

## Deploy Enterprise Data Preparation on the AWS Cloud Marketplace (10.5)

## Abstract

You can deploy Enterprise Data Preparation on the Amazon Web Services (AWS) Marketplace. This deployment reference includes step-by-step instructions for deploying Enterprise Data Preparation on the Amazon Web Services (AWS) Marketplace. It also includes information on prerequisites, and how to troubleshoot common issues.

## Supported Versions

- Enterprise Data Preparation 10.5

## Table of Contents

Overview . . . . .	3
Intended audience. . . . .	3
Specialized knowledge. . . . .	3
Costs and licenses. . . . .	3
Architecture. . . . .	4
AWS Resources in the Deployment . . . . .	5
Informatica Domain. . . . .	7
Informatica Clients. . . . .	9
Before You Begin . . . . .	9
Network Prerequisites. . . . .	9
Account and security prerequisites. . . . .	10
Storage Prerequisite. . . . .	11
Deploy Enterprise Data Preparation on the AWS Marketplace. . . . .	11
Deployment steps. . . . .	12
Step 1. Prepare an AWS account. . . . .	12
Step 2. Get a license for Enterprise Data Preparation. . . . .	13
Step 3. Launch the deployment. . . . .	13
Step 4. Recycle the Informatica Services. . . . .	18
Monitoring Instance Provision and Informatica Domain Creation. . . . .	19
Deployed AWS Resources. . . . .	19
Output Tab Properties. . . . .	21
Deployed Informatica Resources. . . . .	22
Solution Deployment Log Files. . . . .	22
Troubleshooting. . . . .	23
Stack Failure. . . . .	23
Stack Launch Failure When the Pre-validation Check Parameter is Enabled. . . . .	24
Stack Launch Failure When Rollback On Failure Is Enabled. . . . .	24
Frequently asked questions on deployment error messages. . . . .	24

## Overview

Automated reference deployments use AWS CloudFormation templates to launch, configure, and run the AWS compute, network, storage, and other services required to deploy a specific workload on AWS.

Enterprise Data Preparation is a collaborative self-service data discovery and preparation solution for data analysts and data scientists. It enables analysts to rapidly discover and turn raw data into insights and allows IT to ensure quality, visibility, and governance. With Enterprise Data Preparation, analysts to spend more time on analysis and less time on finding and preparing data.

## Intended audience

This document is for users who are responsible for deploying Enterprise Data Preparation on the AWS Marketplace.

As a user with administrator privileges to deploy applications on AWS, it is assumed that you are familiar with AWS resources such as Cloud formation, VPC, EC2, S3, RDS, Internet gateway, NAT Gateway, Route table, Security group, and Elastic IP.

Familiarity with the concepts of IP CIDR and public and private IP addresses is recommended.

You must also be familiar with Enterprise Data Preparation.

For Enterprise Data Preparation documentation, see the [Informatica Documentation Portal](#).

## Specialized knowledge

Before you deploy Enterprise Data Preparation, we recommend that you become familiar with the following AWS services:

- [Amazon VPC](#)
- [Amazon EC2](#)
- [Amazon RDS](#)
- [Amazon S3](#)
- [Elastic IP Addresses](#)

If you are new to AWS, see [Getting Started with AWS](#).

## Costs and licenses

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

The AWS CloudFormation template for this deployment includes configuration parameters that you can customize. Some of these settings, such as instance type, will affect the cost of deployment. See the pricing pages for each AWS service you plan to use, for cost estimates.

This deployment requires a license for Enterprise Data Preparation. To sign up for a license, contact [Informatica](#).

The following table lists the instance types that you can choose based on sizing requirements:

Virtual Machine	Instance Type	Cluster Size
Database	m5.xlarge	Small, Medium, Large
Informatica Domain	m4.2xlarge / m5.2xlarge	Small, Medium, Large
Bastion Server	m4.xlarge / m5.xlarge	Small, Medium, Large
Informatica Embedded Hadoop Cluster	m4.4xlarge / m5.4xlarge m4.2xlarge / m5.2xlarge	Small Medium, Large
Informatica Compute Cluster on EMR	m4.xlarge / m5.xlarge / m4.2xlarge / m5.2xlarge	Small, Medium, Large

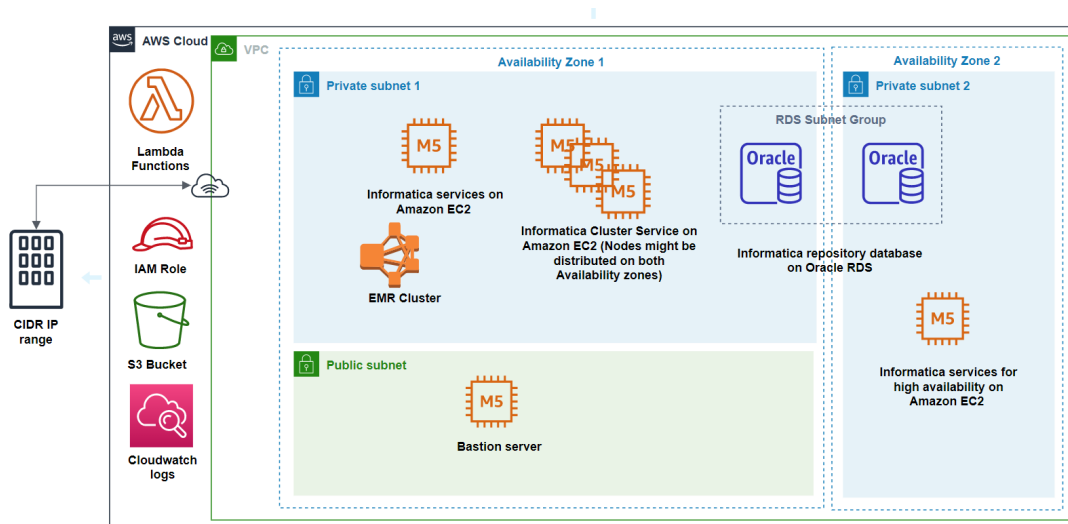
The deployment chooses M4 or M5 type of instances based on instance-type availability in a specific AWS region.

## Architecture

When you deploy the Enterprise Data Preparation solution on a new Virtual Private Cloud (VPC) with default parameters, the deployment builds the following Enterprise Data Preparation environment on the AWS Marketplace:

- An Informatica domain server on an EC2 instance, with additional instances to contain nodes in the Data Integration Service grid.
- An Informatica cluster that Enterprise Data Catalog uses to run metadata processing and profiling jobs.
- Informatica clients on a remote Windows bastion server that runs on a public subnet.
- Amazon S3 storage resources and connections for source and target data in existing Amazon S3 buckets.
- Amazon RDS relational databases for the Informatica repositories.
- AWS security and account management services.
- An Amazon EMR cluster with autoscaling enabled.

