

Security at the Edge: Core Principles

September 24, 2021

This version has been archived.

For the latest version of this document, visit:

<https://docs.aws.amazon.com/whitepapers/latest/security-at-the-edge/security-at-the-edge.html>



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.

This version has been archived.

For the latest version of this document, visit:

<https://docs.aws.amazon.com/whitepapers/latest/security-at-the-edge/security-at-the-edge.html>

Contents

Introduction to edge computing.....	1
Security at the edge.....	2
Secure content delivery	3
Network and application layer protection	4
DDoS mitigation	5
Foundational security principles and best practices.....	7
Compliance and the Shared Responsibility Model	9
Examples of additional security-related edge services	10
Internet of Things	10
Hybrid services at the edge	11
Rugged and disconnected edge.....	12
Conclusion	13
Appendix: AWS services for edge computing.....	14
Edge services.....	14
AWS services that support security at the edge	16
Further reading.....	20
Contributors	20
Document revisions.....	20

This version has been archived.

For the latest version of this document, visit:

<https://docs.aws.amazon.com/whitepapers/latest/security-at-the-edge/security-at-the-edge.html>

Abstract

Today's business leaders know that it is critical to ensure the security of their environments, and the security present in traditional cloud networks is extended to workloads at the edge. This whitepaper provides security executives the foundations for implementing a defense in depth security strategy at the edge by addressing three areas of security at the edge:

- AWS services at AWS edge locations
- How those services and others can be used to implement the best practices outlined in the design principles of the AWS Well-Architected Framework Security Pillar
- The security aspects of additional AWS edge services, which customers can use to help secure their edge environments or expand operations into new, previously unsupported environments

Together, these elements offer core principles for designing a security strategy at the edge, and demonstrate how AWS services can provide a secure environment extending from the core cloud to the edge of the AWS network and out to customer edge devices and endpoints.

This version has been archived.

For the latest version of this document, visit:

<https://docs.aws.amazon.com/whitepapers/latest/security-at-the-edge/security-at-the-edge.html>

Introduction to edge computing

Security is the top priority at AWS. The high security bar set by AWS services covers customers as they expand their use of AWS services to bring workloads out to the edge to use its growing number of capabilities and applications.

Edge computing comprises elements of geography and networking, and brings computing closer to the user. Edge takes place at or near the physical location of either the user or the source of the data. By placing computing services close to these locations, the user benefits from faster, more reliable services.

This paper discusses AWS services that are available to provide a secure environment, from the core cloud to the edge of the AWS network, and out to customer edge devices and endpoints. Many of the AWS services that provide security capabilities to the edge reside at AWS edge locations, or as close to customers' edge devices and endpoints as necessary. AWS edge locations are a worldwide network of data centers that run with AWS at physical locations directly connected to the expanding AWS global infrastructure.

AWS edge services provide infrastructure and software that deliver data processing, analysis, and storage as close to the endpoint as necessary. This includes deploying AWS Managed Services, APIs, and tools to locations outside AWS data centers, and even onto customer-owned infrastructure and devices. AWS enables customers to build high-performance applications that rely on the cloud for data processing and storage, but also need to process or store some data close to where it is generated to deliver ultra-low latency, intelligent, real-time responsiveness, and reduce the amount of data transfer.

This version has been archived.

Every AWS customer is unique, and “edge” can mean something different to different customers. Edge use cases and technology can range from autonomous vehicles, medical devices, and sensors, to user interfaces, retail, and other edge devices. Mobile phones and robot vacuums are also examples of edge devices.

The objective of AWS edge services is to provide consistent capabilities and customer experience from the edge to the cloud. AWS uses the same programming model for the cloud, on-premises infrastructure, and local devices. This gives you the choice of centralized control or de-centralized control, with decentralized implementation. You have access to the same environment to develop, connect, deploy, manage, and secure with the same tools, regardless of where your workloads are located.

Security at the edge

AWS provides services and features you can use to help you create secure architectures, workloads, and services to elevate your security from edge to cloud. Security at AWS starts with core infrastructure, which is built for the cloud and designed to meet the most stringent security requirements in the world. For example, all data flowing across the AWS global network that interconnects data centers and [Regions](#) is automatically encrypted at the physical layer before it leaves AWS secured facilities.

At the edge, AWS offers services that address the different aspects of edge security, including preventive security mechanisms like encryption and access control, continuous monitoring mechanisms like configuration auditing, and physical security like tamper-evident enclosures. Customers that need to store and process data on premises, or in countries where there is no AWS Region, can do so securely with AWS edge services. This capability can help you comply with data handling or data residency requirements.

AWS Cloud security principles are fundamental and apply regardless of where an organization operates. These principles are discussed in detail in a later section of this whitepaper. AWS offerings combine a high security bar with agility to adapt rapidly as needed. AWS customers working at the edge have access to over 200 fully featured, integrated cloud and device services, many of which have specific edge capabilities.

