

White Paper

# How to Measure Your AWS Data Traffic Costs: **The Complete Guide**

- Executive Summary ..... 3
- Breakdown of AWS Data Transfer Costs ..... 3
  - Data Transfer Out of AWS** ..... 4
  - Data Transfer Within AWS** ..... 4
  - Additional Considerations** ..... 5
    - Amazon S3 Transfer Acceleration** ..... 5
    - Content Delivery Network** ..... 5
- Measuring Your AWS Data Transfer Cost ..... 6
  - Native AWS Tools** ..... 6
    - Billing and Cost Management Dashboard** ..... 6
    - Cost Explorer** ..... 8
    - Simple Monthly Calculator** ..... 8
    - Billing Alarms** ..... 9
  - Third-Party Tools** ..... 9
  - Create Your Own Tools** ..... 10
- Optimizing Your Data Transfers ..... 10
  - Incremental Syncing** ..... 10
  - Compression** ..... 10
  - Deduplication** ..... 10
- Summary ..... 11
- About NetApp® ONTAP Cloud for AWS ..... 11

## Executive Summary

Even though cost efficiency was a main reason your company moved to the cloud, staying on top of costs is still a concern and responsibility of IT managers. You think you've planned for every service you're going to consume, but then the monthly bill arrives, and somehow your spending is higher than expected.

What happened? Why were you caught off guard? Often poor planning is the root cause of unnecessary spending, and data transfer costs are what most people overlook or miscalculate.

In this white paper, we review some tools at your disposal that you can use to measure your data traffic costs and understand exactly what is causing the majority of data traffic costs in your cloud infrastructure. We also include tools that can help you not only predict the costs, but also stop overspending before it is too late.

## Breakdown of AWS Data Transfer Costs

With AWS, you have the cost of the services (charged per hour, per reserved instance, per storage unit, and so on, depending on the service), and you have the cost of the data transfer.

When it comes to data transfer, there are specific services that typically generate the majority of the costs. For most AWS users, these costs arise from using services that store a lot of data, such as Amazon S3 storage, RDS, and DynamoDB databases, or even EBS volumes that are attached to EC2 instances. The more data you request from these services, the higher the data consumption traffic—and the higher the cost.

The pricing of data transfer also depends on whether your traffic is flowing within a region (or between availability zones) or moving out of AWS completely. The further data must travel from its place of storage in the cloud, the higher your data egress costs. Keep this fact in mind when architecting your infrastructure.

When we monitor data on AWS, we note the data traffic both in and out. Data traffic into an AWS region from any service in any other region is free. In reviewing outbound traffic, we are interested in the traffic flowing both within AWS and outside of AWS. The latter is much easier to track. Let's take a closer look at how.

## Data Transfer Out of AWS

In this case, data going outside of AWS to the public internet is \$0.090 per GB/month for the first 10TB in the U.S. East 1 (Virginia) region. The next 40TB are \$0.085 per GB/month, then the next 100TB are \$0.070 per GB/month, and so on, as you can see on the [AWS official price list](#) in Table 1.

Table 1: AWS Official Price List

First 1 GB / month	\$0.000 per GB
Up to 10 TB / month	\$0.090 per GB
Next 40 TB / month	\$0.085 per GB
Next 100 TB / month	\$0.070 per GB
Next 350 TB / month	\$0.050 per GB
Next 524 TB / month	<a href="#">Contact Us</a>
Next 4 PB / month	<a href="#">Contact Us</a>
Greater than 5 PB / month	<a href="#">Contact Us</a>

Prices vary by region, as shown in Table 2.

