

# Deploying FlashGrid SkyCluster with Oracle RAC on AWS

*Deployment Process Demonstration*

*rev. 2020-03-16*



## About SkyCluster for AWS

Ensuring high availability of backend relational databases is a critical part of the cloud strategy - whether it is a lift-and-shift migration or a green-field deployment of mission critical applications. FlashGrid SkyCluster is an engineered cloud system designed for database high availability. SkyCluster is delivered as a fully integrated Infrastructure-as-Code template that can be customized and deployed to AWS account with a few mouse clicks. Key components of FlashGrid SkyCluster for AWS include:

- AWS EC2 instances
- AWS EBS and/or local SSD storage
- FlashGrid Storage Fabric software
- FlashGrid Cloud Area Network software
- Oracle Grid Infrastructure software
- Oracle RAC database engine

By leveraging the proven Oracle RAC database engine FlashGrid SkyCluster enables the following use-cases:

- Lift-and-shift migration of existing Oracle RAC databases to AWS.
- Migration of existing Oracle databases from high-end on-premises servers to AWS without reducing availability SLAs.
- Design of new mission critical applications for the cloud based on the industry proven and widely supported database engine.

## About This Demo

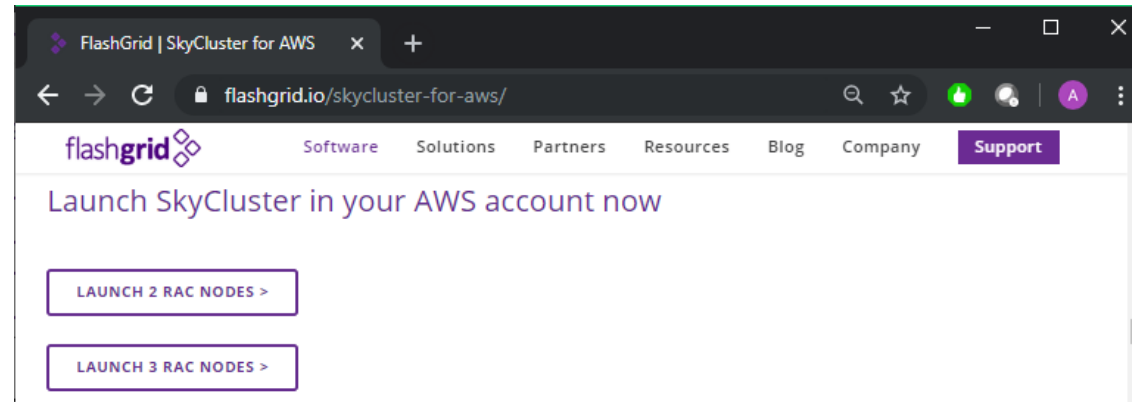
This brief shows the main steps of deploying SkyCluster with Oracle RAC on AWS. The target audience is AWS cloud architects and engineers and database architects and administrators.

More detailed information is available in the following documents:

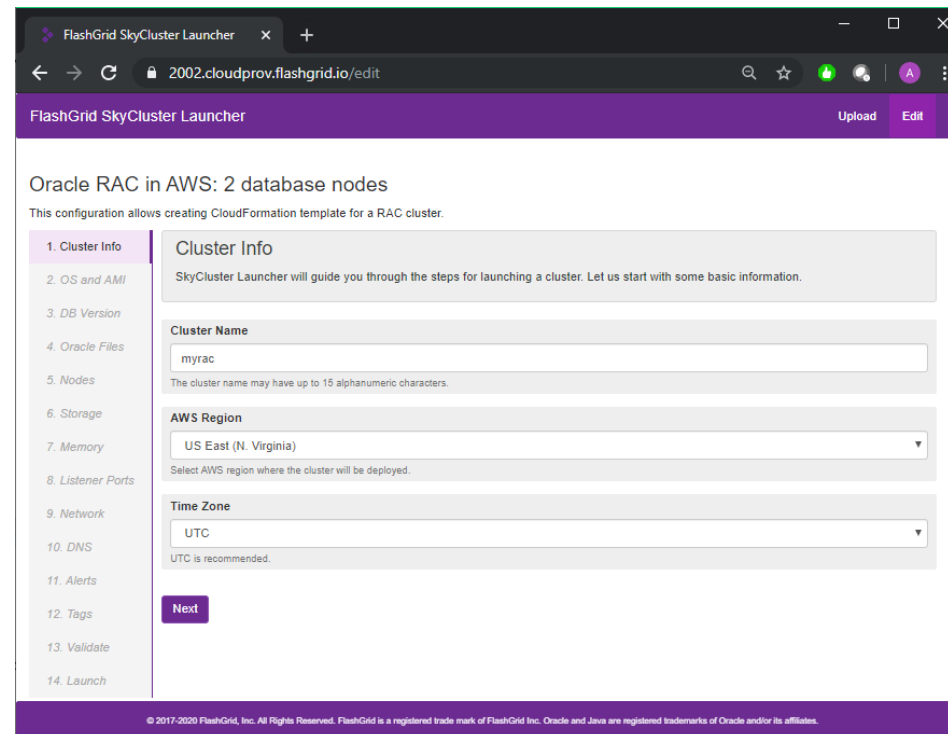
- White paper: [Oracle RAC in AWS Enabled by FlashGrid SkyCluster](#)
- [FlashGrid SkyCluster Deployment Guide for Oracle RAC on AWS](#)

## Step 0: Select one of standard configuration templates

Since majority of deployments have 2 or 3 RAC nodes, these two standard configuration are readily available at <https://www.flashgrid.io/skycluster-for-aws/>



Click on a button that corresponds to either 2, or 3 RAC nodes. It will open SkyCluster Launcher tool. (For configurations other than 2 or 3 RAC nodes contact FlashGrid support.)



## Step 1: Enter basic information about the cluster

At this step you need to enter information such as AWS Region and preferred time zone.

- 1. Cluster Info**
- 2. OS and AMI
- 3. DB Version
- 4. Oracle Files
- 5. Nodes
- 6. Storage
- 7. Memory
- 8. Listener Ports
- 9. Network
- 10. DNS
- 11. Alerts
- 12. Tags
- 13. Validate
- 14. Launch

### Cluster Info

SkyCluster Launcher will guide you through the steps for launching a cluster. Let us start with some basic information.

#### Cluster Name

The cluster name may have up to 15 alphanumeric characters.

#### AWS Region

Select AWS region where the cluster will be deployed.

#### Time Zone

UTC is recommended.

