




Executive Brief

Cloud Adoption —




Introduction

Cloud computing and enterprise use of the cloud has exploded in the past decade. Total spending on public cloud solutions grew 366% between 2009-2019, reaching an estimated \$214B. The growth is expected to continue at a 16% annual rate through 2022 as more companies begin to see the benefits of a cloud-based IT system. Even though the cloud has firmly cemented itself as one of the hot topics within IT, only 50% of global IT infrastructure spending is allocated towards cloud solutions. 58% of businesses face challenges when migrating their infrastructure and services to the cloud, partially due to the complexity of such a process. Understanding the basics of cloud adoption and use are necessary in order to take the first steps towards this new world.

Source: Statista Cloud Computing report 2019, IBM

In this Executive Brief, we want to provide an informative guide to what the cloud is, the benefits and challenges of cloud adoption, and how your company can get started implementing a more flexible cloud solution. Our more than 30 years of experience serving large organizations with top 3% talent means that we have helped solve complex cloud challenges since the very beginning. Now you too can acquire the knowledge we've gained.

Happy reading!



Oslo, Norway —



Contents

2	Introduction
4	What do we talk about when we talk about cloud?
5	Deployment models
6	Service models
8	Why cloud?
13	Establishing a cloud business case
16	Maximizing success with cloud adoption

What do we talk about when we talk about cloud?

Organizations often talk about “the cloud” as if it is a single, simple concept. The reality is that the cloud takes several forms. There is no single approach or solution that is a panacea for every company, as moving to the cloud often is a series of strategic, tactical and operational decisions unique to the given company’s situation. But in this uniqueness there are some commonalities.

With a liquid term as the cloud, and many industries and service providers jumping the bandwagon with “everything-as-a-service” offerings, it can be hard to navigate what actually comprises the cloud. During the last couple of years, the National Institute of Standards and Technology’s definition and characteristics have gained traction¹. They describe 5 common characteristics of cloud computing as being:

- **On-demand self-service:** Users must be able to access services/resources at any level, at any time, with no or limited assistance.
- **Broad network access:** Cloud services and capabilities must be available to the user via network standards through thin clients.
- **Resource pooling:** Computing resources are pooled across several users through multi-tenancy (whether on platform or application level), to ensure efficiency.
- **Rapid elasticity:** Capability can be provisioned quickly to ensure smooth scale in and out, allowing for consistent performance, independent of workload.
- **Measured Service:** Complete transparency in the use of service through metering allowing for both optimization and close control.

Whenever you are able to tick these five “boxes”, you are most likely dealing with a cloud computing solution. Cloud computing solutions are based on three service models and four deployment models. These are the models that have been depicted endless times in some form or the other (this paper no exception). In the following section we will go one step further and define and exemplify these.

The deployment models are the models through which an organization can implement and use cloud services, and there are four types: The Public model, The Private model, The Community model and The Hybrid model. Each determines which cloud providers can be used, the level of openness to your cloud setup, and the types of services that can be built on top of this foundation.

The Service models define how your organization uses the cloud and have parallels to the on-premise setup. Software, infrastructure and platforms are all offered “as a service” in the cloud. Each give varying levels of flexibility and control to cloud users, and the combination of services your organization should use depends on your size, industry, and specific business requirements.



Deployment models

The choice of deployment model dictates how you want a cloud setup to be managed and delivered to your organization. Different approaches are beneficial for different needs. Therefore, it is both important to have an understanding of each, as well as a vision for your cloud journey.

The Public model

The Public model is a completely external option wherein a third party owns and operates the services. This can be a good solution for organizations looking to phase out old infrastructure or software or looking to add new functionality missing from the current setup. The Public model provides greater flexibility, reduced cost, and faster development time than The Private model, although the ability to customize services is reduced.

The Private model

The Private model shares the most similarities with a typical on-premise solution in that everything is based internally on a company's intranet or hosted data center. It can be a strong solution for organizations with existing infrastructure that can still be utilized, and it can be fully custom-built to meet your needs. Back-end work such as maintenance, software updating, and data management remain the responsibility of the company itself.

The Community model

The Community model is constructed of multiple parties that share the same mission, requirements etc. By construct, the community models have similarities to both The Private model and The Public model, and can be seen as a shared private cloud, or a closed public cloud. While parties can benefit from the limited security exposure, legal issues, network bandwidth, the community as a whole is responsible for the creation, management and maintenance of the cloud.

The Hybrid model

The Hybrid model combines aspects from The Private and The Public models. It might make sense for an organization to use public providers for certain services, while retaining the customizability that comes with a private cloud. Using The Hybrid model is also a good way to ease the transition from fully on-premise to complete cloud-based IT. This approach has such an appeal that 58% of enterprises are currently using The Hybrid model.

