

Oracle RAC in AWS Enabled by FlashGrid

Automated Provisioning Guide
rev. 2017-02-18



Table of Contents

1	Introduction	3
2	Compatibility	3
3	Prerequisites	3
4	Provisioning a Cluster.....	4
4.1	Creating cluster with CloudFormation.....	4
4.2	Verifying cluster health	5
4.3	Changing ASM passwords	5
4.4	Configuring NTP service	5
4.5	Deleting a cluster	6
5	Creating ASM Disk Groups	7
6	Installing Database Software	11
7	Creating a Database	11
8	Measuring Performance	12
9	Maintenance Tasks	13
9.1	Rebooting a node.....	13
9.2	Powering off a node.....	13
9.3	Shutting down an entire cluster.....	13
9.4	Adding EBS volumes.....	14
9.5	Removing EBS volumes	14
9.6	Updating FlashGrid software	14
9.7	Updating Linux kernel	15
10	FlashGrid Tools and Commands.....	16
10.1	flashgrid-ca.....	16
10.2	flashgrid-fix-grid-dg-ca	16
10.3	flashgrid-create-dg.....	16
10.4	flashgrid-cluster	16
10.5	flashgrid-node	17
10.6	flashgrid-dg	18
11	Troubleshooting.....	19
12	Contacting FlashGrid Technical Support.....	19

1 Introduction

FlashGrid Cloud Area Network software and FlashGrid Storage Fabric software enable running Oracle RAC clusters in AWS cloud. FlashGrid Cloud Provisioning simplifies the deployment process by automating configuration of multiple components required for an Oracle RAC cluster. This guide provides step-by-step instructions for system and database administrators deploying Oracle RAC in AWS using FlashGrid Cloud Provisioning. Additional information about the solution is available in the following white paper: “*Mission-Critical Databases in the Cloud. Oracle RAC on Amazon EC2 Enabled by FlashGrid®.*”

2 Compatibility

The following versions of software are covered in this guide:

- Oracle Database: ver. 12.1.0.2 or 11.2.0.4 with the latest PSU
- Oracle Grid Infrastructure: ver. 12.1.0.2 with the latest PSU
- Operating System: Oracle Linux 7.3, Red Hat Enterprise Linux 7.3, CentOS 7.3
- FlashGrid Storage Fabric: ver. 17.01
- FlashGrid Cloud Area Network: ver. 17.01
- FlashGrid Cloud Provisioning for AWS: ver. 17.01

The solution can be deployed on the following Amazon EC2 instance types: M4, R4, X1.

3 Prerequisites

The following prerequisites are required for automated provisioning of a FlashGrid cluster:

- **FlashGrid AMI.** FlashGrid will share the AMI with your account. The AMI provided by FlashGrid is based on Oracle Linux 7.3. Please contact your FlashGrid representative if you need to use your own custom AMI.
- **S3 bucket** with FlashGrid and Oracle software files that will be downloaded to the cluster nodes during provisioning. The following files must be placed in the S3 bucket:
 - FlashGrid Storage Fabric RPM: flashgrid-sf-17.01
 - FlashGrid Cloud Area Network RPM: flashgrid-clan-17.01
 - oracle-instantclient12.1-basic-12.1.0.2.0-1.x86_64.rpm
 - Grid Infrastructure: linuxamd64_12102_grid_1of2.zip, linuxamd64_12102_grid_2of2.zip
 - Grid Infrastructure Jan'17 Patch Set Update: p24917825_121020_Linux-x86-64.zip
 - OPatch patch [version 12.2.0.1.8](#): p6880880_121010_Linux-x86-64.zip
 - Oracle Database (EE or SE2):
 - EE: linuxamd64_12102_database_1of2.zip, linuxamd64_12102_database_2of2.zip
 - SE2: linuxamd64_12102_database_se2_1of2.zip, linuxamd64_12102_database_se2_2of2.zip
- **CloudFormation JSON template file.** Your FlashGrid representative will generate the template file based on the required attributes of the cluster. Please provide your FlashGrid representative with a complete *FlashGrid Deployment Checklist* document.

