flashgrid

FlashGrid® Cluster

for Oracle Database and Oracle RAC on Azure Cloud

Deployment Guide

rev. 24.11-2024.12.10

Table of Contents

1 Introduction

FlashGrid Cluster is an engineered cloud system that enables high availability and high-performance shared storage for running Oracle databases in public clouds. This guide provides step-by-step instructions for system and database administrators deploying FlashGrid Cluster with Oracle Database (RAC, failover HA, or standalone) on Azure cloud.

Key components of FlashGrid Cluster on Azure:

- FlashGrid Storage Fabric software
- FlashGrid Cloud Area Network software
- Oracle Database: 19c, 18c, 12.2.0.1, 12.1.0.2, or 11.2.0.4
- Oracle Grid Infrastructure: 19c
- Operating Systems:
  - **Oracle Linux:** 7, 8, or 9
  - **Red Hat Enterprise Linux** (RHEL): 8, or 9
- Azure VMs:
  - **General purpose:** Dsv5
  - **Memory optimized:** Ebsv5, Esv5, Mbsv3, Msv3
- Disks: Premium SSD, Premium SSD v2

FlashGrid Cluster is delivered as an Azure Resource Manager template that automates configuration of multiple components required for a database cluster. FlashGrid Launcher is an online tool that simplifies the deployment process by guiding through the cluster configuration parameters and generating Azure Resource Manager templates.

Additional information about the FlashGrid Cluster architecture for Oracle RAC is available in the following white paper: *"Mission-Critical Databases in the Cloud. Oracle RAC in Microsoft Azure Enabled by FlashGrid®."*

2 Prerequisites

The following prerequisites are required for automated deployment of a FlashGrid Cluster:

- **Azure Storage Blob Container** with Oracle installation files that will be downloaded to the cluster nodes during cluster initialization. The list of files that must be placed in the Storage Container will be shown in FlashGrid Launcher. The corresponding storage account must have access for *'All networks'* enabled in *'Firewall and virtual networks'* settings.
- **Microsoft.Storage** service endpoint configured for the VNet. Having the storage service endpoint allows access to the storage container from the VMs. If Microsoft.Storage service endpoint is not added, and public IPs not assigned then cluster initialization will fail because downloading Oracle files from the VMs will not be possible. See detailed instructions.
- **Azure subscription with sufficient quotas** for creating the required number and type of VMs and sufficient number and size of Premium Managed Disks.
- **SSH key pair** that will be used for accessing the VMs. Use of passwords instead of the key pair is not supported. To create a new key pair use *ssh-keygen* in Linux or *puttygen* in Windows. In the FlashGrid Launcher tool you will need to provide the public key that will be placed on the VMs. Example of a valid public key pair format:

  ```
  ssh-rsa <PublicKeyBody>
  ```

  Keep blank if using Azure managed SSH keys for VMs access.
- **Properly configured Network Security Group (NSG)** when deploying in an existing VNet. You have a choice of attaching an NSG to the VMs or using the NSG attached to the subnet. The following ports should be open:

- Inbound and Outbound: All traffic between the cluster nodes.
- Inbound: 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 FlashGrid Launcher tool.
- Inbound: TCP port 22 for SSH access to the cluster nodes
- Inbound: TCP port 5901 if you choose to use VNC for creating a database using DBCA in GUI mode with direct connection (vs. SSH tunnel)
- Inbound access to the ports listed above must be allowed only from those security groups or IP ranges that require such access. Do not configure *Anywhere* or *0.0.0.0/0* as allowed sources.

FlashGrid recommends configuring the NSG rules by using an Application Security Group (ASG) for the cluster node VMs. You can configure one ASG per cluster or a separate ASG for each cluster.

Note: for deploying FlashGrid Cluster on **dedicated hosts** please refer to the following knowledge base article: https://support.flashgrid.io/hc/en-us/articles/4413508587415-Deploying-on-dedicated-hosts-Azure

3 Deploying a Cluster

The FlashGrid Launcher tool simplifies deployment of database clusters on Azure by automating the following tasks:

- Creating cloud infrastructure: VMs, storage, and optionally network
- Installing and configuring FlashGrid Cloud Area Network
- Installing and configuring FlashGrid Storage Fabric
- Installing, configuring, and patching Oracle Grid Infrastructure
- Installing and patching Oracle Database software
- Creating ASM disk groups

To create a cluster

1. Open FlashGrid Launcher tool with one of the standard configurations:
   - Oracle RAC: https://www.flashgrid.io/products/flashgrid-for-oracle-rac-on-azure
   - Oracle Single-Instance with failover HA: https://www.flashgrid.io/products/flashgrid-for-oracle-failover-ha-on-azure/
   - Oracle Single-Instance with high-throughput storage: https://www.flashgrid.io/products/flashgrid-for-oracle-db-on-azure/

   or, if you have a custom configuration file, upload it at https://2411.cloudprov.flashgrid.io/
2. Configure parameters of the cluster
3. Click *Validate Configuration* button
4. If verification passes then click *Launch Cluster* button, which will take you to Azure Resource Manager
5. Select *Resource group -> Create new*. By having the cluster in a separate resource group, you can later delete the entire cluster by simply deleting the resource group.
6. Enter a name for the new resource group that will contain the cluster. A name matching the cluster name is recommended.
7. Select your target location (region)
8. If you did not provide an SSH key, select *Use existing key stored in Azure* in *SSH public key source* and specify a stored key.
9. Check *'I agree to the terms and conditions state above'*
10. Click *Purchase*
11. Open list of Notifications (bell icon) and click *'Deployment in progress…'*

