Cloud Waste Industry Report

zesty

Executive summary

Standing towards the first quarter of 2021, we see no signs of the rate of enterprise cloud adoption slowing down. More and more organizations are migrating their workloads to the public cloud every day. Companies are increasingly developing applications built to run natively on the public cloud.

With the increase in cloud usage comes the question of spending. The world currently sees enterprise spending for the public cloud in the range of billions of dollars, with the figure estimated to grow dramatically in the coming years. Sadly, most of this spending is caused by “cloud waste” - cloud computing capacity enterprises believe they need but are not used in reality.

This Ebook presents the current picture of the cloud computing market and how organizations spend on their cloud footprint. It introduces the concept of cloud waste, identifies common sources of waste, and the reason for such waste. The later part of the Ebook introduces some means of mitigating this challenge.

Table of contents

The Rise of the Public Cloud	4
- Drivers for Increased Cloud Adoption	4
The Cost of Cloud Adoption: Uncertain Spending and Waste	5
Defining Cloud Waste	7
Major Sources of Cloud Waste	8
- Overprovisioned and Under-utilized Resources	9
- Unscheduled Resources	10
- Idle Resources	11
- On-demand Resources	12
- Old Snapshots and Images	13
Impact of Cloud Waste	14
Mitigating the Challenge of Cloud Waste	15
- Bring Key Stakeholders Together	15
- Define a Cloud Governance Policy	16
- Create Baselines, Forecasts, and Budgets	17
- Build and Manage Solutions for Cost-effectiveness	19
- Follow Architectural Best Practices	19
- Implement Operational Cost Control	20
- Use Cost Allocation and Chargeback	20
- Use Feedback Loop	22
Conclusion	23
References	24

The rise of the public cloud

As we stand in the first quarter of 2021, enterprise adoption of public cloud is still showing an almost unstoppable upward trend. According to a Gartner report published in 2019, the worldwide public cloud services market is forecast to grow to a total of $266.4 billion in 2020 and reach $308.5 billion in 2021. Another Gartner study predicts that by 2025, 80% of organizations will migrate their workloads to colocation, hosting, and the cloud.

2020

worldwide public cloud services market was forecast to grow to

266.4B$

2021

worldwide public cloud services market was forecast to grow to

308.5B$

Drivers for increased cloud adoption

There are few reasons for the enormous scale of cloud adoption. The most obvious reason is the ease of adopting cloud technology, and how quickly you can provision infrastructure resources with code. Companies can sign up and get access to state-of-the-art virtual data centers with almost limitless capacity with a few mouse clicks, and then start building their network. This is made even more attractive with the economics of scaling and pay-as-you-go model. Once deployed, cloud-hosted infrastructure can be dynamically scaled up or down with application demand, and users have to pay for the capacity and computing power they use.

The recent proliferation of serverless and “as-a-service” offerings has also seen enterprises moving their workloads to managed environments in the cloud. This is also driven by the wide-scale adoption of microservice-based distributed applications. Managed platforms easily handle the scalability demands of these applications. As organizations are freed-up from installing, configuring, managing, and upgrading critical applications like databases, firewalls, or container runtimes, their footprint in the cloud is also increased. Another driver for increased cloud adoption is the virtually unlimited storage it offers. As companies of all sizes work in an increasingly connected world, the need for cheap, durable, secured, and scalable storage has exponentially grown over the last decade. Public cloud providers have responded to this with varied choices from object storage mediums to high-performance block storage volumes. The demand resulted in the supply, which in turn has seen further cloud usage.

The cost of cloud adoption: Uncertain spending and waste

With the ever-increasing cloud footprint, many organizations are also seeing a noticeable rise in their cloud spending.

Cost of cloud adoption

20% Of organizations reported spending more than 12$ million every year for the public cloud.

74% Of organizations reported spending more than 1.2$ million every year for their cloud bill.

59% Of enterprises expect their cloud usage to increase due to COVID-19 pandemic.

23% Above the allocated cloud budget spending by companies in 2020.

This increased spending also comes with a more worrying trend - companies are struggling to understand how their IT budget is being spent in the cloud. Worse, they are finding it difficult to forecast future cloud spending.

