



QuickStart Guide to Scylla Cloud

Step-by-Step Instructions to Getting Up and
Running on the Scylla Database as a Service



More and more teams are choosing our Scylla Cloud as their database-as-a-service.

To make this transition even easier for you, this guide takes you step-by-step through how to get your Scylla Cloud cluster up and running quickly.



First Steps

In case you are not able to read this document in full, here are the most important things to remember:

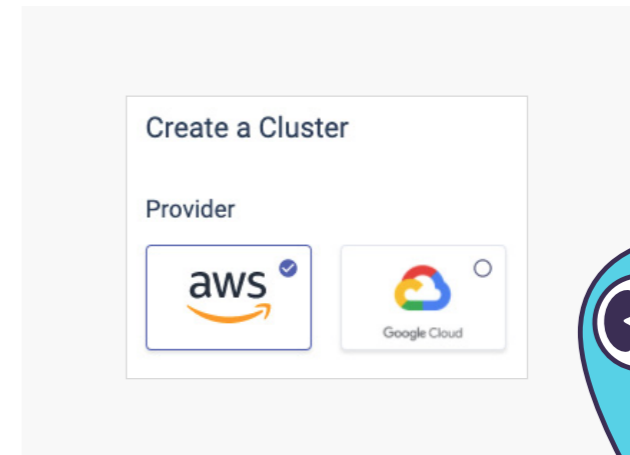
- 1 If you haven't already, [sign up](#) for a new Scylla Cloud account.
- 2 Recommended: Set up Two-Factor Authentication (2FA) for your Scylla Cloud user account.
- 3 Optional: Take advantage of our free trial offering and immediately create your first cluster.
- 4 Recommended: After setting up your free trial, check out our Scylla University [Scylla Cloud lab session](#), which walks you through all the steps to create a simple demo cluster. Then connect to your new cluster and execute basic commands to insert and select sample data.
- 5 Check out the other free [Scylla University courses](#).
- 6 Bookmark the [Scylla Cloud documentation](#) site.
- 7 [Join our community](#) on Slack!

Select Your Cloud Provider


Scylla Cloud is available on both AWS and Google Cloud public clouds, so you can choose your preferred cloud provider to run your cluster. We chose the most optimal compute instances with the best CPU/RAM ratio and local NVMe storage to guarantee predictable performance, high throughput, and low latencies for your applications.


- On the AWS cloud platform, we support high-performance storage optimized i3 and i3en instances.
- On Google Cloud, we are utilizing n2-highmem machines.


Based on your selection, the Scylla Cloud wizard will automatically update the list of geographic regions where you can deploy your cluster and instance types, providing the estimated cost of cloud compute resources per hour per node.











Where to Deploy? ⓘ


Scylla AWS Account


Your AWS Account

 **SCYLLA**

-  New Cluster
-  My Clusters
-  Support 
-  Docs 
-  Scylla University 

BYOA as an Option

Scylla Cloud users on AWS have the option to provision Scylla Cloud EC2 resources directly into servers from their own AWS accounts. We call it “*Bring Your Own Account*” or BYOA. This option allows you to manage all your AWS resources under one account. It also enables you to take advantage of any pre-negotiated AWS rates or cloud credits and apply them towards Scylla Cloud compute resources. Plus, this may help you satisfy strict compliance requirements where sensitive data should remain within your accounts.

Upon selection of your account as a destination for Scylla Cloud deployment, Scylla Cloud provides a wizard that will walk you through all the steps to create a cloud policy and an IAM role for Scylla Cloud. See our documentation for more details. If you have any questions, you can always open a support request through the Scylla Cloud user interface.

