

SkyCluster™ Deployment Guide for Oracle RAC on AWS

rev. 20.06-2020.07.02



Table of Contents

1	Introduction	3
1.1	Key Components	3
1.2	High Availability Architecture	4
1.3	Infrastructure-as-Code Deployment	4
2	Prerequisites	5
2.1	Required Knowledge	5
2.2	Getting access to SkyCluster AMI from AWS Marketplace	5
2.3	Uploading Oracle installation files to S3	6
2.4	Preparing the VPC	7
3	Deploying a Cluster	8
4	After Deploying a Cluster	9
4.1	Verifying cluster status.....	9
4.2	OS user accounts.....	9
4.3	Finalizing cluster configuration	10
4.4	Enabling termination protection.....	10
4.5	Installing database software (standalone or additional RAC db home)	10
4.6	Use of anti-virus software.....	10
4.7	Use of automatic configuration tools	10
4.8	Security hardening	10
5	Monitoring Cluster Health	11
6	Before Going Live	11
7	Deleting a Cluster	12
8	Additional Documentation.....	12
9	Contacting Technical Support	12

1 Introduction

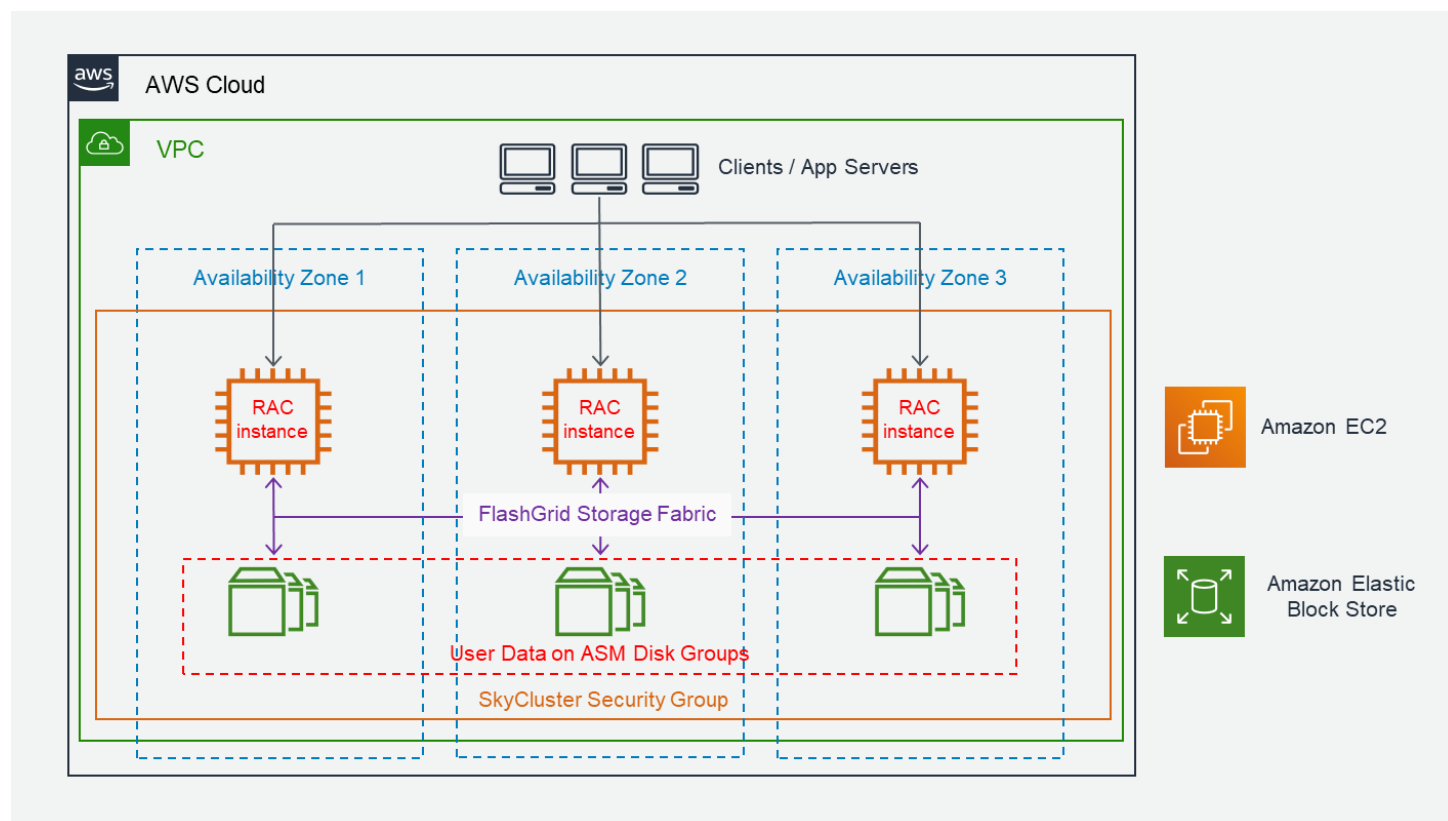
SkyCluster is an engineered cloud system that enables active-active database high availability infrastructure in public clouds. This guide provides step-by-step instructions for system and database administrators deploying SkyCluster with Oracle RAC on AWS cloud.

