AWS Cost Analysis Report

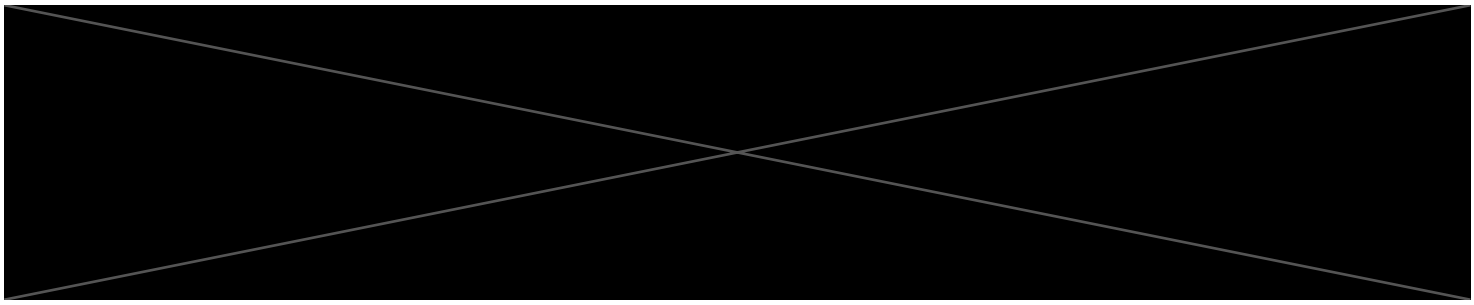

Report Date: [redacted]
Prepared by: [redacted]
Contact: [redacted]

AWS Cost Analysis Report **1**
Executive Summary 3
Introduction 3
Methodology 3
Data Collection 3
Data Analysis 4
Cost Analysis 5
CloudFront Cost Analysis 5
Overview 5
Usage Overview 6
Key Metrics 6
Cost Breakdown 7
Insights and Considerations 7
Benefits of Enabling Compression: 9
S3 Cost Analysis 11
Overview 11
Usage Overview 11
Storage Utilization 11
Cost Breakdown 14
Insights and Considerations 14
EC2 Cost Analysis 15
Overview 15
Usage Overview 15
Instance Types and Distribution 16
Regional Distribution 17
Cost Breakdown 17
Insights and Considerations 17
Additional services 21
Transfer Family 21
Conclusion 22
Appendix 22

Executive Summary

The purpose of this AWS Cost Analysis Report is to provide an overview of the costs associated with the usage of Amazon Web Services (AWS), focusing on the CloudFront, S3, and EC2 services. This report aims to help optimize resource allocation and identify potential cost-saving opportunities.

Introduction

In this report, we will explore the costs associated with these services, highlighting various cost components and their significance. By providing a detailed breakdown of expenses and identifying potential cost-saving strategies, this analysis aims to empower your organization to make financially sound decisions while harnessing the full potential of AWS services. Breakdown of accounts where analysis has been performed:

The analysis is based on data collected during the specified time period (last 3 months, 6 months where applicable), and our recommendations are designed to assist you in maximizing the value of your AWS investments. By gaining insights into how costs are distributed across CloudFront, S3, and EC2, you'll be better equipped to allocate resources efficiently, optimize performance, and achieve operational excellence within your AWS environment.

In the following sections, we will delve into the specifics of the cost breakdown for each service, explore opportunities for cost optimization, and offer actionable recommendations for moving forward. By understanding the nuances of AWS cost management and implementing the insights presented in this report, your organization can achieve a harmonious balance between technology innovation and cost-effectiveness.

Methodology

Data Collection

1. **AWS Cost Explorer**: AWS Cost Explorer is a powerful tool that provides detailed insights into your AWS spending patterns. It offers customizable cost and usage reports, enabling us to analyze historical cost data and identify trends. We leveraged

the AWS Cost Explorer's functionalities to obtain a broad overview of the cost distribution across CloudFront, S3, and EC2 services.

2. **AWS Billing and Cost Management**: The AWS Billing and Cost Management dashboard provides a consolidated view of your AWS billing information. By accessing this dashboard, we collected high-level billing summaries, cost breakdowns, and service-specific cost details. This data served as a foundational element for our analysis.
3. **DoiT Console**: The DoiT Console is a custom in-house platform designed to streamline the data collection process. It extracts granular usage and cost data from AWS services and aggregates it for further analysis. Using the DoiT Console, we collected specific usage metrics and cost data related to CloudFront, S3, and EC2 services.
4. **AWS Service Configuration Check**: In addition to cost data, we conducted a configuration check of the AWS services under analysis. This involved reviewing the configurations of CloudFront distributions, S3 buckets, and EC2 instances. Configuration anomalies or suboptimal settings can sometimes lead to unintended costs, and this step helped us identify potential areas for improvement.
5. **AWS Storage Lens**: An integral addition to our data collection arsenal was the employment of AWS Storage Lens. This tool provided an insightful perspective into our storage usage patterns, enabling us to glean deeper insights into our S3 storage costs. Through metrics, recommendations, and visualizations, AWS Storage Lens enriched our analysis with a focused lens on storage-related expenditures.

Data Analysis

1. **Cost Breakdown**: We analyzed the collected data to break down the costs associated with CloudFront, S3, and EC2. This involved categorizing costs into distinct components such as data transfer, storage, requests, and instance usage.
2. **Comparison with Historical Data**: By comparing the current cost data with historical data, we identified any significant changes in spending patterns. This allowed us to spot anomalies and trends that might require further investigation.
3. **Usage Patterns**: We examined usage patterns for CloudFront data transfer, S3 storage, and EC2 instance hours. This helped us understand how the services were being utilized and identify areas where optimization could lead to cost savings.

