

Best Practices for WordPress on AWS

Reference architecture for scalable
WordPress-powered websites

First Published December 2014

Updated October 19, 2021



Notices

Customers are responsible for making their own independent assessment of the information in this document. This document: (a) is for informational purposes only, (b) represents current AWS product offerings and practices, which are subject to change without notice, and (c) does not create any commitments or assurances from AWS and its affiliates, suppliers or licensors. AWS products or services are provided “as is” without warranties, representations, or conditions of any kind, whether express or implied. 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.

© 2021 Amazon Web Services, Inc. or its affiliates. All rights reserved.

Contents

- Introduction 1
- Simple deployment 1
 - Considerations 1
 - Available approaches 1
 - Amazon Lightsail 2
- Improving performance and cost efficiency 4
 - Accelerating content delivery 4
 - Database caching 7
 - Bytecode caching 7
- Elastic deployment 8
 - Reference architecture 8
 - Architecture components 9
 - Scaling the web tier 9
 - Stateless web tier 11
 - WordPress high availability by Bitnami on AWS Quick Starts 14
- Conclusion 16
- Contributors 16
- Document revisions 16
- Appendix A: CloudFront configuration 17
 - Origins and behaviors 17
 - CloudFront distribution creation 17
- Appendix B: Plugins installation and configuration 20
 - AWS for WordPress plugin 20
 - Static content configuration 26
- Appendix C: Backup and recovery 29
- Appendix D: Deploying new plugins and themes 31

Abstract

This whitepaper provides system administrators with specific guidance on how to get started with WordPress on Amazon Web Services (AWS) and how to improve both the cost efficiency of the deployment and the end user experience. It also outlines a reference architecture that addresses common scalability and high availability requirements.

Introduction

WordPress is an open-source blogging tool and content management system (CMS) based on PHP and MySQL that is used to power anything from personal blogs to high-traffic websites.

When the first version of WordPress was released in 2003, it was not built with modern elastic and scalable cloud-based infrastructures in mind. Through the work of the WordPress community and the release of various WordPress modules, the capabilities of this CMS solution are constantly expanding. Today, it is possible to build a WordPress architecture that takes advantage of many of the benefits of the AWS Cloud.

Simple deployment

For low-traffic blogs or websites without strict high availability requirements, a simple deployment of a single server might be suitable. This deployment isn't the most resilient or scalable architecture, but it is the quickest and most economical way to get your website up and running.

Considerations

This discussion starts with a single web server deployment. There may be occasions when you outgrow it, for example:

- The virtual machine that your WordPress website is deployed on is a single point of failure. A problem with this instance causes a loss of service for your website.
- Scaling resources to improve performance can only be achieved by “vertical scaling;” that is, by increasing the size of the virtual machine running your WordPress website.

Available approaches

AWS has a number of different options for provisioning virtual machines. There are three main ways to host your own WordPress website on AWS:

- Amazon Lightsail
- Amazon Elastic Compute Cloud (Amazon EC2)
- AWS Marketplace

[Amazon Lightsail](#) is a service that enables you to quickly launch a virtual private server (a Lightsail instance) to host a WordPress website. Lightsail is the easiest way to get started if you don't need highly configurable instance types or access to advanced networking features.

[Amazon EC2](#) is a web service that provides resizable compute capacity so you can launch a virtual server within minutes. Amazon EC2 provides more configuration and management options than Lightsail, which is desirable in more advanced architectures. You have administrative access to your EC2 instances and can install any software packages you choose, including WordPress.

[AWS Marketplace](#) is an online store where you can find, buy, and quickly deploy software that runs on AWS. You can use one-click deployment to launch preconfigured WordPress images directly to Amazon EC2 in your own AWS account in just a few minutes. There are a number of AWS Marketplace vendors offering ready-to-run WordPress instances.

This whitepaper covers the Lightsail option as the recommended implementation for a single-server WordPress website.

Amazon Lightsail

Lightsail is the easiest way to get started on AWS for developers, small businesses, students, and other users who need a simple virtual private server (VPS) solution.

The service abstracts many of the more complex elements of infrastructure management away from the user. It is, therefore, an ideal starting point if you have less infrastructure experience, or when you need to focus on running your website and a simplified product is sufficient for your needs.

With Amazon Lightsail, you can choose Windows or Linux/Unix operating systems and popular web applications, including WordPress, and deploy these with a single click from preconfigured templates.

As your needs grow, you have the ability to smoothly step outside of the initial boundaries and connect to additional AWS database, object storage, caching, and content distribution services.

Selecting an Amazon Lightsail pricing plan

A [Lightsail plan](#) defines the monthly cost of the Lightsail resources you use to host your WordPress website. There are a number of plans available to cover a variety of use

cases, with varying levels of CPU resource, memory, solid state drive (SSD) storage, and data transfer. If your website is complex, you may need a larger instance with more resources. You can achieve this by migrating your server to a larger plan [using the web console](#) or as described in the [Amazon Lightsail CLI documentation](#).

Installing WordPress

Lightsail provides templates for commonly used applications such as WordPress. This template is a great starting point for running your own WordPress website, as it comes pre-installed with most of the software you need. You can install additional software or customize the software configuration by using the in-browser terminal or your own SSH client, or via the WordPress administration web interface.

Amazon Lightsail has a partnership with GoDaddy Pro Sites product to help WordPress customers easily manage their instances for free. Lightsail WordPress virtual servers are preconfigured and optimized for fast performance and security, making it easy to get your WordPress site up and running in no time. Customers running multiple WordPress instances find it challenging and time-consuming to update, maintain and manage all of their sites. With this integration, you can easily manage your multiple WordPress instances in minutes with only a few clicks.

For more information about managing WordPress on Lightsail, refer to [Getting started using WordPress from your Amazon Lightsail instance](#). Once you are finished customizing your WordPress website, AWS recommends that you take a snapshot of your instance.

A [snapshot](#) is a way to create a backup image of your Lightsail instance. It is a copy of the system disk and also stores the original machine configuration (that is, memory, CPU, disk size, and data transfer rate). Snapshots can be used to revert to a known good configuration after a bad deployment or upgrade.

This snapshot enables you to recover your server if needed, but also to launch new instances with the same customizations.

Recovering from failure

A single web server is a single point of failure, so you must ensure that your website data is backed up. The snapshot mechanism described earlier can also be used for this purpose. To recover from failure, you can restore a new instance from your most recent

snapshot. To reduce the amount of data that could be lost during a restore, your snapshots must be as recent as possible.

To minimize the potential for data loss, ensure that snapshots are taken on a regular basis. You can schedule automatic snapshots of your Lightsail Linux/Unix instances. For instructions, refer to [Enabling or disabling automatic snapshots for instances or disks in Amazon Lightsail](#).

AWS recommends that you use a static IP—a fixed, public IP address that is dedicated to your Lightsail account. If you need to replace your instance with another one, you can reassign the static IP to the new instance. In this way, you don't have to reconfigure any external systems (such as DNS records) to point to a new IP address every time you want to replace your instance.

Improving performance and cost efficiency

You may eventually outgrow your single-server deployment. In this case, you may need to consider options for improving your website's performance. Before migrating to a multi-server, scalable deployment (discussed later in this whitepaper), there are a number of performance and cost efficiencies you can apply. These are good practices that you should follow anyway, even if you do move to a multi-server architecture.

The following sections introduce a number of options that can improve aspects of your WordPress website's performance and scalability. Some can be applied to a single-server deployment, whereas others take advantage of the scalability of multiple servers. Many of those modifications require the use of one or more WordPress plugins. Although various options are available, [W3 Total Cache](#) is a popular choice that combines many of those modifications in a single plugin.

