

DEPLOYMENT & CONFIGURATION



Table of Contents

1. Deploying Metric Insights	4
1.1 What deployment options are available?	5
1.2 What are the system requirements for a Metric Insights server?	7
1.3 Docker Deployment (Simple Installer)	8
1.4 Deploying Metric Insights in Orchestrated Environment	16
1.5 Deploy Metric Insights as a Virtual Appliance	22
1.6 What are the main components of the Metric Insights stack?.....	32
1.7 Operating system linux package requirements / dependencies for v5.0+.....	34
1.8 Expose Metric Insights on VMware when using NAT networking	38
1.9 Deploy Metric Insights on Amazon EC2	40
1.10 Logging via Splunk (6.x).....	48
1.11 Where are all those Log Files? (6.0+)	55
1.12 Where are all those Log Files? (5.6.1 and lower).....	57
1.13 Accessing the Metric Insights Log files.....	58
1.14 Install or Update MI via Installation Packages	60
1.15 Deploying Metric Insights' Custom Mobile Apps at your organization.....	65
1.16 Deploying Mobile Apps with support for Deep Links (links to the Metric Insights Mobile App).....	69
1.17 MI Slackbot deployment	72
1.18 Operating system Linux package requirements for v6.1+	79
1.19 Deploying Container Orchestration with Docker Swarm	80
1.20 Docker Commands Cheat Sheet.....	84
1.21 Creating Amazon ECS Scheduled Tasks for MI utilities.....	87
1.22 Deploy Metric Insights on Kubernetes.....	91
1.23 Kubectl Commands Cheat Sheet	100
1.24 Administer Metric Insights in Kubernetes	102
1.25 Deploy Metric Insights on OpenShift	107
1.26 Docker Deployment on Windows Host.....	125
2. Using Templates	133
2.1 Setting Defaults for New Users.....	134
2.2 Creating a new Email Digest Template	138

2.3 Applying headers/footers to Email Templates.....	148
2.4 Simplified Email Templates	154
3. Configuring Metric Insights	157
3.1 Using the Branding feature	158
3.2 Setup email.....	162
3.3 How to Set the System Clock.....	165
3.4 Generating Empty Dataset Reports (when there is No Data).....	166
3.5 Generating Legacy Reports When There Is No Data	167
3.6 Debug Logging	169
3.7 Set Timezone (Virtual Appliance or CentOS/Redhat Linux)	170
3.8 Set Timezone (Debian Linux).....	172
3.9 Supported Time and Date Formats	173
3.10 Add new Measurement Time Calculation.....	175
3.11 Default Name/Description Values for Elements	178
3.12 Set Hostname for Metric Insights server	180
3.13 Set up Web SSL Certificate for Metric Insights	182
3.14 How do I specify a domain name when connecting to SQL Server?	187
3.15 Mobile App Security and Encryption Overview.....	188
3.16 Configure Intelligent Global Search	190
4. Troubleshooting Metric Insights	195
4.1 Troubleshooting Email Connectivity - Sending emails	196
4.2 Troubleshooting Email Connectivity - Receiving Email	203
4.3 Getting "ERROR: Can't write to /var/backups/mi-app-backups/mi-full-backup" when trying to backup my Metric Insights Instance.....	207
4.4 Will Metric Insights Run on Windows?	208

1. Deploying Metric Insights

1.1 What deployment options are available?

Overview of the various deployment options for Metric Insights

Overview of Deployment Option

The Metric Insights solution can be deployed in a variety of configurations that will best meet your organization's needs. We currently support the following options:

- Download and run a virtual appliance on a computer of your choice. More information can be found at [Deploy Metric Insights as a Virtual Appliance](#).
- Spin up a Metric Insights Amazon EC2 instance on the cloud. For more information see [Deploy Metric Insights on Amazon EC2](#).
- Provision a virtual or real linux server of your own that meets the Metric Insights hardware requirements. Please contact us if you are interested in this deployment option.

Which method is right for your organization depends on your budget and infrastructure. Here is a list of pros and cons for each method so you can choose which option is right for you.

Deployment option Comparison

Deployment Method	Pros	Cons	Best For
Virtual Appliance	<ul style="list-style-type: none"> • Cheap • Quick & easy • Works on a variety of operating systems 	<ul style="list-style-type: none"> • Networking options can be difficult 	Pilots / Companies with virtualization infrastructure such as VMWare vSphere or Microsoft Hyper-V
Amazon EC2	<ul style="list-style-type: none"> • No infrastructure support cost or waiting for IT to provision servers • Super simple and fast to spin up instances for tests and demos 	<ul style="list-style-type: none"> • EC2 can be expensive • Some organizations do not want data on the cloud • Usually requires Virtual Private Cloud to be setup or a Remote Data Collector to connect to data sources behind your firewall 	Fast implementations
Custom Server	<ul style="list-style-type: none"> • You have full control over your server and where 	<ul style="list-style-type: none"> • Infrastructure costs and wait time provisioning the required server 	Companies with specific infrastructure and security needs

it lives in your
network topology

- Custom installation will usually involve working closely with Metric Insights to run the sometimes complicated installation procedures

1.2 What are the system requirements for a Metric Insights server?

This article explains the minimum system requirements for the Metric Insights server.

System Requirements for Pilot Instances

- 4 Core
- 16 GB RAM
- 500 GB HD

On AWS, this equates to an m5.xlarge EC2 instance type.

System Requirements for Production Instances

- 8 Core
- 32 GB RAM
- 1 TB+ HD

On AWS, this equates to an m5.2xlarge EC2 or instance type.

A Note about Virtual Machines

Metric Insights can be deployed as a virtual appliance, for pilot instances as well as for production instances. (See [Deploy Metric Insights as a Virtual Appliance](#) for more information.) If you choose this option, then your virtual machine should meet the requirements laid out above.

It's important to understand that the Metric Insights virtual appliance *is* a virtual machine. You do not need to create your own virtual machine to serve as a host. In fact, the host machine should be a *physical* machine, not a *virtual* machine. If your IT department offers to provide a virtual server to serve as a host for Metric Insights, then they must certainly be running one of the common virtualization environments. In that case, simply send the Metric Insights virtual appliance to them and ask them to 'import' it.

1.3 Docker Deployment (Simple Installer)

This article will detail how to setup and install the most basic Docker deployment available in 6.+

PREREQUISITES:

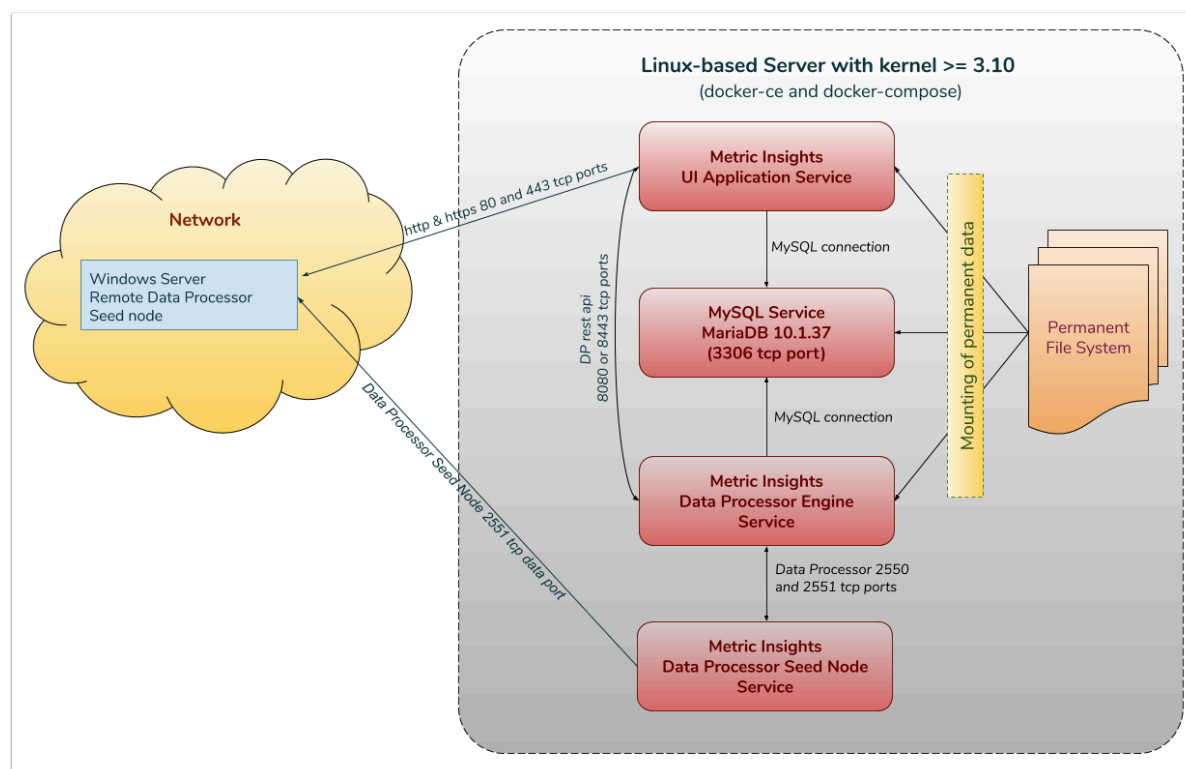
For servers with closed outbound access to the internet, in order to deploy Metric Insights using simple installer, the host machine must have the following packages installed manually:

1. docker-ce
2. docker-compose 1.19.0+

For servers open to internet, these 2 packages will be installed by the MI installation package.

3. Python v2 or v3

Simple installation architecture



💡 If you are using a firewall on the Linux machine, the following ports are necessary for Metric Insights to work correctly:

1. **80, 443** - HTTP and HTTPS ports for the UI Application Service (by default redirection to 443)
2. **2550** - TCP port for the Data Processor cluster
3. **2551** - TCP port for the Seed Node Service
4. **3306** - MySQL port to get access from outside
5. **8080, 8443** - HTTP and HTTPS ports for the REST API Data Processor Service (only one port is enabled at one time)
6. **8081** - TCP port for Monitoring Tool

1. SSH to the server

1. SSH to the host server, as in this example where x.x.x.x represents the IP of the server (or use the cname): `ssh root@x.x.x.x`

Note: The Metric Insights application must be installed as user *root* or as a *user with sudo privileges* (added to the *wheel* or *sudo* group, thereby included in the sudoers list). The docker daemon binds to a unix socket instead of a tcp port, therefore unix sockets are owned by user *root* and accessible by users with *sudo* privileges.

2. Download the MI app installation package

The screenshot shows the Metric Insights website's download page. The URL in the browser is `https://www.metricinsights.com/download/MetricInsights-required version.tar.gz`. The page has a navigation bar with links: Product, Solutions, Learn, Support, About, and a Free Trial button. The main heading is "Download". Below it, a prompt says "Enter your user name and password below". There are two input fields: "Username *" and "Password *". Below these is a section titled "Metric Insights Software Evaluation Agreement" containing a scrollable text area with legal terms. At the bottom, there is a checkbox labeled "I agree with these terms and conditions" and a "Submit" button. Numbered callouts 1 through 4 point to the URL, the login fields, the agreement checkbox, and the submit button respectively.

1. Depending on the version, you will be given a link to the corresponding MI app installation package.
2. Along with the link, **Username** and **Password** will be provided by the MI support team; enter these credentials into the respective fields.
3. Confirm your agreement with the Terms and Conditions of the Metric Insights usage.
4. Click **[Submit]**.

This will initiate the download process of the **.tar.gz** archive (containing all files required to install the system).

3. Package Overview

```

.
├── base
├── build.properties
├── data
├── installer.py
└── utils
    
```

1. **base**
 - subfolder with basic **/opt/mi** directory structure that will be overwritten by next update
2. **build.properties**

- a file with specified components versions

3. data

- a subfolder with binary data and miscellaneous data (such as dependencies list and credentials for pulling docker images from **docker.metricinsights.com** by Installer in case with **Lite** type of package)

4. installer.py

- the main executive file that initializes the Metric Insights installation process

5. utils

- folder containing common libraries used by **installer.py**
- also has third-party libraries


4. Supported Operating Systems

⚠ Metric Insights v6.x can be installed on the following Operating Systems (any OS with Linux kernel ≥ 3.10):

- RHEL 7
- CentOS 7
- Oracle Linux 7
- Ubuntu 16/18
- Debian 7/8



5. Run the installer

-  Please make sure to set the correct **timezone** when running the installer using the `--timezone` and `--dp-mysql-option` parameters (both set timezones for the application and the dataprocessor respectively).

We recommend checking the help to ensure appropriate parameters are defined for the install. To see the help:

```
./installer.py -h
```

```
./installer.py --bind-address 0.0.0.0 --timezone 'US/Pacific' --dp-mysql-option '&serverTimezone=PST8PDT' -vv
```

Here's an example install where the application timezone is being set to Pacific Time. Other timezone examples are:


Eastern Time

- `--timezone 'US/Eastern'`
- `--dp-mysql-option '&serverTimezone=EST5EDT'`

Central Time

- `--timezone 'US/Central'`
- `--dp-mysql-option '&serverTimezone=CST6CDT'`

Once the installation is complete, you will see a *"Metric Insights application installed successfully"* message.

-  The "Waiting while all provision actions are completed" can take a while, so it's not so unusual. While it's running, in a separate terminal, please `tail /opt/mi/log/web.log` to see what it's doing.

5.1. Basic console commands

⚠ Not all tools will be installed with all components, each component has only: **mi-version** and **mi-control**. All of these tools become accessible only if the Web Component is installed.

mi-app-backup	Wrapper for mi-app-backup (/opt/mi/backup is an internal mount point from the web docker container)
mi-app-restore	Wrapper for mi-app-restore
mi-console	Connect to the MI UI docker container directly
mi-control	Docker-compose wrapper helping to manage Metric Insights docker containers
mi-cron	Wrapper for mi-cron manipulating the MI cron tasks inside the docker container
mi-db-dump	Wrapper for mi-db-dump
mi-db-mgr	Wrapper for mi-db-mgr
mi-ldap-usersync	Wrapper for mi-ldap-usersync
mi-maintenance	Wrapper for mi-maintenance
mi-passwd	Wrapper for mi-passwd
mi-version	Show the installed Metric Insights version

5.2. The list of arguments that can be used during installation (from help)

<i>Optional Parameters</i>	
-h, --help	Show this help message and exit
-v, --verbose	Verbose logging (default: False)

-y, --yes	Assume "yes" on interactive requests (default: False)
--path INSTALL_PATH	Base path for Metric Insights (default: /opt/mi)
--dp-hostname DATAPROCESSOR_HOSTNAME	Data Processor service hostname for external services
--skip-packages	Skip installation of any additional packages (default: False)
--show-packages	Show packages required to run Metric Insights (default: False)
--db-ram-config DB_RAM_CONFIG	Amount of RAM (integer value in gigabytes) on the MI Database Server. Options: [12,16,32] (default: None)
--db-hostname DB_HOSTNAME	MySQL hostname of a Database server (default: mysql)
--db-port DB_PORT	MySQL port for db-hostname (default: 3306)
--db-user DB_USER	MySQL admin user name to init a Metric Insights database (default: root)
--db-password DB_PASSWORD	MySQL admin user password
--dp-db-password DP_DB_PASSWORD	MySQL Data Processor user password
--restrict-access RESTRICT_ACCESS	Comma separated list of the MI services that must not have outside access. By default, all services are opened externally. Example: web, mysql, dataprocessor, seed (default: None)
--unprivileged	During installation, skip actions that require root access, like system package installation. CAUTION: system must already have all the required packages installed (default: False)
--seed-hostname DATAPROCESSOR_SEED_HOSTNAME	Data Processor seed service hostname for external services
--dp-url DP_URL	Set Data Processor REST-API URL (default: http://dataprocessor:8080)
--dp-user DP_USER	User to get access to the Data Processor REST-API service (default: dataprocessor)
--dp-password DP_PASSWORD	Password for the Data Processor REST-API user (default: None)
--dp-enable-ssl	Enable HTTPS for the Data Processor REST-API service.

	Ensure that correct SSL certificate and DP URL are provided (default: False)
--components COMPONENTS	Comma separated list of components to deploy. Possible values: web, mysql, dataprocessor, seed (default: web, mysql, dataprocessor, seed)

6. Security

```
./installer.py --restrict-access mysql,dataprocessor
```

This way only UI can be accessible from the outside.

1.4 Deploying Metric Insights in Orchestrated Environment

i Starting from Version 6.x, Metric Insights can be deployed using a container orchestration platform:

- Kubernetes
- Amazon ECS
- Docker Swarm

Metric Insights consists of services running inside of containers:

- Web
- Data Processor
- Seed
- Data Analyzer
- Monitoring

Deployment process:


1. [Prerequisites](#)
2. [Obtain Docker Registry Credentials](#)
3. [Choose a Deployment Method \(Kubernetes, Amazon ECS, Docker Swarm\)](#)
4. [Generate the configuration file to deploy to Amazon ECS](#)
5. [Create the ECS Stack with AWS CloudFormation](#)
6. [Accessing the Metric Insights Deployment](#)
7. [Resources involved in running Metric Insights in ECS](#)

1. Prerequisites


To deploy Metric Insights across the different orchestration platforms, the following architectural pieces are required:

- **kubectl** command-line tool to manage a Kubernetes cluster (not required for ECS and Docker Swarm)
- Remote Database Server to host the application database (e.g., MySQL/MariaDB, Amazon RDS or equivalent)


- Persistent shared storage to store the application file system (e.g., NFS, Portworx, EFS)
- Specific ports open on the network:
 - 80, 443 - HTTP and HTTPS ports for the UI Application Service (by default redirection to 443)
 - 2550 - TCP port for the Data Processor cluster within the kubernetes namespace
 - 2551 - TCP port for the Seed service within the kubernetes namespace
 - 32550 - TCP port for external access to the Data Processor cluster
 - 32551 - TCP port for external access to the Seed service
 - 3306 - MySQL port to get access from outside
 - 8080,8443 - HTTP and HTTPS ports for the REST API Data Processor Service (only one port is enabled at one time)
 - 8081 - TCP port for Monitoring Tool

 For non-orchestrated environments, please see the help article on using [Simple Installer](#)

2. Obtain Docker Registry Credentials

-  Contact MI Support for access to the official Metric Insights Docker Registry. Credentials are needed to pull docker images for each Metric Insights service.
- Note, the default MI Docker Registry address (docker.metricinsights.com:5002) is specified in the deployment configuration file for each orchestration type.

3. Choose Deployment Method (Kubernetes, Amazon ECS, Docker Swarm)

 If deploying to **Kubernetes**, please see [Deploying Metric Insights on Kubernetes](#).

 If deploying to **Amazon ECS**, continue below.

Amazon ECS Prerequisites:

1. Database (RDS or EC2 instance with custom database deployment)
2. EFS or custom NFS shared storage

- Optional: If utilizing a private registry (non-Metric Insights), ensure that you have those credentials available.

4. Generate Configuration File to Deploy to Amazon ECS

The configuration file can be generated using the Metric Insights installer package:

- Download the installer package to a linux system and unpack
- Change into the installer directory then run the installer with the **ecs** command and specify a target filename to generate the configuration file:

```
$ ./installer.py ecs -o <filename>.json
```

- The configuration file can now be used as a template with AWS CloudFormation to create and deploy the Metric Insights environment

5. Create the ECS Stack with AWS CloudFormation

Prepare the following:

- RDS address with root credentials
- EFS address to connect to Metric Insights application

Apply the configuration file through the CloudFormation UI:

- Upload the generated json file as a template

CloudFormation > Stacks > Create stack

Step 1
Specify template

Step 2
Specify stack details

Step 3
Configure stack options

Step 4
Review

Create stack

Prerequisite - Prepare template

Prepare template
Every stack is based on a template. A template is a JSON or YAML file that contains configuration information about the AWS resources you want to include in the stack.

☒ Template is ready ☐ Use a sample template ☐ Create template in Designer

Specify template
A template is a JSON or YAML file that describes your stack's resources and properties.

Template source
Selecting a template generates an Amazon S3 URL, where it will be stored.

☐ Amazon S3 URL ☒ Upload a template file

Upload a template file

Choose file JSON or YAML formatted file

S3 URL: <https://s3.us-east-2.amazonaws.com/cf-templates-20k6o5ky29hn-us-east-2/2019296d75-metricinsights-6.1.0.json>

2. Fill out each field then click [Next] at the bottom of the page. Some key notes:

- To generate passwords for each service, you can either run `echo -n '<pwd>' | base64` to encode a password of your choice, or run something like `openssl rand -base64 8` to auto generate a password for you.
- Use the full RDS address for the field `DBHostName`
- Enter the RDS root user in the field `DBRootUserName`
- Enter the full EFS address in the field `NFSServerAddress`
- Select all Subnet IDs available in the field `SubnetIDs`
- The field `WebReplicationsCount` represents the number of web slave containers (secondary to web master)

CloudFormation > Stacks > Create stack

Step 1
Specify template

Step 2
Specify stack details

Step 3
Configure stack options

Step 4
Review

Specify stack details

Stack name

Stack name

miEC2

Stack name can include letters (A-Z and a-z), numbers (0-9), and dashes (-).

Parameters

Parameters are defined in your template and allow you to input custom values when you create or update a stack.

AllowAccessForSubnet*
Specify the network, the insights application has to listen.

0.0.0.0/0

DBDataProcessorPassword

DBDatasetReadPassword

DBGeneratorPassword

DBHostName
Specify remote database server to connect Metric Insights Application.

dashboard.csblz7yeiqd.us-east-2.rds.amazonaws.com

3. Click [Next] to skip through the subsequent pages until you reach the following window. Click the checkbox to acknowledge that IAM resources might be created on deployment and click the button [Update stack].

Capabilities

The following resource(s) require capabilities: [AWS::IAM::Role]

This template contains Identity and Access Management (IAM) resources that might provide entities access to make changes to your AWS account. Check that you want to create each of these resources and that they have the minimum required permissions. [Learn more](#)

☒ I acknowledge that AWS CloudFormation might create IAM resources.

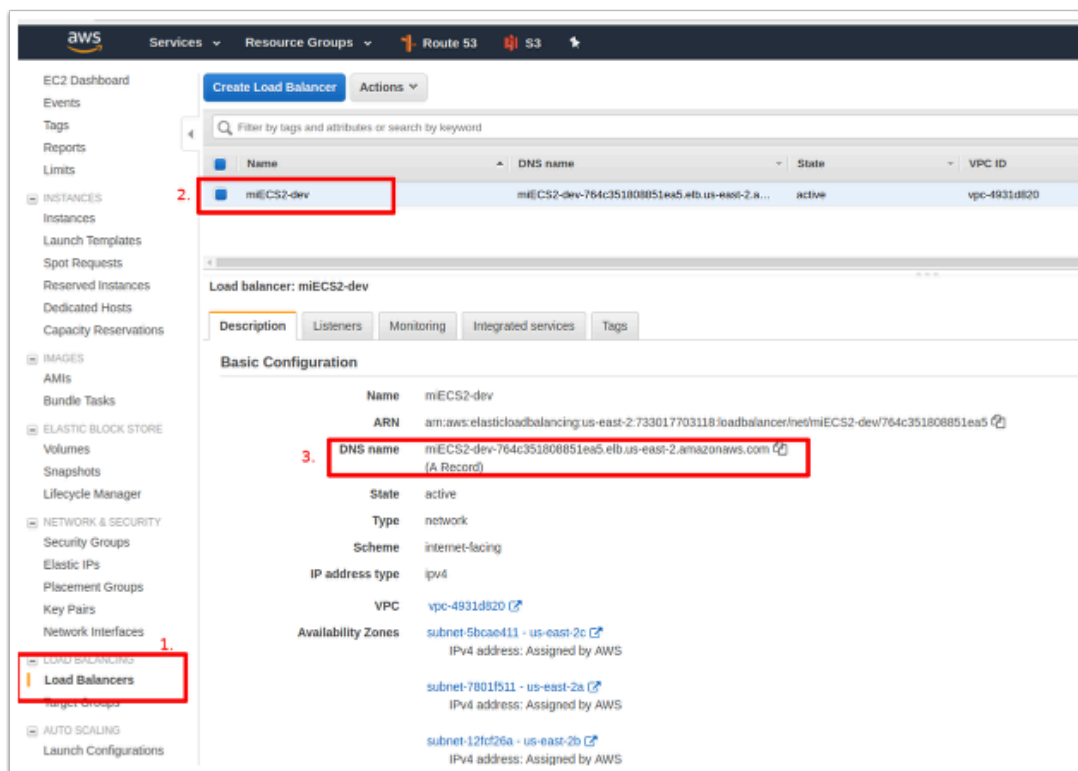
Cancel Previous View change set **Update stack**

4. Allow incoming connections to RDS for new ECS/EC2 security group to complete the deployment

1. As the new ECS Stack is being deployed, go to the EC2 Console and select one of the new EC2s created for ECS
2. Go to the Security Group field and select on the new security group name
3. Copy the Group ID (e.g., `sg-name`)
4. Switch to the RDS Console and select the RDS instance being used for ECS
5. Go to the VPC Security Group field and select the security group name
6. Switch to the `Inbound` tab and click the [Edit] button
7. Add the new EC2 security group to the list then [Save]:
 - Add Rule > All Traffic > Paste Group ID

On adding the group, switch back to CloudFormation to monitor the ECS Stack deployment. The deploy should complete in 5-10 minutes.

6. Accessing Metric Insights Deployment



Once the ECS Stack is deployed:

1. Switch back to the EC2 Console and select Load Balancers in the left menu pane.
2. Identify the Load Balancer DNS name to access the Metric Insights application in a browser.
3. For the best user experience, map the Load Balancer DNS name to a user-friendly name in Amazon Route 53.

Metric Insights is now deployed in ECS and browser ready.

7. Resources Involved in Running Metric Insights in ECS

- AWS ECS Task Definitions
- AWS ECS Cluster
- AWS ECS Services
- AWS EC2 Auto Scaling group
- AWS EC2 Launch Configuration
- AWS EC2 Security Groups
- AWS Target Groups
- AWS Network Load Balancer
- IAM Roles
- AWS Secret Manager
- AWS Cloud Formation (only for deployment and updates)

Non-ECS resources in AWS needed for deployment include:

1. AWS RDS instance based on MariaDB 10.1 (custom parameter group with `log_bin_function_creators` enabled)
2. AWS EFS Shared Storage

1.5 Deploy Metric Insights as a Virtual Appliance

This article shows how to deploy Metric Insights as a virtual appliance. If you have not already done so, you will need to request the Metric Insights software. Please request a copy of the software at http://www.metricinsights.com/request_eval.html

1. Install VMware Player or Another Virtualization Environment

You will need a suitable virtualization environment in which to 'play' the virtual appliance. If you're just trying out Metric Insights and would like to use a desktop machine, you can use VMware Player (Windows) or VMware Fusion (Mac). (We don't recommend desktop machines -- even for pilot projects -- unless they are dedicated machines that will always be running and accessible.) For production installations, you will probably want an enterprise virtualization environment such as VMware vSphere or Microsoft Hyper-V.

It's important to understand that the Metric Insights virtual appliance *is* a virtual machine. You do not need to create your own virtual machine to serve as a host. In fact, the host machine should be a *physical* machine, not a *virtual* machine. If your IT department offers to provide a virtual server to serve as a host for Metric Insights, then they must certainly be running one of the common virtualization environments. In that case, simply send the Metric Insights virtual appliance to them and ask them to 'import' it.

2. Download the Virtual Appliance

Download Form - Metric Ins

www.metricinsights.com/download/eval/MetricInsights-1.0.325_CentOS-6-x86-64.zip

recipes Gmail - Work MI Asana this isn't happiness™ Kinokuniya BookWeb Reading Other Bookmarks

METRIC INSIGHTS HOME | PRODUCT | SOLUTIONS | PARTNERS | VIDEO | BLOG

Email or call 1-800-480-8221

Enter your user name and password below

Login*

Password*

Metric Insights Software Evaluation Agreement

This Metric Insights Customer Agreement (this "**Agreement**") contains the terms and conditions that govern your access to and use of the Service Offerings (as defined below) and is an agreement between Metric Insights, Inc. ("Metric Insights"), a California company, with an address at 123 10th Street, San Francisco, CA 94103, and you or the entity you represent ("**User**"). This Agreement takes effect when an Order is signed, when you click an "I Accept" button or check box presented with these terms or, if earlier, when you use any of the Service Offerings (the "**Effective Date**"). You represent to us that you are lawfully able to enter into contracts (e.g., you are not a minor). If you are entering into this Agreement for an entity, such as the company you work for, you represent to us that you have legal authority to bind that entity.

1. LICENSE.

1.1 Evaluation license. Upon completing this Agreement, Metric Insights grants User and its Affiliates a worldwide, non-exclusive and non-transferable license to use the Software for **30 days** in executable form for the purpose of considering purchasing a commercial license of the software or the purpose of evaluating a distribution relationship with Metric Insights.

☐ I agree with terms and conditions

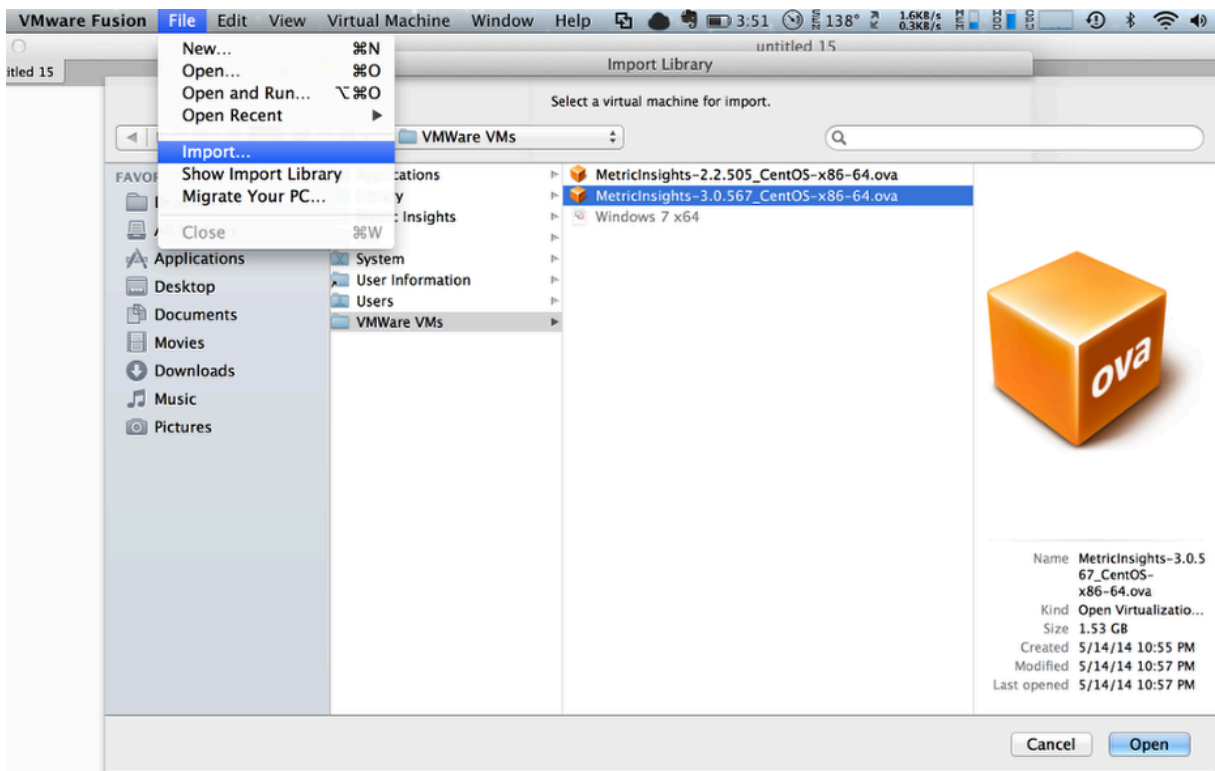
Download license agreement

SUBMIT

After receiving the link to download the virtual appliance, fill out the license agreement and download the ova file. The ova file is around 2GB in size and will require a fast network connection for the download.

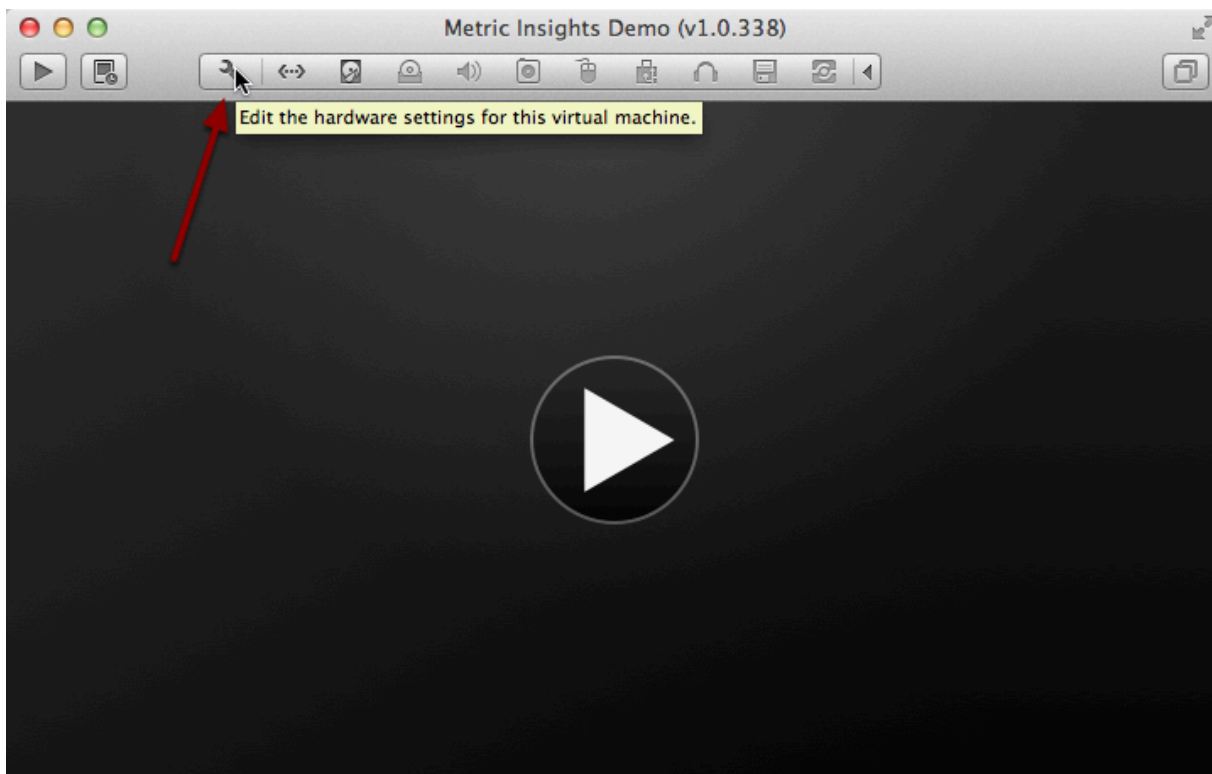
Note: The remaining steps assume that you are planning to deploy Metric Insights on a desktop machine. If you are deploying Metric Insights in an enterprise virtualization environment and encounter any problems, please contact us for assistance at support@metricinsights.com

3. Import the 'ova' File



If you're running VMware Player or VMware Fusion on a desktop machine, 'import' the ova file using the File menu.

4. Set up the Virtual Machine Hardware Settings



To make sure the virtual machine meets the minimum system requirements, edit the virtual hardware settings.

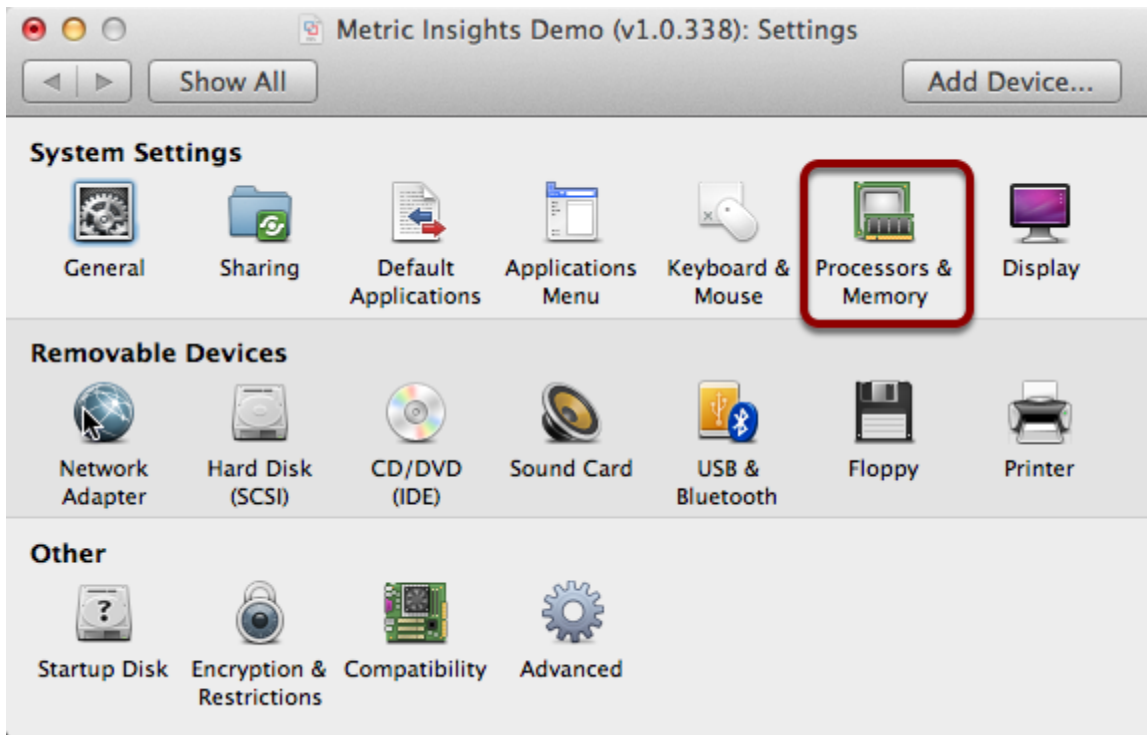
4.1. Set up the Virtual Machine Network Adapter



There are two methods of VMware networking that affect how the Metric Insights server can be accessed by you and others on your network. You will need to choose the method that will work for your network.

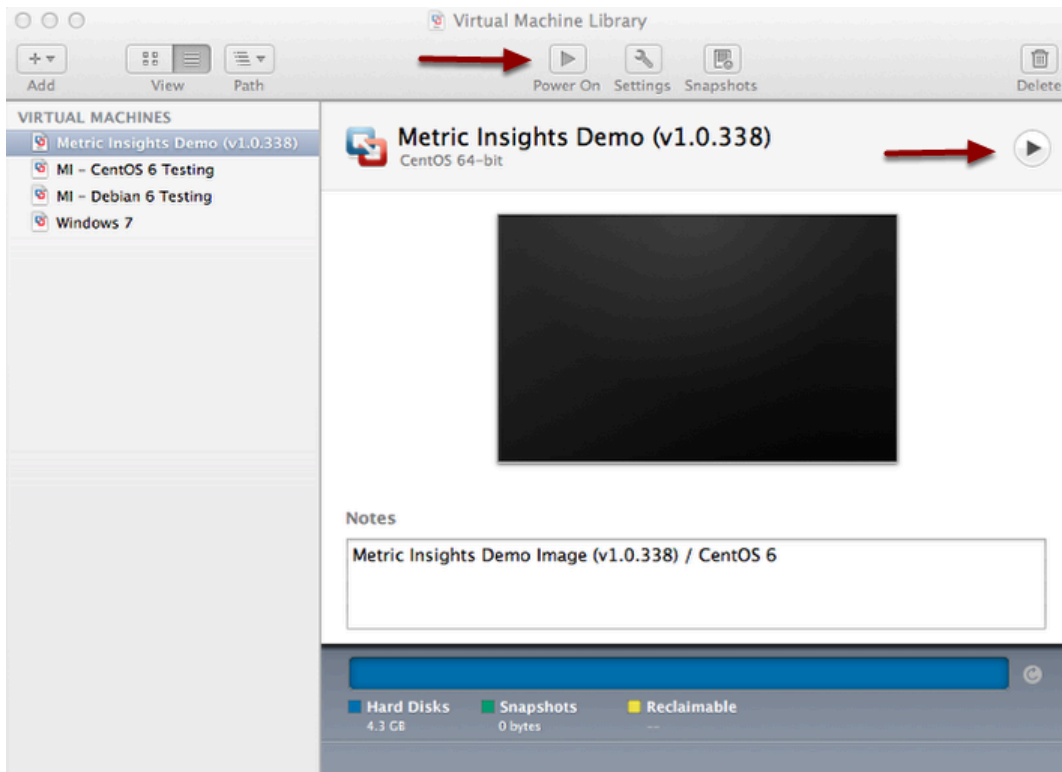
- **Bridged or AutoDetect** - This is the default and recommended setting as it allows the virtual machine to automatically get an IP address from your local network.
- **Network Address Translation (NAT)** - If you do not get an IP address when running under bridged mode, you will need to use NAT networking. Using this method, the virtual machine will share the host operating system's IP address, and the Metric Insights server will only be accessible from the host operating system. If you'd like to be able to access Metric Insights from other machines as well, see [How to expose Metric Insights on VMware when using NAT networking](#).

4.2. Configure the Virtual Machine Resources



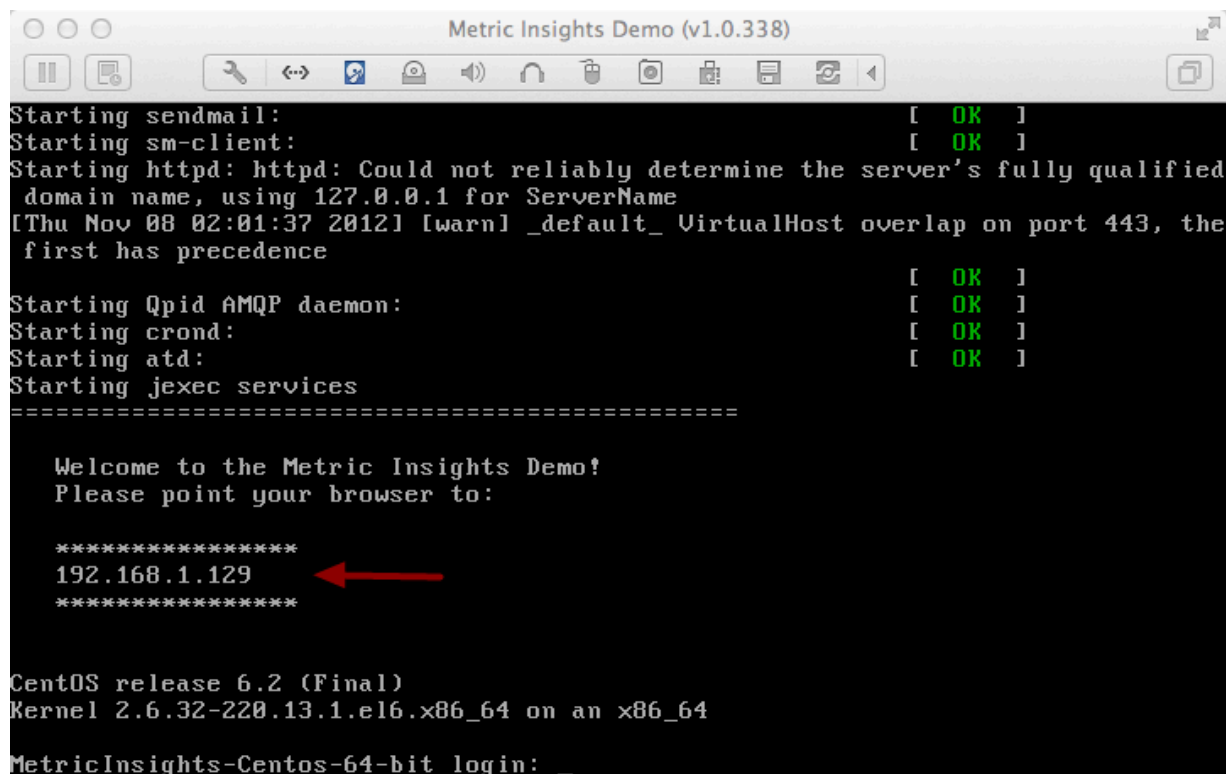
- **Processors** - If you only have two cores on the host machine, keep the number of virtual cores at 1. Otherwise, it's safe to allocate half of the available cores.
- **Memory** - This should be at least 2GB = 2048MB

5. Start the Virtual Machine



After starting up the virtual machine, you'll see a bunch of linux diagnostic messages scroll by the screen. This will take about a minute the first time, and then it should only take 20 seconds or so depending on your operating system speed.

5.1. Get the IP address of your virtual machine



```
Metric Insights Demo (v1.0.338)
Starting sendmail: [ OK ]
Starting sm-client: [ OK ]
Starting httpd: httpd: Could not reliably determine the server's fully qualified
domain name, using 127.0.0.1 for ServerName
[Thu Nov 08 02:01:37 2012] [warn] _default_ VirtualHost overlap on port 443, the
first has precedence
Starting Qpid AMQP daemon: [ OK ]
Starting crond: [ OK ]
Starting atd: [ OK ]
Starting jexec services
=====
Welcome to the Metric Insights Demo!
Please point your browser to:

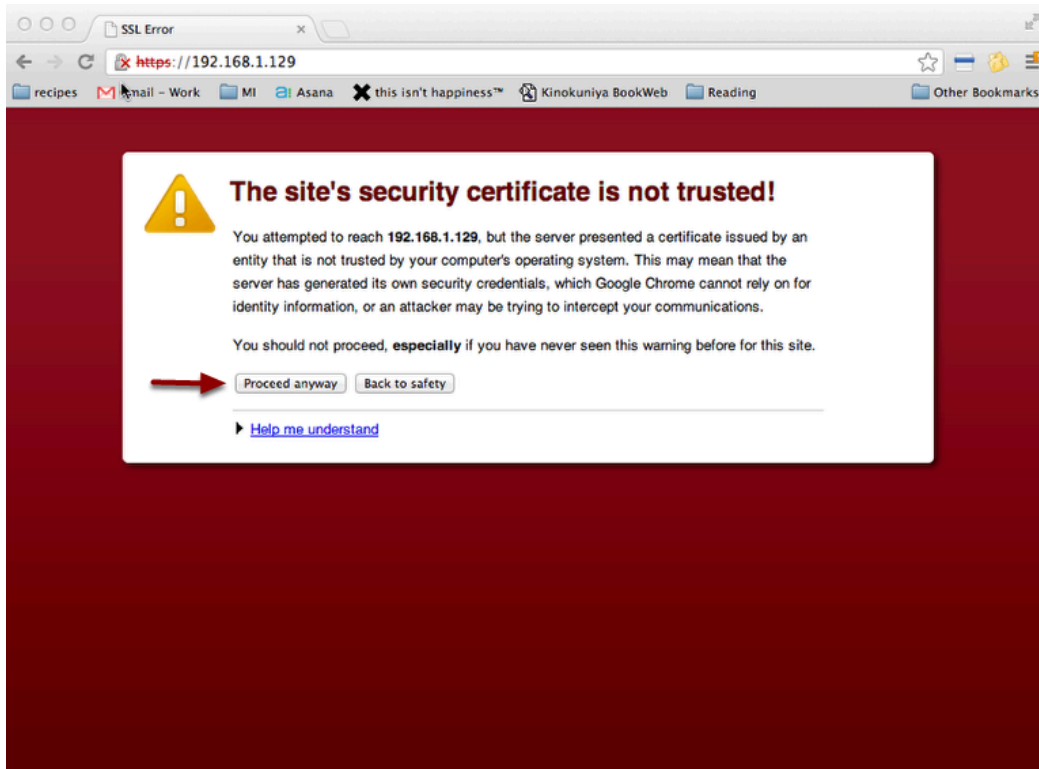
*****
192.168.1.129
*****

CentOS release 6.2 (Final)
Kernel 2.6.32-220.13.1.el6.x86_64 on an x86_64
MetricInsights-Centos-64-bit login: _
```

Once the virtual machine has started up, you should see a message that looks similar to this.

