

Deployment Scenarios for Cisco Web Security Appliance on Amazon Web Services

1. Introduction

Amazon Web Services (AWS) provides several tools and services that can be used by the organization to achieve instant global deployment, cost reduction, or ease of management. Cisco® Web Security Appliance (WSA) technology uses some of these services to offer robust web security for employees and applications. The aim of this document is to provide an overview of various deployment scenarios by using the AWS services with WSA. The motive is to provide customers with a starting point for their WSA deployments on AWS.

The following use cases have been elaborated in detail:

1. **Deploying WSA for workloads running on AWS:** For customers who want to provide web security for applications already running on AWS.
2. **WSA on AWS for corporate data center:** For customers who require a substitute for on-premises WSAs on the public cloud. Customers can choose whether to have a direct link with AWS using direct connect or via a VPN tunnel initiated through Cisco Adaptive Security Appliance or Cisco Cloud Services Router or any VPN gateway.
3. **WSA for roaming users:** For organizations looking to provide consistent web protection for employees whether they are on a corporate network or on public networks.



Contents

1. Introduction

2. WSA for workloads in AWS

2.1 Overview

2.2 Deploying WSA on AWS

2.3 Network load balancer for cloud workloads

2.4 Steps to configure a load balancer

2.5 Next steps

3. Web security for corporate data center with WSA on AWS

3.1 Overview

3.2 AWS Direct Connect

3.3 WSA on AWS for workloads in data center using Direct Connect

3.4 Site-to-site VPN tunnel connecting the data center to AWS

4. Web security for roaming users with WSA on AWS

4.1 Overview

4.2 Users connecting to WSA via AnyConnect

4.3 Users connecting to WSA using a PAC file

4.3.1 About PAC files

4.3.2 About Internet-facing ELBs

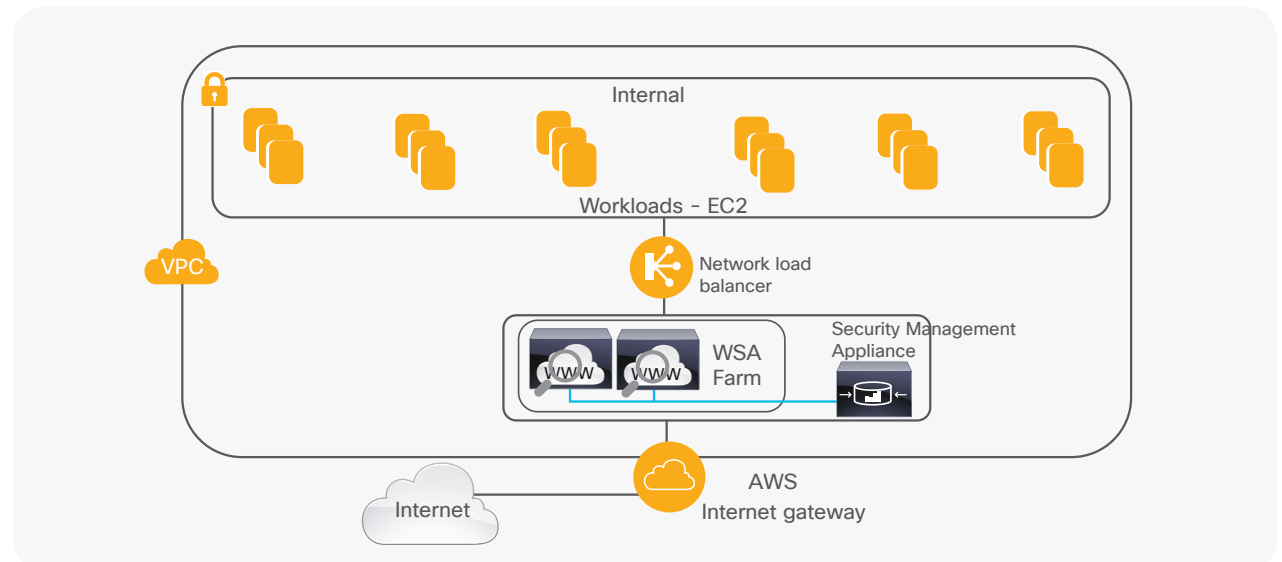
2. WSA for workloads in AWS

2.1 Overview

WSA can be an effective protection against web-based threats for your workloads running in the cloud. Alternatively, if your organization uses on-premises WSA for web security, as a good security practice you can deploy a similar security stack for your cloud-based applications using WSA on AWS. Furthermore, using Security Management Appliance virtual (SMAv) on AWS, you can configure and monitor a large number of WSA farms in your cloud environment itself.

