



EXPRESSCLUSTER X 5.2
HA Cluster Configuration Guide for Oracle Cloud
Infrastructure (Linux)
Release 1

NEC Corporation

Apr 15, 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	12
2.4	Differences between on-premises and OCI	16
3	Mirror Disk Type Cluster Creation Procedure	17
3.1	Creation examples	17
3.2	Configuring OCI	21
3.3	Configuring the EXPRESSCLUSTER settings	23
3.4	Verifying the operations	26
4	Shared Disk Type Cluster Creation Procedure	27
4.1	Creation examples	27
4.2	Configuring OCI	31
4.3	Configuring the EXPRESSCLUSTER settings	33
4.4	Verifying the operations	36
5	Error Messages	37
6	Notes and Restrictions	39
6.1	HA cluster using a load balancer	39
6.2	Notes on creating an instance	40
7	Legal Notice	41
7.1	Disclaimer	41
7.2	Trademark Information	42
8	Revision History	43

1.1 Who Should Use This Guide

The *HA Cluster Configuration Guide for Oracle Cloud Infrastructure (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.

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. *Mirror Disk Type Cluster Creation Procedure*": Describes the procedure to create a mirror disk type cluster.
- "4. *Shared Disk Type Cluster Creation Procedure*": Describes the procedure to create a shared disk type cluster.
- "5. *Error Messages*": Describes the error messages and solutions.
- "6. *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 (hereinafter referred to as EXPRESSCLUSTER) on a cloud service of Oracle Cloud Infrastructure (hereinafter referred to as OCI).

OCI allows virtual machines to be HA-clustered by using regions or availability domains, thus enhancing the operational availability.

- Regions

OCI allows each node to be deployed in physical or logical units called a region (e.g., Tokyo).

It is possible to build all nodes in a single region. However, in such a case, a network failure or a natural disaster can cause all nodes to go down, discontinuing the operation.

To increase availability, distribute nodes to multiple regions.

- Availability domains

OCI allows each node to be deployed in a logical group called an availability domain.

Locating each node in a different availability domain minimizes the impact of planned OCI maintenance or unplanned maintenance due to a physical hardware failure.

For details on regions and availability domains, see the following website:

Regions and Availability Domains:

<https://docs.cloud.oracle.com/en-us/iaas/Content/General/Concepts/regions.htm>

2.2 Basic configuration

This guide assumes an HA cluster using a load balancer (a unidirectional standby cluster).

The following table describes EXPRESSCLUSTER resources to be selected and required OCI services for the HA cluster:

Purpose	EXPRESSCLUSTER resource to be selected	Required OCI services
Accessing from a client by using a virtual IP address (private IP address)	Oracle Cloud virtual IP resource	Private load balancer
Accessing from a client by using a virtual IP address (global IP address)	Oracle Cloud virtual IP resource	Public load balancer

HA clusters using a load balancer

A client application can connect a virtual machine in the OCI environment to a node that constitutes a cluster, by using a virtual IP (VIP) address.

Using the VIP address eliminates the need for clients to be aware of switching between the virtual machines even after a failover or a group migration occurs.

For a private load balancer, an HA cluster built in the OCI environment in [Fig. 2.1 HA cluster using a private load balancer](#) is accessed by specifying the VIP address. This VIP address is a private IP address assigned to the OCI load balancer.




Fig. 2.1: HA cluster using a private load balancer

For a public load balancer, an HA cluster built in the OCI environment in [Fig. 2.2 HA cluster using a public load balancer](#) is accessed by specifying the VIP address. This VIP address is a global IP address assigned to the OCI load balancer.

The active node and the standby node of a cluster are switched through a health check by the OCI load balancer. The health check can be performed through a port provided by the Oracle Cloud virtual IP resource.

