



EXPRESSCLUSTER X 5.2
HA Cluster Configuration Guide for Google Cloud (Linux)
Release 2

NEC Corporation

Nov 29, 2024

TABLE OF CONTENTS:

1	Preface	1
1.1	Who Should Use This Guide	1
1.2	Scope of Application	2
1.3	How This Guide is Organized	3
1.4	EXPRESSCLUSTER X Documentation Set	4
1.5	Conventions	5
1.6	Contacting NEC	6
2	Overview	7
2.1	Functional overview	7
2.2	Basic configuration	8
2.3	Network partition resolution	13
2.4	Differences between on-premises and Google Cloud	14
3	Operating Environments	15
3.1	HA cluster with Cloud DNS	15
4	Cluster Creation Procedure (for an HA Cluster with an Internal TCP Load Balancer)	17
4.1	Creation example	17
4.2	Configuring Google Cloud	22
4.3	Configuring EXPRESSCLUSTER	24
4.4	Verifying the created environment	27
5	Cluster Creation Procedure (for an HA Cluster with Cloud DNS)	29
5.1	Creation example	29
5.2	Configuring Google Cloud	32
5.3	Configuring EXPRESSCLUSTER	35
5.4	Verifying the created environment	38
6	Error Messages	39
7	Notes and Restrictions	41
7.1	HA cluster with a load balancer	41
7.2	HA cluster with Cloud DNS	43
8	Legal Notice	45
8.1	Disclaimer	45
8.2	Trademark Information	46
9	Revision History	47

1.1 Who Should Use This Guide

The *HA Cluster Configuration Guide for Google Cloud (Linux)* is intended for administrators who want to build a cluster system, and for system engineers and maintenance personnel who provide user support.

The software and setup examples introduced in this guide are for reference only, and the software is not guaranteed to run.

This guide is accompanied by sample scripts for Cloud DNS.

Check and modify the sample scripts according to your environment.

1.2 Scope of Application

For information on the system requirements, see "Getting Started Guide" -> "Installation requirements for EXPRESS-CLUSTER".

This guide contains product- and service-related information (e.g., screenshots) collected at the time of writing this guide. For the latest information, which may be different from the content in this guide, refer to corresponding websites and manuals.

1.3 How This Guide is Organized

- "2. *Overview* ": Describes the functional overview.
- "3. *Operating Environments* ": Describes the tested operating environment of this function.
- "4. *Cluster Creation Procedure (for an HA Cluster with an Internal TCP Load Balancer)*": Describes how to create an HA cluster involving an internal TCP load balancer.
- "5. *Cluster Creation Procedure (for an HA Cluster with Cloud DNS)* ": Describes how to create an HA cluster involving Cloud DNS.
- "6. *Error Messages*": Describes the error messages and solutions.
- "7. *Notes and Restrictions* ": Describes the notes and restrictions on creating and operating a cluster.

1.4 EXPRESSCLUSTER X Documentation Set

The EXPRESSCLUSTER X manuals consist of the following five guides. The title and purpose of each guide is described below:

EXPRESSCLUSTER X Getting Started Guide

This guide is intended for all users. The guide covers topics such as product overview, system requirements, and known problems.

EXPRESSCLUSTER X Installation and Configuration Guide

This guide is intended for system engineers and administrators who want to build, operate, and maintain a cluster system. Instructions for designing, installing, and configuring a cluster system with EXPRESSCLUSTER are covered in this guide.

EXPRESSCLUSTER X Reference Guide

This guide is intended for system administrators. The guide covers topics such as how to operate EXPRESSCLUSTER, function of each module and troubleshooting. The guide is supplement to the Installation and Configuration Guide.

EXPRESSCLUSTER X Maintenance Guide

This guide is intended for administrators and for system administrators who want to build, operate, and maintain EXPRESSCLUSTER-based cluster systems. The guide describes maintenance-related topics for EXPRESSCLUSTER.

EXPRESSCLUSTER X Hardware Feature Guide

This guide is intended for administrators and for system engineers who want to build EXPRESSCLUSTER-based cluster systems. The guide describes features to work with specific hardware, serving as a supplement to the Installation and Configuration Guide.

1.5 Conventions

In this guide, Note, Important, See also are used as follows:

Note: Used when the information given is important, but not related to the data loss and damage to the system and machine.

Important: Used when the information given is necessary to avoid the data loss and damage to the system and machine.

See also:

Used to describe the location of the information given at the reference destination.

The following conventions are used in this guide.

Convention	Usage	Example
Bold	Indicates graphical objects, such as text boxes, list boxes, menu selections, buttons, labels, icons, etc.	Click Start. Properties dialog box
Angled bracket within the command line	Indicates that the value specified inside of the angled bracket can be omitted.	<code>clpstat -s[-h <i>host_name</i>]</code>
#	Prompt to indicate that a Linux user has logged on as root user.	<code># clpstat</code>
Monospace	Indicates path names, commands, system output (message, prompt, etc.), directory, file names, functions and parameters.	<code>/Linux</code>
bold	Indicates the value that a user actually enters from a command line.	Enter the following: <code># clpcl -s -a</code>
<i>italic</i>	Indicates that users should replace italicized part with values that they are actually working with.	<code># ping <IP address></code>



In the figures of this guide, this icon represents EXPRESSCLUSTER.

1.6 Contacting NEC

For the latest product information, visit our website below:

<https://www.nec.com/en/global/prod/expresscluster/>

OVERVIEW

2.1 Functional overview

This guide describes how to create an HA cluster based on EXPRESSCLUSTER X (EXPRESSCLUSTER) on the cloud service of Google Cloud.

Google Cloud allows you to use regions and zones to create an HA cluster with virtual machines, increasing the business availability.

- Region

On Google Cloud, a region is a division: a physical and logical unit (like New York and London).

It is possible to build all nodes in a single region. However, a network failure or a natural disaster may make all of them crash to prevent the business from continuing.

To increase availability, distribute nodes to multiple regions.

A region is a group of zones.

- Zone

On Google Cloud, a zone is a logical group to which each node can be distributed.

By distributing each node to a different zone, you can minimize the effects of planned Google Cloud maintenance and those of unplanned maintenance due to a physical hardware failure.

For more information on regions and zones, refer to the following:

Regions and zones:

<https://cloud.google.com/compute/docs/regions-zones/>

2.2 Basic configuration

This guide assumes two HA clusters (uni-directional standby cluster configurations): one with a load balancer and the other with Cloud DNS.