[Next](#)

## Step 2: Select OS

Select which operating system you prefer – Oracle Linux 7 or RHEL 7

1. Cluster Info

2. OS and AMI

3. DB Version

4. Oracle Files

5. Nodes

6. Storage

7. Memory

8. Listener Ports

9. Network

10. DNS

11. Alerts

12. Tags

13. Validate

14. Launch

### OS and AMI

Select an AMI option corresponding to your preferred OS.

Your AWS account must have an active subscription to the selected FlashGrid AMI. Otherwise deployment will fail when creating VM instances. To subscribe or to see FlashGrid software fees click on one of the links below:

- [Oracle Linux 7 based AMI on AWS Marketplace](#)
- [RHEL 7 based AMI on AWS Marketplace](#)

#### Marketplace AMI

Oracle Linux 7 AMI for FlashGrid SkyCluster with Oracle RAC

None

Oracle Linux 7 AMI for FlashGrid SkyCluster with Oracle RAC

RHEL 7 AMI for FlashGrid SkyCluster with Oracle RAC

Prev

Next

## Step 3: Select database version

Select which version of the database you are planning to use, along with Patch Set Update / Release Update version for it.

- 1. Cluster Info
- 2. OS and AMI
- 3. DB Version**
- 4. Oracle Files
- 5. Nodes
- 6. Storage
- 7. Memory
- 8. Listener Ports
- 9. Network
- 10. DNS
- 11. Alerts
- 12. Tags
- 13. Validate
- 14. Launch

### Database Version

Select options for database software installation.

**Install Database Software (RAC mode)**  
Select if you want database software installed. Only RAC mode is currently supported for installation. De-select if RAC One Node or Single-Instance is required.

#### Database Version

19c EE

Select which Database software version will be installed.

#### Database PSU/RU

2020-01-14

Database PSU/RU version to apply.

[Prev](#) [Next](#)

## Step 4: Provide location of Oracle installation files

You need to upload the listed Oracle installation files to an S3 bucket and provide name of the bucket. Cluster initialization script will download and install the files.

- 1. Cluster Info
- 2. OS and AMI
- 3. DB Version
- 4. Oracle Files**
- 5. Nodes
- 6. Storage
- 7. Memory
- 8. Listener Ports
- 9. Network
- 10. DNS
- 11. Alerts
- 12. Tags
- 13. Validate
- 14. Launch

### Oracle Files

We ask you to use your own copy of Oracle installation files. Place the files listed below in an S3 bucket and provide name of the bucket and the folder.

We recommend downloading the files using the hyper-links. If you prefer to download from eDelivery then need to rename the files to the names listed below.

- [LINUX\\_X64\\_193000\\_db\\_home.zip](#) - Oracle Database 19c (19.3) for Linux x86-64
- [LINUX\\_X64\\_193000\\_grid\\_home.zip](#) - Oracle Database 19c Grid Infrastructure (19.3) for Linux x86-64
- [oracle-instantclient19.3-basic-19.3.0.0.0-1.x86\\_64.rpm](#) - Oracle Instant Client Basic 19.3.0.0.0 for Linux x86-64
- [p30484981\\_190000\\_Linux-x86-64.zip](#) - Patch 30484981: OJVM RELEASE UPDATE 19.6.0.0.0. Requires Oracle support subscription.
- [p30501910\\_190000\\_Linux-x86-64.zip](#) - Patch 30501910: GI RELEASE UPDATE 19.6.0.0.0. Requires Oracle support subscription.
- [p68808880\\_190000\\_Linux-x86-64.zip](#) - OPatch 12.2.0.1.19 for DB 19.x releases, Platform: Linux x86-64. Requires Oracle support subscription.

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

### S3 Bucket/Folder with Oracle Installation Files

Provide S3 bucket and folder where Oracle GI and DB installation files are located. Examples: mys3bucket or mys3bucket/myfolder

### S3 Bucket Region

Select S3 region where Oracle files are stored, if it is different from the target deployment region.

### IAM Role for Accessing the S3 Bucket

Leave blank if public access to the files in the S3 bucket is enabled. If cannot enable public access then provide name of an IAM role that allows access to the files in the bucket. Note that this tool cannot verify access to the required files if public access is not enabled. If any file is missing then the cluster initialization process will fail.

Prev **Next**

## Step 5: Configure cluster nodes

Provide hostnames, Availability Zone placement (for regions where AZs are supported), and sizes of the cluster nodes.

1. Cluster Info

2. OS and AMI

3. DB Version

4. Oracle Files

5. Nodes

6. Storage

7. Memory

8. Listener Ports

9. Network

10. DNS

11. Alerts

12. Tags

13. Validate

14. Launch

### Nodes

Configure hostname, availability zone, and size for each node of the cluster.

- Host names must have equal length.
- If placing the nodes across AZs, keeping the default assignment of AZs is recommended unless instructed otherwise by FlashGrid support.
- In regions with two AZs place all nodes in the same AZ.
- Do not change node Roles unless instructed by FlashGrid support.

#### Cluster Nodes

Hostname*	Role*	AZ*	VM type*
<input type="text" value="rac1"/>	<input type="text" value="database"/>	<input type="text" value="a"/>	<input type="text" value="r5.12xlarge: 24 cores, 384 GiB, EBS: max 40000 IOPS, max 875 MB/s"/>
Hostname*	Role*	AZ*	VM type*
<input type="text" value="rac2"/>	<input type="text" value="database"/>	<input type="text" value="b"/>	<input type="text" value="r5.12xlarge: 24 cores, 384 GiB, EBS: max 40000 IOPS, max 875 MB/s"/>
Hostname*	Role*	AZ*	VM type*
<input type="text" value="racq"/>	<input type="text" value="quorum"/>	<input type="text" value="c"/>	<input type="text" value="c5.large, for quorum node only"/>

EC2 Instance Termination Protection (recommended for production deployments)

Enable to protect the EC2 instances from accidental termination. If later you need to delete the cluster then the protection must be disabled for each instance through the EC2 management console.

[Prev](#)

[Next](#)



## Step 6: Configure storage

Specify ASM disk groups that will be created. The corresponding disks will be automatically attached to the nodes.

1. Cluster Info

2. OS and AMI

3. DB Version

4. Oracle Files

5. Nodes

6. Storage

7. Memory

8. Listener Ports

9. Network

10. DNS

11. Alerts

12. Tags

13. Validate

14. Launch

### Storage

Configure ASM disk groups that will be automatically created with EBS disks.