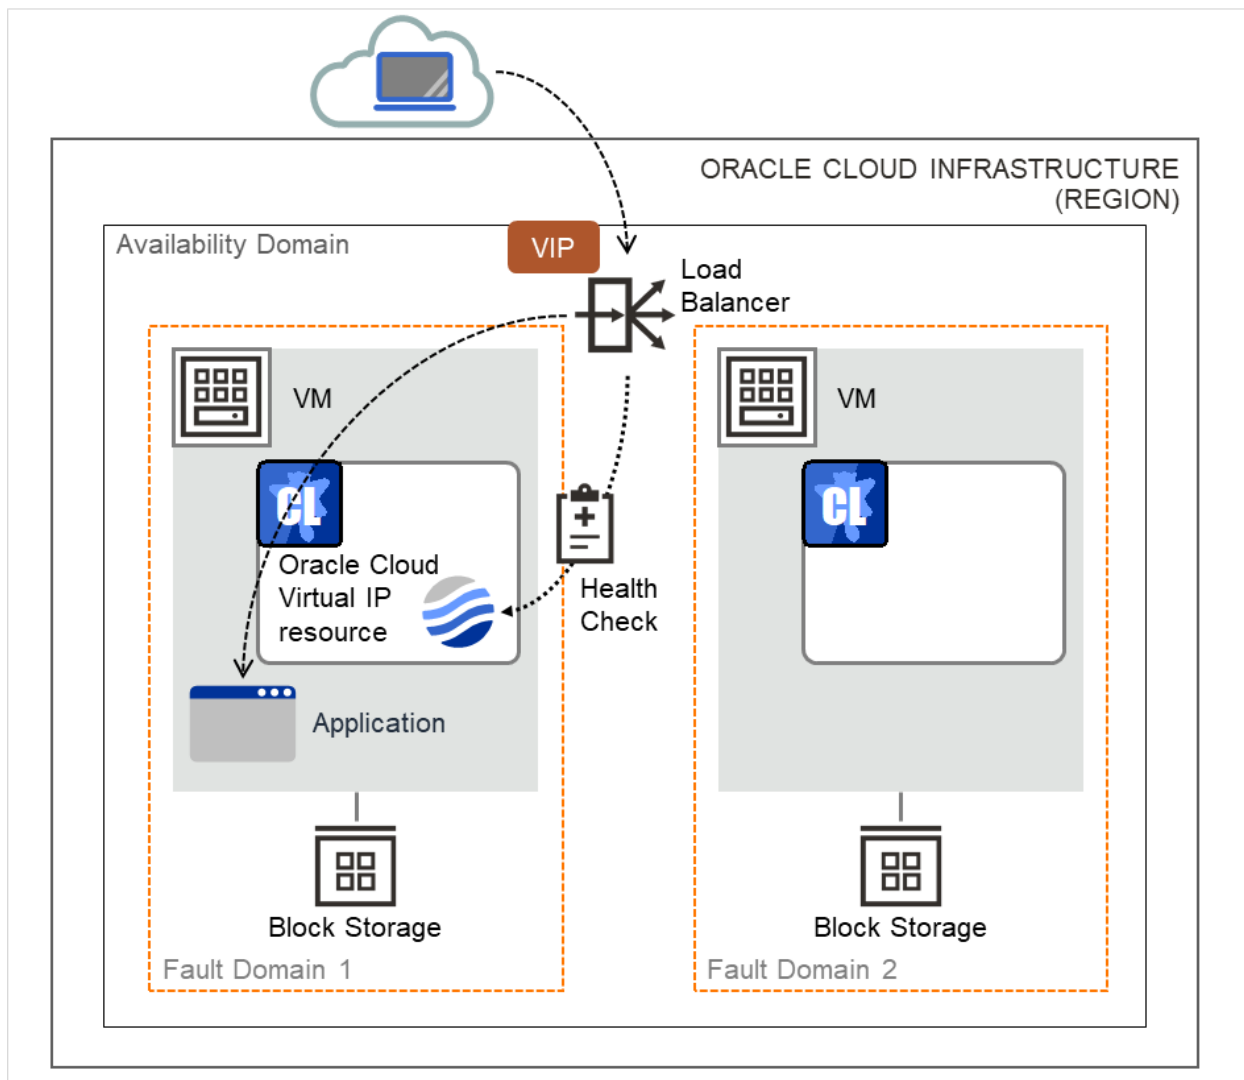


Fig. 2.2: HA cluster using a public load balancer

For details on the load balancer, refer to the following:

Overview of Load Balancing:

<https://docs.cloud.oracle.com/en-us/iaas/Content/Balance/Concepts/balanceoverview.htm>

The following table describes the resources and monitor resources required for an HA cluster configuration using the load balancer:

Resource or monitor resource type	Description	Setting
Oracle Cloud virtual IP resource	<p>Provides a mechanism to wait for the alive monitoring by a load balancer on a specific port of a node where operations are running (wait for an access to the health-check port).</p> <p>Activating the Oracle Cloud virtual IP resource starts the control process on standby for the alive monitoring by the OCI load balancer.</p> <p>Deactivating the Oracle Cloud virtual IP resource stops the control process on standby for the alive monitoring.</p>	Required
Oracle Cloud virtual IP monitor resource	Performs the alive monitoring of the control process, which starts upon the activation of the Oracle Cloud virtual IP resource, for a node where the Oracle Cloud virtual IP resource is running.	Required
Oracle Cloud load balance monitor resource	Monitors whether the same number as the health-check port number has already been used, for a node where an Oracle Cloud virtual IP resource has not been activated.	Required
Other resources and monitor resources	Depends on the configuration of applications, such as mirror disks or shared disks, which are used in an HA cluster.	Optional

2.3 Network partition resolution

Virtual machines constituting an HA cluster mutually perform alive monitoring through a heartbeat communication.

If the virtual machines reside in different subnets, an interruption of the heartbeat causes an undesirable event such as a service starting more than once.

To prevent double startup of the service, it is necessary to identify whether other virtual machines went down or whether the applicable virtual machine was isolated from a network (network partitioning: NP).

The network partition resolution feature (NP resolution) uses a ping command to be sent to a device (access destination) that is always activated and can respond to the ping command etc. If there is no reply to the ping command, it is determined that the device entered the NP status and the specified action (such as a warning, recovery action, or server shutdown) is executed.

For details on the configurations of the heartbeat and NP resolution, refer to the following.