An architectural view of cloud-based workload security is defined in Figure 1. Business applications such as customer relationship management, analytics tools, or collaboration platforms can form internal workloads that operate exclusively from the cloud. These applications can be running into various AWS Elastic Compute Cloud (EC2) instances spread across different availability zones. You can direct the web requests originating from these instances to an internal network load balancer, which will redirect the traffic to a WSA farm.

Figure 1. Architecture for cloud workload security with WSA



Contents

1. Introduction

2. WSA for workloads in AWS

2.1 Overview

2.2 Deploying WSA on AWS

2.3 Network load balancer for cloud workloads

2.4 Steps to configure a load balancer

2.5 Next steps

3. Web security for corporate data center with WSA on AWS

3.1 Overview

3.2 AWS Direct Connect

3.3 WSA on AWS for workloads in data center using Direct Connect

3.4 Site-to-site VPN tunnel connecting the data center to AWS

4. Web security for roaming users with WSA on AWS

4.1 Overview

4.2 Users connecting to WSA via AnyConnect

4.3 Users connecting to WSA using a PAC file

4.3.1 About PAC files

4.3.2 About Internet-facing ELBs

2.2 Deploying WSA on AWS

Deploying WSA on AWS is easy with the availability of the latest Amazon Machine Image (AMI) from the marketplace. You can choose WSA AMIs based on your network load requirement. Once you have the IT workloads up and running in your AWS environment, you can deploy WSA instances in the corporate Virtual Private Cloud (VPC) or a default VPC. Follow the steps defined in the [deployment guide](#) for more details.

2.3 Network load balancer for cloud workloads

By distributing network traffic across multiple virtual WSAs, traffic can be processed faster than in a scenario in which all traffic flowed through a single WSA. The load balancer is designed to achieve this traffic load distribution. A network load balancer is designed to handle tens of millions of requests per second while maintaining high throughput at Ultra low latency. Deploying a load balancer in a VPC comes with two choices: You must choose whether to make it an internal load balancer or an Internet-facing load balancer. The nodes of an Internet-facing load balancer have public IP addresses, whereas the nodes of an internal load balancer have only private IP addresses. The internal load balancers can only route requests from clients with access to the VPC for the load balancer. This implies that all your EC2 instances (or applications) should be part of the same VPC as intended for the load balancer.

2.4 Steps to configure a load balancer

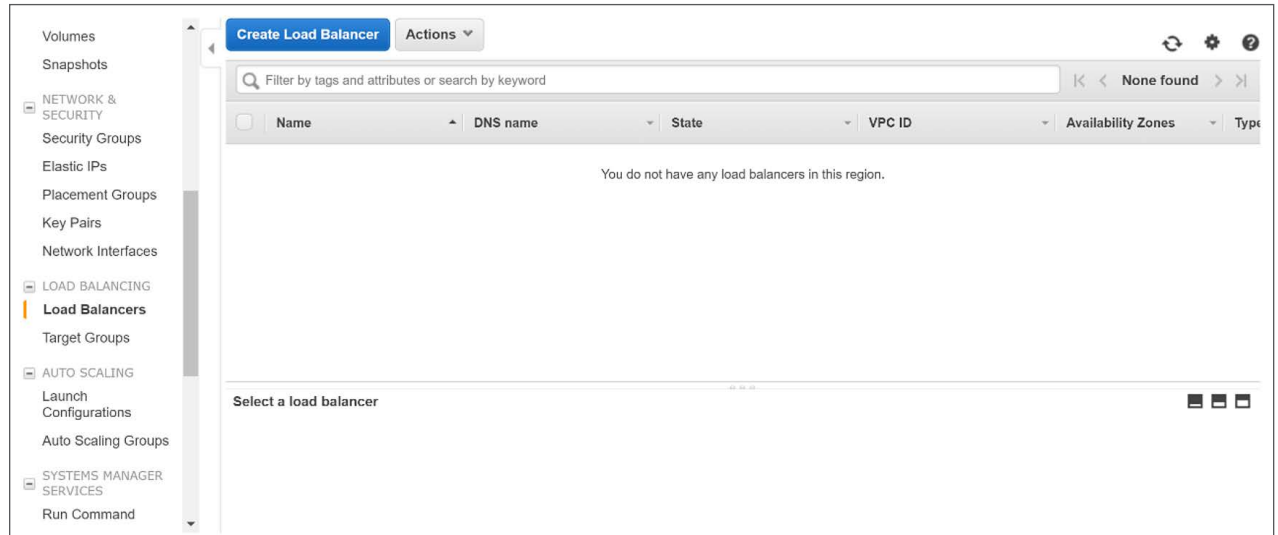
Before configuring the load balancer, make sure you have deployed the WSA instances on the same VPC. It is recommended that you specify a unique name for the WSA running on different EC2 instances for easy identification.

