Check Point
SOFTWARE TECHNOLOGIES LTD.

12 January 2020

CLOUDGUARD IAAS HIGH AVAILABILITY FOR GOOGLE CLOUD

R80.30 AND ABOVE

Deployment Guide

[Classification: Protected]

Check Point Copyright Notice

Important Information

Latest Software

We recommend that you install the most recent software release to stay up-to-date with the latest functional improvements, stability fixes, security enhancements and protection against new and evolving attacks.

Certifications

For third party independent certification of Check Point products, see the Check Point Certifications page.

Latest Version of this Document

Open the latest version of this document in a Web browser.

Download the latest version of this document in PDF format.

Feedback

Check Point is engaged in a continuous effort to improve its documentation.

Please help us by sending your comments.

Revision History

Date	Description
12 January 2020	In *"Known Limitations" on page 30*, removed: "VPN termination is not supported by Check Point CloudGuard IaaS High Availability". S2S VPN termination is supported by GCP HA.
03 December 2019	■ In *"Components of the Check Point Solution" on page 15*, added note "The addresses are not attached to the members' NICs until Step 2 below. ■ Changed the instructions for *"Step 2: Configure Cluster Objects in SmartConsole" on page 16*.
24 November 2019	**Updates to:** *"Step 1: Deploy a Template in GCP" on page 13* ■ Added new parameter - Firewall ■ Edited Cluster network parameter ■ In *"Notes about the template:" on page 15*, removed "A GCP Firewall rule ...no need for a second security mechanism" *"Step 2: Configure Cluster Objects in SmartConsole" on page 16* Step 4 - Added steps for attaching Active Cluster Member's NICO Edited steps 5, 6, and 13
24 September 2019	First release of this document

Table of Contents

Terms

Active

State of a Cluster Member that handles network connections that pass through the cluster. In a cluster deployment, only one Cluster Member is Active and can handle connections.

Check Point WatchDog

A process that launches and monitors critical processes such as Check Point daemons on the local machine, and attempts to restart them if they fail.

Cluster

Two or more Security Gateways that work together in a redundant configuration - High Availability.

Failover

Also, Fail-over. Transferring of a control over traffic (packet filtering) from Cluster Member that suffered a failure to another Cluster Member (based on internal cluster algorithms).

SmartConsole

Check Point main GUI client used to create and manage the Security Policy.

Standby

State of a Cluster Member that is ready to be promoted to Active state (if the current Active Cluster Member fails). Applies only to ClusterXL High Availability Mode.

Check Point CloudGuard IaaS High Availability for GCP

Overview

CloudGuard Security Cluster for Google Cloud Platform provides High Availability through state synchronization. This occurs when a standby CloudGuard Security Gateway Cluster Member, deployed in one Zone, monitors the state of an active member deployed in another Zone. If the active gateway fails, then the standby member assumes active state and performs the necessary changes in your GCP environment so that traffic will be routed through it.

The CloudGuard Security Cluster provides comprehensive enterprise-grade security. It continues to protect your GCP resources even when it encounters a problem, which on a standalone gateway would have resulted in a complete loss of connectivity.

Prerequisites

Before setting up your system, you must be familiar with the following topics:

Vendor	Topics
Google Cloud Platform	■ Virtual Private Cloud Network ■ Virtual Machines ■ Public IP Addresses ■ Routes
Check Point	Check Point R80.30 Check Point with Google Cloud Platform

Setting Up Check Point Clusters in GCP

A cluster is a group of Virtual Machines that work together in High Availability Mode. One Cluster Member is *Active*, and the second Cluster Member is *Standby*. When necessary, the cluster fails over from the Active Cluster Member to the Standby Cluster Member.

- For VPN traffic, Cluster Members use API calls to GCP to communicate the failover from the Active Cluster Member.
- The Standby Cluster Member then promotes itself to Active. During cluster failover, the Standby Cluster Member associates the primary external cluster IP address of the Active Cluster Member with its external interface (eth0).
  The Cluster Member that fails, uses the GCP API to attach the cluster's secondary IP address to itself.

Network

Follow this network diagram to configure your system. Make sure to replace the IP addresses in the sample environment with the IP addresses in your environment.