AWS services with Points of Presence (PoP) at edge locations — globally scaled and connected through the AWS network backbone — provide a more secure, performant, and available experience. AWS also offers services that run on the edge, which enable you to deliver content. AWS edge services, which provide infrastructure and software that deliver data processing, analysis, and storage at endpoints comprise a comprehensive set of cloud services that support the secure deployment and management of edge devices.

For the latest version of this document, visit:
Security at the edge has the same principles as cloud security. By extending cloud services to the edge, AWS gives you a way to operate safely, with strong security infrastructure and safeguards. AWS-owned infrastructure is monitored 24/7 to help safeguard the confidentiality, integrity, and availability of our customers' data. Moving cloud workloads to edge devices or endpoints provides you with more control and visibility, and mitigates risk.

<https://docs.aws.amazon.com/whitepapers/latest/security-at-the-edge/security-at-the-edge.html>

Media and entertainment at the edge

The media and entertainment industry provides natural examples of customers who need to focus on securing their content delivery at the edge. For example, [Amazon CloudFront](#) provides streaming services with low latency, sustained high throughput, lower rebuffering rates, and integration with other AWS services, while also securely distributing content globally. For more information, see [Amazon CloudFront for Media & Entertainment](#).

A defense in depth model (for example, using multiple independent layers of specialized security controls) provides layers of protection. In addition to the design principles of the [AWS Well-Architected Framework's Security Pillar](#), this paper highlights three aspects of edge protection whose PoP is at AWS edge locations. The three highlighted edge protections that help secure the connection points between the origin infrastructure, edge services, and customer edge devices or applications are:

- Secure content delivery
- Network and application layer protection
- Distributed Denial of Service (DDoS) mitigation

The design principles also cover the security of edge devices and applications. A comprehensive defense in depth strategy should include services that account for the security of both AWS edge locations, and edge devices and applications.

Secure content delivery

Secure content delivery provides content, such as data, videos, applications, and APIs, quickly and securely to customers. These should be delivered over secure transport, using the recommended version of Transport Layer Security (TLS) to encrypt communications between endpoints. If necessary, there are a number of methods that you can use to help secure that same content through restricted access, including signed URLs, signed cookies, and token authentication.

<https://docs.aws.amazon.com/whitepapers/latest/security-at-the-edge/security-at-the-edge.html>
[Amazon CloudFront](#), a global content delivery network (CDN) service that securely delivers data, videos, applications, and APIs to viewers with low latency and high transfer speeds, addresses these areas of security when it is deployed at AWS edge locations.

To create a more secure CDN, organizations can gain protection against L3/L4 DDoS attacks using [AWS Shield](#). AWS also offers AWS Shield Advanced, which provides

additional detection and mitigation against large and sophisticated DDoS attacks, near-real-time visibility into attacks, and integration with [AWS WAF](#), a web application firewall service, to protect against application layer (L7) attacks. Together, these services create a flexible, layered security perimeter.

CloudFront offers security capabilities, including field-level encryption and HTTPS support, seamlessly running with AWS Shield Advanced, AWS WAF, and [Amazon Route 53](#) to protect against multiple types of attacks, including network and application layer DDoS attacks. For more details about CloudFront and Route 53, see the [Appendix](#).

Network and application layer protection

Edge networks are architected outside of the security perimeters of traditional cloud. Extending security to edge end devices requires network and application security and continuous monitoring, as well as encryption of data in transit and at rest.

Edge customers should define trust boundaries for networks and accounts, and verify secure system configurations and other policy-enforcement points, including web application firewalls (WAFs) and API gateways. This can be done by blocking well-known exploits, implementing protections specific to applications, responding to new threats, and performing ongoing monitoring.

There are two important aspects to network and application layer protection at the edge:

- Protections from well-known exploits and attacks that could affect an organization's applications
- Visibility and control of workloads

This version has been archived.

Manufacturing at the edge

Edge computing offers manufacturers opportunities to collect, process, and analyze data to enable predictive maintenance, improve quality control, and enhance worker safety with near-real-time alerts, industrial robot fleet management, and simulation. Although these edge applications can increase efficiency and keep costs down, they should be protected against security events. AWS WAF provides security rules to help protect these edge applications against common security attacks. AWS Shield Advanced helps protect against DDoS attacks.

A WAF deployed at AWS edge locations can help to set fundamental protections, customize them to the applications, and help organizations quickly visualize actions so

they can create a dynamic security posture. With AWS WAF, you can use the AWS pre-configured rules (Managed Rules), use Marketplace Rules, or create your own custom rules to protect against common attack vectors. [AWS Managed Rules](#) give you protection against common web application attacks. They are curated by multiple points of intelligence across multiple sources within AWS.

Marketplace Rules are written, updated, and managed by third-party security experts, and can be used on their own or in conjunction with AWS Managed Rules. AWS WAF, which integrates with AWS Shield Advanced at no extra cost, provides easy setup, low operation overhead, minimal latency impact, and customizable security. It also uses advanced automation to analyze web logs, identify malicious requests, and automatically update security rules.