Table 2: Prices by Region

AWS Region	Cost to transfer up to 10 TB OUT from EC2 to internet
US East (N. Virginia)	\$0.09 / GB
US West (Oregon)	\$0.09 / GB
US West (Northern California)	\$0.09 / GB
EU (Ireland)	\$0.09 / GB
EU (Frankfurt)	\$0.09 / GB
Asia Pacific (Singapore)	\$0.12 / GB
Asia Pacific (Tokyo)	\$0.14 / GB
Asia Pacific (Sydney)	\$0.14 / GB
South America (São Paulo)	\$0.25 / GB
AWS GovCloud (US)	\$0.155 / GB

## Data Transfer Within AWS

The traffic flowing within AWS is a bit more complicated: data flows from one region to another, but also within regions. Also, for the data in the same region, there is a difference between the data going from one AZ (Availability Zone) to another and data moving within the same AZ.

### Data Transfer Across Regions

All data across AWS regions is considered internet data traffic and is charged at \$0.02/GB. (Don't forget the hidden costs of provisioning services such as site-to-site VPN required to connect private subnets from the two regions.)

### Data Transfer Within a Region

#### Across availability zones in the same region.

Data that moves from one AZ to another within the same region is considered a regional data transfer. It is charged at \$0.01/GB.

#### Within availability zone.

Data transfer in the same region and within the same AZ is free as long as you use a private IP address. Using a public IP, elastic IP, or IPv6 address causes your data to traverse the public internet and increases your cost to \$0.01/GB.

## Additional Considerations

Aside from data transfer costs incurred for data moving within or outside of AWS, there are some additional expenses that should be considered, particularly for certain services.

### Amazon S3 Transfer Acceleration

If you consume Amazon S3 storage services, and you are also using the Amazon S3 Transfer Acceleration service, there are additional data transfer fees that you should take into account, as shown in Table 3.

**Table 3: Additional Data Transfer Fees**

Data Transfer IN to Amazon S3 from the Internet:	
Accelerated by AWS Edge Locations in the United States, Europe, and Japan	\$0.04/GB
Accelerated by all other AWS Edge Locations	\$0.08/GB
Data Transfer OUT from Amazon S3 to the Internet:	
Accelerated by any AWS Edge Locations	\$0.04/GB
Data Transfer between Amazon S3 and another AWS region:	
Accelerated by any AWS Edge Locations	\$0.04/GB

## Content Delivery Network

If data transfer produces a significant cost for you, you should look into using [CloudFront](#), a content delivery network (CDN) that accelerates delivery of your websites, APIs, video content, and other web assets through CDN caching. It stores a cached version of the content in multiple geographical locations—called edge locations—to minimize the distance between the customer and the server, allowing the customer faster access to the content.

The data traffic out of Amazon S3 to CloudFront is not charged. In addition, in some regions, in high-volume transfers, CloudFront can help lower data transfer costs for the delivery of your content by using the edge locations.

# Measuring Your AWS Data Transfer Cost

There are some useful tools to measure and break down your current costs, as well as predict future ones.

## Native AWS Tools

Aside from data transfer costs incurred for data moving within or outside of AWS, there are some additional expenses that should be considered, particularly for certain services.

## Billing and Cost Management Dashboard

The [Billing and Cost Management Dashboard](#) displays a general overview of your spending. Here you can see how much you pay for each service in a specific month. You can get more details through filtering and pinpoint the source of the spending. For example, under the “Bills” tab, you can choose the specific month and then see the cost breakdown.

Because we are interested in data transfer costs, you can review in greater detail by opening Data Transfer (under Details for AWS service charges) (Table 4) and then select the desired region. Table 5 shows the bandwidth costs for the chosen month.

See tables 4 and 5 on page 7.

## Lower the Cost of Data Transfer to and from the Cloud

Uploading and downloading large files to and from the cloud often take a lot of time, even if your internet connection is fast. When it comes to sending data over the internet, Transmission Control Protocol (TCP) is the method of choice. However, TCP was never designed for moving large files over the open internet, and when you need to move, for example, a large video file, the data transfer speed can vary greatly.

