

# Why Industry Leaders are Leveraging the Cloud to Grow their Business and Drive Innovation

LefeWare Solutions

# Table of Contents


| | | |
|---|---|---|
| 01 | About the Author | P 03 |
| 02 | Introduction | P 03 |
| 03 | 8 of the most common reasons why our clients are moving to the cloud | P 06 |
| 04 | Improve Security | P 08 |
| 05 | Lower Costs | P 10 |
| 06 | Near Infinite Scale | P 14 |
| 07 | High Availability | P 18 |
| 08 | Never Lose Data | P 22 |
| 09 | Improving Efficiency through Automation | P 25 |
| 10 | Drive Innovation | P 28 |
| 11 | Collecting and Processing Data | P 31 |
| 12 | Should we make the Jump to Cloud? | P 34 |
| 13 | References | P 35 |



## About the Author

Joshua Lefebvre is the Director and chief Cloud Architect at LefeWare Solutions.

# Introduction

Here at LefeWare Solutions, one of our goals is to help our clients move their mission-critical applications, files and data to the public cloud in a secure, cost-effective manner which will help them pass security audits, increase their overall team velocity, cut spending, and innovate in their industry. Now, you might be thinking to yourself, "We are already able to accomplish these feats using our private cloud provider or our on-premise servers. So why bother migrating to the public cloud?"

The following paper explores 8 of the most common reasons why our clients are moving to the cloud. It also aims to address several common misconceptions about the public cloud, as well as shine a light on the reasons why so many industries are suddenly shifting the majority of their workloads to public clouds like Microsoft Azure, Amazon Web Service (AWS), or Google Cloud Platform (GCP). In a recent 2021 survey by Harvard Business Review Analytic Services of 300 IT professionals showed *"the vast majority (81%) respondents said that cloud is very or extremely important to their organization's future strategy and growth,"*[1] confirming a clear trend towards cloud adoption amongst innovators.

**81%**

**Responders said Cloud is Extremely Important to their Organizations' Strategy and growth**

## Before jumping in, it's important to first take a moment to clarify the differences that exist between public cloud, private cloud, on-premise, and hybrid:



### Public Cloud:

Servers are owned and operated by a third-party cloud service provider and delivered over the internet. With a public cloud, all hardware, software, and other supporting infrastructure is owned and managed by the cloud provider. Some key examples of this are AWS, Azure, and GCP.



### Private Cloud:

Servers are used exclusively by one business or organization. The private cloud can be physically located at your organization's on-site datacenter, or otherwise hosted by a third-party service provider. All services and infrastructure are constantly maintained on a private network, with the hardware and software being dedicated solely to your organization.



### On-Premise Data Center:

An on-premise data center essentially functions as a private cloud; this means that you are responsible for the hardware inside of a physical data center facility.

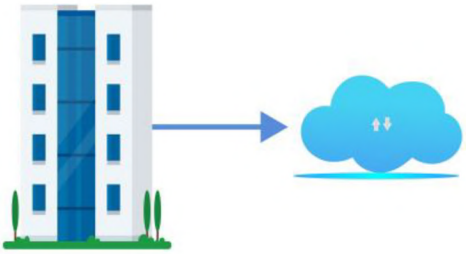


### Hybrid:

Combines on-premises infrastructure (or a private cloud) with a public cloud, allowing data and apps to move between the two environments.

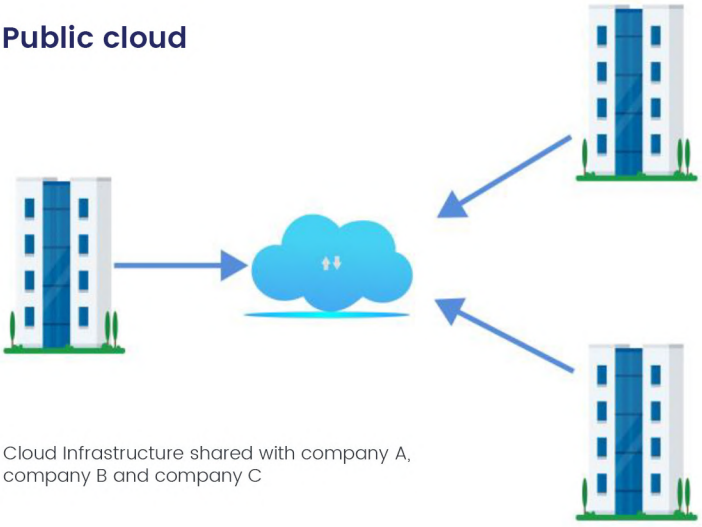


## Private cloud



Cloud infrastructure only
used by company A

## Public cloud



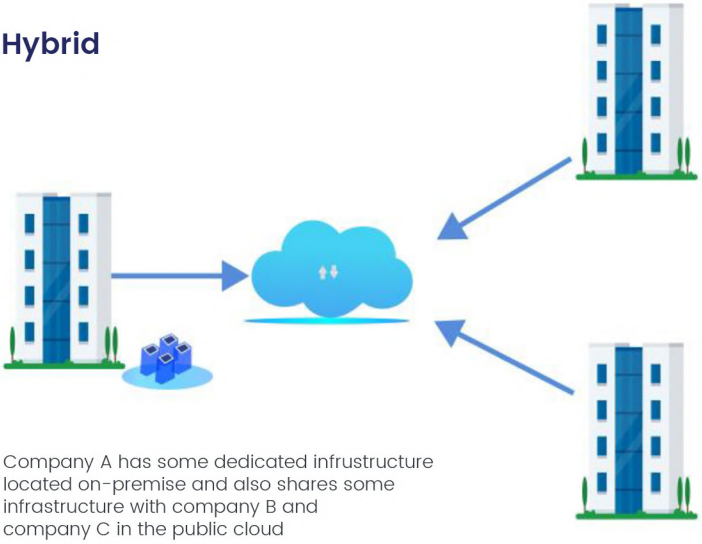
Cloud Infrastructure shared with company A,
company B and company C

## On-Premise



Infrustructure located on-premise
only used by company A

## Hybrid



Company A has some dedicated infrustructure
located on-premise and also shares some
infrastructure with company B and
company C in the public cloud

**Public vs. Private vs. Hybrid Cloud vs. On-Premise**

# 8 of the most common reasons why our clients are moving to the cloud.

## 1. Improve Security

Many of our clients were hesitant to move to the public cloud at first because they believed that their on-premise infrastructure was more secure for hosting their applications and storing sensitive data... In most cases the cloud is more secure due to the number of highly trained professionals and experts that are responsible for running and patching the infrastructure.

## 2. Lower Costs

For many of our clients, cost is often the biggest player in any decision they make - and there is no exception when it comes to IT infrastructure. With the cloud, most services can be optimized for maximum cost efficiency.

## 3. Near Infinite Scale

A good number of our contracts have been launched because our clients have hit bottlenecks with their current IT infrastructure. One of the main benefits that the public cloud has to offer is the ability to quickly provision a near-unlimited number of additional compute resources on demand to run your workloads with minimal delays.

## 4. High Availability

In today's modern IT landscape, a corporation's tolerance for service outages over long periods of time are becoming less and less acceptable. With the cloud it is simple and cheap to set up multiple levels of redundancy so your systems never go offline.




## 5. Never Lose Data

Whether storing data on-premise, in a private or public cloud, we always recommend to our clients that backups should be performed on a regular and frequent basis to a secondary off-site location in the rare event of an unanticipated disaster. Most cloud services provide built in backup and restore functionality with little to no configuration needed.

## 6. Improving Efficiency through Automation

Whether we are working with a startup, mid-size company, or large enterprise, our clients' customers expect continuous improvements in both their products and their services. The cloud's ability to automate most tasks related to IT infrastructure is unparalleled and a massive time saver.

## 7. Drive Innovation

Before jumping into any cloud migration project, we at LefeWare Solutions always try and first build an understanding of what our clients' primary motivations are for moving to the cloud. and lately innovation has been at the forefront... The cloud provides them with the tools they need to be seen as true innovators in their industry.

## 8. Collecting and Processing Data

Organizations are increasingly collecting data from their various IT systems to perform data analysis and train machine learning models, yet don't have the compute needed to do it efficiently... The cloud allows them to easily store the data and rapidly spin up more compute resources to perform the needed processing.

# 1. Improve Security

The primary reason why most of our clients are hesitant to move to the public cloud is because they believe that since they run on-premise or in a private cloud, their applications and data are thus more secure from potential exploits. However, as we will explore, in the majority of incidents, the location of your infrastructure has very little impact on IT security.