In addition to preventing incidents, visibility into traffic coming into and out of a network is a second key aspect of network and application layer protection. There are multiple options available to get insights and metrics: [CloudWatch metrics](#), sampled web requests, and logs.

With CloudWatch, you can monitor web requests and web access control lists (ACLs) and rules. CloudWatch collects and processes raw data from AWS WAF and Shield Advanced into readable, near-real-time metrics. AWS WAF supports full logging of all web requests inspected by the service, which can then be stored in the cloud for compliance and auditing purposes, and used for debugging and additional forensics. You can also integrate the logs with your security information and event management (SIEM) and log analysis tools¹.

For more details about AWS WAF, see the [Appendix](#).

DDoS mitigation

DDoS mitigation as a defense layer is important for organizations operating at the edge with mission-critical operations that cannot afford downtime. DDoS mitigation helps ensure continued availability of those operations and services. DDoS attacks are deliberate attempts to exhaust infrastructure or application resources so they are unavailable to users. Common types of DDoS attacks are [SYN floods](#) that exploit the TCP protocol, [reflected or amplified attacks](#) that use the public IP address of User Datagram Protocol (UDP) for the source, and [HTTP floods](#) that target web servers' capacity to manage requests.

AWS services include basic DDoS protection as a standard feature. All AWS customers using CloudFront, Application Load Balancers, Network Load Balancers, Global

Accelerators, Elastic IPs, or Route 53 receive basic DDoS protection against common network and transport layer attacks.

This protection is always on, but is preconfigured, static, and provides no reporting or analytics. Mitigations are configured with pre-assigned limits based on the service that is being targeted. For example, if your [Elastic Load Balancing](#) (ELB) is targeted by an infrastructure layer DDoS attack, a mitigation that is configured based on ELB service limits to ensure the resource remains operational is placed. This mitigation is effective at blocking many known vectors of attack, and protecting the underlying resource. The nuance is that the limits of your application may differ from the limits of the ELB, resulting in the resource remaining operational, but your application still being impacted.

Shield Advanced is a managed service that builds a customized DDoS protection capability specifically for your applications' needs, based on the resources you specified either in Shield Advanced or through an AWS Firewall Manager Shield Advanced policy. Shield Advanced can be deployed at AWS edge locations, and you get tailored detection based on the specific traffic patterns of your application, protection against Layer 7 DDoS attacks at no additional cost, access to 24x7 specialized support from the Shield Response Team (SRT), centralized management of security policies through [AWS Firewall Manager](#), and cost protection to safeguard against scaling charges resulting from DDoS-related usage spikes. You can also configure AWS WAF to integrate with Shield Advanced to create custom rules.

Some DDoS events can be mitigated by scaling applications to absorb the additional traffic or by using a web application firewall.² Unless encrypted traffic ends with a network-layer device, these devices are generally unable to inspect encrypted requests. This can allow bad actors to use expansive web requests or large volumes of web requests to generate a flood that is challenging to fingerprint, challenging to block or absorb, or both.

Using Amazon CloudFront or [AWS Global Accelerator](#) to distribute request handling across many AWS edge locations and AWS IP addresses that exceed a pre-defined limit can help secure applications targeted by this type of DDoS attack. These events are detected when an Amazon CloudFront distribution or [Application Load Balancer](#) (ALB) is protected by AWS Shield Advanced.³

This version has been archived.
For the latest version of this document, visit:
<https://docs.aws.amazon.com/whitepapers/latest/security-at-the-edge/security-at-the-edge.html>

Foundational security principles and best practices

The services described in the previous section are part of a larger set of best practices outlined in the [design principles](#) of the AWS Well-Architected Framework Security Pillar to secure the edge devices and applications themselves. Together, implementation of these services provides a strong defense in depth strategy to secure edge devices and applications.

- **Implement a strong identity foundation** — Implement the principle of least privilege and enforce separation of duties with appropriate authorization for each interaction with AWS resources. Centralize identity management, and aim to eliminate reliance on long-term static credentials. [AWS Identity and Access Management](#) (AWS IAM) and [Amazon Cognito](#) help configure secure access to edge applications.

AWS IAM provides a seamless process for multi-layer security and identity and access management, either using pre-configured policies or customized policies. IAM Roles and Permissions can be used to limit who can make changes to your network environment, such as CloudFront, AWS WAF, and Route 53. Amazon Cognito enables the addition of user sign up/sign in supporting multi-factor authentication and data encryption. Roles can be defined and users mapped to give applications access to the exact resources authorized for each user. AWS edge services can be configured to verify that only traffic from CloudFront reaches your infrastructure.

- **Enable traceability** — Monitor, alert, and audit actions and changes to an environment in real time. Integrate log and metric collection with systems to automatically investigate and take action. Available and integrated services such as AWS IoT Device Defender continuously audit edge Internet of Things (IoT) configurations to verify latest security best practices, while CloudTrail logs and monitors account activity across your AWS infrastructure.