To create an internal load balancer:

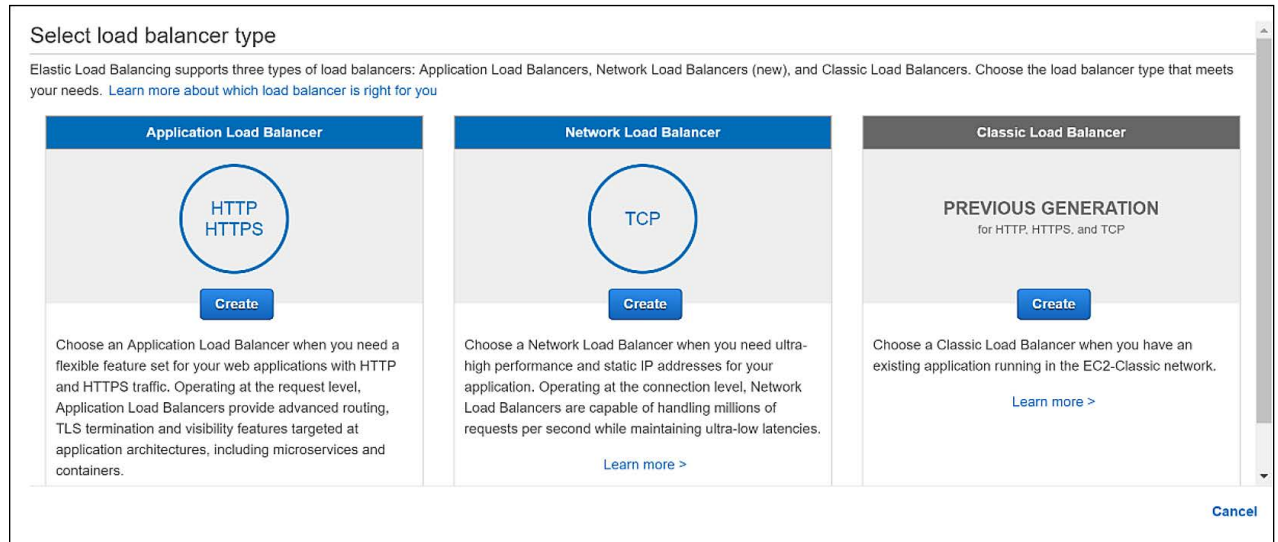
1. Navigate to the AWS dashboard and select EC2 Console.
2. From the left-side navigation pane, under the **Load Balancing** subsection, click on **Load Balancers**.
3. From the Load Balancer menu screen, click on **Create Load Balancer**.

