System

Table of Contents

Deployment Options ... 4

Deploying a 'warm-standby' environment for Metric Insights ... 5

Installation ... 9

How to move the MI app from the /opt directory to a new location ... 10

Metric Insights Virtual Machine disk space expansion ... 11

How often are builds and releases pushed out and how? ... 12

Are all included software packages patched to the latest version? ... 13

Are patch updates for MySQL and other software included with new builds? ... 14

Why are the compilers required and can they be disabled? ... 15

Can Metric Insights be deployed on top of a Linux OS? ... 16

Downgrading Metric Insights to a prior version (Virtual Machine) ... 17

Downgrading Metric Insights to a prior version (AWS) ... 19

Where is the Version number? ... 21

[v6.x] Getting "Requires: container-selinux >= 2.9" error when installing Metric Insights v6.x on a single server ... 22

Administration ... 24

Daily Admin Checklist ... 25

Bulk User Deletion (using Custom Script) ... 28

Upgrades ... 32

Getting a "403 Forbidden" Error when upgrading from 3.2 to 3.3 on a VM ... 33

How do I know the update was successful? ... 34

Upgrading-in-place Metric Insights from v5.6.x to v6.2.x ... 42

Operating Systems ... 46

Getting a "Failed to fetch Backports..." Error when upgrading from Debian 7.4 to latest ... 47

GPG Error when updating Debian or Metric Insights Tools ... 49

Encryption ... 50

How to use Azure Key Vault to encrypt keys ... 51

Authentication ... 54

After upgrading to 3.2.830, some LDAP users can't log in ... 55

What is the difference between 'default-sp' and 'default-signed-sp' for SAML? ... 57

LDAP SSL certificates 58

How to configure Okta for MI SAML SSL setup 60

Database 65

Fine Tuning MySQL Parameters 66

Ports used by Metric Insights to connect to different databases 68

RDS Parameters for MySQL 69

Encrypting database data at rest in Metric Insights 71

Logging 72

How to move the log files from /var/log to a new location 73

SQL Queries for Database statistics 74

MySQL statements for various tasks 75

Network 93

SSH Port Tunneling with Putty 94

How to set up a Static IP for a Virtual Machine (CentOS) 95

"Device eth0 does not seem to be present, delaying initialization" ERROR on VM reboot 98

Deployment Options

Deploying a 'warm-standby' environment for Metric Insights

This article covers how to deploy a warm-standby environment for Metric Insights and the operational requirements to maintain such an environment. By 'warm-standby' we mean an active/passive high-availability deployment. For any additional questions, please contact support@metricinsights.com.

What is a 'warm-standby' deployment?

How to deploy a 'warm-standby' environment

Caveats

What is a 'warm-standby' deployment?

A 'warm-standby' deployment consists of an *active* primary application server and an *inactive* secondary application server. The inactive server essentially serves as a *standby* server to be promoted to active in the event the primary goes down. Both servers maintain their own Metric Insights database locally.