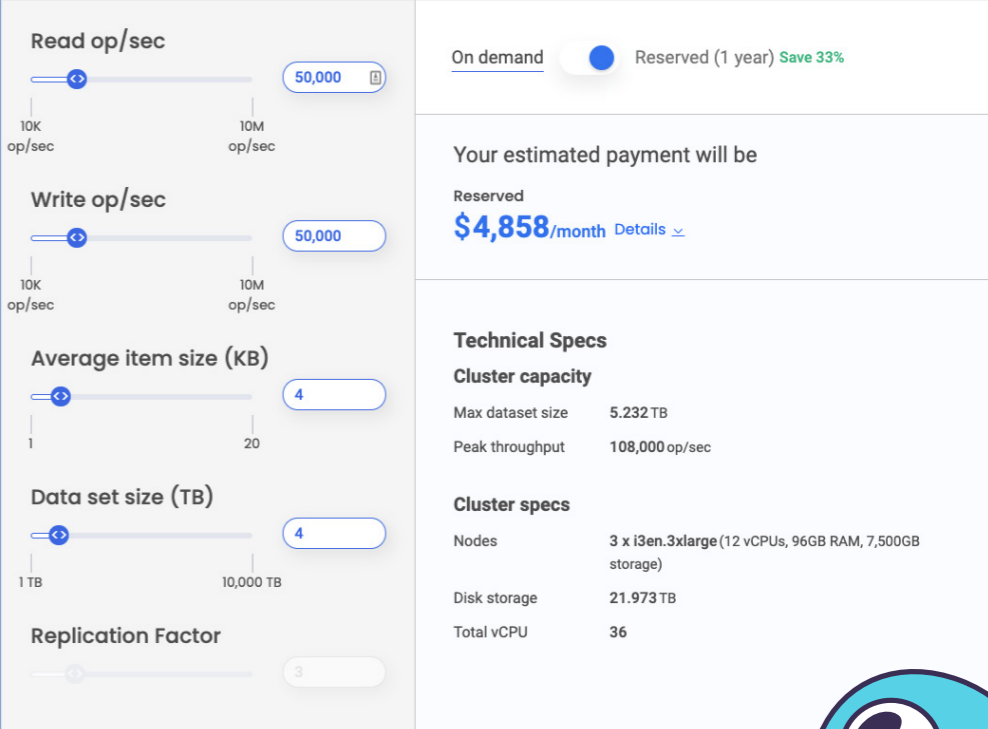
Capacity Planning

After you have settled on your choice of cloud provider and deployment options for Scylla resources, it's time to choose the right instance types to satisfy your workloads. We offer a [capacity planning calculator](#) for your convenience to help you properly size your cluster based on a few properties:

- Read and write throughput
- Average item size
- Projected data set size

The calculator takes those inputs and provides you with a suggested cluster specification, plus a cost estimate for on-demand and reserved capacity. Please note that the calculator cannot factor in more complex types of workloads or advanced features, such as compute or throughput intensive operations. If you have any questions about planning for your own workloads, please reach out to our Solution Architects team, using our [Contact Us](#) page.

To learn more about capacity planning for Scylla Cloud, please read this [blog](#) post.



The screenshot shows the Scylla Cloud capacity planning calculator interface. On the left, there are five sliders for inputting workload characteristics: Read op/sec (set to 50,000), Write op/sec (set to 50,000), Average item size (KB) (set to 4), Data set size (TB) (set to 4), and Replication Factor (set to 3). On the right, the calculator shows the selected option is 'Reserved (1 year)' with a 'Save 33%' indicator. Below this, it states 'Your estimated payment will be Reserved \$4,858/month' with a 'Details' link. The 'Technical Specs' section includes 'Cluster capacity' (Max dataset size: 5.232 TB, Peak throughput: 108,000 op/sec) and 'Cluster specs' (Nodes: 3 x i3en.3xlarge (12 vCPUs, 96GB RAM, 7,500GB storage), Disk storage: 21.973 TB, Total vCPU: 36).



<input type="radio"/>	i3en.2xlarge	8	64 GB	4.88 TB	25,600	\$2,274
<input checked="" type="radio"/>	i3en.3xlarge	12	96 GB	7.32 TB	25,600	\$3,411
<input type="radio"/>	i3en.6xlarge	24	192 GB	14.65 TB	25,600	\$6,822
<input type="radio"/>	i3en.12xlarge	48	384 GB	29.3 TB	51,200	\$13,643
<input type="radio"/>	i3en.24xlarge	96	768 GB	58.59 TB	102,400	\$27,286

Choose your Quantity

Replication Factor: 3 | Number of Nodes: 3 | 3 nodes x \$3,411/hour = \$10,233/hour

Instance Availability
AWS has limited availability for i3en server family. Creating a cluster on i3en depends on available EC2 resources.

