

EXHIBIT B

http://aws.amazon.com/s3/faqs/ **AUG OCT DEC**
 450 captures **05**
 16 Oct 2008 - 4 Jan 2019 **2011 2012 2013**

Amazon Web Services

My Account / Console English

AWS Products & Solutions AWS Product Information **Developers** **Support**

- Amazon S3**
- Amazon S3 Overview
- FAQs**
- Pricing
- Amazon S3 SLA

- Developer Resources**
- Documentation
- Release Notes
- Sample Code & Libraries
- Developer Tools
- Articles & Tutorials
- Community Forum

Amazon Simple Storage Service FAQs

General

- What is Amazon S3?
- What can I do with Amazon S3?
- How can I get started using Amazon S3?
- What are the technical benefits of Amazon S3?
- What can developers do now that they could not before?
- Can I host my static website on Amazon S3?
- What kind of data can I store?
- How much data can I store?
- How can I delete large numbers of objects?
- How can I configure my objects to be deleted after a specific time period?
- What does Amazon do with my data in Amazon S3?
- Does Amazon store its own data in Amazon S3?
- How is Amazon S3 data organized?
- How do I interface with Amazon S3?
- How reliable is Amazon S3?
- What data consistency model does Amazon S3 employ?
- What happens if traffic from my application suddenly spikes?
- What is the BitTorrent™ protocol, and how do I use it with Amazon S3?
- Does Amazon S3 offer a Service Level Agreement (SLA)?

Regions

- Where is my data stored?
- How do I decide which Region to store my data in?
- I'm not in the US or Europe; can I use Amazon S3?

Billing

- How much does Amazon S3 cost?
- Why do prices vary depending on which Amazon S3 Region I choose?
- How will I be charged and billed for my use of Amazon S3?
- How am I charged for accessing Amazon S3 through the AWS Management Console?
- Do your prices include taxes?

Security

- How secure is my data?
- How can I control access to my data stored on Amazon S3?
- Does Amazon S3 support data access auditing?
- What options do I have for encrypting data stored on Amazon S3?
- How does Amazon protect SSE encryption keys?
- Can I comply with EU data privacy regulations using Amazon S3?
- Where can I find more information about security on AWS?

Data Protection

Why should I use Versioning?
How do I start using Versioning?
How does Versioning protect me from accidental deletion of my objects?
How can I ensure maximum protection of my preserved versions?
How am I charged for using Versioning?

Reduced Redundancy Storage (RRS)

What is RRS?
Why would I choose to use RRS?
What is the durability of Amazon S3 when using RRS?
How do I know if I lose an RRS object?
How do I specify that I want to store my data using RRS?
Are my RRS objects backed with the Amazon S3 Service Level Agreement?
How will my performance be impacted as a result of using RRS?
How am I charged for using RRS?

Find more answers in [Amazon S3 Technical FAQs](#)

[↑ Top](#)

General

Q: What is Amazon S3?

Amazon S3 is storage for the Internet. It's a simple storage service that offers software developers a highly-scalable, reliable, and low-latency data storage infrastructure at very low costs.

Q: What can I do with Amazon S3?

Amazon S3 provides a simple web services interface that you can use to store and retrieve any amount of data, at any time, from anywhere on the web. Using this web service, developers can easily build applications that make use of Internet storage. Since Amazon S3 is highly scalable and you only pay for what you use, developers can start small and grow their application as they wish, with no compromise on performance or reliability. It is designed to be highly flexible: Store any type and amount of data that you want; read the same piece of data a million times or only for emergency disaster recovery; build a simple FTP application, or a sophisticated web application such as the Amazon.com retail web site. Amazon S3 frees developers to focus on innovation, not figuring out how to store their data.

Q: How can I get started using Amazon S3?

To sign up for Amazon S3, click the "Sign up for This Web Service" button on the [Amazon S3](#) detail page. You must have an Amazon Web Services account to access this service; if you do not already have one, you will be prompted to create one when you begin the Amazon S3 sign-up process. After signing up, please refer to the Amazon S3 documentation and sample code in the [Resource Center](#) to begin using Amazon S3.

Q: What are the technical benefits of Amazon S3?

Amazon S3 was carefully engineered to meet the requirements for scalability, reliability, speed, low-cost, and simplicity that must be met for Amazon's internal developers. Amazon S3 passes these same benefits onto any external developer. More information about the Amazon S3 design requirements is available on the [Amazon S3](#) detail page.

Q: What can developers do now that they could not before?

