

Hosting Static Websites on AWS

Prescriptive Guidance

May 2017



© 2017, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Notices

This document is provided for informational purposes only. It represents AWS's current product offerings and practices as of the date of issue of this document, which are subject to change without notice. Customers are responsible for making their own independent assessment of the information in this document and any use of AWS's products or services, each of which is provided "as is" without warranty of any kind, whether express or implied. This document does not create any warranties, representations, contractual commitments, conditions or assurances from AWS, its affiliates, suppliers or licensors. The responsibilities and liabilities of AWS to its customers are controlled by AWS agreements, and this document is not part of, nor does it modify, any agreement between AWS and its customers.

Contents

| | |
|---|----|
| Introduction | 1 |
| Static Versus Dynamic Website | 1 |
| Core Architecture | 2 |
| Moving to an AWS Architecture | 4 |
| Use Amazon S3 Website Hosting to Host Without a Single Web Server | 5 |
| Scalability and Availability | 6 |
| Configuration Basics | 7 |
| Evolving the Architecture with Amazon CloudFront | 12 |
| Factors Contributing to Page Load Latency | 12 |
| Speeding Up Your Amazon S3-Based Website Using Amazon CloudFront | 13 |
| Amazon CloudFront Reports | 16 |
| Estimating and Tracking AWS Spend | 19 |
| Estimating AWS Spend | 20 |
| Tracking AWS Spend | 21 |
| Integration with Your Continuous Deployment Process | 22 |
| Access Logs | 23 |
| Analyzing Logs | 23 |
| Archiving and Purging Logs | 24 |
| Securing Administrative Access to Your Website Resources | 25 |
| Managing Administrator Privileges | 26 |
| Auditing API Calls Made in Your AWS Account | 28 |
| Controlling How Long Amazon S3 Content Is Cached by Amazon CloudFront | 28 |
| Conclusion | 33 |
| Contributors | 33 |

Abstract

This whitepaper covers comprehensive architectural guidance for developing, deploying, and managing static websites on Amazon Web Services (AWS) while keeping operational simplicity and business requirements in mind. We also recommend an approach that provides these business benefits: 1) insignificant cost of operation, 2) little or no management required, and 3) a highly scalable, resilient, and reliable website.

This whitepaper first reviews how static websites are hosted in traditional hosting environments. Then we explore a simpler and more cost-efficient approach using Amazon Simple Storage Service (Amazon S3). Finally, we show you how you can enhance the AWS architecture to layer on functionality and improve quality of service by using Amazon CloudFront.

Introduction

As enterprises become more digital operations, we see a proliferation of many kinds of websites. They span a wide spectrum, from mission-critical e-commerce sites to departmental apps, and from business-to-business (B2B) portals to marketing sites. Many factors such as business value, mission criticality, service level agreements (SLAs), quality of service, and information security will drive the choice of architecture and technology stack.

The simplest form of website architecture is the *static website*, where users are served static content (HTML, images, video, JavaScript, style sheets, etc.). Some examples include brand micro sites, marketing websites, and intranet information pages. Static websites are straightforward in one sense, but they still can have demanding requirements in terms of scalability, availability, and service-level guarantees. For example, a marketing site for a consumer brand might need to be prepared for an unpredictable onslaught of visitors when a new product is launched.

Static Versus Dynamic Website

What is the difference between a “static” website and a “dynamic” website? The word *static* refers to the fact that a static website delivers content in the same format in which it is stored. No server-side code execution is required. For example, if a static website consists of HTML documents displaying images, it will deliver the HTML and images as-is to the browser, without altering the contents of the files.

Static websites can be delivered to web browsers on desktops, tablets, or mobile devices. They usually consist of a mix of HTML documents, images, videos, CSS style sheets, and JavaScript files. Static doesn't have to mean boring—static sites can provide client-side interactivity as well. Using HTML5 and client-side JavaScript technologies such as jQuery, AngularJS, React, and Backbone, you can deliver rich user experiences that are engaging and interactive.

Some examples of static sites include:

- Marketing websites
- Product landing pages

- Micro-sites that display the same content to all users
- Team homepages
- Small, self-contained websites for a department, project, or team
- A website that lists available assets (e.g., image files, video files, and press releases) allows the user to download the files as-is
- Proofs-of-concept used in the early stages of web development to test user experience flows and gather feedback

By contrast, *dynamic* websites can display dynamic or personalized content. They usually interact with data sources and web services, and require code development expertise to create and maintain. For example, a sports news site can display information based on the visitor's preferences, and use server-side code to display updated sport scores. Other examples of dynamic sites are e-commerce shopping sites, news portals, social networking sites, finance sites, and most other websites that display ever-changing information.

Static websites load faster than dynamic ones, since content is delivered as-is and can be cached by a content delivery network (CDN), and the web server doesn't need to perform any application logic or database queries. They're also relatively inexpensive to develop and host. However, maintaining large static websites can be cumbersome without the aid of automated tools, and static websites can't deliver personalized information.

Static websites are most suitable when the content is infrequently updated. After the content evolves in complexity or needs to be frequently updated, personalized, or dynamically generated, it's time to consider a dynamic website architecture.

Core Architecture

In a traditional (non-AWS) architecture, web servers serve up static content. Typically, content is managed using a content management system (CMS), and multiple static sites will be hosted on the same infrastructure. The content is stored on local disks, or on a file share on network-accessible storage. A sample file system structure might look like the structure that follows.

```
├─ css/
│  └─ main.css
│    └─ navigation.css
├─ images/
│  └─ banner.jpg
│    └─ logo.jpg
├─ index.html
├─ scripts/
│  └─ script1.js
│    └─ script2.js
├─ section1.html
└─ section2.html
```

A network firewall protects against unauthorized access. It's common to deploy multiple web servers behind a load balancer for high availability (HA) and scalability. Since pages are static, the web servers don't need to maintain any state or session information and the load balancer doesn't need to implement session affinity ("sticky sessions"). In a traditional (non-AWS) hosting environment, a basic architecture looks like the following diagram:

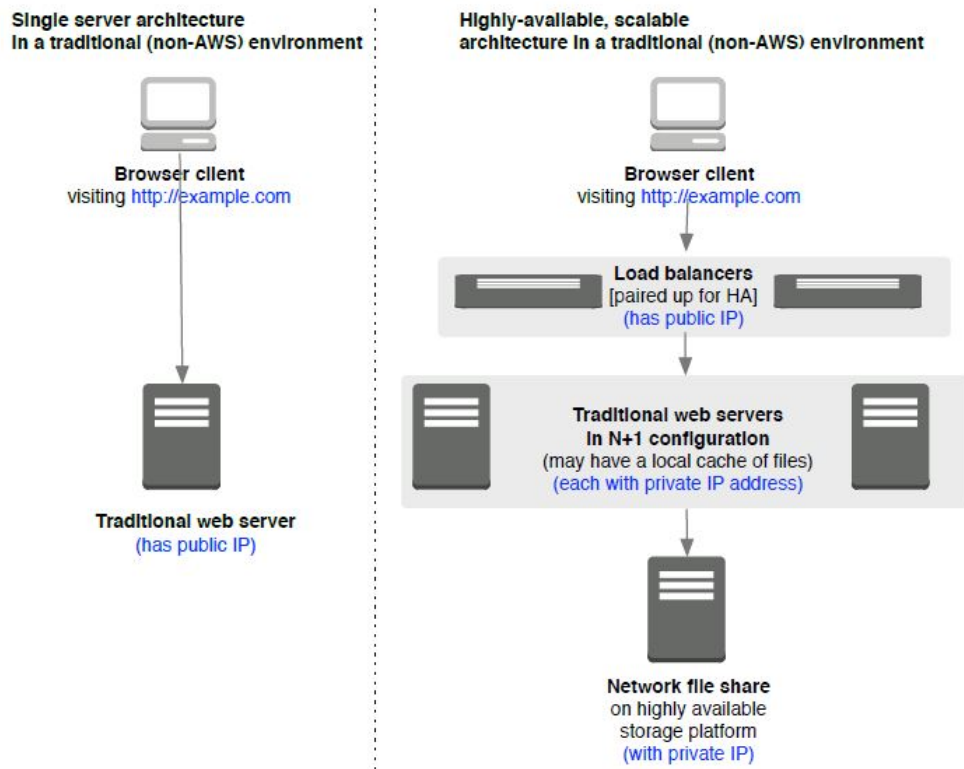


Figure 1: Basic architecture of a traditional hosting environment

Moving to an AWS Architecture

To translate a traditional hosting environment to an AWS architecture you could use a so-called “lift-and-shift” approach. A lift and shift is a like-for-like substitution of AWS services instead of the traditional environment.

You can run Linux or Windows web servers on Amazon Elastic Compute Cloud (Amazon EC2). You can use Elastic Load Balancing (ELB) to load balance and distribute the web traffic. The static content can be stored on Amazon Elastic Block Store (Amazon EBS) or Amazon Elastic File System (Amazon EFS). You can deploy your Amazon EC2 instances in an Amazon Virtual Private Cloud (Amazon VPC), which is your isolated and private virtual network in the AWS Cloud. This gives you full control over the network topology, firewall configuration, and routing rules. The web servers can be spread across multiple Availability Zones for high availability, even if an entire data center were to be down. You can use Auto Scaling to automatically add servers during high traffic periods and then scale back when traffic quiets down.

