



Hyperledger Fabric Developer Guide

Amazon Managed Blockchain (AMB)



Amazon Managed Blockchain (AMB): Hyperledger Fabric Developer Guide

Copyright © 2026 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

Amazon's trademarks and trade dress may not be used in connection with any product or service that is not Amazon's, in any manner that is likely to cause confusion among customers, or in any manner that disparages or discredits Amazon. All other trademarks not owned by Amazon are the property of their respective owners, who may or may not be affiliated with, connected to, or sponsored by Amazon.

Table of Contents

What Is Amazon Managed Blockchain (AMB) Hyperledger Fabric	1
How to Get Started with AMB Access Hyperledger Fabric	1
Key Concepts	3
Networks and Editions	3
Networks, Proposals, and Members	6
Peer Nodes	6
Connecting to Resources	7
Getting Started	9
Prerequisites and Considerations	9
An AWS account	10
A Linux Client (EC2 Instance)	10
A VPC	11
Permissions to Create an Interface VPC Endpoint	11
EC2 Security Groups That Allow Communication on Required Ports	11
Additional Considerations	13
Step 1: Create the Network and First Member	13
Step 2: Create an Endpoint	16
Step 3: Create a Peer Node	17
Step 4: Set Up a Client	18
4.1: Install Packages	19
4.2: Set Up the Fabric CA Client	22
4.3: Clone Samples	24
4.4: Start the Hyperledger Fabric CLI	24
Step 5: Enroll the Member Admin	26
5.1: Create the Certificate File	26
5.2 Enroll the Admin	27
5.3: Copy Certificates	29
Step 6: Create a Channel	29
6.1: Create configtx	30
6.2: Set an Environment Variable for the Orderer	35
6.3: Create the Channel	36
6.4: Join Peer to Channel	36
Step 7: Run Chaincode	36
7.1: Install Vendor Dependencies	37

7.2: Create Chaincode Package	37
7.3: Install Package	37
7.4: Verify Package	37
7.5: Approve Chaincode	38
7.6: Check Commit Readiness	38
7.7: Commit Chaincode	38
7.8: Verify Chaincode	38
7.9: Initialize Chaincode	39
7.10: Query Chaincode	39
7.11: Invoke Chaincode	39
Step 8: Invite a Member and Create a Multi-Member Channel	40
8.1: Create an Invitation Proposal	41
8.2: Vote Yes on the Proposal	42
8.3: Create the New Member	42
8.4: Share Artifacts	44
8.5: Create Artifacts for the MSP	45
8.6: Create configtx	46
8.7 Create the Channel	52
8.8: Get the Genesis Block	52
8.9: Join Peer Nodes to the Channel	53
8.10: Install Chaincode	53
8.11: Query Chaincode	54
8.12: Invoke Chaincode	54
Create a Network	56
Create a Hyperledger Fabric Network	56
Delete a Network	59
Invite or Remove Members	60
Create an Invitation Proposal	60
Create a Removal Proposal	62
Delete a Member in Your AWS Account	63
Accept an Invitation and Create a Member	64
Work with Invitations	65
Create a Member	66
Create an Interface VPC Endpoint	69
Work with Peer Nodes	71
Create a Peer Node	72

View Peer Node Properties	73
Use Peer Node Metrics	75
Viewing Peer Node Metrics	77
Work with Proposals	79
.....	79
View Proposals	82
Vote on a Proposal	84
Create an Invitation Proposal	85
Create a Removal Proposal	86
Automating with CloudWatch Events	87
Example AMB Access Events	88
Work with Hyperledger Fabric	90
Create an Admin	90
Registering an Admin	92
Enrolling an Admin	92
Copying the Admin Certificate	93
Work with Channels	93
Create a Channel	94
Add an Anchor Peer to a Channel	106
Prerequisites and Assumptions	106
Adding a Peer as an Anchor Peer	107
Develop Chaincode	110
Considerations and Limitations When Developing Chaincode for AMB Access	110
Private Data Collections	112
Develop Java Chaincode	113
API Reference	125
Security	126
Data Protection	126
Data Encryption	127
Encryption at Rest	127
Encryption in Transit	135
Authentication and Access Control	135
Identity and Access Management	136
Configuring Security Groups	171
Tagging resources	176
Create and add tags for AMB Access Hyperledger Fabric resources	176

Tag naming and usage conventions	177
Working with tags	177
Add or remove tags	178
Monitoring	181
Considerations and Limitations	182
Enabling and Disabling Logs	182
Enabling and Disabling Peer Node and Chaincode Logs	182
Enabling and Disabling CA Logs	183
Working with Logged Events in the AMB Access Console	183
Searching (Filtering) Logged Events	184
Downloading Logged Events	184
Viewing Different Chaincode Logs	185
Identifying Logs in CloudWatch Logs	186
CloudTrail logs	187
Managed Blockchain information in CloudTrail	187
Understanding Managed Blockchain log file entries	188
Document History	190
AWS Glossary	193

What Is Amazon Managed Blockchain (AMB) Hyperledger Fabric?

Amazon Managed Blockchain (AMB) Access provides you with public blockchain nodes for Ethereum and Bitcoin, and you can also create private blockchain networks with the Hyperledger Fabric framework. Choose from various methods to engage with public blockchains, including fully managed, single-tenant (dedicated), and serverless multi-tenant API operations to public blockchain nodes. For use cases where access controls are important, you can choose from fully managed private blockchain networks. Standardized API operations give you instant scalability on a fully managed, resilient infrastructure, so you can build blockchain applications.

AMB Access gives you two distinct types of blockchain infrastructure services: multi-tenant blockchain network access API operations and dedicated blockchain nodes and networks. With dedicated blockchain infrastructure, you can create and use public Ethereum blockchain nodes and private Hyperledger Fabric blockchain networks for your own use. Multi-tenant, API-based offerings, however, such as AMB Access Bitcoin, are composed of a fleet of Bitcoin nodes behind an API layer where the underlying blockchain node infrastructure is shared among customers.

You can deploy a private, permissioned blockchain with the Hyperledger Fabric framework in minutes without complex configuration tasks. Once your network is created, you can launch one or more blockchain peer nodes for high availability using the AWS Management Console. Whether your use case demands a large consortium of members to join the network or you simply wish to create a development network within your organization, AMB simplifies the process. You can invite other organizations to join your blockchain network through their AWS accounts, or you can create additional members (organizations) in your AWS account to build a simulated network for testing.

This guide covers the fundamentals of creating and working with a Hyperledger Fabric blockchain network using AMB Access. For information about working with AMB Access Ethereum, see [Amazon Managed Blockchain \(AMB\) Access Ethereum Developer Guide](#).

How to Get Started with AMB Access Hyperledger Fabric

We recommend the following resources to get started with Hyperledger Fabric networks and chaincode on AMB Access:

- [Key Concepts: Amazon Managed Blockchain \(AMB\) Networks, Members, and Peer Nodes](#)

This overview helps you understand the fundamental building blocks of a Hyperledger Fabric network on AMB Access. It also tells you how to identify and communicate with network resources.

- [Get Started Creating a Hyperledger Fabric Blockchain Network Using Amazon Managed Blockchain \(AMB\)](#)

Use this tutorial to create your first Hyperledger Fabric network, set up a Hyperledger Fabric client on EC2, and use the open-source Hyperledger Fabric peer CLI to query and update the ledger. You then invite another member to the network. The member can be from a different AWS account, or you can invite a new member in your own account to simulate a multi-account network. The new member then queries and updates the ledger.

- [Hyperledger Fabric Documentation \(v2.2\)](#)

The open-source documentation for Hyperledger Fabric is a starting point for key concepts and the architecture of the Hyperledger Fabric blockchain network that you build using AMB Access. As you develop your blockchain application, you can reference this document for key tasks and code samples. Use the documentation version that corresponds to the version of Hyperledger Fabric that you use.

Key Concepts: Amazon Managed Blockchain (AMB) Networks, Members, and Peer Nodes

A blockchain network is a peer-to-peer network running a decentralized blockchain framework. A Hyperledger Fabric network on Amazon Managed Blockchain (AMB) includes one or more *members*. Members are unique identities in the network. For example, a member might be an organization in a consortium of banks. A single AWS account might have multiple members. Each member runs one or more Hyperledger Fabric *peer nodes*. The peer nodes run chaincode, endorse transactions, and store a local copy of ledger.

Amazon Managed Blockchain (AMB) creates and manages these components for each member in a network. AMB Access also creates components that all network members share, such as the Hyperledger Fabric ordering service and the general networking configuration.

Note

What we call *members* in a Hyperledger Fabric network on AMB Access is very similar to what Hyperledger Fabric calls *organizations*.

AMB Access Hyperledger Fabric Networks and Editions

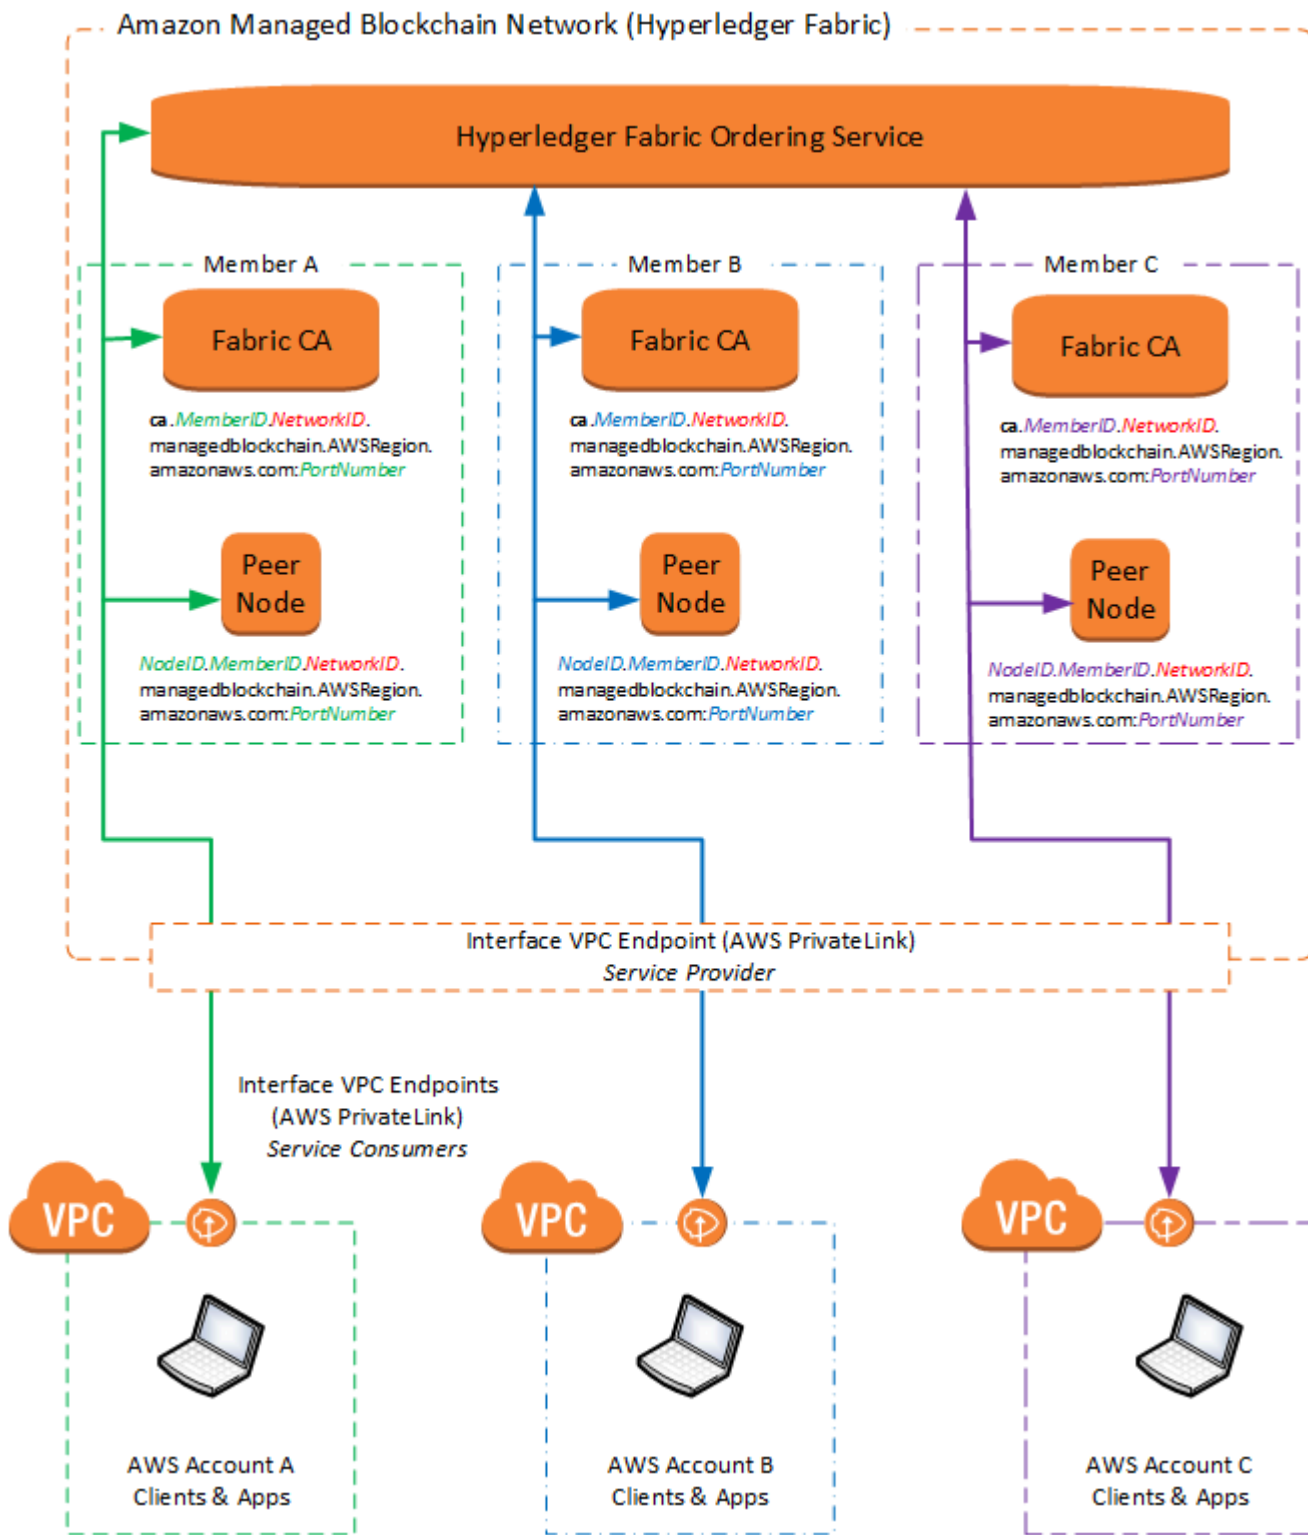
When creating a Hyperledger Fabric network, the creator chooses the framework version and the *edition* of Amazon Managed Blockchain (AMB) to use. The edition determines the capacity and capabilities of the network as a whole.

The creator also must create the first network member. Additional members are added through a proposal and voting process. There is no charge for the network itself, but each member pays an hourly rate (billed per second) for their network membership. Charges vary depending on the edition of the network. Each member also pays for peer nodes, peer node storage, and the amount of data that the member writes to the network. For more information about available editions and their attributes, see [Amazon Managed Blockchain \(AMB\) Pricing](#). For more information about the number of networks that each AWS account can create and join, see [Amazon Managed Blockchain \(AMB\) Limits](#) in the *AWS General Reference*.

A Hyperledger Fabric network on AMB Access remains active as long as there are members. The network is deleted only when the last member deletes itself from the network. No member or AWS

account, even the creator's AWS account, can delete the network until they are the last member and delete themselves.

The following diagram shows the basic components of a Hyperledger Fabric blockchain running on AMB Access.



Inviting and Removing Members

An AWS account initially creates a Hyperledger Fabric network on AMB Access, but the network is not owned by that AWS account or any other AWS account. The network is decentralized, so changes to the network are made by consensus.

To make changes to the network, members make *proposals* that all other members in the network vote on. For another AWS account to join the network, for example, an existing member creates a proposal to invite the account. Other members then vote Yes or No on the proposal. If the proposal is approved, an invitation is sent to the AWS account. The account then accepts the invitation and creates a member to join the network. A similar proposal process is required to remove a member in a different AWS account. A principal in an AWS account with sufficient permissions can remove a member that the account owns at any time by deleting that member directly, without submitting a proposal.

The network creator also defines a *voting policy* for the network during creation. The voting policy determines the basic rules for all proposal voting on the network. The voting policy includes the percentage of votes required to pass the proposal, and the duration before the vote expires.

Peer Nodes

When a member joins the network, one of the first things they must do is create at least one *peer node* in the membership.

Blockchain networks contain a distributed, cryptographically secure ledger that maintains the history of transactions in the network that is immutable—it can't be changed after-the-fact. Each peer node also holds the global state of the network for the channels in which they participate. The global state is updated with each new transaction. When a new peer node in a channel comes online, it fetches the global state and ledger from other peers. Even if there are no other peer nodes on a network, as long as a member exists, ledger data can be restored to a new peer node.

Peer nodes also interact to create and endorse the transactions that are proposed on the network to update the ledger. The members define the rules in the endorsement process based on their business logic. In this way, every member can conduct transactions as allowed by the business logic and independently verify the transaction history without a centralized authority.

Note

Limit transactions to less than 4 MB. Transactions greater than 4 MB result in an error.

To configure Hyperledger Fabric applications on peer nodes and to interact with other network resources, members use a client configured with open-source Hyperledger Fabric tools such as a CLI or SDK. The applications and tools that you choose and your client setup depend on your preferred development environment. For example, in the [Getting Started](#) tutorial, you configure an Amazon EC2 instance in a VPC with open-source Hyperledger Fabric CLI tools.

Identifying AMB Access Resources and Connecting from a Client

Because a Hyperledger Fabric blockchain network is decentralized, members must interact with each other's peer nodes and network-wide resources to make transactions, endorse transactions, verify members, and so on. When a network is created, AMB Access gives the network a unique ID. Similarly, when an AWS account creates a member on the network and peer nodes, AMB Access gives unique IDs to those resources.

Each network resource has a unique, addressable endpoint that AMB Access creates from these IDs. Other members of the network, Hyperledger Fabric chaincode, and other tools use these endpoints to identify and interact with resources on the network.

Resource endpoints for a Hyperledger Fabric network on AMB Access are in the following format:

```
ResourceID.MemberID.NetworkID.managedblockchain.AWSRegion.amazonaws.com:PortNumber
```

For example, to refer to a peer node with ID `nd-6EAJ5VA43JGGNPXOUZP7Y47E4Y`, owned by a member with ID `m-K46ICRRXJRCGRNNS4ES4XUUS5A`, in a Hyperledger Fabric network with ID `n-MWY63ZJZU5HGNCMBQER7IN6OIU`, you use the following peer node endpoint:

```
nd-6EAJ5VA43JGGNPXOUZP7Y47E4Y.m-K46ICRRXJRCGRNNS4ES4XUUS5A.n-MWY63ZJZU5HGNCMBQER7IN6OIU.managedblockchain.us-east-1.amazonaws.com:30003
```

The port that you use with an endpoint depends on the Hyperledger Fabric service that you are calling and your unique network setup. *AWSRegion* is the Region you are using. For a list of supported Regions, see [Amazon Managed Blockchain \(AMB\) Endpoints and Quotas](#) in the *Amazon Web Services General Reference*.

Within the Hyperledger Fabric network, access and authorization for each resource is governed by processes defined in the chaincode and network configurations such as Hyperledger Fabric channels. Outside the confines of the network—that is, from member's client applications and tools—AMB Access uses AWS PrivateLink to ensure that only network members can access required resources. In this way, each member has a private connection from a client in their VPC to the Hyperledger Fabric network on AMB Access. The interface VPC endpoint uses private DNS, so you must have a VPC in your account that is enabled for Private DNS. For more information, see [Create an Interface VPC Endpoint for Amazon Managed Blockchain \(AMB\) Hyperledger Fabric](#).

Get Started Creating a Hyperledger Fabric Blockchain Network Using Amazon Managed Blockchain (AMB)

This tutorial guides you through creating your first Hyperledger Fabric network using Amazon Managed Blockchain (AMB). It shows you how to set up the network and create a member in your AWS account, set up chaincode and a channel, and then invite members from other AWS accounts to join a channel. Instructions for invitees are also provided.

Steps

- [Prerequisites and Considerations](#)
- [Step 1: Create the Network and First Member](#)
- [Step 2: Create and Configure the Interface VPC Endpoint](#)
- [Step 3: Create a Peer Node in Your Membership](#)
- [Step 4: Create an Amazon EC2 Instance and Set Up the Hyperledger Fabric Client](#)
- [Step 5: Enroll an Administrative User](#)
- [Step 6: Create a Hyperledger Fabric Channel](#)
- [Step 7: Install and Run Chaincode](#)
- [Step 8: Invite Another AWS Account to be a Member and Create a Multi-Member Channel](#)

Prerequisites and Considerations

To complete this tutorial, you must have the resources listed in this section. Unless specifically stated otherwise, the requirements apply to both network creators and invited members.

Topics

- [An AWS account](#)
- [A Linux Client \(EC2 Instance\)](#)
- [A VPC](#)
- [Permissions to Create an Interface VPC Endpoint](#)
- [EC2 Security Groups That Allow Communication on Required Ports](#)
- [Additional Considerations](#)

An AWS account

Before you use AMB Access for the first time, you must sign up for an Amazon Web Services (AWS) account.

If you do not have an AWS account, complete the following steps to create one.

To sign up for an AWS account

1. Open <https://portal.aws.amazon.com/billing/signup>.
2. Follow the online instructions.

Part of the sign-up procedure involves receiving a phone call or text message and entering a verification code on the phone keypad.

When you sign up for an AWS account, an *AWS account root user* is created. The root user has access to all AWS services and resources in the account. As a security best practice, assign administrative access to a user, and use only the root user to perform [tasks that require root user access](#).

A Linux Client (EC2 Instance)

You must have a Linux computer with access to resources in the VPC to serve as your Hyperledger Fabric client. This computer must have version 1.16.149 or later of the AWS CLI installed. Earlier versions of the AWS CLI do not have the `managedblockchain` command. We recommend that you use the latest version of the AWS CLI available. For information about updating the AWS CLI, see [Update the AWS CLI version 2 on Linux](#) in the *AWS Command Line Interface User Guide*.

We recommend creating an Amazon Elastic Compute Cloud (Amazon EC2) instance in the same VPC and AWS Region as the VPC endpoint for the Hyperledger Fabric network on AMB Access. This is the setup that the tutorial uses. For instructions to set up a Hyperledger Fabric client using this configuration, see [Step 4: Create an Amazon EC2 Instance and Set Up the Hyperledger Fabric Client](#).

An AWS CloudFormation template to create a Hyperledger Fabric client is available in the [amazon-managed-blockchain-client-templates repository](#) on Github. For more information, see the [readme.md](#) in that repository. For more information about using CloudFormation, see [Getting Started](#) in the *AWS CloudFormation User Guide*.

A VPC

You must have a [VPC](#) with an IPv4 CIDR block, and the `enableDnsHostnames` and `enableDnsSupport` options must be set to `true`. If you will connect to the Hyperledger Fabric client using SSH, the VPC must have an internet gateway, and the security group configuration associated with the Hyperledger Framework client must allow inbound SSH access from your SSH client.

- For more information about creating a suitable network, see [Getting Started with IPv4 for Amazon VPC](#) tutorial in the *Amazon VPC User Guide*.
- For information about using SSH to connect to an Amazon EC2 Instance, see [Connecting to Your Linux Instance Using SSH](#) in the *Amazon EC2 User Guide*.
- For instructions about how to verify if DNS options are enabled, see [Using DNS with Your VPC](#) in the *Amazon VPC User Guide*.

Permissions to Create an Interface VPC Endpoint

The IAM principal (user) identity that you are using must have sufficient IAM permissions to create an interface VPC endpoint in your AWS account. For more information, see [Controlling Access - Creating and Managing VPC Endpoints](#) in the *Amazon VPC User Guide*.

EC2 Security Groups That Allow Communication on Required Ports

The EC2 security groups associated with the Hyperledger Fabric client Amazon EC2 instance and the Interface VPC Endpoint that you create during this tutorial must have rules that allow traffic between them for required Hyperledger Fabric services. EC2 security groups are restrictive by default, so you need to create security group rules that allow required access. In addition, a security group associated with the Hyperledger Fabric client Amazon EC2 instance must have an inbound rule that allows SSH traffic (Port 22) from trusted SSH clients.

For the purposes of simplicity in this tutorial, we recommend that you create an EC2 security group that you associate only with the Hyperledger Fabric client Amazon EC2 instance and the Interface VPC Endpoint. Then create an inbound rule that allows all traffic from within the security group. In addition, create another security group to associate with the Hyperledger Fabric client Amazon EC2 instance that allows inbound SSH traffic from trusted clients.

⚠ Important

This security group configuration is recommended for this tutorial only. Carefully consider security group settings for your desired security posture. For information about the minimum required rules, see [Configuring Security Groups for Amazon Managed Blockchain \(AMB\) Hyperledger Fabric](#).

To create a security group that allows traffic between the Hyperledger Fabric client and the interface VPC endpoint for use in this tutorial

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. Choose **Security groups** in the navigation pane, and then choose **Create security group**.
3. Enter a **Security group name** and **Description** for the security group that helps you find it. For example, **HFClientAndEndpoint**.
4. Make sure that the VPC you select is the default VPC for your account. This is the VPC in which Hyperledger Fabric network resources and the interface VPC endpoint are created.
5. Choose **Create**.
6. Select the security group that you just created from the list, choose **Inbound**, and then choose **Edit**.
7. Under **Type**, select **All traffic** from the list.
8. Under **Source**, leave **Custom** selected, and then begin typing the name or ID of this same security group—for example, **HFClientAndEndpoint**—and then select the security group so that its ID appears under **Source**.
9. Choose **Save**.

You reference this security group later in this tutorial in [Step 2: Create and Configure the Interface VPC Endpoint](#) and [Step 4: Create an Amazon EC2 Instance and Set Up the Hyperledger Fabric Client](#).

To create a security group for the Hyperledger Fabric client that allows inbound SSH connections from the computer that you are working with

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. Choose **Security groups** in the navigation pane, and then choose **Create security group**.

3. Enter a **Security group name** and **Description** for the security group that helps you find it. For example, **HFClientSSH**.
4. Make sure that the VPC you select is the same VPC that you will select for the interface VPC endpoint.
5. Choose **Inbound**, and then choose **Add rule**.
6. Under **Type**, select **SSH** from the list.
7. Under **Source**, select **My IP**. This adds the detected IP address of your current computer. Optionally, you can create additional rules for SSH connections from additional IP addresses or sources if required.
8. Choose **Create**.

You will reference this security group later in this tutorial in [Step 4: Create an Amazon EC2 Instance and Set Up the Hyperledger Fabric Client](#).

Additional Considerations

- All commands in the tutorial assume that you are using an Amazon EC2 instance with an Amazon Linux AMI. Unless noted otherwise, instructions also assume that you are running commands in the default home directory (`/home/ec2-user`). If you have a different configuration, modify instructions to fit your home directory as necessary.
- Hyperledger Fabric 2.2 requires that a channel ID contain only lowercase ASCII alphanumeric characters, dots (`.`), and dashes (`-`). It must start with a letter, and must be fewer than 250 characters.

Step 1: Create the Network and First Member

When you create the network, you specify the following parameters along with basic information such as names and descriptions:

- The open-source framework and version. This tutorial uses Hyperledger Fabric version 2.2.
- The voting policy for proposals on the network. For more information, see [Work with Proposals for a Hyperledger Fabric Network on Amazon Managed Blockchain \(AMB\)](#).
- The first member of the network, including the administrative user and administrative password that are used to authenticate to the member's certificate authority (CA).

⚠ Important

Each member that is created accrues charges according to the membership rate for the network. For more information, see [Amazon Managed Blockchain \(AMB\) Pricing](#).

Create the network using the AWS CLI or AMB Access console according to the following instructions. It takes around 30 minutes for AMB Access to provision resources and bring the network online.

To create a Hyperledger Fabric network using the AWS Management Console

1. Open the AMB Access console at <https://console.aws.amazon.com/managedblockchain/>.
2. Choose **Create private network**.
3. Under **Blockchain frameworks**:
 - a. Select the blockchain framework to use. This tutorial is based on **Hyperledger Fabric version 2.2**.
 - b. Select the **Network edition** to use. The network edition determines attributes of the network, such as the maximum number of members, nodes per member, and transaction throughput. Different editions have different rates associated with the membership. For more information, see [Amazon Managed Blockchain \(AMB\) Pricing](#).
4. Enter a **Network name and description**.
5. Under **Voting Policy**, choose the following:
 - a. Enter the **Approval threshold percentage** along with the comparator, either **Greater than** or **Greater than or equal to**. For a proposal to pass, the Yes votes cast must meet this threshold before the vote duration expires.
 - b. Enter the **Proposal duration in hours**. If enough votes are not cast within this duration to either approve or reject a proposal, the proposal status is EXPIRED, no further votes on this proposal are allowed, and the proposal does not pass.
6. Choose **Next**, and then, under **Create member**, do the following to define the first member for the network, which you own:
 - a. Enter a **Member name** that will be visible to all members and an optional **Description**.

- b. Under **Hyperledger Fabric certificate authority (CA) configuration** specify a username and password to be used as the administrator on the Hyperledger Fabric CA. Remember the user name and password. You need them later any time that you create users and resources that need to authenticate.
 - c. Choose **Next**.
7. Review **Network options** and **Member options**, and then choose **Create network and member**.

The **Networks** list shows the name and **Network ID** of the network you created, with a **Status** of **Creating**. It takes around 30 minutes for AMB Access to create your network, after which the **Status** is **Available**.

To create a Hyperledger Fabric network using the AWS CLI

Use the `create-network` command as shown in the following example. Consider the following:

- The example shows `HYPERLEDGER_FABRIC` as the Framework and `2.2` as the FrameworkVersion. The FrameworkConfiguration properties for `--network-configuration` and `--member-configuration` options might be different for other frameworks and versions.
- The AdminPassword must be at least 8 characters long and no more than 32 characters. It must contain at least one uppercase letter, one lowercase letter, and one digit. It cannot have a single quote('), double quote("), forward slash(/), backward slash(\), @, percent sign (%), or a space.
- The member name must not contain any special characters.
- Remember the user name and password. You need them later any time you create users and resources that need to authenticate.

```
[ec2-user@ip-192-0-2-17 ~]$ aws managedblockchain create-network \
--cli-input-json '{"Name":"OurBlockchainNet", "Description":
"OurBlockchainNetDesc", "Framework":"HYPERLEDGER_FABRIC","FrameworkVersion
": "2.2", "FrameworkConfiguration": {"Fabric": {"Edition": "STARTER
"}}, "VotingPolicy": {"ApprovalThresholdPolicy": {"ThresholdPercentage":
50, "ProposalDurationInHours": 24, "ThresholdComparator": "GREATER_THAN
"}}, "MemberConfiguration": {"Name":"org1", "Description":"Org1 first
member of network", "FrameworkConfiguration":{"Fabric":\n{"AdminUsername":
```

```
\ "MyAdminUser\", \"AdminPassword\": \"Password123\"}}, \"LogPublishingConfiguration\": {\"Fabric\": {\"CaLogs\": {\"Cloudwatch\": {\"Enabled\": true}}}}}'
```

The command returns the Network ID and the Member ID, as shown in the following example:

```
{
  "NetworkId": "n-MWY63ZJZU5HGNCMBQER7IN60IU",
  "MemberId": "m-K46ICRRXJRCGRNNS4ES4XUUS5A"
}
```

The **Networks** page on the console shows a **Status** of **Available** when the network is ready. Alternatively, you can use the `list-networks` command, as shown in the following example, to confirm the network status.

```
aws managedblockchain list-networks
```

The command returns information about the network, including an AVAILABLE status.

```
{
  "Networks": [
    {
      "Id": "n-MWY63ZJZU5HGNCMBQER7IN60IU",
      "Name": "MyTestNetwork",
      "Description": "MyNetDescription",
      "Framework": "HYPERLEDGER_FABRIC",
      "FrameworkVersion": "2.2",
      "Status": "AVAILABLE",
      "CreationDate": 1541497086.888,
    }
  ]
}
```

Step 2: Create and Configure the Interface VPC Endpoint

Now that the network is up and running in your VPC, you set up an interface VPC endpoint (AWS PrivateLink) for your member. This allows the Amazon EC2 instance that you use as a Hyperledger Fabric client to interact with the Hyperledger Fabric endpoints that Amazon Managed Blockchain (AMB) exposes for your member and network resources. For more information, see [Interface VPC Endpoints \(AWS PrivateLink\)](#) in the *Amazon VPC User Guide*. Applicable charges for interface VPC endpoints apply. For more information, see [AWS PrivateLink Pricing](#).

The AWS Identity and Access Management (IAM) principal (user) identity that you use must have sufficient IAM permissions to create an interface VPC endpoint in your AWS account. For more information, see [Controlling Access - Creating and Managing VPC Endpoints](#) in the *Amazon VPC User Guide*.

You can create the interface VPC endpoint using a shortcut in the AMB Access console.

To create an interface VPC endpoint using the AMB Access console

1. Open the AMB Access console at <https://console.aws.amazon.com/managedblockchain/>.
2. Choose **Networks**, select your network from the list, and then choose **View details**.
3. Choose **Create VPC endpoint**.
4. Choose a **VPC**.
5. For **Subnets**, choose a subnet from the list, and then choose additional subnets as necessary.
6. For **Security groups**, choose an EC2 security group from the list, and then choose additional security groups as necessary. We recommend that you select the same security group that your framework client EC2 instance is associated with.
7. Choose **Create**.

Step 3: Create a Peer Node in Your Membership

Now that your network and the first member are up and running, you can use the AMB Access console or the AWS CLI to create a peer node. Your member's peer nodes interact with other members' peer nodes on the blockchain to query and update the ledger, and store a local copy of the ledger.

Use one of the following procedures to create a peer node.

To create a peer node using the AWS Management Console

1. Open the AMB Access console at <https://console.aws.amazon.com/managedblockchain/>.
2. Choose **Networks**, select the network from the list, and then choose **View details**.
3. Select a **Member** from the list, and then choose **Create peer node**.
4. Choose configuration parameters for your peer node according to the guidelines in [Work with Hyperledger Fabric Peer Nodes on AMB Access](#), and then choose **Create peer node**.

To create a peer node using the AWS CLI

- Use the `create-node` command, as shown in the following example. Replace the value of `--network-id`, `--member-id`, and `AvailabilityZone` as appropriate.

```
[ec2-user@ip-192-0-2-17 ~]$ aws managedblockchain create-node \  
--node-configuration '{"InstanceType":"bc.t3.small","AvailabilityZone":"us-  
east-1a"}' \  
--network-id n-MWY63ZJZU5HGNCMBQER7IN60IU \  
--member-id m-K46ICRRXJRCGRNNS4ES4XUUS5A
```

The command returns output that includes the peer node's NodeID, as shown in the following example:

```
{  
  "NodeId": "nd-6EAJ5VA43JGGNPX0UZP7Y47E4Y"  
}
```

Step 4: Create an Amazon EC2 Instance and Set Up the Hyperledger Fabric Client

To complete this step, you launch an Amazon EC2 instance using the Amazon Linux AMI. Consider the following requirements and recommendations when you create the Hyperledger Fabric client Amazon EC2 instance:

- We recommend that you launch the client Amazon EC2 instance in the same VPC and using the same security group as the VPC Endpoint that you created in [Step 2: Create and Configure the Interface VPC Endpoint](#). This simplifies connectivity between the Amazon EC2 instance and the Interface VPC Endpoint.
- We recommend that the EC2 security group shared by the VPC Endpoint and the client Amazon EC2 instance have rules that allow all inbound and outbound traffic between members of the security group. This also simplifies connectivity. In addition, ensure that this security group or another security group associated with the client Amazon EC2 instance has a rule that allows inbound SSH connections from a source that includes your SSH client's IP address. For more information about security groups and required rules, see [Configuring Security Groups for Amazon Managed Blockchain \(AMB\) Hyperledger Fabric](#).