The following image shows the architecture of Enterprise Data Preparation on AWS Marketplace:



The architecture diagram includes the following components:

- A virtual public cloud (VPC) configured across two Availability Zones to contain the Enterprise Data Preparation deployment.
- Availability Zones. The deployment provisions two Availability Zones.
- Subnets to contain specific elements of the deployment. The deployment creates two private subnets, plus one public subnet if you want to use a Windows bastion server for Informatica clients. The deployment creates each of the subnets in a different Availability Zone.
- The Informatica domain where application services run, including the Catalog Service, the Enterprise Data Preparation Service, the Interactive Data Preparation Service, the Model Repository Service and the Data Integration Service.
- An internal Hadoop cluster deployed on Amazon EC2 instances that Enterprise Data Catalog uses to run metadata processing and profiling jobs.
- A compute cluster of Amazon EC2 instances with autoscaling enabled that Enterprise Data Preparation uses to publish prepared data to the data lake.
- An Oracle database on Amazon RDS for the following Informatica repositories:
  - Data Preparation repository
  - Domain repository
  - Model repository
- Informatica clients in a separate EC2 bastion server in a public subnet.
- AWS Lambda functions.
- IAM roles.
- Amazon S3 storage, to act as a temporary location to store Enterprise Data Preparation processing data.
- Amazon CloudWatch.

## *AWS Resources in the Deployment*

The deployment process creates or includes the following components:

- Components created when you specify the solution to create a new VPC.
- Components created, as well as some that already exist, when you specify the solution to use an existing VPC.
- A bastion server, if you choose to create one.
- An internal Hadoop cluster deployed on Amazon EC2 instances that Enterprise Data Catalog uses to run metadata processing and profiling jobs.
- A compute EMR cluster with autoscaling enabled that Enterprise Data Preparation uses to publish prepared data to the data lake.

## *Components in a Deployment on a New VPC*

The deployment creates the following components when you deploy the software in a new virtual private cloud:

<b>Component</b>	<b>Number of Components Created</b>
VPC	1
Internet gateway	1

Component	Number of Components Created
IAM role and instance profile	1
Subnets	4
Route tables	2
Security groups	The following security groups: <ul style="list-style-type: none"> <li>- One security group for RDS.</li> <li>- One security group for the Informatica domain instance.</li> <li>- If you selected the option to create the bastion server, one security group to support remote login to the bastion server.</li> <li>- Two security groups for the embedded Hadoop cluster.</li> <li>- Three security groups for the Informatica compute cluster on EMR.</li> </ul>
Database subnet group	1
Amazon RDS instance	1 Amazon RDS for Oracle instance with 200 GB of storage capacity that runs on a db.m5.xlarge DB instance type.
Amazon EC2 instances	Informatica domain with an EC2 instance and 300 GB of storage space. The instance created is of type m5.2xlarge or m4.2xlarge, depending on availability of the instance size in the selected region. Based on your choice, 1 (one) Informatica high availability domain node on an EC2 instance with 300 GB of storage space. The instance created is of type m5.2xlarge or m4.2xlarge, depending on availability of the instance size in the selected region. Based on your choice, a bastion server with an EC2 instance that has 200 GB of storage space is created. The instance created is of type m5.large or m4.large, depending on availability of the instance size in the selected region. Based on the Deployment Type you select, the deployment launches the following: <ul style="list-style-type: none"> <li>- 1, 3, or 6 EC2 instances with 300 GB of storage space each for the Informatica cluster.</li> <li>- 2 EC2 instances for the Informatica compute cluster on EMR.</li> </ul>
Elastic IP addresses	1 elastic IP address attached to the NAT gateway.

## Components in a Deployment on an Existing VPC

The deployment creates the following components when you deploy the software in an existing virtual private cloud:

Component	Number of Components Created
Security groups	The following security groups: <ul style="list-style-type: none"> <li>- One security group for RDS.</li> <li>- One security group for the Informatica domain instance.</li> <li>- If you selected the option to create the bastion server, one security group to support remote login to the bastion server.</li> <li>- Based on your choice, two security groups for Informatica EMR workflow connections.</li> </ul>
Database subnet group	1
Amazon RDS instance	1 Amazon RDS for Oracle instance with 200 GB of storage capacity that runs on a db.m5.xlarge DB instance type.

Component	Number of Components Created
Amazon EC2 instances	<p>Informatica domain with an EC2 instance and 300 GB of storage space. The instance created is of type m5.2xlarge or m4.2xlarge, depending on availability of the instance size in the selected region.</p> <p>Based on your choice, 1 (one) Informatica high availability domain node on an EC2 instance with 300 GB of storage space. The instance created is of type m5.2xlarge or m4.2xlarge, depending on availability of the instance size in the selected region.</p> <p>Based on your choice, a bastion server with an EC2 instance that has 200 GB of storage space is created. The instance created is of type m5.large or m4.large, depending on availability of the instance size in the selected region.</p> <p>Based on the Deployment Type you select, the deployment launches the following:</p> <ul style="list-style-type: none"> <li>- 1, 3, or 6 EC2 instances with 300 GB of storage space each for the Informatica cluster.</li> <li>- 2 EC2 instances for the Informatica compute cluster on EMR.</li> </ul>
AWS Lambda functions	5 functions to confirm the validity of the VPC, the subnets, and the Informatica license.
IAM role and instance profile	1

## Bastion Server

You can optionally deploy a bastion server as an element in the solution. The bastion server acts as a firewall between the internet and the cloud platform network where the solution deploys. It can also act as a remote Windows server which you can log into to run Informatica clients.

## Informatica Domain

The Informatica domain is a server component that hosts application services, such as the Model Repository Service and the Data Integration Service. These services, together with domain clients, enable you to create and run mappings and other objects to extract, transform, and write data.

### Application Services

#### Enterprise Data Preparation Service

The Enterprise Data Preparation Service is an application service in the Informatica domain that runs the Enterprise Data Preparation application.

#### Interactive Data Preparation Service

The Interactive Data Preparation Service is an application service in the Informatica domain that manages data preparation within the Enterprise Data Preparation application.

#### Catalog Service

The Catalog Service is an application service in the Informatica domain that runs the Enterprise Data Catalog application, and manages connections between service components and external applications.