Additional information about the SkyCluster architecture is available in the following white paper: [“Mission-Critical Databases in the Cloud. Oracle RAC on Amazon EC2 Enabled by FlashGrid® SkyCluster.”](#)

1.1 Key Components

Key components of SkyCluster 20.06 for AWS:

- FlashGrid Storage Fabric: ver. 19.06
- FlashGrid Cloud Area Network: ver. 19.10
- FlashGrid Diagnostics: ver. 20.02
- SkyCluster Health Checker ver. 20.04
- Oracle Database: ver. 19c, 18c, 12.2.0.1, 12.1.0.2, or 11.2.0.4.
- Oracle Grid Infrastructure: ver. 19c. (Versions 18c and 12.2.0.1 available on request)
- Operating System: Oracle Linux 7, Red Hat Enterprise Linux 7
- Amazon EC2 instances: R5, R5D, R4, M5, M5D, M4, i3, i3en, X1, X1E, Z1D, High Memory
- Disks: EBS GP2 volumes or local SSDs



SkyCluster Network Diagram

1.2 High Availability Architecture

By leveraging Oracle RAC active-active database clustering and synchronous data mirroring across nodes and AZs, SkyCluster enables near-zero (seconds) Recovery Time Objective (RTO) and zero Recovery Point Objective (RPO) in case of a failure of a single node instance, of a single EBS volume, or in case of one AZ failure.

If the cluster has 3+ database nodes then two simultaneous database node failures can be tolerated without causing loss of database service.

Standard Oracle Client functionality provides mechanisms for application failover from a failed node, including Transparent Application Failover (TAF)

1.3 Infrastructure-as-Code Deployment

SkyCluster is delivered as an AWS CloudFormation template that automates configuration of multiple components required for a database cluster. SkyCluster Launcher is an online tool that simplifies the deployment process by guiding through the cluster configuration parameters and generating CloudFormation templates.

2 Prerequisites

2.1 Required Knowledge

Working knowledge of the following AWS services is required for successful deployment of SkyCluster on AWS: EC2, VPC, EBS, CloudFormation, S3, IAM, Marketplace

2.2 Getting access to SkyCluster AMI from AWS Marketplace

To be able to create a cluster your AWS account must have an active subscription to the selected SkyCluster AMI. Otherwise deployment will fail when creating VM instances. The SkyCluster AMIs are based on either Oracle Linux 7 or RHEL 7.