For each of the HA clusters, the following EXPRESSCLUSTER resource and Google Cloud service are to be used:

Purpose	EXPRESSCLUSTER resource to be chosen	Necessary Google Cloud service
Accessing the cluster at a virtual IP address from a client	Google Cloud virtual IP resource	Internal TCP load balancing
Accessing the cluster in a DNS name from a client	EXEC resource (to be used with the scripts, which accompanies this guide, for Cloud DNS)	Cloud DNS

HA cluster with a load balancer

For virtual machines in a Google Cloud environment, client applications can use a virtual IP (VIP) address to access nodes that constitute a cluster. Using the VIP address eliminates the need for the clients to be aware of switching between the virtual machines even after a failover or a group migration occurs.

As [Fig. 2.1 HA cluster with an internal TCP load balancer](#) shows, the cluster in the Google Cloud environment can be accessed by specifying the VIP address (front-end IP address for Cloud Load Balancing) of the Google Cloud load balancer (for Cloud Load Balancing).

The Google Cloud load balancer switches between the active server and the standby server, with its health check. The health check is performed through a port provided by the Google Cloud virtual IP resource.

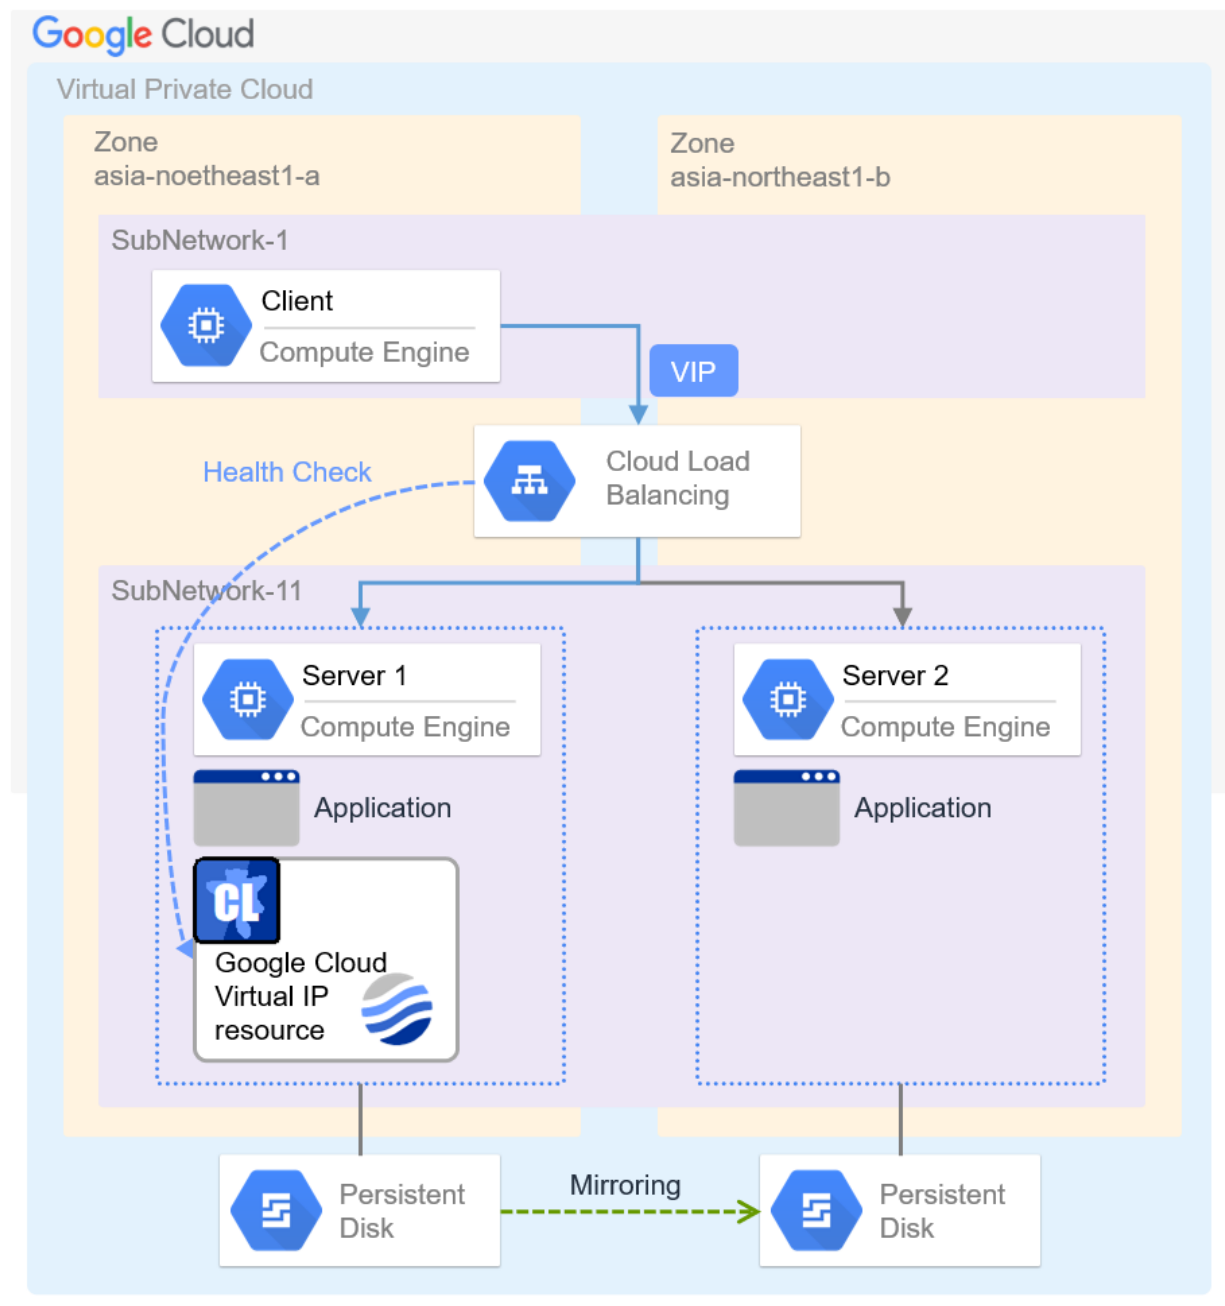


Fig. 2.1: HA cluster with an internal TCP load balancer

SubNetwork-1	10.0.1.0/24
IP Address (Client)	10.0.1.200
Virtual IP Address (VIP)	10.0.1.100
SubNetwork-11	10.0.11.0/24
IP Address (Server1)	10.0.11.101
IP Address (Server2)	10.0.11.102
Health Check Port	12345

For more information on Cloud Load Balancing, refer to the following:

Cloud Load Balancing:

<https://cloud.google.com/load-balancing/>

The following is an example of an HA cluster with a load balancer:

Purpose	Load balancer to be used	Creation procedure
Sharing business within the Google Cloud network	Internal TCP load balancer	See "4. <i>Cluster Creation Procedure (for an HA Cluster with an Internal TCP Load Balancer)</i> " of this guide.

