



# Tungsten MI (TMI)

Continuent Ltd

---

# Tungsten MI (TMI)

Continuent Ltd

Copyright © 2017, 2018, 2019, 2020, 2021, 2022, 2023, 2024 Continuent Ltd

## Abstract

This manual documents Tungsten MI (TMI) 7.1.4.

*Build date:* 2024-12-17 [a004bbb8]

Up to date builds of this document: [Tungsten MI \(TMI\) \[Online\]](#), [Tungsten MI \(TMI\) \[PDF\]](#)

---

# Table of Contents

1. Introduction .....	4
2. Getting Started with Tungsten MI for Tungsten Clustering .....	5
2.1. Launching from the Amazon [AWS] Marketplace .....	5
2.1.1. Launching Manually .....	5
2.1.2. Launching via Cloudformation Templates .....	5
2.2. Launching from the Google [GCP] Marketplace .....	6
3. Getting Started with Tungsten MI for Tungsten Replicator .....	7
3.1. Launching from the Amazon [AWS] Marketplace .....	7
3.2. Launching from the Google [GCP] Marketplace .....	8
4. Configuring Instances through the Launch Wizard .....	10
A. Release Notes .....	11
A.1. Tungsten MI [TMI] 7.1.4 GA [12 December 2024] .....	11
A.2. Tungsten MI [TMI] 7.1.2 GA [9 April 2024] .....	11
A.3. Tungsten MI [TMI] 7.0.1 GA [20 June 2022] .....	11
A.4. Tungsten MI [TMI] 5.0.0 GA [10 May 2020] .....	11
A.5. Tungsten MI [TMI] 4.0.0 GA [17 December 2020] .....	11
A.6. Tungsten MI [TMI] 3.0.0 GA [20 July 2020] .....	12
A.7. Tungsten MI [TMI] 2.0.0 GA [11 November 2019] .....	12
A.8. Tungsten MI [TMI] 1.0.0 GA [6 June 2019] .....	12

---

## Chapter 1. Introduction

Tungsten MI refers to a number of pre-configured TMI's (Tungsten Machine Images) available to subscribe to via the Amazon AWS or Google GCP Marketplaces. A number of TMI's exist for each valid topology/use case, and are grouped into two categories.

- Tungsten MI for Tungsten Clustering consists of a pre-configured TMI containing all the necessary pre-requisites to launch Tungsten Clustering™ and Tungsten Dashboard.
- Tungsten MI for Tungsten Replicator consists of various pre-configured TMI's containing all the necessary pre-requisites to launch a full end-to-end replication pipeline using Tungsten Replicator™

This set of documentation is intended as a guide for launching the TMI's and getting started. For documentation specific to the features, management and administration of Tungsten Cluster, Tungsten Replicator or Tungsten Dashboard please review the product specific documentation.

---

## Chapter 2. Getting Started with Tungsten MI for Tungsten Clustering

The process to launch TMI will depend upon the cloud environment you are launching into. Follow the guides below appropriate the Cloud provider you are using.

- [Section 2.1, "Launching from the Amazon \[AWS\] Marketplace"](#)
- [Section 2.2, "Launching from the Google \[GCP\] Marketplace"](#)
- Microsoft Azure (Coming Soon)

### 2.1. Launching from the Amazon [AWS] Marketplace

You can launch your instances in Amazon AWS either by launching the instances via the Marketplace console, or via pre-configured Cloudformation Templates.

First you will need to subscribe through the Amazon Marketplace by clicking [here](#), or by simply searching for "Continuent" in the AWS Catalog.

You can choose an annual subscription or an hourly based usage subscription. The Amazon Marketplace will provide all the information required to allow you to subscribe accordingly.

Once subscribed you can then start to configure your instance. When choosing the instance size there are a number of factors to take into account

- Which topology are you deploying? For a standard 3-node cluster the minimum instance should be at least a t2.large. For advanced topologies you should consider at least a t2.xlarge [or the equivalent depending on the availability within the region you are deploying]. The larger your anticipated workload, then the larger the instance should be, and should primarily be determined by your MySQL Database needs.
- How large will your database be? By default the instance is launched with a 40Gb partition which should be suitable for most small to medium workloads. You should consider increasing this as required, based on your MySQL Database needs, taking into consideration that binary logging must be enabled and there will be an almost 1:1 relationship between binary logs and THL Logs.

