



E-BOOK

Say Goodbye to Egress Fees

March 2023


Content

- 03** Overview
- 04** Ingress and Egress
- 05** How Do Egress Fees Work?
- 07** The Problem with Egress Fees
 - 07** Padding Provider's Bottom Lines
 - 08** Increasing Lock In
 - 08** Limiting Innovation
- 09** Why Cloudflare Is Different
- 09** How Cloudflare Helps Reduce or Eliminate Egress Fees
 - 09** Bandwidth Alliance
 - 09** Cache Reserve
 - 10** Object Storage with No Egress Fees
- 11** Revisiting Egress Fees

Overview

In recent years, a huge number of businesses and organizations have moved all or some of their workload to the cloud. In late 2021, [DigitalOcean found](#) that twenty-one percent of non-technical, small- and medium-sized businesses were using cloud-based storage solutions, while forty-five percent of enterprises were, and those numbers are increasing every day.

Cloud storage is flexible and scalable, and ensures that you only pay for what you're using. It can also eliminate most, if not all, of the overhead of maintaining your infrastructure, since you don't have to own or manage the servers. Many businesses look to cloud storage as a cost-saving solution—after all, it only costs pennies a gigabyte to store your data, and frees you from time-consuming maintenance and expensive hardware. Hidden in the pricing charts, though, is a line item that many people overlook: data egress fees.



Ingress and Egress


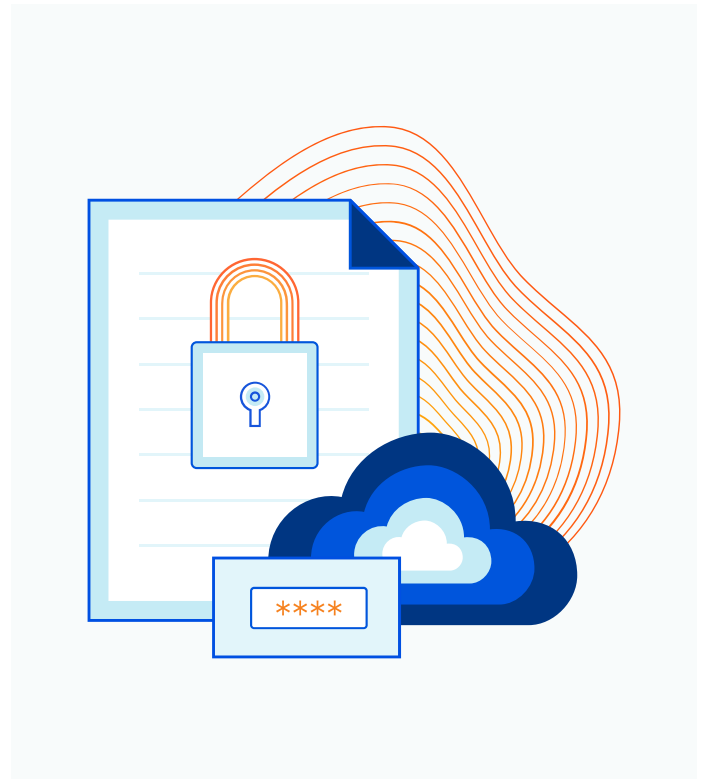
Ingress and egress are fairly straightforward concepts. Ingress is when a cloud service takes your data in, however that happens. It might be automatically loaded through a pipeline, user-generated content in an app, or uploaded directly in bulk. Data gets uploaded to services so you have a centralized area from which you can work with it—it can be served to customers, transformed, analyzed, and otherwise used or manipulated. Egress on the other hand, is getting the data back out of that service, such as by serving it to users, using it to generate a website, or downloading it to a local datastore as a backup.

Ingress fees are usually very simple: there aren't any. Almost every provider allows you to upload your data to their platform without any loading fees, no matter how much data you have or how it's being uploaded. Some providers go a step further, and offer secondary services to make it easy for you to send them as much data as you feel is necessary. Amazon's Snowball service, for example, helps you "easily migrate terabytes of data."