#### Model Repository Service

The Model Repository Service is an application service in the Informatica domain that manages the Model repository. The Model repository stores metadata created by Informatica products in a relational database to enable collaboration among the products. Informatica Developer, the Data Integration Service, and the Administrator tool store metadata in the Model repository.

### **Data Integration Service**

The Data Integration Service is an application service in the Informatica domain that performs data integration tasks for the Developer tool and for external clients.

### **Metadata Access Service**

The Metadata Access Service is an application service that allows the Developer tool to access Hadoop connection information to import and preview metadata. The Metadata Access Service contains information about the Service Principal Name (SPN) and keytab information if the Hadoop cluster uses Kerberos authentication.

### **Content Management Service**

The Content Management Service is an application service in the Informatica domain that manages reference data and is responsible for compiling rule specifications into mapplets. The Content Management Service provides reference data information to the Data Integration Service and to the Developer tool and Analyst tool. The Content Management Service stores reference data in a database that you specify.

### **Analyst Service**

The Analyst Service is an application service in the Informatica domain that runs Informatica Analyst. The Analyst Service manages the connection between the service components and the users who log in to Analyst tool. You can perform column and rule profiling, manage scorecards, and manage bad records and duplicate records in the Analyst tool. The Analyst Service stores profiling, scorecarding, and bad and duplicate record data in databases that you specify.

### **Informatica Cluster Service**

The Informatica Cluster Service is an application service that runs and manages all the associated services that are required to run Enterprise Data Catalog in the Informatica domain. The associated services include Mongo DB, Nomad, Solr, PostgreSQL, and ZooKeeper.

The Informatica domain can run several other services. For more information about Informatica services, see the [Informatica Application Service Guide](#).

## **Repositories**

Informatica repositories, hosted on Oracle or Microsoft SQL Server databases, store metadata about domain objects. Informatica repositories include the following:

### **Domain configuration repository**

The domain configuration repository stores configuration metadata about the Informatica domain. It also stores user privileges and permissions.

### **Model repository**

The Model repository stores metadata for projects and folders and their contents, including all repository objects such as mappings and workflows. The repository also stores rules you apply during data preparation.

### **Data Preparation repository**

The Data Preparation repository stores worksheet metadata created when you use the Enterprise Data Preparation application to prepare data.

In addition to these domain repositories, the solution also requires a repository for Hive metadata. This repository is hosted on an SQL database. It stores Hive table metadata to enable Hadoop operations.

For more information about domain repositories, see the [Informatica Application Service Guide](#).

## **Clusters**

Informatica uses the following EC2 node clusters.



### **Informatica Cluster**

The Informatica cluster deployed on Amazon EC2 instances that the Catalog Service uses to run metadata processing and profiling jobs.

### **Informatica compute cluster on EMR**

A compute cluster of Amazon EC2 instances with autoscaling enabled that Enterprise Data Preparation uses to publish prepared data to the data lake.

## *Informatica Clients*

You can use different clients with Enterprise Data Preparation:

### **Administrator tool**

The Administrator tool enables you to create and administer services, connections, and other domain objects.

### **Analyst tool**

The Analyst tool enables you to analyze, cleanse, standardize, profile, and score data in your enterprise.

### **Developer tool**

The Developer tool enables you to create and run mappings and other objects that enable you to access, transform, and write data to targets.

### **Command line interface**

The command line interface offers hundreds of commands to assist in administering the Informatica domain, creating and running repository objects, administering security features, and maintaining domain repositories.

### **Bastion server**

You can optionally deploy a bastion server as an element in the solution. The bastion server is an AWS Windows instance installed with the Developer tool and the command line interface clients. The bastion server acts as a firewall for access to the VPC. It can also act as a remote Windows server which you can log into to run Informatica clients.

## **Before You Begin**

Before you launch the automated deployment on the AWS Marketplace, verify the prerequisites for each option described in this section.

### *Network Prerequisites*

Before you deploy Enterprise Data Preparation, verify that you have completed the following network prerequisites for the Informatica application services.

Network prerequisites depend on whether you choose to deploy the solution to a new subnet or an existing subnet.

When you deploy to a new VPC, the automated deployment builds a new AWS environment consisting of the VPC, subnets, NAT gateways, security groups, and other infrastructure components. The deployment then deploys Enterprise Data Preparation to the new VPC.

When you deploy to an existing VPC, the automated deployment provisions Enterprise Data Preparation in your existing AWS infrastructure.

### New VPC Prerequisites

Complete the following prerequisites when you choose to deploy into a new VPC:

- Ensure the selected Availability Zone has sufficient capacity to create a new subnet, route table, internet gateway, and NAT gateway.
- Ensure that the VPC creation has not exceeded the supported limit.

### Existing VPC Prerequisites

Complete the following prerequisites when you choose to deploy into an existing VPC:

- Ensure that the VPC has DNS resolution enabled.
- You must also attach the subnet to a route table with a local route to the VPC CIDR.
- Verify that each subnet in the VPC is associated with a route table.
- Verify that the subnet associated with an internet gateway has the Auto-assign Public IPv4 property enabled.
- Verify that each route listed in the route table is defined as one of the following gateway types:
  - Internet gateway
  - NAT gateway
  - Virtual private gateway
- Verify that the subnet is associated with one of the following gateways:
  - Internet gateway
  - NAT gateway
  - Virtual private gateway
- Verify that the service subnets you choose are attached to a route table under the selected VPC.
- Verify that your VPC has a minimum of two subnets. If you choose to create a bastion server, you can deploy it to one of the two subnets or use a separate public subnet.

Subnets require an Internet gateway, an NAT gateway, or a virtual private gateway in their route tables, to allow the instances to download packages and software without exposing them to the Internet. Record the domain name configured in the DHCP options as explained in the [Amazon VPC documentation](#). You provide the domain name in the VPC settings when you configure the deployment.

## Account and security prerequisites

Complete the prerequisites in this section to verify that you obtained the required AWS account and private keys.

### AWS Account and Region

If you do not already have an AWS account, create an account at [aws.amazon.com](https://aws.amazon.com).

Use the region selector in the navigation bar to choose the AWS region where you want to deploy Enterprise Data Preparation on AWS.

**Note:** The region you choose determines the availability of certain sizes of EC2 instances.

If necessary, [request a service limit increase](#) for the EC2 instance type you decided to use for the Informatica domain. You might need to do this if you already have an existing deployment that uses this instance type, and you think you might exceed the [default limit](#) with this reference deployment.

### Private Keys

Create a key pair in your preferred region.

When you log in to an EC2 instance, you use a password file for authentication. The file is called a private key file and has a file name extension of .pem.

If you do not have an existing .pem key to use, follow the instructions in the [AWS documentation](#) to create a key pair.