- Make sure that the client Amazon EC2 instance is configured with an automatically assigned public IP address and that you can connect to it using SSH. For more information, see [Getting Started with Amazon EC2 Linux Instances](#) and [Connect to your Linux instance](#) in the *Amazon EC2 User Guide*.
- Make sure that the service role associated with the EC2 instance allows access to the Amazon S3 bucket where AMB Access certificates are stored and that it has required permissions for working with AMB Access resources. For more information, see [Example IAM Role Permissions Policy for Hyperledger Fabric Client EC2 Instance](#).

Note

An AWS CloudFormation template to create a Hyperledger Fabric client is available in the [amazon-managed-blockchain-client-templates repository](#) on Github. For more information, see the [readme.md](#) in that repository. For more information about using CloudFormation, see [Getting Started](#) in the *AWS CloudFormation User Guide*.

Step 4.1: Install Packages

Your Hyperledger Fabric client needs some packages and samples installed so that you can work with the Hyperledger Fabric resources. In this step, you install Go, Docker, Docker Compose, and some other utilities. You also create variables in the `~/.bash_profile` for your development environment. These are prerequisites for installing and using Hyperledger tools.

While connected to the Amazon EC2 instance using SSH, run the following commands to install utilities, install docker, and configure the Docker user to be the default user for the Amazon EC2 instance:

```
sudo yum update -y
```

```
sudo yum install jq telnet emacs docker libtool libtool-ltdl-devel git -y
```

```
sudo service docker start
```

```
sudo usermod -a -G docker ec2-user
```

Log out and log in again for the usermod command to take effect.

Run the following commands to install Docker Compose.

```
sudo curl -L \
https://github.com/docker/compose/releases/download/1.20.0/docker-compose-`uname \
-s`-`uname -m` -o /usr/local/bin/docker-compose
```

```
sudo chmod a+x /usr/local/bin/docker-compose
```

Run the following commands to install golang.

```
wget https://dl.google.com/go/go1.14.4.linux-amd64.tar.gz
```

```
tar -xzf go1.14.4.linux-amd64.tar.gz
```

```
sudo mv go /usr/local
```

```
sudo yum install git -y
```

Use a text editor to set up variables such as GOROOT and GOPATH in your `~/.bashrc` or `~/.bash_profile` and save the updates. The following example shows entries in `.bash_profile`.

```
# .bash_profile

# Get the aliases and functions
if [ -f ~/.bashrc ]; then
    . ~/.bashrc
fi

# User specific environment and startup programs
PATH=$PATH:$HOME/.local/bin:$HOME/bin

# GOROOT is the location where Go package is installed on your system
export GOROOT=/usr/local/go

# GOPATH is the location of your work directory
export GOPATH=$HOME/go
```

```
# CASERVICEENDPOINT is the endpoint to reach your member's CA
# for example ca.m-K46ICRRXJRCGRNNS4ES4XUUS5A.n-
MWY63ZJZU5HGNCMBQER7IN60IU.managedblockchain.us-east-1.amazonaws.com:30002
export CASERVICEENDPOINT=MyMemberCaEndpoint

# ORDERER is the endpoint to reach your network's orderer
# for example orderer.n-
MWY63ZJZU5HGNCMBQER7IN60IU.managedblockchain.MyRegion.amazonaws.com:30001
export ORDERER=MyNetworkOrdererEndpoint

# Update PATH so that you can access the go binary system wide
export PATH=$GOROOT/bin:$PATH
export PATH=$PATH:/home/ec2-user/go/src/github.com/hyperledger/fabric-ca/bin
```

After you update `.bash_profile`, apply the changes:

```
source ~/.bash_profile
```

After the installation, verify that you have the correct versions installed:

- Docker–17.06.2-ce or later
- Docker-compose–1.14.0 or later
- Go–1.14.x

To check the Docker version, run the following command:

```
sudo docker version
```

The command returns output similar to the following:

```
Client:
 Version: 18.06.1-ce
 API version: 1.38
 Go version: go1.14.4
 Git commit: CommitHash
 Built: Tue Oct 2 18:06:45 2018
 OS/Arch: linux/amd64
 Experimental: false
```

```
Server:
Engine:
Version: 18.06.1-ce
API version: 1.38 (minimum version 1.12)
Go version: go1.14.4
Git commit: e68fc7a/18.06.1-ce
Built: Tue Oct 2 18:08:26 2018
OS/Arch: linux/amd64
Experimental: false
```

To check the version of Docker Compose, run the following command:

```
sudo /usr/local/bin/docker-compose version
```

The command returns output similar to the following:

```
docker-compose version 1.22.0, build f46880fe
docker-py version: 3.4.1
CPython version: 3.6.6
OpenSSL version: OpenSSL 1.1.0f 25 May 2017
```

To check the version of go, run the following command:

```
go version
```

The command returns output similar to the following:

```
go version go1.14.4 linux/amd64
```

Step 4.2: Set Up the Hyperledger Fabric CA Client

In this step, you verify that you can connect to the Hyperledger Fabric CA using the VPC endpoint you configured in [Step 2: Create and Configure the Interface VPC Endpoint](#). You then install the Hyperledger Fabric CA client. The Fabric CA issues certificates to administrators and network peers.

To verify connectivity to the Hyperledger Fabric CA, you need the CAEndpoint. Use the `get-member` command to get the CA endpoint for your member, as shown in the following example. Replace the values of `--network-id` and `--member-id` with the values returned in [Step 1: Create the Network and First Member](#).

```
aws managedblockchain get-member \  
--network-id n-MWY63ZJZU5HGNCMBQER7IN60IU \  
--member-id m-K46ICRRXJRCGRNNS4ES4XUUS5A
```

Use `curl` or `telnet` to verify that the endpoint resolves. In the following example, the value of the variable `$CASERVICEENDPOINT` is the **CAEndpoint** returned by the `get-member` command.

```
curl https://$CASERVICEENDPOINT/cainfo -k
```

The command should return output similar to the following:

```
{"result":  
{"CAName": "abcd1efghijkl1lmn5op3q52rst", "CAChain": "LongStringOfCharacters", "Version": "1.4.7-  
snapshot-"}  
,"errors": [], "messages": [], "success": true}
```

Note that Hyperledger Fabric v2.2 networks should use version 1.4 of the CA client.

Alternatively, you can connect to the Fabric CA using Telnet as shown in the following example. Use the same endpoint in the `curl` example, but separate the endpoint and the port as shown in the following example.

```
telnet CaEndpoint-Without-Port CaPort
```

The command should return output similar to the following:

```
Trying 10.0.1.228...  
Connected to ca.m-K46ICRRXJRCGRNNS4ES4XUUS5A.n-  
MWY63ZJZU5HGNCMBQER7IN60IU.managedblockchain.us-east-1.amazonaws.com.  
Escape character is '^]'.
```

If you are unable to connect to the Fabric CA, double-check your network settings to ensure that the client Amazon EC2 instance has connectivity with the VPC Endpoint. In particular, ensure that the security groups associated with both the VPC Endpoint and the client Amazon EC2 instance have inbound and outbound rules that allow traffic between them.

Now that you have verified that you can connect to the Hyperledger Fabric CA, run the following commands to configure the CA client.

Note

If you are working with Hyperledger Fabric v1.2 networks, you need to install and build the correct client version, which is available at <https://github.com/hyperledger/fabric-ca/releases/download/v1.2.1/hyperledger-fabric-ca-linux-amd64-1.2.1.tar.gz>.

```
mkdir -p /home/ec2-user/go/src/github.com/hyperledger/fabric-ca
```

```
cd /home/ec2-user/go/src/github.com/hyperledger/fabric-ca
```

```
wget https://github.com/hyperledger/fabric-ca/releases/download/v1.4.7/hyperledger-fabric-ca-linux-amd64-1.4.7.tar.gz
```

```
tar -xzf hyperledger-fabric-ca-linux-amd64-1.4.7.tar.gz
```

Step 4.3: Clone the Samples Repository

Note

If you are working with Hyperledger Fabric v1.2 or v1.4 networks, use `--branch v1.2.0` or `--branch v1.4.7` instead of `--branch v2.2.3` in the following commands.

```
cd /home/ec2-user
```

```
git clone --branch v2.2.3 https://github.com/hyperledger/fabric-samples.git
```

Step 4.4: Configure and Run Docker Compose to Start the Hyperledger Fabric CLI

Use a text editor to create a configuration file for Docker Compose named `docker-compose-cli.yaml` in the `/home/ec2-user` directory, which you use to run the Hyperledger Fabric CLI. You use this CLI to interact with peer nodes that your member owns. Copy the following contents into the file and replace the *placeholder values* according to the following guidance:

- *MyMemberID* is the MemberID returned by the `aws managedblockchain list-members` AWS CLI command and shown on the member details page of the AMB Access console—for example, `m-K46ICRRXJRCGRNNS4ES4XUUS5A`.
- *MyPeerNodeEndpoint* is the PeerEndpoint returned by the `aws managedblockchain get-node` command and listed on the node details page of the AMB Access console—for example, `nd-6EAJ5VA43JGGNPXOUZP7Y47E4Y.m-K46ICRRXJRCGRNNS4ES4XUUS5A.n-MWY63ZJZU5HGNCMBQER7IN6OIU.managedblockchain.us-east-1.amazonaws.com:30003`.

When you subsequently use the `cli` container to run commands—for example, `docker exec cli peer channel create`—you can use the `-e` option to override an environment variable that you establish in the `docker-compose-cli.yaml` file.

Note

If you are working with Hyperledger Fabric v1.2 or v1.4 networks, use `image: hyperledger/fabric-tools:1.2` or `image: hyperledger/fabric-tools:1.4` in the following example instead of `image: hyperledger/fabric-tools:2.2.3`. In addition for v1.2, use `CORE_LOGGING_LEVEL=info` instead of `FABRIC_LOGGING_SPEC=info`.

```
version: '2'
services:
  cli:
    container_name: cli
    image: hyperledger/fabric-tools:2.2.3
    tty: true
    environment:
      - GOPATH=/opt/gopath
      - CORE_VM_ENDPOINT=unix:///host/var/run/docker.sock
      - FABRIC_LOGGING_SPEC=info # Set logging level to debug for more verbose logging
      - CORE_PEER_ID=cli
      - CORE_CHAINCODE_KEEPALIVE=10
      - CORE_PEER_TLS_ENABLED=true
      - CORE_PEER_TLS_ROOTCERT_FILE=/opt/home/managedblockchain-tls-chain.pem
      - CORE_PEER_LOCALMSPID=MyMemberID
      - CORE_PEER_MSPCONFIGPATH=/opt/home/admin-msp
      - CORE_PEER_ADDRESS=MyPeerNodeEndpoint
    working_dir: /opt/home
```

```
command: /bin/bash
volumes:
  - /var/run/:/host/var/run/
  - /home/ec2-user/fabric-samples/chaincode:/opt/gopath/src/github.com/
  - /home/ec2-user:/opt/home
```

Run the following command to start the Hyperledger Fabric peer CLI container:

```
docker-compose -f docker-compose-cli.yaml up -d
```

If you restarted or logged out and back in after the `usermod` command in [Step 4.1: Install Packages](#), you shouldn't need to run this command with `sudo`. If the command fails, you can log out and log back in. Alternatively, you can run the command using `sudo`, as shown in the following example:

```
sudo /usr/local/bin/docker-compose -f docker-compose-cli.yaml up -d
```

Step 5: Enroll an Administrative User

In this step, you use a pre-configured certificate to enroll a user with administrative permissions to your member's certificate authority (CA). To do this, you must create a certificate file. You also need the endpoint for the CA of your member, and the user name and password for the user that you created in [Step 1: Create the Network and First Member](#).

Step 5.1: Create the Certificate File

Run the following command to copy the `managedblockchain-tls-chain.pem` to the `/home/ec2-user` directory. Replace *MyRegion* with the AWS Region you are using—for example, `us-east-1`.

```
aws s3 cp s3://MyRegion.managedblockchain/etc/managedblockchain-tls-chain.pem /home/ec2-user/managedblockchain-tls-chain.pem
```

If the command fails with a permissions error, ensure that a service role associated with the EC2 instance allows access to the Amazon S3 bucket location. For more information see [Example IAM Role Permissions Policy for Hyperledger Fabric Client EC2 Instance](#).

Run the following command to test that you copied the contents to the file correctly:

```
openssl x509 -noout -text -in /home/ec2-user/managedblockchain-tls-chain.pem
```

The command should return the contents of the certificate in human-readable format.

Step 5.2: Enroll the Administrative User

AMB Access registers the user identity that you specified when you created the member as an administrator. In Hyperledger Fabric, this user is known as the *bootstrap identity* because the identity is used to enroll itself. To enroll, you need the CA endpoint, as well as the user name and password for the administrator that you created in [Step 1: Create the Network and First Member](#). For information about registering other user identities as administrators before you enroll them, see [Register and Enroll a Hyperledger Fabric Admin](#).

Use the `get-member` command to get the CA endpoint for your membership as shown in the following example. Replace the values of `--network-id` and `--member-id` with the values returned in [Step 1: Create the Network and First Member](#).

```
aws managedblockchain get-member \  
--network-id n-MWY63ZJZU5HGNCMBQER7IN60IU \  
--member-id m-K46ICRRXJRCGRNNS4ES4XUUS5A
```

The command returns information about the initial member that you created in the network, as shown in the following example. Make a note of the `CaEndpoint`. You also need the `AdminUsername` and password that you created along with the network.

The command returns output similar to the following:

```
{  
  "Member": {  
    "NetworkId": "n-MWY63ZJZU5HGNCMBQER7IN60IU",  
    "Status": "AVAILABLE",  
    "Description": "MyNetDescription",  
    "FrameworkAttributes": {  
      "Fabric": {  
        "CaEndpoint": "ca.m-K46ICRRXJRCGRNNS4ES4XUUS5A.n-  
MWY63ZJZU5HGNCMBQER7IN60IU.managedblockchain.us-east-1.amazonaws.com:30002",  
        "AdminUsername": "Example-AdminUser"  
      }  
    },  
    "StatusReason": "Network member created successfully",  
  }  
}
```

```
    "CreationDate": 1542255358.74,  
    "Id": "m-K46ICRRXJRCGRNNS4ES4XUUS5A",  
    "Name": "org1"  
  }  
}
```

Use the CA endpoint, administrator profile, and the certificate file to enroll the member administrator using the `fabric-ca-client enroll` command, as shown in the following example:

```
fabric-ca-client enroll \  
-u 'https://${AdminUsername}:${AdminPassword}@$CASERVICEENDPOINT' \  
--tls.certfiles /home/ec2-user/managedblockchain-tls-chain.pem -M /home/ec2-user/admin-  
msp
```

To use this command, you will need to set the following environment variables:

- **AdminUsername** — The admin username.
- **AdminPassword** — The admin password.

Warning

Always use variables in your code to pass user credentials. For more information, see [Move hard-coded credentials to AWS Secrets Manager](#)

An example command with fictitious administrator name, password, and endpoint is shown in the following example:

```
fabric-ca-client enroll \  
-u https://Example-AdminUser:Example-Password123@ca.m-K46ICRRXJRCGRNNS4ES4XUUS5A.n-  
MwY63ZJZU5HGNCMBQER7IN60IU.managedblockchain.us-east-1.amazonaws.com:30002 \  
--tls.certfiles /home/ec2-user/managedblockchain-tls-chain.pem -M /home/ec2-user/admin-  
msp
```

The command returns output similar to the following:

```
2018/11/16 02:21:40 [INFO] Created a default configuration file at /home/ec2-  
user/.fabric-ca-client/fabric-ca-client-config.yaml
```

```
2018/11/16 02:21:40 [INFO] TLS Enabled
2018/11/16 02:21:40 [INFO] generating key: &{A:ecdsa S:256}
2018/11/16 02:21:40 [INFO] encoded CSR
2018/11/16 02:21:40 [INFO] Stored client certificate at /home/ec2-user/admin-msp/
signcerts/cert.pem
2018/11/16 02:21:40 [INFO] Stored root CA certificate at /home/ec2-user/admin-msp/
cacerts/ca-abcd1efghijklm5op3q52rst-uz2f2xakfd7vcfewqhckr7q5m-managedblockchain-us-
east-1-amazonaws-com-30002.pem
```

Important

It may take a minute or two after you enroll for you to be able to use your administrator certificate to create a channel with the ordering service.

Step 5.3: Copy Certificates for the MSP

In Hyperledger Fabric, the Membership Service Provider (MSP) identifies which root CAs and intermediate CAs are trusted to define the members of a trust domain. Certificates for the administrator's MSP are in `/home/ec2-user/admin-msp` in this tutorial. Because this MSP is for the member administrator, copy the certificates from `signcerts` to `admincerts` as shown in the following example. The example assumes you are in the `/home/ec2-user` directory when running the command.

```
cp -r /home/ec2-user/admin-msp/signcerts admin-msp/admincerts
```

Step 6: Create a Hyperledger Fabric Channel

In Hyperledger Fabric, a ledger exists in the scope of a channel. The ledger can be shared across the entire network if every member is operating on a common channel. A channel also can be privatized to include only a specific set of participants. Members can be in your AWS account, or they can be members that you invite from other AWS accounts.

In this step, you set up a basic channel. Later on in the tutorial, in [Step 8: Invite Another AWS Account to be a Member and Create a Multi-Member Channel](#), you go through a similar process to set up a channel that includes another member.

Wait a minute or two for the administrative permissions from previous steps to propagate, and then perform these tasks to create a channel.

Note

All Hyperledger Fabric networks on AMB Access support a maximum of 8 channels per network, regardless of network edition.

Step 6.1: Create configtx for Hyperledger Fabric Channel Creation

The `configtx.yaml` file contains details of the channel configuration. For more information, see [Channel Configuration \(configtx\)](#) in the Hyperledger Fabric documentation.

This `configtx.yaml` enables application features associated with Hyperledger Fabric 2.2. It is not compatible with Hyperledger Fabric 1.2 or 1.4. For a `configtx.yaml` compatible with Hyperledger Fabric 1.2 or 1.4, see [Work with Channels](#).

Use a text editor to create a file with the following contents and save it as `configtx.yaml` on your Hyperledger Fabric client. Note the following placeholders and values.

- Replace *MemberID* with the MemberID you returned previously. For example *m-K46ICRRXJRCGRNNS4ES4XUUS5A*.
- The `MSPDir` is set to the same directory location, `/opt/home/admin-msp`, that you established using the `CORE_PEER MSPCONFIGPATH` environment variable in the Docker container for the Hyperledger Fabric CLI in [step 4.4](#).

Important

This file is sensitive. Artifacts from pasting can cause the file to fail with marshalling errors. We recommend using emacs to edit it. You can also use VI, but before using VI, enter `:set paste`, press `i` to enter insert mode, paste the contents, press escape, and then enter `:set nopaste` before saving.

```
#####
#
# ORGANIZATIONS
#
# This section defines the organizational identities that can be referenced
# in the configuration profiles.
```

```

#
#####
Organizations:
  # Org1 defines an MSP using the sampleconfig. It should never be used
  # in production but may be used as a template for other definitions.
  - &Org1
    # Name is the key by which this org will be referenced in channel
    # configuration transactions.
    # Name can include alphanumeric characters as well as dots and dashes.
    Name: MemberID
    # ID is the key by which this org's MSP definition will be referenced.
    # ID can include alphanumeric characters as well as dots and dashes.
    ID: MemberID
    # SkipAsForeign can be set to true for org definitions which are to be
    # inherited from the orderer system channel during channel creation. This
    # is especially useful when an admin of a single org without access to the
    # MSP directories of the other orgs wishes to create a channel. Note
    # this property must always be set to false for orgs included in block
    # creation.
    SkipAsForeign: false
    Policies: &Org1Policies
      Readers:
        Type: Signature
        Rule: "OR('Org1.member')"
        # If your MSP is configured with the new NodeOUs, you might
        # want to use a more specific rule like the following:
        # Rule: "OR('Org1.admin', 'Org1.peer', 'Org1.client')"
      Writers:
        Type: Signature
        Rule: "OR('Org1.member')"
        # If your MSP is configured with the new NodeOUs, you might
        # want to use a more specific rule like the following:
        # Rule: "OR('Org1.admin', 'Org1.client')"
      Admins:
        Type: Signature
        Rule: "OR('Org1.admin')"
    # MSPDir is the filesystem path which contains the MSP configuration.
    MSPDir: /opt/home/admin-msp
    # AnchorPeers defines the location of peers which can be used for
    # cross-org gossip communication. Note, this value is only encoded in
    # the genesis block in the Application section context.
    AnchorPeers:
      - Host: 127.0.0.1
        Port: 7051

```

```
#####
#
# CAPABILITIES
#
# This section defines the capabilities of fabric network. This is a new
# concept as of v1.1.0 and should not be utilized in mixed networks with
# v1.0.x peers and orderers. Capabilities define features which must be
# present in a fabric binary for that binary to safely participate in the
# fabric network. For instance, if a new MSP type is added, newer binaries
# might recognize and validate the signatures from this type, while older
# binaries without this support would be unable to validate those
# transactions. This could lead to different versions of the fabric binaries
# having different world states. Instead, defining a capability for a channel
# informs those binaries without this capability that they must cease
# processing transactions until they have been upgraded. For v1.0.x if any
# capabilities are defined (including a map with all capabilities turned off)
# then the v1.0.x peer will deliberately crash.
#
#####
Capabilities:
# Channel capabilities apply to both the orderers and the peers and must be
# supported by both.
# Set the value of the capability to true to require it.
# Note that setting a later Channel version capability to true will also
# implicitly set prior Channel version capabilities to true. There is no need
# to set each version capability to true (prior version capabilities remain
# in this sample only to provide the list of valid values).
Channel: &ChannelCapabilities
# V2.0 for Channel is a catchall flag for behavior which has been
# determined to be desired for all orderers and peers running at the v2.0.0
# level, but which would be incompatible with orderers and peers from
# prior releases.
# Prior to enabling V2.0 channel capabilities, ensure that all
# orderers and peers on a channel are at v2.0.0 or later.
V2_0: true
# Orderer capabilities apply only to the orderers, and may be safely
# used with prior release peers.
# Set the value of the capability to true to require it.
Orderer: &OrdererCapabilities
# V1.1 for Orderer is a catchall flag for behavior which has been
# determined to be desired for all orderers running at the v1.1.x
# level, but which would be incompatible with orderers from prior releases.
# Prior to enabling V2.0 orderer capabilities, ensure that all
# orderers on a channel are at v2.0.0 or later.
```

```

    V2_0: true
# Application capabilities apply only to the peer network, and may be safely
# used with prior release orderers.
# Set the value of the capability to true to require it.
# Note that setting a later Application version capability to true will also
# implicitly set prior Application version capabilities to true. There is no need
# to set each version capability to true (prior version capabilities remain
# in this sample only to provide the list of valid values).
Application: &ApplicationCapabilities
    # V2.0 for Application enables the new non-backwards compatible
    # features and fixes of fabric v2.0.
    # Prior to enabling V2.0 orderer capabilities, ensure that all
    # orderers on a channel are at v2.0.0 or later.
    V2_0: true
#####
#
# CHANNEL
#
# This section defines the values to encode into a config transaction or
# genesis block for channel related parameters.
#
#####
Channel: &ChannelDefaults
    # Policies defines the set of policies at this level of the config tree
    # For Channel policies, their canonical path is
    # /Channel/<PolicyName>
    Policies:
        # Who may invoke the 'Deliver' API
        Readers:
            Type: ImplicitMeta
            Rule: "ANY Readers"
        # Who may invoke the 'Broadcast' API
        Writers:
            Type: ImplicitMeta
            Rule: "ANY Writers"
        # By default, who may modify elements at this config level
        Admins:
            Type: ImplicitMeta
            Rule: "MAJORITY Admins"
    # Capabilities describes the channel level capabilities, see the
    # dedicated Capabilities section elsewhere in this file for a full
    # description
    Capabilities:
        <<: *ChannelCapabilities

```

```
#####
#
# APPLICATION
#
# This section defines the values to encode into a config transaction or
# genesis block for application-related parameters.
#
#####
Application: &ApplicationDefaults
  # Organizations is the list of orgs which are defined as participants on
  # the application side of the network
  Organizations:
  # Policies defines the set of policies at this level of the config tree
  # For Application policies, their canonical path is
  # /Channel/Application/<PolicyName>
  Policies: &ApplicationDefaultPolicies
    LifecycleEndorsement:
      Type: ImplicitMeta
      Rule: "ANY Readers"
    Endorsement:
      Type: ImplicitMeta
      Rule: "ANY Readers"
    Readers:
      Type: ImplicitMeta
      Rule: "ANY Readers"
    Writers:
      Type: ImplicitMeta
      Rule: "ANY Writers"
    Admins:
      Type: ImplicitMeta
      Rule: "MAJORITY Admins"

  Capabilities:
    <<: *ApplicationCapabilities
#####
#
# PROFILES
#
# Different configuration profiles may be encoded here to be specified as
# parameters to the configtxgen tool. The profiles which specify consortiums
# are to be used for generating the orderer genesis block. With the correct
# consortium members defined in the orderer genesis block, channel creation
# requests may be generated with only the org member names and a consortium
# name.
```

```
#
#####
Profiles:
  OneOrgChannel:
    <<: *ChannelDefaults
    Consortium: AWSSystemConsortium
    Application:
      <<: *ApplicationDefaults
    Organizations:
      - <<: *Org1
```

Run the following command to generate the configtx peer block:

```
docker exec cli configtxgen \
-outputCreateChannelTx /opt/home/mychannel.pb \
-profile OneOrgChannel -channelID mychannel \
--configPath /opt/home/
```

Important

Hyperledger Fabric 2.2 requires that a channel ID contain only lowercase ASCII alphanumeric characters, dots (.), and dashes (-). It must start with a letter, and must be fewer than 250 characters.

Step 6.2: Set Environment Variables for the Orderer

Set the \$ORDERER environment variable for convenience. Replace *orderer.n-MWY63ZJZU5HGNCMBQER7IN60IU.managedblockchain.MyRegion.amazonaws.com:30001* with the OrderingServiceEndpoint returned by the `aws managedblockchain get-network` command and listed on the network details page of the AMB Access console. Replace *MyRegion* with the AWS Region you are using—for example, `us-east-1`.

```
export ORDERER=orderer.n-
MWY63ZJZU5HGNCMBQER7IN60IU.managedblockchain.MyRegion.amazonaws.com:30001
```

This variable must be exported each time you log out of the client. To persist the variable across sessions, add the export statement to your `~/ .bash_profile` as shown in the following example.

```
# .bash_profile
```

```
...other configurations
export ORDERER=orderer.n-
MWY63ZJZU5HGNCMBQER7IN60IU.managedblockchain.MyRegion.amazonaws.com:30001
```

After updating `.bash_profile`, apply the changes:

```
source ~/.bash_profile
```

Step 6.3: Create the Channel

Run the following command to create a channel using the variables that you established and the configtx peer block that you created:

```
docker exec cli peer channel create -c mychannel \  
-f /opt/home/mychannel.pb -o $ORDERER \  
--cafile /opt/home/managedblockchain-tls-chain.pem --tls
```

Important

It may take a minute or two after you enroll an administrative user for you to be able to use your administrator certificate to create a channel with the ordering service.

Step 6.4: Join Your Peer Node to the Channel

Run the following command to join the peer node that you created earlier to the channel:

```
docker exec cli peer channel join -b mychannel.block \  
-o $ORDERER --cafile /opt/home/managedblockchain-tls-chain.pem --tls
```

Step 7: Install and Run Chaincode

In this section, you create and install a package for [golang sample chaincode](#) on your peer node. You also approve, commit, and verify the chaincode.

You then use the chaincode's `init` command to initialize values attributed to entities a and b in the ledger, followed by the `query` command to confirm that initialization was successful. Next,

you use the chaincode's `invoke` command to transfer 10 units from a to b in the ledger. Finally, you use the chaincode's `query` command again to confirm that the value attributed to a was decremented by 10 units in the ledger.

Step 7.1: Install Vendor Dependencies

Run the following commands to enable *vendoring* for the Go module dependencies of your example chaincode.

```
sudo chown -R ec2-user:ec2-user fabric-samples/  
cd fabric-samples/chaincode/abstore/go/  
G0111MODULE=on go mod vendor  
cd -
```

Step 7.2: Create the Chaincode Package

Run the following command to create the example chaincode package.

```
docker exec cli peer lifecycle chaincode package ./abstore.tar.gz \  
--path fabric-samples/chaincode/abstore/go/ \  
--label abstore_1
```

Step 7.3: Install the Package

Run the following command to install the chaincode package on the peer node.

```
docker exec cli peer lifecycle chaincode install abstore.tar.gz
```

Step 7.4: Verify the Package

Run the following command to verify that the chaincode package is installed on the peer node.

```
docker exec cli peer lifecycle chaincode queryinstalled
```

The command returns the following if the package is installed successfully.

```
Installed chaincodes on peer:  
Package ID: MyPackageID, Label: abstore_1
```

Step 7.5: Approve the Chaincode

Run the following commands to approve the chaincode definition for your organization. Replace *MyPackageID* with the *Package ID* value returned in the previous step [Step 7.4: Verify the Package](#).

```
export CC_PACKAGE_ID=MyPackageID
docker exec cli peer lifecycle chaincode approveformyorg \
--orderer $ORDERER --tls --cafile /opt/home/managedblockchain-tls-chain.pem \
--channelID mychannel --name mycc --version v0 --sequence 1 --package-id $CC_PACKAGE_ID
```

Step 7.6: Check Commit Readiness

Run the following command to check whether the chaincode definition is ready to be committed on the channel.

```
docker exec cli peer lifecycle chaincode checkcommitreadiness \
--orderer $ORDERER --tls --cafile /opt/home/managedblockchain-tls-chain.pem \
--channelID mychannel --name mycc --version v0 --sequence 1
```

The command returns `true` if the chaincode is ready to be committed.

```
Chaincode definition for chaincode 'mycc', version 'v0', sequence '1' on channel
'mychannel' approval status by org:
m-LVQMIJ75CNCUZATGHLDP24HUHM: true
```

Step 7.7: Commit the Chaincode

Run the following command to commit the chaincode definition on the channel.

```
docker exec cli peer lifecycle chaincode commit \
--orderer $ORDERER --tls --cafile /opt/home/managedblockchain-tls-chain.pem \
--channelID mychannel --name mycc --version v0 --sequence 1
```

Step 7.8: Verify the Chaincode

You might have to wait a minute or two for the commit to propagate to the peer node. Run the following command to verify that the chaincode is committed.

```
docker exec cli peer lifecycle chaincode querycommitted \  
--channelID mychannel
```

The command returns the following if the chaincode is committed successfully.

```
Committed chaincode definitions on channel 'mychannel':  
Name: mycc, Version: v0, Sequence: 1, Endorsement Plugin: escc, Validation Plugin: vscc
```

Step 7.9: Initialize the Chaincode

Run the following command to initialize the chaincode.

```
docker exec cli peer chaincode invoke \  
--tls --cafile /opt/home/managedblockchain-tls-chain.pem \  
--channelID mychannel \  
--name mycc -c '{"Args":["init", "a", "100", "b", "200"]}'
```

The command returns the following when the chaincode is initialized.

```
2021-12-20 19:23:05.434 UTC [chaincodeCmd] chaincodeInvokeOrQuery -> INFO 0ad Chaincode  
invoke successful. result: status:200
```

Step 7.10: Query the Chaincode

You might need to wait a brief moment for the initialization from the previous command to complete before you run the following command to query a value.

```
docker exec cli peer chaincode query \  
--tls --cafile /opt/home/managedblockchain-tls-chain.pem \  
--channelID mychannel \  
--name mycc -c '{"Args":["query", "a"]}'
```

The command should return the value of a, which you initialized with a value of 100.

Step 7.11: Invoke the Chaincode

In the previous steps, you initialized the key a with a value of 100 and queried to verify. Using the invoke command in the following example, you subtract 10 from that initial value.

```
docker exec cli peer chaincode invoke \  
--tls --cafile /opt/home/managedblockchain-tls-chain.pem \  
--channelID mychannel \  
--name mycc -c '{"Args":["invoke", "a", "b", "10"]}'
```

The command returns the following when the chaincode is invoked.

```
2021-12-20 19:23:22.977 UTC [chaincodeCmd] chaincodeInvokeOrQuery -> INFO 0ad Chaincode  
invoke successful. result: status:200
```

Lastly, you again query the value of a using the following command.

```
docker exec cli peer chaincode query \  
--tls --cafile /opt/home/managedblockchain-tls-chain.pem \  
--channelID mychannel \  
--name mycc -c '{"Args":["query", "a"]}'
```

The command should return the new value 90.

Step 8: Invite Another AWS Account to be a Member and Create a Multi-Member Channel

Now that you have a Hyperledger Fabric network set up using Amazon Managed Blockchain (AMB), with an initial member in your AWS account and a VPC endpoint with a service name, you are ready to invite additional members. You invite additional members by creating a proposal for an invitation that existing members vote on. Since the blockchain network at this point consists of only one member, the first member always has the only vote on the invitation proposal for the second member. In the steps that follow, the network creator has an initial member named `org1` and the invited member is named `org2`. For proof of concept, you can create an invitation proposal for an additional member in the same AWS account that you used to create the network, or you can create an invitation proposal for a different AWS account.

After the invitation proposal is approved, the invited account can create a member. Invited members are free to reject the invitation or ignore it until the invitation proposal expires. The invited account needs the network ID and VPC endpoint service name of the blockchain network to create a member. For more information, see [Work with Invitations](#). The invited account also needs to fulfill the prerequisites listed in [Prerequisites and Considerations](#).

Step 8.1: Create an Invitation Proposal

Create a proposal to invite an AWS account to create a member and join the network according to the following procedures. You need the AWS account ID of the member you want to invite. You can also invite your own account to create an additional member. If you are using the CLI, you also need the Network ID and Member ID that you created in [Step 1: Create the Network and First Member](#).

To create an invitation proposal using the AWS Management Console

1. Open the AMB Access console at <https://console.aws.amazon.com/managedblockchain/>.
2. From the navigation pane, choose **Networks**, and then choose the network to which you want to invite an AWS account.
3. Choose **Proposals** and then choose **Propose invitation**.
4. For **Submit proposal as**, choose the member in your account that submits the proposal.

Note

The member who submits the proposal must also vote on it. A Yes vote is not automatically assumed.

5. Enter an optional **Description**. The description appears to other members. It's a good way to communicate key points or a reminder about the proposal before they vote.
6. For each AWS account that you want to invite, enter the account number in the space provided. Choose **Add** to enter additional accounts.
7. Choose **Create**.

To create an invitation proposal using the AWS CLI

- Type a command similar to the following. Replace the value of `Principal` with the AWS account ID that you want to invite. Replace the value of `--member-id` with the value for the member in your account that submits the proposal.

```
[ec2-user@ip-192-0-2-17 ~]$ aws managedblockchain create-proposal \  
--actions Invitations=[{Principal=123456789012}] \  
--network-id n-MWY63ZJZU5HGNCMBQER7IN60IU \  
--member-id m-K46ICRRXJRCGRNNS4ES4XUUS5A
```

The command returns the proposal ID, as shown in the following example:

```
{
  "ProposalId": "p-ZR7KUD2YYNESLNG6RQ33X3FUFE"
}
```

Step 8.2: Vote Yes on the Proposal

After you create the invitation proposal, use the first member that you created to vote Yes and approve the proposal. You must do this within the duration defined by the network voting policy.

1. Open the AMB Access console at <https://console.aws.amazon.com/managedblockchain/>.
2. From the navigation pane, choose **Networks**, and then choose the **Network** for which the proposal was made.
3. Choose **Proposals**.
4. Under **Active**, choose the **Proposal ID** to vote on.
5. Under **Vote on proposal**, select the member in your account to vote as. If your account has multiple members, each member gets a vote.
6. Choose **Yes** to vote to approve the proposal. Voting yes is a requirement for the second member to be created in the next step. Choosing **No** rejects the proposal and an invitation is not created.
7. Choose to **Confirm** your vote.