4 Provisioning a Cluster

The FlashGrid Cloud Provisioning software simplifies provisioning 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

4.1 Creating cluster with CloudFormation

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. Click on *Services* button and select *CloudFormation* under *Management Tools*
3. Click *Create Stack* button
4. Paste the URL to the CloudFormation template JSON file provided to you by FlashGrid and click *Next*
5. Provide a name for the stack, e.g. name of the cluster and click *Next*
6. Optionally specify tags that will be assigned to all resources in the stack
7. Click *Next*
8. Click *Create*
9. Wait until the status of the stack changes to `CREATE_COMPLETE`
10. 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

Note: An error of the following type was observed in rare cases: "*route table rtb-f15b6196 and network gateway igw-f4e5ca90 belong to different networks*". If you encounter such error then delete the failed stack and try creating a new stack again.

11. Wait for the provisioning process to complete
12. After the provisioning process completes, wait 5 more minutes for the node instances to start and all node configuration tasks to complete.
13. Get IP addresses of the newly launched instances on the *Outputs* tab
14. SSH to each of the instance as `root@`
15. If the node and cluster initialization are still in progress or failed then you will see a corresponding welcome message. If there is no welcome message on any of the nodes then the cluster initialization has completed successfully.
16. Wait for cluster initialization (including Oracle software installation and configuration) to complete. You will receive a broadcast message when initialization completes or fails. Cluster initialization may take up to 70 minutes.

4.2 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.

```
[grid@rac1 ~]$ flashgrid-cluster
FlashGrid 16.12.20.50546 #95f2b5603f206af26482ac82386b1268b283fc3c
~~~~~
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
-----
quorum   Good    No        No             Yes           QUORUM
rac1     Good    Yes       Yes            No            RAC1
rac2     Good    Yes       Yes            No            RAC2
-----
-----
GroupName  Status  Mounted  Type    TotalMiB  FreeMiB  OfflineDisks  LostDisks  Resync  ReadLocal
-----
GRID       Good   AllNodes  NORMAL  12588    3376    0              0          No     Enabled
-----
```

4.3 Changing ASM passwords

The automatic configuration process sets a temporary password for ASM users `SYS` and `ASMSNMP`. The temporary password is `"MyPassword2017"`. Use the following commands to set new password(s):

```
# su grid
$ asmcmd passwd sys
$ asmcmd passwd asmsnmp
```

4.4 Configuring NTP service

NTP service is important for keeping clocks of the cluster nodes synchronized. Without active NTP service the clocks are likely to get out of sync. Oracle CTSS service synchronizes system clocks while CRS is running. However, it cannot synchronize clocks before CRS is started or on quorum nodes. Clock synchronization is important for FlashGrid Storage Fabric. Clocks difference of 30 seconds or more will cause warnings. Clocks difference of 60 seconds or more will cause disconnect of remotely attached disks.

On a newly created cluster the NTP service is disabled. In most cases this does not cause immediate problems and short testing can be done without configuring NTP. However, if the cluster is supposed to be up for several days or longer then it is important to configure the NTP service to avoid sudden disruptions.

You can check the current clock difference by running `flashgrid-cluster verify` command.

To configure NTP service, perform the following steps on all nodes including quorum node

1. Stop Oracle cluster services (except quorum node)

```
# crsctl stop cluster
```

2. Move NTP config files to their original locations (they were moved to avoid errors during GI installation)

```
# mv /etc/ntp.bkp /etc/ntp
# mv /etc/ntp.conf.bkp /etc/ntp.conf
# mv /etc/sysconfig/ntpd.bkp /etc/sysconfig/ntpd
# mv /etc/sysconfig/ntpdate.bkp /etc/sysconfig/ntpdate
```

3. Modify `/etc/ntp.conf` with the list of NTP servers that you want to use. Note that public NTP servers cannot be used if public IPs are disabled on the instances and NAT is not configured in the network.
4. Enable and start NTP services

```
# systemctl enable ntpdate; systemctl start ntpdate
# systemctl enable ntpd; systemctl start ntpd
```

5. Check status of the NTP services

```
# systemctl status ntpdate
# systemctl status ntpd
```

6. Start Oracle cluster services (except quorum node)