[Amazon GuardDuty](#) monitors for malicious activity and unauthorized behavior to provide threat detection. [AWS Config](#) enables you to assess, audit, and evaluate the configurations of your AWS resources.

[AWS WAF](#) has full logging of all web requests it inspects. With this feature, logs are stored in S3. You can integrate the logs with SIEM and log analysis tools.

With [Amazon CloudFront](#), you can log the requests that come to your CloudFront distributions, or log the CloudFront service activity in your AWS account. Use [Amazon EventBridge](#) to centralize events into a location for processing, providing more traceability.

- **Apply security at all layers** — Apply a [defense in depth](#)⁴ approach with multiple security controls. Apply to all layers (for example, edge of network, VPC, load balancing, every instance and compute service, operating system, application, and code). At the edge, secure content delivery, application and network protection, and DDoS mitigation strategies cover Layers 3, 4, and 7 for a comprehensive approach.
- **Automate security best practices** — Automation is one of the key security benefits of the cloud, and those benefits extend out to the edge. Automated software-based security mechanisms improve the ability to securely scale more rapidly and cost-effectively.

Create secure architectures, including the implementation of controls that are defined and managed as code in version-controlled templates. AWS WAF, for example, can be completely administered through APIs that make security automation easier, enabling rapid rule propagation and fast incident response.

The AWS WAF Security Automations solution uses AWS CloudFormation to automatically deploy a set of AWS WAF rules designed to filter common web-based attacks. Users can select from preconfigured protective features that define the rules included in an AWS WAF web access control list (web ACL). After the solution deploys, AWS WAF begins inspecting web requests to the user's existing Amazon CloudFront distributions, Amazon ElastiCache Application Load Balancers, and blocks them when applicable⁵.

[AWS Firewall Manager](#) also has automation features, such as automatically enforcing inbound security policies on existing and newly created resources, and automatically protecting against various types of DDoS attacks such as [UDP reflection attacks](#), SYN flood, [DNS query flood](#), and [HTTP flood](#) attacks across accounts.

<https://docs.aws.amazon.com/whitepapers/latest/security-at-the-edge/security-at-the-edge.html>
Firewall Manager enables you to generate policies that ensure any new resources created by developers automatically have the correct Security Group, AWS Shield, or AWS WAF policies. Firewall Manager can integrate with [AWS Security Hub](#), allowing for a centralized view of the policies that apply to your AWS environment.

- **Protect data in transit and at-rest** — Classify data into sensitivity levels and use mechanisms, such as encryption, tokenization, and access control, where appropriate. AWS encryption solutions such as [AWS Key Management Services](#) (KMS) help keep data at the edge encrypted at rest and in motion. AWS KMS enables you to easily create and manage cryptographic keys and control their use in applications across services. [AWS Certificate Manager](#) provisions, manages, and deploys SSL/TLS certificates at the edge for use with AWS services and internal resources. [AWS CloudHSM](#) is useful for managing encryption keys.
- **Keep people away from data** — Use mechanisms and tools to reduce or eliminate the need for direct access or manual processing of data. This reduces the risk of mishandling or modification and human error when handling sensitive data. Machine learning (ML) tools and services, such as [Amazon Macie](#), automate discovery of sensitive data, providing constant visibility into the security of your data and lowering the cost of protecting data.
- **Prepare for security events** — Prepare for an incident by having incident management and investigation policy, and processes that align to organizational requirements. Run incident response simulations and use tools with automation to increase speed for detection, investigation, and recovery.

[Amazon Detective](#) automatically collects log data from your AWS resources, including AWS CloudTrail and Amazon GuardDuty, and uses ML, statistical analysis, and graph theory to build a linked set of data that enables you to easily conduct faster and more efficient security investigations.

This version has been archived.
 AWS Shield Advanced enables proactive engagement from the Security Response Team (SRT) when a DDoS event is detected. With proactive engagement, the SRT will directly contact you if an [Amazon Route 53](#) health check associated with your protected resource becomes unhealthy during an event that is detected by Shield Advanced. Use [Security Hub](#) to collect security data from across AWS accounts, services, and supported third-party products, analyze security trends, and identify the highest priority security issues.

<https://docs.aws.amazon.com/whitepapers/latest/>

Compliance and the Shared Responsibility Model

The [Shared Responsibility Model](#) (SRM) is an important concept applied to the relationship and security responsibilities between AWS and its customers. AWS provides and protects the foundational hardware, infrastructure, and software aspects of the cloud – including the edge – that customers build their applications on. However, it

is the customer's responsibility to choose how they architect their applications on AWS, and how they choose to expose those applications to the internet.⁶ Security at the edge focuses on a defense in depth strategy.

Customers always control their data, including encryption, storage, movement, and retention. AWS services can help customers guard identity and access, protect data, secure applications, and meet their compliance objectives.