## Data Breaches

One of the most potentially harmful, expensive and common scenarios more and more organizations are experiencing as of late, is a data breach. A data breach is a security violation, in which sensitive, protected or confidential data is copied, transmitted, viewed, stolen or used by an individual unauthorized to do so. According to an IBM study titled "Cost of Data Breach Report," *the global average cost of a data breach was $4.24 million, while the average cost of a data breach in the United States has hit an all-time high of $9.05 million in 2021*[2].

**$9.05 M**

**Average Cost of a Data Breach in the United States**



These staggering numbers, and the recent wave of high-profile ransomware attacks (a type of data breach) have many organizations on edge about where they run their IT infrastructure. However, a recent joint study from Stanford University's Professor Jeff Hancock and security firm Tessian revealed that *"88% of data breach incidents are caused by human error, while only 12% were the result of truly malicious behavior"*[3]. In other words, the majority of the time, where you are running your applications and storing your data is irrelevant, if the IT staff responsible for this infrastructure is not well-trained or following proper security protocols.

## Infrastructure Isolation

Another common argument that we frequently hear in favor of private cloud or on-premise infrastructure is that your hardware will be used exclusively by your organization, giving you complete physical isolation and better overall control of your network. On the other hand, in the public cloud, you are using third-party infrastructure at a remote location, shared with multiple customers, where your apps and data are only logically isolated using software, and you are often accessing your resources over public internet.

Those who make this argument might be surprised to find out that even the US Department of Defense (DoD) plans to migrate their existing infrastructure to the cloud over the coming years. Originally the contract code named Joint Enterprise Defense Infrastructure (JEDI) was awarded exclusively to Microsoft Azure, however after being cancelled for numerous reasons, a new contract was issued called Joint Warfighter Cloud Capability (JWCC) that is likely to be awarded to multiple cloud providers. Azure, AWS, IBM, Oracle and Google are all likely to be issued portions of the contract in the DODs plan for a multi cloud strategy. When asked why the Pentagon plans on moving the majority of their infrastructure to the cloud, they stated that it would *"simplify sharing information between headquarters and front lines, enable use of emerging artificial intelligence technology, and gain a tactical edge, or advantage, on the battlefield"*[4].

**$10B over 10 years**

**The amount Microsoft Azure would have been paid by the DOD for project JEDI**

It is important to note that most major cloud providers have data center locations dedicated to these types of top security projects, where they must meet very specific government security and compliance requirements. Furthermore, not any organization will be granted access to these types of special datacenters as they are usually only meant for certain types of entities. For example, "Microsoft Azure Government" requires you to submit an application to verify that you're either a US federal, state, local, or tribal entity, or their partner. The argument could be made that this is not a true public cloud in the sense that only certain government entities are allowed, however, it is still a third party running the infrastructure that has a staff of dedicated security and infrastructure experts.

> If the Pentagon is willing to deem major cloud providers like Azure and AWS secure enough for matters of national security, chances are that the cloud is secure enough for your organization too.

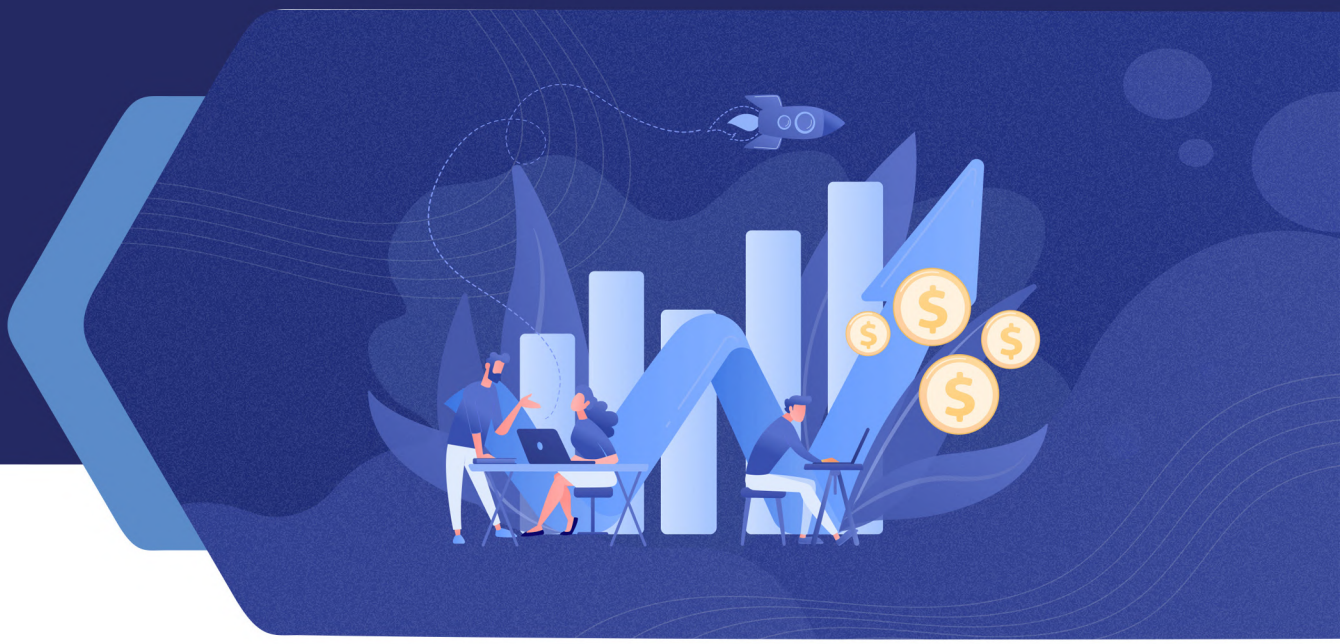
## IT Staff

When evaluating a move to the public cloud based on security, we at LefeWare Solutions strongly believe that the decision should ultimately come down to how well-trained and disciplined your organization's IT staff is. Oftentimes, this means having a dedicated team of IT professionals on staff whose only concern is managing your datacenter. These IT professionals should have deep technical knowledge of the hardware and operating systems that you use, should stay constantly up-to-date with all recent security vulnerabilities, and perform regular patching of these vulnerabilities.

Most small- to mid-size companies we work with do not have this kind of expertise available to them in-house, which results in them not being able to keep up with modern security practices, and therefore vulnerable to many common exploits. Large organizations like Microsoft, Amazon, and Google are able to hire and retain industry-leading experts to maintain their infrastructure. By leveraging their Infrastructure as a Service (IaaS) and Platform as a Service (PaaS) solutions, you are often able to obtain the same level of confidence as an in-house IT staff or your local private cloud provider.

# 2. Lower Costs

For many of our clients, cost is often the biggest player in any decision they make - and there is no exception when it comes to IT infrastructure. In a recent OpsRamp survey of 500 IT professionals *"94% of responders claimed that moving to the cloud would reduce infrastructure setup and maintenance costs with 90% of responders expecting to shave at least 10% or more on their IT budgets, and 55% expected to save 30% or more"*[5].



# 30%

**Majority of IT Professionals Expect to Save 30% or more by Switching to the Cloud**

## Pay-as-you-go (Consumption) Plans

Amongst the most significant IT cost cutting advantages of moving to the cloud, are the resource consumption-based pricing plans that all big cloud providers are now beginning to offer. In other words, only pay for what you use. For many organizations, predicting monthly server compute needs is nearly impossible to get right, and almost always results in organizations over-provisioning their hardware and bringing about unnecessary cost, or under-provisioning and bringing about performance degradation (which can also result in cost). Resources that offer consumption-based pricing plans often fall into the "Serverless" category, meaning that the cloud provider is automatically allocating and deallocating machine resources on demand.

For example, Google Cloud Platforms' (GCP) "Cloud Functions" are priced according to how long your function runs, how many times it's invoked and how many resources you provision for the function. In other words, you pay nothing when your function is idle. For applications that might only need to be used infrequently, this is an excellent way to cut cost. This is only one example of the many consumption-based resource plans that AWS, Azure and GCP offer.

## Predictive Cost

Another benefit of running on the public cloud is that cost is more predictive. When a server crashes, you're not suddenly left with an unexpected surprise expense related to replacing the hardware. In fact, most cloud providers abstract away hardware completely, meaning that even when a cloud server that is running your apps or hosting your data fails, you likely won't even know it died, as it is immediately replaced with a functioning one.

