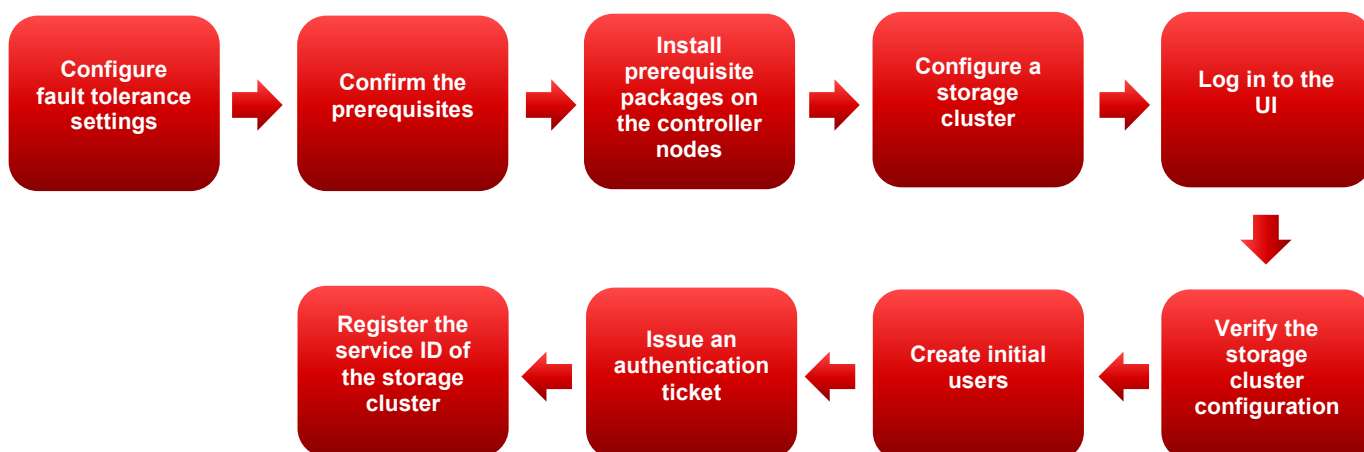


Virtual Storage Platform One SDS Block Deployment on Amazon Web Services

VSP One SDS Block is a Distributed Software-Defined Solution tailored to enhance shared-nothing architectures and modern applications driving digital transformation. VSP One SDS Block adds resiliency, improves performance, and reduces infrastructure complexities of hybrid cloud environments. The following document highlights the general steps needed to deploy the VSP One SDS Block solution on AWS.

For details, see the *Hitachi Virtual Storage Platform One VSP One SDS Block Cloud v1.17 Setup and Configuration Guide*.



The software can be deployed in the following four configurations.

For the single-AZ configuration:

- When configuring a compute node in the same VPC as that for SDS Block
- When configuring a compute node in a different VPC from that for SDS Block

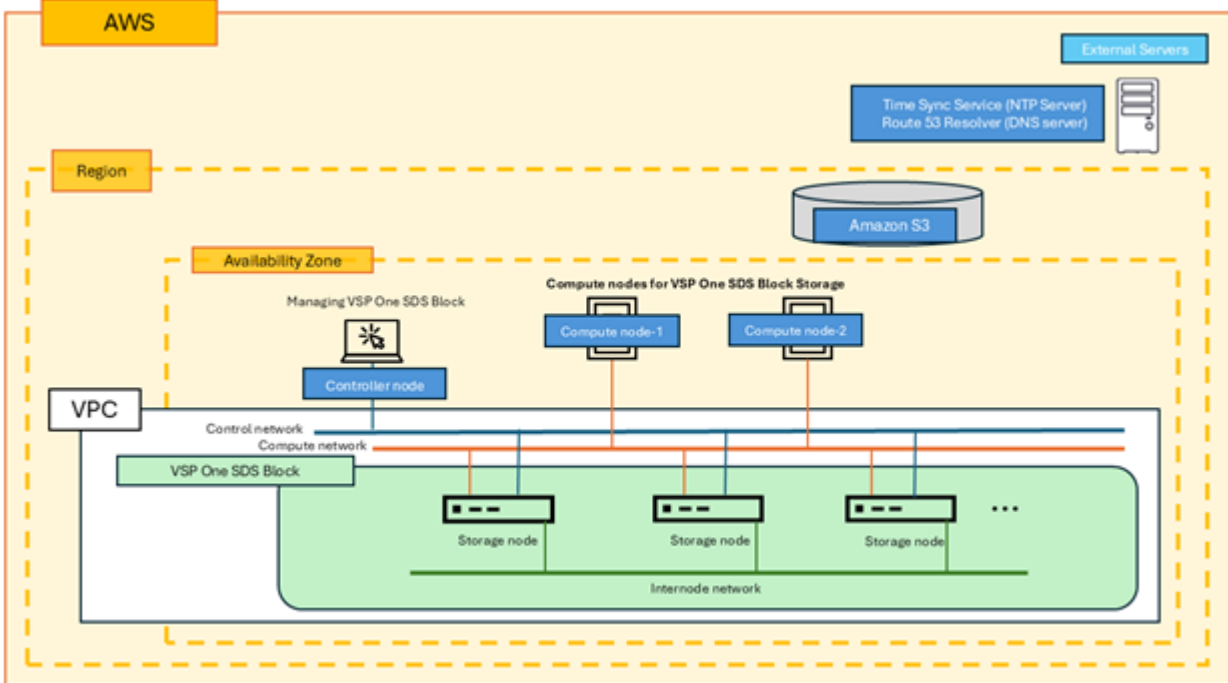
For the multi-AZ configuration:

- When configuring a compute node in the same VPC as that for SDS Block
- When configuring a compute node in a different VPC from that for SDS Block

See the diagrams on the following pages.

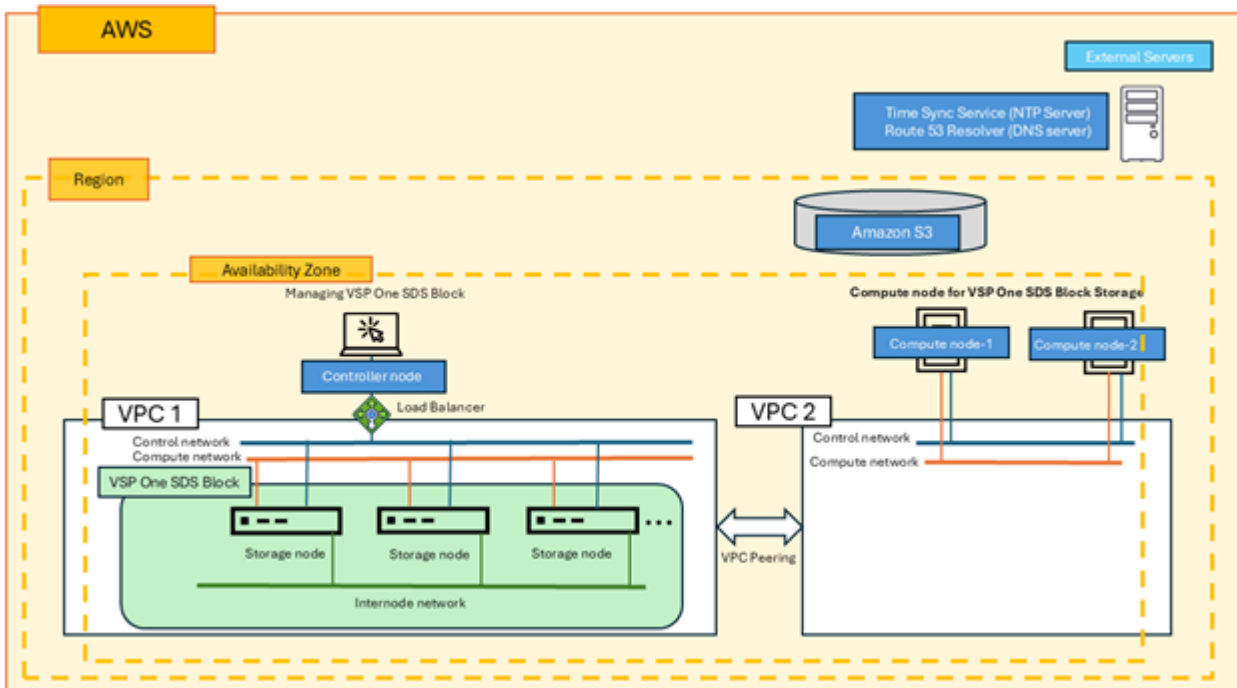
For the single-AZ configuration:

- When configuring a compute node in the same VPC as that for SDS Block



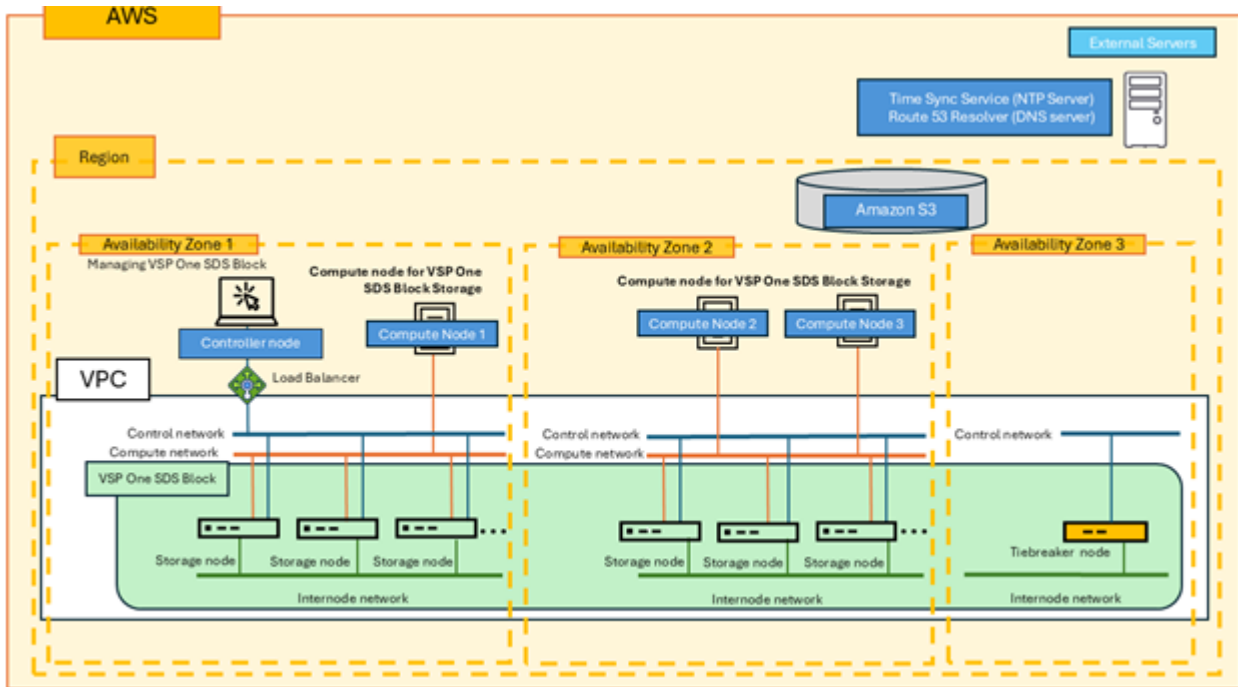
For the single-AZ configuration:

- When configuring a compute node in a different VPC from that for SDS Block



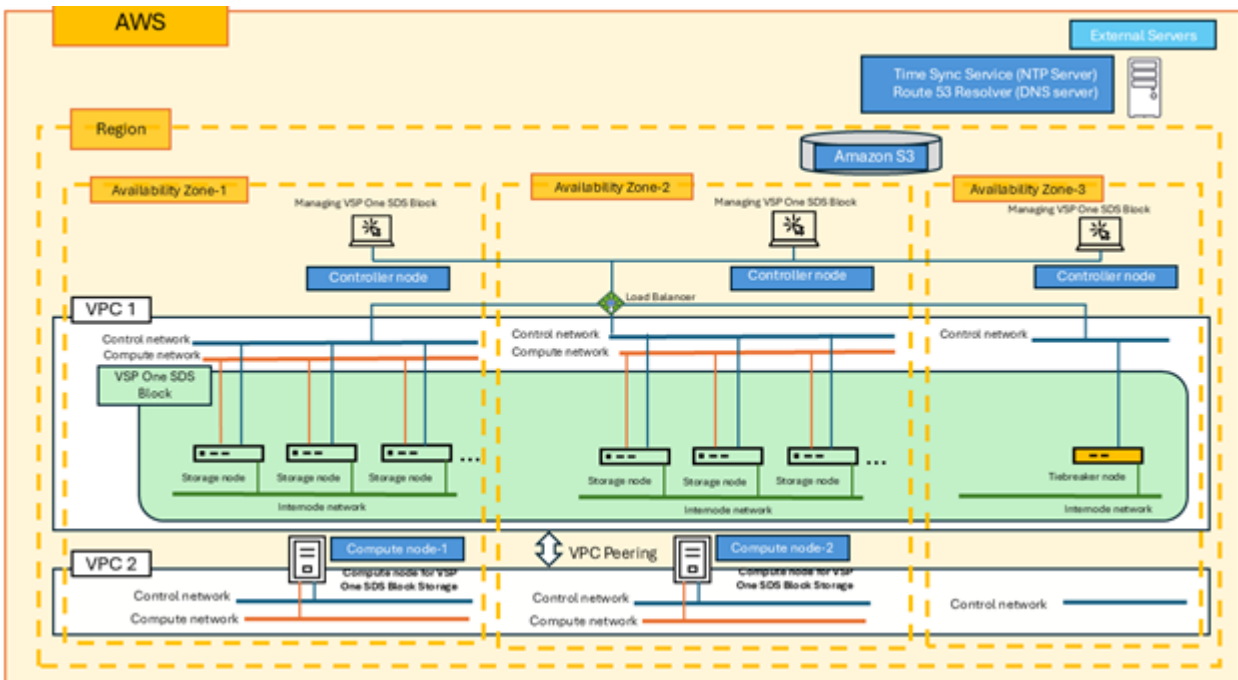
For the multi-AZ configuration:

- When configuring a compute node in the same VPC as that for SDS Block



For the multi-AZ configuration:

- When configuring a compute node in a different VPC from that for SDS Block



Step 1: Configure fault tolerance settings

1. Configure the following functions to create the fault tolerance of the storage cluster, based on whether you have a single-AZ configuration or a multi-AZ configuration:

- User data protection method
- Redundancy of the storage controller
- Redundancy of the cluster primary node
- Spread placement group
- Fault domain
- Shared Nothing Architecture with Data Localization regardless of Plus-1 or Plus-2

The following table shows the redundant configurations:

Configuration	Function
Plus-1 redundant configuration	<ul style="list-style-type: none"> • Higher write performance • Shorter time required to complete rebuilding
Plus-2 redundant configuration settings	<ul style="list-style-type: none"> • Higher fault tolerance

The following table shows the functions for a single-AZ configuration:

Configuration	Function settings			
	User data protection method	Redundancy of the storage controller(1)	Number of cluster master nodes	Number of fault domains
Plus-1 redundant configuration (1)	<ul style="list-style-type: none"> • Mirroring • Duplication 	OneRedundant StorageNode (degree = 3)	3 nodes	1
Plus-2 redundant configuration (2)	<ul style="list-style-type: none"> • HPEC 4D+2P 	TwoRedundant StorageNodes (degree = 3)	3 nodes	1

1. Do not explicitly specify the degree of redundancy of the storage controller, because it is automatically determined depending on the user data protection method. The tiebreaker node does not contribute to the storage capacity of the SDS Block cluster.