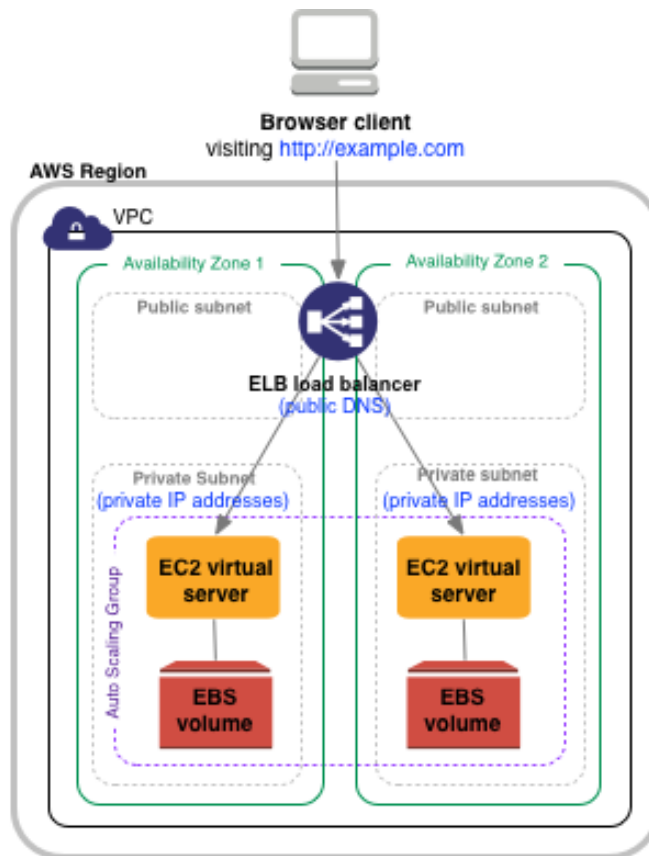


Figure 2: AWS architecture for a “Lift and Shift”

Using this AWS architecture you gain the security, scalability, cost, and agility benefits of running in AWS. This architecture benefits from AWS world-class infrastructure and security operations. By using Auto Scaling, the website is ready for traffic spikes, so you are prepared for product launches and viral websites. With AWS, you only pay for what you use, and there's no need to over-provision for peak capacity. In addition, you gain from increased agility because AWS services are available on demand. (Compare this to the traditional process in which provisioning servers, storage, or networking can take weeks.) You don't have to manage infrastructure, so this frees up time and resources to create business-differentiating value. In many ways, this initial AWS architecture is a major win over the non-AWS one.

However, you can do better.

AWS challenges traditional IT assumptions and enables new “cloud-native” architectures. You can architect a modern static website without needing a single web server.

Use Amazon S3 Website Hosting to Host Without a Single Web Server

Amazon Simple Storage Service (Amazon S3) can be used to host static websites without a need for a web server. The website will be highly performant and scalable at a [fraction of the cost of a traditional web server](#).¹ Amazon S3 is storage for the cloud, providing you with secure, durable, highly scalable object storage. It's easy to use: A simple web services interface allows you to store and retrieve any amount of data from anywhere on the web.²

You can start by creating an Amazon S3 bucket, enabling the Amazon S3 website hosting feature, and configuring access permissions for the bucket. After you upload files, Amazon S3 takes care of serving your content to your visitors.

Amazon S3 provides HTTP web-serving capabilities, and the content can be viewed by any browser. You also need to configure Amazon Route 53, a managed Domain Name System (DNS) service, to point your domain (here we're using *http://example.com*) to your Amazon S3 bucket.

This core architecture is as simple as it sounds. Figure 3 illustrates what it looks like.

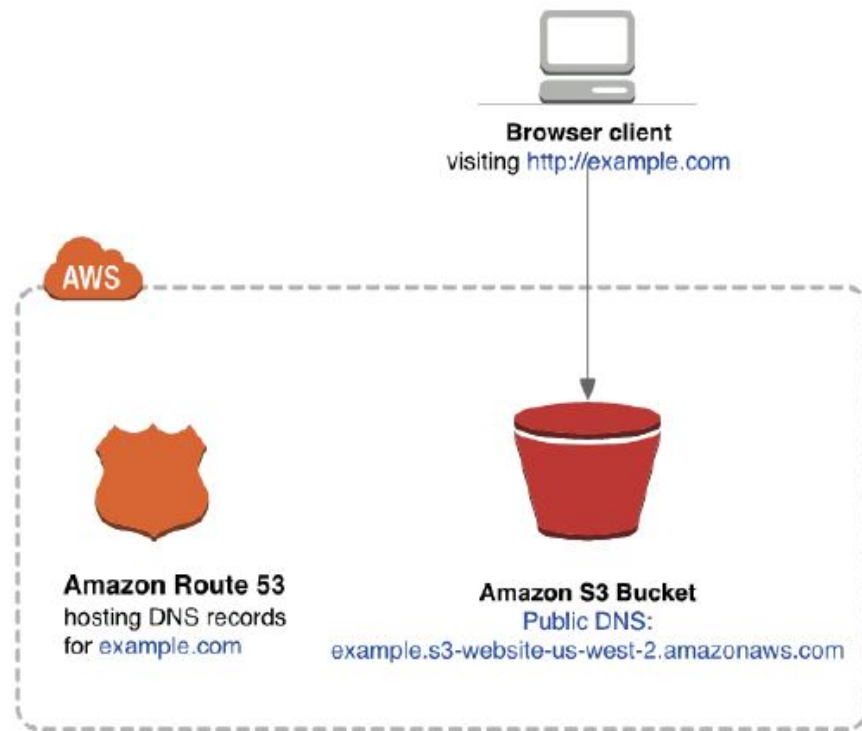


Figure 3: Amazon S3 website hosting

There are no Windows or Linux servers to manage! No need to provision machines, install operating systems, or fine-tune web server configurations. There's also no need to manage storage infrastructure (e.g., SAN, NAS) because Amazon S3 provides practically limitless cloud-based storage. Fewer moving parts means fewer troubleshooting headaches.

Scalability and Availability

Amazon S3 is inherently scalable. For popular websites, the Amazon S3 architecture will scale seamlessly to serve thousands of HTTP requests per second without any changes to the architecture.

In addition, by hosting with Amazon S3, the website is inherently highly available. Amazon S3 is designed for 99.99% availability, and carries a [service level agreement \(SLA\)](#) of 99.9% availability.³ Amazon S3 gives you access to the same highly scalable, reliable, fast, and inexpensive infrastructure that Amazon

uses to run its own global network of websites. As soon as you upload files to Amazon S3, Amazon S3 automatically replicates your content across multiple data centers. Even if an entire AWS data center were to be impaired, your static website would still be running and available to your end users.

Compare this with traditional non-AWS costs for implementing “active-active” hosting for important projects. Active-active, or deploying web servers in two distinct data centers, is prohibitive in terms of server costs and engineering time. As a result, traditional websites are usually hosted in a single data center, because most projects can’t justify the cost of “active-active” hosting.

Configuration Basics

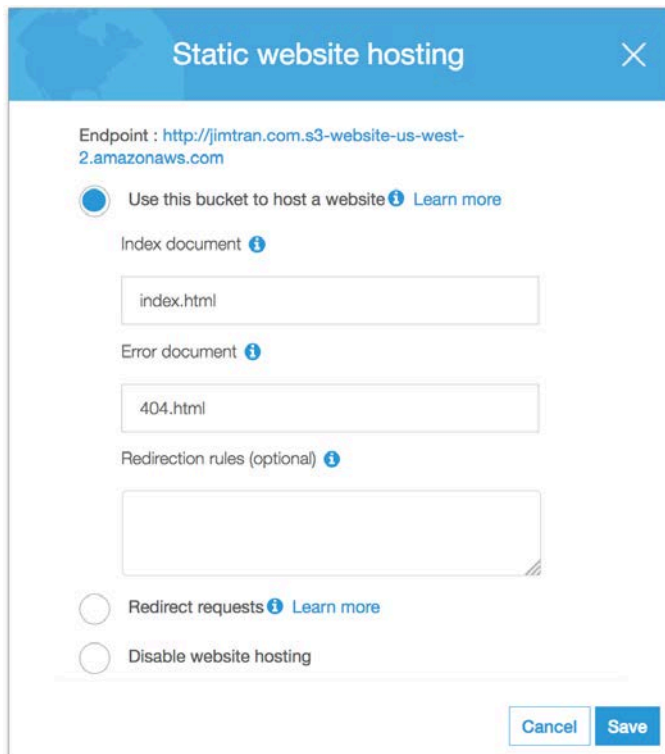
Getting started is easy. First, open the AWS Management Console, where you will create an Amazon S3 bucket using the Amazon S3 console. When you first create an S3 bucket, you select the AWS Region in which the files will be [geographically stored](#).⁴ Select a Region based on its proximity to your visitors, proximity to your corporate data centers, and/or your regulatory or compliance requirements (e.g., some countries have restrictive data residency regulations). Since Amazon S3 supports virtual-host style access to your Amazon S3 buckets, the bucket name must comply with [DNS naming conventions](#).⁵ If you plan to use your own custom domain/subdomain, such as *example.com* or *www.example.com*, your bucket name must be the same as your domain/subdomain. For example, a website available at *http://www.example.com* needs to be in a bucket named *www.example.com*. Each AWS account can have a maximum of 1000 buckets.