—

When choosing either of these deployment models, many larger companies choose to go with a multi-cloud strategy. This due to many factors, but most often: To avoid vendor lock-in – being dependent on single provider makes it difficult to move without high costs, technical compatibility issues or legal issues. To better mitigate risks – spreading workloads across multiple cloud environments secure redundancy, minimize downtimes, outages and data breaches. To gain flexibility – every cloud provider has different features, strengths and weaknesses. Diversifying your deployment let's you reap the advantages of each provider and avoid their weaknesses at the same time.

Source: RightScale 2019 State of the Cloud Report

InfoBox

One of the best analogies to explain the differences in the cloud deployment models, was presented by Mike Lloyd, CTO of RedSeal. By thinking of cloud deployment as office space, we are able to make easy distinctions.

The Private cloud: Your company owns its office space, and hence bare the burden of maintenance, clerical work, having the cantina in place, cleaning etc. You own the building and are responsible for everything to make it work.

The Public cloud: Your company is renting its office and facilities. You are renting on the same terms as all other tenants, and are included in the common cleaning and lunch-offering – for a fee of course.

The Community cloud: You and a couple of other companies decide to buy a shared office space. In doing so, you all agree to have shared rules of conduct and take turns cleaning the offices.

The Hybrid cloud: Your company owns one building but rents some additional space. The company is large and likes to have full control over it's core facilities, but it has some rented space for increased flexibility.

Service models

Though a gradient of a portfolio of increasingly externally managed services, we often see three classifications of cloud service models; Infrastructure as a Service (IaaS), Platform as a service (PaaS), and Software as a Service (SaaS). Though heavily used, it can sometimes be hard to figure out what these terms mean? It's actually all about what degree of "virtual outsourcing" you need.

Infrastructure as a Service

Infrastructure as a Service refers to ownership, management, and operation of computer hardware at a third-party. This includes servers, hard drives, networking, virtualization, storage etc. An organization's access to the IaaS is typically provided through a dashboard or an API. IaaS users are responsible for managing applications, runtime, middleware, and data. This sharing of responsibilities gives users a high degree of control over what is on their cloud system, without dealing with the system itself.

Platform as a Service

A PaaS, or Platform as a Service means that on top of the hardware itself, you can pass to a service provider the burden of owning, managing, and operating systems software. You do not manage or control the underlying cloud infrastructure including network, servers, operating systems, or storage, but you have control over the deployed applications and possibly the application hosting environment configurations.

Software as a Service

SaaS, or Software as a Service, means that you outsource everything related to running an arbitrary piece of software. You don't manage or control the underlying cloud infrastructure including network, servers, operating systems, storage, or even individual application capabilities, with the possible exception of limited user-specific application configuration settings.

InfoBox

One of the best analogies to explain the differences in the cloud service models was presented by Albert Barron at IBM. By thinking of cloud service as pizza, we are able to make easy distinctions.

On-premise: You make the pizza from scratch. You buy all the groceries, make the dough, assemble the pizza, cook it, and eat it in the comfort of your own home.

IaaS: Instead of going through the hassle of buying all the ingredients, you can buy a take & bake pizza – already assembled. You still need to cook it at your home, set the table, buy the soda, and enjoy it in your own home.

PaaS: You order a pizza to be delivered. Though you lose some control over the preparation time, ingredients, etc., you leave it to the specialists to prepare most. You still need to buy the soda and enjoy it in your own home.

SaaS: You go to a pizza restaurant and let them take care of everything. You do not have complete control over the process, or even the selection of soda, but you are completely hassle free, enjoying the pizza out of your home – and you don't have to do the dishes.

Differences between traditional and cloud

Characteristic	Cloud Computing	Traditional	Comments
Time before service can be accessed	Minutes/ Hours	Days/Weeks	Once the cloud computing environment is set up initially, you can gain access faster than in traditional environments where lead time is needed for installation, set-up, and configuration. Furthermore, if the initial setup is done with the "infrastructure-as-code" approach, all subsequent efforts to create cloud-hosted environments will be significantly cheaper to deliver as they will be based on code templates.
Capital Expenditure (CAPEX)	Pay-as-you-go, variable	Upfront cost, fixed	The pay-as-you-go model for cloud computing reduces or eliminates the large upfront costs incurred in procuring hardware and software and standing up traditional environments. However, pay-as-you-go can be more expensive in the long run without the proper reservations, usage forecasting, licensing, and rightsizing.
Economies of scale	Yes for all organizations	Yes for large organizations only	Cloud computing not only provides cost advantages in procurement of hardware and software, it also provides cost advantages from improved productivity. When using cloud-computing it is possible to spin up infrastructure on demand and decommission it when it's not needed - leading to savings.
Multi-tenancy	Yes	Generally no, but can be found in application hosting	Multi-tenancy properly applied to cloud computing services allows providers to host multiple consumers effectively across shared resources. While it is more readily enabled in IaaS through the use of virtualization, PaaS and SaaS providers may need to undertake significant re-architecting of their platforms or applications to apply multi-tenancy to these elements as well as to infrastructure. Where this has not been undertaken, consumers may find that their platforms and applications are not as elastic or cost-effective as anticipated.
Scalability	Elastic and automatic	Manual	Cloud computing resources can often be scaled up or down automatically, whereas human intervention is usually needed to add hardware and software in traditional environments.
Virtualized	Usually	Sometimes	Cloud computing environments are almost always virtualized, whereas traditional environments include a mix of physical and virtualized infrastructure.