Network Diagram

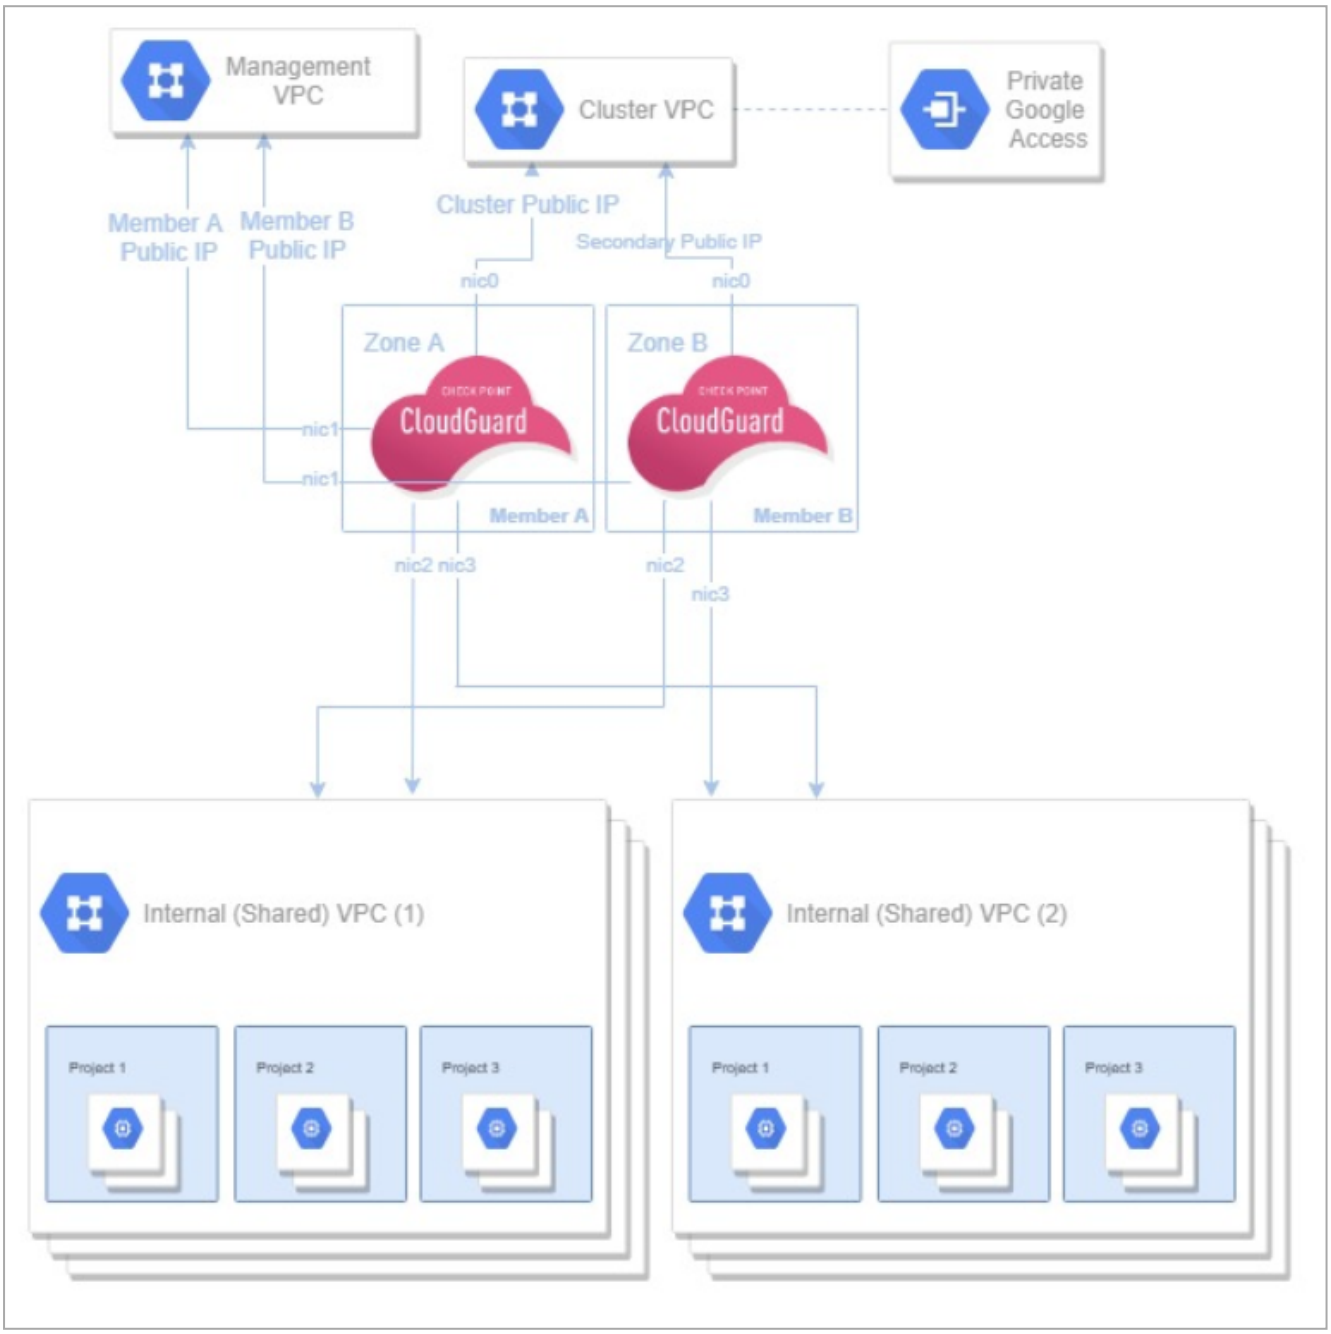

Network routing diagram:

- Each CloudGuard Security Gateway resides in a different GCP Zone.