Step 8.3: Create the New Member

To accept an invitation to create a member and join a network, the steps are similar whether you are creating a member in a AMB Access network in a different AWS account or your own AWS account. You first create the member as shown in the following procedures. If you use the AWS CLI, make sure that you have the relevant information, including the Network ID and the Invitation ID that the network sent to your account. When you create a member, you specify the name that identifies your member on the network. You also specify the admin user and password to authenticate to your member certificate authority (CA).

To accept an invitation to create a member and join a network using the AWS Management Console

1. Open the AMB Access console at <https://console.aws.amazon.com/managedblockchain/>.
2. From the navigation pane, choose **Invitations**.
3. Select the invitation that you want to accept from the list, and then choose **Accept invitation**. To view more information about the network you are invited to join, choose the network **Name** from the list
4. Under **Create member and join network**, configure your network member according to the following guidelines:
 - a. Enter a **Member name** that will be visible to all members and an optional **Description**.
 - b. Under **Hyperledger Fabric certificate authority (CA) configuration** specify a username and password to be used as the administrator on the Hyperledger Fabric CA. Remember the user name and password. You need them later any time that you create users and resources that need to authenticate.
5. Choose **Create member and join network**.

To accept an invitation to create a member and join a network using the AWS CLI

- Use the `create-member` command similar to the example below. Replace the value of `--network-id` with the Network ID that you are joining and `--invitation-id` with the Invitation ID sent to your account from the network.

```
aws managedblockchain create-member \  
--network-id n-MWY63ZJZU5HGNCMBQER7ING0IU \  
--invitation-id i-XL9MDD6LVWWDNA9FF94Y4TFTE \  
--member-configuration 'Name=org2,Description=MyMemberDesc,\  
FrameworkConfiguration={Fabric={AdminUsername=MyAdminUsername,\  
AdminPassword=Example-Password123}}'
```

The command returns output similar to the following:

```
{  
  "MemberId": "m-J46DNSFRTVCCLONS9DT5TTLS2A"  
}
```

Additional Steps to Configure a Member

After you create the member, perform the following steps to configure the member. As you perform the steps, replace values with those specific to your member configuration, including the Member ID returned by the previous command. The Network ID and `OrderingServiceEndpoint` are the same for all members.

- [Step 2: Create and Configure the Interface VPC Endpoint](#)

This step is only required if you are creating the second member in a different AWS account.

- [Step 3: Create a Peer Node in Your Membership](#)

- [Step 4: Create an Amazon EC2 Instance and Set Up the Hyperledger Fabric Client](#)

If you are creating an additional member in the same AWS account, and you already have a Hyperledger Fabric client, you can skip most of these steps. However, you should verify connectivity to the Hyperledger Fabric CA as described in [Step 4.2: Set Up the Hyperledger Fabric CA Client](#), using the new CA endpoint for the new member.

- [Step 5: Enroll an Administrative User](#)

Step 8.4: Share Artifacts and Information with the Network Creator

Before a shared channel can be created, the following artifacts and information need to be shared with `org1` by `org2`:

- **org1 needs the org2 administrative certificate**—This certificate is saved to the `/home/ec2-user/admin-msp/admincerts` directory on `org2`'s Hyperledger Fabric client after [Step 5: Enroll an Administrative User](#). This is referenced in the following steps as `Org2AdminCerts`
- **org1 needs the org2 root CA**—This certificate is saved to `org2`'s `/home/ec2-user/admin-msp/cacerts` directory on `org2`'s Hyperledger Fabric client after the same step as previous. This is referenced in the following steps as `Org2CACerts`
- **org1 needs the Endpoint of the peer node that will join the channel**—This Endpoint value is output by the `get-node` command after [Step 3: Create a Peer Node in Your Membership](#) is complete.

Step 8.5: The Channel Creator (org1) Creates Artifacts for org2's MSP

In the following example, the channel creator is org1. The CA administrator for org1 copies the certificates from the step above to a location on the Hyperledger Fabric client computer. The Membership Service Provider (MSP) uses the certificates to authenticate the member.

On the channel creator's Hyperledger Fabric client, use the following commands to create directories to store the certificates, and then copy the certificates from the previous step to the new directories:

```
mkdir /home/ec2-user/org2-msp
mkdir /home/ec2-user/org2-msp/admincerts
mkdir /home/ec2-user/org2-msp/cacerts

cp Org2AdminCerts /home/ec2-user/org2-msp/admincerts
cp Org2CACerts /home/ec2-user/org2-msp/cacerts
```

Org1 needs org2's member ID. You can get this by running the `list-members` command on org1's Hyperledger Fabric client as shown in the following example:

```
aws managedblockchain list-members \
--network-id n-MWY63ZJZU5HGNCMBQER7IN60IU
```

The channel creator (org1) should verify that the required artifacts for channel creation are saved on the Hyperledger Fabric client as shown in the following list:

- Org1 MSP artifacts:
 - `/home/ec2-user/admin-msp/signcerts/certname.pem`
 - `/home/ec2-user/admin-msp/admincerts/certname.pem`
 - `/home/ec2-user/admin-msp/cacerts/certname.pem`
 - `/home/ec2-user/admin-msp/keystore/keyname_sk`
- Org2 MSP artifacts
 - `/home/ec2-user/org2-msp/admincerts/certname.pem`
 - `/home/ec2-user/org2-msp/cacerts/certname.pem`
- The TLS CA cert used for the Region:
 - `/home/ec2-user/managedblockchain-tls-chain.pem`

- Addresses of all peer nodes to join the channel for both org1 and org2.
- The respective member IDs of org1 and org2.
- A `configtx.yaml` file, which you create in the following step, saved to the `/home/ec2-user` directory on the channel creator's Hyperledger Fabric client.

Note

If you created this `configtx` file earlier, delete the old file, rename it, or replace it.

Step 8.6: Create `configtx` for the Multi-Member Channel

The `configtx.yaml` file contains details of the channel configuration. For more information, see [Channel Configuration \(`configtx`\)](#) in the Hyperledger Fabric documentation.

The channel creator creates this file on the Hyperledger File client. If you compare this file to the file created in [Step 6.1: Create `configtx` for Hyperledger Fabric Channel Creation](#), you see that this `configtx.yaml` specifies two members in the channel.

Use a text editor to create a file with the following contents and save it as `configtx.yaml` on your Hyperledger File client.

- Replace `Org1MemberID` with the MemberID of the first member that you created when you [created the network](#). For example, `m-K46ICRRXJRCGRNNS4ES4XUUS5A`.
- For `&Org1`, the `MSPDir` is set to the same directory location, `/opt/home/admin-msp`, that you established using the `CORE_PEER_MSPCONFIGPATH` environment variable in the Docker container for the Hyperledger Fabric CLI in [step 4.4](#) above.
- Replace `Org2MemberID` with the MemberID of the second member that you created in [step 8.3](#). For example, `m-J46DNSFRTVCCLONS9DT5TTLS2A`.
- For `&Org2`, the `MSPDir` is set to the same directory location, `/opt/home/org2-msp`, that you created and copied artifacts to in [step 8.5](#).

Important

This file is sensitive. Artifacts from pasting can cause the file to fail with marshalling errors. We recommend using `emacs` to edit it. You can also use `VI`, but before using `VI`, enter

:set paste, press i to enter insert mode, paste the contents, press escape, and then enter :set nopaste before saving.

```
#####
#
# ORGANIZATIONS
#
# This section defines the organizational identities that can be referenced
# in the configuration profiles.
#
#####
Organizations:
  # Org1 defines an MSP using the sampleconfig. It should never be used
  # in production but may be used as a template for other definitions.
  - &Org1
    # Name is the key by which this org will be referenced in channel
    # configuration transactions.
    # Name can include alphanumeric characters as well as dots and dashes.
    Name: Org1MemberID
    # ID is the key by which this org's MSP definition will be referenced.
    # ID can include alphanumeric characters as well as dots and dashes.
    ID: Org1MemberID
    # SkipAsForeign can be set to true for org definitions which are to be
    # inherited from the orderer system channel during channel creation. This
    # is especially useful when an admin of a single org without access to the
    # MSP directories of the other orgs wishes to create a channel. Note
    # this property must always be set to false for orgs included in block
    # creation.
    SkipAsForeign: false
    Policies: &Org1Policies
    Readers:
      Type: Signature
      Rule: "OR('Org1.member', 'Org2.member')"
      # If your MSP is configured with the new NodeOUs, you might
      # want to use a more specific rule like the following:
      # Rule: "OR('Org1.admin', 'Org1.peer', 'Org1.client')"
    Writers:
      Type: Signature
      Rule: "OR('Org1.member', 'Org2.member')"
      # If your MSP is configured with the new NodeOUs, you might
      # want to use a more specific rule like the following:
      # Rule: "OR('Org1.admin', 'Org1.client')"
```

```

    Admins:
      Type: Signature
      Rule: "OR('Org1.admin')"
# MSPDir is the filesystem path which contains the MSP configuration.
MSPDir: /opt/home/admin-msp
# AnchorPeers defines the location of peers which can be used for
# cross-org gossip communication. Note, this value is only encoded in
# the genesis block in the Application section context.
AnchorPeers:
  - Host: 127.0.0.1
    Port: 7051
- &Org2
  Name: Org2MemberID
  ID: Org2MemberID
  SkipAsForeign: false
  Policies: &Org2Policies
  Readers:
    Type: Signature
    Rule: "OR('Org2.member', 'Org1.member')"
    # If your MSP is configured with the new NodeOUs, you might
    # want to use a more specific rule like the following:
    # Rule: "OR('Org1.admin', 'Org1.peer', 'Org1.client')"
  Writers:
    Type: Signature
    Rule: "OR('Org2.member', 'Org1.member')"
    # If your MSP is configured with the new NodeOUs, you might
    # want to use a more specific rule like the following:
    # Rule: "OR('Org1.admin', 'Org1.client')"
  Admins:
    Type: Signature
    Rule: "OR('Org2.admin')"
# MSPDir is the filesystem path which contains the MSP configuration.
MSPDir: /opt/home/org2-msp
# AnchorPeers defines the location of peers which can be used for
# cross-org gossip communication. Note, this value is only encoded in
# the genesis block in the Application section context.
AnchorPeers:
  - Host: 127.0.0.1
    Port: 7052
#####
#
# CAPABILITIES
#
# This section defines the capabilities of fabric network. This is a new

```

```

# concept as of v1.1.0 and should not be utilized in mixed networks with
# v1.0.x peers and orderers. Capabilities define features which must be
# present in a fabric binary for that binary to safely participate in the
# fabric network. For instance, if a new MSP type is added, newer binaries
# might recognize and validate the signatures from this type, while older
# binaries without this support would be unable to validate those
# transactions. This could lead to different versions of the fabric binaries
# having different world states. Instead, defining a capability for a channel
# informs those binaries without this capability that they must cease
# processing transactions until they have been upgraded. For v1.0.x if any
# capabilities are defined (including a map with all capabilities turned off)
# then the v1.0.x peer will deliberately crash.
#
#####
Capabilities:
# Channel capabilities apply to both the orderers and the peers and must be
# supported by both.
# Set the value of the capability to true to require it.
# Note that setting a later Channel version capability to true will also
# implicitly set prior Channel version capabilities to true. There is no need
# to set each version capability to true (prior version capabilities remain
# in this sample only to provide the list of valid values).
Channel: &ChannelCapabilities
# V2.0 for Channel is a catchall flag for behavior which has been
# determined to be desired for all orderers and peers running at the v2.0.0
# level, but which would be incompatible with orderers and peers from
# prior releases.
# Prior to enabling V2.0 channel capabilities, ensure that all
# orderers and peers on a channel are at v2.0.0 or later.
V2_0: true
# Orderer capabilities apply only to the orderers, and may be safely
# used with prior release peers.
# Set the value of the capability to true to require it.
Orderer: &OrdererCapabilities
# V1.1 for Orderer is a catchall flag for behavior which has been
# determined to be desired for all orderers running at the v1.1.x
# level, but which would be incompatible with orderers from prior releases.
# Prior to enabling V2.0 orderer capabilities, ensure that all
# orderers on a channel are at v2.0.0 or later.
V2_0: true
# Application capabilities apply only to the peer network, and may be safely
# used with prior release orderers.
# Set the value of the capability to true to require it.
# Note that setting a later Application version capability to true will also

```

```

# implicitly set prior Application version capabilities to true. There is no need
# to set each version capability to true (prior version capabilities remain
# in this sample only to provide the list of valid values).
Application: &ApplicationCapabilities
  # V2.0 for Application enables the new non-backwards compatible
  # features and fixes of fabric v2.0.
  # Prior to enabling V2.0 orderer capabilities, ensure that all
  # orderers on a channel are at v2.0.0 or later.
  V2_0: true
#####
#
# CHANNEL
#
# This section defines the values to encode into a config transaction or
# genesis block for channel related parameters.
#
#####
Channel: &ChannelDefaults
  # Policies defines the set of policies at this level of the config tree
  # For Channel policies, their canonical path is
  # /Channel/<PolicyName>
Policies:
  # Who may invoke the 'Deliver' API
  Readers:
    Type: ImplicitMeta
    Rule: "ANY Readers"
  # Who may invoke the 'Broadcast' API
  Writers:
    Type: ImplicitMeta
    Rule: "ANY Writers"
  # By default, who may modify elements at this config level
  Admins:
    Type: ImplicitMeta
    Rule: "MAJORITY Admins"
  # Capabilities describes the channel level capabilities, see the
  # dedicated Capabilities section elsewhere in this file for a full
  # description
Capabilities:
  <<: *ChannelCapabilities
#####
#
# APPLICATION
#
# This section defines the values to encode into a config transaction or

```

```

# genesis block for application-related parameters.
#
#####
Application: &ApplicationDefaults
  # Organizations is the list of orgs which are defined as participants on
  # the application side of the network
  Organizations:
  # Policies defines the set of policies at this level of the config tree
  # For Application policies, their canonical path is
  # /Channel/Application/<PolicyName>
  Policies: &ApplicationDefaultPolicies
    LifecycleEndorsement:
      Type: ImplicitMeta
      Rule: "ANY Readers"
    Endorsement:
      Type: ImplicitMeta
      Rule: "ANY Readers"
    Readers:
      Type: ImplicitMeta
      Rule: "ANY Readers"
    Writers:
      Type: ImplicitMeta
      Rule: "ANY Writers"
    Admins:
      Type: ImplicitMeta
      Rule: "MAJORITY Admins"

  Capabilities:
    <<: *ApplicationCapabilities
#####
#
# PROFILES
#
# Different configuration profiles may be encoded here to be specified as
# parameters to the configtxgen tool. The profiles which specify consortiums
# are to be used for generating the orderer genesis block. With the correct
# consortium members defined in the orderer genesis block, channel creation
# requests may be generated with only the org member names and a consortium
# name.
#
#####
Profiles:
  TwoOrgChannel:
    <<: *ChannelDefaults

```

```
Consortium: AWSSystemConsortium
Application:
  <<: *ApplicationDefaults
Organizations:
  - *Org1
  - *Org2
```

Run the following command to generate the configtx peer block:

```
docker exec cli configtxgen \
-outputCreateChannelTx /opt/home/ourchannel.pb \
-profile TwoOrgChannel -channelID ourchannel \
--configPath /opt/home/
```

Step 8.7: Create the Channel

The channel creator (org1) uses the following command on their Hyperledger Fabric client to submit the channel to the orderer, which creates the channel `ourchannel`. The command example assumes that Docker environment variables have been configured as described in [Step 4.4: Configure and Run Docker Compose to Start the Hyperledger Fabric CLI](#) and that the `$ORDERER` environment variable has been set on the client.

```
docker exec cli peer channel create -c ourchannel \
-f /opt/home/ourchannel.pb -o $ORDERER \
--cafile /opt/home/managedblockchain-tls-chain.pem --tls
```

Step 8.8: Get Channel Genesis Block

Both org1 and org2 need to run the following command on their respective Hyperledger Fabric clients to get the channel's genesis (oldest) block. For more information about the `peer channel` command, see [peer channel](#) in Hyperledger Fabric documentation.

```
docker exec cli peer channel fetch oldest /opt/home/ourchannel.block \
-c ourchannel -o $ORDERER \
--cafile /opt/home/managedblockchain-tls-chain.pem --tls
```

Step 8.9: Join Peer Nodes to the Channel

Both org1 and org2 need to run the following command on their respective Hyperledger Fabric clients to join their peer nodes to the channel. For more information, see the `peer channel join` command example in [peer channel](#) section of in the Hyperledger Fabric documentation.

```
docker exec cli peer channel join -b /opt/home/ourchannel.block \
-o $ORDERER --cafile /opt/home/managedblockchain-tls-chain.pem --tls
```

Optionally, after you join a peer to a channel, you can set up the peer node as an *anchor peer*. Anchor peers support the gossip protocol, which is required for some features of Hyperledger Fabric, such as private data collections and service discovery. For more information, see [Add an Anchor Peer to a Channel](#).

Step 8.10: Install Chaincode

Both org1 and org2 run the following commands on their respective Hyperledger Fabric clients to install example chaincode on their respective peer nodes:

1. Install the example chaincode package on the peer node.

```
docker exec cli peer lifecycle chaincode install abstore.tar.gz
```

2. Verify that the chaincode package is installed on the peer node.

```
docker exec cli peer lifecycle chaincode queryinstalled
```

The command returns the following if the package is installed successfully.

```
Installed chaincodes on peer:
Package ID: MyPackageID, Label: abstore_1
```

3. Approve the chaincode definition for the organization. Replace *MyPackageID* with the Package ID value returned in the previous step.

```
export CC_PACKAGE_ID=MyPackageID
docker exec cli peer lifecycle chaincode approveformyorg \
--orderer $ORDERER --tls --cafile /opt/home/managedblockchain-tls-chain.pem \
```

```
--channelID ourchannel --name myjointcc --version v0 --sequence 1 --package-id  
$CC_PACKAGE_ID
```

4. Check the commit readiness of the chaincode definition.

```
docker exec cli peer lifecycle chaincode checkcommitreadiness \  
--orderer $ORDERER --tls --cafile /opt/home/managedblockchain-tls-chain.pem \  
--channelID ourchannel --name myjointcc --version v0 --sequence 1
```

The command returns true if the chaincode is ready to be committed.

5. Commit the chaincode definition on the channel.

```
docker exec cli peer lifecycle chaincode commit \  
--orderer $ORDERER --tls --cafile /opt/home/managedblockchain-tls-chain.pem \  
--channelID ourchannel --name myjointcc --version v0 --sequence 1
```

Step 8.11: Query Chaincode

Channel members can run the following command to query the chaincode:

```
docker exec cli peer chaincode query -C ourchannel \  
-n myjointcc -c '{"Args":["query","a"]}'
```

The previous command returns the value of a, which you instantiated to a value of 100.

You may need to wait a brief moment for the instantiation from the previous step to complete before you run the following command to query a value:

Step 8.12: Invoke Chaincode

With the channel created and configured with both members, and the chaincode instantiated with values and an endorsement policy, channel members can invoke chaincode. This example command is similar to the example in [Step 7.11: Invoke the Chaincode](#). However, the command uses the `--peerAddresses` option to specify the endpoints of peer nodes that belong to members in the endorsement policy. The example specifies *Org2PeerNodeEndpoint* in addition to *Org1PeerEndpoint*.

```
docker exec cli peer chaincode invoke \  

```

```
-C ourchannel -n myjointcc -c '{"Args":["invoke","a","b","10"]}' \  
--peerAddresses Org1PeerEndpoint \  
--tlsRootCertFiles /opt/home/managedblockchain-tls-chain.pem \  
--peerAddresses Org2PeerNodeEndpoint \  
--tlsRootCertFiles /opt/home/managedblockchain-tls-chain.pem \  
-o $ORDERER --cafile /opt/home/managedblockchain-tls-chain.pem --tls
```

When we query again using the following command:

```
docker exec cli peer chaincode query -C ourchannel \  
-n myjointcc -c '{"Args":["query","a"]}'
```

The command should return the value of a as the new value 90.

Create a Hyperledger Fabric Blockchain Network on Amazon Managed Blockchain (AMB)

When a user in an AWS account creates a Hyperledger Fabric network on Amazon Managed Blockchain (AMB), they also create the first member in the network. This first member has no peer nodes associated with it until you create them. After you create the network and the first member, you can use that member to create an invitation proposal for other members in the same AWS account or in other AWS accounts. Any member can create an invitation proposal.

When you create the network and the first member in your AWS account, the network exists. However, transactions cannot be conducted and the ledger does not exist because there are no peer nodes. Do the following tasks to make your network functional:

- Create an interface VPC endpoint based on the network's **VPC service name** so that you can privately connect to resources. For more information, see [Create an Interface VPC Endpoint for Amazon Managed Blockchain \(AMB\) Hyperledger Fabric](#).
- Create at least one peer node in your first membership to interact with the network and to create and endorse transactions. For more information, see [Work with Hyperledger Fabric Peer Nodes on AMB Access](#).
- Create an invitation proposal for other AWS accounts to be members of the network, or invite an additional member in your account to simulate a multi-AWS account network. Vote Yes on your own proposal to approve it and create the invitation. For more information about inviting members, see [Create an Invitation Proposal](#).

Create a Hyperledger Fabric Network

You can create a Hyperledger Fabric network using the AWS Management Console, the AWS CLI, or the AMB Access SDK [CreateNetwork](#) action.

To create a Hyperledger Fabric network using the AWS Management Console

1. Open the AMB Access console at <https://console.aws.amazon.com/managedblockchain/>.
2. Choose **Create private network**.
3. Under **Blockchain frameworks**:

- a. Select the blockchain framework to use. This tutorial is based on **Hyperledger Fabric version 2.2**.
 - b. Select the **Network edition** to use. The network edition determines attributes of the network, such as the maximum number of members, nodes per member, and transaction throughput. Different editions have different rates associated with the membership. For more information, see [Amazon Managed Blockchain \(AMB\) Pricing](#).
4. Enter a **Network name and description**.
 5. Under **Voting Policy**, choose the following:
 - a. Enter the **Approval threshold percentage** along with the comparator, either **Greater than** or **Greater than or equal to**. For a proposal to pass, the Yes votes cast must meet this threshold before the vote duration expires.
 - b. Enter the **Proposal duration in hours**. If enough votes are not cast within this duration to either approve or reject a proposal, the proposal status is EXPIRED, no further votes on this proposal are allowed, and the proposal does not pass.
 6. Choose **Next**, and then, under **Create member**, do the following to define the first member for the network, which you own:
 - a. Enter a **Member name** that will be visible to all members and an optional **Description**.
 - b. Under **Hyperledger Fabric certificate authority (CA) configuration** specify a username and password to be used as the administrator on the Hyperledger Fabric CA. Remember the user name and password. You need them later any time that you create users and resources that need to authenticate.
 - c. Choose **Next**.
 7. Review **Network options** and **Member options**, and then choose **Create network and member**.

The **Networks** list shows the name and **Network ID** of the network you created, with a **Status** of **Creating**. It takes around 30 minutes for AMB Access to create your network, after which the **Status** is **Available**.

To create a Hyperledger Fabric network using the AWS CLI

Use the `create-network` command as shown in the following example. Consider the following:

- The example shows HYPERLEDGER_FABRIC as the Framework and 2.2 as the FrameworkVersion. The FrameworkConfiguration properties for --network-configuration and --member-configuration options might be different for other frameworks and versions.
- The AdminPassword must be at least 8 characters long and no more than 32 characters. It must contain at least one uppercase letter, one lowercase letter, and one digit. It cannot have a single quote('), double quote("), forward slash(/), backward slash(\), @, percent sign (%), or a space.
- The member name must not contain any special characters.
- Remember the user name and password. You need them later any time you create users and resources that need to authenticate.

```
[ec2-user@ip-192-0-2-17 ~]$ aws managedblockchain create-network \
--cli-input-json '{"Name":"OurBlockchainNet", "Description":
"OurBlockchainNetDesc", "Framework":"HYPERLEDGER_FABRIC", "FrameworkVersion
": "2.2", "FrameworkConfiguration": {"Fabric": {"Edition": "STARTER
"}, "VotingPolicy": {"ApprovalThresholdPolicy": {"ThresholdPercentage":
50, "ProposalDurationInHours": 24, "ThresholdComparator": "GREATER_THAN
"}, "MemberConfiguration": {"Name":"org1", "Description":"Org1 first
member of network", "FrameworkConfiguration":{"Fabric":\n{"AdminUsername":
"MyAdminUser", "AdminPassword":"Password123"}}}}, "LogPublishingConfiguration":
{"Fabric":{"CaLogs":{"Cloudwatch":{"Enabled": true}}}}}'
```

The command returns the Network ID and the Member ID, as shown in the following example:

```
{
  "NetworkId": "n-MWY63ZJZU5HGNCMBQER7IN60IU",
  "MemberId": "m-K46ICRRXJRCGRNNS4ES4XUUS5A"
}
```

Delete a Hyperledger Fabric Network on Amazon Managed Blockchain (AMB)

A Hyperledger Fabric network on Amazon Managed Blockchain (AMB) remains active as long as there are members. A network is deleted only when the last member deletes itself from the network. No member or AWS account, even the creator's AWS account, can delete the network until they are the last member and delete themselves. When you delete the last member, all resources for that member and the blockchain network are deleted. For more information, see [Delete a Member in Your AWS Account](#).

Invite or Remove Hyperledger Fabric Network Members on Amazon Managed Blockchain (AMB)

To invite and remove other Hyperledger Fabric network members, any member can create a *proposal* that is submitted for a vote to all network members. If a proposal is approved within the duration and with the percentage of Yes votes specified in the voting policy for the network, the appropriate action is carried out.

A member can only join the network through an approved invitation proposal. The exception is the first member, which is created along with the network. The first member then submits a proposal and is the sole voter on the proposal to invite the second member.

Note

Each member that is created accrues charges according to the membership rate for the network. For more information, see [Amazon Managed Blockchain \(AMB\) Pricing](#).

An AWS account can delete members from the network that they own directly. A proposal is not required. To delete a member in a different AWS account, a proposal to remove the member is required. Information about all proposals, including the member who created the proposal, the current vote count, and more is available to all network members.

This topic provides basic information for creating proposals to invite or remove members, and to delete a member that your AWS account owns. For more detailed information about proposals, including how to vote on a proposal, see [Work with Proposals for a Hyperledger Fabric Network on Amazon Managed Blockchain \(AMB\)](#).

Create a Proposal to Invite an AWS Account to the Network

You can use the AWS Management Console, the AWS CLI, or the AMB Access API to create an invitation proposal. When a proposal to invite a member is approved, an invitation is sent to the specified AWS accounts. An administrator with the appropriate permissions in that account can then choose to either create a member and join the network or reject the invitation.

To create an invitation proposal using the AWS Management Console

1. Open the AMB Access console at <https://console.aws.amazon.com/managedblockchain/>.
2. From the navigation pane, choose **Networks**, and then choose the network to which you want to invite an AWS account.
3. Choose **Proposals** and then choose **Propose invitation**.
4. For **Submit proposal as**, choose the member in your account that submits the proposal.

Note

The member who submits the proposal must also vote on it. A Yes vote is not automatically assumed.

5. Enter an optional **Description**. The description appears to other members. It's a good way to communicate key points or a reminder about the proposal before they vote.
6. For each AWS account that you want to invite, enter the account number in the space provided. Choose **Add** to enter additional accounts.
7. Choose **Create**.

To create an invitation proposal using the AWS CLI

- Type a command similar to the following. Replace the value of `Principal` with the AWS account ID that you want to invite. Replace the value of `--member-id` with the value for the member in your account that submits the proposal.

```
[ec2-user@ip-192-0-2-17 ~]$ aws managedblockchain create-proposal \  
--actions Invitations=[{Principal=123456789012}] \  
--network-id n-MWY63ZJZU5HGNCMBQER7IN60IU \  
--member-id m-K46ICRRXJRCGRNNS4ES4XUUS5A
```

The command returns the proposal ID, as shown in the following example:

```
{  
  "ProposalId": "p-ZR7KUD2YYNESLNG6RQ33X3FUFE"  
}
```

Create a Proposal to Remove a Member From the Network

You can use the AWS Management Console, the AWS CLI, or the AMB Access API to create a proposal to remove a member owned by another AWS account.

To create a proposal to remove a member using the AWS Management Console

1. Open the AMB Access console at <https://console.aws.amazon.com/managedblockchain/>.
2. From the navigation pane, choose **Networks**, and then choose the network.
3. Choose **Proposals** and then choose **Propose removal**.
4. For **Submit proposal as**, choose the member in your account that submits the proposal.

Note

The member who submits the proposal must also vote on it. A Yes vote is not automatically assumed.

5. Enter an optional **Description**. The description appears to other members. It's a good way to communicate key points or a reminder about the proposal before they vote.
6. For each member that you want to remove, enter the member ID in the space provided. Choose **Add** to enter additional members.

To create a removal proposal using the AWS CLI

- Type a command similar to the following. Replace the value of `Principal` with the AWS account ID that you want to invite. Replace the value of `--member-id` with the value for the member in your account that submits the proposal.

```
[ec2-user@ip-192-0-2-17 ~]$ aws managedblockchain create-proposal \  
--actions Removals=[{MemberID=m-K46ICRRXJRCGRNNS4ES4XUUS5A}] \  
--network-id n-MWY63ZJZU5HGNCMBQER7ING0IU \  
--member-id m-J46DNSFRTVCCLONS9DT5TTLS2A
```

The command returns the proposal ID, as shown in the following example:

```
{  
  "ProposalId": "p-ZR7KUD2YYNESLNG6RQ33X3FUFE"
```

```
}
```

Delete a Member in Your AWS Account

You can use the AWS Management Console, the AWS CLI, or the AMB Access API to directly remove members that your AWS account owns from a network.

Warning

Removing a member in your account deletes all associated resources, such as peer nodes. For your AWS account to rejoin the network, an existing member must create a proposal to invite your AWS account, and the proposal must be approved.

To delete a member in your account using the AWS Management Console

1. Open the AMB Access console at <https://console.aws.amazon.com/managedblockchain/>.
2. Choose **Networks**, choose the network **Name**, and then choose **Members**.
3. Under **Members owned by you**, select a member and then choose **Delete member**.
4. Choose **Delete** when prompted to confirm.

To delete a member in your account using the AWS CLI

- Use the `delete-member` command as shown in the following example. Replace the values of `--network-id` and `--member-id` as appropriate.

```
aws managedblockchain delete-member --network-id n-MWY63ZJZU5HGNCMBQER7IN60IU --  
member-id m-J46DNSFRTVCCLONS9DT5TTLS2A
```

Accept an Invitation to Create a Member and Join a Hyperledger Fabric Network on Amazon Managed Blockchain (AMB)

In Amazon Managed Blockchain (AMB), a member is a distinct identity within the Hyperledger Fabric network associated with an AWS account. An AWS account can have multiple members on the network. Every member in a network must be invited to participate through a proposal made by an existing member and approved according to the network's voting policy. The exception is the first member, which is created along with the network. For more information, see [Work with Proposals for a Hyperledger Fabric Network on Amazon Managed Blockchain \(AMB\)](#). After the invitation is approved, the invited AWS account can create a member and join the network using the invitation ID.

Each member pays an hourly rate, billed per second, for their network membership, peer nodes, and peer node storage. Charges also apply to the amount of data written to the network. Charges may vary depending on the network edition selected when the network was created. For more information, see [Amazon Managed Blockchain \(AMB\) Pricing](#). The resources associated with a member's account depend on the specific blockchain framework and application requirements, but each member must have the following resources:

- **An interface VPC endpoint in the account**—AMB Access is a PrivateLink-powered service, so you must have an interface VPC endpoint in your account to communicate with the service endpoint that the AMB Access network makes available. For more information, see [Create an Interface VPC Endpoint for Amazon Managed Blockchain \(AMB\) Hyperledger Fabric](#) and [Key Concepts: Amazon Managed Blockchain \(AMB\) Networks, Members, and Peer Nodes](#).
- **One or more peer nodes**—Each member must have at least one peer node to actively participate in the blockchain network. When you create a member it has no peer nodes by default. You create peer nodes after you create the member. Peer nodes run on AMB Access instances. Custom Amazon EC2 instances or on-premises instances cannot participate as peer nodes on a Hyperledger Fabric network on AMB Access. For more information, see [Work with Hyperledger Fabric Peer Nodes on AMB Access](#).

Topics

- [Work with Invitations](#)

- [Create a Member and Join a Network](#)

Work with Invitations

If you are invited to join a Hyperledger Fabric network on Amazon Managed Blockchain (AMB), you can accept the invitation by creating a member using the invitation ID. You can also reject the invitation. After you reject an invitation, the invitation ID is no longer valid. A new invitation proposal must be approved, and a new invitation ID is required to create a member. If don't accept or reject an invitation before it expires, the invitation lapses. As with a rejected invitation, a new invitation ID is required.

You can see all pending, accepted, and rejected invitations for your AWS account in the AWS Management Console. Alternatively, you can use the AWS CLI or the AMB Access SDK [ListInvitations](#) action.

You can set up Amazon CloudWatch Events along with Amazon Simple Notification Service so that you receive an alert when there is an invitation for your account. For more information, see [Automating AMB Access Proposal Notifications with CloudWatch Events](#).

To list blockchain network member invitations for your AWS account using the console

1. Open the AMB Access console at <https://console.aws.amazon.com/managedblockchain/>.
2. Choose **Invitations**, and then do one of the following:

To...	Do this...
View details about the network, such as the network ID, the description, endpoints, voting policy details, and current members.	Select the invitation from the list and choose View details .
Use the invitation to create a member and join the network.	Select the invitation from the list and choose Accept Invitation . For next steps, see Create a Member and Join a Network
Reject the invitation.	Select the invitation from the list and choose Reject Invitation .

To list blockchain network member invitations for your AWS account using the AWS CLI

- Use the following command:

```
aws managedblockchain list-invitations
```

The command returns a list of invitations, along with detail for each invitation, as shown in the following example for an invitation in the PENDING status:

```
{
  "Invitations": [
    {
      "CreationDate": "2019-04-08T23:40:20.628Z",
      "ExpirationDate": "2019-04-09T23:40:20.628Z",
      "InvitationId": "i-XL9MDD6LVWWDNA9FF94Y4TFTE",
      "NetworkSummary": {
        "CreationDate": "2019-04-03T13:15:22.345Z",
        "Description": "Test network for supply chain blockchain.",
        "Framework": "HYPERLEDGER_FABRIC",
        "FrameworkVersion": "2.2",
        "Id": "n-MWY63ZJZU5HGNCMBQER7ING0IU",
        "Name": "Example Corp.",
        "Status": "AVAILABLE"
      },
      "Status": "PENDING"
    }
  ]
}
```

You can use the InvitationID with the create-member command to create a member and join the network. For next steps, see [Create a Member and Join a Network](#).

Create a Member and Join a Network

You can use the AMB Access console, the AWS CLI, or the AMB Access SDK [CreateMember](#) action to create a member in a network that your account is invited to. If you created the network, you create the first member when you create the network. All subsequent members must be invited to join by way of a member proposal.

After you create the member, for the member to be functional on the network, your account must have a VPC endpoint associated with the VPC endpoint service name published by the network. For more information, see [Create an Interface VPC Endpoint for Amazon Managed Blockchain \(AMB\) Hyperledger Fabric](#). You also must create at least one peer node in your membership. For more information, see [Work with Hyperledger Fabric Peer Nodes on AMB Access](#).

To accept an invitation to create a member and join a network using the AWS Management Console

1. Open the AMB Access console at <https://console.aws.amazon.com/managedblockchain/>.
2. From the navigation pane, choose **Invitations**.
3. Select the invitation that you want to accept from the list, and then choose **Accept invitation**. To view more information about the network you are invited to join, choose the network **Name** from the list
4. Under **Create member and join network**, configure your network member according to the following guidelines:
 - a. Enter a **Member name** that will be visible to all members and an optional **Description**.
 - b. Under **Hyperledger Fabric certificate authority (CA) configuration** specify a username and password to be used as the administrator on the Hyperledger Fabric CA. Remember the user name and password. You need them later any time that you create users and resources that need to authenticate.
5. Choose **Create member and join network**.