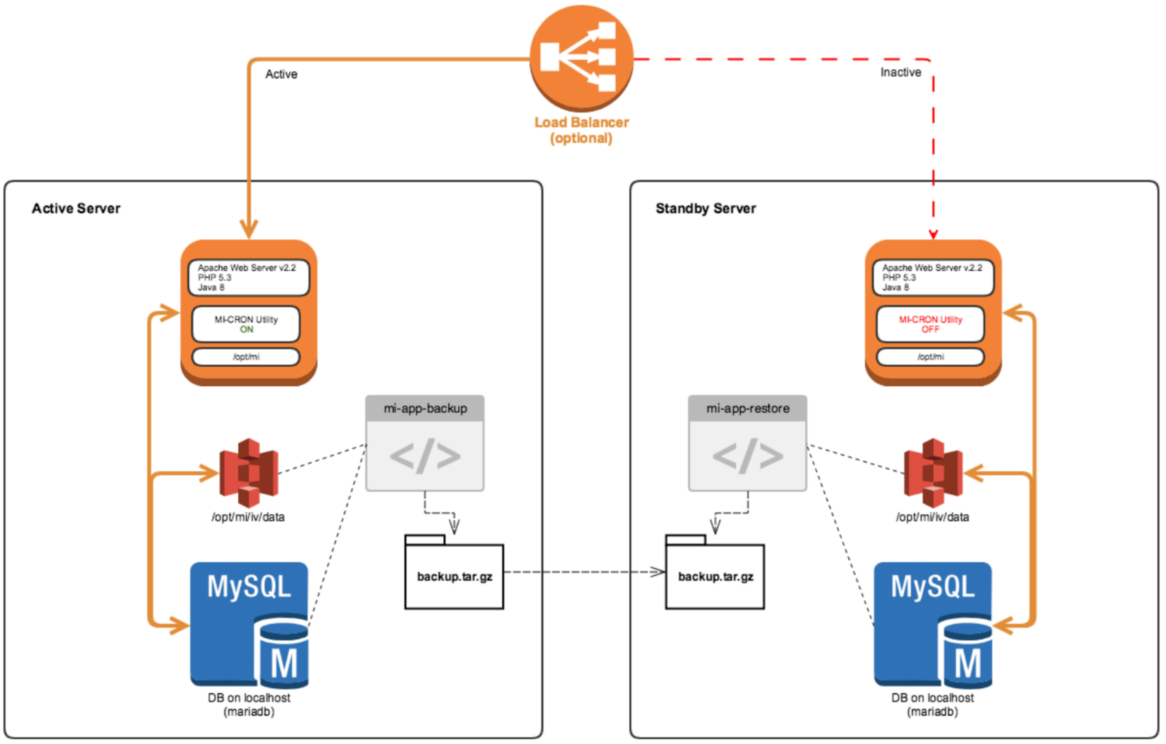

In the diagram above, notice the Active and the Standby servers mirror each other in terms of architecture:

- Both servers consists of a basic LAMP stack (Linux, Apache, MySQL, PHP) to support the Metric Insights application
- Metric Insights is installed to the default /opt/mi location (this can be changed of course)
- MySQL is running locally for the Metric Insights database
- All data files (json, png, etc.) sit in /opt/mi/iv/data

There are *two key differences* that distinguish the Active from the Standby however:

1. Only the primary serves user sessions
2. Only the primary has cron enabled (the mi-cron utility)

This means all user traffic is directed to the primary server and all Metric Insights jobs run on the primary server only. Jobs include data collections and distributions.

Load Balancer Optional

You can choose to direct users directly to the primary server, or you can have the primary sit behind a load balancer.

How to deploy a 'warm-standby' environment

To deploy a warm-standby environment, you need at the very least a second server to serve as a 'mirror' for the primary server. The servers should have the same hardware and software specs:

- Same number of CPU cores
- Identical RAM
- Identical Storage
- Identical Linux OS
- Identical Linux kernel
- Identical Apache, MySQL, PHP installations
- Identical MySQL parameters
- Identical Apache, Syslog, Cron configurations

On the active server, install Metric Insights with **mi-cron enabled**.

On the standby server, install Metric Insights with **mi-cron disabled**.

How to enable/disable mi-cron from the command line

```
mi-cron enable
mi-cron disable
```

Create scheduled backups

On the active primary server, create scheduled backups of Metric Insights using the **mi-app-backup** utility. Make sure to create full backups (full backups include the database and data files). Please see http://help.metricinsights.com/m/MI_System_Maintenance/l/104530-backup-your-metric-insights-instance for more.

Copy the backup file to the standby server

As soon as the backup file is created on the active server, copy the backup to the standby server. There are a number of ways to do this including:

- Using a tool like *rsync* to copy the backup file (as seen in the diagram)
- Saving the backup file to a shared network drive (mounted on both the active and standby servers)

Restore the backup on the standby server

On the standby server, create scheduled restores of Metric Insights using the **mi-app-restore** utility. Make sure the utility restores the latest backup from the active server. Please see http://help.metricinsights.com/m/MI_System_Maintenance/l/104531-restore-your-metric-insights-instance for more.

On restore completion, make sure to disable cron by running **mi-cron disable**. This ensures triggers are not running on the standby server to collect data as well as notification schedules from delivering content.

A simple shell script can be created to both run the restore and disable cron on restore completion.

Caveats

- Our general recommendation is to create backups daily. Therefore, any backups restored on the standby server will always be a day behind. Should the standby server have to be promoted to active, simply running data collections for the day will bring the data to current.
- In the event the active server goes down and users have created new content since the most recent backup, then the new content will be lost when the standby server is promoted.
- Should the standby server be promoted to active, make sure to *enable mi-cron* so data collections and distributions can run.
- If the server hostname is mapped directly to the active server (active server's IP), and the active goes down, network administrators must remap the hostname to the standby server (DNS update).

- If using a Load Balancer instead, you can configure the LB to auto-detect server failure and auto-switch to the standby server.
- Both active and standby servers cannot serve user sessions. In other words, the standby server must truly be on *standby*.

For any additional questions, please contact support@metricinsights.com.