Source: The Open Group

Why cloud?

While the cloud definitely is the current trend and on everybody's lips these days, there can be thousands of reasons to both go cloud and thousands of reasons not to. Whether moving to cloud is valuable to you depends on your organization's size, needs, required customizability, security considerations, etc. There is no single truth, but at least there are some key considerations that can guide your decisions.


Most companies face many considerations when discussing whether the cloud is beneficial. Who owns the data? Who's responsible for security and compliance? Can we live with internet connection dependency?

Do we have the competence to manage multiple vendors? How do we handle interfaces and data integrations? – to our on-prem core?

Concern about data security

43% see cloud data security as a challenge during adoption. There is a perception that off-site storage is less secure than on-site, especially with the use of large public-cloud providers.

This perception is not always reflected in reality. Companies that implement appropriate cloud security measures in cooperation with providers experience 33% fewer security failures than they did with an on-premises solution.

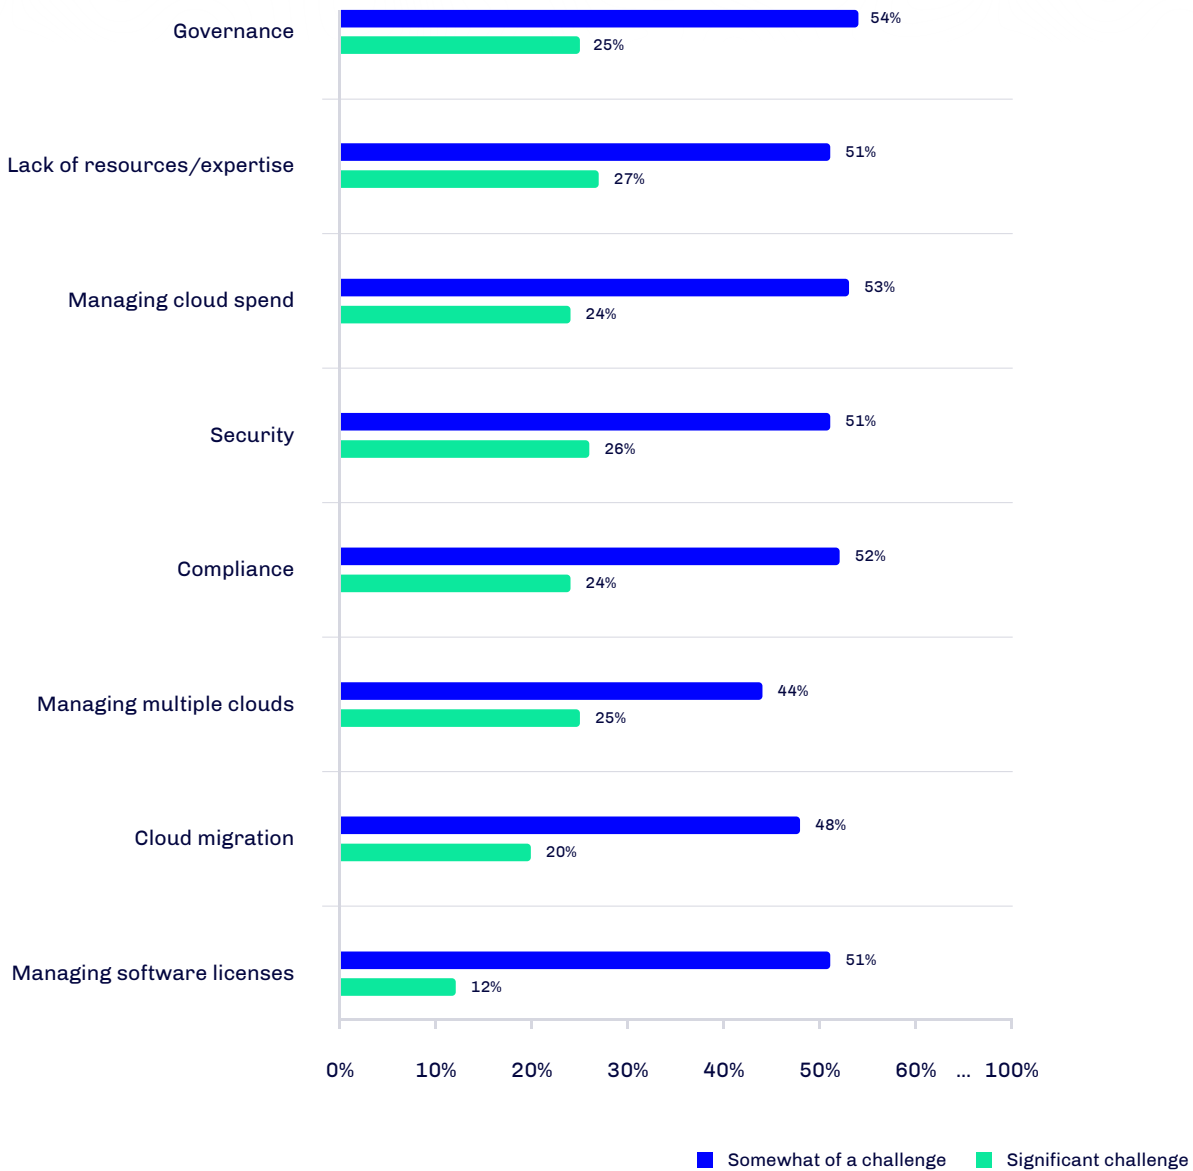


Cloud challenges

These concerns are all relevant and should be part of every assessment of if and how the cloud can add value to your operation. In a recent study by Flexera², companies reported a diverse spectrum of cloud

challenges that is both a representation of the uncertainty revolving the relatively “new” adoption of cloud and the pressure to “get on the journey”.

Challenges of using cloud computing worldwide as of 2019



Sources: RightScale, Flexera Software © Statista 2019


Additional Information: Worldwide; January 2019; 786 respondents; Technical executives, managers and practitioners of cloud technology

Main cloud goals

While the uncertainty of moving to the Cloud is driving the challenges and concerns regarding the cloud move, there are also significant benefits to be reaped. With a still increasing number of companies moving to the cloud, we get a better insight to the realized

benefits – the ones that work in reality, and not just on paper. Looking at the IDG Cloud Computing Study, there are many different drivers of the cloud computing journey.

What business goals or objectives are driving investments for cloud initiatives at your organization?



Source: IDG Research Services © Statista 2019

Additional Information: Worldwide; IDG Research Services; 2018; 550 respondents; Involved in the purchasing process for cloud computing

Drivers of the cloud

In general, the drivers of cloud can be divided into four different categories as proposed by the open group³: Agility, Productivity, Quality and Cost.

Agility

Rooted in the construct of cloud computing is both the on-demand self service and the rapid elasticity. This allows for higher levels of agile operation, as resources can be deployed much faster than traditionally, and also with flexibility as the elasticity allows for provisional demand balancing. Agility drivers hereby comprised reduces time to market, flexible business process support, near-instant access to development resources, smooth resource scaling, etc.

Productivity

Cloud computing allows for a unified business ecosystem, which significantly eases transparency, information distribution, common processes, and tools. A central “pro” for cloud solutions is the limited dependability of geographical closeness. Hereby a cloud adoption can be the lever for simplification, collaboration, work distribution, shared logic, and ease of access – all drivers of productivity value.