- Each gateway has a network interface in a subnetwork (in the Cluster VPC). The CloudGuard IaaS High Availability solution will inspect inbound traffic from this VPC. Therefore, the GCP Firewall rules and routes should be configured to allow all traffic.
- Private access to Google APIs must be enabled for this subnetwork. This is to allow the CloudGuard IaaS High Availability VM instances access to Google APIs to execute a failover process.
- The cluster's public IP address will be attached to the network interface of the active gateway in this network.
- Each gateway has a network interface in a subnetwork in the Management VPC. This is the network that manages the gateways.
- The gateways have a network interface in each of the internal, optionally shared VPC networks that the cluster secures. In each of the internal VPC networks, a GCP Route routing all outbound traffic (0.0.0.0/0) to the active member will be automatically created.

In the diagram:

- The cluster protects two web applications.

Each web application has:

- Web server
- Application server

You must manually configure these components:

Backend hosts

Name	Attached to	Use
Cluster primary public address	The external interface (eth0) of the Active Cluster Member.	VPN and publishing services. Do not disable or delete this resource.
Cluster secondary public IP address	The external interface (eth0) of the Secondary Cluster Member.	Gets important Check Point updates. Do not disable or delete this resource.
Member *A* public address	Management Network Interface of Member A (eth1).	External management of Member A Do not disable or delete this resource.
Member *B* public address	Management Network Interface of Member B (eth1).	External management of Member B Do not disable or delete this resource.

Static IP Addresses:

Use the primary public IP address of the Active member (eth0) to forward traffic that comes from the Internet.

Important - You cannot use the following ports:

- 80
- 443
- 444

- 8082
- 8880

Diagram Components

The diagram shows:

- Four Virtual Networks in GCP, each contains four subnets:
  - Cluster
  - Management
  - Web
  - App

Check Point High Availability consists of two Cluster Members - *Member 1* and *Member 2*.

Cluster VPC Routes

3	**Destination**	**Nexthop**
	`0.0.0.0/0`	Default internet gateway (**Default**)
	`10.0.0.0/24`	Cluster Virtual Network (**Default**)

Management VPC

4	**Destination**	**Nexthop**
	`0.0.0.0/0`	Default internet gateway (**Default**)
	`172.20.0.0/24`	Management Virtual Network

Routes for each internal VPC

Internal VPC (1)

5	**Frontend**	**Nexthop**
	`10.2.1.0/24`	Virtual Network (Internal VPC 1)
	`0.0.0.0/0`	Default internet gateway (with High Priority 1000)
	`0.0.0.0/0` x-chkp-XXX-to-member-b	Instance Cluster Member-b (with High Priority 2)
	`0.0.0.0/0` x-chkp-XXX-to-member-a	Instance Cluster Member-a (with High Priority 1)

Internal VPC (2)

6	**Frontend**	**Nexthop**
	`10.3.1.0/24`	Virtual Network (Internal VPC2)
	`0.0.0.0/0`	Default internet gateway (Priority 1000)
	x-chkp-XXX-to-member-b	Instance Cluster Member-b (Priority 2)
	x-chkp-XXX-to-member-a	Instance Cluster Member-a (Priority 1)

Failover

The following occurs during cluster failover:

- When the active gateway fails, the standby member will identify the failure, and then do the following:
  - Detach the cluster's public IP address from the failed member.
  - Detach the cluster's secondary public IP address from the standby member. When the failed members returns to standby state, it will attach this address to itself.
  - Attach the cluster's public IP address to itself.
  - For each of the internal VPC networks:
    - If *Member A* becomes the active member, it will create high priority routes that will route all outbound traffic in the internal networks to itself.
      Note - A lower priority number equals a higher priority for the route.
    - If *Member B* becomes the active member, it will remove the high priority routes created by *Member A*. The result is that existing, lower priority routes will take effect, and route all outbound traffic in the internal networks to itself.

Note - This usually happens in less than 40 seconds. This affects East-West, inbound outbound and VPN tunnel failover.

These are the expected failover times - based on use case:

Use Case	Expected Failover Time	Comments
Site-to-Site VPN	Less than 40 seconds	Depends on GCP API.
Inbound inspection through primary public IP	Less than 40 seconds	Depends on GCP API.
Outbound inspection through primary public IP	Less than 40 seconds	Depends on GCP API.
East-West inspection	Less than 30 seconds	Depends on GCP API.

Traffic Flows

Note - Other Virtual Machines cannot be deployed in the Check Point solution subnets.

Inbound Traffic

- Traffic travels into the cluster's primary public address (Cluster VIP, attached to the eth0 interface of the Active cluster member).
- The Active Cluster Member inspects the traffic, and forwards it to the destination.