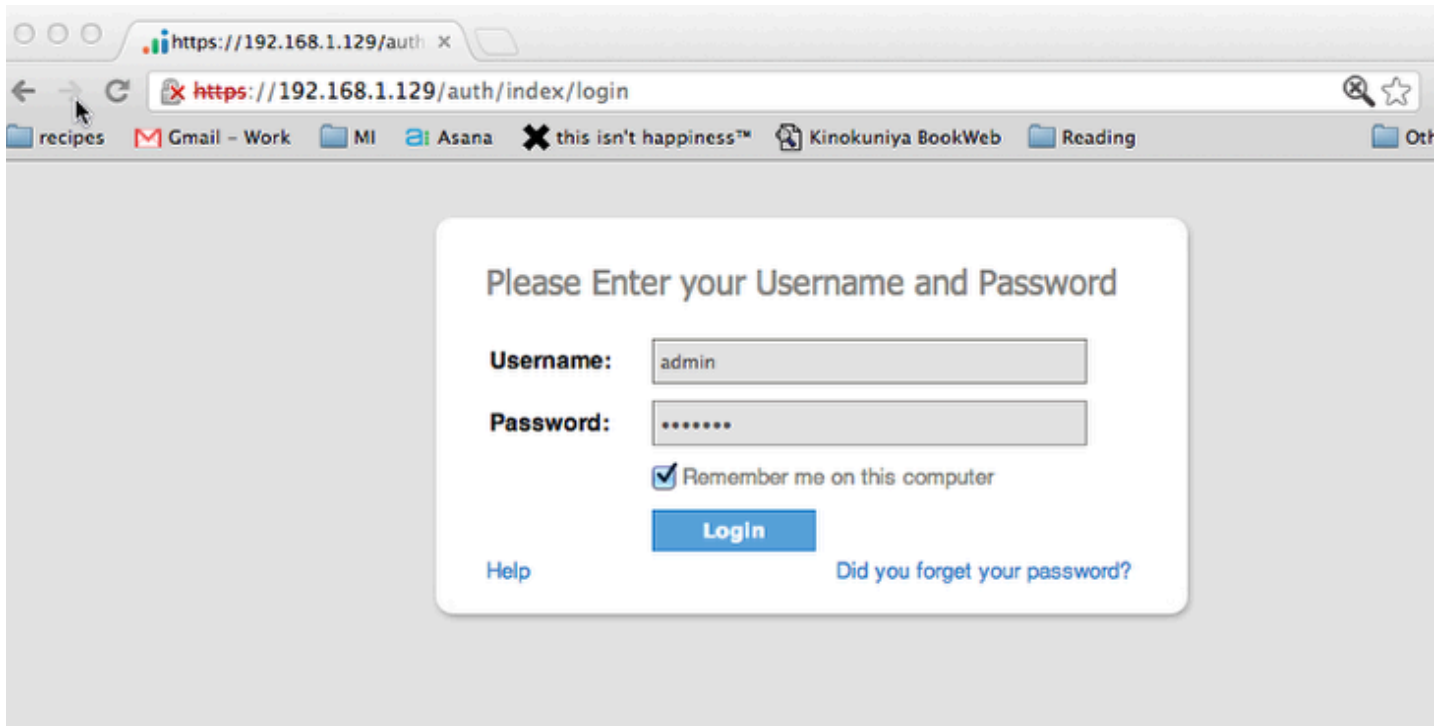
NOTE: If you do not get an IP address here, the Metric Insights server was not able to get an ip address from your DHCP server. If this happens you will need to set the Network Adapter to use NAT addressing or talk with your local Network Administrator.

5.2. Point a browser to the Metric Insights IP and accept the SSL warning



Open a web browser and put the ip address obtained above into the address bar. Depending on your browser you will likely get an SSL warning. This is because the Metric Insights server ships with a default self-signed certificate. You can just accept the warning and you will get to the login prompt

5.3. Log into Metric Insights



The screenshot shows a web browser window with the address bar displaying `https://192.168.1.129/auth/index/login`. The browser's tab bar shows several open tabs: "recipes", "Gmail - Work", "MI", "Asana", "this isn't happiness™", "Kinokuniya BookWeb", "Reading", and "Otr". The main content area displays a login form titled "Please Enter your Username and Password". The form contains the following fields and elements:

- Username:** A text input field containing the value "admin".
- Password:** A password input field containing seven dots (password masked).
- ☒ Remember me on this computer
- Login** button (blue)
- [Help](#) (blue text)
- [Did you forget your password?](#) (blue text)

Log in with the username and password given to you by our Support Group (support@metricinsights.com)

1.6 What are the main components of the Metric Insights stack?

The main components of the Metric Insights stack are listed below. Versions listed are the minimum required.

** For system and base linux package requirements, please see the following articles:*

- [System Requirements](#)
- [Package Requirements for v6.1+](#)

Versions 6.1+ till 6.1.4 (Application runs inside Docker containers)

- [Linux]
 - CentOS 7/8
 - RHEL 7/8
 - Debian 9/10
 - Ubuntu 16/18
- Docker-CE 17.09.0+
- Docker-Compose 1.19.0+
- [Inside our Docker Containers]
 - Apache 2.4.25
 - PHP 7
 - Python 2.7 & 3
 - Java 11
- MySQL 5.6, 5.7 / MariaDB 10.1
 - With @@sql_mode = ' ' for MySQL 5.7

NOTE: for v6.2.0+ ONLY MySQL v8 is acceptable with @@sql_mode = ' ' OR sql_mode = NO_ENGINE_SUBSTITUTION

Versions 5.0+ (last release in 2019, v5.6.3)

- [Linux]
 - CentOS 7
 - RHEL 7
 - Debian 7 (kernel v3.10+) / Debian 8
 - Ubuntu 14

- Apache 2.2
- MySQL 5.5, 5.6, 5.7 / MariaDB 10.1
 - With @@sql_mode = ' ' for MySQL 5.7
- PHP 5.6
- Python 2.7
- Java 8

NOTE: Support for version v5.6.3 will end in 2020.

Versions pre 5.0

- [Linux] CentOS 6, Debian 6/7
- Apache 2.2
- MySQL Server 5.1+
- PHP 5.4
- Python 2.7
- Java 7/8

Notes on SSL/SSH

Encryption to and from the MI virtual appliance (if using an MI provided OVA image) makes use of OpenSSL version 1.0.2k (this version does not suffer from the Heartbleed vulnerability). SSH Communication with the server uses OpenSSH 7.4p1.

1.7 Operating system linux package requirements / dependencies for v5.0+

Metric Insights can be installed on the following Linux operating systems using a package installer provided by MI Support:

- [RHEL](#)
- [CentOS](#)
- [Debian](#)

If installing on an internally provisioned server, **please note the following list of required linux packages needed to support the Metric Insights application.** If the server does not have outbound access and must use internal linux repositories instead, then please be aware that linux packages have interdependencies between them. All interdependent packages must be on the same version.

Note:

* *This list is current for Metric Insights version 5.3.0+*

* *For Debian deployments, you must have a kernel version 3.10+ for docker support*

* *We recommend provisioning RHEL7 / CentOS7 or Debian8*

Linux packages required for RHEL and CentOS

RedHat / CentOS	RHEL:	RHEL7 specific:	Sourced from these RPMs:
	<ul style="list-style-type: none">• postfix• httpd• mod_ssl• mod_wsgi• man• ntp• git• curl• wget• unzip• vim• less• screen	<ul style="list-style-type: none">• mariadb• mariadb-server• mariadb-devel• docker• docker-compose• baekmuk-ttf-batang-fonts• baekmuk-ttf-dotum-fonts• libXScrnSaver• at-spi2-atk• gtk3	<ul style="list-style-type: none">• RedHat Enterprise Linux 7 Server• RedHat Enterprise Linux 7 Server - Optional• Extra Packages for Enterprise Linux 7 (EPEL)

- php
- php-ldap
- php-mysql
- php-curl
- php-cli
- php-pear
- php-gd
- php-xml
- php-mbstring
- php-bcmath
- php-mcrypt
- php-xcache
- php-soap
- php-posix
- php-pecl-imagick
- php-pecl-ssh2
- gcc
- gcc-c++
- python-pip
- python-devel
- python-boto
- junit4
- java-1.7.0-openjdk
- java-1.7.0-openjdk-devel
- libicu
- dejavu-sans-fonts
- dejavu-sans-mono-fonts
- net-tools
- libffi-devel
- openssl-devel
- fontconfig
- openldap-devel
- cronie
- cjkuni-ukai-fonts
- cjkuni-uming-fonts
- ipa-mincho-fonts
- ipa-gothic-fonts
- lksctp-tools
- python-virtualenv
- python-devel
- openssh
- e2fsprogs
- xfsprogs
- openldap-clients

Linux packages required for Debian

Debian	Both Debian 7/8:	Debian7 specific:	Debian8 specific:	Kernel required:
	<ul style="list-style-type: none"> • postfix • apache2 • libmysqlclient-dev • libapache2-mod-wsgi • libapache2-mod-proxy-html • man-db • ntp • git • unzip • curl • wget • vim • less • screen • sudo • php5 • php5-ldap • php5-mysql • php5-curl • php5-cli • php-pear • php5-mcrypt • php5-xcache • php5-gd • php5-xmlrpc • gcc • python-pip • python-dev • python-boto • junit4 • openjdk-7-jdk • tomcat7 • tomcat7-admin • tomcat7-common • libtomcat7-java • libffi-dev • libssl-dev • fontconfig 	<ul style="list-style-type: none"> • mysql-client • mysql-server • libicu48 • libssh2-php • ttf-dejavu 	<ul style="list-style-type: none"> • mariadb-client • mariadb-server • libicu52 • php5-ssh2 • fonts-dejavu • docker • docker-compose 	<ul style="list-style-type: none"> • kernel 3.10+

	<ul style="list-style-type: none">• libldap2-dev• libsasl2-dev• rsyslog• fonts-arphic-ukai• fonts-arphic-uming• fonts-ipafont-mincho• fonts-ipafont-gothic• fonts-unfonts-core• libyajl2• apt-transport-https• ca-certificates• e2fsprogs• xfsprogs• docker-ce• ldap-utils			
--	------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	--	--	--

1.8 Expose Metric Insights on VMware when using NAT networking

When using Network Address Translation (NAT) networking on VMware, you will need to configure port forwarding to access the Metric Insights virtual server in order to reach the Metric Insights server from other machines. Each VMware product unfortunately has a different method for configuring NAT forwarding, but the requirements are the same: ports 80, 443, and 8443 need to be port forwarded from your host operating system to the Metric Insights server virtual machine.

1. Verify that no services are running on ports 80, 443, or 8443 on the host OS

In order for port forwarding to work correctly, verify that the host OS (which ever machine you're running VMware Player on) does not have any other services that are using ports 80, 443, or 8443. If you do already have a webserver or something running on ports 80 or 443, you will need to turn them off or choose other ports to map to.

Note, however, if you do something like

```
1080 = 192.168.220.129:80
10443 = 192.168.220.129:443
18443 = 192.168.220.129:8443
```

You will need to point your browser to <https://<your-ip>:10443> instead of <http://<your-ip>:1080> (This is because Metric Insights has built in port forwarding and will try to forward HTTP on port 80 to HTTPS over port 443 automatically.)

2. Setup VMware Port Forwarding

NAT port forwarding is different for each VMware product. Please follow the links below for more information:

Windows

- **VMware Workstation** - In product, use Edit -> Virtual Network Settings -> NAT menu. http://www.vmware.com/support/ws55/doc/ws_net_nat_advanced.html
- **VMware Player** - In VMware directory in cmd prompt, run rundll32.exe vmnetui.dll VMNetUI_ShowStandalone. <http://www.dbuggr.com/smallwei/quick-start-vmware-player-4-5-workstation-advance-networking-gui-tool/>

Mac

- **VMware Fusion** - Edit the incomingtcp section of /Library/Preferences/VMware Fusion/vmnet8/nat.conf and restart nat networking with sudo "/Applications/VMware Fusion.app/Contents/Library/vmnet-cli" --start (older versions of VMware Fusion reference /Library/Application Support/VMware Fusion/vmnet8/nat.conf and sudo "/Library/Application Support/VMware Fusion/boot.sh" --restart as mentioned at <http://www.sharepointjohn.com/vmware-fusion-nat-port-forwarding/>)
- The [incomingtcp] section of the vmnet8/nat.conf should look like:

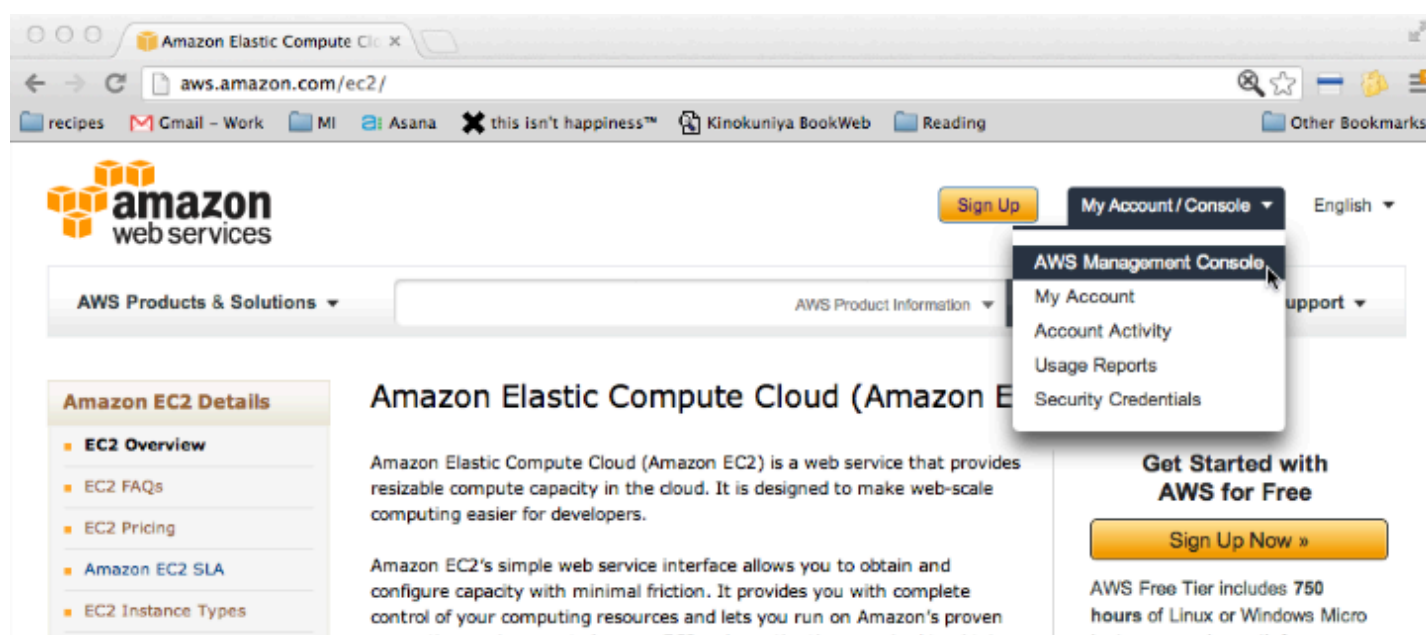
```
80 = 192.168.220.129:80
443 = 192.168.220.129:443
8443 = 192.168.220.129:8443
```


1.9 Deploy Metric Insights on Amazon EC2

This article will detail how to setup and install Metric Insights on Amazon EC2. If you do not already have an Amazon EC2 account, sign up for one at <http://aws.amazon.com/ec2/>. In addition, if you haven't already done so, please request the Metric Insights software at http://www.metricinsights.com/request_eval.html.

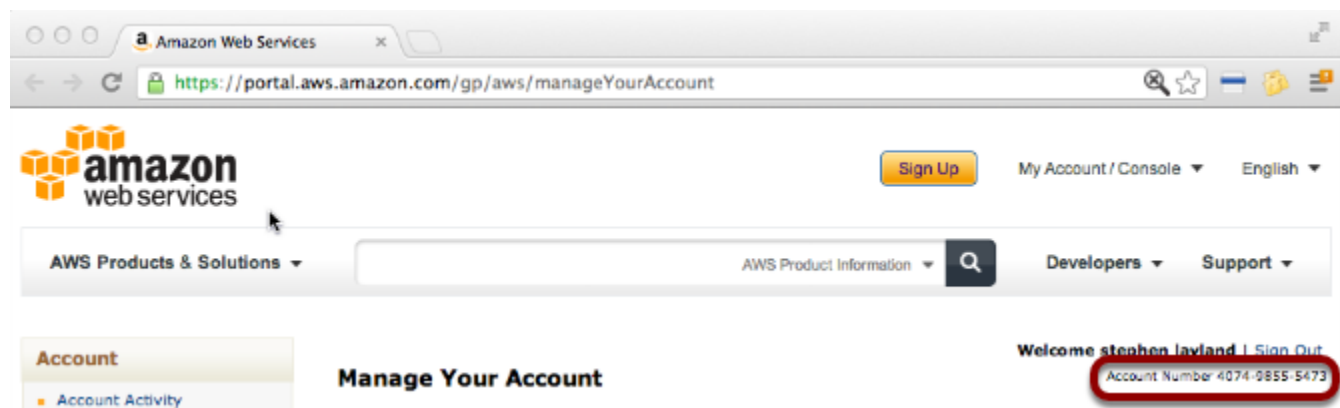
1. Get your AWS Account ID

1.1. Log into your AWS account



Log into your AWS account at <http://aws.amazon.com/ec2>

1.2. Note your AWS account number



You can usually find this under My Account under the AWS console. It is a 12 digit number in the form XXXX-XXXX-XXXX

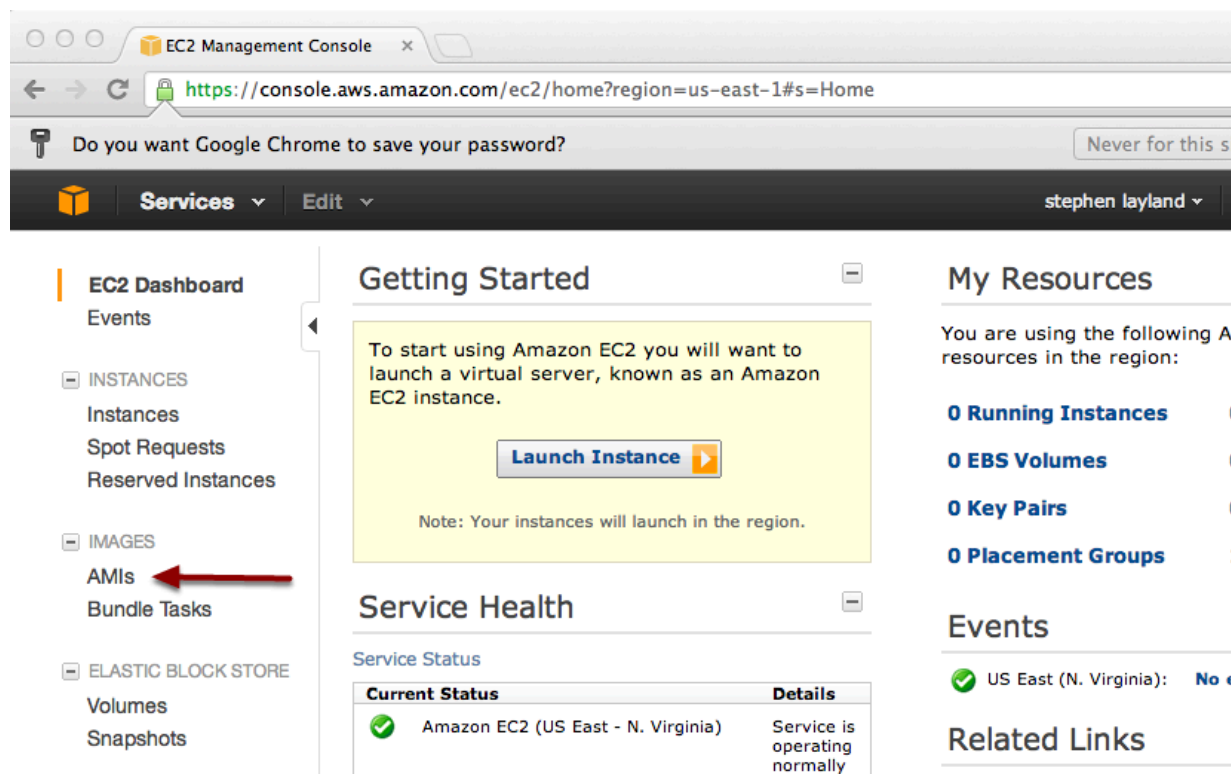
2. Get the Metric Insights AMI

After requesting the Metric Insights software from http://www.metricinsights.com/request_eval.html you will receive an email asking for your AWS Account ID and desired EC2 region. Once you share that information via email, we will respond by sharing an AMI-ID (Amazon Machine Image Id) that you can use to launch your EC2 image.

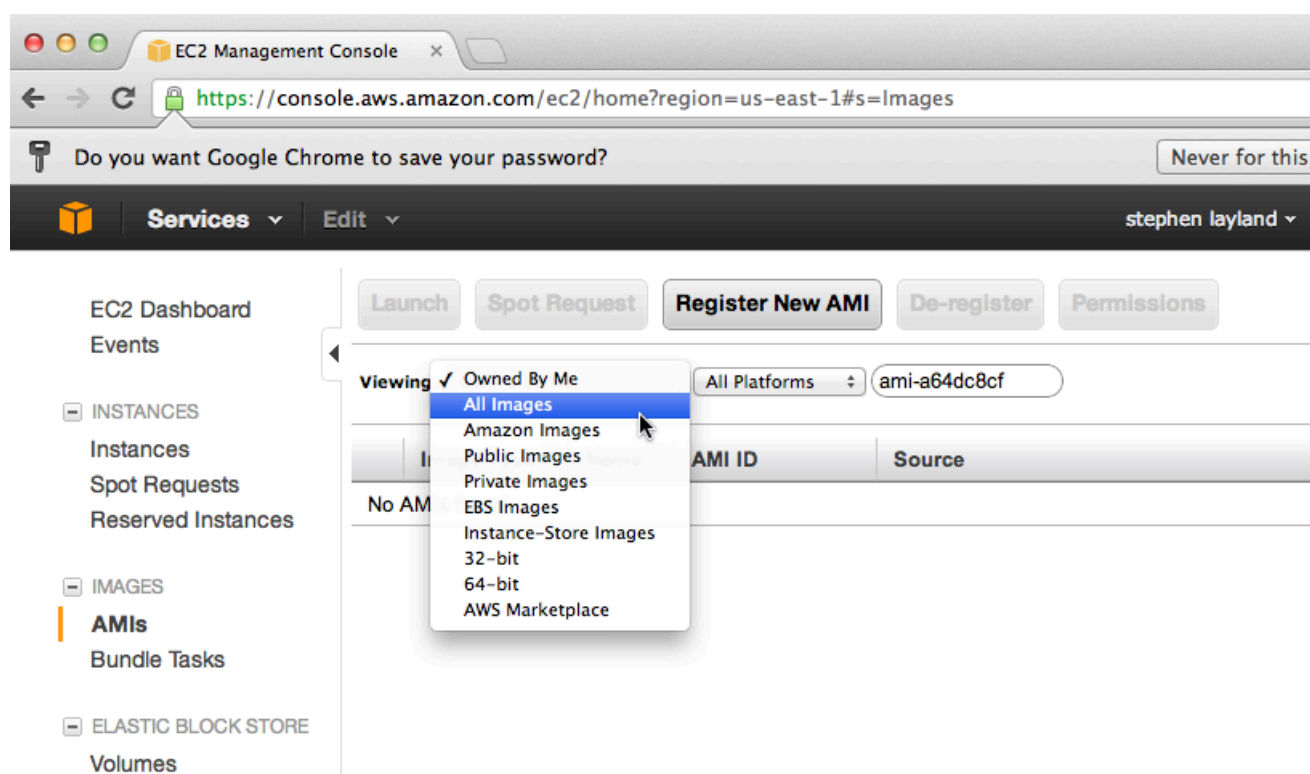
3. Launch a Metric Insights EC2 Instance

At this point, you should have received an AMI-ID (Amazon Machine Image ID) for the Metric Insights server. Use this AMI-ID to launch an EC2 instance

3.1. From the EC2 AWS Console, select the AMIs menu item

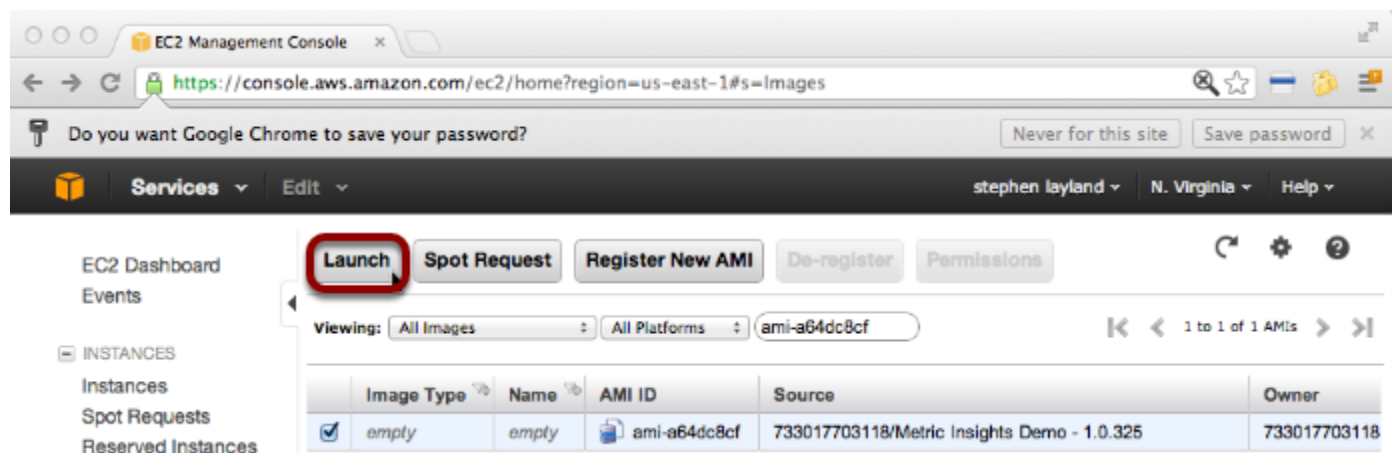


3.2. Find the Metric Insights AMI

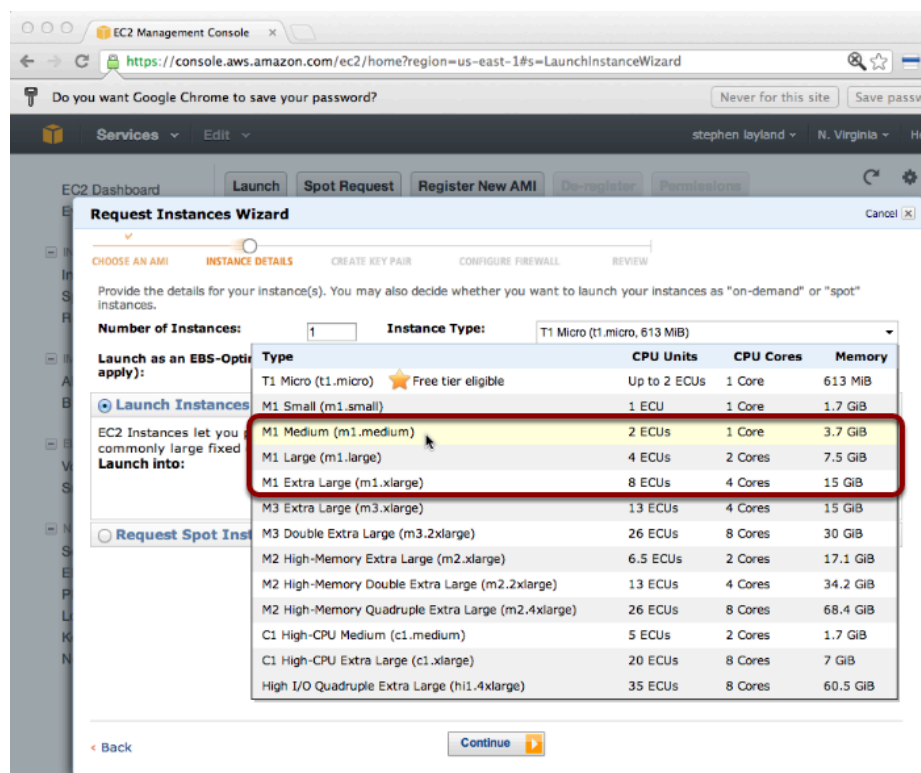


In the AMI list, insert the AMI-ID in the search bar and select 'All Images'. The search will churn for a while (maybe 30 seconds)

3.3. Select the Metric Insights AMI and Launch it



3.4. Launch an EC2 instance with enough resources



See the [minimum system requirements](#) for the Metric Insights Server and select an appropriate image size. M1. Mediums are the minimum recommended size.

3.5. Open the necessary ports for Metric Insights to work

Request Instances Wizard Cancel

CHOOSE AN AMI INSTANCE DETAILS CREATE KEY PAIR **CONFIGURE FIREWALL** REVIEW

Security groups determine whether a network port is open or blocked on your instances. You may use an existing security group, or we can help you create a new security group to allow access to your instances using the suggested ports below. Add additional ports now or update your security group anytime using the Security Groups page.

☐ Choose one or more of your existing Security Groups

☒ **Create a new Security Group**

Group Name

Group Description

Inbound Rules

Create a new rule: Custom TCP rule

Port range:
(e.g., 80 or 49152-65535)

Source:
(e.g., 192.168.2.0/24, sg-47ad482e, or 1234567890/default)

+ Add Rule

TCP	Port (Service)	Source	Action
	80 (HTTP)	0.0.0.0/0	Delete
	443 (HTTPS)	0.0.0.0/0	Delete

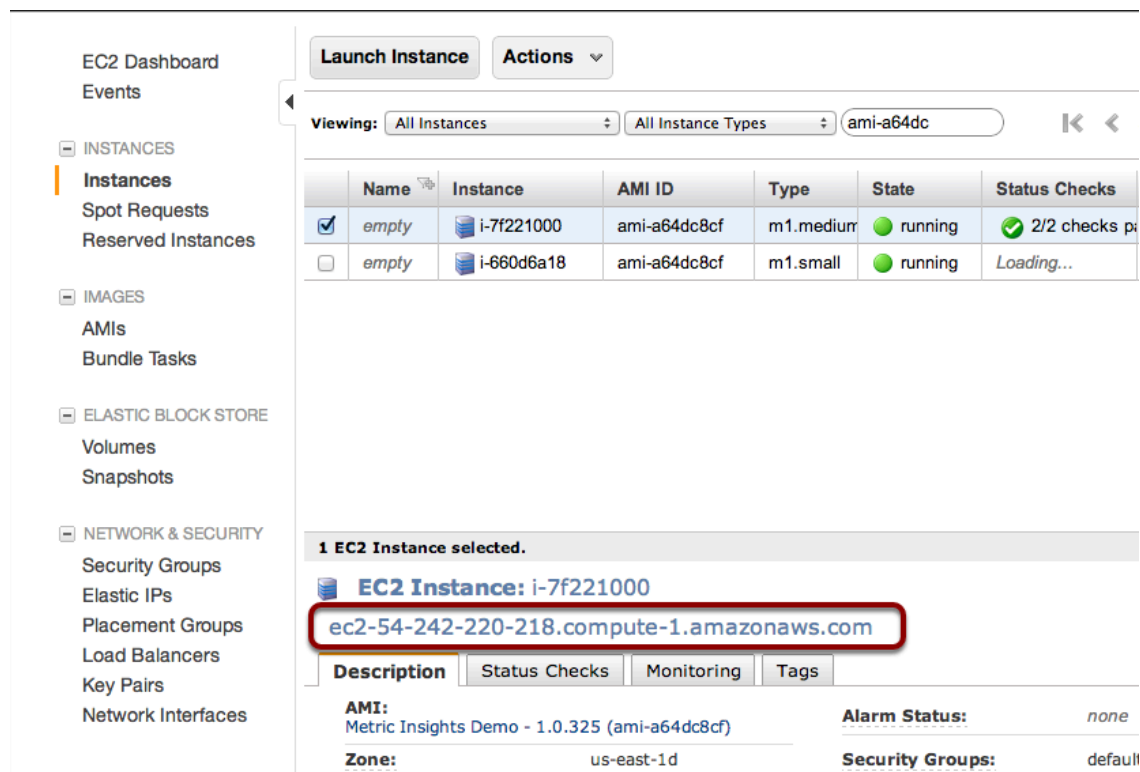
[< Back](#) [Continue >](#)

Add inbound access the following ports to an EC2 Security Group (EC2's firewall)

- 80
- 443
- 8443

If you want to lock down the subnet from which inbound connections are allowed, modify the Source ip range as well. More information can be found in the EC2 documentation.

3.6. Launch the image!



EC2 Dashboard
Events

Launch Instance Actions

Viewing: All Instances All Instance Types ami-a64dc

	Name	Instance	AMI ID	Type	State	Status Checks
<input checked="" type="checkbox"/>	empty	i-7f221000	ami-a64dc8cf	m1.medium	running	2/2 checks passed
<input type="checkbox"/>	empty	i-660d6a18	ami-a64dc8cf	m1.small	running	Loading...

1 EC2 Instance selected.

EC2 Instance: i-7f221000


ec2-54-242-220-218.compute-1.amazonaws.com

Description	Status Checks	Monitoring	Tags
AMI: Metric Insights Demo - 1.0.325 (ami-a64dc8cf)			
Zone: us-east-1d			
Alarm Status: none			
Security Groups: default			

After launching the image, go back to the Instances menu and get the new image's URL. It will appear in a display similar to the example above

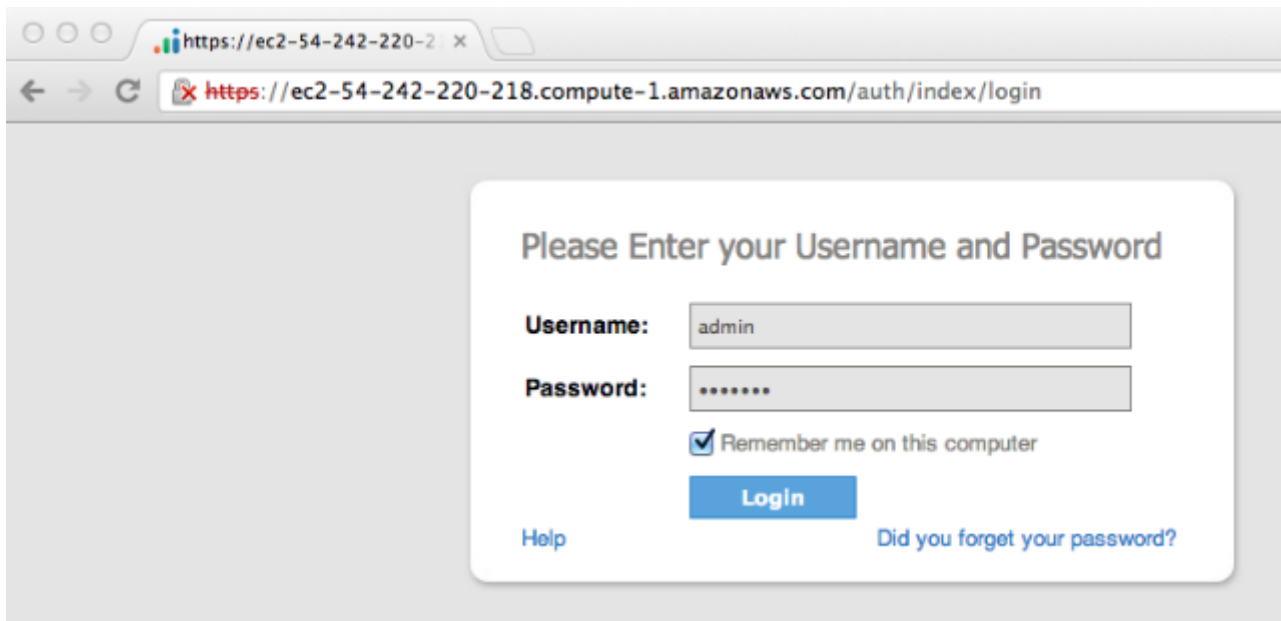
4. Point your web browser to your new EC2 URL

4.1. Accept the SSL warning



Depending on your browser, you might get an SSL warning because the instance is using a self-signed certificate

4.2. Log in to Metric Insights!



The screenshot shows a web browser window with the address bar displaying `https://ec2-54-242-220-218.compute-1.amazonaws.com/auth/index/login`. The page content is a login form titled "Please Enter your Username and Password". The form includes a "Username:" label with a text input field containing "admin", a "Password:" label with a password input field containing seven dots, a checked checkbox labeled "Remember me on this computer", a blue "Login" button, a "Help" link, and a "Did you forget your password?" link.

Log in with the username and password supplied to you via email.


1.10 Logging via Splunk (6.x)

By default, Metric Insights 6.1.0 deploys a logger service (**rsyslog**) that collects all logs and saves them to `/opt/mi/log`. Optionally, instead of the default, you can use the **Splunk logging driver** that sends container logs to *HTTP Event Collector* in Splunk Enterprise and Splunk Cloud.

- For more information on default logging, refer to [Where are all those Log Files? \(6.0+\)](#)
- To learn about Splunk logging options, see [Splunk logging driver](#)

To configure logging via Splunk, use the process below to replace the default with the custom Splunk logging driver.


1. Disable the default logging driver

-  To disable the logger, install Metric Insights application with the option **--disable-logger**.

```
$ ./installer.py --disable-logger
```

2. Configure the HTTP Event Collector

The [HTTP Event Collector](#) is an endpoint allowing developers to send application events directly to the Splunk software via HTTP or HTTPS using a token-based authentication model.

-  An Administrator role is required to configure a HEC (HTTP Event Collector).

-  For each of the Metric Insights services, repeat all the steps described below.

PREREQUISITES:

- Log into Splunk Enterprise (or Splunk Cloud)

2.1. Access the HTTP Event Collector

The screenshot shows the Splunk Enterprise web interface. At the top, the 'Settings' menu is highlighted with a red circle labeled '1'. In the left sidebar, the 'Data inputs' link under the 'DATA' section is highlighted with a red circle labeled '2'. The main content area shows the 'Data inputs' configuration page. A table titled 'Local inputs' lists various input types. The 'HTTP Event Collector' row is highlighted with a red circle labeled '3'. A red arrow points from the 'HTTP Event Collector' link in the sidebar to the corresponding row in the table.

Type	Inputs	Actions
Files & Directories Index a local file or monitor an entire directory.	12	+ Add new
HTTP Event Collector Receive data over HTTP or HTTPS.	3	+ Add new
TCP Listen on a TCP port for incoming data, e.g. syslog.	0	+ Add new
UDP Listen on a UDP port for incoming data, e.g. syslog.	0	+ Add new
Scripts Run custom scripts to collect or generate more data.	33	+ Add new

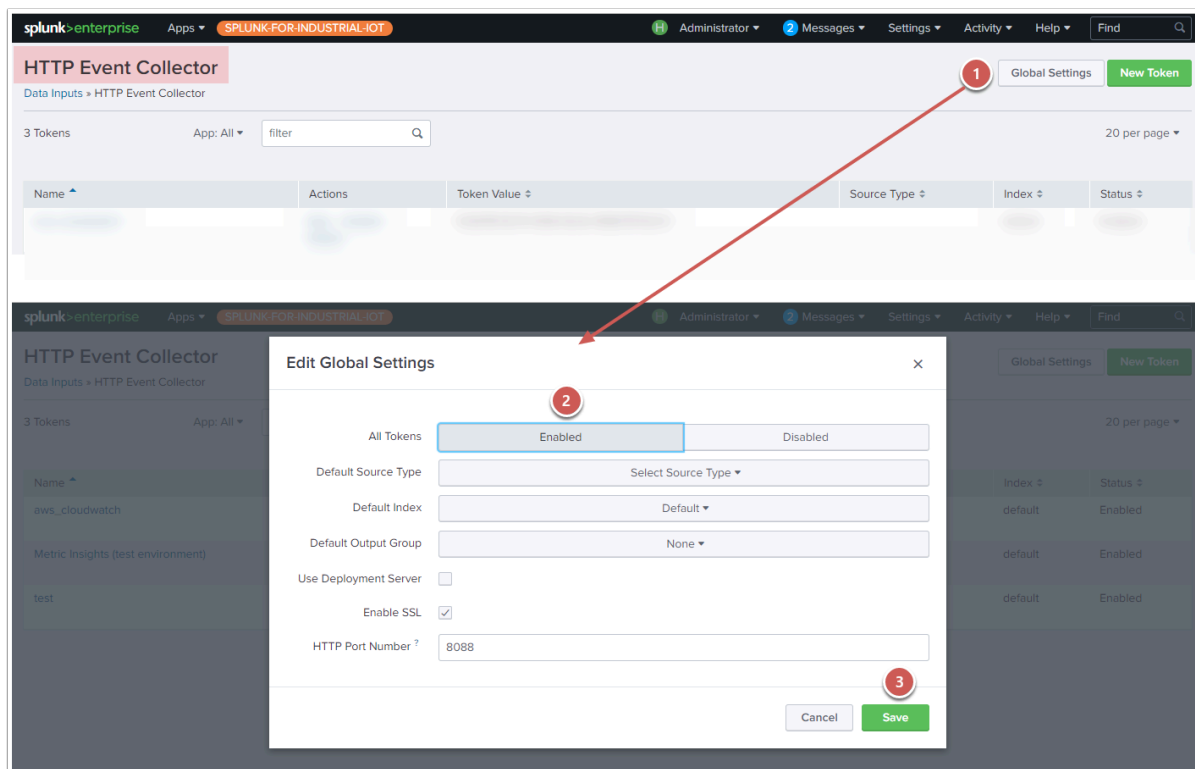
To configure an HTTP Event Collector:

1. Go to **Settings** menu
2. DATA section > select **Data inputs**
3. Click **HTTP Event Collector**

2.2. Enable Token use on the instance

⚠ NOTE:

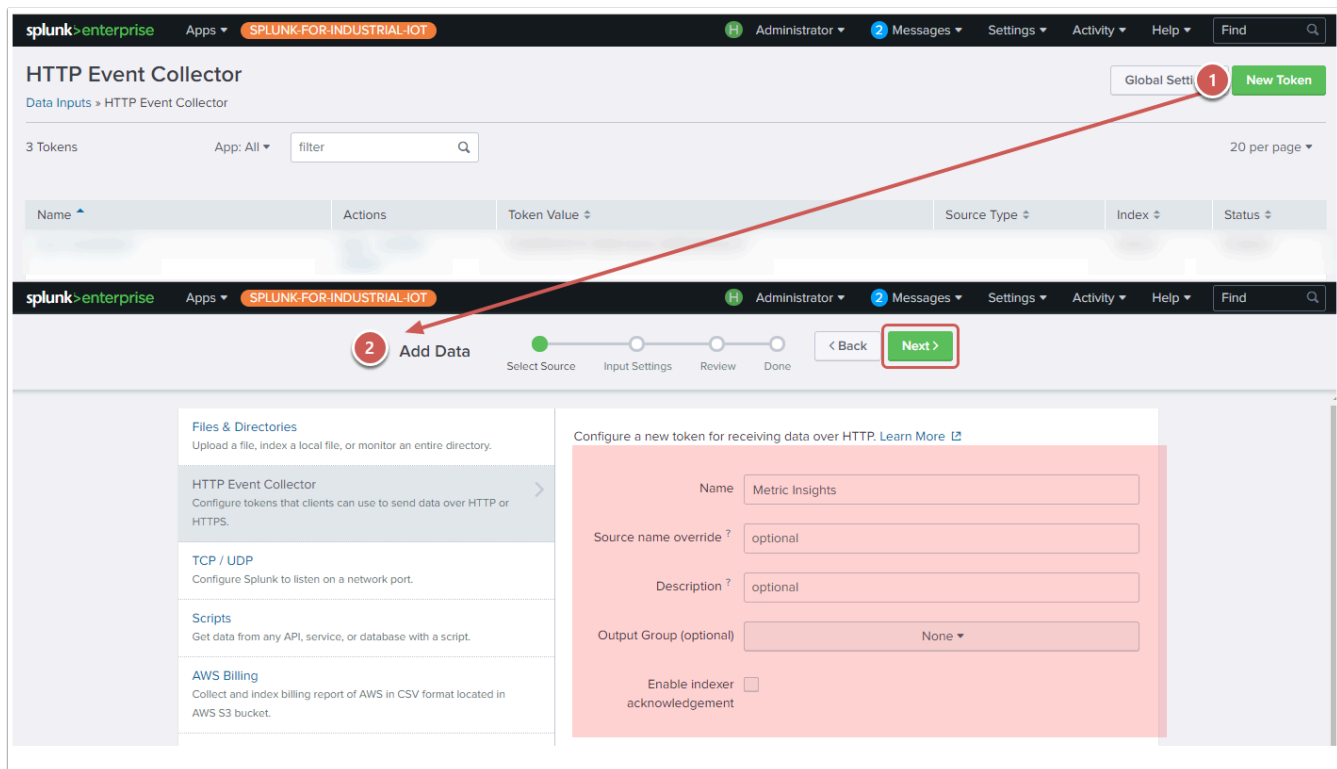
- Events sent without a Token are rejected.
- Before HEC (HTTP Event Collector) can be properly used, Token usage must be enabled on the instance.



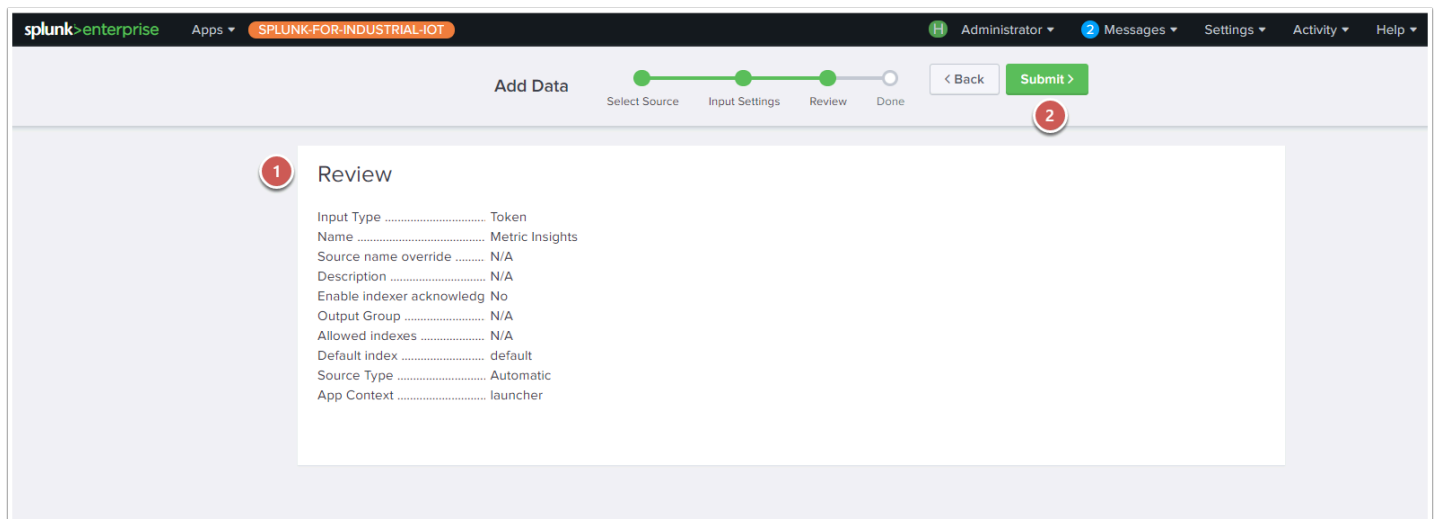
Having accessed the HTTP Event Collector:

1. Click **[Global Settings]**
2. **Enable** Tokens
3. **Save**

2.3. Add a new token for receiving data over HTTP

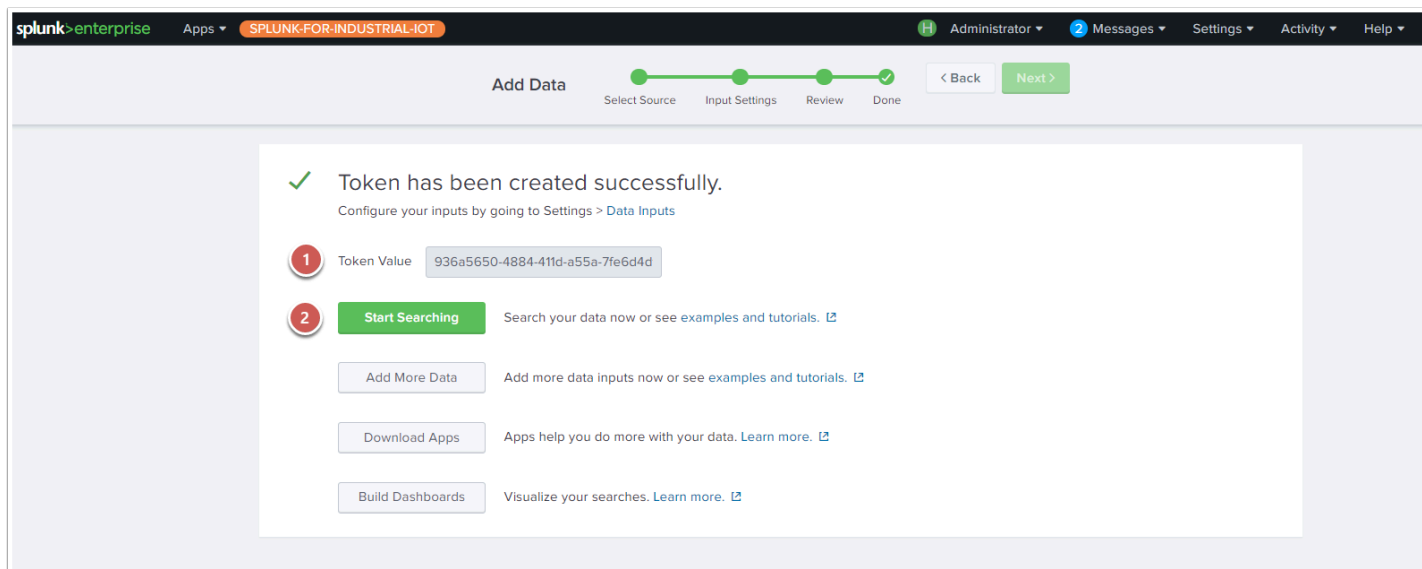


2.4. Review the settings and submit



1. **Review:** keep the default settings
2. Click **[Submit]** to proceed

2.5. Copy the Splunk Token



1. Copy the generated **Token Value**
2. **Start Searching**

3. Reconfigure logging in Metric Insights Docker Containers

To do the logging reconfiguration, go to the Metric Insights deployment subfolder and create the `logging.yml` file inside `/opt/mi/config/deployment`

PREREQUISITES:

1. Determine Splunk HEC URL `http(s)://<splunk_instance>:8088` (where 8088 is the HEC default port number determined in [Global Settings](#) that can be customized if needed)
2. Generate and determine HEC Tokens for each service running inside the Docker Containers

NOTE:

- Even though, it is possible to use one Token for all Docker Container Services, we recommend generating separate Tokens for each Container.