Preview


US-EAST-1A | US-EAST-1B | US-EAST-1C

Details

Cluster Name: ⓘ

Allowed IPs: ADD ⓘ
Please enter at least one IP address

Enable VPC Peering
Allows you to set-up VPC peering with your application's VPC ([What is VPC peering?](#)).
This configuration is mandatory for a multi DC cluster.
Cluster Network
172.31.0.0/16 ⓘ



Deployment

Now use the cluster specifications to select instance types within your desired geographic region, the number of nodes you need, and the replication factor you want for your data. Scylla Cloud will automatically provision your cluster across multiple availability zones to ensure the high availability and resilience of your cluster.

You can learn more about Scylla's high availability design [here](#).

Don't forget to name your cluster! We take security seriously, so you have to provide a list of IP addresses that will be able to communicate with your cluster on Scylla Cloud.

We highly recommend you avoid routing your traffic over the open internet and instead directly connect your Virtual Private Cloud (VPC) to Scylla Cloud. Read how you can enable [VPC peering](#) on your cluster.



Connecting to Your Cluster

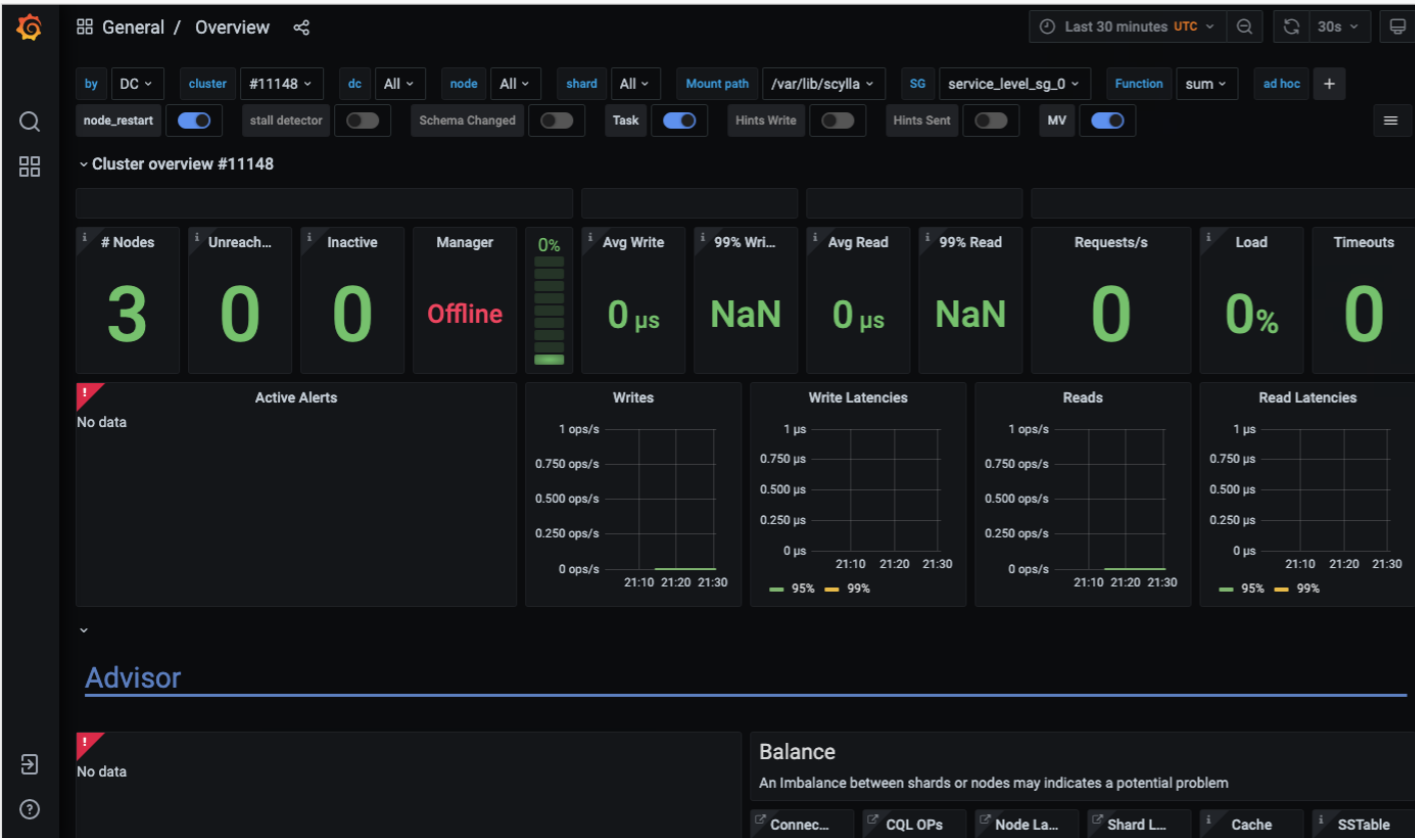
After your VPC peering is set up, provisioning is done, and your cluster is ready for use, you can immediately connect to your cluster. Scylla Cloud provides you with default credentials and detailed instructions on how to connect to your cluster with different clients and drivers.