Similar to its high bar for security, AWS has high standards for compliance that extend to the edge. AWS regularly achieves third-party validation for thousands of global compliance requirements that are continually monitored to help customers meet standards for finance, retail, healthcare, government, and beyond.

AWS customers also receive access to tools they can use to reduce the cost and time to run their own specific security assurance requirements. For example, some of the AWS key edge services, such as CloudFront, have many security standards and certifications, including PCI DSS Level 1, HIPAA, FedRamp, ISO 9001, 27001, 27017, 2701.

Examples of additional security-related edge services

AWS offers a range of other security-related edge services, which customers can use to help secure their individual edge environments. These include IoT and hybrid services, as well as services that can be used at the rugged and disconnected edge.

Internet of Things **This version has been archived.**

If edge takes computing closer to where the data is generated, AWS IoT services enable the user to enable devices to take actions, aggregate data, and filter it locally on the device.⁷ AWS IoT offers integrated edge services for all layers of security, including preventive security mechanisms, like encryption and access control to device data, and a service that continuously monitors and audits configurations through [AWS IoT Device Defender](#). **For the latest version of this document, visit:** <https://docs.aws.amazon.com/whitepapers/latest/security-at-the-edge/security-at-the-edge.html>

Other IoT services support customers to connect their devices and operate them at the edge. For example, [AWS IoT Greengrass](#) seamlessly extends AWS to edge devices so they can act locally on the data they generate, while still using the cloud for management, analytics, and durable storage. AWS IoT Greengrass is an IoT open-

source edge runtime and cloud service that helps build, deploy, and manage device software.

AWS IoT Greengrass authenticates and encrypts device data for both local and cloud communications, so that data is never exchanged between devices and the cloud without proven identity. Another example is [FreeRTOS](#). FreeRTOS is an open-source, real-time operating system for microcontrollers that makes small, low-power edge devices easy to program, deploy, secure, connect, and manage. FreeRTOS provides the kernel to run low-power devices as well as software libraries that make it easy to connect securely to the cloud or other edge devices, so you can collect data for IoT applications and take action.

FreeRTOS includes support for Transport Layer Security (TLS v1.2) and PKCS #11 to help your devices connect securely to AWS. FreeRTOS also includes an over-the-air (OTA) update library to remotely update devices with feature enhancements or security patches and a code signing feature to ensure your device code is not compromised during deployment and OTA updates.

AWS Outposts: Security, low latency, and data residency

Sometimes data is required to remain in a specific geographical location for regulatory, contractual, or security reasons. Additionally, some industries, such as financial services, require business applications with single digit millisecond latencies. Customers in the financial services industry use [Outposts](#) to deliver high-frequency trading, banking, payments processing, and risk management services while meeting data locality requirements.

This version has been archived.

Hybrid services at the edge

- [AWS Outposts](#) — While on your cloud adoption journey, you may find that certain workloads are better suited for on-premises management, whether for lower latency or other local processing needs, and require a hybrid cloud approach. For these workloads, AWS Outposts extends AWS infrastructure and services to your environments. This enables you to support workloads, including sensitive workloads, which need to remain on-premises, while leveraging the security and operational capabilities of commercial cloud services.⁸ With Outposts, you can control where your workloads run and where your data resides, while using local operational tooling for things like monitoring and stability.

- **[AWS Wavelength](#)** — AWS Wavelength is an AWS Infrastructure offering which minimizes latency. AWS Wavelength enables developers to build applications that deliver single-digit millisecond latencies to mobile devices and end users. AWS developers can deploy their applications to [Wavelength Zones](#), AWS infrastructure deployments that embed AWS compute and storage services within the telecommunications providers' data centers at the edge of the 5G networks, and seamlessly access the breadth of AWS services in the Region. This enables developers to deliver applications that require single-digit millisecond latencies, such as game and live video streaming, ML inference at the edge, and augmented reality/virtual reality (AR/VR).
- **[AWS Storage Gateway](#)** — For customers in hybrid environments, AWS Storage Gateway seamlessly connects and extends on-premises applications to AWS Cloud storage, caching data locally for low-latency access and optimizing data transfers to AWS.

By integrating with AWS services such as Amazon CloudWatch, Storage Gateway enables secure access to AWS services, easy management and monitoring, and tracking of user activity on AWS resources. Customers with data in the cloud can distribute the data to multiple edge locations, or capture data from multiple edge locations, perform in-cloud processing and analytics, and provide access to endpoints in distributed locations.

Customers with hybrid environments can use [AWS Direct Connect](#), a VPN, or the public internet to connect their on-premises environment to the core AWS Cloud.

Rugged and disconnected edge

The edge is continually expanding, even into austere environments without data centers, and in locations without consistent network connectivity. These environments are called the rugged and disconnected edge. For customers running workloads at the rugged and disconnected edge, there is the [AWS Snow Family](#). The AWS Snow Family is comprised of a number of highly secure, portable devices and capacity points, most with built-in computing capabilities, which help you run operations. These services help physically transport up to exabytes of data into and out of AWS. The Snow Family devices are owned and managed by AWS, and may be used with AWS security, monitoring, storage management, and computing capabilities.