Cost Analysis

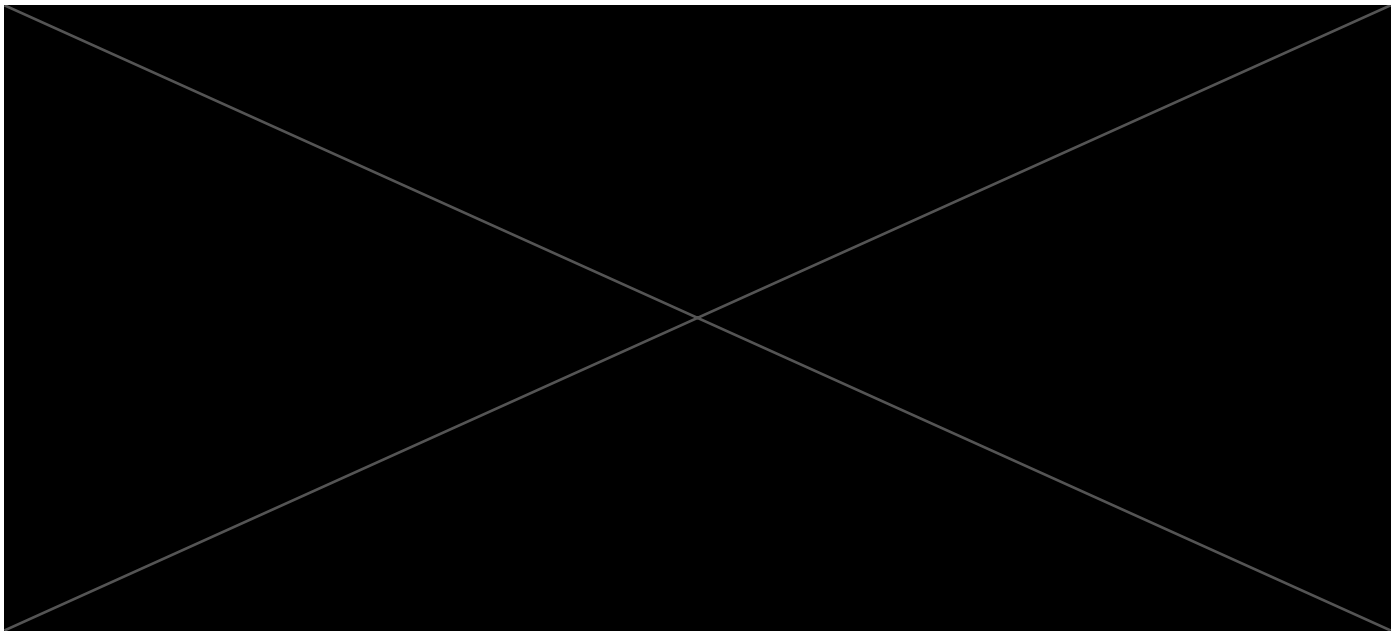

Picture 1. Breakdown of service costs in the last 3 months (Source: DoiT Console)

CloudFront Cost Analysis

Overview

Amazon CloudFront, a globally distributed content delivery network (CDN), plays a pivotal role in ensuring low-latency and high-performance delivery of content to users across the globe. As an integral part of many AWS-powered applications, understanding the cost dynamics of CloudFront is essential for achieving optimal performance while managing expenses.

Usage Overview

During the analysis period, we observed substantial CloudFront usage, with a significant volume of data being delivered to end-users.

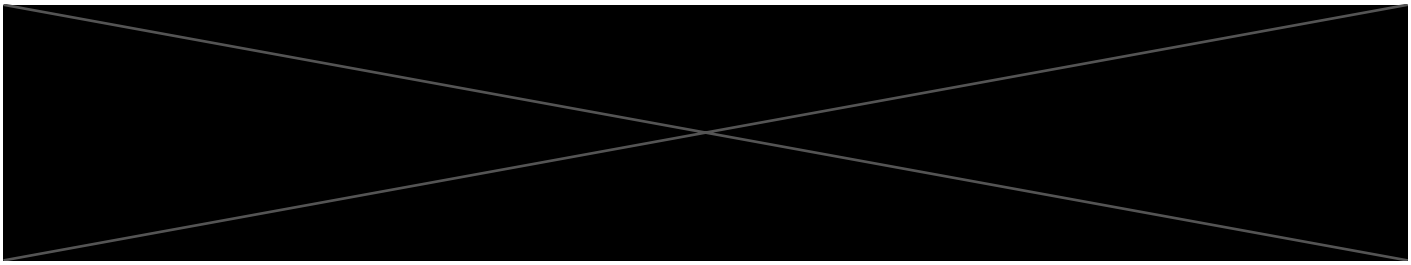

Picture 2. Breakdown of CloudFront usage for June 2023 (Source: DoiT Console)

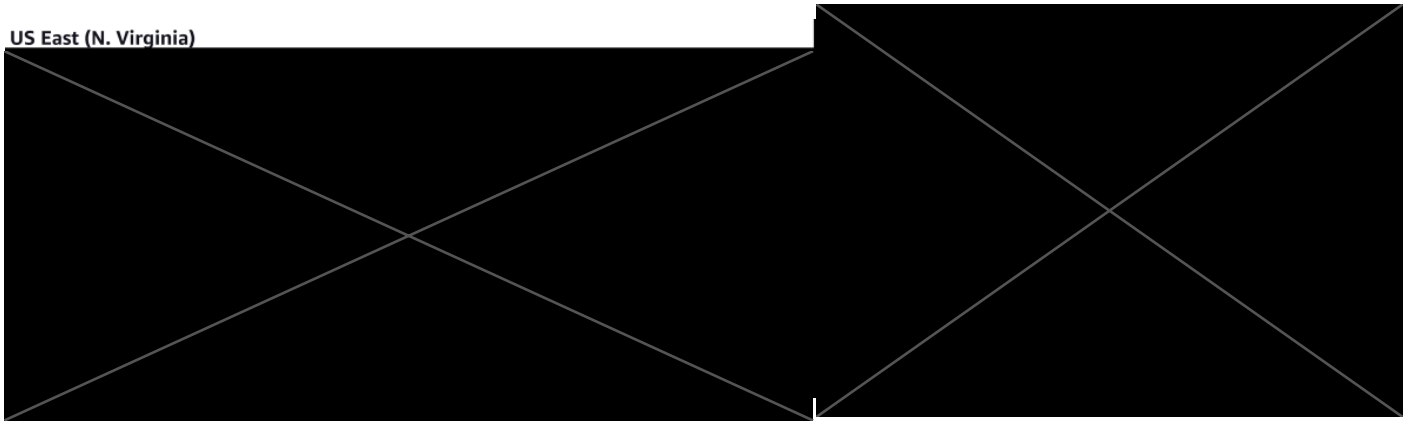

Picture 3. CloudFront usage for [redacted] US-EAST (Source: AWS Billing)

Key Metrics

- **Data Transferred**: The total volume of data transferred through CloudFront during the analysis period provides insights into the scale of content delivery.
- **Request Count**: The number of requests made to CloudFront distributions indicates the level of user engagement and the popularity of the content being delivered.
- **Geographical Distribution**: Analysis of the geographical distribution of CloudFront traffic reveals the regions with the highest demand for content (US-EAST), helping to optimize content placement.