In the event that you do experience any downtime, most cloud services have Service Level Agreements (SLAs) that guarantee a certain percentage of uptime – and if that SLA commitment is breached, then a penalty or credit fee is issued in order to compensate for the commitment breach. For AWS's Elastic Compute (EC2) service, in the event any Single EC2 Instance does not meet the Instance-Level SLA, you will be eligible to receive a Service Credit as described below:



| Instance-Level Uptime Percentage | Service Credit Percentage |
|---|---|
| Less than 99.5% but equal to or greater than 99.0% | 10% |
| Less than 99.0% but equal to or greater than 95.0% | 30% |
| Less than 95.0% | 100% |

## Hidden Costs of an On-Premise Datacenter

Despite the many obvious cost savings that come from migrating to the cloud, when our clients see their first month's public cloud bill, they tend to be a bit surprised, as we had sold them on the idea that moving to the cloud would save them money. However, there are many hidden costs that cloud newcomers who

have recently moved away from an on-premise data-center forget to factor in to their calculations.

When setting up an on-premise data center you will have to purchase all physical hardware and networking equipment needed to run your servers. The recent chip shortages of 2020 and 2021 have only served to drastically exuberate these prices. Add in the licenses you will likely need to purchase for certain software's such as VMWare, Windows Server or SQL Server and you are left with some quite significant setup costs. For startups who might lack that early capital, or businesses that need to rapidly scale their infrastructure, these upfront costs likely aren't feasible, thus making a very strong argument for a private or public cloud provider where the hardware already exists.

In addition to upfront costs, there are also several monthly costs many forget to factor in to their monthly overhead of running a datacenter. Although these monthly costs do contribute to your monthly cloud bill, you can be assured that most popular cloud providers are running their datacenters at peak efficiency to save you, the consumer unnecessary costs. Some of the monthly costs include the cost electricity to run and cool your hardware and your internet service provider costs if your servers are being exposed over the internet. However, the most expensive cost with running an on-premise data center is by far the IT Staff you will need to hire to manage it. According to Glassdoor, "*in 2021, the average cost of an IT staff in America is $72,532*"[6], and chances are you will need several highly-trained IT staff to properly manage an in-house data center. The objective of moving to the cloud is not to eliminate your IT staff completely, but rather to free them up from having to work on trivial tasks, and instead focus on higher value tasks.

**$72,532**

**Average Cost of an IT Staff in the United States as of 2021**



# 3. Near Infinite Scale



One of the main benefits that the public cloud has to offer is the ability to quickly provision a near-unlimited number of additional compute resources on demand to run your workloads with minimal delays – otherwise referred to as “scaling”. Scaling comes in two forms:

- **Scaling up:** adding more resources (CPU cores, Memory, etc.) to an existing server.
- **Scaling out:** adding additional servers to help balance the server load.

The cloud is an excellent solution for handling both types of operations.

For most of our clients, the ability to scale up/out quickly during periods of high demand is absolutely critical for their day-to-day business operations. After all, being unable to stand up the servers they need within a reasonable amount of time, can result in application slowdowns, unexpected crashes, loss of sales, and the inability to onboard new customers. Each of these problems can result in serious financial losses and harm to their reputation.

## Time to Provision Hardware

With an on-premise data center, it can take anywhere from hours to days to stand up a single server – and if you don't already have the right hardware on hand, you may need to order parts, add with it the occasional hardware part shortages and shipment delays, and it could turn out to be weeks before a new server is stood up. Even if you leverage a private cloud, having them provision another dedicated server rack isn't usually as easy as clicking a few buttons, and usually requires a series of back-and-forth emails or calls. Of course, you could have servers that are already set up in your data center that are always ready to go, but then you are likely paying for hardware and taking up space in your data center for resources you might never actually use.

On the other hand, most major public clouds allow you to "rent" infrastructure in a pay-as-you-go basis, known as Infrastructure as a Service (IaaS). IaaS allows you to easily spin up resources such as Virtual Machine (VM) in seconds to minutes, or add additional resources to an existing server with little-to-no interruption. The process of adding these resources is also extremely simple in the cloud and often requires just a few simple buttons clicks or commands and can even be automated using Infrastructure as Code (IaC) scripts.

## Autoscaling

Another useful feature that many public cloud providers offer (and one that we suggest to all our clients) is infrastructure autoscaling solutions. As the name suggests, infrastructure autoscaling allows your systems to listen and react to live ongoing events by automatically provisioning additional resources once a certain threshold is reached. For example, you might create a trigger in your cloud provider that automatically provisions a new VM when all current VMs under your server farm or Kubernetes cluster reach 80% CPU usage.



One of our largest customers in the healthcare industry recently approached us wanting help to accommodate fluctuating traffic volume on their website and meeting customer demands. Their on-premise data center didn't have the ability to scale up and down in real time based on customer demand and thus they wanted to explore AWS public cloud and its automatic scaling solutions. We ended up landing on Amazon Elastic Kubernetes Service (EKS) due to its ability to automatically scale resources up or down to meet changing demands, a functionality that would otherwise require extensive human resources to perform manually. The benefits they saw:

- *Improved hourly patient file processing by 200%*
- *Reduced site outages by 65%*
- *Saved 42% annually using EKS over original plan to migrate to EC2 severs*
- *Decreased time to procure and setup infrastructure from 4–5 weeks to under 30 seconds*
- *Improved network security by running a private cluster and adding firewalls*[7]

**200%**

**Increase in hourly patient file processing by using EKS**

**65%**

**Reduced site outages by 65%**

## Seamless Scaling using Serverless Solutions

Getting your resources to scale flawlessly can be extremely challenging and time consuming for your team, which is why most cloud providers offer preconfigured serverless resources that completely abstract underlying infrastructure to run your code and manage your data at peak efficiency. By moving to the cloud and utilizing serverless technologies, you might not ever need to worry about scale again, allowing you to sleep well at night knowing your resources will continue to run smoothly without any intervention. Of course, when using autoscaling serverless solutions in the cloud, it's always a good idea to set an upper limit to the amount your resources can scale, thus ensuring you don't get surprised with a high bill for compute time.

Recently, LefeWare Solutions completed a project for a client in the manufacturing industry that wanted to begin collecting and analyzing data from various embedded systems they had in plants all across North America. This data needed to be processed, and sent to the correct system for reporting, dashboards and other analysis. The client's development team was small and most of them were embedded systems engineers with little knowledge of infrastructure management or the cloud. We decided to use Azure Functions, a serverless, event driven solution that offered a consumption plan that scaled as the number of requests grew. With Azure Functions, all the underlying infrastructure and scaling capabilities were completely abstracted away, allowing the engineers to focus more on working on new features and less on infrastructure management. *At the time of writing this, the new system is in 13 plants, and the serverless cloud functions are able to process over 100,000 incoming events from the Azure IoT Hub per hour.*[8]

**~100,000/hour**

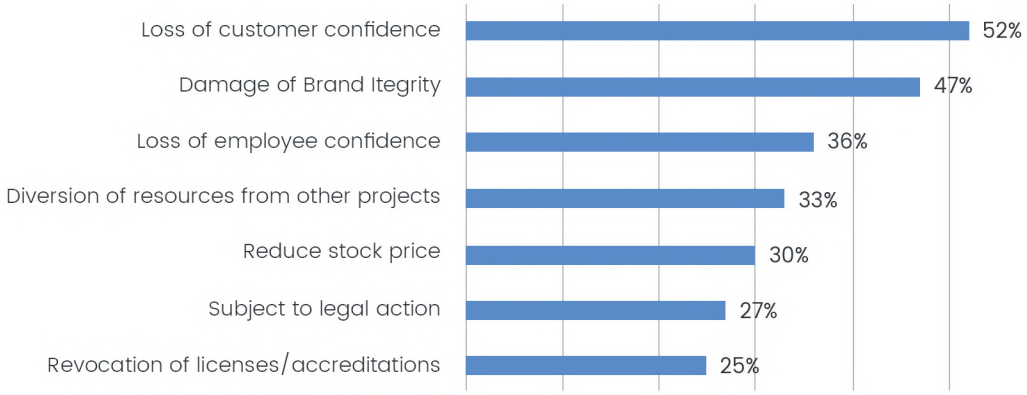
**The number of incoming events from embedded systems needing to be processed**



# 4. High Availability

In today's modern IT landscape, a corporation's tolerance for service outages over long periods of time are becoming less and less acceptable. In other words, their maximum tolerable period of disruption (MTPD) is expected to be reduced year over year. Although hardware failure is just one cause for an outage, one thing is for certain, whether hosting your apps and data on private, public, or on-premise servers, all servers are bound to fail eventually. If you're organization is not prepared for these outages, it is likely to cost you in many different ways. According to Veeam's 2021 Data protection report, *"In 2020, the average downtime for an outage was 79 minutes with the cost of downtime now averaging $84,650 US per hour"*[9].