Accelerating content delivery

Any WordPress website needs to deliver a mix of static and dynamic content. Static content includes images, JavaScript files, or style sheets. Dynamic content includes anything generated on the server side using the WordPress PHP code; for example, elements of your site that are generated from the database or personalized to each viewer.

An important aspect of the end user experience is the network latency involved when delivering the previous content to users around the world. Accelerating the delivery of the previous content improves the end user experience, especially users geographically

spread across the globe. This can be achieved with a Content Delivery Network (CDN) such as Amazon CloudFront.

[Amazon CloudFront](#) is a web service that provides an easy and cost-effective way to distribute content with low latency and high data transfer speeds through multiple edge locations across the globe. Viewer requests are automatically routed to a suitable CloudFront [edge location](#) to lower the latency.

If the content can be cached (for a few seconds, minutes, or even days) and is already stored in a particular edge location, CloudFront delivers it immediately. If the content should not be cached, has expired, or isn't currently in that edge location, CloudFront retrieves content from one or more sources of truth, referred to as the origin(s) (in this case, the Lightsail instance) in the CloudFront configuration. This retrieval takes place over optimized network connections, which work to speed up the delivery of content on your website. Apart from improving the end user experience, the model discussed also reduces the load on your origin servers and has the potential to create significant cost savings.

Static content offload

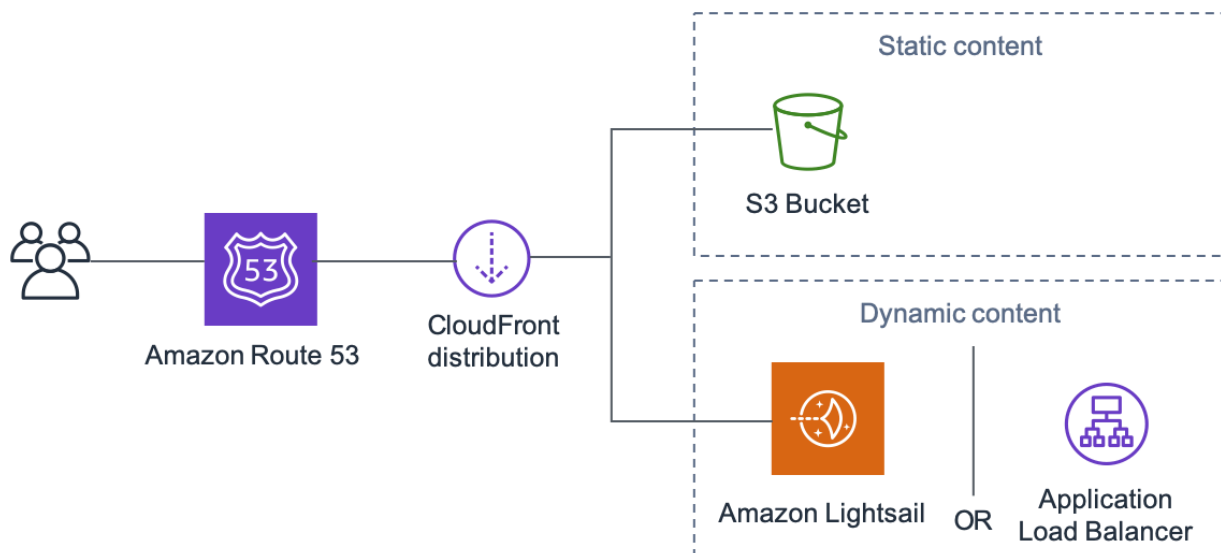
This includes CSS, JavaScript, and image files—either those that are part of your WordPress themes or those media files uploaded by the content administrators. All these files can be stored in Amazon Simple Storage Service (Amazon S3) using a plugin such as W3 Total Cache and served to users in a scalable and highly available manner. [Amazon S3](#) offers a highly scalable, reliable, and low-latency data storage infrastructure at low cost, which is accessible via REST APIs. Amazon S3 redundantly stores your objects, not only on multiple devices, but also across multiple facilities in an AWS Region, providing exceptionally high levels of durability.

This has the positive side effect of offloading this workload from your Lightsail instance and letting it focus on dynamic content generation. This reduces the load on the server and is an important step towards creating a stateless architecture (a prerequisite before implementing automatic scaling).

You can subsequently configure Amazon S3 as an origin for CloudFront to improve delivery of those static assets to users around the world. Although WordPress isn't integrated with Amazon S3 and CloudFront out of the box, a variety of plugins add support for these services (for example, W3 Total Cache).

Dynamic content

Dynamic content includes the output of server-side WordPress PHP scripts. Dynamic content can also be served via CloudFront by configuring the WordPress website as an origin. Since dynamic content includes personalized content, you need to configure CloudFront to forward certain HTTP cookies and HTTP headers as part of a request to your custom origin server. CloudFront uses the forwarded cookie values as part of the key that identifies a unique object in its cache. To ensure that you maximize the caching efficiency, configure CloudFront to forward only those HTTP cookies and HTTP headers that actually vary the content (not cookies that are only used on the client side or by third-party applications, for example, for web analytics).



Whole website delivery via Amazon CloudFront

The preceding figure includes two origins: one for static content and another for dynamic content. For implementation details, refer to [Appendix A: CloudFront configuration](#) and [Appendix B: Plugins installation and configuration](#).

CloudFront uses standard cache control headers to identify if and for how long it should cache specific HTTP responses. The same cache control headers are also used by web browsers to decide when and for how long to cache content locally for a more optimal end user experience (for example, a `.css` file that is already downloaded will not be redownloaded every time a returning visitor views a page). You can configure cache control headers on the web server level (for example, via `.htaccess` files or modifications of the `httpd.conf` file) or install a WordPress plugin (for example, W3 Total Cache) to dictate how those headers are set for both static and dynamic content.

Database caching

Database caching can significantly reduce latency and increase throughput for read-heavy application workloads like WordPress. Application performance is improved by storing frequently accessed pieces of data in memory for low-latency access (for example, the results of input/output (I/O)-intensive database queries). When a large percentage of the queries is served from the cache, the number of queries that need to hit the database is reduced, resulting in a lower cost associated with running the database.

Although WordPress has limited caching capabilities out-of-the-box, a variety of plugins support integration with [Memcached](#), a widely adopted memory object caching system. The W3 Total Cache plugin is a good example.

In the simplest scenarios, you install Memcached on your web server and capture the result as a new snapshot. In this case, you are responsible for the administrative tasks associated with running a cache.

Another option is to take advantage of a managed service such as [Amazon ElastiCache](#) and avoid that operational burden. ElastiCache makes it easy to deploy, operate, and scale a distributed in-memory cache in the cloud. You can find information about how to connect to your ElastiCache cluster nodes in the [Amazon ElastiCache documentation](#).

If you are using Lightsail and wish to access an ElastiCache cluster in your AWS account privately, you can do so by using VPC peering. For instructions to enable VPC peering, refer to [Set up Amazon VPC peering to work with AWS resources outside of Amazon Lightsail](#).

Bytecode caching

Each time a PHP script is run, it gets parsed and compiled. By using a PHP bytecode cache, the output of the PHP compilation is stored in RAM so the same script doesn't have to be compiled again and again. This reduces the overhead related to running PHP scripts, resulting in better performance and lower CPU requirements.

A bytecode cache can be installed on any Lightsail instance that hosts WordPress and can greatly reduce its load. For PHP 5.5 and later, AWS recommends the use of [OPcache](#), a bundled extension with that PHP version.

Note that OPcache is enabled by default in the Bitnami WordPress Lightsail template, so no further action is required.