Ingress is nearly frictionless, and once your data has been ingested, you'll pay relatively minimal charges for storage, as shown in the chart in the next section. These charges are often around \$.02 USD per GB per month, which works out to around \$20 USD per terabyte. Storage tends to become less expensive, on a per-gig basis, as the amount you're storing increases. For example, storage of up to 250 TB could cost \$.024 GB per month, and if you're storing more than 250 TB, the price for additional storage might be \$.021 GB per month.

You'll also have to pay for operations performed on the data, whether that's simply reading it or using functions as a service, such as Cloudflare Workers, to perform complex data transformations.

So far, so good. But there's another major factor to consider, because with most providers, what you pay for isn't just storage and compute, but also egress, sometimes referred to as bandwidth or data transfer. Egress is getting data back out of the service to which it was uploaded, and this is where things get complicated.



How Do Egress Fees Work?

Most providers charge you for taking data out of their system, whether that's using it to fulfill a user request, downloading the data that you've uploaded, or even just moving it to another storage region with the same provider. All of these are common requests, and essential components of a wide range of use cases, but many cloud providers will charge you up to \$0.09 USD per gigabyte transferred.

While \$0.09 per gigabyte may not sound like much, hopefully you have either an expansive budget or minimal data, because egress charges can run up very quickly. There are plenty of [people trying to understand](#) why their egress fees are so expensive, and others who [forgot about them completely](#)—a real problem when egress is often substantially more expensive than storage.

For ease of comparison, the chart below looks at the standard or default storage options that each provider offers. Unless otherwise stated, all regions are US-east, which is generally the least expensive region, and prices are in US dollars. It's also worth noting that this chart is only considering storage and egress fees, and doesn't include charges that you might incur for reading, writing, or transforming your data.

	AWS S3	Google Cloud	Azure Blob Storage	IBM Cloud
Ingress	Free	Free	Free	Free
Storage price per GB	\$.023	\$.02	\$.021	\$.022
Egress price per GB	\$.09	\$.12	\$.087	\$.09
Region transfer price per GB	\$.02	\$.05	\$.05	\$.09

*prices as of January 2023

To put those prices in perspective, let's consider a scenario where you have a 5 TB archive. This archive is stored in US-east, but you'd like to download a copy for local storage, and for availability reasons, you'd like to have a duplicate copy stored in the EU region. **Region transfer** is sending data from one geographical region to another. It helps to ensure data availability, since if [one region goes down](#), your data will still be accessible. This type of redundancy is often recommended, if not required, for security certifications and regulatory compliance. Hosting data in multiple regions can also reduce latency, especially if you have a far-flung userbase.

	AWS S3	Google Cloud	Azure Blob Storage	IBM Cloud
Storage (5 TB)	\$117.76	\$102.40	\$107.52	\$112.64
Egress (5TB)	\$460.80	\$614.40	\$445.44	\$460.80
Region transfer, US-east to EU (5TB)	\$102.40	\$256.00	\$256.00	\$460.80
Total cost	\$680.96	\$972.80	\$808.96	\$1034.24

*prices as of January 2023

As you can see, storing your data in the cloud is easy and inexpensive—until you try to use it. In some instances, it costs six times more to move a gigabyte of your data than it does to store it!

It's not all bad, though. Many providers offer free or discounted egress fees if you're keeping the data within that provider's ecosystem. For example, if you're storing data in Amazon S3, you'll pay \$.09 per GB for data transfer if you want to use Google Cloud for your compute. If you want to use Amazon's Elastic Compute in the same region, though, and connect to it through the service endpoint, you can move your data free of charge; if you'd like to use EC2 in a different region, it's only the \$.02 per GB region transfer fee seen above.

Other providers offer similar arrangements, and assuming that you plan to stay within the same ecosystem, they can seem like a pretty good deal. But what if you don't want to stay there forever? What if you want to pursue a multicloud strategy, to move your data to somewhere that's more optimized for your specific use case, or to create an on-site backup? What if you'd just like the ability to use whatever services you'd like—without worrying that you're going to be charged a small fortune in egress fees?