The cost of cloud adoption: Uncertain spending and waste

A 2020 Gartner research identifies six reasons why companies face challenges in managing cloud expenses. These include:

- Pricing structures from cloud providers often have thousands of combinations and options. This complexity makes it significantly challenging to choose the most cost-effective option.
- Billing reports can be too granular, which makes chargeback to internal stakeholders difficult.
- Ease of provisioning cloud services sometimes results in large fleets of unmanaged resources.
- A constant flow of new services, features, and changes to pricing models make it difficult to adopt the best solution for the same price.
- An abundance of similar or alternate architecture for the same solution often results in selecting higher-priced cloud components.
- The lack of standardization in APIs, services, billing models, and billing reports from major cloud providers also contribute to inefficient spending.

Defining cloud waste

Simply put, “cloud-waste” is cloud computing capacity purchased upfront or currently being paid for that is not being used optimally.
Organizations usually follow a stricter procurement policy, budgeting process, and Return-on-Investment (ROI) calculation when purchasing compute and storage capacity for their on-premise networks.
With public clouds, much of these upfront calculations are gone.
A much simpler no-upfront, no lock-in, pay-as-you-go model makes cloud migrations an attractive option, and as a result, the waste in cloud capacity is not sometimes evident.

There are two ways organizations can migrate applications to the cloud:

Lift-and-shift: With this approach, an existing application is migrated to a public cloud environment to run on a like-for-like or very similar infrastructure. There is no application re-architecture. For example, if an on-premise ELK stack (Elasticsearch, Logstash, Kibana) ran on a five-node cluster with 64 GB RAM in each server, the same architecture, the same number of nodes, and the same server specifications are used in its cloud implementation.

Re-architecting:
Here, an existing on-premise application is re-architected for the public cloud. Each application component maps to a relevant service from the public cloud provider, often resulting in a broader, sometimes more distributed footprint.

In both cases, there are chances of over-provisioning infrastructure, storage, or compute capacity. For example, depending on the current volume and growth of data, the ELK stack in the cloud may not need large-sized nodes. Similarly, a re-architected application could be running a message queuing application on a VM, but using a managed service could be cost-effective.
A cloud resource is occasionally provisioned but then swapped for another resource when the application is re-architected, leaving the previously-provisioned one unused.

As these unused, under-utilized, over-provisioned resources slowly add up, an organization keeps spending more and more than it should be for its cloud footprint

Cloud waste top contributors

Overprovisioned and Under-utilized Resources

Overprovisioning is when a cloud resource is assigned more capacity than needed. Under-utilization is when a resource is consistently used well below its available capacity.
For example, if a database instance has four CPUs, but the overall CPU utilization is always less than 20% during peak production hours, it can be easily considered both overprovisioned and under-utilized. In this scenario, the organization will be paying for a capacity it does not need.

A 2019 Flexera report found that about 40% of cloud-hosted instances are provisioned at least one size larger than needed. The same report also found that, on average, a statistically representative group of AWS customers was spending 39% of their EC2 on under-utilized VMs, and most of those instances were running only on 20% utilization.

Unscheduled Resources

Non-production environments usually emulate the production workload. Typically, these environments are used for development, testing, QA, demonstration, or training purposes. Although resources in these environments are generally not required to run 24x7, VMs, databases, or multi-node clusters are often left to run after-hours and weekends. Without scheduled shutdowns, the unused capacity gets wasted, but the per-hour cost still applies.

To illustrate this, an AWS customer running an on-demand m5.2xlarge RedHat Enterprise Linux EC2 instance in the Sydney region will be paying $0.61 per hour (as of 2020). Running the instance 24x7 for one year would cost the company $5,343.60 ($0.61 x 24 x 365). If the same instance was running for 12 hours during weekdays and scheduled during after hours and weekends, the cost would be $1903 ($0.61 x 12 x5 x 52) for one year, a 64% saving.

Idle resources

Unlike unscheduled resources where there is some level of usage, idle resources are never used. These resources are often provisioned during temporary projects, migrations, troubleshooting, or Proof-of-Concept (PoC) initiatives, but never removed afterwards. Examples of such resources can be un-attached block storage volumes, Elastic IP addresses, load balancers, or VMs.