To get access to the SkyCluster AMI

1. Open SkyCluster product page in AWS Marketplace:
 - [Oracle Linux 7 based AMI](#)
 - [RHEL 7 based AMI](#)
2. Click **Continue** button
3. Select **Manual Launch** tab
4. Click **Accept Software Terms** button

Software fees charged through AWS Marketplace include SkyCluster software license and 24x7 Mission-Critical support plan. The fees are charged per cluster node instance and depend on the selected EC2 instance type and size. *Hourly* and *Annual* subscription models are available. Pricing information is available on the AWS Marketplace product pages – see the two links above.

2.3 Uploading Oracle installation files to S3

During cluster initialization Oracle installation files will be downloaded from an S3 bucket. The list of files that must be placed in the S3 bucket will be shown by the SkyCluster Launcher tool. The same S3 bucket can be used for deploying multiple clusters.

Two options are available for allowing access to the files in the S3 bucket for the cluster node instances:

- Enabling public access to each file for the duration of cluster deployment
OR
- Assigning the cluster node instances an IAM role that has permissions for accessing files in the bucket

To allow public access to the files in S3

1. Create an S3 bucket/folder for uploading the installation files
2. Upload the required files to the S3 bucket/folder
3. In S3 Management Console navigate to the bucket and the folder to see the list of files
4. Select all files
5. Click *More* -> *Make Public*
6. You can disable public access after the cluster completes initialization

To use an IAM role for access to the files in S3

1. Create an S3 bucket/folder for uploading the installation files
2. Upload the required files to the S3 bucket/folder
3. In *IAM Management Console* create a new policy named **GetOracleFilesFromS3** that allows **s3:GetObject** action on all uploaded files. See an example below.
4. In *IAM Management Console* create a new role named **GetOracleFilesFromS3** and attach the **GetOracleFilesFromS3** policy to it.
5. Use the **GetOracleFilesFromS3** role when configuring cluster parameters in the SkyCluster Launcher tool.