The Problem with Egress Fees

While consumers tend to dislike egress fees, most providers insist that they're necessary, even beneficial. Unfortunately, this claim doesn't hold up to scrutiny.

Padding Provider's Bottom Lines

At first glance, egress fees might seem like a reasonable charge, and many providers are quick to assure you that all they're doing is passing on their costs. After all, storing data is one service, and moving it is another, and the infrastructure that transports information around the world has to be maintained by someone. The problem with this argument is that egress is actually very inexpensive on the provider side. It doesn't cost exponentially more than ingress or storage—in fact, egress fees are often marked up hundreds of times more than the cost that's actually incurred by the provider.

When a provider tells you that they're just passing on the fees for data egress, they do so hoping that you'll assume that bandwidth is similar to many other consumer products. Stores purchase goods **wholesale**: they agree to purchase a large quantity in exchange for a lower price. The store then increases (or **marks up**) the price and sells it to the consumer.

Markup varies by industry. Staples like groceries often have markups in single-digit percents, while clothing and furniture often sees markups of one hundred to two hundred and fifty percent. It's difficult to get definite numbers for electronics, such as phones, but they're often estimated to come with a markup of between seventy-five and one hundred and twenty-five percent.

After all that, \$0.09 per gigabyte for data transfer might seem like a bargain—until you discover that the best estimate for the markup on it is as much as **8,000 percent** (for North American-based customers). This is possible because of how bandwidth is sold at the wholesale level. In the

examples above, money is being exchanged for an object, whether that's a phone or a jug of milk. This is how retail bandwidth works, as well: for your \$0.09, you get one gigabyte of data transfer.

With wholesale bandwidth, what's being purchased is the capacity to transfer data, otherwise known as **transit**. The service provider pays for bandwidth not based on how much data is actually transferred, but based on how much data, in Mb per second, they transfer during their peak use of the network. Wholesale bandwidth isn't like the store buying some milk which is then sold to you. Instead, they're paying a relatively low price to maintain a supply of milk that will meet their needs at peak demand, regardless of how much demand there currently is.

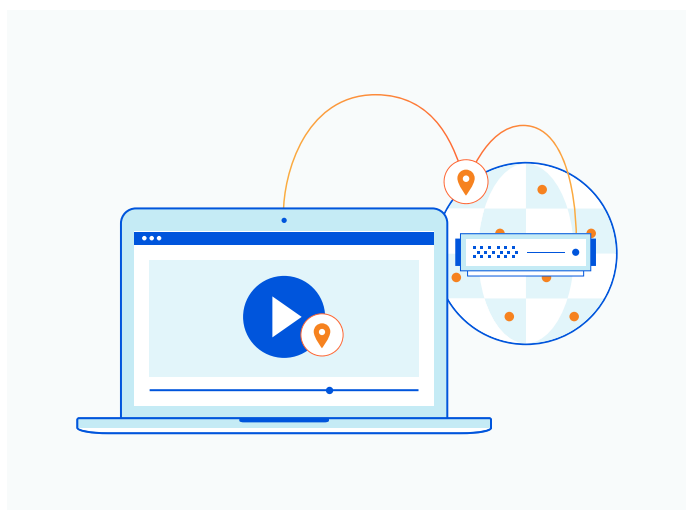
Wholesale bandwidth has an enormous markup, but that assumes that the service provider is paying for bandwidth in the first place. The internet is often described as a "network of networks." These networks can be owned by major telecommunications companies, as in the case of transit, but they can also be privately owned by smaller ISPs, or by cloud service providers such as Cloudflare, Google, Amazon, or IBM. Corporations and organizations of this size often come to an agreement known as **settlement-free peering**, which means that the organizations have agreed to a reciprocal arrangement in which they can use each other's networks free of charge. This allows the participating networks to route things quickly and efficiently to end users, frees organizations from needing to maintain networks in every corner of the globe, and makes services more stable by eliminating single points of failure.