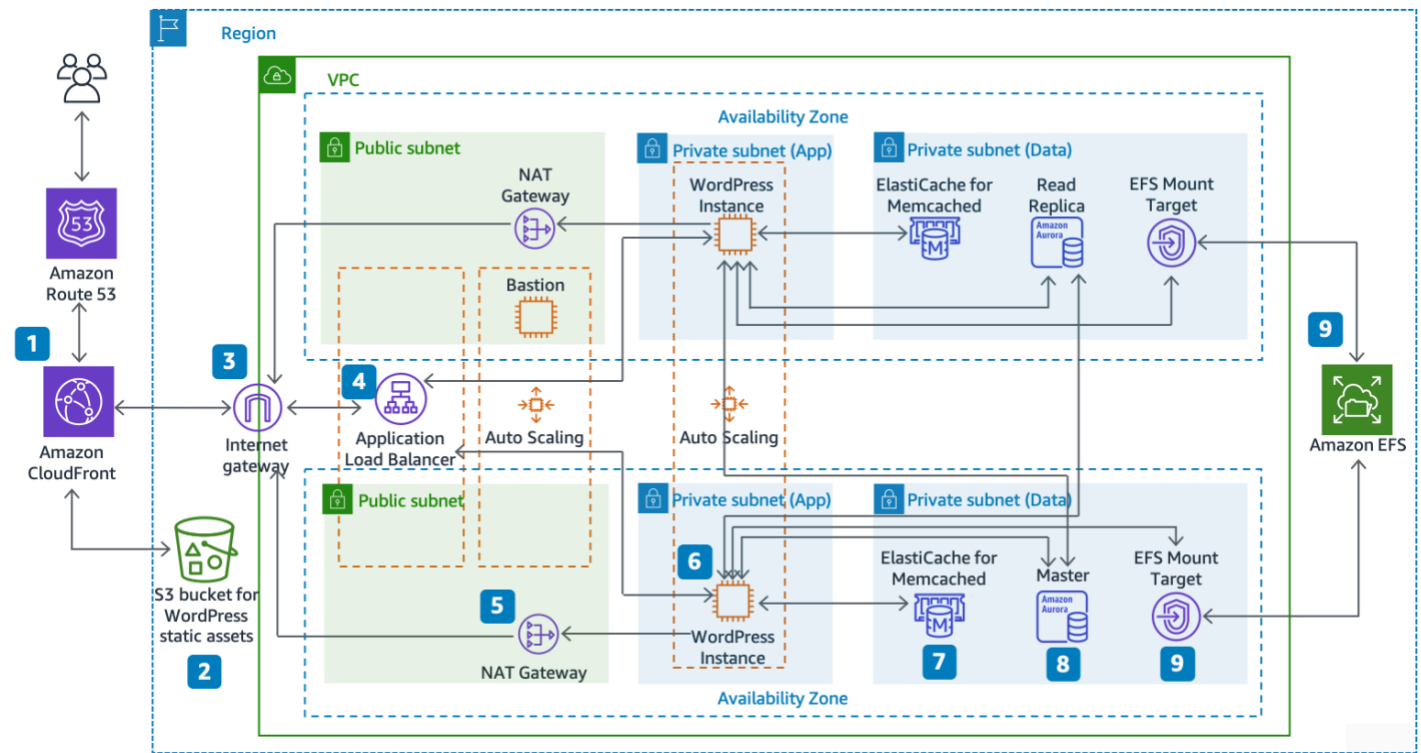
Elastic deployment

There are many scenarios where a single-server deployment may not be sufficient for your website. In these situations, you need a multi-server, scalable architecture.

Reference architecture

The [Hosting WordPress on AWS reference architecture](#) available on GitHub outlines best practices for deploying WordPress on AWS and includes a set of AWS CloudFormation templates to get you up and running quickly. The following architecture is based on that reference architecture. The rest of this section reviews the reasons behind the architectural choices.

The based AMI in the GitHub was changed from Amazon Linux1 to Amazon Linux2 in July 2021. However, deployment templates at S3 were not changed yet. It is recommended to use templates at GitHub if there is an issue to deploy the reference architecture with templates at S3.



Reference architecture for hosting WordPress on AWS

Architecture components

The preceding reference architecture illustrates a complete best practice deployment for a WordPress website on AWS.

- It starts with edge caching in **Amazon CloudFront** (1) to cache content close to end users for faster delivery.
- CloudFront pulls static content from an **S3 bucket** (2) and dynamic content from an **Application Load Balancer** (4) in front of the web instances.
- The web instances run in an **Auto Scaling group of Amazon EC2 instances** (6).
- An **ElastiCache** cluster (7) caches frequently queried data to speed up responses.
- An **Amazon Aurora** MySQL instance (8) hosts the WordPress database.
- The WordPress EC2 instances access shared WordPress data on an **Amazon EFS** file system via an **EFS Mount Target** (9) in each Availability Zone.
- An **Internet Gateway** (3) enables communication between resources in your VPC and the internet.
- **NAT Gateways** (5) in each Availability Zone enable EC2 instances in private subnets (App and Data) to access the internet.

Within the **Amazon VPC** there exist two types of subnets: public (**Public Subnet**) and private (**App Subnet** and **Data Subnet**). Resources deployed into the public subnets will receive a public IP address and will be publicly visible on the internet. The **Application Load Balancer** (4) and a bastion host for administration are deployed here. Resources deployed into the private subnets receive only a private IP address and are not publicly visible on the internet, improving the security of those resources. The **WordPress web server instances** (6), **ElastiCache cluster instances** (7), **Aurora MySQL database instances** (8), and **EFS Mount Targets** (9) are all deployed in private subnets.

The remainder of this section covers each of these considerations in more detail.

Scaling the web tier

To evolve your single-server architecture into a multi-server, scalable architecture, you must use five key components:

- Amazon EC2 instances
- Amazon Machine Images (AMIs)
- Load balancers
- Automatic scaling
- Health checks

AWS provides a wide variety of EC2 instance types so you can choose the best server configuration for both performance and cost. Generally speaking, the compute-optimized (for example, C4) instance type may be a good choice for a WordPress web server. You can deploy your instances across multiple Availability Zones within an AWS Region to increase the reliability of the overall architecture.

Because you have complete control of your EC2 instance, you can log in with root access to install and configure all of the software components required to run a WordPress website. After you are done, you can save that configuration as an AMI, which you can use to launch new instances with all the customizations that you've made.

To distribute end user requests to multiple web server nodes, you need a load balancing solution. AWS provides this capability through [Elastic Load Balancing](#), a highly available service that distributes traffic to multiple EC2 instances. Because your website is serving content to your users via HTTP or HTTPS, we recommend that you make use of the Application Load Balancer, an application-layer load balancer with content routing and the ability to run multiple WordPress websites on different domains, if required.

Elastic Load Balancing supports distribution of requests across multiple Availability Zones within an AWS Region. You can also configure a health check so that the Application Load Balancer automatically stops sending traffic to individual instances that have failed (for example, due to a hardware problem or software crash). AWS recommends using the WordPress admin login page (`/wp-login.php`) for the health check because this page confirms both that the web server is running and that the web server is configured to serve PHP files correctly.

You may choose to build a custom health check page that checks other dependent resources, such as database and cache resources. For more information, refer to [Health checks for your target groups](#) in the *Application Load Balancer Guide*.

Elasticity is a key characteristic of the AWS Cloud. You can launch more compute capacity (for example, web servers) when you need it and run less when you don't. [AWS Auto Scaling](#) is an AWS service that helps you automate this provisioning to scale your Amazon EC2 capacity up or down according to conditions you define with no need for manual intervention. You can configure AWS Auto Scaling so that the number of EC2 instances you're using increases seamlessly during demand spikes to maintain performance and decreases automatically when traffic diminishes, so as to minimize costs.