After your Amazon S3 bucket is created, toggle on the [static website hosting feature](#) for the bucket.⁶ This will generate an Amazon S3 website endpoint. You will be able to access your Amazon S3-hosted website at the following URL:

```
http://<bucket-name>.s3-website-<AWS-region>.amazonaws.com
```

For small, non-public websites, the Amazon S3 website endpoint is probably adequate. You can also use internal DNS to point to this endpoint. For a public-facing website, we recommend using a custom domain name instead of the provided Amazon S3 website endpoint. This way users can see user-friendly URLs in their browsers. As mentioned earlier, if you plan to use a custom

domain name, your bucket name must match the domain name. For custom root domains (such as *example.com*), only Amazon Route 53 can configure a DNS record to point to the Amazon S3-hosted website. For non-root subdomains (such as *www.example.com*), any DNS service (including Amazon Route 53) will be able to create a CNAME entry to the subdomain. See the Amazon S3 documentation for more [details on how to associate domain names with your website](#).⁷



The screenshot shows the 'Static website hosting' configuration window in the Amazon S3 console. At the top, the endpoint is displayed as `http://jimtran.com.s3-website-us-west-2.amazonaws.com`. Below this, there are three main sections: 1) A radio button selected for 'Use this bucket to host a website', with a 'Learn more' link. 2) An 'Index document' field containing 'index.html'. 3) An 'Error document' field containing '404.html'. Below these is a 'Redirection rules (optional)' text area. At the bottom, there are two radio buttons: 'Redirect requests' (unselected) and 'Disable website hosting' (unselected), both with 'Learn more' links. 'Cancel' and 'Save' buttons are located at the bottom right of the window.

Figure 4: Configuring static website hosting using the Amazon S3 console

The Amazon S3 website hosting configuration screen in the Amazon S3 console presents additional options to configure. Some of the key options are as follows:

- You can configure a default page that users will see if they visit the domain name directly (without specifying a specific page).⁸ You can also specify a custom “404 - Page Not Found” error page if the user stumbles onto a non-existent page.
- By default, logging is disabled. You can enable it to give you access to the raw web access logs.⁹

- Add tags to your Amazon S3 bucket. These tags will be useful later on when you want to analyze your AWS spend by project.¹⁰

A quick note about Amazon S3 object naming: In Amazon S3, a bucket is a flat container of objects. It doesn't provide a hierarchical organization the way the file system on your computer does. However, there is a straightforward mapping between a file system's folders/files to Amazon S3 objects. The example that follows shows how folders/files are mapped to Amazon S3 objects. Most third-party tools, as well as the AWS Management Console and AWS Command Line Interface (AWS CLI), will [handle this mapping transparently](#) for you. However, it's good to be aware of this fact.¹¹ It's also a good habit to use lower-case characters for file and folder names. Case is important and simply easier to manage when all characters are lowercase.

| <i>Hierarchical file system</i> | <i>S3 object names in the S3 bucket</i> |
|---------------------------------|---|
| └─ css/ | css/ |
| └─ main.css | css/main.css |
| └─ navigation.css | css/navigation.com |
| └─ images/ | images/ |
| └─ banner.jpg | images/banner.jpg |
| └─ logo.jpg | images/logo.jpg |
| └─ index.html | index.html |
| └─ scripts/ | scripts/ |
| └─ script1.js | scripts/script1.js |
| └─ script2.js | scripts/script2.js |
| └─ section1.html | section1.html |
| └─ section2.html | section2.html |

Uploading Content

On AWS, uploading content is straightforward. Continue to design your static website using your website authoring tool of choice. Most web design and authoring tools can save the static content on your local hard drive. Then simply upload the HTML, images, JavaScript files, CSS files, and other static assets into your Amazon S3 bucket. The deployment process consists of a single step: copy any new or modified files to the Amazon S3 bucket. You can use the AWS API, SDKs, or CLI to script this step for a fully automated deployment.

You can upload files using the AWS Management Console. You can also use AWS partner offerings such as CloudBerry, S3 Bucket Explorer, S3 Fox, and

other visual management tools. The easiest way, however, is to use the [AWS CLI](#).¹² The `s3 sync` command will recursively upload files and [synchronize your Amazon S3 bucket](#) with your local folder.¹³

```
aws s3 sync $LOCAL_FOLDER s3://$S3_BUCKET_NAME/ --delete
```

Making Your Content Publicly Accessible

For your visitors to access content at the Amazon S3 website endpoint, the Amazon S3 objects need to have the appropriate permissions. Amazon S3 enforces a security-by-default policy. New objects in a new bucket are private by default. For example, you'll see an Access Denied error when trying to view a newly uploaded file using your web browser. To fix this, configure the content as publicly accessible. It's possible to set object-level permissions for every individual object, but that quickly becomes tedious. It's far more convenient to define an Amazon S3 bucket-wide policy. The following sample Amazon S3 bucket policy enables everyone to view all objects in a bucket:

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Sid": "PublicReadGetObject",
    "Effect": "Allow",
    "Principal": "*",
    "Action": ["s3:GetObject"],
    "Resource": ["arn:aws:s3:::S3_BUCKET_NAME_GOES_HERE/*"]
  }]
}
```

This policy defines who can view the contents of your Amazon S3 bucket. In the “Managing Administrative Access to Your AWS Resources” section of this whitepaper, there’s a description of the use of AWS Identity and Access Management (IAM) policies to manage permissions for your team members.

Together, Amazon S3 bucket policies and IAM policies give you fine-grained control over who can manage and view your website.

Low Costs Encourage Experimentation

So how much will this website cost to run? Amazon S3 costs are storage plus bandwidth. The actual costs will depend on your asset file sizes, and your site's popularity (the number of visitors making browser requests). There's no minimum charge and no setup costs.

When you use Amazon S3, you pay for what you use. You're only [charged for the actual Amazon S3 storage required to store the site assets](#).¹⁴ These assets include HTML files, images, JavaScript files, CSS files, videos, audio files, and any other downloadable files. Your bandwidth charges will depend on the actual site traffic.¹⁵ Small websites with few visitors will have minimal hosting costs. Popular websites that serve up large videos and images will incur higher bandwidth charges. The [Estimating and Tracking AWS Spend](#) section of this whitepaper describes how you can estimate and track your costs.

With Amazon S3, experimenting with new ideas is easy and cheap. If a website idea turns out to be a dud, the costs will be minimal. This should encourage you to experiment more frequently. As they say, "fail fast, fail often!" This is great for micro-sites – publish many independent micro-sites at once, run A/B tests, and keep only the winners.

You've learned the basics of hosting a static website in Amazon S3. The architecture is refreshingly simple: no servers—just a single Amazon S3 bucket. It's ready to meet the real-world demands of a high-traffic website. It's highly available, resilient, scalable, and cost-efficient.

Now you'll learn how to iterate on this core architecture and make it faster, improve its performance for international visitors, and lower its cost even further.

Evolving the Architecture with Amazon CloudFront

When it comes to websites, speed matters. Web visitors enjoy and expect a fast browsing experience. Even slight page-loading delays can hurt business. In today's global economy, your site needs to be responsive and deliver page loads with low latency. Major search engines penalize slow websites by burying their search results. The Amazon CloudFront content delivery web service integrates with other AWS products to give you an easy way to distribute content to users with low latency, high data transfer speeds, and no minimum usage commitments.

Factors Contributing to Page Load Latency

To explore factors that contribute to latency, we use the example of a user in Singapore visiting a web page hosted from an Amazon S3 bucket in the US West (Oregon) Region in the United States. From the moment the user visits a web page to the moment it shows up in the browser, several factors contribute to latency:

- **FACTOR (1)** Time it takes for the browser (Singapore) to request the web page from Amazon S3 (US West (Oregon) Region).
- **FACTOR (2)** Time it takes for Amazon S3 to retrieve the page contents and serve up the page.
- **FACTOR (3)** Time it takes for the page contents (US West (Oregon) Region) to be delivered across the Internet to the browser (Singapore).
- **FACTOR (4)** Time it takes for the browser to parse and display the web page.

This latency is illustrated in Figure 5.

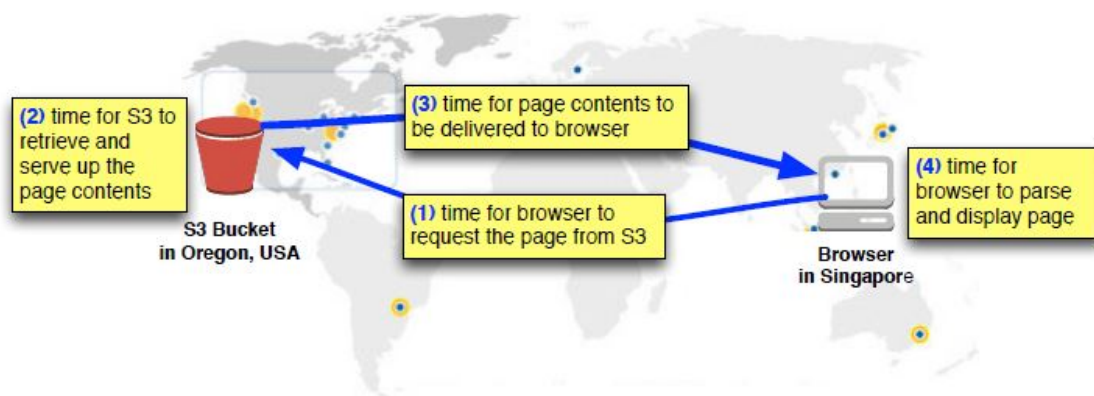


Figure 5: Factors affecting page load latency

The good news is that AWS has already taken care of **FACTOR (2)** by optimizing Amazon S3 to serve up content as quickly as possible. You can improve **FACTOR (4)** by optimizing the actual page content (e.g., minifying CSS and JavaScript, using efficient image and video formats). However, page-loading studies consistently show that most latency is due to **FACTOR (1)** and **FACTOR (3)**. Most of the delay in accessing pages over the Internet is due to the round-trip delay associated with establishing TCP connections (the infamous three-way TCP handshake) and the time it takes for TCP packets to be delivered across long Internet distances.

To sum up: serve content as close to your users as possible. In our example, American users will experience relatively fast page load times, whereas Singaporean users will experience slower page loads. Ideally, for the Singaporean users, you would want to serve up content as close to Singapore as possible.

Speeding Up Your Amazon S3-Based Website Using Amazon CloudFront

To reduce latency and provide a better user experience, use Amazon CloudFront. Amazon CloudFront is a CDN that uses a global network of edge locations for content delivery. A side benefit is that Amazon CloudFront also provides reports to help you understand how users are using your website.

As a CDN, Amazon CloudFront can distribute content with low latency and high data transfer rates.¹⁶ As of March 2017, there are over [70 CloudFront edge locations](#) all around the world.¹⁷ Therefore, no matter where a visitor lives in the world, there is an Amazon CloudFront edge location that is relatively close (from an Internet latency perspective).

The Amazon CloudFront edge locations will cache content from an origin server and deliver that cached content to the user. When creating an [Amazon CloudFront distribution](#), specify your Amazon S3 bucket as the origin server.¹⁸ The Amazon CloudFront distribution itself will have a DNS. You can refer to it using a CNAME if you have a custom domain name. To point the A record of a root domain to an Amazon CloudFront distribution, you can use Amazon Route 53 alias records, as illustrated in Figure 6.

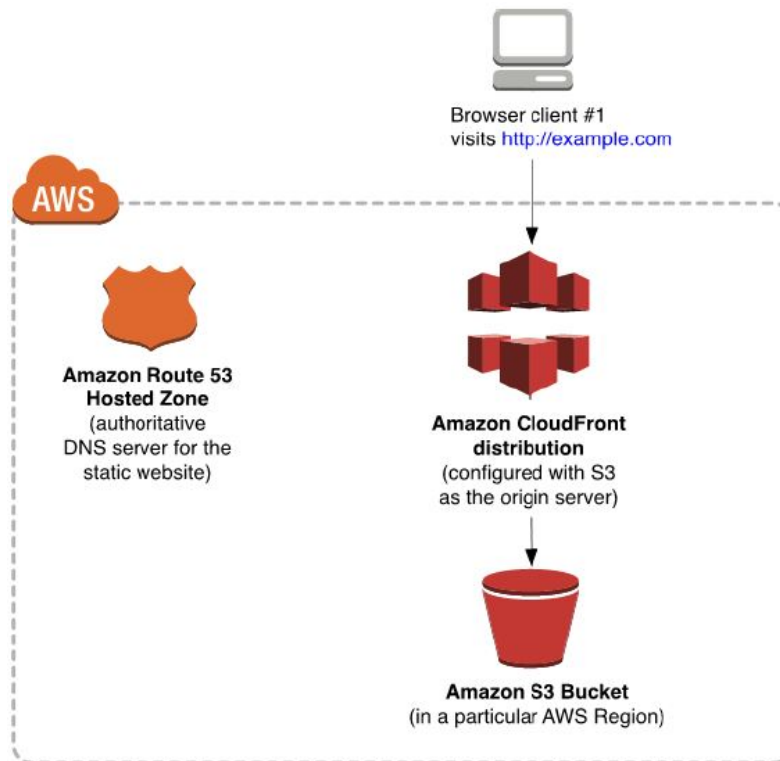


Figure 6: Using Amazon Route 53 alias records with an Amazon CloudFront distribution

To understand how CloudFront accelerates websites, it's helpful to understand what is happening under the covers. When an end user requests a web page using that domain name, CloudFront determines the best edge location to serve the content. If an edge location doesn't yet have a cached copy of the requested

content, CloudFront will pull a copy from the Amazon S3 origin server and hold it at the edge location to fulfill future requests. Subsequent users requesting the same content from that edge location will experience faster page loads because that content is already cached. Figure 7 shows the flow in detail. In the “Controlling How Long Amazon S3 Content Is Cached by Amazon CloudFront” section later in this whitepaper, we cover advanced strategies.

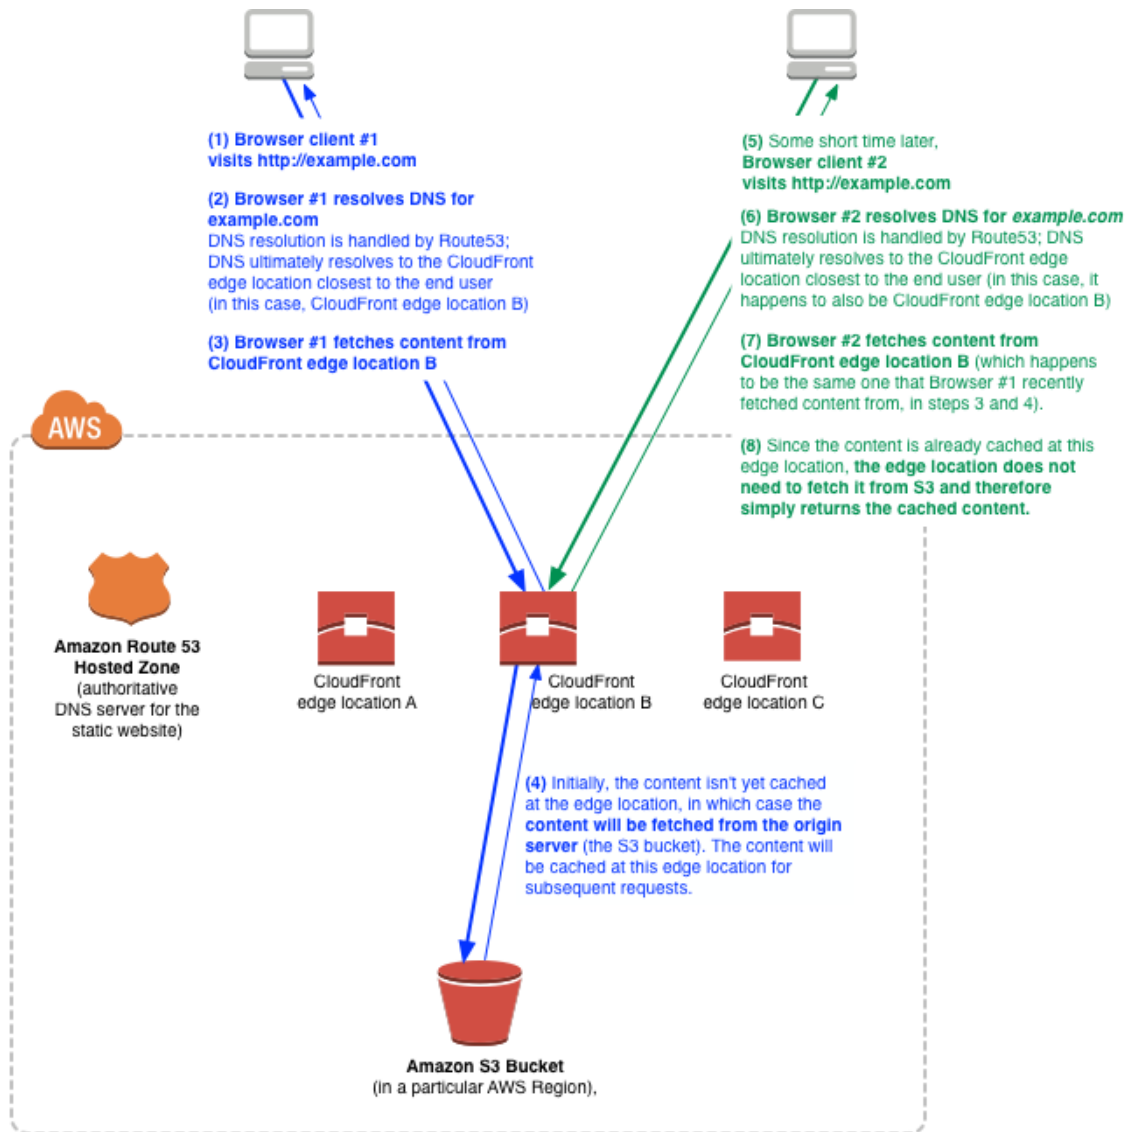



Figure 7: How Amazon CloudFront caches content

