting to emerge, especially for complex or cloud-intensive capabilities such as AI and big data.

In addition, we encounter some confusion, if not strictly speaking a myth, about the scope of services that are delivered in a cloud environment.

**Myth 3**

**Cloud computing is the same as Software as a Service (SaaS)**

**Reality**

Cloud computing encompasses several deployment models (private, public, hybrid, community) and several service models (infrastructure as a service or IaaS, platform as a service or PaaS, and infrastructure as a service or IaaS). SaaS is therefore just one flavor of cloud computing.

**Discussion**

Many observers of the cloud phenomenon attribute its inception to the launch of the Salesforce CRM application in 1999, even though the notion of "utility computing" predates this by decades. NIST proposed a widely recognized reference model [7] in 2011, in which the three service models mentioned above (IaaS, PaaS, SaaS) are defined. Each of these models is in use today and created a significant market. Therefore, equating the cloud with SaaS is a mistake.

In addition, new "X as a service" terms (often dubbed "XaaS") have emerged. Some are used to designate offerings that do not fit neatly into the IaaS/PaaS/SaaS boxes, while others are marketing terms for services that are usually a specialization of SaaS to a given industry. This is discussed in detail in OMG's *XaaS Glossary* [8].

Finally, in this category, we need to dispel another myth about the nature of cloud offerings:

<b>Myth 4</b>	<b>Cloud services and hosted services are the same thing (sometimes phrased as “a CSP and an MSP (Managed Service Provider) are the same”)</b>
<b>Reality</b>	Cloud computing is not characterized by where and how the service is hosted, but by such properties as access from anywhere, rapid elasticity, metered usage (pay as you go), and resource self-provisioning. Traditional outsourcing (hosting services) lacks most of those.
<b>Discussion</b>	<p>The wave of data center outsourcing that occurred in the 1990s and early 2000s can only be confused with cloud computing if you incorrectly believe that all the cloud does is to move your hardware and applications out of your data center.</p> <p>In reality, traditional hosting services offered by MSPs lack several or all of the defining characteristics of the cloud:</p> <ul style="list-style-type: none"><li>• The outsourced resources are connected back to users in the enterprise, but not necessarily through the Internet. This may provide faster connections, but only from the customer’s offices.</li><li>• The customer usually pays a fixed monthly fee, regardless of usage.</li><li>• The amount of available resources (number of servers, storage space, number of application licenses) is fixed for a certain duration and can only be changed by renegotiating the contract.</li><li>• Therefore, there is no facility for self-provisioning of resources.</li></ul> <p>The other reason why this is a myth is that cloud services are not necessarily hosted externally. In a private cloud, the corporate IT department offers resources to units, business lines, branch offices, stores, etc., with the same properties mentioned above (network access, elasticity, metered usage, self-provisioning).</p>

## Myths about Costs

<b>Myth 5</b>	<b>Cloud solutions are (always) cheaper than on-premises solutions, and this is the main reason to move to the cloud</b>
<b>Reality</b>	Cloud solutions may or may not be cheaper depending on multiple factors. The main financial impact is the shift from capital expenses to operational expenses (CAPEX to OPEX), and there are more important justifications to move to the cloud.
<b>Discussion</b>	<p>The main argument behind the myth is that in the cloud, you only pay for what you use (the “utility pricing” model). If your usage decreases for any reason, you can downscale your service (e.g., release some servers or reduce the number of user accounts) and reduce your bill. If you decide after one year that the software, hardware, or platform you chose was the wrong one, or if you must switch systems following an acquisition or merger, you are not burdened by the remaining depreciation of something you paid full price for but no longer need.</p> <p>Another cost reduction factor is that an on-premises solution must be sized for the expected maximum resource usage over a certain time horizon, whereas a cloud</p>



solution can be sized according to the initial demand and scaled up later as needed, or can accommodate short periods of excess usage for a price.