2. In the following cases, only one failure is assumed to have occurred irrespective of the number of failures:

- One or more drive failures occurred on a faulty storage node.
- Drive failures occurred on a single storage node.

You can use a fixed reserve rebuild policy to protect against drive failures of EBS in the above two use cases.

The following table shows the settings for a multi-AZ configuration:

Configuration	Function settings			
	User data protection method	Redundancy of the storage controller(1)	Number of cluster primary nodes	Number of fault domains
Plus-1 redundant configuration	<ul style="list-style-type: none"> • Mirroring • Duplication 	OneRedundant StorageNode (degree = 2)	3 nodes	3
<p>1. Do not explicitly specify the degree of redundancy of the storage controller, because it is automatically determined depending on the user data protection method.</p> <p>2. In the following cases, only one failure is assumed to have occurred, irrespective of the number of failures:</p> <ul style="list-style-type: none"> • One or more drive failures occurred on a faulty storage node. • Drive failures occurred on a single storage node. <p>You can use a fixed reserve rebuild policy to protect against drive failures of EBS.</p>				

Step 2: Confirm the prerequisites

Work with your administrator to set up these privileges.

1. If you want to perform remote copy, establish communication between the compute network subnet and the network in which the external storage system is installed.
2. Locate the files required for setup:
 - SDS Block installer package: `hsds_installer-<version>.<number>-py3-none-any.whl`
 - CLI package file: `hsds_cli-<version>.<number>-py3-none-any.whl`
 - EULA documentation on SUSE Linux Enterprise and Cavium SPDK FC Target Driver used by the storage node: `EULA_for_SLE_for_StorageNode.txt` and `EULA_for_SPDK.txt`
 - If you are communicating with SDS Block in SSL/TLS: server certificate to import into the storage node
3. Create an AWS account.
4. Create an Amazon S3 bucket or folder in the AWS region where you will configure a storage cluster. The bucket or folder is for storing error message files or dump log files.
5. Create an IAM user with the AWS management policy `AWSMarketplaceFullAccess` and the specific permissions, or an IAM user with administrative privileges. For details, see *Confirming prerequisites for startup* in the setup and configuration guide.
6. Create an IAM role with specific privileges. Set EC2 with the trusted entity of the IAM role. For details, see *Confirming prerequisites for startup* in the setup and configuration guide.
7. Verify that the VPC has sufficient IP address range to meet specific conditions. For details, see *Confirming prerequisites for startup* in the setup and configuration guide.
8. Enable the DNS resolution setting and DNS host name.
9. Create VPC endpoints that meet the requirements.
10. Create the VPC endpoint for AWS S3 as a Gateway endpoint.
11. (PayGo option) Create the following items for the AWS License Manager:
 - VPC endpoint for AWS License Manager
 - (Single-AZ configuration): 1 VPC endpoint
 - (Multi-AZ configuration): 3 VPC endpoints {1 VPC endpoint for each Availability Zone for redundancy}
12. Create the subnets that meet the requirements for the control network, internode network, and compute network.
13. Establish communication between the control network subnet and its outside network.
14. Verify that there is no communication between the internode network subnet and its outside network.

15. (Optional) Enable EBS encryption.
16. If you do not use VPC endpoints for the AWS License Manager, ensure that the control network can access the internet.
17. (Multi-AZ) Create the following subnets in each Availability Zone:
 - For an AZ in which you will install storage nodes: Control network subnet, internode network subnet, compute network subnet
 - For an AZ in which you will install a tiebreaker node: Control network subnet, internode network subnet

Step 3: Install prerequisite packages on the controller nodes, and configure and deploy controller nodes

A controller node is required to perform cluster maintenance.

1. Log in to the EC2 instance for a controller node.
2. For Linux, switch to a root user. For Windows, switch to an admin user.
3. Download the installer package or copy the package to the controller node.
4. Download the CLI package.
5. Install the following packages on the controller node:
 - Python: 3.11 or later
 - pip: 20.0.2 or later
 - unzip: 6.00 or later
6. Upgrade the installed pip.
7. Install the SDS Block installer package.
8. Verify that the version of the SDS Block installer matches that of the configured storage cluster by running the **hdsinstall -version** command.
9. Install the CLI package.
10. If you plan to import a server certificate into a storage node and use your own root CA certificate to verify the server certificate, add your own root CA certificate according to the *VSP One SDS Block System Administrator Operation Guide*.
11. Configure one controller node for each storage cluster.
12. Set up a console to run on a controller node so that you can access it in the following tasks.
13. Add AWS Key Management Service access rights to the IAM role.
14. Create an EC2 instance for the controller node in each cluster.
15. Deploy the EC2 instance in the subnet for the control network created earlier.
16. Set a security group for the control network.
17. Confirm the security group for the control network by opening the stack window of NetworkResources, and then referring to ControlSecurityGroup in the **Resources** tab.
18. Grant the IAM role specific privileges to the controller node.
19. Set the AWS credentials on the controller node. Run the **aws configure** command on the controller node with the ec-2 user, and then specify the region and output format.

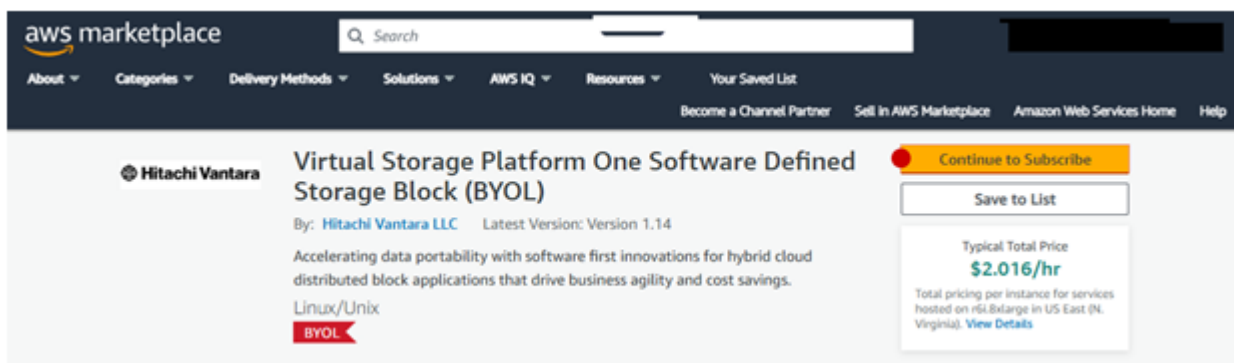
Storage node requirements

The following table lists the requirements for the SDS Block storage node.

Item	Requirement	Remarks
Instance type	<ul style="list-style-type: none"> m7i.8xlarge m6i.8xlarge r7i.8xlarge r6i.8xlarge 	<ul style="list-style-type: none"> You can change the setting value only at the time of initial configuration. Do not change the setting value after initial configuration. The system might become inoperable. When you create a volume with data reduction enabled, set the instance type r7i.8xlarge or r6i.8xlarge. In the case of Multi-AZ configuration, r5.xlarge is always set for tiebreaker node.
User data drive	Number of units	6 to 24
		In the case of Multi-AZ configuration, no user data drive is created for tiebreaker node.