Cost Breakdown

The cost structure of CloudFront is multi-faceted, comprising several key components that contribute to the overall expenses associated with content delivery. The primary cost drivers include:

1. **Data Transfer Costs**: CloudFront charges for the amount of data transferred from edge locations to end-users. Data transfer costs vary based on the AWS region and the destination of the traffic.
2. **HTTP/HTTPS Requests**: CloudFront charges for the number of requests made to your distributions. These requests encompass both HTTP and HTTPS requests, and the pricing is influenced by request type and volume.
3. **Distribution Charges**: CloudFront distributions, which determine how your content is cached and delivered, can incur costs based on the configuration and usage of these distributions.

Over the past three months, the peak billing occurred in June, reaching a total of 10,094.46 US$.

- **US-Requests-Tier2-HTTPS**: With a substantial count of 1,534,248,375 requests, the US-Requests-Tier2-HTTPS category contributed to a cost of 1,534.25 US$. This metric indicates the volume of advanced HTTPS requests made within the US region, reflecting the scale of interactions between users and CloudFront distributions.
- **US-DataTransfer-Out-Bytes**: The US-DataTransfer-Out-Bytes metric reached 124,490.35, resulting in a cost of 8,544.62 US$. This measurement represents the volume of data transferred out from CloudFront edge locations to end-users in the US region.

Note: US-DataTransfer-Out-Bytes represent data transferred out to the internet.

Insights and Considerations

The analysis of CloudFront costs revealed several key insights that merit consideration:

- **Caching**: While Implementing caching strategies can reduce the load on your origin server and improve content delivery speeds, it can also significantly reduce data transfer costs and enhance user experience, in your case scenario, caching statistics are showing good values as seen in the Picture 4, therefore changing caching will not make an impact on cost savings in your case.

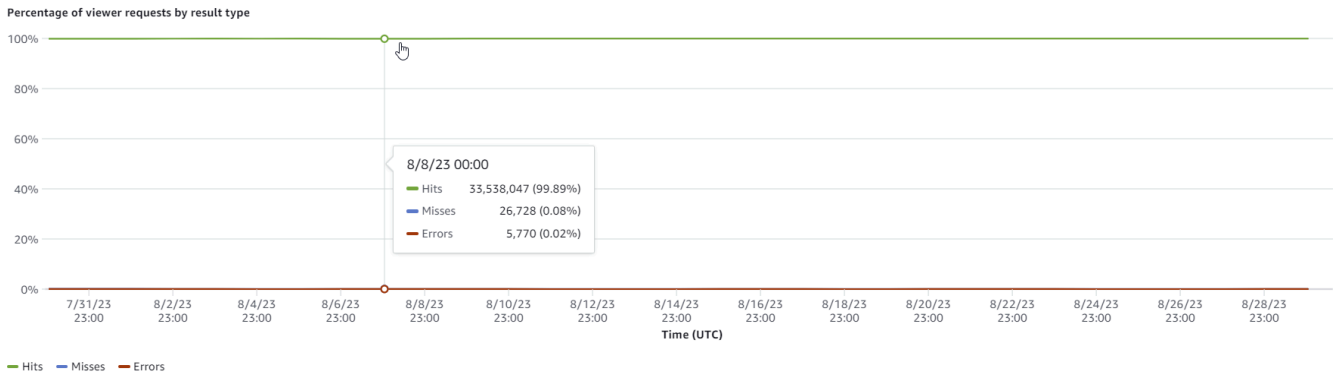

Picture 4. CloudFront caching statistics (Source: AWS CloudFront Reporting)

- **Utilize Content Compression**: Enabling content compression to reduce the amount of data transferred can lead to lower data transfer costs and faster content delivery.

  I checked the configuration of the biggest distros that you have configured and I noticed that compression is not enabled.

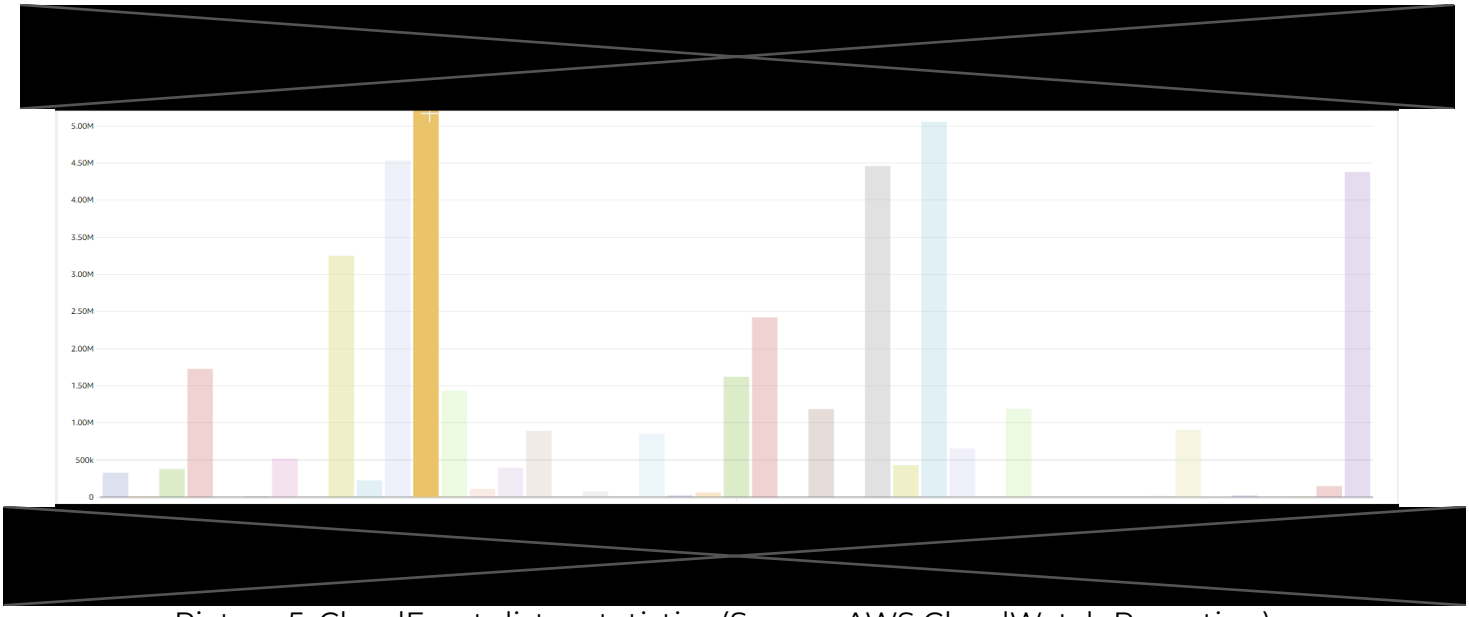

Picture 5. CloudFront distro statistics (Source: AWS CloudWatch Reporting)

