MemSQL on the AWS Cloud

Quick Start Reference Deployment

June 2019

MemSQL Inc.
AWS Quick Start team

Visit our GitHub repository for source files and to post feedback, report bugs, or submit feature ideas for this Quick Start.

Contents

This Quick Start was created by MemSQL in collaboration with Amazon Web Services (AWS).

Quick Starts are automated reference deployments that use AWS CloudFormation templates to deploy key technologies on AWS, following AWS best practices.

Overview

This Quick Start reference deployment guide provides step-by-step instructions for deploying MemSQL on the AWS Cloud.

This Quick Start is for anyone who would like to evaluate or just play around with MemSQL, including database administrators, data architects, data scientists, and data analysts.

MemSQL on AWS

MemSQL is a distributed, highly scalable structured query language (SQL) database that can run anywhere (in the cloud and on premises). MemSQL delivers maximum performance for transactional and analytical workloads with familiar relational data structures.

MemSQL has the following capabilities:

- **High-speed ingest:** Fast bulk loading or stream ingestion with real-time data pipelines
- **Memory-optimized tables**: Ultra-low latency for scalable transactions and analytics
- **Disk-optimized tables:** Fast petabyte scale analytics with column-store compression

MemSQL ingests data continuously to perform operational analytics on billions of rows of data in relational SQL, JavaScript Object Notation (JSON), geospatial, and full-text search formats. This Quick Start also includes MemSQL Studio, a visual user interface that is supported on Google Chrome and Mozilla Firefox browsers.

Together, MemSQL and AWS provide a compelling platform for building real-time applications. MemSQL can handle both database workloads and data warehouse workloads, meeting transactional and analytical requirements. AWS provides useful services such as Elastic Load Balancing and Amazon Elastic Block Store (Amazon EBS), as well as a comprehensive infrastructure for standing up powerful solutions in the cloud. For details, see the MemSQL blog post Using MemSQL within the AWS Ecosystem.

Costs and licenses

You are responsible for the cost of the AWS services used while running this Quick Start reference deployment. There is no additional cost for using the Quick Start.

The AWS CloudFormation template for this Quick Start includes configuration parameters that you can customize. Some of these settings, such as Amazon Elastic Compute Cloud (Amazon EC2) instance type, will affect the cost of deployment. For cost estimates, see the pricing pages for each AWS service you will be using. Prices are subject to change.

> **Tip** After you deploy the Quick Start, we recommend that you enable the AWS Cost and Usage Report to track costs associated with the Quick Start. This report delivers billing metrics to an Amazon Simple Storage Service (Amazon S3) bucket in your account. It provides cost estimates based on usage throughout each month, and finalizes the data at the end of the month. For more information about the report, see the AWS documentation.

There are no licensing costs for using MemSQL up to 128 GB of RAM across four nodes. To request a **Free Cloud Trial**, visit the MemSQL website. If you need a cluster with more RAM, you can create an account and request an **Enterprise License**. Fill out the form, and you will receive a license key in email.

Architecture

Deploying this Quick Start for a new virtual private cloud (VPC) with **default parameters** builds the following MemSQL environment in the AWS Cloud. **Note that only one cluster will be created** (shown on the left).