```
# crsctl start cluster
```

After configuring and starting NTP services on all nodes, run `flashgrid-cluster verify` to check that the clocks are synchronized.

4.5 Deleting a cluster

To delete a cluster, open AWS CloudFormation Manager console and delete the corresponding stack.

5 Creating ASM Disk Groups

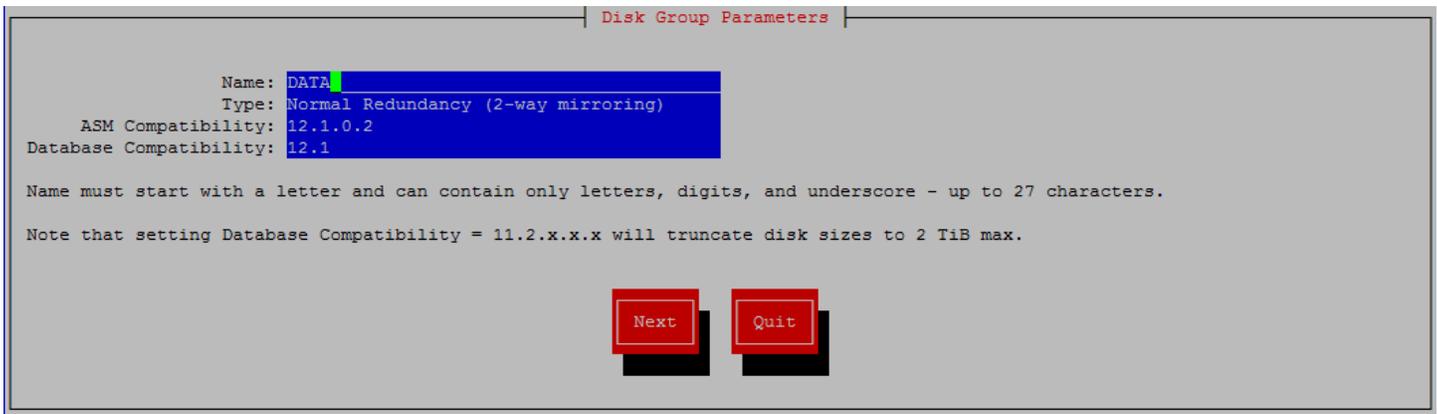
Use of FlashGrid Disk Group Configuration Assistant is recommended for creating ASM disk groups on FlashGrid disks. This wizard-like tool helps avoid mistakes by automating the following tasks:

- Enforces correct disk group type and compatibility attributes
- Assigns each disk to a correct failgroup based on its physical location
- Checks disk number and size restrictions
- Helps determine correct number of quorum disks
- Configures preferred read settings

To create an ASM disk group on FlashGrid disks

1. Make sure that ASM is running on all nodes that are configured to run ASM.
2. On any ASM node as the Grid Infrastructure owner user run `flashgrid-create-dg`
3. Complete all steps of the FlashGrid Disk Group Configuration Assistant following the instructions on each screen.
 - If using EBS volumes, on the *Repair time* screen set *2400h* for both Disk and Failgroup repair time. EBS volumes have built-in protection and physical failures are very unlikely. Setting long repair time helps avoid disks being dropped by ASM in case a VM is down for long time.
4. Run `'flashgrid-dg show -G <dgname>'` command to check status of the newly created disk group.

FlashGrid Disk Group Configuration Assistant: Disk Group Attributes



FlashGrid Disk Group Configuration Assistant: Selecting Disks

Select Disks

Select FlashGrid disks that you want to include in the disk group DATA

```
[*] xvdk 2 GiB N/A
[*] xvdl 2 GiB N/A
[ ] xvdm 10 GiB N/A
<-> RAC2
[ ] grid-grid 6 GiB N/A !GRID member header!
[ ] grid-gridtmp1 3 GiB N/A
[ ] grid-gridtmp2 3 GiB N/A
[*] xvdk 2 GiB N/A
[*] xvdl 2 GiB N/A
[ ] xvdm 10 GiB N/A
```

Total capacity available to ASM: 8 GiB | Capacity after mirroring: 4 GiB

No errors

Next

Back

Quit

FlashGrid Disk Group Configuration Assistant: Selecting Quorum Disks

Quorum disks

For the selected disk group configuration the recommended number of quorum disks is 1.

Quorum disks store copies of ASM metadata and must be located independently of the storage nodes. Possible locations are separate Quorum Servers, separate database nodes configured as Quorum Servers, or external iSCSI, FC, or NFS storage (quorum disks for disk groups with ACFS cannot reside on NFS). No user data is stored on the quorum disks. Recommended quorum disk size is 128 MiB.