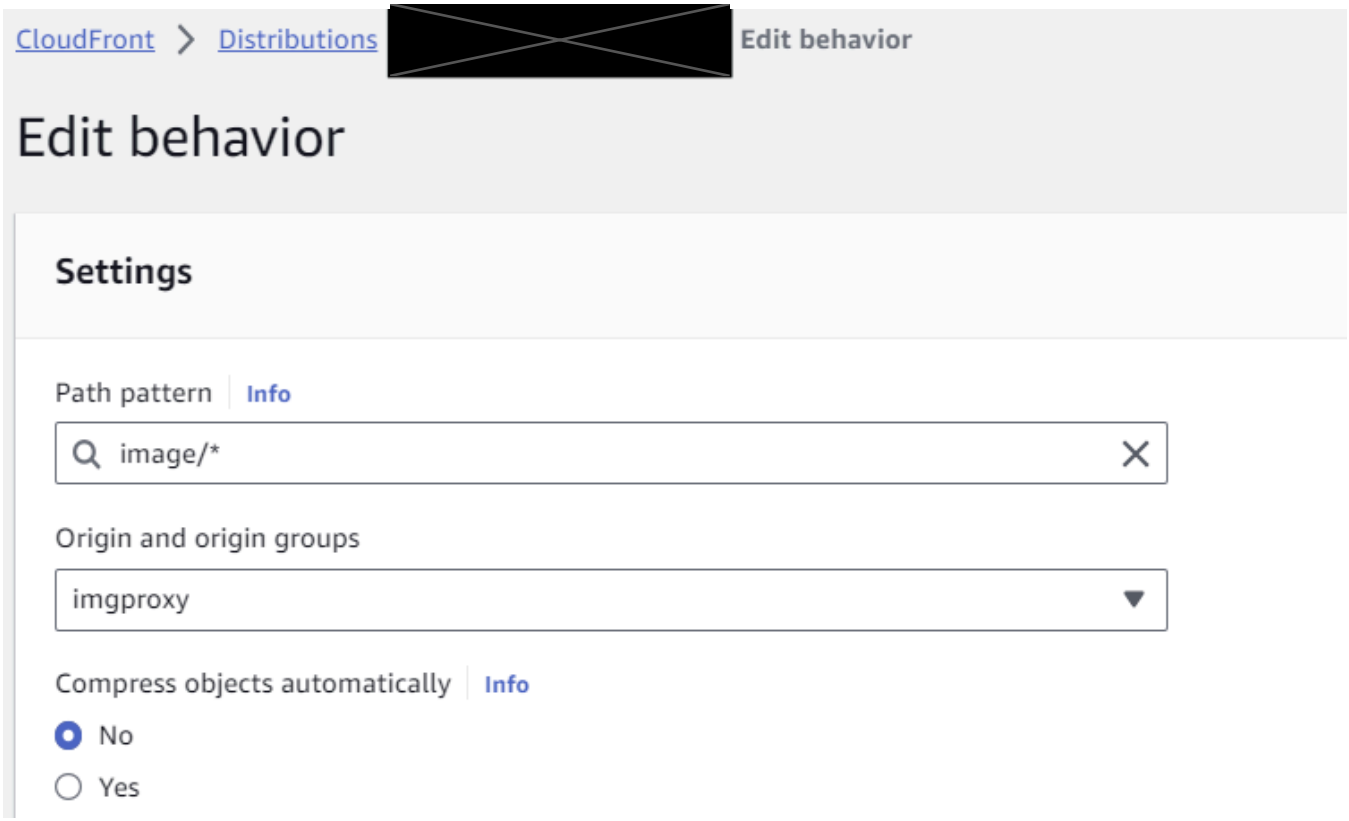

Picture 6. Compression (Source: AWS CloudFront settings)

Benefits of Enabling Compression:

1. **Faster Content Delivery**: Compressed content is smaller in size, resulting in quicker download times for users. This improves the overall user experience by reducing waiting times and enhancing website or application performance.
2. **Bandwidth Conservation**: Compressed content requires less data to be transferred between CloudFront edge locations and end-users' devices. This conserves bandwidth, which is particularly beneficial when dealing with large files or high-traffic scenarios.
3. **Cost Savings**: Reduced data transfer between CloudFront and end-users translates to lower data transfer costs. Enabling compression contributes to optimizing your AWS usage and managing expenses effectively.

Here is an overview of your data transfer over the past six months, as displayed in the Picture 7 below.

Picture 7. Transfer-Out Data (Source: DoiT Console)

This summary indicates substantial data transfer quantities for the following months:

March	April	May	June	July	August
52TB	59TB	53TB	124TB	74TB	80TB

Table 1. Transfer-Out Data in TB

Considering the volumes involved, the potential for optimizing costs becomes a central consideration. The most impactful avenue for savings, given these significant data amounts, is the adoption of a Private Pricing Agreement (PPA). By implementing a PPA, you stand to achieve substantial cost reductions, with potential savings ranging from 40% to 60% in certain scenarios.

DoiT does provide PPA CloudFront service which can give you lower prices. The whole process may take 10-30 days to implement. Your DoiT Account Manager can give you more details of the same.

S3 Cost Analysis

Overview

Amazon S3 (Simple Storage Service) is a foundational storage service that provides scalable and durable object storage for a wide range of use cases. Understanding the cost structure of S3 is crucial for managing data storage costs effectively while ensuring data availability and durability.

Usage Overview

During the analysis period, our examination of S3 usage revealed the following key patterns:

Storage Utilization

Total Storage: The total amount of data stored within Amazon S3 provides an overview of the volume of content managed within the service.

Metric name	Metric category	Total for Aug 29, 2023
Total storage	Summary	141.6 TB
Object count	Summary	210.1 M
Average object size	Summary	723.5 KB
Active buckets	Summary	296
Accounts	Summary	1
Buckets	Summary	363

Picture 8. S3 (Source: AWS Storage Lens)

Storage Classes Breakdown: Distribution of data across various storage classes offered by Amazon S3, including Standard, Intelligent-Tiering, One Zone-IA, Glacier, and others, shows that Standard Tier is used the most.