**Note:** Your administrator might ask you to use an existing key pair.

When you create a key pair, you save the .pem file to your desktop system. Simultaneously, AWS saves the key pair to your account. Make a note of the key pair that you want to use for the deployment, so that you can provide the key pair name during network configuration.

### Account User Policies

Choose to grant the AWS user one of the following privileges:

- Administrator privileges
- The following minimum policies attached to the user account:
  - arn:aws:iam::aws:policy/AWSCloudFormationFullAccess
  - arn:aws:iam::aws:policy/AmazonEC2FullAccess
  - arn:aws:iam::aws:policy/CloudWatchFullAccess
  - arn:aws:iam::aws:policy/IAMFullAccess
  - arn:aws:iam::aws:policy/AmazonS3FullAccess
  - arn:aws:iam::aws:policy/AmazonRDSFullAccess
  - arn:aws:iam::aws:policy/AmazonVPCFullAccess
  - arn:aws:iam::aws:policy/AWSLambdaFullAccess
  - arn:aws:iam::aws:policy/AmazonElasticMapReduceFullAccess

### Storage Prerequisite

Create an Amazon S3 bucket for the Enterprise Data Preparation license.

Amazon S3 stores objects, consisting of one or more files and accompanying file metadata, in "buckets." For information about Amazon S3, see the [AWS documentation](#).

To use the automated solution on the cloud marketplace, you must create an S3 bucket to store the license. You can store the license on the bucket level or in a folder within the S3 bucket.

Note the location of the license key and supply it when you configure the solution in the **License Key S3 Bucket URI** parameter.

## Deploy Enterprise Data Preparation on the AWS Marketplace

The automated deployment of Enterprise Data Preparation on AWS Marketplace uses the AWS CloudFormation console to guide your choices and launch the solution deployment.

You can select one of the following deployment options:

- Deploy Enterprise Data Preparation into a new VPC (end-to-end deployment). This option builds and deploys a new AWS environment. The environment that it builds consists of the VPC, subnets, NAT gateways, security groups, and other infrastructure components, and then deploys Enterprise Data Preparation into this new VPC.
- Deploy Enterprise Data Preparation into an existing VPC. This option provisions Enterprise Data Preparation in your existing AWS infrastructure.

The deployment also lets you configure additional settings such as CIDR blocks, instance types, and Enterprise Data Preparation settings, as discussed later in this guide.

Use the AWS CloudFormation console to specify the stack name and to enter parameter values to the template that you use to deploy the application. A stack is a logical representation of an AWS deployment.

Specify the identifiers for the VPC, for the subnet that contains the Amazon EC2 instances where the deployment process deploys the application components, and for the subnet where the deployment process creates the Model repository. The deployment process provisions the application components and creates the Model repository based on the values that you specify.

### Availability Zones and Subnets

The automated solution deploys the components and services for Enterprise Data Preparation into an Amazon VPC configured across two Availability Zones. Each Availability Zone is provisioned with two subnets.

In the first subnet, the deployment process provisions Enterprise Data Preparation application services on an Amazon EC2 instance within the VPC:

- Data Integration Service
- Model Repository Service
- Metadata Access Service
- Informatica Cluster Service
- Catalog Service
- Enterprise Data Preparation Service
- Interactive Data Preparation Service

For more information about domain components, see the [Informatica Administrator Guide](#) and [Informatica Application Service Guide](#).

In the second subnet, the solution deploys an Oracle database on Amazon RDS for use as the Model repository. The Model repository stores metadata created by Informatica services.

## Deployment steps

You must perform multiple tasks to deploy the Enterprise Data Preparation 10.5 AWS Marketplace solution.

You configure multiple parameters, such as availability zones and subnets. The parameters differ depending on whether you deploy the Enterprise Data Preparation on a new VPC or an existing VPC.

### Step 1. Prepare an AWS account

Perform the following steps to prepare an AWS account:

1. If you do not have an AWS account, create an account at <http://aws.amazon.com> by following the on-screen instructions.
2. Use the region selector in the navigation bar to choose the AWS Region where you want to deploy Enterprise Data Preparation on AWS.
3. Create a key pair in your preferred region.
4. When you log in to an EC2 instance, you use a password file for authentication. The file is called a private key file and has a file name extension of .pem.
5. If you do not have an existing .pem key to use, follow the instructions in the [AWS documentation](#) to create a key pair.  
When you create a key pair, you save the .pem file to your desktop system. Simultaneously, AWS saves the key pair to your account. Make a note of the key pair that you want to use for the Enterprise Data Preparation instance, so that you can provide the key pair name during network configuration.

**Note:** Your administrator might ask you to use an existing key pair.

6. If necessary, [request a service limit increase](#) for the EC2 instance type that you want to use for the Informatica domain. You might need to do this if an existing deployment uses this instance type and you think you might exceed the [default limit](#) with this reference deployment.

## Step 2. Get a license for Enterprise Data Preparation

This deployment requires a license for Enterprise Data Preparation. To sign up for a license, contact [Informatica](#).

1. In your AWS account, create an S3 bucket.
2. Optionally, create a directory under the S3 bucket.
3. Place the Enterprise Data Preparation license key file for the software in the S3 bucket or directory.

## Step 3. Launch the deployment

You are responsible for the cost of the AWS services used while running this deployment. There is no additional cost for using this deployment. See the pricing pages for each AWS service you will be using for full details. Each deployment takes approximately one to two hours to complete.

Perform the following steps to launch the deployment:

1. Log in to AWS Marketplace using the AWS account.
2. Type Enterprise Data Preparation in the search box and press **Enter**.
3. Click **Enterprise Data Preparation**. The **Product Overview** page appears.
4. Click **Continue to Subscribe**. The **Subscribe to this software** page appears.
5. Click **Continue to Configuration**. The **Configure this software** page appears.
6. Based on your requirement, select one of the following options from the **Delivery Method** drop-down box to specify where you want to deploy Enterprise Data Preparation:
  - Enterprise Data Preparation (New VPC). Select this option to deploy Enterprise Data Preparation on a new VPC.
  - Enterprise Data Preparation (Existing VPC). Select this option to deploy Enterprise Data Preparation on an existing VPC.
7. Choose **10.5.0** from the **Software Version** list.
8. Click **Continue to Launch**. The **Launch this software** page appears.
9. Select **Launch CloudFormation** from the **Choose Action** drop-down list and click **Launch**. The **Create stack** page appears.
10. Check the region for the deployment displayed in the upper-right corner of the navigation bar, and change it if necessary. This is where the network infrastructure for Enterprise Data Preparation will be built. The template is launched in the U.S. West (Oregon) Region by default.
11. On the **Specify template** section, keep the default setting for the template URL, and then choose **Next**. The default option is **Amazon S3 URL**.
12. On the **Specify stack details** page, change the stack name if needed. Review the parameters for the template. Provide values for the parameters that require input.
13. Under **Capabilities**, select the check box to acknowledge that the template will create IAM resources.
14. To deploy the stack, click **Create**.