To accept an invitation to create a member and join a network using the AWS CLI

- Use the `create-member` command similar to the example below. Replace the value of `--network-id` with the Network ID that you are joining and `--invitation-id` with the Invitation ID sent to your account from the network.

```
aws managedblockchain create-member \  
--network-id n-MWY63ZJZU5HGNCMBQER7ING0IU \  
--invitation-id i-XL9MDD6LVWWDNA9FF94Y4TFTE \  
--member-configuration 'Name=org2,Description=MyMemberDesc,\  
FrameworkConfiguration={Fabric={AdminUsername=MyAdminUsername,\  
AdminPassword=Example-Password123}}'
```

The command returns output similar to the following:

```
{  
  "MemberId": "m-J46DNSFRTVCCL0NS9DT5TTLS2A"  
}
```

After you create the member, you can use the `get-member` command to return important details about the member configuration.

Create an Interface VPC Endpoint for Amazon Managed Blockchain (AMB) Hyperledger Fabric

Each member of a Hyperledger Fabric network on AMB Access needs to privately access resource endpoints from their client applications and tools. Amazon Managed Blockchain (AMB) uses [Interface VPC Endpoints \(AWS PrivateLink\)](#) to accomplish this.

AMB Access creates a *VPC service name* for each network when it is created. Each network is a unique *endpoint service* with its own VPC service name. Each member then uses the VPC service name to create an interface VPC endpoint in their account. This interface VPC endpoint lets you access resources on the Hyperledger Fabric network on AMB Access through their endpoints. AWS accounts that are not invited to the network don't have access to the VPC service name and cannot set up an interface VPC endpoint for access.

The IAM principal (user) identity that you are using must have sufficient IAM permissions to create an interface VPC endpoint in your AWS account. For more information, see [Controlling Access - Creating and Managing VPC Endpoints](#) in the *Amazon VPC User Guide*.

Any blockchain framework clients that access resources on the network need access to the interface VPC endpoint. For example, if you use an Amazon EC2 instance as a Hyperledger Fabric client, you can create it in a subnet and security group that are shared with the interface VPC endpoint.

Applicable charges for interface VPC endpoints apply. For more information, see [AWS PrivateLink Pricing](#).

The interface VPC endpoint that you set up to access a network must be enabled for private DNS names. This requires that you create the endpoint in a VPC that has the `enableDnsHostnames` and `enableDnsSupport` options set to `true`.

To create an interface VPC endpoint using the AMB Access console

1. Open the AMB Access console at <https://console.aws.amazon.com/managedblockchain/>.
2. Choose **Networks**, select your network from the list, and then choose **View details**.
3. Choose **Create VPC endpoint**.
4. Choose a **VPC**.
5. For **Subnets**, choose a subnet from the list, and then choose additional subnets as necessary.

6. For **Security groups**, choose an EC2 security group from the list, and then choose additional security groups as necessary. We recommend that you select the same security group that your framework client EC2 instance is associated with.
7. Choose **Create**.

To create an interface VPC Endpoint for a network

1. Find the **VPC endpoint service name** of the network. This value is returned by `get-network` command using the AMB Access CLI, and is available on the network **Details** page using the AMB Access console (choose **Networks**, select the network from the list, and then choose **View details**).
2. Open the Amazon VPC console at <https://console.aws.amazon.com/vpc/>.
3. Choose **Endpoints, Create Endpoint**.
4. Choose **Find service by name**. For **Service Name**, enter the VPC Endpoint Service Name from step 1.
5. Choose **Verify** and then choose **Create endpoint**.
6. Make sure that **Enable Private DNS Name** is selected. This option is only available if the VPC you selected has **Enable DNS hostnames** and **Enable DNS support** set to true for the VPC. This is a requirement for the VPC.
7. We recommend that the EC2 security group that you specify for the VPC endpoint is the same as the EC2 security group for the blockchain client that you create to work with the network.

Work with Hyperledger Fabric Peer Nodes on AMB Access

Hyperledger Fabric peer nodes do the work for your member on the network. They keep a local copy of the shared ledger, let you query the ledger, and interact with clients and other peer nodes to perform transactions. A new member has no peer nodes. Create at least one peer node per member.

Each peer node runs on a *AMB Access instance type*. You cannot add a custom Amazon EC2 instance to your member, nor can you connect an on-premises machine. The number of peer nodes and the AMB Access instance type of peer nodes available to each member depends on the network edition specified when the network was created. For more information, see [Amazon Managed Blockchain \(AMB\) Pricing](#).

When you create a peer node, you select the following characteristics:

- **AMB Access instance type**

This determines the computational and memory capacity allocated to this node for the blockchain workload. You can choose more CPU and RAM if you anticipate a more demanding workload for each node. For example, your nodes may need to process a higher rate of transactions. Different instance types are subject to different pricing. You can monitor CPU and memory utilization to determine if your AMB Access instance type is appropriate. For more information, see [Use Hyperledger Fabric Peer Node Metrics on Amazon Managed Blockchain \(AMB\)](#)

- **Availability Zone**

You can select the Availability Zone to launch the peer node in. The ability to distribute peer nodes in a member across different Availability Zones allows you to design your blockchain application for resiliency. For more information, see [Regions and Availability Zones](#) in the *Amazon EC2 User Guide*.

- **State DB configuration**

This determines the type of state database that the peer node uses. **LevelDB** stores chaincode data as simple key-value pairs. It supports only key, key range, and composite key queries.

- **Logging configuration**

You can enable peer node logs, chaincode logs, or both for a member. Peer node logs allow you to debug timeout errors associated with proposals and identify rejected proposals. Chaincode logs help you analyze and debug business logic. For more information, see [Monitoring AMB Access Hyperledger Fabric Using CloudWatch Logs](#).

Create a Hyperledger Fabric Peer Node on Amazon Managed Blockchain (AMB)

You can create a Hyperledger Fabric peer node in a member that is in your AWS account using the AWS Management Console, the AWS CLI, or the AMB Access SDK [CreateNode](#) action.

To create a peer node using the AWS Management Console

1. Open the AMB Access console at <https://console.aws.amazon.com/managedblockchain/>.
2. Choose **Networks**, select the network from the list, and then choose **View details**.
3. Select a **Member** from the list, and then choose **Create peer node**.
4. Choose configuration parameters for your peer node according to the guidelines in [Work with Hyperledger Fabric Peer Nodes on AMB Access](#), and then choose **Create peer node**.

To create a peer node using the AWS CLI

- Use the `create-node` command, as shown in the following example. Replace the value of `--network-id`, `--member-id`, and `AvailabilityZone` as appropriate.

```
[ec2-user@ip-192-0-2-17 ~]$ aws managedblockchain create-node \  
--node-configuration '{"InstanceType":"bc.t3.small","AvailabilityZone":"us-  
east-1a"}' \  
--network-id n-MWY63ZJZU5HGNCMBQER7IN60IU \  
--member-id m-K46ICRRXJRCGRNNS4ES4XUUS5A
```

The command returns output that includes the peer node's NodeID, as shown in the following example:

```
{  
  "NodeId": "nd-6EAJ5VA43JGGNPX0UZP7Y47E4Y"
```

}

View Hyperledger Fabric Peer Node Properties on Amazon Managed Blockchain (AMB)

You can view information about each Hyperledger Fabric peer node that belong to your member using the AWS Management Console, the AWS CLI or the AMB Access API [GetNode](#) command. Details include basic information like the AMB Access instance type, Availability Zone, and creation date, along with the following important properties:

- **Status**

- **Creating**

- AMB Access is provisioning and configuring the AMB Access instance for the peer node.

- **Available**

- The peer node is running and available on the network.

- **Failed**

- The peer node has an issue that has caused AMB Access to add it to the deny list on the network. This usually indicates that the peer node has reached memory or storage capacity. As a first step, we recommend that you delete the instance and provision an instance with more capacity.

- **Create Failed**

- The node could not be created with the AMB Access instance type and the Availability Zone specified. We recommend trying another availability zone, a different instance type, or both.

- **Deleting**

- The node is being deleted. This can happen because the node was deleted by the member, the member was deleted by the AWS account, or the member was deleted through an approved removal proposal.

- **Deleted**

- The node has been deleted. See the previous item for possible reasons.

- **Endpoints**

Hyperledger Fabric uses endpoints associated with each peer node to identify the peer node on the network for different processes. AMB Access assigns unique peer node endpoints to each peer node on each network when the peer node is created. The peer node endpoint consists of the applicable port and the domain name of the peer node derived from the network ID, member ID, and peer node ID. For more information, see [Identifying AMB Access Resources and Connecting from a Client](#). Do not assume that the ports for a service are the same among members; different members may use different ports for the same service. Conversely, peer nodes in different networks may use the same ports, but their endpoints are always unique.

- **Peer endpoint**

Use this endpoint, including the port, within Hyperledger Fabric to address the peer node when using all services other than peer channel-based event services.

- **Peer event endpoint**

Use this endpoint, including the port, within Hyperledger Fabric to address the peer node for peer channel-based event services.

You can also view and monitor **Metrics** related to peer node performance. For more information, see [Use Hyperledger Fabric Peer Node Metrics on Amazon Managed Blockchain \(AMB\)](#).

You can check the peer node status using the `get-node` command, as shown in the following example:

```
aws managedblockchain get-node \  
--network-id n-MWY63ZJZU5HGNCMBQER7IN60IU \  
--member-id m-K46ICRRXJRCGRNNS4ES4XUUS5A \  
--node-id nd-6EAJ5VA43JGGNPXOUZP7Y47E4Y
```

The command returns output that includes the peer node's `PeerEndpoint` and `PeerEventEndpoint`, as shown in the following example. You need this endpoint and port when communicating with the node using your blockchain framework client or addressing the node within an application.

```
{  
  "Node": {  
    "AvailabilityZone": "us-east-1a",  
    "CreationDate": "2019-04-08T23:40:20.628Z",  
    "FrameworkAttributes": {
```

```

    "Fabric": {
      "PeerEndpoint": "nd-6EAJ5VA43JGGNPXOUZP7Y47E4Y.m-
K46ICRRXJRCGRNNS4ES4XUUS5A.n-MWY63ZJZU5HGNCMBQER7IN60IU.managedblockchain.us-
east-1.amazonaws.com:30003",
      "PeerEventEndpoint": "nd-6EAJ5VA43JGGNPXOUZP7Y47E4Y.m-
K46ICRRXJRCGRNNS4ES4XUUS5A.n-MWY63ZJZU5HGNCMBQER7IN60IU.managedblockchain.us-
east-1.amazonaws.com:30004"
    }
  },
  "Id": "nd-6EAJ5VA43JGGNPXOUZP7Y47E4Y",
  "InstanceType": "bc.t3.small",
  "LogPublishingConfiguration": {
    "Fabric": {
      "ChaincodeLogs": {
        "Cloudwatch": {
          "Enabled": true
        }
      },
      "PeerLogs": {
        "Cloudwatch": {
          "Enabled": true
        }
      }
    }
  },
  "StateDB": "CouchDB"
  "Status": "AVAILABLE"
}
}

```

Use Hyperledger Fabric Peer Node Metrics on Amazon Managed Blockchain (AMB)

You can use peer node metrics to track the activity and health of Hyperledger Fabric peer nodes on Amazon Managed Blockchain (AMB) that belong to your member. You can use the AMB Access console to view the metrics for a peer node. AMB Access also reports metrics to Amazon CloudWatch. You can use CloudWatch to set up dashboards, receive alarms, and view log files for peer node metrics. For more information, see [Using Amazon CloudWatch Metrics](#) in the *Amazon CloudWatch User Guide*.

In addition to using peer node metrics, you optionally can enable CloudWatch Logs for peer nodes and for instances of chaincode running on a peer node to view **Peer node logs** and **Chaincode logs**. These logs are useful for troubleshooting and analysis of chaincode activity. For more information, see [Monitoring AMB Access Hyperledger Fabric Using CloudWatch Logs](#).

AMB Access collects the following metrics for each peer node in the `aws/managedblockchain` namespace. Available metrics in AMB Access correspond to [Hyperledger Fabric metrics](#).

Metric name	Description
Fabric metrics	
ChaincodeExecuteTimeouts	The number of chaincode executions (Init or Invoke) that have timed out. Units: Count
EndorserProposalDuration	For each proposal, the time to complete the proposal. Units: Seconds
EndorserProposalValidationFailures	The number of proposals that have failed initial validation. Units: Count
EndorserProposalsReceived	The number of proposals received. Units: Count
EndorserSuccessfulProposals	The number of successful proposals. Units: Count
Transactions	The number of transactions that a peer node receives per minute. Units: Count
Utilization metrics	

Metric name	Description
CPUUtilization	The percentage of total CPU capacity used on the peer node's AMB Access instance at any given instant. Units: Percent
MemoryUtilization	The percentage of total available memory used on the peer node's AMB Access instance at any given instant. Units: Percent

Viewing Peer Node Metrics


You can use the Amazon Managed Blockchain (AMB) console to view graphs for peer node metrics. Metrics are available on the peer node details page.

To view metrics using the AMB Access console

1. Open the AMB Access console at <https://console.aws.amazon.com/managedblockchain/>.
2. Under **Network**, choose the **Name** of the network.
3. Choose **Members**. Under **Members owned by you**, choose the **Name** of the member to which the node belongs.
4. Under **Peer Nodes**, choose the **Node ID** you want to view.

Under **Metrics**, tabs for **Channel Metrics** and **Utilization Metrics** are available.

5. For **Channel Metrics**, choose the channels you want to view or compare from the list.
6. Choose a chart and then use **Statistics**, **Time Range**, and **Period** to customize the chart.



Work with Proposals for a Hyperledger Fabric Network on Amazon Managed Blockchain (AMB)

To make a change to a Hyperledger Fabric network on Amazon Managed Blockchain (AMB) that requires consensus among network members, network members create *proposals*. For example, a member can create a proposal to invite another AWS account to become a member, to invite multiple accounts, or to remove one or more members in different AWS accounts.

A proposal is submitted to all network members to cast a Yes or No vote. If the proposal is approved within the duration and with the percentage of Yes votes specified in the voting policy for the network, the proposed action is carried out. The *voting policy* is established when the network is created and governs votes on all proposals. It can't be updated after the network is created. For more information, see [Create a Hyperledger Fabric Blockchain Network on Amazon Managed Blockchain \(AMB\)](#).

Understanding the Proposal Lifecycle

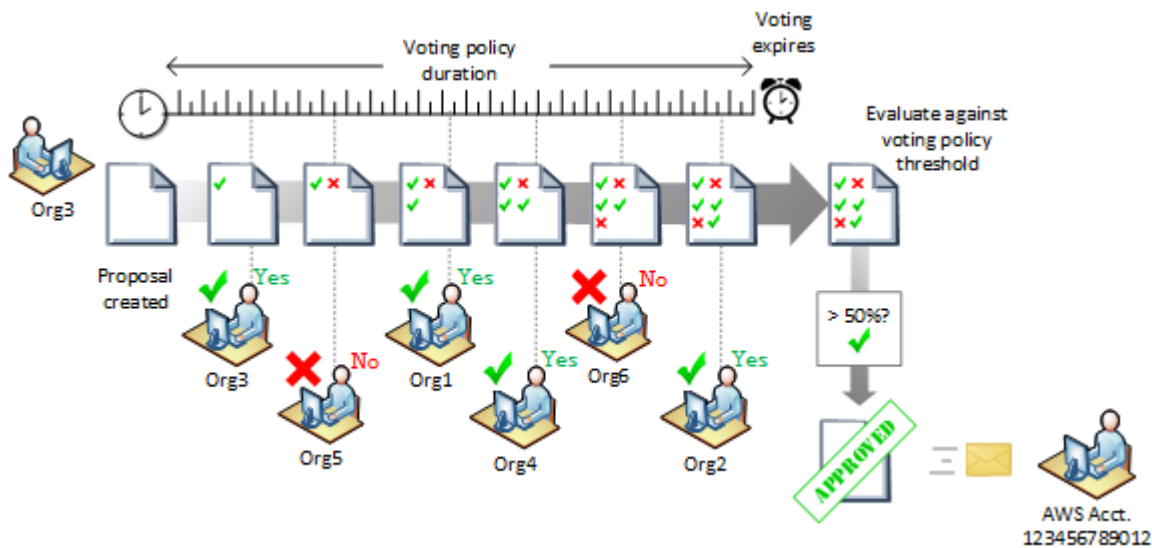
To understand the proposal lifecycle, consider a hypothetical proposal to invite AWS account 123456789012 to join a Hyperledger Fabric network on AMB Access made by a member named Org3. The network currently has six members: Org1, Org2, Org3, and so on. The network was created by Org1, who specified a voting policy with a **50% approval threshold**, a **greater than** comparator, and a proposal duration of 24 hours.

The following flow diagrams depict the possible outcomes of a proposal using this example:

- [Approved with Full Vote](#)
- [Approved with Partial Vote](#)
- [Rejected with Full Vote](#)
- [Rejected with Partial Vote](#)
- [Expired, Does Not Pass](#)

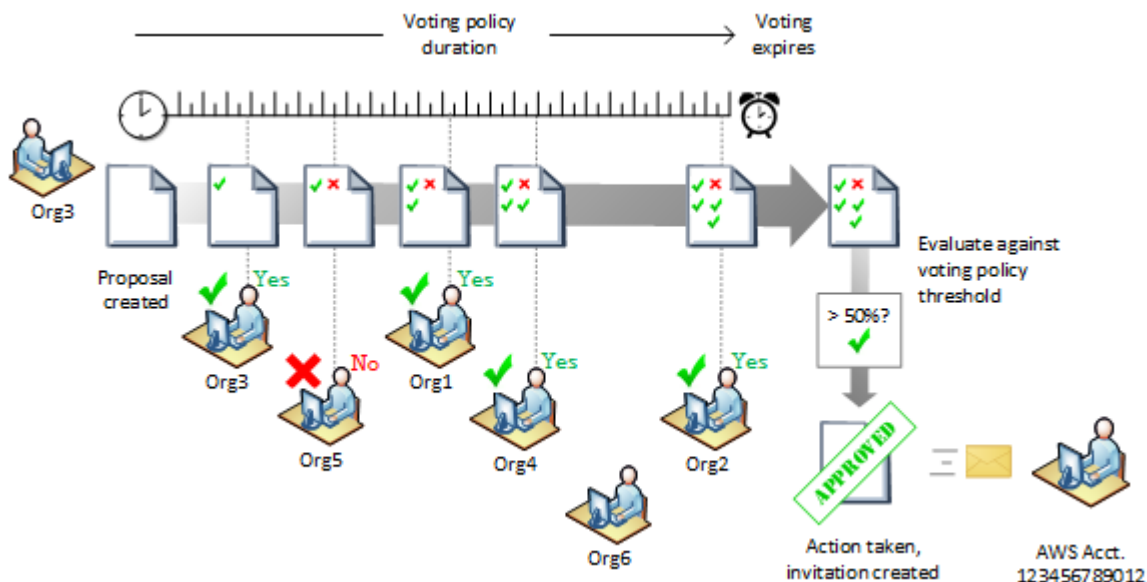
Example– Proposal approved with full member vote

For the following proposal, all members cast a vote before the duration expired. The proposal is APPROVED, and an invitation is extended to the AWS account.




Example– Proposal approved with partial member vote

For the following proposal, not all members cast a vote before the duration expired. However, enough Yes votes were cast to approve the proposal according to the voting policy. The proposal is APPROVED, and an invitation is extended to the AWS account.



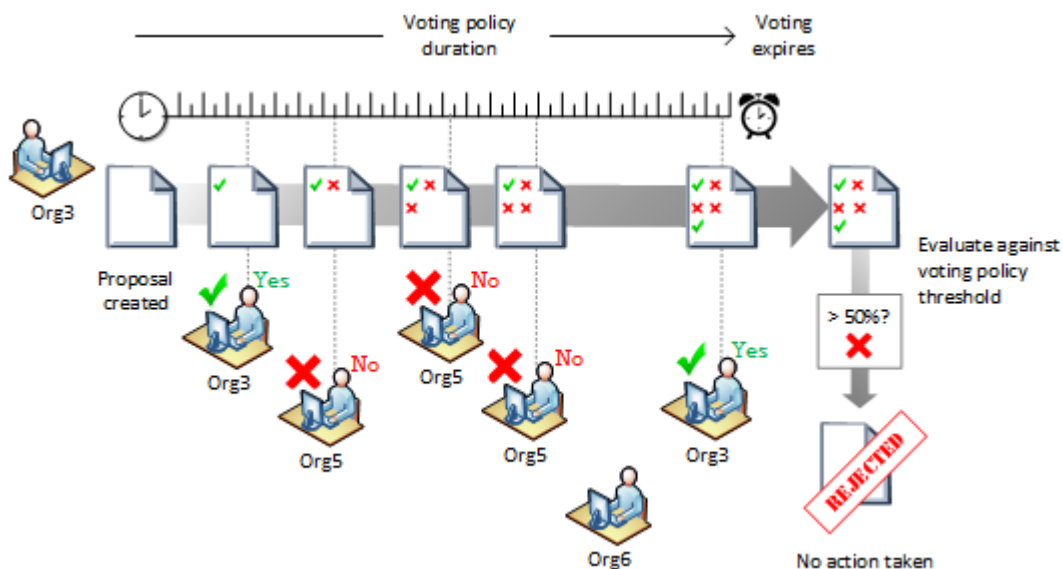
Example– Proposal rejected with full member vote

For the following proposal, all members cast a vote before the duration expired. Because the comparator in the voting policy is **greater than**, a three-to-three vote does not pass the threshold for approval. The proposal is REJECTED, and an invitation is not extended to the AWS account.




Example– Proposal rejected with partial member vote

For the following proposal, not all members cast a vote before the duration expired. However, enough No votes were cast to reject the proposal according to the voting policy. The proposal is REJECTED, and an invitation is not extended to the AWS account.



Example– Proposal expires and is not approved

For the following proposal, not all members cast a vote before the duration expired, and neither the number of Yes nor No votes were cast to determine the outcome of the proposal. The proposal is EXPIRED, and an invitation is not extended to the AWS account.



View Proposals

All proposals made on a network are shown on the **Proposals** page for a network. Both **Active** proposals and **Completed** proposals are listed. Active proposals are still open for voting. You can also list proposals from the AWS CLI using the `list-proposals` command, or using the [ListProposals](#) action with the AMB Access API.

The **Proposals** page for a network shows both **Active** and **Completed** proposals, listing the **Proposal ID**, the name of the member that created the proposal, and the **Expiration Date (UTC)**, which is the creation time plus the proposal duration specified in the network's voting policy. You can choose a **Proposal ID** to vote on active proposals and to see more detail about any proposal, including the actions proposed and a voting summary by member.

Proposals have one of the following statuses:

- **IN_PROGRESS** - The proposal is active and open for member voting.
- **APPROVED** - The proposal was approved with sufficient YES votes among members according to the VotingPolicy specified for the Network. The specified proposal actions are carried out.
- **REJECTED** - The proposal was rejected with insufficient YES votes among members according to the VotingPolicy specified for the Network. The specified ProposalActions are not carried out.
- **EXPIRED** - Members did not cast the number of votes required to determine the proposal outcome before the proposal expired. The specified ProposalActions are not carried out.

- **ACTION_FAILED** - One or more of the specified `ProposalActions` in a proposal that was approved could not be completed because of an error. The **ACTION_FAILED** status occurs even if only one proposal action fails and other actions are successful.

To view proposals for a network using the AWS Management Console

1. Open the AMB Access console at <https://console.aws.amazon.com/managedblockchain/>.
2. Choose **Networks**, choose a network **Name**, and then choose **Proposals**.
3. Choose a **Proposal ID** from the list to view more detailed information, such as the description, a summary of **Actions**, and a **Voting Summary**.
4. Under **Voting Summary**, expand **Votes** to see individual member's votes on the proposal to date.

To view proposals for a network using the AWS CLI

- Enter a command similar to the following example. Replace *n-MWY63ZJZU5HGNCMBQER7IN60IU* with the network ID for which you want to list proposals.

```
aws managedblockchain list-proposals --network-id n-MWY63ZJZU5HGNCMBQER7IN60IU
```

The command returns output similar to the following:

```
{
  "Proposals": [
    {
      "CreationDate": 2019-04-08T23:40:20.628Z,
      "Description": "Proposal to add Example Corp. member",
      "ExpirationDate": 2019-04-09T23:40:20.628Z,
      "ProposalId": "p-ZR7KUD2YYNESLNG6RQ33X3FUFE",
      "ProposedByMemberId": "m-J46DNSFRTVCCLONS9DT5TTLS2A",
      "ProposedByMemberName": "org1",
      "Status": "IN_PROGRESS"
    }
  ]
}
```

To view the details of a proposal using the AWS CLI

- Enter a command similar to the following. Replace *n-MWY63ZJZU5HGNCMBQER7IN60IU* with the network ID and *p-ZR7KUD2YYNESLNG6RQ33X3FUFE* with the proposal ID to view.

```
aws managedblockchain get-proposal --network-id n-MWY63ZJZU5HGNCMBQER7IN60IU --proposal-id p-ZR7KUD2YYNESLNG6RQ33X3FUFE
```

The command returns output similar to the following:

```
{
  "Proposal": {
    "Actions": {
      "Invitations": [
        {
          "Principal": "0123456789012"
        }
      ],
      "CreationDate": "2019-04-08T23:40:20.628Z",
      "Description": "Proposal to invite AWS Acct 0123456789012",
      "ExpirationDate": "2019-04-08T23:40:20.628Z",
      "NetworkId": "n-MWY63ZJZU5HGNCMBQER7IN60IU",
      "NoVoteCount": 1,
      "OutstandingVoteCount": 3,
      "ProposalId": "p-ZR7KUD2YYNESLNG6RQ33X3FUFE",
      "ProposedByMemberId": "m-J46DNSFRTVCCLONS9DT5TTLS2A",
      "ProposedByMemberName": "org1",
      "Status": "IN_PROGRESS",
      "YesVoteCount": 2
    }
  }
}
```

Vote on a Proposal

You can use the AWS Management Console, the AWS CLI `vote-on-proposal` command, or the [VoteOnProposal](#) action of the AMB Access API to vote Yes or No on an active proposal. You cannot change a vote after you make it.

To vote on a proposal using the AWS Management Console

1. Open the AMB Access console at <https://console.aws.amazon.com/managedblockchain/>.
2. Choose **Networks**, choose a network **Name**, and then choose **Proposals**.
3. From the list of **Active** proposals, choose a **Proposal ID**.
4. Under **Vote on proposal**, choose the member to vote as from the list, and then choose **Yes** or **No**.
5. When prompted, choose **Confirm**.

To vote on a proposal using the AWS CLI

- Use the `vote-on-proposal` command as shown in the following example. Replace the values of `--network-id`, `--member-id`, and `--vote` as appropriate.

```
aws managedblockchain vote-on-proposal --network-id n-MWY63ZJZU5HGNCMBQER7IN60IU --  
proposal-id p-ZR7KUD2YYNESLNG6RQ33X3FUFE --member-id m-K46ICRRXJRCGRNNS4ES4XUUS5A  
--vote YES
```

Create an Invitation Proposal

You can use the AWS Management Console, the AWS CLI, or the AMB Access API to create an invitation proposal.

To create an invitation proposal using the AWS Management Console

1. Open the AMB Access console at <https://console.aws.amazon.com/managedblockchain/>.
2. From the navigation pane, choose **Networks**, and then choose the network to which you want to invite an AWS account.
3. Choose **Proposals** and then choose **Propose invitation**.
4. For **Submit proposal as**, choose the member in your account that submits the proposal.

Note

The member who submits the proposal must also vote on it. A Yes vote is not automatically assumed.

5. Enter an optional **Description** . The description appears to other members. It's a good way to communicate key points or a reminder about the proposal before they vote.
6. For each AWS account that you want to invite, enter the account number in the space provided. Choose **Add** to enter additional accounts.
7. Choose **Create**.

To create an invitation proposal using the AWS CLI

- Type a command similar to the following. Replace the value of `Principal` with the AWS account ID that you want to invite. Replace the value of `--member-id` with the value for the member in your account that submits the proposal.

```
[ec2-user@ip-192-0-2-17 ~]$ aws managedblockchain create-proposal \  
--actions Invitations=[{Principal=123456789012}] \  
--network-id n-MWY63ZJZU5HGNCMBQER7IN60IU \  
--member-id m-K46ICRRXJRCGRNNS4ES4XUUS5A
```

The command returns the proposal ID, as shown in the following example:

```
{  
  "ProposalId": "p-ZR7KUD2YYNESLNG6RQ33X3FUFE"  
}
```

Create a Proposal to Remove a Network Member

To create a proposal to remove a member using the AWS Management Console

1. Open the AMB Access console at <https://console.aws.amazon.com/managedblockchain/>.
2. From the navigation pane, choose **Networks**, and then choose the network.
3. Choose **Proposals** and then choose **Propose removal**.
4. For **Submit proposal as**, choose the member in your account that submits the proposal.

Note

The member who submits the proposal must also vote on it. A Yes vote is not automatically assumed.

5. Enter an optional **Description** . The description appears to other members. It's a good way to communicate key points or a reminder about the proposal before they vote.
6. For each member that you want to remove, enter the member ID in the space provided. Choose **Add** to enter additional members.

To create a removal proposal using the AWS CLI

- Type a command similar to the following. Replace the value of `Principal` with the AWS account ID that you want to invite. Replace the value of `--member-id` with the value for the member in your account that submits the proposal.

```
[ec2-user@ip-192-0-2-17 ~]$ aws managedblockchain create-proposal \  
--actions Removals=[{MemberID=m-K46ICRRXJRCGRNNS4ES4XUUS5A}] \  
--network-id n-MWY63ZJZU5HGNCMBQER7ING0IU \  
--member-id m-J46DNSFRTVCCLONS9DT5TTLS2A
```

The command returns the proposal ID, as shown in the following example:

```
{  
  "ProposalId": "p-ZR7KUD2YYNESLNG6RQ33X3FUFE"  
}
```

Automating AMB Access Proposal Notifications with CloudWatch Events

Amazon CloudWatch Events enables you to automate your AWS services and respond automatically to system events. Events from AWS services are delivered to CloudWatch Events in near real time on a best-effort basis. You can write simple rules to indicate which events are of interest to you, and what automated actions to take when an event matches a rule. With AMB Access, you can monitor CloudWatch Events events to respond to proposals, including invitations

sent to your AWS account to join a network, and notification that proposals are APPROVED or REJECTED. Some examples include notifying an Amazon SNS topic or an AWS SNS queue when an invitation is sent or when a proposal made by a member in your account changes status.

For more information, see the [Amazon CloudWatch Events User Guide](#).

Example AMB Access Events

AWS Account Received an Invitation Event

The detail-type of these messages is Managed Blockchain Invitation State Change.

```
{
  "version": "0",
  "id": "abcd1234-eeee-4321-a1a2-123456789012",
  "detail-type": "Managed Blockchain Invitation State Change",
  "source": "aws.managedblockchain",
  "account": "123456789012",
  "time": "2019-04-08T23:40:20.628Z",
  "region": "us-east-1",
  "resources": [],
  "detail": {
    "invitationId": "i-XL9MDD6LVWWDNA9FF94Y4TFTE",
    "networkId": "n-MWY63ZJZU5HGNCMBQER7IN60IU",
    "networkName": "ExampleCorpNetwork",
    "status": "PENDING",
    "expirationDate": "2019-04-09T23:40:20.628Z",
    "message": "You have received invitation i-XL9MDD6LVWWDNA9FF94Y4TFTE for Amazon Managed Blockchain (AMB) Network n-MWY63ZJZU5HGNCMBQER7IN60IU and it will expire at 2016-12-16 20:42 UTC."
  }
}
```

Proposal State Change Event

The detail-type of these messages is Managed Blockchain Proposal State Change. The following example shows an event for a proposal that changed state to APPROVED.

```
{
  "version": "0",
  "id": "abcd1234-eeee-4321-a1a2-123456789012",
  "detail-type": "Managed Blockchain Proposal State Change",
```

```
"source": "aws.managedblockchain",
"account": "123456789012",
"time": "2019-04-08T23:40:20.628Z",
"region": "us-east-1",
"resources": [],
"detail": {
  "proposalId": "p-ZR7KUD2YYNESLNG6RQ33X3FUFE",
  "networkId": "n-MWY63ZJZU5HGNCMBQER7IN60IU",
  "status": "APPROVED",
  "proposedByMemberId": "m-K46ICRRXJRCGRNNS4ES4XUUS5A",
  "proposedByMemberName": "NetworkMember1",
  "expirationDate": "2019-04-09T23:40:20.628Z",
  "description": "Proposal to remove AnyCompany from supply chain blockchain
network.",
  "message": "Voting on proposal p-ZR7KUD2YYNESLNG6RQ33X3FUFE in Amazon
Managed Blockchain (AMB) Network n-MWY63ZJZU5HGNCMBQER7IN60IU completed
at 2016-19-16T20:10:50Z UTC and the proposal was approved."
}
}
```

Work with Hyperledger Fabric Components and Chaincode

You access services and applications in AMB Access Hyperledger Fabric from a framework client machine. The client runs tools and applications that you install for the version of Hyperledger Fabric on the network.

The client accesses AMB Access endpoints for components such as the CA, the orderer, and peer nodes through an interface VPC endpoint that you set up in your AWS account. The client machine must have access to the interface VPC endpoint to access resource endpoints. For more information, see [Create an Interface VPC Endpoint for Amazon Managed Blockchain \(AMB\) Hyperledger Fabric](#).

You can get the endpoints that networks, members, and peer nodes make available using the AWS Management Console, or using get commands and actions with the AMB Access AWS CLI or SDK.

An AWS CloudFormation template to create a Hyperledger Fabric client is available in the [amazon-managed-blockchain-client-templates repository](#) on Github. For more information, see the [readme.md](#) in that repository. For more information about using CloudFormation, see [Getting Started](#) in the *AWS CloudFormation User Guide*.

Topics

- [Register and Enroll a Hyperledger Fabric Admin](#)
- [Work with Channels](#)
- [Add an Anchor Peer to a Channel](#)
- [Develop Hyperledger Fabric Chaincode](#)

Register and Enroll a Hyperledger Fabric Admin

Only identities who are admins within a Hyperledger Fabric member can install and query chaincode.

Creating an admin in Hyperledger Fabric is a three-step process.

1. You *register* the identity with the Hyperledger Fabric CA. Registering stores the user name and password in the CA database as an admin.

2. After you register, you *enroll* the identity. This sends the CA a Certificate Signing Request (CSR). The CA validates that the identity is registered and otherwise valid, and returns a signed certificate. The certificate is stored in the Hyperledger Fabric client machine's local Membership Service Provider (MSP).
3. You then copy the certificate to the `admincerts` subdirectory, and the certificate validates the role of the identity as an admin. Similarly, the CA updates the local MSP for the member's peer nodes and the ordering service so that the admin is recognized.

For more information, see [Fabric CA User's Guide](#) and [Membership](#) in Hyperledger Fabric documentation.

When you first create a member in a Hyperledger Fabric network on AMB Access, you specify the member's first user. AMB Access registers the identity of this user automatically with the Hyperledger Fabric CA. This is called a *bootstrap identity*. Even though the identity is registered, the remaining steps need to be performed. The identity must then enroll itself as an admin and certificates must be updated. After the steps are complete, the identity can install chaincode and can be used to enroll additional identities as admins.

After you enroll an identity as an admin, it may take a minute or two until the identity is able to use the admin certificate to perform tasks.

 **Important**

AMB Access does not support revoking user certificates. After an admin is created, the admin persists for the life of the member.