3.1. Create logging.yml

```
$ cd /opt/mi/config/deployment/
$ nano logging.yml
...
$ cat logging.yml
version: "3.4"
services:

  web:
    logging:
      driver: "splunk"
      options:
        splunk-token: xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx
        splunk-url: https://<splunk_instance>:8088
        splunk-insecureskipverify: "true"
        labels: web
    ...

  dataprocessor:
    ...

  seed:
    ...

  data-analyzer:
    ...

  mysql:
    ...
```

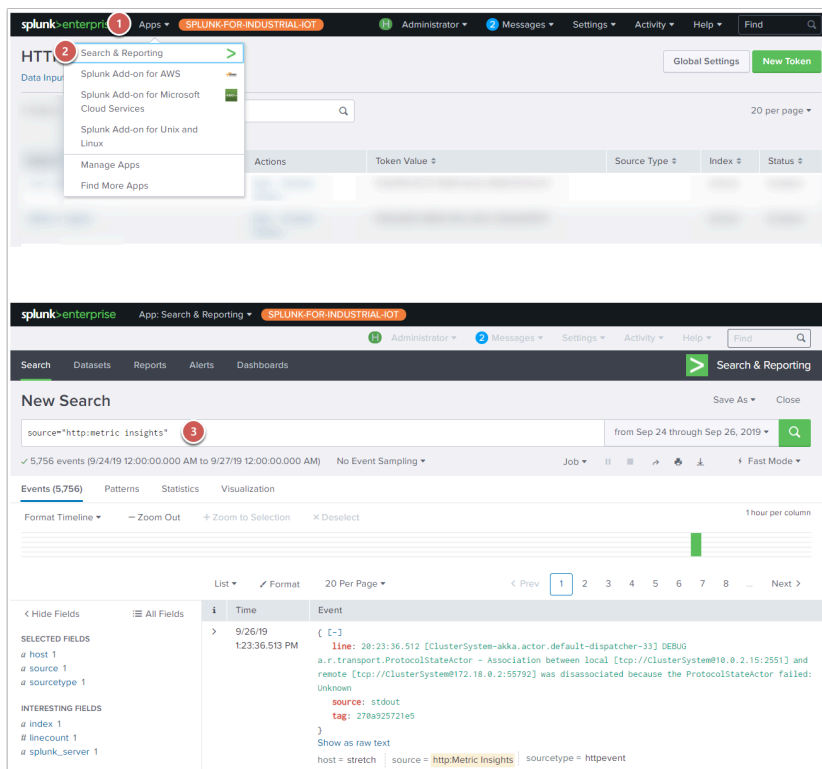
3.2. Edit the .env deployment config file

```
$ nano .env
$ cat .env | grep COMPOSE_FILE
COMPOSE_FILE=docker-compose.web.yml:docker-compose.mysql.yml:docker-compose.
dataprocessor.yml:docker-compose.seed.yml:docker-compose.data-analyzer.yml:logging.yml
```


3.3. Apply the new configuration

```
$ mi-control up -d
Recreating mi_mysql_1 ...
Recreating mi_seed_1 ...
```

4. Check if Metric Insights logs are sent to Splunk



To verify that logging is working correctly:

1. Open the **Apps** menu
2. Select **Search & Reporting**
3. Enter your **search query**

If all configured properly, logs will be displayed below.

1.11 Where are all those Log Files? (6.0+)

As of Release 6.0.0, Metric Insights becomes a dockerized application. All logs from Docker containers are stored in one directory **/opt/mi/log**.

- Deployment details are discussed in [Docker Deployment \(Simple Installer\)](#)

All logs are redirected to the **rsyslog service** that will store them in files. Alternatively, [Splunk logging service](#) can be used.

Log Files and their description

Types of Logs	Description
mi.log, mi.debug, mi.error	General logs from the UI and generator (migration utility scripts)
dataprocessor.log	Logs from the Data Processor service
seed.log	Logs from the Data Processor Seed Node service
mysql.log	Logs from MySQL service
web.log	Logs from the Web Service

Additional Logging options

 Logging can be configured using System Variables that are available from the UI.

1. To access, go to **Admin > System > System Config**
2. Click the **Edit (Gear)** icon to see the Variable's description

System Config

New...ContentAdminYana

System Variables

Change advanced system settings. Changes are not applied until they are committed.

Alllogging

System Variables

Uncommitted Changes

Variable Name	Assigned Value	Valid Values	Description	
ALLOW_EXT_LINK_TRACKING	Y	Y,N	Set to Y to track links to 3rd party reporting tool...	
DATA_COLLECTION_RESULT_LOGGING_EN...	Y	Y,N	Set to Y for all data collected during Trigger ex...	
EMAIL_LOGGING	N	Y,N	Y = generate a log file of all e-mail digest activi...	
ENABLE_APPLICATION_ERROR_LOGGING_I...	Y	Y,N	Enables logging of application errors into db ta...	
ENABLE_METRIC_VALUE_CHANGE_LOGGING	N	Y,N	Enables logging of changes for metric points	
LOG_CALC_TIME		true,false	"Broker logging, only use for profiling. Files can...	
MAX_SIZE_OF_DATA_COLLECTION_RESULT...	32		Failsafe for situations when large data inserts c...	
SAML_ASSERTION_ATTRIBUTES_LOGGING	N	N,Y		
UI_DATA_FETCH_LOGGING_ENABLED	Y	Y,N	Set to Y to log data fetches initiated from the U...	
USER_SESSION_LOGGING_ENABLED	Y	Y,N	Log user sessions	

Discard ChangesCommit Changes

1.12 Where are all those Log Files? (5.6.1 and lower)

Metric Insights has a number of different pieces to the application. The log files generated by each of these programs help us to diagnose any issues you may be having with your installation of Metric Insights. Below you'll find a list of the log files we use and where they are located.

There is also additional logging that can be turned ON/OFF through the Admin > Utilities > [Config Variables](#) page. See below for more details.

Log Files and Their Locations

1. Application Logging
 - /var/log/mi.error
 - /var/log/mi.log
 - /var/log/mi.debug
2. Data Fetch Request Broker Log: most recent file matching:
 - /opt/mi/iv/data/temp/broker/*.log
3. Data Collector Logs:
 - /opt/mi/datacollector/logs/insight.log
 - /opt/mi/datacollector/logs/query.log
4. E-mail:
 - /opt/mi/iv/data/temp/email/*.log

Turning ON and OFF additional debugging logs

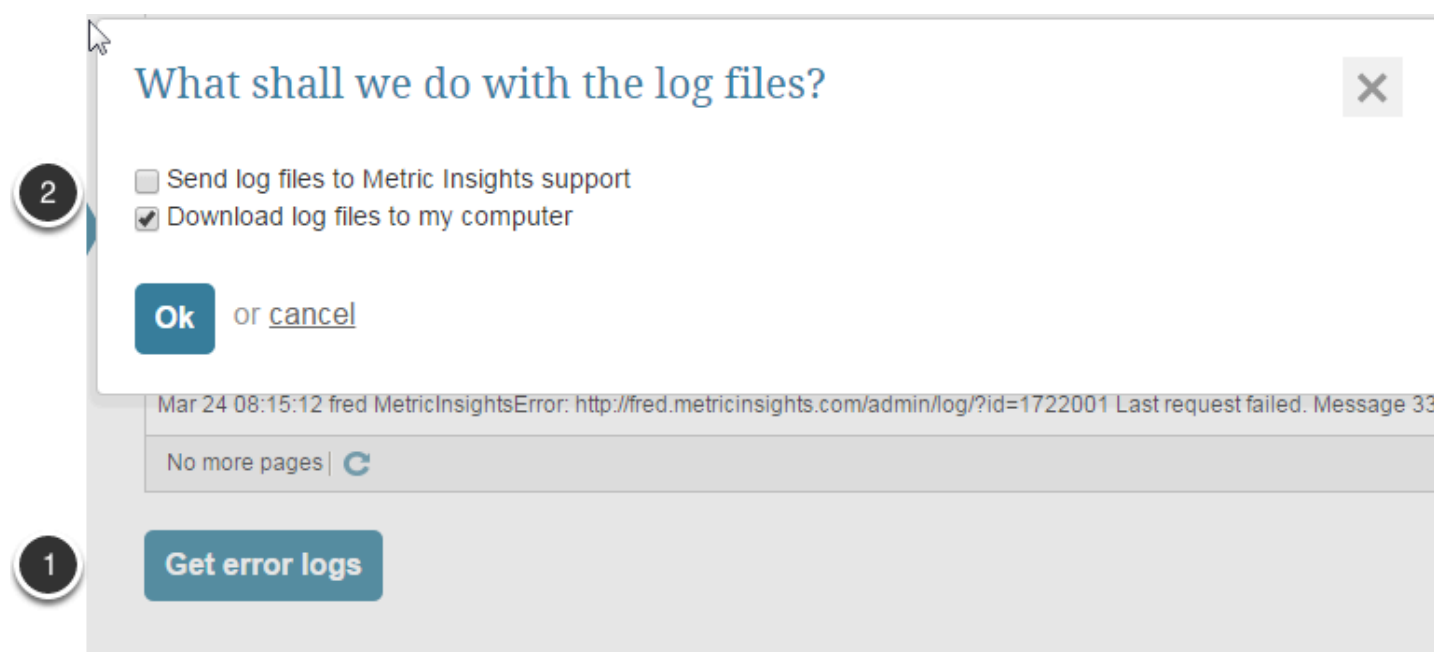
There are two additional logging switches that can be enabled for debugging. These can be turned ON and OFF in the [Config Variables](#) page:

1. UI_DATA_FETCH_LOGGING_ENABLED and DATA_COLLECTION_RESULT_LOGGING_ENABLED - writes full data collection queries and outputs to files in /var/www/iv/data/temp/broker/
2. EMAIL_LOGGING - writes email headers and attempts to /var/www/iv/data/temp/email/

1.13 Accessing the Metric Insights Log files

Metric Insights has a number of different pieces to the application. The log files generated by each of these programs help us to diagnose any issues you may be having with your installation of Metric Insights. Below you'll find a list of the log files we use and where they are located.

1. Easiest way to access error Log files - via the Status Monitor



1. There is a **Get error logs** button on the Status Monitor > Application Errors tab page.
2. From the pop-up, you can download the logs (zipped) to your local machine, send them to our support team for assistance, or both.

2. Manually access Log Files

Logs can exist in different locations depending on the Operating System:

- Debian
- CentOS/RHEL

2.1. Debian log files

1. Apache error log
 - /var/log/apache2/error.log
 - /var/log/apache2/pyinsight-error.log
 - /var/log/apache2/ssl_access.log
2. Report Generator Log – most recent file matching:
 - /var/log/mi.debug
3. Data Fetch Request Broker Log: most recent file matching:
 - /opt/mi/iv/data/temp/broker/*.log
4. Data Collector Logs:
 - /opt/mi/datacollector/logs/insight.log
 - /opt/mi/datacollector/logs/query.log
5. Mysql:
 - /var/log/syslog

2.2. CentOS / RHEL log files

1. Apache error log
 - /var/log/httpd/error_log
 - /var/log/httpd/pyinsight-error.log
 - /var/log/httpd/ssl_access.log
2. Report Generator Log – most recent file matching:
 - /var/log/mi.debug
3. Data Fetch Request Broker Log: most recent file matching:
 - /opt/mi/iv/data/temp/broker/*.log
4. Data Collector Logs:
 - /opt/mi/datacollector/logs/insight.log
 - /opt/mi/datacollector/logs/query.log
5. Mysql:
 - /var/lib/mysql/hostname.err [where hostname is the name of the server.]

1.14 Install or Update MI via Installation Packages

This article provides a step-by-step instruction on how to install Metric Insights on your virtual machine and describes the ways available to update the system to a newer version.

NOTE: Metric Insights by default installs to **/opt**. MySQL by default installs to **/var**. The installer and the application makes use of **/tmp**. If the three directories live as *separate partitions* on the host server, please ensure there is adequate disk space to support the application and database.

```
Metric Insights app directory = /opt/mi
MySQL default directory = /var/lib/mysql
Temp files = /tmp
```

For disk space requirements, please see [What are the system requirements for a Metric Insights server?](#)

1. Connect to the server by ssh

1. Use **ssh root@000.000.000.000** and replace the zeros with the server IP address
2. When prompted, enter the root password provided by your IT specialist

2. Download MI installation package

1

2

3

4

required version

METRIC INSIGHTS

Product Solutions Learn Support About Free Trial

Download

Enter your user name and password below

Username *

Password *

Metric Insights Software Evaluation Agreement

This Metric Insights Customer Agreement (this "Agreement") contains the terms and conditions that govern your access to and use of the Service Offerings (as defined below) and is an agreement between Metric Insights, Inc. ("Metric Insights"), a California company, with an address at 123 10th Street, San Francisco, CA 94103, and you or the entity you represent ("User"). This Agreement takes effect when an Order is signed, when you click an "I Accept" button or check box presented with these terms, when this agreement is signed between Metric Insights and User, or, if earlier, when you use any of the Service Offerings (the "Effective Date"). You represent to us that you are lawfully able to enter into contracts (e.g., you are not a minor). If you are entering into this Agreement for an entity, such as the company you work for, you represent to us that you have legal authority to bind that entity.

☐ Download license agreement ☐ I agree with these terms and conditions

Submit

1. Depending on the system version your organization plans to install, Metric Insights team will provide you with a link to a corresponding installation package version
2. Along with a link, **Username** and **Password** should be provided by the MI support team.; enter credentials into corresponding fields
3. Check the box to confirm your agreement with terms and conditions of Metric Insights usage
4. Click **Submit**

The download process of the .zip archive containing all files required to install the system will begin.

3. Operating Systems Supported

 redhat.	 CentOS	 debian
Versions: 6.6, 7.2	Versions: 6.7, 6.9, 7.0, 7.2	Versions: Wheezy, Jessie

3.1. Additional Actions for Operating Systems

Different operating systems require specific additional steps to be performed before and after installation process. Refer to the table below to learn additional commands that should be run for successful package installation.

If you are using a firewall on the linux machine, inbound TCP ports 443 and 8443 are necessary for Metric Insights to work correctly. To open them use the following commands:

Examples for **iptables**:

```
iptables -A INPUT -p tcp --dport 80 -m state --state NEW,ESTABLISHED -j ACCEPT
iptables -A OUTPUT -p tcp --sport 80 -m state --state ESTABLISHED -j ACCEPT
iptables -A INPUT -p tcp --dport 443 -m state --state NEW,ESTABLISHED -j ACCEPT
iptables -A OUTPUT -p tcp --sport 443 -m state --state ESTABLISHED -j ACCEPT
iptables -A INPUT -p tcp --dport 8443 -m state --state NEW,ESTABLISHED -j ACCEPT
iptables -A OUTPUT -p tcp --sport 8443 -m state --state ESTABLISHED -j ACCEPT
```

Examples for **ufw**:

```
ufw allow 80/tcp
ufw allow 443/tcp
ufw allow 8443/tcp
```

Examples for **firewall-cmd**:

```
firewall-cmd --zone=public --add-port=80/tcp --permanent
firewall-cmd --zone=public --add-port=443/tcp --permanent
firewall-cmd --zone=public --add-port=8443/tcp --permanent
firewall-cmd --reload
```

Operating System	Additional steps
Centos 6, 7	<ol style="list-style-type: none"> 1. <code>yum -y update</code> 2. <code>sudo yum install epel-release</code>
Debian 7	<p>If you cannot get access to ssh root@000.000.000.000, run the following command to verify whether the ssh service has been previously installed:</p> <pre>service ssh status</pre> <p>If it is installed, the system will return the OK status, informing you that ssh is running. If the system perceives this command as unknown, install ssh service on Debian:</p> <pre>apt-get install ssh</pre>

RedHat 7

1. `yum -y update`
2. `cd /tmp/`
3. `wget https://dl.fedoraproject.org/pub/epel/epel-release-latest-7.noarch.rpm`
4. `sudo yum install epel-release-latest-7.noarch.rpm`
5. `subscription-manager repos --enable=rhel-7-server-optional-rpms`

4. Ensure required Linux packages to support the application are available

See this article for more details: [Operating system linux package requirements / dependencies](#)

5. Unpack the downloaded zip archive using the following command

```
tar -xvf MetricInsights-Installer-<version_number>.tar.gz
```

6. Change directory to the new install location

```
cd MetricInsights-Installer-<version_number>
```

7. Initiate the installation process

```
./installer.py -vv
```

NOTE: The installer must be run as root or a user with sudo access.

For additional parameters to use with the installer, please see `./installer.py --help`

If you encounter any errors, please contact support@metricinsights.com.

8. Post-checks once the install is complete

Once the installation is complete, confirm the version of Metric Insights from the command line by running this command:

```
mi-version
```

Also confirm that major system services for the application are running without issue:

- MySQL/MariaDB service
- Apache service
- Cron service
- NTP service

Next, log into the UI and confirm the version number in the lower left-hand corner of the homepage. Also check Admin > Status Monitor to get an overview of the application's status.

9. If upgrading by running the installer

First, back up your application before running the installer to upgrade. Run this utility from the command line:

```
mi-app-backup -vv
```

By default, the backup file will save to `/var/backups/mi-app-backups/`. To specify a different path and for additional parameters to use with the utility, please see `mi-app-backup --help`.

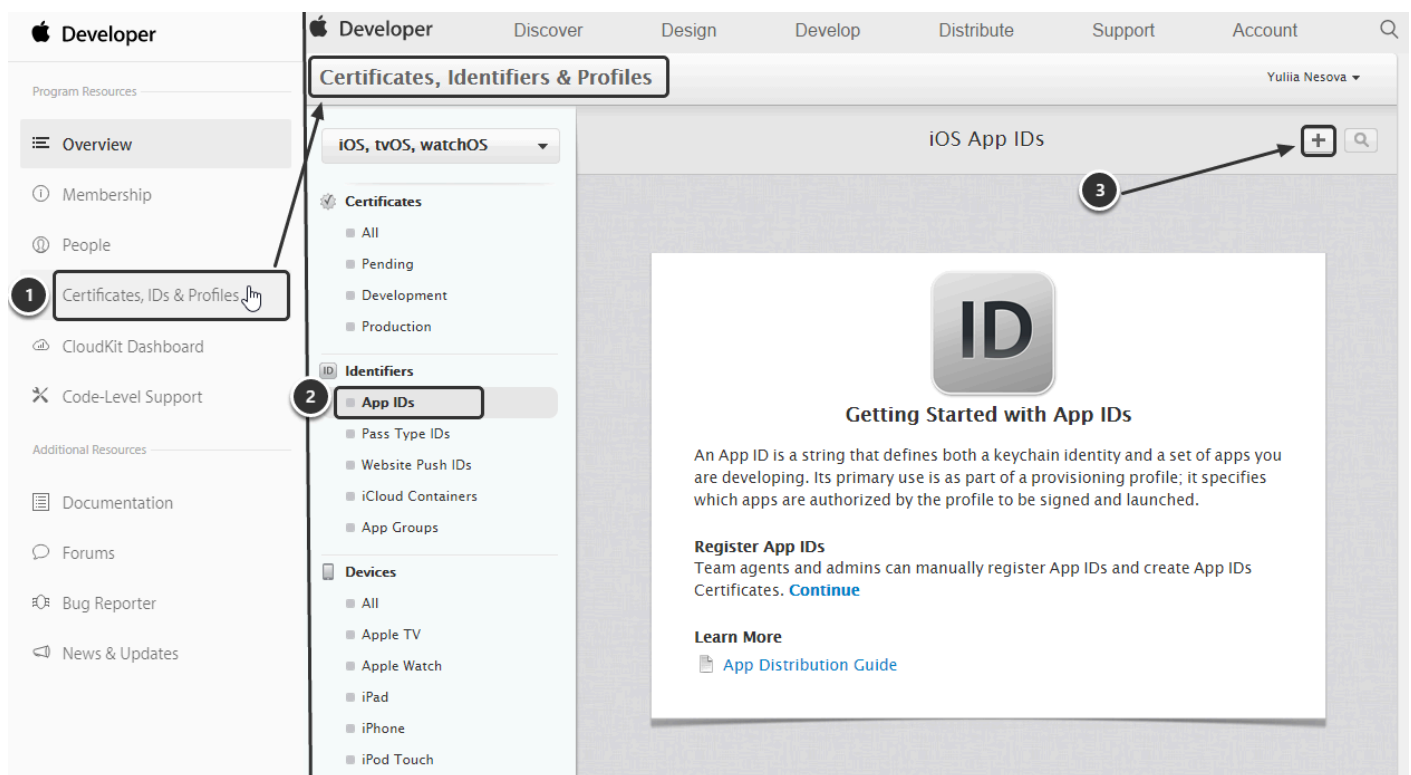
1.15 Deploying Metric Insights' Custom Mobile Apps at your organization

This document will describe the steps needed to deploy the Metric Insights mobile app on iOS; Android is much more simple—a small section at the bottom is dedicated to it.

PREREQUISITES

- Submit your information via the Custom Mobile Build form (please request it from support@metricinsights.com)
- Receive the Xcode project and APK for your custom mobile build from Metric Insights team
- Acquire an Apple Developer Enterprise account

Create a new 'Metric Insights' App ID




Log into the [iOS Dev Center](#) and open your **Account** page:

1. From the **Program Resources** Navigation menu to the left of the screen select **Certificates, IDs & Profiles**
2. In the opened page select **App IDs** from the **Identifiers** section to the left of the screen

3. At the **iOS App IDs** screen click **Add** (a plus button)

Register an App ID


Registering an App ID

The App ID string contains two parts separated by a period (.) — an App ID Prefix that is defined as your Team ID by default and an App ID Suffix that is defined as a Bundle ID search string. Each part of an App ID has different and important uses for your app. [Learn more](#)

App ID Description

1

Name:

You cannot use special characters such as @, &, " , ' , "

App ID Prefix

Value:

App ID Suffix

* **Explicit App ID**

If you plan to incorporate app services, such as Game Center, In-App Purchase, Data Protection, and iCloud, or want a provisioning profile unique to a single app, you must register an explicit App ID for your app.

To create an explicit App ID, enter a unique string in the Bundle ID field. This string should match the Bundle ID of your app.

2

Bundle ID:

We recommend using a reverse-domain name style string (i.e., com.domainname.appname). It cannot contain an asterisk (*).

○ **Wildcard App ID**

This allows you to use a single App ID to match multiple apps. To create a wildcard App ID, enter an asterisk (*) as the last digit in the Bundle ID field.

Bundle ID:

Example: com.domainname.*

App Services

Select the services you would like to enable in your app. You can edit your choices after this App ID has been registered.

Enable Services:

☐ App Groups

☐ Associated Domains

☐ Data Protection

☐ Complete Protection

☐ Protected Unless Open

☐ Protected Until First User Authentication

☒ Game Center

☐ HealthKit

☐ HomeKit

☐ Wireless Accessory Configuration

☐ iCloud

☐ Compatible with iCloud

☐ Include CloudKit support (requires iCloud)

☒ In-App Purchase

☐ In-App Apple

☐ Wallet

☒ Push Notifications

☐ Remote View

☐ Service

Cancel

Continue

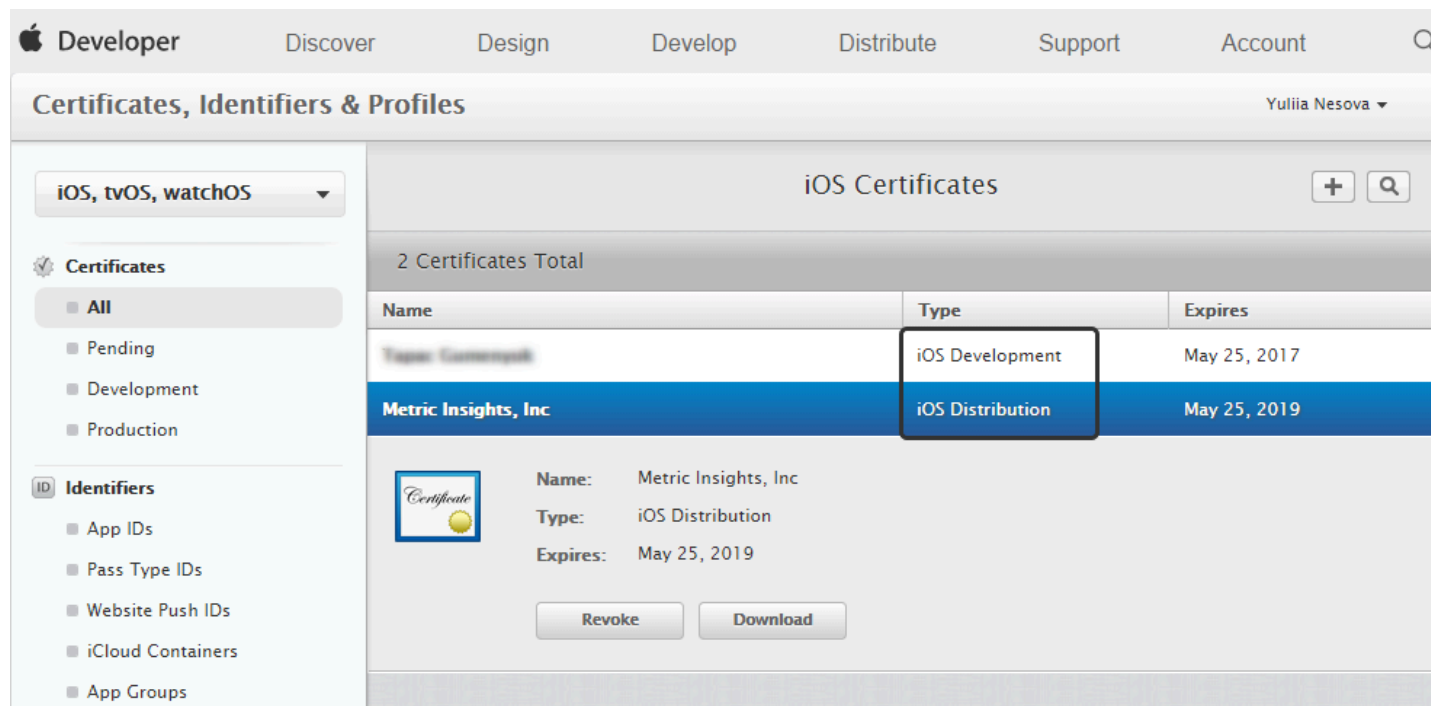
1. Define the **Name** of a new App ID
2. Select **Explicit App ID** option and define Bundle ID
3. In the **App Services** section make sure to select **Push Notifications** to enable them for your application

You should now see the info about App ID that you have created. Click **Submit** to continue.

Deployment & Configuration

Page 66

Create a "Metric Insights" Certificate for Push Notifications



Generate a certificate for the and upload it the '**certificates**' directory on your Metric Insights instance at:

```
/opt/mi/iv/data/certificates
```

Certificate has 2 variants: **[app_id].pem** (prooduction) and **[app_id].dev.pem** (development).

For example, 'com.metricinsights.viewer.pem' or 'com.metricinsights.viewer.dev.pem'.

For this reason, we have a Config Variable in Metric Insights that determines which certificate should be used (see next step).

Define the type of certificate to be used in Metric Insights

1. Go to *Admin > Utilities > Config Variables*
2. Search for a **MOBILE_PUSH_ENVIRONMENT** variable and define which certificate should be used and click **Save** in the pop-up
3. Do not forget to click **Save changes** at the bottom of the *Config Variables* page to apply the changes

Once the certificates have been created and are in the appropriate locations, Push Notifications should begin to work.

NOTE: The certificates should be added to all supported instances (dev, prod, etc.).

If you have created a passphrase while building the Certificate, copy it into Metric Insights

1. Go to *Admin > Utilities > Config Variables*
2. Find **MOBILE_PUSH_PRIVATE_KEY_PASSPHRASE** variable and click a **Gear** button to edit it
3. Define a passphrase for this variable and click **Save** in the pop-up
4. Do not forget to click **Save changes** at the bottom of the *Config Variables* page to apply the changes
5. Log out of **Flex Alerts** on the mobile device, and log back in
6. Now go back to **My Settings** page in Metric Insights and try resending the test alert (described in the step below)

NOTE: You'll need to go through these steps for every server that needs to function with the Dashmonitor app.

Testing Alerts for an individual user

1. In the drop-down menu below your name select **My Settings**
2. In the opened screen click **Send a test Alert**. Assuming the device is logged in to the instance from which you are sending the test, you should receive it on your phone

Resending an existing alert

To test Alerts more thoroughly, you can create an Alert and then re-trigger it:

To get the Alert ID:

1. Find the required Alert
2. Go the *Metric Editor* of the Metric where this Alert has fired
3. Open the *Alerting* tab; find and copy **Alert ID** in the **Alert History** grid
4. Go to **[INSTANCE_URL]/service/ns/alert**
5. Enter **Alert ID** in the corresponding field and click **Send** at the bottom of the screen
6. System informs you whether the alert has been successfully resent

Android Requirements

On Android, there is no additional setup required for receiving Push Notifications. There is a key pair that comes with the MI instance and the Android app, and as there is no third party gatekeeper that manages notifications (e.g. Apple Push Notification Service), it should work as-is.

1.16 Deploying Mobile Apps with support for Deep Links (links to the Metric Insights Mobile App)

This document describes the basic process for deploying a mobile app with support for Deep Links (e.g. a link is clicked in an email from Metric Insights, and the Metric Insights mobile app is opened rather than the web browser).

Note This cannot be done without a custom build from the Metric Insights Engineering Team.

Provide List of Supported Domains

In order for deep linking to work, the mobile app (iOS or Android) needs to know which domains to route to the Metric Insights app instead of a web browser.

This is done by including a list of **Associated Domains**, e.g:

- metricinsights.company.com
- metricinsights-qa.company.com
- metricinsights-dev.company.com

This would support 3 separate metric insights environments (product, development and QA)

Associated Domains need to be applied both in the custom build prepared by Metric Insights, and in the customer's own Apple Developer Enterprise account:

[← All Identifiers](#)

Edit your App ID Configuration

Remove

Platform

iOS, macOS, tvOS, watchOS

App ID Prefix

H2QP6695AT (Team ID)

Description

Metrics

Bundle ID



com.metricinsights.metrics (explicit)

You cannot use special characters such as @, &, *, ' , "

Capabilities

ENABLED

NAME _____

 [Access WiFi Information](#) 

☒ App Groups

Configure

Enabled App Groups (0)

Associated Domains

☐ AutoFill Credential Provider  ClassKit 

Important You will need to provide the Metric Insights Engineering team with your list of Associated Domains for a custom build with Deep Linking to be prepared.

To set the [site association](#), go to `https://<metric insights domain>/apple-app-site-association` (e.g. `https://metricinsights.company.com//apple-app-site-association`). It should prompt you to download a file. Proceed and open the file in a text editor.

You'll need to replace the X's in `XXXXXXXXXX.com.metricinsights.mobile.ios` with your own App ID:

```
{
  "applinks": {
    "apps": [],
    "details": [
      {
        "appID": "XXXXXXXXXX.com.metricinsights.mobile.ios",
        "paths": [ "/deep_link*" ]
      }
    ]
  },
}
```



```

"webcredentials": {
  "apps": [
    "XXXXXXXXXX.com.metricinsights.mobile.ios"
  ]
}
}

```

Once the file has been edited, it needs to be uploaded to the server.

Optional: Set Deep Link Prefix

i This information is provided for your understanding, but we recommend sticking to the system default ("deep_link") and not making any modifications.

There is a single system Config Variable, **MOBILE_DEEP_LINKING_PREFIX** that defines the prefix to place before the trailing URL, e.g. `https://<domain>/<deep link prefix>/<remainingurl>` or, expressed more clearly, `https://metricinsights.company.com/deep_link/report/index/index/element/123/segment/456`

MOBILE_DEEP_LINKING_PREFIX is set to `deep_link` by default and **does not need to be changed**.

All links in Bursts and other email distributions that include an Element will use this URL pattern.

Scope of objects that support Deep Linking

Deep linking will work for the following Element types:

1. Reports (including Dimensions)
2. Metrics (including Dimensions)
3. External Reports (including Dimensions)

Technically, each Element supports the following URL pattern

- `https://<domain>/deep_link/report/index/index/element/<int value>/segment/<int value>`
- `https://<domain>/deep_link/chart/index/index/element/<int value>/segment/<int value>`
- `https://<domain>/deep_link/extreport/index/index/element/<int value>/segment/<int value>`

1.17 MI Slackbot deployment

This article provides a step-by-step instruction on how to install the **MI Slackbot** on your server and update the bot to a newer version.

The **MI Slackbot** is realized as part of the **MI Slack App functionality** and works in conjunction with the **Slack App** built as bot on the Slack Side. After the MI Slackbot has been installed on your Server and the Slack App with bot functions has been built in Slack, the two have to be connected (*MI Slack App = MI Slackbot + Slack App*).

The whole process comprises the following steps:

1. MI Slackbot deployment
2. [Building a bot app for Slack](#)
3. [Connecting Slack App to Metric Insights](#)

1. Connect to the server by ssh

1. Use **ssh root@000.000.000.000** and replace the zeros with the server IP address.
2. When prompted, enter the root password provided by your IT specialist.

2. Download the MI Slackbot installation package

The screenshot shows a web browser at the URL <https://www.metricinsights.com/download/MetricInsights-Slackbot-v5.3.2-Full.tar.gz>. The browser's address bar highlights the URL with a callout box labeled "required version". The page features the Metric Insights logo and a navigation menu with links for Product, Solutions, Learn, Support, and About, along with a "Free Trial" button. The main heading is "Download", followed by the instruction "Enter your user name and password below". A form contains two input fields: "Username *" and "Password *", both marked with an asterisk. Below the form is a section titled "Metric Insights Software Evaluation Agreement" with a scrollable text area containing the terms of service. At the bottom, there is a checkbox labeled "I agree with these terms and conditions" and a "Submit" button. Four numbered callouts are present: 1 points to the URL, 2 points to the Username field, 3 points to the agreement checkbox, and 4 points to the Submit button.

1. Depending on the version, you will be given a link to the corresponding MI Slackbot installation package.
2. Along with the link, **Username** and **Password** will be provided by the MI support team; enter these credentials into the respective fields.
3. Confirm your agreement with the Terms and Conditions of the MI Slackbot usage.
4. Click **[Submit]**.

This will initiate the download process of the **.tar.gz** archive (containing all files required to install the system).

3. Operating Systems Supported



The list of **supported Operating Systems**:

1. Debian 7 (*wheezy*)
2. Debian 8 (*jessie*)
3. Red Hat Enterprise Linux 7
4. CentOS 6 (deprecated in Release 5.5)
5. CentOS 7
6. Ubuntu 14.04

3.1. Additional Actions for Operating Systems

If you are using a firewall on the Linux machine, inbound TCP ports 80 and 443 are necessary for the MI Slackbot to work correctly. If they need to be opened, use the following commands:

Examples for **iptables**:

```
iptables -A INPUT -p tcp --dport 80 -m state --state NEW,ESTABLISHED -j ACCEPT
iptables -A OUTPUT -p tcp --sport 80 -m state --state ESTABLISHED -j ACCEPT
iptables -A INPUT -p tcp --dport 443 -m state --state NEW,ESTABLISHED -j ACCEPT
iptables -A OUTPUT -p tcp --sport 443 -m state --state ESTABLISHED -j ACCEPT
```

Examples for **ufw**:


```
ufw allow 80/tcp  
ufw allow 443/tcp
```

Examples for **firewall-cmd**:

```
firewall-cmd --zone=public --add-port=80/tcp --permanent  
firewall-cmd --zone=public --add-port=443/tcp --permanent
```


4. Unpack the downloaded .tar.gz archive with the following command


```
tar -xvf MetricInsights-Slackbot-<version_number>.tar.gz
```

5. Enter the new unpacked folder

```
cd MetricInsights-Slackbot-<version_number>
```

6. Initiate the installation process

-  • The MI Slackbot must have an open access to the web.
- The installation commands below can only be run by a root user.

-  • The MI Slackbot can be installed on a server with the pre-installed MI application (Version 5.3.0. and beyond).
- The Slackbot installer will not work with an MI application prior to 5.3.0

6.1. Install the MI Slackbot on the same server as the main MI application

```
./installer.py
```


6.2. Install the MI Slackbot on a separate server

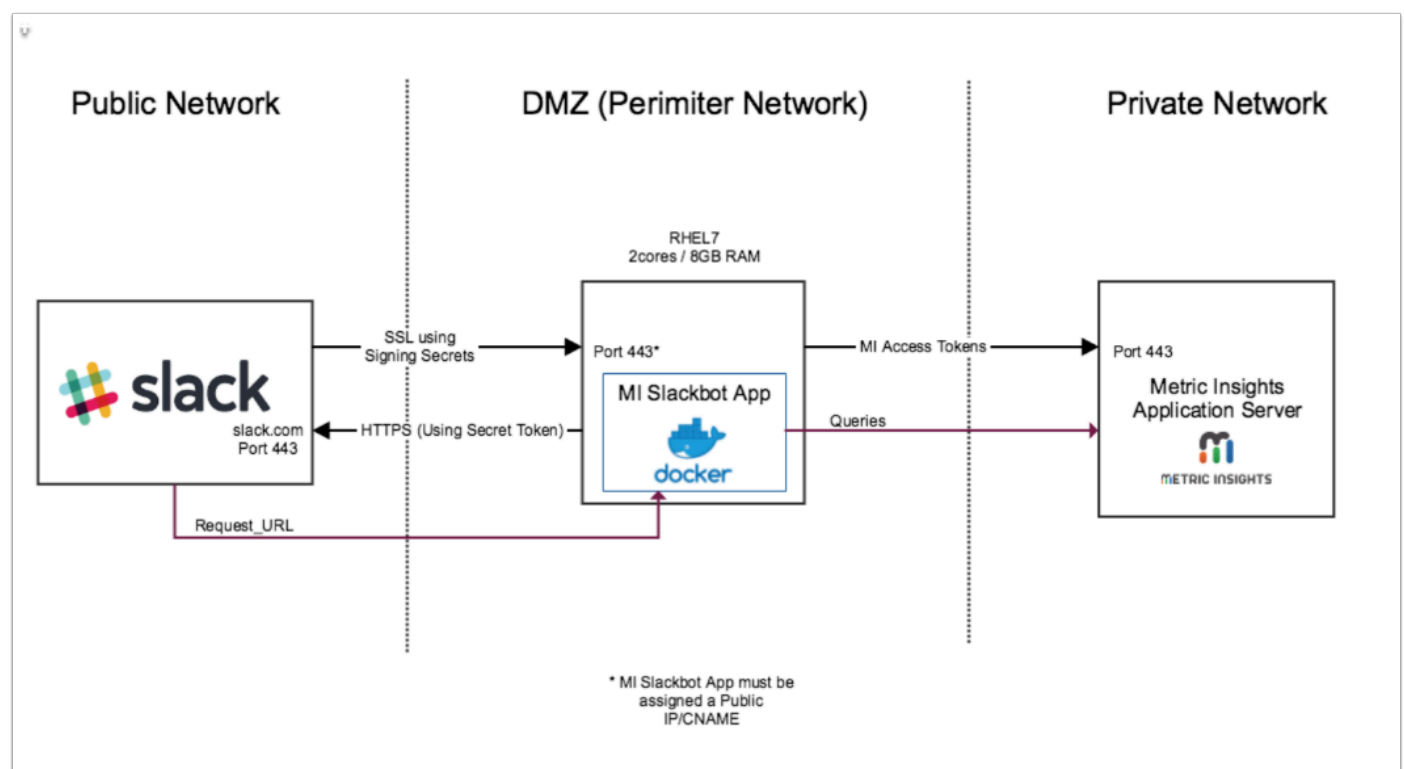
```
./installer.py --standalone
```

6.3. Install the MI Slackbot on a DMZ server

The **MI Slackbot** can also be deployed within a **DMZ (Perimeter) network**.

As an external-facing app, the MI Slackbot must be accessible from the internet. DMZ deployment allows the Bot app to be reachable from the Public Network, while the data it queries from Metric Insights remains securely locked within a Private Network.

For details, view the Architecture diagram below.



⚠ In case there is a *connection issue* or a *WebSocket failure* (check for errors with [Slack connection test](#)), you will need to:

- **Whitelist the following domains:** *.slack.com *.slack-msgs.com *.slack-files.com *.slack-imgs.com *.slack-edge.com *.slack-core.com *.slack-redir.net slackb.com edgeapi.slack.com wss-primary.slack.com wss-backup.slack.com wss-mobile.slack.com app.slack.com a.slack-imgs.com b.slack-imgs.com
- **Check if your proxy is running SSL decryption.** If it is, the proxy must either support *WebSockets*, or you'll need to exempt the following domains: *.slack-msgs.com wss-primary.slack.com wss-backup.slack.com wss-mobile.slack.com

For more information, see [Manage Slack connection issues](#)

6.4. The list of arguments that can be used during installation

Optional Parameters	
-h, --help	Show this help message and exit.
-v, --verbose	Output a diagnostic information for every file processed.
--path=INSTALL_PATH	Base path for Metric Insights.
-y, --yes	Assume yes on interactive requests.
-r, --remove	Remove mi-slackbot from current system.
-p, --purge	Purge mi-slackbot data folders and config files.
--standalone	Install Metric Insight Slackbot as standalone service.

7. The MI Slackbot update via packages

This method only works if the MI Slackbot has been initially installed via the installation package (as described above).

To update the MI Slackbot:

1. Download the new package
2. Extract the archive

3. Install it over the old version

8. What's next?

- 💡. [Building a bot app for Slack](#)
2. [Connecting Slack App to Metric Insights](#)

1.18 Operating system Linux package requirements for v6.1+

Metric Insights can be installed on the following Linux operating systems using a package installer provided by MI Support:

- RHEL
- CentOS
- Debian
- Ubuntu

If installing on an internally provisioned server, **please note the following list of required linux packages needed to support the Metric Insights application.** The application runs from inside Docker containers deployed to the host system. The packages required are primarily for supporting Docker and the installer package used for the initial deployment.

Note:

** This list is current for Metric Insights version 6.1.0+*

** For Debian deployments, you must have a kernel version 3.10+ for docker support*

Linux packages required on host system

- apt-transport-https
- ca-certificates
- curl
- gnupg2
- gzip
- ntp
- wget
- software-properties-common
- docker-ce v17.09.0+
- docker-compose v1.19.0+

1.19 Deploying Container Orchestration with Docker Swarm

Docker Swarm is an excellent option for container orchestration if you meet the following conditions:

- Need to utilize existing hardware procured for Metric Insights v5.x
- Don't have an AWS account to deploy to ECS
- Don't have an Azure account to deploy to AKS
- Don't have an on-premise Kubernetes cluster to deploy

In terms of hardware, for Docker Swarm to work, we would need to deploy Metric Insights across a minimum of *two application servers* (with each server serving as nodes for the containers to run on).*

Please see [Operating system Linux package requirements for v6.1+](#) for host server linux package requirements.

** Note, for best fault tolerance, three application nodes are recommended.*

 [This article will cover the following](#)


1. [Requirements](#)
2. [Preparing for the deployment](#)
3. [Deployment](#)
4. [Architectural Diagram of Docker Swarm](#)
5. [Upgrading Metric Insights to a newer version](#)

Requirements

1. Docker-CE with swarm mode enabled
2. Connect the required list of nodes to the Docker Swarm cluster
3. Credentials for your private Docker Registry to pull Metric Insights docker images
 - You can also pull from Metric Insights' official docker registry (contact support@metricinsights.com for credentials)
4. NFS Share mounted to all nodes to store the Metric Insights file system

- This ensures continuous operations where critical application and data files survive and containers across the cluster continue to run should any nodes go down.
5. Ensure the following ports are open in the load balancer:
 - **80, 443** - HTTP and HTTPS ports for the Web UI Service (default redirection to 443)
 - **2550** - TCP port for the Data Processor cluster
 - **2551** - TCP port for the Seed Node Service
 - **3306** - MySQL port for external access
 - **8080, 8443** - HTTP and HTTPS ports for the REST API Data Processor Service
 - **8081** - TCP port for Monitoring Tool

Preparing for the deployment

 See [Docker Commands Cheat Sheet](#) to understand the Docker commands used below.

1. Install all required packages to both application nodes to support Docker
2. Provision a remote MySQL server database to store the application database
 - MySQL 5.6/5.7 and MariaDB 10.1/10.2 are currently supported
 - MySQL root user required for the deploy
3. Provision an NFS share to mount to all nodes to store the MI filesystem (*default path is `/opt/mi`*)
4. Identify and label the nodes as *master* and *slave*:
 - For master node, enter the following docker command to label it as master:


```
docker node update --label-add type=master node_master
```
 - For slave node, enter the following docker command to label it as slave:


```
docker node update --label-add type=worker node_slave
```
5. Download the Metric Insights installer then unpack to the master node (contact support@metricinsights.com for the installer tarball). This will create an installer directory from where we can run the installer to generate a Deployment Manifest for Docker Swarm
 - ```
tar xf MetricInsights-Installer-v6.x.x-Full.tar.gz
```
  - ```
cd MetricInsights-Installer-v6.x.x-Full
```
6. Prepare Docker Secret files for each MI service
 - In the installer directory, you can find the template files to generate the secrets in `../utils/orchestration/swarm/secrets`
 - Copy the `*.env.example` files there to your local directory and rename to `*.env`
 - Each template file represents a different MI service:
 - web.env
 - seed.env
 - dataprocessor.env
 - data-analyzer.env
 - monitoring.env

- Edit each template file and update each parameter. Make sure the passwords you generate are consistent throughout the files
7. Generate the deployment manifest to use with Docker Swarm by running the following from the installer directory: `./installer.py swarm -o mi-swarm.yml`
 8. Finally, update the NFS share address in the deployment manifest file:
 - Look for this section and change the address `nfs-server.metricinsights.com` to the appropriate NFS share:

```
data:
  driver_opts:
    type: "nfs"
    o: "addr=nfs-server.metricinsights.com,nolock,soft,rw"
    device: ":/opt/mi/data/data"
```

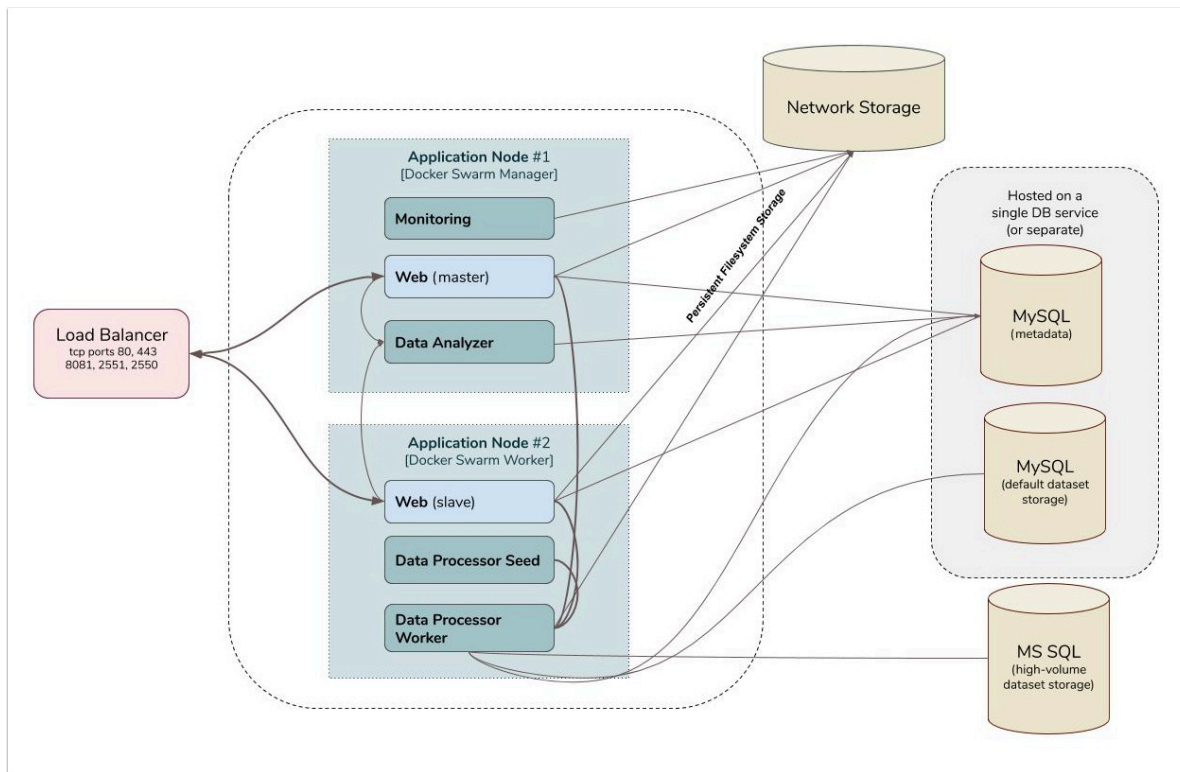
Deployment

To deploy docker swarm using the deployment manifest, use the **docker stack deploy** command:

```
$ docker stack deploy -c mi-swarm.yml --with-registry-auth mi
Creating secret mi_mysql
Creating secret mi_web
Creating secret mi_dataprocessor
Creating secret mi_seed
Creating secret mi_monitoring
Creating service mi_dataprocessor
Creating service mi_seed
Creating service mi_web
Creating service mi_monitoring
```

- Note, the `docker stack deploy` command will create the docker secrets required for each service based on the `*.env` files. There is no need to create the secrets yourself.
- Please see [Docker Commands Cheat Sheet](#) for a list of docker commands to use to manage the deployment.