Once you have decided on your instance size you can launch following the steps in the next sections for the method you choose, either Manual, or Cloudformation.

#### 2.1.1. Launching Manually

Quantity of instances and their Location

- Standard Clusters : You will need a minimum of 3-nodes, for any greater, there must be an odd number to maintain quorum, for example 3,5,7,9 etc. All nodes must be in the same region however they can, and probably should, be in different availability zones.
- Composite Active/Passive & Composite Active/Active Clusters : These topologies consist of multiple clusters connected together, each cluster must conform to the same node count and rules for a standard cluster, and there must be a minimum of 2 clusters. Therefore for these topologies you will need a minimum of 6 nodes [3 per cluster]. Each cluster may reside in a different region, allowing you to configure full geo-distribution.
- For more details and to understand the different topologies, see [Deployment: MySQL Topologies](#). Whilst the launch wizard on the launched VM will assist you to pre-configure either a Standard cluster, a Composite Active/Passive cluster or a Composite Active/Active cluster, users with advanced knowledge can adjust and configure the nodes to suit any required clustering topology available.

By default, the instances only launch with the default SSH security group. Based on your business rules, you will need to configure the appropriate security group to allow communication between all nodes within the cluster, and to allow your applications to communicate with the proxy nodes. A full breakdown of the network requirements can be found here: [Configuring Network and SSH Environment](#)

After launching your instances, you can now proceed to configure them, see [Chapter 4, Configuring Instances through the Launch Wizard](#) for instructions

#### 2.1.2. Launching via Cloudformation Templates

There are three cloudformation templates available that will configure one of the three available topologies. They are preconfigured to launch the nodes and configure all of the necessary networking and security groups for you based on the parameters that you provide.

The list below outlines the template configuration. To launch the cloudformation template, simply click on the topology name. Be sure to check the region from the AWS dashboard before launching as this will determine the destination region for your deployment.

- **Single Standalone Cluster** : This cluster will be configured with 3 nodes. The primary and one of the replica nodes will reside in one AZ, and the remaining replica will be in a different AZ. All 3 nodes must be, and will be, in the same region.
- **Composite Active/Passive Cluster** : This topology will exist of 2 clusters, containing 3 nodes within each cluster. One cluster will be the primary Read/Write cluster, the second will be Read-Only. Each cluster will reside in the same region. Within the cluster, the nodes will be distributed across all 3 AZ's.
- **Composite Active/Active Cluster** : This topology will exist of 2 clusters, containing 3 nodes within each cluster. Both clusters will be Read/Write. Each cluster will reside in the same region. Within the cluster, the nodes will be distributed across all 3 AZ's.

#### Note

The node count in the clusters when deploying via cloudformation cannot be changed at this time, therefore if you wish to deploy more than 3 nodes per cluster, you will need to deploy via the manual method outlined above.

#### Note

For the Composite Active/Passive and Composite Active/Active deployments, all nodes and clusters will reside in the same region. If you wish to deploy a geo-distributed topology, with each cluster in a different region, you will need to deploy via the manual method outlined above.

After the launch process has complete you will be able to connect to your instances and access all of the features of Tungsten Clustering. For more full details on the various command-line tools and cluster operation, review the complete set of documentation [here](#)

To connect to the dashboard, using the browser of your choice connect via http to the public-ip of one of the instances. For more information on the Dashboard, review the documentation [here](#)

## 2.2. Launching from the Google [GCP] Marketplace

You can launch your instances in GCP via the [Marketplace](#). Simply search for "Continuent" to view all the listings available. For the Clustering solution you will need to subscribe to "Tungsten Clustering for MySQL", and then clicking the "Get Started" button.