An HA cluster configuration with a load balancer requires the following resources and monitor resources. For a network partition resolution resource, which is not set in this guide, see "2.3. *Network partition resolution*" to determine how it should be depending on your system configuration.

Resource/monitor resource	Description	Setting
Google Cloud virtual IP resource	Provides a mechanism for awaiting access from the load balancer to a specific port for the alive monitoring (health check)--over the port, an application works on a node. At activation, this resource starts up a control process to await access from the Google Cloud load balancer for its alive monitoring. At deactivation, the resource stops the control process.	Required
Google Cloud virtual IP monitor resource	Performs alive monitoring of a control process to be started in activating a Google Cloud virtual IP resource, for the node where the Google Cloud virtual IP resource is started.	Required
Google Cloud load balance monitor resource	Checks whether the same port as that for the health check is opened, for a node where the Google Cloud virtual IP resource is not started.	Required
Other resources and monitor resources	Depends on the configuration of the application (such as a mirror disk) to be used in the HA cluster.	Optional

HA cluster with Cloud DNS

This configuration uses Cloud DNS to make an HA cluster accessible via a DNS name.

Each of the two virtual machines is placed in a different zone in order to minimize the effects of planned Google Cloud maintenance and those of unplanned maintenance due to a physical hardware failure.

In Fig. 2.2 HA cluster with Cloud DNS, the cluster can be accessed by specifying its DNS name.

EXPRESSCLUSTER manages the record sets (A records) of Cloud DNS so that a given IP address can be found according to the DNS name.

This eliminates the need for the client to be aware of switching between the virtual machines even after a failover or a group migration occurs.

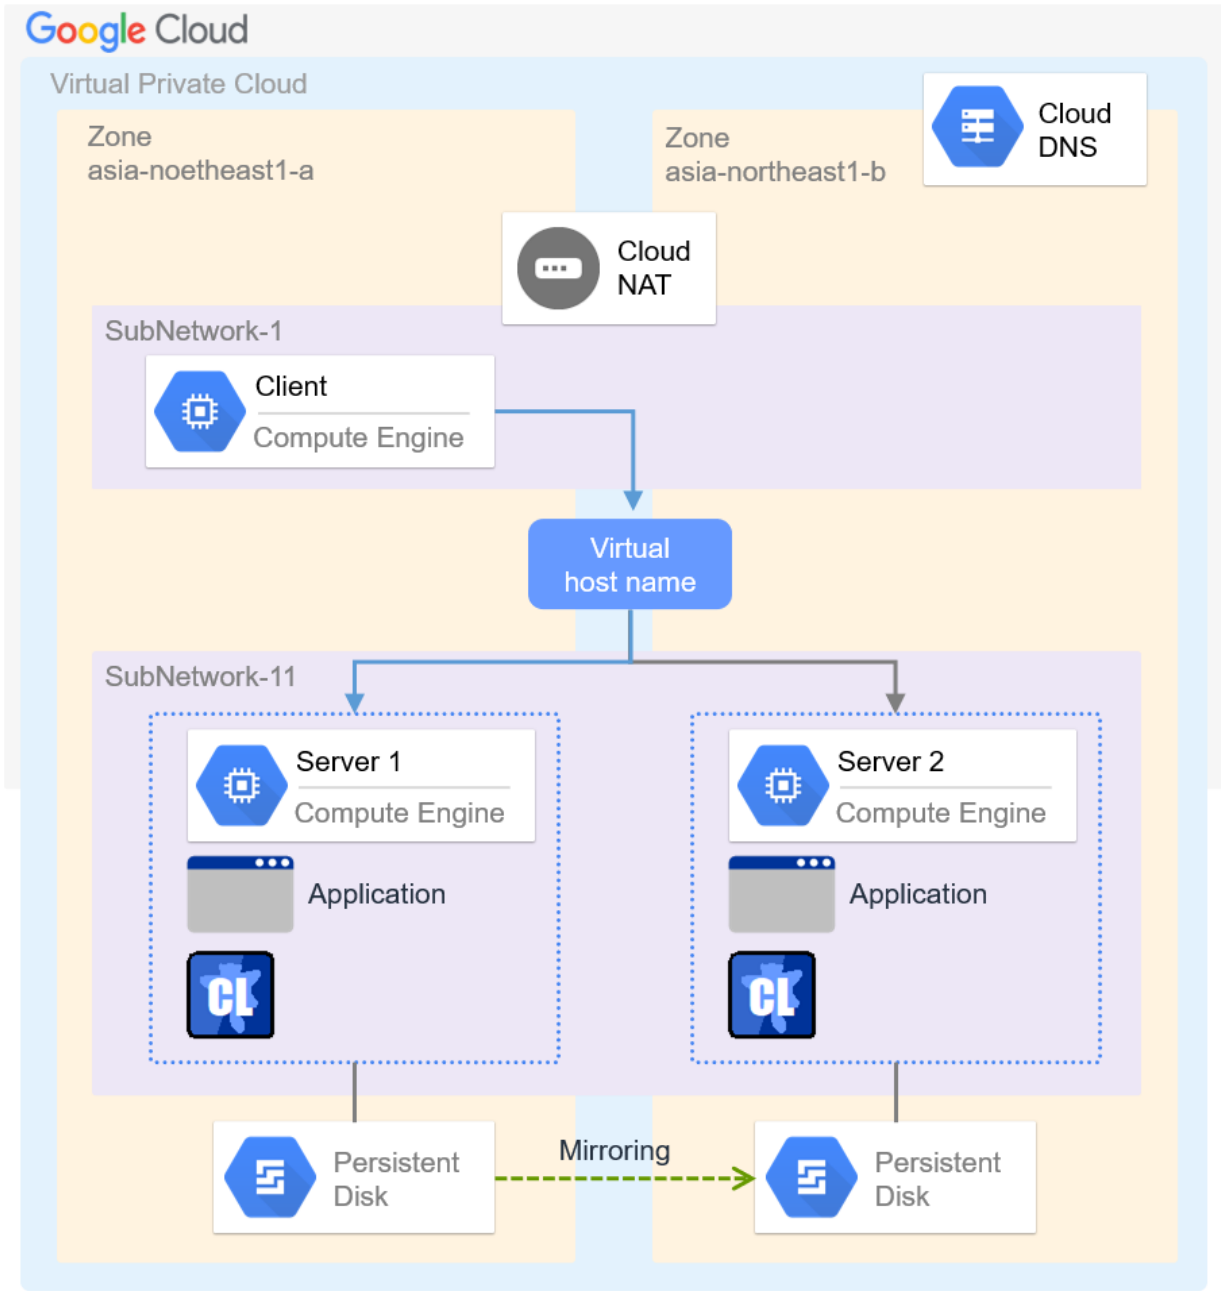


Fig. 2.2: HA cluster with Cloud DNS

SubNetwork-1	10.0.1.0/24
IP Address (Client)	10.0.1.200
SubNetwork-11	10.0.11.0/24
IP Address (Server1)	10.0.11.101
IP Address (Server2)	10.0.11.102
Virtual host name	test-cluster.example.com

1. Before updating the record (Cloud DNS)

Name	Type	Data
test-cluster.example.com	A	10.0.11.101

2. After updating the record (Cloud DNS)

Name	Type	Data
test-cluster.example.com	A	10.0.11.102

This guide describes a configuration in which instances constituting the HA cluster within the VPC network can access the Internet.