12. Wait until the deployment status changes to *Succeeded*
13. If the deployment fails:
    a) Check for the cause of the failure in the *Operation details*
    b) Correct the cause of the error
    c) Delete the failed resource group
    d) Repeat the steps for creating a new resource group
14. SSH to the first (as it was specified on the cluster configuration page) cluster node VM as user ***az-admin@***
    Note: If you selected to create a new VNet and connecting through a public IP address then need to edit VPC Network configuration attached to the database nodes. In the Azure Virtual Machine settings select Networking, and for the *allow-tcp22* inbound port rule set *Source* to your client system IP.
15. The welcome message will show the current initialization status of the cluster: in progress, failed, or completed.
16. 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 for RAC/HA cluster or 30 to 40 minutes for a single-instance with no HA.

Note: for deploying FlashGrid Cluster with ***SELinux*** please refer to the following knowledge base article:
https://support.flashgrid.io/hc/en-us/articles/26368224225687-How-to-enable-disable-SELinux

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 18.07.15.48564 #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:

- *az-admin* - the user account used to SSH to the VMs with the SSH key that was selected when creating the cluster configuration. The user has sudo rights.
- *fg* - can be used for running FlashGrid Storage Fabric or FlashGrid Cloud Area Network utilities. The user has sudo rights. The user has key-based SSH configured between *all* nodes of the cluster.
- *grid* - Grid Infrastructure (GI) owner. GI environment variables are preconfigured. The user has key-based SSH configured between all *database* nodes of the cluster.
- *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. The user has key-based SSH configured between all *database* nodes of the cluster.

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.

Users *az-admin* and *fg* have 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. Creating a database: https://support.flashgrid.io/hc/en-us/articles/1500011215081
2. Connecting clients to a database: https://support.flashgrid.io/hc/en-us/articles/1500011176122

Note: ACFS support on RHEL may require the additional Oracle Clusterware patch. Please refer to Oracle Doc ID 1369107.1 for ACFS patch information.

4.4 Adding a protection lock for the cluster

It is strongly recommended to add a lock to the cluster resource group to protect it against accidental deletion or modification.

4.5 Installing an additional database home

In most cases manual installation of database software is not required. However, if you need to install an additional database home, then follow Oracle Database documentation for installing the database software.

4.6 Use of anti-virus and other third-party software

If anti-virus software must 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.

Any proprietary kernel modules installed by third-party software create risks to reliable operation of the system. Such proprietary kernel modules are not tested or supported by FlashGrid, Red Hat, or Oracle Linux. Proprietary kernel modules may consume kernel resources and may create instability, especially under high load. Symptoms may include kernel crashes, network disruptions, storage i/o disruptions, node evictions, and cluster brown-out. If such reliability issue is encountered and no other root cause can be readily identified, FlashGrid support reserves the right to request removal of all proprietary kernel modules before continuing investigation.

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

For applying security hardening to the OS using CIS aligned security profiles, see https://support.flashgrid.io/hc/en-us/articles/5883226799639

For applying a different hardening profile, the following steps are recommended:

1) Request FlashGrid support to review the list of required changes.
2) Back up all cluster nodes: https://support.flashgrid.io/hc/en-us/articles/1500011214501
3) Implement the required changes on all nodes.
4) Restart the entire cluster: https://support.flashgrid.io/hc/en-us/articles/4404886778519

5) Verify health of the cluster as user fg:

```
$ flashgrid-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:

- *flashgrid-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 as user *fg*:

   $ `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 (run commands as user *fg*):

1. Verify health of the cluster: `$ flashgrid-health-check`
2. Confirm that email alerts are configured and delivered: `$ flashgrid-node test-alerts`
3. Upload diags to FlashGrid support: `$ flashgrid-diags upload-all`
4. Stop the cluster and back up all cluster nodes: https://support.flashgrid.io/hc/en-us/articles/4404887112215-Shutting-down-entire-cluster-Azure-
5. Start the cluster and do final check of the cluster health: `$ flashgrid-health-check`

7 Deleting a cluster

To delete a cluster

1. Delete any protection lock(s) for the resource group
2. Delete the resource group corresponding to the cluster

8 Additional Documentation

Knowledge Base: https://support.flashgrid.io/hc/en-us/categories/1500001520222-FlashGrid-Cluster-on-Azure

Backup Best Practices in Azure: https://support.flashgrid.io/hc/en-us/articles/1500011214501

FlashGrid Storage Fabric CLI Reference Guide: https://support.flashgrid.io/hc/en-us/articles/1500011214681

FlashGrid Cloud Area Network CLI Reference Guide: https://support.flashgrid.io/hc/en-us/articles/1500011214661

9 Contacting Technical Support

For technical help with FlashGrid Cluster 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 as user *fg*:

```
$ 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.