Once subscribed you can then start to configure your instance. When choosing the instance size there are a number of factors to take into account

- Which topology are you deploying? For a standard 3-node cluster the minimum instance should be at least a c3-standard-4. For advanced topologies you should consider at least a c3-standard-8 [or the equivalent depending on the availability within the region you are deploying]. Ideally, you should also consider ssd for the disk. The larger your anticipated workload, then the larger the instance should be, and should primarily be determined by your MySQL Database needs.
- How large will your database be? By default the instance is launched with a 40Gb partition which should be suitable for most small to medium workloads. You should consider increasing this as required, based on your MySQL Database needs, taking into consideration that binary logging must be enabled and there will be an almost 1:1 relationship between binary logs and THL Logs.

Quantity of instances and their Location

- **Standard Clusters** : You will need a minimum of 3-nodes, for any greater, there must be an odd number to maintain quorum, for example 3,5,7,9 etc. All nodes must be in the same region however they can, and probably should, be in different zones.
- **Composite Active/Passive & Composite Active/Active Clusters** : These topologies consist of multiple clusters connected together, each cluster must conform to the same node count and rules for a standard cluster, and there must be a minimum of 2 clusters. Therefore for these topologies you will need a minimum of 6 nodes [3 per cluster]. Each cluster may reside in a different region, allowing you to configure full geo-distribution.
- For more details and to understand the different topologies, see [Deployment: MySQL Topologies](#). Whilst the launch wizard on the launched VM will assist you to pre-configure either a Standard cluster, a Composite Active/Passive cluster or a Composite Active/Active cluster, users with advanced knowledge can adjust and configure the nodes to suit any required clustering topology available.

Based on your business rules, you will need to configure the appropriate network rules to allow you to access the host via your chosen terminal application, additionally you will need to allow communication between all nodes within the cluster, and to allow your applications to communicate with the proxy nodes. A full breakdown of the network requirements can be found here: [Configuring Network and SSH Environment](#)

After launching your instances, you can now proceed to configure them, see [Chapter 4, Configuring Instances through the Launch Wizard](#) for instructions

## Chapter 3. Getting Started with Tungsten MI for Tungsten Replicator

The process to launch TMI will depend upon the cloud environment you are launching into. Follow the guides below appropriate the Cloud provider you are using.

- [Section 3.1, "Launching from the Amazon \[AWS\] Marketplace"](#)
- [Section 3.2, "Launching from the Google \[GCP\] Marketplace"](#)
- Microsoft Azure (Coming Soon)

A number of different MI's are available for different use cases, in all cases you will need to launch at least one "Extractor" and the appropriate applier for your use case. The table below lists all of the images available in all cloud environments.

Name	Description	AWS	GCP
Tungsten Replicator for MySQL Extraction	Required for all use-cases. Can be configured to extract from a remote MySQL Instance, a local MySQL Instance (must be manually configured first), Amazon Aurora, Google MySQL or Azure Cloud SQL. All flavours and version of MySQL supported up to and including MySQL 8.4	Y	Y
Tungsten Replicator for MySQL Appliers	Can be configured to apply to a remote MySQL Instance, a local MySQL Instance (must be manually configured first), Amazon Aurora, Google MySQL or Azure Cloud SQL. All flavours and version of MySQL supported up to and including MySQL 8.4	Y	Y
Tungsten Replicator for PostgreSQL Appliers	Can be configured to apply to a remote PostgreSQL Instance, a local Instance (must be manually configured first).	Y	Y
Tungsten Replicator for Oracle Appliers	Can be configured to apply to a remote Oracle Instance, a local Instance (must be manually configured first).	Y	Y
Tungsten Replicator for Kafka Appliers	Can be configured to apply to a remote Kafka Instance, a local Instance (must be manually configured first). Currently requires zookeeper.	Y	Y
Tungsten Replicator for MongoDB	Can be configured to apply to a remote MongoDB Instance, a local Instance (must be manually configured first) or MongoDB Atlas.	Y	Y
Tungsten Replicator for Hadoop Appliers	Can be configured to apply to a remote Hadoop cluster.	Y	Y
Tungsten Replicator for Vertica Appliers	Can be configured to apply to a remote Vertica cluster.	Y	Y
Tungsten Replicator for Redshift Appliers	Can be configured to apply to a Amazon Redshift. Requires S3 access and only available as an AWS AMI.	Y	N
Tungsten Replicator for S3 Appliers	Coming Soon...	N	N

### 3.1. Launching from the Amazon [AWS] Marketplace

You can launch your instances in Amazon AWS by launching the instances via the Marketplace console.

