# LEARNING AMAZON WEB SERVICES (AWS)

## A Hands-On Guide to the Fundamentals of AWS Cloud

MARK WILKINS

# Learning Amazon Web Services (AWS)

applied to a single network adapter, providing an effective set of allow rules mandating what ports and protocols are allowed for both inbound and outbound traffic.

The main characteristics of security group rules are as follows:

- Security groups define allow rules. (You can't create rules that explicitly deny access.)
- Security group rules allow you to direct traffic outbound from one security group and inbound to another security group within the same VPC.
- Security groups don't deny traffic explicitly; instead, they deny traffic implicitly by only stating what is allowed.
- Security groups are defined as stateful; if requests are allowed in, response traffic is allowed out regardless of defined outbound rules.
- For each rule, you define the protocol, the port, or port range, and the source inbound rules or destination outbound rules for the traffic.
- The protocols allowed are TCP, UDP, or ICMP.
- Port Range: for either TCP or UDP, or a custom protocol, this is the range of ports to allow. You can specify a single port such as port 22, or a range of ports if you're dealing with outbound dynamic port mapping. The ports that can be assigned are defined by RFC 5237 and 7045, which define the standard TCP/IP protocols.

One important concept to grasp about security groups is that they don't deny traffic flow. Instead, they allow traffic flow. Another equally important concept is the "direction of traffic flow" allowed by a security group. As we know, the network traffic that is allowed in by a security group rule is also allowed out. However, a defined inbound port request does not use the same port number for the outbound response. For example, if there's a rule defined allowing inbound HTTP traffic across port 80 inbound, the outbound traffic response is allowed out, but the response traffic does not use port 80 outbound. In most cases, outbound traffic uses a dynamically assigned port called an ephemeral port, determined by the operating system of the server making the response. We will explore ephemeral ports when we cover network ACLs later in this chapter.

When a VPC is created, a default security group is also created, as shown in Figure 3-23. Note that all outbound traffic is allowed between the EC2 instances that are assigned the security group; however, any other traffic is implicitly denied. Therefore no one can get into any of the EC2 instances from the outside world because the security group did not allow any external inbound traffic. It's also important to understand that you can't delete a default security group; you also don't have to use the default security groups; you can ignore them and create and assign custom security groups.

If you don't pay attention and don't specify a custom security group when an EC2 instance is created, at launch, the default security group is associated automatically. As we now know,

the default security group allows all outbound traffic from the instance associated with the default security group, but accessing the EC2 instance from an external location will not allowed.

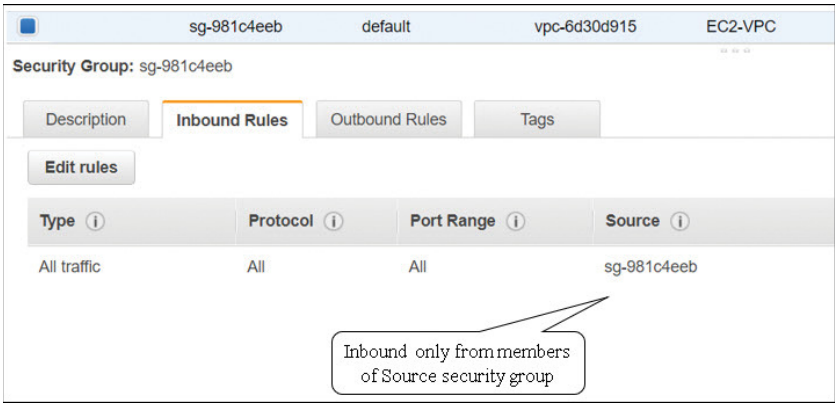


Figure 3-23 Default security group in newly created VPC

## Custom Security Groups

What happens when you create a custom security group? First, you must associate the security group with a specific VPC. After it is first created, a custom security group allows no inbound traffic but allows all outbound traffic by default. After the initial creation of the custom security group and selecting the security group properties, from the Inbound Rules tab you create the inbound rules defining the network traffic that's allowed inbound. On the Outbound Rules tab, define the network traffic or security group that is allowed to communicate outbound.