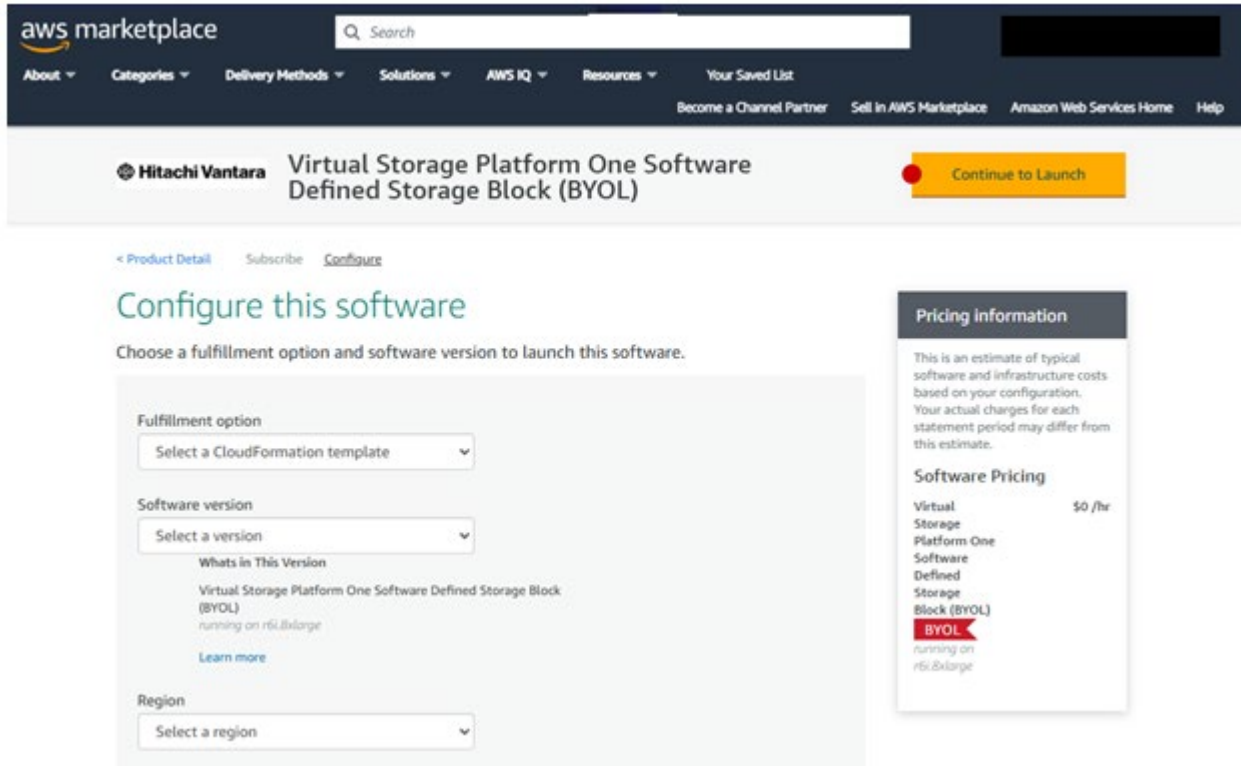
Step 4: Configure a storage cluster

- As the IAM user, access the SDS Block product page.
- Click Continue to subscribe.

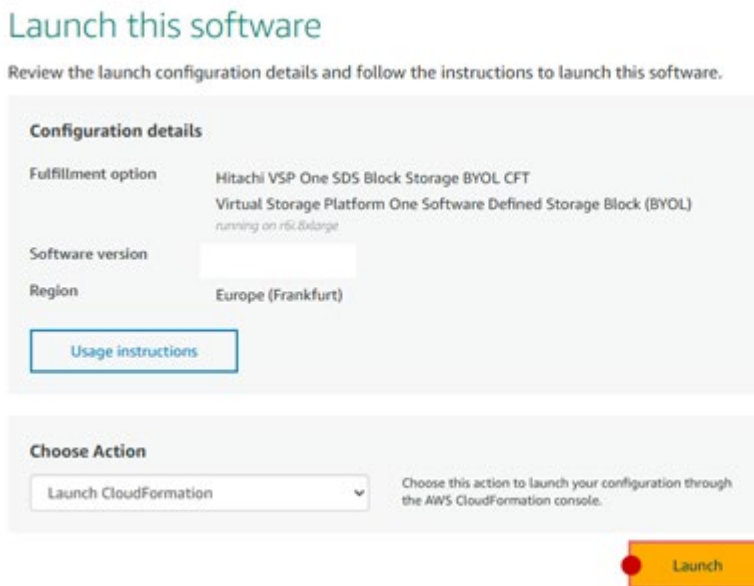


- (PayGo) Specify the license period, licensed capacity, and auto-update settings managed by AWS Manager. If you signed a license contract in the AWS account, go to step 4.
- (PayGo) Confirm the license fee, and then click **Create Contract**.
- (PayGo or BYOL) In the Subscribe to this Software window, confirm and agree to the Terms and Conditions, and then click **Accept Terms**. Allow time for the process to finish.
You only need to perform this step the first time you log on.
- Click **Continue to Configuration**.

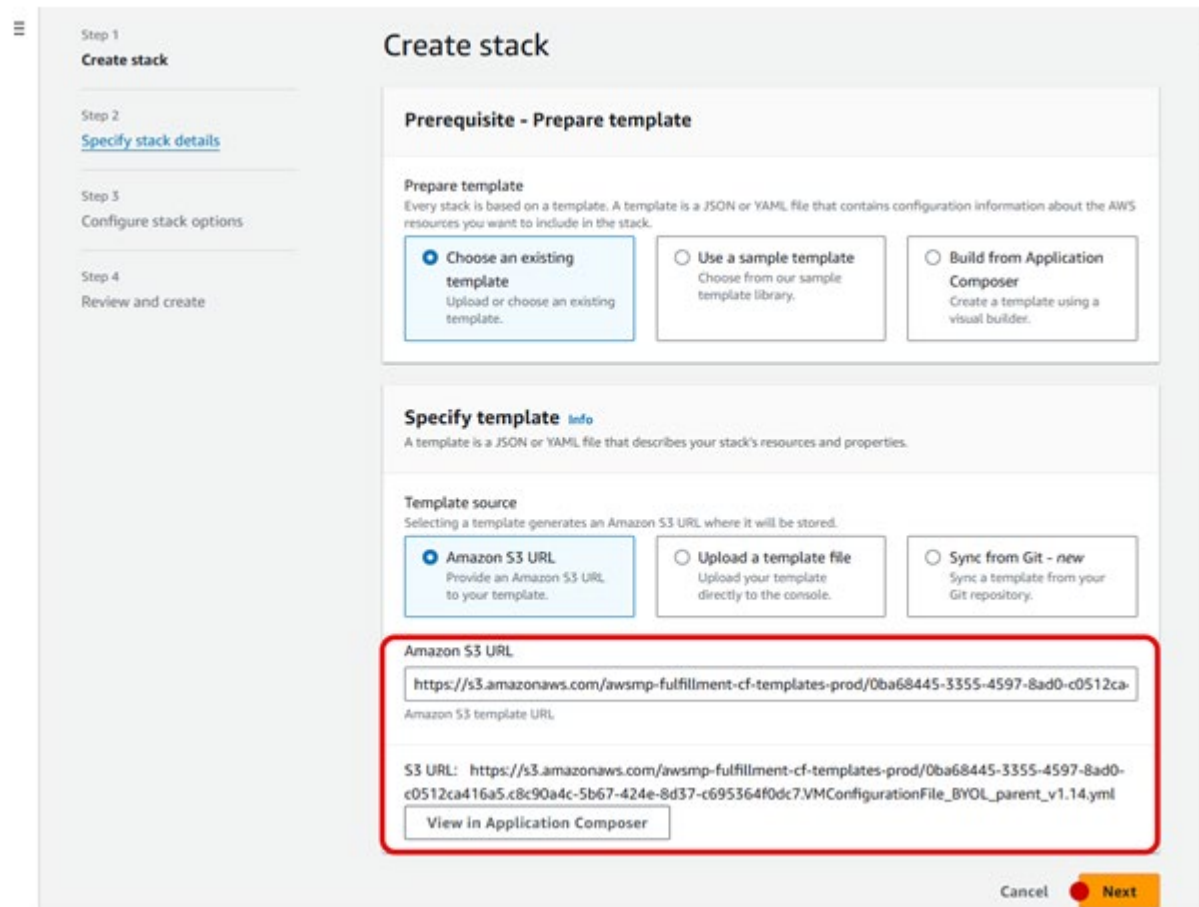
7. Select a product configuration:
 - Fulfillment option: Single-AZ configuration or multi-AZ configuration
 - Software Version: Version of SDS Block to be set up
 - Region: Region of AWS for configuring a storage cluster