The following heartbeat and NP resolution types correspond to the figures below:

Heartbeat or NP resolution type	Type
(1)	Kernel mode LAN heartbeat resource
(2)	Kernel mode LAN heartbeat resource
(3)	Witness heartbeat resource HTTP network partition resolution resource
(4)	Disk heartbeat resource

Heartbeat and NP resolution configuration (mirror disk type cluster)




Fig. 2.3: Heartbeat and NP resolution configuration (mirror disk type cluster)

Heartbeat and NP resolution configuration (shared disk type cluster)

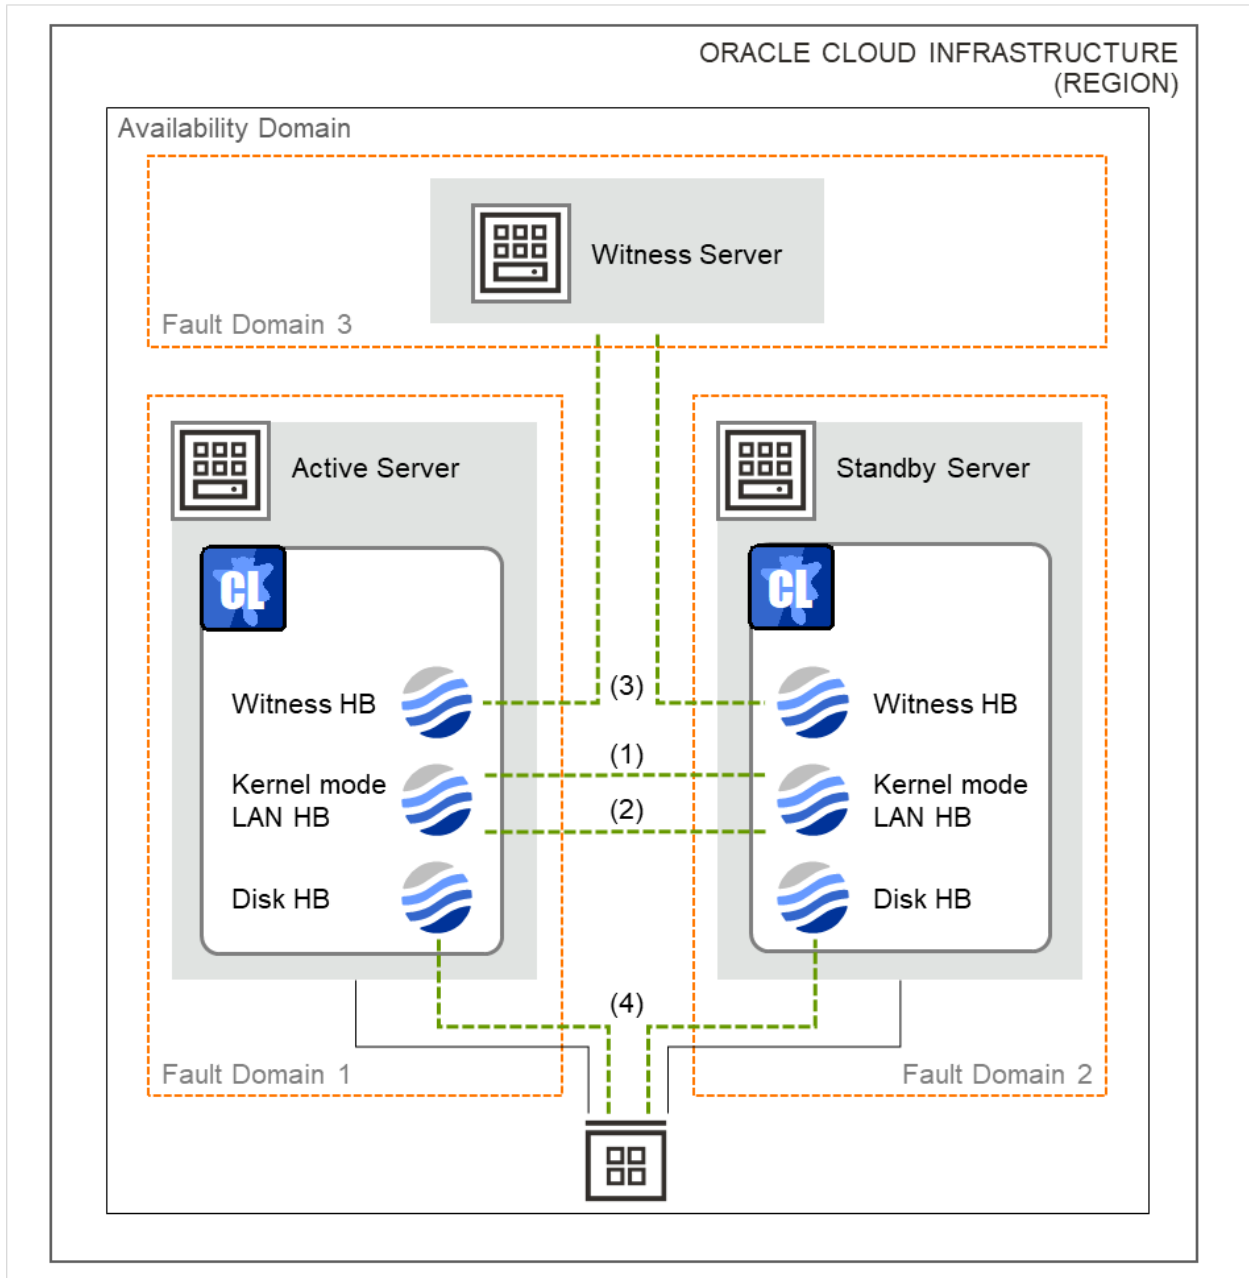


Fig. 2.4: Heartbeat and NP resolution configuration (shared disk type cluster)

Using the Witness heartbeat resource allows the aliveness of the other server to be confirmed based on the information on an access to the Witness server.