If, on the other hand, you have made the right choice and use the same solution for more than (typically) two or three years, the accumulated monthly payments for the solution may well exceed what you would have paid for an outright purchase, whether the system is installed on-premises or at a traditional hosting service provider. But you will only know after the fact that the cost curves have crossed and that your total cost of ownership (TCO) has become greater in the cloud. You may also realize that some TCO studies are biased in favor of cloud solutions: they may include overinflated data center savings, operational savings that never become realized, or neglect to factor in all costs such as software licensing.

An additional cost of moving to the cloud has to do with the cost of the transition itself, if the cloud solution is replacing a previous system (especially a large and complex one, such as an ERP package). This includes all the labor involved in analyzing the need, selecting the solution, preparing the data, performing test migrations, training the users, and the cost of operating two systems in parallel until the migration is complete.

Your CFO may love the utility pricing model even if the eventual cost is higher, because you did not have to put up the cash up front. Whether this move from CAPEX to OPEX is advantageous is a necessary discussion between the CIO and the CFO.

Because it is uncertain whether a cloud solution will be cheaper or not in the end, it is wrong and dangerous to use this as the major argument to “sell” the cloud to executives. There are many other advantages of the cloud that can justify such a move, especially agility – the speed of deployment of a solution in response to new or changing business needs. Selecting and procuring the hardware and software to deploy an ERP or a CRM system in-house typically takes several months; while a rash decision should not be made for a cloud solution either (this could be another myth in itself), it is possible to subscribe to a cloud-based solution and start using it within weeks, if not days.

## The Infinite Resource Myth

<b>Myth 6</b>	<b>The cloud will automatically scale up and deploy whatever resources we need, essentially offering infinite resources to the customer</b>
<b>Reality</b>	There are technical limitations to the scaling-up capabilities of the provider, contractual limitations to what you can do, and financial consequences to excess usage.
<b>Discussion</b>	There is no such thing as infinite resources. In fact, the NIST cloud reference model says nothing about capacity, although the “rapid elasticity” and “on demand provisioning” characteristics imply that there is at least some spare capacity to satisfy scale-up requests. The hyperscale providers (such as Amazon Web Services, Microsoft Azure, Google Cloud Platform, and others) promise some reasonable degree of extra available capacity, and their business model depends on it. But this is still only an illusion of infinite capacity.

There are certain usage patterns that can cause the same bottlenecks and slow-downs in a cloud environment as they do on premises. For example, meetings start at the top of the hour (so that a cloud-based teleconferencing service may be slower then), and Black Friday sales affect cloud-based payment systems.

A well-designed application should be able to take advantage of scaling features within a cloud platform. But even then, the cloud is still a shared environment all the way down to the physical electrical and network interconnects into and out of the data centers.

In addition, certain applications may require specific hardware or network configurations to deliver maximum performance. The cloud environment is usually a “common denominator” that works well for most customers, but is not specifically optimized for any of them. Most cloud environments include virtualization layers in order to emulate the operating system required by the application on top of a common base kernel, and those layers consume resources. To avoid this for a high-performance application, the customer may need to configure the entire stack from scratch on top of a bare metal Infrastructure-as-a-Service (IaaS) offering. The more esoteric or specialized resources – such as high-performance computing (HPC) servers equipped with GPUs, used for visualization, scientific computations, or AI and machine learning – will typically be offered in more limited numbers. A customer who depends on such resources may need to pay for “reserved resources” to ensure that they will be available when needed.

On the contractual side, the customer may have rented a certain capacity with a limit to how much can be added. If a burst in demand can be predicted, it is possible to use the cloud’s self-service provisioning features to increase the available resources (at a cost, of course). Otherwise, the cloud customer may bump into the contracted limit. Some contracts allow usage to exceed the contracted level, but in that case the excess charges can be steep.

## Myths About Exiting the Cloud

<b>Myth 7</b>	<b>Migrating out of the cloud is as straightforward as it was to migrate into the cloud</b>
<b>Reality</b>	Migrating away from a cloud service is usually much more difficult – and more costly – than it was to migrate to it.
<b>Discussion</b>	<p>You are not married to a cloud provider forever, and there may be several reasons why you decide to migrate away from a service:</p> <ul style="list-style-type: none"><li>• the provider did not give you satisfaction</li><li>• a better provider has been identified</li><li>• you have decided to insource the service (whether for technical or financial reasons, or by executive fiat).</li></ul> <p>When your organization adopted a cloud solution, the provider was likely very supportive. When you decide to move away, you will generally find the cloud provider somewhat less supportive. The account executive who was looking for their commission</p>

when you were selecting their service now has no incentive to do more than the minimum specified in the Cloud Service Agreement (CSA).