8. Click **Continue to Launch**.
9. Click **Launch** in the Launch this software window.



10. (Optional, BYOL) If you did not previously enable EBS encryption, enable Hitachi Data at rest encryption. Data at rest encryption is available only for BYOL users.
11. In the Create Stack window, verify that the Amazon S3 URL is displayed, and then click **Next**.



Step 1
Create stack

Step 2
[Specify stack details](#)

Step 3
Configure stack options

Step 4
Review and create

Create stack

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.

- Choose an existing template**
Upload or choose an existing template.
- Use a sample template**
Choose from our sample template library.
- Build from Application Composer**
Create a template using a visual builder.

Specify template Info

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**
Provide an Amazon S3 URL to your template.
- Upload a template file**
Upload your template directly to the console.
- Sync from Git - new**
Sync a template from your Git repository.

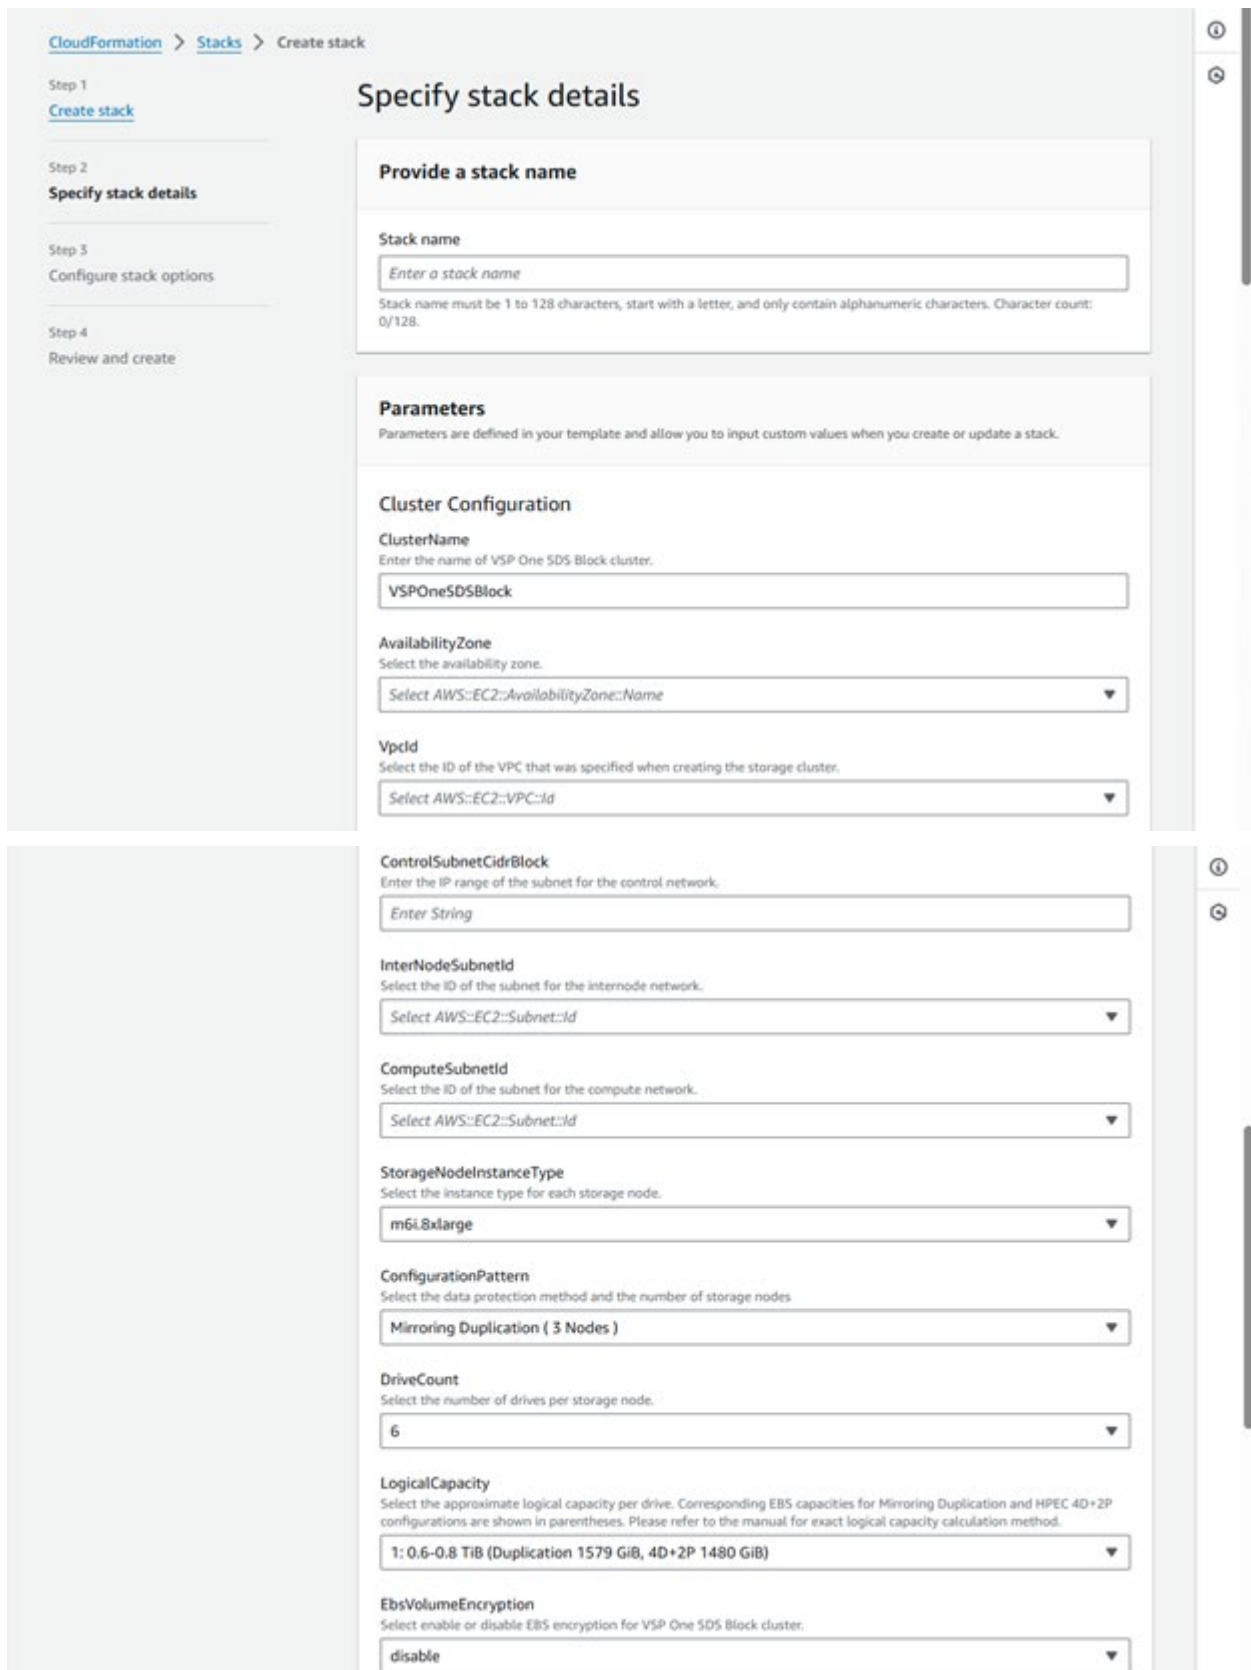
Amazon S3 URL

Amazon S3 template URL

S3 URL: `https://s3.amazonaws.com/awssmp-fulfillment-cf-templates-prod/0ba68445-3355-4597-8ad0-c0512ca416a5.c8c90a4c-5b67-424e-8d37-c695364f0dc7.VMConfigurationFile_BYOL_parent_v1.14.yml`

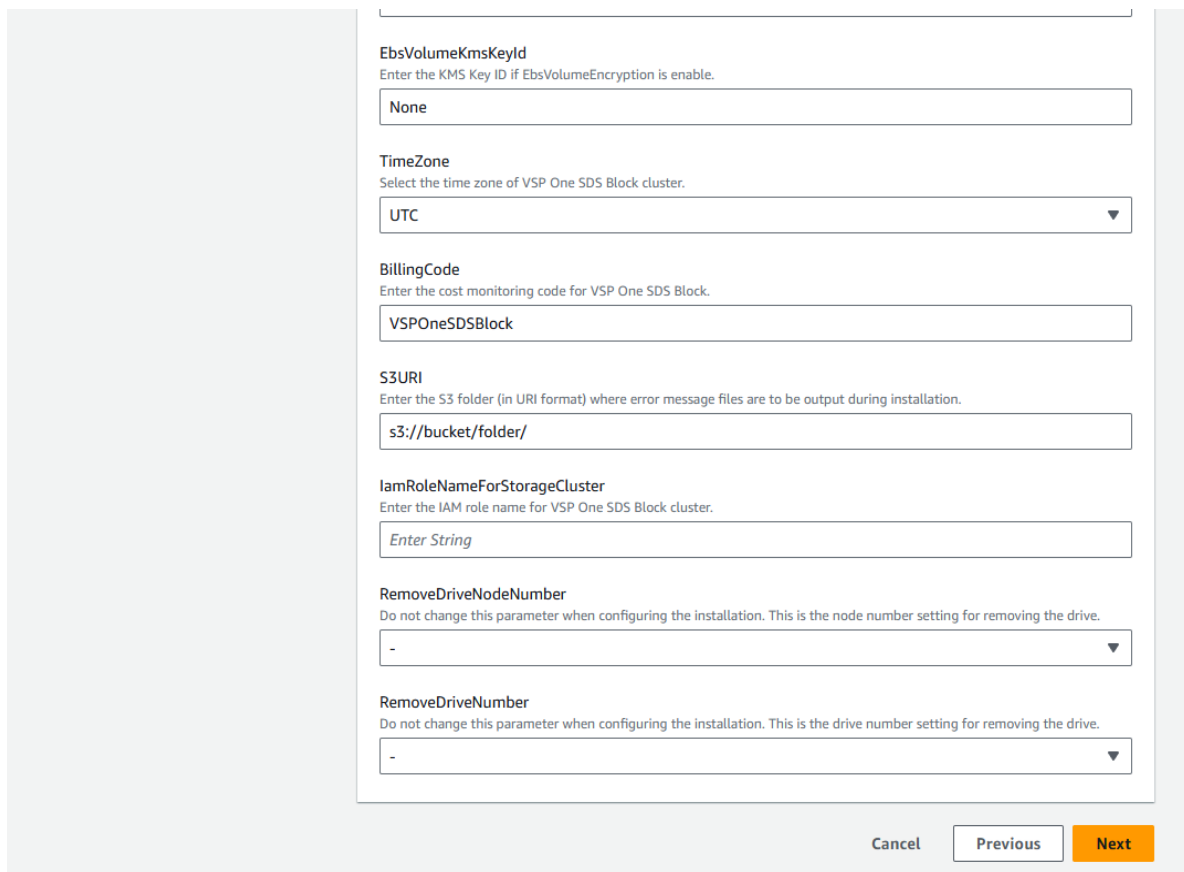
Cancel **Next**

12. In the Specify Stack Details window, specify the stack name and parameters, and then click **Next**. For details, see *Configuring a cluster* in the setup guide.



The screenshot shows the 'Specify stack details' page in the AWS CloudFormation console. The page is divided into several sections for configuring the stack:

- Provide a stack name:** A text input field with the placeholder 'Enter a stack name'. Below it, a note states: 'Stack name must be 1 to 128 characters, start with a letter, and only contain alphanumeric characters. Character count: 0/128.'
- Parameters:** A section titled 'Parameters' with a sub-section 'Cluster Configuration'.