With the Witness heartbeat resource in combination with the HTTP network partition resolution resource, when a failure occurs in all network channels (heartbeats) and network partitioning occurs, emergency shutdown takes place to protect data.

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).

For details on the heartbeat resource and 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 OCI

The following table describes the functional differences of EXPRESSCLUSTER between on-premises and OCI. "Y" indicates that the relevant function can be used and "N" indicates that the relevant function cannot be used.

Function	On-premise	OCI
Creating a shared disk type cluster	Y	Y
Creating a mirror disk type cluster	Y	Y
Using the management group	Y	N
Using a floating IP resource	Y	N
Using a virtual IP resource	Y	N
Using an Oracle Cloud virtual IP resource	N	Y
Using an Oracle Cloud DNS resource	N	Y

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

MIRROR DISK TYPE CLUSTER CREATION PROCEDURE

3.1 Creation examples

This guide introduces the procedure for creating a two-node unidirectional standby cluster in OCI using EXPRESSCLUSTER.

An HA cluster to be created is accessible from clients in the same virtual cloud network (hereafter described as VCN) in OCI.

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

The following tables describe the parameters that do not have default values and the parameters whose values have been changed from the default values.

- OCI settings (common to each instance)

Setting item	Setting value
VCN settings	
– Name	test-vcn
Load balancer settings	
– Load balancer name	test-loadbalancer
– Visibility type	Private / public (*)
– Virtual cloud network	test-vcn
Load balancer settings (backend set settings)	
– Added backend: name	server1, server2
– Added backend: port	8080 (the number of the port through which the application is available: cluster side)

Continued on next page

Table 3.1 – continued from previous page

Setting item	Setting value
– Health check policy: protocol	TCP
– Health check policy: port	12345
– Health check policy: interval (in milliseconds)	5000
Load balancer settings (listener settings)	
– Traffic type	TCP
– Port through which a listener monitors	80 (the number of the port through which the application is available: client side)

(*) Select Private when using a private load balancer, and Public when using a public load balancer.

- OCI settings (individually configured for each instance)

Setting item	Setting value		
	server1	server2	witness-server
Compute instance settings			
– Instance name	server1	server2	witness-server
– Availability domain	LhRE:AP-TOKYO-1-AD-1	LhRE:AP-TOKYO-1-AD-1	LhRE:AP-TOKYO-1-AD-1
– Instance type	Virtual machine	Virtual machine	Virtual machine
– Virtual cloud network	test-vcn	test-vcn	test-vcn
– Fault domain	FAULT-DOMAIN-1	FAULT-DOMAIN-2	FAULT-DOMAIN-3
Block volume settings			
– Name	server1-datadisk-0	server2-datadisk-0	-
– Availability domain	LhRE:AP-TOKYO-1-AD-1	LhRE:AP-TOKYO-1-AD-1	-

Continued on next page

Table 3.2 – continued from previous page

Setting item	Setting value		
	server1	server2	witness-server
- Device path	/dev/oracleoci/oracleovdb	/dev/oracleoci/oracleovdb	-
Network settings			
- Private IP address	10.0.0.2	10.0.0.3	10.0.0.4
- Private IP address	10.0.1.2	10.0.1.3	-
- (*) Public IP address	140.238.54.236	158.101.136.208	164.92.39.211

(*) Set when using a public load balancer.

- EXPRESSCLUSTER settings (cluster properties)

Setting item	Setting value	
	server1	server2
- Cluster name	cluster1	cluster1
- Server name	server1	server2
Interconnect		
- Kernel mode	10.0.0.2	10.0.0.3
- Kernel mode	10.0.1.2	10.0.1.3
- Witness	Used	Used
NP resolution		
- HTTP	Used	Used

- EXPRESSCLUSTER settings (failover group)

Resource name	Setting item	Setting value
Mirror disk resource	Resource name	md1
	Details tab - Mirror partition device name	/dev/NMP1
	Details tab - Mount point	/mnt/md1

Continued on next page

Table 3.4 – continued from previous page

Resource name	Setting item	Setting value
	Details tab - Data partition device name	/dev/oracleoci/oraclevd2
	Details tab - Cluster partition device name	/dev/oracleoci/oraclevd1
Oracle Cloud virtual IP resource	Resource name	ocvip1
	Port number	12345 (as specified in Health check policy: port)

- EXPRESSCLUSTER settings (monitor resource)

Monitor resource name	Setting item	Setting value
Mirror disk monitor resource	Monitor resource name	mdw1
Mirror disk connect monitor resource	Monitor resource name	mdnw1
Oracle Cloud virtual IP monitor resource	Monitor resource name	ocvipw1
	Recovery target	ocvip1
Oracle Cloud load balance monitor resource	Monitor resource name	oclbw1
	Recovery target	ocvip1

3.2 Configuring OCI

1. Creating a VCN

Log in to OCI Console (<https://console.us-ashburn-1.oraclecloud.com/>).

Create a VCN and a subnet.