**Veeam's 2021 Data protection report – Survey of impacts organizations could experience from downtime.**






Even though your on-premise datacenter might have uninterruptible power supplies (UPS) configured in case of power outages, a load balancer to serve content from multiple servers in case one faults and your data is copied and stored on multiple hard disks for redundancy, it's still very difficult to beat some of the high availability features that big cloud providers have to offer.



# Cloud Redundancy Levels

If your organization has applications that are truly “mission critical”, such as an emergency response system, or a piece of network connected medical equipment where people's lives might be on the line, then ensuring that your apps are always online and your data is available is crucial. The cloud’s ability to protect against local and regional outages is unmatched. All big cloud providers (AWS, Azure, and GCP) have datacenters around the globe that offer several tiers of redundancy in case of an isolated outage in a particular region:

### Local Redundancy

Protect against failures within a single datacenter in the event of hardware malfunctions or software update cycles.

### Zonal Redundancy

Protect against datacenter failures within a single region in the event of power, cooling, or networking issues.

### Regional Redundancy

Protect against entire-region failures beyond a single region in the event of a tornado, earthquake, or other large-scale events

With the cloud, configuring these levels of redundancy is usually quite straight forward and the costs are reasonable if configured correctly. With each form of redundancy your organization implements, the uptime of your applications and the durability of your data are drastically improved.

# Service Level Agreements

When a cloud provider's customer hosts their mission critical applications and data on the provider's cloud platform, they are entrusting this third-party provider to manage their infrastructure needs with the utmost care. The customer expects an unparalleled level of reliability and durability, unmatched by even the most well-run on-premise datacenters. Cloud providers like AWS, Azure and GCP are often so confident in their abilities to maintain a certain percentage uptime, that they will offer their customers well defined service level agreements (SLAs) for many of their cloud resources. If for any reason the SLA commitment is breached, then a penalty or credit fee is issued by the cloud provider to the customer in order to compensate for the commitment breach.

Below is a table showing the SLAs for managed virtual machines running in each major cloud provider:

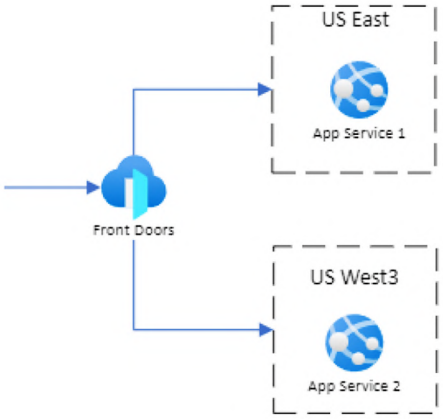
| Resource | Single Instance Uptime | Two Instances Spread Over Availability Zones Uptime |
|---|---|---|
| AWS EC2 | 99.95% | 99.99% |
| GCP Compute Engine | 99.95% | 99.99% |
| Azure VM | 99.95% | 99.99% |
| | | |

As you can see, no matter which cloud provider you end up choosing, when configured correctly the chances of a service outage or data loss are astronomically low.

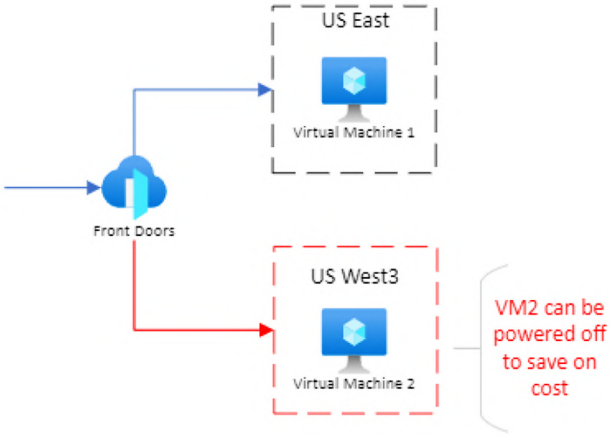
# Modern High Availability Architecture

At LefeWare Solutions, we recently worked with a client in the healthcare industry that was operating across the continental United States and wanted us to ensure their mission critical patient file processing application never had any downtime. We decided on an "Active/Active" architecture, using Azure Front Door, where requests would be load balanced and served to two separate regions (West US 3 and US East) depending on where the end user was closest to. If one the two regions became unavailable, it would be taken out of rotation and all traffic would be directed to the functioning region.

### Active/Active Architecture




### Active/Passive Architecture




### Active/Active vs Active/Passive Architecture

In this particular case, paying for infrastructure in two separate regions was justified as it also resulted in less network latency for incoming requests by serving from the closest region. We also decided on Azure CosmosDB as the persistence store due to its ability to replicate data across multiple regions and ease of setting up failovers in the event of an outage. For projects with a tighter budget, where running and paying for two active instances of the app at all times might not be justified, an Active/Passive architecture might be more appropriate. In this case only one region is processing requests, while the other region is on standby in the rare event a failure occurs.