On-demand resources

Public cloud providers often have two types of resource purchasing options. With the on-demand option, users only pay for the hours or minutes a resource is running but are charged a higher hourly rate.
With the reserved option, users commit to running a resource for a specified period but charged a much lower per-hour or a per-minute rate.
For production workloads, buying long-running resources at an n-demand rate contributes to unnecessary higher spending. Epsilon, a global network service provider, mentions in a 2020 report that 47% of AWS customers use AWS Reserved Instances while this number is only 23% for Microsoft Azure users.

Old snapshots and images

The 2016 Flexera report also found that unused cloud storage (including unattached volumes and old snapshots) represented 7 percent of cloud spend. Disk volume snapshots and machine images are created primarily for disaster recovery (DR) purposes. Usually, block storage volumes, entire VMs, database instances, or even multi-node clusters are candidates for snapshots or imaging. While there is undoubtedly a need for disaster readiness, sometimes there are no cycling policies attached to the automated snapshot processes. As a result, images and snapshots can accumulate over time, with the older ones becoming useless for recovery purposes. A large number of old snapshots can significantly impact the bill.

Other sources of cloud waste include:

- Use of suboptimal containerization architecture.
- Use of legacy instance types.
- Large volumes of unnecessary data in objects storage media.
- Large numbers of unconsolidated database instances.
- Frequently running but failing serverless functions.
- Suboptimal network design causing large outbound traffic over the internet.

The impact of cloud waste

The impact of cloud waste is multi-faceted.

With cloud expenses rising, enterprises see a rise in their Total Cost of Ownership (TCO) along with a decreasing Return on Investment (ROI). According to IDC, by 2020, cloud spending will reach 60% of all IT infrastructure and 60-70% of all software, services, and technology spending.

This rise in expenses affects a company's balance sheet and, in turn, affects the investors in lower or no dividends and falling share prices. According to Gartner, 35% of Cloud spend is wasted.Cloud service providers are also affected by this waste. The vendors' business model depends on increasing the number of customers who sign up for their service.

When existing customers' idle resources occupy a large part of the infrastructure capacity, cloud providers must expand the underlying network and hardware capacity for new customers. This expansion does not happen overnight - and the vendor has to bear the cost of the growth initially.
This cost is then passed on to new and existing customers.

With a growing data center footprint comes the question of waste.
Increased power consumption, non-recyclable water for cooling systems, building maintenance waste, and discarded hardware components negatively impact the environment.

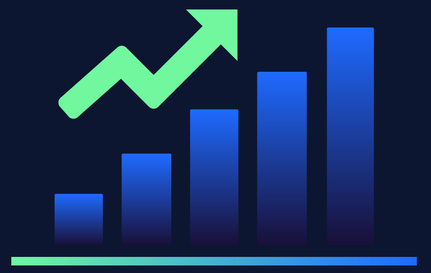

By 2020, cloud spending will reach **60%** of all IT infrastructure and **60-70%** of all software, services, and technology spending.

According to Gartner, 35% **of cloud spend is wasted**

Mitigating the challenge of cloud waste

Mitigating the Challenge of Cloud Waste Naturally, optimizing cloud spending is now viewed as a top priority by many companies. The Flexera 2020 report finds 73% of organizations now consider optimizing cloud spend as their top priority.

This whitepaper provides some best-practice recommendations for optimizing cloud costs. Some of these recommendations are functional and business-centric, while others are more technology-oriented. This split is because implementing a cloud cost optimization practice is not the responsibility of a single discipline or department but should be an organization-wide effort. Enterprises can implement these recommendations as new practices or integrate them with their existing initiatives. Some of these processes are iterative, where continuous feedback and adjustments will lead to optimal results.

Bring key stakeholders together

Although the benefit of cloud computing is organization-wide, different teams and departments have different perspectives of cloud costs. When it comes to planning strategy, each can bring a rich set of expertise and knowledge. Identifying these teams and bringing them under a working body is the first step in cloud cost optimization.

For example, IT development and infrastructure teams have first-hand knowledge of the most expensive cloud components currently running in the network. They can recommend more efficient application architectures to cut down cloud costs. Similarly, the operations team will have the necessary knowledge of various critical applications and their utilization profiles.

Platform owners and project sponsors can bring a wealth of knowledge about future application requirements and demands, as well as opportunities for consolidating or scaling back. The Finance department can provide input about current and projected future spending. They will also be aware of any budget in place, cost allocations against teams or applications, and how the spending matches with the current budget. Together, the stakeholders will discuss the objectives of cloud cost optimization and form a strategy. One of the key deliverables of this strategy will be formulating a cloud governance policy.