In addition, exiting the cloud is usually done under time pressure. Once a termination clause has been invoked, the clock is ticking. When moving to the cloud, if you forget to import some records, you can usually fix this later. When moving out of the cloud, mistakes and omissions are less forgiving.

The CSA needs to be very clear about what happens during the exit process, including how many months' notice you need to give (and to whom and how) to avoid penalties or out-of-contract month-to-month costs. See *OMG's Practical Guide to Cloud Service Agreements* [9] for more details. Note that the incumbent cloud provider will sometimes offer additional services to assist during the exit process, but at a cost – and since you are abandoning them, they have no incentive to make this easy or cheap for you. In fact, you may need to hire a third party to support the exit process.

A special note about the handling of *data* during the exit: during the use of a cloud solution, which will usually last several years, it is likely that the data will have become more massive and more complex. Columns will often have been added to databases. Proprietary features and services offered by the cloud provider (e.g., validation rules and triggers) may have been leveraged, making the resulting data harder to extract and reload onto a different platform. It is easy to believe that you can stay away from such non-standard add-ons, but in practice they are impossible to resist. As a result, recovering all your data in a way that makes it easy to transfer to another service is more complicated and costly than most customers realize.

## Myths About Security

*“Some people are saying the cloud is less secure than having all your personal data on a server in the basement of your doctor’s office — a server that is maintained once a month by a 17-year-old who gets paid minimum wage. It works better in the cloud.”*

—Michael Nelson, then in charge of public policy at Cloudflare, Inc. as quoted in the Proceedings of Global Forum 2017 (Winnipeg).

*“From a technology perspective, placing data in a public cloud may actually improve security rather than weaken it.”*

—Claude Baudoin, in “Cloud Lessons Learned” [10]

### Myth 8

#### The cloud is less secure than your own data center

### Reality

As the above quotes imply, your data may be more secure in the cloud than on your premises. Keeping the data in your own data center for security reasons will sometimes be justified by regulatory and management reasons, but rarely by the technology itself.

### Discussion

The reasons why the cloud may be more secure, not less, include:

- The number and skills of security professionals that a CSP employs to protect its systems.

- The business imperative for a CSP to ensure security is as high, if not higher, as its customers'. A breach at a CSP is usually discovered, discussed in the media, and can have devastating business consequences.
- Multi-tenancy provides a level of obfuscation about where a particular customer's data resides.
- Most security breaches are committed by insiders, who can often exfiltrate their company's data without leaving a trace, while stealing data from the cloud is harder to conceal.

Comparisons between reported security incidents affecting cloud providers vs. end user companies are misleading, as end users often keep quiet about incidents, unless required by law to divulge data breaches, therefore those numbers are understated.

This said, security concerns may lead to a decision stay away from the cloud for legal, regulatory compliance, or management reasons:

- A CEO or CIO may fear that a decision to move to the cloud will cost them their job if there is a data breach at the cloud provider. Of course, they could still be fired if there is a data breach on premises, but at least they will not be told "this wouldn't have happened if you hadn't moved our data to the cloud."
- Data protection regulations (GDPR in Europe, CCPA in California, PIPEDA in Canada, etc.) have instilled additional fear in customers.
- Data residency laws may also interfere if there is no suitable domestic cloud provider, or if a cloud provider cannot guarantee that the data will remain in the country as they "load balance" their servers. In this case, security concerns are what caused the data residency requirements (or they were used as a pretext when the real motivation was government oversight), and this in turn may prevent a company from migrating to the cloud.