Elastic Load Balancing also supports dynamic addition and removal of Amazon EC2 hosts from the load balancing rotation. Elastic Load Balancing itself also dynamically increases and decreases the load balancing capacity to adjust to traffic demands with no manual intervention.

Stateless web tier

To take advantage of multiple web servers in an automatic scaling configuration, your web tier must be *stateless*. A stateless application is one that needs no knowledge of previous interactions and stores no session information. In the case of WordPress, this means that all end users receive the same response, regardless of which web server processed their request. A stateless application can scale horizontally since any request can be serviced by any of the available compute resources (web server instances). When that capacity is no longer required, any individual resource can be safely terminated (after running tasks have been drained). Those resources do not need to be aware of the presence of their peers—all that is required is a way to distribute the workload to them.

When it comes to user session data storage, the WordPress core is completely stateless because it relies on cookies that are stored in the client's web browser. Session storage isn't a concern unless you have installed any custom code (for example, a WordPress plugin) that instead relies on native PHP sessions.

However, WordPress was originally designed to run on a single server. As a result, it stores some data on the server's local file system. When running WordPress in a multi-server configuration, this creates a problem because there is inconsistency across web servers. For example, if a user uploads a new image, it is only stored on one of the servers.

This demonstrates why we need to improve the default WordPress running configuration to move important data to shared storage. The best practice architecture

has a database as a separate layer outside the web server and makes use of shared storage to store user uploads, themes, and plugins.

Shared storage (Amazon S3 and Amazon EFS)

By default, WordPress stores user uploads on the local file system and so isn't stateless. Therefore, you need to move the WordPress installation and all user customizations (such as configuration, plugins, themes, and user-generated uploads) into a shared data platform to help reduce load on the web servers and to make the web tier stateless.

[Amazon Elastic File System](#) (Amazon EFS) provides scalable network file systems for use with EC2 instances. Amazon EFS file systems are distributed across an unconstrained number of storage servers, enabling file systems to grow elastically and enabling massively parallel access from EC2 instances. The distributed design of Amazon EFS avoids the bottlenecks and constraints inherent to traditional file servers.

By moving the entire WordPress installation directory onto an EFS file system and mounting it into each of your EC2 instances when they boot, your WordPress site and all its data is automatically stored on a distributed file system that isn't dependent on any one EC2 instance, making your web tier completely stateless. The benefit of this architecture is that you don't need to install plugins and themes on each new instance launch, and you can significantly speed up the installation and recovery of WordPress instances. It is also easier to deploy changes to plugins and themes in WordPress, as outlined in the [Deployment considerations](#) section of this document.

To ensure optimal performance of your website when running from an EFS file system, check the recommended configuration settings for Amazon EFS and OPcache on the [AWS Reference Architecture for WordPress](#).

You also have the option to offload all static assets, such as image, CSS, and JavaScript files, to an S3 bucket with CloudFront caching in front. The mechanism for doing this in a multi-server architecture is exactly the same as for a single-server architecture, as discussed in the [Static content](#) section of this whitepaper. The benefits are the same as in the single-server architecture—you can offload the work associated with serving your static assets to Amazon S3 and CloudFront, enabling your web servers to focus on generating dynamic content only and serve more user requests per web server.

Data tier (Amazon Aurora and Amazon ElastiCache)

With the WordPress installation stored on a distributed, scalable, shared network file system, and static assets being served from Amazon S3, you can focus your attention on the remaining stateful component: the database. As with the storage tier, the database should not be reliant on any single server, so it cannot be hosted on one of the web servers. Instead, host the WordPress database on Amazon Aurora.

[Amazon Aurora](#) is a MySQL and PostgreSQL compatible relational database built for the cloud that combines the performance and availability of high-end commercial databases with the simplicity and cost-effectiveness of open-source databases. Aurora MySQL increases MySQL performance and availability by tightly integrating the database engine with a purpose-built distributed storage system, backed by SSD. It is fault-tolerant and self-healing, replicates six copies of your data across three Availability Zones, is designed for greater than 99.99% availability, and nearly continuously backs up your data in Amazon S3. Amazon Aurora is designed to automatically detect database crashes and restart without the need for crash recovery or to rebuild the database cache.

Amazon Aurora provides a number of [instance types](#) to suit different application profiles, including memory-optimized and burstable instances. To improve the performance of your database you can select a large instance type to provide more CPU and memory resources.

Amazon Aurora automatically handles failover between the primary instance and [Aurora Replicas](#) so that your applications can resume database operations as quickly as possible without manual administrative intervention. Failover typically takes less than 30 seconds.

After you have created at least one Aurora Replica, connect to your primary instance using the cluster endpoint to enable your application to automatically fail over in the event the primary instance fails. You can create up to 15 low-latency read replicas across three Availability Zones.

As your database scales, your database cache will also need to scale. As discussed previously in the [Database caching](#) section of this document, ElastiCache has features to scale the cache across multiple nodes in an ElastiCache cluster, and across multiple Availability Zones in a Region for improved availability. As you scale your ElastiCache cluster, ensure that you configure your caching plugin to connect using the configuration endpoint so that WordPress can use new cluster nodes as they are added and stop

using old cluster nodes as they are removed. You must also set up your web servers to use the [ElastiCache Cluster Client for PHP](#) and update your AMI to store this change.

WordPress high availability by Bitnami on AWS Quick Starts

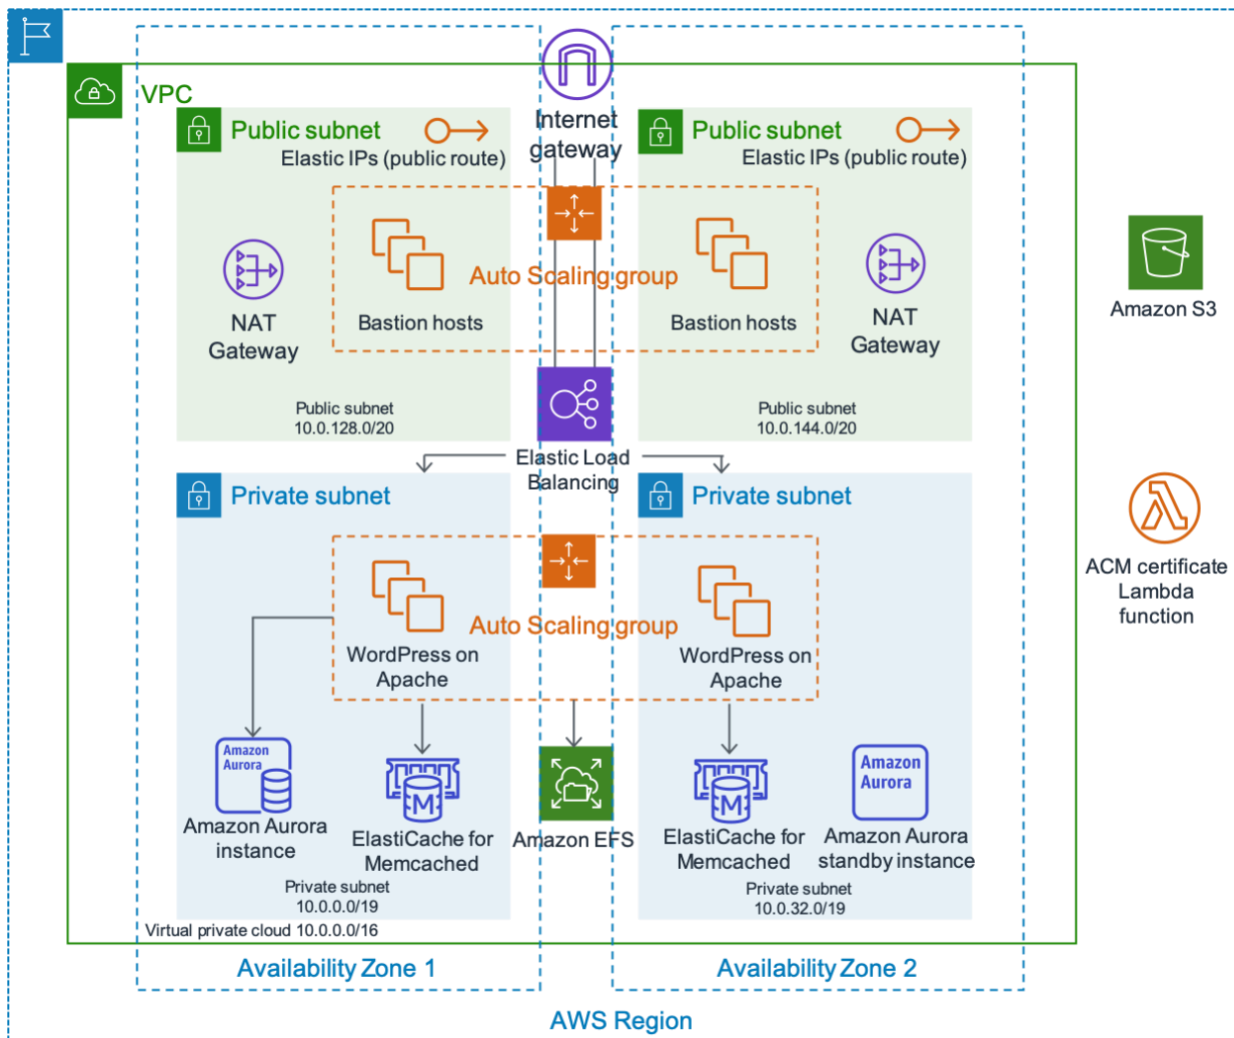
Quick Starts are built by AWS solutions architects and partners to help you deploy popular technologies on AWS, based on AWS best practices for security and high availability. These accelerators reduce hundreds of manual procedures into just a few steps, so you can build your production environment quickly and start using it immediately. Each Quick Start includes AWS CloudFormation templates that automate the deployment and a guide that discusses the architecture and provides step-by-step deployment instructions.