Define a Cloud Governance Policy

Many organizations do not have a written or updated cloud governance policy. Although it can take many iterations to create, a cloud governance policy provides the core blueprint for cost optimization.

Here are some of the cost-related processes, guidelines, and methods a cloud governance policy should include or refer to:

- The primary stakeholding teams or roles responsible for financing, managing, and operating the cloud infrastructure.
- The business process for selecting, commissioning and approving cloud migration of legacy applications.
- The business process for commissioning and approving new cloud applications or enhancing existing ones.
- The business process for retiring cloud-hosted applications.
- The business process for allocating extra resources to existing cloud applications.
- A guideline for creating cloud budgets at organization and cost-center level.
- A guideline for measuring cloud costs at both granular and consolidated level.
- Cost allocation and chargeback policies for different cloud applications.
- A range of cost optimization methods for different types of resources and the means of measuring the results.
- The method of comparing actual spendings against budgets and identifying areas that need optimization.

Create Baselines, Forecasts And Budgets 1 | 2

An organization cannot be sure if it is spending more than it should unless there is a baseline to compare against. This baseline is usually a spending forecast or a budget. This principle also applies to cloud computing.

Instead of budgeting for the entire organization's cloud spending, it is best to start with an incremental approach. With this method, each department, team, or project performs an audit of their applications and the cloud resources the applications are using. The resource spendings are then obtained from the cloud provider's bill and form an initial baseline.

The process is often tricky as multiple teams can share the same applications or cloud resources. To work around this problem, a company can assign best-guess weighting for each stakeholder based on their usage time or data storage. For example, if a data warehouse cluster services the application development, data analytics, and the marketing team, their usage can be split as 20% (development), 50% (data analytics), and 30% (marketing). This usage split can be based on how often, or how long connections, queries, and processes from each team is running on the cluster. Similarly, the spending on an object storage media like Amazon S3 can be split between different teams based on how much data each team is storing there. As mentioned later, using resource tags and naming standards can greatly simplify the cost allocation process.

And trends, they can start forecasting for the next few months - ideally a quarter or half-year. Forecasts are also necessary when migrating applications to the cloud or building new cloud-native applications.
Most public cloud vendors and some third-party companies offer cloud cost calculators that can help create forecasts. Most of these calculators will allow users to specify cloud resource types, upfront or on-demand

Purchasing options, approximate computing time, storage needs, and estimated network traffic. Based on the details provided, the calculator would show an expected spending figure. This spending amount will be based on the cloud vendor's current price list.

When using cost calculators, organizations should use multiple combinations of resource types and usage scenarios to see which values are optimal or realistic. For example, it may be worth checking if a managed NoSQL service is cheaper than a customer-managed cluster, or perhaps a dedicated NAT gateway will be better than a managed solution. Although some solutions will be less expensive than others, those may not be architecturally sound. Therefore it is necessary to balance the pricing estimates with architectural best practices. It is also important to factor in future data growth, performance requirements, and increased traffic volume. Also, the calculators will be as good as the numbers given, and for most teams, the numbers will be approximate at the beginning. Therefore, there should be a margin of error.

Create Baselines, Forecasts And Budgets 2 | 2

For new cloud initiatives, application teams or business stakeholders can present their case to the cloud governance body once an estimate is ready. The cloud governance body can then conduct needs and benefit analysis. The needs analysis will also take into consideration similar application spendings or previous failures. This practice creates a culture of transparency and accountability.

Once the initiatives and forecasts are approved, teams can start budgeting for their cloud spending. Once again, some cloud vendors offer budgeting tools that can compare running costs against the current monthly budget. Organizations should adopt native cost-estimation and reporting tools where possible. Such tools can automatically update their values with pricing changes and report with better granularity.

Build and Manage Solutions for Cost-effectiveness 1 | 3

The primary goal of creating a budget for the near and long term is to keep costs within the forecasted limit. The primary means of ensuring that is to design and develop cost-effective cloud solutions and optimizing existing solutions for cost.

There are several recommended methods for ensuring both. These methods are best developed and implemented by enterprise architects, business analysts, DevOps teams, application developers, database administrators, and operations personnel due to the technical nature.