Speaking of data residency, the following myth has a basis in truth, but in the absolute form in which it is usually expressed, it is indeed a myth:

**Myth 9**

**It does not matter where a cloud service is hosted**

**Reality**

While the ability to locate the servers and the data wherever it is cheaper, or to load balance the infrastructure by moving services around multiple data centers, is appealing, there can be serious disadvantages in doing so. Customers should always know where their data is located.

**Discussion**

Large CSPs maintain data centers in multiple locations and often in multiple countries. This provides additional redundancy, reduced network latency, and allows the provider to reduce costs by choosing locations where energy or labor costs are lower.

When a peak in demand occurs in one location (e.g., Black Friday sales in the U.S.), the CSP can balance the load on its servers by moving services around its network. This is often sold to the customer as a benefit, and in many cases, it is indeed a positive thing.

However, if the cloud customer's data contains personal information such as bank accounts and balances, patient health data, or other personally identifiable information (PII), there may be laws and regulations that severely restrict where that data can be stored. A common false belief is that this is a problem for the cloud provider to solve.

However, regarding data protection laws, the CSP is only the custodian of the data, and the customer is the owner and the party legally responsible for any data residency violation. See [6].

The following myth relates specifically to the management of encryption keys, which are an important security feature in certain applications, especially for financial institutions and government entities.

<b>Myth 10</b>	<b>Managing encryption is easy thanks to the providers' BYOK (Bring Your Own Key) offering</b>
<b>Reality</b>	Encryption key management is critical for certain applications, especially for financial institutions, and the BYOK offerings are often more limited or more complex than what the brochure claims.
<b>Discussion</b>	<p>Some key questions to consider are (a) whether the provider actually lets the customer bring its own keys, or simply generates a key for the provider based on some data such as the customer ID number, (b) how that key is transmitted between the parties, (c) whether the customer can store that key where they want (and knows how to safeguard it). In summary, this is not as easy as it may seem.</p> <p>For specific considerations regarding key management for financial institutions, see the upcoming ANSI X9.125 standard.</p>

## Myth About Availability

<b>Myth 11</b>	<b>There are more outages in the cloud than in your data center</b>
<b>Reality</b>	Cloud outages impact many users at one time and therefore make the headline news. By contrast, most IT organizations do not publish their availability numbers, which are often worse.
<b>Discussion</b>	<p>The same argument applies to availability as to security: a CSP has a lot at stake in providing a stable environment, and they have more skilled resources to manage their systems, and restore them to normal in case of an incident, than any of their customers can.</p> <p>A common discussion item during the negotiation of a cloud contract is, "you only guarantee 99.5% availability, that's not enough, we need 99.9%." Some CIOs who demand this level of uptime may be unable to answer the question, "what is your current in-house uptime?" They may be surprised to learn that their current numbers are not as high as what the provider offers in their standard service level agreement.</p> <p>Obviously, different cloud providers perform differently. Generally, the larger and well-known providers will ensure higher availability, if only because they have more redundant systems and more system administrators. Looking at a provider's history is a must before selecting it.</p> <p>Availability commitments, and the remedies to be provided in case those commitments are not met, need to be specified in the cloud service agreement. For example, many high availability and disaster recovery (HADR) requirements can often be met by</p>

leveraging replication methods and availability zones, but it is up the customer to implement and test this.

While the CSA is often presented as a standard document that cannot be changed, some clauses may be negotiable. Even if they are not, the customer needs to understand the exact implications of the numbers that are mentioned. See the OMG Cloud Working group's guides on this subject. [11] [12] [13]

## Myths About Management

<b>Myth 12</b>	<b>Once you move to the cloud, the Cloud Service Provider does everything for you</b>
<b>Reality</b>	Management of the entire end-to-end environment is a shared responsibility between the cloud provider and the cloud customer.
<b>Discussion</b>	<p>In this shared responsibility model, the boundary between what the customer and the provider are responsible for is highly variable, depending in particular on:</p> <ul style="list-style-type: none"><li>• the service model (the customer maintains more responsibility in an IaaS environment, and less in a SaaS one)</li><li>• the terms of the Customer Service Agreement</li><li>• additional services that may be offered by the CSP at an extra cost.</li></ul> <p>This myth is often rooted in a misunderstanding of the difference between a Cloud Service Provider and an outsourcing provider. They are not equivalent. For example, an IaaS CSP does not eliminate the need for a systems administrator at the operating system level.</p> <p>The provider shoulders, at best, partial responsibility – for the pieces of the puzzle that are totally under their control. The customer is responsible for the rest, and especially for its data. The customer “owns” the interactions with its partners and customers. It is important to remember that an organization can never outsource responsibility for governance.</p>