To register and enroll an identity as an admin, you must have the following:

- The member CA endpoint
- The user name and password either of the bootstrap identity or of an admin with permissions to register and enroll identities
- A valid certificate file and the path to the MSP directory of your identity, which will register the new administrator

Registering an Admin

The following example uses a [Fabric-CA Client CLI](#) `register` command to register an admin with these options:

- `--url` specifies the endpoint of the CA along with an existing user name of an admin with permissions to register, such as the bootstrap identity. The example uses a user name of `Example-AdminUser` with password `Example-Password123`.
- `--id.name` and `--id.secret` parameters establish the user name and password for the new admin.
- `--id.type` is set to `user` and `--id.affiliation` is set to the member name to which the admins belong. The example member name is `org1`.
- `--id.attrs` is set to `'hf.admin=true'`. This is a property specific to AMB Access that registers the identity as an admin.
- The `--tls.certfiles` option specifies the location and file name of the AMB Access TLS certificate that you copied from Amazon S3 (see [Step 5.1: Create the Certificate File](#)).
- `--mspdir` specifies the MSP directory on the local machine where certificates are saved. The example uses `/home/ec2-user/admin-msp`.

```
fabric-ca-client register \  
--url https://Example-AdminUser:Example-Password123@ca.m-K46ICRRXJRCGRNNS4ES4XUUS5A.n-  
MWY63ZJZU5HGNQMBQER7IN60IU.managedblockchain.us-east-1.amazonaws.com:30002 \  
--id.name AdminUser2 --id.secret Password456 \  
--id.type user --id.affiliation org1 \  
--id.attrs 'hf.Admin=true' --tls.certfiles /home/ec2-user/managedblockchain-tls-  
chain.pem \  
--mspdir /home/ec2-user/admin-msp
```

Enrolling an Admin

After registering an identity as an admin, or creating a member along with the bootstrap identity, you can use the [Fabric-CA Client CLI](#) `enroll` command to enroll that identity as an admin. This is shown in the following example using these options:

- `-u` (an alternative for `--url`) specifies the endpoint of the CA along with the user name and password of the identity that you are enrolling.

- `tls.certfiles` specifies the location and file name of the AMB Access TLS certificate that you copied from Amazon S3 (see [Step 5.1: Create the Certificate File](#)).
- `-M` (an alternative for `--mspdir`) specifies the MSP directory on the local machine where certificates are saved. The example uses `/home/ec2-user/admin-msp`.

```
fabric-ca-client enroll \  
-u https://Example-AdminUser:Example-Password123@ca.m-K46ICRRXJRCGRNNS4ES4XUUS5A.n-  
MWY63ZJZU5HGNCMBQER7IN60IU.managedblockchain.us-east-1.amazonaws.com:30002 \  
--tls.certfiles /home/ec2-user/managedblockchain-tls-chain.pem \  
-M /home/ec2-user/admin-msp
```

Copying the Admin Certificate

After you enroll the admin, copy the certificates from the `signcerts` directory to the `admincerts` directory as shown in the following example. The MSP directory `/home/ec2-user/admin-msp` is used in the example, and the example assumes that you are running the command in the `/home/ec2-user` directory.

```
cp -r admin-msp/signcerts admin-msp/admincerts
```

Work with Channels

A *channel* is a private communication pathway between two or more members of a Hyperledger Fabric network on Amazon Managed Blockchain (AMB). Each transaction on a Hyperledger Fabric network occurs on a channel. A network must contain at least one channel and a network on AMB Access can have up to eight channels.

Each channel has its own transaction ledger. The ordering service on the network manages transactions according to a channel configuration that a member creates. The channel configuration specifies which members participate in the channel, the peer nodes that can transact on the channel, and the policies that govern channel updates.

To participate on a channel, a member must join at least one peer node to the channel. A peer node can belong to multiple channels, but ledger data from one channel can't pass to ledger data in another channel. A member shares their admin certificates and root CA with the channel creator when they join the channel. These artifacts allow the ordering service to authenticate and authorize members of the channel when they join their peer nodes to the channel.

Optionally, a channel also can contain private data collections. A private data collection allows a subset of channel members to share ledger data independently of the ordering service. Private data collections are supported only with Hyperledger Fabric 1.4 or later. Private data collections require that at least one peer node for a member is set up as an anchor peer. For more information, see [Add an Anchor Peer to a Channel](#) and [Work with Private Data Collections in AMB Access Hyperledger Fabric](#).

Create a Channel

A channel creator is a member that uses their Hyperledger Fabric client to create a channel configuration file. The channel creator then runs a command that outputs a channel transaction file based on that configuration file. Finally, the channel creator uses the transaction file to run a command that creates the channel with the Hyperledger Fabric ordering service on AMB Access. The channel creator must get security artifacts from channel members to be able to add the members to the channel.

After the channel is created, channel members, including the channel creator, download the channel genesis block to their respective Hyperledger Fabric client machines and join peer nodes to the channel.

To create and join a Hyperledger Fabric channel on AMB Access

The following steps specify file locations in terms of both the Hyperledger Fabric client machine's file system and the Hyperledger Fabric CLI Docker container file system, as appropriate. When the CLI is installed, the Docker Compose configuration file maps these locations to one another. For more information, see [step 4.4](#) in the [Getting Started](#) tutorial of this guide.

In the following examples, the client machine directory `/home/ec2-user` is mapped to the container directory `/opt/home`. In steps where you save artifacts, the steps use the client machine file system. In steps where you reference artifacts in CLI commands and configuration files, the steps use the container file system.

1. Each member that joins a channel shares its security artifacts with the channel creator. The channel creator saves the artifacts on the Hyperledger Fabric client machine. The channel creator then references the location of these artifacts in the next step. For more information about sharing artifacts, see [Step 8.4](#) and 8.5 in the getting started tutorial of this guide. Sharing artifacts is a prerequisite for the steps below.
2. On your Hyperledger Fabric client machine, create a configuration file in yaml format named `configtx.yaml`. Save this file to a file location on your Hyperledger Fabric client machine

that is mapped to the CLI container file system. The following steps use `/home/ec2-user`, which is mapped to `/opt/home`.

Use the examples below as a starting point. For more information about configuration files, see [Channel configuration \(configtx\)](#) in the Hyperledger Fabric documentation. Replace the following configuration parameters as appropriate for your channel, and add sections as required for additional members.

- *Org<n>MemberID* is the respective member ID for each channel participant, for example, `m-K46ICRRXJRCGRNNS4ES4XUUS5A`.
- *Org<n>AdminMSPDir* is the Docker container file system directory where security artifacts for each corresponding member are saved. For example, `/opt/home/Org2AdminMSP` references artifacts saved to `/home/ec2-user/Org2AdminMSP` on the channel creator's client machine.

The `configtx` file for a network on Hyperledger Fabric 1.4 or later includes application settings to enable expanded features. These settings are not supported in Hyperledger Fabric 1.2 channels. Examples for version 1.2, 1.4, and 2.2 networks are provided below.

Version 2.2

```
#####
#
# ORGANIZATIONS
#
# This section defines the organizational identities that can be referenced
# in the configuration profiles.
#
#####
Organizations:
  # Org1 defines an MSP using the sampleconfig. It should never be used
  # in production but may be used as a template for other definitions.
  - &Org1
    # Name is the key by which this org will be referenced in channel
    # configuration transactions.
    # Name can include alphanumeric characters as well as dots and dashes.
    Name: Org1MemberID
    # ID is the key by which this org's MSP definition will be referenced.
    # ID can include alphanumeric characters as well as dots and dashes.
    ID: Org1MemberID
```

```

# SkipAsForeign can be set to true for org definitions which are to be
# inherited from the orderer system channel during channel creation.
This
# is especially useful when an admin of a single org without access to
the
# MSP directories of the other orgs wishes to create a channel. Note
# this property must always be set to false for orgs included in block
# creation.
SkipAsForeign: false
Policies: &Org1Policies
  Readers:
    Type: Signature
    Rule: "OR('Org1.member', 'Org2.member')"
    # If your MSP is configured with the new NodeOUs, you might
    # want to use a more specific rule like the following:
    # Rule: "OR('Org1.admin', 'Org1.peer', 'Org1.client')"
  Writers:
    Type: Signature
    Rule: "OR('Org1.member', 'Org2.member')"
    # If your MSP is configured with the new NodeOUs, you might
    # want to use a more specific rule like the following:
    # Rule: "OR('Org1.admin', 'Org1.client')"
  Admins:
    Type: Signature
    Rule: "OR('Org1.admin')"
# MSPDir is the filesystem path which contains the MSP configuration.
MSPDir: /opt/home/Org1AdminMSP
# AnchorPeers defines the location of peers which can be used for
# cross-org gossip communication. Note, this value is only encoded in
# the genesis block in the Application section context.
AnchorPeers:
  - Host: 127.0.0.1
    Port: 7051
- &Org2
  Name: Org2MemberID
  ID: Org2MemberID
  SkipAsForeign: false
  Policies: &Org2Policies
    Readers:
      Type: Signature
      Rule: "OR('Org2.member', 'Org1.member')"
      # If your MSP is configured with the new NodeOUs, you might
      # want to use a more specific rule like the following:
      # Rule: "OR('Org1.admin', 'Org1.peer', 'Org1.client')"

```

```

    Writers:
      Type: Signature
      Rule: "OR('Org2.member', 'Org1.member')"
      # If your MSP is configured with the new NodeOUs, you might
      # want to use a more specific rule like the following:
      # Rule: "OR('Org1.admin', 'Org1.client')"
    Admins:
      Type: Signature
      Rule: "OR('Org2.admin')"
    # MSPDir is the filesystem path which contains the MSP configuration.
    MSPDir: /opt/home/Org2AdminMSP
    # AnchorPeers defines the location of peers which can be used for
    # cross-org gossip communication. Note, this value is only encoded in
    # the genesis block in the Application section context.
    AnchorPeers:
      - Host: 127.0.0.1
        Port: 7052
#####
#
# CAPABILITIES
#
# This section defines the capabilities of fabric network. This is a new
# concept as of v1.1.0 and should not be utilized in mixed networks with
# v1.0.x peers and orderers. Capabilities define features which must be
# present in a fabric binary for that binary to safely participate in the
# fabric network. For instance, if a new MSP type is added, newer binaries
# might recognize and validate the signatures from this type, while older
# binaries without this support would be unable to validate those
# transactions. This could lead to different versions of the fabric binaries
# having different world states. Instead, defining a capability for a channel
# informs those binaries without this capability that they must cease
# processing transactions until they have been upgraded. For v1.0.x if any
# capabilities are defined (including a map with all capabilities turned off)
# then the v1.0.x peer will deliberately crash.
#
#####
Capabilities:
  # Channel capabilities apply to both the orderers and the peers and must be
  # supported by both.
  # Set the value of the capability to true to require it.
  # Note that setting a later Channel version capability to true will also
  # implicitly set prior Channel version capabilities to true. There is no
  need
  # to set each version capability to true (prior version capabilities remain

```

```

# in this sample only to provide the list of valid values).
Channel: &ChannelCapabilities
  # V2.0 for Channel is a catchall flag for behavior which has been
  # determined to be desired for all orderers and peers running at the
v2.0.0
  # level, but which would be incompatible with orderers and peers from
  # prior releases.
  # Prior to enabling V2.0 channel capabilities, ensure that all
  # orderers and peers on a channel are at v2.0.0 or later.
  V2_0: true
# Orderer capabilities apply only to the orderers, and may be safely
# used with prior release peers.
# Set the value of the capability to true to require it.
Orderer: &OrdererCapabilities
  # V1.1 for Orderer is a catchall flag for behavior which has been
  # determined to be desired for all orderers running at the v1.1.x
  # level, but which would be incompatible with orderers from prior
releases.
  # Prior to enabling V2.0 orderer capabilities, ensure that all
  # orderers on a channel are at v2.0.0 or later.
  V2_0: true
# Application capabilities apply only to the peer network, and may be safely
# used with prior release orderers.
# Set the value of the capability to true to require it.
# Note that setting a later Application version capability to true will also
# implicitly set prior Application version capabilities to true. There is no
need
# to set each version capability to true (prior version capabilities remain
# in this sample only to provide the list of valid values).
Application: &ApplicationCapabilities
  # V2.0 for Application enables the new non-backwards compatible
  # features and fixes of fabric v2.0.
  # Prior to enabling V2.0 orderer capabilities, ensure that all
  # orderers on a channel are at v2.0.0 or later.
  V2_0: true
#####
#
# CHANNEL
#
# This section defines the values to encode into a config transaction or
# genesis block for channel related parameters.
#
#####
Channel: &ChannelDefaults

```

```

# Policies defines the set of policies at this level of the config tree
# For Channel policies, their canonical path is
# /Channel/<PolicyName>
Policies:
  # Who may invoke the 'Deliver' API
  Readers:
    Type: ImplicitMeta
    Rule: "ANY Readers"
  # Who may invoke the 'Broadcast' API
  Writers:
    Type: ImplicitMeta
    Rule: "ANY Writers"
  # By default, who may modify elements at this config level
  Admins:
    Type: ImplicitMeta
    Rule: "MAJORITY Admins"
# Capabilities describes the channel level capabilities, see the
# dedicated Capabilities section elsewhere in this file for a full
# description
Capabilities:
  <<: *ChannelCapabilities
#####
#
# APPLICATION
#
# This section defines the values to encode into a config transaction or
# genesis block for application-related parameters.
#####
Application: &ApplicationDefaults
  # Organizations is the list of orgs which are defined as participants on
  # the application side of the network
  Organizations:
  # Policies defines the set of policies at this level of the config tree
  # For Application policies, their canonical path is
  # /Channel/Application/<PolicyName>
  Policies: &ApplicationDefaultPolicies
    LifecycleEndorsement:
      Type: ImplicitMeta
      Rule: "ANY Readers"
    Endorsement:
      Type: ImplicitMeta
      Rule: "ANY Readers"
    Readers:

```

```

        Type: ImplicitMeta
        Rule: "ANY Readers"
    Writers:
        Type: ImplicitMeta
        Rule: "ANY Writers"
    Admins:
        Type: ImplicitMeta
        Rule: "MAJORITY Admins"

    Capabilities:
        <<: *ApplicationCapabilities
#####
#
#  PROFILES
#
#  Different configuration profiles may be encoded here to be specified as
#  parameters to the configtxgen tool. The profiles which specify consortiums
#  are to be used for generating the orderer genesis block. With the correct
#  consortium members defined in the orderer genesis block, channel creation
#  requests may be generated with only the org member names and a consortium
#  name.
#
#####
Profiles:
  TwoOrgChannel:
    <<: *ChannelDefaults
    Consortium: AWSSystemConsortium
    Application:
      <<: *ApplicationDefaults
      Organizations:
        - *Org1
        - *Org2

```

Version 1.4

```

#####
#
#  Section: Organizations
#
#  - This section defines the different organizational identities which will
#  be referenced later in the configuration.
#
#####

```

Organizations:

```

- &Org1
  # member id defines the organization
  Name: Org1MemberID
  # ID to load the MSP definition as
  ID: Org1MemberID
  #msp dir of org1 in the docker container
  MSPDir: /opt/home/Org1AdminMSP
  # AnchorPeers defines the location of peers which can be used
  # for cross org gossip communication. Note, this value is only
  # encoded in the genesis block in the Application section context
  AnchorPeers:
    - Host:
      Port:
- &Org2
  Name: Org2MemberID
  ID: Org2MemberID
  #msp dir of org2 in the docker container
  MSPDir: /opt/home/Org2AdminMSP
  AnchorPeers:
    - Host:
      Port:

```

```
#####
```

```
#
```

```
# CAPABILITIES
```

```
#
```

```

# This section defines the capabilities of fabric network. This is a new
# concept as of v1.1.0 and should not be utilized in mixed networks with
# v1.0.x peers and orderers. Capabilities define features which must be
# present in a fabric binary for that binary to safely participate in the
# fabric network. For instance, if a new MSP type is added, newer binaries
# might recognize and validate the signatures from this type, while older
# binaries without this support would be unable to validate those
# transactions. This could lead to different versions of the fabric binaries
# having different world states. Instead, defining a capability for a channel
# informs those binaries without this capability that they must cease
# processing transactions until they have been upgraded. For v1.0.x if any
# capabilities are defined (including a map with all capabilities turned off)
# then the v1.0.x peer will deliberately crash.

```

```
#
```

```
#####
```

Capabilities:

```

# Channel capabilities apply to both the orderers and the peers and must be
# supported by both.

```

```
# Set the value of the capability to true to require it.
# Note that setting a later Channel version capability to true will also
# implicitly set prior Channel version capabilities to true. There is no
need
# to set each version capability to true (prior version capabilities remain
# in this sample only to provide the list of valid values).
Channel: &ChannelCapabilities
    # V1.4.3 for Channel is a catchall flag for behavior which has been
    # determined to be desired for all orderers and peers running at the
v1.4.3
    # level, but which would be incompatible with orderers and peers from
    # prior releases.
    # Prior to enabling V1.4.3 channel capabilities, ensure that all
    # orderers and peers on a channel are at v1.4.3 or later.
V1_4_3: true
    # V1.3 for Channel enables the new non-backwards compatible
    # features and fixes of fabric v1.3
V1_3: false
    # V1.1 for Channel enables the new non-backwards compatible
    # features and fixes of fabric v1.1
V1_1: false
# Application capabilities apply only to the peer network, and may be safely
# used with prior release orderers.
# Set the value of the capability to true to require it.
# Note that setting a later Application version capability to true will also
# implicitly set prior Application version capabilities to true. There is no
need
# to set each version capability to true (prior version capabilities remain
# in this sample only to provide the list of valid values).
Application: &ApplicationCapabilities
    # V1.4.2 for Application enables the new non-backwards compatible
    # features and fixes of fabric v1.4.2
V1_4_2: true
    # V1.3 for Application enables the new non-backwards compatible
    # features and fixes of fabric v1.3.
V1_3: false
    # V1.2 for Application enables the new non-backwards compatible
    # features and fixes of fabric v1.2 (note, this need not be set if
    # later version capabilities are set)
V1_2: false
    # V1.1 for Application enables the new non-backwards compatible
    # features and fixes of fabric v1.1 (note, this need not be set if
    # later version capabilities are set).
V1_1: false
```

```
#####
#
# SECTION: Application
#
# - This section defines the values to encode into a config transaction or
# genesis block for application related parameters
#
#####
Application: &ApplicationDefaults
  # Organizations is the list of orgs which are defined as participants on
  # the application side of the network
  Organizations:
  Capabilities:
    <<: *ApplicationCapabilities
#####
#
# Profile
#
# - Different configuration profiles may be encoded here to be specified
# as parameters to the configtxgen tool
#
#####
Profiles:
  TwoOrgChannel:
    Consortium: AWSSystemConsortium
    Application:
      <<: *ApplicationDefaults
      Organizations:
        - *Org1
        - *Org2
```

Version 1.2

```
#####
#
# Section: Organizations
#
# - This section defines the different organizational identities which will
# be referenced later in the configuration.
#
#####
Organizations:
  - &Org1
```

```

        # member id defines the organization
Name: Member1ID
        # ID to load the MSP definition as
ID: Member1ID
        #msp dir of org1 in the docker container
MSPDir: /opt/home/Org1AdminMSP
        # AnchorPeers defines the location of peers which can be used
        # for cross org gossip communication. Note, this value is only
        # encoded in the genesis block in the Application section context
AnchorPeers:
    - Host:
      Port:
- &Org2
    Name: Member2ID
    ID: Member2ID
    MSPDir: /opt/home/Org2AdminMSP
    AnchorPeers:
        - Host:
          Port:

#####
#
# SECTION: Application
#
# - This section defines the values to encode into a config transaction or
# genesis block for application related parameters
#
#####
Application: &ApplicationDefaults
    # Organizations is the list of orgs which are defined as participants on
    # the application side of the network
    Organizations:

#####
#
# Profile
#
# - Different configuration profiles may be encoded here to be specified
# as parameters to the configtxgen tool
#
#####
Profiles:
    TwoOrgChannel:
        Consortium: AWSSystemConsortium

```

```

Application:
  <<: *ApplicationDefaults
Organizations:
  - *Org1
  - *Org2

```

- The channel creator uses the Hyperledger Fabric CLI [configtxgen](#) command as shown in the following example to write a channel creation transaction.

Replace the following configuration parameters as appropriate for your channel.

- `/opt/home/ourchannel.pb` is the channel creation transaction file that the command creates. Specify the file location using the container file system.
- `TwoOrgChannel` is the profile specified in the `configtx.yaml` from the previous step.
- `ourchannel` is the channel ID that will be used for the channel.
- `/opt/home/` for the `--configPath` option is the container location where the `configtx.yaml` file is saved.

```

docker exec cli configtxgen \
  -outputCreateChannelTx /opt/home/ourchannel.pb \
  -profile TwoOrgChannel -channelID ourchannel \
  --configPath /opt/home/

```

- The channel creator uses the Hyperledger Fabric CLI [channel create](#) command to create a channel and write the genesis block to the orderer. The `-f` option specifies the transaction file that you created in the previous step, and the `$ORDERER` environment variable in the example has been set to the endpoint of the network orderer for simplicity—for example, `orderer.n-MWY63ZJZU5HGNCMBQER7IN6OIU.managedblockchain.MyRegion.amazonaws.com:30001`.

```

docker exec cli peer channel create -c ourchannel \
  -f /opt/home/ourchannel.pb -o $ORDERER \
  --cafile /opt/home/managedblockchain-tls-chain.pem --tls

```

- To join a peer node to the channel, each member must fetch the channel genesis block from the orderer, write it to a file, and then reference that genesis block when they join their peer or peers.
 - To get the channel genesis block, each member uses the Hyperledger Fabric CLI [peer channel fetch oldest](#) command. In the following example, the genesis block is written

to a file in the container file system, `/opt/home/ourchannel.block`. The `$ORDERER` variable is used in the same way as the previous step, and the `ourchannel` channel ID is specified.

```
docker exec cli peer channel fetch oldest /opt/home/ourchannel.block \  
-c ourchannel -o $ORDERER \  
--cafile /opt/home/managedblockchain-tls-chain.pem --tls
```

- b. Using the channel genesis block file created above, each member uses the Hyperledger Fabric [peer channel join](#) command to join their member's peer node to the channel.

```
docker exec cli peer channel join -b /opt/home/ourchannel.block \  
-o $ORDERER --cafile /opt/home/managedblockchain-tls-chain.pem --tls
```

Add an Anchor Peer to a Channel

Anchor peers within an organization enable cross-organization communication among peer nodes using the Hyperledger Fabric gossip protocol. In general, the gossip protocol identifies other available member peers, disseminates ledger data across peers on a channel, and brings new peers up to date quickly.

Although anchor peers are not strictly required for an organization on a channel, at least one anchor peer is required for an organization to participate in features that use the gossip protocol, such as [private data collections](#) and service discovery. In addition, we recommend that an organization creates redundant anchor peers on a channel for resiliency. For more information, see [Anchor Peer](#) and [Gossip data dissemination protocol](#) in the Hyperledger Fabric documentation.

The steps in this section demonstrate how you can update a channel configuration to add a peer as an anchor peer for your organization in Amazon Managed Blockchain (AMB). The peer that you want to designate as an anchor peer for your organization must first be joined to the channel.

Prerequisites and Assumptions

The steps in this section are based on the assumption that you are using a network, client machine, member, peer, and joint channel similar to those that you create during the [getting started tutorial](#) in this guide.

The example commands shown presume that the environment variables, the directory locations on the client machine, the locations in the Docker container, the properties of members and peer

nodes, and the properties of the joint channel from that tutorial are being used. You might find it useful to review that information before running these commands.

For more information, see the following resources:

- [Create a client](#)
- [Enroll an administrator](#)
- [Create the peer node](#)
- [Create a multi-member channel](#)
- [Join the peer node to the channel](#)

Adding a Peer as an Anchor Peer

A peer node must already be joined to a channel to be set up as an anchor peer. To update a peer to be an anchor peer, you download a configuration file from the orderer on the Amazon Managed Blockchain (AMB) Hyperledger Fabric network into a directory that you create on the Hyperledger Fabric client. Next, you use the [configtxlator](#) command and the [jq](#) tool to manipulate the original protobuf configuration file between protobuf and JSON. Then you use your final protobuf file to update the channel configuration in the orderer.

1. Create a directory on your Hyperledger Fabric client machine, *channel-artifacts*, that you can use to save channel configuration artifacts that you create.

```
cd /home/ec2-user
```

```
mkdir channel-artifacts
```

2. Use the Hyperledger Fabric [peer channel](#) fetch sub-command to get the channel configuration and save it locally, as shown in the following command. In this example, *ourchannel* is the Channel ID. This Channel ID is the same one established in the [step for creating the multi-member channel](#) in the [Getting Started](#) tutorial of this guide. The command saves the configuration file to the directory you created in a protobuf file named *ourchannelCfg.pb*.

```
docker exec cli peer channel fetch config \  
/opt/home/channel-artifacts/ourchannelCfg.pb \  
-c ourchannel -o $ORDERER \  

```

```
--cafile /opt/home/managedblockchain-tls-chain.pem --tls
```

3. Decode the `ourchannelCfg.pb` configuration from protobuf into a JSON object as shown in the following command.

```
docker exec cli configtxlator proto_decode \
--input /opt/home/channel-artifacts/ourchannelCfg.pb --type common.Block \
--output /opt/home/channel-artifacts/config_block.json
```

4. Switch directories to the `channel-artifacts` directory.

```
cd channel-artifacts
```

5. Convert the JSON object that you created from the protobuf file into a streamlined JSON file named `config.json`, and then copy it to `config_copy.json`, as shown in the following commands.

```
sudo jq .data.data[0].payload.data.config config_block.json > config.json
```

```
cp config.json config_copy.json
```

6. Use the `jq` tool to modify the `config_copy.json` configuration file by adding an anchor peer and to save the modifications to a new JSON configuration file named `modified_config.json` as shown in the following command. Replace the items below as appropriate for your AMB Access member, network, and peer configuration.
 - *OrgMemberID* is the ID of a AMB Access member that belongs to the channel and owns the peer node that will be an anchor peer. For example, `m-K46ICRRXJRCGRNNS4ES4XUUS5A`.
 - *PeerNodeDNS* is the endpoint, minus the port identifier, of the peer node. For example, `nd-6EAJ5VA43JGGNPXOUZP7Y47E4Y.m-K46ICRRXJRCGRNNS4ES4XUUS5A.n-MWY63ZJZU5HGNCMBQER7IN60IU.managedblockchain.us-east-1.amazonaws.com`.
 - *PortNumber* is the port number included in the **peer node endpoint** in the console or the `PeerEndpoint` in the Amazon Managed Blockchain (AMB) API for the peer. For example `30003`.

```
jq '.channel_group.groups.Application.groups["OrgMemberID"].values +=
{"AnchorPeers":{"mod_policy": "Admins","value":{"anchor_peers": [{"host":
```

```
"PeerNodeDNS", "port": PortNumber]]], "version": "0"]}]' config_copy.json >
modified_config.json
```

- Return the configuration JSON files to protobuf format, calculate the difference between them, and then generate a new `config_update.json` configuration file as shown in the following commands.

```
docker exec cli configtxlator proto_encode \
--input /opt/home/channel-artifacts/config.json \
--type common.Config \
--output /opt/home/channel-artifacts/config.pb
```

```
docker exec cli configtxlator proto_encode \
--input /opt/home/channel-artifacts/modified_config.json \
--type common.Config \
--output /opt/home/channel-artifacts/modified_config.pb
```

```
docker exec cli configtxlator compute_update \
--channel_id ourchannel \
--original /opt/home/channel-artifacts/config.pb \
--updated /opt/home/channel-artifacts/modified_config.pb \
--output /opt/home/channel-artifacts/config_update.pb
```

```
docker exec cli configtxlator proto_decode \
--input /opt/home/channel-artifacts/config_update.pb \
--type common.ConfigUpdate \
--output /opt/home/channel-artifacts/config_update.json
```

- Change permissions for the `config_update.json` file you created above so that you can pipe its contents along with transaction envelope parameters to the `jq` tool and create the `config_update_in_envelope.json` file, as shown in the following commands.

```
sudo chmod 755 config_update.json
```

```
echo '{"payload":{"header":{"channel_header":{"channel_id":"ourchannel",
"type":2}}, "data":{"config_update":'$(cat config_update.json)'}}}' | jq . >
config_update_in_envelope.json
```

9. Convert the `config_update_in_envelope.json` file to a `config_update_in_envelope.pb` configuration file in protobuf format as shown in the following command.

```
docker exec cli configtxlator proto_encode --input /opt/home/channel-artifacts/  
config_update_in_envelope.json --type common.Envelope --output /opt/home/channel-  
artifacts/config_update_in_envelope.pb
```

10. Specify the `config_update_in_envelope.pb` configuration file to update the channel configuration on the orderer for the AMB Access Hyperledger Fabric network as shown in the following command.

```
docker exec cli peer channel update \  
-f /opt/home/channel-artifacts/config_update_in_envelope.pb \  
-c ourchannel -o $ORDERER \  
--cafile /opt/home/managedblockchain-tls-chain.pem --tls
```

Develop Hyperledger Fabric Chaincode

Smart contracts in Hyperledger Fabric are known as *chaincode*.

- For a conceptual overview of chaincode, see [Smart Contracts](#) and [Developing Applications](#) in the Hyperledger Fabric documentation.
- For links to Hyperledger Fabric SDKs, see [Getting Started](#) in the Hyperledger Fabric documentation.
- Hyperledger Fabric 2.x introduces decentralized governance for smart contracts. This includes a new Fabric chaincode lifecycle that enables multiple organizations in a network to come to agreement on the parameters of a chaincode. For more information, see [Decentralized governance for smart contracts](#) in the Hyperledger Fabric documentation.

Considerations and Limitations When Developing Chaincode for AMB Access

- All Hyperledger Fabric networks on AMB Access support a maximum of 8 channels per network, regardless of network edition.
- **Hyperledger Fabric Shims for Node.js**

To simplify chaincode development, AMB Access includes versions of the [fabric-shim library](#). This library provides a low-level chaincode interface between applications, peers, and the Hyperledger Fabric system for chaincode applications developed with Node.js. The library version is specified using the `dependencies` object in the `package.json` file bundled with your chaincode. You can specify a version explicitly or use the semantic versioner (semver) for NPM to specify a version range. The following library versions are available without bundling.

- **Hyperledger Fabric 1.4**

- [1.4.0](#)
- [1.4.1](#)
- [1.4.2](#)
- [1.4.4](#)
- [1.4.5](#)

- **Hyperledger Fabric 1.2**

- [1.2.0](#)
- [1.2.2](#)
- [1.2.3](#)
- [1.2.4](#)

Dependencies on other versions of `fabric-shim` or other library packages require that you bundle them with your chaincode because peer nodes do not have internet access to the NPM repository.

AMB Access doesn't currently include `fabric-shim` for **Hyperledger Fabric 2.2**. You need to bundle this version of the library in your Node.js application.

- The default limit for the size of a transaction payload is 1MB. To request a limit increase, create a case using the [AWS Support Center](#).
- AMB Access doesn't support the [External chaincode launcher](#) in Hyperledger Fabric 2.2. This feature requires that you modify a peer's `core.yaml` file, which you don't currently have access to.

Topics

- [Work with Private Data Collections in AMB Access Hyperledger Fabric](#)
- [Develop Hyperledger Fabric Chaincode Using Java on Amazon Managed Blockchain \(AMB\)](#)

Work with Private Data Collections in AMB Access Hyperledger Fabric

Network members that participate on a Hyperledger Fabric channel on AMB Access may want to keep some data private from other members on the same channel. Hyperledger Fabric *private data collections* allow you to accomplish this. Private data collections allow a subset of members on a channel to endorse, commit, and query private data without having to create a separate channel to exclude other members. Private data reduces the overhead associated with channel creation and maintenance. In addition, private data is disseminated only between authorized peers using the Hyperledger Fabric *gossip protocol*, so the data is kept confidential from the ordering service.

To create a private data collection, a peer creates a *private data definition*. The definition establishes the members who have access to the private data collection. It also defines a policy that governs how transactions are endorsed and committed.

Private data is stored on each peer node in a way that is accessible to the peer chaincode and logically distinct from other channel transaction data. This helps ensure that channel members who are not specified in the private data definition cannot access the private data.

For more information about private data collections, see [Private Data](#) and the [Private Data architecture](#) topic in the Hyperledger Fabric documentation. Hyperledger Fabric version 2.x introduces collection-level endorsement policies that can override the endorsement policy set at the chaincode level. To learn more, see [Sharing private data](#) in the Hyperledger Fabric documentation.

Private data collections with Amazon Managed Blockchain (AMB) require the following:

- A network created using Hyperledger Fabric version 1.4 or later.
- A channel configuration that supports the latest features of your version. The [multi-member channel configuration](#) featured in the [Getting Started](#) tutorial of this guide enables the latest features.
- At least one anchor peer for each member that participates in the private data collection. This is because Hyperledger Fabric uses the gossip protocol to distribute private data to authorized peers. For more information, see [Add an Anchor Peer to a Channel](#).

Develop Hyperledger Fabric Chaincode Using Java on Amazon Managed Blockchain (AMB)

Hyperledger Fabric 2.2 networks on Amazon Managed Blockchain (AMB) are based on the following components:

- [JDK 8 or higher](#)
- [Hyperledger Fabric v2.2.4](#)
- [Hyperledger Fabric CA v1.4.7](#)
- [fabric-chaincode-java](#)

AMB Access Dependency on Javaenv Docker Image

AMB Access uses the `fabric-chaincode-java` package to build the `hyperledger/fabric-javaenv` Docker image. This image is the base layer of the Java chaincode container. It has built-in support to build the Java chaincode using Gradle or Maven at the following versions:

- `gradle@5.6.2`
- `maven@3.8.1`

You can verify this by examining the [Docker file in the fabric-chaincode-java package](#).

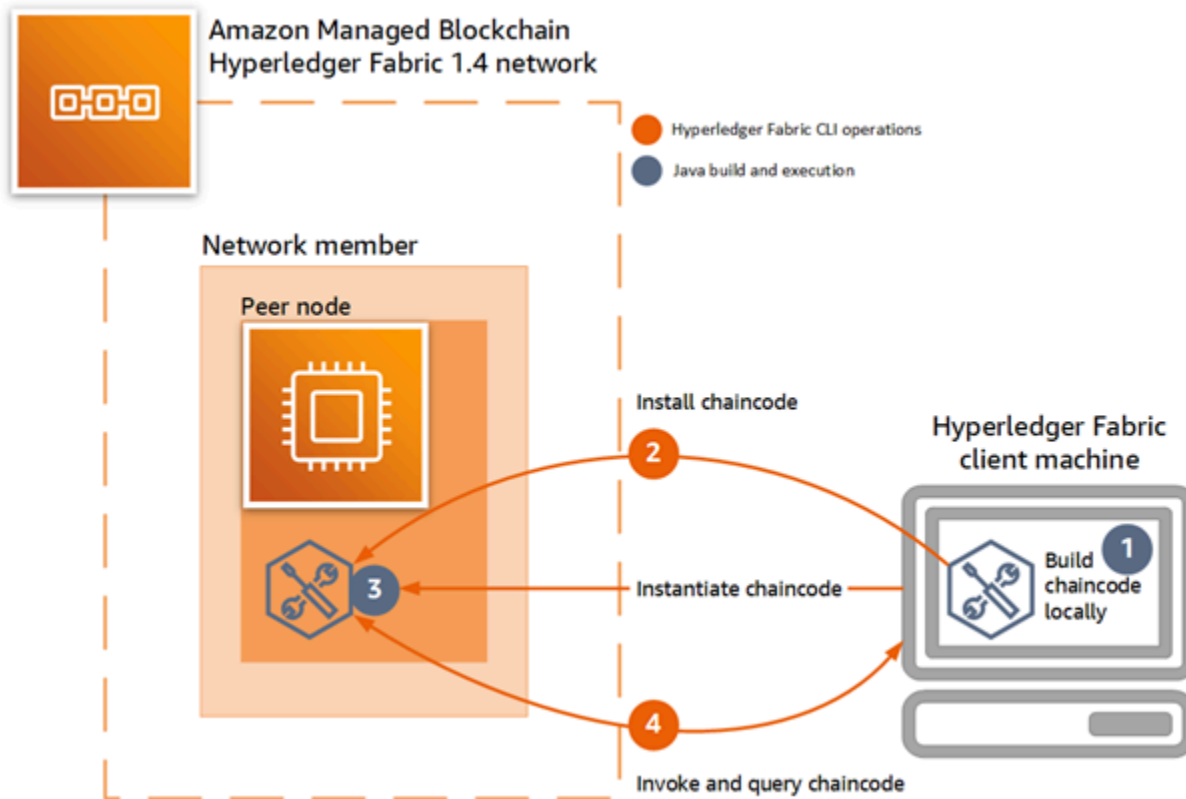
The default settings support Gradle or Maven for building Java chaincode. Once the `javaenv` Docker image is deployed, it is used to provision all new peer nodes and install Java chaincode in those peer nodes. You cannot modify these settings or update versions on the `javaenv` Docker image because AMB Access manages these.