Architectural Diagram of Docker Swarm



Upgrading Metric Insights to a newer version

To upgrade the deployment to a newer version of Metric Insights, simply generate a new config file using the installer for the new release, then redeploy using docker stack deploy.

1.20 Docker Commands Cheat Sheet

Here you'll find a list of **Docker** and **Docker Swarm** commands you can use to manage Metric Insights, whether it's a *simple deployment* or a *swarm* (orchestrated environment).

1. [Deploying](#)
2. [Getting Information](#)
3. [Getting into container](#)
4. [Getting Logs](#)
5. [Removing](#)

1.) Deploying

- Deploying the swarm:

```
docker stack deploy -c <deployment manifest>.yaml --with-registry-auth  
<user used for docker login to pull docker images>
```

2.) Getting Information

- Get list of Docker Services:

```
docker services ls
```

- Check processes for specific Docker Service:

```
docker service ps <service ID>
```

- Get more information about a Docker Service:

```
docker service inspect <service ID>
```

```
docker service inspect <service name>
```

- Get list of running Docker containers (Simple installation):

```
docker ps
```

```
docker ps -a
```


- Get more information about a Docker Container:

```
docker inspect <container ID>
```

```
docker inspect <container ID>
```

- Get list of Docker Volumes:

```
docker volume ls
```

- Get list of Docker Networks:

```
docker network ls
```

3.) Getting into container

- For Simple installation:

```
mi-console
```

```
docker exec -it <container name or ID> sh
```

- For Docker Swarm and AWS ECS (while being on ECS worker):

```
docker exec -it <container ID> bash
```

4.) Getting Logs

- See full process logs for specific Docker Service:

```
docker service ps --no-trunc <service ID>
```

- See Docker Container logs:

```
docker logs <container ID>
```

```
docker logs <container name>
```


5.) Removing

- Remove Docker Volume:

```
docker volume rm <volume ID>
```

```
docker volume rm <volume name>
```

- Remove specific Docker Service:

```
docker service rm <service ID>
```

```
docker service rm <service ID>
```

- Remove specific Docker Network:

```
docker network rm <network ID>
```

```
docker network rm <network name>
```

- Remove all Docker Services / entire Stack:

```
docker stack rm <stack name>
```


1.21 Creating Amazon ECS Scheduled Tasks for MI utilities

Amazon ECS supports the ability to schedule tasks on either a cron-like schedule or in a response to CloudWatch Events. This is supported for Amazon ECS tasks using both the Fargate and EC2 launch types.

If you have tasks to run on a schedule in your cluster, such as a backup operation or a log scan, you can use the Amazon ECS console to create a CloudWatch Events rule. Your scheduled event rule can be set to either a specific interval (run every N minutes, hours, or days), or for more complicated scheduling, you can use a cron expression.

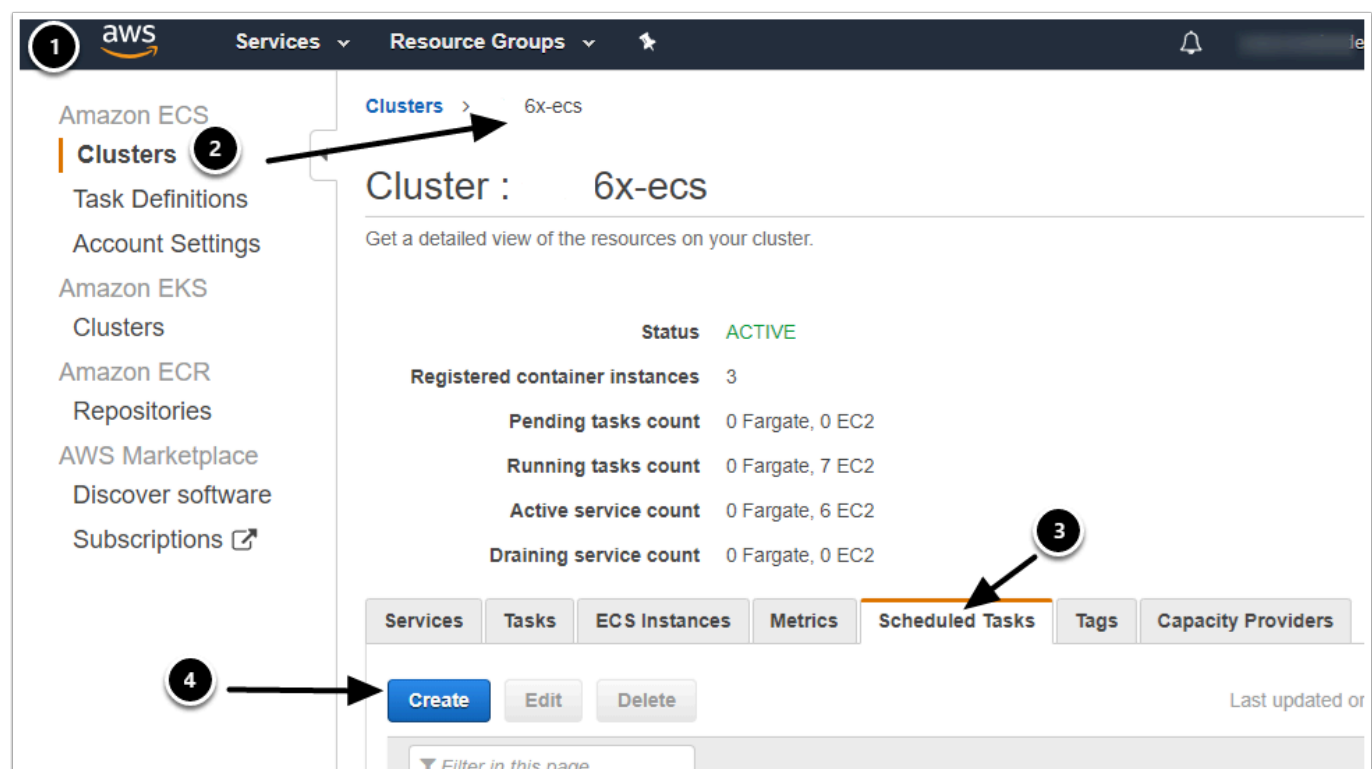
Metric Insights application has its own utilities to create backups of your database and all critical data of your application server (mi-app-backup tool), as well as synchronize users with the LDAP server (ldap-user-sync tool). For more information, please refer to:

[Backup Your Metric Insights Instance](#)

[Syncing Groups and Users from LDAP/AD using 'mi-ldap-usersync' script](#)

If it is not possible to set these 2 tools on the MI cron job, we strongly recommend to configure ECS Scheduled Tasks to run these utilities. Follow the below steps to set AWS ECS Scheduled Tasks.

Create scheduled task in cluster



1. Open Amazon ECS console <https://console.aws.amazon.com/ecs/>
2. Select the cluster you want to create a scheduled task in
3. On the cluster page click on **Scheduled Tasks Tab**
4. Click Create button

Configure scheduled task

Create scheduled task

Run Amazon ECS tasks on a cron-like schedule using CloudWatch Events rules and targets.

1

Schedule rule name*

mi-app-backup

i

Schedule rule enabled*

☒

i

Schedule rule description

i

2

Schedule rule type

☒ Run at fixed interval

☐ Cron expression

i

3

Run at fixed interval*

24

hour(s)

i

Specify the following:

1. *Schedule rule name* identify your task in scheduled tasks list
2. *Schedule rule type* choose whether to use a fixed interval schedule or a cron expression for your scheduled rule
3. *Run at fixed interval* indicate the interval your task should run at, or set cron expression if it was selected in the previous step. For more information, see [Schedule Expressions for Rules](#) in the Amazon CloudWatch Events User Guide.

Specify scheduled targets

Schedule targets

Create one or more task definition targets that run Amazon ECS tasks when the schedule rule is triggered.

Target name

mi-app-backup

1

Target id*

mi-app-backup

i

2

Launch type

☐ FARGATE
 ☒ EC2

i

3

Task Definition

Family

metricinsights-web

Enter a value

Revision

1 (latest)

Cluster*

6x-ecs

i

Number of tasks*

1

i

Current task role

None

i

Task role override

Select role...

i

1. *Target id* specify unique identifier for your target
2. Select *Launch type* - ECS2
3. *Task Definition Family* - choose the name of your web slave container from the list

Configure CloudWatch Events IAM role

CloudWatch Events IAM role

CloudWatch Events makes calls to Amazon ECS to run tasks on your behalf, and it uses this IAM role to get the appropriate permissions to do so. If you do not have the ecsEventsRole already, we can create one for you.

CloudWatch Events IAM role for this target

ecsEventsRole

i

1

Container overrides

web

2

3

Command override

mi-app-backup

Environment variable overrides

Key

Value

Add Environment Variable

target

4

Create

1. Leave `ecsEventsRole` for *CloudWatch Events IAM role for this target*
2. expand the *web* section under *Container overrides* header
3. *Command override*: specify the command that should be executed according to the schedule

Command Examples:

`mi-app-backup /some/dir` - saving tarball in a predefined folder (by default tarball created under `/opt/mi/backup`)

`mi-app-backup -k2` - keep only last two backups

`mi-ldap-usersync provision`

`"CN=some_group,OU=groups,OU=business_org,DC=global,DC=myorg,DC=com" --ldap-host="ldaps://ldap.internal.myorg.com" --ldap-`

`user="cn=miadmin,cn=users,dc=global,dc=myorg,dc=com" --ldap-`

`pass="password" --verbose` - basic provision command for syncing LDAP/AD groups and users in Metric Insights.

To run a custom script, please save the script in the directory `/opt/mi/custom/scripts` and indicate the full path to script in Command override field. For example: `./opt/mi/custom/scripts/MIAppBackup.sh`

4. After specifying the command, click Create.

Command line can contain optional parameters for running command. For more information on the options available, please refer to:

[Backup Your Metric Insights Instance](#)

[Syncing Groups and Users from LDAP/AD using 'mi-ldap-usersync' script](#)

1.22 Deploy Metric Insights on Kubernetes

Leveraging Kubernetes (K8) for highly scalable container orchestration is a great option. It's an open-source framework so while there are many Kubernetes providers (Azure, GCP, TKE, etc.), deployment is generally the same across said providers. In summary, initial deployment requires a deployment configuration file and docker images for each of the Metric Insights services. Let's walk through the process together below.

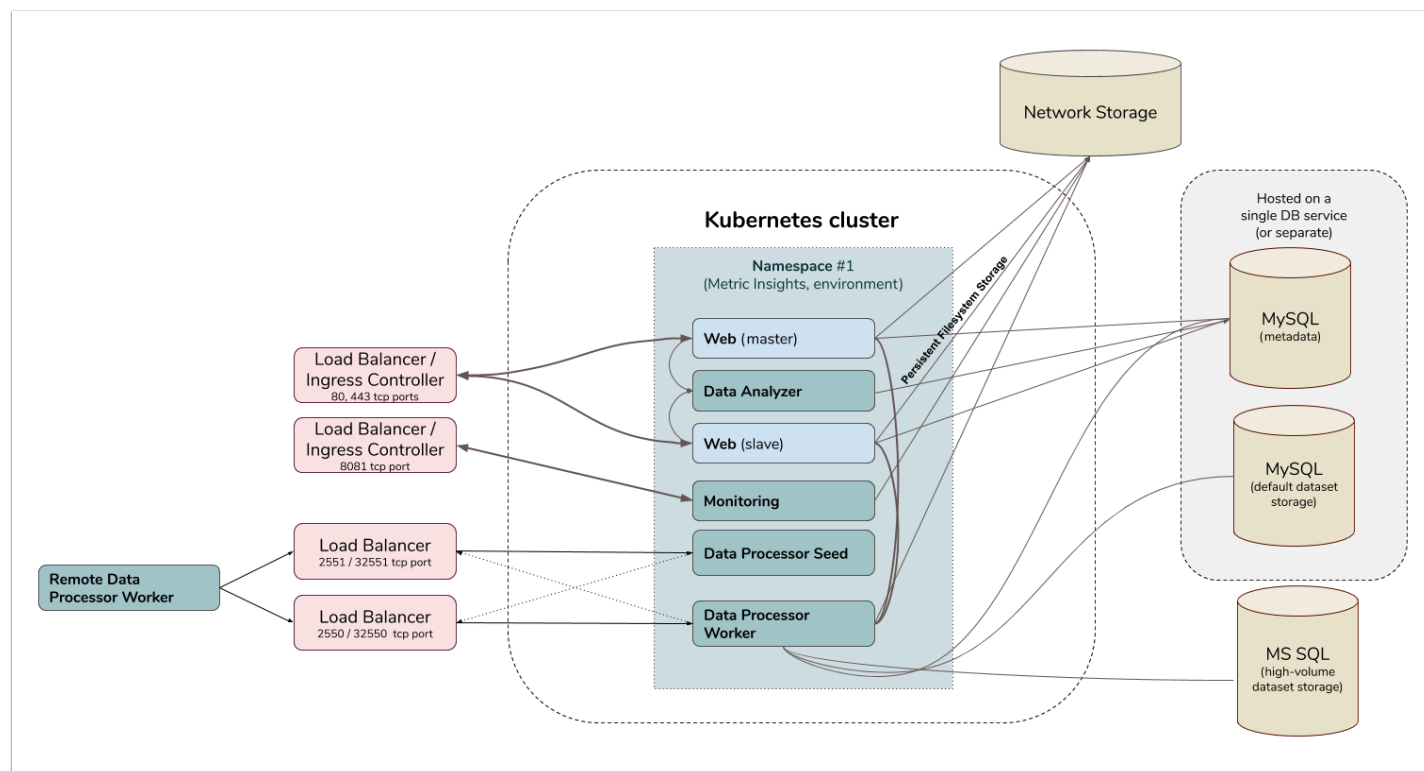
Deployment process:

1. [Understanding Application Architecture in Kubernetes](#)
2. [Prerequisites](#)
 1. [Select Storage Class for Persistent Volumes](#)
 2. [Choose Ingress Controller Type](#)
3. [Obtain Docker Registry Credentials](#)
4. [Create Secrets for Docker Registry](#)
5. [Create & Upload Secrets for Each MI Service](#)
6. [Deploy Load Balancers for External Communications to Data Processor & Seed](#)
7. [Create Deployment Configuration File to Deploy Metric Insights Application](#)
8. [Confirm Metric Insights is Deployed in Kubernetes Dashboard](#)
9. [Access Metric Insights Application in Browser](#)

For non-orchestrated environments, see the help article on using [Simple Installer](#).

1. Understanding Application Architecture in Kubernetes

Below is an architectural diagram of Metric Insights deployed in Kubernetes Namespace. A namespace is a virtual cluster that consists of several Nodes (servers). The nodes host Pods which is essentially a Container. Metric Insights consists of services that run inside their own container, rooted to a shared file system for persistence.



The deployment scheme consists of the following services deployed in individual pods (1 service per pod):

Required services:

1. Web Master
2. Web Slave
3. Data Analyzer
4. Data Processor
5. Seed
6. Monitoring

Optional services:

1. Remote Data Processor 1
2. Remote Data Processor 2

Additional items of note:

- A minimum of 3 Nodes are required for automatic failover
- MySQL is required to host the Metric Insights application database, and it should run on a remote server
- Persistent storage is required for the shared file system

2. Prerequisites

The following is required to deploy to Kubernetes:


- Access to the Kubernetes Dashboard (Web UI)
- Kubectl command-line tool to manage the Kubernetes Cluster
- Remote database server to host the application database (e.g., MySQL/MariaDB)
- Persistent shared storage (e.g., NFS, Portworx)

The following ports must also be open on the network:

- 80, 443 - HTTP/HTTPS ports for UI access
- 2550 - TCP port for the Data Processor cluster within the kubernetes namespace
- 2551 - TCP port for the Seed service within the kubernetes namespace
- 32550 - TCP port for external access to the Data Processor cluster
- 32551 - TCP port for external access to the Seed service
- 3306 - MySQL port
- 8080, 8443 - HTTP/HTTPS ports for the REST API Data Processor service
- 8081 - TCP port for the Monitoring Tool

Kubernetes namespace requirements:

1. Create a unique namespace for Metric Insights:
 - How to deploy a Kubernetes cluster: <https://kubernetes.io/docs/setup/>
 - How to create a unique namespace: <https://kubernetes.io/docs/tasks/administer-cluster/namespaces-walkthrough/#create-new-namespaces>
 - Azure AKS: <https://docs.microsoft.com/en-us/azure/aks/>
 - Amazon EKS: <https://aws.amazon.com/eks/>
2. Mount persistent storage volumes to the namespace. See [Select a Storage Class for Persistent Volumes](#) below for more.
3. Configure an ingress controller to allow incoming connections to the namespace (UI access in a browser). See [Choose an Ingress Controller type](#) below for more.

 The main services (web, data analyzer, data processor, seed, monitoring) can be installed separately depending on the needs of your deployment.

2.1. Select a Storage Class for Persistent Volumes

Metric Insights supports the following [Storage Classes](#) for the application file system, which is shared across pods as persistent volumes. Supported types are NFS and Portworx.

If using an NFS server hosted on Linux, please configure the mounted share as follows:

- In `/etc/exports`, set the mounted share to `/opt/mi` with the following options:
 - `rw,fsid=1,crossmnt,no_subtree_check,no_root_squash,async)`

If using Portworx, identify the exact class name to use for the deployment.

- The class name must also be enabled for multi-node volume attachment.

Running the installer to set the storage class for the deployment configuration file:

```
$ ./installer.py kubernetes --storage-class <nfs/portworx>
```

2.2. Choose an Ingress Controller type

Metric Insights supports the following [Ingress Controller Types](#) for incoming traffic from outside of the namespace. For example, users opening the application in a browser will be routed through the Ingress Controller. Supported types are:


- Nginx
- Traefik v1.7

Note, for Traefik v1.7, TCP ports are not supported. So for external communications to Seed and Data Processor, additional load balancers must be deployed. See [Deploy Load Balancers for external communications to Data Processor & Seed](#) for more.

Running the installer to set the ingress controller type for the deployment configuration file:

```
$ ./installer.py kubernetes --ingress-controller-type <nginx/traefik>
```

3. Obtain Docker Registry Credentials

 Contact MI Support for access to the official Metric Insights Docker Registry. Credentials are needed to pull docker images for each Metric Insights service.

- Note, the MI Docker Registry address (docker.metricinsights.com:5002) is specified in the deployment configuration file: `metricinsights-6.x.x.yml`.

4. Create Secrets for Docker Registry

- i** Before deploying to Kubernetes, Docker Registry credentials must be registered as Secret for K8 to reference. Metric Insights uses a secret labeled `docker-registry` to authenticate against a docker registry to pull docker images.
- Note, Secret is Kubernetes Object for storing and managing sensitive information like passwords and OAuth tokens. See [Kubernetes Secrets](#) for more.

Create Secret for Docker Registry credentials using kubectl:

```
kubectl --namespace <MI-namespace> create secret docker-registry <secret-name> --docker-server <docker-registry-url> --docker-username <username> --docker-password <password> --docker-email <email-address>
```

5. Create & Upload Secrets for Each MI Service

Metric Insights provides template files (environment variables) to create *secrets* for each service:

- web.env
- dataprocessor.env
- seed.env
- mysql.secret
- data_analyzer.env
- monitoring.env

Templates for each are available in the installer directory: `MetricInsights-Installer-v6.x.x-Full/utils/orchestration/kubernetes/secrets/`

1. Copy each template file from the installer directory to your working directory
2. The template files are saved with the extension .example (e.g., web.env.example). Rename each file by removing the .example extension.
3. Open each file in an editor and update all fields. Ensure passwords are consistent between files.
 - You can encode passwords of our choice or create random passwords.
 - To encode, run `echo -n '<password>' | base64`
 - To generate a random password, run `openssl rand -base64 8 | tr -d /+= | cut -c -11`
4. Create the secrets for each service by uploading each file to the namespace using kubectl:

```
$ kubectl --namespace <MI-namespace> create secret generic metricinsights-web --from-
```



```

file web.env
$ kubectl --namespace <MI-namespace> create secret generic metricinsights-dataprocessor
--from-file dataprocessor.env
$ kubectl --namespace <MI-namespace> create secret generic metricinsights-seed --from-
file seed.env
$ kubectl --namespace <MI-namespace> create secret generic metricinsights-mysql-root-
password --from-file mysql.secret
$ kubectl --namespace <MI-namespace> create secret generic metricinsights-data-analyzer
--from-file data-analyzer.env
$ kubectl --namespace <MI-namespace> create secret generic metricinsights-monitoring --
from-file monitoring.env

```

6. Deploy Load Balancers for External Communications to Data Processor & Seed

This step is critical if you plan on integrating with BI Tools that require a Remote Data Processor to run on a Windows machine (e.g., Power BI, Qlik Sense, Tibco Spotfire, etc.)

For the Remote Data Processor to communicate with the Local Data Processor in Kubernetes, deploy additional Load Balancers to handle external TCP traffic. Make sure the following ports are open on the load balancers:

- 32550 (data processor)
- 32551 (seed)

Note, additional load balancers are needed because K8 Ingress Controller types like Traefik only handle HTTP/HTTPS traffic and not TCP. Comparatively, for an ingress controller like Nginx, you can use a combination of the ingress controller + nodeports for external communications, or deploy load balancers as their own service (separate pods).

For the public hostname, we recommend using the hostname of the K8 master node and adding "dataprocessor" or "seed" to the name to keep things simple. For example, if the hostname of the master node (the hostname by which users will access the UI in a browser) is customName.companyName.com then you can map the following for the Data Processor LB: customName-**dataprocessor**.companyName.com.

Once the load balancers are deployed and a public hostname is set, update the following parameters in either the Data Processor secrets file (**dataprocessor.env**) or in the K8 deployment manifest file (**metricinsights-6.x.x.yml**):

- DATAPROCESSOR_HOSTNAME
- DATAPROCESSOR_SEED_HOSTNAME

Here's an example for the deployment configuration file:

...


```
env:
  - name: SPRING_PROFILES_ACTIVE
    value: "singleton,swagger,admin"
  - name: DATAPROCESSOR_HOSTNAME
    value: "metricinsights-dataprocessor"
  - name: DATAPROCESSOR_SEED_HOSTNAME
    value: "metricinsights-seed"
  - name: DATAPROCESSOR_SEED_BIND_PORT
    value: "2551"
...
```


7. Create Deployment Configuration File to Deploy Metric Insights Application

To create a Kubernetes Deployment configuration file (also called a manifest), use the MI installer to generate a yaml file. Here's an example where we are setting the following values:

- Storage Type = NFS
- Ingress Controller = Nginx
- Private Docker Registry
- Data Processor Hostname

```
$ ./installer.py kubernetes --storage-class nfs --nfs-server-address <nfs.example.com> -
-ingress-controller-type nginx --hostname <MI-hostname> --dp-hostname
<dataprocessor_hostname> --registry <registry-url> -o metricinsights-6.1.3.yml
```

The key here is using the **-o** option for the output file and then specifying a file name with a *.yaml extension (yaml file).

 See `./installer.py kubernetes -h` for more options.

Next, deploy the Metric Insights application to the K8 namespace using the newly created deployment configuration file:

```
$ kubectl --namespace <MI-namespace> apply -f metricinsights-6.1.3.yml
```

You should see the services and pods being created as soon as the yaml file is applied. The kubectl apply output should look something like this:

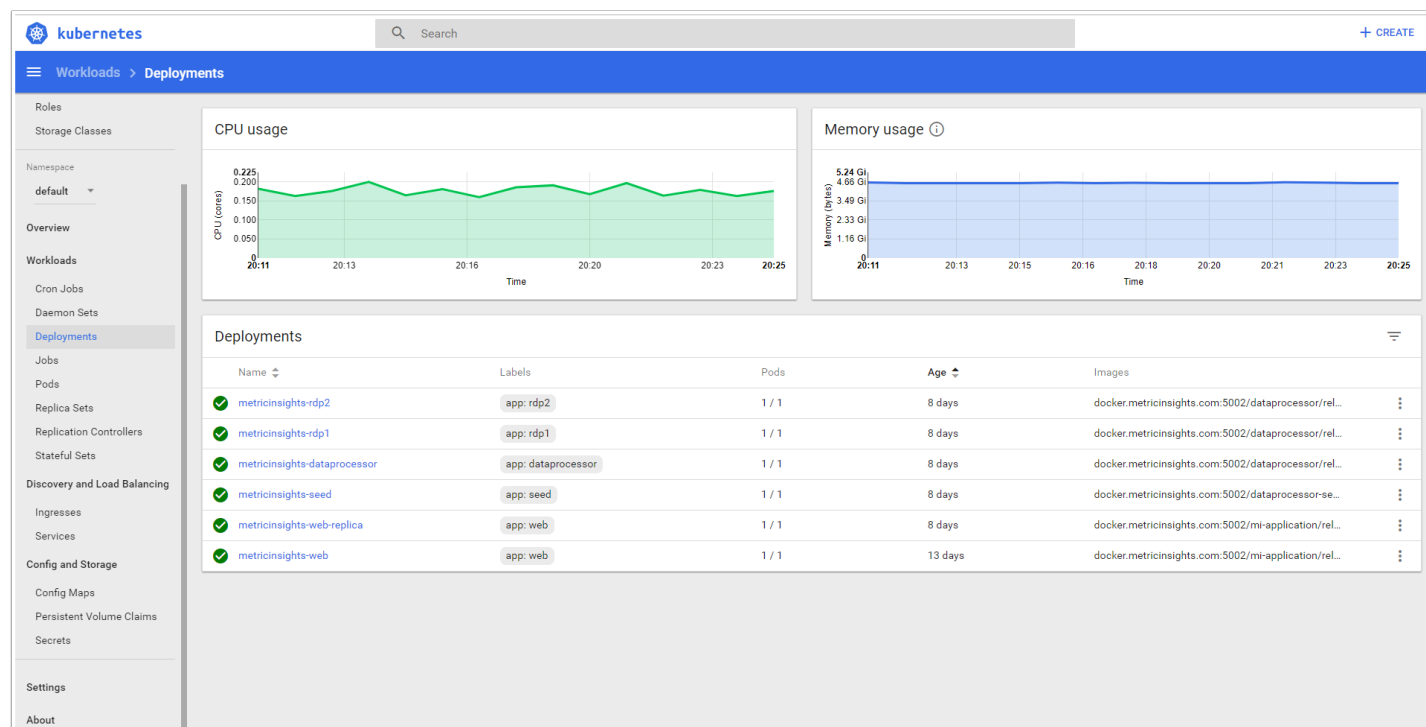
```
$ kubectl --namespace <MI-namespace> apply -f metricinsights-6.1.3.yml
service/metricinsights-web created
deployment.apps/metricinsights-web-master created
deployment.apps/metricinsights-web-slave created
```



```
service/metricinsights-seed created
deployment.apps/metricinsights-seed created
service/metricinsights-dataprocessor created
service/metricinsights-data-analyzer created
deployment.apps/metricinsights-data-analyzer created
persistentvolume/metricinsights-v613-data created
persistentvolumeclaim/metricinsights-v613-data created
persistentvolume/metricinsights-v613-ssl created
persistentvolumeclaim/metricinsights-v613-ssl created
persistentvolume/metricinsights-v613-external-libraries created
persistentvolumeclaim/metricinsights-v613-external-libraries created
ingress.extensions/metricinsights-ingress-nginx created
.
.
.
```

8. Confirm Metric Insights is Deployed in Kubernetes Dashboard

Once the deployment is complete, log into the [Kubernetes Dashboard](#) to confirm Metric Insights is available. See the image below as a reference.



9. Access Metric Insights Application in Browser

Now that Metric Insights is up and running in Kubernetes, the web UI is accessible through the requested Hostname or IP Address:

- The hostname and ip address can be for the Ingress Controller or the Kubernetes worker node + port 80/443 on which the Web pod is running.

 To learn how to administer Metric Insights in Kubernetes, see [Administering Metric Insights in Kubernetes](#)

1.23 Kubectl Commands Cheat Sheet

Here you'll find a list of **kubectl** commands you can use to manage Metric Insights if deployed in Kubernetes.

1. [1.\) Deploying](#)
2. [2.\) Getting Information](#)
3. [3.\) Getting into Pods with containers](#)
4. [4.\) Getting Logs](#)
5. [5.\) Deleting](#)

1.) Deploying

- Deploy Metric Insights in the namespace using a configuration file:

```
kubectl --namespace <MI-namespace> apply -f <configuration-file>.yaml
```

2.) Getting Information

- Get list of all services in the namespace:

```
kubectl get services -n <MI-namespace>
```

- Get list of pods in the namespace:

```
kubectl get pods -n <MI-namespace>
```

```
kubectl get pods -n <MI-namespace> -o wide
```

```
kubectl get pods --all-namespaces
```

```
kubectl get pods --sort-by='.status.containerStatuses[0].restartCount'
```

- Get detailed information about each pod:

```
kubectl describe pods -n <MI-namespace> -o wide
```

- Get list of events sorted by timestamp:

```
kubectl get events --sort-by=.metadata.creationTimestamp
```


3.) Opening a Bash Shell inside a Pod

- Open a bash shell for a specific pod:

```
kubectl exec -it <pod-ID> -n <MI-namespace> bash
```

4.) Getting Logs

- See logs for a specific pod:

```
kubectl logs <pod-ID>
```

```
kubectl logs -f <pod-ID>
```

5.) Deleting

- Deleting pods and services (to generate a new pod):

```
kubectl -n <MI-namespace> delete pod <pod-ID>
```

```
kubectl -n <MI-namespace> delete svc <service-ID>
```

```
kubectl -n <MI-namespace> delete pod/svc --all
```

- Deleting all resources in the namespace

```
kubectl -n <MI-namespace> delete -f <configuration-file>.yaml
```


1.24 Administer Metric Insights in Kubernetes

There are two primary options for administering the Metric Insights Kubernetes cluster:

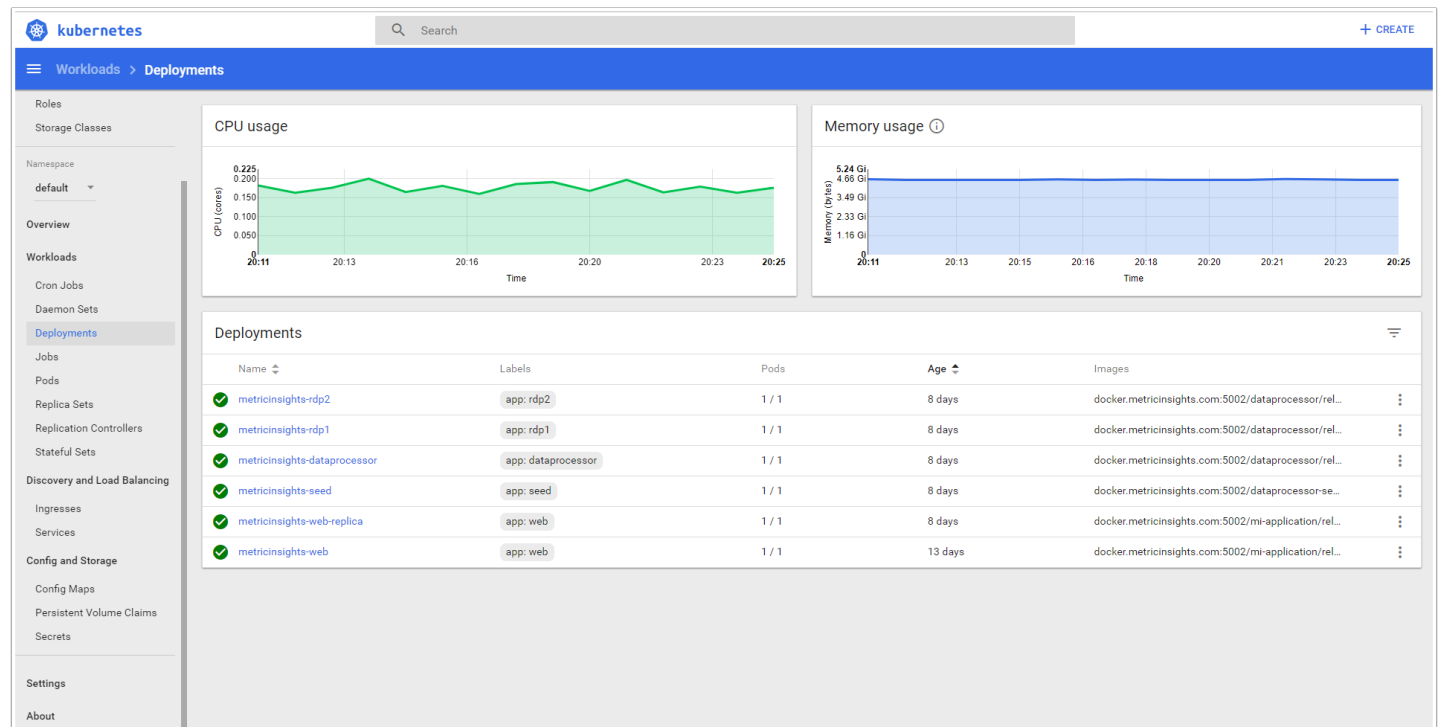
- Kubernetes Dashboard (a web-based user interface)
- Kubectl (command-line tool):
 - Can be installed in both Linux and Windows environments
 - Must be able to reach the Kubernetes cluster

This article covers the following:

1. [Kubernetes Dashboard](#)
 1. [Accessing Metric Insights Deployment from Kubernetes Dashboard](#)
 2. [Accessing Container Logs in Kubernetes](#)
 3. [Accessing Container Logs via Splunk](#)
2. [Kubectl](#)

1. Kubernetes Dashboard

Accessing Metric Insights Deployment from Kubernetes Dashboard



In the image above, you can see the list of Metric Insights services under the Deployments menu option. One service corresponds to one pod in the namespace. Let's take a closer look at the pods below.

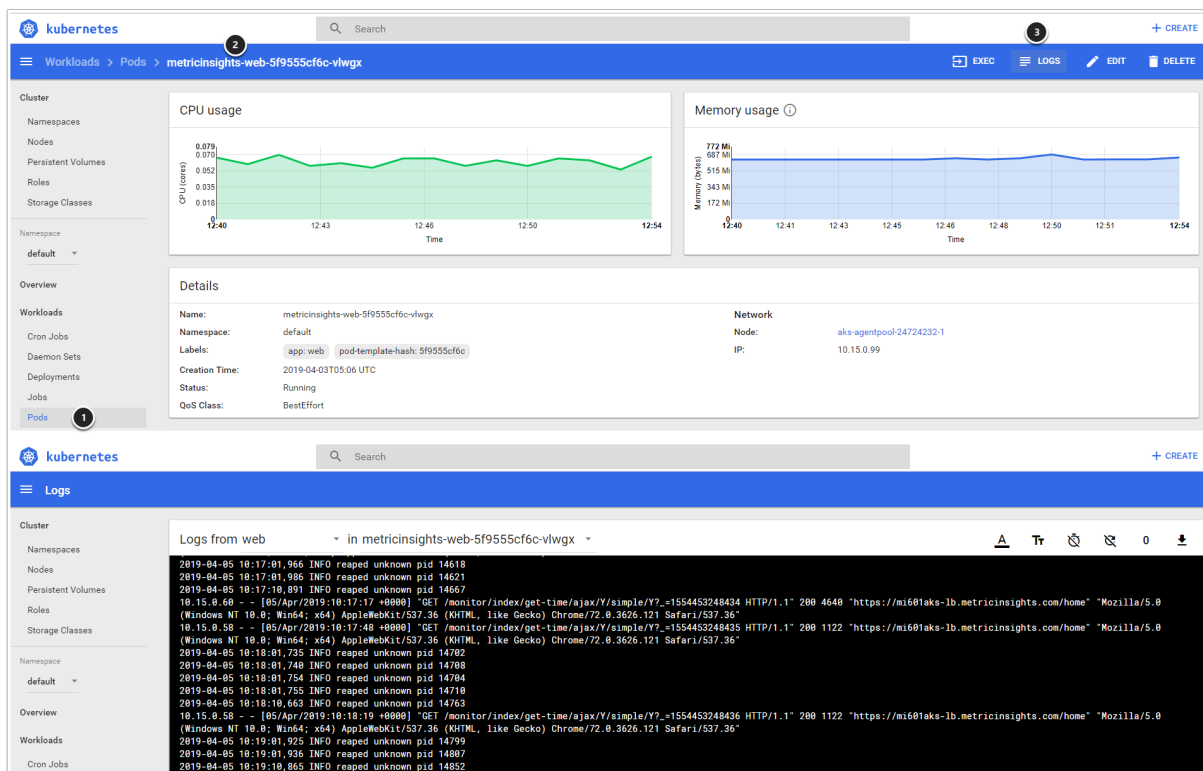
The image displays three sequential screenshots of the Kubernetes dashboard interface, illustrating the process of accessing a pod's shell console.

- Top Screenshot:** Shows the 'Pods' tab selected in the left-hand navigation menu. The main area displays a table of pods. A red circle highlights the 'EXEC' button in the top right corner of the pod list table.
- Middle Screenshot:** Shows the 'Details' view for a specific pod. The 'EXEC' button is highlighted with a red circle, and a red arrow points from this button to the 'EXEC' button in the top screenshot.
- Bottom Screenshot:** Shows the 'Shell' view, which is a terminal window where commands can be executed within the pod's container. The terminal shows the prompt 'root@metricinsights-web-5f955cf6c-vlwgx:~#'.

A pod is a unit of deployment, and in the case of Metric Insights, it equals 1 application service (packaged inside a container). Accessing the Pods menu allows you to get information about each pod as well as enter the container to view the file system, run commands, and get logging.

1. Go to Navigation Menu > Pods
2. Select Pod to open Pod's Detail Page
3. Click [EXEC] to open the shell console
4. Use the shell console to view the file system and run commands

Accessing Container Logs in Kubernetes



By default, each of the Metric Insights Services prints its own logs to stdout.

To view the logs:

1. Go to the Navigation Menu > Pods
2. Select Pod to open Pod's Detail Page
3. Click [LOGS]
4. Logs generally show the boot process of the container and the running application service
 - e.g., for the Web service: apache logs, https requests, and other provisioning info
 - e.g., for the Data Processor service: logs from the Java application

i You can also enter the container to view the application logs directly. The logs can also be downloaded from the Metric Insights UI.

Accessing Container Logs via Splunk

Instead of using the default logging service, you can use the Splunk logging driver to send container logs to Splunk Enterprise and Cloud. The logs are sent to Splunk HTTP Event Collector.

- See [Logging via Splunk \(6.x\)](#) on how to configure it.

2. Kubectl

To use kubectl, you must install the tool on a machine that can access the Kubernetes cluster. The tool can be installed on Linux, Windows, and MacOS. See [Install Kubectl](#) for steps on how to install and configure the tool.

As for using the tool itself, here's an example of a kubectl command to get a list of pods for a namespace named MI:

```
# kubectl --namespace mi get pods
```

For more useful kubectl commands to manage Metric Insights, see [Kubectl Commands Cheat Sheet](#).

1.25 Deploy Metric Insights on OpenShift

OpenShift is Red Hat's version of Kubernetes. It is a viable option for container orchestration, especially if your organization already has a Red Hat subscription.

Deploying to OpenShift requires a *deployment configuration file* and *docker images* for each of the Metric Insights services. Let's walk through the process below.

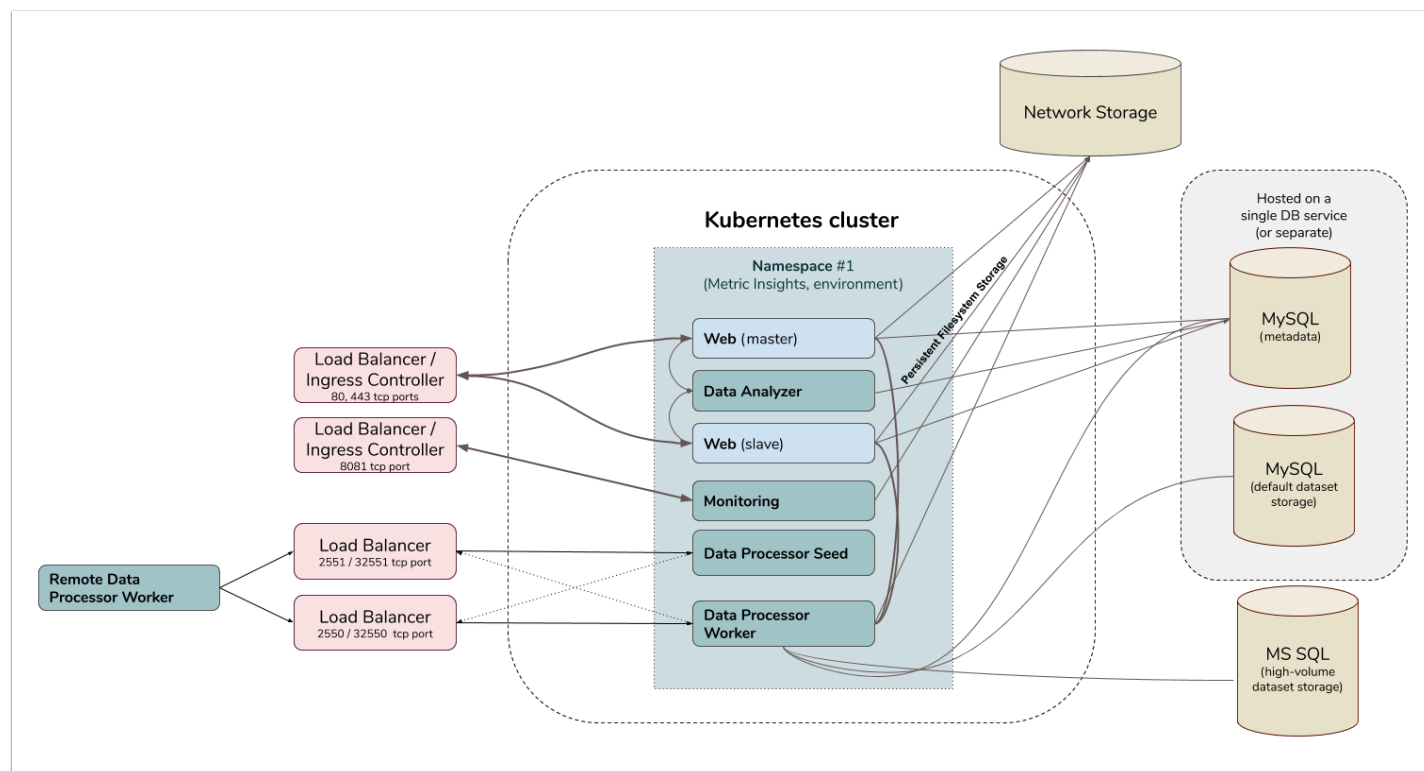
Deployment process:

1. [Understanding Application Architecture](#)
2. [Prerequisites](#)
 1. [Storage Class for Persistent Volumes](#)
3. [Obtain Docker Registry Credentials](#)
4. [Create Secret for Docker Registry](#)
5. [Create Secrets for Each Metric Insights Service](#)
6. [Create Configuration File and Deployment of Metric Insights Application](#)
7. [Create Routes](#)
8. [Access Metric Insights Application in Browser](#)

For non-orchestrated environments, see the help article on using [Simple Installer](#).

1. Understanding Application Architecture

Below is an architectural diagram of Metric Insights deployed in Kubernetes Namespace. A namespace is a virtual cluster that consists of several Nodes (servers). The nodes host Pods which is essentially a Container. Metric Insights consists of services that run inside their own container, rooted to a shared file system for persistence.



The deployment scheme consists of the following services deployed in individual pods (1 service per pod):

Required services:

1. Web Master
2. Data Analyzer
3. Data Processor
4. Seed
5. Monitoring

Optional services:

1. Web Slave
2. Remote Data Processor 1
3. Remote Data Processor 2

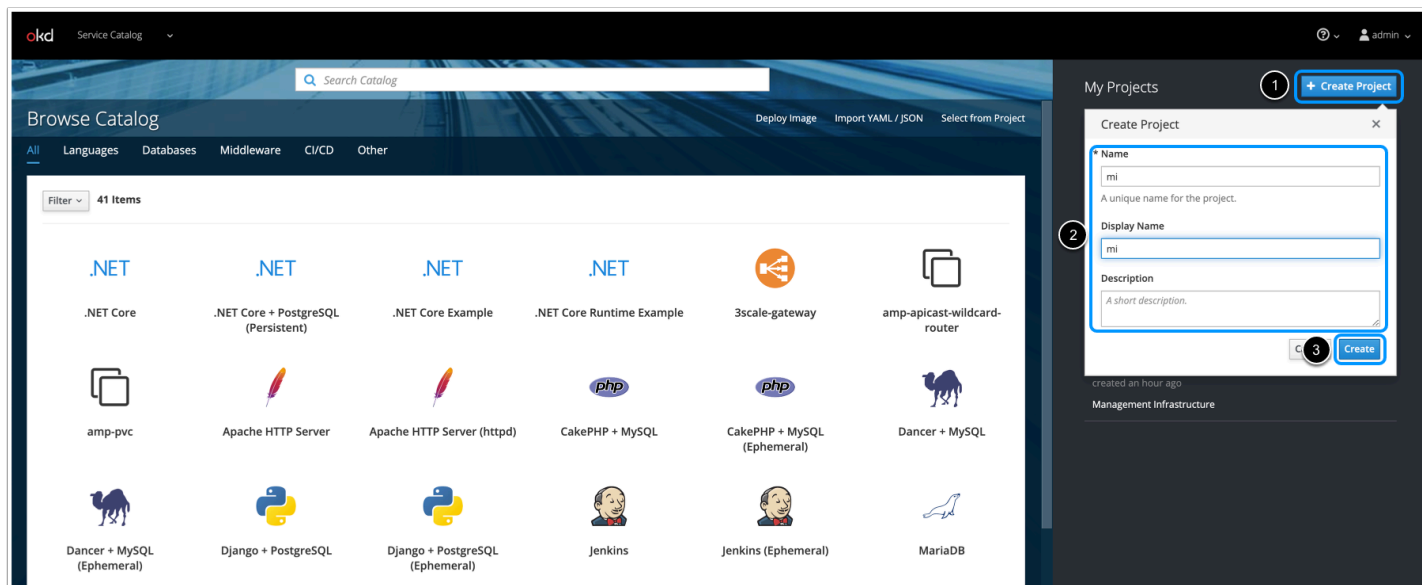
Additional items of note:

- A minimum of 3 Nodes are required for automatic failover
- MySQL is required to host the Metric Insights application database, and it should run on a remote server
- Persistent storage is required for the shared file system

2. Prerequisites

The following is required to deploy to OpenShift:

- Access to the OpenShift Dashboard (Web UI)
- OpenShift Container Platform command line interface (CLI)
- Remote database server to host the application database (e.g., MySQL 8)
- Persistent shared storage (e.g., NFS)
- OpenShift Project for Metric Insights application:
to create new project in OpenShift Dashboard click on **Create Project** (1) → fill text boxes (2) → click **Create**



2.1. Storage Class for Persistent Volumes

Openshift deployment supports only NFS type for the application file system, which is shared across pods as persistent volumes.