 - ClusterName:** A text input field with the value 'VSPOneSDSBlock'. Description: 'Enter the name of VSP One SDS Block cluster.'
 - AvailabilityZone:** A dropdown menu with the selected value 'Select AWS::EC2::AvailabilityZone::Name'. Description: 'Select the availability zone.'
 - VpcId:** A dropdown menu with the selected value 'Select AWS::EC2::VPC::Id'. Description: 'Select the ID of the VPC that was specified when creating the storage cluster.'
- ControlSubnetCidrBlock:** A text input field with the placeholder 'Enter String'. Description: 'Enter the IP range of the subnet for the control network.'
- InterNodeSubnetId:** A dropdown menu with the selected value 'Select AWS::EC2::Subnet::Id'. Description: 'Select the ID of the subnet for the internode network.'
- ComputeSubnetId:** A dropdown menu with the selected value 'Select AWS::EC2::Subnet::Id'. Description: 'Select the ID of the subnet for the compute network.'
- StorageNodeInstanceType:** A dropdown menu with the selected value 'm6i.xlarge'. Description: 'Select the instance type for each storage node.'
- ConfigurationPattern:** A dropdown menu with the selected value 'Mirroring Duplication (3 Nodes)'. Description: 'Select the data protection method and the number of storage nodes.'
- DriveCount:** A dropdown menu with the selected value '6'. Description: 'Select the number of drives per storage node.'
- LogicalCapacity:** A dropdown menu with the selected value '1: 0.6-0.8 TiB (Duplication 1579 GiB, 4D+2P 1480 GiB)'. Description: 'Select the approximate logical capacity per drive. Corresponding EBS capacities for Mirroring Duplication and HPEC 4D+2P configurations are shown in parentheses. Please refer to the manual for exact logical capacity calculation method.'
- EbsVolumeEncryption:** A dropdown menu with the selected value 'disable'. Description: 'Select enable or disable EBS encryption for VSP One SDS Block cluster.'



EbsVolumeKmsKeyId
Enter the KMS Key ID if EbsVolumeEncryption is enable.

None

TimeZone
Select the time zone of VSP One SDS Block cluster.

UTC

BillingCode
Enter the cost monitoring code for VSP One SDS Block.