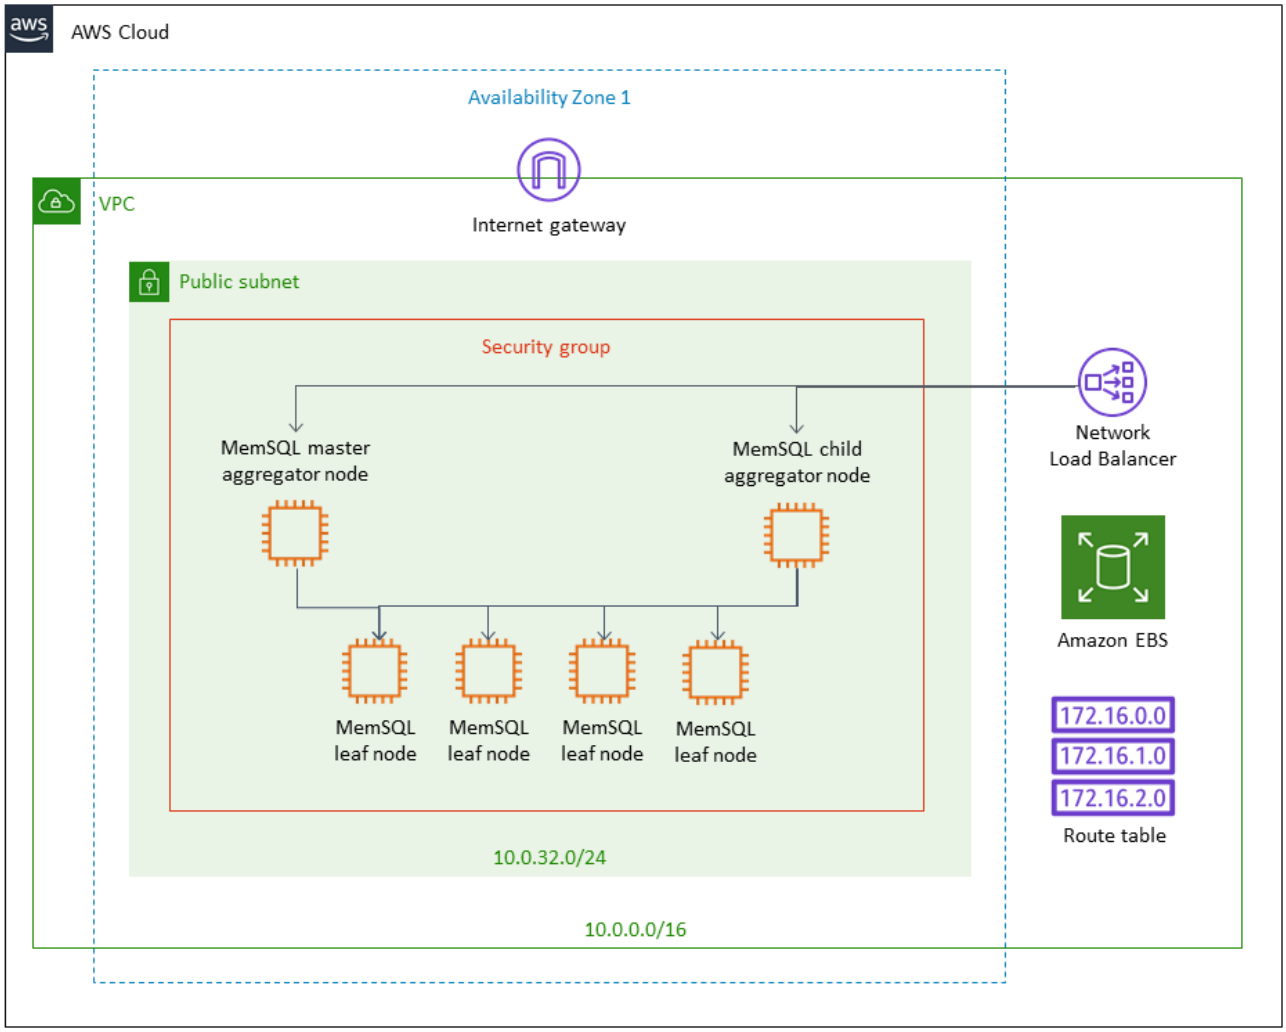

Figure 1: Quick Start architecture for MemSQL on AWS

The Quick Start sets up the following:

- A high-durability architecture in a single Availability Zone.*
- A VPC configured with a public subnet according to AWS best practices, to provide you with your own virtual network on AWS.*
- An internet gateway to allow access to the internet. This gateway is used by hosts to send and receive traffic.*
- In the public subnet, one master aggregator node, one child aggregator node, and leaf nodes; the diagram shows four leaf nodes, the default is two, and the maximum is six.
- A network load balancer that controls access to the MemSQL cluster's aggregator nodes.
- One attached Amazon EBS volume in each Amazon EC2 instance.

You can connect by using Secure Shell (SSH) through port 22 to log in to the EC2 instances for any maintenance activities as long as they are within the VPC.

In the MemSQL cluster, the aggregators and leaves communicate over the private interfaces.

* If you are deploying into an existing VPC, the template skips the components marked by asterisks. If you are deploying into an existing VPC, you also have the option of creating another cluster using the same MemSQL template in another Availability Zone and then setting up MemSQL native replication by following instructions in the MemSQL documentation.

Planning the deployment

Specialized knowledge

Before you deploy this Quick Start, we recommend that you become familiar with the following AWS services. (If you are new to AWS, see Getting Started with AWS.)