Installing and Running Java Chaincode

The following is a high-level summary of the steps to install and run Java Chaincode using AMB Access with Hyperledger Fabric 2.2

1. **Marshall dependencies locally** – This step uses Gradle and the Gradle shadow plugin to download dependencies to the client machine and package them for installation to the peer node.
2. **Install the chaincode** – This step packages the Java chaincode project and copies it to the Hyperledger Fabric peer node.

3. **Invoke and query the chaincode** – After, you can send query transactions or invoke transactions to modify the ledger.



Building Java Chaincode with AMB Access: Step-by-step

In a traditional installation of Hyperledger Fabric chaincode built with Java, Gradle automatically resolves dependencies by loading them locally in addition to downloading them from remote repositories. AMB Access does not allow downloaded dependencies to help ensure the security and reliability of the network. For this reason, to run Java chaincode, or any chaincode, in a Hyperledger Fabric network on AMB Access, you manually download dependencies to the client machine and add them into the chaincode project.

The steps in this section demonstrate how to use Gradle to build and deploy a [sample Java chaincode](#) project with fabric-shim@2.2.1 to a Hyperledger Fabric network on AMB Access. The [Gradle Shadow plugin](#) is needed to combine all jar files into a single jar file.

Prerequisites and Assumptions

The steps below require that you have the following:

- A Hyperledger Fabric client machine with access to the internet and access to a member's peer node, ordering service, and CA on an AMB Access network running Hyperledger Fabric v2.2 or later. For more information, see the following:
 - [Create a Hyperledger Fabric Blockchain Network on Amazon Managed Blockchain \(AMB\)](#)
 - [Create an Interface VPC Endpoint for Amazon Managed Blockchain \(AMB\) Hyperledger Fabric](#)
 - [Work with Hyperledger Fabric Peer Nodes on AMB Access](#)
 - [Step 4: Create an Amazon EC2 Instance and Set Up the Hyperledger Fabric Client](#). Optionally, modify [step 4.4](#) when launching the Docker container to establish CLI variables demonstrated in the install steps below.
- You must be a Hyperledger Fabric admin to install chaincode. For more information, see [Register and Enroll a Hyperledger Fabric Admin](#).

Step 1: Download the Gradle Shadow plugin and add it to the chaincode project

Because AMB Access peer nodes don't have internet access, the [Gradle Shadow plugin](#) needs to be downloaded manually to the client and added to the project. The plugin is downloaded from the [Gradle plugin repository](#).

1. Create a sub-directory in the Java chaincode working directory on the Hyperledger Fabric client machine where you want to download the plugin, and then switch to that directory. The example below uses a directory named `plugin`.

```
mkdir plugin
cd plugin
```

2. Download the plugin.

```
wget https://plugins.gradle.org/m2/com/github/jengelman/gradle/plugins/shadow/5.1.0/shadow-5.1.0.jar
```

3. Verify that the plugin was downloaded as expected.

```
ls

shadow-5.1.0.jar
```

4. After downloading the shadow plugin jar file, use an editor to modify the `build.gradle` file so that it references the local directory and plugin.

Make sure to remove the id 'com.github.johnrengelman.shadow' version '5.1.0' entry from the plugins section.

The following example shows an excerpt from an edited build.gradle file, with *relevant sections emphasized*:

```
buildscript {
    dependencies {
        classpath fileTree(dir: 'plugin', include:['*.jar'])
    }
}

plugins {
    id 'java'
}

group 'org.hyperledger.fabric-chaincode-java'
version '1.0-SNAPSHOT'

sourceCompatibility = 1.8

repositories {
    mavenLocal()
    mavenCentral()
}

dependencies {
    implementation 'org.hyperledger.fabric-chaincode-java:fabric-chaincode-shim:2.+
}

apply plugin: 'com.github.johnrengelman.shadow'

shadowJar {
    baseName = 'chaincode'
    version = null
    classifier = null

    manifest {
        attributes 'Main-Class': 'org.hyperledger.fabric_samples.ABstore'
    }
}
```

Step 2: Add fabric-chaincode-shim and the getDeps task to the chaincode project

Update `build.gradle` to add [fabric-chaincode-shim](#) as a dependency in the `dependencies` section. In addition, reference the task `getDeps` to download dependencies to the `libs` sub-directory on the client machine. The following example shows an excerpt from a `build.gradle` file with *relevant sections emphasized*.

```
buildscript {
    dependencies {
        classpath fileTree(dir: 'plugin', include:['*.jar'])
    }
}

plugins {
    id 'java'
}

group 'org.hyperledger.fabric-chaincode-java'
version '1.0-SNAPSHOT'

sourceCompatibility = 1.8

repositories {
    mavenLocal()
    mavenCentral()
    maven {
        url "https://hyperledger.jfrog.io/hyperledger/fabric-maven"
    }
    maven {
        url 'https://jitpack.io'
    }
}

dependencies {
    compile group: 'org.hyperledger.fabric-chaincode-java', name: 'fabric-chaincode-
shim', version: '2.2.1'
    testCompile group: 'junit', name: 'junit', version: '4.12'
}

apply plugin: 'com.github.johnrengelman.shadow'

shadowJar {
    baseName = 'chaincode'
}
```

```
version = null
classifier = null

manifest {
    attributes 'Main-Class': 'org.hyperledger.fabric_samples.ABstore'
}
}

task getDeps(type: Copy) {
    from sourceSets.main.compileClasspath
    into 'libs/'
}
```

Step 3: Download dependencies and build the project locally

Run `gradle build` and `gradle getDeps` as shown in the following examples to build the chaincode locally and download dependencies. When the `gradle getDeps` task runs, Gradle creates the `lib` sub-directory on the client machine and downloads dependencies to it.

1. Download chaincode project dependencies locally.

```
gradle getDeps
```

The command creates output similar to the following example.

```
Deprecated Gradle features were used in this build, making it incompatible with
Gradle 7.0.
Use '--warning-mode all' to show the individual deprecation warnings.
See https://docs.gradle.org/6.5/userguide/
command_line_interface.html#sec:command_line_warnings

BUILD SUCCESSFUL in 1s
1 actionable task: 1 executed
```

2. Verify that dependencies were downloaded to the libs sub-directory.

```
ls libs/

bcpkix-jdk15on-1.62.jar          error_prone_annotations-2.3.4.jar  grpc-protobuf-
lite-1.31.1.jar                jsr305-3.0.2.jar
```

```

bcprov-jdk15on-1.62.jar          fabric-chaincode-protos-2.2.1.jar  grpc-
stub-1.31.1.jar                 listenablefuture-9999.0-empty-to-avoid-conflict-with-
guava.jar
checker-compat-qual-2.5.5.jar   fabric-chaincode-shim-2.2.1.jar    gson-2.8.6.jar
                                org.everit.json.schema-1.12.1.jar
classgraph-4.8.47.jar           failureaccess-1.0.1.jar            guava-29.0-
android.jar                     protobuf-java-3.12.0.jar
commons-cli-1.4.jar             grpc-api-1.31.1.jar               handy-uri-
templates-2.1.8.jar            protobuf-java-util-3.11.1.jar
commons-collections-3.2.2.jar   grpc-context-1.31.1.jar           j2objc-
annotations-1.3.jar           proto-google-common-protos-1.17.0.jar
commons-digester-1.8.1.jar     grpc-core-1.31.1.jar              javax.annotation-
api-1.3.2.jar  re2j-1.3.jar
commons-logging-1.2.jar        grpc-netty-shaded-1.31.1.jar      joda-
time-2.10.2.jar
commons-validator-1.6.jar      grpc-protobuf-1.31.1.jar          json-20190722.jar

```

3. Build the project using Gradle.

```
gradle build
```

The command creates output similar to the following example.

```

> Task :compileJava
Note: /home/ec2-user/fabric-samples/chaincode/abstore/java/src/main/java/org/
hyperledger/fabric-samples/ABstore.java uses or overrides a deprecated API.
Note: Recompile with -Xlint:deprecation for details.

BUILD SUCCESSFUL in 5s

```

Step 4: Add dependencies downloaded in the previous step to the project

Use a text editor to update the `build.gradle` file again so that it reference the dependencies in the `libs` sub-directory. This adds the dependencies that you downloaded in the previous step to the project.

The following example shows an excerpt from a `build.gradle` file with *relevant sections emphasized*.

```

buildscript {
    dependencies {

```

```
        classpath fileTree(dir: 'plugin', include:['*.jar'])
    }
}

plugins {
    id 'java'
}

group 'org.hyperledger.fabric-chaincode-java'
version '1.0-SNAPSHOT'

sourceCompatibility = 1.8

repositories {
    mavenLocal()
    mavenCentral()
}

dependencies {
    compile fileTree(dir:'libs',includes:['*.jar'])
    testCompile group: 'junit', name: 'junit', version: '4.12'
}

apply plugin: 'com.github.johnrengelman.shadow'

shadowJar {
    baseName = 'chaincode'
    version = null
    classifier = null

    manifest {
        attributes 'Main-Class': 'org.hyperledger.fabric_samples.ABstore'
    }
}

task getDeps(type: Copy) {
    from sourceSets.main.compileClasspath
    into 'libs/'
}
```

Step 5: Use the Hyperledger Fabric CLI to install the chaincode on the peer node

Use the Hyperledger Fabric CLI to package the Java chaincode project and copy it to the peer node. The `-l java` option is required for the package command to specify that the chaincode is Java-based. The path to the Java chaincode working directory is an absolute path, which is different from chaincode written in golang.

1. Before running the following chaincode commands, establish variables for the Hyperledger Fabric CLI Docker container as shown in the following example. Alternatively, you can configure a Docker compose file to establish these variables. For more information, see [Step 4.4: Configure and Run Docker Compose to Start the Hyperledger Fabric CLI](#) in the getting started tutorial. Replace the values with those appropriate for your network. Set `CORE_PEER_ADDRESS` to the endpoint of the peer node for which you want to install the chaincode, set `CORE_PEER_LOCALMSPID` to the ID of the member that owns the peer node, and set `CORE_PEER_MSPCONFIGPATH` as shown, which is the location of the membership service provider directory on the peer node.

```
docker exec -e "CORE_PEER_TLS_ENABLED=true" \  
-e "CORE_PEER_TLS_ROOTCERT_FILE=/opt/home/managedblockchain-tls-chain.pem" \  
-e "CORE_PEER_ADDRESS=nd-6EAJ5VA43JGGNPXOUZP7Y47E4Y.m-K46ICRRXJRCGRNNS4ES4XUUS5A.n-MWY63ZJZU5HGNCMBQER7IN60IU.managedblockchain.us-east-1.amazonaws.com:30003" \  
-e "CORE_PEER_LOCALMSPID=m-K46ICRRXJRCGRNNS4ES4XUUS5A" \  
-e "CORE_PEER_MSPCONFIGPATH=/opt/home/admin-msp"
```

2. Running the Hyperledger Fabric CLI package command on the client packages the chaincode and writes the package to a file. The following example demonstrates a package command using the following flags:
 - The `-p` flag specifies the location of the chaincode on the client machine. When installing Java chaincode, this must be an absolute path.
 - The `-l` flag specifies that the chaincode language is java.
 - The `--label` flag specifies the package label, which is a human-readable description of the package.

For more information about options, see [peer lifecycle chaincode package](#) in Hyperledger Fabric documentation.

```
cli peer lifecycle chaincode package ./abstorejava.tar.gz \  
-p /opt/home/fabric-samples/chaincode/abstore/java/ \  
-l java --label abstorejava
```

```
-l java \  
--label MyLabel
```

3. The following example installs the chaincode package on the peer node.

```
cli peer lifecycle chaincode install abstorejava.tar.gz
```

4. The following example queries the installed chaincodes on the peer node.

```
cli peer lifecycle chaincode queryinstalled
```

5. Running the Hyperledger Fabric CLI `approveformyorg` command on the client approves the chaincode definition for your organization. The following example demonstrates an `approveformyorg` command.

- The `-o` flag specifies the ordering service endpoint for the member.
- The `--tls` flag specifies that communication with the ordering service uses TLS.
- The `--cafile` flag specifies the location of the certificate for the ordering service that you copied when you set up the Hyperledger Fabric admin. For more information, see [step 5.1](#) in the [Getting Started](#) tutorial.
- The `-C` flag specifies the channel on which to approve the chaincode.
- The `-n` and `-v` options establish the name and version of the chaincode.
- The `--sequence` flag specifies the sequence number of the chaincode definition for the channel.
- The `--package-id` flag specifies the identifier of the chaincode install package. This value is returned by the `queryinstalled` command in the previous step.

For more information about options, see [peer lifecycle chaincode approveformyorg](#) in Hyperledger Fabric documentation.

```
cli peer lifecycle chaincode approveformyorg \  
-o orderer.n-  
MWY63ZJZU5HGNCMBQER7IN60IU.managedblockchain.MyRegion.amazonaws.com:30001 \  
--tls --cafile /opt/home/managedblockchain-tls-chain.pem \  
-C MyHLFChannelID -n MyJavaChaincodeName -v MyCCVerNumber --sequence MySeqNumber --  
package-id MyPackageID
```

6. The following example checks whether the chaincode definition is ready to be committed on the channel.

```
cli peer lifecycle chaincode checkcommitreadiness \
-o orderer.n-
MWY63ZJZU5HGNCMBQER7IN60IU.managedblockchain.MyRegion.amazonaws.com:30001 \
--tls --cafile /opt/home/managedblockchain-tls-chain.pem \
-C MyHLFChannelID -n MyJavaChaincodeName -v MyCCVerNumber --sequence MySeqNumber
```

7. The following example commits the chaincode definition on the channel.

```
cli peer lifecycle chaincode commit \
-o orderer.n-
MWY63ZJZU5HGNCMBQER7IN60IU.managedblockchain.MyRegion.amazonaws.com:30001 \
--tls --cafile /opt/home/managedblockchain-tls-chain.pem \
-C MyHLFChannelID -n MyJavaChaincodeName -v MyCCVerNumber --sequence MySeqNumber
```

8. The following example queries the committed chaincode definitions by channel on the peer node.

```
cli peer lifecycle chaincode querycommitted \
-C MyHLFChannelID
```

Step 6: Invoke the sample chaincode to perform a transaction

Precede the following example command with the variable overrides for the CLI container as shown in the previous example.

The following example invokes the sample chaincode. The parameters are the same as the previous example, except for the JSON constructor. This constructor invokes the [the sample chaincode](#) to transfer a value of 10 from a to b.

```
cli peer chaincode invoke \
-o orderer.n-MWY63ZJZU5HGNCMBQER7IN60IU.managedblockchain.MyRegion.amazonaws.com:30001
\
--cafile /opt/home/managedblockchain-tls-chain.pem --tls \
-C one-org-channel-java \
-c '{"Args":["invoke","a", "b", "10"]}' -n javacc
```

The command should print output similar to the following.

```
2020-06-26 23:49:43.895 UTC [chaincodeCmd] chaincodeInvokeOrQuery -> INFO 003 Chaincode
invoke successful. result: status:200 message:"invoke finished successfully"
payload:"a: 90 b: 210"
```

Step 7: Invoke the sample chaincode to query

Precede the following command with the variable overrides for the CLI container as shown in the previous example.

The following example queries the ledger for the value attributed to a. The query should print 90 as output.

```
cli peer chaincode query \  
-o orderer.n-MWY63ZJZU5HGNCMBQER7IN60IU.managedblockchain.MyRegion.amazonaws.com:30001 \  
\  
--cafile /opt/home/managedblockchain-tls-chain.pem --tls \  
-C one-org-channel-java \  
-c '{"Args":["query","a"]}' -n javacc
```

Amazon Managed Blockchain (AMB) API Reference for AMB Access Hyperledger Fabric

Amazon Managed Blockchain (AMB) provides API operations for creating and managing networks, peer nodes, invitations, members, and proposals for AMB Access Hyperledger Fabric. For more information, see the [AMB Access API Reference Guide](#).

Amazon Managed Blockchain (AMB) Hyperledger Fabric Security

Amazon Managed Blockchain (AMB) utilizes AWS features and the features of the open-source framework running on AMB Access to provide data protection as well as authentication and access control.

This chapter covers security information specific to AMB Access Hyperledger Fabric. For security information specific to AMB Access Ethereum, see [AMB Access Ethereum Security](#) in the *Amazon Managed Blockchain (AMB) Access Ethereum Developer Guide*.

Topics

- [Data Protection for Amazon Managed Blockchain \(AMB\) Hyperledger Fabric](#)
- [Authentication and Access Control for AMB Access Hyperledger Fabric](#)
- [Configuring Security Groups for Amazon Managed Blockchain \(AMB\) Hyperledger Fabric](#)

Data Protection for Amazon Managed Blockchain (AMB) Hyperledger Fabric

The AWS [shared responsibility model](#) applies to data protection in Amazon Managed Blockchain (AMB) Hyperledger Fabric. As described in this model, AWS is responsible for protecting the global infrastructure that runs all of the AWS Cloud. You are responsible for maintaining control over your content that is hosted on this infrastructure. You are also responsible for the security configuration and management tasks for the AWS services that you use. For more information about data privacy, see the [Data Privacy FAQ](#). For information about data protection in Europe, see the [AWS Shared Responsibility Model and GDPR](#) blog post on the *AWS Security Blog*.

For data protection purposes, we recommend that you protect AWS account credentials and set up individual users with AWS IAM Identity Center or AWS Identity and Access Management (IAM). That way, each user is given only the permissions necessary to fulfill their job duties. We also recommend that you secure your data in the following ways:

- Use multi-factor authentication (MFA) with each account.
- Use SSL/TLS to communicate with AWS resources. We require TLS 1.2 and recommend TLS 1.3.

- Set up API and user activity logging with AWS CloudTrail. For information about using CloudTrail trails to capture AWS activities, see [Working with CloudTrail trails](#) in the *AWS CloudTrail User Guide*.
- Use AWS encryption solutions, along with all default security controls within AWS services.
- Use advanced managed security services such as Amazon Macie, which assists in discovering and securing sensitive data that is stored in Amazon S3.
- If you require FIPS 140-3 validated cryptographic modules when accessing AWS through a command line interface or an API, use a FIPS endpoint. For more information about the available FIPS endpoints, see [Federal Information Processing Standard \(FIPS\) 140-3](#).

We strongly recommend that you never put confidential or sensitive information, such as your customers' email addresses, into tags or free-form text fields such as a **Name** field. This includes when you work with AMB Access Hyperledger Fabric or other AWS services using the console, API, AWS CLI, or AWS SDKs. Any data that you enter into tags or free-form text fields used for names may be used for billing or diagnostic logs. If you provide a URL to an external server, we strongly recommend that you do not include credentials information in the URL to validate your request to that server.

Data Encryption for AMB Access Hyperledger Fabric

Data encryption helps prevent unauthorized users from reading data from a blockchain network and the associated data storage systems. This includes data saved to persistent media, known as *data at rest*, and data that may be intercepted as it travels the network, known as *data in transit*.

Topics

- [Encryption at Rest for AMB Access Hyperledger Fabric](#)
- [Encryption in Transit for AMB Access Hyperledger Fabric](#)

Encryption at Rest for AMB Access Hyperledger Fabric

Amazon Managed Blockchain (AMB) offers fully managed encryption at rest. AMB Access *encryption at rest* provides enhanced security by encrypting all data at rest on networks, members, and peer nodes using keys in AWS Key Management Service (AWS KMS). This functionality helps reduce the operational burden and complexity involved in protecting sensitive data. With encryption at rest, you can build security-sensitive blockchain applications that meet strict encryption compliance and regulatory requirements.

Encryption at rest integrates with AWS KMS for managing the encryption key that is used to protect your Hyperledger Fabric data. For more information about AWS KMS, see [AWS Key Management Service concepts](#) in the *AWS Key Management Service Developer Guide*.

In AMB Access Hyperledger Fabric, you can specify the type of AWS KMS key at the member level. This encryption key is then inherited by any peer nodes that the member creates. Data at the network level is encrypted using AWS owned keys by default at no additional cost.

When you create a new member, you can choose one of the following types of KMS keys to protect your member and its peer nodes:

- *AWS owned key* – The default encryption type, if no KMS key ARN is specified. The key is owned by AMB Access (no additional charge, and no configuration required).
- *Customer managed key* – The key is stored in your AWS account and is created, owned, and managed by you. You have full control over the key (AWS KMS charges apply).

When you access networks, members, and peer nodes, AMB Access decrypts the data transparently. You don't have to change any code or applications to use or manage encrypted data. All AMB Access operations work seamlessly on your encrypted data.

You can specify a KMS key when you create a new member by using the AWS Management Console, the AMB Access API, or the AWS Command Line Interface (AWS CLI). For more information, see the [MemberConfiguration](#) data type in the *Amazon Managed Blockchain (AMB) API Reference*. You can pass this data type as an input parameter to the [CreateMember](#) or [CreateNetwork](#) API operations.

Note

By default, Amazon Managed Blockchain (AMB) automatically enables encryption at rest using AWS owned keys at no additional charge. However, AWS KMS charges apply for using a customer managed key. For information about pricing, see [AWS Key Management Service pricing](#).

Topics

- [Encryption at Rest: How It Works](#)
- [How AMB Access Uses Grants in AWS KMS](#)

- [Encryption at Rest Considerations](#)
- [Using Customer Managed Keys in AMB Access](#)

Encryption at Rest: How It Works

AMB Access Hyperledger Fabric encryption at rest encrypts your data using 256-bit Advanced Encryption Standard (AES-256), which helps secure your data from unauthorized access to the underlying storage.

Encryption at rest integrates with AWS Key Management Service (AWS KMS) for managing the encryption key that is used to protect your Hyperledger Fabric data. When creating a new member, you can choose one of the following types of AWS KMS keys:

- [AWS owned key](#)
- [Customer managed key](#)

AWS owned key

AWS owned keys aren't stored in your AWS account. They are part of a collection of KMS keys that AWS owns and manages for use in multiple AWS accounts. AWS services can use AWS owned keys to protect your data.

You don't need to create or manage AWS owned keys. However, you can't view, use, track, or audit them. You aren't charged a monthly fee or a usage fee for AWS owned keys, and they don't count against the AWS KMS quotas for your account.

For more information, see [AWS owned keys](#) in the *AWS Key Management Service Developer Guide*.

Customer managed key

Customer managed keys are KMS keys in your AWS account that you create, own, and manage. You have full control over these KMS keys. AMB Access supports symmetric customer managed keys only.

Use a customer managed key to get the following features:

- Setting and maintaining key policies, IAM policies, and grants to control access to the key
- Enabling and disabling the key

- Rotating cryptographic material for the key
- Creating key tags and aliases
- Scheduling the key for deletion
- Importing your own key material or using a custom key store that you own and manage
- Using AWS CloudTrail and Amazon CloudWatch Logs to track the requests that AMB Access sends to AWS KMS on your behalf

For more information, see [Customer managed keys](#) in the *AWS Key Management Service Developer Guide*.

Customer managed keys [incur a charge](#) for each API call, and AWS KMS quotas apply to these KMS keys. For more information, see [AWS KMS resource or request quotas](#).

When you specify a customer managed key as the member-level KMS key, the member and its peer nodes are protected with the same customer managed key. Network-level data is encrypted using AWS owned keys by default at no additional cost.

Inaccessible Customer Managed Keys

If you disable your customer managed key, schedule the key for deletion, or revoke the grants on the key, the status of your member and its peer nodes becomes `INACCESSIBLE_ENCRYPTION_KEY`. In this state, the member and its peer nodes are impaired and might not function as expected. An inaccessible key prevents all users and the AMB Access service from encrypting or decrypting data—and from performing read and write operations on the nodes. AMB Access must have access to your KMS key to ensure that you can continue to access your nodes and to prevent data loss.

The effect of disabling or deleting a key or of revoking a grant is not immediate. It might take some time for the member and peer node resources to discover that the key is inaccessible. When a resource is in this state, we recommend deleting and recreating the resource, and reconfiguring the KMS key. To check the encryption status of a member or peer node resource, use the [GetMember](#) or [GetNode](#) API operation respectively.

How AMB Access Uses Grants in AWS KMS

AMB Access requires *grants* to use your customer managed key. When you create a member that is protected with a customer managed key, AMB Access creates grants on your behalf by sending [CreateGrant](#) requests to AWS KMS. Grants in AWS KMS are used to give AMB Access access to

a KMS key in a customer AWS account. For more information, see [Using Grants](#) in the *AWS Key Management Service Developer Guide*.

AMB Access requires the grants to use your customer managed key for the following API operations:

- [Decrypt](#)
- [DescribeKey](#)
- [Encrypt](#)
- [GenerateDataKey](#)
- [GenerateDataKeyWithoutPlaintext](#)
- [ReEncryptFrom](#)
- [ReEncryptTo](#)

AMB Access uses Amazon Elastic Compute Cloud (Amazon EC2), Amazon Elastic Block Store (Amazon EBS), and AWS Secrets Manager to call these AWS KMS operations on its behalf.

You can revoke a grant to remove the service's access to the customer managed key at any time. If you do, the key becomes inaccessible, and AMB Access loses access to any of the data protected by the customer managed key. In this state, the member and its peer node resources are impaired and might not function as expected. When a resource is in this state, we recommend deleting and recreating the resource, and reconfiguring the KMS key.

Encryption at Rest Considerations

Consider the following when you are using encryption at rest in AMB Access Hyperledger Fabric.

- Server-side encryption at rest is enabled on all AMB Access network, member, and peer node data and can't be disabled. You can't encrypt only a subset of data in these resource types.
- Encryption at rest only encrypts data while it is static (at rest) on a persistent storage media. If data security is a concern for data in transit or data in use, you might need to take additional measures as follows:
 - *Data in transit*: All your data in AMB Access is encrypted in transit. By default, communications to and from AMB Access use the HTTPS protocol, which protects network traffic by using Secure Sockets Layer (SSL)/Transport Layer Security (TLS) encryption.
 - *Data in use*: Use client-side encryption to protect your data before sending it to AMB Access.

- AMB Access currently doesn't support encryption context for AWS KMS cryptographic operations.

Using Customer Managed Keys in AMB Access

You can use the AMB Access console, the AMB Access API, or the AWS CLI to specify the AWS KMS key for new members in AMB Access. The following topics describe how to manage and monitor the usage of your customer managed keys in AMB Access.

Topics

- [Prerequisites](#)
- [Specifying a Customer Managed Key](#)
- [Monitoring Your Customer Managed Keys](#)

Prerequisites

Before you can protect AMB Access resources with a customer managed key, you must first create the key in AWS KMS. You must also specify a key policy that allows AMB Access to create grants on that AWS KMS key on your behalf.

Creating a Customer Managed Key

To create a customer managed key, follow the steps in [Creating symmetric encryption KMS keys](#) in the *AWS Key Management Service Developer Guide*. AMB Access doesn't support [asymmetric keys](#).

Setting a Key Policy

Key policies are the primary way to control access to customer managed keys in AWS KMS. Every customer managed key must have exactly one key policy. The statements in the key policy document determine who has permission to use the AWS KMS key and how they can use it. For more information, see [Using key policies in AWS KMS](#) in the *AWS Key Management Service Developer Guide*.

You can specify a key policy when you create your customer managed key. To change a key policy for an existing customer managed key, see [Changing a key policy](#).

To allow AMB Access to use your customer managed key, the key policy must include permissions for the [CreateGrant](#) API operation. This operation adds a [grant](#) to a customer managed key. Grants control access to a specified KMS key. AMB Access creates grants that

allow access to the [grant operations](#) that it requires, which are listed in [How AMB Access Uses Grants in AWS KMS](#).

Key Policy Example

The following is a key policy example that you can use for AMB Access. This policy allows principals that are authorized to use AMB Access from the account 111122223333 to call the `CreateGrant` operation on all resources.

To use this policy, replace *aws-region* and *111122223333* in the example with your own information.

Specifying a Customer Managed Key

You can specify the Amazon Resource Name (ARN) of a customer managed KMS key when you create a new member by using the AMB Access console, the AMB Access API, or the AWS CLI. The following is an example of a KMS key ARN.

```
arn:aws:kms:us-east-1:111122223333:key/1234abcd-12ab-34cd-56ef-1234567890ab
```

For more information, see the [MemberConfiguration](#) data type in the *Amazon Managed Blockchain (AMB) API Reference*. You can pass this data type as an input parameter to the [CreateMember](#) or [CreateNetwork](#) API operations.

Monitoring Your Customer Managed Keys

If you use a customer managed key to protect your Amazon Managed Blockchain (AMB) resources, you can use [AWS CloudTrail](#) or [Amazon CloudWatch Logs](#) to track the requests that AMB Access sends to AWS KMS on your behalf. For more information, see [Monitoring AWS KMS keys](#) in the *AWS Key Management Service Developer Guide*.

The following example is a CloudTrail log entry for the `CreateGrant` API operation. In addition to this event from AMB Access, you can expect AWS KMS calls for your customer managed key from Amazon EC2, Amazon EBS, and AWS Secrets Manager. These AWS services call other AWS KMS operations such as `GenerateDataKey`, `Decrypt`, and `Encrypt` on behalf of AMB Access.

CreateGrant

When you specify a customer managed key to protect your member, AMB Access sends `CreateGrant` requests to AWS KMS on your behalf to allow access to your KMS key.

The grants that AMB Access creates are specific to a member and its peer nodes. The principal in the CreateGrant request is the user who created the member or peer node.

The event that records the CreateGrant operation is similar to the following example event. The parameters include the Amazon Resource Name (ARN) of the customer managed key, the grantee principal and retiring principal (the AMB Access service), and the operations that the grant covers.

```
{
  "eventVersion": "1.08",
  "userIdentity": {
    "type": "IAMUser",
    "principalId": "AKIAIOSFODNN7EXAMPLE",
    "arn": "arn:aws:iam::111122223333:user/sample-user",
    "accountId": "111122223333",
    "accessKeyId": "AKIAI44QH8DHBEXAMPLE",
    "userName": "sample-user",
    "sessionContext": {
      "sessionIssuer": {},
      "webIdFederationData": {},
      "attributes": {
        "mfaAuthenticated": "false",
        "creationDate": "2021-06-08T16:36:03Z"
      }
    }
  },
  "invokedBy": "managedblockchain.amazonaws.com",
},
"eventTime": "2021-06-08T16:36:04Z",
"eventSource": "kms.amazonaws.com",
"eventName": "CreateGrant",
"awsRegion": "us-east-1",
"sourceIPAddress": "managedblockchain.amazonaws.com",
"userAgent": "managedblockchain.amazonaws.com",
"requestParameters": {
  "keyId": "arn:aws:kms:us-east-1:111122223333:key/1234abcd-12ab-34cd-56ef-1234567890ab",
  "granteePrincipal": "managedblockchain.us-east-1.amazonaws.com",
  "retiringPrincipal": "managedblockchain.us-east-1.amazonaws.com",
  "operations": [
    "Decrypt"
  ],
  "constraints": {
    "encryptionContextSubset": {}
  }
}
```

```
    }
  },
  "responseElements": {
    "grantId":
"64690d7d6ad23d39edd0c8fdea6400f297fae673ad84ce3563b6ca3a1e685ba2"
  },
  "requestID": "55b9f158-f537-43c2-acec-f336f41866b1",
  "eventID": "944aa59c-3bf9-4092-9ca2-435f585b41f6",
  "readOnly": false,
  "resources": [
    {
      "accountId": "111122223333",
      "type": "AWS::KMS::Key",
      "ARN": "arn:aws:kms:us-
east-1:111122223333:key/1234abcd-12ab-34cd-56ef-1234567890ab"
    }
  ],
  "eventType": "AwsApiCall",
  "managementEvent": true,
  "eventCategory": "Management",
  "recipientAccountId": "111122223333"
}
```

Encryption in Transit for AMB Access Hyperledger Fabric

The Hyperledger Fabric certificate authority (CA) in each membership provides a TLS certificate authority to secure Hyperledger Fabric communication channels in the Amazon Managed Blockchain (AMB) network. For more information, see the [Fabric CA's User Guide](#) in Hyperledger Fabric documentation.

Authentication and Access Control for AMB Access Hyperledger Fabric

AWS Identity and Access Management (IAM) permissions policies, VPC endpoint services powered by AWS PrivateLink, and Amazon EC2 security groups provide the primary means for you to control access to Amazon Managed Blockchain (AMB). In addition to these AWS services, open-source frameworks that run on AMB Access have authentication and access control features that you can configure.

IAM permissions policies are associated with AWS users in your account and determine who has access to what. Permissions policies specify the actions that each user can perform using AMB Access and other AWS services. VPC endpoint services allow each AMB Access network member to connect privately to AMB Access resources. Amazon EC2 security groups act as virtual firewalls and determine the inbound and outbound network traffic that is allowed between AMB Access resources and other Amazon EC2 resources. In AMB Access, these security groups are associated with the VPC endpoint in your account and with any framework clients that run on AWS, such as a Hyperledger Fabric client running on an Amazon EC2 instance.

Before you configure authentication and access control using AWS services and open-source features, we recommend that you review the following resources:

- For more information about IAM and IAM permissions policies, see [Identity and Access Management for Amazon Managed Blockchain \(AMB\) Hyperledger Fabric](#). We also recommend [What is IAM?](#) and [IAM JSON Policy Reference](#) in the *IAM User Guide*.
- For more information about VPC endpoints, see [Create an Interface VPC Endpoint for Amazon Managed Blockchain \(AMB\) Hyperledger Fabric](#) and [VPC Endpoints](#) in the *Amazon VPC User Guide*.
- For more information about Amazon EC2 security groups, see [Configuring Security Groups for Amazon Managed Blockchain \(AMB\) Hyperledger Fabric](#) and [Amazon EC2 Security Groups for Linux Instances](#) in the *Amazon EC2 User Guide*.
- For more information about the Hyperledger Fabric Certificate Authority (CA), see [Certificate Authority \(CA\) Setup](#) in the Hyperledger Fabric documentation.
- For more information about the supported Hyperledger Fabric 2.2 application access control lists, see [Application Access Control Lists](#) in the Hyperledger Fabric documentation.

Identity and Access Management for Amazon Managed Blockchain (AMB) Hyperledger Fabric

AWS Identity and Access Management (IAM) is an AWS service that helps an administrator securely control access to AWS resources. IAM administrators control who can be *authenticated* (signed in) and *authorized* (have permissions) to use AMB Access Hyperledger Fabric resources. IAM is an AWS service that you can use with no additional charge.

Topics

- [Audience](#)
- [Authenticating with identities](#)
- [Managing access using policies](#)
- [How Amazon Managed Blockchain \(AMB\) Hyperledger Fabric works with IAM](#)
- [Troubleshooting Amazon Managed Blockchain \(AMB\) Hyperledger Fabric identity and access](#)
- [Amazon Managed Blockchain \(AMB\) Hyperledger Fabric Identity-Based Policy Examples](#)
- [Example IAM Role Permissions Policy for Hyperledger Fabric Client EC2 Instance](#)
- [Using Service-Linked Roles for AMB Access](#)

Audience

How you use AWS Identity and Access Management (IAM) differs based on your role:

- **Service user** - request permissions from your administrator if you cannot access features (see [Troubleshooting Amazon Managed Blockchain \(AMB\) Hyperledger Fabric identity and access](#))
- **Service administrator** - determine user access and submit permission requests (see [How Amazon Managed Blockchain \(AMB\) Hyperledger Fabric works with IAM](#))
- **IAM administrator** - write policies to manage access (see [Amazon Managed Blockchain \(AMB\) Hyperledger Fabric Identity-Based Policy Examples](#))

Authenticating with identities

Authentication is how you sign in to AWS using your identity credentials. You must be authenticated as the AWS account root user, an IAM user, or by assuming an IAM role.