VSPOneSDSBlock

S3URI
Enter the S3 folder (in URI format) where error message files are to be output during installation.

s3://bucket/folder/

IamRoleNameForStorageCluster
Enter the IAM role name for VSP One SDS Block cluster.

Enter String

RemoveDriveNodeNumber
Do not change this parameter when configuring the installation. This is the node number setting for removing the drive.

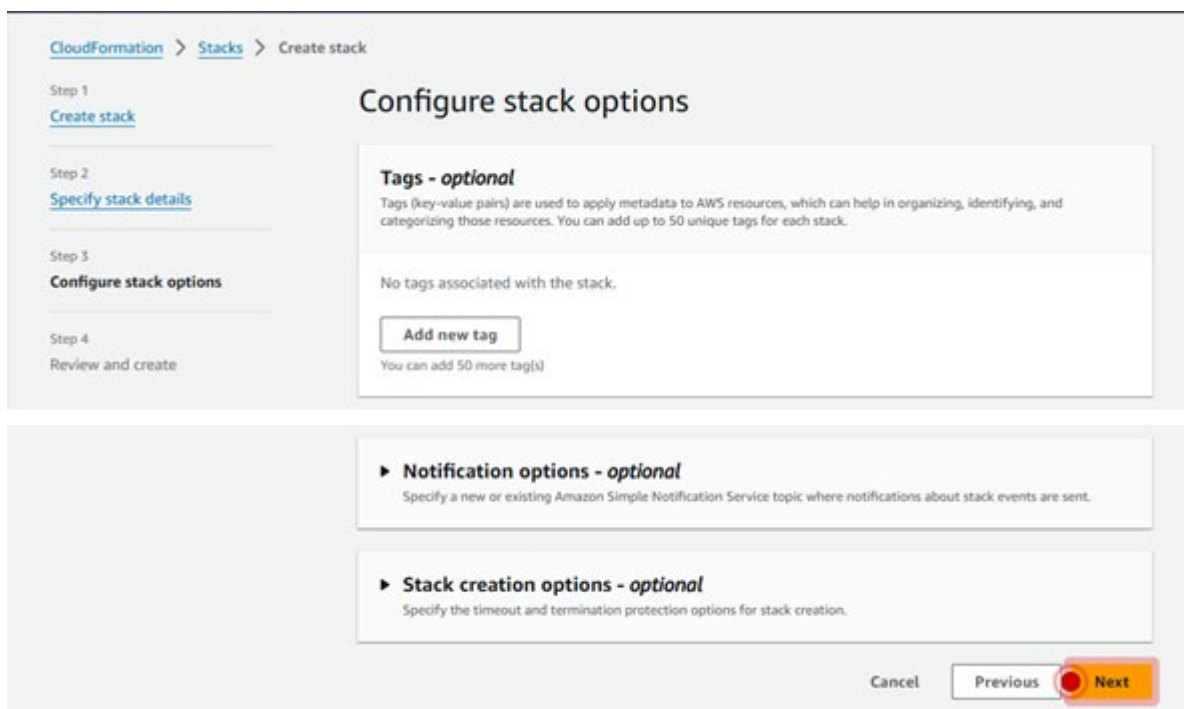
-

RemoveDriveNumber
Do not change this parameter when configuring the installation. This is the drive number setting for removing the drive.

-

Cancel Previous **Next**

13. In the Configure stack options window, click **Next**.



CloudFormation > Stacks > Create stack

Step 1
[Create stack](#)

Step 2
[Specify stack details](#)

Step 3
Configure stack options

Step 4
[Review and create](#)

Configure stack options

Tags - optional
Tags (key-value pairs) are used to apply metadata to AWS resources, which can help in organizing, identifying, and categorizing those resources. You can add up to 50 unique tags for each stack.

No tags associated with the stack.

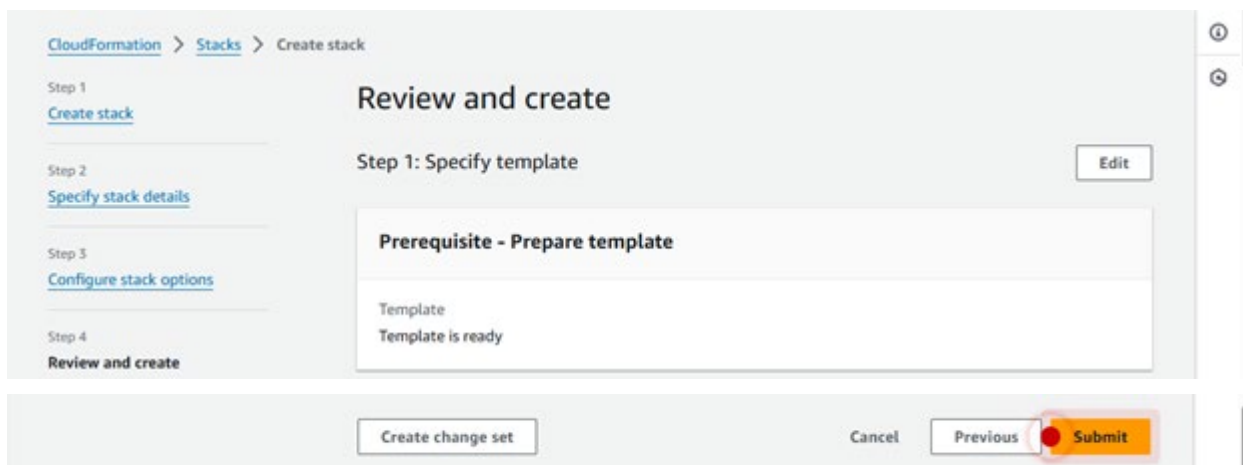
[Add new tag](#)
You can add 50 more tag(s)

Notification options - optional
Specify a new or existing Amazon Simple Notification Service topic where notifications about stack events are sent.

Stack creation options - optional
Specify the timeout and termination protection options for stack creation.

Cancel Previous **Next**

- In the Review and create window, confirm the settings, click **Submit**, and then verify that the SDS Block installation finishes.



The approximate time required until processing completes differs depending on the configuration, but it takes approximately 30 minutes in a configuration of three or six storage nodes.

- When the message "The following resources require capabilities" is displayed, select the checkbox.
- Confirm the parameters of the target group, IP address, and DNS name.

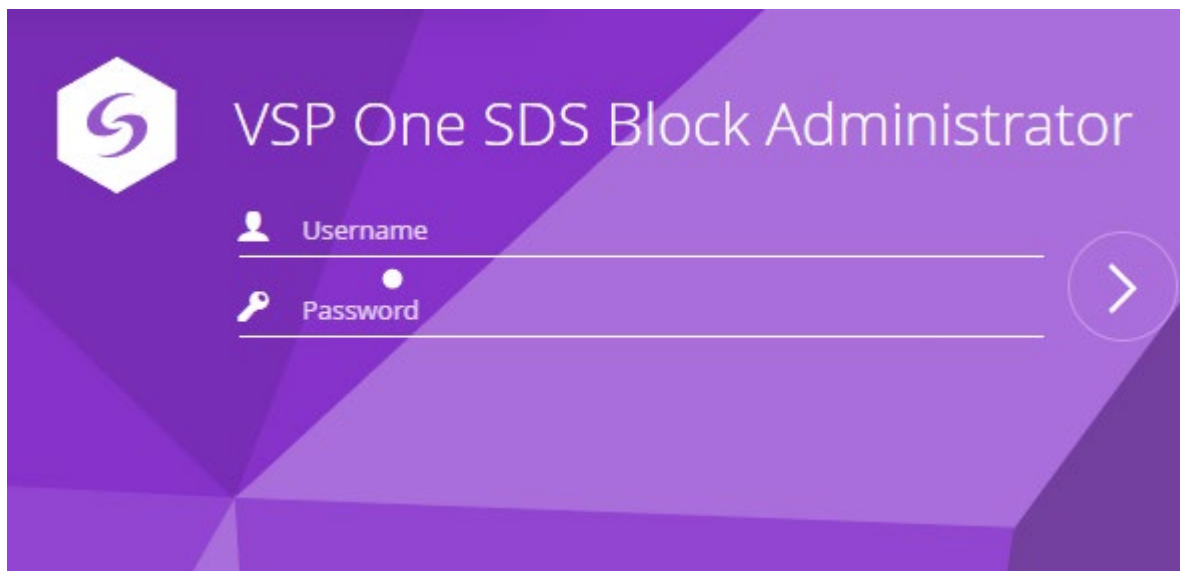
Step 5: Log in to the UI

- In the web browser, enter the following URL:

`https://<<IP-address-or-corresponding-FQDN>>[:443]/hsds/`

Specify the IP address of the load balancer confirmed in *Step 4: Configure a storage cluster* or the FQDN that corresponds to that IP address.