For details on the procedure, refer to the following:

Networking Overview:

<https://docs.cloud.oracle.com/en-us/iaas/Content/Network/Concepts/overview.htm>

Networking Scenarios:

<https://docs.cloud.oracle.com/en-us/iaas/Content/Network/Concepts/scenarios.htm>

2. Creating an instance

Create as many instances as there are virtual machines constituting a cluster, by using platform images.

For details on the procedure, refer to the following:

Creating an instance:

<https://docs.cloud.oracle.com/en-us/iaas/Content/Compute/Tasks/launchinginstance.htm>

Configure a secondary VNIC.

For details on the procedure, refer to the following:

<https://docs.cloud.oracle.com/en-us/iaas/Content/Network/Tasks/managingVNICs.htm>

3. Setting an instance

Connect to each of the created instances and log in.

For details on the procedure, refer to the following:

Connecting to an Instance:

<https://docs.cloud.oracle.com/en-us/iaas/Content/Compute/Tasks/accessinginstance.htm>

4. Creating a block volume

Add and attach a block volume to be used for mirror disks (cluster partitions and data partitions).

Access: select Read/Write.

For details on the procedure, refer to the following:

Overview of Block Volume:

<https://docs.cloud.oracle.com/en-us/iaas/Content/Block/Concepts/overview.htm>

Create a partition for mirror disk resources.

For details on the settings of partitions for mirror disk resources, 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)"

5. Adding security lists for a load balancer

In creating a load balancer as shown below, the addition of a backend server allows the Load Balancing service to automatically create security list rules.

For details on the procedure, refer to the following:

Security Lists:

<https://docs.cloud.oracle.com/en-us/iaas/Content/Network/Concepts/securitylists.htm>

Managing Load Balancers:

<https://docs.cloud.oracle.com/en-us/iaas/Content/Balance/Tasks/managingloadbalancer.htm>

6. Creating a load balancer

Create a load balancer.

Click **Add Backends** to add server1 and server2.

Set **Interval in ms** for the health check policy to be shorter than **Health Check Timeout** for the Oracle Cloud virtual IP resource.

For details on the procedure, refer to the following:

Managing Load Balancers:

<https://docs.cloud.oracle.com/en-us/iaas/Content/Balance/Tasks/managingloadbalancer.htm>

Configure a route table and a security list as required.

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 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"

3.3 Configuring the EXPRESSCLUSTER settings

1. Installing EXPRESSCLUSTER

For information on the installation procedure, refer to the following document. Restart the OS upon the completion of the installation.

- "Installation and Configuration Guide"

2. Registering the EXPRESSCLUSTER license

For details on the license registration procedure, refer to the following:

- "Installation and Configuration Guide"

3. Configuring EXPRESSCLUSTER

For Cluster WebUI setup and connection procedures, refer to the following:

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

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

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

For the settings of other resources and monitor resources, refer to the following:

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

4. Creating a cluster

Start the Cluster generation wizard to create a cluster.

- Creating a cluster
 1. Access Cluster WebUI, and click **Cluster generation wizard**.
 2. **Cluster of Cluster generation wizard** is displayed.
Enter a desired name in **Cluster Name**.
Select an appropriate language in **Language**. Click **Next**.
 3. **Basic Settings** is displayed.
The instance connected to Cluster WebUI is displayed as a registered master server.
Click **Add** to add the remaining instances (by specifying the private IP address of each instance). Click **Next**.
 4. The **Interconnect** window is displayed.
Specify the IP addresses (the private IP address of each instance) to be used for interconnect, and a Witness heartbeat.
For **MDC**, select mdc1 as the communication path of a mirror disk resource to be created later.
Click **Next**.
For more information, refer to the following:
- "Reference Guide" -> "Understanding Witness heartbeat resources"

5. The **NP Resolution** window is displayed.
Specify HTTPNP and click **Next**.
For more information, refer to the following:
- "Reference Guide" -> "Understanding network partition resolution by HTTP method"

5. Adding group resources

- Defining a group

Create a failover group.

1. The **Group List** window is displayed.
Click **Add**.
2. The **Group Definition** window is displayed.
In **Name**, enter failover1 as a failover group name. Click **Next**.
3. The **Startup Servers** window is displayed.
Click **Next** without specifying anything.
4. The **Group Attributes** window is displayed.
Click **Next** without specifying anything.
5. The **Group Resource List** window is displayed.
On this page, add a group resource following the procedure below.

- Mirror disk resource

Create a mirror disk resource.

For more information, refer to the following:

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

1. Click **Add** on the **Group Resource List** page.
2. The **Resource Definition of Group | failover1** window is displayed.
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** window is displayed.
Click **Next** without specifying anything.
4. The **Recovery Operation** window is displayed.
Click **Next**.
5. The **Details** window is displayed.
In **Data Partition Device Name** and **Cluster Partition Device Name**, enter the device name of the partition created in "4. Creating a block volume". Specify **Mount Point** and **File System**. Click **Finish** to finish the settings.

- Oracle Cloud virtual IP resource

In using EXPRESSCLUSTER in OCI, provides a mechanism to wait for the alive monitoring by a load balancer on a specific port of a node where operations are running.