Quality

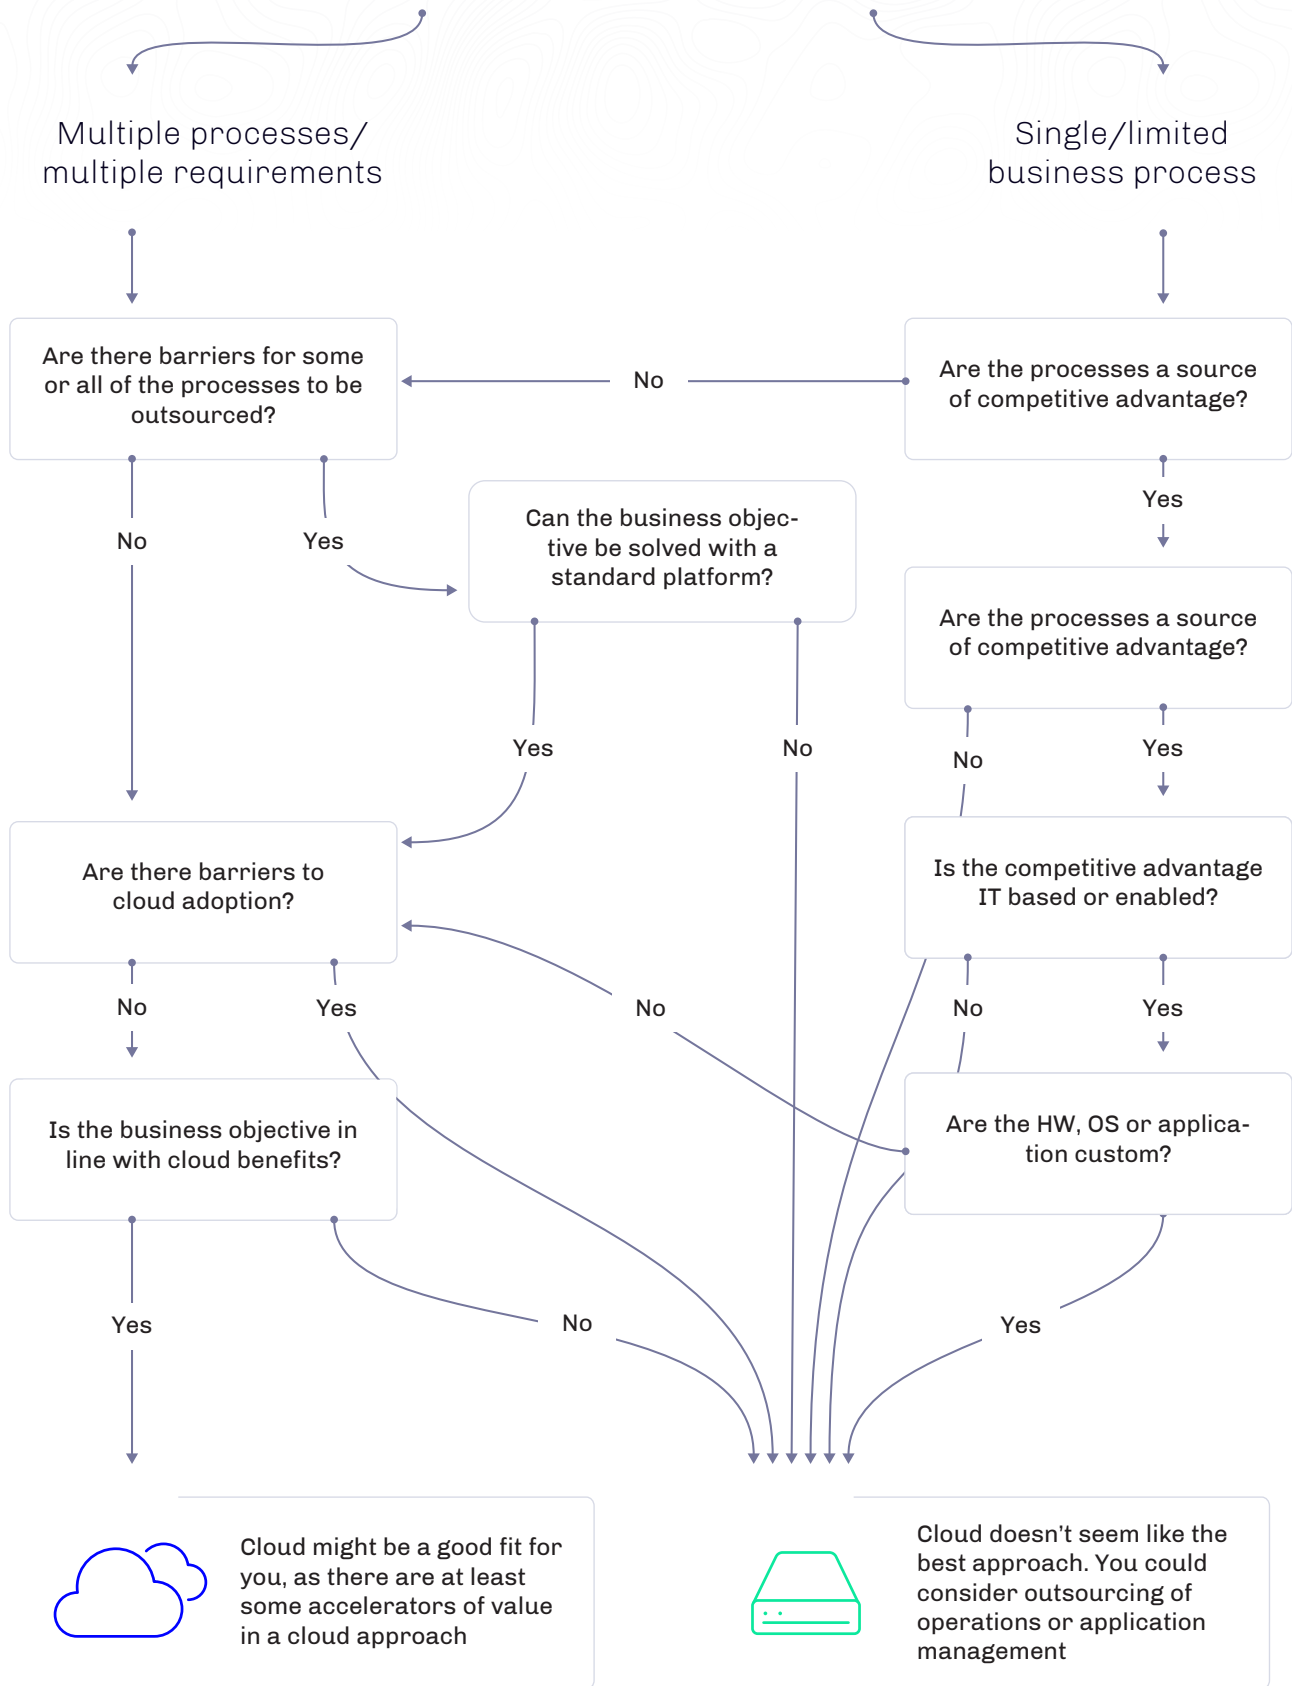
Since measurability is one of the core pillars of cloud computing, one can easily get a better overview of the actual use across systems. This is an integral aspect of the quality drivers, as the better measurability allows for a general better use of information, utilization optimization, better manageability through central provision, better responsiveness, reduced idle time, etc.