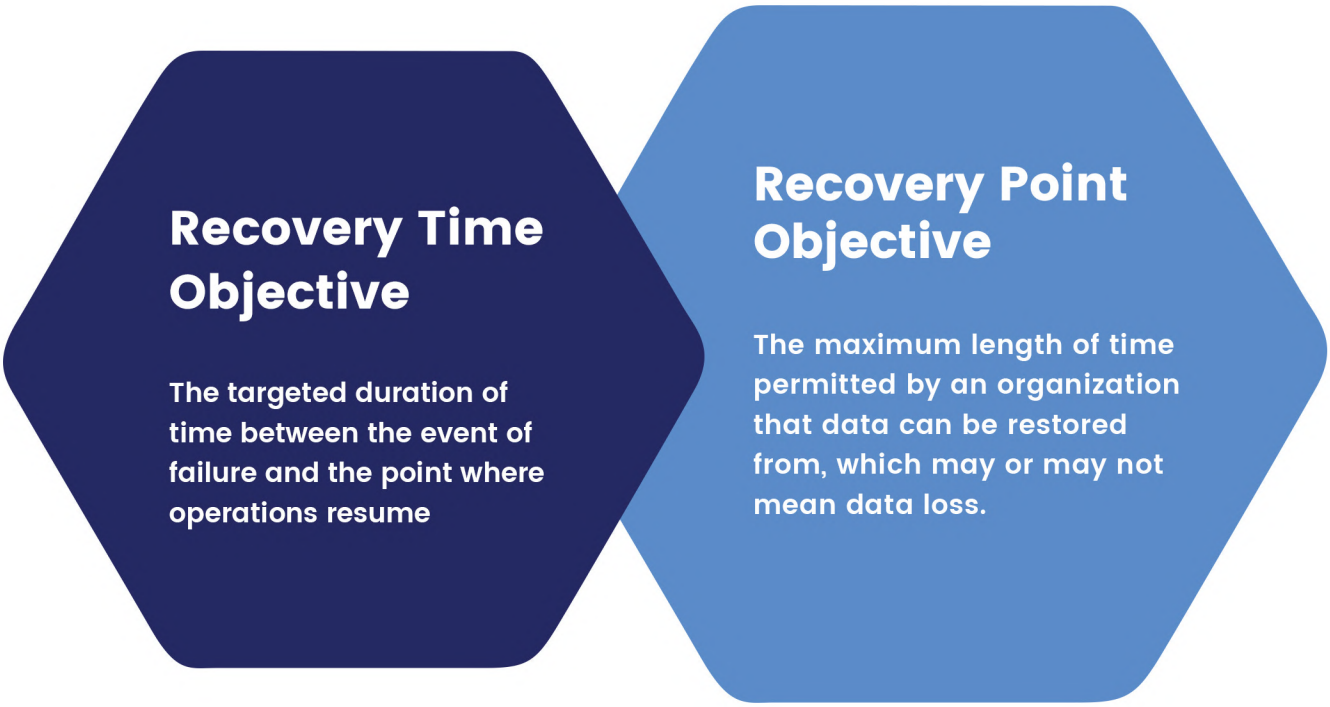
# 5. Never Lose Data

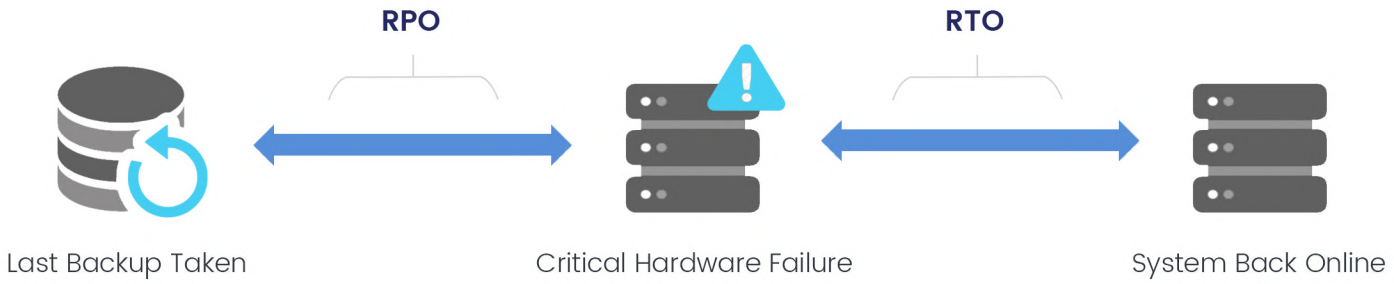
Whether you're storing your data on-premise, in a private or public cloud, we always recommend to our clients that backups should be performed on a regular and frequent basis to a secondary off-site location in the rare event of an unanticipated disaster. In a survey of small- and mid-sized businesses conducted by UNITRENDS, *"45% said their business had experienced data loss, with an average of $9000 in data recovery fees. 54% of them said the data loss was due to a hardware failure"*[10].

**54%**

**The amount of Data Loss Caused by Hardware Failure**

When data loss does occur, it's important that your organization has a process in place to be able to quickly recover from these events and avoid further downtime. The processes are usually part of a "Business Continuity Plan" and require defining the organization's Recovery Point Objective (RPO) and Recovery Time Objective (RTO).





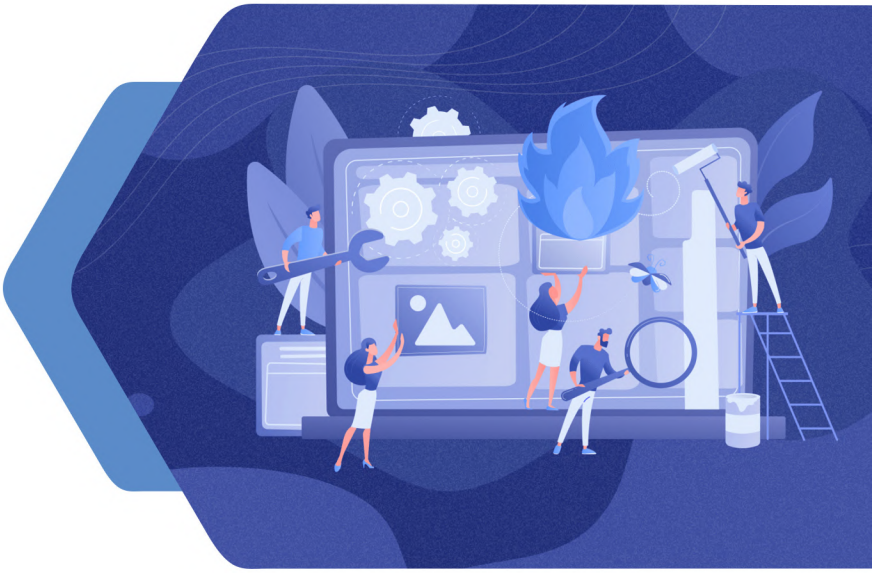

**RPO vs RTP Time in the Event of a Critical Failure**

## Backup as a Service

Despite the cloud's amazing forms of redundant data storage, data redundancy will not protect your data from ransomware attacks or human error, in the event that records are either intentionally or unintentionally deleted or corrupted. This is why all of the leading public cloud providers offer several backups as a service (BaaS) options that are built into existing resources, such as Virtual Machines and Databases. These preconfigured cloud solutions are likely far superior to anything that most organizations can set up in their own on-premise data centers, and can serve to save your team the trouble of having to set them up.

For example, the default option of Azure's Cosmos DB is to automatically take backups of your data every 4 hours. This means that if your organization is using Cosmos DB with the default settings, in the unlikely event that an attacker somehow gains access to your data, encrypts or corrupts the data as part of a ransomware attack, the Azure support team will be able to restore your data to a point from 4 hours prior. This value can be adjusted if your organization requires a shorter RPO

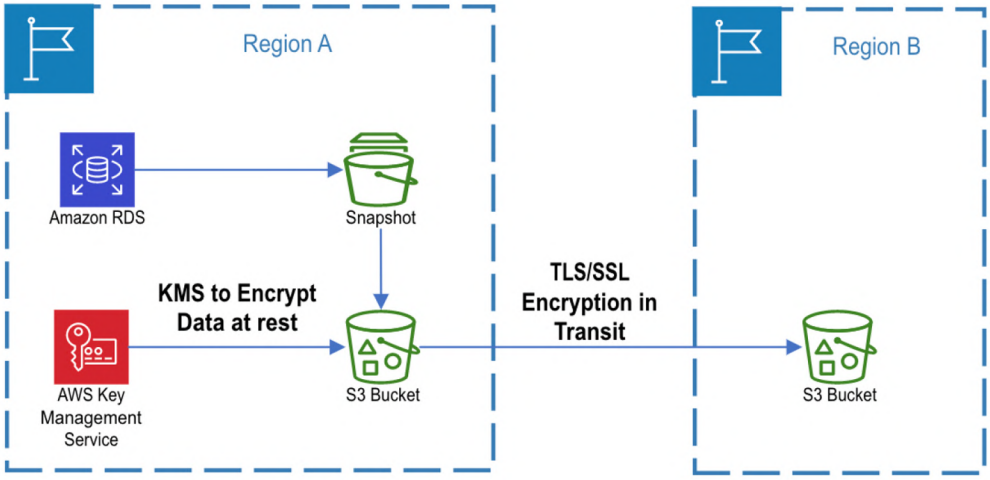
Of course, you could set up similar backup processes for your data in your on-premise servers; however,performing frequent backups can be extremely taxing on your hardware and network resources,depending on the amount of data – and frequency – of your backups. The cloud's ability to easily scale can enable you to store months' or even years' worth of backups, allows you to easily perform backups and restore your services without affecting performance and availability.



# Backup Security

Just as it is important to backup your data at regular intervals, it is also extremely important to secure your backups. The cloud offers some of the best security options and policies around ensuring that your data is safe and secure, and follows proper industry compliance standards for data storage.

Amazon Relation Database Service (RDS) backups and snapshots are automatically encrypted with industry standard AES-256 encryption algorithm at rest using AWS Key Management Service (KMS). These backups are stored separately in a storage service, using TLS/SSL encryption in transit and are globally replicated for resiliency against regional disasters.