For NFS server hosted on Linux, please configure the mounted share as follows:

- In /etc/exports, set the mounted share to /opt/mi with the following options:
 - `rw,fsid=1,crossmnt,no_subtree_check,no_root_squash,async)`

3. Obtain Docker Registry Credentials

! Contact Metric Insights Support for access to the official Metric Insights Docker Registry. Credentials are needed to pull docker images for each Metric Insights service.

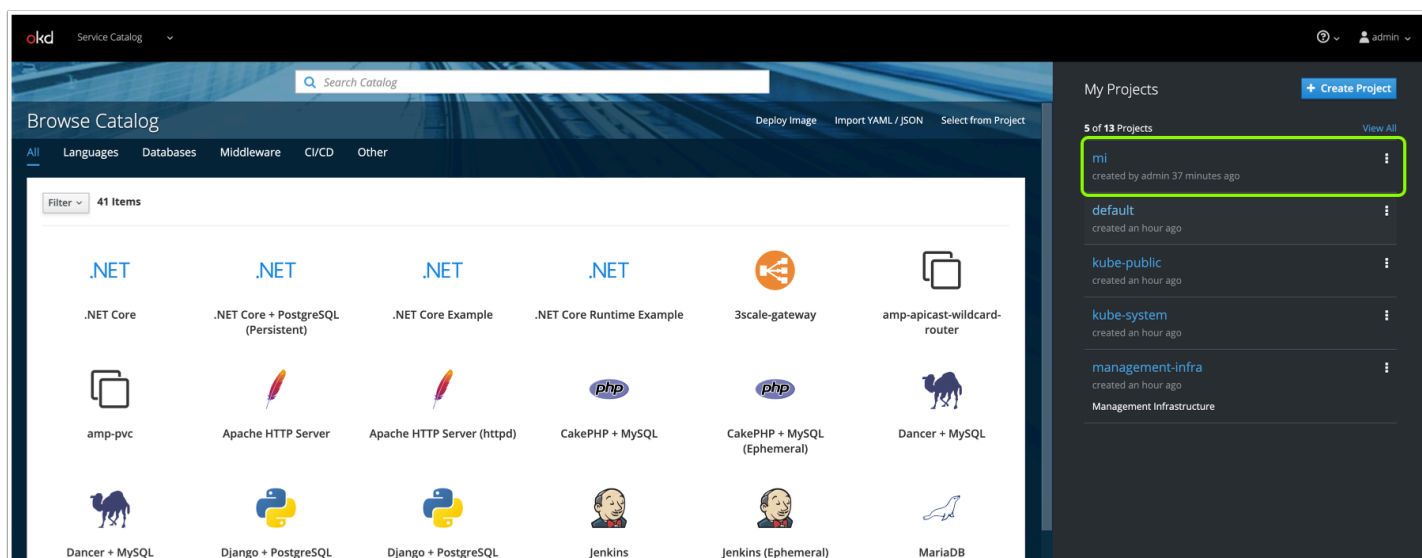
- Note, the Metric Insights Docker Registry address (docker.metricinsights.com:5002) is specified in the deployment configuration file.

4. Create Secret for Docker Registry

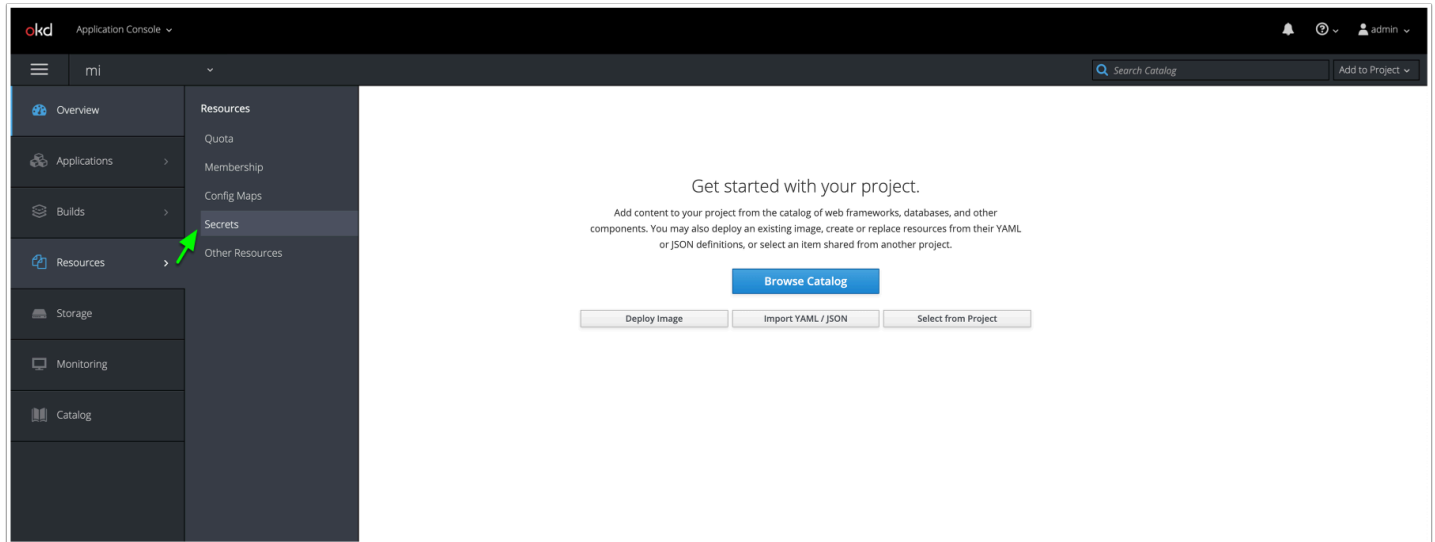
- i Before deploying to OpenShift, Docker Registry credentials must be registered as a Secret for K8 to reference. Metric Insights uses a secret labeled `docker-registry` to authenticate against a docker registry to pull docker images.
- Note, a Secret is an object for storing and managing sensitive information like passwords and OAuth tokens. See Kubernetes/OpenShift Secrets for more.

Create a Secret for Docker Registry credentials using the OpenShift Dashboard:

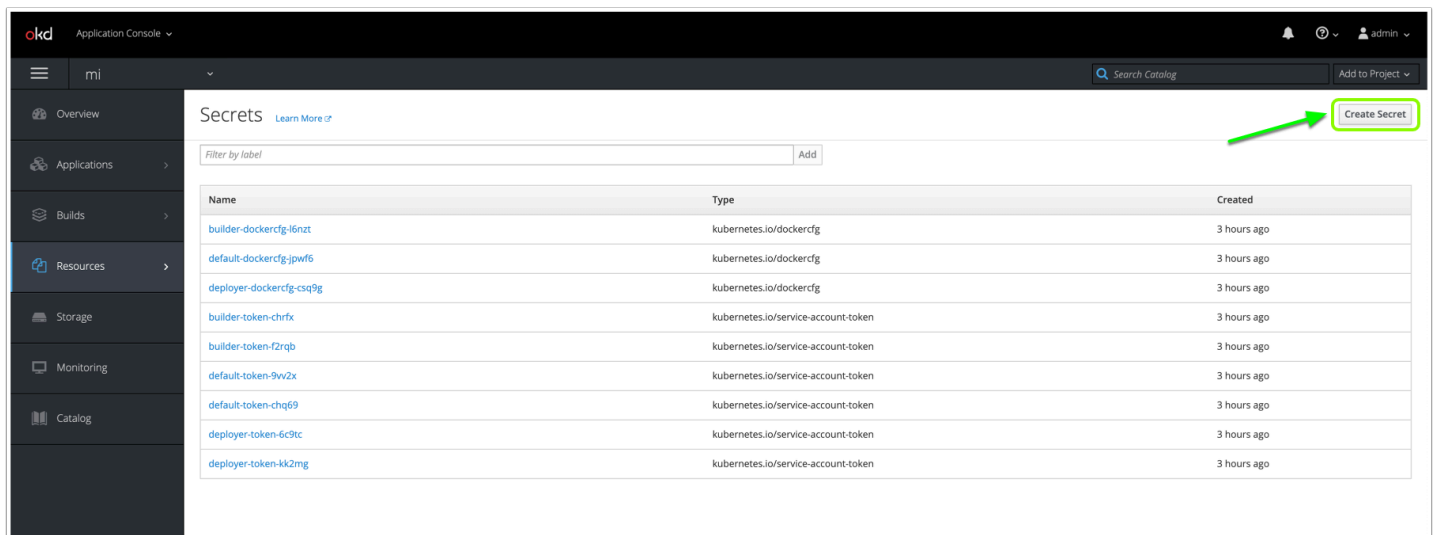
- Select the Metric Insights project:



- Navigate to **Resources** → **Secrets**:



- Click on **Create Secret** button:



- Edit fields with the following values and click on **Create** button:

5. Create Secrets for Each Metric Insights Service

Metric Insights provides template files (environment variables) to create *secrets* for each service:

- web.env
- dataprocessor.env
- seed.env
- mysql.secret
- data_analyzer.env
- monitoring.env

Templates for each are available in the installer directory: `MetricInsights-Installer-v6.x.x-Full/utils/orchestration/kubernetes/secrets/`

1. Copy each template file from the installer directory to your working directory
2. The template files are saved with the extension `.j2` (e.g., `web.env.j2`). Rename each file by removing the `.j2` extension.
3. Open each file in an editor and update all fields. Ensure passwords are consistent between files.
 - You can encode passwords of our choice or create random passwords.
 - To encode, run `echo -n '<password>' | base64`
 - To generate a random password, run `openssl rand -base64 8 | tr -d /+= | cut -c -11`
4. Create the secrets for each service with **Secret Type: Generic Secret** by uploading each file to the form (1) or copy-paste content to the appropriate field (2).
Here's an example for **data-analyzer.env**:

okd Application Console

mi

Search Catalog Add to Project

Overview Applications Builds Resources Storage Monitoring Catalog

Secrets > Create Secret

Create Secret

Secrets allow you to authenticate to a private Git repository or a private image registry.

Secret Type
Generic Secret

*** Secret Name**
metricsights-data-analyzer
Unique name of the new secret.

*** Key**
data-analyzer.env
A unique key for this secret entry.

Value
Enter a value for the secret entry or use the contents of a file. [Browse...](#)

[Clear Value](#)

```

1 MYSQL_HOSTNAME=192.168.33.10
2 MYSQL_PORT=3306
3 MYSQL_DATABASE=dashboard
4 MYSQL_PASSWORD=FUS0KMHr3hqk9Nt5L
5 MYSQL_USER=mi_read
6 MONITORING_ACCESS_TOKEN=ZKPSH8stueSchBp2G

```

[Remove Item](#) | [Add Item](#)

[Create](#) [Cancel](#)

Example for **dataprocessor.env**:

okd Application Console

mi

Search Catalog Add to Project

Overview Applications Builds Resources Storage Monitoring Catalog

Secrets > Create Secret

Create Secret

Secrets allow you to authenticate to a private Git repository or a private image registry.

Secret Type
Generic Secret

*** Secret Name**
metricsights-dataprocessor
Unique name of the new secret.

*** Key**
dataprocessor.env
A unique key for this secret entry.

Value
Enter a value for the secret entry or use the contents of a file. [Browse...](#)

[Clear Value](#)

```

1 MYSQL_HOSTNAME=192.168.33.10
2 MYSQL_PORT=3306
3 MYSQL_DATABASE=dashboard
4 MYSQL_DATAPROCESSOR_PASSWORD=wp18ik821X5g7Qh9V
5 MYSQL_USERNAME=dataprocessor
6 DP_PASSWORD=u8kz5M67HPS2wfh1W
7 MONITORING_ACCESS_TOKEN=ZKPSH8stueSchBp2G

```

[Remove Item](#) | [Add Item](#)

[Create](#) [Cancel](#)

Example for **monitoring.env**:

okd Application Console

mi

Search Catalog Add to Project

Overview Applications Builds Resources Storage Monitoring Catalog

Secrets > Create Secret

Create Secret

Secrets allow you to authenticate to a private Git repository or a private image registry.

Secret Type
Generic Secret

*** Secret Name**
metricsights-monitoring
Unique name of the new secret.

*** Key**
monitoring.env
A unique key for this secret entry.

Value
Enter a value for the secret entry or use the contents of a file.
Clear Value

```
1 MONITORING_READ_TOKEN=66pf4vazvlgmW7C
2 MONITORING_WRITE_TOKEN=ZKPSH8sive5ch8p2G
3 MONITORING_USERNAME=admin
4 MONITORING_PASSWORD=r10b0PtmdjKT2T11C
```

Remove Item Add Item

Create Cancel

Example for **mysql.secret**:

okd Application Console

mi

Search Catalog Add to Project

Overview Applications Builds Resources Storage Monitoring Catalog

Secrets > Create Secret

Create Secret

Secrets allow you to authenticate to a private Git repository or a private image registry.

Secret Type
Generic Secret

*** Secret Name**
metricsights-mysql-root-password
Unique name of the new secret.

*** Key**
mysql.secret
A unique key for this secret entry.

Value
Enter a value for the secret entry or use the contents of a file.
Clear Value

```
1 passwd
```

Remove Item Add Item

Create Cancel

Example for **seed.env**:

okd Application Console

mi

Search Catalog Add to Project

Overview Applications Builds Resources Storage Monitoring Catalog

Secrets > Create Secret

Create Secret

Secrets allow you to authenticate to a private Git repository or a private image registry.

Secret Type
Generic Secret

*** Secret Name**
metricsights-seed
Unique name of the new secret.

*** Key**
seed.env
A unique key for this secret entry.

Value
Enter a value for the secret entry or use the contents of a file.
Browse...

Clear Value

```
1 MONITORING_ACCESS_TOKEN=ZKPSH8s1ue5ch8p2G
```

Remove Item Add Item

Create Cancel

Example for **web.env**:

okd Application Console

mi

Search Catalog Add to Project

Overview Applications Builds Resources Storage Monitoring Catalog

Secrets > Create Secret

Create Secret

Secrets allow you to authenticate to a private Git repository or a private image registry.

Secret Type
Generic Secret

*** Secret Name**
metricsights-web
Unique name of the new secret.

*** Key**
web.env
A unique key for this secret entry.

Value
Enter a value for the secret entry or use the contents of a file.
Browse...

Clear Value

```
1 MYSQL_HOSTNAME=192.168.33.10
2 MYSQL_PORT=3306
3 MYSQL_DATABASE=dashboard
4 MYSQL_DATASET_HOSTNAME=192.168.33.10
5 MYSQL_DATASET_PORT=3306
6 MYSQL_ROOT_USER=root
7 MYSQL_DATASET_READ_PASSWORD=VQdGL6Jcx1Re0Bw4H
8 MYSQL_GENERATOR_PASSWORD=659AnL0n1FgMek00
9 MYSQL_READ_ONLY_PASSWORD=FUSOKMhr3hqk9Nt5L
10 MYSQL_SETUP_PASSWORD=xJztU9i3PeE63b3c6Z
11 MYSQL_WEB_PASSWORD=gdeso11NX94vePt7P
12 MYSQL_DATAPROCESSOR_PASSWORD=wp181k821Xsg7Qh9V
```

Remove Item Add Item

Create Cancel

Note, you can also generate secret files and the application deployment manifest using the installer wizard option.

Secret files are saved to a separate folder, which must be created before running the installer wizard:

```
$ ./installer.py kubernetes --wizard
```


Installer wizard example:

```

nicobillon@ip-192-168-1-105 MetricInsights-Installer-v6.2.1-Full % mkdir openshift
nicobillon@ip-192-168-1-105 MetricInsights-Installer-v6.2.1-Full % ./installer.py kubernetes --wizard

WELCOME to Metric Insights Kubernetes Config Generator Wizard!

ATTENTION! Leave empty to use default value.

1. Enter output folder for target artifacts. [ Default: ./ ] ./openshift
2. Set scenario type for kubernetes. Possible values: openshift, k8s [ Default: k8s ] openshift
3. Enter docker registry address. [ Default: docker.metricinsights.com:5002 ]
4. Enter login to get access to registry docker.metricinsights.com:5002. [ Default: installer ]
5. Enter email address to get access to registry docker.metricinsights.com:5002. [ Default: your@email.com ]
6. Enter password to get access to registry docker.metricinsights.com:5002.
7. Enter docker image pull secret name. metricinsights-docker-registry
8. Enter desired name for namespace. [ Default: default ] mi
9. Enter desired number of web replicas. [ Default: 1 ]
10. Set hostname of Dataprocessor. Default values: metricinsights-dataprocessor [ Default: metricinsights-dataprocessor ]
11. Data Processor seed service hostname for external services. Default values: metricinsights-seed [ Default: metricinsights-seed ]
12. Enter ingress controller type. Possible values: nginx, traefik [ Default: None ]
13. Enable specification for shared drive from host to web container. Type "Yes" to enabled [ Default: None ]
14. Enter storage class type. Possible values: nfs, portworx [ Default: nfs ]
15. Enter NFS server address. Example: 192.168.34.10, 192.168.33.10
16. Enter NFS shared folder. Example: /opt/mi/data. [ Default: /opt/mi/data ]
17. Enter service type. Possible values: ClusterIP, LoadBalancer, NodePort [ Default: LoadBalancer ]
18. Enter load balancer type. Possible values: azure-internal, external [ Default: external ]
19. Enter MySQL hostname address. 192.168.33.10
20. Enter MySQL hostname port. [ Default: 3306 ]
21. Enter MySQL dataset hostname address. [ Default: 192.168.33.10 ]
22. Enter MySQL dataset hostname port. [ Default: 3306 ]
23. Enter MySQL root username. [ Default: root ]
24. Enter MySQL root password.
2020-09-01 15:02:15,824 [INFO] Generated OpenShift config file of Metric Insights Metric Insights, v6.2.1 and secrets
  
```


6. Create Configuration File and Deployment of Metric Insights Application

To create a deployment configuration file (also called a manifest), use the Metric Insights installer to generate a yaml file. Here's an example where we are setting the following values:

- Kubernetes Type = openshift
- Storage Type = NFS
- NFS server address = 192.168.33.10
- Private Docker Registry

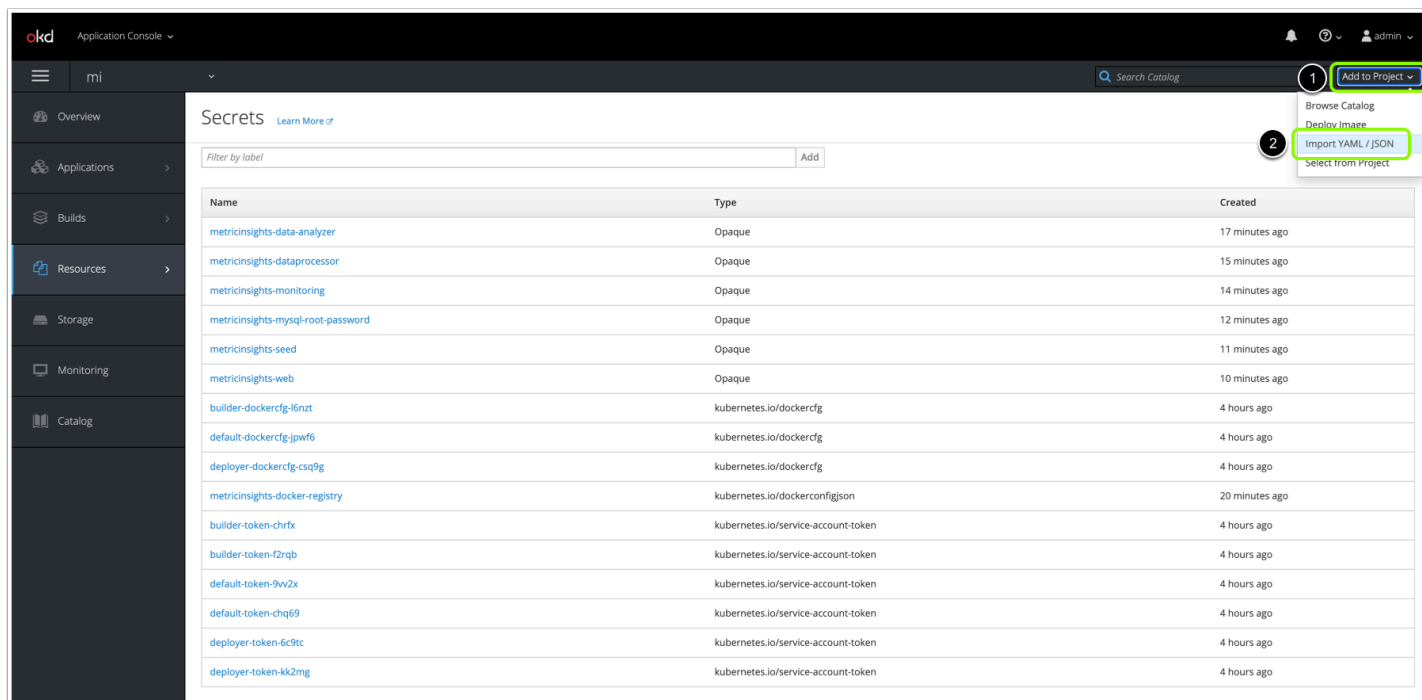
```
$ ./installer.py kubernetes --type openshift --nfs-server-address 192.168.33.10 --
images-pull-secret-name metricinsights-docker-registry -o application.yml
```

The key here is using the **-o** option for the output file and then specifying a file name with a *.yaml extension (yaml file).

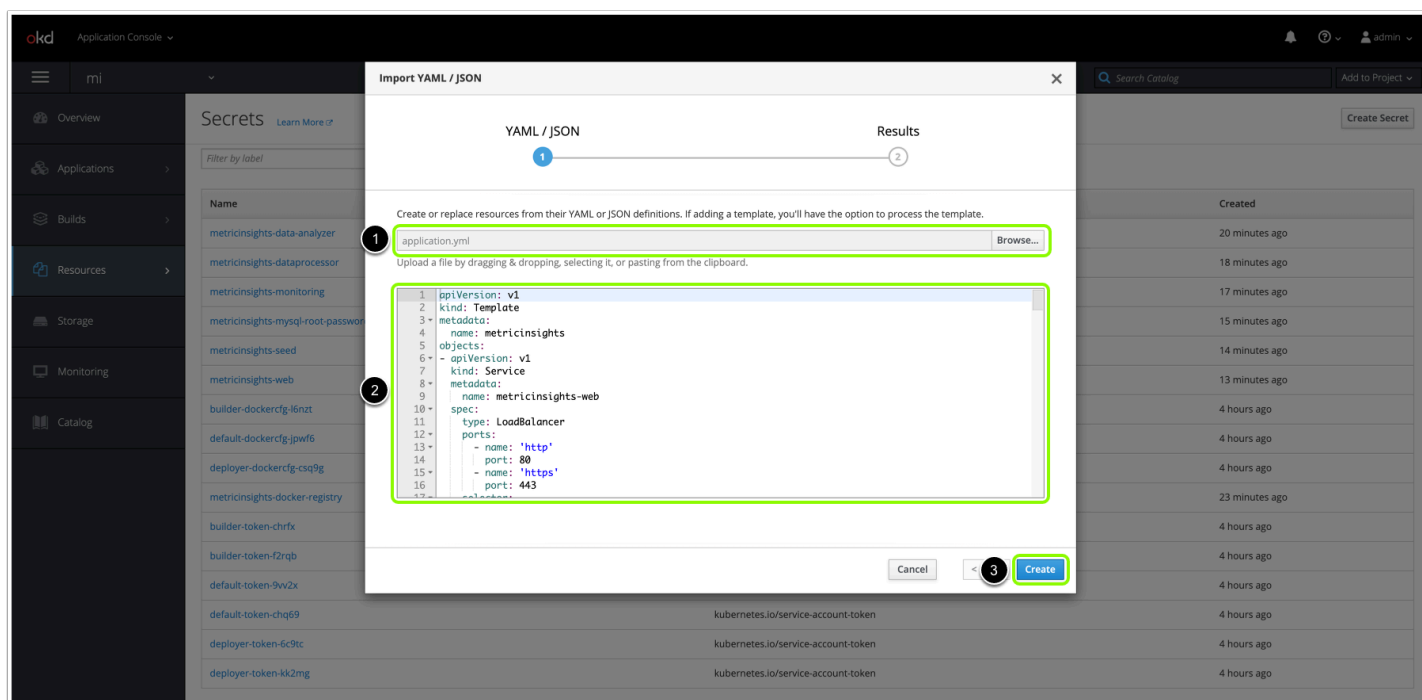
 See `./installer.py kubernetes -h` for more options.

Next, apply the OpenShift manifest in the OpenShift Web Console using the newly created deployment configuration file:

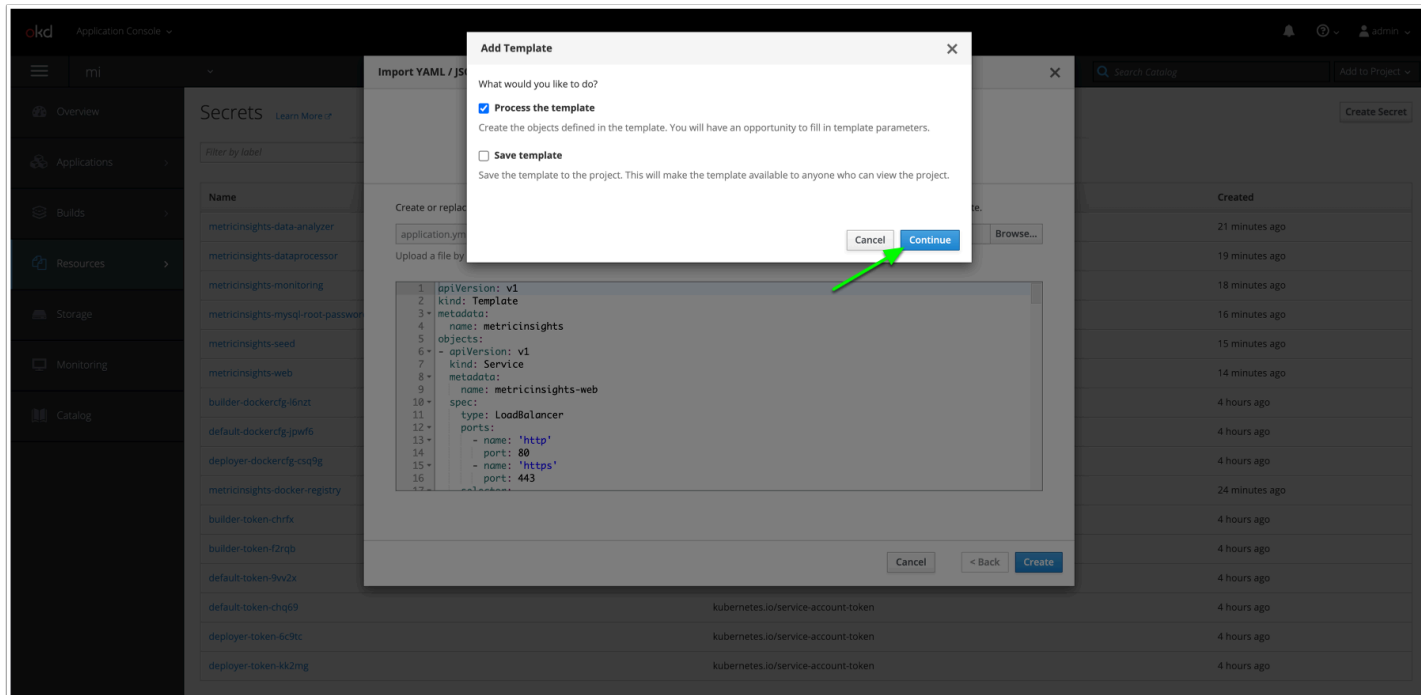
Click the **Add to Project** dropdown menu (1) then select **Import YAML / JSON** (2):



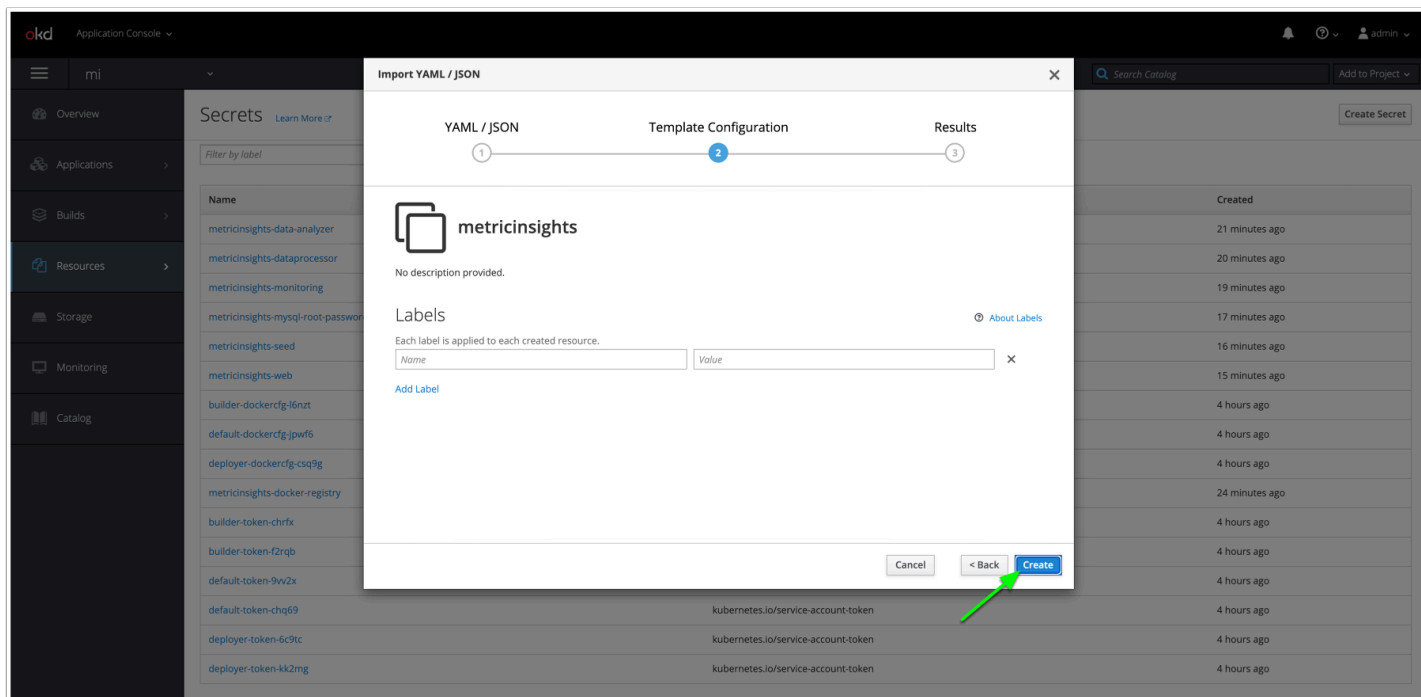
Use **Browse** to upload the manifest file (1), or copy the file content directly to appropriate fields (2), then click on **Create** button (3):



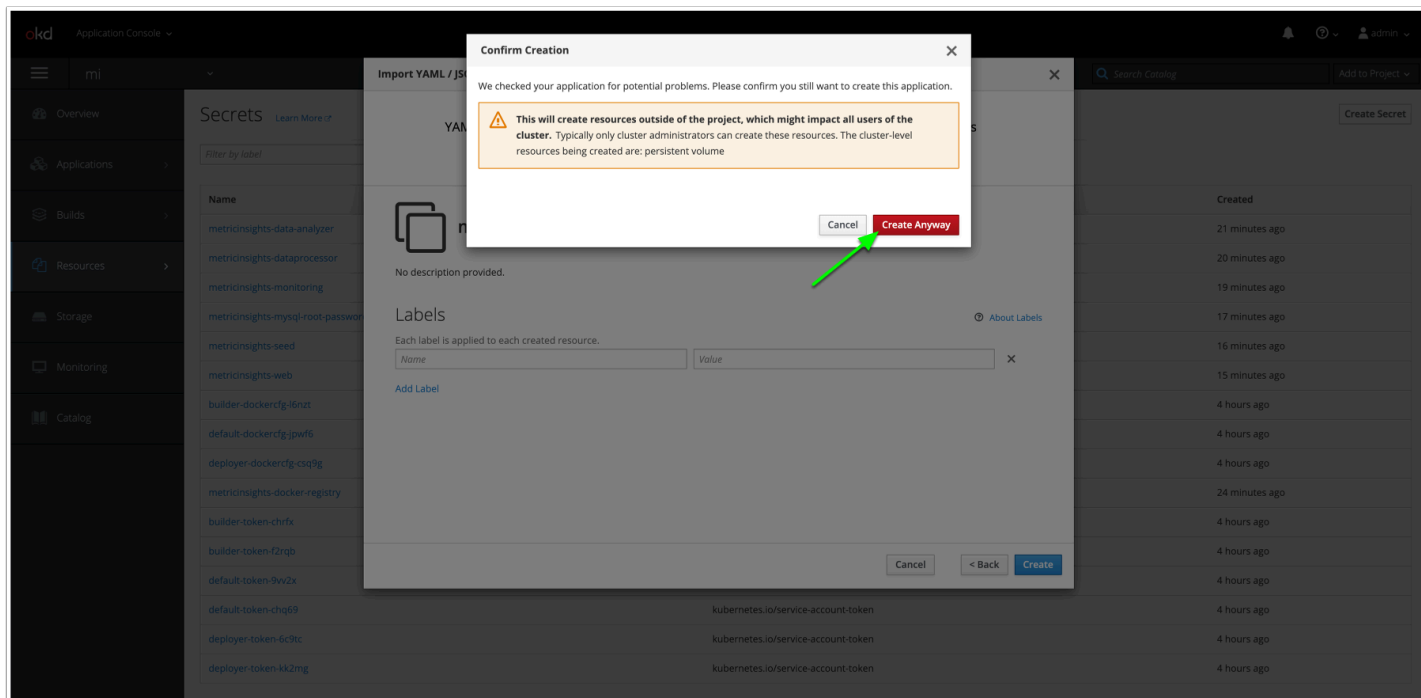
Next, click the **Continue** button:



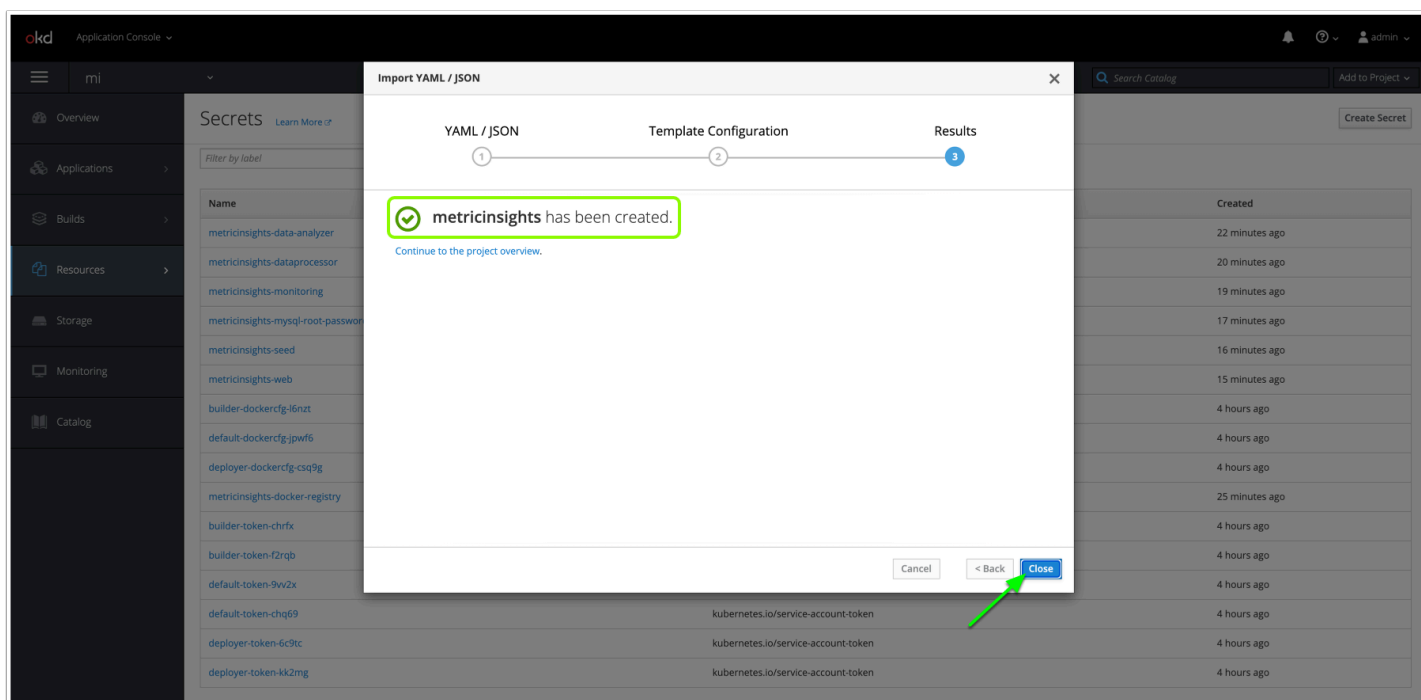
Click the **Create** button:



Confirm Creation by clicking on **Create Anyway:**



A confirmation message should appear confirming the deployment after which we can click on **Close**:



💡 To deploy/update Metric Insights from the command line, rather than through the OpenShift Web Console:


```
$ oc apply -f application.yml
template.template.openshift.io/metricinsights unchanged

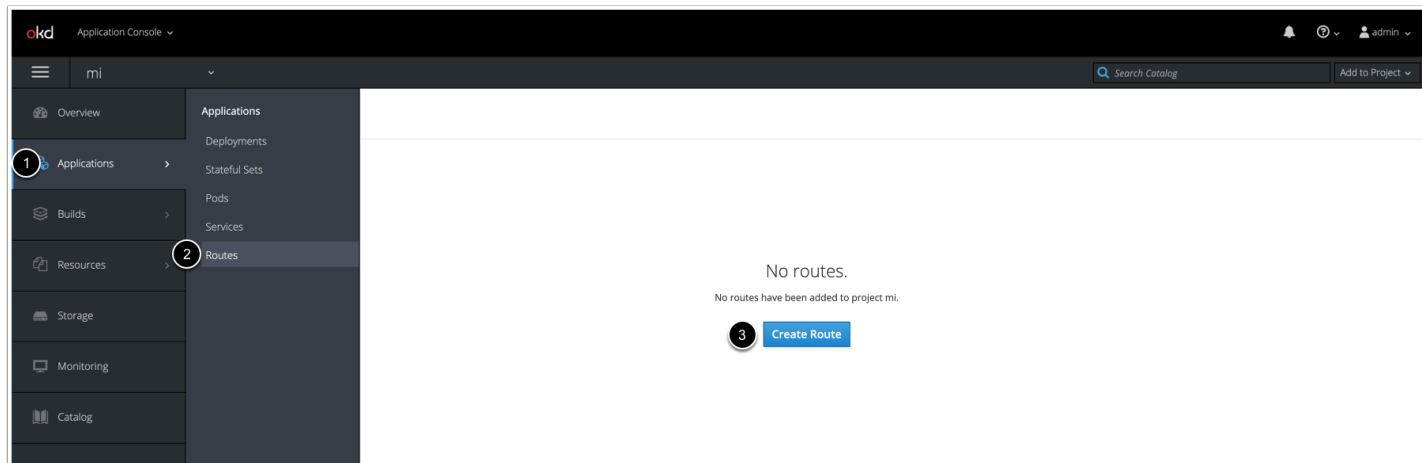
$ oc process -f manifest.yml | oc apply -f -
service/metricinsights-web unchanged
deployment.apps/metricinsights-web-master unchanged
deployment.apps/metricinsights-web-slave unchanged
service/metricinsights-seed unchanged
deployment.apps/metricinsights-seed unchanged
service/metricinsights-dataprocessor unchanged
deployment.apps/metricinsights-dataprocessor unchanged
service/metricinsights-data-analyzer unchanged
deployment.apps/metricinsights-data-analyzer unchanged
service/metricinsights-monitoring unchanged
deployment.apps/metricinsights-monitoring unchanged
persistentvolume/metricinsights-default-data configured
persistentvolumeclaim/metricinsights-default-data unchanged
```

7. Create Routes

Once Metric Insights is deployed, we need to create routes for ***MetricInsights Application*** and ***Monitoring Tool***

Go to **Application** → **Routes** → **Create Route**

Example for creating Web Route:



Example for Metric Insights Web UI Route:

Deployment & Configuration

Example for Metric Insights Monitoring Route:

Deployment & Configuration

okd

Application Console

mi

Search Catalog

Add to Project

Overview

Applications

Builds

Resources

Storage

Monitoring

Catalog

TLS Termination

Re-encrypt

Learn More

Insecure Traffic

Redirect

Policy for traffic on insecure schemes like HTTP.

Certificates

TLS certificates for edge and re-encrypt termination. If not specified, the router's default certificate is used.

Certificate

Browse...

The PEM format certificate. Upload file by dragging & dropping, selecting it, or pasting from the clipboard.

-----BEGIN CERTIFICATE-----
MIIGvTCCBaWgAwIBAgIQcus5DRTKJFn5mT217JJZMjANBgkqhkiG9w0BAQsFADBe
MQswCQYDVQQGEwJVVzEVMHMGAlUECHMHRGlnaUNlcnQ5SjYMRkwFwYDVQQLExB3
d3cuZGlnaUNlcnQ5SjYMRkwFwYDVQDEwRXYXBPZnNTTCBSU0EgQ0EgMjAxODAw
Fw0xODAwMjAxMDAwMDBaFw0yMTA2MjQxMjAxMDAwMDBaMB8xHTABBgNVBAMPFouBwV0
Clear Value

Private Key

Browse...

The PEM format key. Upload file by dragging & dropping, selecting it, or pasting from the clipboard.

-----BEGIN RSA PRIVATE KEY-----
MIIEowIBAAKCAQEAzVTCSeLmc7bx75o9Sgh4JJFhnBRtbXGtmG1wXsyR5vfyYI
D3pmK9yfJN1B65rrcPPytVH8fVJ/GWCjaXaXpAFAfhppM02p0aZ8GUT/wptmp0a
/fRLKuv68tDr5Pa8y/o3tP/paz2MfUoG8LEdchcIH7Y2x9g/RemKk+07BE9M5c
zNaLpKrcmRra+1Ghwz30IH1jBkge1S8ratvxE1096CoRnubZzLv0eOR0ADMPa98b
Clear Value

CA Certificate

Browse...

The PEM format CA certificate chain. Upload file by dragging & dropping, selecting it, or pasting from the clipboard.

-----BEGIN CERTIFICATE-----
MIIGvTCCBaWgAwIBAgIQcus5DRTKJFn5mT217JJZMjANBgkqhkiG9w0BAQsFADBe
MQswCQYDVQQGEwJVVzEVMHMGAlUECHMHRGlnaUNlcnQ5SjYMRkwFwYDVQQLExB3
d3cuZGlnaUNlcnQ5SjYMRkwFwYDVQDEwRXYXBPZnNTTCBSU0EgQ0EgMjAxODAw
Fw0xODAwMjAxMDAwMDBaFw0yMTA2MjQxMjAxMDAwMDBaMB8xHTABBgNVBAMPFouBwV0
Clear Value

okd

Application Console

mi

Search Catalog

Add to Project

Overview

Applications

Builds

Resources

Storage

Monitoring

Catalog

-----BEGIN RSA PRIVATE KEY-----
MIIEowIBAAKCAQEAzVTCSeLmc7bx75o9Sgh4JJFhnBRtbXGtmG1wXsyR5vfyYI
D3pmK9yfJN1B65rrcPPytVH8fVJ/GWCjaXaXpAFAfhppM02p0aZ8GUT/wptmp0a
/fRLKuv68tDr5Pa8y/o3tP/paz2MfUoG8LEdchcIH7Y2x9g/RemKk+07BE9M5c
zNaLpKrcmRra+1Ghwz30IH1jBkge1S8ratvxE1096CoRnubZzLv0eOR0ADMPa98b
Clear Value

CA Certificate

Browse...

The PEM format CA certificate chain. Upload file by dragging & dropping, selecting it, or pasting from the clipboard.

-----BEGIN CERTIFICATE-----
MIIGvTCCBaWgAwIBAgIQcus5DRTKJFn5mT217JJZMjANBgkqhkiG9w0BAQsFADBe
MQswCQYDVQQGEwJVVzEVMHMGAlUECHMHRGlnaUNlcnQ5SjYMRkwFwYDVQQLExB3
d3cuZGlnaUNlcnQ5SjYMRkwFwYDVQDEwRXYXBPZnNTTCBSU0EgQ0EgMjAxODAw
Fw0xODAwMjAxMDAwMDBaFw0yMTA2MjQxMjAxMDAwMDBaMB8xHTABBgNVBAMPFouBwV0
Clear Value

Destination CA Certificate

Browse...

The PEM format CA certificate chain to validate the endpoint certificate for re-encrypt termination. Upload file by dragging & dropping, selecting it, or pasting from the clipboard.

-----BEGIN CERTIFICATE-----
MIIF3zCCABegAwIBAgIUIGt4QK/x+ww9ldKTy0dSxba1cd8sw0QYJKoZIhvcNAQEL
BQAwfzELMAkGA1UEBmMCQVlxEzARBgNVBAGNCUNvbmWlU3RhdGUxITAFBgNVBAoM
GEudGybnV0IFdpZ2dpdHMqUHR5IE80ZDE4MDYGC5G5I3D0QJARYpbWV0cmlj
aW5zaWdodH0td2VlLW1hc3Rlc11jNTZmZmY1YjU1ZzZnbnYwHicMjAxMzE4NTUy
Clear Value

Labels

About Labels

Labels for this route.