Amazon AWS begins to create the stack and displays the CloudFormation dashboard. When the status is CREATE\_COMPLETE, the deployment is complete. You can use the URL displayed on the **Output tab** to view the stack properties.

- For all other parameters, review the default settings and customize them as necessary. When you finish reviewing and customizing the parameters, choose **Next**.

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

- [“Option 1. Deploy into a new VPC” on page 14](#)
- [“Option 2. Deploy into an existing VPC” on page 16](#)

**Note:** If you are deploying Enterprise Data Preparation into an existing VPC, make sure that your VPC has a minimum of two subnets. If you choose a remote desktop, you can either use one of the two subnets or use a separate public subnet. These subnets require NAT gateways in their route tables, to allow the instances to download packages and software without exposing them to the Internet. You also need the domain name option configured in the DHCP options, as explained in the [Amazon VPC documentation](#). When you launch the deployment, the wizard prompts you for the VPC settings.

### Option 1. Deploy into a new VPC

Specify the identifiers for the VPC, for the subnet containing the Amazon EC2 instances where the deployment process deploys the application components, and for the subnets used by the Informatica services. The deployment process provisions the application components and creates the required repository databases based on the values you specify.

- Enter the name of the stack to create. The stack name can include letters (A-Z and a-z), numbers (0-9), and dashes (-).
- Enter the parameters required to deploy the application.

The following table lists the parameters in the **Network Configuration** section:

Parameter	Default	Description
Availability Zones	Requires input	The list of Availability Zones to use for the subnets in the VPC. Make sure that you select two zones.
VPC CIDR	10.0.0.0/16	Classless Inter-Domain Routing IP address range associated with the Amazon VPC. The CIDR range must have permissions to access the public subnet where the application runs. Format the string as <IP address>/<IP network prefix>. For example, to specify a range from 10.20.30.40 to 10.20.30.47, enter 10.20.30.40/29.
Public Subnet CIDR	Requires input	CIDR block for the public subnet. Ensure that you include the subnet in the VPC CIDR.
Private Subnet 1 CIDR	10.0.0.0/19	CIDR block for the private subnet 1. Ensure that you include the subnet in the VPC CIDR.

Parameter	Default	Description
Private Subnet 2 CIDR	10.0.32.0/19	CIDR block for the private subnet 2. Ensure that you include the subnet in the VPC CIDR.
CIDR IP Range	Requires input	The CIDR IP range that is permitted to access Enterprise Data Preparation. To reduce the potential of inbound attacks from unknown IP addresses, you can use a constrained CIDR range to reduce. For example, to specify the range of 10.20.30.40 to 10.20.30.47, enter 10.20.30.40/29.

The following table lists the parameters in the **Bastion Server for Informatica** section:

Parameter	Default	Description
Deploy Bastion Server	Requires input	Select <b>Yes</b> to deploy a bastion server to access other resources in the VPC.

The following table lists the parameters in the **Amazon EC2 Configuration** section:

Parameter label (name)	Default	Description
Key Pair Name	Requires input	Name of an existing Amazon EC2 key pair to enable SSH and Remote Desktop Protocol (RDP) access to Enterprise Data Preparation.
Keypair S3 URI	Requires input	S3 Bucket URI for the keypair selected. Use the following format: s3://bucket-name/Keypair.pem or s3://bucket-name/sub-folder/Keypair.pem

The following table lists the parameters in the Informatica Enterprise Data Preparation section:

Parameter label (name)	Default	Description
License Key S3 Bucket URI	Requires input	The S3 bucket URI for the Enterprise Data Preparation license key in the following format: s3://<bucket-name>/<directory>/<license key>
High Availability Deployment	Requires input	Enable or disable high availability for the deployment. For information about high availability for the Informatica domain, see the <a href="#">Informatica documentation</a> .
Deployment Type	Requires input	Choose from the following options for the Informatica cluster: <ul style="list-style-type: none"> <li>- Small: one m*.4xlarge node</li> <li>- Medium: three m*.2xlarge nodes</li> <li>- Large: six m*.2xlarge nodes</li> </ul> Choose from the following options for the Informatica compute cluster on EMR: <ul style="list-style-type: none"> <li>- Small: one m*.4xlarge node</li> <li>- Medium: one to three m*.2xlarge nodes</li> <li>- Large: one to six m*.2xlarge nodes</li> </ul> <b>Note:</b> The instance types you can select might depend on your AWS region.

Parameter label (name)	Default	Description
S3 Bucket Name	N/A	Specify the name of an existing S3 bucket. <b>Note:</b> This bucket is not the same as the bucket that contains the Enterprise Data Preparation license key. To use an existing bucket that contains source and target data, specify the name of a bucket in your account. If you do not specify a value, the deployment process creates an S3 bucket.
Cluster ID for Amazon EMR	N/A	Specify the identifier for an existing Amazon EMR cluster to use for the Informatica compute cluster. If you do not specify a value, the deployment process creates a cluster with autoscaling enabled. The number of EC2 nodes in the cluster depends on the Deployment Type value you set. For more information about cluster workflows in Intelligent Data Streaming, see <a href="#">Cluster Workflows</a> .
Use S3 files	No	Use configuration files stored on an Informatica-owned S3 bucket. Selecting Yes downloads configuration files and then configures. May contain latest fixes if any, else the deployment uses local configuration files stored in the EC2 instance.

- Click **Next**.
- On the Options page, expand the **Advanced** tab.
- Set the **Rollback on Failure** property to **No**.
- Click **Next**.
- Click **Create**.
- When the status for the stack reads CREATE\_COMPLETE, click the stack name.
- Click the **Outputs** tab to find the URLs for the Informatica Administrator, Catalog, and Enterprise Data Preparation.

## Option 2. Deploy into an existing VPC

- Enter a name for the stack. The stack name can include letters (A-Z and a-z), numbers (0-9), and dashes (-).
- Enter the parameters required to deploy the application.  
The following table lists the parameters in the **Network Configuration** section:

Parameter	Default	Description
VPC	Requires input	ID of your existing VPC where you want to deploy Enterprise Data Preparation.
Subnet 1	Requires input	Select the subnet ID where you want to deploy Enterprise Data Preparation. Make sure that the subnet is included in the VPC and associated with an Internet Gateway or a NAT Gateway. Ensure that service subnet ID1 and service subnet are not the same.
Subnet 2	Requires input	If you want to deploy Enterprise Data Preparation with high availability enabled, provide the subnet ID that you want to use for high availability. Make sure that subnet 1 and subnet 2 are in different availability zones.