The instances need to access the Internet when executing the gcloud CLI used with a script for Cloud DNS.

The Internet can be accessed through a NAT gateway or a NAT instance. This guide uses the former (Cloud NAT).

For more information on the procedures for creating the VPC network, configuring the NAT gateway/NAT instance, and configuring the firewall rules, refer to the documents of Google Cloud.

An HA cluster configuration with Cloud DNS requires the following resources and monitor resources.

For a network partition resolution resource, which is not set in this guide, see "[Network partition resolution](#)" to determine how it should be depending on your system configuration.

Resource/monitor resource	Description	Setting
EXEC resource	Manages the record sets (A records) of Cloud DNS so that a given IP address can be found according to the DNS name.	Required
Other resources and monitor resources	Depends on the configuration of the application (such as a mirror disk) to be used in the HA cluster.	Optional

2.3 Network partition resolution

Virtual machines constituting an HA cluster mutually perform alive monitoring with heartbeat.

If heartbeat ceases with each virtual machine existing on a different subnet, an undesirable event (such as double-launching a service) occurs.

Its prevention requires determining whether any of the virtual machines has been isolated from the network, in a network partition (NP) state, or another virtual machine has crashed.

The NP resolution feature pings a device which always operates and is expected to respond. If there is no reply, the feature considers the device to be in an NP state and takes a specified measure (such as warning, recovery, and server shutdown).

The target and method of NP resolution needs to be individually considered, in accordance with the locations of clients accessing the cluster system and with the conditions for connecting to an on-premise environment (e.g. using a leased line).

There is no recommended target or method of NP resolution.

For more information on NP resolution, refer to the following:

- "Installation and Configuration Guide" -> "Understanding network partition resolution resources"
- "Reference Guide" -> "Heartbeat resources details"
- "Reference Guide" -> "Network partition resolution resources details"

2.4 Differences between on-premises and Google Cloud

The following table describes the functional differences of EXPRESSCLUSTER between on-premises and Google Cloud. "✓" indicates that the relevant function can be used and "n/a" indicates that the relevant function cannot be used.

Function	On-premise	Google Cloud
Creating a shared disk type cluster	✓	n/a
Creating a mirror disk type cluster	✓	✓
Using the management group	✓	n/a
Using the floating IP resource	✓	n/a
Using the virtual IP resource	✓	n/a
Using the Google Cloud virtual IP resource	n/a	✓
Using the Google Cloud DNS resource	n/a	✓

There is no difference in the procedure for creating a cluster between an on-premise environment and an Google Cloud environment except that Google Cloud needs to be configured in advance.

OPERATING ENVIRONMENTS

3.1 HA cluster with Cloud DNS

The following configuration has been confirmed to operate:

x86_64

OS	Red Hat Enterprise Linux 7.6
CLUSTERPRO	CLUSTERPRO X 4.2 for Linux (internal version: 4.2.0-1) CLUSTERPRO X Replicator 4.2 for Linux
Region	asia-northeast1 (Tokyo)
Mirror disk size	Disk size: 20 GB (1 GB for the cluster partition; 19 GB for the data partition)
Cloud SDK(gcloud CLI)	245.0.0

The scripts for Cloud DNS are to be used with the Cloud DNS service.

For more information on Cloud DNS, refer to the following:

Cloud DNS:

<https://cloud.google.com/dns/>

The scripts for Cloud DNS uses the gcloud CLI, which is part of Cloud SDK, to control Cloud DNS.

For more information on Cloud SDK, refer to the following:

Google Cloud SDK documentation:

<https://cloud.google.com/sdk/docs/>

For more information on the system requirements of Cloud SDK, refer to the following:

Installing Google Cloud SDK:

<https://cloud.google.com/sdk/install/>

CLUSTER CREATION PROCEDURE (FOR AN HA CLUSTER WITH AN INTERNAL TCP LOAD BALANCER)

4.1 Creation example

This guide describes how to create a two-node, uni-directional standby cluster with EXPRESSCLUSTER.

Through this procedure, you are to create an HA cluster accessible from clients within the same VPC network on Google Cloud.

This procedure is intended to create a mirror-disk type configuration in which server1 is used as the active server.

The following tables list parameters which do not have their default values, and parameters whose values are changed from their default values.

Of the firewall rules, **IP range** is necessary for allowing communication from the Google Cloud health check system (130.211.0.0/22, 35.191.0.0/16).

- Google Cloud settings (common to server1 and server2)

Item	Value
Configuration of the VPC network	
- Name	test-vpc
- New subnet (name)	subnetwork-1, subnetwork-11
- New subnet (region)	asia-northeast1
- New subnet (IP address range)	10.0.1.0/24, 10.0.11.0/24
Configuration of the firewall rules	
- Name	test-allow-health-check

Continued on next page

Table 4.1 – continued from previous page

Item	Value
– Network	test-vpc
– Traffic direction	Upstream
– Action in response to agreement	Allow
– Target	Specified target tag
– Target tag	test-allow-health-check
– Source filter	IP range
– Source IP range	130.211.0.0/22, 35.191.0.0/16
– Specified protocol and port	Allow all
Configuration of the load balancer	
– Type	TCP load balancing
– For internet connection or private use	Only between VMs
– Multi- or mono-region	Only mono-region
– Name	test-lb
Configuration of the load balancer (back end)	
– Region	asia-northeast1
– Network	test-vpc
– Instance group	test-ig-a, test-ig-b
– Health check (name)	test-health-check

Continued on next page

Table 4.1 – continued from previous page

Item	Value
– Health check (protocol)	TCP
– Health check (port)	12345
– Health check (proxy protocol)	None
– Session affinity	None
Configuration of the load balancer (front end)	
– Name	test-frontend
– Subnetwork	subnetwork-1
– Internal IP address	10.0.1.100
– Port	80 (number of the port through which the application is available)

- Google Cloud settings (to be set separately on server1 and server2)