Until now, a sophisticated and scalable data storage infrastructure like Amazon's has been beyond the reach of small developers. Amazon S3 enables any developer to leverage Amazon's own benefits of massive scale with no up-front investment or performance compromises. Developers are now free to innovate knowing that no matter how successful their businesses become, it will be inexpensive and simple to ensure their data is quickly accessible, always available, and secure.

Q: Can I host my static website on Amazon S3?

Yes, you can host your entire static website on Amazon S3. Customers commonly store images, video, and other content for their website on Amazon S3. You can also take advantage of root and index document support as well as custom HTML errors to host your entire website from your Amazon S3 bucket. To learn more about hosting your website on Amazon S3, please visit the [Amazon S3 Developer Guide](#).

Q: What kind of data can I store?

You can store virtually any kind of data in any format. Please refer to the [Amazon Web Services Licensing Agreement](#) for details.

Q: How much data can I store?

The total volume of data and number of objects you can store are unlimited. Individual Amazon S3 objects can range in size from 1 byte to 5 terabytes. The largest object that can be uploaded in a single PUT is 5 gigabytes. For objects

Q: How can I delete large numbers of objects?

You can use [Multi-Object Delete](#) to delete large numbers of objects from Amazon S3. This feature allows you to send multiple object keys in a single request to speed up your deletes. Amazon does not charge you for using Multi-Object Delete.

Q: How can I configure my objects to be deleted after a specific time period?

You can use the Object Expiration feature to remove objects from your buckets after a specified number of days. You can define the expiration rules for a set of objects in your bucket through the Lifecycle Configuration policy that you apply to the bucket. Each Object Expiration rule allows you to specify a prefix and an expiration period. The prefix field identifies the objects subject to the rule. To apply the rule to an individual object, specify the key name. To apply the rule to a set of objects, specify their common prefix (e.g. "logs/"). For expiration period, provide the number of days from creation date (i.e. age) after which you want your objects removed. You may create multiple rules for different prefixes. For example, you could create a rule that removes all objects with the prefix "logs/" 30 days from creation, and a separate rule that removes all objects with the prefix "backups/" 90 days from creation.

After an Object Expiration rule is added, the rule is applied to objects that already exist in the bucket as well as new objects added to the bucket. Once objects are past their expiration date, they are identified and queued for removal. You will not be billed for storage for objects on or after their expiration date, though objects may still be accessible while they are in queue before they are removed. As with standard delete requests, Amazon S3 doesn't charge you for removing objects using Object Expiration. For more information on using Object Expiration, please refer to the [Object Expiration](#) topic in the [Amazon S3 Developer Guide](#).

Q: What does Amazon do with my data in Amazon S3?

Amazon will store your data and track its associated usage for billing purposes. Amazon will not otherwise access your data for any purpose outside of the Amazon S3 offering, except when required to do so by law. Please refer to the [Amazon Web Services Licensing Agreement](#) for details.

Q: Does Amazon store its own data in Amazon S3?

Yes. Developers within Amazon use Amazon S3 for a wide variety of projects. Many of these projects use Amazon S3 as their authoritative data store, and rely on it for business-critical operations.

Q: How is Amazon S3 data organized?

Amazon S3 is a simple key-based object store. When you store data, you assign a unique object key that can later be used to retrieve the data. Keys can be any string, and can be constructed to mimic hierarchical attributes.

Q: How do I interface with Amazon S3?

Amazon S3 provides simple, standards-based REST and SOAP web services interfaces that are designed to work with any Internet-development toolkit. The operations are intentionally made simple to make it easy to add new distribution protocols and functional layers.

Q: How reliable is Amazon S3?

Amazon S3 gives any developer access to the same highly scalable, reliable, fast, inexpensive data storage infrastructure that Amazon uses to run its own global network of web sites. The service was designed for 99.99% availability, and carries a service level agreement providing service credits if a customer's availability falls below 99.9%.

Q: What data consistency model does Amazon S3 employ?

Amazon S3 buckets in the US West (Oregon), US West (Northern California), EU (Ireland), Asia Pacific (Singapore), Asia Pacific (Tokyo), and South America (Sao Paulo) Regions provide read-after-write consistency for PUTS of new objects and eventual consistency for overwrite PUTS and DELETES. Amazon S3 buckets in the US Standard Region provide eventual consistency.

Q: What happens if traffic from my application suddenly spikes?

Amazon S3 was designed from the ground up to handle traffic for any Internet application. Pay-as-you-go pricing and unlimited capacity ensures that your incremental costs don't change and that your service is not interrupted. Amazon S3's massive scale enables us to spread load evenly, so that no individual application is affected by traffic spikes.