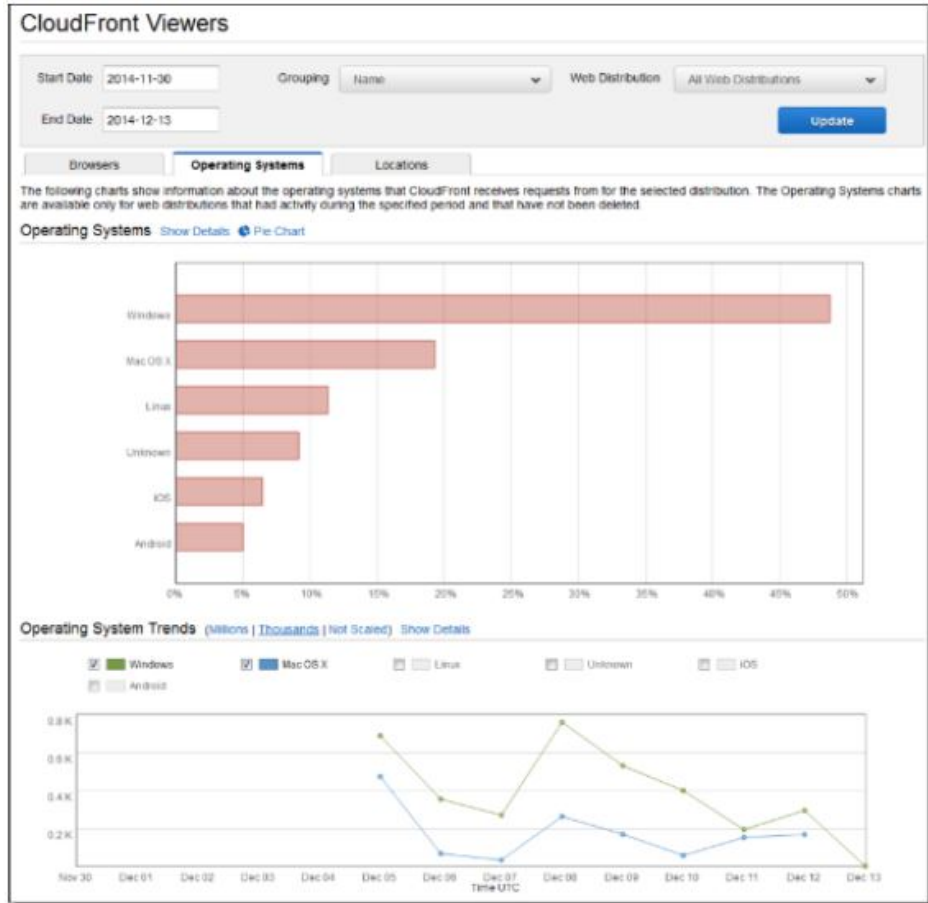
Amazon CloudFront Reports

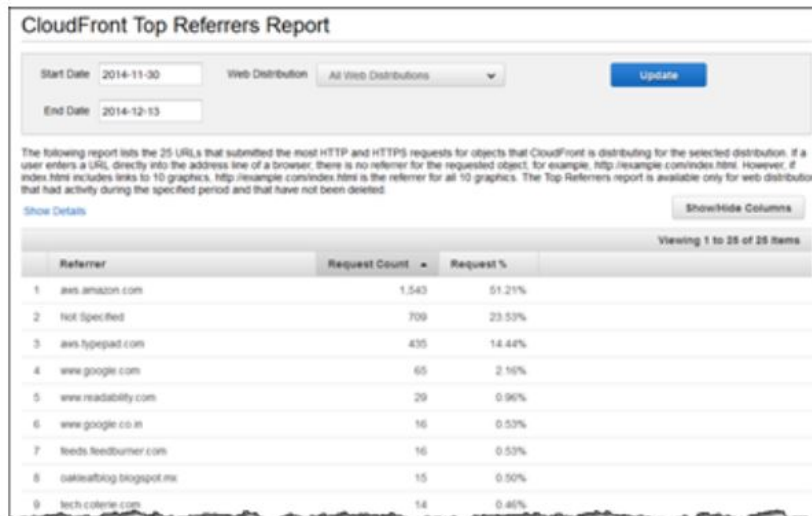
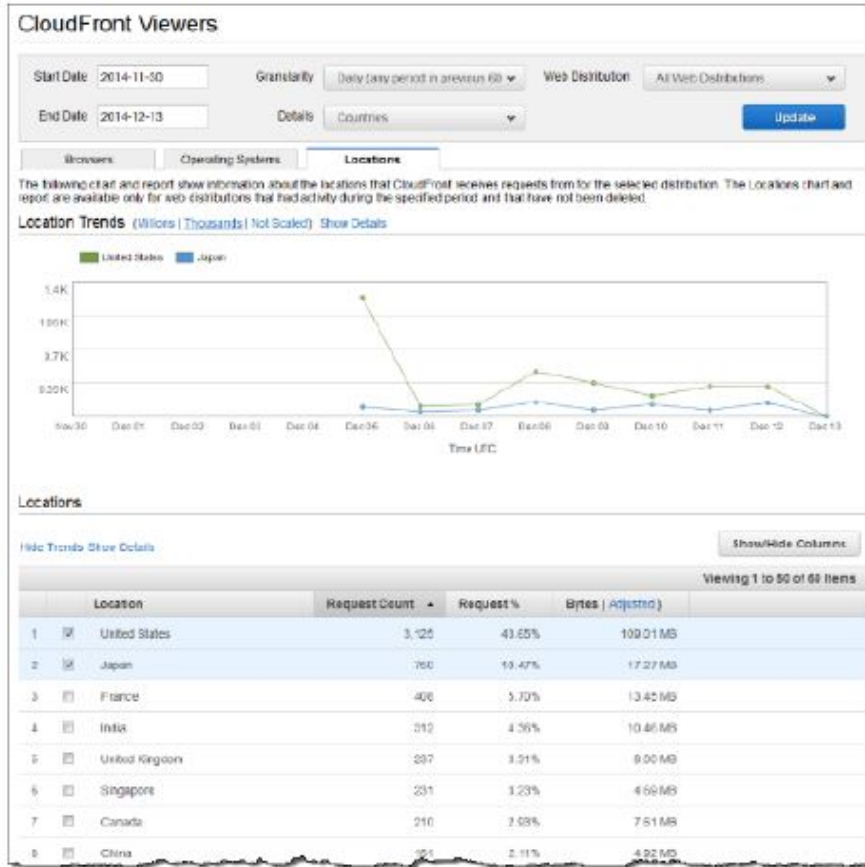
Amazon CloudFront provides you with a rich set of reports. The reports help you understand how your visitors are using your website. You can track trends and keep a close eye on your website's performance and effectiveness. The reports are great for gaining insights by answering questions such as the following:

- What is the overall health of my website?
- How many visitors are viewing my website?
- Which browsers, devices, and operating systems are they using?
- Which countries are they coming from?
- Which websites are the top referrers to my site?
- What assets are the most popular ones on my site?
- How often is CloudFront caching taking place?

Amazon CloudFront reports can be used alongside other online analytics tools, and we encourage the use of multiple reporting tools. Note that some analytics tools might require you to embed client-side JavaScript in your HTML pages. Amazon CloudFront reporting does not require any changes to your web pages. Figure 8 shows a sampling of types of Amazon CloudFront reports.










Figure 8: Various types of Amazon CloudFront reports

For any website serving a global user base, we highly recommend using Amazon CloudFront.¹⁹ Amazon S3 is solid on its own, and adding Amazon CloudFront will help you deliver a higher-quality experience to your users.

Let's now turn our attention to advanced topics. We cover the following in the remaining sections:

- Estimating and tracking your AWS spend
- Integration with your continuous deployment process
- Analyzing and archiving web access logs
- Securing administrative access to your AWS resources
- Auditing administrative actions performed in your AWS account

Estimating and Tracking AWS Spend

With AWS, there is no upper limit to the amount of Amazon S3 storage or network bandwidth you can consume. You pay as you go and only pay for actual usage. Because you're not using web servers in this architecture, you have no licensing costs or concern for server scalability or utilization.

Estimating AWS Spend

To estimate your monthly costs, you can use the [AWS Simple Monthly Calculator](#).²⁰ Pricing sheets for Amazon Route 53, Amazon S3, and Amazon CloudFront are all available online, with costs clearly itemized.²¹ The following table provides a breakdown of the key cost components for a website served out of Amazon CloudFront and an Amazon S3 bucket in the US West Oregon Region (us-west-2):

| Component cost | Monthly cost | Notes |
|--|---|---|
| Amazon Route 53 hosted zone | \$0.50/month | |
| DNS queries to the Amazon Route 53 hosted zone | \$0.400 per million queries for the first 1 Billion queries | Most static websites will be well below 1 billion queries / month. |
| Amazon S3 storage costs | \$0.023 per GB/month, for the first TB of data stored | Most static websites will be well below a TB of storage. Therefore, unless your website contains thousands of large video files, your Amazon S3 monthly storage costs will \$0.023 for each GB in your Amazon S3 bucket each month. |
| GET requests to Amazon S3 | Every 10,000 GET is \$0.004 | There is a cost to using Amazon CloudFront, but because it also cuts down on the number of Amazon S3 requests, it saves on S3 request costs. ²² For popular websites, using CloudFront can reduce overall AWS costs. |
| HTTP requests to CloudFront | Every 10,000 HTTP requests is \$0.0075 | Popular websites will cost more to run, of course. Note that if a web page references other assets such as images, JavaScript scripts, and CSS files this will result in multiple HTTP requests to CloudFront/S3. |
| Data transfer from Amazon CloudFront to end users | \$0.085/month for the first 10 TB | Most static websites will be well below 10 TB of outbound data transfer. Popular websites will cost more to run, of course. |
| Transferring data from Amazon S3 to Amazon CloudFront | No charge | |
| API costs to deploy the website | Negligible | |

** As of March 2017 pricing. Also keep in mind that AWS pricing is Region-specific.

Plug your own assumptions into the AWS Simple Monthly Calculator. You might be pleasantly surprised. Most static websites will cost less than *a dollar a month* to host using Amazon S3 with Amazon CloudFront and Amazon Route 53.

Tracking AWS Spend

It's a good idea to track your AWS spend on a regular basis. The [AWS Cost Explorer](#) can help you track cost trends by service type.²³ It's integrated in the AWS Billing console and runs in your browser. The Monthly Spend by Service chart allows you to see where your money is going. The Daily Spend report helps you track your spending as it happens. If you configured tags for your Amazon S3 bucket, you can filter your reports against specific tags for [cost allocation purposes](#).²⁴

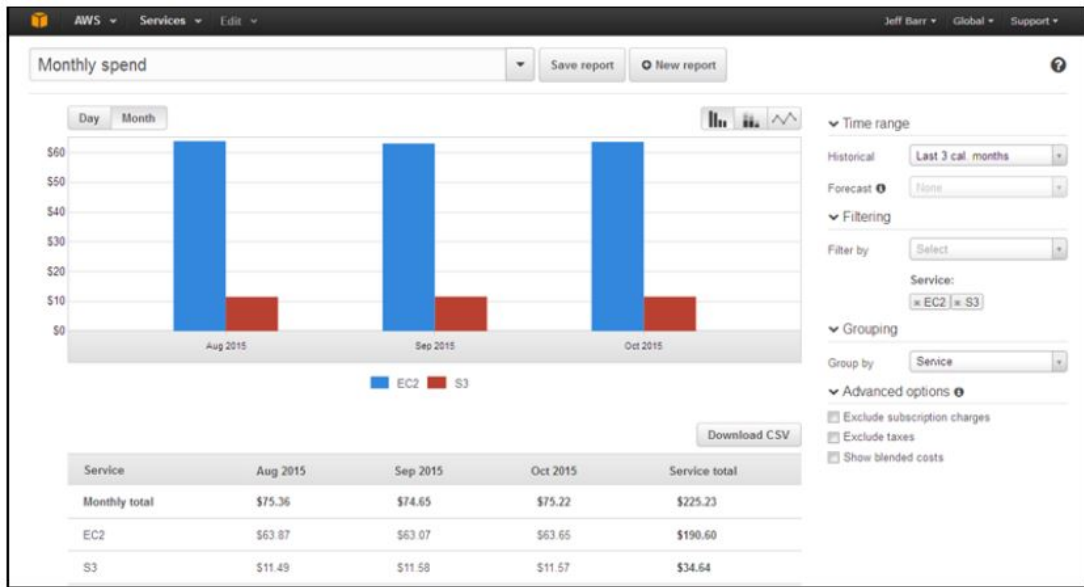


Figure 9: Sample from AWS Cost Explorer

Generally speaking, your AWS spend should be in the range of a few dollars at most, and the vast majority of static websites will not require AWS spend optimization.²⁵

Are you curious how your costs changed since your last marketing blitz? Simply define a custom filter based on time and services. After you configure it just the way you want, remember to [save the report for future use](#).²⁶

Integration with Your Continuous Deployment Process

Fresh content helps to keep your website attractive. A carefully thought out deployment process will make it easy to publish fresh content. For example, earlier you learned how to use the AWS command line tool to upload content from your local hard drive to Amazon S3.

Your website content should be managed using version control software (such as Git, Subversion, or Microsoft Team Foundation Server) to make it possible to revert to older versions of your files.²⁷ AWS offers a managed source control service called AWS CodeCommit that makes it easy to host secure and private Git repositories. Regardless of which version control system your team uses, consider tying it to a continuous build/integration tool so that your website will automatically update whenever the content changes.

For example, if your team is using a Git-based repository for version control, a Git post-commit hook can notify your continuous integration tool (e.g., Jenkins) of any content updates. At that point, your continuous integration tool can perform the actual deployment to synchronize the content with Amazon S3 (using either the AWS CLI or the Amazon S3 API), and notify the user of the deployment status. Setting up continuous deployment will streamline the process for keeping content fresh, and we highly recommend it as a best practice.

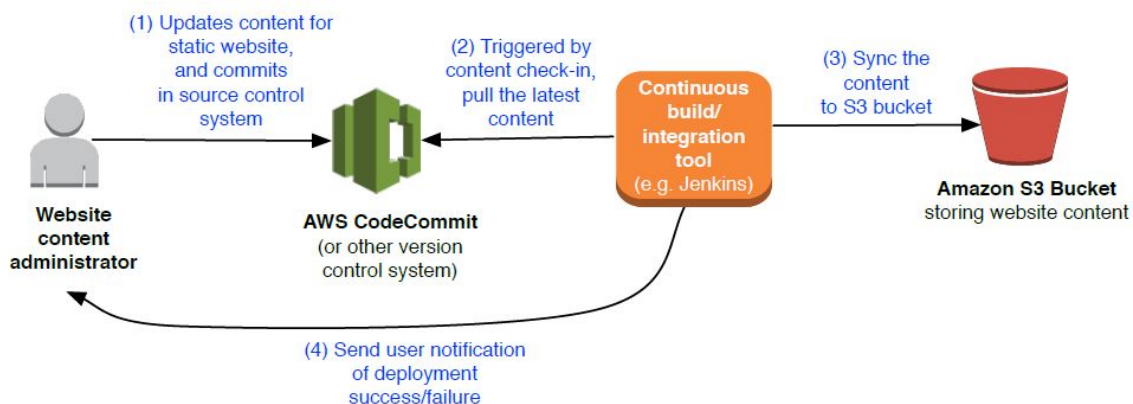


Figure 10: Example continuous deployment process

If you don't want to use version control, then be sure to periodically download your website and back up that snapshot. The AWS CLI lets you download your entire website with a single command:

```
aws s3 sync s3://bucket /my_local_backup_directory
```

Access Logs

Access logs can help you troubleshoot or analyze traffic coming to your site. Both Amazon CloudFront and Amazon S3 give you the option of turning on access logs. There's no extra charge to enable logging, other than the storage of the actual logs. The access logs are delivered on a best-effort basis; they are usually delivered within a few hours after the events are recorded.

Analyzing Logs

Amazon S3 access logs are deposited in your Amazon S3 bucket as plain text files. Each record in the log files provides details about a single Amazon S3 access request, such as the requester, bucket name, request time, request action, response status, and error code, if any. You could open individual log files in a text editor.

However, it's generally far better and easier to use one of the many third-party tools that can interpret the Amazon S3 access log format.

CloudFront logs are deposited in your Amazon S3 bucket as GZIP-compressed text files. CloudFront logs follow the standard W3C extended log file format and can be analyzed using any of the myriad log analyzers that are available.

If you prefer to build out a custom analytics solution, AWS Lambda and Amazon Elasticsearch Service (Amazon ES) can help. AWS Lambda functions can be [hooked to an Amazon S3 bucket](#) to detect when new log files are available for processing.²⁸ AWS Lambda function code can process the log files and send the data to an Amazon ES cluster. Users can then analyze the logs by querying Amazon ES or using the Kibana visual dashboard. Both AWS Lambda and Amazon ES are managed services, and there are no servers to manage.

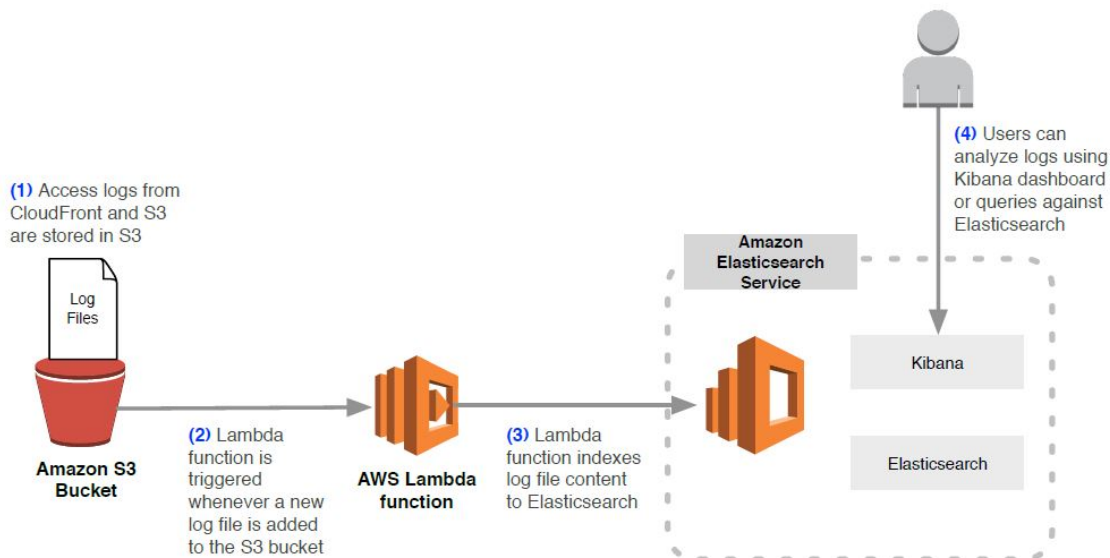


Figure 11: Using AWS Lambda to send logs from Amazon S3 to Amazon Elasticsearch Service

Archiving and Purging Logs

So how long should you preserve logs? Historical logs can help you track multi-year trends or perform security audits. Amazon S3 buckets don't have a storage cap, and you're free to retain logs for as long as you want. As a practical matter, however, you should archive and/or purge older logs.

An AWS best practice is to archive files into Amazon Glacier, which you can think of as a lower-cost distant cousin of Amazon S3. Amazon Glacier is suitable for long-term storage of infrequently accessed files. Like Amazon S3, Amazon Glacier is also designed for 99.999999999% data durability, and you have practically unlimited storage. The difference is in retrieval time. Amazon S3 supports immediate file retrieval. With Amazon Glacier, after you initiate a file retrieval request, there will be a delay before you can start downloading the files. As long as you understand that tradeoff, Amazon Glacier is a terrific choice for archiving old logs. In the US West (Oregon) Region (us-west-2) for example, Amazon Glacier storage costs are just \$0.004 per GB each month—significantly cheaper than S3, disks, or tape drives (pricing as of March 2017).

The easiest way to archive data into Amazon Glacier is to use [Amazon S3 lifecycle policies](#).²⁹ The lifecycle policies can be applied to an entire Amazon S3 bucket or to specific objects within the bucket (e.g., only the log files). A minute

of configuration in the Amazon S3 console can reduce your storage costs significantly in the long run. Here's an example of setting up data tiering using lifecycle policies:

- Lifecycle policy #1: After X days, automatically move logs from Amazon S3 into Amazon Glacier.
- Lifecycle policy #2: After Y days, automatically delete logs from Amazon Glacier.

Data tiering is illustrated in Figure 12.

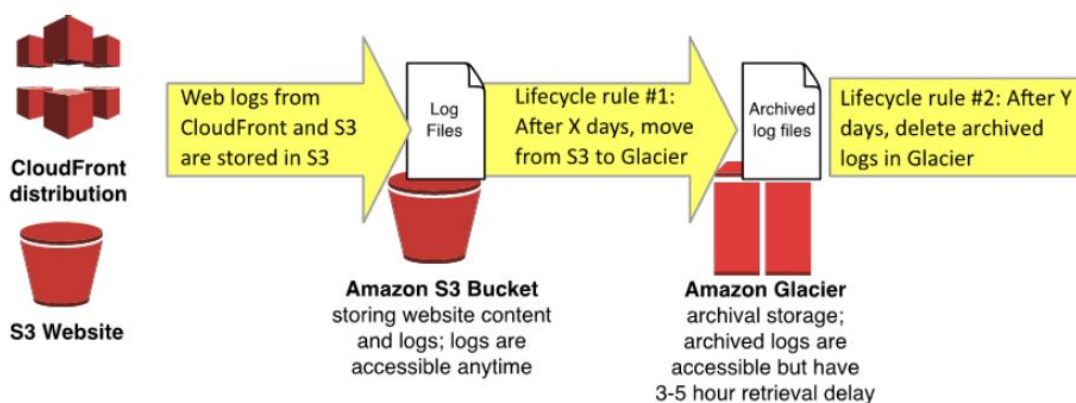


Figure 12: Data tiering using Amazon S3 lifecycle policies

Securing Administrative Access to Your Website Resources

Under the AWS shared responsibility model, the responsibility for a secure website is shared between AWS and the customer (you). AWS provides a global secure infrastructure and foundation compute, storage, networking, and database services, as well as higher level services. AWS also provides a range of security services and features that you can use to secure your assets.

As an AWS customer, you're responsible for protecting the confidentiality, integrity, and availability of your data in the cloud, and for meeting your specific business requirements for information protection. We strongly recommend working closely with your Security and Governance teams to implement the recommendations in this whitepaper as well as the ones covered in the [AWS Security Best Practices](#) whitepaper.³⁰

A benefit of using Amazon S3 and Amazon CloudFront as a serverless architecture is that the security model is also simplified. You have no operating system to harden, servers to patch, or software vulnerabilities to generate concern. Also, Amazon S3 offers security options such as [server-side data encryption](#) and access control lists.³¹ This results in a significantly reduced surface area for potential attacks.

Managing Administrator Privileges

Which team members should be able to modify your Amazon S3 bucket? Who should be able to tune your Amazon CloudFront distribution or manage Amazon Route 53?

Enforcing the principle of *least privilege* is a critical part of a [security and governance strategy](#).³² In most organizations, the team in charge of DNS configurations is separate from the team that manages web content. You should grant users appropriate levels of permissions to access the resources they need, but no more than that. In AWS, you can use IAM to lock down permissions.

You can create multiple IAM users under your AWS account, each with their own login and password.³³ An IAM user can be a person, service, or application that needs access to your AWS resources (in this case, Amazon S3 buckets, Amazon CloudFront distributions, and Amazon Route 53 hosted zones).³⁴ You can then organize them into IAM groups based on functional roles. When an IAM user is placed in an IAM group, it will inherit the group's permissions.

The fine-grained policies of IAM allow you to grant administrators the minimal privileges needed to accomplish their tasks. The permissions can be scoped to specific Amazon S3 buckets and Amazon Route 53 hosted zones.

For example, the *separation of duties* might look like this:

IAM configuration can be managed by:

- *Super_Admins*

Amazon Route 53 configuration can be managed by:

- *Super_Admins*
- *Network_Admins*

CloudFront configuration can be managed by:

- *Super_Admins*
- *Network_Admins*
- *Website_Admin*

Amazon S3 configuration can be managed by:

- *Super_Admins*
- *Website_Admin*

Amazon S3 content can be updated by:

- *Super_Admins*
- *Website_Admin*
- *Website_Content_Manager*

An IAM user can belong to more than one IAM group. For example, if someone needs to manage both Amazon Route 53 and Amazon S3, that user can belong to both the *Network_Admins* and the *Website_Admins* groups.

The best practice is to require all IAM users to rotate their IAM passwords periodically. Multi-factor authentication (MFA) should be enabled for any IAM user account with administrator privileges.

Auditing API Calls Made in Your AWS Account

A single misconfiguration error of Amazon S3, Amazon CloudFront, or Amazon Route 53 can knock your website offline. When changes are made that have an impact on your AWS environment, it's important to be able to identify who made the change, when, and what the change was.

You can use AWS CloudTrail to see an audit trail for API activity in your AWS account. Toggle it on for all AWS Regions, and the audit logs will be deposited to an Amazon S3 bucket. You can use the AWS Management Console to search against API activity history. Or you can use a third-party log analyzer to analyze and visualize the CloudTrail logs.

Do you want to build a custom analyzer instead? You can start by configuring CloudTrail to send the data to Amazon CloudWatch Logs. CloudWatch Logs allows you to create automated rules that notify you of suspicious API activity. CloudWatch Logs also has seamless integration with Amazon ES, and you can configure the data to be automatically streamed over to a managed Amazon ES cluster. Once the data is in Amazon ES, users can query against that data directly or visualize the analytics using a Kibana dashboard.

Controlling How Long Amazon S3 Content Is Cached by Amazon CloudFront

It is important to control how long your Amazon S3 content is cached at the CloudFront edge locations. This helps make sure that website updates appear correctly. If you're ever confused by a situation in which you've updated your website, but you are still seeing stale content when visiting your CloudFront powered website, one likely reason is that CloudFront is still serving up cached content. You can [control CloudFront caching behavior](#) with a combination of Cache-Control HTTP headers, CloudFront Minimum Time-to-Live (TTL) specifications, Maximum TTL specifications, content versioning, and CloudFront Invalidation Requests.³⁵ Using these correctly will help you manage website updates.

So how long does a CloudFront edge location keep an item in its cache before expiring it? CloudFront will typically serve cached content from an edge location until the content expires. After it expires, the next time that content is requested

by an end user, CloudFront will go back to the Amazon S3 origin server to fetch the content and then cache it. CloudFront edge locations automatically expire content after Maximum TTL seconds elapse (by default, this is 24 hours). However, it could be sooner because CloudFront reserves the flexibility to expire content if it needs to, before the Maximum TTL is reached. By default, the Minimum TTL is set to 0 (zero) seconds, and that's configurable as well. So to be more accurate, CloudFront may expire content anytime between the Minimum TTL (default is 0 seconds) and Maximum TTL (default is 24 hours). For example, if Minimum TTL=60s and Maximum TTL=600s, then content will be cached for *at least* 60 seconds and *at most* 600 seconds.

For example, say you deploy updates to your marketing website, with the latest and greatest product images. After uploading your new images to Amazon S3, you immediately browse to your website DNS, and you still see the old images! It is likely that one and possibly more CloudFront edge locations are still holding onto cached versions of the older images and serving the cached versions up to your website visitors. If you're the patient type, you can wait for CloudFront to expire the content, but it could take up to Maximum TTL seconds for that to happen. Often a stale website is simply not acceptable.

There are several ways to tackle this, each with its pros and cons. The first approach is that you can set the Maximum TTL to be a relatively low value. The tradeoff is that cached content expires faster because of the low Maximum TTL value. This results in more frequent requests to your Amazon S3 bucket because the CloudFront caches need to be repopulated more often. In addition, the Maximum TTL setting applies across the board to all CloudFront files, and for some websites you might want to control cache expiration behaviors based on file types.

The second approach is to implement content versioning. Every time you update website content, embed a version identifier in the file names. It can be a timestamp, a sequential number, or any other way that allows you to distinguish between two versions of the same file. For example, instead of `banner.jpg`, call it `banner_20170401_v1.jpg`. When you [update the image](#), name the new version `banner_20170612_v1.jpg` and update all files that need to link to the new image.³⁶

In the following example, the banner and logo images are updated and receive new file names. However, because those images are referenced by the HTML

files, the HTML markup also needs to be updated to reference the new image file names. Note that the HTML file names shouldn't have version identifiers in order to provide stable URLs for end users.

Stale website updated April 4, 2017