- Amazon EC2
- Amazon EBS
- Amazon VPC
- AWS CloudFormation

This deployment guide also requires a moderate level of familiarity with AWS services. If you're new to AWS, visit the Getting Started Resource Center and the AWS Training and Certification website for materials and programs that can help you develop the skills to design, deploy, and operate your infrastructure and applications on the AWS Cloud.

In addition, we recommend becoming famliar with MemSQL.

AWS account

If you don't already have an AWS account, create one at https://aws.amazon.com by following the on-screen instructions. Part of the sign-up process involves receiving a phone call and entering a PIN using the phone keypad.

Your AWS account is automatically signed up for all AWS services. You are charged only for the services you use.

Technical requirements

License information for MemSQL is mentioned in the previous section.

Before you launch the Quick Start, your account must be configured as specified in the following table. Otherwise, deployment might fail.

Resources	If necessary, request service limit increases for the following resources. You might need to do this if you already have an existing deployment that uses these resources, and you think you might exceed the default limits with this deployment. For default limits, see the AWS documentation. AWS Trusted Advisor offers a service limits check that displays your usage and limits for some aspects of some services. **Resource** / **This deployment uses** **VPCs** / 1 **Network Load Balancers** / 1 **m4.2xlarge instances** / 3-9
Key pair	Make sure that at least one Amazon EC2 key pair exists in your AWS account in the Region where you are planning to deploy the Quick Start. Make note of the key pair name. You'll be prompted for this information during deployment. To create a key pair, follow the instructions in the AWS documentation. If you're deploying the Quick Start for testing or proof-of-concept purposes, we recommend that you create a new key pair instead of specifying a key pair that's already being used by a production instance.
IAM permissions	To deploy the Quick Start, you must log in to the AWS Management Console with AWS Identity and Access Management (IAM) permissions for the resources and actions the templates will deploy. The *AdministratorAccess* managed policy within IAM provides sufficient permissions, although your organization may choose to use a custom policy with more restrictions.
S3 buckets	Unique S3 bucket names are automatically generated based on the account number and AWS Region. If you delete a stack, **the logging buckets are not deleted** (to support security review). If you plan to re-deploy this Quick Start in the same Region, you must first manually delete the S3 buckets that were created during the previous deployment; **otherwise, the re-deployment will fail**.

Deployment options

This Quick Start provides two options for creating a MemSQL cluster.

- **Deploy MemSQL into a new VPC (end-to-end deployment)**. This option builds a new AWS environment consisting of the VPC, subnets, NAT gateways, security groups, bastion hosts, and other infrastructure components, and then deploys MemSQL into this new VPC.
- **Deploy MemSQL into an existing VPC**. This option provisions MemSQL in your existing AWS infrastructure.

Deployment steps

Step 1. Sign in to your AWS account

1. Sign in to your AWS account at https://aws.amazon.com with an IAM user role that has the necessary permissions. For details, see Planning the deployment earlier in this guide.
2. Make sure that your AWS account is configured correctly, as discussed in the Technical requirements section.

Step 2. Get a license for MemSQL

MemSQL is free up to 128 GB of Memory across all the nodes combined with no time limit. To try MemSQL absolutely free for 8 hours, choose the **Free Cloud Trial** option on the MemSQL website.

If your usage or capacity requirements will exceed this limitation, visit the MemSQL portal to create a new account and request an Enterprise License for MemSQL. You will need to provide the license key during the deployment.

Step 3. Launch the Quick Start

> **Notes** The instructions in this section reflect the older version of the AWS CloudFormation console. If you're using the redesigned console, some of the user interface elements might be different.
>
> You are responsible for the cost of the AWS services used while running this Quick Start reference deployment. There is no additional cost for using this Quick Start. For full details, see the pricing pages for each AWS service you will be using in this Quick Start. Prices are subject to change.

1. Sign in to your AWS account, and choose one of the following options to launch the AWS CloudFormation template. For help choosing an option, see deployment options earlier in this guide.