Process data locally with AWS Snowball

[AWS Snowball Edge](#) computing applications enable you to collect and process data that is continuously generated by sensors or machines in hospitals, factory floors, or other edge locations, before transferring the data back to AWS. For example, by using tamper-evident enclosures, encryption, and other methods designed to ensure full chain of custody for your data, Snowball can provide a secure path for health customers to migrate their HIPAA-compliant data to the cloud, where they can centrally manage the configuration and operation of Snowball devices deployed across worldwide customers and organizations.

For more details about the services mentioned in this section, see the [Appendix](#).

Conclusion

As AWS customers continue moving workloads to the edge to gain the benefits of the edge's low latency and take advantage of the resiliency and scalability of the cloud, security remains a top priority. By extending cloud services to the edge, AWS provides you with the tools to operate safely utilizing a defense in depth strategy.

The strategy focuses on implementing the best practices outlined in the design principles of the [AWS Well-Architected Framework Security Pillar's design principles](#) to secure edge devices and applications, and to create an integrated, layered security perimeter at AWS edge locations to secure content delivery, protect the application layer, and mitigate DDoS attacks.

Regardless of your environment — in the cloud, at the edge, or hybrid — basic security principles remain the same. Applying the best practices found in the design principles of the AWS Well-Architected Framework Security Pillar helps set a firm security foundation for your network.

For the latest version of this document, visit:

<https://docs.aws.amazon.com/whitepapers/latest/security-at-the-edge/security-at-the-edge.html>

Appendix: AWS services for edge computing

This appendix provides additional information on the AWS services described in the paper. Further detail can also be found at the AWS websites provided below.

Edge services

Amazon CloudFront

[Amazon CloudFront](#) is a fast content delivery network (CDN) service that securely delivers data, videos, applications, and APIs to customers globally with low latency, high transfer speeds, all within a developer-friendly environment.

Amazon CloudFront offers the most advanced security capabilities, including field level encryption and HTTPS support, seamlessly running with [AWS Shield](#), [AWS Web Application Firewall](#), and [Amazon Route 53](#) to protect against multiple types of attacks, including network and application layer DDoS attacks. These services co-reside at AWS edge locations – globally scaled and connected via the AWS network backbone – providing a more secure, performant, and available experience for your users.

FreeRTOS

[FreeRTOS](#) is an open source, real-time operating system for microcontrollers that makes small, low-power edge devices easy to program, deploy, secure, connect, and manage. Distributed freely under the MIT open-source license, FreeRTOS includes a kernel and a growing set of software libraries suitable for use across industry sectors and applications. This includes security, connectivity, and power management capabilities for low-power devices to AWS Cloud services like [AWS IoT Core](#), or to more powerful edge devices running [AWS IoT Greengrass](#).

For the latest version of this document, visit:
AWS IoT Greengrass

[AWS IoT Greengrass](#) is an IoT open-source edge runtime and cloud service that helps you build, deploy, and manage device software. Customers use AWS IoT Greengrass for their IoT applications on millions of devices in homes, factories, vehicles, and businesses. You can program your devices to act locally on the data they generate, make predictions based on ML models, filter and aggregate device data, and transmit only necessary information to the cloud.

AWS Snow Family

The [AWS Snow Family](#) helps customers that need to run operations in austere, non-data center environments, and in locations where there's a lack of consistent network connectivity. The AWS Snow Family, comprised of [AWS Snowcone](#), [AWS Snowball](#), and [AWS Snowmobile](#), offers a number of physical devices and capacity points, most with built-in computing capabilities. These services help physically transport up to exabytes of data into and out of AWS. AWS Snow Family devices are owned and managed by AWS and integrate with AWS security, monitoring, storage management, and computing capabilities.

AWS Storage Gateway

[AWS Storage Gateway](#) is a hybrid cloud storage service that gives you on-premises access to virtually unlimited cloud storage. Customers use AWS Storage Gateway to simplify storage management and reduce costs for key hybrid cloud storage use cases. These include moving backups to the cloud, using on-premises file shares backed by cloud storage, and providing low latency access to data in AWS for on-premises applications.

AWS Outposts

[AWS Outposts](#) is a fully managed service that offers the same AWS infrastructure, AWS services, APIs, and tools to virtually any data center, co-location space, or on-premises facility for a truly consistent hybrid experience. AWS Outposts is ideal for workloads that require low latency access to on-premises systems, local data processing, data residency, and migration of applications with local system interdependencies.

This version has been archived.

AWS Wavelength

[AWS Wavelength](#) is an AWS Infrastructure offering optimized for mobile edge computing applications. [Wavelength Zones](#) are AWS infrastructure deployments that embed AWS compute and storage services within communications service providers' (CSP) data centers at the edge of the 5G network, so application traffic from 5G devices can reach application servers running in Wavelength Zones without leaving the telecommunications network.