First you will need to subscribe through the Amazon Marketplace by clicking [here](#), or by simply searching for "Continuent" in the AWS Catalog.

You can choose an annual subscription or an hourly based usage subscription. The Amazon Marketplace will provide all the information required to allow you to subscribe accordingly.

Once subscribed you can then start to configure your instance. When choosing the instance size there are a number of factors to take into account

- Images do not deploy with a database installed, in most cases a t3.medium will be sufficient, however should you choose to install a database locally, then you should increase the instance size to suits the needs of the database you are configuring.
- How much data are you replicating? Each instance launches with a 40Gb partition. Tungsten Replicator will store THL logs on this partition. A THL log is similar to a MySQL binary log and contains all the transactions and data we are replicating. These files are kept by default for 7 days. If you are replicating significant amounts of data you may need to increase this volume and/or consider reducing the THL retention which can be done during the instance configuration.

Quantity of instances and their Location

- As mentioned earlier, you will need at least one "Tungsten Replicator for MySQL Extraction". Ideally this should reside in the same region and zone as the source database that you are extracting from.
- You will need at least one "Tungsten Replicator for xxxx Appliers" where xxxx relates to the target type. Ideally this should reside in the same region and zone as the target database.
- The source and target replicator instances do not have to reside in the same region, and can replicate globally. They can even be used to replicate into and out of AWS between on premise instances and other cloud providers.
- For more details and to understand the different topologies, see [Deploying MySQL Extractors](#) and [Deploying Appliers](#). Whilst the launch wizard on the launched VM will assist you to pre-configure a single extractor or applier service, users with advanced understanding of Tungsten Replicator can manually configure for more advanced topologies such as Fan-In and Fan-Out.

By default, the instances only launch with the default SSH security group. Based on your business rules, you will need to configure the appropriate security group to allow communication between all nodes within the topology, and to allow your hosts to communicate with the source and target database instances. A full breakdown of the network requirements can be found here: [Configuring Network and SSH Environment](#)

After launching your instances, you can now proceed to configure them, see [Chapter 4, Configuring Instances through the Launch Wizard](#) for instructions

## 3.2. Launching from the Google [GCP] Marketplace

You can launch your instances in GCP via the [Marketplace](#). Simply search for "Continuent" to view all the listings available.

Once subscribed you can then start to configure your instance. When choosing the instance size there are a number of factors to take into account

- Images do not deploy with a database installed, in most cases a t3.medium will be sufficient, however should you choose to install a database locally, then you should increase the instance size to suits the needs of the database you are configuring.
- How much data are you replicating? Each instance launches with a 40Gb partition. Tungsten Replicator will store THL logs on this partition. A THL log is similar to a MySQL binary log and contains all the transactions and data we are replicating. These files are kept by default for 7 days. If you are replicating significant amounts of data you may need to increase this volume and/or consider reducing the THL retention which can be done during the instance configuration.

Quantity of instances and their Location

- As mentioned earlier, you will need at least one "Tungsten Replicator for MySQL Extraction". Ideally this should reside in the same region and zone as the source database that you are extracting from.
- You will need at least one "Tungsten Replicator for xxxx Appliers" where xxxx relates to the target type. Ideally this should reside in the same region and zone as the target database.
- The source and target replicator instances do not have to reside in the same region, and can replicate globally. They can even be used to replicate into and out of GCP between on premise instances and other cloud providers.
- For more details and to understand the different topologies, see [Deploying MySQL Extractors](#) and [Deploying Appliers](#). Whilst the launch wizard on the launched VM will assist you to pre-configure a single extractor or applier service, users with advanced understanding of Tungsten Replicator can manually configure for more advanced topologies such as Fan-In and Fan-Out.