[WordPress High Availability by Bitnami on AWS Quick Starts](#) sets up the following configurable environment on AWS:

- A highly available architecture that spans two Availability Zones.*
- A virtual private cloud (VPC) configured with public and private subnets according to AWS best practices. This provides the network infrastructure for your deployment.*
- An internet gateway to provide access to the internet. This gateway is used by the bastion hosts to send and receive traffic.*
- In the public subnets, managed NAT gateways to allow outbound internet access for resources in the private subnets.*
- In the public subnets, Linux bastion hosts in an Auto Scaling group to allow inbound Secure Shell (SSH) access to EC2 instances in public and private subnets.*
- Elastic Load Balancing to distribute HTTP and HTTPS requests across multiple WordPress instances.
- In the private subnets, EC2 instances that host the WordPress application on Apache. These instances are provisioned in an Auto Scaling group to ensure high availability.
- In the private subnets, Amazon Aurora DB instances administered by Amazon Relational Database Service (Amazon RDS).

- In the private subnets, Amazon Elastic File System (Amazon EFS) to share assets (such as plugins, themes, and images) across WordPress instances.
- In the private subnets, Amazon ElastiCache for Memcached nodes for caching database queries.

* The template that deploys the Quick Start into an existing VPC skips the tasks marked by asterisks and prompts you for your existing VPC configuration.



WordPress high availability architecture by Bitnami

A detailed description of deploying WordPress High Availability by Bitnami on AWS is beyond the scope of this document. For configuration and options, refer to [WordPress High Availability by Bitnami on AWS](#).

Conclusion

AWS presents many architecture options for running WordPress. The simplest option is a single server installation for low traffic websites. For more advanced websites, site administrators can add several other options, each one representing an incremental improvement in terms of availability and scalability. Administrators can select the features that most closely match their requirements and their budget.

Contributors

Contributors to this document include:

- Paul Lewis, Solutions Architect, Amazon Web Services
- Ronan Guilfoyle, Solutions Architect, Amazon Web Services
- Andreas Chatzakis, Solutions Architect Manager, Amazon Web Services
- Jibril Touzi, Technical Account Manager, Amazon Web Services
- Hakmin Kim, Migration Partner Solutions Architect, Amazon Web Services

Document revisions

Date	Description
October 19, 2021	Updated to modify Reference Architecture and AWS for WordPress plugin.
October 2019	Updated to include new deployment approaches and AWS for WordPress plugin.
February 2018	Updated to clarify Amazon Aurora product
December 2017	Updated to include AWS services launched since first publication.
December 2014	First publication.

Appendix A: CloudFront configuration

To get optimal performance and efficiency when using Amazon CloudFront with your WordPress website, it's important to configure the website correctly for the different types of content being served.

Origins and behaviors

An [origin](#) is a location where CloudFront sends requests for content that it distributes through the edge locations. Depending on your implementation you can have one or two origins. One for dynamic content (the Lightsail instance in the [single-server deployment option](#), or the Application Load Balancer in the [elastic deployment option](#)) using a custom origin. You may have a second origin to direct CloudFront to for your static content. In the preceding [reference architecture](#), this is an S3 bucket. When you use Amazon S3 as an origin for your distribution, you need to use a [bucket policy](#) to make the content publicly accessible.

[Behaviors](#) enable you to set rules that govern how CloudFront caches your content, and, in turn, determine how effective the cache is. Behaviors enable you to control the protocol and HTTP methods your website is accessible by. They also enable you to control whether to pass HTTP headers, cookies, or query strings to your backend (and, if so, which ones). Behaviors apply to specific URL path patterns.

CloudFront distribution creation

Create a CloudFront web distribution by following the Distribution, the default Origin and Behavior automatically created will be used for dynamic content. Create four additional behaviors to further customize the way both static and dynamic requests are treated. The following table summarizes the configuration properties for the five behaviors. You can also skip this manual configuration and use the AWS for WordPress plugin covered in [Appendix B: Plugins Installation and Configuration](#), which is the easiest way to configure CloudFront to accelerate your WordPress site.

Table 1: Summary of configuration properties for CloudFront behaviors

Property	Static	Dynamic (admin)	Dynamic (front end)
Paths (Behaviors)	wp-content/* wp-includes/*	wp-admin/* wp-login.php	default (*)
Protocols	HTTP and HTTPS	Redirect to HTTPS	HTTP and HTTPS
HTTP methods	GET, HEAD	ALL	ALL
HTTP headers	NONE	ALL	Host CloudFront-Forwarded-Proto CloudFront-Is-Mobile-Viewer CloudFront-Is-Tablet-Viewer CloudFront-Is-Desktop-Viewer
Cookies	NONE	ALL	comment_* wordpress_* wp-settings-*
Query Strings	YES (invalidation)	YES	YES

For the default behavior, AWS recommends the following configuration:

- Allow the Origin Protocol Policy to Match Viewer, so that if viewers connect to CloudFront using HTTPS, CloudFront connects to your origin using HTTPS as well, achieving end-to-end encryption. Note that this requires you install a trusted SSL certificate on the load balancer. For details, refer to [Requiring HTTPS for Communication Between CloudFront and Your Custom Origin](#).
- Allow all HTTP methods since the dynamic portions of the website require both GET and POST requests (for example, to support POST for the comment submission forms).