For the latest version of this document, visit:

<https://docs.aws.amazon.com/whitepapers/latest/security-at-the-edge/security-at-the-edge.html>

AWS services that support security at the edge

Amazon API Gateway

[Amazon API Gateway](#) is a fully managed service that makes it easy for developers to publish, maintain, monitor, secure, and operate APIs at any scale. It's a pay-as-you-go service that takes care of all of the undifferentiated heavy lifting involved in securely and reliably running APIs at scale.

AWS Certificate Manager

[AWS Certificate Manager](#) enables you to easily provision, manage, and deploy public and private Secure Sockets Layer/Transport Layer Security (SSL/TLS) certificates for use with AWS services and your internal connected resources. SSL/TLS certificates are used to secure network communications and establish the identity of websites over the internet as well as resources on private networks. AWS Certificate Manager removes the time-consuming manual process of purchasing, uploading, and renewing SSL/TLS certificates.

AWS CloudHSM

[AWS CloudHSM](#) is a cloud-based hardware security module (HSM) that enables you to easily generate and use your own encryption keys on the AWS Cloud. With CloudHSM, you can manage your own encryption keys using FIPS 140-2 Level 3 validated HSMs.

AWS CloudTrail

This version has been archived.
For the latest version of this document, visit:
<https://docs.aws.amazon.com/whitepapers/latest/security-at-the-edge/security-at-the-edge.html>

[AWS CloudTrail](#) enables governance, compliance, operational auditing, and risk auditing of your AWS account. With AWS CloudTrail, you can log, continuously monitor, and retain account activity related to actions across your AWS infrastructure. AWS CloudTrail provides event history of your AWS account activity, including actions taken through the [AWS Management Console](#), [AWS SDKs](#), command line tools, and other AWS services. This event history simplifies security analysis, resource change tracking, and troubleshooting. In addition, you can use AWS CloudTrail to detect unusual activity in your AWS accounts. These capabilities help simplify operational analysis and troubleshooting.

Amazon Cognito

[Amazon Cognito](#) is an access control service that enables you to add user sign-up, sign-in, and access control to your web and mobile apps quickly and easily. Amazon

Cognito scales to millions of users and supports sign-in with social identity providers, such as Apple, Facebook, Google, and Amazon, and enterprise identity providers via SAML 2.0 and OpenID Connect.

AWS Config

[AWS Config](#) enables you to assess, audit, and evaluate the configurations of your AWS resources. AWS Config continuously monitors and records your AWS resource configurations and allows you to automate the evaluation of recorded configurations against desired configurations.

Amazon Detective

[Amazon Detective](#) makes it easy to analyze, investigate, and quickly identify the root cause of potential security issues or suspicious activities. Amazon Detective automatically collects log data from your AWS resources and uses ML, statistical analysis, and graph theory to build a linked set of data that enables you to easily conduct faster and more efficient security investigations.

AWS Firewall Manager

[AWS Firewall Manager](#) is a security management service that enables you to centrally configure and manage firewall rules across your accounts and applications in [AWS Organizations](#). As new applications are created, Firewall Manager makes it easy to bring new applications and resources into compliance by enforcing a common set of security rules. This single service can build firewall rules, create security policies, and enforce them in a consistent, hierarchical manner across your entire infrastructure, from a central administrator account.

AWS Global Accelerator

[AWS Global Accelerator](#) is a networking service that sends your user's traffic through the AWS global network infrastructure, improving your internet user performance by up to 60%. When the internet is congested, the AWS Global Accelerator automatic routing optimizations help keep your packet loss, jitter, and latency consistently low.

For the latest version of this document, visit:
<https://docs.aws.amazon.com/whitepapers/latest/security-at-the-edge/security-at-the-edge.html>

Amazon GuardDuty

[Amazon GuardDuty](#) is a threat detection service that continuously monitors for malicious activity and unauthorized behavior to protect your AWS accounts, workloads, and data stored in S3. With the cloud, the collection and aggregation of account and

network activities is simplified, but it can be time consuming for security teams to continuously analyze event log data for potential threats. With GuardDuty, you now have an intelligent and cost-effective option for continuous threat detection in AWS. The service uses ML, anomaly detection, and integrated threat intelligence to identify and prioritize potential threats.

AWS Identity and Access Management

[AWS Identity and Access Management](#) (IAM) enables you to manage access to AWS services and resources securely. Using AWS IAM, you can create and manage AWS users and groups, and use permissions to allow and deny their access to AWS resources. AWS IAM is a feature of your AWS account offered at no additional charge. You will be charged only for use of other AWS services by your users.

AWS Key Management Service

[AWS Key Management Service](#) (KMS) makes it easy for you to create and manage cryptographic keys and control their use across a wide range of AWS services and in your applications. AWS KMS is a secure and resilient service that uses hardware security modules that have been validated under FIPS 140-2, or are in the process of being validated, to protect your keys. AWS KMS runs with [AWS CloudTrail](#) to provide you with logs of all key usage to help meet your regulatory and compliance needs.