Based on your business rules, you will need to configure the appropriate network rules to allow you to access the host via your chosen terminal application, additionally you will need to allow communication between all nodes within the deployment, and to allow your instances to communicate with the source and target database instances. A full breakdown of the network requirements can be found here: [Configuring Network and SSH Environment](#)

After launching your instances, you can now proceed to configure them, see [Chapter 4, Configuring Instances through the Launch Wizard](#) for instructions

---

## Chapter 4. Configuring Instances through the Launch Wizard

### Note

These steps do not apply if you launched a Tungsten Cluster via Cloudformation

Once you have launched all of the required instances, you can now start to configure them.

A very simple CLI wizard will be launched when you connect to the instance and prompt you for information to configure the instances.

- For Tungsten Clustering instances, you only need to run the wizard on one node, all other nodes will retrieve the configuration ensuring consistency across the cluster.
  - For Tungsten Replicator instances, you will need to run the wizard on every instance, ensuring the extractor instance has been completed first, or is already running from an existing installation.
  - If you are planning to replicate from an existing Tungsten Cluster, you only need a Tungsten Replicator Applier instance, and when prompted, configure as a "cluster-slave"
1. First, you will need to connect to your instance, the method will vary depending on your cloud provider and also on your choice of Terminal Application. For AWS you will initially connect to the instance as the "rocky" OS user. For Google, this will depend on your method of authentication. An example of each is shown below:

```
AWS> ssh -i .ssh/yoursshkey.pem rocky@public-ip-of-host
```

```
GCP> gcloud compute ssh your-instance-name --zone=instance-zone --project=your-gcp-project
```

2. Next, sudo as the tungsten user:

```
shell> sudo su - tungsten
```

3. The wizard will now launch, first showing a readme outlining any requirements that may need to be completed manually such as installing MySQL if required.

Step through the wizard providing answers to various prompts

4. If you are configuring Tungsten Clustering, repeat steps 1 & 2 on all hosts. They will pickup the configuration from the first host and automatically configure themselves, only then proceed to the next step
5. When complete, you can let the wizard begin the installation, or wait and execute this last step manually. The instructions will be displayed on the screen should you wish to run manually, these are also shown below as an example:

```
shell> cd /opt/continuent/software/tungsten-replicator-for-mysql-extraction-7.1.4-11  
shell> tools/tpm install
```

6. After installation, to ensure the environment variables are setup correctly, issue:

```
shell> source /opt/continuent/share/env.sh
```

### Note

This step only needs executing once, for subsequent terminal sessions the environment will be loaded when you connect to the terminal as the tungsten OS user.

7. You can now start the software by issuing the command `startall`
8. For Tungsten Clustering users, you will also now be able to access the Tungsten Dashboard by connecting to the public IP via a web-browser, providing you have opened the HTTP port within your security firewall

For advanced configurations, review the full product documentation by clicking the relevant product name in the top toolbar

For Tungsten Replicator installations, the replicators will begin replicating from the current binary log position at the point at which the replicator was started. To start the replicator from an earlier binary log position and for more advanced options, review [The trepctl Command](#)

For heterogeneous targets, you will need to ensure that your database structure has been manually created first, alternatively you can use the [ddlscan](#) tool to help reverse engineer the database structure. More information on this can be found at [The ddlscan Command](#)

Each heterogeneous applier may have different requirements, these are explained in the relevant sections at [Deploying Appliers](#)

---

## Appendix A. Release Notes

### A.1. Tungsten MI [TMI] 7.1.4 GA [12 December 2024]

Version End of Life. Not Yet Set

This release marks a change in name from Tungsten AMI which was specific to Amazon AWS, to Tungsten MI [TMI] which coincides with the launch of the images to the Google GCP Marketplace.

To ensure consistency across cloud providers, the AWS images now launch on Rocky Linux 9.5, previously Amazon Linux 2. GCP images also launch with the same.

These images launch with version 7.1.4-11 of the Tungsten suite of products. The product specific release notes for this release can be found at [Tungsten Clustering 7.1.4 Release Notes](#) and [Tungsten Replicator 7.1.4 Release Notes](#)

#### Important

The "Tungsten Clustering for MySQL 5.7" images on AWS are no longer available

### A.2. Tungsten MI [TMI] 7.1.2 GA [9 April 2024]

Version End of Life. Not Yet Set

This AMI release fixes a few issues with the launch wizard, which now allows easy configuration with IPv6 hosts.