Parameter	Default	Description
CIDR IP Range	Requires input	The CIDR IP range that is permitted to access Enterprise Data Preparation. To reduce the potential of inbound attacks from unknown IP addresses, you can use a constrained CIDR range. For example, to specify the range of 10.20.30.40 to 10.20.30.47, enter 10.20.30.40/29.
Pre-validation Check	Requires input for a	Select to validate the VPC and subnets before deploying Enterprise Data Preparation.  If you choose to have the automated solution perform a validation check before launching the deployment, the check validates your parameter inputs. If the deployment as specified fails the check, the deployment process stops. The deployment logs record the problem.

The following table lists the parameters in the **Bastion Server for Informatica** section:

Parameter	Default	Description
Deploy Bastion Server	Requires input	Select <b>Yes</b> to deploy a bastion server to access other resources in the VPC.
Bastion Subnet	Requires input	Select the subnet for the bastion server if you selected <b>Yes</b> for the previous option. Ensure that the subnet is included in the VPC that you use and attached to an Internet Gateway.

The following table lists the parameters in the **Amazon EC2 Configuration** section:

Parameter label (name)	Default	Description
Key Pair Name	Requires input	Name of an existing Amazon EC2 key pair to enable SSH and Remote Desktop Protocol (RDP) access to Enterprise Data Preparation.
Keypair S3 URI	Requires input	S3 Bucket URI for the keypair selected. Use the following format: s3://bucket-name/Keypair.pem or s3://bucket-name/sub-folder/Keypair.pem

The following table lists the parameters in the Informatica Enterprise Data Preparation Configuration section:

Parameter label (name)	Default	Description
License Key S3 Bucket URI	Requires input	The S3 bucket URI for the Enterprise Data Preparation license key in the following format: s3://<bucket-name>/<directory>/<license key>
High Availability Deployment	Requires input	Enable or disable high availability for the deployment.  For information about high availability for the Informatica domain, see the <a href="#">Informatica documentation</a> .

Parameter label (name)	Default	Description
Deployment Type	Requires input	<p>You can select one of the following sizes:</p> <ul style="list-style-type: none"> <li>- Select the Amazon EC2 instance type where you deploy the Informatica cluster. Choose from the following options for the Informatica cluster: <ul style="list-style-type: none"> <li>- Small: one m*.4xlarge node</li> <li>- Medium: three m*.2xlarge nodes</li> <li>- Large: six m*.2xlarge nodes</li> </ul> </li> <li>- Choose from the following options for the Informatica compute cluster on EMR: <ul style="list-style-type: none"> <li>- Small: one m*.4xlarge node</li> <li>- Medium: one to three m*.2xlarge nodes</li> <li>- Large: one to six m*.2xlarge nodes</li> </ul> </li> </ul> <p><b>Note:</b> The instance types you can select might depend on your AWS region.</p>
S3 Bucket Name	N/A	<p>Specify the name of an existing S3 bucket.</p> <p><b>Note:</b> This bucket is not the same as the bucket that contains the Enterprise Data Preparation license key.</p> <p>If you do not specify a value, the deployment process creates an S3 bucket.</p>
Cluster ID for Amazon EMR	N/A	<p>Specify the identifier for an existing Amazon EMR cluster to use for the Informatica compute cluster.</p> <p>If you do not specify a value, the deployment process creates a cluster with autoscaling enabled.</p> <p>The number of EC2 nodes in the cluster depends on the Deployment Type value you set.</p> <p>For more information about cluster workflows in Intelligent Data Streaming, see <a href="#">Cluster Workflows</a>.</p>
Use S3 files	No	<p>Use configuration files stored on an Informatica-owned S3 bucket. Selecting Yes downloads configuration files and then configures. May contain latest fixes if any, else the deployment uses local configuration files stored in EC2 instance.</p>

3. Click **Next**.
4. On the **Configure stack options** page, you can specify tags (key-value pairs) for resources in your stack and set advanced options. When you are done, choose **Next**.
5. On the **Review** page, review and confirm the template settings. Under Capabilities, select the check box to acknowledge that the template will create IAM resources.
6. Choose **Create** to deploy the stack.
7. Monitor the status of the stack. When the status is `CREATE_COMPLETE`, the deployment is complete.
8. Click the **Outputs** tab to find the URLs for the Informatica Administrator, Catalog, and Enterprise Data Preparation applications.

See the [Enterprise Data Preparation User Guide](#) for information about logging in to Enterprise Data Preparation.

## Step 4. Recycle the Informatica Services

After the deployment completes, use the Administrator tool to recycle the Informatica application services.

1. Log in to the Administrator tool.
2. Click the **Services and Nodes** tab.
3. In the Domain Navigator, select the Model Repository Service.
4. From the **Actions** menu, select **Search Index > Re-index**.

5. From the **Search Index > Re-index** menu, select **Recycle Service**.
6. Select the Data Integration Service, and then select **Actions > Recycle Service**.
7. Select the Content Management Service, and then select **Actions > Recycle Service**.
8. Select the Interactive Data Preparation Service, and then select **Actions > Recycle Service**.
9. Select the Enterprise Data Preparation Service, and then select **Actions > Recycle Service**.
10. Log in to the Enterprise Data Preparation application.
11. Log out of the Enterprise Data Preparation application, and then log in again.

## Monitoring Instance Provision and Informatica Domain Creation

You can use cloud platform dashboards, logs, or other artifacts to see whether cluster creation succeeded and how to locate and identify the Informatica domain and Enterprise Data Preparation on the cloud platform.

### Deployed AWS Resources

You can monitor the progress of the solution deployment in the Create Stack section of the AWS portal. Each row of the display shows stack components and creation status. The stack is complete when the status reads CREATE\_COMPLETE.

When you deploy the solution to an existing VPC, the console displays the following stack components:

Component	Description
StackName/EDPStack	The main stack. Created by default.
StackName-LicenseCheckLambdaStack	Checks whether the license key file is valid. Created by default.
StackName-KeypairCheckLambdaStack	Checks whether the keypair file is valid. Created by default. <b>Note:</b> Applicable only for Linux deployments.
StackName-BastionSubnetCheckLambdaStack	Checks whether the bastion server subnet is valid. Created by default. Lambda function will be created only when: - Pre-validation check parameter set to Yes - Deploy Bastion Server set to Yes
StackName-SubnetsAZCheckLambdaStack	Checks whether the subnets are in different Availability Zones. Created by default.
StackName-PublicIpDnsCheckLambdaStack	Checks whether an auto-assign public IPv4 address is enabled for a subnet. Created by default.
StackName-Subnet2ChecklambdaStack	Checks whether the second subnet is valid. Created by default. Creates the Lambda function if the Pre-validation Check parameter is set to Yes.