For details on the Oracle Cloud virtual IP resource, refer to the following:

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

1. Click **Add** on the **Group Resource List** page.
2. The **Resource Definition of Group | failover1** window is displayed.
In the **Type** box, select **Oracle Cloud virtual IP resource** as a group resource type. In the **Name** box, enter the resource name. Click **Next**.
3. The **Dependency** window is displayed. Click **Next** without specifying anything.
4. The **Recovery Operation** window is displayed. Click **Next**.
5. In **Port Number**, enter the value specified in **Health check policy: port** in the load balancer settings (the backend set settings).
6. Click **Finish**.

6. Adding monitor resources

- Oracle Cloud virtual IP monitor resource

Provides a mechanism for monitoring the alive-monitoring port for the node where an Oracle Cloud virtual IP resource has been activated.

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

For details on the Oracle Cloud virtual IP monitor resource, refer to the following:

- "Reference Guide" -> "Understanding Oracle Cloud Virtual IP monitor resources"

- Oracle Cloud load balance monitor resource

For the node where an Oracle Cloud virtual IP resource has not been activated, provides a mechanism for monitoring whether the same port number as the health-check port number has already been used.

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

For details on the Oracle Cloud load balance monitor resource, refer to the following:

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

7. Applying the settings and starting the cluster

Refer to the following:

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

3.4 Verifying the operations

Verify whether the created environment works properly by generating a monitoring error to fail over a failover group. If the cluster is running normally, the verification procedure is as follows:

1. Start the failover group (failover1) on the active node (server1).
In the **Status** tab on Cluster WebUI, confirm that the status of failover1 is **Online** at server1.
Access the IP address of the front-end from the client to confirm that the connection to the active node is available.
2. Change **Operation mode** to **Verification mode** from the Cluster WebUI pull-down menu.
3. In the **Status** tab on Cluster WebUI, click the **Enable dummy failure** icon of ocvip1.
4. Upon three times of reactivation of the Oracle Cloud virtual IP resource (ocvip1), the failover group (failover1) enters an error status and fails over to node server2.
In the **Status** tab on Cluster WebUI, confirm that the status of failover1 is **Online** at server2.
Also make sure that, after the failover, the front-end IP address of the load balancer can be normally accessed.

Verifying the failover operation in case of a dummy failure is now complete. If necessary, perform operation checks for other failures.

SHARED DISK TYPE CLUSTER CREATION PROCEDURE

4.1 Creation examples

This guide introduces the procedure for creating a two-node unidirectional standby cluster in OCI using EXPRESSCLUSTER.

An HA cluster to be created is accessible from clients in the same virtual cloud network (hereafter described as VCN) in OCI.

This procedure is intended to create a shared disk type configuration in which server1 serves as the active server.

The following tables describe the parameters that do not have default values and the parameters whose values have been changed from the default values.

- OCI settings (common to each instance)

Setting item	Setting value
VCN settings	
– Name	test-vcn
Load balancer settings	
– Load balancer name	test-loadbalancer
– Visibility type	Private / public (*)
– Virtual cloud network	test-vcn
Load balancer settings (backend set settings)	
– Added backend: name	server1, server2
– Added backend: port	8080 (the number of the port through which the application is available: cluster side)

Continued on next page

Table 4.1 – continued from previous page

Setting item	Setting value
– Health check policy: protocol	TCP
– Health check policy: port	12345
– Health check policy: interval (in milliseconds)	5000
Load balancer settings (listener settings)	
– Traffic type	TCP
– Port through which a listener monitors	80 (the number of the port through which the application is available: client side)

(*) Select Private when using a private load balancer, and Public when using a public load balancer.

- OCI settings (individually configured for each instance)

Setting item	Setting value		
	server1	server2	witness-server
Compute instance settings			
– Instance name	server1	server2	witness-server
– Availability domain	LhRE:AP-TOKYO-1-AD-1	LhRE:AP-TOKYO-1-AD-1	LhRE:AP-TOKYO-1-AD-1
– Instance type	Virtual machine	Virtual machine	Virtual machine
– Virtual cloud network	test-vcn	test-vcn	test-vcn
– Fault domain	FAULT-DOMAIN-1	FAULT-DOMAIN-2	FAULT-DOMAIN-3
Block volume settings			
– Name	shared-datadisk	shared-datadisk	-
– Availability domain	LhRE:AP-TOKYO-1-AD-1	LhRE:AP-TOKYO-1-AD-1	-

Continued on next page

Table 4.2 – continued from previous page

Setting item	Setting value		
	server1	server2	witness-server
– Device path	/dev/oracleoci/oracleovdb	/dev/oracleoci/oracleovdb	-
Network settings			
– Private IP address	10.0.0.2	10.0.0.3	10.0.0.4
– Private IP address	10.0.1.2	10.0.1.3	-
– (*) Public IP address	140.238.54.236	158.101.136.208	164.92.39.211

(*) Set when using a public load balancer.