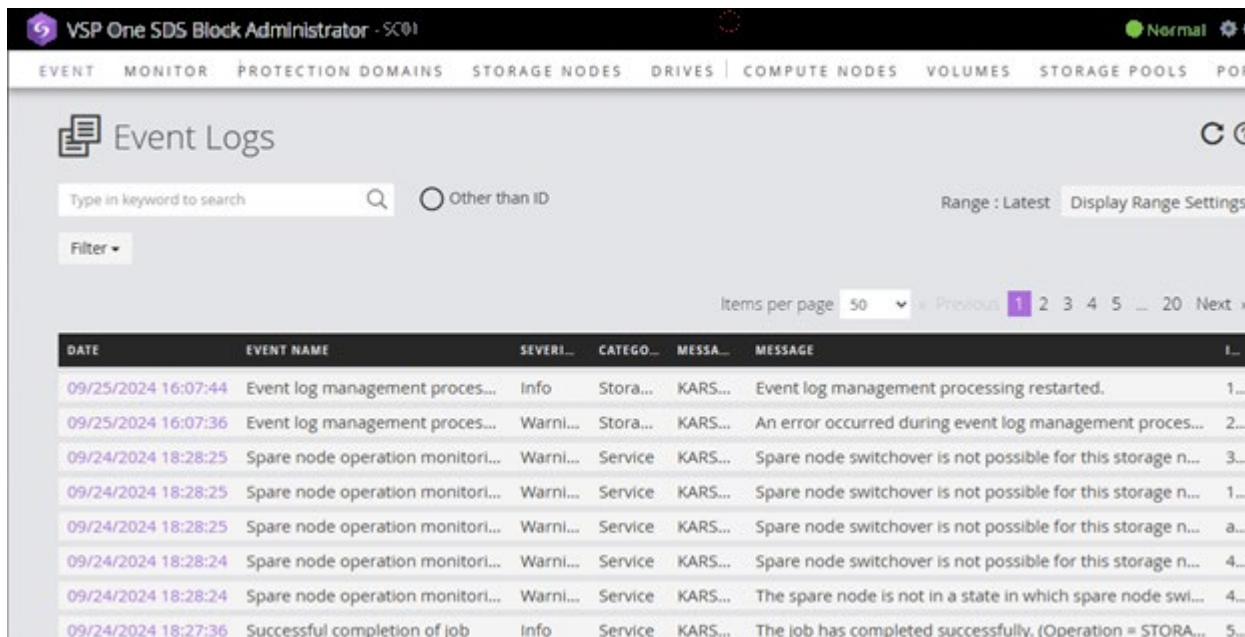
- When the login window is displayed, enter your username and password.



- In the Change Password window, specify a new password, and click Submit.

Step 6: Verify the storage cluster configuration

1. Obtain a list of event logs, and then verify the storage cluster has started up.



DATE	EVENT NAME	SEVERITY	CATEGORY	MESSAGE	MESSAGE	...
09/25/2024 16:07:44	Event log management proces...	Info	Stora...	KARS...	Event log management processing restarted.	1...
09/25/2024 16:07:36	Event log management proces...	Warni...	Stora...	KARS...	An error occurred during event log management proces...	2...
09/24/2024 18:28:25	Spare node operation monitori...	Warni...	Service	KARS...	Spare node switchover is not possible for this storage n...	3...
09/24/2024 18:28:25	Spare node operation monitori...	Warni...	Service	KARS...	Spare node switchover is not possible for this storage n...	1...
09/24/2024 18:28:25	Spare node operation monitori...	Warni...	Service	KARS...	Spare node switchover is not possible for this storage n...	a...
09/24/2024 18:28:24	Spare node operation monitori...	Warni...	Service	KARS...	Spare node switchover is not possible for this storage n...	4...
09/24/2024 18:28:24	Spare node operation monitori...	Warni...	Service	KARS...	The spare node is not in a state in which spare node swi...	4...
09/24/2024 18:27:36	Successful completion of job	Info	Service	KARS...	The job has completed successfully. (Operation = STORA...	5...

2. Verify the configuration after obtaining information about each of the following items:

- Compute ports
- Drives
- Storage clusters
- Storage nodes
- Domains
- Storage controllers
- Storage pools
- Storage cluster time
- Storage cluster networks
- Protection domains
- Control ports
- Internode ports
- Storage node networks

Step 7: Create initial users

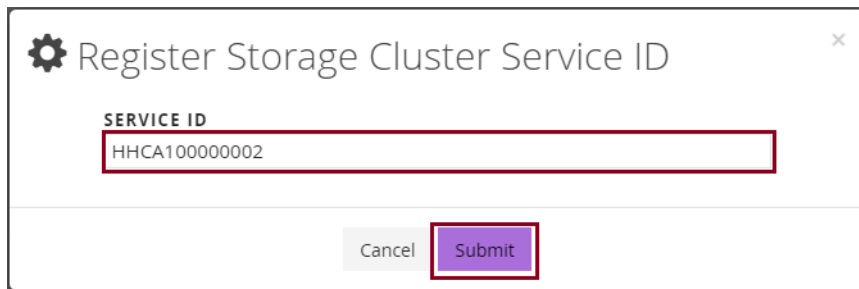
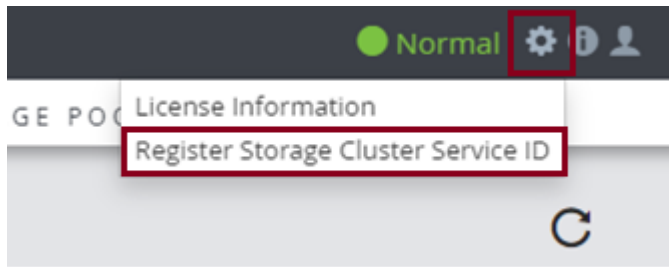
1. Configure SSL/TLS, and then import server certificates.
2. Create a user that belongs to the SecurityAdministrators group and a user that belongs to the ServiceAdministrators group. Or create one user that belongs to both groups.
3. Change the default passwords for those users.
4. For security purposes, invalidate the default user (userid: admin).
5. As the ServiceAdministrators user, verify that the storage cluster has started up.

Step 8: Issue an authentication ticket

1. As the ServiceAdministrators user, issue the authentication ticket for dumping in case of failure.
2. Create an authentication ticket file.

Step 9: Register the service ID of the storage cluster

1. If you are performing the task from the controller node, start the terminal.
2. Register the service ID of the storage cluster.



3. Verify the state of the job.
4. Verify the registered ID.

Reference

[VSP One SDS Block](#)