Example of an IAM policy:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "Stmt1508867055000",
      "Effect": "Allow",
      "Action": [
        "s3:GetObject"
      ],
      "Resource": [
        "arn:aws:s3:::mybucket/mydirectory/*"
      ]
    }
  ]
}
```

Additional information about IAM and IAM best practices is available at:

<https://docs.aws.amazon.com/IAM/latest/UserGuide/introduction.html>

<https://docs.aws.amazon.com/IAM/latest/UserGuide/best-practices.html>

2.4 Preparing the VPC

When creating a new cluster you have two options:

- **Automatically create a new VPC.**
This option is usually used for test clusters isolated in their own sandbox VPCs. A VPC will be created together with the required subnets, placement group(s), and security groups. By default the VPC will be created with CIDR 10.100.0.0/16
- **Create the cluster in an existing VPC.**
This option is used for majority of production deployments where other systems (e.g. app servers) share the same VPC as the cluster. You will need to provide the VPC ID in the SkyCluster Launcher tool and subnet IDs and security group IDs in the CloudFormation Manager.

If using an existing VPC then make sure that the following pre-requisites are met before creating a cluster:

- The VPC may have any CIDR that does not overlap with 192.168.0.0/16, for example 10.100.0.0/16. If you have to use VPC with CIDR that overlaps with 192.168.0.0/16 then please request a customized configuration file from FlashGrid support.
- The VPC has a subnet in each of the availability zones used for the cluster nodes.
- The VPC has an S3 endpoint configured (required unless public IPs can be enabled for access to S3)
- If you choose to enable Public IPs on the VM instances then the VPC must have Internet Gateway configured.
- The VPC has a security group with the following ports open for inbound traffic:
 - UDP ports 4801, 4802, 4803 between any members of the security group
 - TCP ports 1521, 1522 for SCAN and Local Listener access to the database nodes from app servers and other database clients. These are default port numbers that can be changed in the SkyCluster Launcher tool.
 - TCP port 22 for SSH access to the cluster nodes
 - TCP port 5901 if you choose to use VNC for creating a database using DBCA in GUI mode
- Create a *Spread Placement Group* that will be used for the VM instances. The use of the spread placement group is not required only if each VM is placed in a separate availability zone, e.g. 2 db nodes + 1 quorum node across three AZs, or 3 db nodes + 2 quorum nodes across five AZs. It is required in all other cases to guarantee that all nodes are running on separate physical hosts.

3 Deploying a Cluster

The FlashGrid SkyCluster Launcher tool simplifies deployment of Oracle RAC clusters in AWS by automating the following tasks:

- Creating and configuring EC2 VPC, subnets, security group (optional)
- Creating EBS volumes and launching EC2 instances for all nodes in the cluster
- Installing and configuring FlashGrid Cloud Area Network
- Installing and configuring FlashGrid Storage Fabric
- Installing and patching Oracle Grid Infrastructure software
- Configuring Grid Infrastructure cluster
- Installing and patching Oracle Database software
- Creating ASM disk groups

To create a cluster

1. Log in to AWS Management Console with a user account that has the following privileges:
 - AWSCloudFormationFullAccess
 - AmazonEC2FullAccess
 - AmazonVPCFullAccess (required only if creating a new VPC)
2. Open SkyCluster Launcher tool:
 - Start with one of the standard configurations at <https://www.flashgrid.io/skycluster-for-aws>
 - or, if you have a custom configuration file, upload it at <https://2006.cloudprov.flashgrid.io>
3. Configure parameters of the cluster
4. Click *Validate Configuration* button
5. If verification passes then click *Launch Cluster* button, which will take you to AWS CloudFormation Manager
6. Click *Next*
7. Select your SSH key
8. If using an existing VPC, then select a spread placement group, subnet(s), and security group.
9. Click *Next*
10. On the *Options* page:
 - If you added tags in SkyCluster Launcher then **do not** add the same tags in CloudFormation Manager
 - If the cluster is for production use then expand the *Advanced* options and enable *Termination Protection*
11. Click *Next*
12. Click *Create*
13. Wait until the status of the stack changes to *CREATE_COMPLETE*
14. If creating the stack fails:
 - a) Check for the cause of the failure on the *Events* tab
 - b) Correct the cause of the error
 - c) Delete the failed stack
 - d) Repeat the steps for creating a new stack
15. Use EC2 Management Console to get IP addresses of the cluster node instances
16. SSH to the first (as it was specified on the cluster configuration page) cluster node as user `fg@`
17. The welcome message will show the current initialization status of the cluster: in progress, failed, or completed.
18. If initialization is still in progress then wait for it to complete (this includes Oracle software installation and configuration). You will receive a broadcast message when initialization completes or fails. Cluster initialization takes 1 to 2 hours depending on configuration.

4 After Deploying a Cluster

4.1 Verifying cluster status

On any of the cluster nodes run `flashgrid-cluster` command to verify that the cluster status is *Good* and all checks are passing.

```
[fg@rac1 ~]$ flashgrid-cluster
FlashGrid 18.07.10.46032 #95f2b5603f206af26482ac82386b1268b283fc3c
License: via Marketplace Subscription
Support plan: 24x7
~~~~~
FlashGrid running: OK
Clocks check: OK
Configuration check: OK
Network check: OK

Querying nodes: quorum, rac1, rac2 ...

Cluster Name: myrac
Cluster status: Good
-----
Node      Status  ASM_Node  Storage_Node  Quorum_Node  Failgroup
-----
rac1     Good   Yes      Yes           No           RAC1
rac2     Good   Yes      Yes           No           RAC2
racq     Good   No       No            Yes          QUORUM
-----
-----
GroupName  Status  Mounted  Type      TotalMiB  FreeMiB  OfflineDisks  LostDisks  Resync  ReadLocal  Vote
-----
GRID      Good   AllNodes  NORMAL   12588     3376     0              0          No     Enabled    3/3
DATA      Good   AllNodes  NORMAL   2048000   2048000  0              0          No     Enabled    None
FRA       Good   AllNodes  NORMAL   1024000   1024000  0              0          No     Enabled    None
-----
```