Cost

One of the main driving categories of cloud adoption is cost. At face value, and with the outsourcing of responsibility, one might expect that cloud is more expensive. While this is often true, it doesn't have to be, as the elasticity and measurability enables cost savings – if done correctly. Server consolidation, rightsizing, thin clients, and cost sharing are the main drivers of direct cost savings, while the transition from OPEX to CAPEX might be beneficial for companies with either high cost of capital or need for cash-flow.

In the evaluation of different options, thorough analysis of requirements, strategy, commitment, and cost are always needed. Cloud adoption is not done overnight and it should never be done without careful consideration, and close involvement of your business. To get you started with these considerations of both challenges and benefits, we have created a decision tree, helping you start the process. This tree is borrowed from the open group that has a comprehensive guide for the decision process.

What are you planning on moving to the cloud?



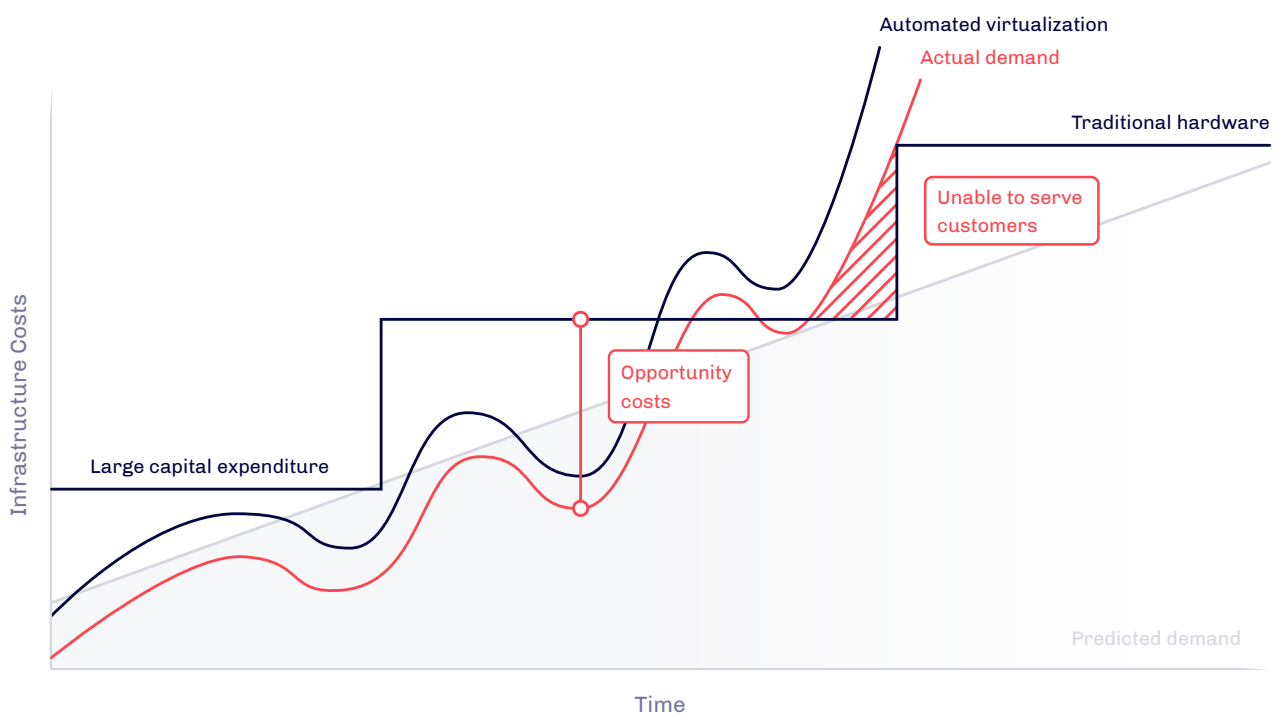
Establishing the cloud business case

An integral part of any cloud decision should also encompass the business case. As with any migration project, most often additional cost will occur in the transition phase, but with cloud it is not “just” a transition of system, it is a transition of the cost structure as a whole. In order to reap the complete benefits of the cloud, you also need to be prepared and equipped to manage this new reality.