Don't have a driver? Now's the time to download one. We have links to our CQL drivers right inside the Scylla Cloud user interface.

While we generate a very strong randomized password for you, you can take this opportunity to change the default password of "scylla" user for security purposes!



The screenshot displays the Scylla Cloud user interface. At the top, there are tabs for 'General', 'Connect', 'Monitoring', and 'Actions'. On the right, status indicators show 'Op/Sec 0.00', 'Storage 1%', and 'Load 0%'. The 'Connect' tab is active, showing a sidebar with 'Instructions' and a list of client languages: Cqlsh, Java, Go, Python, Ruby, JavaScript, and Other. The main content area is titled 'Instructions' and shows the 'Data Center: AWS_US_EAST_1'. It lists the 'NAME' as 'AWS_US_EAST_1' and provides instructions on how to use this name. The 'PUBLIC' IP addresses are listed as '35.153.81.200', '54.209.10.27', and '184.73.96.98'. Below this, the 'Default Credentials for Password Authentication' section states that the cluster has Password Authentication enabled and provides the default 'Username: scylla' and 'Password: *****'. A note recommends changing the default password for security, and a code block shows the CQL command: `ALTER ROLE scylla WITH PASSWORD = 'new_password';`



Monitoring Your Cluster

Scylla Cloud comes integrated with [Scylla Monitoring Stack](#) for greater visibility into your cluster's health and performance. While our team of engineers is taking care of all cluster management and health monitoring tasks, we also give you a real-time dashboard that provides a transparent view of your cluster health, lets you explore CQL metrics, and proactively identifies potential issues in the [Advisor section](#).





Testing the Limits

While we claim Scylla is the most performant NoSQL database, it's best to see the objective evidence for yourself. How about putting your cluster to a stress test to go [from zero to 2M OPS just in 5 minutes?](#)

Note that due to your own cluster's specification, you may see more or less operations per second compared to the test setup.

Alternator

It's important to note that this guide was written with Scylla's CQL interface in mind. However, if you are using our Amazon DynamoDB-compatible API, known as [Project Alternator](#), you can check out the documentation [here](#).



Conclusion

Thank you again for trusting your organization's data and daily operations to Scylla Cloud. We strive to improve the product every day, so if there is a feature you don't see, or have questions we haven't already answered in this guide, please let us know!

Feel free to ask a question on our [Slack channel](#)

Or if your question is more of a confidential nature, [contact us](#) privately.

ABOUT SCYLLADB

Scylla is the real-time big data database. API-compatible with Apache Cassandra and Amazon DynamoDB, Scylla embraces a shared-nothing approach that increases throughput and storage capacity as much as 10X.

Comcast, Discord, Disney+ Hotstar, Grab, Medium, Starbucks, Ola Cabs, Samsung, IBM, Investing.com and [many more leading companies](#) have adopted Scylla to realize order-of-magnitude performance improvements and reduce hardware costs. Scylla's database is available as an open source project, an enterprise edition, and a fully managed database as a service. ScyllaDB was founded by the team responsible for the KVM hypervisor. For more information: [ScyllaDB.com](https://scylladb.com)



United States Headquarters
2445 Faber Place, Suite 200
Palo Alto, CA 94303 U.S.A.
Email: info@scylladb.com

Israel Headquarters
11 Galgalei Haplada
Herzeliya, Israel