## Transferring Large Files from an Amazon S3 Bucket to a Data Center

If we consider the example of transmitting a large 30GB video file using a 100Mbps connection, we can estimate the transfer time at close to 42 minutes. However, this estimate is only in theory, because we are ignoring the impact of not only latency, but also the potential packet loss, which would lead to retransmission of the lost packets, therefore increasing the transfer time.

Latency itself can be a huge factor, and you can expect your transfer speed to be reduced by 10% to 15% for traffic within the same geographical area and probably around 20% to 25% for traffic crossing the continents. The greater the distance between your Amazon S3 bucket location and your data center, the greater the latency.

We also have to take into consideration that our data needs to be packed into a TCP/IP packet, which increases the total size of the data sent, increasing the time as well.

## Billing and Cost Management Dashboard

Table 4: AWS Service Charges Under Details

Details	
<b>AWS Service Charges</b>	\$31,980.26
CloudTrail	\$0.00
CloudWatch	\$0.00
Data Pipeline	\$0.04
Data Transfer	\$2,304.09
DynamoDB	\$5,044.71
Elastic Compute Cloud	\$5,858.07
Elastic File System	\$0.74
Elastic MapReduce	\$6.51
Glacier	\$887.24
Key Management Service	\$0.00
RDS Service	\$19.65
Simple Notification Service	\$0.01
Simple Queue Service	\$40.59
Simple Storage Service	\$15,718.22
SimpleDB	\$0.00
Support (Developer)	\$49.00
Virtual Private Cloud	\$74.40

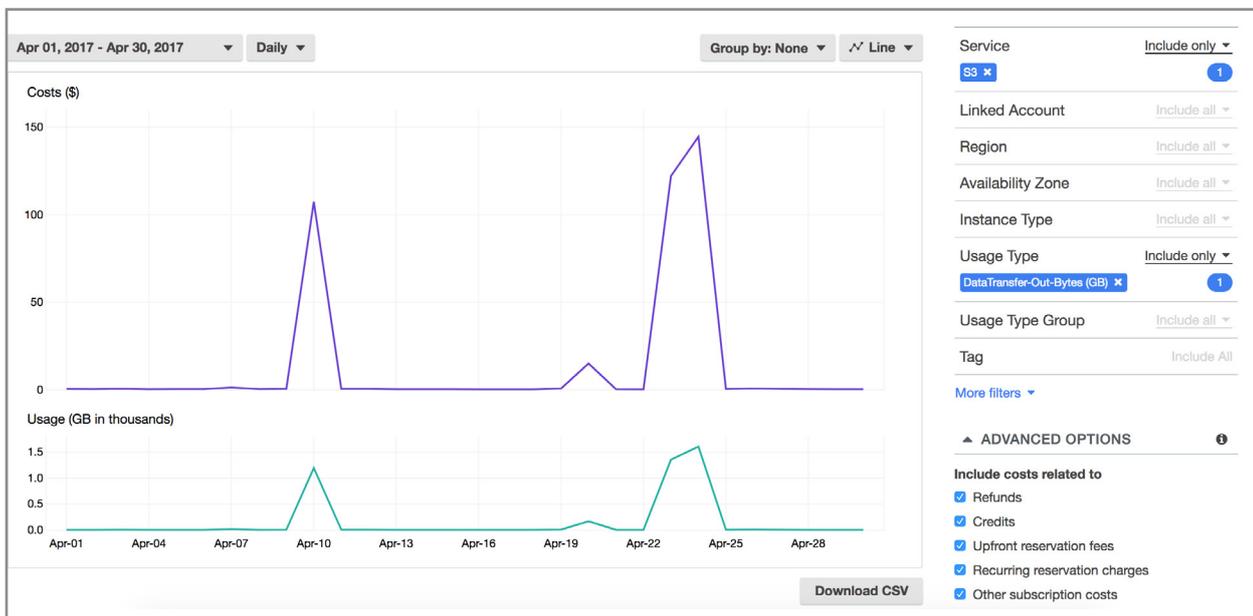
Table 5: Bandwidth Costs for the Chosen Month