In the previous parts of this report we have already touched upon the “soft savings” of moving to the cloud. The added value agility and productivity can be seen as an alternative cost to a traditional setup and should not be neglected when establishing the business case.

Beyond these soft cost savings there are also some more hard savings that we will discuss in this chapter. The general notion is that you over time need significant capital investments to be able to meet increas-

ing demand. This block (or stair) scalation approach is often not favorable as you either underserve the customers or pay for unused capacity. This is due to the incremental nature of capital investments, that is not very suitable for dynamically increasing demand. What the cloud offers is a way to “pay-as-you-go”, which not only mimics the predicted demand better, but also allows for a real-time allocation to meet actual demand.



“So” – you might think – “perfect, I’ll just move to the cloud and I will not pay too much or too little”. Unfortunately, this is not the case. In order to get a better and more comprehensive appreciation of the cost

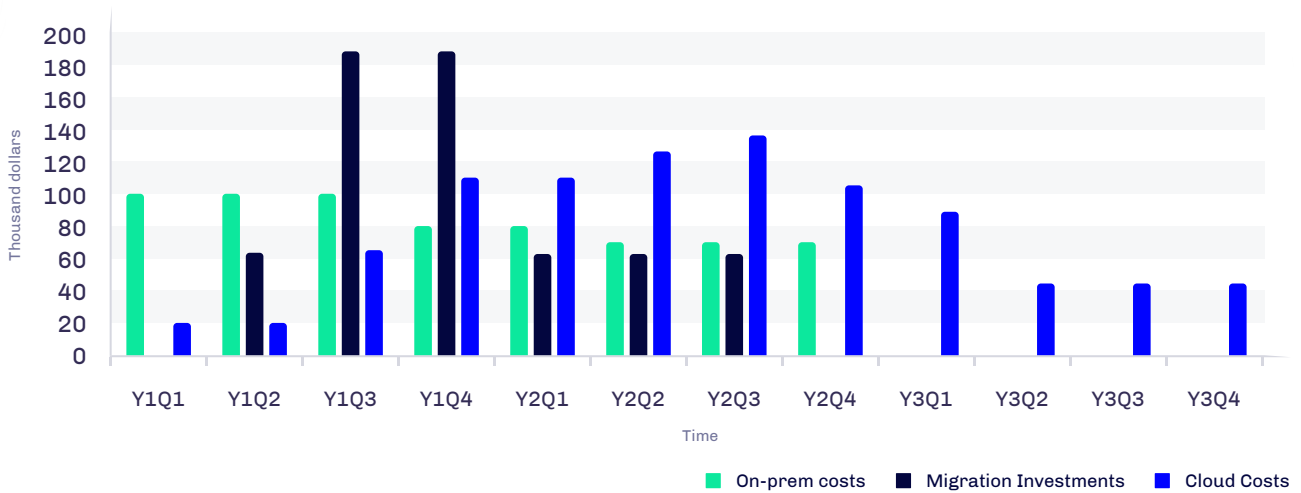
structure related to cloud adoption, we need to look at two things – one is the migration cost, and then the operating costs.

Migration costs

Moving from one setup to another is not cheap. One of the most commonly overlooked problems in relation to cloud adoption is that we tend to be overly focused on the endpoints – that is the delta between the traditional maintenance cost, and the cost of maintaining a perfectly tweaked and optimized cloud setup.


The period between these two end points are of huge importance, as a lot of investment and double cost will occur – costs that if not accounted for, might jeopardize the entire migration process.

In the below example borrowed from Gartner, this is clearly illustrated.



It becomes painfully evident, that there is a lot more happening than just the end points, and that securing the investment is of absolute importance. If we com-

pare the total cost over the course of the three years above, it is almost 85% more expensive to migrate than to just service the traditional on-premise setup.



As seen above on the aggregated cost over time, this suggests an isolated pay-back time of more than seven and a half year! This is of course not the whole story. Most companies face the cloud journey with

a clear ambition of increasing the value delivery and to get a complete overview of your business case, these needs both to be quantified and incorporated in the calculation.

Source: Gartner

Operating costs

Apart from the migration cost, you also need to have a good understanding of a completely new cost scheme. If you look at the illustration of the 3-year migration period in the previous section, you might notice that the cloud cost at some point is almost 40% higher than the initial on-prem cost. This is because at face value, the cloud IS more expensive than traditional deployment. If you don't tweak your setup, the 1:1 operating cost does not justify any migration at all from a cost perspective. How you operate your cloud is what will enable you to see a declining cost curve on your cloud expenditure.