Follow Architectural Best Practices

Application architecture best practices should not only cater for high availability, performance, and security in applications, these should also help choose the most cost-effective building blocks.
For example, during the design phase, solution architecture may decide using managed services with built-in fault tolerance instead of creating highly-available clusters. Network design may stipulate keeping all resources in a single geographical location to minimize cross-region outbound traffic. Yet another design choice may opt for open-source platforms to reduce licensing fees.
Other design choices that can positively impact costs include using horizontal auto-scaling instead of vertical, reusing and consolidating non-production environments, throttling serverless function calls, and using managed runtimes for containerized applications instead of VMs.

Provisioning storage is directly dependent on the volume of data, and architecture best practices must match capacity and performance requirements with the best cost-effective solutions. For example, managed database service from some cloud providers like AWS allows users to provision storage manually or let the database instance automatically add additional storage when needed. The latter type of instances cost more than the former, but architecture considerations may find it cheaper in the long run.

Other storage-related architectural considerations can include the type of storage classes to use. Usually, “hot” data needs low-latency storage, which costs higher. Cheaper storage tiers or magnetic mediums can be suitable for infrequently accessed data. Another area to consider is data archival and retention. Archived data should reside in the cheapest storage tier. The retrieval time from this tier is much longer, but it can offer significant savings over regular storage.

One of the most critical design decisions includes the use of reserved instances. Reserved instances are suitable for steady workloads running over long periods. The customer pays a reduced hourly rate in exchange for committing to use a computing resource for a specific period. The discount model also depends on the type of resource, the specified term, whether upfront, partial upfront, or no upfront payment are to be made, and the flexibility of choosing geographical locations for the resource or the operating system it will run.

Reserved Instances are a feature of all three major cloud service providers and play a significant role in minimizing cost. One of the core tasks of operations teams is to monitor unused reserved instances that are negatively affecting cloud costs. Amazon Web Service also offers other

Build and Manage Solutions for Cost-effectiveness 2 | 3

Cost-saving pricing models like Savings Plans and hosts a Reserved Instance Marketplace where Ris can be purchased and sold.

Implement Operational Cost Control

Operational cost control involves proactive and reactive monitoring of cloud resources for their cost profiles and taking remedial measures as necessary. However, an enterprise cloud setup can include multiple accounts across different regions and often contain hundreds or even thousands of resources. Manually checking resources for their cost-effectiveness can be a tedious and error-prone process in such cases. Organizations should, therefore, adopt automated scanning where possible.

Some automated processes can scan for unattached block storage volumes, load balancers, and static IP addresses and remove if any of these are found. Other such processes can shut down instances that have not been logged into for a configurable period, remove snapshots and images older than a threshold date, or unsubscribe from unused managed resources like Notebooks, message queues, or serverless functions.

There are cloud cost management tools in the market that can schedule instances. These applications can shut down non-production instances during specific periods like after hours or weekends, thereby saving costs. Many of these tools can also show the savings made daily.
Proactive monitoring and alerting also play a significant role in keeping track of non-optimal cloud resources. Such processes are often custom written by cloud engineers. Nevertheless, they can be used to report about consistently under-utilized resources, object storage buckets without lifecycle policies, untagged resources, long-running on-demand instances, and unused reserved instances.

Automated cost-saving processes should be designed after careful analysis of resource utilization patterns and usage forecasts.

For example, a cluster of VM's resource consumption may be high only during the month-end processing. In such a case, it will make sense to move the workload to a fewer number of lower-specification nodes and using auto-scaling during the peak period.

Use Cost Allocation and Chargeback

Cloud cost allocation or a chargeback is a process where an organization bills its internal teams, projects, or departments for using its cloud resources. The charged amount is based on the costs incurred by the billed entity. These costs are directly affected by the processing time, compute power, and storage capacity the entity is using.