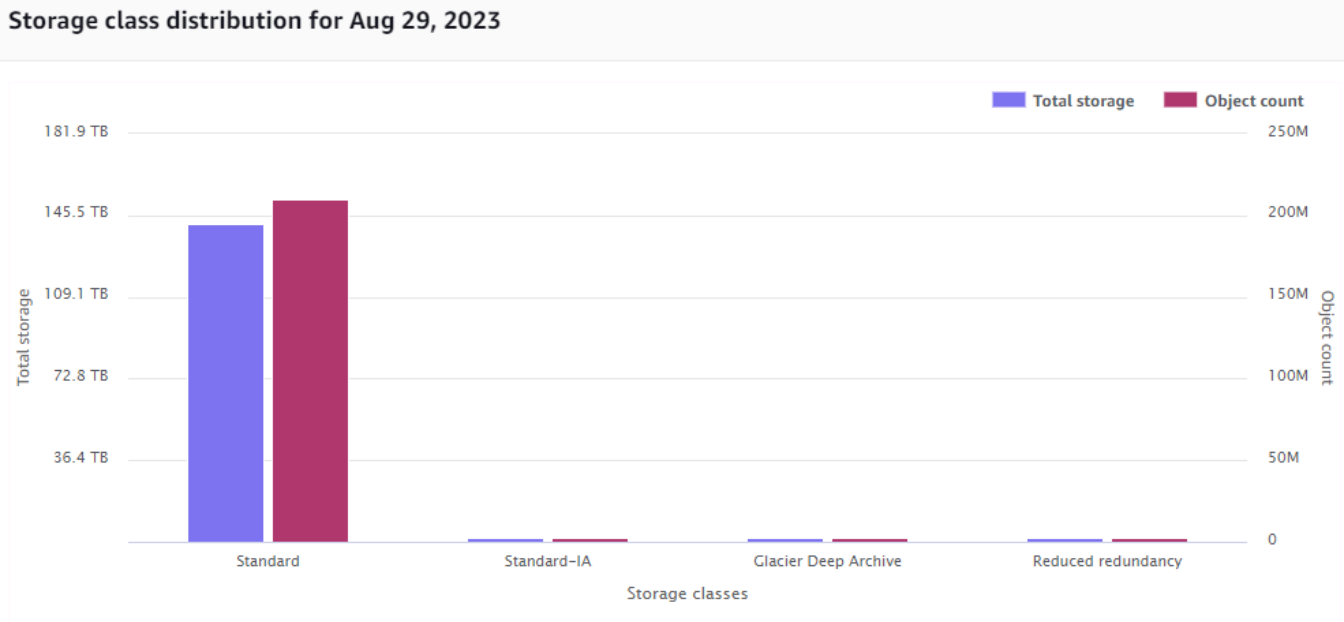

Picture 9. S3 Storage class distribution (Source: AWS Storage Lens)

Picture 10. S3 Cost by Storage Operation (Source: AWS Storage Lens)

Data Distribution Across Regions: Our investigation encompassed the geographic dispersion of data across different AWS regions. The data are mostly sitting in EU and US regions.

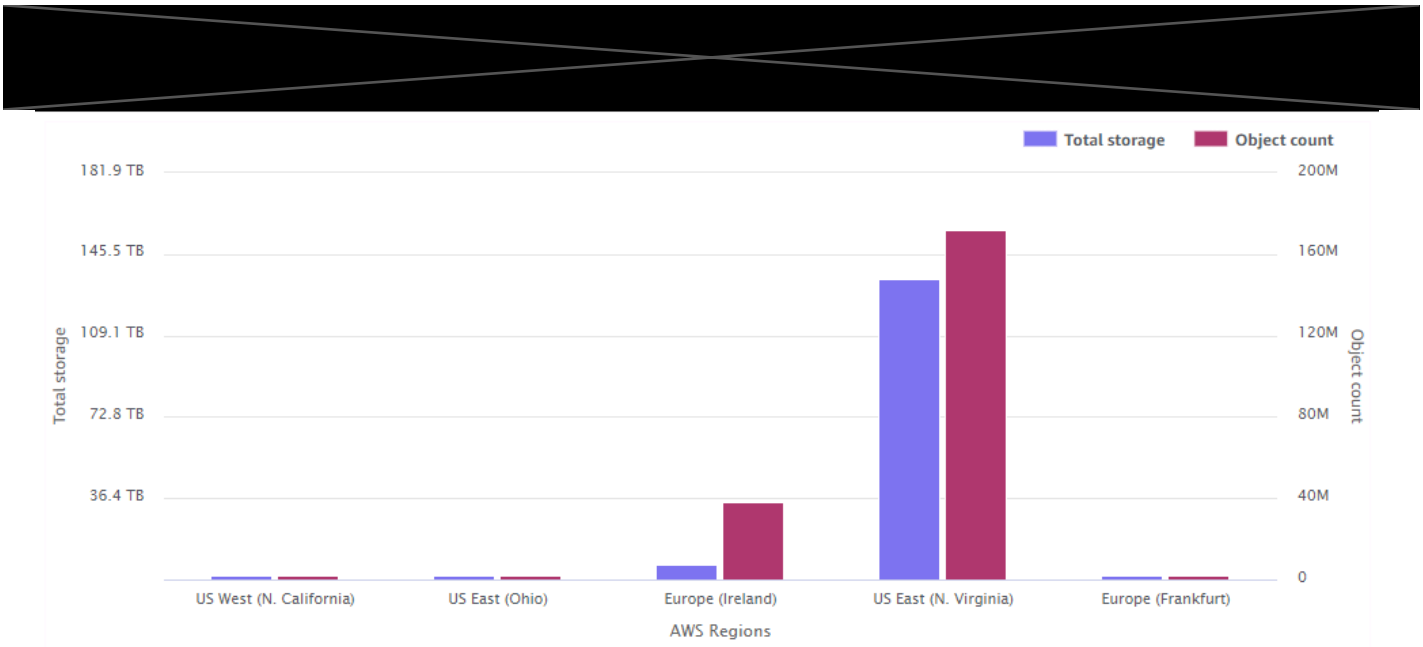

Picture 11. S3 Region distribution (Source: AWS Storage Lens)

Primary Data Repositories: By identifying and analyzing the most prominent S3 buckets in terms of data volume, we gained an understanding of your primary data repositories.