Once the customer realizes that they are still responsible for a number of management tasks, the question shifts to *how* to perform this management. This gives rise to another myth:

<b>Myth 13</b>	<b>You can apply your existing policies and procedures to manage the cloud environment consistently with your on-premises environment.</b>
<b>Reality</b>	You will be expected to use the management tools provided for you by the cloud service provider. Anything else will require more work.
<b>Discussion</b>	<p>The cloud provider offers a good set of cloud-enabled tools they expect the customer to learn to use, and they build management applications to meet the general administration needs of all their cloud tenants. But if you have specific policies for password management, cryptographic algorithms, or certificates, you may have to build exceptions for your cloud environment – especially related to the SaaS service model.</p> <p>If you are currently using a specific best-in-class third-party tool, you may have to periodically export data from the cloud and use that tool in your on-premises</p>

environment. This will allow non-real-time analysis, but will not help troubleshoot transient performance or security issues. For those, you may have to be satisfied with your cloud provider's canned queries and reports. Those will only show you information related to your own usage of the cloud service, even though information from other tenants might be important to diagnose issues ("no, you cannot see those logs").

In addition, the cloud customer's IT policies and procedures that relate to systems currency (operating system and software version levels, security updates, etc.) must be aligned with the CSP's policies and procedures, update cycles and standardization, since mismatches could create interoperability or security problems.

**Myth 14**      **Disaster recover (DR) is built-in to cloud services.**

**Reality**      Moving to the cloud does not mean leaving behind disaster preparation. Cloud customers are ultimately responsible for their data and for the services they provide to their own customers, and cannot abdicate the responsibility to own their DR plan. As in Myth 12, DR should be planned and managed as a shared responsibility.

**Discussion**      A cogent argument against this myth is provided in [14]. The key points are as follows:

- The cloud customer is ultimately responsible for its data, regardless of the fact that it is paying the provider for a service.
- The cloud is not infallible.
- The cloud provider's backup and DR solutions may be incomplete or offer excessive recovery times.

In addition, consider that if an accident happens and the cloud provider's disaster recovery fails, the provider's liability is certain to be limited and the customer will not be properly compensated for the damage suffered, including interruption of business.

## Conclusion and Recommendations

Myths can be dangerous because they take a life of their own – they get repeated and amplified even when cooler heads have denounced them. In the case of cloud computing case, myths can lead to a wide range of inappropriate decisions, such as:

- rushing to adopt a cloud solution without the necessary analysis and planning,
- "overselling" the cloud to executives,
- conversely, denying your organization the benefits of cloud solutions,
- abandoning governance, management, and security to the cloud provider,
- turning a blind eye to data protection issues.

Over time, customers will learn to discern the reality behind some old myths, but new ones will emerge. It is important to stay informed and to use multiple reliable sources when forming a cloud strategy, rather than superficial analyses or brochures from vendors.

OMG's Cloud Working Group will do its best to help by periodically revising this document as well as its other Practical Guides and Discussion Papers.

## Appendix A: ChatGPT's Answer to the Myths Question

Generative pre-trained transformers, the best known of which is OpenAI's ChatGPT, burst on the scene in 2022. These large language models (LLMs) can be fed millions of documents during the training phase, can correlate the language tokens found in a user question to select elements of a response, and can generate high-quality textual answers that mimic what a reasonably educated user with an excellent command of grammar and style might write.

It was therefore tempting to ask ChatGPT what it knew about the subject of this paper. Its answers were close enough to our own understanding of the topic that it seemed justified to include them as an Appendix to this paper. If nothing else, the reader will appreciate what can be derived by a powerful (but still dumb) tool from the mass of available documents.