- The number of disks is specified per node.
- Usable\_Capacity = Number\_of\_Disks\_per\_Node x Disk\_Size (because of mirroring between the nodes)
- Each GP2 disk provides 3 IOPS/GiB, e.g. 3000 IOPS for 1000 GiB, up to 16 000 IOPS per disk.
- GP2 disks larger than 334 GiB provide 250 MB/s of throughput per disk.
- GP2 disks smaller than 334 GiB are not recommended for production use or performance testing.
- GRID disk group is configured automatically for Vote+OCR.
- Disk groups on local SSDs (e.g. with i3 instance type) must be configured after cluster initialization.

#### ASM Disk Groups

##### Storage profile

database

Disk Group Name*	# Disks per Node*	Disk Size, GiB*
DATA	3	1000

Disk Group Name*	# Disks per Node*	Disk Size, GiB*
FRA	2	1000

×

×

+

Encrypted EBS Volumes

Select to enable built-in encryption on EBS volumes. ATTENTION! For encrypted root volume must use a custom encrypted AMI. With standard AMI the root volume will not be encrypted.

Prev

Next

## Step 7: Specify memory allocation

If needed, customize database memory allocation percentages. These percentages are used for automatic configuration of HugePages when database node boots up based on the total memory.

- 1. Cluster Info
- 2. OS and AMI
- 3. DB Version
- 4. Oracle Files
- 5. Nodes
- 6. Storage
- 7. Memory**
- 8. Listener Ports
- 9. Network
- 10. DNS
- 11. Alerts
- 12. Tags
- 13. Validate
- 14. Launch

### Memory

Select whether HugePages will be automatically configured for SGA.

- Enabling HugePages is recommended for reducing CPU utilization.
- The number of HugePages will be updated automatically when VM size changes.

**Automatically configure HugePages**  
Recommended except when Oracle AMM must be used.

#### % of System Memory allocated for Databases (SGA+PGA)

Percentage of system memory that will be allocated for use by all databases. 80% recommended. Ignored if automatic configuration of HugePages is disabled.

#### % of the Database Memory allocated for SGA

This parameter is used for automatically calculating the number of required HugePages. Ignored if automatic configuration of HugePages is disabled.

[Prev](#) [Next](#)

## Step 8: Specify listener ports

If needed, customize SCAN and Local listener port numbers.

1. Cluster Info

2. OS and AMI

3. DB Version

4. Oracle Files

5. Nodes

6. Storage

7. Memory

8. Listener Ports

9. Network

10. DNS

11. Alerts

12. Tags

13. Validate

14. Launch

### Listener Ports

Select listener port numbers. The SCAN listener and Local listener port numbers must be different.

#### SCAN Listener Port

Default: 1521. Must be different from Local Listener port.

#### Local Listener Port

Default: 1522. Must be different from SCAN Listener port.

[Prev](#)

[Next](#)

## Step 9: Provide information about target VPC

Provide ID of the VPC where you want the cluster deployed.

1. Cluster Info

2. OS and AMI

3. DB Version

4. Oracle Files

5. Nodes

6. Storage

7. Memory

8. Listener Ports

9. Network

10. DNS

11. Alerts

12. Tags

13. Validate

14. Launch

### Network

Provide ID for an existing VPC or keep blank to create a new VPC. If using an existing VPC then it must have a subnet in each of the used availability zones and a security group with the following ports open:

- UDP 4801, 4802, 4803 between the cluster node VMs (cluster initialization will fail if any of these ports are not open)
- TCP ports 1521 (or customized SCAN Listener port) and 1522 (or customized Local Listener port) for client and app server access
- TCP port 22 for SSH access

#### VPC ID

Provide VPC ID to use an existing VPC. Leave blank to create a new VPC.

#### Assign Public IPs to Cluster Nodes

ATTENTION! Enabling Public IPs is NOT recommended on production clusters for security reasons. If disabled, then the VPC must have either an S3 endpoint configured and the S3 bucket with Oracle files located in the same region, or NAT for accessing the Internet must be configured for the VPC. Enable if SSH access via Internet without VPN is required.

Prev

Next

## Step 10: Configure DNS

Specify domain name that will be assigned to cluster nodes. You can also replace the default Route53 DNS server with your own list of DNS servers.

1. Cluster Info

2. OS and AMI

3. DB Version

4. Oracle Files

5. Nodes

6. Storage

7. Memory

8. Listener Ports

9. Network

10. DNS

11. Alerts

12. Tags

13. Validate

14. Launch

### DNS

Within the cluster host name resolution is performed by DNSMASQ service configured locally.

For resolving cluster node names on clients or app servers need to add corresponding records to your DNS servers. If using Route53 then this can be done automatically.

For resolving host names that are outside of the cluster (e.g. S3 endpoint) on the cluster nodes, you can use Route53 (default) or custom DNS servers.

#### Domain Name

The domain name will be configured on all nodes of the cluster. The domain must be in a zone hosted on your DNS servers.

#### Route53 DNS Hosted Zone Used by Database Clients / App Servers (optional)

If using Route53 for DNS on the database clients, enter the hosted zone ID to automatically add DNS records for SCAN address and DB nodes. Leave blank to add the DNS records manually.

#### Custom DNS Server(s)



Keep the list empty if using Route53 or creating a new VPC. Add the DNS server(s) if using custom DNS servers in existing VPC.

[Prev](#)

[Next](#)

## Step 11: Configure email alerts

Specify list of emails where alerts will be sent for errors that may happen during operation of the cluster.

1. Cluster Info

2. OS and AMI

3. DB Version

4. Oracle Files

5. Nodes

6. Storage

7. Memory

8. Listener Ports

9. Network

10. DNS

11. Alerts

12. Tags

13. Validate

14. Launch

### Alerts

Specify email addresses where alerts will be sent in case of a failure during cluster operation. Optionally, enable sending alerts directly to FlashGrid support.

#### Email Addresses

root@localhost



Send SkyCluster operation error alerts to FlashGrid support (preview)

Must have outbound HTTPS traffic to <https://alerts.support.flashgrid.io> allowed. The alert information consists of cluster name, host name, source IP, license status, error details.

[Prev](#)

[Next](#)

## Step 12: Specify tags

Optionally, specify the list of tags that will be assigned to the EC2 instances and disks.

1. Cluster Info

2. OS and AMI

3. DB Version

4. Oracle Files

5. Nodes

6. Storage

7. Memory

8. Listener Ports

9. Network

10. DNS

11. Alerts

12. Tags

13. Validate

14. Launch

### Tags for Instances and Disks

Optionally, specify tags that will be attached to the EC2 instances and to the EBS volumes. Do not add *name* or *cluster* tags, these tags are configured automatically.