Item	Value	
	server1	server2
Configuration of the instances		
– Region	asia-northeast1	asia-northeast1
– Zone	asia-northeast1-a	asia-northeast1-b
– New disk	server1-datadisk-0	server2-datadisk-0
Configuration of the instance groups		
– Name	test-ig-a	test-ig-b
– Group type	Unmanaged instance group	Unmanaged instance group

Continued on next page

Table 4.2 – continued from previous page

Item	Value	
	server1	server2
– Region	asia-northeast1	asia-northeast1
– Zone	asia-northeast1-a	asia-northeast1-b
– Network	test-vpc	test-vpc
– Subnetwork	subnetwork-11	subnetwork-11
– VM instance	server1	server2
Network configuration		
– Network	test-vpc	test-vpc
– Subnetwork	subnetwork-11	subnetwork-11
– Internal IP address	10.0.11.101	10.0.11.102

- EXPRESSCLUSTER settings (cluster properties)

Item	Value	
	server1	server2
– Cluster name	Cluster1	Cluster1
– Server name	server1	server2

- EXPRESSCLUSTER settings (failover group)

Resource name	Item	Value
Mirror disk resource	Resource name	md1
Mirror disk resource	Details tab - mirror partition device name	/dev/NMP1
Mirror disk resource	Details tab - mount point	/mnt/md1
Mirror disk resource	Details tab - data partition device name	/dev/sdb2
Mirror disk resource	Details tab - cluster partition device name	/dev/sdb1

Continued on next page

Table 4.4 – continued from previous page

Resource name	Item	Value
Mirror disk resource	Details tab - file system	ext4
Google Cloud virtual IP resource	Resource name	gcvip1
Google Cloud virtual IP resource	Port number	12345: as specified for Health check (port)

- EXPRESSCLUSTER settings (monitor resource)

Monitor resource name	Item	Value
Mirror disk monitor resource	Monitor resource name	mdw1
Mirror disk connect monitor resource	Monitor resource name	mdnw1
Google Cloud virtual IP monitor resource	Monitor resource name	gcvipw1
Google Cloud virtual IP monitor resource	Recovery target	gcvip1
Google Cloud load balance monitor resource	Monitor resource name	gclbw1
Google Cloud load balance monitor resource	Recovery target	gcvip1

4.2 Configuring Google Cloud

1. Creating the VPC network

Access Google Cloud Console (<https://console.cloud.google.com/>).

Create the VPC network and subnets.

For more information on the procedure, refer to the following:

Using VPC networks:

<https://cloud.google.com/vpc/docs/using-vpc/>

2. Creating the instances

Create each of the instances based on a public image.

During this creation, add a secondary disk for the mirror disk (cluster partition and data partition).

Create as many instances as the number of virtual machines constituting the cluster.

For more information on the procedure, refer to the following:

How-to guides:

<https://cloud.google.com/compute/docs/how-to/>

3. Configuring the instances

Access and log in to the created instances (server1 and server2).

For more information on the procedure, refer to the following:

Connecting to Linux VMs using the Cloud Console and the gcloud tool:

<https://cloud.google.com/compute/docs/instances/connecting-to-instance/>

Next, set the partitions for the mirror disk resource.

Create the cluster partition and data partition on the secondary disk added to the instance.

For more information on setting the partitions for the mirror disk resource, refer to the following:

- "Installation and Configuration Guide" -> "Determining a system configuration" -> "Settings after configuring hardware" -> "Partition settings for Mirror disk resource (when using Replicator)"

4. Creating the firewall rules

Create the firewall rules for allowing communication from the Google Cloud health check system (130.211.0.0/22, 35.191.0.0/16), where the health check is to be performed by the load balancer to the instance.

In addition, add the target tag to the network tag of the instances (server1 and server2).

For more information on the procedure, refer to the following:

Using firewall rules:

<https://cloud.google.com/vpc/docs/using-firewalls/>

Creating health checks:

<https://cloud.google.com/load-balancing/docs/health-checks/>

5. Creating the instance groups

Create the instance groups to be specified as the back ends of Cloud Load Balancing. Add the instances (server1 and server2).

For more information on the procedure, refer to the following:

Creating groups of unmanaged instances:

<https://cloud.google.com/compute/docs/instance-groups/creating-groups-of-unmanaged-instances>

6. Creating the load balancer

Create the load balancer. Select **TCP Load Balancing**.

For more information on the procedure, refer to the following:

Setting up a network load balancer with a target pool:

<https://cloud.google.com/load-balancing/docs/network/setting-up-network/>

Next, configure the back end and the front end.

For **Ports** of the front end, specify the number of the port through which the application is available.

For more information on the procedure, refer to the following:

Setting up Internal TCP/UDP Load Balancing:

<https://cloud.google.com/load-balancing/docs/internal/setting-up-internal>

7. Adjusting the OS startup time, verifying the network settings, verifying the root file system, verifying the firewall settings, synchronizing the server clock, and verifying the SELinux settings

For information on each of the procedures, refer to the following:

- "Installation and Configuration Guide" -> "Determining a system configuration" -> "Settings after configuring hardware"

8. Installing EXPRESSCLUSTER

For information on the procedure, refer to the following document. After completing the installation, reboot the OS.

- "Installation and Configuration Guide"

9. Registering the EXPRESSCLUSTER license

For information on the procedure, refer to the following document:

- "Installation and Configuration Guide"

4.3 Configuring EXPRESSCLUSTER

For information on how to set up and access Cluster WebUI, refer to the following:

- "Installation and Configuration Guide" -> "Creating the cluster configuration data"

This section describes how to add the following resources and monitor resources:

- Mirror disk resource
- Google Cloud virtual IP resource
- Google Cloud virtual IP monitor resource
- Google Cloud load balance monitor resource

For information on other settings, refer to the following:

- "Installation and Configuration Guide"
- "Reference Guide"

1. Creating a cluster

To create a cluster, start the cluster generation wizard first.

- Creating a cluster
 1. Access Cluster WebUI, and click **Cluster generation wizard**.
 2. Of **Cluster generation wizard**, **Cluster** appears.
In **Cluster Name**, enter a cluster name.
From **Language**, select an appropriate language. Click **Next**.
 3. **Basic Settings** appears.
The instance connected to Cluster WebUI appears as a registered master server.
Click **Add** to add the remaining instances (by specifying their internal IP addresses). Click **Next**.
 4. The **Interconnect** screen appears.
Specify the instance's internal IP address to be used for the interconnect. For **MDC**, select **mdc1** as the communication path of a mirror disk resource to be created later. Click **Next**.
 5. The **Fencing** screen appears.
With nothing specified, click **Next**.

2. Adding group resources

- Defining a group

Create a failover group.