Q: What is the BitTorrent™ protocol, and how do I use it with Amazon S3?

BitTorrent is an open source Internet distribution protocol. Amazon S3's bandwidth rates are inexpensive, but BitTorrent allows developers to further save on bandwidth costs for a popular piece of data by letting users download from Amazon and other users simultaneously. Any publicly available data in Amazon S3 can be downloaded via the BitTorrent protocol, in addition to the default client/server delivery mechanism. Simply add the ?torrent parameter at the end of your GET request in the REST API.

Q: Does Amazon S3 offer a Service Level Agreement (SLA)?

Yes. The Amazon S3 SLA provides for a service credit if a customer's monthly uptime percentage is below our service commitment in any billing cycle. More information can be found [here](#).

[↑ Top](#)

Q: Where is my data stored?

Amazon S3 offers storage in the US Standard, US West (Oregon), US West (Northern California), EU (Ireland), Asia Pacific (Singapore), Asia Pacific (Tokyo), South America (Sao Paulo), and AWS GovCloud (US) Regions. You specify a Region when you create your Amazon S3 bucket. Within that Region, your objects are redundantly stored on multiple devices across multiple facilities.

Q: How do I decide which Region to store my data in?

There are several factors to consider based on your specific application. You may want to store your data in a Region that...

- ...is near to your customers, your data centers, or your other AWS resources in order to reduce data access latencies.

- ...is remote from your other operations for geographic redundancy and disaster recovery purposes.

- ...enables you to address specific legal and regulatory requirements.

- ...allows you to reduce storage costs. You can choose a lower priced Region to save money. Please see the pricing section on the [S3 detail page](#).

Q: I'm not in the US or Europe; can I use Amazon S3?

Anyone can use Amazon S3. You just have to decide which Region you want Amazon S3 to store your data in.

[↑ Top](#)

Billing

Q: How much does Amazon S3 cost?

With Amazon S3, you pay only for what you use. There is no minimum fee. You can estimate your monthly bill using the [AWS Simple Monthly Calculator](#).

We charge less where our costs are less. Some prices vary across Amazon S3 Regions and are based on the location of your bucket. There is no Data Transfer charge for data transferred within an Amazon S3 Region via a COPY request. Data transferred via a COPY request between Regions is charged at regular rates. There is no Data Transfer charge for data transferred between Amazon EC2 and Amazon S3 within the same Region or for data transferred between the Amazon EC2 Northern Virginia Region and the Amazon S3 US Standard Region. Data transferred between Amazon EC2 and Amazon S3 across all other Regions (i.e. between the Amazon EC2 Northern California and Amazon S3 US Standard Regions) will be charged at Internet Data Transfer rates on both sides of the transfer.

For S3 pricing information, please visit the pricing section on the [S3 detail page](#).

Q: Why do prices vary depending on which Amazon S3 Region I choose?

We charge less where our costs are less. For example, our costs are lower in the US Standard Region than in the US West (Northern California) Region.

Q: How will I be charged and billed for my use of Amazon S3?

There are no set-up fees or commitments to begin using the service. At the end of the month, your credit card will automatically be charged for that month's usage. You can view your charges for the current billing period at any time on the Amazon Web Services web site, by logging into your Amazon Web Services account, and clicking "Account Activity" under "Your Web Services Account".

With the [AWS Free Usage Tier*](#), you can get started with Amazon S3 for free in all regions except the AWS GovCloud Region. Upon sign-up, new AWS customers receive 5 GB of Amazon S3 standard storage, 20,000 Get Requests, 2,000 Put Requests, 15GB of data transfer in, and 15GB of data transfer out each month for one year.

Amazon S3 charges you for the following types of usage:

Note: The calculations below assume there is no AWS Free Tier in place.

Storage Used:

Amazon S3 storage pricing is summarized on the [Amazon S3 Pricing Chart](#).

The volume of storage billed in a month is based on the average storage used throughout the month. This includes all object data and metadata stored in buckets that you created under your AWS account. We measure your storage usage in "TimedStorage-ByteHrs," which are added up at the end of the month to generate your monthly charges.

Storage Example:

Assume you store 100GB (107,374,182,400 bytes) of standard Amazon S3 storage data in your bucket for 15 days in March, and 100TB (109,951,162,777,600 bytes) of standard Amazon S3 storage data for the final 16 days in March.

At the end of March, you would have the following usage in Byte-Hours:

Total Byte-Hour usage

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts



Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research



With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips



Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.