You can sign in as a federated identity using credentials from an identity source like AWS IAM Identity Center (IAM Identity Center), single sign-on authentication, or Google/Facebook credentials. For more information about signing in, see [How to sign in to your AWS account](#) in the *AWS Sign-In User Guide*.

For programmatic access, AWS provides an SDK and CLI to cryptographically sign requests. For more information, see [AWS Signature Version 4 for API requests](#) in the *IAM User Guide*.

AWS account root user

When you create an AWS account, you begin with one sign-in identity called the AWS account *root user* that has complete access to all AWS services and resources. We strongly recommend that you

don't use the root user for everyday tasks. For tasks that require root user credentials, see [Tasks that require root user credentials](#) in the *IAM User Guide*.

Federated identity

As a best practice, require human users to use federation with an identity provider to access AWS services using temporary credentials.

A *federated identity* is a user from your enterprise directory, web identity provider, or Directory Service that accesses AWS services using credentials from an identity source. Federated identities assume roles that provide temporary credentials.

For centralized access management, we recommend AWS IAM Identity Center. For more information, see [What is IAM Identity Center?](#) in the *AWS IAM Identity Center User Guide*.

IAM users and groups

An *IAM user* is an identity with specific permissions for a single person or application. We recommend using temporary credentials instead of IAM users with long-term credentials. For more information, see [Require human users to use federation with an identity provider to access AWS using temporary credentials](#) in the *IAM User Guide*.

An *IAM group* specifies a collection of IAM users and makes permissions easier to manage for large sets of users. For more information, see [Use cases for IAM users](#) in the *IAM User Guide*.

IAM roles

An *IAM role* is an identity with specific permissions that provides temporary credentials. You can assume a role by [switching from a user to an IAM role \(console\)](#) or by calling an AWS CLI or AWS API operation. For more information, see [Methods to assume a role](#) in the *IAM User Guide*.

IAM roles are useful for federated user access, temporary IAM user permissions, cross-account access, cross-service access, and applications running on Amazon EC2. For more information, see [Cross account resource access in IAM](#) in the *IAM User Guide*.

Managing access using policies

You control access in AWS by creating policies and attaching them to AWS identities or resources. A policy defines permissions when associated with an identity or resource. AWS evaluates these policies when a principal makes a request. Most policies are stored in AWS as JSON documents. For

more information about JSON policy documents, see [Overview of JSON policies](#) in the *IAM User Guide*.

Using policies, administrators specify who has access to what by defining which **principal** can perform **actions** on what **resources**, and under what **conditions**.

By default, users and roles have no permissions. An IAM administrator creates IAM policies and adds them to roles, which users can then assume. IAM policies define permissions regardless of the method used to perform the operation.

Identity-based policies

Identity-based policies are JSON permissions policy documents that you attach to an identity (user, group, or role). These policies control what actions identities can perform, on which resources, and under what conditions. To learn how to create an identity-based policy, see [Define custom IAM permissions with customer managed policies](#) in the *IAM User Guide*.

Identity-based policies can be *inline policies* (embedded directly into a single identity) or *managed policies* (standalone policies attached to multiple identities). To learn how to choose between managed and inline policies, see [Choose between managed policies and inline policies](#) in the *IAM User Guide*.

Resource-based policies

Resource-based policies are JSON policy documents that you attach to a resource. Examples include IAM *role trust policies* and Amazon S3 *bucket policies*. In services that support resource-based policies, service administrators can use them to control access to a specific resource. You must [specify a principal](#) in a resource-based policy.

Resource-based policies are inline policies that are located in that service. You can't use AWS managed policies from IAM in a resource-based policy.

Other policy types

AWS supports additional policy types that can set the maximum permissions granted by more common policy types:

- **Permissions boundaries** – Set the maximum permissions that an identity-based policy can grant to an IAM entity. For more information, see [Permissions boundaries for IAM entities](#) in the *IAM User Guide*.

- **Service control policies (SCPs)** – Specify the maximum permissions for an organization or organizational unit in AWS Organizations. For more information, see [Service control policies](#) in the *AWS Organizations User Guide*.
- **Resource control policies (RCPs)** – Set the maximum available permissions for resources in your accounts. For more information, see [Resource control policies \(RCPs\)](#) in the *AWS Organizations User Guide*.
- **Session policies** – Advanced policies passed as a parameter when creating a temporary session for a role or federated user. For more information, see [Session policies](#) in the *IAM User Guide*.

Multiple policy types

When multiple types of policies apply to a request, the resulting permissions are more complicated to understand. To learn how AWS determines whether to allow a request when multiple policy types are involved, see [Policy evaluation logic](#) in the *IAM User Guide*.

How Amazon Managed Blockchain (AMB) Hyperledger Fabric works with IAM

Before you use IAM to manage access to AMB Access Hyperledger Fabric, learn what IAM features are available to use with AMB Access Hyperledger Fabric.

IAM features you can use with Amazon Managed Blockchain (AMB) Hyperledger Fabric

IAM feature	AMB Access Hyperledger Fabric support
Identity-based policies	Yes
Resource-based policies	No
Policy actions	Yes
Policy resources	Yes
Policy condition keys (service-specific)	No
ACLs	No
ABAC (tags in policies)	Yes
Temporary credentials	Yes

IAM feature	AMB Access Hyperledger Fabric support
Principal permissions	Yes
Service roles	Yes
Service-linked roles	Yes

To get a high-level view of how AMB Access Hyperledger Fabric and other AWS services work with most IAM features, see [AWS services that work with IAM](#) in the *IAM User Guide*.

Identity-based policies for AMB Access Hyperledger Fabric

Supports identity-based policies: Yes

Identity-based policies are JSON permissions policy documents that you can attach to an identity, such as an IAM user, group of users, or role. These policies control what actions users and roles can perform, on which resources, and under what conditions. To learn how to create an identity-based policy, see [Define custom IAM permissions with customer managed policies](#) in the *IAM User Guide*.

With IAM identity-based policies, you can specify allowed or denied actions and resources as well as the conditions under which actions are allowed or denied. To learn about all of the elements that you can use in a JSON policy, see [IAM JSON policy elements reference](#) in the *IAM User Guide*.

Identity-based policy examples for AMB Access Hyperledger Fabric

To view examples of AMB Access Hyperledger Fabric identity-based policies, see [Amazon Managed Blockchain \(AMB\) Hyperledger Fabric Identity-Based Policy Examples](#).

Resource-based policies within AMB Access Hyperledger Fabric

Supports resource-based policies: No

Resource-based policies are JSON policy documents that you attach to a resource. Examples of resource-based policies are IAM *role trust policies* and Amazon S3 *bucket policies*. In services that support resource-based policies, service administrators can use them to control access to a specific resource. For the resource where the policy is attached, the policy defines what actions a specified principal can perform on that resource and under what conditions. You must [specify a principal](#)

in a resource-based policy. Principals can include accounts, users, roles, federated users, or AWS services.

To enable cross-account access, you can specify an entire account or IAM entities in another account as the principal in a resource-based policy. For more information, see [Cross account resource access in IAM](#) in the *IAM User Guide*.

Policy actions for AMB Access Hyperledger Fabric

Supports policy actions: Yes

Administrators can use AWS JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The **Action** element of a JSON policy describes the actions that you can use to allow or deny access in a policy. Include actions in a policy to grant permissions to perform the associated operation.

To see a list of AMB Access Hyperledger Fabric actions, see [Actions defined by Amazon Managed Blockchain \(AMB\) Hyperledger Fabric](#) in the *Service Authorization Reference*.

Policy actions in AMB Access Hyperledger Fabric use the following prefix before the action:

```
managedblockchain:
```

For example, to grant someone permission to create a node with the AMB Access CreateNode API operation, you include the `managedblockchain:CreateNode` action in their policy. Policy statements must include either an **Action** or **NotAction** element. AMB Access Hyperledger Fabric defines its own set of actions that describe tasks that you can perform with this service.

To specify multiple actions in a single statement, separate them with commas.

```
"Action": [  
  "managedblockchain:action1",  
  "managedblockchain:action2"  
]
```

You can specify multiple actions using wildcards (*). For example, to specify all actions that begin with the word `Describe`, include the following action:

```
"Action": "managedblockchain::List*"
```

To view examples of AMB Access Hyperledger Fabric identity-based policies, see [Amazon Managed Blockchain \(AMB\) Hyperledger Fabric Identity-Based Policy Examples](#).

Policy resources for AMB Access Hyperledger Fabric

Supports policy resources: Yes

Administrators can use AWS JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The Resource JSON policy element specifies the object or objects to which the action applies. As a best practice, specify a resource using its [Amazon Resource Name \(ARN\)](#). For actions that don't support resource-level permissions, use a wildcard (*) to indicate that the statement applies to all resources.

```
"Resource": "*"
```

AMB Access resource types that can be used in IAM permissions policy statements for resources on Ethereum networks include the following:

- network
- member
- node
- proposal
- invitation

Members, nodes, and invitations are associated with your account. Networks and proposals, on the other hand, are scoped to the entire blockchain network and are not associated with a particular account.

For example a network resource on AMB Access has the following ARN.

```
arn:${Partition}:managedblockchain:${Region}::networks/${NetworkId}
```

For example, to specify the n-MWY63ZJZU5HGNCMBQER7IN60IU network in your statement, use the following ARN.

```
"Resource": "arn:aws:managedblockchain:us-east-1::networks/n-  
MWY63ZJZU5HGNCMBQER7IN60IU"
```

To specify any network that is visible to your account, use the wildcard (*).

```
"Resource": "arn:aws:managedblockchain:us-east-1::networks/*"
```

Some AMB Access Hyperledger Fabric actions, such as `CreateNetwork`, `ListInvitations`, and `ListNetworks` cannot be performed on a specific resource. In those cases, you must use the wildcard (*).

```
"Resource": ""
```

To see a list of AMB Access Hyperledger Fabric resource types and their ARNs, see [Resources defined by Amazon Managed Blockchain \(AMB\) Hyperledger Fabric](#) in the *Service Authorization Reference*. To learn with which actions you can specify the ARN of each resource, see [Actions defined by Amazon Managed Blockchain \(AMB\) Hyperledger Fabric](#).

To view examples of AMB Access Hyperledger Fabric identity-based policies, see [Amazon Managed Blockchain \(AMB\) Hyperledger Fabric Identity-Based Policy Examples](#).

Policy condition keys for AMB Access Hyperledger Fabric

Supports service-specific policy condition keys: No

Administrators can use AWS JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The `Condition` element specifies when statements execute based on defined criteria. You can create conditional expressions that use [condition operators](#), such as equals or less than, to match the condition in the policy with values in the request. To see all AWS global condition keys, see [AWS global condition context keys](#) in the *IAM User Guide*.

Note

AMB Access Hyperledger Fabric does not provide any service-specific condition keys, but it does support using some AWS global condition keys.

To see a list of the AWS global condition keys supported, see [Condition keys for Amazon Managed Blockchain \(AMB\) Hyperledger Fabric](#) in the *Service Authorization Reference*. To learn with which actions and resources you can use a condition key, see [Actions defined by Amazon Managed Blockchain \(AMB\) Hyperledger Fabric](#).

To view examples of AMB Access Hyperledger Fabric identity-based policies, see [Amazon Managed Blockchain \(AMB\) Hyperledger Fabric Identity-Based Policy Examples](#).

ACLs in AMB Access Hyperledger Fabric

Supports ACLs: No

Access control lists (ACLs) control which principals (account members, users, or roles) have permissions to access a resource. ACLs are similar to resource-based policies, although they do not use the JSON policy document format.

ABAC with AMB Access Hyperledger Fabric

Supports ABAC (tags in policies): Yes

Attribute-based access control (ABAC) is an authorization strategy that defines permissions based on attributes called tags. You can attach tags to IAM entities and AWS resources, then design ABAC policies to allow operations when the principal's tag matches the tag on the resource.

To control access based on tags, you provide tag information in the [condition element](#) of a policy using the `aws:ResourceTag/key-name`, `aws:RequestTag/key-name`, or `aws:TagKeys` condition keys.

If a service supports all three condition keys for every resource type, then the value is **Yes** for the service. If a service supports all three condition keys for only some resource types, then the value is **Partial**.

For more information about ABAC, see [Define permissions with ABAC authorization](#) in the *IAM User Guide*. To view a tutorial with steps for setting up ABAC, see [Use attribute-based access control \(ABAC\)](#) in the *IAM User Guide*.

To control access based on tags, you provide tag information in the [condition element](#) of a policy using the `managedblockchain::ResourceTag/key-name`, `aws:RequestTag/key-name`, or `aws:TagKeys` condition keys. For more information about tagging AMB Access Hyperledger Fabric resources, see [Tagging Amazon Managed Blockchain \(AMB\) resources](#).

To view example identity-based policies for allowing or denying access to resources and actions based on tags, see [Controlling access using tags](#).

Using temporary credentials with AMB Access Hyperledger Fabric

Supports temporary credentials: Yes

Temporary credentials provide short-term access to AWS resources and are automatically created when you use federation or switch roles. AWS recommends that you dynamically generate temporary credentials instead of using long-term access keys. For more information, see [Temporary security credentials in IAM](#) and [AWS services that work with IAM](#) in the *IAM User Guide*.

Cross-service principal permissions for AMB Access Hyperledger Fabric

Supports forward access sessions (FAS): Yes

Forward access sessions (FAS) use the permissions of the principal calling an AWS service, combined with the requesting AWS service to make requests to downstream services. For policy details when making FAS requests, see [Forward access sessions](#).

Service roles for AMB Access Hyperledger Fabric

Supports service roles: Yes

A service role is an [IAM role](#) that a service assumes to perform actions on your behalf. An IAM administrator can create, modify, and delete a service role from within IAM. For more information, see [Create a role to delegate permissions to an AWS service](#) in the *IAM User Guide*.

Warning

Changing the permissions for a service role might break AMB Access Hyperledger Fabric functionality. Edit service roles only when AMB Access Hyperledger Fabric provides guidance to do so.

Service-linked roles for AMB Access Hyperledger Fabric

Supports service-linked roles: Yes

A service-linked role is a type of service role that is linked to an AWS service. The service can assume the role to perform an action on your behalf. Service-linked roles appear in your AWS

account and are owned by the service. An IAM administrator can view, but not edit the permissions for service-linked roles.

For details about creating or managing service-linked roles, see [AWS services that work with IAM](#). Find a service in the table that includes a Yes in the **Service-linked role** column. Choose the **Yes** link to view the service-linked role documentation for that service.

Troubleshooting Amazon Managed Blockchain (AMB) Hyperledger Fabric identity and access

Use the following information to help you diagnose and fix common issues that you might encounter when working with AMB Access Hyperledger Fabric and IAM.

Topics

- [I am not authorized to perform an action in AMB Access Hyperledger Fabric](#)
- [I am not authorized to perform iam:PassRole](#)
- [I want to allow people outside of my AWS account to access my AMB Access Hyperledger Fabric resources](#)

I am not authorized to perform an action in AMB Access Hyperledger Fabric

If you receive an error that you're not authorized to perform an action, your policies must be updated to allow you to perform the action.

The following example error occurs when the `mateojackson` IAM user tries to use the console to view details about a fictional `my-example-widget` resource but doesn't have the fictional `managedblockchain::GetWidget` permissions.

```
User: arn:aws:iam::123456789012:user/mateojackson is not authorized to perform:
managedblockchain::GetWidget on resource: my-example-widget
```

In this case, the policy for the `mateojackson` user must be updated to allow access to the `my-example-widget` resource by using the `managedblockchain::GetWidget` action.

If you need help, contact your AWS administrator. Your administrator is the person who provided you with your sign-in credentials.

I am not authorized to perform iam:PassRole

If you receive an error that you're not authorized to perform the `iam:PassRole` action, your policies must be updated to allow you to pass a role to AMB Access Hyperledger Fabric.

Some AWS services allow you to pass an existing role to that service instead of creating a new service role or service-linked role. To do this, you must have permissions to pass the role to the service.

The following example error occurs when an IAM user named `marymajor` tries to use the console to perform an action in AMB Access Hyperledger Fabric. However, the action requires the service to have permissions that are granted by a service role. Mary does not have permissions to pass the role to the service.

```
User: arn:aws:iam::123456789012:user/marymajor is not authorized to perform:
iam:PassRole
```

In this case, Mary's policies must be updated to allow her to perform the `iam:PassRole` action.

If you need help, contact your AWS administrator. Your administrator is the person who provided you with your sign-in credentials.

I want to allow people outside of my AWS account to access my AMB Access Hyperledger Fabric resources

You can create a role that users in other accounts or people outside of your organization can use to access your resources. You can specify who is trusted to assume the role. For services that support resource-based policies or access control lists (ACLs), you can use those policies to grant people access to your resources.

To learn more, consult the following:

- To learn whether AMB Access Hyperledger Fabric supports these features, see [How Amazon Managed Blockchain \(AMB\) Hyperledger Fabric works with IAM](#).
- To learn how to provide access to your resources across AWS accounts that you own, see [Providing access to an IAM user in another AWS account that you own](#) in the *IAM User Guide*.
- To learn how to provide access to your resources to third-party AWS accounts, see [Providing access to AWS accounts owned by third parties](#) in the *IAM User Guide*.

- To learn how to provide access through identity federation, see [Providing access to externally authenticated users \(identity federation\)](#) in the *IAM User Guide*.
- To learn the difference between using roles and resource-based policies for cross-account access, see [Cross account resource access in IAM](#) in the *IAM User Guide*.

Amazon Managed Blockchain (AMB) Hyperledger Fabric Identity-Based Policy Examples

By default, IAM users and roles don't have permission to create or modify AMB Access Hyperledger Fabric resources. They also can't perform tasks using the AWS Management Console, AWS CLI, or AWS API. An IAM administrator must create IAM policies that grant users and roles permission to perform specific API operations on the specified resources they need. The administrator must then attach those policies to the IAM users or groups that require those permissions.

To learn how to create an IAM identity-based policy using these example JSON policy documents, see [Creating Policies on the JSON Tab](#) in the *IAM User Guide*.

Topics

- [Policy Best Practices](#)
- [Allow Users to View Their Own Permissions](#)
- [Using the AMB Access Hyperledger Fabric Console](#)
- [Performing All AMB Access Actions on All Accessible Networks for an AWS Account](#)
- [Controlling access using tags](#)

Policy Best Practices

Identity-based policies determine whether someone can create, access, or delete AMB Access Hyperledger Fabric resources in your account. These actions can incur costs for your AWS account. When you create or edit identity-based policies, follow these guidelines and recommendations:

- **Get started with AWS managed policies and move toward least-privilege permissions** – To get started granting permissions to your users and workloads, use the *AWS managed policies* that grant permissions for many common use cases. They are available in your AWS account. We recommend that you reduce permissions further by defining AWS customer managed policies

that are specific to your use cases. For more information, see [AWS managed policies](#) or [AWS managed policies for job functions](#) in the *IAM User Guide*.

- **Apply least-privilege permissions** – When you set permissions with IAM policies, grant only the permissions required to perform a task. You do this by defining the actions that can be taken on specific resources under specific conditions, also known as *least-privilege permissions*. For more information about using IAM to apply permissions, see [Policies and permissions in IAM](#) in the *IAM User Guide*.
- **Use conditions in IAM policies to further restrict access** – You can add a condition to your policies to limit access to actions and resources. For example, you can write a policy condition to specify that all requests must be sent using SSL. You can also use conditions to grant access to service actions if they are used through a specific AWS service, such as CloudFormation. For more information, see [IAM JSON policy elements: Condition](#) in the *IAM User Guide*.
- **Use IAM Access Analyzer to validate your IAM policies to ensure secure and functional permissions** – IAM Access Analyzer validates new and existing policies so that the policies adhere to the IAM policy language (JSON) and IAM best practices. IAM Access Analyzer provides more than 100 policy checks and actionable recommendations to help you author secure and functional policies. For more information, see [Validate policies with IAM Access Analyzer](#) in the *IAM User Guide*.
- **Require multi-factor authentication (MFA)** – If you have a scenario that requires IAM users or a root user in your AWS account, turn on MFA for additional security. To require MFA when API operations are called, add MFA conditions to your policies. For more information, see [Secure API access with MFA](#) in the *IAM User Guide*.

For more information about best practices in IAM, see [Security best practices in IAM](#) in the *IAM User Guide*.

Allow Users to View Their Own Permissions

This example shows how you might create a policy that allows IAM users to view the inline and managed policies that are attached to their user identity. This policy includes permissions to complete this action on the console or programmatically using the AWS CLI or AWS API.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "ViewOwnUserInfo",
```

```

    "Effect": "Allow",
    "Action": [
        "iam:GetUserPolicy",
        "iam:ListGroupsWithUser",
        "iam:ListAttachedUserPolicies",
        "iam:ListUserPolicies",
        "iam:GetUser"
    ],
    "Resource": ["arn:aws:iam::*:user/${aws:username}"]
},
{
    "Sid": "NavigateInConsole",
    "Effect": "Allow",
    "Action": [
        "iam:GetGroupPolicy",
        "iam:GetPolicyVersion",
        "iam:GetPolicy",
        "iam:ListAttachedGroupPolicies",
        "iam:ListGroupPolicies",
        "iam:ListPolicyVersions",
        "iam:ListPolicies",
        "iam:ListUsers"
    ],
    "Resource": "*"
}
]
}

```

Using the AMB Access Hyperledger Fabric Console

To access the Amazon Managed Blockchain (AMB) Hyperledger Fabric console, you must have a minimum set of permissions. These permissions must allow you to list and view details about the AMB Access Hyperledger Fabric resources in your AWS account. If you create an identity-based policy that is more restrictive than the minimum required permissions, the console won't function as intended for entities (IAM users or roles) with that policy.

To ensure that those entities can still use the AMB Access Hyperledger Fabric console, also attach the following AWS managed policy to the entities.

```
AmazonManagedBlockchainConsoleFullAccess
```

For more information, see [Adding Permissions to a User](#) in the *IAM User Guide*.

You don't need to allow minimum console permissions for users that are making calls only to the AWS CLI or the AWS API. Instead, allow access to only the actions that match the API operation that you're trying to perform.

Performing All AMB Access Actions on All Accessible Networks for an AWS Account

The following example shows you how to grant an IAM user in AWS account 123456789012 access to all the network and member resources in the account in the us-east-1 Region. This example shows how the user is permitted to list and create; manage networks, proposals, members, and nodes; and reject invitations to join other networks.

Note

- Deny and Allow actions on a resource (for example Network) only affect actions on the resource itself and *do not* apply to child resources such as Member and Node.
- Child resources (such as nodes, members and invitations) must have AWS account ids in their policy so they can be enforced as shown in the following example.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "ManageNetworkResources",
      "Effect": "Allow",
      "Action": [
        "managedblockchain:CreateNetwork",
        "managedblockchain:GetNetwork",
        "managedblockchain:CreateProposal",
        "managedblockchain:ListProposals",
        "managedblockchain:CreateMember",
        "managedblockchain:ListMembers"
      ],
      "Resource": [
        "arn:aws:managedblockchain:us-east-1::networks/*"
      ]
    }
  ]
}
```

```
    "Sid": "ManageMemberResources",
    "Effect": "Allow",
    "Action": [
        "managedblockchain:GetMember",
        "managedblockchain>DeleteMember",
        "managedblockchain>CreateNode",
        "managedblockchain:ListNodes"
    ],
    "Resource": [
        "arn:aws:managedblockchain:us-east-1:123456789012:members/*"
    ]
},

{
    "Sid": "ManageNodeResources",
    "Effect": "Allow",
    "Action": [
        "managedblockchain:GetNode",
        "managedblockchain>DeleteNode",
        "managedblockchain:UpdateNode"
    ],
    "Resource": [
        "arn:aws:managedblockchain:us-east-1:123456789012:nodes/*"
    ]
},

{
    "Sid": "ManageProposalResources",
    "Effect": "Allow",
    "Action": [
        "managedblockchain:GetProposal",
        "managedblockchain:ListProposalVotes",
        "managedblockchain:VoteOnProposal"
    ],
    "Resource": [
        "arn:aws:managedblockchain:us-east-1:proposals/*"
    ]
},

{
    "Sid": "ManageInvitationResources",
    "Effect": "Allow",
    "Action": [
        "managedblockchain:RejectInvitation"
```

```

    ],
    "Resource": [
      "arn:aws:managedblockchain:us-east-1:123456789012:invitations/*"
    ]
  }
]
}

```

Controlling access using tags

The following example policy statements demonstrate how you can use tags to limit access to AMB Access Hyperledger Fabric resources and actions performed on those resources.

To tag resources during creation, policy statements that allow the create action for the resource as well as the TagResource action must be attached to the IAM principal performing the operation.

Note

This topic includes examples of policy statements with a Deny effect. Each policy statement assumes that a statement with a broader Allow effect for the same actions exists; the Deny policy statement restricts that otherwise overly-permissive allow statement.

Example– Require a tag key with either of two values to be added during network creation

The following identity-based policy statements allow the IAM principal to create a network only if the network is created with the tags specified.

The statement with the Sid set to RequireTag specifies that network creation is allowed only if the network is created with a tag that has a tag key of department and a value of either sales or marketing; otherwise, the create network operation fails.

The statement with Sid set to AllowTagging allows the IAM principal to tag networks, which are not associated with an AWS account ID. It also allows tagging of members in the specified AWS account, 111122223333.

JSON

```
{
```

```

"Version": "2012-10-17",
"Statement": [
  {
    "Sid": "RequireTag",
    "Effect": "Allow",
    "Action": [
      "managedblockchain:CreateNetwork"
    ],
    "Resource": [
      "*"
    ],
    "Condition": {
      "StringLike": {
        "aws:RequestTag/department": [
          "sales",
          "marketing"
        ]
      }
    }
  },
  {
    "Sid": "AllowTagging",
    "Effect": "Allow",
    "Action": [
      "managedblockchain:TagResource"
    ],
    "Resource": [
      "arn:aws:managedblockchain:us-east-1::networks/*",
      "arn:aws:managedblockchain:us-east-1:111122223333:members/*"
    ]
  }
]
}

```

Example– Deny access to networks that have a specific tag key

The following identity-based policy statement denies the IAM principal the ability to retrieve or view information for networks that have a tag with a tag key of restricted.

JSON

```
{
```

```

"Version": "2012-10-17",
"Statement": [
  {
    "Sid": "DenyRestrictedNetwork",
    "Effect": "Deny",
    "Action": [
      "managedblockchain:GetNetwork"
    ],
    "Resource": [
      "*"
    ],
    "Condition": {
      "StringLike": {
        "aws:ResourceTag/restricted": [
          "*"
        ]
      }
    }
  }
]
}

```

Example– Deny member creation for networks with a specific tag key and value

The following identity-based policy statement denies the IAM principal from creating a member on any network that has a tag with a tag key of department with the value of accounting.

JSON

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "DenyMemberCreation",
      "Effect": "Deny",
      "Action": [
        "managedblockchain:CreateMember"
      ],
      "Resource": [
        "*"
      ]
    }
  ]
}

```

```

    ],
    "Condition": {
      "StringEquals": {
        "aws:ResourceTag/department": "accounting"
      }
    }
  }
]
}

```

Example– Allow member creation only if a specific tag key and value are added

The following identity-based policy statements allow the IAM principal to create members in the account 111122223333 only if the member is created with a tag that has the tag key of department and a tag value of accounting; otherwise, the create member operation fails.

JSON

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "RestrictMemberCreation",
      "Effect": "Allow",
      "Action": [
        "managedblockchain:CreateMember"
      ],
      "Resource": [
        "*"
      ],
      "Condition": {
        "StringEquals": {
          "aws:RequestTag/department": "accounting"
        }
      }
    },
    {
      "Sid": "AllowTaggingOfMembers",
      "Effect": "Allow",
      "Action": [

```

```

        "managedblockchain:TagResource"
    ],
    "Resource": [
        "arn:aws:managedblockchain:us-east-1:111122223333:members/*"
    ]
}
]
}

```

Example– Allow access only to members that have a specified tag key

The following identity-based policy statement allows the IAM principal to retrieve or view information about members only if the member has a tag with a tag key of `unrestricted`.

JSON

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "",
      "Effect": "Allow",
      "Action": [
        "managedblockchain:GetMember"
      ],
      "Resource": [
        "*"
      ],
      "Condition": {
        "StringLike": {
          "aws:ResourceTag/unrestricted": [
            "*"
          ]
        }
      }
    }
  ]
}

```

Example– Deny node creation if a specific tag key and value are added during creation

The following identity-based policy statement denies the IAM principal the ability to create a node if a tag with the tag key of `department` and a value of `analytics` is added during creation.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "DenyCreateNodeWithTag",
      "Effect": "Deny",
      "Action": [
        "managedblockchain:CreateNode"
      ],
      "Resource": [
        "*"
      ],
      "Condition": {
        "StringEquals": {
          "aws:RequestTag/department": [
            "analytics"
          ]
        }
      }
    }
  ]
}
```

Example– Deny node creation for members with a specific tag key and value

The following identity-based policy statement denies the IAM principal the ability to create a node if the member to which the node will belong has a tag with a tag key of `department` and a tag value of `analytics`.

JSON

```
{
  "Version": "2012-10-17",
```

```

    "Statement": [
      {
        "Sid": "DenyCreateNodeForTaggedMember",
        "Effect": "Deny",
        "Action": [
          "managedblockchain:CreateNode"
        ],
        "Resource": [
          "*"
        ],
        "Condition": {
          "StringEquals": {
            "aws:ResourceTag/department": [
              "analytics"
            ]
          }
        }
      }
    ]
  }
}

```

Example– Allow access only to nodes with a specific tag key and value

The following identity-based policy statement allows the IAM principal to retrieve or view information about nodes only if the node has a tag with a tag key of department and a tag value of sales.

JSON

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowTaggedNodeAccess",
      "Effect": "Allow",
      "Action": [
        "managedblockchain:GetNode"
      ],
      "Resource": [
        "*"
      ],
      "Condition": {
        "StringEquals": {
          "aws:ResourceTag/department": [
            "sales"
          ]
        }
      }
    }
  ]
}

```

```

        "Condition": {
            "StringEquals": {
                "aws:ResourceTag/department": [
                    "sales"
                ]
            }
        }
    ]
}

```

Example– Deny access to nodes with a specific tag key

The following identity-based policy statement denies the IAM principal the ability to retrieve or view information about a node if the node has a tag with a tag key of `restricted`.

JSON

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "DenyAccessToTaggedNodes",
      "Effect": "Deny",
      "Action": [
        "managedblockchain:GetNode"
      ],
      "Resource": [
        "*"
      ],
      "Condition": {
        "StringLike": {
          "aws:ResourceTag/restricted": [
            "*"
          ]
        }
      }
    }
  ]
}

```

Example– Allow proposal creation only for networks with a specific tag key and value

The following identity-based policy statement allows the IAM principal to create a proposal only if the network for which the proposal is made has a tag with a tag key of `department` and a value of `accounting`.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowCreateProposalOnlyForTaggedNetwork",
      "Effect": "Allow",
      "Action": [
        "managedblockchain:CreateProposal"
      ],
      "Resource": [
        "*"
      ],
      "Condition": {
        "StringEquals": {
          "aws:ResourceTag/department": [
            "accounting"
          ]
        }
      }
    }
  ]
}
```

Example– Allow proposal creation only if a specific tag key and value are added

The following identity-based policy statements allow the IAM principal to create a proposal for inviting the specified account `123456789012` to the network only if a tag with a tag key of `consortium` and a tag value of `exampleconsortium` is added during proposal creation.

JSON

```
{
```

```

"Version": "2012-10-17",
"Statement": [
  {
    "Sid": "RequireTagWithProposal",
    "Effect": "Allow",
    "Action": [
      "managedblockchain:CreateProposal"
    ],
    "Resource": [
      "*"
    ],
    "Condition": {
      "StringEquals": {
        "aws:RequestTag/consortium": [
          "exampleconsortium"
        ]
      }
    }
  },
  {
    "Sid": "AllowProposalAndInvitationTagging",
    "Effect": "Allow",
    "Action": [
      "managedblockchain:TagResource"
    ],
    "Resource": [
      "arn:aws:managedblockchain:us-east-1:proposals/*",
      "arn:aws:managedblockchain:us-east-1:123456789012:invitations/*"
    ]
  }
]
}

```

Example– Allow access only to proposals with a specific tag key and value

The following identity-based policy statement allows the IAM principal to retrieve and view proposals only if the proposal has a tag with a tag key of `consortium` and a tag value of `exampleconsortium`.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowOnlyTaggedProposalAccess",
      "Effect": "Allow",
      "Action": [
        "managedblockchain:GetProposal"
      ],
      "Resource": [
        "*"
      ],
      "Condition": {
        "StringEquals": {
          "aws:ResourceTag/consortium": [
            "exampleconsortium"
          ]
        }
      }
    }
  ]
}
```

Example– Deny the ability to vote on proposals based on a specific tag key and value

The following identity-based policy statement denies the IAM principal the ability to vote on proposals that have a tag with the tag key of `consortium` and a tag value of `exampleconsortium`.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "DenyVoteOnTaggedProposal",
      "Effect": "Deny",
      "Action": [
```

```

        "managedblockchain:VoteOnProposal"
    ],
    "Resource": [
        "*"
    ],
    "Condition": {
        "StringEquals": {
            "aws:ResourceTag/consortium": [
                "exampleconsortium"
            ]
        }
    }
}

```

Example– Deny the ability to reject an invitation with a specific tag and value

The following identity-based policy statement denies the IAM principal the ability to reject on invitation that has a tag with the tag key of `consortium` and a tag value of `exampleconsortium`.

JSON

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "1",
      "Effect": "Deny",
      "Action": [
        "managedblockchain:RejectInvitation"
      ],
      "Resource": [
        "*"
      ],
      "Condition": {
        "StringEquals": {
          "aws:ResourceTag/consortium": [
            "exampleconsortium"
          ]
        }
      }
    }
  ]
}

```

```

    }
  ]
}

```

Example IAM Role Permissions Policy for Hyperledger Fabric Client EC2 Instance

When you administer, develop, and deploy chaincode using an EC2 instance as a Hyperledger Fabric client, the permissions policies attached to the AWS Identity and Access Management *instance profile* and *instance role* associated with the instance determine its permissions to interact with other AWS resources, including AMB Access. The permissions policy shown in the following procedure provides sufficient privileges when it is attached to the IAM role of the instance.

The procedure demonstrates how to create a role with only this permissions policy attached and then attach that role to an EC2 instance. If you have an existing service role and instance profile attached to your EC2 instance, you can create an additional policy using the following example and then attach it to the existing role. Only one role can be attached to an EC2 instance.

For more information, see [IAM roles for Amazon EC2](#) in the *Amazon EC2 User Guide*.

To create a permissions policy, attach it to an IAM role, and attach the role to a Hyperledger Fabric client EC2 instance

1. Open the IAM console at <https://console.aws.amazon.com/iam/>.
2. In the navigation pane, choose **Policies**, and then **Create policy**.
3. Choose the JSON tab, and then copy and paste the following policy statement.