1. The **Group List** screen appears.
Click **Add**.
2. The **Group Definition** screen appears.
In **Name**, enter failover1 as a failover group name. Click **Next**.
3. The **Startup Servers** screen appears.
With nothing specified, click **Next**.
4. The **Group Attributes** screen appears.
With nothing specified, click **Next**.

5. The **Group Resource List** screen appears.
Here, add group resources as below.

- Mirror disk resource

Create a mirror disk resource.

For more information, refer to the following:

- "Reference Guide" -> "Understanding Mirror disk resources"

1. In **Group Resource List**, click **Add**.
2. The **Resource Definition of Group | failover1** screen appears.
From the **Type** box, select **Mirror disk resource** as a group resource type. In the **Name** box, enter the resource name. Click **Next**.
3. The **Dependency** screen appears.
With nothing specified, click **Next**.
4. The **Recovery Operation** screen appears.
Click **Next**.
5. The **Details** screen appears.
In each of **Data Partition Device Name** and **Cluster Partition Device Name**, enter the device name of the partition created through "3. **Configuring the instances**". In **Mount Point** and **File System**, enter their respective values. Click **Finish** to complete the setting.

- Google Cloud virtual IP resource

With EXPRESSCLUSTER used on Google Cloud, this resource provides a mechanism for awaiting access from the load balancer to a specific port for the alive monitoring--over the port, an application works on a node. For more information on the Google Cloud virtual IP resource, refer to the following:

- "Reference Guide" -> "Understanding Google Cloud Virtual IP resources"

1. In **Group Resource List**, click **Add**.
2. The **Resource Definition of Group | failover1** screen appears.
From the **Type** box, select **Google Cloud Virtual IP resource** as a group resource type. In the **Name** box, enter the resource name. Click **Next**.
3. The **Dependency** screen appears. With nothing specified, click **Next**.
4. The **Recovery Operation** screen appears. Click **Next**.
5. In **Port Number**, enter the same value as specified in **Health check (port)** during the configuration of the load balancer (back end).
6. Click **Finish**.

3. Adding monitor resources

- Google Cloud virtual IP monitor resource

This monitor resource provides a mechanism for monitoring the alive-monitoring port, for the node where the Google Cloud virtual IP resource is started.

Adding one Google Cloud virtual IP resource automatically creates one Google Cloud virtual IP monitor resource.

For more information on the Google Cloud virtual IP monitor resource, refer to the following:

- "Reference Guide" -> "Understanding Google Cloud Virtual IP resources"

- Google Cloud load balance monitor resource

This monitor resource provides a mechanism for checking whether the same port as that for the health check is opened, for a node where the Google Cloud virtual IP resource is not started.

Adding one Google Cloud virtual IP resource automatically creates one Google Cloud load balance monitor resource.

For more information on the Google Cloud load balance monitor resource, refer to the following:

- "Reference Guide" -> "Understanding Google Cloud load balance monitor resources"

4. Applying the settings and starting the cluster

Refer to the following:

- "Installation and Configuration Guide" -> "How to create a cluster"

4.4 Verifying the created environment

Verify whether the created environment works properly, by producing a monitoring error for a failover of the failover group.

With the cluster running normally, the verification procedure is as follows:

1. On the active node (server1), start the failover group (failover1).
In the **Status** tab of Cluster WebUI, make sure that the status of failover1 is **Online** on server1.
From a client, access the front-end IP address to make sure of being able to connect to the active node.
2. In the pull-down menu of Cluster WebUI, change the mode option from **Operation mode** to **Verification mode**.
3. In the **Status** tab of Cluster WebUI, select the **Enable dummy failure** icon of gcvip1.
4. The Google Cloud virtual IP resource (gcvip1) is reactivated three times. Then the failover group (failover1) fails, being failed over to the corresponding node (server2).
In the **Status** tab of Cluster WebUI, make sure that the status of failover1 is **Online** on server2.
Also make sure that, after the failover, the front-end IP address of the load balancer can be normally accessed.

That is all for testing the failover through a dummy failure. If necessary, perform operation checks for other failures as well.

CLUSTER CREATION PROCEDURE (FOR AN HA CLUSTER WITH CLOUD DNS)

5.1 Creation example

This guide describes how to create a two-node, uni-directional standby cluster with EXPRESSCLUSTER.

Through this procedure, you are to create an HA cluster accessible from clients within the same VPC network on Google Cloud.

This procedure is intended to create a mirror-disk type configuration in which server1 is used as the active server.

The following tables list the parameters which do not have their default values, and the parameters whose values are changed from their default values.

- Google Cloud settings (common to server1 and server2)

Item	Value
Configuration of the VPC network	
- Name	test-vpc
- New subnet (name)	subnetwork-1, subnetwork-11
- New subnet (region)	asia-northeast1
- New subnet (IP address range)	10.0.1.0/24, 10.0.11.0/24
Configuration of Cloud DNS (DNS zone)	
- Zone type	Undisclosed
- Zone name	test-zone
- DNS name	example.com

Continued on next page

Table 5.1 – continued from previous page

Item	Value
- Option	Default (limitedly disclosed)
- Network	test-vpc
Configuration of Cloud DNS (record set)	
- DNS name	test-cluster.example.com
- Resource record type	A
- TTL	5
- TTL unit	Seconds
- IPv4 address	10.0.11.101

- Google Cloud settings (to be set separately on server1 and server2)

Item	Value	
	server1	server2
Configuration of the instances		
• Region	asia-northeast1	asia-northeast1
• Zone	asia-northeast1-a	asia-northeast1-b
• New disk	server1-datadisk-0	server2-datadisk-0
Network configuration		
• Network	test-vpc	test-vpc
• Subnetwork	subnetwork-11	subnetwork-11
• Internal IP address	10.0.11.101	10.0.11.102

- EXPRESSCLUSTER settings (cluster properties)

Item	Value	
	server1	server2
- Cluster name	Cluster1	Cluster1
- Server name	server1	server2

- EXPRESSCLUSTER settings (failover group)

Resource name	Item	Value
Mirror disk resource	Resource name	md1
Mirror disk resource	Details tab - device name of the mirror partition	/dev/NMP1
Mirror disk resource	Details tab - mount point	/mnt/md1
Mirror disk resource	Details tab - device name of the data partition	/dev/sdc2
Mirror disk resource	Details tab - device name of the cluster partition	/dev/sdc1
Mirror disk resource	Details tab - file system	ext4
EXEC resource	Resource name	exec1

- EXPRESSCLUSTER settings (monitor resource)

Monitor resource name	Item	Value
Mirror disk monitor resource	Monitor resource name	mdw1
Mirror disk connect monitor resource	Monitor resource name	mdnw1

5.2 Configuring Google Cloud

1. Creating the VPC network

Access Google Cloud Console (<https://console.cloud.google.com/>).