Inbound Traffic Reply

1. The traffic travels from the Web Server to the Active Cluster Member.
2. The Active Cluster Member forwards it to the destination.

Inbound VPN Traffic

1. Packet enters the frontend network interface (eth0) of the Active Cluster Member.
2. The Active Cluster Member decrypts the packet.
3. The Active Cluster Member forwards the packet to its destination.

Outbound Traffic

1. Traffic travels to an Active Cluster Member based on the high priority route.
2. The Active Cluster Member inspects the traffic and forwards it to the destination.

East-West Traffic

1. Traffic travels from one of the internal servers to the Active Cluster Member.
2. The Active Cluster Member forwards the traffic to the destination.

Intra-Subnet Traffic

Traffic travels freely in the subnet without inspection.

Workflow for Setting Up a High Availability Cluster in GCP

Step 1: Deploy a Template in GCP

Deploy this solution throughout the GCP Portal:

- Check Point CloudGuard High Availability BYOL (Bring Your Own License)

  OR

- Check Point CloudGuard High Availability PAYG (Pay as You Go)

When the template appears, enter information for these parameters:

Parameter	Description
Deployment name	The name of the deployment.
Member A Zone	The zone in which to deploy Member A. The zone determines what computing resources are available and where your data is stored and used. See GCP Regions and Zones documentation for more information.
Member B Zone	The zone in which to deploy Member B, must be in the same region as Member A's zone.
Machines type	The machine type of both Cluster Members. Machine types determine the specifications of your machines, such as the amount of memory, virtual cores, and persistent disk limits an instance will have.
Disk type	The type of disk with which the Cluster Members will be deployed.
Disk size in GB	The size of disk with which the Cluster Members will be deployed.
Public SSH key for the user 'admin'	The public SSH key used to connect to both Cluster Member with the user 'admin'.
Enable Stackdriver monitoring	Select this box if you wish to enable Stackdriver Monitoring.

Parameter	Description
Security Management Server address	The public address of your Security Management Server, in CIDR notation. VPN peers addresses cannot be in this CIDR block, so this value cannot be the zero-address. The Check Point CloudGuard for Google Cloud Platform version R80.30 and above can be managed by a Check Point Security Management Server running version R80.30 and above.
SIC key	The Secure Internal Communication key creates trusted connections between Check Point components. Trust is required to install policies on gateways and to send logs between gateways and Management Servers.
Automatically generate an administrator password	To manage the environment's security, administrators can connect to the Management Server with SmartConsole clients and via the Gaia WebUI using this password. For additional security, it is recommended that you change the password after the deployment is complete.
Allow download from/upload to Check Point	Do not check this box if you do not want to automatically download Software Blade Contracts and other important data. You can improve your experience with the product by sending data to Check Point.
Admin shell	Change the admin shell to enable advanced command line configuration.
Cluster network	The Cluster public IP address will be translated to a private address assigned to the Active member in this external network. Provide an RFC 1918 CIDR block in the Cluster external subnet CIDR field, or select a Network and a Subnetwork below the CIDR field. **Note** - Google Private Access must be enabled for this selected network.
Firewall	These GCP rules will determine if ICMP / TCP / UDP / SCTP / ESP traffic is enabled by default for the provided IP CIDR sources.
Management network	The public IP address used to manage each member. It will be translated to a private address in this external network. Provide an RFC 1918 CIDR block in the Management external subnet CIDR field, or select a Network and a Subnetwork below the CIDR field.
Number of internal networks	This number will determine how many of the networks specified below this field will be used in this deployment.
Internal network (s)	Provide a RFC 1918 CIDR block in the internal subnet CIDR field, or select a Network and a Subnetwork below the CIDR field.

The deployment will take about five minutes to complete. Once completed, useful information will be displayed in the deployment details page, such as the public IP addresses created and the network used for Primary Cluster Synchronization used later in this guide.

Components of the Check Point Solution

The Check Point deployed solution has these components:

- Virtual Private Cloud Networks and subnets:
  - Cluster
  - Management
  - Internal
- Two Virtual Machines configured as a Check Point cluster.
- Public IP addresses for each Cluster Member in Cluster VPC:
  - Primary public IP address
  - Secondary public IP address

  Note - The addresses are not attached to the member's NICs until Step 2 below.
- Public IP addresses for Management to each Cluster Member in Management VPC:
  - Member A address
  - Member B address

Important - No other Virtual Machines can be deployed in the solution's subnet.