- EXPRESSCLUSTER settings (cluster properties)

Setting item	Setting value	
	server1	server2
– Cluster name	cluster1	cluster1
– Server name	server1	server2
Interconnect		
– Kernel mode	10.0.0.2	10.0.0.3
– Kernel mode	10.0.1.2	10.0.1.3
– Witness	Used	Used
– DISK	/dev/oracleoci/oracleovdb1	/dev/oracleoci/oracleovdb1
NP resolution		
– HTTP	Used	Used

- EXPRESSCLUSTER settings (failover group)

Resource name	Setting item	Setting value
Disk resource	Resource name	disk1
	Details tab - Disk type	disk
	Details tab - File system	ex4
	Details tab - Device name	/dev/oracleoci/oraclevd2
	Details tab - Mount point	/mnt/disk1
Oracle Cloud virtual IP resource	Resource name	ocvip1
	Port number	12345 (as specified in Health check policy: port)

- EXPRESSCLUSTER settings (monitor resource)

Monitor resource name	Setting item	Setting value
Oracle Cloud virtual IP monitor resource	Monitor resource name	ocvipw1
	Recovery target	ocvip1
Oracle Cloud load balance monitor resource	Monitor resource name	oclbw1
	Recovery target	ocvip1

4.2 Configuring OCI

1. Creating a VCN

Log in to OCI Console (<https://console.us-ashburn-1.oraclecloud.com/>).

Create a VCN and a subnet.

For details on the procedure, refer to the following:

Networking Overview:

<https://docs.cloud.oracle.com/en-us/iaas/Content/Network/Concepts/overview.htm>

Networking Scenarios:

<https://docs.cloud.oracle.com/en-us/iaas/Content/Network/Concepts/scenarios.htm>

2. Creating an instance

Create as many instances as there are virtual machines constituting a cluster, by using platform images.

For details on the procedure, refer to the following:

Creating an instance:

<https://docs.cloud.oracle.com/en-us/iaas/Content/Compute/Tasks/launchinginstance.htm>

Configure a secondary VNIC.

For details on the procedure, refer to the following:

<https://docs.cloud.oracle.com/en-us/iaas/Content/Network/Tasks/managingVNICs.htm>

3. Setting an instance

Connect to each of the created instances and log in.

For details on the procedure, refer to the following:

Connecting to an Instance:

<https://docs.cloud.oracle.com/en-us/iaas/Content/Compute/Tasks/accessinginstance.htm>

4. Creating a block volume

Add and attach a block volume to be used for shared disks (disk heartbeat and partitions for shared disks).

Access: select Read/Write and sharable.

For details on the procedure, refer to the following:

Overview of Block Volume:

<https://docs.cloud.oracle.com/en-us/iaas/Content/Block/Concepts/overview.htm>

Create a partition for disk resources.

For details on the settings of partitions for disk resources, refer to the following:

- "Installation and Configuration Guide" -> "Determining a system configuration" -> "Settings after configuring hardware" -> "Shared disk settings for disk resource (Required for disk resource)"

5. Adding security lists for a load balancer

In creating a load balancer as shown below, the addition of a backend server allows the Load Balancing service to automatically create security list rules.

For details on the procedure, refer to the following:

Security Lists:

<https://docs.cloud.oracle.com/en-us/iaas/Content/Network/Concepts/securitylists.htm>

Managing Load Balancers:

<https://docs.cloud.oracle.com/en-us/iaas/Content/Balance/Tasks/managingloadbalancer.htm>

6. Creating a load balancer

Create a load balancer.

Click **Add Backends** to add server1 and server2.

Set **Interval in ms** for the health check policy to be shorter than **Health Check Timeout** for the Oracle Cloud virtual IP resource.

For details on the procedure, refer to the following:

Managing Load Balancers:

<https://docs.cloud.oracle.com/en-us/iaas/Content/Balance/Tasks/managingloadbalancer.htm>

Configure a route table and a security list as required.

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 SELinux settings

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

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

4.3 Configuring the EXPRESSCLUSTER settings

1. Installing EXPRESSCLUSTER

For details on the installation procedure, refer to the following. Restart the OS upon the completion of the installation.

- "Installation and Configuration Guide"

2. Registering the EXPRESSCLUSTER license

For details on the license registration procedure, refer to the following:

- "Installation and Configuration Guide"

3. Configuring EXPRESSCLUSTER

For Cluster WebUI setup and connection procedures, refer to the following:

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

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

- Disk resource
- Oracle Cloud virtual IP resource
- Oracle Cloud virtual IP monitor resource
- Oracle Cloud load balance monitor resource

For the settings of other resources and monitor resources, refer to the following:

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

4. Creating a cluster

Start the Cluster generation wizard to create a cluster.

- Creating a cluster
 1. Access Cluster WebUI, and click **Cluster generation wizard**.
 2. **Cluster of Cluster generation wizard** is displayed.
Enter a desired name in **Cluster Name**.
Select an appropriate language in **Language**. Click **Next**.
 3. **Basic Settings** is displayed.
The instance connected to Cluster WebUI is displayed as a registered master server.
Click **Add** to add the remaining instances (by specifying the private IP address of each instance). Click **Next**.
 4. The **Interconnect** window is displayed.
Specify the IP addresses (private IP address of each instance) for interconnect, a Witness heartbeat, and a disk heartbeat.
For more information, refer to the following:
 - "Reference Guide" -> "Understanding Witness heartbeat resources"
 - "Reference Guide" -> "Understanding disk heartbeat resources"
 5. The **NP Resolution** window is displayed.

Specify HTTPNP and click **Next**.

For more information, refer to the following:

- "Reference Guide" -> "Understanding network partition resolution by HTTP method"

5. Adding group resources

- Defining a group

Create a failover group.