- **Inbound rules**—Define the source of the traffic—that is, where it is coming from, and what the destination port or port range is. The actual source of the traffic could be a single IP address (IPv4 or IPv6), a range of addresses, or another security group.
- **Outbound rules**—Define the destination of the traffic—that is, where it is going to, and the destination port, or port range. The actual destination of the traffic could be a single IP address (IPv4 or IPv6), a range of addresses, or another security group.
  - A prefix list ID for a specific AWS network service, such as an Internet gateway
  - Another security group—This feature allows instances that are associated with one security group to access instances associated with another security group. The security group can reside in the same VPC or can be from a VPC that has been peered together through a VPC peering connection.

> **Note**
>
> Security groups "allow" access; however, security groups are also said to "deny by default." If an incoming port is not allowed, technically access is denied.

The best practice when designing security groups is to restrict what ports need to be opened. If you place a load balancer in front of your EC2 instances, the only ports that need to be allowed by the security group protecting the load balancer are the port or ports your application requires.

For your production applications in your design, it's a good idea to consider picking unique ports based on your production server's tier placement. For example, in a three-tier design, Web servers could be assigned 3020 and application servers 3030, and database instances could have custom ports chosen. That's just an extra level of security to consider. In addition, be prepared to plan for a clean security audit, or for successful troubleshooting by designing and documenting your security group's naming scheme up front. Putting some real thought into all naming conventions pays off when you are under pressure. We have all been in this situation; fill in the blank "what's the blasted name of the --------. Who named it that?" Here are some examples to consider for your security group configurations and initial setup.

### App Server Inbound Ports

This security group has rules to allow inbound access of HTTP and HTTPS from any IPv4 address as the traffic is arriving inbound from the public Internet. As shown in Table 3-3, the source IP address of 0.0.0.0/0 indicates the traffic is allowed from any Internet location.

Table 3-3 **Security Group Allowing HTTP Access**

| **Inbound Rules** | | | | |
|---|---|---|---|---|
| **Protocol** | **Number** | **Port** | **Source IP** | **Comment** |
| TCP | 6 | 80 (HTTP) | 0.0.0.0/0 | Inbound from anywhere (IPv4) |
| **Outbound Rules** | | | | |
| **Protocol** | **Number** | **Port** | **Destination IP** | **Comment** |
| ALL | 6 | 80 (HTTP) | 0.0.0.0/0 | Outbound IPv4 traffic |

### Database Server Inbound Ports

At AWS, there are a number of managed database server options available with RDS; during configuration, you can choose to use the default port address assigned based on the database product's default port access. You could change this to a custom port number to add an additional element of security. The source IP address could be specified as a single IP address or a range of IP addresses from your subnet of application servers that need to query the database. Default RDS security group database port options are listed in Table 3-4.

Table 3-4 **Default Database Inbound Ports**

| Protocol | Number | Port | Source IP | Product | Comment |
|---|---|---|---|---|---|
| TCP | 6 | 1433 | Single IP address, or, A range of IP addresses | Microsoft SQL | Default port to access selected database instance |
| TCP | 6 | 3306 | | Aurora / MySQL | |
| TCP | 6 | 5432 | | PostgreSQL | |
| TCP | 6 | 1521 | | Oracle | |
| TCP | 6 | 5439 | | Redshift | |

### Administration Access

If you need to connect to an EC2 instance to perform direct administration, you must associate a security group with the network interface card that has inbound rules allowing Secure Shell (SSH) or Remote Desktop Protocol (RDP) access depending on the host operating system of the instance (see Table 3-5). A database instance may be hosted on a private subnet and not directly accessible from the Internet. However, setting up a bastion host would allow access for administrators to first authenticate to the bastion host and then "jump" to the associated EC2 instance in the private subnet.

Table 3-5 **Security Groups Allowing Administrative Access**

| Protocol | Number | Port | Operating system | Comment |
|---|---|---|---|---|
| TCP | 6 | 22 | Linux | Public IPv4 address, or IPv6 address or range of addresses |
| TCP | 6 | 3389 | Windows | |

### Pinging an EC2 Instance

Pinging an instance requires access to ICMP traffic. It's almost an automatic process; drop to a command prompt and ping the public IP address of the newly created EC2 instance. In this case, the instance didn't get the ping request because the security group is not allowing ICMP. The ping traffic request is inbound, so you must add an inbound ICMP rule, as shown in Table 3-6.