Notes about the template:

- Network interfaces can only be configured when instances are created.
- By default, subnets/CIDR are automatically suggested. If you do not delete it, a new VPC network and a subnet within it, in the same region as the cluster members' Zones, will be created.
- If you want to use existing networks instead of specifying IP CIDR blocks, create them with a subnet in the same region as the cluster members Zones before starting this deployment. When deploying, delete the suggested subnet's address in “Cluster/Management/Internal subnet CIDR”, and then choose from the subnet CIDR in your existing VPC. If you specify both, the CIDR field will take precedence.
- The Cluster Members will each have a network interface in each of the subnetworks specified in this deployment. The IP ranges of those subnets must not overlap.
- The Security Cluster will manipulate the routing in the networks you defined as internal in this deployment. The result is that all outbound traffic will go through the cluster member that is currently active.
- It does not deploy any other Virtual Machines in the solution's frontend and backend subnets.
- Virtual Machines that are launched in the backend subnets, may require Internet access to final provisioning. Launch these Virtual Machines only after you have applied Hide NAT rules on the cluster object to support this type of connectivity.
- The Check Point First Time Configuration Wizard automatically deploys after you have set up the cluster object. The cluster object is configured based on the parameters you apply.
- After the First Time Configuration Wizard completes, the Virtual Machines automatically reboot.
- The member operating mode (A or B) is decided independently from the deployment or the order in which the members were added to the cluster. Instead, it is decided by the private IP addresses of their main network interface.

- The cluster's secondary address is used for internet access to the member to which it is attached, while it is in standby mode, so that it will be able to receive important updates. The address is not used in the configuration of the cluster.

Step 2: Configure Cluster Objects in SmartConsole

To configure objects in SmartConsole:

Step	Description
1	If the Security Management Server is deployed in GCP and manages a Cluster Member in a different VPC, then modify the Security Management IP object in SmartConsole to be the public IP of the Management Server. Click **Publish** to apply the change.
2	Click the **Objects** menu > **New** > **More** > **Network Object** > **Gateways & Servers** > **Cluster** > **New Cluster**.
3	Select **Wizard Mode**. The Check Point Installed Gateway Cluster wizard window opens.
4	Enter a Cluster Name. Example: `checkpoint-cluster`
5	In the Cluster IPv4 Address field, enter the cluster IP address (VIP). You can find the cluster IP address in the GCP portal: a. Browse to the **Deployments** page. b. Locate and select the deployment of the HA solution. c. Use the **Cluster IP external address** property.
6	Click **Next**. The **Gateway Cluster Properties** window opens.

Step	Description
7	Click **Add new cluster member**. a. In the **Name** field, enter the first Cluster Member's name. Example: `member1` b. In the IPv4 address field: Enter the Cluster Member's public IP address from Management VPC (Member A external IP and Member B external IP found in the deployment details page). c. In the **Activation Key** field, enter the SIC key (set up in GCP). d. In the **Confirm Activation Key** field, enter the SIC key again. e. Click **Initialize**. If the Activation Key is confirmed, the **Trust State** field shows **Trust Established**. f. Click **OK**.
8	Repeat the Step 7 to add the second Cluster Member.
9	Click **Next**. The **Cluster Topology** window opens.
10	Select the subnetwork provided in the **Management network** field during the deployment as the **Primary Cluster Synchronization network**.
11	Configure the other subnetworks as **Private use** for each member.
12	Complete the wizard, and then click **Publish** to save the settings.
13	Open the Cluster object.
14	In the **Network management** tab, disable **Anti-Spoofing** for all interfaces by editing those interfaces in the cluster object.
15	Click **OK**.
16	The IPsec VPN blade is automatically enabled. To use the VPN blade, see *"Step 6: Configure VPN" on page 20*. Otherwise, disable the VPN blade.
17	Install the applicable Access Control Policy on the cluster object.

A few minutes after the applicable Access Control Policy is installed, the following changes occur automatically in GCP:

1. The following public IP addresses for each Cluster Member in Cluster VPC will be attached:
   - Primary public IP address.
   - Secondary public IP address.
2. In each of the internal VPC networks, a GCP Route routes all outbound traffic (`0.0.0.0/0`) to the Active member with high priority (1) and to the secondary member with lower priority (2).

Note - For the failover process to function, each Cluster Member initiates outbound HTTP and HTTPS traffic which is allowed by the gateway's implied rules. Do not override these implied rules.

Step 3: Enable Outbound Traffic

To enable outbound traffic:

Step	Description
1	From SmartConsole, connect to the Security Management Server.
2	Find the Security Cluster object in the **Gateways & Servers** tab.
3	Select the **NAT** tab.
4	Check the **Hide internal networks behind the Gateway's external IP** check box.
5	Click **OK**.
6	Install policy.

Note - NAT does not support Connection synchronization during failover. If you configure the cluster to always hide the internal networks (by selecting to automatically add address translation rules - instead of the option described above) this will prevent connection synchronization in additional use cases, such as East-West traffic between internal VPCs.

Step 4: Create Object `LocalGatewayExternal` in SmartConsole

In SmartConsole, create the Dynamic object called `LocalGatewayExternal`.

This object represents the private Cluster Member's IP addresses.

Note - You will use this Dynamic object in step 6.