Contents

1. Introduction
2. WSA for workloads in AWS
 - 2.1 Overview
 - 2.2 Deploying WSA on AWS
 - 2.3 Network load balancer for cloud workloads
 - 2.4 Steps to configure a load balancer
 - 2.5 Next steps
3. Web security for corporate data center with WSA on AWS
 - 3.1 Overview
 - 3.2 AWS Direct Connect
 - 3.3 WSA on AWS for workloads in data center using Direct Connect
 - 3.4 Site-to-site VPN tunnel connecting the data center to AWS
4. Web security for roaming users with WSA on AWS
 - 4.1 Overview
 - 4.2 Users connecting to WSA via AnyConnect
 - 4.3 Users connecting to WSA using a PAC file
 - 4.3.1 About PAC files
 - 4.3.2 About Internet-facing ELBs



4. Choose **Network Load Balancer** from the Select load balancer type menu.



Contents

1. Introduction

2. WSA for workloads in AWS

2.1 Overview

2.2 Deploying WSA on AWS

2.3 Network load balancer for cloud workloads

2.4 Steps to configure a load balancer

2.5 Next steps

3. Web security for corporate data center with WSA on AWS

3.1 Overview

3.2 AWS Direct Connect

3.3 WSA on AWS for workloads in data center using Direct Connect

3.4 Site-to-site VPN tunnel connecting the data center to AWS

4. Web security for roaming users with WSA on AWS

4.1 Overview

4.2 Users connecting to WSA via AnyConnect

4.3 Users connecting to WSA using a PAC file

4.3.1 About PAC files

4.3.2 About Internet-facing ELBs

5. Provide a name for your load balancer, and under the **Scheme**, select **internal**.
6. Specify the load balancer protocol and the port. For WSA, you can specify the protocol as TCP and the ports as 80 and 443.
7. Next, under **Availability Zones**, select the relevant VPC with the correct subnet. An internal load balancer requires a public IP, and it will be automatically assigned by AWS.

1. Configure Load Balancer 2. Configure Routing 3. Register Targets 4. Review

Step 1: Configure Load Balancer

Basic Configuration

To configure your load balancer, provide a name, select a scheme, specify one or more listeners, and select a network. The default configuration is an Internet-facing load balancer in the selected network with a listener that receives TCP traffic on port 80.

Name ⓘ WSALB

Scheme ⓘ Internet-facing internal

Listeners

A listener is a process that checks for connection requests, using the protocol and port that you configured.

Load Balancer Protocol	Load Balancer Port	
TCP	80	✕
TCP	443	✕

Add listener

8. In the following section, define the target group for your load balancer and give it a name.

1. Configure Load Balancer 2. Configure Routing 3. Register Targets 4. Review

Step 2: Configure Routing

Your load balancer routes requests to the targets in this target group using the protocol and port that you specify, and performs health checks on the targets using these health check settings. Note that each target group can be associated with only one load balancer.

Target group

Target group ⓘ New target group

Name ⓘ wsalbgrou

Protocol ⓘ TCP

Port ⓘ 80

Target type ⓘ instance

Health checks

Protocol ⓘ TCP

▶ Advanced health check settings

Contents

1. Introduction

2. WSA for workloads in AWS

- 2.1 Overview
- 2.2 Deploying WSA on AWS
- 2.3 Network load balancer for cloud workloads
- 2.4 Steps to configure a load balancer
- 2.5 Next steps

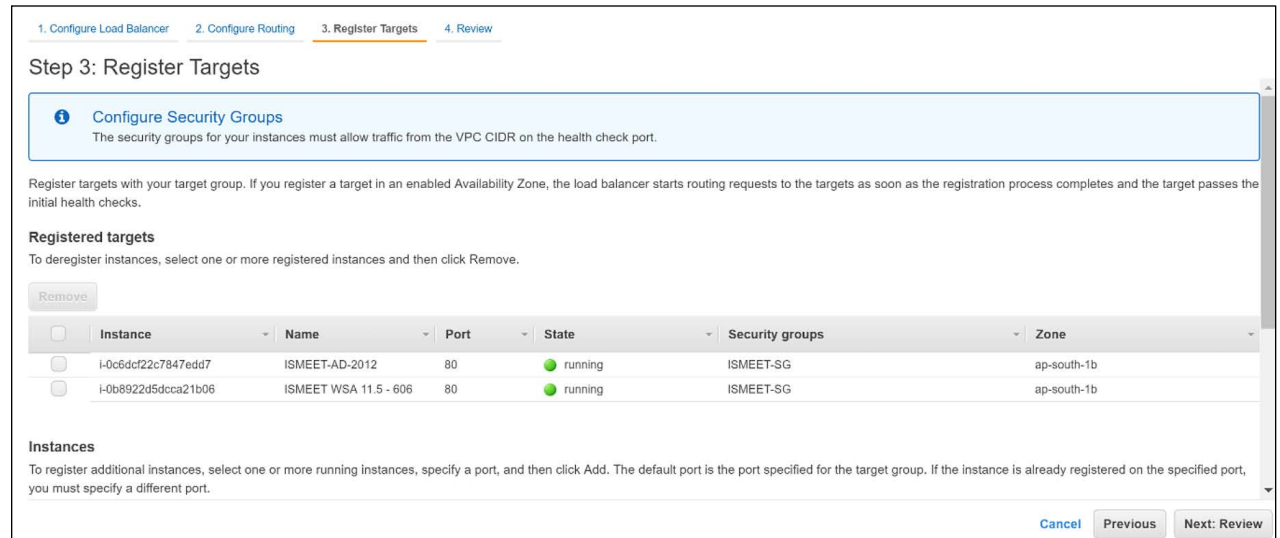
3. Web security for corporate data center with WSA on AWS