Wholesale bandwidth might be marked up by a percentage of thousands, but the markup on data transfer over peered connections is effectively infinite. There's no charge to the service providers, but the "cost" is still passed on to the end users, padding the service provider's pocket and the user's bill.

Increasing Lock In

By making it easy for you to consolidate all your data on their servers, and making it quite expensive for you to get that data back out, cloud providers are effectively using egress fees as a way to create lock in. This helps to ensure that even if you might be served better by another provider or another toolset, you'll think twice about moving. After all, your application is working well enough and making a profit, despite the inflated egress charges that you're paying. No one wants to see a sudden line item for tens of thousands of dollars on their balance sheet, and providers hope that the threat of it is enough of a disincentive that you'll stick with them.

The internet is full of egress fee horror stories. Many users move to the cloud expecting to save on costs, only to discover that their careful planning and well-architected solutions didn't account for egress fees. Users can easily run up enormous charges before they even realize it's a charge to worry about, and because egress is billed in arrears, the bills are unpredictable, making it difficult to budget for them. Estimate it too low, and you'll find that you're quickly over budget; estimate it too high, and money that could have been used for core areas of your business is tied up unnecessarily. In either case, a burst of popularity or a misconfigured deployment can result in an unwanted and unexpected bill, and users are left with little recourse.



Limiting Innovation

Other costs are more subtle, and even easier to overlook until it's too late. One example of this is the way that egress fees hinder your ability to build multicloud architectures.

There are a number of advantages to a multicloud approach. It helps you avoid the lock-in mentioned above, which is a major concern for many businesses, and can help reduce costs by allowing you to take advantage of favorable pricing models for specific tasks. It offers improved scalability and resilience, as well; spreading your workload across multiple cloud providers reduces the likelihood that an outage at one of them will take your service offline with it. In some cases, you may need a multicloud approach to comply with consumer protection regulation, or to meet the criteria for security certifications.

Multicloud strategies also allow you to take advantage of the most exciting features being offered by each provider, which helps to ensure that you can use the technologies that are the best fit for your application rather than having to use whatever your cloud provider offers.

The catch, of course, is that to take advantage of a multicloud approach, you have to be able to move data freely between those clouds. Egress fees are designed to inhibit that ability, and create artificial constraints on your ability to innovate and take advantage of new functionality.

When you're paying egress fees, you shortchange your own products. The fees funnel money away from what should be core areas of focus, such as innovation and development, and into the pockets of service providers, slowing your progress and inflating your costs. Not only do you see a lower return on the money you invest in your application, but you also deliver less value to your customers.

Why Cloudflare Is Different

Cloudflare's long-standing goal has been to help build a better internet, one that's fast, safe, and affordable, and where information is shared freely. To further this goal, we have a deep commitment to building products that benefit everyone, such as offering a free global content delivery network (CDN) and domain name services.

In light of this, it shouldn't surprise you to learn that Cloudflare has chosen to take a different approach to egress fees. Egress fees aren't the best version of the internet. They aren't reflective of real costs, and they artificially limit consumer choice. They also place an unnecessary burden on developers, slowing innovation and introducing unwanted—and unneeded—friction to the development process. No one should be trapped in the walled garden of a cloud provider's ecosystem, or be forced to choose between doing what's right for their application and what's right for their accountant.

How Cloudflare Helps Reduce or Eliminate Egress Fees

As part of our commitment to building a better internet, Cloudflare has taken on egress fees, and is doing our best to reduce or eliminate them through a number of innovative programs.

Bandwidth Alliance

In 2018, Cloudflare launched the [Bandwidth Alliance](#), which is a group of cloud providers that have banded together to say "enough" to egregious egress charges. It now boasts twenty members, and includes providers such as Azure, Google Cloud, Oracle, and Alibaba Cloud.

To help combat egress charges, members agreed that data being transferred between any of the participating providers could be moved for a deeply discounted rate. Currently, eighteen of the providers allow for free data transfer to other participating services, while Google Cloud and Azure offer substantial discounts.