Step 5: Configure Inbound Protection

Overview

- You will need to configure Access Control and NAT rules for North-South inbound traffic.
- You can configure the Cluster's External IP to listen on the TCP port 443, and forward this traffic to the Active Cluster Member. The Active Cluster Member will then inspect the traffic and forward it to the Application server on the TCP port 8084.
- The Active Cluster Member uses NAT to forward traffic, that belongs to the two web applications, to the appropriate web server.
- NAT rules are defined with the special Dynamic Object.

- The Dynamic object `LocalGatewayExternal` represents the private IP addresses of the external interface of Member 1 and Member 2.
- For more information, see *"Step 4: Create Object LocalGatewayExternal in SmartConsole" on the previous page*.

To configure Inbound Protection:

Step	Description
1	Connect with SmartConsole to your Security Management Server.
2	Create a host object to represent: ■ The specific host that you want to access through the Internet.
3	Create a new TCP service. Do the following for each internal port, such as port 8081. Follow these steps: a. Click the **Objects** menu > **More object types** > **Service** > **New TCP**. b. Enter a descriptive name. For example: `http-8081` c. In the **Protocol** field, select the applicable protocol (such as, HTTP or HTTPS). d. In the **Port** field, select **Customize** and enter the port number. For example: `8081` e. Click **OK**.

Configure Access Control and NAT rules for North-South inbound traffic by using the following NAT rules:

Create a NAT rule with these values.

NAT Rule	Value
Rule No	1
Original Source	`All_Internet` (do not use `*Any`)
Original Destination	LocalGatewayExternal
Original Services	The service object that represents the internal port
Translated Source	`Original`
Translated Destination	The Host object that represents your web server

NAT Rule	Value
Translated Services	The service object that represents the port on which the Active member listens (for example, http)
Install On	`*Policy Targets`

About this NAT rule:

- Matches any traffic that arrives at the CloudGuard Security Gateway on the applicable internal port.
- Translates the destination IP address to the IP address of the Web Servers.

Step 6: Configure VPN

For more information, see the Check Point Security Management Administration Guide for your Management Server version R80.30.

To configure a VPN:

Step	Description
1	Create a Network Group object to represent the encryption domain of the cluster: a. In SmartConsole, click the **Objects** menu > **Object Explorer**. b. From the top toolbar, select **New** > **Network Group**. c. In the **Enter Object Name** field, enter the desired name. d. Click the **+** icon and select the applicable network objects. e. Click **OK**. f. Close the **Object Explorer**.
2	Edit the cluster object: a. In SmartConsole, from the left navigation panel, click **Gateways & Servers**. b. Double-click the cluster object. The **Gateway Cluster Properties** window shows.
3	Define your Network Group as the encryption domain of the cluster object: a. In SmartConsole, from the left navigation panel, click **Gateways & Servers**. b. Double-click the cluster object. The **Gateway Cluster Properties** window shows. c. In the cluster object left tree, click **Network Management** > **VPN Domain**. d. Select **Manually defined**. e. In the right corner of this field, click the **[...]** button and select the Network Group object created in Step 1.

Step	Description
4	Define the VPN community: a. In the cluster object left tree, click **IPsec VPN**. b. In the section **This Security Gateway participates in the following VPN Communities**, select the applicable VPN community.
5	Define the outgoing VPN interface: a. In the cluster object left tree, click **IPsec VPN** > **Link Selection.** b. In the **IP Selection by Remote Peer** section, select **Always use this IP address** > **Main IP**. c. In the Outgoing Route Selection section: i. Click **Source IP address settings**. ii. Select **Manual**. iii. Choose **IP address of chosen interface**. d. In the **Tracking** field, select the desired option. e. To close the **Gateway Cluster Properties** window, click **OK**.
6	Configure the VPN Community to use Permanent Tunnels: a. In SmartConsole, click the **Objects** menu > **Object Explorer**. b. In the left tree, clear all boxes except for **VPN Communities**. c. Double-click the VPN community, in which this cluster object participates. The **VPN Community** window shows. d. In the left tree, click **Tunnel Management**. e. Select **Set Permanent Tunnels**. f. Select the applicable option. g. To close the VPN Community properties window, click **OK**. h. Close the **Object Explorer**.
7	Install the applicable Access Control Policy on the cluster object.

Additional Information

Testing and Troubleshooting

You can use the APIs to retrieve information about the cluster resource group.

Use these commands on each Cluster Member to confirm that the cluster operates correctly.

Note - Run these commands from Gaia Clish, or the Expert Mode.

```
cphaprob state
cphaprob -a if
```

Example:

```
[Expert@HostName:0]# cphaprob state
Cluster Mode: High Availability (Active Up) with IGMP Membership
Number Unique Address Assigned Load State
1 (local) 10.0.1.10 0% Active
2 10.0.1.20 100% Standby
```

Use the cluster configuration test script on each Cluster Member to confirm that it is configured correctly.