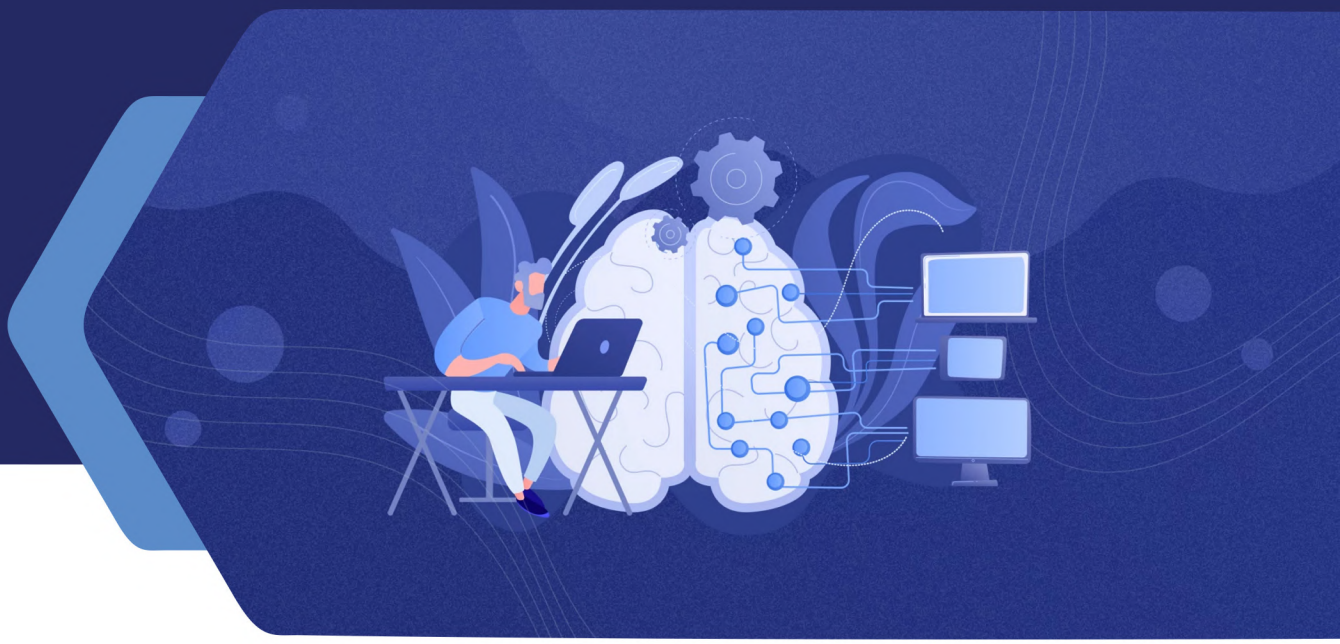



**AWS RDS Backups**

# Disaster Recovery

To decrease the RTO, many cloud services have automated failover and failback processes built into their resources to further enable business continuity. This includes things like built-in alerts to be notified of outages, automated failover to a secondary region, and easy data restoration from a previous point.

A Google Cloud Provider SQL instance can easily be configured to be made up of a primary instance and a standby instance within a region. Through synchronous replication to each zone's persistent disk, all writes made to the primary instance are replicated to disks in both zones before a transaction is reported as committed. In the event of an instance or zone failure, the standby instance becomes the new primary instance. Users are then rerouted to the new primary instance. This process is called a regional failover and can be done almost instantaneously, resulting in very short RTOs.



# 6. Improving Efficiency through Automation

Whether we are working with a startup, mid-size company, or large enterprise, our clients' customers expect continuous improvements in both their products and their services. The ability to release new features and apply patches to their applications and to their IT environment in a quick and efficient manner is critical for continuous business growth.

## The Rise of DevOps

If you are in the IT space, then you have no doubt heard the term "DevOps" and are likely already well aware of some of the benefits it can bring to an organization. DevOps is a methodology practiced by many innovative companies which uses a series of tools and processes centered around automation in order to streamline common IT activities and shorten delivery times in releasing software. Implementing DevOps methodologies in your organization ultimately frees up your developers and IT staff from having to perform trivial, repetitive tasks, so that they can instead focus on tasks that provide higher value to your organization. In Puppet's yearly State of DevOps Report, *"an overwhelming majority of organizations that identify as mid- and high-evolution groups (<97%) agree that DevOps automation is critical to meet the rising pressure to innovate faster in their respective industries"*[11].

# 97%

## Agree DevOps is Critical To Innovate Faster

Now obviously, you don't need to be in the cloud to embrace automation and DevOps methodologies. Automation can be achieved by simply writing scripts to perform common tasks, such as automated backups of your on-premise databases, snapshots of your VMs, or compiling and deploying a piece of software to one of your on-premise servers. But why go through the trouble of writing these automation tools and scripts when most cloud providers already have these features built-in? By default, most cloud-based databases already offer built-in automated backups, cloud-native technologies like Docker and Platform as a Service make VM snapshots a thing of the past, and cloud-based tools like Jenkins and Azure DevOps Pipelines make compiling and releasing your software a breeze.



# Automating Software Deployments to the Cloud

Adding automation around the process of compiling and releasing software is a popular choice amongst many organizations. Without automation, releasing software usually entails a series of manual steps performed by your staff that are prone to errors. Many choose to automate these steps by developing a series of scripts that are run in a sequential order, otherwise known as a "pipeline". Cloud-based tools like Azure DevOps Pipelines and Jenkins offer many prebuilt pipeline solutions for automating the compilation and deployment of software that will run automatically when code is checked in. This process is known as "Continuous Integration/Continuous Deployment" (CI/CD Pipelines).

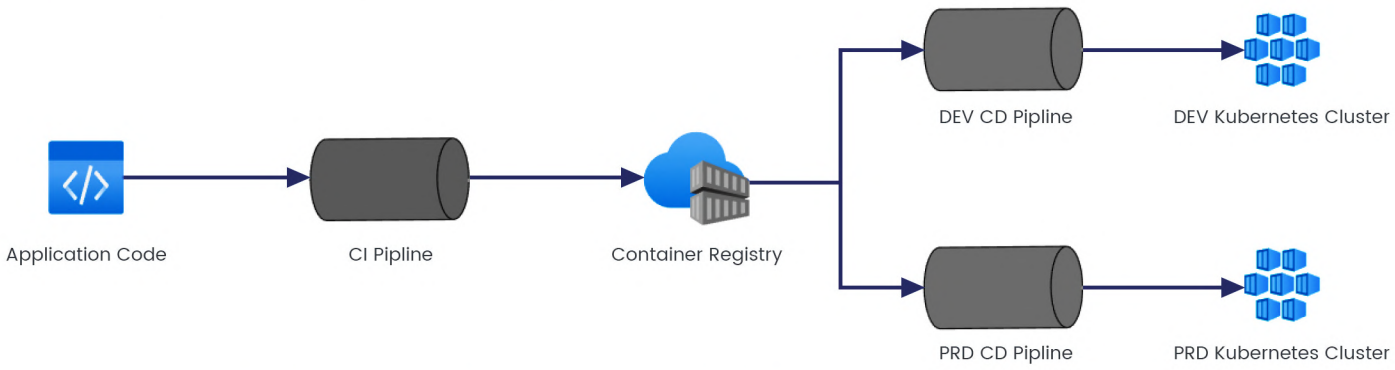
In the 2021 "Accelerate State of DevOps Report" by the DevOps Research and Assessment (DORA) team at Google Cloud, they surveyed over 32,000 IT professional and created software delivery and operational (SDO) performance metrics to place the teams into four performance clusters: elite, high, medium, and low. It was found the *"performers in the elite category deployed code 973x more frequently, and recovered from incidents 6570x faster than those in the low categories"*[12]. This is of course in large parts possible by automating their software code deployments using CI/CD pipelines

**6570x**

**The frequency of times faster To Recover from Live Incidents**

Many of these tools have offerings that allow you to run these pipelines in your own data center; however, unless there is a good reason for doing so (i.e., strict security regulations), then this will result in more infrastructure for your team to manage, which ultimately defeats many of the benefits of automation such as wasting time on trivial tasks.

At LefeWare Solutions, we bake in DevOps methodologies to every single project we work on. For example, any time we are building a web application for a client, we set up a build pipeline to compile the code, run automated tests, perform code analysis and publish the artifacts. We also set up release pipelines for releasing to various cloud platform environments where the application runs. Of course, this results in a small upfront investment to get these pipelines set up and running, but the overall benefits over the project's lifespan always pay off several times over.




**CI/CD Application Deployment to Multiple Cloud Environments**

## Automating Cloud Infrastructure Deployments

Automation can be taken one step further, with many of the public cloud providers offering the ability to write code to manage and provision your cloud-based infrastructure. This is known as "Infrastructure as Code" (IaC), and the potential of using it to automate processes around your infrastructure needs is truly profound. Each cloud provider has their own default IaC language, but there are a few languages – such as HashiCorp's Terraform, for example – that have plugins for all the popular cloud providers (AWS, Azure, GCP) which make it possible for your team to only have to learn one IaC language. This is great for companies who might have a multi-cloud strategy.