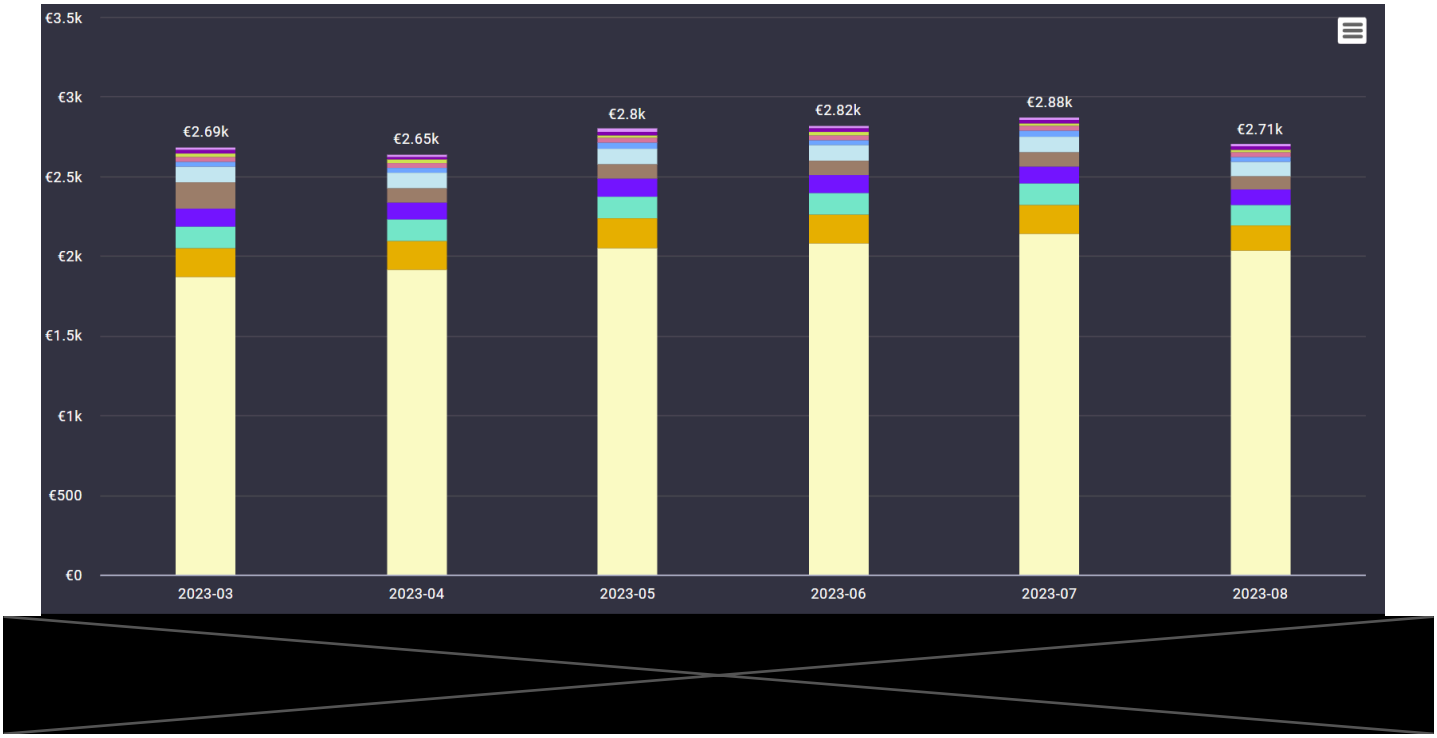

Picture 12. Top buckets spenders (Source: DoiT Console)

Cost Breakdown

Amazon S3 expenses are influenced by a range of cost components, each of which contributes to the overall expenditure. These components span across diverse elements of storage and data transfer, collectively shaping the financial landscape of S3 usage:

1. **Storage Costs**: The cost of storing objects in S3 is influenced by the total volume of data stored, the storage class used, and any data management operations performed.
2. **Data Transfer Costs**: S3 charges for data transferred out of the storage. Outbound data transfer costs are determined by the volume of data retrieved by users and applications.

Insights and Considerations

Key insights emerged from our S3 cost analysis, offering valuable considerations for managing storage costs effectively:

- **Lifecycle Policies**: Implementing lifecycle policies can help transition objects to lower-cost storage classes as their access patterns change over time, thus optimizing

storage costs. A well-crafted lifecycle policy allows you to automatically transition objects to more cost-effective storage classes as they age. By moving data from Standard to Infrequent Access or Glacier storage, you can significantly reduce storage costs without compromising accessibility.
With a lifecycle policy, you can strategically tier your data based on its access frequency and relevance. Frequently accessed data remains in high-performance storage, while less frequently accessed data is moved to lower-cost storage classes. This optimization ensures that you're paying for the appropriate level of storage for each data type.

- **Data Deletion**: Regularly auditing and deleting unnecessary objects can prevent unnecessary storage costs. During our call you mentioned that you will look into doing cleanup of S3 unused buckets, this might be perfect opportunity to do it while saving costs

To cut down Amazon S3 costs, focus on two strategies:

1. **Lifecycle Policies**: Automate data transitions to cheaper storage classes based on access frequency, optimizing costs over time.
2. **File Deletion**: Implement systematic deletion practices to remove outdated or redundant files, freeing up storage and preventing unnecessary expenses.

EC2 Cost Analysis

Overview

Amazon EC2 (Elastic Compute Cloud) provides resizable compute capacity in the cloud, enabling organizations to deploy and manage virtual servers, known as instances. Understanding the cost structure of EC2 is crucial for optimizing compute resources, ensuring performance, and managing expenses effectively.

Usage Overview

In our previous discussion during the discovery phase, you highlighted the utilization of EC2 for your Kubernetes cluster, along with MongoDB and OpenSearch as integral components of the legacy configuration.

Here is the breakdown of EC2 costs in the last 6 months:

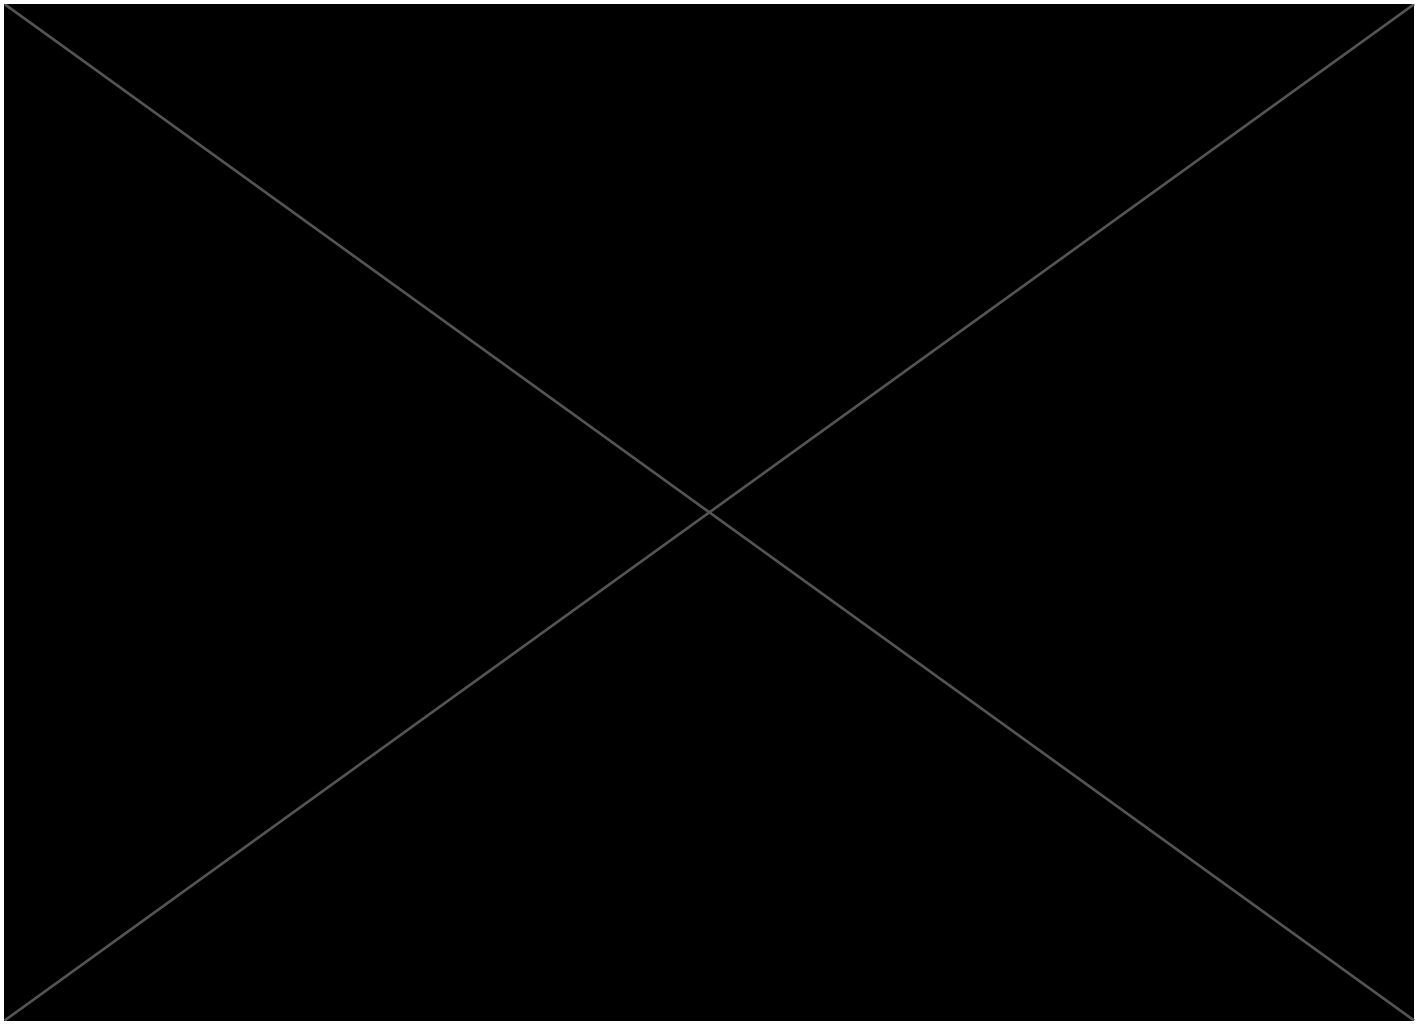

Picture 13. EC2 costs (Source: DoiT Console)

Instance Types and Distribution

- **Instance Types**: A variety of EC2 instance types are being used at the moment including:
  R5.2xlarge, c5.2xlarge, t3.2xlarge, c5.xlarge, c5.xlarge, t3.large

- **Disk Types**: EBS:VolumeUsage.gp2 and EBS:VolumeUsage.gp3
- **NAT**: Single NAT instance

Regional Distribution

- **AWS Regions**: I focused mostly on US-EAST region due to fact that most costs are coming from that region

Cost Breakdown

The cost components contributing to EC2 expenses encompass a range of factors related to instance provisioning, usage, and data transfer:

1. **Instance Costs**: The primary cost driver is the hourly charge for each running instance, determined by the chosen instance type and any additional features like Elastic IP addresses.
2. **Data Transfer Costs**: EC2 instances incur data transfer costs based on inbound and outbound data traffic. These costs vary depending on data movement across AWS regions and availability zones.

Insights and Considerations

Our EC2 cost analysis uncovered several insights to consider for cost optimization and resource management:

- **Instance Types:**
  If your workload allows, consider switching to different instance types. Please note that I included graviton processors as they tend to have the lowest prices. Here is an overview of potential savings for your environment's big spenders.

    1. **R5.2xlarge (Current Generation)**:
        - Newer Generation Equivalent: **R6g.2xlarge**
        - Benefits: The R6g instance type is powered by AWS Graviton2 processors, offering improved performance, better price-to-performance ratio, and potential cost savings. It's designed for memory-intensive workloads.

Instance name	vCPUs	Memory	Network Performance	Storage	On-Demand Hourly Cost
r5.2xlarge	8	64 GiB	Up to 10 Gigabit	EBS only	0.504

Instance name	vCPUs	Memory	Network Performance	Storage	On-Demand Hourly Cost
r6g.2xlarge	8	64 GiB	Up to 10 Gigabit	EBS only	0.4032

Picture 14. EC2 instance Type costs 1 (Source: AWS)

2. **C5.2xlarge (Current Generation)**:
    - Newer Generation Equivalent: **C6g.2xlarge**
    - Benefits: The C6g instance type uses AWS Graviton2 processors, providing higher performance at a lower cost compared to the C5 instances. It's well-suited for compute-intensive applications.

Instance name	vCPUs	Memory	Network Performance	Storage	On-Demand Hourly Cost
c5.2xlarge	8	16 GiB	Up to 10 Gigabit	EBS only	0.34
Instance name	**vCPUs**	**Memory**	**Network Performance**	**Storage**	**On-Demand Hourly Cost**
c6g.2xlarge	8	16 GiB	Up to 10 Gigabit	EBS only	0.272

Picture 15. EC2 instance Type costs (Source: AWS)

3. **T3.2xlarge (Current Generation)**:
    - Newer Generation Equivalent: **T4g.2xlarge**
    - Benefits: The T4g instance type, powered by AWS Graviton2, delivers better performance and cost efficiency. It's suitable for burstable workloads with balanced CPU and memory needs.