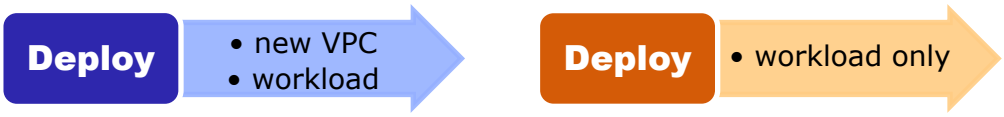

Deploy MemSQL into a new VPC on AWS

Deploy MemSQL into an existing VPC on AWS

> **Important** If you're deploying MemSQL into an existing VPC, make sure that your VPC has at least one public subnet, with the domain name option configured in the DHCP options as explained in the Amazon VPC documentation. You will be prompted for your VPC settings when you launch the Quick Start.

   Depending on the number of instances chosen (aggregators and leaves), each deployment may take 5-15 minutes to complete.

2. Check the AWS Region that's displayed in the upper-right corner of the navigation bar and change it if necessary. This is where the network infrastructure for MemSQL will be built. The template is launched in the US East (N. Virginia) Region by default.
3. On the **Select Template** page, keep the default setting for the template URL, and then choose **Next**.
4. On the **Specify Details** page, change the stack name if needed. Review the parameters for the template. Provide values for the parameters that require input. For all other parameters, review the default settings and customize them as necessary.

   In the following tables, parameters are listed by category and described separately for the two deployment options:

   - Parameters for deploying MemSQL into a new VPC
   - Parameters for deploying MemSQL into an existing VPC

Option 1: Parameters for deploying MemSQL into a new VPC

View template

Basic Setup (Required):

Parameter label (name)	Default	Description
Key pair name (KeyPairName)	*Requires input*	The EC2 key pair to allow SSH access to the nodes.
License (License)	*Requires input*	A license you received from the MemSQL Customer Portal.
MemSQL Password (RootPassword)	*Requires input*	A MemSQL root password that will be set for all the nodes in your cluster.

Advanced Configuration:

Parameter label (name)	Default	Description
Aggregator Instance Type (AggInstanceType)	m4.2xlarge	The EC2 instance type for master and child aggregators. For information on instance types, see the AWS documentation.
Number of Child Aggregators (NumAggregators)	1	The number of child aggregators in the cluster (between 0 and 3). Separately, the master aggregator will automatically be created.
Leaf Instance Type (LeafInstanceType)	m4.2xlarge	The EC2 instance type for leaves.
Number of Leaves (NumLeaves)	2	The number of leaves in the cluster (between 1 and 6).
Enable High Availability (EnableHighAvailability)	true	If **true**, specify an even number of leaves, as an extra leaf will not be utilized. To learn more about High Availability, see the MemSQL documentation.
Availability Zone (AvailabilityZone)	*Requires input*	The Availability Zone to use for the subnet in the VPC.
CIDR Range for Remote Access (RemoteAccessCIDR)	0.0.0.0/0	The IP CIDR range that is allowed to access the nodes (including SSH access, connecting to MemSQL on port 3306, viewing MemSQL Studio).

Network Configuration:

Parameter label (name)	Default	Description
VPC CIDR (VPCCIDR)	10.0.0.0/16	The CIDR block for the VPC.
Public subnet CIDR (PublicSubnetCIDR)	10.0.32.0/24	The CIDR range used in the public subnet located in Availability Zone 1.

AWS Quick Start Configuration:

Parameter label (name)	Default	Description
Quick Start S3 bucket name (`QSS3BucketName`)	aws-quickstart	The S3 bucket you have created for your copy of Quick Start assets, if you decide to customize or extend the Quick Start for your own use. The bucket name can include numbers, lowercase letters, uppercase letters, and hyphens, but should not start or end with a hyphen.
Quick Start S3 key prefix (`QSS3KeyPrefix`)	quickstart-memsql/	The S3 key name prefix used to simulate a folder for your copy of Quick Start assets, if you decide to customize or extend the Quick Start for your own use. This prefix can include numbers, lowercase letters, uppercase letters, hyphens, and forward slashes.

OPTION 2: PARAMETERS FOR DEPLOYING MEMSQL INTO AN EXISTING VPC

View template

Basic Setup (Required):

Parameter label (name)	Default	Description
Key Name (`KeyPairName`)	*Requires input*	The EC2 key pair to allow SSH access to the nodes.
License (`License`)	*Requires input*	A license you received from the MemSQL Customer Portal.
MemSQL Password (`RootPassword`)	*Requires input*	A MemSQL root password that will be set for all the nodes in your cluster.