Bandwidth		\$1,584.16
\$0.000 per GB - data transfer in per month	8,105.089 GB	\$0.00
\$0.000 per GB - first 1 GB of data transferred out per month	1 GB	\$0.00
\$0.010 per GB - regional data transfer - in/out/between EC2 AZs or using elastic IPs or ELB	12.260 GB	\$0.12
\$0.085 per GB - next 40 TB / month data transfer out	7,794.525 GB	\$662.53
\$0.090 per GB - first 10 TB / month data transfer out beyond the global free tier	10,239 GB	\$921.51

## Cost Explorer

[Cost Explorer](#) is a great tool for viewing charts of your costs for the desired period of time. Cost Explorer offers many options for filtering: by time range, for instance. With usage type filtering, you can find the DataTransfer-Out-Bytes filter. You can also add more filters, for example, Amazon S3 (service filter), so that you look only at those specific costs for the past month, as shown in Figure 1.

Figure 1: Costs by Month



More advanced filtering options are also available.

If you have multiple projects in your AWS infrastructure, you can open the charts based on the user-defined reports and then filter out only the data you require from specific projects.

Cost Explorer requires billing rights, which are by default only reserved for root accounts. Don't forget to delegate this permission to users who should view billing information. If you use consolidated billing (several AWS accounts under one "master" through which you aggregate costs), you can select which user with billing rights can view specific accounts or all of them.

## Simple Monthly Calculator

AWS Simple Monthly Calculator (Figure 2) is a tool to predict the costs of your data transfers (if they're constant). There are multiple options, and you can choose interregion data transfer, data transfer outside the AWS, and others. You can specify how many gigabytes per day/week/month or terabytes per month you expect to transfer and receive an estimate for your monthly bill.

Figure 2: AWS Simple Monthly Calculator

**Services** Estimate of your Monthly Bill (\$ 49.65)

Choose region: US-East / US Standard (Virginia) Inbound Data Transfer is Free and Outbound Data Transfer is 1 GB free per region per month

Amazon Elastic Compute Cloud (Amazon EC2) is a web service that provides resizable compute capacity in the cloud. It is designed to make web-scale computing easier for developers. Amazon Elastic Block Store (EBS) provides persistent storage to Amazon EC2 instances.

**Compute: Amazon EC2 Instances:**

Description	Instances	Usage	Type	Billing Option	Monthly Cost
+ Add New Row					

**Compute: Amazon EC2 Dedicated Hosts:**

Description	Number of Hosts	Usage	Type	Billing Option
+ Add New Row				

**Storage: Amazon EBS Volumes:**

Description	Volumes	Volume Type	Storage	IOPS	Baseline Throughput	Snapshot Storage
+ Add New Row						

**Elastic IP:**

Number of Additional Elastic IPs:

Elastic IP Non-attached Time:  Hours/Month

Number of Elastic IP Remaps:  Per Month

**Data Transfer:**

Inter-Region Data Transfer Out:  GB/Month

Data Transfer Out:  GB/Month

Data Transfer In:  GB/Month

VPC Peering Data Transfer:  GB/Month

Intra-Region Data Transfer:  GB/Month

Public IP/Elastic IP Data Transfer:  GB/Month

Of course, you can use this calculator for more than just data transfer cost estimates, and it can play a huge role in the planning of your overall AWS infrastructure.

## Billing Alarms

[Billing alarms](#) can be very beneficial if you have a predefined monthly budget you don't want to exceed. This kind of alarm is also useful if you are in the AWS free tier and do not want to exceed its limits.

## Third-Party Tools

Although AWS tools allow you to monitor and track your data traffic, they are often not enough to determine the source of the majority of your costs, especially when you manage hundreds or thousands of resources in the cloud.