<b>Component</b>	<b>Description</b>
StackName-Subnet1CheckLambdaStack	Checks whether the first subnet is valid. Created by default. Creates the Lambda function if the Pre-validation Check parameter is set to Yes.
StackName-IAMStack	Creates IAM roles and profiles used by EC2 instances. Created by default.
StackName-SecurityGroupStack	Creates security groups and network interfaces used by EC2 instances. Created by default.
StackName-RDSStack	Creates RDS resources. Created by default.
StackName-BastionServerStack	Creates bastion server on an EC2 instance. Created when you choose to create the bastion server.
StackName-S3BucketStack	Creates an Amazon S3 bucket. Created when you leave the S3 Bucket Name value for Enterprise Data Preparation and Amazon EMR empty.
StackName-EMRStack	Creates an Amazon EMR cluster. Created when you leave the Cluster ID for Amazon EMR value empty.
StackName-ICSStack	Creates an Informatica cluster. Created by default.
StackName-ICSStack-ICSGatewayStack	Creates an Informatica cluster. Created when you set Deployment Type to Small. Medium, or Large.
StackName-ICSStack-MultiNodeICNode2Stack	Creates an Informatica cluster. Created when you set Deployment Type to Medium or Large.
StackName-ICSStack-MultiNodeICNode3Stack	Creates an Informatica cluster. Created when you set Deployment Type to Medium or Large.
StackName-ICSStack-MultiNodeICNode4Stack	Creates an Informatica cluster. Created when you set Deployment Type to Large.
StackName-ICSStack-MultiNodeICNode5Stack	Creates an Informatica cluster. Created when you set Deployment Type to Large.
StackName-ICSStack-MultiNodeICNode6Stack	Creates an Informatica cluster. Created when you set Deployment Type to Large.

Component	Description
StackName-EDPServerStack	Creates an Informatica gateway node on which Enterprise Data Preparation runs on an EC2 instance. Created by default.
StackName-EDPHANodeStack	Creates a high availability node on which Enterprise Data Preparation runs on an EC2 instance. Created when you set High Availability Deployment to Enabled.

When you deploy to a new VPC, the console displays the following stack components:

Component	Condition for Creation
StackName-VPCStack	Creates network-related resources including the VPC, subnets, internet gateway, and the NAT gateway. Created by default.
EDPStack	The main stack. Created by default.

## Output Tab Properties

Click the **Output** tab to view the list of resources that the deployment process created in the VPC. The following table lists the outputs:

Output	Description
AdminConsoleUser	User name of the Administrator tool user.
AdminConsoleUserPwd	Password for the Administrator tool user.
BastionPrivateDNS	DNS of the private bastion server. For example: ip-<IPv4 address>-<region designation>.compute.internal
BastionPublicIPDNS	DNS of the public bastion server. For example: ip-<IPv4 address>-<region designation>.compute.amazonaws.com
DBUserPwd	Password for the database users of the domain databases created in the deployment.
DBUsers	User IDs for the database users of the domain databases created in the deployment.
HANodeIP	IPv4 address of the high availability or failover node for the domain.
InformaticaAdminConsolePrivateURL	URL:port of the Administrator tool hosted on a private server.
InformaticaAdminConsolePublicURL	Public URL:port of the Administrator tool.

Output	Description
InformaticaAdminConsoleServerLog	Path to the domain installation log.
InformaticaEDPPrivateURL	URL:port of the Enterprise Data Preparation application user interface hosted on a private server.
InformaticaEDPPublicURL	Public URL:port of the of the Enterprise Data Preparation application user interface.
InstanceSetupLogs	Path to the domain EC2 host setup log.

Use the URL, user ID, and password to access the Administrator tool to perform administrator tasks.

For a description of domain and Enterprise Data Preparation administrator tasks, see the [Informatica Administrator Guide](#) and the [Informatica Enterprise Data Preparation Administrator Guide](#).

## Deployed Informatica Resources

The list of deployed resources includes the **InformaticaAdminConsolePublicURL**. Click the linked value of this property.

The following image shows an example of this property, with the IP address portion redacted:

Key	Value
InformaticaAdminConsolePublicURL	<a href="https://ec2-██████████.ap-south-1.compute.amazonaws.com:8443">https://ec2-██████████.ap-south-1.compute.amazonaws.com:8443</a>

The Administrator tool opens in a browser window.

Use the Informatica Administrator user name and password listed in the **Outputs** list to log in to the Administrator tool.

You can use the Administrator tool to view and administer the Informatica application services and resources that the automated deployment created.

## Solution Deployment Log Files

You can review the following log files to see the solution deployment as it happened.

Review the following log files in the `/var/log` directory:

- `cfn-init.log`
- `cfn-init-cmd.log`
- `cfn-wire.log`
- `cloud-init.log`
- `messages`

Review the following log file in the `/opt` directory:

- `Infa_OneClick_Solution.log`

## Informatica Domain Logs

You can access the following logs on the Informatica domain VM:

## Command execution log

This log records the following events:

- Creation of Informatica connections, cluster configurations, and services.
- Population of the data warehouse and SQL databases.
- Data Integration Service recycling to register all changes.

At the top of the log file is a summary section that lists automated tasks and their status. Beneath the summary section are detailed sections about each task. If any of the tasks failed to complete successfully, you can look at the detailed section for the task to troubleshoot the task.

File name: `Infa_OneClick_Solution.log`

Location: `/opt`

Node preparation log file: `/opt/Infa_node_configure.log`

## CloudWatch logs for Informatica

CloudWatch is an AWS service that enables you to monitor all of the events in your AWS deployments. CloudWatch provides a central location from which to monitor EC2 instances and other resources. For more information, see the [AWS documentation](#).

To monitor Enterprise Data Preparation events in CloudWatch, go to the CloudWatch page and search for "Informatica."

If no log files are returned by the search, log into the VM where the Informatica domain is installed and check for logs with the filename "`cfm*.log`"

Access the Informatica deployment logs through the AWS console:

1. In the Amazon Web Services Console, select CloudWatch from the Services menu.
2. In the navigation pane, choose **Logs > Log groups**.
3. Enter `InformaticaEDPLogGroup` in the search pane.