JSON

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "ListNetworkMembers",
      "Effect": "Allow",
      "Action": [
        "managedblockchain:GetNetwork",
        "managedblockchain:ListMembers"
      ],
      "Resource": [

```

```

        "arn:aws:managedblockchain:*:123456789012:networks/*"
    ]
},
{
    "Sid": "AccessManagedBlockchainBucket",
    "Effect": "Allow",
    "Action": [
        "s3:GetObject"
    ],
    "Resource": "arn:aws:s3:::us-east-1.managedblockchain/*"
},
{
    "Sid": "ManageNetworkResources",
    "Effect": "Allow",
    "Action": [
        "managedblockchain:CreateProposal",
        "managedblockchain:GetProposal",
        "managedblockchain>DeleteMember",
        "managedblockchain:VoteOnProposal",
        "managedblockchain:ListProposals",
        "managedblockchain:GetNetwork",
        "managedblockchain:ListMembers",
        "managedblockchain:ListProposalVotes",
        "managedblockchain:RejectInvitation",
        "managedblockchain:GetNode",
        "managedblockchain:GetMember",
        "managedblockchain>DeleteNode",
        "managedblockchain:CreateNode",
        "managedblockchain:CreateMember",
        "managedblockchain:ListNodes"
    ],
    "Resource": [
        "arn:aws:managedblockchain:*::networks/*",
        "arn:aws:managedblockchain:*::proposals/*",
        "arn:aws:managedblockchain:*:123456789012:members/*",
        "arn:aws:managedblockchain:*:123456789012:invitations/*",
        "arn:aws:managedblockchain:*:123456789012:nodes/*"
    ]
},
{
    "Sid": "WorkWithNetworksForAcct",
    "Effect": "Allow",
    "Action": [
        "managedblockchain:ListNetworks",

```

```
        "managedblockchain:ListInvitations",
        "managedblockchain:CreateNetwork"
    ],
    "Resource": "*"
}
]
```

4. Replace `123456789012` with your AWS Account ID. If you are working in a different Region, replace `us-east-1` with the appropriate Region, and then choose **Review policy**.
5. Enter a **Name** for the policy, for example, `HyperledgerFabricClientAccess`. Enter an optional **Description**, and then choose **Create policy**.

You now have a permissions policy that you can attach to an IAM role for an EC2 instance.

6. From the navigation pane, choose **Roles, Create role**.
7. Under **Select type of trusted entity**, leave **AWS service** selected. Under **Common use cases**, choose **EC2 - Allows EC2 instances to call AWS services on your behalf**, and then choose **Next: Permissions**.
8. Under **Attach permissions policies**, start typing the name of the permissions policy you created in the previous steps. Select that policy from the list, and then choose **Next: Tags**.
9. Leave the key-value fields blank or type key-value pairs for any tags that you want to apply to this role, and then choose **Next: Review**.
10. Enter a **Role name** that helps you identify this role, for example, `ServiceRoleForHyperledgerFabricClient`. Enter an optional **Description**, and then choose **Create role**.

You now have a service role with appropriate permissions that you can associate with your Hyperledger Fabric EC2 instance.

11. Open the EC2 console at <https://console.aws.amazon.com/ec2/>.
12. From the navigation pane, choose **Instances**.
13. From the list of instances, select the instance that you are using as a Hyperledger Fabric client.
14. Choose **Actions, Security, Modify IAM role**.
15. For **IAM role** begin typing the name of the role you created, for example, `ServiceRoleForHyperledgerFabricClient`. Select it from the list, and then choose **Apply**.

Using Service-Linked Roles for AMB Access

Amazon Managed Blockchain (AMB) uses AWS Identity and Access Management (IAM) [service-linked roles](#). A service-linked role is a unique type of IAM role that is linked directly to AMB Access. Service-linked roles are predefined by AMB Access and include all the permissions that the service requires to call other AWS services on your behalf.

A service-linked role can make setting up AMB Access easier because you don't have to manually add the necessary permissions. AMB Access defines the permissions of its service-linked roles, and, unless defined otherwise, only AMB Access can assume its roles. The defined permissions include the trust policy and the permissions policy. The permissions policy cannot be attached to any other IAM entity.

You can delete a service-linked role only after first deleting its related resources. This protects your AMB Access resources because you can't inadvertently remove permission to access the resources.

For information about other services that support service-linked roles, see [AWS Services That Work with IAM](#) and look for the services that have **Yes** in the **Service-Linked Role** column. Choose a **Yes** with a link to view the service-linked role documentation for that service.

Service-Linked Role Permissions for AMB Access

AMB Access uses the service-linked role named **AWSServiceRoleForAmazonManagedBlockchain**. This role enables access to AWS Services and Resources used or managed by Amazon Managed Blockchain.

The **AWSServiceRoleForAmazonManagedBlockchain** service-linked role trusts the following services to assume the role:

- `managedblockchain.amazonaws.com`

The role permissions policy allows AMB Access to complete actions on the specified resources shown in the following example policy.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
```

```
{
  "Action": [
    "logs:CreateLogGroup"
  ],
  "Effect": "Allow",
  "Resource": "arn:aws:logs:*:*:log-group:/aws/managedblockchain/*"
},
{
  "Effect": "Allow",
  "Action": [
    "logs:CreateLogStream",
    "logs:PutLogEvents",
    "logs:DescribeLogStreams"
  ],
  "Resource": [
    "arn:aws:logs:*:*:log-group:/aws/managedblockchain/*:log-
stream:*"
  ]
}
]
```

You must configure permissions to allow an IAM entity (such as a user, group, or role) to create, edit, or delete a service-linked role. For more information, see [Service-Linked Role Permissions](#) in the *IAM User Guide*.

Creating a Service-Linked Role for AMB Access

You don't need to manually create a service-linked role. When you create a network, a member, or a peer node, AMB Access creates the service-linked role for you. It doesn't matter if you use the AWS Management Console, the AWS CLI, or the AWS API. The IAM entity performing the action must have permissions to create the service-linked role. After the role is created in your account, AMB Access can use it for all networks and members.

If you delete this service-linked role, and then need to create it again, you can use the same process to recreate the role in your account. When you create a network, member, or node, AMB Access creates the service-linked role for you again.

Editing a Service-Linked Role for AMB Access

AMB Access does not allow you to edit the `AWSServiceRoleForAmazonManagedBlockchain` service-linked role. After you create a service-linked role, you cannot change the name of the role because various entities might reference the role. However, you can edit the description of the role using IAM. For more information, see [Editing a Service-Linked Role](#) in the *IAM User Guide*.

Deleting a Service-Linked Role for AMB Access

If you no longer need to use a feature or service that requires a service-linked role, we recommend that you delete that role. That way you don't have an unused entity that is not actively monitored or maintained. However, you must clean up the resources for your service-linked role before you can manually delete it.

Note

If the AMB Access service is using the role when you try to delete the resources, then the deletion might fail. If that happens, wait for a few minutes and try the operation again.

To manually delete the service-linked role

Use the IAM console, the AWS CLI, or the AWS API to delete the `AWSServiceRoleForAmazonManagedBlockchain` service-linked role. For more information, see [Deleting a Service-Linked Role](#) in the *IAM User Guide*.

Supported Regions for AMB Access Service-Linked Roles

AMB Access supports using service-linked roles in all of the Regions where the service is available. For more information, see [AWS Regions and Endpoints](#).

Configuring Security Groups for Amazon Managed Blockchain (AMB) Hyperledger Fabric

Security groups act as virtual firewalls. They control inbound and outbound traffic between your Hyperledger Fabric client and network resources on AMB Access through the VPC endpoint in your account. By default, security group rules are restrictive, so you must add rules that allow traffic for any resources, such as client computers, that must access the network. The following tables list the minimum required security group rules that must be associated with the VPC endpoint and the Hyperledger Fabric client.

VPC Endpoint Security Group Minimum Rules

Inbound/Outbound	Type	Source/Destination	Purpose
Outbound	All traffic	0.0.0/0 (Anywhere)	Default. Allows unrestricted outbound traffic from the interface VPC endpoint to all recipients.
Inbound	Custom TCP, Port for Ordering Service (ranging between 30000 and 34000) —for example, 30001. The port is available within the Ordering service endpoint on the network details page using the console and returned within the <code>OrderingServiceEndpoint</code> property using the <code>get-network</code> command from the AWS CLI or using the <code>GetNetwork</code> API action.	The IPv4 address, an address range, or a security group that includes all members' Hyperledger Fabric clients.	Allows the Hyperledger Fabric ordering service to receive traffic from Hyperledger Fabric clients.
Inbound	Custom TCP, Port for the CA Service for a member (ranging between 30000 and 34000)—for example, 30002. This is unique	The IPv4 address, an address range, or a security group that includes all members' Hyperledger Fabric clients.	Allows the Hyperledger Fabric certificate authority (CA) for each member to receive traffic from

Inbound/Outbound	Type	Source/Destination	Purpose
	<p>to each member, and each member only needs access to their own CA. The port is available within the Fabric certificate authority endpoint on the member details page using the console and returned within the <code>CaEndpoint</code> property using the <code>get-member</code> command from the AWS CLI or using the <code>GetMember</code> API action.</p>		<p>respective Hyperledger Fabric clients.</p>

Inbound/Outbound	Type	Source/Destination	Purpose
Inbound	Custom TCP, Ports, or Range of Ports for Peer Event Services on Peer Nodes (ranging between 30000 and 34000). The port is available within the Peer node endpoints on the member details page using the console and returned as the <code>PeerEventPort</code> property using the <code>get-node</code> command from the AWS CLI or using the <code>GetNode</code> API action.	The IPv4 address, an address range, or a security group that includes all members' Hyperledger Fabric clients.	Allows the network to receive traffic from peer nodes as required. Each node in each membership has a unique port associated with its peer event service. Any node that might be a participant in an endorsement policy, regardless of membership, must be allowed communications in order to endorse transactions.

Hyperledger Fabric Client Security Group Minimum Rules

Inbound/Outbound	Type	Source/Destination	Purpose
Outbound	All traffic	0.0.0/0 (Anywhere)	Default. Allows unrestricted outbound traffic from the Hyperledger Fabric client to all recipients. If necessary, you can limit the destination to the interface VPC endpoint.

Inbound/Outbound	Type	Source/Destination	Purpose
Inbound	SSH (Port 22)	The IP address, address range, or security group that includes trusted SSH clients that connect to the Hyperledger Fabric client.	Allows trusted clients to use SSH to connect to the Hyperledger Fabric client to interact—for example, to query and run chaincode.

Tagging Amazon Managed Blockchain (AMB) resources

A *tag* is a custom attribute label that you assign or that AWS assigns to an AWS resource. Each tag has two parts:

- A *tag key*, such as `CostCenter`, `Environment`, or `Project`. Tag keys are case-sensitive.
- An optional field known as a *tag value*, such as `111122223333` or `Production`. Omitting the tag value is the same as using an empty string. Like tag keys, tag values are case-sensitive.

Tags help you do the following:

- Identify and organize your AWS resources. Many AWS services support tagging, so you can assign the same tag to resources from different services to indicate that the resources are related. For example, you could assign the same tag to an Amazon Managed Blockchain (AMB) node and an EC2 instance that you use as a client for the AMB Access framework.
- Track your AWS costs. You activate these tags on the AWS Billing and Cost Management dashboard. AWS uses the tags to categorize your costs and deliver a monthly cost allocation report to you. For more information, see [Using cost allocation tags](#) in the [AWS Billing User Guide](#).
- Control access to your AWS resources with AWS Identity and Access Management (IAM). For information, see [Controlling access using tags](#) in this developer guide and [Control access using IAM tags](#) in the *IAM User Guide*.

For more information about tags, see the [Tagging Best Practices](#) guide.

The following sections provide more information about tags for AMB Access.

Create and add tags for AMB Access Hyperledger Fabric resources

You can tag the following resources:

- Networks
- Members
- Nodes

- **Proposals and invitations**

Members can assign tags to proposals when they create them. If a proposal is approved and an invitation is sent, the invitation inherits the tags added during proposal creation. Members who receive invitations or vote on proposals can subsequently add or remove tags for proposals and invitations. These subsequent tag edits are scoped only to the AWS account that made the edits.

Tags that you create—including those for network-wide resources like networks, proposals, and invitations—are scoped only to the account in which you create them. Tags created during proposal creation are the exception. Other AWS accounts participating on the network cannot access the tags.

Tag naming and usage conventions

The following basic naming and usage conventions apply to tags used with AMB Access resources:

- Each resource can have a maximum of 50 tags.
- For each resource, each tag key must be unique, and each tag key can have only one value.
- The maximum tag key length is 128 Unicode characters in UTF-8.
- The maximum tag value length is 256 Unicode characters in UTF-8.
- Allowed characters are letters, numbers, spaces representable in UTF-8, and the following characters: . : + = @ _ / - (hyphen).
- Tag keys and values are case-sensitive. As a best practice, decide on a strategy for capitalizing tags, and consistently implement that strategy across all resource types. For example, decide whether to use `Costcenter`, `costcenter`, or `CostCenter`, and use the same convention for all tags. Avoid using similar tags with inconsistent case treatment.
- The `aws :` prefix is reserved for AWS use. You can't edit or delete a tag's key or value when the tag has a tag key with the `aws :` prefix. Tags with this prefix do not count against your limit of tags per resource.

Working with tags

You can use the AMB Access console, the AWS CLI, or the AMB Access API to add, edit, or delete tag keys and tag values. You can assign tags when you create a resource, or you can apply tags after the resource is created.

For more information about AMB Access API actions for tagging, see the following topics in the *Amazon Managed Blockchain (AMB) API Reference*:

- [ListTagsForResource](#)
- [TagResource](#)
- [UntagResource](#)

Add or remove tags

You can add a tag to AMB Access Hyperledger Fabric resources when you create or work with them.

Topics

- [Add or remove tags for networks](#)
- [Add or remove tags for members](#)
- [Add or remove tags for nodes](#)
- [Add or remove tags for proposals](#)
- [Add or remove tags for invitations](#)

Add or remove tags for networks

For information about adding a tag when you create a network, see [Create a Hyperledger Fabric Blockchain Network on Amazon Managed Blockchain \(AMB\)](#).

To add or remove a tag for a Hyperledger Fabric network on AMB Access using the AWS Management Console

1. Open the AMB Access console at <https://console.aws.amazon.com/managedblockchain/>.
2. Choose **Networks** and then choose a Hyperledger Fabric network from the list.
3. Under **Tags**, choose **Edit tags**, and then do one of the following:
 - To add a tag, choose **Add new tag**, enter a **Key** and optional **Value**, and then choose **Save**.
 - To remove a tag, choose **Remove** next to the **Tag** you want to remove, and then choose **Save**.

Add or remove tags for members

For information about adding a tag when you create a member, see [Create a Member and Join a Network](#).

To add or remove a tag for a member

1. Open the AMB Access console at <https://console.aws.amazon.com/managedblockchain/>.
2. Choose **Networks** and then choose a Hyperledger Fabric network from the list.
3. Choose **Members**.
4. Under **Members owned by you**, choose a member from the list.
5. Under **Tags**, choose **Edit tags**, and then do one of the following:
 - To add a tag, choose **Add new tag**, enter a **Key** and optional **Value**, and then choose **Save**.
 - To remove a tag, choose **Remove** next to the **Tag** you want to remove, and then choose **Save**.

Add or remove tags for nodes

For information about adding a tag when you create a node, see [Create a Hyperledger Fabric Peer Node on Amazon Managed Blockchain \(AMB\)](#).

To add or remove a tag for a node

1. Open the AMB Access console at <https://console.aws.amazon.com/managedblockchain/>.
2. Choose **Networks** and then choose a Hyperledger Fabric network from the list.
3. Choose **Members**.
4. Under **Members owned by you**, choose a member from the list.
5. Under peer nodes, choose a **Node ID** from the list.
6. Choose **Tags**, choose **Edit tags**, and then do one of the following:
 - To add a tag, choose **Add new tag**, enter a **Key** and optional **Value**, and then choose **Save**.
 - To remove a tag, choose **Remove** next to the **Tag** you want to remove, and then choose **Save**.

Add or remove tags for proposals

For information about adding a tag when you create a proposal, see [Create an Invitation Proposal](#) and [Create a Proposal to Remove a Network Member](#).

You can add and remove tags from active and completed proposals.

To add or remove a tag for a proposal

1. Open the AMB Access console at <https://console.aws.amazon.com/managedblockchain/>.
2. Choose **Networks** and then choose a network from the list.
3. Choose **Proposals**.
4. Under **Active** or **Completed**, choose a **Proposal ID** from the list.
5. Under **Tags**, choose **Edit tags**, and then do one of the following:
 - To add a tag, choose **Add new tag**, enter a **Key** and optional **Value**, and then choose **Save**.
 - To remove a tag, choose **Remove** next to the **Tag** you want to remove, and then choose **Save**.

Add or remove tags for invitations

If the member who created a proposal for an invitation tagged the proposal, the invitation inherits the tag key and value from the proposal.

To add or remove a tag for an invitation

1. Open the AMB Access console at <https://console.aws.amazon.com/managedblockchain/>.
2. Choose **Invitations** and then choose a **Network name** from the list.
3. Under **Tags**, choose **Edit tags**, and then do one of the following:
 - To add a tag, choose **Add new tag**, enter a **Key** and optional **Value**, and then choose **Save**.
 - To remove a tag, choose **Remove** next to the **Tag** you want to remove, and then choose **Save**.

Monitoring AMB Access Hyperledger Fabric Using CloudWatch Logs

Amazon Managed Blockchain (AMB) Hyperledger Fabric supports publishing peer node, chaincode, and certificate authority (CA) logs to Amazon CloudWatch Logs. You can use these logs to troubleshoot during chaincode development and to monitor network activity and errors.

You can enable and view logs in the AMB Access management console, in the CloudWatch Logs console, and using AWS CLI commands for CloudWatch Logs. In addition, you can set up *metric filters* in CloudWatch Logs to turn log data into numerical CloudWatch metrics that you can graph and set an alarm on. For each member that has logging enabled, AMB Access creates a log group in CloudWatch Logs. For more information about CloudWatch Logs, see the [Amazon CloudWatch Logs User Guide](#). For more information about creating metric filters, see [Searching and Filtering Log Data](#) in the *Amazon CloudWatch Logs User Guide*.

- **Peer node logs** help you debug timeout errors associated with proposals and identify rejected proposals that do not meet the endorsement policies. Peer node logs contain messages generated when your client submits transaction proposals to peer nodes, requests to join channels, enrolls an admin peer, and lists the chaincode instances on a peer node. Peer node logs also contain the results of chaincode installation. You can enable and disable logs on individual peer nodes.
- **Chaincode logs** help you analyze and debug the business logic and execution of chaincode on a peer node. They contain the results of invoking, and querying the chaincode. A peer can run multiple instances of chaincode. When you enable chaincode logging, individual log streams are created for each and every chaincode on the peer.
- **CA logs** help you determine when a member in your account joins the network, or when new peers register with a member CA. You can use CA logs to debug problems related to certificates and enrollment. CA logging can be enabled and disabled for each member. A single log stream for the CA exists for each member.

Note

AMB Access gathers CloudWatch metrics for peer nodes automatically and separately from CloudWatch Logs for CAs, peer nodes, and chaincode. For more information, see [Use Hyperledger Fabric Peer Node Metrics on Amazon Managed Blockchain \(AMB\)](#).

Considerations and Limitations

Consider the following before you enable and view CloudWatch Logs for AMB Access Hyperledger Fabric.

- CA logs can be enabled only for members created after April 6, 2020. Peer node logs and chaincode logs can be enabled only for peer nodes created after April 6, 2020.
- Log entries are updated every five seconds.
- Logging requires the service-linked role for AMB Access. The role is created automatically when an IAM principal (user or role) in your account with permissions to create the service-linked role creates a network, member, or peer. For more information, see [Using Service-Linked Roles for AMB Access](#).
- Logging currently does not support CloudWatch Logs encryption.
- Logging currently does not support CloudWatch Logs Insights.
- To log chaincode events, you must first configure Go chaincode for logging. For more information, see [Logging Control for Chaincode](#) in the Hyperledger Fabric documentation.

Enabling and Disabling Logs

You can enable logs using the AMB Access management console when you create a member or node, or at any time after a member or node is created. You can disable CA logging, peer node logging, and chaincode logging at any time. When a log is disabled, log entries are not published to CloudWatch Logs. Previous log entries are still viewable and available in CloudWatch Logs.

Enabling and Disabling Peer Node and Chaincode Logs

You can enable peer node logs, chaincode logs, or both when you create a peer node or when viewing information about a peer node. You can also disable logs while viewing information about a peer node.

To enable peer node or chaincode logs when you create a node

1. On the **Members** tab of the network that you are working with, under **Members owned by you**, choose the name of the member from the list, and then choose **Create peer node**.
2. Under **Logging configuration**, choose **Enable peer node logs**, **Enable chaincode logs**, or both, and then choose **Create peer node**.

To enable or disable peer node and chaincode logs for an existing member

1. On the **Members** tab of the network that you are working with, under **Members owned by you**, choose the name of the member from the list.
2. Under **Peer nodes**, choose the **Node ID** of the peer.
3. Under monitoring, choose **Peer node logs** or **Chaincode logs**.
4. Choose **Enable logging**, or choose **Actions** and then **Disable Logging**.

Enabling and Disabling CA Logs

You can enable CA logs when you create a member or when viewing information about a member. There are two ways you can create members. One is when you create a new network, you must also create the first member. The other is when you receive an invite to create a member to join an existing network. You can only disable CA logs while viewing information about a member.

To enable CA logs when you create a member

- Under the **Logging configuration** section in the member creation form, choose **Enable certificate authority logs**.

To enable or disable CA logs for an existing member


1. On the **Members** tab of the network that you are working with, under **Members owned by you**, choose the name of the member from the list.
2. Under **CA logged events**, choose the **Enable logging**, or choose **Actions** and then **Disable logging**.

Working with Logged Events in the AMB Access Console

AMB Access publishes logged events to CloudWatch Logs every five seconds. By default, logged events are updated every five seconds in the AMB Access console under **Logged events** on member and node information pages. When viewing logged events in AMB Access, you can choose **Actions** and then view **View in CloudWatch** to open the CloudWatch Logs management console focused on the log stream that you are viewing. Choose the gear icon to configure the logging interval and other details.

Searching (Filtering) Logged Events

While viewing events for any log in AMB Access, you can enter a keyword or phrase in the **Search events** box to show only those events that contain the search term. For example, you can enter a date, a date and time, or a log level such as CRITICAL, DEBUG, or WARNING. When you download a log after searching, only the events filtered by your search term are downloaded.



The screenshot shows the AMB Access interface for viewing logged events. At the top, there is a header for 'CA logged events (44/100+)'. Below this is a search bar containing the text 'DEBUG'. To the right of the search bar is an 'Actions' dropdown menu. Below the search bar is a table with two columns: 'Timestamp (UTC)' and 'Message'. The table contains 12 rows of event data, all of which are filtered to show only 'DEBUG' level events. The messages include 'Cleaning up expired nonces for CA' and 'Received request for /cainfo'.

Timestamp (UTC)	Message
2/3/2020, 12:49:34 AM	2020/02/03 00:49:34 [DEBUG] Cleaning up expired nonces for CA 'm-TRNVTNOEISERPFMOYSLJVVXJUA'
2/3/2020, 12:45:56 AM	2020/02/03 00:45:56 [DEBUG] Received request for /cainfo
2/3/2020, 12:40:56 AM	2020/02/03 00:40:56 [DEBUG] Received request for /cainfo
2/3/2020, 12:35:56 AM	2020/02/03 00:35:56 [DEBUG] Received request for /cainfo
2/3/2020, 12:34:34 AM	2020/02/03 00:34:34 [DEBUG] Cleaning up expired nonces for CA 'm-TRNVTNOEISERPFMOYSLJVVXJUA'
2/3/2020, 12:30:56 AM	2020/02/03 00:30:56 [DEBUG] Received request for /cainfo
2/3/2020, 12:25:56 AM	2020/02/03 00:25:56 [DEBUG] Received request for /cainfo
2/3/2020, 12:20:56 AM	2020/02/03 00:20:56 [DEBUG] Received request for /cainfo
2/3/2020, 12:19:34 AM	2020/02/03 00:19:34 [DEBUG] Cleaning up expired nonces for CA 'm-TRNVTNOEISERPFMOYSLJVVXJUA'
2/3/2020, 12:15:56 AM	2020/02/03 00:15:56 [DEBUG] Received request for /cainfo
2/3/2020, 12:10:56 AM	2020/02/03 00:10:56 [DEBUG] Received request for /cainfo
2/3/2020, 12:05:56 AM	2020/02/03 00:05:56 [DEBUG] Received request for /cainfo

Downloading Logged Events

Choose **Actions** and then **Download** while viewing any log in AMB Access to save the events that are loaded to the default download directory on your local machine with a `.log` extension. Logged events are listed along with a header that contains log information as shown in the following example.

```
Event Count: 100
Filtered: Yes
Filtered Event Count: 44
Filter: "DEBUG"
```

```
2020-02-03T01:04:34.548Z 2020/02/03 01:04:34 [DEBUG] Cleaning up expired nonces for CA
'm-J46DNSFRTVCCLONS9DT5TTLS2A'
2020-02-03T01:00:56.382Z 2020/02/03 01:00:56 [DEBUG] Received request for /cainfo
2020-02-03T00:55:56.308Z 2020/02/03 00:55:56 [DEBUG] Received request for /cainfo
2020-02-03T00:50:56.208Z 2020/02/03 00:50:56 [DEBUG] Received request for /cainfo
2020-02-03T00:49:34.544Z 2020/02/03 00:49:34 [DEBUG] Cleaning up expired nonces for CA
'm-J46DNSFRTVCCLONS9DT5TTLS2A'
```

```

2020-02-03T00:45:56.282Z 2020/02/03 00:45:56 [DEBUG] Received request for /cainfo
2020-02-03T00:40:56.111Z 2020/02/03 00:40:56 [DEBUG] Received request for /cainfo
2020-02-03T00:35:56.026Z 2020/02/03 00:35:56 [DEBUG] Received request for /cainfo
2020-02-03T00:34:34.539Z 2020/02/03 00:34:34 [DEBUG] Cleaning up expired nonces for CA
'm-J46DNSFRTVCCLONS9DT5TTLS2A'
2020-02-03T00:30:56.081Z 2020/02/03 00:30:56 [DEBUG] Received request for /cainfo
2020-02-03T00:25:56.123Z 2020/02/03 00:25:56 [DEBUG] Received request for /cainfo
2020-02-03T00:20:56.197Z 2020/02/03 00:20:56 [DEBUG] Received request for /cainfo

```

Viewing Different Chaincode Logs

When viewing chaincode logs for a peer node with multiple chaincodes, you can choose the chaincode to view by choosing the chaincode name from the list next to **Logged events**. When you download a log, only the logged events for the chaincode that you are viewing are downloaded.

The screenshot shows the 'Monitoring' section of the Amazon Managed Blockchain console. The 'Chaincode logs' tab is active. Under 'Logged events (53/53)', there is a search bar and a dropdown menu. The dropdown menu is open, showing two options: 'a_chaincode' and 'b_chaincode'. The dropdown menu is highlighted with a red box. Below the dropdown is a table of logged events.

Timestamp (UTC)	Message
2/11/2020, 12:15:49 AM	Error starting Simple chaincode: error sending: rpc error: code = Unavailable desc = transport is closing
2/4/2020, 7:31:17 PM	Aval = 100, Bval = 200
2/4/2020, 7:31:17 PM	ex02 Invoke
2/4/2020, 7:31:04 PM	Aval = 110, Bval = 190
2/4/2020, 7:31:04 PM	ex02 Invoke
2/4/2020, 7:30:52 PM	Aval = 90, Bval = 210
2/4/2020, 7:30:52 PM	ex02 Invoke
2/4/2020, 7:30:46 PM	Aval = 70, Bval = 230
2/4/2020, 7:30:46 PM	ex02 Invoke
2/4/2020, 7:30:45 PM	Aval = 100, Bval = 200
2/4/2020, 7:30:45 PM	ex02 Invoke
2/4/2020, 7:30:42 PM	Aval = 80, Bval = 220
2/4/2020, 7:30:42 PM	ex02 Invoke
2/4/2020, 7:30:38 PM	Aval = 90, Bval = 210
2/4/2020, 7:30:38 PM	ex02 Invoke

Identifying Logs in CloudWatch Logs

Each set of log events in AMB Access corresponds to a *log stream* in CloudWatch Logs. The easiest way to access a log stream in the CloudWatch Logs management console is to choose **Actions** and then **View in CloudWatch** while viewing a log in the AMB Access management console.

All log streams associated with a member and peer nodes that a member owns are in a log group with the naming pattern shown below—for example, **aws/managedblockchain/n-MWY63ZJZU5HGNCMBQER7IN6OIU/m-J46DNSFRTVCCLONS9DT5TTLS2A**.

```
aws/managedblockchain/NetworkID/MemberID
```

Log streams in the log group are named according to the patterns in the table below.

Type of log	Stream name
Peer node logs	<i>PeerNodeID</i> , for example, nd-6EAJ5VA43JGGNPXOUZP7Y47E4Y .
Chaincode logs	<i>MemberID-PeerNodeID -ChaincodeName -ChaincodeVersion</i> , for example, m-J46DNSFRTVCCLONS9DT5TTLS2A-nd-6EAJ5VA43JGGNPXOUZP7Y47E4Y-MyChaincode-v0 .
CA logs	ca

Logging Amazon Managed Blockchain API calls using AWS CloudTrail

Amazon Managed Blockchain is integrated with AWS CloudTrail, a service that provides a record of actions taken by a user, role, or an AWS service in Managed Blockchain. CloudTrail captures all API calls for Managed Blockchain as events. The calls captured include calls from the Managed Blockchain console and code calls to the Managed Blockchain API operations.

If you create a trail, you can enable continuous delivery of CloudTrail events to an Amazon S3 bucket, including events for Managed Blockchain. If you don't configure a trail, you can still view the most recent events in the CloudTrail console in **Event history**. Using the information collected by CloudTrail, you can determine the request that was made to Managed Blockchain, the IP address from which the request was made, who made the request, when it was made, and additional details.

To learn more about CloudTrail, see the [AWS CloudTrail User Guide](#).

Managed Blockchain information in CloudTrail

CloudTrail is enabled on your AWS account when you create the account. When activity occurs in Managed Blockchain, that activity is recorded in a CloudTrail event along with other AWS service events in **Event history**. You can view, search, and download recent events in your AWS account. For more information, see [Viewing events with CloudTrail Event history](#).

For an ongoing record of events in your AWS account, including events for Managed Blockchain, create a trail. A *trail* enables CloudTrail to deliver log files to an Amazon S3 bucket. By default, when you create a trail in the console, the trail applies to all AWS Regions. The trail logs events from all Regions in the AWS partition and delivers the log files to the Amazon S3 bucket that you specify. Additionally, you can configure other AWS services to further analyze and act upon the event data collected in CloudTrail logs. For more information, see the following:

- Overview for [Creating a trail](#)
- [CloudTrail supported services and integrations](#)
- [Configuring Amazon SNS notifications for CloudTrail](#)
- [Receiving CloudTrail log files from multiple Regions](#) and [Receiving CloudTrail log files from multiple accounts](#)

All Managed Blockchain actions are logged by CloudTrail and are documented in the [Amazon Managed Blockchain API Reference](#). For example, calls to the `CreateNode`, `GetNode` and `DeleteNetwork` actions generate entries in the CloudTrail log files.

Every event or log entry contains information about who generated the request. The identity information helps you determine the following:

- Whether the request was made with root or AWS Identity and Access Management (IAM) user credentials.
- Whether the request was made with temporary security credentials for a role or federated user.
- Whether the request was made by another AWS service.

For more information, see the [CloudTrail userIdentity element](#).

Understanding Managed Blockchain log file entries

A trail is a configuration that enables delivery of events as log files to an Amazon S3 bucket that you specify. CloudTrail log files contain one or more log entries. An event represents a single request from any source and includes information about the requested action, the date and time of the action, request parameters, and so on. CloudTrail log files aren't an ordered stack trace of the public API calls, so they don't appear in any specific order.

The following example shows a CloudTrail log entry that demonstrates the `GetNode` action.

```
{
  "eventVersion": "1.05",
  "userIdentity": {
    "type": "AssumedRole",
    "principalId": "ABCD1EF23G4EXAMPLE56:carlossalazar",
    "arn": "arn:aws:sts::111122223333:assumed-role/Admin/carlossalazar",
    "accountId": "111122223333",
    "accessKeyId": "AKIAIOSFODNN7EXAMPLE",
    "sessionContext": {
      "sessionIssuer": {
        "type": "Role",
        "principalId": "ABCD1EF23G4EXAMPLE56",
        "arn": "arn:aws:iam::111122223333:role/Admin",
        "accountId": "111122223333",
        "userName": "Admin"
      }
    }
  },
```

```
    "webIdFederationData": {},
    "attributes": {
      "mfaAuthenticated": "false",
      "creationDate": "2020-12-10T05:36:38Z"
    }
  },
  "eventTime": "2020-12-10T05:50:48Z",
  "eventSource": "managedblockchain.amazonaws.com",
  "eventName": "GetNode",
  "awsRegion": "us-east-1",
  "sourceIPAddress": "198.51.100.1",
  "userAgent": "aws-cli/2.0.7 Python/3.7.3 Linux/5.4.58-37.125.amzn2int.x86_64
botocore/2.0.0dev11",
  "requestParameters": {
    "networkId": "n-MWY63ZJZU5HGNCMBQER7IN60IU",
    "nodeId": "nd-6EAJ5VA43JGGNPXOUZP7Y47E4Y"
  },
  "responseElements": null,
  "requestID": "1e2xa3m4-56p7-819e-0ex1-23456a78m90p",
  "eventID": "ex12345a-m678-901p-23e4-567ex8a9mple",
  "readOnly": true,
  "eventType": "AwsApiCall",
  "recipientAccountId": "111122223333"
}
```

Document History for Amazon Managed Blockchain (AMB) Hyperledger Fabric Developer Guide

The following table describes important additions to the Amazon Managed Blockchain (AMB) Management Guide. For notification about updates to this documentation, you can subscribe to the RSS feed.

Change	Description	Date
Amazon Managed Blockchain (AMB) Access	Updated terminology to change Amazon Managed Blockchain to Amazon Managed Blockchain (AMB) Access.	July 27, 2023
Updated for Hyperledger Fabric v2.2.	Updated channel configuration file examples, Java chaincode examples, and getting started tutorial for Hyperledger Fabric v2.2.	January 4, 2022
Updated for support of customer managed AWS KMS keys.	Updated for support of encryption at rest using customer managed keys in AWS Key Management Service (AWS KMS) for new member resources.	June 10, 2021
Guide updated to be specific to AMB Access Hyperledger Fabric.	Updated title from Amazon Managed Blockchain (AMB) Management Guide to Amazon Managed Blockchain (AMB) Hyperledger Fabric Developer Guide in conjunction with release of Ethereum framework	November 15, 2020

support. Slight changes to content throughout, clarifying applicability to Hyperledger Fabric networks, members, peer nodes, and chaincode.

[Updated for Hyperledger Fabric v1.4, including anchor peers, private data collections, Java chaincode.](#)

Added sections on setting up anchor peers, setting up channels, using private data collections, and running Java chaincode. Updated getting started tutorial for Hyperledger Fabric v1.4. Added new peer node metrics.

September 14, 2020

[Updated for support of Amazon CloudWatch Logs and streamlined Getting Started experience.](#)

Added Monitoring chapter for publishing peer node, chaincode, and CA logs to CloudWatch Logs. Miscellaneous improvements to getting started tutorial to streamline and simplify.

April 6, 2020

[Major updates for new proposal and voting workflow for member invitations and removals](#)

Updated Getting Started tutorial, conceptual information, and procedures for new voting proposal design.

April 8, 2019

[Added security group configuration guidance](#)

Added prescriptive guidance for configuring security groups for the tutorial. Added references for minimum inbound and outbound security group rules required for Hyperledger Fabric client and interface VPC endpoint for reference and customization.

February 28, 2019

[Updates to getting started steps](#)

Removed redundant steps in 3.2. The step to update `.bash_profile` with path to `fabric-ca` was already covered in step 3.1.

December 3, 2018

[Initial release of Amazon Managed Blockchain \(AMB\) \(Preview\)](#)

Initial documentation for Amazon Managed Blockchain (AMB).

November 28, 2018

AWS Glossary

For the latest AWS terminology, see the [AWS glossary](#) in the *AWS Glossary Reference*.