In theory there are endless ways to optimize your cloud cost; multiple providers with different offerings and prices, complete overview of data enabling continuous optimizations, and so on. Cloud economics is a topic for a whole white-paper in itself, but here we try to present you with the most common processes you will need to incorporate in your operating model; rightsizing, allocating, scheduling and reserving.

Rightsizing

We have traditionally been used to oversize our servers for many different reasons. Whether it is to ensure room for growth, avoidance of request processes, seasonality, or uncertainty regarding load it seem that we are prone to oversizing – especially in traditional environments. With the measurability of cloud computing, we are able to assimilate the actual need to a much better degree and we can ensure automatic procedures to account for most shifts in workload – ensuring continuous performance.

Allocation

We have not been used being able to “move” our computing capacity. With the cloud your computing can happen from datacenters anywhere in the world. Due to differences in demand and seasonality etc., computing can be seen as electricity and other “commodities” – you can buy it from where it is cheapest, even from a single provider. Cloud providers themselves are trying to optimize at all times and an effective way of doing so is to incentivize computing on underused assets.

Sheduling

With a PAYG setup, you are paying by the minute/hour of runtime – so it might seem simple, that you turn off resources when they are not used. If you are scheduling to shut down and restart instances, so you don't leave them running overnight or on the weekend when nobody is using the resources you can bet significant cost reductions. This often seems easier than it actually is; as an example one could argue that dev servers shouldn't be live during weekends or nights, but then you risk a situation where remote work or working from home is not possible – limiting progress and often leading to annoyance. Scheduling is a careful skill that needs to be mastered and continuously optimized.

Reserving

It is actually quite simple – if you book resources in the future and commit to payment on the long term, then you will be able to acquire your resources cheaper. All cloud providers are working with uncertainty as one of their enemies. If they can have a larger degree of predictability, they can offer lower rates (often significant on a 2-3 year horizon). To minimize cost, it is therefore important that you analyze and commit to your base load.

Maximizing success with cloud adoption

Success with cloud adoption (in terms of cost savings, time to implementation, and value-delivered to the organization) is dependent on thorough planning and a deep understanding of your organization's capabilities, needs, and ambitions. We know that every company and every journey is unique, and that's why we at 7N don't have any "shelf products". We are cloud and technology agnostic, and we focus on delivering the specialists, that with knowledge of your industry and tech stack can enable your own organization to drive a successful cloud adoption journey.




Cloud adoption is not a checkbox

Embarking on a cloud adoption journey is not a straightforward process. Even though the market is saturated with different out-of-the-box offerings, the reality is that your specific setup is unique and hence the implementation should be as well. We at 7N do not deliver anything to you. We deliver with you, so that the knowledge accumulation and know-how is anchored in your business. We carefully select every single consultant to match the exact profile and needs

of your company. By providing the most experienced consultants the industry have to offer, you get the best possible guidance to succeed with your cloud journey – no matter the stage or maturity. With more than 1.400 specialists working across Europe, North America, and Asia, and more than 30 years of proven delivery, we know the difference it makes to not settle for good – but to find the perfect match.

Sources

- Flexare. (2019). *RightScale® State of the Cloud Report*.
- IDG Research Services. (2018). *IDG Cloud Insights 2018*.
- KPMG; Harvey Nash. (2018). *CIO Survey*.
- Meinardi, M. (2017, November 27). *New Research: How To Manage Public Cloud Costs on Amazon Web Services and Microsoft Azure*. Retrieved from Gartner Blog Network: <https://blogs.gartner.com/marco-meinardi/2017/11/27/new-research-how-manage-public-cloud-costs-amazon-web-services-microsoft-azure/>
- Meinardi, M. (2018, November 30). *Is Public Cloud Cheaper Than Running Your Own Data Center?* Retrieved from Gartner Blog Network: <https://blogs.gartner.com/marco-meinardi/2018/11/30/public-cloud-cheaper-than-running-your-data-center/>
- Mell, P., & Grance, T. (2011). *The NIST Definition of Cloud - Recommendations of the National Institute*. U.S. Department of Commerce.
- RightScale, Flexera Software © Statista 2019.
- Statista Cloud Computing report 2019, IBM.
- The Open Group. (n.d.). *Cloud Computing for Business*. Retrieved from The Open Group: http://www.opengroup.org/cloud/cloud_for_business/index.htm



Contact



www.7n.com/contact



7N A/S is a global, elite IT consultancy and agency with 30 years market experience in serving all aspects of critical IT projects both within the public and private sector both clients and consultants.

We have dedicated ourselves to finding the right match between our consultants and the companies we serve – we believe that is how the best results are created. At 7N, we have built a professional community of extraordinary people. A community dedicated to achieving professional and personal development. A place where the best gets to play with the best.