- Forward only the cookies that vary the WordPress output, for example, `wordpress_*`, `wp-settings-*` and `comment_*`. You must extend that list if you have installed any plugins that depend on other cookies not in the list.
- Forward only the HTTP headers that affect the output of WordPress, for example, `Host`, `CloudFront-Forwarded-Proto`, `CloudFront-is-Desktop-Viewer`, `CloudFront-is-Mobile-Viewer`, and `CloudFront-is-Tablet-Viewer`:
 - `Host` allows multiple WordPress websites to be hosted on the same origin.
 - `CloudFront-Forwarded-Proto` allows different versions of pages to be cached depending on whether they are accessed via HTTP or HTTPS.
 - `CloudFront-is-Desktop-Viewer`, `CloudFront-is-Mobile-Viewer`, and `CloudFront-is-Tablet-Viewer` allow you to customize the output of your themes based on the end user's device type.
- Forward all the query strings to cache based on their values because WordPress relies on these, they can also be used to invalidate cached objects.

If you want to serve your website under a custom domain name (not `*.cloudfront.net`), enter the appropriate URIs under **Alternate Domain Names** in the Distribution Settings. In this case, you also need an SSL certificate for your custom domain name. You can [request](#) SSL certificates via the AWS Certificate Manager and configure them against a CloudFront distribution.

Now, create two more cache behaviors for dynamic content: one for the login page (path pattern: `wp-login.php`) and one for the admin dashboard (path pattern: `wp-admin/*`). These two behaviors have the exact same settings, as follows:

- Enforce a Viewer Protocol Policy of HTTPS Only.
- Allow all HTTP methods.
- Cache based on all HTTP headers.
- Forward all cookies.
- Forward and cache based on all query strings.

The reason behind this configuration is that this section of the website is highly personalized and typically has just a few users, so caching efficiency isn't a primary concern. The focus is to keep the configuration simple to ensure maximum compatibility with any installed plugins by passing all cookies and headers to the origin.

The [AWS for WordPress plugin](#) covered in Appendix B automatically creates a CloudFront distribution that meets the preceding configuration.

By default, WordPress stores everything locally on the web server, which is block storage (Amazon EBS) for single server deployment, and file storage (Amazon EFS) for elastic deployment. In addition to reducing storage and data transfer costs, moving static assets to Amazon S3 offers scalability, data availability, security and performance. There are several plugins that make it easy to move static content to Amazon S3; one of them is [W3 Total Cache](#), also covered in Appendix B.

Appendix B: Plugins installation and configuration

AWS for WordPress plugin

The AWS for WordPress plugin is the only WordPress plugin written and actively maintained by AWS. It enables customers to easily configure [Amazon CloudFront](#) and [AWS Certificate Manager](#) (ACM) to WordPress websites for enhanced performance and security. The plugin uses [Amazon Machine Learning](#) (ML) services to translate content into one or more languages, produces audio versions of each translation, and reads WordPress websites through Amazon Alexa devices.

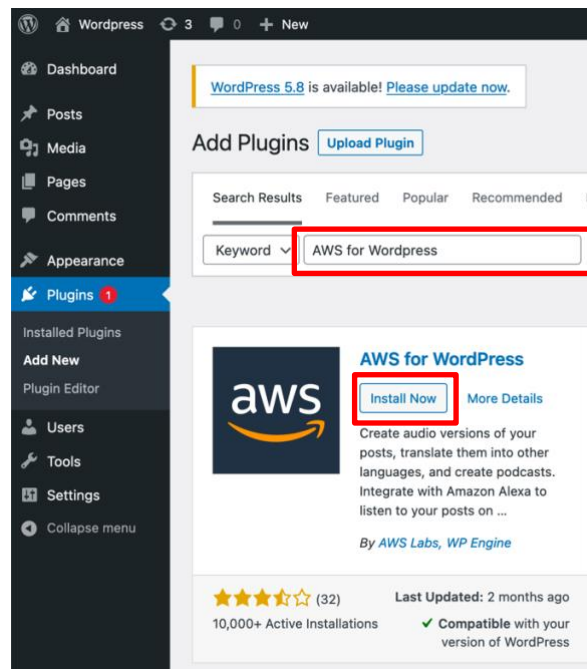
The plugin is installed already in [WordPress High Availability by Bitnami on AWS Quick Start](#).

Plugin installation and configuration

To install the plugin:

1. To use the AWS for WordPress plugin, you must create an *IAM user* for the plugin. An IAM user is a person or application under an AWS account that has permission to make API calls to AWS services.

2. You need an [AWS Identity and Access Management](#) (IAM) role or an IAM user to control authentication and authorization for your AWS account. To prevent unauthorized users from gaining these permissions, protect the IAM user's credentials. Treat the secret access key like a password; store it in a safe place, and don't share it with anyone. Like a password, [rotate the access key periodically](#). If the secret access key is accidentally leaked, [delete it immediately](#). Then you can create a new access key to use with the AWS for WordPress plugin.
3. In the **Plugins** menu of the WordPress admin panel, search **AWS for WordPress** and choose **Install Now**.



4. If the plugin installation is not working, there may be a user permission problem. Connect to WordPress web server and complete the following instructions to solve the issue.
 - a. Open The `wp-config.php` file in the WordPress install directory and write the following code at the end of the `wp-config.php` file:

```
define('FS_METHOD', 'direct');
```

- b. Launch the following command to give writing permission:

```
chmod 777 <WordPress install directory>/wp-content
```

Warning: Keeping the writing permission as 777 is risky. If the permission is kept as 777, anyone can edit or delete this folder. Change the writing permission into 755 or below after completing the plugin work.

- c. If the [reference architecture](#) is used, the WordPress install directory is ``/var/www/wordpress/<site directory>``.

A detailed description of all AWS for WordPress settings is beyond the scope of this document. For configuration and options, refer to [Getting started with the AWS for WordPress plugin](#).

Amazon CloudFront and AWS Certificate Manager

To set up CloudFront and AWS Certificate Manager:

1. On the plugin menu, choose **CloudFront** and enter the following parameters:
 - **Origin domain name:** DNS domain of the HTTP origin server where CloudFront gets your website's content (such as `example.com`).
 - **Alternate domain name (CNAME):** domain name that your visitors use for the accelerated website experience. AWS recommends using 'www' in front of the domain (such as `www.example.com`).
2. Choose **Initiate Setup** to start the configuration.

The plugin automatically requests an SSL certificate for the CNAME via ACM, once you validate the ACM token by [updating the DNS records](#) with the CNAME entries, the plugin will create a CloudFront distribution that meets the best practices defined in [Appendix A](#).

Note: AWS for WordPress plugin requires HTTPS for communication between CloudFront and your custom origin. Make sure your origin has an SSL certificate valid for the Origin domain name. For more information, refer to [Using HTTPS with CloudFront](#).

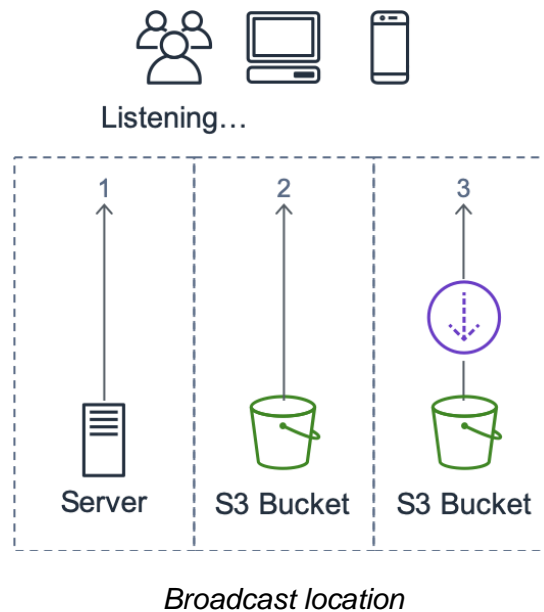
Translate and vocalize your content

The AWS for WordPress plugin enables you to automatically translate text in different languages, and convert the written content into multilingual audio formats. These features are powered by Amazon Machine Learning services.