#### Tags / Labels (optional)

Key / Name

Org

Value

MyOrg

✕

+

Do not add Name or Cluster tags.

Prev

Next

## Step 13: Validate configuration

Click *Validate* to confirm that the provided configuration is consistent.

The screenshot shows a configuration wizard interface. On the left is a vertical sidebar with 14 steps: 1. Cluster Info, 2. OS and AMI, 3. DB Version, 4. Oracle Files, 5. Nodes, 6. Storage, 7. Memory, 8. Listener Ports, 9. Network, 10. DNS, 11. Alerts, 12. Tags, 13. Validate (highlighted in purple), and 14. Launch (highlighted in light gray). The main content area is a light gray box with the heading 'Validate' and the text 'We are ready to do final validation of the entire configuration.' Below this text are two buttons: 'Prev' and 'Validate Configuration' (a dark purple button).



## Step 14: Generate CloudFormation template

When you click *Launch SkyCluster*, a CloudFormation template will be generated and CloudFormation Manager portal will open.

- 1. Cluster Info
- 2. OS and AMI
- 3. DB Version
- 4. Oracle Files
- 5. Nodes
- 6. Storage
- 7. Memory
- 8. Listener Ports
- 9. Network
- 10. DNS
- 11. Alerts
- 12. Tags
- 13. Validate
- 14. Launch**

### Launch

Click *Launch SkyCluster* to generate a CloudFormation template and open it in CloudFormation Manager for deploying.

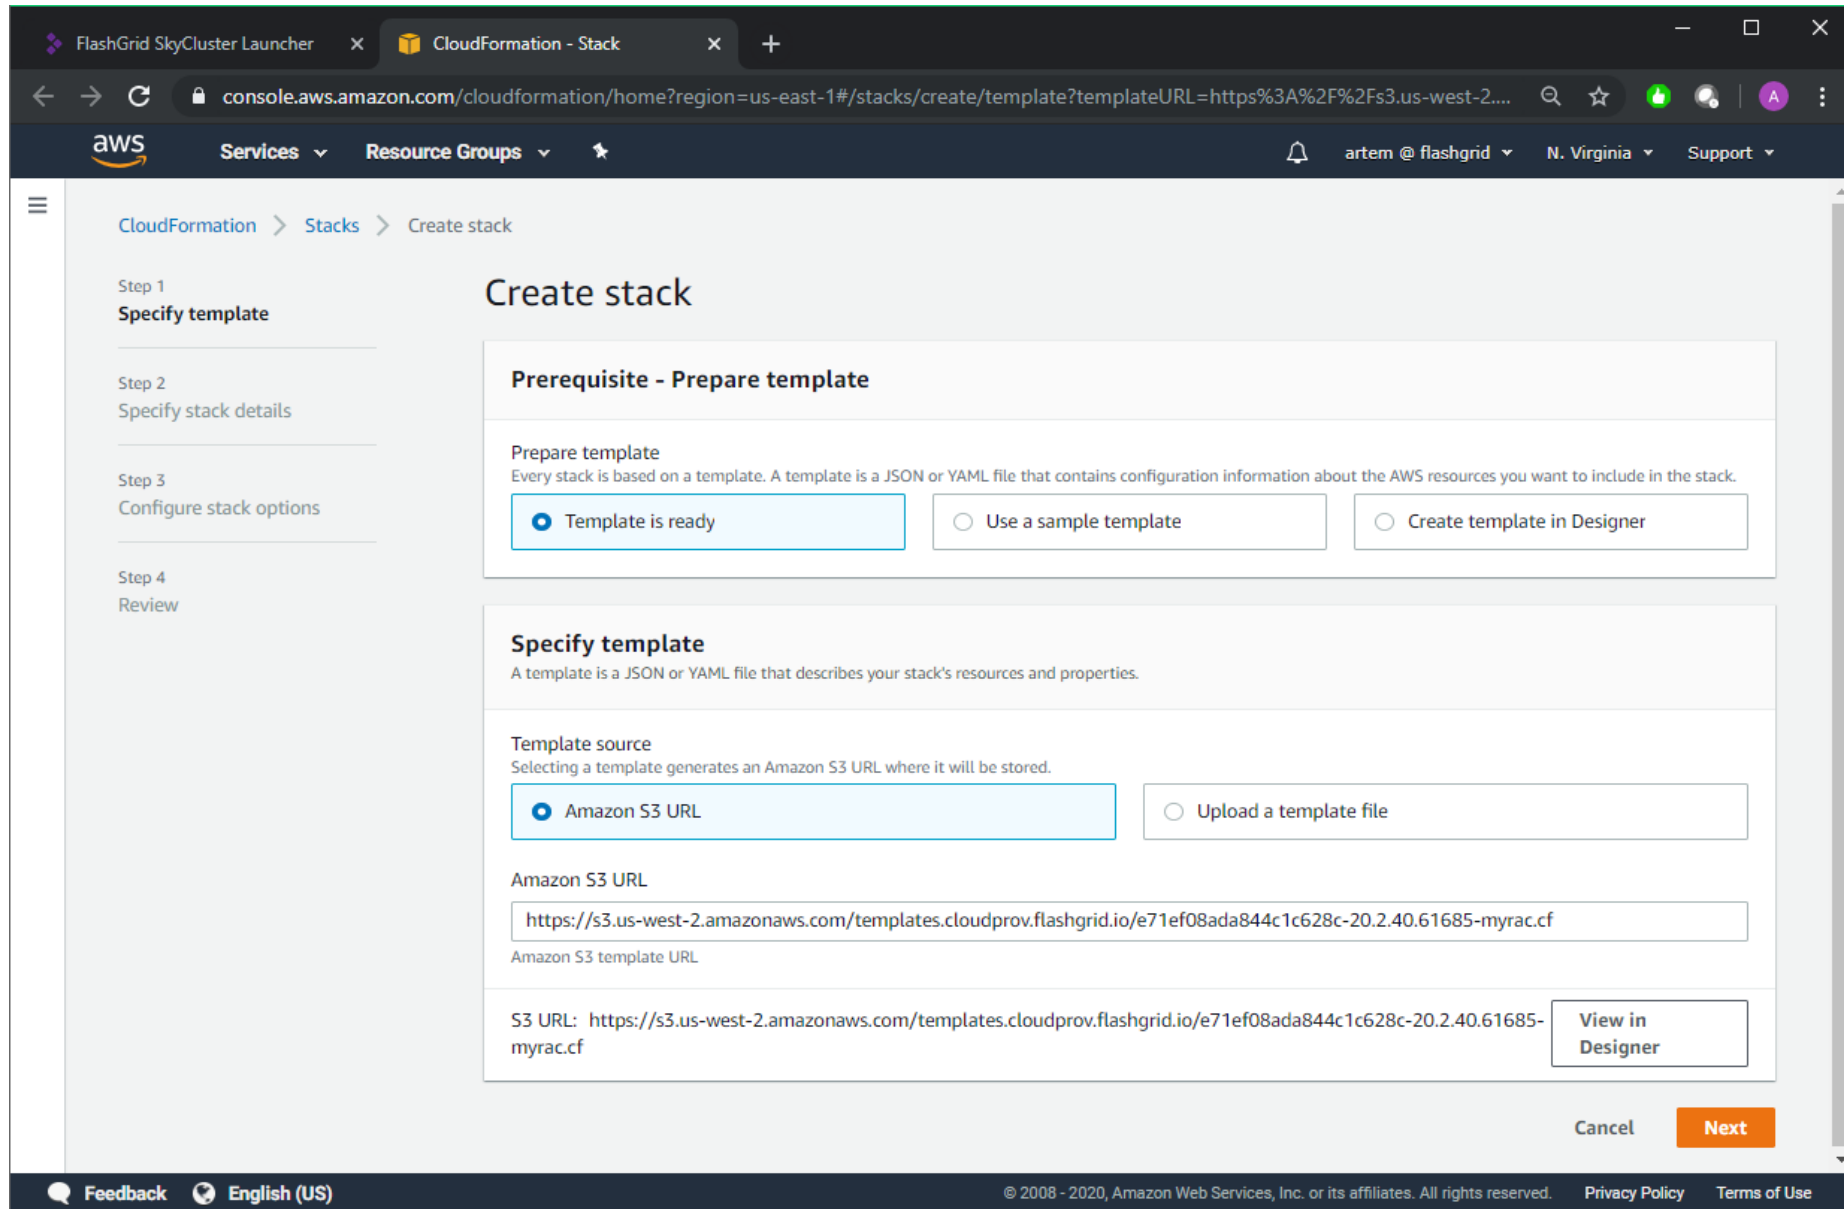
- You can keep defaults on the Options page in CloudFormation Manager.
- No need to add tags again in CloudFormation Manager if you have already added tags in this tool.
- You can see estimated costs of the EC2 and EBS infrastructure on the Review page in CloudFormation Manager. Note that these costs do not include FlashGrid software fees.
- After the CloudFormation template is successfully deployed, cluster initialization process starts and takes about 90 minutes.
- Connect to the first database node via SSH as user fg@.
- See FlashGrid knowledge base for troubleshooting [CloudFormation errors](#) or [errors during cluster initialization](#).