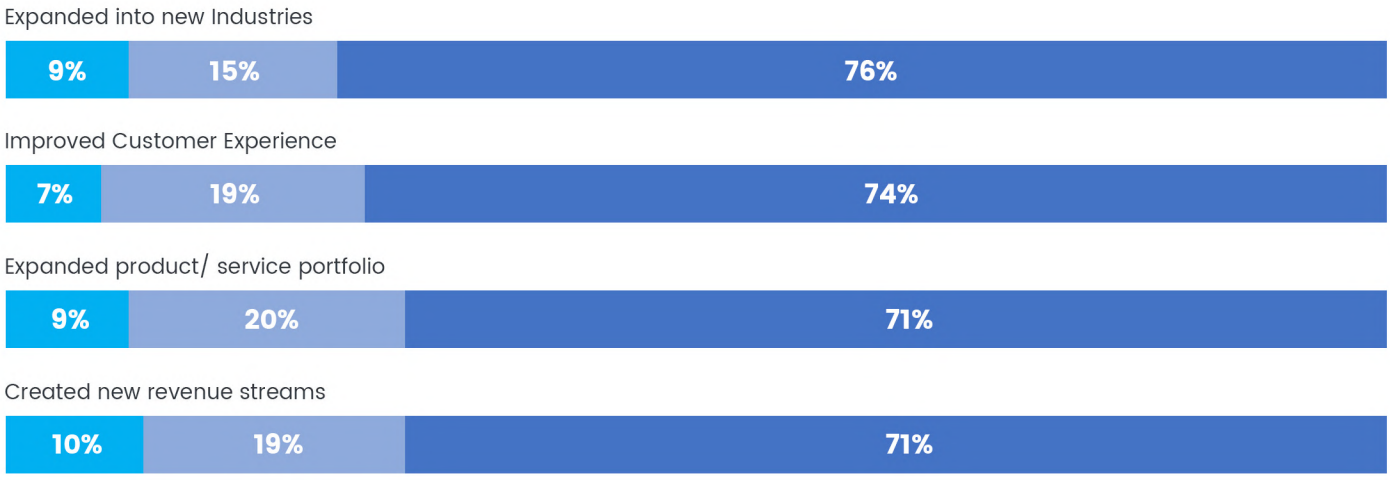

**IaC Cloud Pipeline Deployment**

We recently had a client here at LefeWare Solutions that needed to spin up a total of 12 virtual machines to perform some heavy data processing in parallel. It only took us a handful of hours to write the IaC code and set up an IaC pipeline to automate the deployment of all 12 VMs into Microsoft Azure's public cloud. Once these processing tasks were complete, we were easily able to deprovision the VMs, since they were no longer needed. The amount of physical hardware and work that would have been needed to set up these exact same 12 VMs in an on-premise data center would have been extreme.

# 7. Drive Innovation

Before jumping into any cloud migration project, we at LefeWare Solutions always try and first build an understanding of what our clients' primary motivations are for moving to the cloud. Only a couple of years ago, the main reason why our clients were moving to the cloud was to cut costs and streamline their IT infrastructure management. More recently, however, many have stated that they are moving to the cloud out of necessity, so that they can continue to innovate and stay ahead of their competition.

IBM recently surveyed over 1,000 business leaders globally, asking them about their most successful cloud initiative. Based on their responses, it was shown that the top reasons these businesses are using cloud is to *"move into new industries, transform customer experiences and forge new business models with tremendous results"*[13].




**Innovative Outcomes Achieved by Most Successful Cloud Initiatives**[7]

## What does it mean to be an Innovator?

But what exactly does it mean to "innovate in the IT space"? Is it just another buzzword that companies like to throw around for bragging rights? Or has innovation become one of the most important arguments for moving to the cloud? Below is an excellent quote by author and speaker Stephen Shapiro explaining how he describes innovation:

*Very simply put, innovation is about staying relevant. We are in a time of unprecedented change. As a result, what may have helped an organization be successful in the past could potentially the cause of their failure in the future. Companies need to adapt and evolve to meet the ever-changing needs of their constituents*[14].



## Innovation Creates Disruption

The movie rental store Blockbuster is a prime example of a company who failed to innovate and adapt to the ever-changing landscape of the digital age. At its peak in the Late 90s, *"Blockbuster owned over 9,000 video-rental stores in the United States, employed 84,000 people worldwide, and had 65 million registered customers"*[15]. Like so many other businesses, however, they failed to adapt and listen to their consumers. Their approach that had worked for them at one time no longer worked for them anymore, and thus the movie rental business model became a prime candidate for disruption.

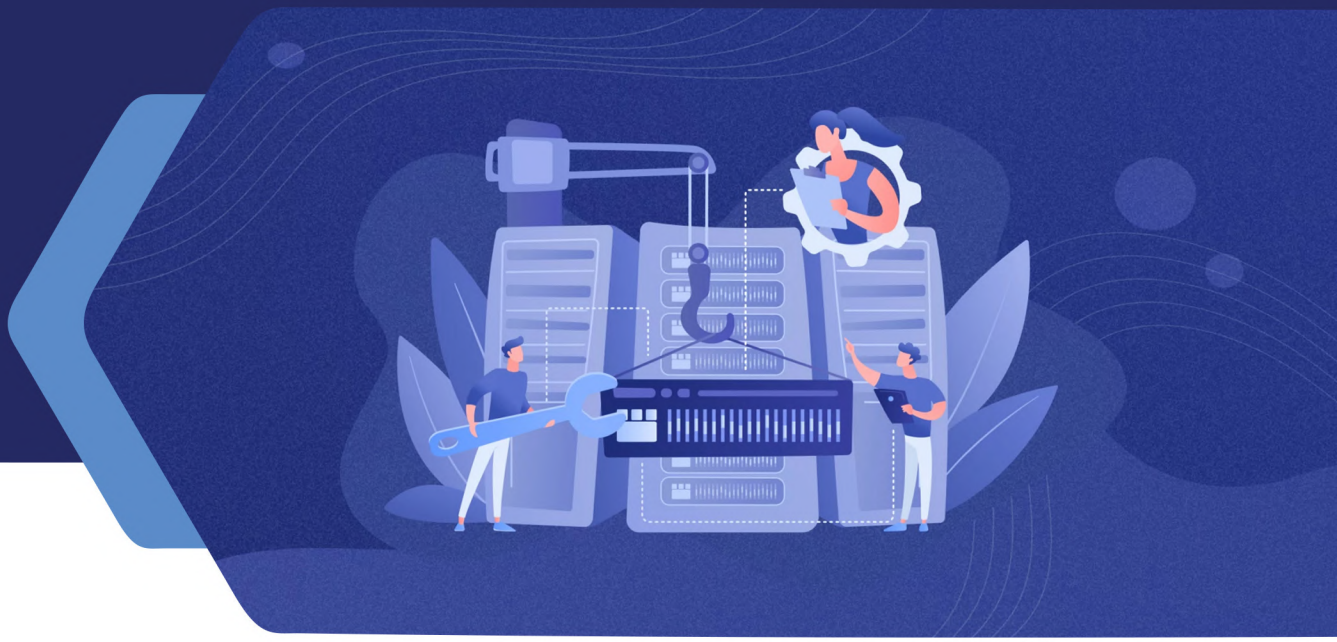
Soon enough, Netflix came along offering two extremely innovative features that consumers had long been asking for: a subscription-based model that gave customers the ability to watch unlimited TV shows and movies (without late fees), and a platform for watching movies online so that they no longer had to leave the comfort of their home.

## Leveraging the Cloud to Drive Innovation

Netflix's user count and content library both grew extremely quickly, and as a result, they were unable to scale their on-premise infrastructure to handle the sudden influx – sometimes resulting in multi-day outages. But because Netflix were innovators, and wanted to continue to grow and bring their customers a reliable service, they became an early adapter of the public cloud, beginning their cloud migration journey to AWS public cloud in 2008.



The same type of disruption is currently occurring in many other industries, with *"cloud adoption expecting to completely revolutionize the banking, retail, and manufacturing sectors in 2022"*[16]. The public clouds' abilities to easily scale and process enormous amounts of data are simply unmatched, and any company that wants to experience exponential growth should strongly consider making the leap. Just because others in your industry haven't adopted the cloud to its full potential yet, doesn't mean that you shouldn't lead the charge. If you don't, then you could very well end up as yet another cautionary tale just like Blockbuster!



# 8. Collecting and Processing Data

Organizations are increasingly leveraging the cloud to collect data from their various IT systems to perform data analysis and train machine learning models. These practices allow them to gain useful insights into their operations, drive decision making, and provide better experiences for their customers. For example, an e-commerce site might be collecting data on a customer's shopping habits to provide recommended items, and a manufacturing company might collect and aggregate data from their various factory machines in order to run them more efficiently and save on cost.

## Why You Should be Collecting Data

Even if you think your organization might not yet be in a place to leverage and analyze data, you should still begin collecting it. A 2017 article by the economist titled *"The world's most valuable resource is no longer oil, but data"*[17] goes into detail on how big tech companies like Google and Facebook leverage their customer and systems data to maintain their dominance and make billions of dollars in profits. Whether you believe that big tech collecting and storing massive amounts of data from their customers is ethical or not, chances are that your competitors are leveraging the cloud to collect and analyze data to continuously innovate and gain a competitive edge over you.