```
<-> QUORUM
[*] quorum.quorum-data 300 MiB
[ ] quorum.quorum-fra 300 MiB
```

No errors

Next

Back

Reload

Quit

FlashGrid Disk Group Configuration Assistant: Setting Disk and FailGroup Repair Time

Repair time

If needed, customize the disk and failgroup repair times. The disk repair time defines how long ASM waits for an offline disk to come back online before it drops the disk and starts rebalancing. The failgroup repair time is similar, but used only when an entire node (failgroup) goes offline.

These parameters are mainly for unexpected disk or node failure scenarios. For planned maintenance the parameters can be overridden in the manual offline command for a disk or a failgroup.

Disk repair time: 2400h
Failgroup repair time: 2400h

FlashGrid Disk Group Configuration Assistant: Allocation Unit Size

Allocation Unit Size

If needed, customize the Allocation Unit Size (AU_SIZE) for the disk group.

Allocation Unit Size: 4MB

FlashGrid Disk Group Configuration Assistant: Finish

Creating Disk Group

Do you want to create the new disk group now or just save the SQL code in a file?

Create the new disk group and save the SQL code in a file
 Only save the SQL code in a file

Example of a cluster status summary after configuring one disk group

```
[grid@rac1 ~]$ flashgrid-dg show -G DATA
FlashGrid 16.12.20.50546 #95f2b5603f206af26482ac82386b1268b283fc3c
~~~~~
Querying nodes: quorum, rac1, rac2 ...
-----
Disk Group:      DATA
Status:          Good
Redundancy:      NORMAL
Total:           8492
Free:            8192
ReadLocal:      Enabled
Resync:          No
Offline disks:   0
Lost disks:     0
-----
Node  Mounted
-----
rac1  Yes
rac2  Yes
-----
FailGroup  ASM_Disk_Name      Drive                               Phys_GiB  ASM_GiB  Status
-----
QUORUM     QUORUM$QUORUM_DATA   /dev/flashgrid/quorum.quorum-data  0         0        ONLINE
RAC1       RAC1$XVDK             /dev/flashgrid/rac1.xvdk           2         2        ONLINE
RAC1       RAC1$XVDL             /dev/flashgrid/rac1.xvdl           2         2        ONLINE
RAC2       RAC2$XVDK             /dev/flashgrid/rac2.xvdk           2         2        ONLINE
RAC2       RAC2$XVDL             /dev/flashgrid/rac2.xvdl           2         2        ONLINE
-----
```