Cost allocation ensures a culture of responsibility and accountability.
It also helps different teams and departments align their own cloud budgets and forecasts with the ones created for the organization.
Cost allocation is usually the domain of FinOps. There are many factors that affect it. Its calculation can also be a complex process depending

Build and Manage Solutions for Cost-effectiveness 3 | 3

On company size and structure, shared applications, and commonly used infrastructure or backbone services like network link, directory service, or firewalls.However, there are some simple ways to help apportion cloud costs between different organizational entities. The simplest one is to create individual cloud accounts for each department or team. Various cost centers in the organization can have their cloud account as needed, with every account linking back to a “master account” owned and managed by the finance department. The cloud service provider bills this master account for its sub-accounts’ spending. The account owner is responsible

For paying the total bill and assigning each account’s cost to its owning cost center for payback. Designing an effective multi-account strategy can take some time. Still, this method is the easiest way to keep track of spending. account’s cost to its owning cost center for payback. Designing an effective multi-account strategy can take some time. Still, this method is the easiest way to keep track of spending.
Although it can be challenging to allocate costs in single accounts, some level of separation is possible by grouping each cost center’s resources in separate virtual networks. Effective naming standards can also identify resources belonging to applications, environments, projects, or teams.

Another method to help cost allocation is to use tags or labels. Tags are name-value pairs that contain user-specified properties for a cloud resource. For example, a tag can include a resource’s name.

Another can indicate its environment (DEV, TST, PRD), yet another can contain the cost center or project code, and so on.

The benefit of using tags is that it helps identify and group the resources belonging to each environment, application, cost center, team, or project. There may be common resources like storage gateways, VPN connections, or CDNs shared across all parties. Simple tagging mechanisms are not sufficient in these cases, and the organization must devise a strategy to apportion these costs.

Designing an effective tagging strategy can take time. This is where the cloud governance policy can provide guidance, and architecture best practices can specify a set of common hierarchical tags for all resources. The use of hierarchy in tagging ensures smaller entities within larger groups are billed correctly. Cost allocation reports can then easily break down the charged amount. For example, there can be a higher level tag.

For all HR and Payroll department resources. The payroll department could be using two servers - one for accounts receivable and another for accounts payable. Each of these servers could be appropriately tagged to show costs incurred by each workload.
Naming standards and tagging are often enforced with automation.
Most organizations deploy cloud infrastructure with automated code; these codes can check for correct naming and mandatory tags to ensure resources adhere to established standards.

Use Feedback Loop

To ensure early intervention when running costs are going above the budgeted values, enterprises should adopt an alerting mechanism.
This works as a feedback loop where deviations are reported early in the billing cycle. Most organizations implement this in the form of a billing alert where a notification is sent to the operations or finance team when
a spending threshold is reached. Although billing alerts are certainly encouraged, it does not report the particular cloud resource causing the breach. Also, a billing alert is triggered by the current total spending - it cannot be customized for individual resources or groups of resources exceeding their allocated spending quota.

Therefore, it is essential to adopt a more granular alerting method.
For example, alerts can notify when

- The current spending for a group of resources exceeds their budgeted value.
- A new resource is added to a pool - for example, an auto-scaling event.
- Automated cost control methods find unused or low-utilized resources or resources.
- One or more Reserved Instances become idle
- A large volume of outbound traffic is detected

With continuous feedback, operations teams can proactively investigate the reasons for cost spikes.

Summary

The biggest challenge in cost optimization is often the culture and mindset shift. However, with an inclusive approach and user awareness, the process becomes much easier. **Also, enterprise cloud cost optimization is an ongoing practice. A single person or team cannot do it, nor can it be sustained without guidelines and processes in place and company-wide adoption.** Organizations should therefore create a framework that uses the methods discussed here, and make it flexible for future improvements.

References

1. Ashok, A. (2018). Four Trends In Cloud Computing CIOs Should Prepare For In 2019. [online] Forbes. Available at: **https://www.forbes.com/sites/forbestechcouncil/2018/07/05/four-trends-in-cloud-computing-cios-should-prepare-for-in-2019/#2faf33544dc2** [Accessed Oct. 2020].

2. Axworthy, J. (2020). Why firms are wasting their cloud spending? [online] DZone. Available at: **https://www.raconteur.net/technology/cloud-business-2020/cloud-waste-digital-transformation** [Accessed Oct. 2020].

3. C, Prasanna. (2020). WASTED CLOUD SPEND: WHERE IS THE SILVER LINING?. [online] Epsilon. Available at: **https://www.epsilontel.com/blogposts/wasted-cloud-spend-where-is-the-silver-lining/**

4. Chapel, J. (2019). Cloud Waste To Hit Over $14 Billion in 2019. [online] DevOps.com. Available at: **https://devops.com/cloud-waste-to-hit-over-14-billion-in-2019/** [Accessed Oct. 2020].