Name

Value

Add Label

Create

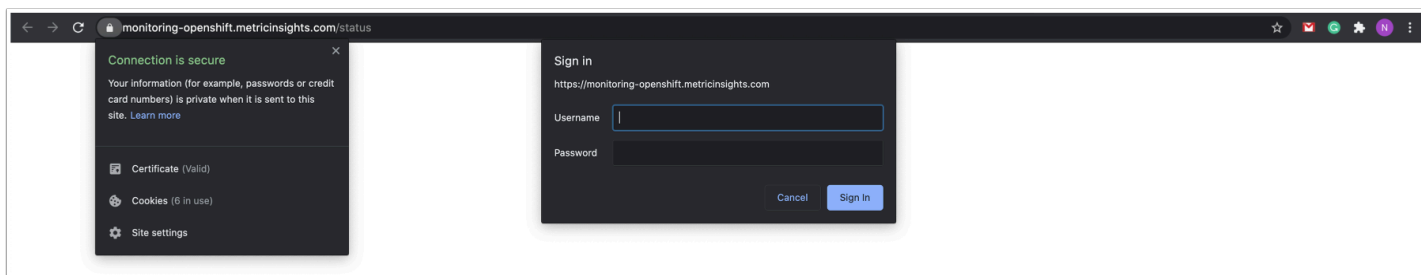
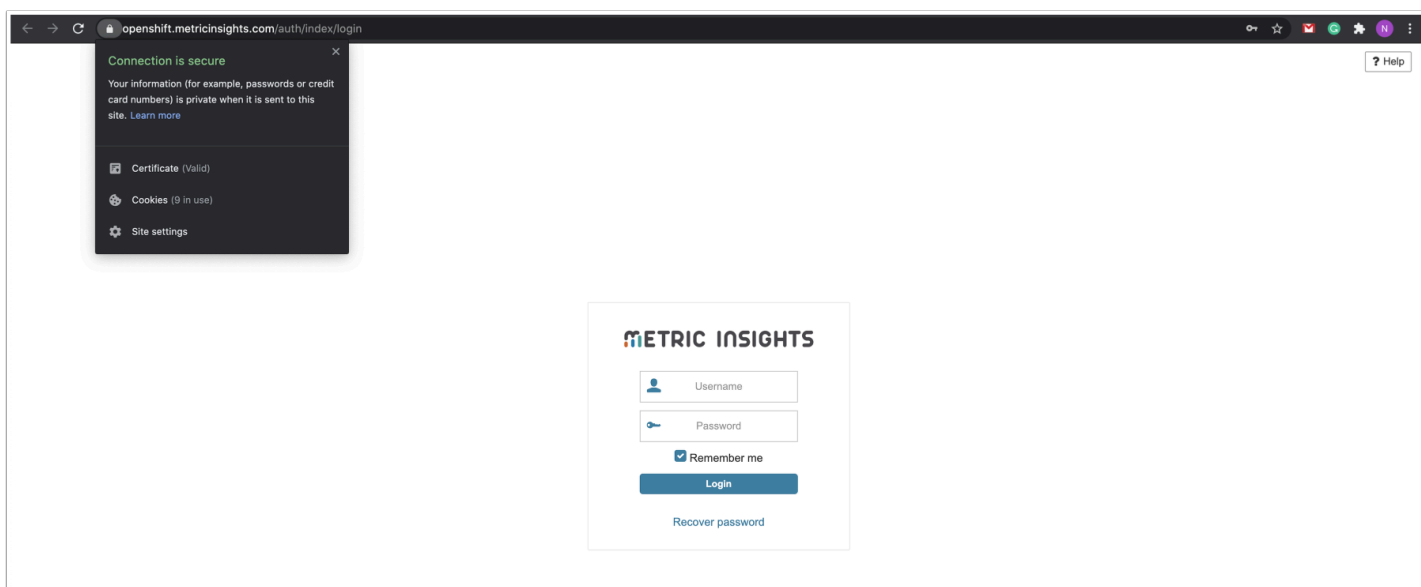
Cancel

8. Access Metric Insights Application in a Browser

After creating the routes, we can access the Metric Insights application in a browser. Navigate to the **Routes** menu and click on the appropriate link:

Name	Hostname	Service	Target Port	TLS Termination
monitoring	https://monitoring-openshift.metricinsights.com	metricinsights-monitoring	http	Re-encrypt
mi	https://openshift.metricinsights.com	metricinsights-web	https	Re-encrypt

Each link should open the Metric Insights application and monitoring page in a new browser tab:



1.26 Docker Deployment on Windows Host

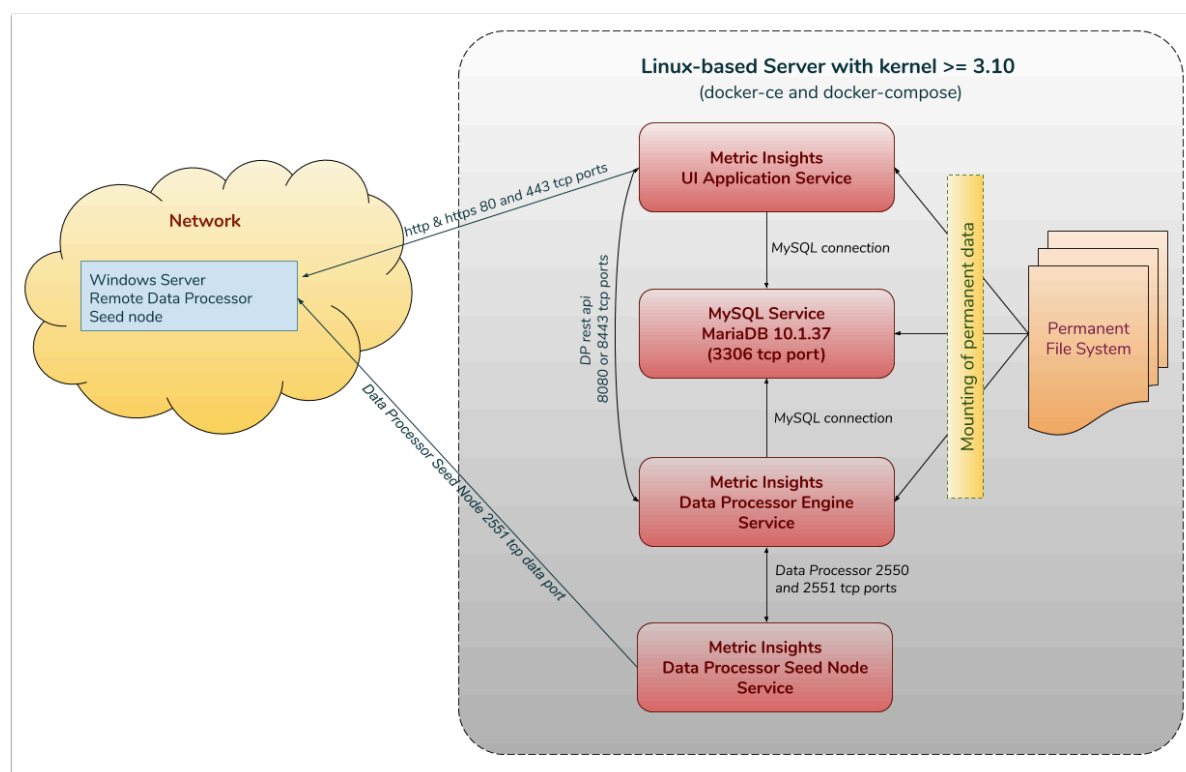
PREREQUISITES:

We recommend using a bare-metal Windows system.

The Windows Subsystem for Linux (WSL) version 2 introduces a significant architectural change (it is now a full Linux kernel built by Microsoft), allowing Linux containers to run natively without emulation.

1. Windows hypervisor should have a minimum 16Gb of RAM
2. If the hypervisor is virtualized, nested virtualisation should be enabled
3. Windows host should support Docker Desktop WSL 2 backend → <https://docs.docker.com/docker-for-windows/wsl/>
4. Follow the steps described in official Docker and Microsoft documentation → <https://docs.docker.com/docker-for-windows/wsl/#prerequisites>
5. To run the Metric Insights installer, the Linux subsystem needs to be enabled and one of our supported Linux distributions must be installed → [Supported Operating Systems](#)

Simple installation architecture

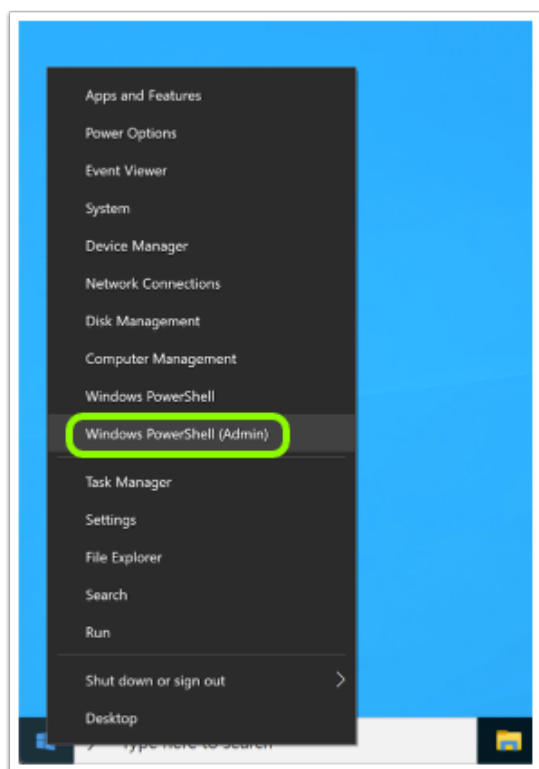


💡 If you are using a firewall on the Windows machine, the following ports are necessary for Metric Insights to work correctly:

1. **80, 443** - HTTP and HTTPS ports for the UI Application Service (by default redirection to 443)
2. **2550** - TCP port for the Data Processor cluster
3. **2551** - TCP port for the Seed Node Service
4. **3306** - MySQL port to get access from outside
5. **8080, 8443** - HTTP and HTTPS ports for the REST API Data Processor Service (only one port is enabled at one time)
6. **8081** - TCP port for Monitoring Tool

1. Connect to the server

1. Connect to the host server and run Windows PowerShell as Admin user.
Click **Right** mouse button on **Start** button → **Windows PowerShell (Admin)**:



2. In PowerShell terminal for accessing Linux Subsystem run `bash` command and for switching to **root** user run `sudo -i` :


```
root@doker-windows-test: ~  
Copyright (C) Microsoft Corporation. All rights reserved.  
Try the new cross-platform PowerShell https://aka.ms/pscore6  
PS C:\WINDOWS\system32> bash  
Welcome to Ubuntu 20.04.1 LTS (GNU/Linux 4.19.128-microsoft-standard x86_64)  
  
* Documentation:  https://help.ubuntu.com  
* Management:    https://landscape.canonical.com  
* Support:        https://ubuntu.com/advantage  
  
System information as of Wed Sep 16 04:40:08 PDT 2020  
  
System load:  0.83          Processes:      21  
Usage of /:   0.7% of 250.9GB Users logged in:  0  
Memory usage: 54%          IPv4 address for eth0: 192.168.97.223  
Swap usage:   0%  
  
49 updates can be installed immediately.  
12 of these updates are security updates.  
To see these additional updates run: apt list --upgradable  
  
This message is shown once once a day. To disable it please create the  
/home/mi/.hushlogin file.  
mi@doker-windows-test: /mnt/c/windows/system32$ sudo -i  
[sudo] password for mi:  
Welcome to Ubuntu 20.04.1 LTS (GNU/Linux 4.19.128-microsoft-standard x86_64)  
  
* Documentation:  https://help.ubuntu.com  
* Management:    https://landscape.canonical.com  
* Support:        https://ubuntu.com/advantage  
  
System information as of Wed Sep 16 04:40:57 PDT 2020  
  
System load:  0.56          Processes:      23  
Usage of /:   0.7% of 250.9GB Users logged in:  0  
Memory usage: 54%          IPv4 address for eth0: 192.168.97.223  
Swap usage:   0%  
  
49 updates can be installed immediately.  
12 of these updates are security updates.  
To see these additional updates run: apt list --upgradable  
  
This message is shown once once a day. To disable it please create the  
/root/.hushlogin file.  
root@doker-windows-test:~#
```

Note: The Metric Insights application must be installed as user *root* or as a *user with sudo privileges* (added to the *wheel* or *sudo* group, thereby included in the sudoers list). The docker daemon binds to a unix socket instead of a tcp port, therefore unix sockets are owned by user *root* and accessible by users with *sudo* privileges.

2. Download the MI app installation package

The screenshot shows the Metric Insights download page. Step 1 points to the URL bar showing the download link. Step 2 points to the Username and Password input fields. Step 3 points to the 'I agree with these terms and conditions' checkbox. Step 4 points to the 'Submit' button.

1. <https://www.metricinsights.com/download/MetricInsights-required-version.tar.gz>

METRIC INSIGHTS Product Solutions Learn Support About Free Trial

Download

Enter your user name and password below

2. Username *

Password *

Metric Insights Software Evaluation Agreement

This Metric Insights Customer Agreement (this "Agreement") contains the terms and conditions that govern your access to and use of the Service Offerings (as defined below) and is an agreement between Metric Insights, Inc. ("Metric Insights"), a California company, with an address at 123 10th Street, San Francisco, CA 94103, and you or the entity you represent ("User"). This Agreement takes effect when an Order is signed, when you click an "I Accept" button or check box presented with these terms, when this agreement is signed between Metric Insights and User, or, if earlier, when you use any of the Service Offerings (the "Effective Date"). You represent to us that you are lawfully able to enter into contracts (e.g., you are not a minor). If you are entering into this Agreement for an entity, such as the company you work for, you represent to us that you have legal authority to bind that entity. [Download license agreement](#)

3. ☐ I agree with these terms and conditions

4. Submit

1. Depending on the version, you will be given a link to the corresponding MI app installation package.
2. Along with the link, **Username** and **Password** will be provided by the MI support team; enter these credentials into the respective fields.
3. Confirm your agreement with the Terms and Conditions of the Metric Insights usage.
4. Click **[Submit]**.

This will initiate the download process of the **.tar.gz** archive (containing all files required to install the system).

3. Preparing Application Installer Package

In terminal go to archive location and untar it.

Example for *v6.2.1 Full Installer Package*:

```
cd /mnt/c/MetricInsights/

tar -xvzf MetricInsights-Installer-v6.2.1-Full.tar.gz

cd ./MetricInsights-Installer-v6.2.1-Full
```


4. Supported Linux Subsystems

! Metric Insights v6.x installer can work on the following Operating Systems (any OS with Linux kernel ≥ 3.10):

- RHEL 7/8
- CentOS 7/8
- Oracle Linux 7
- Ubuntu 14/16/18
- Debian 8/9/10



5. Run the installer

i Normally, the application can be installed by simply running `./installer.py -v` but we recommend checking the help to ensure appropriate parameters are defined for the install. To see the help:

```
./installer.py -h
```


Also, for installing under Windows hypervisor need to run installer with options `--skip-packages` and `--bind-address 0.0.0.0`

```
./installer.py --skip-packages --bind-address 0.0.0.0 -v
```

Once the installation is complete, you will see a "Metric Insights application installed successfully" message.

5.1. Basic console commands

! Not all tools will be installed with all components, each component has only: **mi-version** and **mi-control**. All of these tools become accessible only if the Web Component is installed.

mi-app-backup	Wrapper for mi-app-backup (/opt/mi/backup is an internal mount point from the web docker container)
mi-app-restore	Wrapper for mi-app-restore
mi-console	Connect to the MI UI docker container directly
mi-control	Docker-compose wrapper helping to manage Metric Insights docker containers
mi-cron	Wrapper for mi-cron manipulating the MI cron tasks inside the docker container
mi-db-dump	Wrapper for mi-db-dump
mi-db-mgr	Wrapper for mi-db-mgr
mi-ldap-usersync	Wrapper for mi-ldap-usersync
mi-maintenance	Wrapper for mi-maintenance
mi-passwd	Wrapper for mi-passwd
mi-version	Show the installed Metric Insights version

6. Checking deployment

To check container statuses we can run `docker ps` command from **PowerShell** or **Linux Subsystem** terminal console:

```
Administrator: Windows PowerShell
root@docker-windows-test:/mnt/c/MetricInsights/MetricInsights-Installer-v6.2.1-Full docker ps
CONTAINER ID        IMAGE                                     COMMAND                  CREATED          STATUS          PORTS                               NAMES
bf50161d9b3        docker-metricinsights.com:5080/mi-application/release_6_2/442  "/docker-entrypoint..." 2 days ago       Up 2 days (healthy) 0.0.0.0:8080->8080/tcp, 0.0.0.0:443->443/tcp  ni_web_1
79933769396a        docker-metricinsights.com:5080/monitoring/release_6_2/311      "/bin/sh -c 'docker..." 2 days ago       Up 2 days (healthy) 0.0.0.0:8081->8081/tcp  ni_monitoring_1
582797237465        docker-metricinsights.com:5080/eval/release_6_2/2              "/docker-entrypoint..." 2 days ago       Up 2 days (healthy) 127.0.0.1:3306->3306/tcp, 33060/tcp  ni_mysql_1
8421850e2026b5      docker-metricinsights.com:5080/dataprocessor/release_6_2/73     "/docker-entrypoint..." 2 days ago       Up 2 days (healthy) 0.0.0.0:2550->2550/tcp, 0.0.0.0:8080->8080/tcp, 0.0.0.0:8443->8443/tcp  ni_dataprocessor_1
8a2d3a5ee2d         docker-metricinsights.com:5080/data-analyzer/release_6_2/99     "/docker-entrypoint..." 2 days ago       Up 2 days (healthy) 0.0.0.0:2551->2551/tcp  ni_data_analyzer_1
80e72a2f7ba        docker-metricinsights.com:5080/dataprocessor-seed-mode/release_6_2/24  "/docker-entrypoint..." 2 days ago       Up 2 days (healthy) 127.0.0.1:5134->5134/tcp  ni_seed_1
8421850e2026b5      docker-metricinsights.com:5080/logger/release_6_2/48           "/docker-entrypoint..." 2 days ago       Up 2 days (healthy) 127.0.0.1:5134->5134/tcp  ni_logger_1
8086a679fb47a       docker-metricinsights.com:5080/logger/release_6_2/48           "/docker-entrypoint..." 2 days ago       Up 2 days (healthy) 127.0.0.1:5134->5134/tcp  logger_logrotate_1

root@docker-windows-test:/mnt/c/MetricInsights/MetricInsights-Installer-v6.2.1-Full logot
root@docker-windows-test:/mnt/c/MetricInsights/MetricInsights-Installer-v6.2.1-Full logot
PS C:\WINDOWS\system32> docker ps
CONTAINER ID        IMAGE                                     COMMAND                  CREATED          STATUS          PORTS                               NAMES
bf50161d9b3        docker-metricinsights.com:5080/mi-application/release_6_2/442  "/docker-entrypoint..." 2 days ago       Up 2 days (healthy) 0.0.0.0:8080->8080/tcp, 0.0.0.0:443->443/tcp  ni_web_1
79933769396a        docker-metricinsights.com:5080/monitoring/release_6_2/311      "/bin/sh -c 'docker..." 2 days ago       Up 2 days (healthy) 0.0.0.0:8081->8081/tcp  ni_monitoring_1
582797237465        docker-metricinsights.com:5080/eval/release_6_2/2              "/docker-entrypoint..." 2 days ago       Up 2 days (healthy) 127.0.0.1:3306->3306/tcp, 33060/tcp  ni_mysql_1
8421850e2026b5      docker-metricinsights.com:5080/dataprocessor/release_6_2/73     "/docker-entrypoint..." 2 days ago       Up 2 days (healthy) 0.0.0.0:2550->2550/tcp, 0.0.0.0:8080->8080/tcp, 0.0.0.0:8443->8443/tcp  ni_dataprocessor_1
8a2d3a5ee2d         docker-metricinsights.com:5080/data-analyzer/release_6_2/99     "/docker-entrypoint..." 2 days ago       Up 2 days (healthy) 0.0.0.0:2551->2551/tcp  ni_data_analyzer_1
80e72a2f7ba        docker-metricinsights.com:5080/dataprocessor-seed-mode/release_6_2/24  "/docker-entrypoint..." 2 days ago       Up 2 days (healthy) 127.0.0.1:5134->5134/tcp  ni_seed_1
8421850e2026b5      docker-metricinsights.com:5080/logger/release_6_2/48           "/docker-entrypoint..." 2 days ago       Up 2 days (healthy) 127.0.0.1:5134->5134/tcp  ni_logger_1
8086a679fb47a       docker-metricinsights.com:5080/logger/release_6_2/48           "/docker-entrypoint..." 2 days ago       Up 2 days (healthy) 127.0.0.1:5134->5134/tcp  logger_logrotate_1

PS C:\WINDOWS\system32>
```

In browser navigate to Windows hypervisor address to check Web UI and Monitoring UI:

⚠ Not Secure | 10.8.0.162/monitor#top_tab_nodes
🔍 ☆

Status Monitor
+ New... 📄 Content ⚙ Admin 👤 System

☒ Notifications ON system-wide · Wednesday - September 16, 2020 System ⌚ 13:13 UTC(UTC +00:00)

[Status](#)
[Nodes](#)
[Errors & Logs](#)
[Bulk Change History](#)

System Nodes						
ID	Name	IP	Uptime	CPU	RAM	Disk
4	data-analyzer	172.19.0.5	Up 2 days	4.3 %	519.6Mb of 14.06Gb (3.61%)	2 Volumes
5	dataprocessor	172.19.0.6	Up 2 days	4.9 %	2.73Gb of 14.06Gb (19.38%)	2 Volumes
1	monitoring	172.19.0.3	Up 2 days	3.7 %	88.45Mb of 14.06Gb (0.61%)	2 Volumes
3	mysql	172.19.0.4	Up 2 days	4.4 %	3.24Gb of 14.06Gb (23.02%)	2 Volumes
2	seed	172.19.0.2	Up 2 days	4.3 %	360.35Mb of 14.06Gb (2.50%)	2 Volumes
6	web	172.19.0.7	Up 2 days	3.6 %	483.97Mb of 14.06Gb (3.36%)	2 Volumes

Data Storage Nodes		
Name	Status	Connections
Default MySQL dataset storage	UP	27

▲ Not Secure | 10.8.0.162:8081/status

☆

#	Name	Role	Services	Uptime	Status
1.	data-analyzer	Data Analyzer	monitor, python	Up 2 days	✓
2.	dataprocessor	Data Processor	java, monitor	Up 2 days	✓
3.	monitoring	Monitoring Service	monitor, python	Up 2 days	✓
4.	mysql	MySQL	mysqld, monitor	Up 2 days	✓
5.	seed	Data Processor (Seed Node)	java, monitor	Up 2 days	✓
6.	web	UI (master)	supervisord, cron, rsyslogd, apache2, monitor	Up 2 days	✓

Recent events

1. 2020-09-14 08:58:29. Attention: The role *UI (master)* has been assigned to node **web**.

2. 2020-09-14 08:57:44. Success! Service *supervisord* on node **web** is available.

3. 2020-09-14 08:57:44. Success! Service *cron* on node **web** is available.

4. 2020-09-14 08:57:44. Success! Service *rsyslogd* on node **web** is available.

5. 2020-09-14 08:57:44. Success! Service *apache2* on node **web** is available.

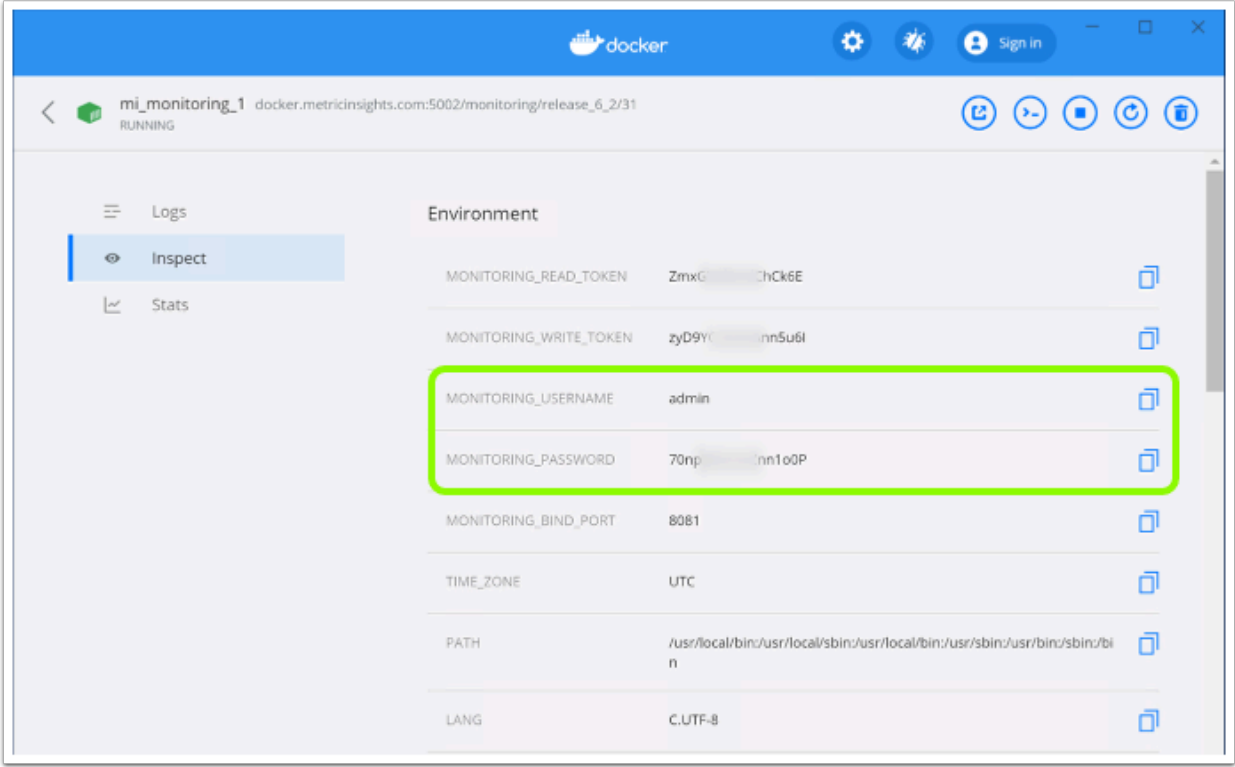
6. 2020-09-14 08:57:44. Success! Service *monitor* on node **web** is available.

7. 2020-09-14 08:57:34. Attention: The node **web** is now available.

8. 2020-09-14 08:57:28. Attention: The node **web** is registered with address **172.19.0.7**

9. 2020-09-14 08:57:28. Service *supervisord* on node **web** is registered.

Note: username and password for Monitoring UI can be accessed from Docker Dashboard:



2. Using Templates

2.1 Setting Defaults for New Users

Templates allow you to set a number of User attributes automatically **when creating new Users**. Most of the templates will allow different defaults to be set based on User's Authority Level (currently defined for either Admin or Regular User).

NOTE: The values set in the templates **will not** change any existing Users - and any of these fields can be independently changed by accessing the associated screens from **User Editor**, either while creating users or after.

After the first login a user can set up his own defaults at **My Preferences** menu at the top right corner of the screen below his Username.

Access Admin > Advanced > New User Defaults

Define the Preferences settings applied to each new Admin and Regular/Power User		
User Preference Templates		
Name		
User Defaults		
Administrator Defaults		
Alert Subscriptions Templates		
Name		
User Defaults		
Administrator Defaults		
Email Notifications Templates		
Name		
User Defaults		
Administrator Defaults		

There are currently 3 sets of **Templates** for available for setting User-level defaults:

- [User Preference Templates](#)
- [Alert Subscriptions Templates](#)
- [Email Notifications Templates](#)

Each these will have separate screens for defining defaults for **Admin** and **Regular Users**.

User Preference Templates

User Preference Templates > User Defaults > Profile Information tab

The screenshot shows the 'New User Defaults / User Defaults' page. The top navigation bar includes a home icon, the page title, and buttons for 'New...', 'Content', 'Admin', and a user profile 'Alex'. Below the navigation bar, there are two tabs: 'Profile Information' (selected) and 'Homepage Display'. The 'Profile Information' tab contains two sections: 'Report Viewer' and 'Data Entry Preferences'. The 'Report Viewer' section has a numbered callout '1' pointing to a text input field labeled 'Max data table display lines' with the value '20'. The 'Data Entry Preferences' section has two numbered callouts: '2' pointing to a radio button group for 'Warn me when i attempt to leave a web page that has pending changes?' with 'yes' selected, and '3' pointing to a radio button group for 'Validate fetch command variables?' with 'yes' selected. At the bottom right, there is a 'Saved' button.

New User Defaults / User Defaults

New... Content Admin Alex

Profile Information Homepage Display

Report Viewer

1 Max data table display lines 20

Data Entry Preferences

2 Warn me when i attempt to leave a web page that has pending changes? ☒ yes ☐ no

3 Validate fetch command variables? ☒ yes ☐ no

Saved

1. Define the standard maximum rows of the tables listing system entities (Elements, Bursts, Dimensions, Data Sources, etc.)
2. Set preference for warning message generation
3. This setting is applied to validating substitution parameters in SQL statements at *Metric/Report Editor > Data Collection*.

User Preference Templates > User Defaults > Home Page Display tab

New User Defaults / User Defaults

Profile Information | **Homepage Display**

1 Color scheme: Default 2-color

Outperforming Value: facebook.com Monthly Active Users, Saturday 26/07/2016, +3.5%, 256.0M

Performing Value: facebook.com Monthly Active Users, Saturday 29/07/2016, +0.25%, 256.0M

Underperforming Value: facebook.com Monthly Active Users, Saturday 26/07/2016, -1.3%, 256.0M

2 Show preview Charts: yes | no

Homepage refresh frequency: 600 seconds

When grouping tiles by a Dimension: show totals first | show totals last | omit totals group

3 Start Homepage in:

Layout: Tile

Folder: Last Active Category

Topic: All Topics

Filter: Show All

Group: Favorites & Categories

4 Start Mobile Application in:

Folder: My Favorites

Homepage Flags for Alerts, Annotations and Notes

5 Clear from tile after hovering: 3 seconds

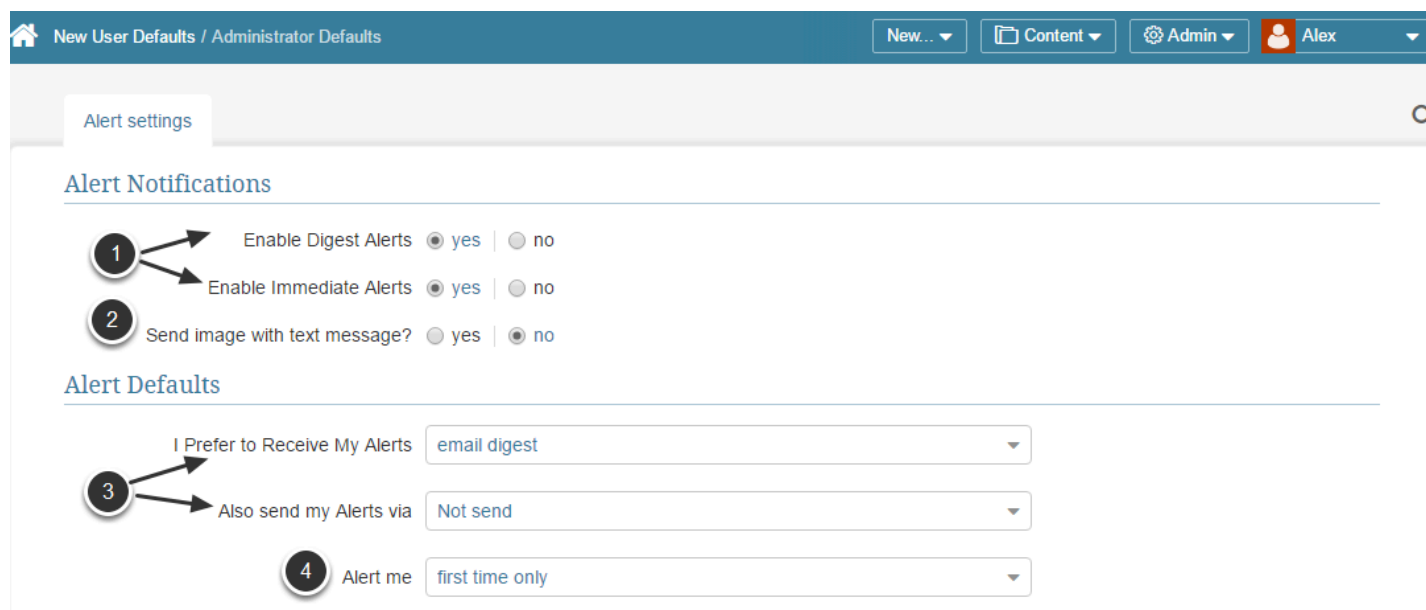
Preview: New High (previous high: \$5.49M on 09/17/2016), \$5.33M, \$28.77M

Saved

You can setup the default *Homepage* display for new Regular Users or New Administrators.

1. **Tile color scheme:** Choose a combination of colors that are going to define Performing, Underperforming and Outperforming values. This powerful tool makes the Homepage easier to scan and grasp a bigger picture within seconds.
2. **Show preview Charts:** If this option is set to 'no', upon single-clicking a tile user is sent directly to an element Viewer. If it is set to 'yes', single-clicking leads to a Preview page and double clicking leads to the Viewer.
3. Define a a Layout, Folder, Topic, Filter or even Favorites Folder the Homepage should start in.
4. You may define the Folder or Category the Mobile application should start in.
5. **Clear from tile after hovering:** The time after which Alert, Annotation or Note icons should be hidden when hovering over them

Alert Subscriptions Template > Administrator Defaults



New User Defaults / Administrator Defaults

Alert settings

Alert Notifications

1 Enable Digest Alerts ☒ yes | ☐ no

Enable Immediate Alerts ☒ yes | ☐ no

2 Send image with text message? ☐ yes | ☒ no

Alert Defaults

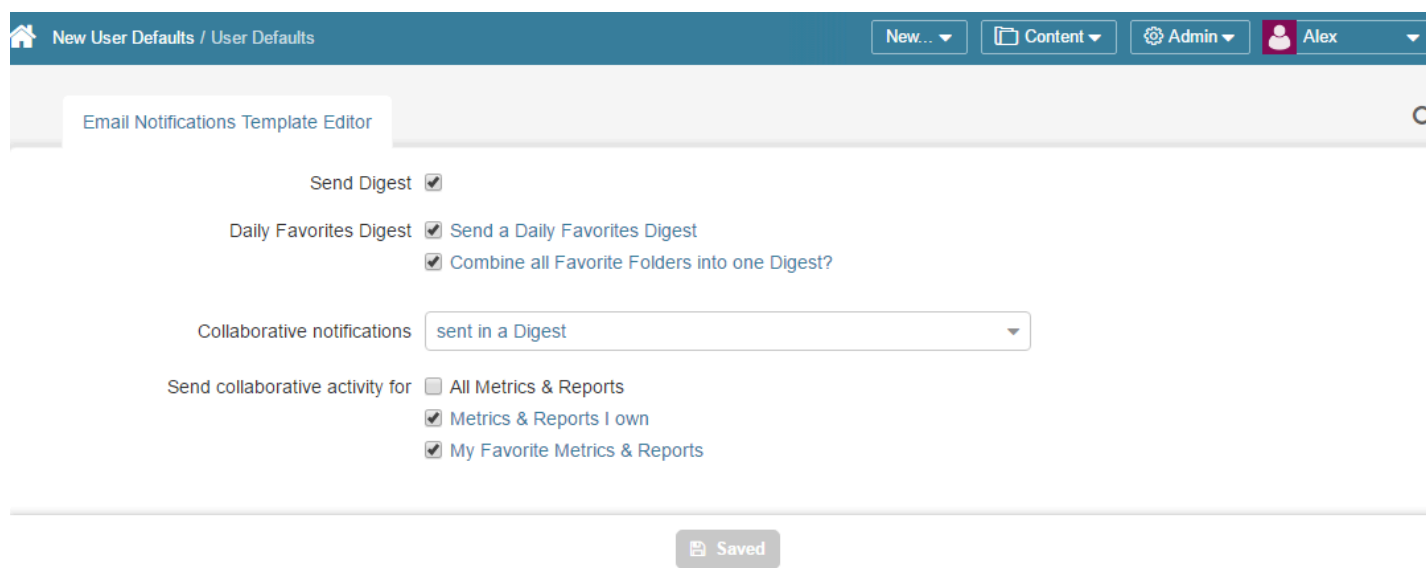
3 I Prefer to Receive My Alerts

Also send my Alerts via

4 Alert me

1. When a new User is created, he is usually assigned to some existing User Group providing him with certain Permissions and automatically subscribing him to Alerts this group is subscribed to. Define whether these alerts should be enabled.
2. Send image with text message?: relates to SMS/MMS functionality.
3. Define the primary and secondary methods of Alerts receiving.
4. Configure alert persistence

Email Notifications Template > User Defaults



New User Defaults / User Defaults

Email Notifications Template Editor

Send Digest ☒

Daily Favorites Digest ☒ Send a Daily Favorites Digest

☒ Combine all Favorite Folders into one Digest?

Collaborative notifications

Send collaborative activity for ☐ All Metrics & Reports

☒ Metrics & Reports I own

☒ My Favorite Metrics & Reports

Saved

2.2 Creating a new Email Digest Template

Email templates in Metric Insights are created using HTML, CSS, and JavaScript, allowing great control over how the email body of your company's digests will look.

If your business requires a custom template, you can create one following the instructions below, or contact support@metricinsights.com to submit a request for a template with specific requirements.

1. Open the list of Email Templates

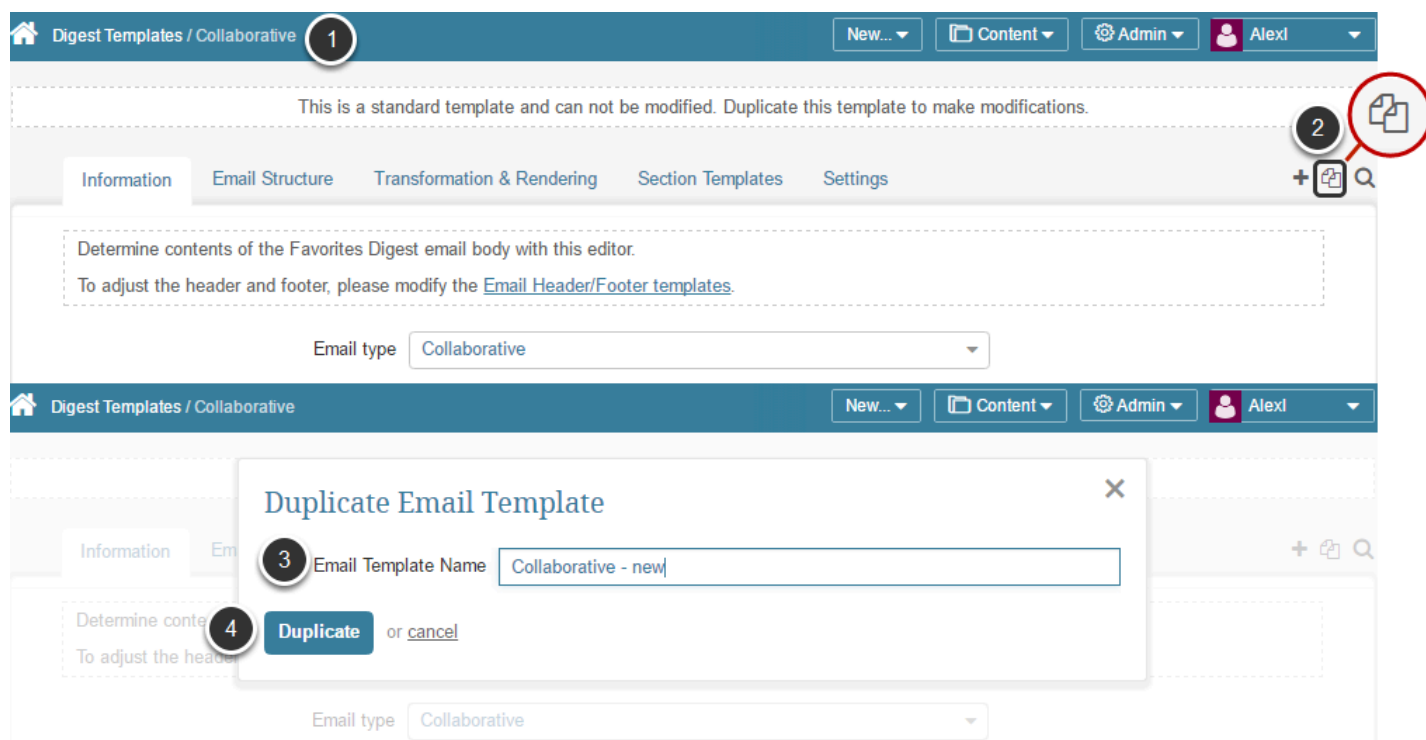
The screenshot shows the Metric Insights Admin interface. On the left, the 'Admin' menu is open, and 'Email Templates' is highlighted with a circled '1'. The main content area shows the 'Email Templates' list with a circled '2' indicating the list itself. The list contains 16 templates, each with a 'Name' and a 'Default Template' column. The templates are: Alert (locked), Alert Simplified (locked), Burst - Any Image Size (locked), Burst - Any Image, Burst - Fixed Image Size (locked), Collaborative (locked), Data Trust (locked), Examples for header/footer templates (locked), Favorite (locked), Favorite Simplified (locked), Favorites with Headlines (locked), Out of disk space (locked), Remove collaboration item (locked), Request Permission (locked), Revisions to Favorites Digest, Share Element/Discussion/Comment (locked), and Welcome/Recovery (locked). A '+ New Template' button is at the bottom left of the list.

Name	Default Template
Alert (locked)	Alert
Alert Simplified (locked)	Alert
Burst - Any Image Size (locked)	
Burst - Any Image	
Burst - Fixed Image Size (locked)	
Collaborative (locked)	Collaborative
Data Trust (locked)	
Examples for header/footer templates (locked)	
Favorite (locked)	Digest
Favorite Simplified (locked)	
Favorites with Headlines (locked)	
Out of disk space (locked)	
Remove collaboration item (locked)	
Request Permission (locked)	
Revisions to Favorites Digest	
Share Element/Discussion/Comment (locked)	
Welcome/Recovery (locked)	

1. Go to *Admin > Email Templates*
2. The list of email templates available in the system opens. Your Metric Insights instance comes preloaded with 16 default templates shown at the screen above. Each of such templates is marked either as **locked** or **default** and can be neither edited, nor deleted

2. Creating a new template

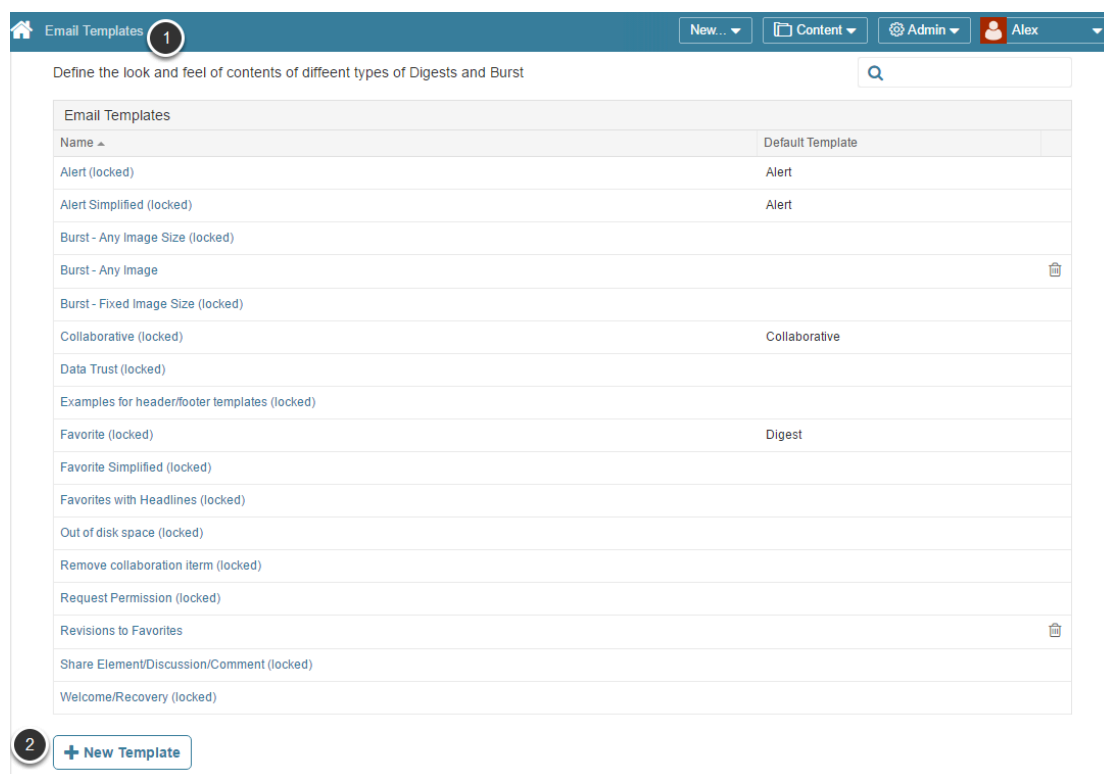
[Option 1] Duplicate on of the default templates



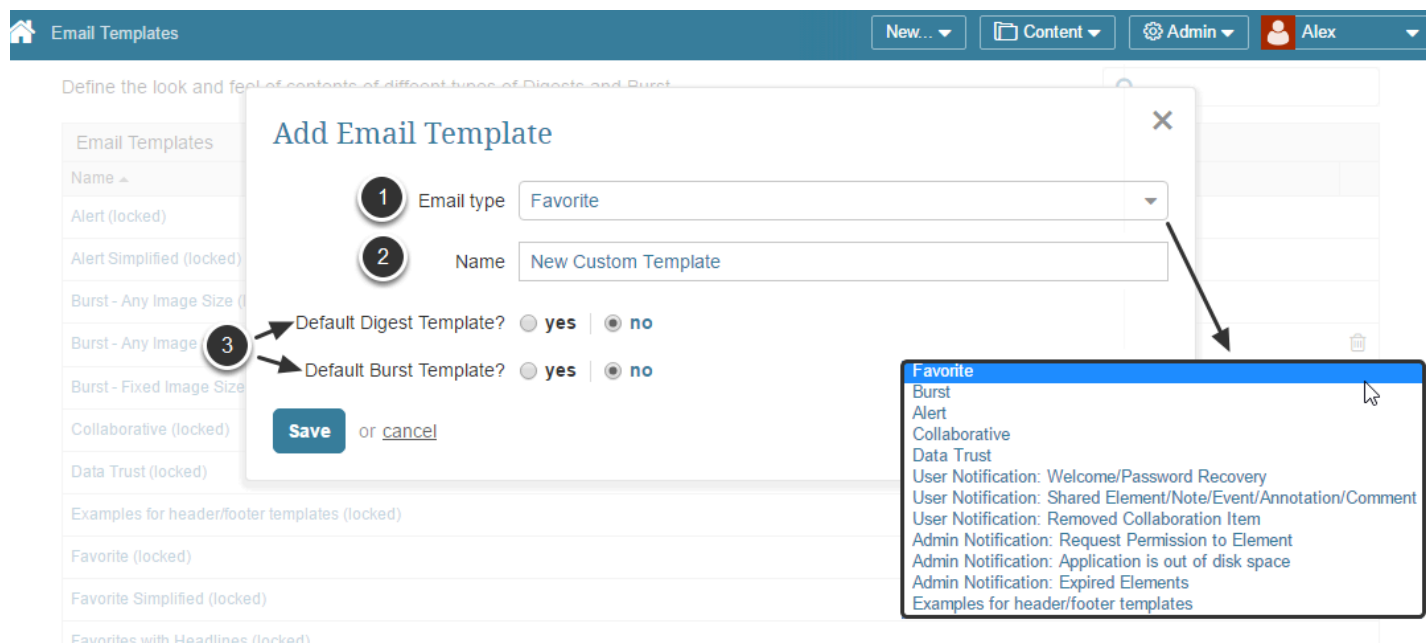
1. Open the *Digest Template* that should serve as a basis of a new editable template
2. Click the **Duplicate** icon at the top right corner of the screen
3. Provide a name for a new template in the opened pop-up
4. Click **Duplicate**

You can edit this template according to your requirements. For more details on Email Template properties and settings, refer to Step 3.

[Option 2] Create a new Template based on an existing one



1. Open the list of Email Templates
2. At the page bottom click **[+ New Template]**



1. Select the email type to which a new template should be applied
2. Enter unique template **Name**

3. This section changes dynamically depending on the Email type in the first field. Specify if the template will be the default for either Digests or Bursts

Save your entry.

3. Defining Digest Template settings in the Editor

Information tab

Information tab allows you to set Default templates and to specify datasets for inclusion in Preview and Test emails.

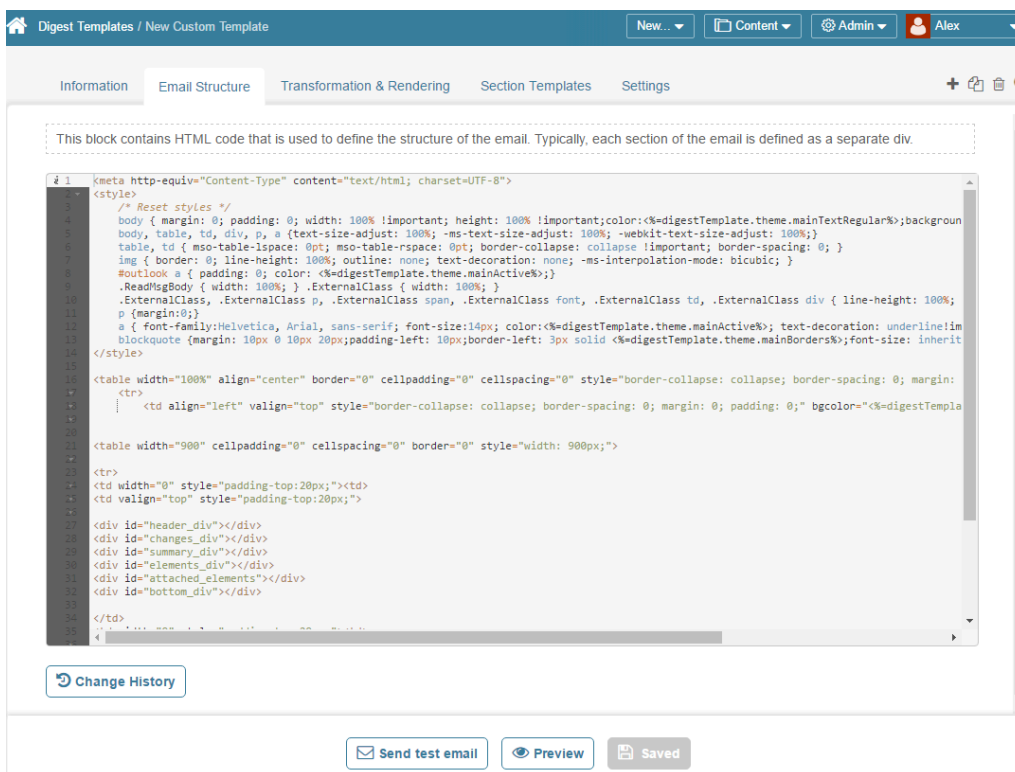
The *Attached Datasets* section allows you to add a set of data from MI elements to be available in the email digest. In order to specify a dataset:

1. Click **[+ New Dataset]**
2. Select a dataset you want to use
3. **Save** the changes made

You can use this option to specify MI elements you want to see in your digest for the specified template. E.g. When you specify Report 1 in Dataset for a particular Digest Template and set users to receive the template emails, they will receive the Report 1 dataset in emails.