Note that if one reruns this experiment today, the results are likely to be different – hopefully even better – since the tool itself is constantly evolving and its training material continues to grow.

Here is (verbatim) the dialog we had with ChatGPT in January 2023.

### **Q: What are some of the myths of cloud computing?**

**A:**

1. Cloud computing is only for large businesses: While many large companies use cloud computing, it is also accessible and beneficial for small and medium-sized businesses.
2. Cloud computing is not secure: Cloud providers use advanced security measures to protect data, and in many cases, it is more secure than traditional on-premises data storage.
3. Cloud computing is expensive: While the costs of cloud computing can vary, it can often be more cost-effective than maintaining on-premises IT infrastructure.
4. Cloud computing is only for certain types of workloads: Cloud computing can be used for a wide range of workloads, including storage, backups, disaster recovery, and more.
5. Cloud computing eliminates the need for IT staff: Cloud computing can require different skills and roles for IT staff, but it does not eliminate the need for them altogether.

### **Q: What are those different skills and roles?**

**A:** The different skills and roles that may be required for cloud computing include:

1. Cloud architects: These professionals design and plan the overall structure and integration of a company's cloud environment.
2. Cloud engineers: These professionals implement and maintain the cloud infrastructure, including deploying, configuring, and managing cloud resources.
3. Cloud administrators: These professionals manage and monitor the day-to-day operations of a cloud environment, including security, performance, and scalability.
4. Cloud security specialists: These professionals are responsible for securing a company's cloud environment, including identifying and mitigating security risks.
5. Cloud developers: These professionals write and test code for cloud-based applications, and ensure that they are optimized for the cloud.



6. Cloud cost optimization specialist: These professionals are responsible for optimizing the cost of the cloud infrastructure and services
7. Cloud migration specialist: These professionals are responsible for planning, executing, and validating the migration of workloads and data to the cloud.
8. Cloud DevOps: These professionals are responsible for automating the cloud infrastructure deployments and managing the cloud infrastructure as code.

Note that the specific roles and responsibilities may vary depending on the organization and the specific cloud platform being used.

**Q: Does moving to the cloud always reduce costs, and how?**

**A:** Moving to the cloud can potentially reduce costs in several ways, but it is important to note that it is not always the case. It depends on the specific situation and usage of the organization. Some ways that cloud computing can potentially reduce costs include:

1. Pay-as-you-go pricing: With cloud computing, organizations only pay for the resources they use, rather than having to invest in and maintain their own IT infrastructure.
2. Scalability: Cloud computing allows organizations to easily scale up or down their resources as needed, which can help to avoid overspending on IT infrastructure.
3. Reduced maintenance costs: Cloud providers typically handle the maintenance and updates of the infrastructure, which can help to reduce costs for the organization.
4. Increased efficiency: Cloud computing can help organizations to streamline their operations and become more efficient, which can lead to cost savings.
5. Reduced capital expenses: By using cloud computing, organizations can avoid the upfront costs of purchasing and maintaining IT infrastructure, and instead pay for the resources they use on a subscription basis.

However, it is important to evaluate the costs of cloud computing and compare them to the costs of maintaining an on-premises IT infrastructure. There can be additional costs associated with cloud such as data egress charges, security, compliance, training and specialized skills, and so on.

It is also important to consider the organization's specific needs, including regulatory compliance, security, and performance requirements, and ensure that the chosen cloud platform can provide the necessary features and services to meet those needs.

**Q: How prevalent are some of those myths?**

Some of the myths of cloud computing, such as the belief that cloud computing is not secure or is only for large businesses, are less prevalent now than in the past. As more and more organizations have adopted cloud computing and have had positive experiences with it, these myths have been dispelled. However, there are still some misconceptions that persist in the industry:

1. Cloud computing is only for certain types of workloads: While this was true in the early days of cloud computing, now cloud providers offer a wide range of services and solutions that can be used for different types of workloads, from storage to big data, from machine learning to gaming.