Table 3-6 **Allowing PING Access**

| Protocol | Number | Port | Source IP | Comment |
|---|---|---|---|---|
| ICMP | 1 | 8 (ECHO) | Admin workstation network | Public IPv4 address or range of addresses |

### Elastic Load Balancing (ELB)

Using the ELB service to "protect" your Web tier with a public facing load balancer also requires a security group with rules that allow communication with the targeted groups of EC2 instances or containers, as detailed in Table 3-7.

Table 3-7 **Controlling ELB Traffic Flow**

**Inbound ELB Rule**

| Protocol | Number | Port | Source IP | Comment |
|---|---|---|---|---|
| TCP | 6 | Defined listener port | 0.0.0.0/0 for Internet-facing load balancer<br>IPv4 CIDR block if Internal load balancer | Inbound from anywhere (IPv4) |

**Outbound ELB Rule**

| Protocol | Number | Port | Source IP | Comment |
|---|---|---|---|---|
| TCP | 6 | Instance listener port | ID of the instance security group | Outbound traffic to instance's listener port |
| TCP | 6 | Health check port | ID of the instance security group | Outbound traffic to instance health check port |

Load balancers also can perform health checks on the EC2 instances that are registered with the load balancer. Health check communication from the load balancer on both the listener and the health check ports must be defined, as shown in Table 3-8.

Table 3-8 **ELB Communication**

**Inbound Instance Rule**

| Protocol | Number | Port | Source IP | Comment |
|---|---|---|---|---|
| TCP | 6 | Instance listener port | ID of the load balancer security group | Traffic allowed from the load balancer traffic to the instance's listener port |
| TCP | 6 | Health check port | ID of the load balancer security group | Traffic allowed from the load balancer to the instance health check port |

> **Note**
> Security group entries for IPv4 and IPv6 are added as separate rules.

The initial limit for the number of security groups that can be applied to a network interface is 5; you can request an increase from AWS support increasing the limit up to 16. How many rules are allowed per security group? Again, soft limits decree that every custom security group can have up to 50 inbound, and 50 outbound, IPv4 or IPv6 rules. You can increase the number of rules per security group by contacting AWS.

AWS is right to have these limits; imagine a situation in which there were no limits on the number of security groups and rules. Is it possible to slow down an instance by having too many rules to calculate and determine what is allowed? I think we know the answer to that question.

### Security Group Summary

- Plan your security groups with names that make sense.
- Create security groups for administrative tasks.
- Create a security group for your public-facing load balancer that sends traffic to your Web tier security group.
- Create a security group for your Web tier that sends traffic to your app tier or app tier load balancer security group.
- Create a security group for your app tier that sends traffic to your DB tier security group.
- Deploy and test everything on a test VPC before deploying anything in production.

> Make sure to watch the companion video "Creating Security Groups."

## Network ACLs

The network access control list (NACL) is an optional software firewall that controls inbound and outbound traffic for each subnet within the VPC. Each VPC is associated with a default network ACL that is really a placeholder; the default network ACL allows all inbound and outbound IPv4 traffic and IPv6 traffic at the subnet level. Custom NACLs can be created and, just like security groups, are a reusable security template associated with your AWS account. NACLs, once created, can be associated with one or multiple subnets.

Each NACL contains a set of inbound and outbound subnet traffic rules that are listed in order from a starting number rule to the highest number rule, as shown in Figure 3-24. Rules are processed in order and evaluated to determine if traffic is allowed or denied inbound or outbound on each subnet.

If you are working in a networking environment where it's a mandated practice to control network traffic at the subnet layer through zone controls, you'll feel right at home with NACLs. If you have properly set up security groups to allow communication on the specific ports and protocols required for your applications, you may not feel it is necessary to add NACLs to your security design. However, NACLs operate at the lower subnet level and provide an additional layer of defense. A single NACL can protect multiple application servers at the subnet level. Rules can target an entire subnet or an individual IP address.