- 3.1 Overview
- 3.2 AWS Direct Connect
- 3.3 WSA on AWS for workloads in data center using Direct Connect
- 3.4 Site-to-site VPN tunnel connecting the data center to AWS

4. Web security for roaming users with WSA on AWS

- 4.1 Overview
- 4.2 Users connecting to WSA via AnyConnect
- 4.3 Users connecting to WSA using a PAC file
 - 4.3.1 About PAC files
 - 4.3.2 About Internet-facing ELBs

9. In the **Register Targets** section, you will get the list of instances running in your VPC. Select the WSA instances to which you want to point the load balancer and click on **Add to Registered**.



1. Configure Load Balancer 2. Configure Routing 3. Register Targets 4. Review

Step 3: Register Targets

Configure Security Groups
The security groups for your instances must allow traffic from the VPC CIDR on the health check port.

Register targets with your target group. If you register a target in an enabled Availability Zone, the load balancer starts routing requests to the targets as soon as the registration process completes and the target passes the initial health checks.

Registered targets
To deregister instances, select one or more registered instances and then click Remove.

Remove

Instance	Name	Port	State	Security groups	Zone
<input type="checkbox"/> i-0c6dcf22c7847edd7	ISMEET-AD-2012	80	running	ISMEET-SG	ap-south-1b
<input type="checkbox"/> i-0b8922d5dcca21b06	ISMEET WSA 11.5 - 606	80	running	ISMEET-SG	ap-south-1b

Instances
To register additional instances, select one or more running instances, specify a port, and then click Add. The default port is the port specified for the target group. If the instance is already registered on the specified port, you must specify a different port.

Cancel Previous Next: Review

10. Click on **Review** to revisit your settings and then click **Create**.
11. You can see your load balancer getting provisioned in the **Load Balancing** section in the navigation pane.
12. Once the load balancer is provisioned, you can review the target groups within the load balancing section to check whether correct WSA instances are listed and healthy.

2.5 Next steps

Once the load balancer is provisioned, you can redirect all the web traffic of your cloud workload to the private IP of the load balancer. This will make the load balancer direct all the traffic to the different WSAs based on the scale of web traffic.

Contents

1. Introduction

2. WSA for workloads in AWS

- 2.1 Overview
- 2.2 Deploying WSA on AWS
- 2.3 Network load balancer for cloud workloads
- 2.4 Steps to configure a load balancer
- 2.5 Next steps

3. Web security for corporate data center with WSA on AWS

- 3.1 Overview
- 3.2 AWS Direct Connect
- 3.3 WSA on AWS for workloads in data center using Direct Connect
- 3.4 Site-to-site VPN tunnel connecting the data center to AWS

4. Web security for roaming users with WSA on AWS

- 4.1 Overview
- 4.2 Users connecting to WSA via AnyConnect
- 4.3 Users connecting to WSA using a PAC file
 - 4.3.1 About PAC files
 - 4.3.2 About Internet-facing ELBs

3. Web security for corporate data center with WSA on AWS

3.1 Overview

For organizations that are already invested in on-premises installations, complete cloud adoption may not be possible. Complete migration to the cloud takes time. Additionally, it is possible that you may require the on-premises systems as an executive mandate. In this scenario, a hybrid approach that can seamlessly integrate both on-premises and cloud workloads, without relying on heavy new investments, works the best. With WSA deployed on AWS, you get a leading web security product in a hybrid deployment, making it easier for your organization to migrate towards the public cloud.