Instance name	vCPUs	Memory	Network Performance	Storage	On-Demand Hourly Cost
t3.2xlarge	8	32 GiB	Up to 5 Gigabit	EBS only	0.3328
Instance name	**vCPUs**	**Memory**	**Network Performance**	**Storage**	**On-Demand Hourly Cost**
t4g.2xlarge	8	32 GiB	Up to 5 Gigabit	EBS only	0.2688

Picture 16. EC2 instance Type costs 2 (Source: AWS)

4. **C5.xlarge (Current Generation)**:
    - Newer Generation Equivalent: **C6g.xlarge**
    - Benefits: Similar to the C6g.2xlarge, the C6g.xlarge instance type provides improved performance and cost savings, ideal for compute-intensive tasks.

Instance name	vCPUs	Memory	Network Performance	Storage	On-Demand Hourly Cost
c5.xlarge	4	8 GiB	Up to 10 Gigabit	EBS only	0.17

Instance name	vCPUs	Memory	Network Performance	Storage	On-Demand Hourly Cost
c6g.xlarge	4	8 GiB	Up to 10 Gigabit	EBS only	0.136

Picture 17. EC2 instance Type costs 3 (Source: AWS)

5. **T3.large (Current Generation)**:
    - Newer Generation Equivalent: **T4g.large**
    - Benefits: The T4g.large instance type offers better performance and cost-effectiveness compared to its T3 counterpart. It's a great choice for general-purpose workloads.

Instance name	vCPUs	Memory	Network Performance	Storage	On-Demand Hourly Cost
t3.large	2	8 GiB	Up to 5 Gigabit	EBS only	0.0832

Instance name	vCPUs	Memory	Network Performance	Storage	On-Demand Hourly Cost
t4g.large	2	8 GiB	Up to 5 Gigabit	EBS only	0.0672

Picture 18. EC2 instance Type costs 4 (Source: AWS)

It is worth mentioning that you should also try and run the latest generation if possible. The reason being is that older instance types are running on Nitro v1 virtualization platform and newer ones on v2. The second-generation hypervisor offers further enhancements in performance, security, and feature set compared

to the first generation. Basically, you want your instance to run on the newer hardware.

The reasons why you might consider Graviton processors; They offer several sustainability benefits due to their unique architecture and design. Here's how Graviton processors contribute to sustainability:

1. **Energy Efficiency**: Graviton processors are designed to be highly energy-efficient. They use the ARM architecture, which is known for its power efficiency. Graviton2 processors are built using the 7nm manufacturing process, which helps reduce power consumption while maintaining performance.
2. **Reduced Carbon Footprint**: Because Graviton processors consume less power, they contribute to reducing the carbon footprint associated with data centers and cloud computing. Lower power consumption translates to less energy used for computation, which in turn leads to lower greenhouse gas emissions.
3. **Better Performance-Per-Watt**: Graviton2 processors provide improved performance-per-watt compared to traditional x86 processors. This means that you can achieve higher computational output with the same or even less energy consumption, resulting in increased efficiency and sustainability.

- **Instance Right-Sizing**: Matching instance types with workload requirements is crucial. Utilizing the appropriate instance type can enhance performance while preventing overprovisioning. You mentioned during the discovery call that you are aware that EC2 instances are not right sized at the moment, I would highly advise to look into compute optimiser to rightsize instances.
- **Autoscaling**: Implement auto scaling to automatically adjust the number of instances based on demand. Set up Auto Scaling Groups with multiple instance types to take advantage of the most cost-effective options while maintaining availability.
- **Suspend instances** in off-business hours if not needed using the instance scheduler.
  All the instances that do not need to run during certain periods of time, consider switching them off. This is especially useful for development and testing environments.
- **GP3 volumes** offer potential cost savings compared to GP2 volumes due to their more efficient IOPS provisioning, higher baseline performance, and lower cost per gigabyte. Consider migration which could lead savings up to 20%
- **Reserved Instance Planning**: Utilizing AWS Reserved Instances can provide substantial cost savings over time. Properly identifying suitable instances for reservation is essential for realizing these benefits.

- **Spot Instances/Spot fleets**: Leverage Spot Instances for fault-tolerant, flexible, and cost-effective non-critical workloads. Spot Instances can provide substantial discounts compared to On-Demand instances, but they are subject to availability.
- **Instance Lifecycle Management**: Regularly reviewing and terminating underutilized or idle instances can prevent unnecessary costs and resource waste.
- **CloudWatch Alarms**: Set up CloudWatch alarms to receive notifications when your instances are underutilized or overutilized. This helps you make informed decisions about scaling and instance type adjustments.

Moving from traditional EC2 instances running Kubernetes and MongoDB to Amazon ECS/EKS (Elastic Container Service/Elastic Kubernetes Service) and AWS MongoDB (Amazon DocumentDB) can offer several potential ways to save money while also gaining benefits in terms of scalability, managed services, and reduced operational overhead.
It will all depend on specific needs and requirements of your application, but looking into current setup it is definitely worthwhile to investigate and perform a detailed cost analysis and workload assessment.

Additional services

Transfer Family

Looking into June costs, I noticed that you have a Transfer Family endpoint setup, with 141 MB of use, which you paid 216 US$ due to the fact that you are paying hourly service for the endpoint. You may consider removing this, if that is not in use that significantly and find an alternative.

Transfer Family		USD 200.89
US West (N. California)		**USD 216.01**
AWS Transfer Family SFTP:S3		USD 0.01
$0.04 per GigaByte uploaded over SFTP to S3 in US West (N. California)	0.141 GigaBytes	USD 0.01
AWS Transfer Family USW1-ENDPOINT		USD 216.00
$.30 per hour - SFTP Endpoint fee	720 Hourly	USD 216.00

Conclusion

In conclusion, our analysis of AWS cost utilization across CloudFront, S3, and EC2 services has revealed valuable insights and opportunities for substantial cost savings. By leveraging a combination of data-driven strategies, service optimization, and architectural improvements, we have identified actionable recommendations to optimize resource allocation and reduce unnecessary expenditure.

Appendix

AWS Cost Explorer:: AWS Cost Explorer

AWS Billing and Cost Management: AWS Billing and Cost Management

AWS Storage Lens: AWS Storage Lens

AWS Billing: AWS Billing

Doit: DoiT Console