Advanced Configuration:

Parameter label (name)	Default	Description
Aggregator Instance Type (`AggInstanceType`)	m4.2xlarge	The EC2 instance type for master and child aggregators. For information on instance types, see the AWS documentation.
Number of Child Aggregators (`NumAggregators`)	1	The number of child aggregators in the cluster (between 0 and 3). Separately, the master aggregator will automatically be created.
Leaf Instance Type (`LeafInstanceType`)	m4.2xlarge	The EC2 instance type for leaves.
Number of Leaves (`NumLeaves`)	2	The number of leaves in the cluster (between 1 and 6).

Parameter label (name)	Default	Description
Enable High Availability (`EnableHighAvailability`)	true	If **true**, specify an even number of leaves, as an extra leaf will not be utilized. To learn more about High Availability, see the MemSQL documentation.

VPC Configuration:

Parameter label (name)	Default	Description
VPC ID (`VPCID`)	*Requires input*	The ID of the VPC to launch into.
Subnet ID (`PublicSubnetID`)	*Requires input*	The ID of the public subnet to launch into. The subnet must be in the specified VPC.
Security Group IDs (`SecurityGroups`)	*Requires input*	A comma-delimited list of security group IDs for each instance. The security groups must be in the specified VPC.

5. On the **Options** page, you can specify tags (key-value pairs) for resources in your stack and set advanced options. When you're done, choose **Next**.
6. On the **Review** page, review and confirm the template settings. Under **Capabilities**, select the two check boxes to acknowledge that the template will create IAM resources and that it might require the capability to auto-expand macros.
7. Choose **Create** to deploy the stack.
8. Monitor the status of the stack. When the status is **CREATE_COMPLETE**, the MemSQL cluster is ready.
9. Use the URLs displayed in the **Outputs** tab for the stack to view the resources that were created.

Step 4. Post-launch steps

For instructions on using MemSQL, see the following tutorials:

- Connect to MemSQL
- Run queries in MemSQL
- Load data into MemSQL

Troubleshooting

Q. I encountered a CREATE_FAILED error when I launched the Quick Start.

A. If AWS CloudFormation fails to create the stack, we recommend that you relaunch the template with **Rollback on failure** set to **No**. (This setting is under **Advanced** in the AWS CloudFormation console, **Options** page.) With this setting, the stack's state will be retained and the instance will be left running, so you can troubleshoot the issue. (For Windows, look at the log files in `%ProgramFiles%\Amazon\EC2ConfigService` and `C:\cfn\log`.)

> **Important** When you set **Rollback on failure** to **No**, you will continue to incur AWS charges for this stack. Please make sure to delete the stack when you finish troubleshooting.

For additional information, see Troubleshooting AWS CloudFormation on the AWS website.

Q. I encountered a size limitation error when I deployed the AWS CloudFormation templates.

A. We recommend that you launch the Quick Start templates from the links in this guide or from another S3 bucket. If you deploy the templates from a local copy on your computer or from a non-S3 location, you might encounter template size limitations when you create the stack. For more information about AWS CloudFormation limits, see the AWS documentation.

Send us feedback

To post feedback, submit feature ideas, or report bugs, use the **Issues** section of the GitHub repository for this Quick Start. If you'd like to submit code, please review the Quick Start Contributor's Guide.

Additional resources

AWS resources

- Getting Started Resource Center
- AWS General Reference
- AWS Glossary

AWS services

- AWS CloudFormation

- Amazon EBS
- Amazon EC2
- IAM
- Amazon VPC

MemSQL documentation

- MemSQL documentation
- MemSQL best practices
- MemSQL security
- MemSQL FAQ

Other Quick Start reference deployments

- AWS Quick Start home page

Document revisions

Date	Change	In sections
June 2019	Initial publication	—

Notices

This document is provided for informational purposes only. It represents AWS's current product offerings and practices as of the date of issue of this document, which are subject to change without notice. Customers are responsible for making their own independent assessment of the information in this document and any use of AWS's products or services, each of which is provided "as is" without warranty of any kind, whether express or implied. This document does not create any warranties, representations, contractual commitments, conditions or assurances from AWS, its affiliates, suppliers or licensors. The responsibilities and liabilities of AWS to its customers are controlled by AWS agreements, and this document is not part of, nor does it modify, any agreement between AWS and its customers.

The software included with this paper is licensed under the Apache License, Version 2.0 (the "License"). You may not use this file except in compliance with the License. A copy of the License is located at http://aws.amazon.com/apache2.0/ or in the "license" file accompanying this file. This code is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.