Create the VPC network and subnets.

For more information on the procedure, refer to the following:

Using VPC networks:

<https://cloud.google.com/vpc/docs/using-vpc/>

2. Creating the instances

Create each of the instances based on a public image.

During this creation, add a secondary disk for the mirror disk (cluster partition and data partition).

Create as many instances as the number of virtual machines constituting the cluster.

For more information on the procedure, refer to the following:

How-to guides:

<https://cloud.google.com/compute/docs/how-to/>

3. Configuring the instances

Access and log in to the created instances (server1 and server2).

For more information on the procedure, refer to the following:

Connecting to Linux VMs using the Cloud Console and the gcloud tool:

<https://cloud.google.com/compute/docs/instances/connecting-to-instance/>

Next, set the partitions for the mirror disk resource.

Create the cluster partition and data partition on the secondary disk added to the instance.

For more information on setting the partitions for the mirror disk resource, refer to the following:

- "Installation and Configuration Guide" -> "Determining a system configuration" -> "Settings after configuring hardware" -> "Partition settings for Mirror disk resource (when using Replicator)"

4. Creating the DNS zone

Create and configure the DNS zone.

For more information on the procedure, refer to the following:

Managing zones:

<https://cloud.google.com/dns/zones/>

From **Zone type**, select **Undisclosed**. For more information on **Zone type**, refer to the following:

Cloud DNS overview:

<https://cloud.google.com/dns/docs/overview/>

5. Setting up the gcloud CLI

Log in to server1 and server2, install the gcloud CLI, and then initialize the SDK.

If already installed, the gcloud CLI does not have to be reinstalled.

Before using the gcloud CLI, you as the root user must authorize Cloud SDK tools.

For more information on this procedure, refer to the following:

Authorizing Cloud SDK tools:

<https://cloud.google.com/sdk/docs/authorizing/>

For more information on other procedures, refer to the following:

Installing Google Cloud SDK:

<https://cloud.google.com/sdk/install/>

Initializing Cloud SDK:

<https://cloud.google.com/sdk/docs/initializing/>

Quickstarts:

<https://cloud.google.com/sdk/docs/quickstarts/>

6. Configuring Cloud IAM

Cloud IAM allows you to authorize the user account or service account used for authorizing Cloud SDK tools, to manage the record sets of Cloud DNS.

Use Cloud IAM to give authority for the following:

- dns.changes.create
- dns.managedZones.get
- dns.resourceRecordSets.create
- dns.resourceRecordSets.update
- dns.resourceRecordSets.delete
- dns.resourceRecordSets.list

Or give the /roles/dns.admin role.

For more information on Cloud IAM, refer to the following:

Cloud Identity and Access Management (IAM):

<https://cloud.google.com/iam/>

Granting, changing, and revoking access to resources:

<https://cloud.google.com/iam/docs/granting-changing-revoking-access/>

Access control:

<https://cloud.google.com/dns/docs/access-control/>

7. Adjusting the OS startup time, verifying the network settings, verifying the root file system, verifying the firewall settings, synchronizing the server clock, and verifying the SELinux settings

For information on each of the procedures, refer to the following:

- "Installation and Configuration Guide" -> "Determining a system configuration" -> "Settings after configuring hardware"

8. Installing EXPRESSCLUSTER

For information on the procedure, refer to the following document. After completing the installation, reboot the OS.

- "Installation and Configuration Guide"

9. Registering the EXPRESSCLUSTER license

For information on the procedure, refer to the following:

- "Installation and Configuration Guide"

5.3 Configuring EXPRESSCLUSTER

For information on how to set up and access Cluster WebUI, refer to the following:

- "Installation and Configuration Guide" -> "Creating the cluster configuration data"

This section describes how to add the following resources and monitor resources:

- Mirror disk resource
- EXEC resource

For information on other settings, refer to the following:

- "Installation and Configuration Guide"
- "Reference Guide"

1. Creating a cluster

To create a cluster, start the cluster generation wizard first.

- Creating a cluster
 1. Access Cluster WebUI, and click **Cluster generation wizard**.
 2. Of **Cluster generation wizard**, **Cluster** appears.
In **Cluster Name**, enter a cluster name.
From **Language**, select an appropriate language. Click **Next**.
 3. **Basic Settings** appears.
The instance connected to Cluster WebUI appears as a registered master server.
Click **Add** to add the remaining instances (by specifying their internal IP addresses). Click **Next**.
 4. The **Interconnect** screen appears.
Specify the instance's internal IP address to be used for the interconnect. For **MDC**, select **mdc1** as the communication path of a mirror disk resource to be created later. Click **Next**.
 5. The **Fencing** screen appears.
With nothing specified, click **Next**.

2. Adding group resources

- Defining a group

Create a failover group.

1. The **Group List** screen appears.
Click **Add**.
2. The **Group Definition** screen appears.
In **Name**, enter failover1 as a failover group name. Click **Next**.
3. The **Startup Servers** screen appears.
With nothing specified, click **Next**.
4. The **Group Attributes** screen appears.
With nothing specified, click **Next**.
5. The **Group Resource List** screen appears.
Here, add group resources as below.

- Mirror disk resource

Create a mirror disk resource.

For more information, refer to the following:

- "Reference Guide" -> "Understanding Mirror disk resources"

1. In **Group Resource List**, click **Add**.
2. The **Resource Definition of Group | failover1** screen appears.
From the **Type** box, select **Mirror disk resource** as a group resource type. In the **Name** box, enter the resource name. Click **Next**.
3. The **Dependency** screen appears.
With nothing specified, click **Next**.
4. The **Recovery Operation** screen appears.
Click **Next**.
5. The **Details** screen appears.
In each of **Data Partition Device Name** and **Cluster Partition Device Name**, enter the device name of the partition created through "3. **Configuring the instances**". In **Mount Point** and **File System**, enter their respective values. Click **Finish** to complete the setting.

- EXEC resource

This resource provides a mechanism for (when activated) registering records in and (when deactivated) removing them from Cloud DNS.

For more information, refer to the following:

- "Reference Guide" -> "Understanding EXEC resources"

1. In **Group Resource List**, click **Add**.
2. The **Resource Definition of Group | failover1** screen appears.
From the **Type** box, select **EXEC resource** as a group resource type. In the **Name** box, enter the resource name. Click **Next**.
3. The **Dependency** screen appears.
With nothing specified, click **Next**.
4. The **Recovery Operation** screen appears.
Click **Next**.
5. Edit the start script (start.sh of the scripts for Cloud DNS) and the stop script (stop.sh of them).
Set the parameter values of start.sh and stop.sh, according to your environment and using the following list:

Parameter name	Description	Value (example for this guide)
SERVER1_NAME	Server name of server1	"server1"
SERVER1_IP	IP address of server1	"10.0.11.101"
SERVER2_NAME	Server name of server2	"server2"
SERVER2_IP	IP Address of server2	"10.0.11.102"
ZONE_NAME	Zone name	"test-zone"
DOMAIN_NAME	DNS name	"test-cluster.example.com"
DOMAIN_TYPE	Record type	"A"
DOMAIN_TTL	TTL value	"5"

6. Click **Finish**.

3. Adding monitor resources

The scripts for Cloud DNS do not provide a means for checking whether records registered by the EXEC resource exist or whether name resolution is possible.