Amazon Macie

[Amazon Macie](#) is a fully managed data security and data privacy service that uses ML and pattern matching to discover and protect your sensitive data in AWS.

This version has been archived.

AWS Network Firewall

[AWS Network Firewall](#) is a managed service that makes it easy to deploy essential network protections for all of your [Amazon Virtual Private Clouds](#) (VPCs). The service can be set up with just a few clicks and scales automatically with your network traffic, so you don't have to worry about deploying and managing any infrastructure.

For the latest version of this document, visit:
<https://docs.aws.amazon.com/whitepapers/latest/security-at-the-edge/security-at-the-edge.html>

The AWS Network Firewall flexible rules engine enables you to define firewall rules that give you fine-grained control over network traffic, such as blocking outbound Server Message Block (SMB) requests to prevent the spread of malicious activity. You can also import rules you've already written in common open-source rule formats as well as enable integrations with managed intelligence feeds sourced by AWS Partners.

Amazon Route 53

[Amazon Route 53](#) is a highly available and scalable cloud Domain Name System (DNS) web service. It is designed to give developers and businesses an extremely reliable and cost-effective way to route end users to internet applications by translating names such as “www.example.com” into the numeric IP addresses, such as 192.0.2.1, that computers use to connect to each other. Route 53 is fully compliant with IPv6 as well.

Amazon Route 53 effectively connects user requests to infrastructure running in AWS, such as Amazon EC2 instances, Elastic Load Balancing load balancers, or S3 buckets, and can be used to route users to infrastructure outside of AWS. You can use Amazon Route 53 to configure DNS health checks to route traffic to healthy endpoints or to independently monitor the health of your application and its endpoints.

AWS Shield Advanced

[AWS Shield Advanced](#) is a managed service that builds a customized DDoS protection capability specifically for your applications needs. You receive tailored detection based on the specific traffic patterns of your application, protection against Layer 7 DDoS attacks including AWS WAF at no additional cost, access to 24x7 specialized support from the AWS Shield Response Team (SRT), centralized management of security policies through AWS Firewall Manager, and cost protection to safeguard against scaling charges resulting from DDoS-related usage spikes. For more information on Shield Advanced, see [AWS Shield Advanced documentation](#).

AWS WAF

[AWS WAF](#) is a web application firewall that helps protect your web applications or APIs against common web exploits that may affect availability, compromise security, or consume excessive resources. AWS WAF gives you control over how traffic reaches your applications by enabling you to create security rules that block common attack patterns, such as SQL injection, cross-site scripting, and others. You can also create specific traffic patterns you define.

You can get started quickly using [Managed Rules for AWS WAF](#), a pre-configured set of rules managed by AWS or AWS Marketplace sellers. The Managed Rules for AWS WAF address issues like the OWASP Top 10 security risks. These rules are regularly updated as new issues emerge. AWS WAF includes a full-featured API that you can use to automate the creation, deployment, and maintenance of security rules.

Further reading

For additional information, see:

- [GE Power Case Study: How GE Power Uses AWS to Monitor Power Plants and Save Its Customers Millions](#)
- [Vector and AWS join forces to accelerate the future of energy](#) (blog)
- [Formula 1 Case Study](#)
- [AWS IoT Customers: From emerging start-ups to large enterprises, learn why our customers choose AWS IoT](#)

Contributors

Contributors to this document include:

- Jana Kay, Cloud Security Strategist, AWS Security
- Maddie Bacon, Technical Writer, AWS Security

Document revisions

Date	Description
September 24, 2021	First publication

This version has been archived.

Notes

For the latest version of this document, visit:

- ¹ [AWS WAF Launches New Comprehensive Logging Functionality](#)
- ² [AWS Best Practices for DDoS Resiliency](#)
- ³ [AWS Shield Threat Landscape Report – Q1 2020](#)
<https://docs.aws.amazon.com/whitepapers/latest/security-at-the-edge/security-at-the-edge.html>
- ⁴ *Defense in depth is an information security concept that involves building multiple, independent layers of security.* [AWS re:Invent 2019: A defense-in-depth approach to building web applications](#)

⁵ [AWS WAF Security Automations](#)

⁶ [AWS re:Invent 2019: A defense-in-depth approach to building web applications](#)

⁷ [AWS IoT for the Edge](#)

⁸ AWS encourages customers to assess their data classification approach and hone in on which data needs to stay within their country or Region, and why. For more information on data residency, see the AWS whitepaper [Data Residency: AWS Policy Perspectives](#).

For more information on data classification, see the AWS whitepaper [Data Classification: Secure Cloud Adoption](#).

This version has been archived.

For the latest version of this document, visit:

<https://docs.aws.amazon.com/whitepapers/latest/security-at-the-edge/security-at-the-edge.html>