It was launched with version 7.1.2-42 of the Tungsten suite of products. The product specific release notes for this release can be found at [Tungsten Clustering 7.1.2 Release Notes](#) and [Tungsten Replicator 7.1.2 Release Notes](#)

### A.3. Tungsten MI [TMI] 7.0.1 GA [20 June 2022]

Version End of Life. Not Yet Set

For this release, the version number of the AMI was adjusted to be in line with the version of the Tungsten Products that are shipped with the images.

This AMI release was launched with version 7.0.1-96 of the Tungsten suite of products. The product specific release notes for this release can be found at [Tungsten Clustering 7.0.1 Release Notes](#) and [Tungsten Replicator 7.0.1 Release Notes](#)

### A.4. Tungsten MI [TMI] 5.0.0 GA [10 May 2020]

Version End of Life. Not Yet Set

This AMI release resolves Amazon Linux OS vulnerabilities identified by AWS and a number of small improvements in the deployment Wizard for manual configurations.

It was launched with version 6.1.12-53 of the Tungsten suite of products. The product specific release notes for this release can be found at [Tungsten Clustering 6.1.12 Release Notes](#) and [Tungsten Replicator 6.1.12 Release Notes](#)

#### Bug Fixes

- Operating System/3rd Party Packages
  - Addresses OS security vulnerability identified by AWS (More info at <https://alas.aws.amazon.com/AL2/ALAS-2019-1367.html>)
  - Deployment Wizard now correctly configures the `tungsten.ini` for the MongoDB Atlas applier.

*Issues:* CT-1161

- When configuring a cluster-extractor to apply into AWS Redshift, the `s3.json` configuration file would be named incorrectly.

*Issues:* CT-1219

### A.5. Tungsten MI [TMI] 4.0.0 GA [17 December 2020]

Version End of Life. Not Yet Set

This AMI release brings the Tungsten Replicator and Tungsten Clustering AMI's in line with the latest release of the core products, in addition new Cloudformation Templates were released to allow easier deployment of Tungsten Clustering.

Improvements, new features and functionality

- Installation and Deployment
  - The MongoDB applier now supports replication into MongoDB Atlas.

Bug Fixes

- Operating System/3rd Party Packages
  - Addresses OS security vulnerability identified by AWS (More info at <https://alas.aws.amazon.com/AL2/ALAS-2019-1367.html>)

## A.6. Tungsten MI [TMI] 3.0.0 GA [20 July 2020]

Version End of Life. 17 December 2020

The v3.0 AMI release introduces two new AMI's for Tungsten Clustering.

Known Issues

- Installation and Deployment
  - Cloudformation support for Composite Active/Active and Composite Active/Passive clusters will only allow all clusters to be deployed within the same AWS region. This will be addressed in a future release.

## A.7. Tungsten MI [TMI] 2.0.0 GA [11 November 2019]

Version End of Life. 17 December 2020

The v2.0 AMI release includes additional Targets and a more simplified deployment wizard.

Improvements, new features and functionality

- Behavior Changes
  - Additional support added for the following Targets:
    - Oracle
    - Cassandra
    - Clickhouse
    - Elasticsearch
    - Hadoop
    - MongoDB
  - Support added to allow easy configuration of extraction from, or replication into, an existing Tungsten Cluster.
  - Installation Wizard now makes it easier to enable additional properties, such as SSL and basic filtering
  - Ships with the latest v6.1.1 release of Tungsten Replicator

## A.8. Tungsten MI [TMI] 1.0.0 GA [6 June 2019]

Version End of Life. 11 November 2019

The first release of the Tungsten AMI on the Amazon Marketplace is available and supports replication out of MySQL [Community Edition, Enterprise Edition, Percona, MariaDB, Amazon RDS/Aurora] into the following targets:

- MySQL [Community Edition, Enterprise Edition, Percona, MariaDB, Amazon RDS/Aurora]
- AWS Redshift
- Vertica

- PostgreSQL
- Kafka

Known Issues

- Installation and Deployment
  - Additional Targets will be supported in a future release.