The script verifies:

- The configuration file is defined in `$FWDIR/conf/gcp-ha.json`.
- A Primary DNS server is configured and works.
- The machine is set up as a Cluster Member.
- IP forwarding is enabled on all network interfaces of the Cluster Member.
- Calibration of ClusterXL configuration for GCP.

To verify the configuration, run the following script on each Cluster Member:

Step	Description
1	Connect to the command line.
2	Log in to Expert mode.

Step	Description
3	Run the script with this command (do not change the syntax): `# $FWDIR/scripts/gcp_ha_test.py` If all tests were successful, you will see the following message: *All tests were successful!* Otherwise, an error message is displayed with information about how to troubleshoot the problem.

Common configuration errors:

Message	Recommendation
`The attribute (ATTRIBUTE) is missing in the configuration`	Verify that the configuration file is correct.
`Primary DNS server is not configured Failed to resolve (host)`	The Cluster Member is not configured with a DNS server.
`Failed in DNS resolving test`	Confirm that DNS resolution on the Cluster Member works.
`You do not seem to have a valid cluster configuration`	Make sure that the Cluster Member configuration on the Check Point Security Management Server is complete, and that the Security Policy is installed
`IP forwarding is not enabled on Interface (Interface-name)`	Use PowerShell to enable IP forwarding on all the network interfaces of the Cluster Member.
`Failed to read configuration file: $FWDIR/conf/gcp-ha.json`	The GCP Cluster Member configuration is not up to date, or is written incorrectly.
`Testing credentials`	Failed to login with the credentials provided. See the exception text to understand why.
`Testing authorization (Exception)`	Make sure the GCP daemon has access to GCP.

Simulate a cluster failover:

For example, shut down the internal interface of the Active Cluster Member.

- On the current Active Cluster Member, run from the Expert Mode:
  ```
  # ip link set dev eth1 down
  ```

- After a few seconds, the second Cluster Member has to report itself as the Active Cluster Member. Examine the cluster state on each Cluster Member. Run from Gaia Clish, or Expert Mode:

```
cphaprob state
```

- On the former Active Cluster Member, run from Expert Mode:

```
# ip link set dev eth1 up
```

If you experience issues:

To make the networking changes automatically, the Cluster Members have to communicate with GCP. This requires HTTPS connections over TCP port 443 to the GCP end points. Make sure that the Security Policy that is installed on the Cluster Members allows for this type of communication.

Using the GCP High Availability Daemon

The cluster solution in GCP uses the daemon to make API calls to GCP when a cluster failover takes place. This daemon uses a configuration file, `$FWDIR/conf/gcp-ha.json`, on each Cluster Member.

When you deploy the above solution from the supplied template , a configuration file is automatically created .

The configuration file is in JSON format and contains these attributes:

Attribute's name	Type	Value
`debug`	Boolean	True or False
`public ip`	String	Name of the cluster's external, primary public IP address
`secondary public ip`	String	Name of the cluster's external, secondary public IP address
`dest ranges`	String	IP range for updating

You can confirm that the daemon in charge of communicating with GCP runs on each Cluster Member.

From Expert Mode, run:

```
# cpwd_admin list | grep -E "PID|GCP_HAD"
```

The output should be similar to this example:

APP	PID	STAT	#START	START_TIME	MON	COMMAND
GCP_HAD	3663	E	1	[12:58:48] 15/1/2016	N	python /opt/CPsuite
R80.30/fw1/scripts/gcp_had.py						

Notes:

- The script appears in the output.
- The `STAT` column should show **E** (executing).

• The `#START` column should show **1** (the number of times this script was started by the Check Point Watchdog).

To troubleshoot issues related to this daemon, generate debug.

From Expert Mode, run:

```
# vi $FWDIR/conf/gcp-ha.json
```

To enable debug printouts, edit the file according to the following:

```
#{

  "debug": true,
   "public_ip": "XXX-primary-cluster-address",
   "secondary_public_ip": "XXX-secondary-cluster-address",
   "dest_ranges": ["0.0.0.0/0"]

}
```

To disable debug printouts, edit the file according to the following:

```
#{

   "debug": False,
   "public_ip": "XXX-primary-cluster-address",
   "secondary_public_ip": "XXX-secondary-cluster-address",
   "dest_ranges": ["0.0.0.0/0"]

}
```

To load the configuration:

Step	Description

1	Kill the GCP daemon by running:
	`# ps aux | grep had`
	`# kill HAD process`

2	Make sure the process is running.

3	From Expert mode run:
	`# cpwd_admin list |`
	`grep -E "PID|GCP_HAD"`

The debug output is written to `$FWDIR/log/gcp_had.elg*` files.

Creating Objects in SmartConsole

For more information, see the Check Point Security Management Administration Guide for your Management Server version (R80.30).

Important - After you create an object, you must publish the session to save the changes in the management database.

To create a Host object:

Step	Description
1	From the top right **Objects Pane**, click **New** > **Host**. The **New Host** window shows.
2	In the **Machine** field, enter the private IP address of the machine.