To download the CloudFormation template without opening CloudFormation Manager click [here](#).

[Launch SkyCluster](#) [Download Configuration](#) [Edit Configuration](#)

## Step 15: Create CloudFormation stack – specify template

Just click Next to confirm use of the generated template.



The screenshot shows the AWS CloudFormation console interface for creating a stack. The browser tabs include 'FlashGrid SkyCluster Launcher' and 'CloudFormation - Stack'. The URL is 'console.aws.amazon.com/cloudformation/home?region=us-east-1#/stacks/create/template?templateURL=https%3A%2F%2Fs3.us-west-2...'. The navigation bar shows 'aws', 'Services', 'Resource Groups', and user information 'artem @ flashgrid' in 'N. Virginia'.

The main content area is titled 'Create stack' and shows a progress indicator for four steps: 'Step 1: Specify template' (active), 'Step 2: Specify stack details', 'Step 3: Configure stack options', and 'Step 4: Review'.

**Prerequisite - Prepare template**  
Prepare template  
Every stack is based on a template. A template is a JSON or YAML file that contains configuration information about the AWS resources you want to include in the stack.

Template is ready  Use a sample template  Create template in Designer

**Specify template**  
A template is a JSON or YAML file that describes your stack's resources and properties.

**Template source**  
Selecting a template generates an Amazon S3 URL where it will be stored.

Amazon S3 URL  Upload a template file

Amazon S3 URL  
  
Amazon S3 template URL

S3 URL: `https://s3.us-west-2.amazonaws.com/templates.cloudprov.flashgrid.io/e71ef08ada844c1c628c-20.2.40.61685-myrac.cf`

Footer: Feedback English (US) © 2008 - 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved. Privacy Policy Terms of Use

## Step 16: Create CloudFormation stack – specify stack details

Specify additional details, such as your SSH key name, subnet IDs, and security group ID.

The screenshot shows the AWS CloudFormation console interface. The browser address bar displays the URL: `console.aws.amazon.com/cloudformation/home?region=us-east-1#/stacks/create/parameters`. The page title is "Specify stack details".

On the left side, there is a navigation pane with the following steps:

- Step 1: Specify template
- Step 2: **Specify stack details** (current step)
- Step 3: Configure stack options
- Step 4: Review

The main content area contains the following sections:

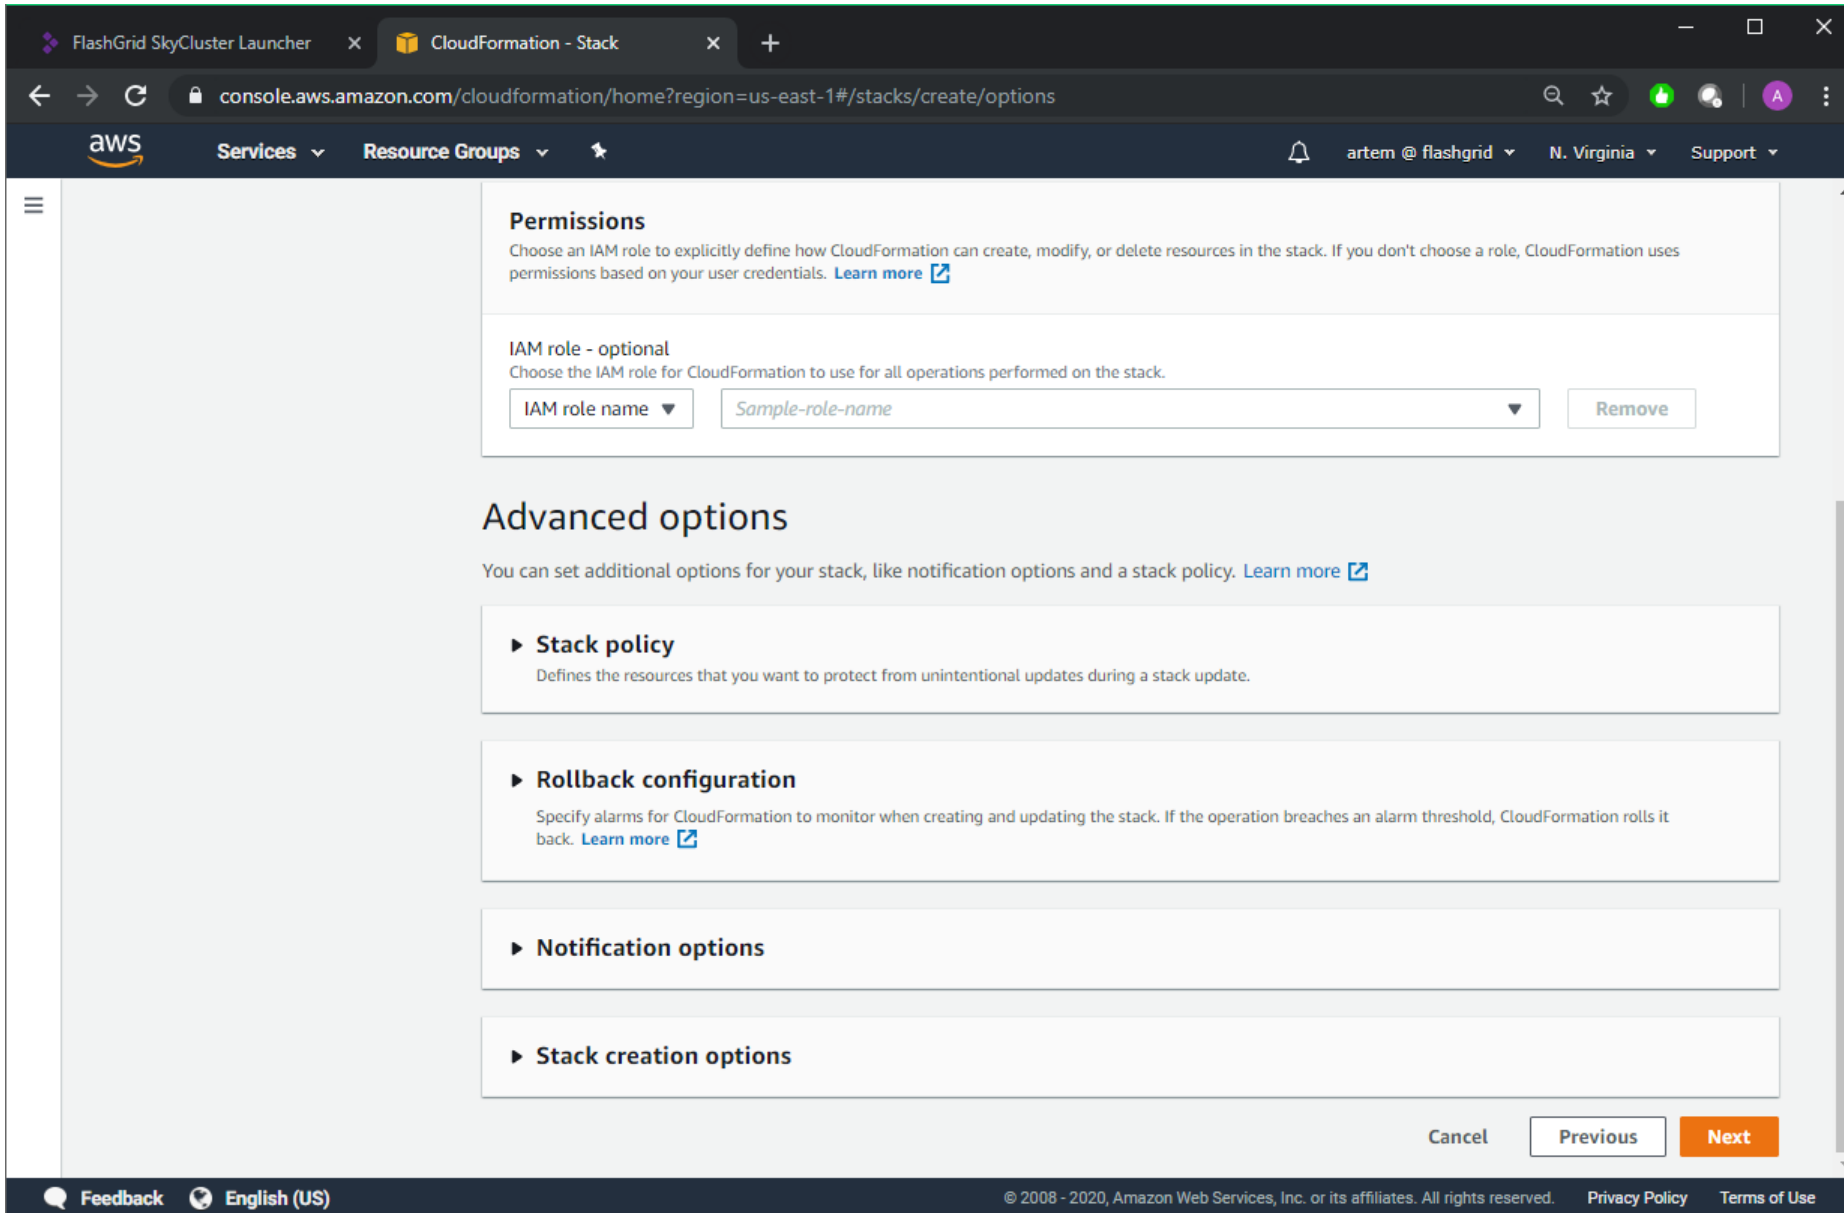
- Stack name**: A text input field containing "myrac". Below the field, a note states: "Stack name can include letters (A-Z and a-z), numbers (0-9), and dashes (-)." The label "Stack name" is also present above the input field.
- Parameters**: A section with the heading "Parameters" and a sub-heading "Cluster Nodes". It contains two input fields:
  - AMI**: A text input field containing "ami-065cb8343236ef9e2". The label "AMI" and the description "AMI for cluster nodes" are positioned above the field.
  - SSH key**: A dropdown menu containing "flashgrid-amazon". The label "SSH key" and the description "SSH key for cluster nodes" are positioned above the dropdown.

At the bottom right of the form, there are three buttons: "Cancel", "Previous", and "Next".

The footer of the page includes a "Feedback" link, "English (US)" language selection, and copyright information: "© 2008 - 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved." along with "Privacy Policy" and "Terms of Use" links.

## Step 17: Create CloudFormation stack – advanced options

You can specify additional options for creating a stack or just click Next.