Example of a node status summary after configuring one disk group

```
[grid@rac1 ~]$ flashgrid-node
FlashGrid 16.12.20.50546 #95f2b5603f206af26482ac82386b1268b283fc3c
~~~~~
rac1 node status: Good
-----
Local NICs:
-----
Address      Status  Iface      HW      Speed  MTU
-----
192.168.3.1  Good    fg-storage ethernet N/A    8951
-----
Local Drives:
-----
DriveName      Status  SizeGiB  Slot  WritesUsed  ASMName      ASMSizeGiB  DiskGroup  ASMStatus
-----
rac1.grid-grid      Good    6         N/A  N/A          RAC1$GRID_GRID      6           GRID       ONLINE
rac1.grid-gridtmp1  Good    3         N/A  N/A          RAC1$GRID_GRIDTMP1  N/A         N/A        N/A
rac1.grid-gridtmp2  Good    3         N/A  N/A          RAC1$GRID_GRIDTMP2  N/A         N/A        N/A
rac1.xvdk           Good    2         N/A  N/A          RAC1$XVDK           2           DATA      ONLINE
rac1.xvdl           Good    2         N/A  N/A          RAC1$XVDL           2           DATA      ONLINE
rac1.xvdm           Good    10        N/A  N/A          RAC1$XVDM           N/A         N/A        N/A
-----
Remote Drives:
-----
DriveName      Status  SizeGiB  CfgPaths  ActPaths  ASMName      ASMSizeGiB  DiskGroup  ASMStatus
-----
quorum.quorum-data  Good    0         1         1         QUORUM$QUORUM_DATA  0           DATA      ONLINE
quorum.quorum-fra   Good    0         1         1         QUORUM$QUORUM_FRA   N/A         N/A        N/A
quorum.quorum-grid  Good    0         1         1         QUORUM$QUORUM_GRID  0           GRID       ONLINE
rac2.grid-grid      Good    6         1         1         RAC2$GRID_GRID      6           GRID       ONLINE
rac2.grid-gridtmp1  Good    3         1         1         RAC2$GRID_GRIDTMP1  N/A         N/A        N/A
rac2.grid-gridtmp2  Good    3         1         1         RAC2$GRID_GRIDTMP2  N/A         N/A        N/A
rac2.xvdk           Good    2         1         1         RAC2$XVDK           2           DATA      ONLINE
rac2.xvdl           Good    2         1         1         RAC2$XVDL           2           DATA      ONLINE
rac2.xvdm           Good    10        1         1         RAC2$XVDM           N/A         N/A        N/A
-----
```

6 Installing Database Software

Current version of FlashGrid Cloud Provisioning provides automated installation and patching of Oracle Database 12.1.0.2 Enterprise Edition in RAC mode. For installing and patching other database software versions use standard procedures described in corresponding Oracle documentation.

7 Creating a Database

Use standard procedures for creating a database. Select ASM storage and the disk group created for the database.

Note: Do not use the GRID disk group as storage for a database. The GRID disk group is dedicated for OCR files, voting files, and MGMTDB and must not be used for any other purposes.

For running DBCA in GUI mode use X-forwarding or VNC.

To run DBCA via VNC, perform the following steps on database node 1

1. Make sure VNC port 5901 (TCP) is either open via the security group settings or forwarded via SSH tunnel.
2. As root, run `vncserver` command and set VNC password
3. Connect your VNC client (e.g. RealVNC) to the node. Use one of the following connection strings:
 - `PublicIP:5901` - if connecting via public IP
 - `PrivateIP:5901` - if connecting via VPN to the VPC private IP
 - `localhost:5901` - if using port forwarding via SSH
4. Run the following commands in the x-terminal inside VNC client

```
# xhost localhost
# su oracle
$ dbca
```

8 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

5. Create or load a database on the corresponding ASM disk group
6. 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.
7. 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.
8. Connect to the database with sqlplus and run the customized SQL code.
9. 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 four io1 volumes (20,000 IOPS each) per node

```
SET SERVEROUTPUT ON;
DECLARE
  lat INTEGER;
  iops INTEGER;
  mbps INTEGER;
BEGIN DBMS_RESOURCE_MANAGER.CALIBRATE_IO (8, 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 = 154864
Latency = 0
Max_MB/s = 2219

PL/SQL procedure successfully completed.
```

9 Maintenance Tasks

9.1 Rebooting a node

To reboot a database or quorum node in a running cluster

1. Make sure there are no other nodes that are in offline or re-syncing state. All disk groups must have zero offline disks and *Resync = No*.

```
# flashgrid-cluster
```

2. Reboot the node using flashgrid-node command. It will gracefully put the corresponding failure group offline.

```
# flashgrid-node reboot
```

3. Wait until the node boots up and re-synchronization completes before rebooting or powering off any other node.

9.2 Powering off a node

To power off a database or quorum node in a running cluster

1. Make sure there are no other nodes that are in offline or re-syncing state. All disk groups must have zero offline disks and *Resync = No*.

```
# flashgrid-cluster
```

2. Power off the node using flashgrid-node command. It will gracefully put the corresponding failure group offline.

```
# flashgrid-node poweroff
```

3. After rebooting the node, wait until the node boots up and re-synchronization completes before rebooting or powering off any other node.

9.3 Shutting down an entire cluster

To shut an entire cluster down

1. Stop Oracle cluster services on all nodes.