## CloudWatch logs for Lambda functions

If the automated deployment fails during execution, CloudWatch logs display information about process errors.

Access the CloudWatch logs through the AWS console:

1. In the Amazon Web Services Console, select CloudWatch from the Services menu.
2. In the navigation pane, choose **Logs > Log groups**.
3. Enter `/aws/lambda/<stack name>` in the search pane.
4. Select the log group associated with the name of the stack that failed.

# Troubleshooting

Review the AWS log files to troubleshoot issues with your Enterprise Data Preparation deployment.

## Stack Failure

Review the following log files if an EC2 instance where Informatica services run is running but the stack fails, or if the Administrator tool fails to launch:

- Node preparation log file: `/opt/Infa_node_configure.log`
- Installation log file: `/home/Infa_OnceClick_Solution.log`

- Command Initialization log file: `/var/log/cfn-init-cmd.log`
- Service Initialization log file: `/var/log/cfn-init.log`
- Cloud log file: `/var/log/cloud-init.log`

## Stack Launch Failure When the Pre-validation Check Parameter is Enabled

Review the AWS CloudWatch log files if a stack failure occurs when the Pre-validation Check option is enabled.

1. In the AWS Management Console, select **CloudWatch** from the **Services** menu.
2. Select **Logs**, then and then select **Log Groups**.
3. Enter the following string in the search bar:  
`/aws/lambda/<stack name>`
4. Click the log group associated with the stack.

## Stack Launch Failure When Rollback On Failure Is Enabled

Review the AWS CloudWatch log files if a stack failure occurs when the **Rollback on Failure** option is enabled.

1. In the AWS Management Console, select **CloudWatch** from the **Services** menu.
2. Select **Logs**, then and then select **Log Groups**.
3. Enter the following string in the search bar: `InformaticaEDPLogGroup`
4. Click a log file in the log group associated with the stack.

You might want to review the following log files:

- `installation.log`
- `cfn-init`
- `InformaticaClusterService`
- `EDCService`
- `EDPService`

## Frequently asked questions on deployment error messages

### I encountered a **CREATE\_FAILED** error when I launched the deployment.

If AWS CloudFormation fails to create the stack, we recommend that you relaunch the template with **Rollback on failure** set to **No**. (This setting is under **Advanced** in the AWS CloudFormation console, **Options** page.) With this setting, the stack's state will be retained, and the instance will be left running, so you can troubleshoot the issue. (You will want to look at the log files in `/var/log/cfn-init-cmd.log`)

**Note:** When you set **Rollback on failure** to **No**, you will continue to incur AWS charges for this stack. Please make sure to delete the stack when you have finished troubleshooting.

For additional information, see [Troubleshooting AWS CloudFormation](#) on the AWS website or contact us on the [AWS Discussion Forum](#).

### I encountered errors while installing the Informatica domain and services. Where can I find more information?

View the `/opt/Infa_OneClick_Solution.log` log file to see more information about the errors that you encountered.



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

We recommend that you launch the deployment templates from the location we provided or from another S3 bucket. If you deploy the templates from a local copy on your computer or from a non-S3 location, you might encounter template size limitations when you create the stack. For more information about AWS CloudFormation limits, see the [AWS documentation](#).

### What should I do before I relaunch the template with Rollback on failure?

Ensure that you disable rollback when you launch the stack in case of a stack failure. If Enterprise Data Preparation server stack is created, check the following logs to identify the issues:

- Node preparation log file: /opt/Infa\_node\_configure.log
- /opt/Infa\_OneClick\_Solution.log
- /var/log/cfn-init.log
- /opt/Informatica/<Infa\_Oracle>/logs/InfaNode/services/InfaHadoopService/Informatica\_Cluster\_Service/ICS.log
- /opt/Informatica/<Infa\_Oracle>/logs/InfaNode/services/CatalogService/Catalog\_Service/LDM.log

If the Enterprise Data Preparation server stack is not created, you can check the CloudWatch logs to identify the issues. For more information, see the Verifying AWS CloudWatch Logs for Failures section.

### I am facing Model Repository Service connectivity issues with the following error messages:

```
MRS_50015 "The Repository Service operation failed. ['[PRSVCSHARED_01707] Connection issues with the configured database. Request failed with the error message [[PERSISTENCEAPI_0307] [DBPERSISTER_1005] Failed to process requested operation. This was caused by [informatica][Oracle JDBC Driver]No more data available to read.].']"
```

```
MRS_50015 "The Repository Service operation failed. ['[PRSVCSHARED_01707] Connection issues with the configured database. Request failed with the error message [[PERSISTENCEAPI_0307] [PERSISTENCECOMMON_0001] Internal error. The request processing failed. This was caused by [informatica][Oracle JDBC Driver]Exception generated during deferred local transaction handling. See next exception via SQLException.getNextException for details.].']"
```

Applicable for Oracle. Perform the following steps to resolve the Model repository database connectivity issue:

1. Launch the AWS RDS page.
2. Select the RDS instance.
3. Click the **Configuration** tab.
4. Scroll down the tab page and click the link under the **Parameter group** section.
5. Search for the following parameters in the list of parameters that appear:
  - sqlnetora.sqlnet.send\_timeout
  - sqlnetora.sqlnet.recv\_timeout
  - sqlnetora.sqlnet.inbound\_connect\_timeout
  - sqlnetora.sqlnet.outbound\_connect\_timeout
  - sqlnetora.tcp.connect\_timeout

**Note:** If the listed parameters are not present in the list, make sure that add the parameters to the list.

6. Change the value of each of the parameters listed in the previous step to 0 as shown in the following sample image:

The following image shows the `sqlnetora.sqlnet.send_timeout` parameters

Parameters Cancel editing Preview changes Reset Save changes

Q sqlnetora.sqlnet.recv\_timeout X < 1 > ⌘

<input type="checkbox"/>	Name	Values	Allowed values	Modifiable	Source	Apply type	Data type	Description
<input type="checkbox"/>	sqlnetora.sqlnet.recv_timeout	0	0, 10-7200	true	user	dynamic	integer	The buffer space limit for receive operations of sessions, supported by TCP/IP, TCP/IP with SSL, and SDP protocols.

7. Click **Save changes**.

## Additional resources

### AWS Services

- [AWS CloudFormation](#)
- [Amazon EBS](#)
- [Amazon EC2](#)
- [Amazon VPC](#)

### Enterprise Data Preparation Documentation

- [Product Documentation](#)

### Reference Deployments

- [AWS Deployment home page](#)

## Author

**Saurabh Sugandh**