NOTE: It's not obligatory to add datasets for creating digest templates.

Email Structure



The *Change history* button allows you to view revisions to your code over time.

The screenshot above displays the default Favorites template. You can use this as a base example of the code from which to build your own template.

NOTE: If you duplicated one of the default templates, this tab is going to contain the HTML code that is used to define the structure of the email body. If you create a template from scratch, it is going to be empty.

Transformation & Rendering

The screenshot shows the 'Transformation & Rendering' tab in the Metric Insights interface. The top navigation bar includes a home icon, 'Digest Templates / New Custom Template', and buttons for 'New...', 'Content', 'Admin', and a user profile 'Alex'. Below the navigation bar, there are tabs for 'Information', 'Email Structure', 'Transformation & Rendering' (which is active), 'Section Templates', and 'Settings'. A dashed box contains the instruction: 'Manipulate input data (append report data to the core digest data object (digestTemplate.data), apply transformation functions), initiate rendering of email sections defined in sections code.' The main area is a code editor with a light gray background and a dark gray sidebar on the left. The code is as follows:

```
1 if(digestTemplate.data && digestTemplate.data[0]) {  
2   $('#header_div').html(digestTemplate.renderSection('headerSection', digestTemplate.data[0]));  
3   } else $('#header_div').remove();  
4  
5   if(digestTemplate.data && digestTemplate.data[0] && digestTemplate.data[0]['changes'] && digestTemplate.data[0]['changes'].length>0) {  
6     $('#changes_div').html(digestTemplate.renderSection('changesSection', digestTemplate.data[0])+'<br />');  
7     } else $('#changes_div').remove();  
8  
9   if(digestTemplate.data && digestTemplate.data[0] && digestTemplate.data[0]['summary']) {  
10    $('#summary_div').html(digestTemplate.renderSection('summarySection', digestTemplate.data[0])+'<br />');  
11    } else $('#summary_div').remove();  
12  
13   if(digestTemplate.data && digestTemplate.data[0] && digestTemplate.data[0]['elements']) {  
14    $('#elements_div').html(digestTemplate.renderSection('elementsSection', digestTemplate.data[0])+'<br />');  
15    } else $('#elements_div').remove();  
16  
17   if(digestTemplate.data && digestTemplate.data[0]) {  
18    $('#bottom_div').html(digestTemplate.renderSection('bottomSection', digestTemplate.data[0]));  
19    } else $('#bottom_div').remove();  
20 }
```

Below the code editor is a 'Change History' button. At the bottom of the interface, there are three buttons: 'Send test email', 'Preview', and 'Saved'.

This tab contains the rendering template for each section identified in the sections tab. This code specifies how the data will be rendered.

NOTE: If you duplicated one of the default templates, this tab is going to contain the HTML code that is used to define the structure of the email body. If you create a template from scratch, it is going to be empty.

Section Templates

Digest Templates / New Custom Template

New... Content Admin Alex

Information Email Structure Transformation & Rendering **Section Templates** Settings

This contains the rendering template for each section identified in the sections tab. For each section, this code specifies how the data will be rendered. Conditionally build HTML based on the dataset contents.

```
1 <script type="text/html" id="headerSection">
2   <% if (digestTemplate.getSetting('HEADER_IMAGE')) { %>
3     <!-- PREHEADER -->
4     <!-- Set text color to background color -->
5     <div style="display: none; visibility: hidden; overflow: hidden; opacity: 0; font-size: 1px; line-height: 1px; height: 0; max-height: 0; margin: 0; padding: 0;">
6       <!-- LOGO -->
7       
9     <!-- WRAPPER for Header -->
10    <%digestTemplate.embedHeaderFooter.header%>
11    <!-- End WRAPPER for Header -->
12  } %>
13 </script>
14
15 <script type="text/html" id="summarySection">
16   <% var $summary = $data.summary;
17   if ($summary) for (var $i in $summary) if ($summary.hasOwnProperty($i)) { var $favorite = $summary[$i]; if (!($favorite['showSummary'])) conti
18   <table width="100%" cellpadding="0" cellspacing="0" border="0" style="clear: both;">
19     <tbody>
20       <tr>
21         <td style="padding-bottom: 10px; border-bottom: 1px solid (%digestTemplate.theme.mainBorders%); padding-top: 15px;">
22           <p style="text-align: left; font-family: Helvetica, Arial, sans-serif; margin: 0px 4px 0px; font-size: 14px; color: (%digestTem
23         </td>
24       </tr>
25     </tbody>
26   </table>
27   <% if ($favorite['header']) { %>
28   <table width="100%" cellpadding="0" cellspacing="0" border="0" style="margin-top: 15px; margin-bottom: 10px; clear: both;">
29     <tbody>
30       <tr>
31         <td width="27" height="30px">
32           
33         </td>
34       </tr>
35     </tbody>
36   </table>
```

Change History

Send test email Preview Saved

From this tab, you can manipulate the input data, such as appending report data (**datasets**) to the core digest data object (**digestTemplate.data**), applying transformation functions, and initiating the rendering of email sections as defined in the sections code.

NOTE: If you duplicated one of the default templates, this tab is going to contain the HTML code that is used to define the structure of the email body. If you create a template from scratch, it is going to be empty.

Settings

This contains the rendering template settings.

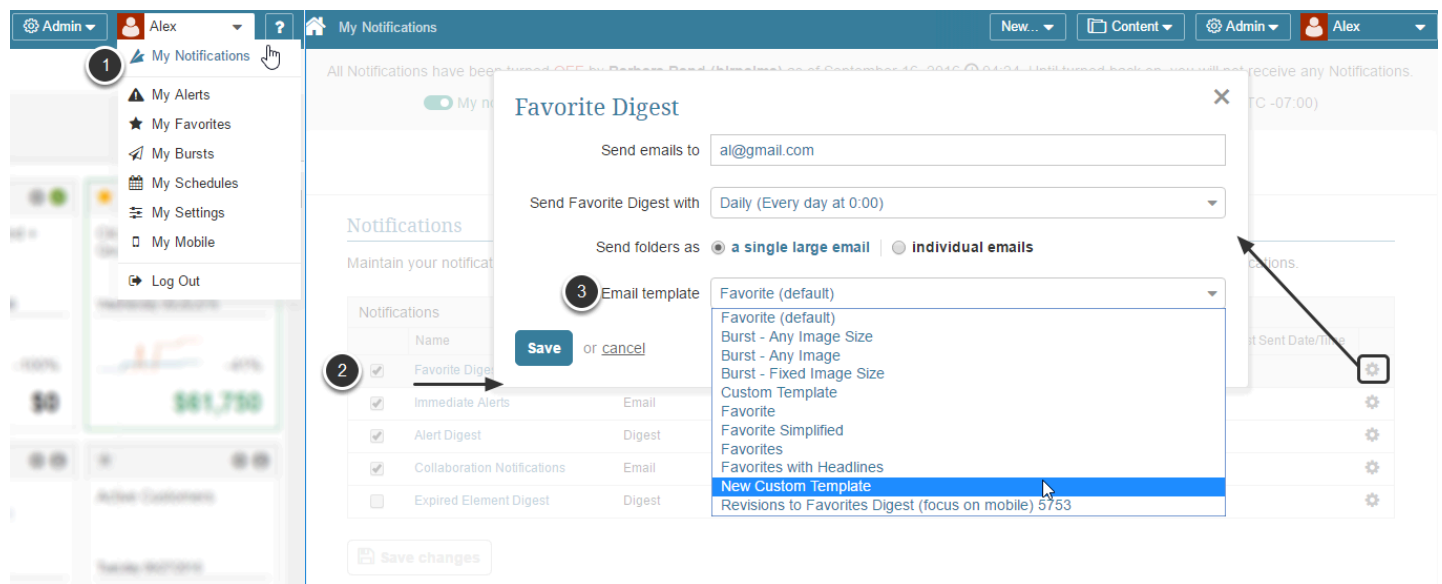
Assigned Settings				
Name	Display Name	Default Value	Available Values	
DIGEST_SHOW_UNSUBSCRIBE_LINK	Digest show unsubscribe link	Y	Y,N	⚙️ 🗑️
DIGEST_WIDTH_EXTERNAL_REPORT	Digest width external report	1200		⚙️ 🗑️
FAVORITE_DIGEST_EXTERNAL_REPORT_LINK_ON	Favorite digest external report link on	Chart_view	Chart_view,...	⚙️ 🗑️
FOOTER_IMAGE	Footer image	footer-imag...		⚙️ 🗑️
HEADER_IMAGE	Header image	header-ima...		⚙️ 🗑️
INCLUDE_ALERTS	Include alerts	Y	Y,N	⚙️ 🗑️
INCLUDE_ALL_XTD_VARIANCES	Include all xtd variances	N	Y,N	⚙️ 🗑️
INCLUDE_ANNOTATIONS	Include annotations	Y	Y,N	⚙️ 🗑️
INCLUDE_DEFAULT_HEADER_AND_FOOTER	Include default header and footer	N	Y,N	⚙️ 🗑️
INCLUDE_METRIC_SUMMARY_DETAIL	Include metric summary detail	Y	Y,N	⚙️ 🗑️
INCLUDE_METRIC_SUMMARY_TABLE	Include metric summary table	Y	Y,N	⚙️ 🗑️
MAX_DIGEST_ALERTS_DISPLAY	Max digest alerts display	3		⚙️ 🗑️
MAX_DIGEST_ANNOTATION_DISPLAY	Max digest annotation display	3		⚙️ 🗑️
MAX_DIGEST_EVENTS_DISPLAY	Max digest events display	3		⚙️ 🗑️
MAX_DIGEST_NOTES_DISPLAY	Max digest notes display	3		⚙️ 🗑️

The variables in the Settings tab allow you to customize as well as turn on/off features for your Digest and Burst Emails. Some examples are:

1. SHOW_UNSUBSCRIBE_LINK will determine whether or not an unsubscribe link will be included in the emails
2. HEADER_IMAGE can be changed to a file uploaded via the [Branding page](#)
3. MAX_DIGEST_ANNOTATION_DISPLAY will set the maximum number of annotations displayed within the Digest

Applying custom email digest templates

Favorites Digest



In order to use your newly created custom Email Digest Template in Favorites Digest you need to:

1. From the User menu, select "Favorites"
2. Choose your custom Email Digest Template from the Email template drop down list

To learn more about Favorites Digests, visit: [What are Favorites Digests?](#)

Bursts

Admin Alex ? My Bursts / New Burst (7) New... Content Admin Alex

My Notifications

- My Alerts
- My Favorites
- My Bursts
- My Schedules
- My Settings
- My Mobile
- Log Out

Content Schedule Subscriptions Send now

Name
New Burst

Send...

☒ favorite folders ☐ selected tiles

Send folders as ☒ a single large email ☐ individual emails

Click below to add your first folder
[+ New Folders](#)

Include ☒ only what's been updated since last email ☐ everything all the time

Render Overlays on Burst Images ☒ yes ☐ no

Email template

- Burst - Any Image Size
- Burst - Any Image
- Burst - Fixed Image Size
- Custom Template
- Favorite
- Favorite Simplified
- Favorites
- Favorites with Headlines
- New Custom Template**
- Revisions to Favorites Digest

Save Save & Review Save & Create Cancel

Templates are applied to your Bursts in the same manner as favorites.

1. From the User menu, select "Bursting"
2. Choose your custom Email Digest Template from the Email template drop down list

To learn more about Bursts, visit: [Setting up a new report burst](#)

2.3 Applying headers/footers to Email Templates

Metric Insights provides an ability to add various custom headers and footers to email templates. This article provides step-by-step instructions on the procedure of applying images for custom headers and footers to the email digest templates.

NOTE: Images for custom headers and footers can be uploaded only at the [Branding](#) page.

Headers and footers uploaded at the **Branding** page may be applied to email digest templates in several ways:

- in the **Email Digest Templates** page at the **Admin** menu
- via settings in the **Favorites** Editor

IMPORTANT: Settings defined in the **Favorites** Editor have a higher priority than other settings defined at **Email Digest Templates**

1. Email Digest Templates header/footer settings

The screenshot shows the Metric Insights Admin interface. The top navigation bar includes 'Content', 'Admin', and a user profile 'Alex'. The left sidebar contains a menu with 'Email Templates' highlighted. The main content area is titled 'Email Templates' and contains a table with the following data:

Name	Default Template
Alert (locked)	Alert
Alert Simplified (locked)	
Burst - Any Image Size (locked)	
Burst - Any Image	
Burst - Fixed Image Size (locked)	
Collaborative (locked)	Collaborative
Data Trust (locked)	
Examples for header/footer templates (locked)	
Favorite (locked)	Digest
Favorite Simplified (locked)	
Favorites with Headlines (locked)	
Out of disk space (locked)	
Remove collaboration item (locked)	
Request Permission (locked)	
Revisions to Favorites Digest	
Share Element/Discussion/Comment (locked)	
Welcome/Recovery (locked)	

A '+ New Template' button is located at the bottom of the table.

Go to *Admin > Email Templates*

1. Every kind of notification has its own default template, which is marked as **locked**. You are free to create your own templates and apply them to different type of email notifications.
2. Select one of the templates

1.1. Default header and footer

The screenshot shows the 'Digest Templates / New Custom Template' interface. A table lists various templates. The row 'INCLUDE_DEFAULT_HEADER_AND_FOOTER' is highlighted. A pop-up window titled '2 Edit Digest Template Setting' is open, showing the 'Name' as 'Include default header and footer' and the 'Default Value' as 'Y'. A red arrow points from the highlighted row to the pop-up. The pop-up has 'Save' and 'cancel' buttons.

Name	Description	Y	N	Y,N	Settings	Delete
DIGEST_SHOW_UNSUB						
DIGEST_WIDTH_EXTEN						
FAVORITE_DIGEST_EX						
FOOTER_IMAGE						
HEADER_IMAGE	Header image					
INCLUDE_ALERTS	Include alerts			Y,N		
INCLUDE_ALL_XTD_VARIANCES	Include all xtd variances	N		Y,N		
INCLUDE_ANNOTATIONS	Include annotations	Y		Y,N		
INCLUDE_DEFAULT_HEADER_AND_FOOTER	Include default header and footer	Y		Y,N		
INCLUDE_METRIC_SUMMARY_DETAIL	Include metric summary detail	Y		Y,N		
INCLUDE_METRIC_SUMMARY_TABLE	Include metric summary table	Y		Y,N		
MAX_DIGEST_ALERTS_DISPLAY	Max digest alerts display	3				

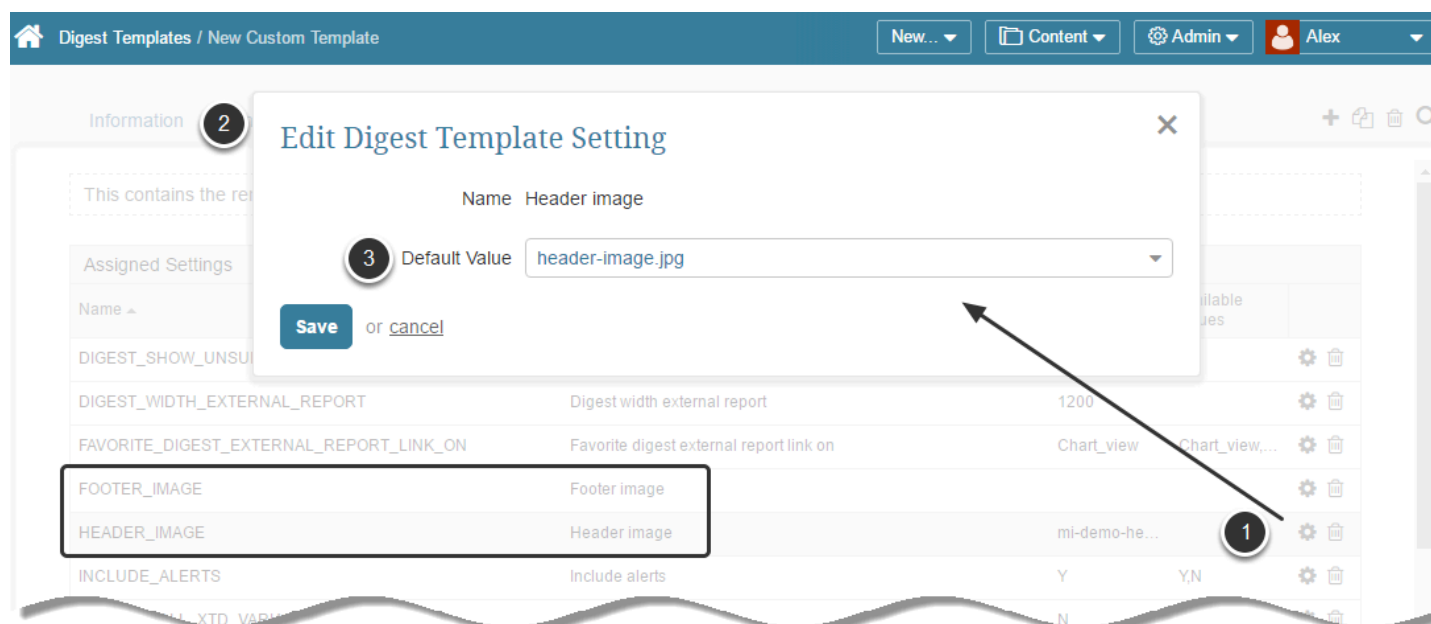
Open the *Settings* tab.

1. You can assign default Metric Insight header and footer to the selected template. To do that, click the **Edit Setting (Gear)** icon in the corresponding row
2. The *Edit Digest Template Setting* pop-up opens
3. Set the **Default Value** field to 'Y'

Save your settings.

NOTE: Click **Preview** to make sure your email is going to look the way you expect it to.

1.2. Custom header and footer



You can specify custom header and footer for any template.

NOTE: Images for custom headers and footers can be uploaded only via the **Branding** page at *Admin > Utilities > Branding*.

To assign a custom footer image:

1. Click the **Edit Setting (Gear)** icon in the FOOTER_IMAGE row
2. The *Edit Digest Template Setting* pop-up opens
3. Type in/paste the name of the header image file previously uploaded in the Branding page

Click **Save** to apply the image file.

NOTE: If you manually type in the name of a picture, double-check whether it is spelled correctly to avoid typos/mistakes, otherwise default email header and footer are going to be applied.

1.3. Once you have added header and footer images, click 'Preview' to see the changes made

This contains the rendering template settings.

Name	Display Name	Default Value	Available Values	
DIGEST_SHOW_UNSUBSCRIBE_LINK	Digest show unsubscribe link	Y	Y,N	
DIGEST_WIDTH_EXTERNAL_REPORT	Digest width external report	1200		
FAVORITE_DIGEST_EXTERNAL_REPORT_LINK_ON	Favorite digest external report link on	Chart_view	Chart_view,...	
FOOTER_IMAGE	Footer image	footer-imag...		
HEADER_IMAGE	Header image	header-ima...		
INCLUDE_ALERTS	Include alerts	Y	Y,N	
INCLUDE_ALL_XTD_VARIANCES	Include all xtd variances	N	Y,N	
INCLUDE_ANNOTATIONS	Include annotations	Y	Y,N	
INCLUDE_DEFAULT_HEADER_AND_FOOTER	Include default header and footer	N	Y,N	
INCLUDE_METRIC_SUMMARY_DETAIL	Include metric summary detail	Y	Y,N	
INCLUDE_METRIC_SUMMARY_TABLE	Include metric summary table	Y	Y,N	

Buttons: Send test email Preview Saved

NOTE: Make sure that the names of images for header and footer are displayed in the **Default Values** column and click **Preview** to find out how the email is going to look.

2. Favorites Editor email digest settings

1 My Notifications

2 Settings

3 Send Favorite Digest with: Daily (Every day at 0:00)

4 Email template: Favorite (default)

Buttons: Save or cancel

Name	Type	Frequency	Email	Settings
<input checked="" type="checkbox"/> Favorite Digest	Digest	Daily (Every day at 0:00)	al@noemail.com	
<input checked="" type="checkbox"/> Immediate Alerts	Email	Daily (Every day at 0:00)	al@noemail.com	
<input checked="" type="checkbox"/> Alert Digest	Digest	Daily (Every day at 0:00)	al@noemail.com	
<input checked="" type="checkbox"/> Collaboration Notifications	Email	Daily (Every day at 0:00)	al@noemail.com	
<input type="checkbox"/> Expired Element Digest	Digest	Daily (Every day at 0:00)	al@noemail.com	

Buttons: Save changes

Individual Users may override the Digest Template for any of their Favorite folders. In order to do that, follow these steps:

1. Go to *My Notifications* page
2. In the Notifications grid find Favorite Digest row and click the **Gear** icon to edit settings
3. **Send folder as** field should be set to 'individual emails' value to allow sending individual emails with customized settings for each selected folder
4. Define general **Email template** for all digest emails

Save your changes.

NOTE: Your Administrator must have already added a new image in the Branding page and given you its name in order to follow these steps:

2.1. Select the Favorite folder that should be customized

My Favorites

New... Content Admin Alex

My notifications are ON for emails, SMS and Phone Alerts • Friday - September 30, 2016 System 05:30 PDT(UTC -07:00)

Notifications Alerts Favorites Burst Schedules Settings Mobile

Maintain content, schedule and sharing of your Favorite folders

Favorite Folders

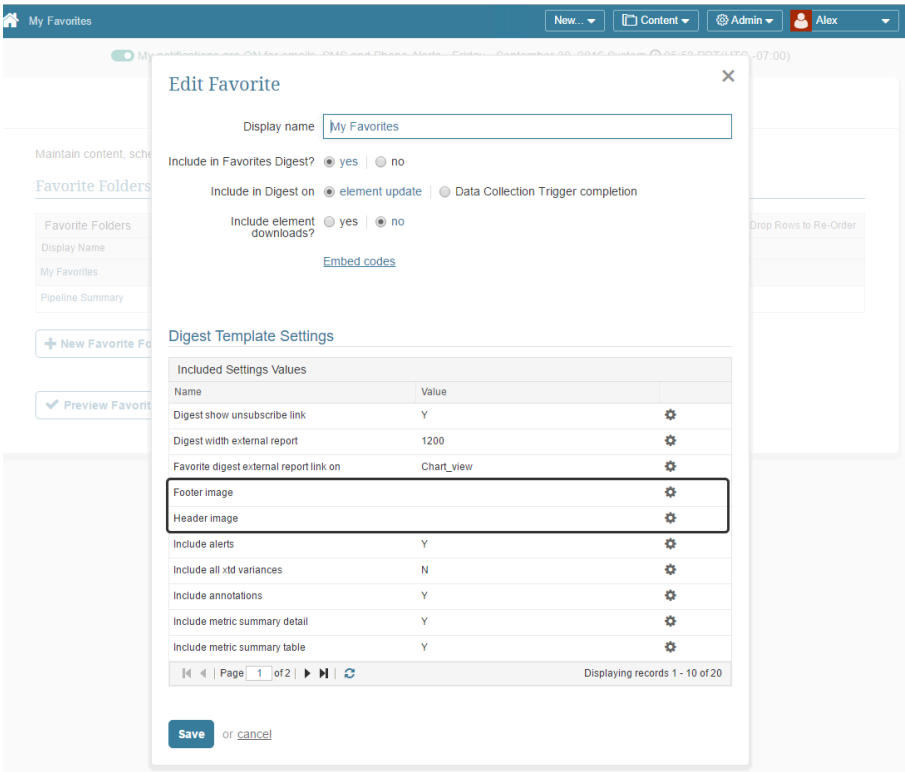
Display Name	Number of Elements	Include in Digest	Actions
My Favorites	0	✓	⚙️ 📄 📧 🗑️
Pipeline Summary	9		⚙️ 📄 📧 🗑️

+ New Favorite Folder

✓ Preview Favorites Digest Send now This will send an email to al@noemail.com

1. Open the *Favorites* tab
2. Select the Favorite folder that should be customized and click the **Gear** icon to change its settings

2.2. Once saved, header and footer with their provided values will be added to the 'Digest Template Settings' grid



To choose the pictures that should serve as a footer and header, click the **Gear** icon in the respective row and select the required file from the drop-down list.

2.4 Simplified Email Templates

If you have a legacy email client like Lotus Notes or have any problem rendering our emails, Metric Insights team have taken care of it. You can enable simplified email templates for:

- [Alerts](#): set by Admin and not available for User selection
- [Favorites](#)
- [Bursts](#)

This article gives step-by-step instructions on how to do it for each type of notification emails.

Alerts - set the Default template to Alert Simplified (Admin only)

The screenshot shows the 'Email Templates' configuration page in the Metric Insights Admin interface. The page is titled 'Email Templates' and has a search bar. Below the header, there is a table of templates. The table has two columns: 'Name' and 'Default Template'. The 'Alert' template is highlighted with a red box and a red circle with the number 1. The 'Alert Simplified' template is also highlighted with a red box and a red circle with the number 2. Below the table, there is a section for 'Digest Templates / Alert Simplified'. This section has a message: 'This is a standard template and can not be modified. Duplicate this template to make modifications.' Below this message, there are tabs for 'Information', 'Email Structure', 'Transformation & Rendering', 'Section Templates', and 'Settings'. The 'Email Structure' tab is selected. The 'Email Structure' tab shows a form for editing the template. The form has a dropdown for 'Email type' set to 'Alert', a text field for 'Name' set to 'Alert Simplified', and a radio button for 'Default Alert Template?' set to 'yes'.

Access *Admin > Email Templates*

1. In the *Email Templates* list find alternatives for Alert templates: **Alert** and **Alert Simplified** or [custom ones](#).
2. The one that is marked as 'Alert' a default one. Only one alert template can serve as default.

The new **Alert Simplified** email template is a safe option that can be used for different legacy email clients.

- To make **Alert Simplified** a default template, open it and make sure that **Default Alert Template?** field is set to 'yes'.

Save the changes made.

Favorites

The screenshot shows the 'My Notifications' interface. At the top, there's a header with 'My Notifications', 'New...', 'Content', 'Admin', and a user profile 'Alex'. Below the header, a status bar indicates 'My notifications are ON for emails, SMS and Phone Alerts - Thursday - November 17, 2016 System 14:04 UTC(UTC +00:00)'. A navigation bar contains icons for Notifications, Alerts, Favorites, Bursts, Schedules, Settings, and Mobile. The main content area is titled 'Notifications' and includes a sub-header 'Maintain your notifications and Burst subscriptions; use the checkboxes to toggle on/off, and gears to configure. [Learn more](#) about Notifications.' Below this is a table with columns: Name, Type, Schedule, Send To, and Last Sent Date/Time. The table lists two notifications: 'Favorite Digest' and 'Immediate Alerts'. A red circle with the number '1' points to the 'Favorite Digest' row. A red circle with the number '2' points to the gear icon in the 'Last Sent Date/Time' column for 'Favorite Digest'. A red circle with the number '3' points to the 'Favorite Digest' pop-up configuration window. A red circle with the number '4' points to the 'Email template' dropdown menu in the pop-up, which is set to 'Favorite Simplified'. The pop-up also includes fields for 'Send emails to', 'Send Favorite Digest with', 'Send folders as', and 'Save' or 'cancel' buttons.

A new template has been also added to Favorites Digest. To enable it go to *My Notifications*:

- Select **Favorite Digest** in the *Notifications* grid.
- Click the **Gear** icon in the respective row.
- The **Favorite Digest** pop-up opens.
- In the **Email template** field select "Favorite Simplified" from the drop-down list.

Save the changes made.

Bursts

My Bursts

New... Content Admin Alex

My notifications are ON for emails, SMS and Phone Alerts • Thursday - November 17, 2016 System 14:38 UTC(UTC +00:00)

Notifications Alerts Favorites **Bursts** Schedules Settings Mobile

Bursts

Maintain your Bursts and subscribe to other Users' distribution of a group of tiles or folders

Name	Schedule	Recipients	Last Sent on	Disabled Schedule	Disabled Bursts
Burst	Every 12hrs, after 1st min & after 10-minute...	6			
DT20-digest	DT20-notification-schedule	4	2016-11-17 06:04:18		New Data Trust Flow (20)

+ New Burst Send now

My Bursts / Burst

New... Content Admin Alex

Content Schedule Subscriptions Send now

Name

Burst

Send...

favorite folders selected tiles

Selected Tiles

Display Name
Monthly Unsatisfied Survey Responses

+ Add tiles

Include only what's been updated since last email everything all the time

Render Overlays on Burst Images yes no

Email template

Favorite Simplified

An email template compatible with legacy email clients, such as Lotus Notes has been set up for Bursts as well. To apply it to Burst emails:

1. Go to *My Bursts*
2. Select an existing burst or create a new one
3. In the *Content* tab, **Email template** field select 'Favorite Simplified' from the drop-down list

Do not forget to **Save** your changes to apply new settings.

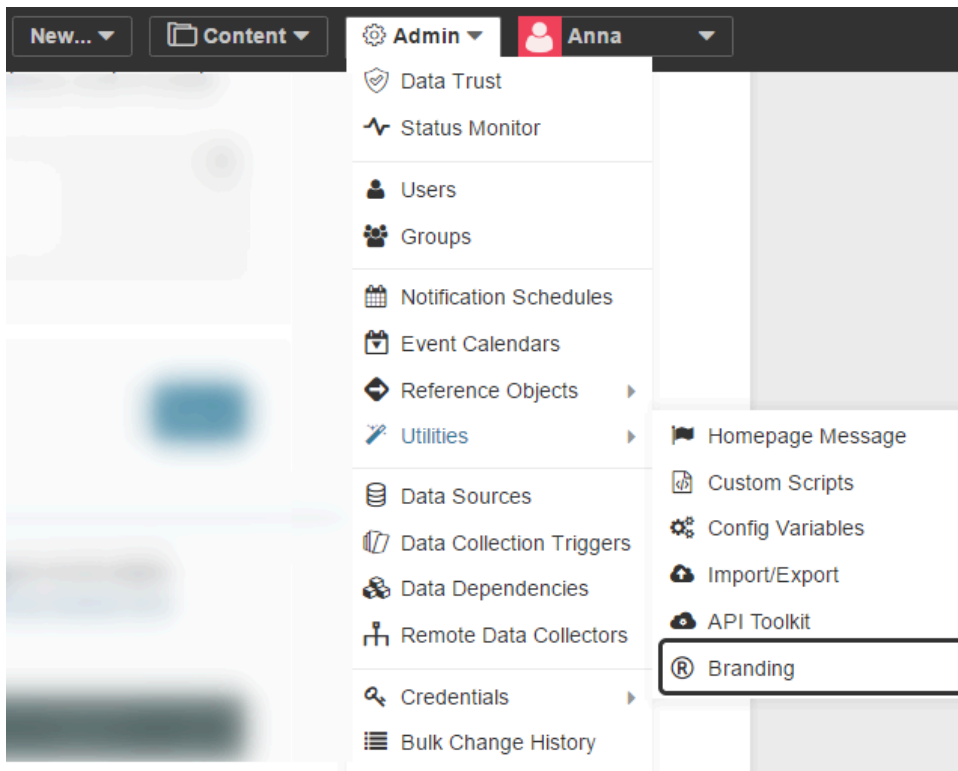
3. Configuring Metric Insights

3.1 Using the Branding feature

The **Branding** feature serves to customize your Metric Insights instance by adding:

- your own logo to the *Home Page*, slide shows and certain emails
- multiple custom headers and footers to your email templates

1. Access Admin > Utilities > Branding



2. Click 'Upload logo'

Branding

New... Content Admin Anna

Branding

Your logo will appear on the Homepage, in slideshows, and in emails.

Logo

1 Upload logo

2 Email from name Your reports team Save

1. Click 'Upload logo' to access the *New Logo* pop-up
2. Enter the **name** you would like to include as the 'digest or burst sent from' on your emails

2.1. Click 'Browse' and select a picture with your logo from your local machine

New Logo

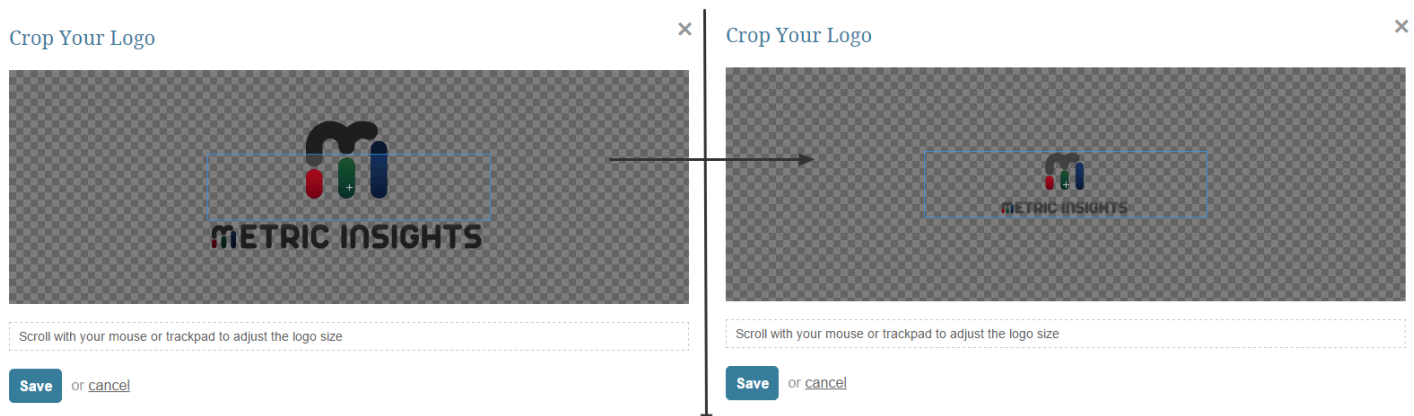


Upload your custom logo

Choose file Browse

The logo file size must be less than 5mb. Supported file types: JPEG, PNG, GIF, BMP.

2.2. Adjust the logo size/position to fit in the cropping section as shown in the picture below



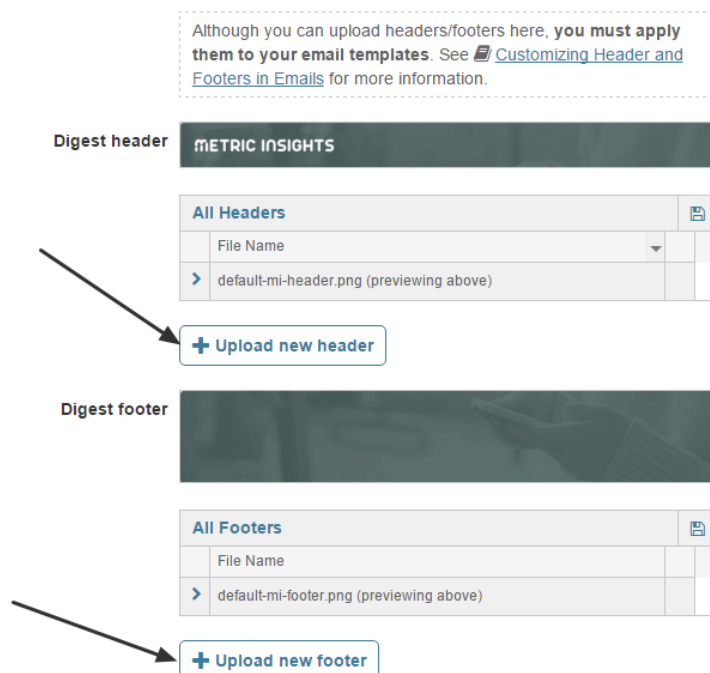
- Move your logo to position it in the preferred part of the cropping section
- To adjust the picture size, scroll your mouse or trackpad

NOTE: Your logo may or may not have a background.

Click **Save** to keep the changes made.

3. Providing custom Email headers & footers

Email headers & footers




You can upload and preview custom headers and footers directly on this page

3.1. Upload header/footer image

Digest headers & footers

Although you can upload headers/footers here, **you must apply them to email templates** via [Email Headers/Footers](#).

Digest header




All Headers

File Name		
default-mi-header.png		
> header_image.jpg (previewing above)		

+ Upload new header

Digest footer



All Footers

File Name		
default-mi-footer.png		
> footer_image.jpg (previewing above)		

+ Upload new footer

- Footer and header images are uploaded in the same way as a logo
- To preview the header or footer with a new picture, click the row with uploaded picture in **All Headers** list or **All Footers** list

NOTE: Your header image should be 960px wide - if it's bigger or smaller it will be resized, maintaining your original proportions. There are no constraints on height. Supported file types: JPEG, PNG, GIF, BMP

4. Apply newly added headers/footers to email templates

To learn how to apply your uploaded custom headers/footers to certain email templates, refer to this article: [Applying custom headers/footers to email templates](#)

3.2 Setup email

Previously, an Administrator was required to ssh into the Metric Insights server and configure the following variables in the `/var/www/iv/engine/config/const.php` file in order to have the system communicate with your organization's mail server. Today this can be done via the *Config Variables table*: See [here](#) for general information on updating Config variables

The sample values shown below must be replaced with the correct values for your email servers.

1. Configure outbound messages for notifications and digests

Key: Changed

Email Connectivity ▼

Config Variables			
Name	Assigned value	Valid values	
SEND_MAIL_SERVER	mailer.metricinsights.com		⚙
SEND_MAIL_AUTH	plain	,plain,login,crammd5,ntlm	⚙
SEND_MAIL_LOGIN	adam		⚙
SEND_MAIL_PASSWORD	*****		⚙
SEND_MAIL_SSL	tls	,ssl,tls	⚙
SEND_MAIL_PORT	587		⚙
SEND_MAIL_FROM	adam@app.metricinsights.com		⚙
SEND_MAIL_FROM_NAME	MI UAT		⚙
SEND_MAIL_REPLY	adam@app.metricinsights.com		⚙

- **SEND_MAIL_SERVER** -the SMTP server you would like to use
- **SEND_MAIL_AUTH** - the authentication method that the server expects. **If authentication is not required, use ''**; otherwise, the method should be one of: **'plain'**, **'login'**, **'crammd5'**. If your server requires authentication, most of the time it will be **'login'**. *Note: these values are case sensitive.*
- **SEND_MAIL_LOGIN** - if authentication is required, the username to use
- **SEND_MAIL_PASSWORD** - if authentication is required, the password to use
- **SEND_MAIL_SSL** - the method of SSL encryption to use. Set to **'ssl'** if you are connecting to SMTPS, otherwise use **'tls'** if the smtp server requires you to start tls over an unencrypted

smtp connection. An empty value will cause SSL to not be used. Note: these values are case sensitive.

- **SEND_MAIL_PORT** - the SMTP port that the server listens on. The default ports are 25 for SMTP, and 465 for SMTPS.
- **SEND_MAIL_FROM** - the address from which to send notifications
- **SEND_MAIL_FROM_NAME** - A sender name from whom notifications are sent
- **SEND_MAIL_REPLY** - the reply to email address to be used For collaboration, it's recommended to set this to the RECV_MAIL_LOGIN address so replies to comments are automatically logged

If you encounter any problems you can [troubleshoot your email connectivity for outbound](#)

2. Configure an incoming mailbox for sharing / collaboration

Config Variables

Key: Changed Email Connectivity

Name	Assigned value	Valid values	
SEND_MAIL_SERVER	mailer.metricinsights.com		⚙
SEND_MAIL_AUTH	plain	,plain,login,crammd5,ntlm	⚙
SEND_MAIL_LOGIN	adam		⚙
SEND_MAIL_PASSWORD	*****		⚙
SEND_MAIL_SSL	tls	,ssl,tls	⚙
SEND_MAIL_PORT	587		⚙
SEND_MAIL_FROM	adam@app.metricinsights.com		⚙
SEND_MAIL_FROM_NAME	MI UAT		⚙
SEND_MAIL_REPLY	adam@app.metricinsights.com		⚙
RECV_MAIL_SERVER	mailer.metricinsights.com		⚙
RECV_MAIL_AUTH	plain	,plain,login,crammd5	⚙
RECV_MAIL_LOGIN	adam		⚙
RECV_MAIL_PASSWORD	*****		⚙
RECV_MAIL_SSL	SSL	,SSL,TLS	⚙
RECV_MAIL_PORT	995		⚙

Page 1 of 1

Displaying records 1 - 15 of 15

Generate const.php

In order to make use of the collaboration features of Metric Insights, you will need an account to automatically process replies.

- **RECV_MAIL_SERVER** - the POP3 server to use to get email from
- **RECV_MAIL_LOGIN** - the user name that will receive the notifications
- **RECV_MAIL_PASSWORD** - the password to use to get to the POP3 mailbox
- **RECV_MAIL_SSL** - what type of SSL encryption to use. Set to '**SSL**' if you are connecting to POP3S, otherwise use '**TLS**' if the POP3 server requires you to start TLS over an unencrypted POP3 connection. An empty value will cause SSL to not be used. (note: these values are case sensitive)

- **RECV_MAIL_PORT** - the port on which the POP3 server is listening. The default is 110 for pop3, and 995 for pop3s.

Be sure to Generate const.php file!; i.e., Save changes

If you encounter any problems you can [troubleshoot your email connectivity for inbound](#)

3.3 How to Set the System Clock

This article explains the commands necessary to set your system's clock. Login to Server and run one of the following commands

1. When Metric Insights has outbound access to the internet, use NTP

1.1. Commands for CentOS instances:

```
service ntpd stop
ntpd -gq
service ntpd start
```

1.2. Commands for Debian instances:

```
service ntp stop
ntp -gq
service ntp start
```

2. When Metric Insights does not have outbound access to the internet

Use a Shell prompt:

```
date -s "OCT 7 2015 18:00:00"
```

inserting the correct data and time

3.4 Generating Empty Dataset Reports (when there is No Data)

By default, Reports containing no data are not included in Bursts and Digests. Sometimes, especially when you are reporting changes or defects in data, you will want to display *a report that contains no data*.

If you like, Metric Insights can generate an empty report in these circumstances. This article explains how to do this.

If **"Include Empty Report in Distribution" = yes** then an empty Report can be sent via:

- Burst to email
- Burst to chat apps (Slack/Teams)
- Favorites Digest
- Share Now

Access Dataset Report Editor > Distribution tab

The screenshot shows the 'Report Editor / Tableau Dataset Report' interface. The top navigation bar includes a home icon, a search bar, a notification bell with '107', a tab count '2', a 'New...' dropdown, and a 'Copy' icon. Below this is a secondary navigation bar with tabs: 'Info', 'Data', 'Content', 'Distribution' (highlighted in green), 'Associations', 'Documents', and 'Advanced'. The main content area is titled 'Distribution Content' and contains a green toggle switch for 'Include Empty Report in Distribution' with 'yes' selected. Below this is a dropdown menu for 'Attach to Email' currently set to 'None'.

Note: The following are exceptions for sending empty Reports:

- Options **"Shared Drive"** and **"FTP"** are not supported for sending empty Reports
- Requested attachments will **not** be sent with empty Reports

3.5 Generating Legacy Reports When There Is No Data

Sometime the report that you want to generate contains no data. If you like, Metric Insights can generate an empty report in these circumstance. This article explains how.

💡 For newer versions of Reporting (Dataset Reports) see [Generating Dataset Reports When There Is No Data](#)

1. Configure the Report's Advanced Settings

Reports / Day Order Volume by Channel

Advanced

Ownership & Access

Permissions

Business owner: Alex Kennedy (bkADMINOFF@nomail.com)

Technical owner: Alex Kennedy (bkADMINOFF@nomail.com)

This Report is: ☐ public | ☒ internal

Display Options

Show numeric NULL values as: ☐ blank | ☒ zero

Flip rows and columns: ☐ yes | ☒ no

Drill to Report for date not generated: ☐ shows missing instance message | ☒ generates Report instance

1 If no data fetched for Report instance: ☐ skip generation | ☒ generate empty Report

2 Empty instance message: No data available for this report date.

Other

Expire if latest data is more than: 2 day

Embed codes

In the *Report Editor / Advanced* tap:

1. Toggle **If no data fetched for Report Instance** to **"generate empty report"** to ensure that Reports are generated even when there is no data.
2. Enter the message that you would like to be displayed in **Empty instance message** NOTE: A system-wide default message is set in **EMPTY_REPORT_INSTANCE_MESSAGE** variable in **Config Variables** screen

IMPORTANT: If you set **If no data fetched for report instance** to “**skip generation**”, then no report instance will be generated when there is no data. If the report is a single-instance report (i.e., if **Keep-History** is set to "No"), then there will simply be no report. If the report has **Keep History** set to "Yes", then the last non-empty instance of the report will be displayed.

3.6 Debug Logging

FETCH_REQUEST_BROKER_LOGGING

The `FETCH_REQUEST_BROKER_LOGGING` parameter will tell the application to log information related to fetch requests in the directory named **`/var/www/iv/data/temp/broker/`**. This information is useful for debugging problems with data collection or plugin errors. Here's what the entry in `const.php` looks like:

```
define('FETCH_REQUEST_BROKER_LOGGING', 0);
```

The '0' indicates that logging is currently turned off. If you want to turn it on, just change the value to '1'. Log files will be named according to the date, i.e., `YYYY-MM-DD.log`

EMAIL_LOGGING

The `EMAIL_LOGGING` parameter will tell the application to log details of automatic email messages in the directory named **`/var/www/iv/data/temp/email/`**. This information is useful for debugging problems related to automatic emailing. Here's what the entry in `const.php` looks like:

```
define('EMAIL_LOGGING', 'N');
```

The 'N' indicates that logging is currently turned off. If you want to turn it on, just change the value to 'Y'. Log files will be named according to the date, i.e., `YYYY-MM-DD.log`

Note that a user should only be delivered favorite content once per day. Therefore, you may see debug entries similar to this example: "No Digest sent to <email address> since no favorite content has been updated since prior digest was sent."

3.7 Set Timezone (Virtual Appliance or CentOS/Redhat Linux)

In order for alerts to work properly, you will need to set your Metric Insights server to the same timezone as your data source. (If that is not possible, please contact us for assistance. We will show you how to adjust the Metric Insights config parameters so that your alerts work properly.)

This article assumes that you are running Metric Insights as a virtual appliance obtained from our website. It also covers situations where you have installed Metric Insights on your own server or in your own Amazon EC2 environment, using either CentOS or Redhat Linux.

Note: If your instance of Metric Insights is hosted within our EC2 environment, we will adjust your timezone settings ourselves -- you don't need to do anything here.

1. Login to the Metric Insights server

Enter the following command from a shell prompt:

```
ssh root@000.000.000.000
```

...replacing '000.000.000.000' with the IP address for your Metric Insights server.

Note: This step assumes that you are running Metric Insights as a virtual appliance. If you have installed Metric Insights on your own server or in your own Amazon EC2 environment, please contact your network administrator for access to the server.

2. Create symlink between /etc/localtime and timezone file

Run the following commands from a shell prompt:

```
rm /etc/localtime
```

```
ln -s /usr/share/zoneinfo/Your/Timezone /etc/localtime
```

Example (Pacific time):

```
ln -s /usr/share/zoneinfo/US/Pacific /etc/localtime
```


3. Restart MySQL

You will need to restart MySQL in order to synchronize the application database with the server. Run the following command from a shell prompt:

```
service mysqld restart
```

4. Create permanent timezone entry in appropriate location (so it will survive reboot)

The preceding steps will re-set the timezone for your Metric Insights server. In order for your adjustments to survive a reboot, however, you will need to create a permanent entry in the file named **/etc/sysconfig/clock**. Edit the file and create (or modify) the following entry:

```
ZONE="Your/Timezone"
```

Example (Pacific time):

```
ZONE="US/Pacific"
```


3.8 Set Timezone (Debian Linux)

In order for alerts to work properly, you will need to set your Metric Insights server to the same timezone as your data source. (If that is not possible, please contact us for assistance.)

DEFAULT_TIMEZONE is now set from `/etc/mi/insight.conf`. It does not need to be modified in `const.php` or `config_variable`, anymore. It may come a time where we want users to be able to set the timezone from the web UI. For now, however, that is not possible.

For now, a customer can set the timezone either by **manually editing `insight.conf`**,

OR

by running ``mi-db-mgr timezone <tz_name>``.

The last method should be used as it takes care of modifying **`/etc/localtime`** as well as sets the timezone on boot.

3.9 Supported Time and Date Formats

Metric Insights supports various types of Date formats. If a column of uploaded data contains dates in one of these formats, the system recognizes it and can perform different aggregations based on year, month, quarter, etc.