The screenshot shows the AWS CloudFormation console interface for the 'Advanced options' step of creating a stack. The browser address bar shows the URL: `console.aws.amazon.com/cloudformation/home?region=us-east-1#/stacks/create/options`. The page title is 'CloudFormation - Stack'. The user is logged in as 'artem @ flashgrid' in the 'N. Virginia' region. The main content area is titled 'Advanced options' and includes a 'Permissions' section with an 'IAM role - optional' dropdown menu. Below this, there are four expandable sections: 'Stack policy', 'Rollback configuration', 'Notification options', and 'Stack creation options'. At the bottom right, there are three buttons: 'Cancel', 'Previous', and 'Next'.

**Permissions**  
Choose an IAM role to explicitly define how CloudFormation can create, modify, or delete resources in the stack. If you don't choose a role, CloudFormation uses permissions based on your user credentials. [Learn more](#)

**IAM role - optional**  
Choose the IAM role for CloudFormation to use for all operations performed on the stack.

IAM role name ▼  ▼

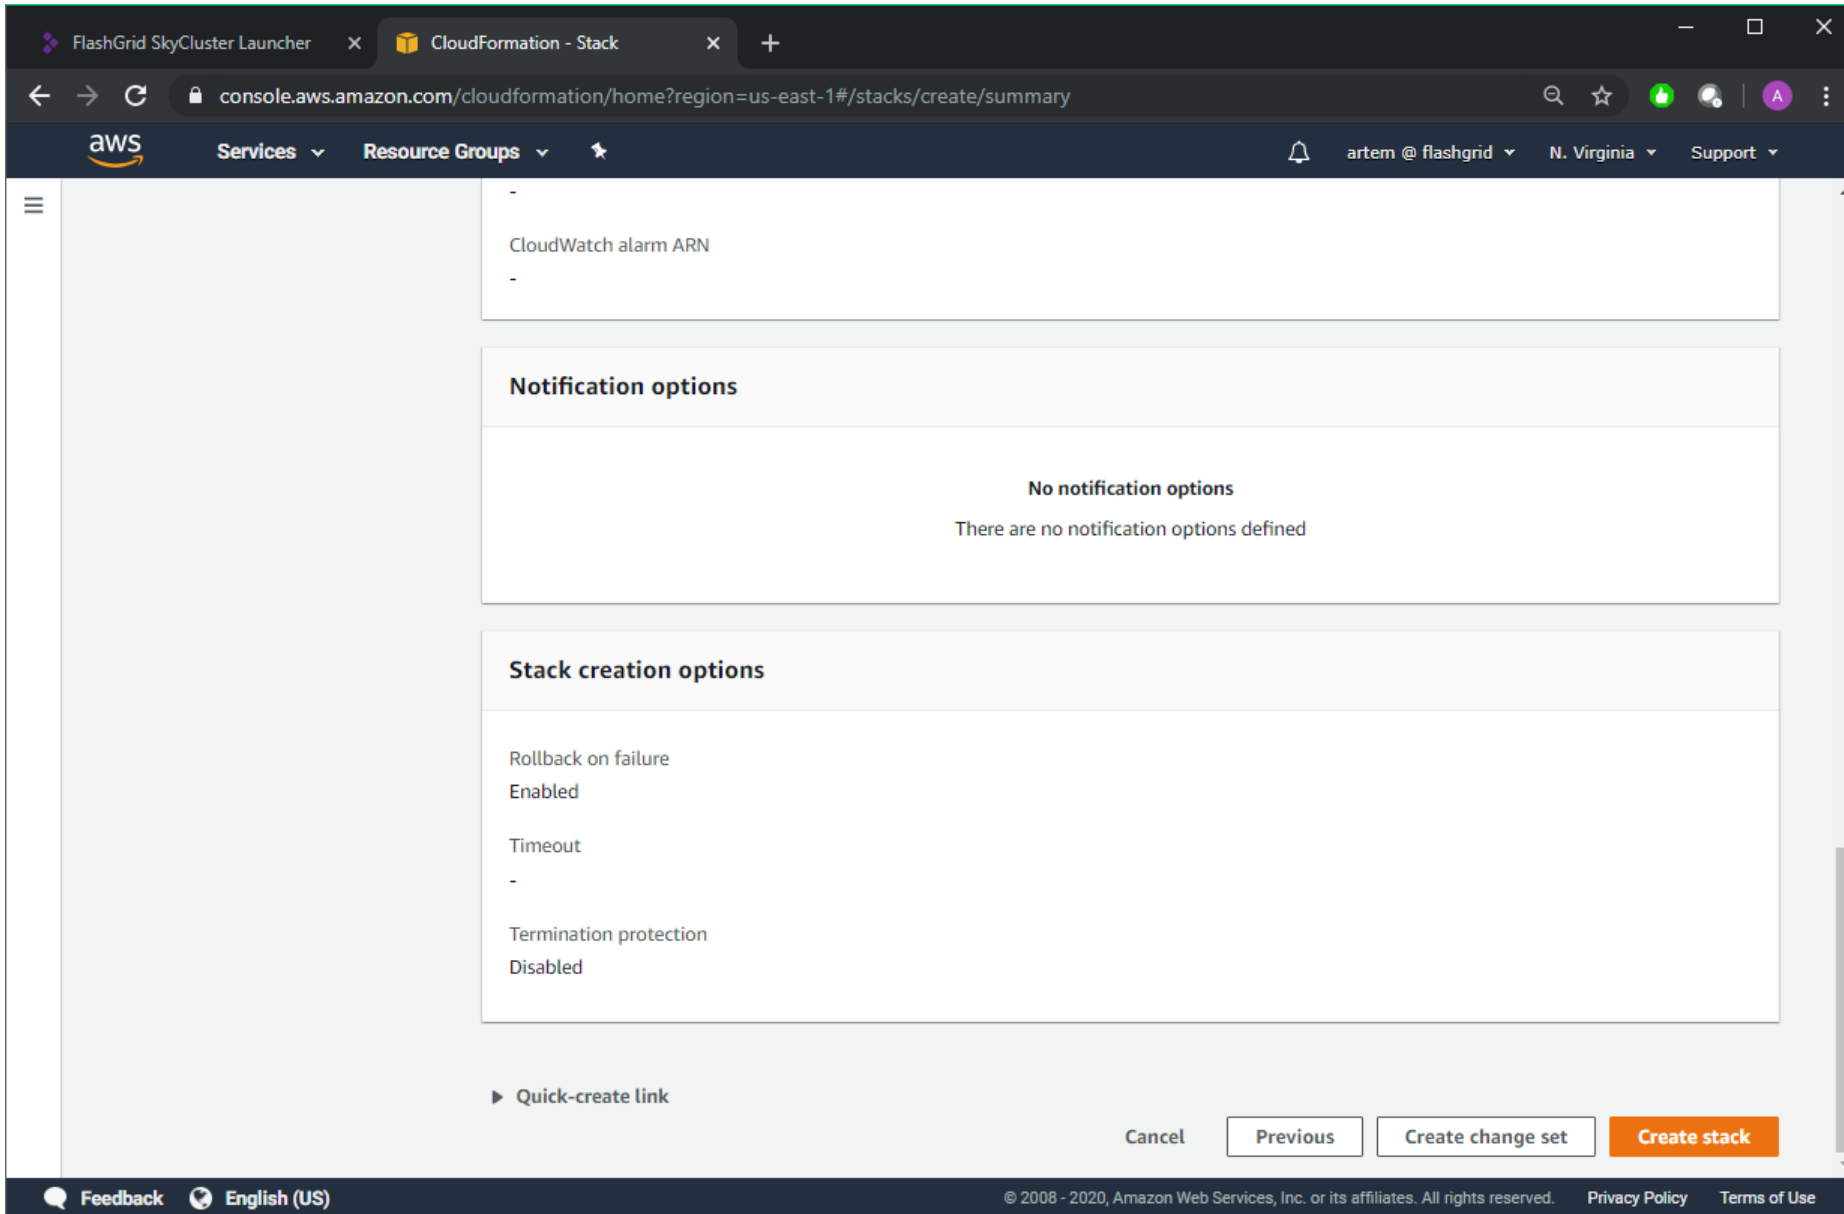
### Advanced options

You can set additional options for your stack, like notification options and a stack policy. [Learn more](#)

- ▶ **Stack policy**  
Defines the resources that you want to protect from unintentional updates during a stack update.
- ▶ **Rollback configuration**  
Specify alarms for CloudFormation to monitor when creating and updating the stack. If the operation breaches an alarm threshold, CloudFormation rolls it back. [Learn more](#)
- ▶ **Notification options**
- ▶ **Stack creation options**

## Step 17: Create CloudFormation stack – review and create

Review stack options and click *Create stack*.



The screenshot shows the AWS CloudFormation console in the 'Create stack' summary view. The browser address bar shows the URL: `console.aws.amazon.com/cloudformation/home?region=us-east-1#/stacks/create/summary`. The page header includes the AWS logo, 'Services', 'Resource Groups', and user information for 'artem @ flashgrid' in 'N. Virginia'.

The main content area is divided into several sections:

- CloudWatch alarm ARN:** A text field containing a hyphen (-).
- Notification options:** A section with the heading 'No notification options' and the text 'There are no notification options defined'.
- Stack creation options:** A section with the following settings:
  - Rollback on failure: Enabled
  - Timeout: -
  - Termination protection: Disabled

At the bottom of the main content area, there is a 'Quick-create link' with a right-pointing triangle icon.

The bottom navigation bar contains the following buttons: 'Cancel', 'Previous', 'Create change set', and 'Create stack' (highlighted in orange).

The footer of the page includes a 'Feedback' icon, 'English (US)' language selection, and copyright information: '© 2008 - 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved. Privacy Policy Terms of Use'.

## Step 18: SSH in to the first node

After the CloudFormation stack is created, use SSH to connect to the first node of the cluster as user *fg*. If the cluster init already finished (takes 60-90 minutes) then you will see the following message.