```
# crsctl stop cluster -all
```

2. Power all nodes off.

9.4 Adding EBS volumes

When adding new volumes make sure that each disk group has disks of the same size and that the number of disks per node is the same.

To add new EBS volumes in a running cluster

1. Create and attach new volumes to the database node instances.
2. Use 'flashgrid-cluster drives' command to determine FlashGrid names of the drives, e.g. *rac2.xvdm*
3. To create a new disk group, run 'flashgrid-create-dg'
4. To add the disks to an existing disk group:

```
$ flashgrid-dg add-disks -G MYDG -d /dev/flashgrid/rac1.xvdm
/dev/flashgrid/xvdm
```

9.5 Removing EBS volumes

To remove EBS volumes from a running cluster

1. Use 'flashgrid-cluster drives' command to determine FlashGrid names of the drives, e.g. *rac2.xvdm*
2. If the drives are members of an ASM disk group then drop the drives from the disk group. Example:

```
SQL> alter diskgroup MYDG
drop disk RAC1$XVDM
drop disk RAC2$XVDM
rebalance wait;
```

3. Prepare the drives for removal. Example:

```
[root@rac1 ~] # flashgrid-node stop-target /dev/flashgrid/rac1.xvdm
[root@rac2 ~] # flashgrid-node stop-target /dev/flashgrid/rac2.xvdm
```

4. Detach the EBS volumes from the instances.

9.6 Updating FlashGrid software

The following procedure applies to minor updates. Minor updates are those that have the same first two numbers in the version number, for example, from 16.12.22 to 16.12.40. However, update from 16.12 to 17.03 is considered major and may require a different procedure. Contact FlashGrid support for assistance if you need to do a major version update.

To update FlashGrid software on a running cluster repeat the following steps on each node, one node at a time

1. If the node has ASM running then stop Oracle CRS:

```
# crsctl stop crs
```

2. Stop the FlashGrid Storage Fabric services:

```
# flashgrid-node stop
```

3. Stop the FlashGrid Cloud Area Network service:

```
# systemctl stop flashgrid-clan
```

4. Update the *flashgrid* and *flashgrid-clan* RPMs using yum or rpm tool

5. Start the FlashGrid Cloud Area Network service:

```
# systemctl start flashgrid-clan
```

6. Start the FlashGrid Storage Fabric service

```
# flashgrid-node start
```

7. If the node has ASM installed on it then start Oracle CRS:

```
# crsctl start crs -wait
```

4. Wait until all disks are back online and resyncing operations complete on all disk groups before updating the next node. All disk groups must have zero offline disks and *Resync = No*.