5. Chapel, J. (2020). The Cloud Is Booming — But so Is Cloud Waste. [online] [online] DevOps.com. Available at: **https://devops.com/the-cloud-is-booming-but-so-is-cloud-waste/** [Accessed Oct. 2020].

6. Chapel, J. (2020). The Leading Cloud Computing Challenge and How to Address It. [online] NETWORK Computing. Available at: **https://www.networkcomputing.com/cloud-infrastructure/leading-cloud-computing-challenge-and-how-address-it** [Accessed Oct. 2020].

7. Gartner.com. (2020). Gartner Forecasts Worldwide Public Cloud Revenue to Grow 17% in 2020. [online] Gartner Press Releases. Available at: **https://www.gartner.com/en/newsroom/press-releases/2019-11-13-gartner-forecasts-worldwide-public-cloud-revenue-to-grow-17-percent-in-2020** [Accessed Oct. 2020].

8. Gartner.com. (2018). Gartner Identifies the Top 10 Trends Impacting Infrastructure and Operations for 2019. [online] Gartner Press Releases. Available at: **https://www.gartner.com/en/newsroom/press-releases/2018-12-04-gartner-identifies-the-top-10-trends-impacting infras#:~:text=Gartner%20predicts%20that%20by%202025, down%20their%20traditional%20data%20center** [Accessed Oct. 2020].

9. Juneza, A. (2019). Why Is There So Much Cloud Waste? [online] DZone. Available at: **https://dzone.com/articles/why-is-cloud-waste-so-huge** [Accessed Oct. 2020].

10. hostingtribunal.com. (2020). 25 Must-Know Cloud Computing Statistics in 2020. [online] Available at: **https://hostingtribunal.com/blog/cloud-computing-statistics/#gref** [Accessed Oct. 2020].

References

11. hostingtribunal.com. (2020). The Latest Cloud Computing Trends – 31+ Telling Facts. [online] Available at: **https://hostingtribunal.com/blog/cloud-computing-trends/** [Accessed Oct. 2020].

12. Meinardi, M. and Clayton T. (2020). How to Manage and Optimize Costs of Public Cloud IaaS and PaaS. [online] Gartner. Available at: **https://www.gartner.com/en/documents/3982411/how-to-manage-and-optimize-costs-of-public-cloud-iaas-an** [Accessed Oct. 2020].

13. Raza M. and Kidd C. (2020). Cloud Growth in 2020: Trends & Outlook. [online] BMC Blogs. Available at: **https://www.bmc.com/blogs/cloud-growth-trends/** [Accessed Oct. 2020].

14. Rimol, M. (2020). 4 Trends Impacting Cloud Adoption in 2020. [online] Gartner. Available at: **https://www.gartner.com/smarterwithgartner/4-trends-impacting-cloud-adoption-in-2020/** [Accessed Oct. 2020].

15. Stalcup, K. (2020). Wasted Cloud Spend to Exceed $17.6 Billion in 2020, Fueled by Cloud Computing Growth. [online] business2community.com. Available at: **https://www.business2community.com/cloud-computing/wasted-cloud-spend-to-exceed-17-6-billion-in-2020-fueled-by-cloud-computing-growth-02292542** [Accessed Oct. 2020].

16. Weins, K. (2016). AWS Costs: How Much Are You Wasting? [online] Flexera. Available at: **https://www.flexera.com/blog/cloud/aws-costs-how-much-are-you-wasting/**

17. Weins, K. (2017). Where Is the $10B in Waste in Public Cloud Costs? [online] Flexera. Available at: **https://www.flexera.com/blog/cloud/where-is-the-10b-in-waste-in-public-cloud-costs/** [Accessed Oct. 2020].

18. Weins, K. (2020). Cloud Computing Trends: 2020 State of the Cloud Report. [online] Flexera. Available at: **https://www.flexera.com/blog/industry-trends/trend-of-cloud-computing-2020/documents/3982411/how-to-manage-and-optimize-costs-of-public-cloud-iaas-an** [Accessed Oct. 2020].

19. Wickizer, D. (2017). The Cloud Waste Problem That's Killing Your Business (and What To Do About It). [online] parkmycloud.com. Available at: **https://www.parkmycloud.com/blog/cloud-waste/** [Accessed Oct. 2020].