

FlashGrid® SkyCluster™

Deployment Guide

for Oracle RAC in AWS

rev. 19.02-2019.02.25



Table of Contents

1	Introduction.....	3
2	Prerequisites.....	4
2.1	Getting Access to FlashGrid AMI.....	4
2.2	Creating AMI with Encryption of Boot Volume.....	4
2.3	Uploading Oracle Installation Files to S3.....	4
2.4	Preparing the VPC.....	4
3	Deploying a Cluster.....	6
4	After Deploying a Cluster.....	7
4.1	Verifying cluster health.....	7
4.2	OS User Accounts.....	7
4.3	Finalizing Cluster Configuration.....	8
4.4	Enabling instance termination protection.....	8
4.5	Deleting a cluster.....	8
5	Measuring Performance.....	9
6	Monitoring Cluster Health.....	10
7	Additional Documentation.....	10
8	Contacting FlashGrid Technical Support.....	10

1 Introduction

FlashGrid 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 FlashGrid SkyCluster with Oracle RAC in AWS cloud.

Key components of FlashGrid SkyCluster 19.02 for AWS:

- FlashGrid Storage Fabric: ver. 19.02
- FlashGrid Cloud Area Network: ver. 18.11
- FlashGrid Diagnostics: ver. 19.02
- Oracle Database: ver. 18c, 12.2.0.1, 12.1.0.2, or 11.2.0.4
- Oracle Grid Infrastructure: ver. 18c or 12.2.0.1
- Operating System: Oracle Linux 7, Red Hat Enterprise Linux 7
- Amazon EC2 instances: R5, R5D, R4, M5, M5D, M4, i3, X1, X1E, Z1D
- Disks: EBS GP2 volumes or local SSDs

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

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

2 Prerequisites

2.1 Getting Access to FlashGrid AMI

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

To get access to the FlashGrid AMI

1. Open FlashGrid product page at 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

2.2 Creating AMI with Encryption of Boot Volume

The FlashGrid SkyCluster Launcher tool has an option for enabling encryption on EBS volumes. However, this option does not cover the boot volume. If EBS encryption must be enabled on the boot volume too then you need to create a new encrypted AMI based on the FlashGrid SkyCluster AMI.

To create an encrypted AMI

1. Launch a single instance of t2.micro (or any other) type using the FlashGrid AMI.
2. Stop the instance.
3. Create a new AMI from the instance and name it "*FlashGrid SkyCluster AMI not encrypted YY-MM-DD*".
4. Copy the "*FlashGrid SkyCluster AMI not encrypted YY-MM-DD*" AMI to "*FlashGrid SkyCluster AMI Encrypted YY-MM-DD*" and enable **Encryption** option when copying.

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. For instructions on how to enable access to the files see <https://kb.flashgrid.io/uploading-s3>

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 FlashGrid 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://1902.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 approximately 90 minutes.

4 After Deploying a Cluster

4.1 Verifying cluster health

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:

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

4.5 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 Measuring Performance

DBMS_RESOURCE_MANAGER.CALIBRATE_IO procedure provides an easy way for measuring storage performance including maximum bandwidth, random IOPS, and latency. The CALIBRATE_IO procedure generates I/O through the database stack on actual database files. The test is read-only and it is safe to run it on any existing database. It is also a good tool for directly comparing performance of two storage systems because the CALIBRATE_IO results do not depend on any non-storage factors, such as memory size or the number of CPU cores.

To measure storage performance with CALIBRATE_IO

1. Create or load a database on the corresponding ASM disk group
2. Make sure the total size of the database files is larger than 5 GB per disk. If needed, create an additional large table space / data file.
3. Customize the first parameter in the SQL code below with the number of disks corresponding to your storage setup. Keep the second parameter (max latency) with the minimum allowed value of 10 milliseconds.
4. Connect to the database with sqlplus and run the customized SQL code.
5. Wait for the CALIBRATE_IO to complete. This may take 10 to 30 minutes.

Example of running CALIBRATE_IO on a 2-node cluster with i3.16xlarge instances and eight NVMe SSDs per node

```
SET SERVEROUTPUT ON;
DECLARE
  lat INTEGER;
  iops INTEGER;
  mbps INTEGER;
BEGIN DBMS_RESOURCE_MANAGER.CALIBRATE_IO (16, 10, iops, mbps, lat);
DBMS_OUTPUT.PUT_LINE ('Max_IOPS = ' || iops);
DBMS_OUTPUT.PUT_LINE ('Latency = ' || lat);
DBMS_OUTPUT.PUT_LINE ('Max_MB/s = ' || mbps);
end;
/

Max_IOPS = 1375694
Latency = 0
Max_MB/s = 27338

PL/SQL procedure successfully completed.
```

6 Monitoring Cluster Health

The following methods of monitoring cluster health are available:

- The *flashgrid-cluster* utility displays status of the cluster and its main components.
- 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.

7 Additional Documentation

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

Backup and Restore Best Practices in 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>

8 Contacting FlashGrid Technical Support

For help with troubleshooting an issue on an existing FlashGrid cluster please use Technical Support Request form located 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 24x7 telephone hotline: +1-650-641-2421 ext 7. Please note that use of the 24x7 hotline is reserved for emergency situations only.

Copyright © 2016-2019 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. 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.