```
# flashgrid-cluster
```

9.7 Updating Linux kernel

To update Linux kernel on a running cluster repeat the following steps on each node, one node at a time

1. If using instances with Elastic Network Adapter (ENA), make sure you have ENA driver for the new kernel available
2. Install the new kernel
3. If using instances with ENA, install the ENA driver for the new kernel
4. Follow the steps for rebooting a node

10 FlashGrid Tools and Commands

10.1 flashgrid-ca

FlashGrid Configuration Assistant helps configure a new FlashGrid cluster in a few easy steps.

Usage: flashgrid-ca [-h] [-f]

Options:

-h Show this help message and exit
-f Ignore terminal size check

10.2 flashgrid-fix-grid-dg-ca

Assistant tool for replacing temporary disks with permanent disks after GI installation.

Usage: flashgrid-fix-grid-dg-ca [-h] [-f]

Options:

-h Show this help message and exit
-f Ignore terminal size check

10.3 flashgrid-create-dg

FlashGrid Disk Group Configuration Assistant helps configure ASM disk groups with FlashGrid disks in a few easy steps.

Usage: flashgrid-create-dg [-h] [-f]

Options:

-h Show this help message and exit
-f Ignore terminal size check

10.4 flashgrid-cluster

CLI tool for cluster-level monitoring and management of all cluster nodes and disk groups.

Usage: flashgrid-cluster [command]

Commands:

show	Show cluster status (default if no command provided)
drives	List all drives in the cluster
net	List all NICs in the cluster
verify	Verify FlashGrid cluster configuration
deploy-config	Deploy configuraiton from /etc/flashgrid.cfg on the local node to all nodes
fix-readlocal	Fix Read-Local settings on all ASM instances
fix-disk-names	Generate ASM SQL for fixing ASM disk names if any are invalid

strict-read-local-enable

Enable strict Read-Local mode. Database auto-start is delayed until all local disks are online.

strict-read-local-disable

Disable strict Read-Local mode and restore original database startup dependencies.

strict-read-local-show

For all databases show whether strict Read-Local mode is enabled or disabled.

help

Show this help message and exit

10.5 flashgrid-node

CLI tool for monitoring and management of the local node.

Usage: flashgrid-node [-f] [command]

Commands:

show	Show node details (default if no command provided)
test-alerts	Send test email alert.
collect-diags	Create diagnostics archive file for the local node.
stop	Gracefully offline the local failgroup and stop flashgrid services. '-f' option forces stop even if graceful offline fails.
start	Start flashgrid service and online local disks in ASM
reboot	Gracefully offline the local failgroup and reboot the node. '-f' option forces reboot even if graceful offline fails.
poweroff	Gracefully offline the local failgroup and power the node off. '-f' option forces reboot even if graceful offline fails.
offline	Gracefully offline the local failgroup.
online	Bring the local failgroup online for each mounted disk group.
stop-waiting	Stop waiting for all storage nodes and allow Oracle services to start.
stop-target	Gracefully offline the corresponding disk in ASM and stop the target. '-f' option forces target target stop even if graceful offline fails.
reload-config	Reload flashgrid config without restart. Note that network configuration is not reloaded.
restart-services	Gracefully offline the local failgroup, restart flashgrid service and online local disks in ASM. '-f' option forces restart even if graceful offline fails.

restart-targets Gracefully offline the local failgroup, restart flashgrid targets service, and online local disks in ASM. '-f' option forces target service restart even if graceful offline fails.

restart-initiators If there are no mounted ASM disk groups using FlashGrid disks then restart flashgrid initiators service. '-f' option forces restart even if there are mounted groups.

restart-asm-connector
Restart the FlashGrid ASM connector service.

restart-cluster-connector
Restart FlashGrid cluster connector service.

verify-config Verify validity of the local config file /etc/flashgrid.cfg

shell Start shell (for internal and tech support use).

asm-state Show ASM state in JSON format (for internal use).

help Show this help message and exit

10.6 flashgrid-dg

CLI tool for monitoring and management of individual disk groups.

Usage: flashgrid-dg [command]

Commands:

list List all FlashGrid disk groups (default if no command provided).

show Show details of the disk group.

create Create a new ASM disk group.

add-disks Add disks to the disk group.

replace-disk Replace a disk with a new one.

help Show this help message and exit

11 Troubleshooting

The following troubleshooting steps are recommended in case of any issues with FlashGrid cluster configuration or operation:

1. Check status of all FlashGrid nodes, network, and disk groups by running `'flashgrid-cluster'` on any node
2. If network verification fails then run `'flashgrid-cluster verify'` to get more detailed information
3. On any node that has a *Warning*, *Critical*, or *Inaccessible* status:
 - a. Check whether the FlashGrid service is active:
`# systemctl status flashgrid`
 - b. Check status of NICs, local disks, and remote disks:
`# flashgrid-node`
 - c. Check that the configuration has no errors:
`# flashgrid-node verify-config`
4. Check detailed status information of various FlashGrid components on the affected nodes:
 - a. Run `'flashgrid-node shell'`
 - b. Navigate through the tree of FlashGrid objects using `'ls'` and `'cd'` commands
 - c. For example, display details of all physical drives: `'ls /hw'`
5. Check FlashGrid logs in the following log files on the affected nodes:
 - a. `/opt/flashgrid/log/fgridd-all.log`
 - b. `/opt/flashgrid/log/fgridd-error.log`
 - c. `/opt/flashgrid/log/iamback-all.log`
 - d. `/opt/flashgrid/log/iamback-error.log`

12 Contacting FlashGrid Technical Support

If you need help with troubleshooting an issue on an existing FlashGrid cluster please collect diagnostic data by running `'flashgrid-node collect-diags'` command on each node and email it to FlashGrid technical support along with a summary of the issue.

Users with technical support subscription: support@flashgrid.io

New users: info@flashgrid.io

Phone: +1-650-641-2421

Copyright © 2017 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.