You can deploy WSAs on AWS to provide web security for your corporate data center. Or, you can even consider this deployment for your branch offices that do not have the infrastructure for on-premises WSA installations. There are two possible scenarios for deploying WSA for your corporate data center or branch offices:

1. Using AWS Direct Connect from a data center to AWS
2. Creating a site-to-site VPN tunnel for connecting a data center to AWS

3.2 AWS Direct Connect

Using AWS Direct Connect, you can connect directly to your AWS VPC using a private network connection. This connection uses industry-standard 802.1q VLANs, and it can be partitioned into multiple virtual interfaces. The multiple interfaces allow you to access public as well as private resources within the AWS environment. The AWS Direct Connect provides 1-Gbps and 10-Gbps connections, and you can easily provision multiple connections if you need more capacity. For more information, please visit <https://aws.amazon.com/directconnect/>.

The user guide for AWS Direct Connect can be accessed from here:

<https://docs.aws.amazon.com/directconnect/latest/UserGuide/dc-ug.pdf>

3.3 WSA on AWS for workloads in data center using Direct Connect

Figure 2 describes scenario 1, where customers can provision AWS Direct Connect to establish a dedicated network connection from on-premises systems to AWS. Once your Direct Connect infrastructure is in place, you can connect it to your private VPC. To achieve this, you can create a private virtual interface to connect to your VPC, or you can create a public virtual interface to connect to public AWS services that aren't in a VPC. More details are described in step 4 of the Direct Connect user guide.

Inside the VPC, you can have internal workloads that connect with the network load balancer and then to the WSA instances as described in section 1 of this deployment guide. The WSA instances would terminate the server-side connections on the AWS Internet gateway for Internet access.

Contents

1. Introduction

2. WSA for workloads in AWS

- 2.1 Overview
- 2.2 Deploying WSA on AWS
- 2.3 Network load balancer for cloud workloads
- 2.4 Steps to configure a load balancer
- 2.5 Next steps

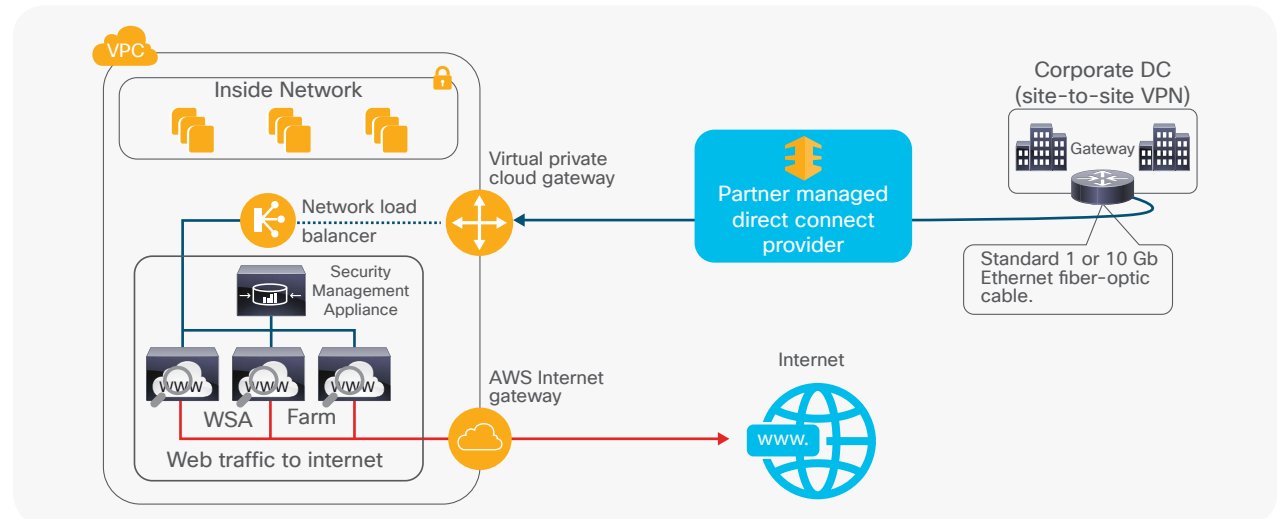
3. Web security for corporate data center with WSA on AWS