## Why You Should be Collecting Data

Even if you think your organization might not yet be in a place to leverage and analyze data, you should still begin collecting it. A 2017 article by the economist titled, *"The world's most valuable resource is no longer oil, but data"*[18] goes into detail on how big tech companies like Google and Facebook leverage their customer and systems data to maintain their dominance and make billions of dollars in profits. Whether you believe that big tech collecting and storing massive amounts of data from their customers is ethical or not, chances are that your competitors are leveraging the cloud to collect and analyze data to continuously innovate and gain a competitive edge over you.

> *The world's most valuable resource is no longer oil, but data.*

However, similar to oil, if you do not have the ability to refine the data, then it is just a commodity. These days data storage is cheap and small, so although we don't advise it, storing all of your collected data in your on-premise data center is likely possible. Where many organizations struggle is having the financial means, the in-house expertise, and the physical space needed for performing the processing of all of this data.




**May 6th 2017 issue of "The Economist"**

## Data Refinement to Gain Useful Insights

One of the most popular, and potentially profitable, use cases for an organization's collected data is to use it in training machine learning models, so that it can be used in operations such as predictive analysis, image/video processing, sentiment analysis, etc. The computation needed for training these models (refine) is extremely computationally heavy, and often requires the use of many Graphical Processing Units (GPUs).

GPUs are an optimal tool for training AI and deep learning models, as they can process multiple computations simultaneously, thus resulting in shorter and more efficient processing. Unfortunately, GPUs are extremely popular and in-demand at the moment, as they are also a popular choice amongst gamers and bitcoin miners. Nvidia, one of the top GPU producers in the entire world, warned earlier this year *"that the great GPU shortage would last throughout 2021, and now the company expects supply issues to continue well into 2022"*[19] This has resulted in a nationwide GPU shortage that has driven up costs significantly. In other words, many companies can't afford the infrastructure they need to perform these types of computations on-premise. So, what's the solution?

Throughout the latter part of the 20th Century, many innovative companies and academic institutions would have a central mainframe computer, in which individuals could reserve time slots to "remotely" connect to the mainframe through their department's terminals (remotely, as in usually still within the same office building or campus). Users of the mainframe could share the mainframe's infrastructure to perform computationally-heavy research or data business processing. Sharing a mainframe was a great way for these organizations to cut costs and save on space, as the typical IBM mainframe at the time occupied anywhere from *"2,000 to 10,000 square feet (600 to 3000 square meters)"*[20].

> *The great GPU shortage would last throughout 2021, and now the company expects supply issues to continue well into 2022*

In many ways, early mainframes offered similar benefits that today's cloud providers offer to businesses. Instead of every department/faculty having to have their own mainframe (in this case, a data center), they could share resources similar to how we share resources in the public cloud, thus saving space and cost. The cloud is essentially an evolution of the mainframe computer, which brings with it several key advantages.

One such advantage of today's modern public cloud usage over traditional mainframeis that there is usually no need to over reserved time slots in order to use the computing resources, thanks to the cloud's vast array of resources that are available at any given time. Another advantage is that most public clouds are not "centralized" like a mainframe; rather, the cloud computing resources are "distributed" across multiple regions, giving you additional redundancy against failures and outages.




**IBM System 360 Mainframe**

## Using the Cloud for On Demand Data Processing

A business being able to allocate resources on demand to perform some heavy computations associated with data analysis and then deallocate those resources once the task is done has proven to be a monumental leap in innovation. This concept has made it possible for many early-stage AI-based startups to materialize in the last decade, and has allowed massive enterprises to easily scale up in order to analyze the petabytes of data that they might collect each month. The cloud also offers some of the most advanced and powerful machines that money can buy for performing your data processing in an efficient, time-effective manner – something most startups certainly couldn't afford. On top of this, Azure, AWS, and GCP all offer several Machine Learning as a Service (MLaaS) options for many common scenarios and, in some cases, save you the trouble of having to hire hordes of data scientists.

# Should we make the Jump to Cloud?

Every organization is going to have its own set of unique infrastructure needs for running their apps, storing their files, and hosting their data that is defined by what their risk tolerance is, how innovative they are, how many customers they're serving, how much data they're processing in a day, and a million other factors. There's no one size that fits all when it comes to cloud computing, and there is certainly no reason why you should jump headfirst into moving your entire IT environment to the cloud all at once.

Here at LefeWare Solutions, we work closely with many of our clients who use hybrid cloud solutions, in which they might maintain certain legacy applications on their on-premise servers, all while also developing new applications in the cloud, or perhaps only migrating their mission-critical apps to the cloud, for better scale and redundancy. If there is a single commonality that exists amongst all the organizations we have helped move to the cloud, it's that they all wished that they had started the move sooner. When we reach out to them after the migration is complete, they almost always agree that the cloud has eased their lives significantly and helped them to grow their businesses exponentially.

# Need help Scaling up your Infrastructure or Migrating to the Cloud?

LefeWare Solutions' team of cloud migration experts has helped numerous organizations to leverage the cloud in order to pass industry-specific security assessments, cut costs by up to 40% by switching to more cost-efficient auto scaling cloud infrastructure, and increase team agility by 32% using proven DevOps automation strategies. Curious about how you can obtain similar results? Visit our webiste https://LefeWareSolutions.com and set up a free consultation with us now, and we will help assess your current needs.

## References


1. Accelerating Forward: The State of Cloud-Driven Transformation - https://www.splunk.com/en_us/pdfs/resources/analyst-report/hbr-the-state-of-cloud-driven-transformation-2021.pdf
2. Cost of a Data Breach Report - https://www.ibm.com/security/data-breach
3. The Psychology of Human Error -https://www.tessian.com/research/the-psychology-of-human-error/
4. Microsoft, Amazon to Split Contract as $10 Billion Deal Canceled - https://about.bgov.com/news/microsoft-amazon-to-split-contract-as-10-billion-deal-canceled/
5. Forever 21 moves to Microsoft Azure, registers record-breaking holiday revenues - https://customers.microsoft.com/en-ca/story/724149-forever-21-retailers-azure
6. Five Trends Reveal the Emergence of Cloud-First Enterprises - https://www.opsramp.com/resources/insights/cloud-first-enterprise/
7. How much does a IT Staff make? - https://www.glassdoor.ca/Salaries/it-staff-salary-SRCH_KO0,8.htm?countryRedirect=true#:~:text=How%20much%20does%20an%20IT,States%20is%20%2448%2C452%20per%20year.
8. Veaam Data Protection Report 2021 - https://www.veeam.com/wp-2021-data-protection-trends.html
9. Don't Let Hardware Failures Cripple Your Business - https://www.unitrends.com/blog/dont-let-hardware-failures-cripple-business
10. 2021 State of DevOps Report - https://puppet.com/resources/report/2021-state-of-devops-report
11. DevOps Research and Assessment (DORA) team at Google Cloud 2021 Accelerate State of DevOps Report - https://services.google.com/fh/files/misc/state-of-devops-2021.pdf
12. 5 ways cloud drives enterprise innovation - https://www.ibm.com/blogs/cloud-computing/2017/05/17/cloud-drives-enterprise-innovation/
13. https://www.ideatovalue.com/inno/nickskillicorn/2016/03/innovation-15-experts-share-innovation-definition/
14. The Downfall of Blockbuster - https://medium.com/an-idea/the-downfall-of-blockbuster-da69f6c8a536
15. Cloud Adoption Statistics for 2022 - https://webtribunal.net/blog/cloud-adoption-statistics
16. May 6th 2017 issue of "The Economist" - https://www.economist.com/leaders/2017/05/06/the-worlds-most-valuable-resource-is-no-longer-oil-but-data
17. Nvidia expects GPU supply constraints for the 'vast majority' of 2022 - https://www.theverge.com/2021/8/19/22632047/nvidia-gpu-supply-rtx-30-series-constraints-2022
18. IBM Mainframe Concepts - https://www.ibm.com/docs/en/zos-basic-skills?topic=today-what-is-mainframe-its-style-computing

LefeWare Solutions

# Report

Exploring cloud computing market trends and drivers in 2022 as businesses transform their IT infrastructure to cut cost, improve security, and innovate.

**Phone:**
727-729-1288

**Email:**
info@lefewaresolutions.com

**Website:**
LefeWareSolutions.com