[Amazon Polly](#) is a service that turns text into lifelike speech. With dozens of voices across a variety of languages, you can select the ideal voice and build engaging speech-enabled applications that work in many different countries. Use the plugin to create audio files in any of the voices and languages supported by Amazon Polly. Your visitors can stream the audio at their convenience using inline audio players and mobile applications.

By default, the plugin stores new audio files on your web server. You can choose to store the files on Amazon S3 or on Amazon CloudFront. Users have the same listening experience regardless of where you store your audio files. Only the broadcast location changes:

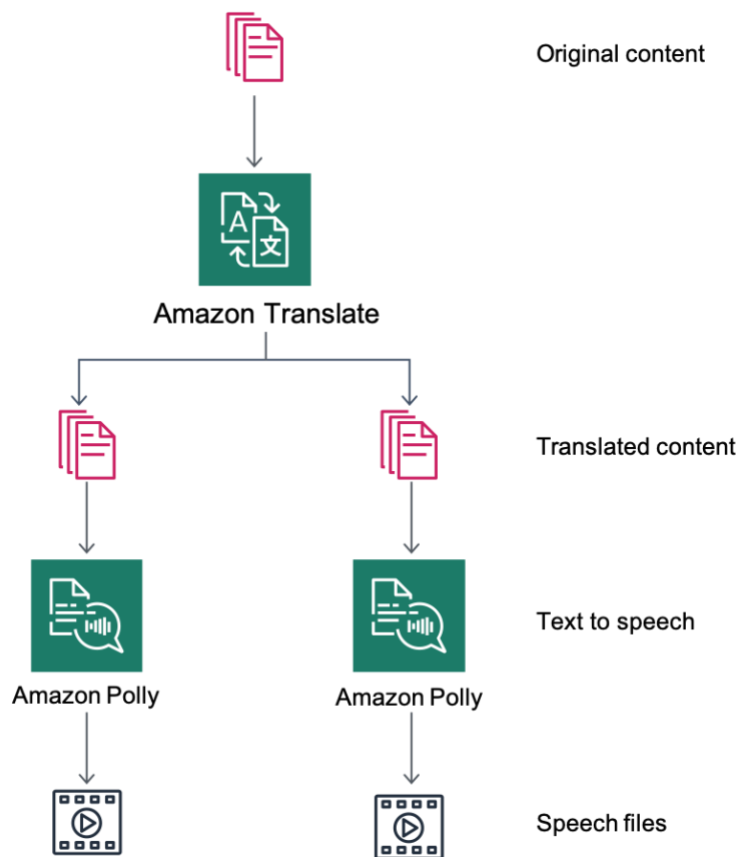
- For audio files stored on the WordPress server, files are broadcast directly from the server.
- For files stored in an S3 bucket, files are broadcast from the bucket.
- If you use CloudFront, the files are stored on Amazon S3 and are broadcast with CloudFront.



[Amazon Translate](#) is a machine translation service that delivers fast, high-quality, and affordable language translation. Providing multilingual content represents a great opportunity for site owners. Although English is the dominant language of the web, native English speakers are a mere 26% of the total online audience.

By offering written and audio versions of your WordPress content in multiple languages, you can meet the needs of a larger international audience. You can configure the plugin to do the following:

- Automatically translate into different languages and create audio recordings of each translation for new content upon publication, or choose to translate and create recordings for individual posts
- Translate into different languages and create audio recordings for each translation of your archived content
- Use the Amazon Pollycast RSS feed to podcast audio content



Overview of content translation and text to speech

Podcasting with Amazon Pollycast

With Amazon Pollycast feeds, your visitors can listen to your audio content using standard podcast applications. RSS 2.0-compliant Pollycast feeds provide the XML data needed to aggregate podcasts by popular mobile podcast applications, such as iTunes, and podcast directories.

When you install the AWS for WordPress plugin, you will find option to enable generation of XML feed in the **Podcast** configuration tab. There, you will also find option to configure multiple optional properties. After enabling the functionality, you will receive a link do the feed.

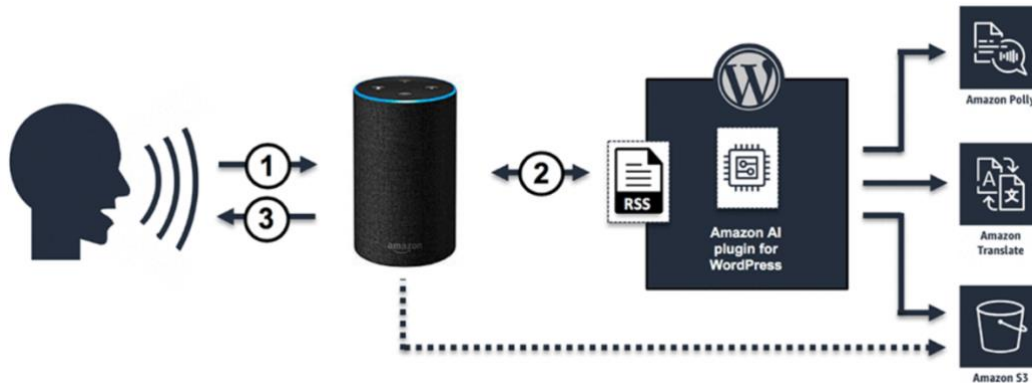
Reading your content through Amazon Alexa devices

You can extend WordPress websites and blogs through Alexa devices. This opens new possibilities for the creators and authors of websites to reach an even broader audience. It also makes it easier for people to listen to their favorite blogs by just asking Alexa to read them.

To expose the WordPress website to Alexa, you must enable:

- AWS for WordPress plugin.
- The text-to-speech and Amazon Pollycast functionalities. These functionalities generate an RSS feed on your WordPress site which is consumed by Amazon Alexa.
- Amazon S3 as the default storage for your files in text-to-speech, it's important that your website uses a secure HTTPS connection to expose its feed to Alexa.

The following diagram presents the flow of interactions and components that are required to expose your website through Alexa.



Flow of interactions required to expose WordPress websites through Alexa

1. The user invokes a new Alexa skill, for example by saying: “Alexa, ask Demo Blog for the latest update”. The skill itself is created using one of the Alexa Skill Blueprints. This enables you to expose your skill through Alexa devices even if you don’t have deep technical knowledge.
2. The Alexa skill analyzes the call and RSS feed that was generated by the AWS for WordPress plugin, and then returns the link to the audio version of the latest article.
3. Based on the link provided by the feed, Alexa reads the article by playing the audio file saved on Amazon S3.

Refer to the [plugin page](#) on WordPress marketplace for a detailed step-by-step guide for installing and configuring the plugin and its functionalities.

Static content configuration

By default, WordPress stores everything locally on the web server, which is block storage (Amazon EBS) for single server deployment, and file storage (Amazon EFS) for elastic deployment. In addition to reducing storage and data transfer costs, moving static asset to Amazon S3 offers scalability, data availability, security and performance.

In this example, the W3 Total Cache (W3TC) plugin is used to store static assets on Amazon S3. However, there are other plugins available with similar capabilities. If you want to use an alternative you can adjust the following steps accordingly. The steps only refer to features or settings relevant to this example. A detailed description of all settings is beyond the scope of this document. Refer to the [W3 Total Cache plugin page](#) at wordpress.org for more information.

IAM user creation

You need to create an IAM user for the WordPress plugin to store static assets in Amazon S3. For instructions, refer to [Creating an IAM User in Your AWS Account](#).