2. Cloud computing eliminates the need for IT staff: As mentioned before, cloud computing can require different skills and roles for IT staff, but it does not eliminate the need for them altogether.
3. Cloud computing is always cheaper: While cloud computing can be more cost-effective than maintaining on-premises IT infrastructure, it is not always the case. As mentioned before, it is important to evaluate the costs of cloud computing and compare them to the costs of maintaining an on-premises IT infrastructure.
4. Cloud computing is only for modern companies: Some people might think that cloud computing is only for startups, digital natives, and high-tech companies but it is not true, companies of all sizes and industries can benefit from cloud computing.

It is important to understand that cloud computing is a complex field with many options and it is important to evaluate the best options and solutions for a specific company or organization. It is always good to gather information, ask experts or to consult with a cloud provider before making any decision.

Note: The right to reproduce ChatGPT-generated content, in particular who owns the copyright to such content, is still unclear. OpenAI, the company that created the software, recommends the following, which the above Appendix complies with:

*“Creators who wish to publish their first-party written content (e.g., a book, compendium of short stories) created in part with the OpenAI API are permitted to do so under the following conditions:*

- *The published content is attributed to your name or company.*
- *The role of AI in formulating the content is clearly disclosed in a way that no reader could possibly miss, and that a typical reader would find sufficiently easy to understand.*
- *Topics of the content do not violate OpenAI’s Content Policy or Terms of Use, e.g., are not related to adult content, spam, hateful content, content that incites violence, or other uses that may cause social harm.*
- *We kindly ask that you refrain from sharing outputs that may offend others.”*

# Appendix B: Stories from the Trenches

(Edited from contributions by Troy Anderson)

The following examples are provided in support of the preceding discussions of the myths.

## About the inevitability of moving to the cloud

One large company I worked for (Fortune 50) initially moved one of their e-commerce applications to the cloud because the costs could come from a different part of their budget (OPEX instead of CAPEX). Another reason was that it solved internal political issues. Finally, they could consolidate and standardize on a specific technology stack supported by the cloud provider. Once these issues were resolved, they could move the standardized and consolidated e-commerce solution back into their own data center and in fact save costs.

## About the “infinite resources” myth

A start-up specialized in attracting dissatisfied customers in the retail industry who were currently using a certain SaaS application, hosted by a certain cloud vendor. All customers of this application were sharing the same SaaS database, which would regularly freeze and require restarting during any significant retail event (such as post-holiday sales). One of the underlying issues was that many customers had been somewhat trapped by a licensing agreement where they were only being charged for a “single socket, four-core” license, but the vendor noted they would not enforce those usage limits, so they could use as much of the vendor’s instance as they needed. But of course, that is what the vendor told every customer...

## About policies and procedures

A popular SaaS collaboration offering is Microsoft Teams. It is hosted in Microsoft’s Azure cloud, and individual user access is based on unprotected user access tokens, which are saved in the file system of the user’s client machine. This SaaS application also requires user certificates with a lifetime set to 10 years. Microsoft calls this certificate lifetime an “industry standard” – referring to the provider’s industry, not the customer’s – and does not accommodate shorter certificate lifetimes required to meet regulated industry policy requirements. It also does not offer an option to treat user access tokens more like passwords (i.e., hashing or encrypting). So even a large bank, for example, may have to adjust its policies and procedures to meet the SaaS provider’s expectations, and may even need to customize its cybersecurity scanning tools to ignore these policy violations, which would normally produce alerts.

## About cloud costs

A start-up company with no data center initially served a single customer with an e-commerce environment using external hosting services, then started adding more customers – first quarterly, then monthly. The fees charged to customers had to be increased when the hosting provider raised their prices. If a customer had left, it would also have been necessary to cut down on the infrastructure in order to match costs to revenue. This rental model remained valid during the company’s expansion. But when the startup was purchased by a larger international corporation, customer environments were migrated into their wholly-owned and licensed infrastructure to begin maximizing their profits.