```
├─ css/
│   ├── main_20170401_v1.css
│   └─ navigation_20170401_v1.css
├─ images/
│   ├── banner_20170401_v1.jpg
│   └─ logo_20170401_v1.jpg
├─ index.html /
├─ scripts/
│   ├── script1_20170204_v1.js
│   └─ script2_20170204_v1.js
├─ section1.html
└─ section2.html
```

Website with images updated on June 15, 2017

```
├─ css/
│   ├── main_20170401_v1.css
│   └─ navigation_20170401_v1.css
├─ images/
│   ├── banner_20170612_v1.jpg
│   └─ logo_20170612_v1.jpg
├─ index.html
├─ scripts/
│   ├── script1_20170204_v1.js
│   └─ script2_20170204_v1.js
├─ section1.html
└─ section2.html
```

Content versioning has a clear benefit: it sidesteps CloudFront expiration behaviors altogether. Since new file names are involved, CloudFront will immediately fetch the new files from Amazon S3 (and afterwards, cache them).

Non-HTML website changes are reflected immediately. Additionally, you can roll back and roll forward between different versions of your website.

The main challenge is that content update processes will need to be “version-aware.” File names will need to be versioned. Files with references to changed files will also need to be updated. For example, if an image is updated, the following need to be updated as well:

- The image file name
- Content in any HTML, CSS, and JavaScript files referencing the older image file name
- The file names of any referencing files (with the exception of HTML files)

A few static site generator tools can automatically rename file names with version identifiers, but most tools aren’t version-aware. Manually managing version identifiers can be cumbersome and error-prone. If your website would benefit from content versioning, it might be worth investing in a few automation scripts to streamline your content update process.

A third approach to managing CloudFront expiration behavior is [to specify Cache-Control headers](#) for your website content.³⁷ Remember how CloudFront reserves the flexibility to expire content anytime between the Minimum TTL and Maximum TTL seconds? There’s a way you can override that behavior. If you keep the Minimum TTL at the default 0 seconds, then CloudFront will honor any `Cache-Control: max-age` HTTP header that is individually set for your content. So if an image is configured with a `Cache-Control: max-age=60` header, then CloudFront will expire it at the 60 second mark. This gives you more precise control over content expiration for individual files.

You can configure Amazon S3 to return a `Cache-Control` HTTP header with the value of `max-age=<seconds>` when S3 serves up the content. This setting is on a file-by-file basis, and we recommend using different values depending on the file type (HTML, CSS, JavaScript, images, etc.). Since HTML files won’t have version identifiers in their file names, we recommend using smaller `max-age` values for HTML files so that CloudFront will expire the HTML files sooner than other content. You can set this by editing the Amazon S3 object metadata using the AWS Management Console.

Figure 13: Setting Cache-Control Values in the console

In practice, you should automate this as part of your Amazon S3 upload process. If you're using the AWS CLI, as we recommended earlier, a small tweak to your deployment scripts will do the trick:

```
aws s3 sync /path s3://yourbucket/ --delete --recursive \
    --cache-control max-age=60
```

Lastly, a fourth approach for managing CloudFront expiration behavior is to use [CloudFront invalidation requests](#).³⁸ Invalidation requests are a way to force CloudFront to expire content. Invalidation requests aren't immediate. It takes several minutes from the time you submit one to the time that CloudFront actually expires the content. For the occasional requests, you can submit them using the AWS Management Console. Otherwise, use the AWS CLI or AWS APIs to script the invalidation. In addition, CloudFront lets you specify which content should be invalidated: You can choose to invalidate your entire Amazon S3 bucket, individual files, or just those matching a wildcard pattern. For example, to invalidate only the images directory, issue an invalidation request for: `/images/`.

The best practice is to understand and use the four approaches together. If possible, implement content versioning. It allows you to immediately review changes and gives you the most precise control over the CloudFront and Amazon S3 experience. Set the Minimum TTL to be 0 seconds and the Maximum TTL to be a relatively low value. Also, use `Cache-Control` headers for individual pieces of content. If your website is infrequently updated, then set a large value for `Cache-Control :max-age=<seconds>` and then issue

CloudFront invalidation requests every time your site is updated. If the website is updated more frequently, use a relatively small value for `Cache-Control:max-age=<seconds>` and then issue CloudFront invalidation requests only if the `Cache-Control:max-age=<seconds>` settings exceeds several minutes.

Conclusion

This whitepaper started with a look at traditional (non-AWS) architectures for static websites. We then showed you an AWS Cloud-native architecture based on Amazon S3, Amazon CloudFront, and Amazon Route 53.

The AWS architecture is highly available and scalable, secure, and provides for a responsive user experience at very low cost. By enabling and analyzing the available logs, you can understand your visitors and how well the website is performing. Fewer moving parts means less maintenance is required. In addition, the architecture costs only a few dollars a month to run.

Contributors

The following individuals and organizations contributed to this document:

- Jim Tran, AWS Principal Enterprise Solutions Architect
- Bhushan Nene, Senior Manager, AWS Solutions Architecture

Notes

¹ <http://docs.aws.amazon.com/AmazonS3/latest/dev/WebsiteHosting.html>

² Each S3 object can be 1 byte to 5 TB in file size, and there's no limit to the number of Amazon S3 objects you can store.

³ <http://aws.amazon.com/s3/sla/>

⁴ For a list of available AWS Regions, visit: <http://aws.amazon.com/about-aws/global-infrastructure/>. For an overview of AWS Regions and Availability Zones, visit <http://aws.amazon.com/about-aws/global-infrastructure/>. If your high-availability requirements require that your website must remain available even in the case of a failure of an entire AWS Region, you might want to

explore the Amazon S3 Cross-Region Replication capability to automatically replicate your S3 data to another S3 bucket in a second AWS Region. See:

<http://docs.aws.amazon.com/AmazonS3/latest/UG/cross-region-replication.html>

5 <http://docs.aws.amazon.com/AmazonS3/latest/dev/BucketRestrictions.html>

6 <http://docs.aws.amazon.com/gettingstarted/latest/swh/website-hosting-intro.html>

7 <http://docs.aws.amazon.com/gettingstarted/latest/swh/getting-started-configure-route53.html>

8 If you're familiar with Microsoft IIS web servers, this is equivalent to "default.html"; for Apache web servers, this is equivalent to "index.html"

9 <http://docs.aws.amazon.com/AmazonS3/latest/dev/ServerLogs.html>

10 <http://docs.aws.amazon.com/AmazonS3/latest/dev/BucketBilling.html>

11

<http://docs.aws.amazon.com/AmazonS3/latest/dev/IndexDocumentSupport.html>

12 <http://aws.amazon.com/cli/>

13 Refer to <http://docs.aws.amazon.com/cli/latest/reference/s3/sync.html> for more details on the AWS CLI command. For moving very large quantities of data, there are alternative methods of moving large numbers of large files into S3: <https://aws.amazon.com/s3/cloud-data-migration/>

14 <http://docs.aws.amazon.com/AmazonS3/latest/dev/BucketBilling.html>

15 More specifically, the number of bytes that are delivered to the website visitor in the HTTP responses.

16 In addition to ensuring that end-user requests are served by the closest edge location, Amazon CloudFront also keeps persistent connections with your origin servers so that those files can be fetched from the origin servers as quickly as possible. Finally, Amazon CloudFront uses additional optimizations (e.g., wider TCP initial congestion window) to provide higher performance while delivering your content to viewers.

17 <http://aws.amazon.com/cloudfront/details/>

18

<http://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/MigrateS3ToCloudFront.html>

19 “What about replicating our website content to Amazon S3 buckets in the other AWS Regions, and then serving up websites from the regional S3 bucket closest to the user?” There’s no way to route all requests to a single domain to one of several Amazon S3 buckets. An Amazon S3 bucket name must be the same as the CNAME, and S3 bucket names are globally unique, so no two buckets can have the same CNAME. For use cases where it’s fine to direct different sets of users to distinct websites (with different domains/subdomains), you can use two different S3 buckets and leverage the Amazon S3 Cross Region Replication capability:

<http://docs.aws.amazon.com/AmazonS3/latest/UG/cross-region-replication.html>

20 <http://calculator.s3.amazonaws.com/>

21 Any pricing information included in this document is provided only as an estimate of usage charges for AWS services based on the prices effective at the time of this writing. Monthly charges will be based on your actual use of AWS services, and may vary from the estimates provided.

22 <https://aws.amazon.com/cloudfront/pricing/>

23 <http://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/cost-explorer-what-is.html>

24 <http://docs.aws.amazon.com/AmazonS3/latest/dev/BucketBilling.html>

25 For static websites that host very large numbers of large files that are infrequently accessed (such as static video content), one potential cost optimization to consider is to set up an Amazon S3 Lifecycle Rule to move content from the “S3 Standard” storage class to the “S3 Standard - Infrequently- Accessed” storage class in order to take advantage of a lower-priced storage tier for infrequently-accessed files. This blog post has more details: <https://aws.amazon.com/blogs/aws/aws-storage-update-new-lower-cost-s3-storage-option-glacier-price-reduction/>

26 <https://aws.amazon.com/blogs/aws/the-new-cost-explorer-for-aws/>

27 If version control software is not in use at your organization, one alternative approach is to look at the Amazon S3 object versioning feature for versioning your critical files. Note that object versioning introduces storage costs for each version, and requires you to programmatically manage the different versions. For more information, see

<http://docs.aws.amazon.com/AmazonS3/latest/dev/Versioning.html>.

28 <http://docs.aws.amazon.com/lambda/latest/dg/with-s3.html>

29 <https://aws.amazon.com/blogs/aws/amazon-s3-lifecycle-management-update/>

30 <https://aws.amazon.com/whitepapers/aws-security-best-practices/>

31

<http://docs.aws.amazon.com/AmazonS3/latest/dev/UsingServerSideEncryption.html>

32 <http://blogs.aws.amazon.com/security/post/Tx4BUZIS3E2QG2/Make-a-New-Year-s-Resolution-Adhere-to-IAM-Best-Practices>

33 The AWS account is the account that you create when you first sign up for AWS. Each AWS account has root permissions to all AWS resources and services. The best practice is to enable multi-factor authentication (MFA) for your root account, and then lock away the root credentials so that no person or system uses the root credentials directly for day-to-day operations. Instead, create IAM groups and IAM users for the day-to-day operations.

34 through the AWS Management Console, command line tools, or APIs

35

<http://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/Expiration.html>

36

<http://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/ReplacingObjects.html>

37

<http://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/Expiration.html>

38

<http://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/Invalidation.html>