Note: IAM roles provide a better way of managing access to AWS resources, but at the time of writing, the W3 Total Cache plugin does not support [IAM roles](#).

Take a note of the user security credentials and store them in a secure manner – you need these credentials later.

Amazon S3 bucket creation

1. First, create an Amazon S3 bucket in the AWS Region of your choice. For instructions, refer to [Creating a bucket](#). Enable static website hosting for the bucket by following the guide for [Configuring a static website on Amazon S3](#).
2. Create an IAM policy to provide the IAM user created previously access to the specified S3 bucket, and attach the policy to the IAM user. For instructions to create the following policy, refer to [Managing IAM Policies](#).

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "Stmt1389783689000",
      "Effect": "Allow",
      "Principal": "*",
      "Action": [
        "s3:DeleteObject",
        "s3:GetObject",
        "s3:GetObjectAcl",
        "s3:ListBucket",
        "s3:PutObject",
        "s3:PutObjectAcl"
      ],
      "Resource": [
        "arn:aws:s3:::wp-demo",
        "arn:aws:s3:::wp-demo/*"
      ]
    }
  ]
}
```

```
]
}
```

3. Install and activate the W3TC plugin from the WordPress admin panel.
4. Browse to the General Settings section of the plugin's configuration, and ensure that both **Browser Cache** and **CDN** are enabled.
5. From the dropdown list in the CDN configuration, choose **Origin Push: Amazon CloudFront** (this option has Amazon S3 as its origin).
6. Browse to the Browser Cache section of the plugin's configuration and enable the **expires**, **cache control**, and **entity tag (ETag)** headers.
7. Also activate the **Prevent caching of objects after settings change** option so that a new query string is generated and appended to objects whenever any settings are changed.
8. Browse to the CDN section of the plugin's configuration and enter the security credentials of the IAM user you created earlier, as well as the name of the S3 bucket.
9. If you are serving your website via the CloudFront URL, enter the distribution domain name in the relevant box. Otherwise, enter one or more CNAMEs for your custom domain name(s).
10. Finally, export the media library and upload the wp-includes, theme files, and custom files to Amazon S3 using the W3TC plugin. These upload functions are available in the **General** section of the **CDN** configuration page.

Static origin creation

Now that the static files are stored on Amazon S3, go back to the CloudFront configuration in the CloudFront console, and configure Amazon S3 as the origin for static content. To do that, add a second origin pointing to the S3 bucket you created for that purpose. Then create two more cache behaviors, one for each of the two folders (`wp-content` and `wp-includes`) that should use the S3 origin rather than the default origin for dynamic content. Configure both in the same manner:

- Serve HTTP GET requests only.
- Amazon S3 does not vary its output based on cookies or HTTP headers, so you can improve caching efficiency by not forwarding them to the origin via CloudFront.

- Despite the fact that these behaviors serve only static content (which accepts no parameters), you will forward query strings to the origin. This is so that you can use query strings as version identifiers to instantly invalidate, for example, older CSS files when deploying new versions. For more information, refer to the [Amazon CloudFront Developer Guide](#).

Note: After adding the static origin behaviors to your CloudFront distribution, check the order to ensure the behaviors for `wp-admin/*` and `wp-login.php` have higher precedence than the behaviors for static content. Otherwise, you may see strange behavior when accessing your admin panel.

Appendix C: Backup and recovery

Recovering from failure in AWS is faster and easier to do compared to traditional hosting environments. For example, you can launch a replacement instance in minutes in response to a hardware failure, or you can make use of automated failover in many of our managed services to negate the impact of a reboot due to routine maintenance.

However, you still need to ensure you are backing up the right data in order to successfully recover it. To re-establish the availability of a WordPress website, you must be able to recover the following components:

- Operating system (OS) and services installation and configuration (Apache, MySQL, and so on)
- WordPress application code and configuration
- WordPress themes and plugins
- Uploads (for example, media files for posts)
- Database content (posts, comments, and so on.)

AWS provides a variety of methods for backing up and restoring your web application data and assets.

This whitepaper previously discussed making use of Lightsail snapshots to protect all data stored on the instance's local storage. If your WordPress website runs off the Lightsail instance only, regular Lightsail snapshots should be sufficient for you to recover your WordPress website in its entirety. However, you will still lose any changes

applied to your website since the last snapshot was taken if you do restore from a snapshot.

In a multi-server deployment, you need to back up each of the components discussed earlier using different mechanisms. Each component may have a different requirement for backup frequency, for example, the OS and WordPress installation and configuration will change much less frequently than user-generated content and, therefore, can be backed up less frequently without losing data in the event of a recovery.

To back up the OS and services installation and configuration, and the WordPress application code and configuration, you can create an AMI of a properly configured EC2 instance. AMIs can serve two purposes: to act as a backup of instance state, and to act as a template when launching new instances.

To back up the WordPress application code and configuration, you need to make use of AMIs and also Aurora backups.

To back up the WordPress themes and plugins installed on your website, back up the Amazon S3 bucket or the Amazon EFS file system they are stored on.

- For themes and plugins stored in an S3 bucket, you can enable [Cross-Region Replication](#) so that all objects uploaded to your primary bucket are automatically replicated to your backup bucket in another AWS Region. Cross-Region Replication requires that [Versioning](#) is enabled on both your source and destination buckets, which provides you with an additional layer of protection and enables you to revert to a previous version of any given object in your bucket.
- For themes and plugins stored on an EFS file system, you can create an AWS Data Pipeline to copy data from your production EFS file system to another EFS file system, as outlined in the documentation page [Using AWS Backup with Amazon EFS](#). You can also back up an EFS file system using any backup application you are already familiar with.
- To back up user uploads you should follow the steps outlined earlier for backing up the WordPress themes and plugins.

- To back up database content you need to make use of [Aurora backup](#). Aurora backs up your cluster volume automatically and retains restore data for the length of the *backup retention period*. Aurora backups are nearly continuous and incremental so you can quickly restore to any point within the backup retention period. No performance impact or interruption of database service occurs as backup data is being written. You can specify a backup retention period from 1 to 35 days. You can also create [manual database snapshots](#), which persist until you delete them. Manual database snapshots are useful for long-term backups and archiving.

Appendix D: Deploying new plugins and themes

Few websites remain static. In most cases, you will periodically add publicly available WordPress themes and plugins or upgrade to a newer WordPress version. In other cases, you will develop your own custom themes and plugins from scratch.

Any time you are making a structural change to your WordPress installation there is a certain risk of introducing unforeseen problems. At the very least, take a backup of your application code, configuration, and database before applying any significant change (such as installing a new plugin). For websites of business or other value, test those changes in a separate staging environment first. With AWS, it's easy to replicate the configuration of your production environment and run the whole deployment process in a safe manner. After you are done with your tests, you can simply tear down your test environment and stop paying for those resources. Later, this whitepaper discusses some WordPress-specific considerations.

Some plugins write configuration information to the `wp_options` database table (or introduce database schema changes), whereas others create configuration files in the WordPress installation directory. Because we have moved the database and storage to shared platforms, these changes are immediately available to all of your running instances without any further effort on your part.

When deploying new themes in WordPress, a little more effort may be required. If you are only making use of Amazon EFS to store all your WordPress installation files, then new themes will be immediately available to all running instances. However, if you are offloading static content to Amazon S3, you must process a copy of these to the right bucket location. Plugins like W3 Total Cache provide a way for you to manually initiate that task. Alternatively, you could automate this step as part of a build process.

Because theme assets can be cached on CloudFront and at the browser, you need a way to invalidate older versions when you deploy changes. The best way to achieve this is by including some sort of version identifier in your object. This identifier can be a query string with a date-time stamp or a random string. If you use the W3 Total Cache plugin, you can update a media query string that is appended to the URLs of media files.