```
+-----+
|                                     |
| CLUSTER INITIALIZATION COMPLETED SUCCESSFULLY |
|                                     |
+-----+
| Please follow the steps below to finalize cluster configuration: |
|                                     |
| 1. Run 'flashgrid-cluster' to verify status of the cluster. |
|                                     |
| 2. Protect the cluster from accidental deletion: |
|    - in AWS/GCP enable instance termination protection for each node |
|    - in Azure add a lock to the cluster resource group |
|                                     |
| 3. Add records to the DNS servers used by clients and app servers: |
|                                     |
|    rac2.example.com 10.100.101.30 |
|    racl.example.com 10.100.101.145 |
|                                     |
| 4. Test email alerts from each node: $ flashgrid-node test-alerts |
|                                     |
| 5. Disable database features that you do not have a license for. |
|                                     |
| 6. See FlashGrid Knowledge Base for instructions for the following tasks: |
|    - Changing temporary ASM password: kb.flashgrid.io/asm-password |
|    - Creating a database: kb.flashgrid.io/createdb |
|    - Connecting clients to a database: kb.flashgrid.io/connect-clients |
|    - Maintenance procedures (reboot, etc.): kb.flashgrid.io/maintenance |
|                                     |
| 7. Before putting cluster in production, upload diags for review by support |
|    $ sudo flashgrid-diags upload-all |
|                                     |
| Submit support requests at flashgrid.io/support |
|                                     |
| To stop seeing this message after login, run 'sudo rm /etc/motd' |
+-----+
| [fg@racl ~]$ █ |
```

## Step 19: Check status of the cluster

```
# sudo flashgrid-cluster
```

```
[fg@rac1 ~]$ flashgrid-cluster
FlashGrid 19.6.266.68520 #ec4afccc6e8189e7fd7d53feld899757d54bf7bc
License: Active, Expires 2020-06-03
Licensee: demo
Support plan: Demo
-----
FlashGrid running: OK
Clocks check: OK
Configuration check: OK
Network check: OK

Querying nodes: rac1, rac2, racq ...

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           RACQ
-----

GroupName  Status  Mounted  Type    TotalMiB  FreeMiB  OfflineDisks  LostDisks  Resync  ReadLocal  Vote
-----
DATA       Good    AllNodes  NORMAL  12288     11896    0              0          No     Enabled   None
FRA        Good    AllNodes  NORMAL  12288     11936    0              0          No     Enabled   None
GRID       Good    AllNodes  NORMAL  10240     9496     0              0          No     Enabled   3/3
-----
[fg@rac1 ~]$
```

The cluster is ready. You can now create your database using DBCA.

## Additional Information

For more information see <https://www.flashgrid.io/skycluster-for-aws> or email [info@flashgrid.io](mailto:info@flashgrid.io)

Copyright © 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. Oracle and Java are registered trademarks of Oracle and/or its affiliates. Red Hat is a registered trademark of Red Hat Inc. Microsoft and Azure are registered trademarks of Microsoft Corporation. Other names may be trademarks of their respective owners.