Various third-party tools that are not free provide you with a more granular look at your monthly spending. All of these tools rely on analysis of the AWS billing report, which is updated daily and sent to your bucket of choice.

For example, using these tools, you can pinpoint which of your Amazon S3 buckets are the most expensive or even how much you paid for transfers of a specific file within a bucket. These tools also allow you to search for costly untagged resources and access other important information about your monthly spending.

These alternatives can assist you in the detection and control of unexpected data transfers. A high level overview doesn't suffice when you are on the lookout for the precise sources of data traffic costs. With the filters—such as costs separated by project or team—provided by many of these tools, you can get an accurate breakdown, giving you clear insight into details such as which service is transferring the most data.

## Create Your Own Tools

Sometimes the tools offered, either native or third-party, are not exactly what you need. Instead, you require a tool customized for your environment. If so, you can create your own tools or find one that was already created by someone else and modify it to fit your needs.

One thing you can use when creating your own tools is [AWS Lambda](#). Lambda is a compute service that allows you to run code without provisioning servers and with zero administration. Instead, you only need to provide your code in one of the supported languages (Node.js, Java, C#, and Python), and Lambda runs it on a high-availability compute

infrastructure for you, taking care of capacity provisioning, code monitoring, and logging. Using some Lambda functions and the [AWS price list API](#), you are able to create your own tools that predict costs for you as you go. If you're using CloudWatch, there would be at least a five-minute delay. Keep in mind that this delay would represent an approximate spending, because the data transfer cannot be extracted 100% from CloudWatch metrics alone, and all data would be considered as data going out to the internet. The only precise measurement of your pricing comes from analysis of the AWS billing reports.

## Optimizing Your Data Transfers

Optimizing your data in one of the following ways can greatly affect your data storage as well as your data transfers.



### Incremental Syncing

Incremental syncing is a technique in which, after the initial full synchronization of the data, each subsequent synchronization only focuses on the differences in files. Only files that were created or modified since the last sync are copied to the destination. This approach greatly reduces the amount of data transferred.



### Compression

Data compression is a very useful tool that saves disk space by packing the data into a smaller size (often significantly). When looking at data transfer, compression is often used, because the reduced data size reduces the used bandwidth.



### Deduplication

Data deduplication is a specialized compression technique that eliminates duplicate copies of repeating data. Deduplication uses analysis to find patterns and replaces the duplicates with a small reference that points back to the stored copy of that pattern. It can be used to improve storage utilization, but is also very important in data transfer, because it can reduce the necessary number of bytes sent.

## Summary

Moving to cloud can be very beneficial for companies today. The change in pricing model, the ability to easily scale on demand, and the access to various services that are offered provide a lot of options for companies of various sizes. However, fully utilizing the benefits of the cloud, while not exceeding the desired budget, requires not only good planning, but also good data management and staying on top of your spending each month.

NetApp offers NetApp® [ONTAP®Cloud](#) and [Cloud Sync](#) products, which utilize the mentioned techniques and can help you streamline and optimize data transfer costs, allowing you to maintain a cost-efficient cloud infrastructure.

## About NetApp ONTAP Cloud for AWS

[ONTAP Cloud](#), an offering in AWS of the leading enterprise storage operating system, is deployed using [OnCommand® Cloud Manager](#) to deliver secure, proven NFS, CIFS, and iSCSI data management for AWS EBS and Amazon S3 storage.

A software-only storage service running NetApp ONTAP software, ONTAP Cloud combines data control with enterprise-class storage features such as data deduplication and compression to minimize your EBS storage footprint. You can take Snapshot™ copies of your data without requiring additional storage or affecting your application's performance.

And ONTAP Cloud can tie your Amazon cloud storage to your data center using the leading NetApp replication protocol, SnapMirror®.

To find out more, go to [cloud.netapp.com](http://cloud.netapp.com).

© 2017 NetApp, Inc. All Rights Reserved. NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.