4.2 OS user accounts

During cluster initialization the following OS user accounts are created:

- *fg* - the user account used to SSH to the VMs with the SSH key that was selected when creating the cluster configuration. It can also be used for running FlashGrid Storage Fabric or FlashGrid Cloud Area Network utilities. The user *fg* has sudo rights.
- *grid* - Grid Infrastructure owner. GI environment variables are preconfigured.
- *oracle* - Database home owner. Database environment variables, except `ORACLE_SID` and `ORACLE_UNQNAME`, are preconfigured. After creating a database you can configure `ORACLE_SID` and `ORACLE_UNQNAME` by editing `/home/oracle/.bashrc` file on each database node.

Note that no passwords are configured for any users. Also password-based SSH authentication is disabled in `/etc/ssh/sshd_config`. Key-based authentication is recommended for better security. Creating passwords for any user is not recommended.

User *fg* has sudo rights and allows switching to any other user without requiring a password (which is not configured by default). Example:

```
$ sudo su - grid
```

Users *fg*, *grid*, and *oracle* have key-based SSH access configured between the nodes of the cluster. The corresponding key pairs are generated automatically during cluster initialization. For example, if you are logged in to *node1* as user *fg* then you can SSH into *node2* by simply running `'ssh node2'` without entering a password or providing a key.

4.3 Finalizing cluster configuration

See knowledge base articles for performing the following steps:

1. Changing temporary ASM passwords: <https://kb.flashgrid.io/asm-password>
2. Creating a database: <https://kb.flashgrid.io/createdb>
3. Connecting clients to a database: <https://kb.flashgrid.io/connect-clients>

4.4 Enabling termination protection

If termination protection was not enabled when creating the cluster and if the cluster is for production use then it is strongly recommended to enable termination protection:

- Enable instance termination protection for each cluster node instance.
- Enable termination protection for the CloudFormation stack.

4.5 Installing database software (standalone or additional RAC db home)

In most cases manual installation of database software is not required. However, if you need a standalone (non-RAC) database or an additional RAC database home then follow Oracle Database documentation for installing the database software.

4.6 Use of anti-virus software

If anti-virus software has to be used then it is recommended to configure it in a way that avoids putting any files in quarantine. Automatic quarantine of files creates risk of the cluster downtime in case of a false positive detection on a critical system file on multiple nodes of the cluster.

4.7 Use of automatic configuration tools

Automatic configuration tools (e.g. Ansible, Salt, etc.) must be used with extra care. Incorrect modification of a critical system file (e.g. `/etc/resolv.conf`) on multiple cluster nodes may cause cluster downtime. Note that many critical system configuration files are protected with immutable attribute and have warnings in them. Do not remove the immutable attribute or allow automatic modification of such files unless absolutely necessary.

4.8 Security hardening

Cluster nodes are deployed using RHEL 7 or Oracle Linux 7 images that have main security best practices implemented by default. The following steps are recommended, in case additional security hardening is required:

- 1) Request FlashGrid support to review the list of required changes.
- 2) Back up all cluster nodes: <https://kb.flashgrid.io/backup-restore/backup-and-restore-in-aws>
- 3) Implement the required changes on all nodes.
- 4) Restart the entire cluster: <https://kb.flashgrid.io/maintenance/maintenance-aws#restarting-the-entire-cluster>
- 5) Verify health of the cluster: `$ sudo skycluster-health-check`
- 6) In case of errors, roll back the changes or restore the nodes from backup.