## Appendix C: References

- [1] Object Management Group: *Practical Guide to Cloud Computing, v3.0*. Dec. 2017. [www.omg.org/cloud/deliverables/CSCC-Practical-Guide-to-Cloud-Computing.pdf](http://www.omg.org/cloud/deliverables/CSCC-Practical-Guide-to-Cloud-Computing.pdf)
- [2] Gartner: *The Top 10 Cloud Myths*. Nov. 2019. [www.gartner.com/smarterwithgartner/the-top-10-cloud-myths](http://www.gartner.com/smarterwithgartner/the-top-10-cloud-myths)
- [3] McKinsey: *Debunking seven common myths about cloud*. Oct. 2020. [www.mckinsey.com/capabilities/mckinsey-digital/our-insights/debunking-seven-common-myths-about-cloud](http://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/debunking-seven-common-myths-about-cloud)
- [4] IBM: *Myths and Truths in Today's Hybrid Cloud World*. Feb. 2023. [www.ibm.com/cloud/blog/myths-and-truths-in-todays-hybrid-cloud-world](http://www.ibm.com/cloud/blog/myths-and-truths-in-todays-hybrid-cloud-world)
- [5] Linthicum, David: *3 common myths about sustainability and cloud computing*. InfoWorld, Apr. 2023. [www.infoworld.com/article/3692532/3-common-myths-about-sustainability-and-cloud-computing.html](http://www.infoworld.com/article/3692532/3-common-myths-about-sustainability-and-cloud-computing.html)
- [6] Object Management Group: *Data Residency Challenges*. May 2017. [www.omg.org/cloud/deliverables/CSCC-Data-Residency-Challenges.pdf](http://www.omg.org/cloud/deliverables/CSCC-Data-Residency-Challenges.pdf)
- [7] National Institute of Standards and Technology (NIST): *NIST Cloud Computing Reference Architecture*. Special Publication 500-292, Sep. 2011. <https://doi.org/10.6028/NIST.SP.500-292>
- [8] Object Management Group: *XaaS (Anything as a Service) Glossary*. June 2022. [www.omg.org/cloud/Anything-as-a-Service-Glossary-22-06-08.pdf](http://www.omg.org/cloud/Anything-as-a-Service-Glossary-22-06-08.pdf)
- [9] Object Management Group: *Practical Guide to Cloud Service Agreements, v3.0*. Feb. 2019. [www.omg.org/cloud/deliverables/Practical-Guide-to-Cloud-Service-Agreements.pdf](http://www.omg.org/cloud/deliverables/Practical-Guide-to-Cloud-Service-Agreements.pdf)
- [10] Baudoin, Claude: *Cloud Lessons Learned*. Cutter Business Technology Journal, Vol. 31 No. 10, Dec. 2018. [www.cutter.com/article/cloud-lessons-learned-501386](http://www.cutter.com/article/cloud-lessons-learned-501386)
- [11] Object Management Group: *Practical Guide to Cloud Service Agreements, v3.0*. Feb. 2019. [www.omg.org/cloud/deliverables/Practical-Guide-to-Cloud-Service-Agreements.pdf](http://www.omg.org/cloud/deliverables/Practical-Guide-to-Cloud-Service-Agreements.pdf)
- [12] Object Management Group: *Cloud Service Agreements: What to Expect and What to Negotiate, v3.0*. Sep. 2019. [www.omg.org/cloud/deliverables/cloud-service-agreements-what-to-expect-3.0.pdf](http://www.omg.org/cloud/deliverables/cloud-service-agreements-what-to-expect-3.0.pdf)
- [13] Object Management Group: *Proposal for a Standard Cloud Service Agreement Template*. Dec. 2022. [www.omg.org/cgi-bin/doc?mars/22-12-14.pdf](http://www.omg.org/cgi-bin/doc?mars/22-12-14.pdf)
- [14] Paquette, David: *Cloud Alone Is Not Enough. You Still Need Disaster Recovery*. Nov. 2021. [www.zerto.com/blog/cloud-data-management/cloud-alone-is-not-enough-you-still-need-disaster-recovery](http://www.zerto.com/blog/cloud-data-management/cloud-alone-is-not-enough-you-still-need-disaster-recovery)