Data is delivered via Cloudflare's global network to data centers owned by Bandwidth Alliance partners, dramatically reducing data transfer fees. Cloudflare has one of the most extensive networks in the world, with a presence in over a hundred countries, and is constantly expanding. This allows for data to be transferred without the need for anyone to incur substantial expense, which means that there's no need to pass the negligible costs on to the end users.

Our research has found that reducing or eliminating egress fees [saved customers](#) between 7.5 percent and twenty-seven percent of their total monthly bill.

Cache Reserve

Caching is when content is stored at a location between the server from which it originated and the end user. Once the content has been cached, it can be served to the user from the cache rather than from the origin server. This is beneficial to users because the closer proximity of the content means that they can access it more quickly on subsequent views. It's also beneficial to site owners, because delivering content from the cache, rather than the server, allows them to avoid paying egress charges for the same piece of content repeatedly.

Cloudflare has always offered CDN caching, caching content in our global data centers to speed page loads and minimize requests to the server. The content is then held in the cache until it either reaches the expiration that was set in its headers or it's evicted. Evictions happen when the cache's storage space is full, and some of the previously cached content is removed to make room for newer content to be cached. We use an algorithm called "least recently used" to determine what gets evicted when the storage is full.

While this solution does an excellent job of ensuring that the most-requested content is served from a cache as often as possible, the downside is sites with a long-tail of content that isn't frequently accessed, or low traffic volumes can be negatively impacted.

To mitigate this, Cloudflare has introduced [Cache Reserve](#).

Cache Reserve is a persistent data store that's designed to serve as a complement to standard CDN caching, and serves as a sort of cache of last resort. If a user requests something, Cloudflare will first look to our standard cache. If the content is there, which is called a **cache hit**, the content will be served to the user. If the content isn't there, called a **cache miss**, it will then be requested from Cache Reserve rather than the server of origin. It's only if the content is missing from both the CDN cache and Cache Reserve that the server of origin will be hit.

Content can stay in Cache Reserve for significantly longer than it can stay in a standard CDN cache. By default, content that's still fresh will be retained for thirty days, and if it's requested in that time, another thirty days is added. This allows you to serve as much content as possible from the cache, improving performance and reducing egress fees.

Object Storage with No Egress Fees

The biggest challenge that Cloudflare has made to the egress fee landscape, though, is [R2](#): S3-compatible object storage hosted on Cloudflare's global network. Unlike most object storage, R2 is regionless—by default, there's no need to select a region, and files are distributed to the best region on Cloudflare's network based on where they're being uploaded from, though restrictions on data transfer between regions can be enabled to comply with data governance regulations. Regionless storage helps ensure that both developers and end users can access content from somewhere that's close to them, rather than hoping that everyone can be served quickly from US-east.

To empower developers to choose the tools and services that work best for them, R2 has no egress fees. Your data is yours, and your ability to use it shouldn't be paywalled. There's not a "gotcha" in the form of increased storage fees, either: S3 charges \$.023 USD per GB of storage, while R2 charges only \$.015 USD. This means you'll have not just a lower bill, but a more predictable one.

Part of the reason Cloudflare has been able to do this is our unique position within the market.

Since Cloudflare has an extensive CDN, its cache sits in front of an object storage bucket. The cache has to exist whether it sees any egress or not, but since it has to exist anyhow, we decided that we'd rather not pass made-up fees on to customers just to pad their bottom line.

R2 offers flexible, performant object storage, and is well integrated with other Cloudflare products, such as Workers, Durable Objects, and the Cache API. It's also fully compatible with the S3 APIs, which offers a mature toolkit and allows for R2 to be easily integrated with your existing workflows, applications, libraries, and tools. If what you're using works with S3, it will work with R2, as well.

This compatibility also means that migration from one service to the other is simple, and can be done incrementally, allowing you to start realizing cost savings immediately. While you can migrate manually or use a Cloudflare Worker to manage it for you, Cloudflare has also recently released the [R2 Super Slurper](#) that will handle all of the heavy lifting for you. R2 Super Slurper has two modes: sipping and slurping.