- 3.1 Overview
- 3.2 AWS Direct Connect
- 3.3 WSA on AWS for workloads in data center using Direct Connect
- 3.4 Site-to-site VPN tunnel connecting the data center to AWS

4. Web security for roaming users with WSA on AWS

- 4.1 Overview
- 4.2 Users connecting to WSA via AnyConnect
- 4.3 Users connecting to WSA using a PAC file
 - 4.3.1 About PAC files
 - 4.3.2 About Internet-facing ELBs

Figure 2. Connecting corporate data center to AWS using Direct Connect



3.4 Site-to-site VPN tunnel connecting the data center to AWS

An alternate way of connecting the data center with the AWS is via a virtual private network. Cisco Adaptive Security Appliance (ASA) or Cisco Cloud Services Router, deployed either on-premises or virtually, can be used to create a VPN tunnel over the public network between your data center and AWS. ASA also allows you to create multiple tunnels for redundancy. The Cisco ASA 5500 Series can offer 100 Mbps to 1 Gbps of throughput depending upon your requirement and number of users.

More details on ASA deployment can be found here:

https://docs.aws.amazon.com/AmazonVPC/latest/NetworkAdminGuide/Cisco_ASA.html

The VPN can be terminated at the AWS virtual private gateway. You can create the AWS VPN and attach it to the VPC, which houses your workloads and WSA farms. The AWS VPN is not the initiator of the VPN tunnel. The tunnel comes up when the traffic originates from the ASA, as per the requests made by the on-premises users. For details on how to configure the AWS VPN and attach it to your VPC, please follow the steps mentioned here:

https://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_VPN.html

Once the VPN infrastructure is in place, you can deploy the WSA farms in a similar fashion as mentioned in section 1 of this deployment guide.

Contents

1. Introduction

2. WSA for workloads in AWS

2.1 Overview

2.2 Deploying WSA on AWS

2.3 Network load balancer for cloud workloads

2.4 Steps to configure a load balancer

2.5 Next steps

3. Web security for corporate data center with WSA on AWS

3.1 Overview

3.2 AWS Direct Connect

3.3 WSA on AWS for workloads in data center using Direct Connect

3.4 Site-to-site VPN tunnel connecting the data center to AWS

4. Web security for roaming users with WSA on AWS

4.1 Overview

4.2 Users connecting to WSA via AnyConnect

4.3 Users connecting to WSA using a PAC file

4.3.1 About PAC files

4.3.2 About Internet-facing ELBs

More information can be obtained here:

<https://www.cisco.com/c/en/us/support/security/anyconnect-secure-mobility-client/tsd-products-support-series-home.html>

Furthermore, as described in section 3.3, ASA should form a VPN tunnel and terminate it on the AWS VPN gateway.

4.3 Users connecting to WSA using a PAC file

As an alternative to using AnyConnect, you can deploy PAC files on client machines to redirect their web traffic to WSAs. A PAC file is a language to inform web browsers how to leverage proxies on their network. The PAC file checks the local IP subnet address of the client's machine and then makes a decision based on IF/ELSE statement/s. If the client machine is located in a subnet (for example, corporate network) that matches the one specified in the PAC file, a proxy server is used. If the client machine is on any other subnet (for example, public Wi-Fi), the connection is established to an Internet-facing Elastic Load Balancer (ELB) on AWS. The ELB redirects the web requests to the WSAs located across different availability zones in the AWS.

4.3.1 About PAC files

Admins can host the PAC file on WSA under GUI > Security Services > PAC file hosting. By default, the proxy PAC file would be hosted on port 9001. When using WSA to host PAC files, by default, users need to point the browser to the following location while remaining in the corporate network.

`http://WSA_IP:9001/pacfile.pac`

Additionally PAC files can be hosted on an external web server as well.

4.3.2 About Internet-facing ELBs

AWS provides options for three types of ELBs: an application load balancer, a network load balancer, and a classic load balancer.

More information on creating an Internet-facing ELB can be found here: <https://docs.aws.amazon.com/elasticloadbalancing/latest/classic/elb-internet-facing-load-balancers.html>