If necessary, create a custom monitor resource.

For more information, refer to the following:

- "Reference Guide" -> "Understanding custom monitor resources"

4. Applying the settings and starting the cluster

Refer to the following:

- "Installation and Configuration Guide" -> "How to create a cluster"

5.4 Verifying the created environment

Verify whether the created environment works properly, through a failover of the failover group.

With the cluster running normally, the verification procedure is as follows:

1. On the active node (server1), start the failover group (failover1). In the **Status** tab of Cluster WebUI, make sure that the status of failover1 is **Online** on server1.

2. From a client, access test-cluster.example.com to make sure of being able to connect to the active node.

```
$ nslookup test-cluster.example.com <DNS server>
```

3. In Cluster WebUI, manually move the failover group from the active node to the standby node. In the **Status** tab of Cluster WebUI, make sure that the status of failover1 is **Online** on server2.

4. From a client, access test-cluster.example.com to make sure of being able to connect to the standby node.

```
$ nslookup test-cluster.example.com <DNS server>
```

That is all for testing the connection from the client to the HA cluster by using the DNS name of registered records. If necessary, perform operation checks for other failures as well.

ERROR MESSAGES

For information on error messages of the resources/monitor resources, refer to the following:

- "Reference Guide" -> "Error messages"

NOTES AND RESTRICTIONS

7.1 HA cluster with a load balancer

7.1.1 Notes on Google Cloud

- In designing a performance-oriented system, keep this in mind: Google Cloud tends to increase its performance deterioration rate in multi-tenant cloud environments, compared with that in physical environments or general and virtualized (non-cloud) environments.

7.1.2 Notes on EXPRESSCLUSTER

- Google Cloud's specification requires HTTP-protocol-based legacy health checks for external TCP network load balancing.
The Google Cloud virtual IP resource, which supports only TCP-protocol-based health checks, cannot respond to health checks by an external TCP network load balancer.
Therefore, use an internal TCP load balancer instead of an external TCP network load balancer, with which and the Google Cloud virtual IP resource the HA cluster cannot be used.
Refer to the following:
Health checks overview:
<https://cloud.google.com/load-balancing/docs/health-check-concepts/>
- For an HA cluster configuration with an internal TCP load balancer, the HA cluster cannot be accessed from any client which belongs to a region different from that of the HA cluster--due to Google Cloud's specification.
Refer to the following:
Internal TCP/UDP Load Balancing overview:
<https://cloud.google.com/load-balancing/docs/internal/#architecture>
- Make the OS startup time longer than the time of **Heartbeat Timeout**.
- Going to **Cluster Properties** -> the **Monitor** tab -> **Shutdown Monitor Timeout**, you can change the default value (Use **Heartbeat Timeout**) if necessary. Then make the value equal to or less than that of **Heartbeat Timeout**.
Refer to the following:
 - "Reference Guide" -> "Cluster properties" -> "Timeout tab"
 - "Reference Guide" -> "Cluster properties" -> "Monitor tab"
 - "Getting Started Guide" -> "Notes and Restriction" -> "Adjusting OS startup time"

See also:

- "Getting Started Guide" -> "Notes and Restriction" -> "Communication port number"

- "Getting Started Guide" -> "Notes and Restriction" -> "Setting up Google Cloud virtual IP resources"
- "Getting Started Guide" -> "Notes and Restriction" -> "Setting up Google Cloud load balance monitor resources"
- "Reference Guide" -> "Notes on Google Cloud Virtual IP resources"
- "Reference Guide" -> "Notes on Google Cloud Virtual IP monitor resources"
- "Reference Guide" -> "Notes on Google Cloud load balance monitor resources"

7.2 HA cluster with Cloud DNS

7.2.1 Notes on Google Cloud

- In designing a performance-oriented system, keep this in mind: Google Cloud tends to increase its performance deterioration rate in multi-tenant cloud environments, compared with that in physical environments or general and virtualized (non-cloud) environments.
- If a client takes time to resolve a name, its possible causes are as below. Check their corresponding settings:
 - The OS-side resolver takes time.
 - The TTL value is high.
 - If a name resolution is tried prior to the completion of Cloud DNS change propagation, the DNS returns NXDOMAIN (non-existing domain). In this case, the name resolution fails until the valid period of the negative cache expires.

7.2.2 Notes on EXPRESSCLUSTER

- Have only one EXEC resource with the scripts for Cloud DNS.
If multiple actions are simultaneously taken to the record sets of Cloud DNS, some of the actions may fail. This is due to Google Cloud's specification.
- Make the OS startup time longer than the time of **Heartbeat Timeout**.
- Going to **Cluster Properties** -> the **Monitor** tab -> **Shutdown Monitor Timeout**, you can change the default value (Use **Heartbeat Timeout**) if necessary. Then make the value equal to or less than that of **Heartbeat Timeout**.
Refer to the following:
 - "Reference Guide" -> "Cluster properties" -> "Timeout tab"
 - "Reference Guide" -> "Cluster properties" -> "Monitor tab"
 - "Getting Started Guide" -> "Notes and Restriction" -> "Adjusting OS startup time"

LEGAL NOTICE

8.1 Disclaimer

- Information in this document is subject to change without notice.
- NEC Corporation is not liable for technical or editorial mistakes in or omissions from this document.
In addition, whether the customer achieves the desired effectiveness by following the introduction and usage instructions in this document is the responsibility of the customer.
- No part of this document may be reproduced or transmitted in any form by any means, electronic or mechanical, for any purpose, without the express written permission of NEC Corporation.

8.2 Trademark Information

- EXPRESSCLUSTER X is a registered trademark of NEC Corporation.
- Linux is a registered trademark of Linus Torvalds in the United States and other countries.
- Google Cloud is a trademark or a registered trademark of Google LLC.
- Other product names and slogans written in this manual are trademarks or registered trademarks of their respective companies.

REVISION HISTORY

Edition	Revised Date	Description
1st	Apr 15, 2024	New Guide
2nd	Nov 29, 2024	Corrected typographical errors and other mistakes.

© Copyright NEC Corporation 2024. All rights reserved.