In slurp mode, the R2 Super Slurper just needs to be pointed to an S3 compatible bucket, and it'll slurp up the contents of either a folder within the bucket or the whole bucket, then deposit the data into an R2 bucket. It's fast and easy, but if you prefer a more gradual approach, you can use sip mode, instead. In sip mode, objects are moved from S3 to R2 as they're requested. When an object is requested, Cloudflare first checks if the object is available on R2. If it is, it's served to the user; if it isn't, it's requested from S3, then served to the user and also saved in R2, from which it will be served on subsequent requests.

The gradual approach has the benefit of making sure that your most frequently requested objects are served from R2 as soon as possible, helping you maximize savings during transition. It also enables you to consolidate data from multiple buckets, or even multiple providers, as it can sip from any S3 compatible source. The R2 Super Slurper works with S3 compatible buckets now, and more sources are expected to become available later in 2023.

Revisiting Egress Fees

Let's look again at the chart illustrating storage fees versus egress fees and see what it looks like when R2 is added.

	AWS S3	Google Cloud	Azure Blob Storage	IBM Cloud	Cloudflare R2
Ingress	Free	Free	Free	Free	Free
Storage price per GB	\$.023	\$.02	\$.021	\$.022	\$.015
Egress price per GB	\$.09	\$.12	\$.087	\$.09	Free
Region transfer price per GB	\$.02	\$.05	\$.05	\$.09	Free

*prices as of January 2023

As before, let's also look at what it would cost to back up a 5 TB archive.

	AWS S3	Google Cloud	Azure Blob Storage	IBM Cloud	Cloudflare R2
Storage (5 TB)	\$117.76	\$102.40	\$107.52	\$112.64	\$76.80
Egress (5TB)	\$460.80	\$614.40	\$445.44	\$460.80	Free
Region transfer, US-east to EU (5TB)	\$102.40	\$256.00	\$256.00	\$460.80	Free
Total cost	\$680.96	\$972.80	\$808.96	\$1034.24	\$78.80

*prices as of January 2023

The savings from moving to R2 are immediately clear. Your storage costs are twenty-five percent lower than the next lowest-priced option, and your data transfer fees are reduced to nothing.

The other fee schedule you encounter with object storage was mentioned briefly above: operations performed on your data. Though looking at these fees closely is outside the scope of this piece, let's look at a comparison between the providers to see how R2 compares.

All providers divide data operations, sometimes called requests, into two tiers, usually referred to as Class A and Class B. Class A operations usually mutate the state of the data, and include actions such as `put` and `copy`, while Class B operations usually read the data, and include `get` and similar operations.

	AWS S3	Google Cloud	Azure Blob Storage	IBM Cloud	Cloudflare R2
Class A (per 10,000)	\$0.05	\$0.05	\$0.065	\$0.05	\$0.045
Class B (per 10,000)	\$0.004	\$0.004	\$0.005	\$0.004	\$0.0036

*prices as of January 2023

Though not as dramatic as with storage fees, R2 would be cost saving based on operational requests alone. When combined with the lower data storage fees and the lack of egress fees, moving to R2 could dramatically reduce the cost of your cloud hosting, especially for workloads that see large amounts of data transfer.

Cloudflare offers a generous free tier for most of our products, and R2 is no exception. With ten GB of storage, a million class A requests, and ten million class B requests per month, there's more than enough room for you to try it out and see how easy it can be to use R2—or other Cloudflare products—in your projects.

Building with Cloudflare offers you incredible speed, easy scaling, low prices, top-notch security, and a frictionless development experience that empowers you to make the choices that work best for your architecture, and use your data the ways that work best for you. Sign up for [R2](#) and start saving today.





© 2023 Cloudflare Inc. All rights reserved. The Cloudflare logo is a trademark of Cloudflare. All other company and product names may be trademarks of the respective companies with which they are associated.

1 888 99 FLARE | enterprise@cloudflare.com | www.cloudflare.com