To create a Network object:

Step	Description
1	From the top right **Objects Pane**, click **New** > **Network**. The **New Network** window shows.
2	Enter the **Object Name** (specifically the subnet name).
3	Enter the **Network address** and **Net mask**.

To create a Service (port) object:

Step	Description
1	From the top right **Objects Pane**, click **New** > **More** > **Service**.
2	Select your TCP/UDP service.
3	Enter the **Object name**.
4	In the **Enter Object Comment** field, enter the port name.
5	In the **General** field, select your **Protocol**.
6	In the **Match By** field, select the **Port** number.
7	Click **OK**.

Upgrading a Check Point CloudGuard IaaS High Availability Solution to a Newer Version

Use the following instructions to upgrade a deployed Check Point CloudGuard IaaS High Availability solution to a newer version.

Note:

During the upgrade process, a new Check Point CloudGuard IaaS High Availability solution is deployed. The upgrade will maintain the network configurations used in the original Check Point CloudGuard IaaS High Availability solution.

Two key terms to remember:

Source - The original solution (with the lower version)

Target - The new deployed solution (with the higher version)

Step-by-step instructions for upgrading to a new version:

Step	Description
1	Log in to the GCP portal.
2	Open the source CloudGuard High Availability instances (member-a and member-b): a. In the active member page: Locate the primary cluster address (nic0 External IP) and copy its name for future reference ('XXX-primary-cluster-address'). b. In the stand-by member page: Locate the secondary cluster address (nic0 External IP) and copy its name for future reference ('XXX-secondary-cluster-address').
3	Deploy a new Check Point CloudGuard IaaS High Availability solution (this is the "target solution"). a. Under **High Availability Version**, choose the version. b. Under **Instance Configuration**, choose the same configurations as in the Source solution. c. Under **Check Point**, choose the same configurations as in the Source solution. d. Under **Networking**, choose the same network configurations as in the Source solution, such as Cluster external subnet, Management external subnet, and internal networks.

Step	Description

4	Adjust the configuration file of the target solution instances to match the Source solution's external IP addresses.
	
	For both instances of the target solution:
	
	a. Log in to SSH.
	b. From Expert Mode, run:
	`# vi $FWDIR/conf/gcp-ha.json`
	c. Edit the file so to match the following lines:
	`"public_ip": "<primary cluster address name(copied in 2.a)>", "secondary_public_ip": "<secondary cluster address name(copied in 2.b)>",`
	
	Keep the other lines in the file the same.
	
	Note - the separating commas at the end of each line.
	
	d. Save the file and exit.

	Important - Connectivity loss will occur during the next steps.
5	Delete routes from the cluster's internal networks manually: Go to the **Navigation menu** > **NETWORKING** > **VPC network** > **VPC networks**. Do the following for each internal network in the solution: a. Choose the internal network. b. Choose **Routes**. c. Delete these routes: ■ Start with **"x-chkp"** and ends with **"to-member-a"** (if exists, this depends on the identity of the current active member). ■ Start with **"x-chkp"** and ends with **"to-member-b"**.
6	Release the primary and secondary IP addresses. From the **Navigation** menu > **NETWORKING** > **VPC network** > **External IP addresses**. a. Locate the primary cluster address name (see 2.a above) and the secondary cluster address name (see 2.b above). b. For both IP addresses: i. Select **Change** The Attach IP address window will appear ii. Under **Attach to**, choose **None**, and then clear the **Assign a new ephemeral IP address** box.

Step	Description
7	Configure in SmartConsole In **Gateways & Servers**, click twice on the cluster object and edit the following: a. Under **General Properties**, choose the new version (the version of the Target solution created in step 3). b. Under **Cluster Members**, update members to match the members of the Target solution: For each member update the **IPv4 Address** (management - the network's external IP). c. Under **Network management** - modify the interfaces to match the Target solution members. d. Install policy on the cluster.
	Note - At this point, and after the all new routes and IP addresses configurations are completed, the *Target* CloudGuard IaaS High Availability handles all the traffic in the environment (such as, inbound, outbound, E-W, and VPN tunneling). Verify that all the traffic flows work as expected (you can also check for failover) before proceeding.
8	Delete your source CloudGuard IaaS High Availability instances and release redundant IP addresses. **Important - Do not delete the entire deployment of the source solution since the Target solution uses the primary and secondary IP addresses.**

Related Solutions

- sk109360 - Check Point Reference Architecture for Google Cloud Platform

Known Limitations

- Setting service principal that uses a certificate credential is not supported on Check Point CloudGuard IaaS High Availability.
- Only two Cluster Members in a cluster are supported.
- Only High Availability Mode (Active/Standby) is supported. Load Sharing Mode is not supported.
- VRRP Cluster is not supported.
- For outbound and traffic, you cannot delete or disable the public IP addresses of Cluster Members.
- Working with a Proxy is not supported.
- East-West traffic between different subnets in the same VPC is not supported.