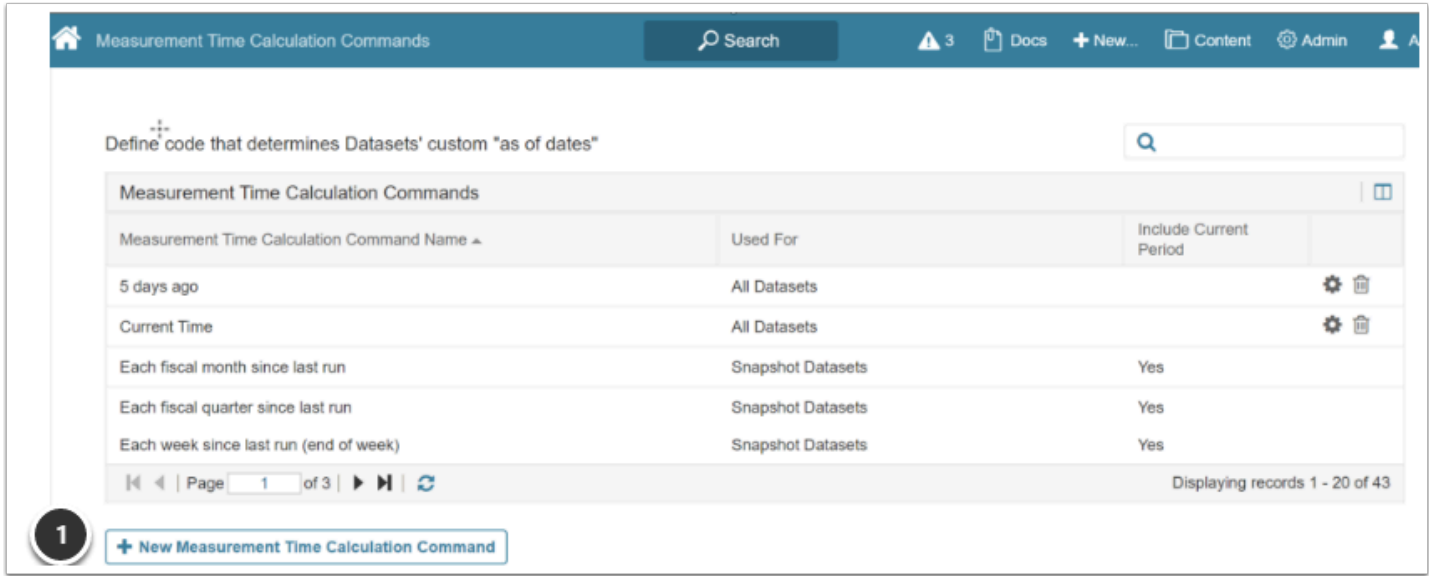
Date			Time		
Directive	Description	Example	Directive	Description	Example
%Y	Upper case: 4 digit year	1997, 2003, 2016	%H	Hour (24-hour clock)	00, 01, ..., 23
%y	Lower case: 2 digit year	97, 03, 16	%I	Hour (12-hour clock)	01, 02, ..., 12
%m	Month as a decimal value.	07, 12	%M	Minute	
%B	Month as a full name.	June, December	%S	Second	00, 01, ..., 59
%b	Month as an abbreviated name.	Jun, Dec	%p	AM or PM. Is used with 12-hour clock.	AM, PM
%d	Day of the month as a decimal number.	03, 15, 24	%Z	Time zone name	UTC, EST

Format	Example
'%Y-%m-%d %H:%M:%S' (or '%Y-%m-%d %H:%M')	2017-03-27 15:25:31 (2017-03-27 15:25)
'%Y/%m/%d %H:%M:%S' (or '%Y/%m/%d %H:%M')	2017/03/27 15:25:31 (2017/03/27 15:25)
'%Y.%m.%d %H:%M:%S' (or '%Y.%m.%d %H:%M')	2017.03.27 15:25:31 (2017.03.27 15:25)
'%m/%d/%Y %H:%M:%S' (or '%m/%d/%Y %H:%M')	03/27/2017 15:25:31 (03/27/2017 15:25)
'%m/%d/%y %H:%M:%S' (or '%m/%d/%y %H:%M')	03/27/17 15:25:31 (03/27/17 15:25)
'%Y-%B-%d %H:%M:%S' (or '%Y-%B-%d %H:%M')	2017-March-27 15:25:31 (2017/03/27 15:25)
'%m/%d/%Y %l:%M:%S %p' (or '%m/%d/%Y %l:%M %p')	03/27/2017 3:25:31 PM (03/27/2017 3:25 PM)
'%B %d, %Y %l:%M %p'	March 27, 2017 3:25 PM
'%B %d, %Y %l %p'	March 27, 2017 3 PM
'%a %b %d %H:%M:%S %Z %Y'	Mon Mar 27 15:25:31 EST 2017
'%Y-%m-%d'	2017-03-27

'%Y/%m/%d'	2017/03/27
'%Y.%m.%d'	2017.03.27
'%Y-%b-%d'	2017-Mar-27
'%m-%d-%Y'	03-27-2017
'%m-%d-%y'	03-27-17
'%m/%d/%Y'	03/27/2017
'%m/%d/%y'	03/27/17
'%d %b %Y'	27 Mar 2017
'%d-%b-%y'	27-Mar-17
'%B %d %Y'	March 27 2017
'%B %d,%Y'	March 27, 2017
'%Y-%m'	2017-03
'%b, %Y'	Mar, 2017
'%b,%Y'	Mar,2017
'%b %Y'	Mar 2017
'%B, %Y'	March, 2017
'%B,%Y'	March,2017
'%B %Y'	March 2017

3.10 Add new Measurement Time Calculation

1. Access via Admin > Reference Objects > Meas. Time Calcs



1. Select [+ New Measurement Time Calculation Command] to open Editor

2. Input new Measurement Time Calculation Command data

Measurement Time Calc Commands / 5 days ago

Search

0 Docs

Info Measurement Interval Datasets

Name 5 days ago

Description 5 days ago

1 Data Collected Using Dashboard DB (SQL)

2 SQL statement SELECT NOW() - INTERVAL 5 DAY

3 Validate

Creator support

1. Data Collected Using - use the default value of Dataset DB (SQL)
2. SQL Statement - see example above
3. Validate the SQL statement

If Validate is successful, the SQL statement box will turn green, and you can [Save]

Your new option will now appear for Datasets and Metrics

Dataset Columns

Column Name	Reference Name	Type	Length	Display Mask
calendar_date	calendar_date	datetime		
name	name	text	700	
product_category	product_category	text	700	
product_subcategory	product_subcategory	text	700	
units	units	decimal		
sales	sales	decimal		

Advanced: Manage Indexes

Snapshot Dataset? (keep history) ☒ yes ☐ no

Can historical instances be backfilled? ☒ Today ☐ Yesterday ☒ 5 days ago

Set :measurement_time value 5 days ago

:measurement_time would be set to 2020-04-11 22:10:34 if data fetch is run now.

1. Select an option for "Set :measurement_time value"
2. ":measurement_time:" will reflect the setting as if update were run now

3.11 Default Name/Description Values for Elements

For more information see [Setting the Configuration Variables](#).

NOTE: The format of the *Metric Creation Wizard* is slightly different in Version 3.

UNDIMENSIONED_METRIC_NAME_PREFIX

Metric Insights defaults element's **Name** and **Description** based on the **Measurement Interval** and **Measure Of** settings selected when defining a new element; for example, 'Daily Incidents'.

The configuration parameter **UNDIMENSIONED_METRIC_NAME_PREFIX** in *Config Variables* is initially set to 'Total' and the Undimensioned Metrics are automatically named as follows:

Name = "Total Daily Incidents"


Description = "Total Daily Incidents Metric"

If the element is dimensioned, the **UNDIMENSIONED_METRIC_NAME_PREFIX** is removed and the defaults are as follows:

Name = "Daily Incidents"


Description = "Daily Incidents for a Hotels Metric"


NOTE: You can change the **UNDIMENSIONED_METRIC_NAME_PREFIX** to any values, such as 'Overall' or set it to the empty string.

 Metric Creation Wizard

New ▾

Admin ▾

 Alex ▾



Define the basics

Collect data

Preview & publish

1

2

3

Measurement Interval

Daily ▾

Measure of

Incidents ▾

Name this Metric

Total Daily Incidents

Category

Uncategorized ▾

Create dimensioned Metric?

☐ yes

☒ no

Next: collect data

or [cancel](#)

☒ Continue in our full (i.e. advanced) interface for Metric creation.

3.12 Set Hostname for Metric Insights server

This article will show you how to set the Hostname for your Metric Insights server.

Note: If your instance of Metric Insights is hosted within our EC2 environment, we will set the Hostname ourselves -- you don't need to do anything here.

1. Login to the Metric Insights server

In order to set the Hostname, you will need direct access to the Metric insights server. Please contact your network administrator for access. Once you are properly set up, you will need to connect to the server by ssh. If you are using a Mac, you can create an ssh connection from a terminal window. If you are using a Windows machine, you can use an ssh tunneling application such as PuTTY.

2. Set the Hostname for the server

Run the following command from a shell prompt:

Hostname your.Hostname.com

Example (demo.metricinsights.com):

Hostname demo.metricinsights.com

3. Make the server Hostname permanent

Step #2 will set the Hostname. In order for the new Hostname to survive a system reboot, however, you will need to create permanent Hostname entry in an appropriate location. This will depend upon your operating system.

3.1. Debian Linux

If you are running Metric Insights on Debian Linux, you can create a permanent Hostname entry by running the following command from a shell prompt:

echo 'your.Hostname.com' > /etc/Hostname

3.2. Virtual Appliance or CentOS / Redhat Linux

If you are running Metric Insights:

1. as a virtual appliance; or...
2. on CentOS Linux or Redhat Linux

...then you can create a permanent Hostname entry by editing the file named **/etc/sysconfig/network** and editing the **Hostname** parameter as follows:

Hostname=your.Hostname.com

4. Update Hostname config variable in Metric Insights

The screenshot shows the 'Config Variables' section of the Metric Insights interface. At the top, there's a navigation bar with 'Config Variables', 'New', 'Admin', and 'Bk' buttons. Below this, a 'Key: Changed' status is shown. A search bar contains 'hostname' and a dropdown menu is set to 'All'. A table titled 'Config Variables' lists the 'HOSTNAME' variable with its assigned value 'https://uma.metricinsights.com'. A gear icon next to the value indicates an edit option. At the bottom, a 'Save changes' button is highlighted.

Name	Assigned Value	Valid Values
HOSTNAME	https://uma.metricinsights.com	

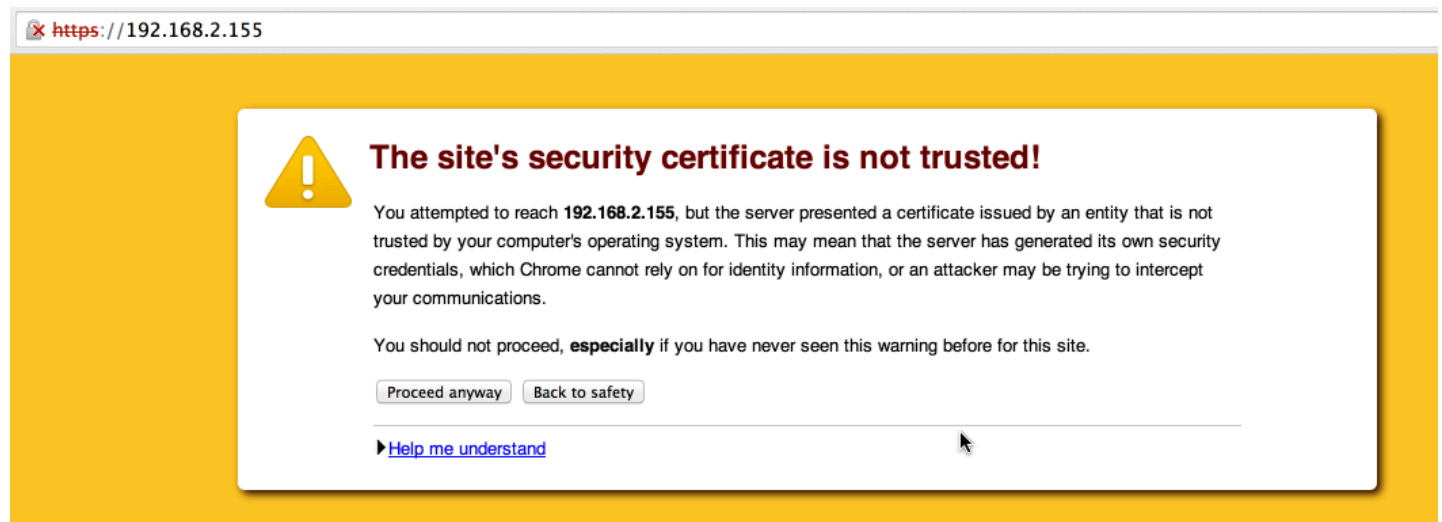
1. Locate the Config Variable for HOSTNAME using the search field
2. Click the edit icon to change hostname
3. Save changes

3.13 Set up Web SSL Certificate for Metric Insights

This applies to a Metric Insights instance hosted on VMware (CentOS linux) where the web SSL certificate has not been signed by certificate authority (CA). For Metric Insights instance hosted on Amazon EC2 (debian linux) see [this article](#).

This article discusses how to set up web SSL certificate for Metric Insights.

1. Web SSL Certificate warning



When running Metric Insights in VMware you will typically see this upon reaching the site for the first time. Depending on the browser you can just accept this SSL certificate warning and continue using Metric Insights. However, if you need to have a valid SSL certificate then this article walks you through the steps.

2. Create private key and CSR for the Metric Insights server

A private **key** and **certificate signing request** are necessary in order to generate an SSL certificate. You can create a private key / CSR pair with the following command on any linux machine:

(**Hint:** if you run this command directly on the MI server, you won't need to copy your private key file to the machine later)


```
openssl req -out your-machine.csr -new -newkey rsa:2048 -nodes -keyout your-machine.key
```

If you already have a private key that you'd like to use, please run the following command instead:

```
openssl req -out your-machine.csr -new -key your-machine.key
```

This command will ask you a bunch of questions about the machine. Most of the questions are self explanatory, however, please pay attention to the following points:

**** 'Common Name'**. The common name must be the fully qualified domain name for your server (the same that people will be putting in the address bar of their browser to access Metric Insights.)

**** 'Challenge Password'** - leave the challenge password blank, otherwise someone will need to enter a password whenever they start up Apache. Not good for automation.

Once you have answered all the prompts, your private key (your-machine.key) and CSR (your-machine.csr) files will have been created in the directory you ran the command in.

3. Request certificate from signing certificate authority

Provide the CSR file to your certificate authority. Many large companies will have their own internal signing authority while others will use one of the many commercial public trusted CAs on the market. Ask the IT or System Administrator in your organization if you are unsure.

Once the signing authority has approved / verified your request, they will issue you your SSL certificate. This can usually be downloaded in multiple file formats. Please get the certificate in Base64 encoded form (sometimes this is called the PEM format). It should look something like this:

```
-----BEGIN CERTIFICATE-----
MIIEVzCCA6egAwIBAgIDAxmRMA0GCSqGSIb3DQEBCwUAMEcxChAJBgNVBAYTA1VT
MRYwFAYDVQQKEw1HZW9UcnVzdCBJbmMuMSAwHgYDVQQDEXdSYXBpZFNNTTCTBTSEEy
NTYgQ0EgLSBHMzAeFw0xNTAzMTgwMDM2NDNaFw0xODAzMjAwMDI2NDJaMIGYMRMw
EQYDVQQLEwphVHdka4MTA5NzglMTEwLWYDVQQLEyhTZWUgd3d3LnJhcGlkc3NsLmNv
bS9yZXNvdXJjZXMvY3BzICHjKTE1MS8wLQYDVQQLEyZEB21haW4gQ29udHJvbCBW
YWxpZGF0ZWQgLSBSYXBpZFNNTTChSKTEdMBsGA1UEAwUKi5tZXRYaWNpbnNpZ2h0
cy5jb20wggeiMA0GCSqGSIb3DQEBAQUAA4IBDwAwggEKAoIBAQDhi0T6Rp8QxxYH
L2kGZDw0NJwBsCV+3ktj fMph8BWL6R1yG2YJJfp6WVZHJPqAzfX+JApqa9o6f7h8
Tw00CBXlZ0mm6KgHv18EN2U+IMu2pn8WmdTap6+D680ZmRtknbDYaxDyU/QJjV3u
/f7a/2X2uNCajfMtidhycubl4rvY8Mh96IDX8o5umM5PN4Fk43mjncbuWJPByzGk
kWiKgcBfJLUyywOYRirWN5lXZe8v3PE31KCjt1VIYB8071ru0y1Y8aodYCZMpoSo
rG3ec8mDIkwMeGcb6jxu2Hig8RLXcgM/CB8/Cob2UArSsYsJIJt4Bfy7/bFBRoZip
0s19f5AZAgMBAAGjggFgMIIBXDAfBgNVHSMEGDAWgBTDnPP800YINLvORn+gfFvz
4gjLWTBxBggrBgEFBQcBAQRLMEkwHwYIKwYBBQUHMAGGE2h0dHA6Ly9ndi5zeW1j
```



```
ZC5jb20wJgYIKwYBBQUHMAKGGmh0dHA6Ly9ndi5zeW1jYi5jb20vZ3YuY3J0MA4G
A1UdDwEB/wQEAWIfoDAdBgNVHSUEFjAUBggrBgEFBQcDAQYIKwYBBQUHAWIwMwYD
VR0RBCwwKoIUKi5tZXRYaWNpbnNpZ2h0cy5jb22CEm1ldHJpY2luc2lnaHRzLmNv
bTArBgNVHR8EJDAiMCCGhQACHhpodHRwOi8vZ3Yuc3ltY2IuY29tL2d2LmNybdAM
BgNVHRMBAf8EAjAAMEEGA1UdIAQ6MDgwNgYGGZ4EMAQIBMCwwKgYIKwYBBQUHAgEW
Hmh0dHBzOi8vd3d3LnJhcGlkc3NsLmNvbS9sZWdhbDANBgkqhkiG9w0BAQsFAAOC
AQEALfCsMphtWwz5XBWL0C5rwjsyNOBLwD9NEMNJc4ZQ2m9XrfF2TkL6Yb6UGqHZ
HY0yYfxxIjFyE3lcAafT/gjvsFr0uk/P3UNWUK+F1DEz5vpB4s8gFRlo/hq/hif0
f7B2iYqeRnhdMtG0GRKMY9kuXEOpqKK3A2soK5hecFBOWUK02D4/EfF0/Agxignw
nknfiuwAlOvh8obM1ldm2os8w15l8Lcmaw2658TwJyqDDBQeA4+0+1nM0mWknzA6
RCRo19zuMF7iRqTJgXL2zfFz8MbyxmgcUvSIMnYMNZM/+XYZJmlkuK8zG8D/Qt/W
kMBLr4VkmfXhSl093lSxA3tBwg==
-----END CERTIFICATE-----
```

4. Place your private key and certificate files on the server

Copy your new certificate file onto the Metric Insights server into the `/opt/mi/ssl` directory. Metric Insights expects these files to be named in the following way:

- `server.crt` - your server's public certificate you received from the CA
- `server.key` - the private key file you generated

```
cp <your-server-public-certificate-file>.crt /opt/mi/ssl/server.crt
cp <your-server-private-key>.key /opt/mi/ssl/server.key
```

5. Optional: update the intermediate certificate chain file

If you require any intermediate certificates, download all intermediate certificates from your signing authority (or ask your local Network / IT admin) and place them on the MI server:

```
/opt/mi/ssl/ca.crt
```

If you do not know, whether or not you need this, you can safely ignore it.

NOTE: The `/opt/mi/ssl/ca.crt` file is required to exist and be an actual certificate file. If you do not need a certificate file and accidentally deleted `ca.crt`, please create a link to your server's public certificate with the following command:

```
ln -s /opt/mi/ssl/server.crt /opt/mi/ssl/ca.crt
```

6. Restart

For Metric Insights v 6+, restart the web container with **mi-control restart web**


```
[root@ip-10-9-1-210 ec2-user]# mi-control restart web
Restarting mi_web_1 ... done
[root@ip-10-9-1-210 ec2-user]#
```

For Metric Insights v 5.x and below:

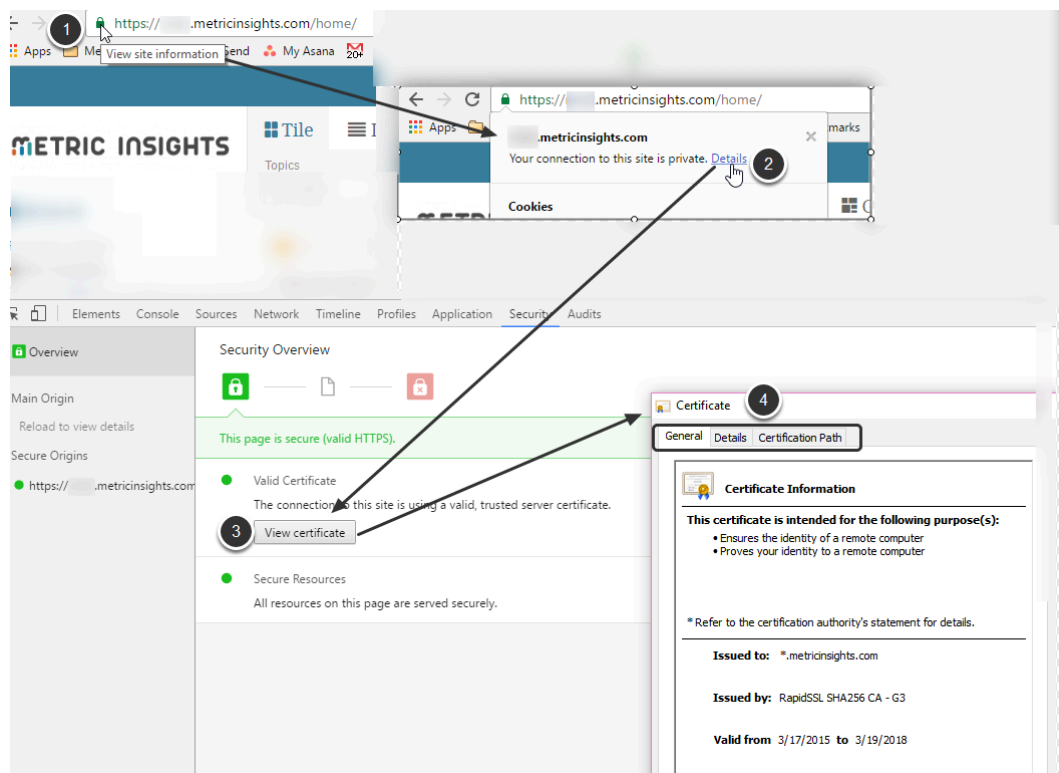
Debian:

```
service apache2 restart
```

CentOS / RedHat:

```
service httpd restart
```

7. Verify that the new certificate is in place



Access your instance, and:

1. Click the **Lock** icon next to view your site information
2. Click **Details**
3. Select **View certificate**
4. Verify **Certificate** information

7.1. Optionally, you could run the following command

```
echo | openssl s_client -connect your.metricinsights.host:443 2>/dev/null | openssl  
x509 -noout -dates
```

Replace 'your.metricinsights.host' with the hostname or IP address for your Metric Insights server. (This command can be run from a Linux shell prompt on any machine that is able to connect to your Metric Insights server, including the Metric Insights server itself.)

You should get output similar to the following:

```
notBefore=Mar 18 00:36:43 2015 GMT  
notAfter=Mar 20 00:26:42 2018 GMT
```


3.14 How do I specify a domain name when connecting to SQL Server?

The setup for SQL Server connections is different from other SQL connections. Users almost always have 'domain' names that accompany their usernames. The standard Windows protocol is to combine the domain and username like this:

domain\username

...and enter the combination where it says 'Username: '. That's not the correct approach when setting up data source connections in MI. The correct approach is to enter the username as-is and to specify the domain by appending it to the jdbc string, as follows:

jdbc:jtds:sqlserver://[servername]:1433/[DBName];domain=[DomainName]

3.15 Mobile App Security and Encryption Overview

The Metric Insights mobile apps follow all best practices related to security and encryption. Both apps store as little information on the phone as possible: Metadata about each object is stored to generate the list in the "Analyst Insights" homepage view, and Metric data is stored for rendering Metrics in more recent versions. Images and Dataset Reports (stored as PDF files) are stored locally in sandboxed directories that cannot be accessed by other applications.

Both iOS and Android support biometric authentication via Touch ID / fingerprint scanning.

Communication and Authentication

iOS and Android apps use the same basic approaches for communication and authentication:

1. The apps use HTTPS for all our requests to the Metric Insights server
2. The apps do not save any passwords, only authorization tokens with a limited lifetime
3. The apps will work with common mobile VPN clients

Data Storage

Across both iOS and Android, images and cached data are stored locally in the application's sandboxed (isolated) documents directory that cannot be accessed by other applications or the phone's owner. The application stores the user's:

- email
- name
- server-url
- token
- list of all objects available to the logged in user (only what the user has been specifically granted access to)
- data from any Element (Metric Insights object) that has been opened/viewed, including:
 - all chart data for Metrics
 - PDFs generated to display Reports and External Reports
 - thumbnail preview images for all Elements

There is no purge or garbage collection process for periodically clearing old data. When a user opens the same Element again, the data is updated.

More information on [secured/sandboxed directory in the Android system folder](#)



Using an MDM/EMM will provide an additional layer of security (data encryption, remote wiping, etc) depending on the platform and its specific features.

3.16 Configure Intelligent Global Search

Intelligent Global Search was introduced in Release 6.1.0. The feature uses Hypertext Preprocessor (PHP) that is a general purpose scripting language and MySQL Natural Language Understanding (NLU) search facilities. Searching is based on Python, the engine that receives a user request. With the help of Natural Language, the engine combines different forms and parts of speech of words in the search argument to produce an accurate and contextual base word ; i.e., a lemma. The lemma is then used to return the most relevant results to meet the user's request. This supports a search that is very flexible and produces more results than our previous searches.

Examples of words and their base form (lemma) that typically have a similar contextual meaning:

- *going* > *go*
- *sale* > *sale*
- *is* > *be*

In Release 6.2.1, additional functionality is available:

- "Recent requests and opened items" are displayed in the Search suggestion form.
- Users are informed when indexing is in progress.
- Search automatically switches/reverts to MySQL-based search whenever the Data Analyzer is indexing or re-indexing.
- A Voice Search option is supported.
- Documents are included in the Search and can be located by Document Name, file name or URL
- The Data Source tab contains External References and External Reports only.

Intelligent Search works with different types of Search Entities; i.e., objects in Metric Insights. Search results can be grouped by:

- Default: "All", "Tiles", "Datasets", "Bursts", Data Sources and Datasets tabs
- Category: "All", existing Categories and Datasets are shown as tabs
- Content Type: "All", "Metric", "Dataset Report", other existing content types and Datasets that are displayed as tabs

The search includes metadata about Search Entities: Information is searched across Names, Descriptions, Categories, Tags but not the data that they contain. The system relies on metadata when locating Elements and Bursts that meet the search criteria.

For Datasets, additional information about columns to be included is defined during the Search Setup process. It is important to remember that **only MySQL is supported for searching across Datasets**.

To understand Searching, see this [article](#).

1. Configure Intelligent Global Search: Admin > System > System Variables

i System Variables are list of configuration parameters in the system that allow a particular value; e.g., Y or N, to be assigned in order to enable or disable a selected feature of the system.

The screenshot displays the 'System Variables' configuration page in the Metric Insights Admin interface. The page title is 'Change advanced system settings. Changes are not applied until they are committed.' Below this is a table of system variables:

Variable Name	Assigned Value	Valid Values	Description
ENABLE_GLOBAL_SEARCH	N	Y/N	Enable Global Search
GLOBAL_SEARCH_DATA_ANALYZER_ENDPOINT	https://docs.metricinsights.com/data...		Global Search Data Analyzer Endpoint
GLOBAL_SEARCH_USE_DATA_ANALYZER	Y	Y/N	Use Data Analyzer for Global Search

Four numbered steps are indicated by arrows:

- Step 1: Select the 'ENABLE_GLOBAL_SEARCH' variable in the table.
- Step 2: Click the gear icon (edit) for the 'ENABLE_GLOBAL_SEARCH' variable.
- Step 3: In the 'Edit Variable' dialog, change the 'Assigned value' from 'N' to 'Y'.
- Step 4: Click the 'Commit Changes' button at the bottom of the table.

The 'Edit Variable' dialog for 'ENABLE_GLOBAL_SEARCH' shows the 'Assigned value' dropdown set to 'Y'. The 'Edit Variable' dialog for 'GLOBAL_SEARCH_USE_DATA_ANALYZER' shows the 'Assigned value' dropdown set to 'Y'.

1. Set the Configuration Variable: **ENABLE_GLOBAL_SEARCH**
2. Click [gear] icon to Global Search,
3. Assign "Y" value and **[Save]** and **[Commit Changes]**
4. Repeat Steps 2 and 3 for the **GLOBAL_SEARCH_USE_DATA_ANALYZER** Configuration Variable

1.1. Search Setup: Admin > System > Search Setup

The screenshot shows the 'Search Setup / Search Setup Editor' interface. At the top, there's a navigation bar with a search icon and a search input field. Below the navigation bar, there's a tab labeled 'Info' with a callout '5' pointing to the 'Update All Search Indexes' and 'Save' buttons. The main content area is divided into three sections: 'Objects', 'Data Sources', and 'Datasets for Search'. In the 'Objects' section, there's a 'Include in Search' section with checkboxes for 'Elements', 'Datasets', and 'Bursts', with a callout '1' pointing to the 'Include in Search' label. Below this is a 'Group Results by' section with radio buttons for 'Default', 'Category', and 'Content Type', with a callout '2' pointing to the 'Group Results by' label. The 'Data Sources' section has a search input field and a list of 'Data Sources to include in search', with a callout '3' pointing to the 'Add Data Source' button. The 'Datasets for Search' section has a search input field and a list of 'Datasets to include in search', with a callout '4' pointing to the 'Add Dataset' button. A callout '5' also points to the 'Update All Search Indexes' button.

On the Search Setup Editor:

1. Select one or all types of entities to be included in the Search
2. Choose the manner by which Search Results are to be grouped:
 1. Default: "All", "Tiles", "Datasets", "Bursts", Data Sources and Datasets tabs
 2. Category: "All", existing Categories and Datasets are shown as tabs
 3. Content Type: "All", "Metric", "Dataset Report", other existing content types and Datasets are displayed as tabs
3. Add all Data Sources to be searched across External Reports
4. Specify Datasets to be searched across metadata
5. **[Save]** and **[Update All Search Indexes]** to commit changes

1.2. Define Search of a specific Dataset

The system looks for lemmas in Objects; i.e., Search Entities themselves and Objects associated with them, such as Tags, Categories, Data Sources, Dimension Values, and Dataset columns.

Search Entities group results and order the entities by determining relevance (rating) of each item. This rating depends on a number of "hits" and a weight of a field where they are detected; e.g., a Name of a field weighs more than a field containing its Description.

Datasets for Search

Search across metadata (typically a title, description and link)

1

+ Add Dataset

Title	Name	
CNN Daily News	CNN Daily News	
Daily Active Users ds	Daily Active Users ds	
Search Dataset	Search Dataset	
VS_dataset_for_search	VS_dataset_for_search	

2

Dataset 0412-MSSQL

3

Save or cancel

To add new Dataset to search:

1. **[+Add Dataset]**
2. Select desired Dataset
3. **[Save]**

1. Define the Datasets:

- Title Column: specifies the field where the search is performed
- Link Column: allows attaching links to resources, if any
- Thumbnail Link Column: identifies the preview to be displayed in the search results

2. **[+Add Dataset Column]** adds data to the search
3. Select a Column and **[Save]**
4. **[Back to Search Setup]**
5. **[Update All Search Indexes]**
6. **[Proceed]**
7. Indexing in Progress message appears

Continue adding Data Sources or Datasets to the Search criteria as necessary by repeating steps above.

4. Troubleshooting Metric Insights

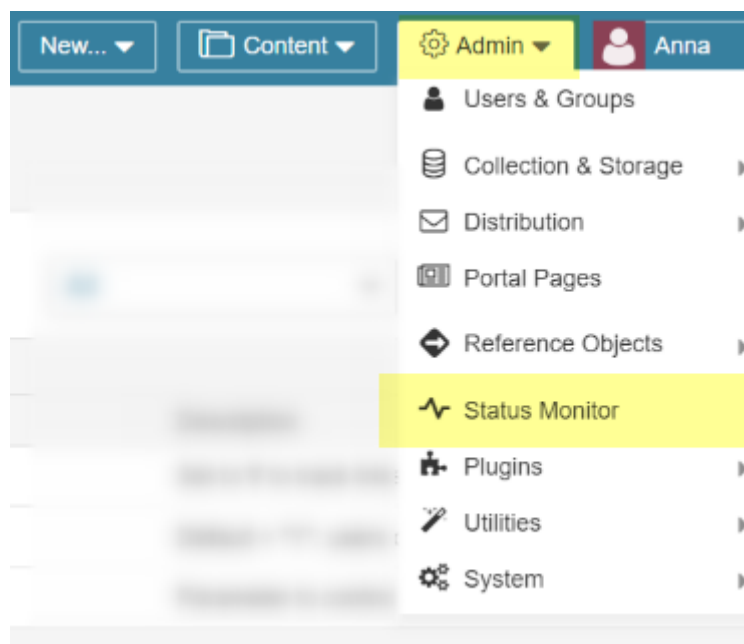
4.1 Troubleshooting Email Connectivity - Sending emails

Once you have email up and running you might run into a situation where users are not receiving emails. This article gives some tips on how to troubleshoot email issues if this occurs.

1. Check Status Monitor screen for Sending Email

The Status Monitor screen provides information on your system such as usage and errors. If you have errors in your email setup then the Status Monitor provides information on that.

1.1. Admin menu > Status Monitor



1.2. Status Monitor for Sending email error

The Status Monitor displays error information for your email setup. Each time you refresh this page, the system attempts to connect to the email server.

1. If an error occurs then it displays in email server section. Clicking on Fail will display error
2. Example error
3. Any emails not sent should reside in the Queue and can be corrected in that section

1.3. Email error and possible causes

The following is a list of email errors with possible causes

1. phpnetwork_getaddresses: getaddrinfo failed: Name or service not known

```
Error encountered when Sending mail: phpnetwork_getaddresses: getaddrinfo failed: Name
or service not known Check setting of SEND_MAIL_ variables in /var/www/iv/engine/config/
const.php
```

Cause: Incorrect value for SEND_MAIL_SERVER

2. Error: authentication failed: authentication failure

Error encountered when Sending mail: 5.7.8 Error: authentication failed: authentication failure Check setting of SEND_MAIL_ variables in /var/www/iv/engine/config/const.php.

Cause: Incorrect value for SEND_MAIL_LOGIN, SEND_MAIL_PASSWORD

3. Username and Password not accepted

Error encountered when Sending mail: 5.7.8 Username and Password not accepted. Learn more at 5.7.8 <http://support.google.com/mail/bin/answer.py?answer=14257>
yn4sm25783434pbb.95 - gsmtplib Check setting of SEND_MAIL_ variables in /var/www/iv/engine/config/const.php.

Cause: Incorrect value for SEND_MAIL_LOGIN, SEND_MAIL_PASSWORD

4. Network is unreachable

Error encountered when Sending mail: Network is unreachable Check setting of SEND_MAIL_ variables in /var/www/iv/engine/config/const.php.

Cause: Incorrect value for SEND_MAIL_SERVER, SEND_MAIL_PORT

5. Failed to parse address

Error encountered when Sending mail: Failed to parse address "smtp.gmail.com:" Check setting of SEND_MAIL_ variables in /var/www/iv/engine/config/const.php.

Cause: Incorrect or no value for SEND_MAIL_PORT

6. Error: authentication not enabled

Error encountered when Sending mail: 5.5.1 Error: authentication not enabled Check setting of SEND_MAIL_ variables in /var/www/iv/engine/config/const.php.

Cause: Incorrect value for SEND_MAIL_PORT

7. Must issue a STARTTLS command first

Error encountered when Sending mail: 5.7.0 Must issue a STARTTLS command first.
sm8sm20393998pbc.92 - gsmtplib Check setting of SEND_MAIL_ variables in /var/www/iv/engine/config/const.php.

Cause: Incorrect value for SEND_MAIL_SSL

8. Last request failed

Error encountered when Sending mail: Last request failed. Check setting of SEND_MAIL_variables in /var/www/iv/engine/config/const.php.

Cause: Value for SEND_MAIL_SERVER and SEND_MAIL_PORT are valid but email server at that location is not able to process request

9. Connection timed out

Error encountered when Sending mail: Connection timed out Check setting of SEND_MAIL_variables in /var/www/iv/engine/config/const.php.

Cause: Incorrect value for SEND_MAIL_SERVER, SEND_MAIL_PORT.

At machine confirm that you have connectivity to the email server. For example, use telnet to test

```
[root@MetricInsights-Centos-64-bit ~]# telnet smtp.gmail.com 587
Trying 173.194.79.109...
Connected to smtp.gmail.com.
Escape character is '^]'.
220 mx.google.com ESMTP lr3sm64684873pab.4 - gsmtpl
```

Upon successful telnet connection, then exit via "control]" keys followed by "quit"

```
^]
telnet> quit
Connection closed.
[root@MetricInsights-Centos-64-bit ~]#
```

10. Unrecognized command

Error encountered when Sending mail: 5.3.3 Unrecognized command. Check setting of SEND_MAIL_variables in /var/www/iv/engine/config/const.php.

Cause: Value for SEND_MAIL_SSL should be empty ("")

11. No specific error information supplied.


```
Error encountered when Sending mail. Check setting of SEND_MAIL_ variables in /var/www/  
iv/engine/config/const.php.
```

Many possible reasons can cause this message. For example, certain values for the SEND_MAIL_ parameters are not completed, such as SEND_MAIL_SERVER, SEND_MAIL_LOGIN

2. If no errors are shown on the Status Monitor, try these steps:

2.1. Test sending an email

Try sending an email. There are various places in the Metric Insights application where you can send an email. Below are just a few examples.

1. [Requesting a 'Daily' Favorites Digest](#)
2. [The Share Function](#)

2.2. Test sending emails to various accounts

If users within an organization are not getting emails then try sending emails to accounts outside the organization. For example, try sending to your personal email account.



2.3. Check whitelist

If you find that some accounts do not receive emails but others do, then ask your email administrator to verify that the Metric Insights is on the whitelist for receiving email.





2.4. Check spam folder

If you find that some accounts do not receive emails but others do, then check their spam folder. Emails could have been blocked as spam by your system.


3. Check email log files


Config Variables
Admin 

Key: Changed
All
logging

Config Variables			
Name	Assigned value	Valid values	
UI_DATA_FETCH_LOGGING_ENABLED	Y	Y,N	
DATA_COLLECTION_RESULT_LOGGING_EN...	Y	Y,N	
EMAIL_LOGGING	Y	Y,N	
LOGGING_SYSLOG_FILE_NAME	/var/log/mi.log		

⏪ ⏩ | Page 1 of 1 | ⏴ ⏵ | 🔄
Displaying records 1 - 4 of 4

 **Generate const.php**
Generate a new file for your changes to take effect

If you have access to Metric Insights application server then you can check the system logs for any errors.

1. View the contents of **/var/mail/root**. This file contains information on emails with which the system had trouble.
2. View the contents of log files in all email digest activity for troubleshooting. The log file is saved to: `/var/www/iv/data/temp/email/YYYY-MM-DD.log` to see when email was sent. You will need to have set **EMAIL_LOGGING** to 'Y' in the [Setting the Configuration Variables](#)

4. Troubleshoot mail server via telnet

You can use telnet to do some basic querying of the mail server. Depending on the mail server, the commands and order of the commands will differ. What follows is just a few of the first simple SMTP commands via telnet

```
me@myserver:~$ telnet exchange2010.server.com 25
Trying 192.168.1.4...
Connected to exchange2010.server.com.
Escape character is '^]'.
220 EXCHANGE2010.server.com Microsoft ESMTP MAIL Service ready at Sun, 18 May 2014
23:36:01 -0400
ehlo
250-EXCHANGE2010.server.com Hello
250-SIZE
250-PIPELINING
250-DSN
250-ENHANCEDSTATUSCODES
```


250-STARTTLS
250-X-ANONYMOUSTLS
250-AUTH LOGIN
250-X-EXPS GSSAPI NTLM
250-8BITMIME
250-BINARYMIME
250-CHUNKING
250-XEXCH50
250-XRDST
250 XSHADOW

AUTH LOGIN

4.2 Troubleshooting Email Connectivity - Receiving Email

Once you have email up and running you might run into a situation where emails sent to the system by users are not showing up. This article gives some tips on how to troubleshoot email issues if this occurs.

1. Check email log files

If you have access to Metric Insights application server then you can check the system logs for any errors.

1. View the contents of `/var/mail/root`. This file contains information on emails that the system had trouble with.

2. Checklist of common items

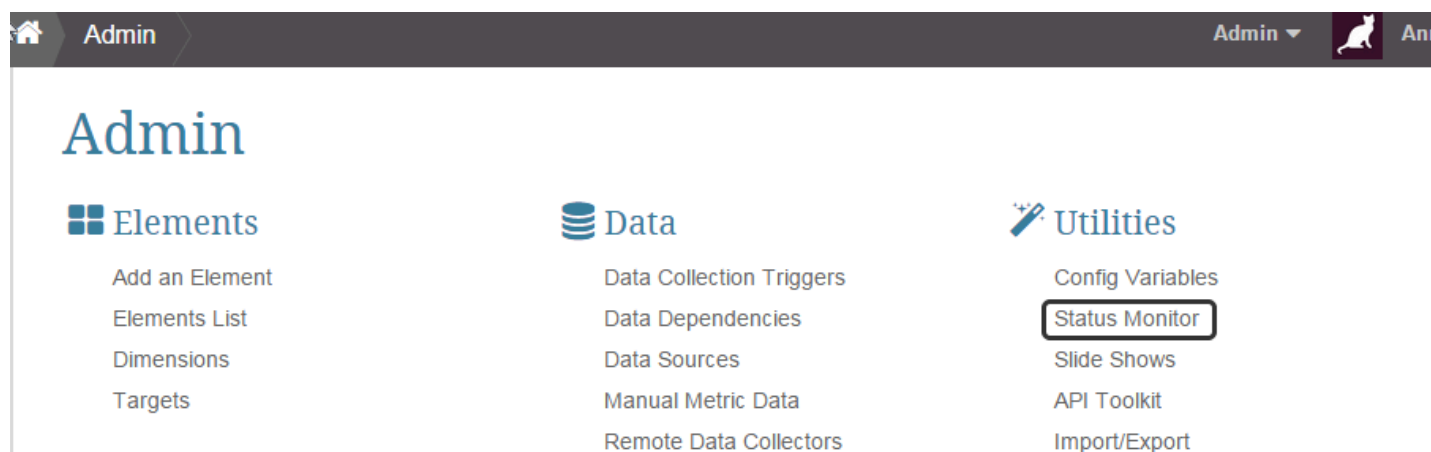
[Verify your settings](#). Common things to look for

1. Are you using SSL or TLS encryption? Value for **RECV_MAIL_SSL**
2. Are you using the correct port for listening? Value for **RECV_MAIL_PORT**
3. Are there any firewall rules preventing the Metric Insights server from contacting your email server's pop listening port?

3. Check Status Monitor screen for Receiving Email

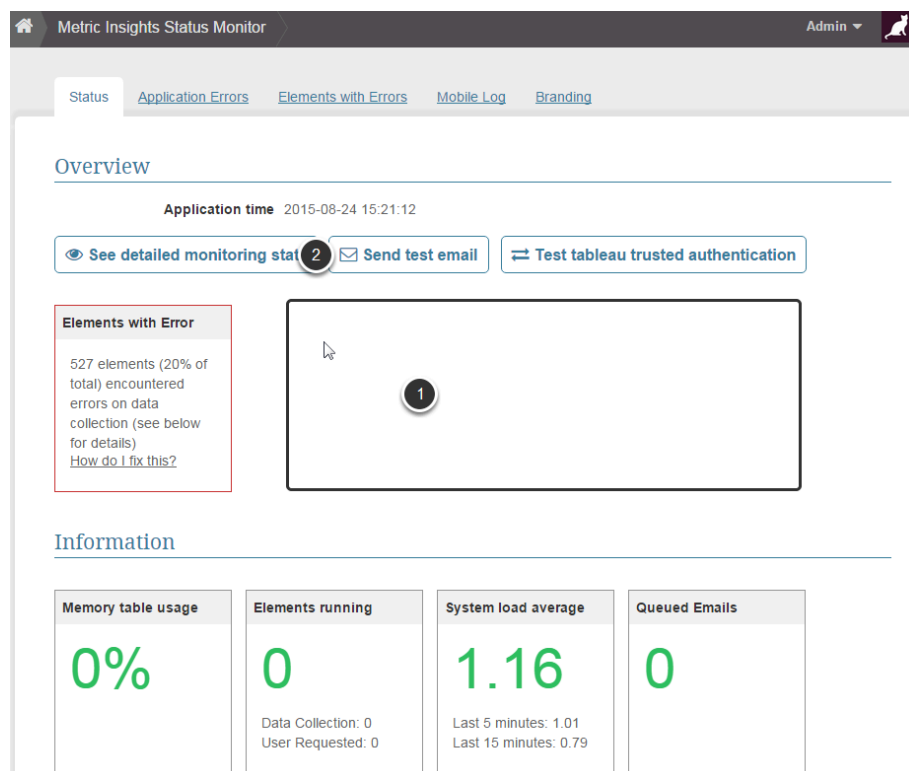
The Status Monitor screen provides information on your system such as usage and errors. If you have errors in your email setup then the Status Monitor provides information on that.

3.1. Status Monitor via Admin menu



Select **Status Monitor** from Admin menu

3.2. Status Monitor for Receiving Email error



The Status Monitor displays error information for your email setup.

1. Each time you refresh this page, the system attempts to connect to the email server. If an error occurs then it displays here.
2. You can also send a test email from this screen

3.3. Email error and possible causes

The following is a list of email errors with possible causes

1. read failed - connection closed?

```
Error encountered when Receiving mail: read failed - connection closed? Check setting of RECV_MAIL_ variables in /var/www/iv/engine/config/const.php.
```

Cause: Incorrect value for RECV_MAIL_SSL

2. cannot connect to host; error = Connection timed out (errno = 110)

```
Error encountered when Receiving mail: cannot connect to host; error = Connection timed out (errno = 110 ) Check setting of RECV_MAIL_ variables in /var/www/iv/engine/config/const.php.
```

Cause: Incorrect values for RECV_MAIL_SERVER, RECV_MAIL_PORT

At machine confirm that you have connectivity to the email server. For example, use telnet to test

```
[root@MetricInsights-Centos-64-bit ~]# telnet pop.gmail.com 995
Trying 74.125.25.108...
Connected to smtp.gmail.com.
Escape character is '^]'.
Connected to pop.gmail.com.
```

Upon successful telnet connection, then exit via "control]" keys followed by "quit"

```
^]
telnet> quit
Connection closed.
[root@MetricInsights-Centos-64-bit ~]#
```

3. cannot connect to host; error = (errno = 0)

```
Error encountered when Receiving mail: cannot connect to host; error = (errno = 0 )
Check setting of RECV_MAIL_ variables in /var/www/iv/engine/config/const.php.
```

Cause: Incorrect values for RECV_MAIL_SERVER, RECV_MAIL_PORT

4. last request failed

Error encountered when Receiving mail: last request failed Check setting of RECV_MAIL_variables in /var/www/iv/engine/config/const.php.

Cause: Incorrect value for RECV_MAIL_LOGIN, RECV_MAIL_PASSWORD

5. Last request failed. Authentication failed

Error encountered when Receiving mail: Last request failed. Authentication failed. Check setting of RECV_MAIL_variables in /var/www/iv/engine/config/const.php.

Cause: Incorrect value for RECV_MAIL_LOGIN, RECV_MAIL_PASSWORD

6. No specific error information supplied.

Error encountered when Receiving mail. Check setting of RECV_MAIL_variables in /var/www/iv/engine/config/const.php.

Many possible reasons can cause this message. For example, certain values for the RECV_MAIL_parameters are not filled in, such as RECV_MAIL_SERVER, RECV_MAIL_LOGIN

4.3 Getting "ERROR: Can't write to /var/backups/mi-app-backups/mi-full-backup" when trying to backup my Metric Insights Instance.

This is a permission issue that may occur when you try to run *mi-app-backup* on your Metric Insights Virtual Machine with user that don't have root privileges.

Resolution

In order to get root privileges to run *mi-app-backup*, do the following:

1. Find where on your vm *mi-app-backup* binary file is located by running:

```
which mi-app-backup
```

You should get something like:

```
/usr/local/bin/mi-app-backup
```

Note: If this path can't be found then you should add `"/usr/local/bin"` to path in user's `".bashrc"` file.

2. Then run:

```
sudo /usr/local/bin/mi-app-backup
```

Note: If it asks password, please ask it from your VM administrator.

4.4 Will Metric Insights Run on Windows?

Yes. Metric Insights is designed to run on Linux, so if you want to run Metric Insights with a Windows operating system you will need to use a virtualization environment. Metric Insights can then be deployed as a virtual appliance. Here is [how to install](#).