5 Monitoring Cluster Health

The following methods of monitoring cluster health are available:

- *skycluster-health-check* utility checks multiple items including database configuration, storage, OS kernel, config file modifications, errors in the logs, and other items that may affect health of the cluster or could help with troubleshooting. It is recommended for manual checks only.
- *flashgrid-cluster* utility displays status of the storage subsystem (FlashGrid Storage Fabric and ASM) and its main components. The utility can be used in monitoring scripts. It returns a non-zero value if status of the cluster is *Warning* or *Critical*.
- Alerts about failures are recorded in system log and can be analyzed by 3rd-party tools.
- Email alerts can be sent to one or several email addresses.
- ASM disk group monitoring and alerting via Oracle Enterprise Manager.

To test email alerts

1. On all nodes (including quorum node) run

```
$ flashgrid-node test-alerts
```

2. Check that test alert emails were received from all cluster nodes at each of the configured email addresses.

To modify the list of email alert recipients

As user *fg@* on any database node run

```
$ flashgrid-cluster set-email-alerts name1@host1 name2@host2 ...
```

Note that by default the *From* address is set to *flashgrid@localhost.localdomain*. This will ensure that delivery failure notifications are sent to root's mailbox on the originating node, which can help with troubleshooting delivery issues. It is recommended to add this address to the whitelist of senders on the receiving email server and in the email clients.

6 Before Going Live

Before switching the cluster to live use:

1. Verify health of the cluster: `$ sudo skycluster-health-check`
2. Confirm that email alerts are configured and delivered: `$ flashgrid-node test-alerts`
3. Upload diags to FlashGrid support: `$ sudo flashgrid-diags upload-all`
4. Stop the cluster and back up all cluster nodes: <https://kb.flashgrid.io/backup-restore/backup-and-restore-in-aws>
5. Start the cluster and do final check of the cluster health: `$ sudo skycluster-health-check`

7 Deleting a Cluster

To delete a cluster

1. Disable instance termination protection for each cluster node if it was enabled
2. Open AWS CloudFormation Manager console
3. Disable termination protection for the corresponding CloudFormation stack if it was enabled
4. Delete the stack corresponding to the cluster
5. If any EBS volumes were added after deploying the cluster, those volumes must be deleted separately
6. If any AMI images or volume snapshots were created after deploying the cluster, those AMIs and snapshots must be deleted separately

8 Additional Documentation

Maintenance Tasks on AWS: <https://www.kb.flashgrid.io/maintenance/maintenance-aws>

Backup and Restore Best Practices on AWS: <https://www.kb.flashgrid.io/backup-restore/backup-and-restore-in-aws>

Troubleshooting: <https://www.kb.flashgrid.io/troubleshooting>

FlashGrid Storage Fabric CLI Reference Guide: <https://www.kb.flashgrid.io/cli-ref/sf-cli>

FlashGrid Cloud Area Network CLI Reference Guide: <https://www.kb.flashgrid.io/cli-ref/clan-cli>

9 Contacting Technical Support

For technical help with SkyCluster please open a support request at <https://www.flashgrid.io/support/>

To expedite troubleshooting please also collect and upload diagnostic data to the secure storage used by FlashGrid support by running the following command:

```
$ sudo flashgrid-diags upload-all
```

For reporting *emergency* type of issues that require immediate attention please also use the 24/7 telephone hotline: +1-650-641-2421 ext 7. Please note that use of the 24/7 hotline is reserved for emergency situations only.

Copyright © 2016-2020 FlashGrid Inc. All rights reserved.

This document is provided for information purposes only, and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document, and no contractual obligations are formed either directly or indirectly by this document.

FlashGrid is a registered trademark of FlashGrid Inc. SkyCluster is a trademark of FlashGrid Inc. Amazon and Amazon Web Services are registered trademarks of Amazon.com Inc. and Amazon Web Services Inc. Oracle and Java are registered trademarks of Oracle and/or its affiliates. Red Hat is a registered trademark of Red Hat Inc. Other names may be trademarks of their respective owners.