1. The **Group List** window is displayed.
Click **Add**.
2. The **Group Definition** window is displayed.
In **Name**, enter failover1 as a failover group name. Click **Next**.
3. The **Startup Servers** window is displayed.
Click **Next** without specifying anything.
4. The **Group Attributes** window is displayed.
Click **Next** without specifying anything.
5. The **Group Resource List** window is displayed.
On this page, add a group resource following the procedure below.

- Disk resource

Create a disk resource.

For more information, refer to the following:

- "Reference Guide" -> "Understanding Disk resource"

1. Click **Add** on the **Group Resource List** page.
2. The **Resource Definition of Group | failover1** window is displayed.
From the **Type** box, select **Disk resource** as a group resource type. In the **Name** box, enter the resource name. Click **Next**.
3. The **Dependency** window is displayed.
Click **Next** without specifying anything.
4. The **Recovery Operation** window is displayed.
Click **Next**.
5. The **Details** window is displayed.
In **Device Name**, enter the device name of the partition created in "4. Creating a block volume". Specify **Mount Point** and **File System**. Click **Finish** to finish the settings.

- Oracle Cloud virtual IP resource

In using EXPRESSCLUSTER in OCI, provides a mechanism to wait for the alive monitoring by a load balancer on a specific port of a node where operations are running.

For details on the Oracle Cloud virtual IP resource, refer to the following:

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

1. Click **Add** on the **Group Resource List** page.

2. The **Resource Definition of Group | failover1** window is displayed.
In the **Type** box, select **Oracle Cloud virtual IP resource** as a group resource type. In the **Name** box, enter the resource name. Click **Next**.
3. The **Dependency** window is displayed. Click **Next** without specifying anything.
4. The **Recovery Operation** window is displayed. Click **Next**.
5. In **Port Number**, enter the value specified in **Health check policy: port** in the load balancer settings (the backend set settings).
6. Click **Finish**.

6. Adding monitor resources

- Oracle Cloud virtual IP monitor resource

Provides a mechanism for monitoring the alive-monitoring port for the node where an Oracle Cloud virtual IP resource has been activated.

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

For details on the Oracle Cloud virtual IP monitor resource, refer to the following:

- "Reference Guide" -> "Understanding Oracle Cloud Virtual IP monitor resources"

- Oracle Cloud load balance monitor resource

Provides a mechanism for monitoring whether the same port number as the health-check port number has already been used, for the node where an Oracle Cloud virtual IP resource has not been activated.

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

For details on the Oracle Cloud load balance monitor resource, refer to the following:

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

7. Applying the settings and starting the cluster

Refer to the following:

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

4.4 Verifying the operations

Verify whether the created environment works properly by generating a monitoring error to fail over a failover group. If the cluster is running normally, the verification procedure is as follows:

1. Start the failover group (failover1) on the active node (server1).
In the **Status** tab on Cluster WebUI, confirm that the status of failover1 is **Online** at Server1.
Access the IP address of the front-end from the client to confirm that the connection to the active node is available.
2. Change **Operation mode** to **Verification mode** from the Cluster WebUI pull-down menu.
3. In the **Status** tab on Cluster WebUI, click the **Enable dummy failure** icon of ocvip1.
4. Upon three times of reactivation of the Oracle Cloud virtual IP resource (ocvip1), the failover group (failover1) enters an error status and fails over to node server2.
In the **Status** tab on Cluster WebUI, confirm that the status of failover1 is **Online** at server2.
Also make sure that, after the failover, the front-end IP address of the load balancer can be normally accessed.

Verifying the failover operation in case of a dummy failure is now complete. If necessary, perform operation checks for other failures.

ERROR MESSAGES

For details on the error messages related to resources and monitor resources, refer to the following:

- "Reference Guide" -> "Error messages"

NOTES AND RESTRICTIONS

6.1 HA cluster using a load balancer

6.1.1 Notes on OCI

- In designing a performance-oriented system, keep this in mind: OCI 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.

6.1.2 Notes on EXPRESSCLUSTER

- Set the OS startup time to be longer than **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"

- If a network failure occurs with the attachment method for a block volume set to be iSCSI, the deactivation of mirror disk resources or disk resources may fail.

Therefore, if iSCSI is specified as the attachment method, set the operation in response to a resource deactivation failure to be a cluster service stop and OS shutdown.

Refer to the following:

- "Getting Started Guide" -> "Notes and Restriction" -> "Communication port number"
- "Getting Started Guide" -> "Notes and Restriction" -> "Setting up Oracle Cloud virtual IP resources"
- "Getting Started Guide" -> "Notes and Restriction" -> "Setting up Oracle Cloud load balance monitor resources"
- "Reference Guide" -> "Notes on Oracle Cloud Virtual IP resources"
- "Reference Guide" -> "Notes on Oracle Cloud Virtual IP monitor resource"
- "Reference Guide" -> "Notes on Oracle Cloud load balance monitor resources"

6.2 Notes on creating an instance

6.2.1 Required packages

The following package is not included in the platform images (Oracle Linux 9.0 or later) provided by OCI.

- glibc-gconv-extra

Install the package after creating an instance.

LEGAL NOTICE

7.1 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.

7.2 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.
- Oracle and Java are registered trademarks of Oracle Corporation and its affiliates in the United States and other countries.
- Oracle Cloud Infrastructure is a registered trademark of Oracle Corporation and its affiliates in the United States and other countries.
- 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

© Copyright NEC Corporation 2024. All rights reserved.