Installation

How to move the MI app from the /opt directory to a new location

If you need move the Metric Insights app from the /opt directory to a new location then run the following commands:

```
mv /opt/mi /new/path
```

```
ln -sf /new/path /opt/mi
```

Next time when you will need to run the installer for Metric Insights application upgrade, it should upgrade the application in the new location.

Metric Insights Virtual Machine disk space expansion

If the Metric Insights server is starting to run low on disk space then you can follow this link with instruction to increase the size of Linux LVM (logical volume manager) by expanding the MI virtual machine disk:
https://www.rootusers.com/how-to-increase-the-size-of-a-linux-lvm-by-expanding-the-virtual-machine-disk/

How often are builds and releases pushed out and how?

Question

I was told new a build of Metric Insights is pushed out every two weeks and that a new version is released on a yearly basis - can you please confirm the differences and frequency of both types of updates? Additionally as we have no internet connectivity, will these both need to be done via internal mirrors?

Answer

New *builds* generally contain fixes to any issues uncovered in the application and feature updates. New *releases* contain new features and major feature changes.

If updating Metric Insights using our *install packages* and having outbound access is not a possibility, then the MI server must have access to internal mirrors (for any linux package updates to support the application, like python).

If using our OVAs to deploy a new virtual machine of the latest MI build, then the internal mirrors aren't needed but of course you'd be deploying a new VM every time (and must copy the content over from the old).

To request the latest build or release, please contact support@metricinsights.com for access to the Metric Insights downloads page.

Are all included software packages patched to the latest version?

Question

With each new Metric Insights update, are all included software packages patched to the latest available version?

Answer

The Metric Insights application itself will be updated along with any linux packages and modules needed to support the application. These are not always the latest available version however - only what's needed at the time of development and to address any security issues.

Are patch updates for MySQL and other software included with new builds?

Question

If a new build release is expected on a bi-weekly basis, are patch updates for MySQL and other software included?

Answer

Metric Insight updates are only for the MI application and any supporting linux packages and modules. MySQL, Linux OS, etc. are not touched and must be maintained by your internal systems team.

Why are the compilers required and can they be disabled?

Question

Why are the compilers required and can they be permanently disabled to stop then re-enabling after a Metric Insights update?

Answer

The compilers are needed to run the Metric Insights install packages when updating MI. You can disable the compilers, but they will be installed again during the next update of MI.

Can Metric Insights be deployed on top of a Linux OS?

Question

Can Metric Insights be deployed in an alternative format to an appliance so we can install it on top of a Linux OS and have the underlying OS administered separately?

Answer

Yes, you can install Metric Insights in a Linux OS provisioned by your internal system team. We support Debian, CentOS, and Redhat (see links below). You would then use our install package to install MI. An OVA is provided when a Linux OS cannot be provisioned by your systems team or a request is made to stand up MI quickly.

http://help.metricinsights.com/m/Deployment_and_Configuration/l/554106-install-or-update-mi-via-installation-packages

http://help.metricinsights.com/m/Deployment_and_Configuration/l/608858-operating-system-requirements-dependencies

Downgrading Metric Insights to a prior version (Virtual Machine)

Metric Insights (MI) does not officially recommend downgrading. However, if you are facing unforeseen issues after upgrading MI to a newer version, here are the steps to take to downgrade to the previous version. Note, the steps below apply to MI running on a Virtual Machine:

[If snapshots are taken of the VM periodically]

Restore a snapshot of the VM hosting a prior version of MI. Ideally, a snapshot is taken right before an upgrade in case a rollback is necessary.

[If snapshots are not taken of the VM periodically]

Run the MI installer for the previous version, then restore the most recent backup of MI.

For example:

Metric Insights was running version 4.2 prior to the upgrade to 5.0. Upon upgrading to 5.0, you realize that a downgrade is necessary. *Before attempting to downgrade, please ensure you have the most recent backup available in /var/backups/mi-app-backups (the default location for the standard backup job that runs in MI).*

1.) SSH to the MI server and locate the directory where the prior installation package was unpacked (in this example, MI version 4.2). Locate and run the *installer.py* script:

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

Make sure the installation completes successfully, then check the version is 4.2 by running this command:

```
mi-version
```

For any additional questions about running the installer, please see this help doc first: Install or Update MI via Installation Packages

2.) Now, restore the most recent backup of version 4.2. Ideally, you'll have created a backup prior to the 5.0 upgrade. The restore command is as follows:

```
mi-app-restore /PATH/TO/version4.2_BACKUP_FILE -vv
```

You can also reference this help doc on restoring backups: Restore Your Metric Insights Instance

Make sure the restore completes without errors.

Once the restore is complete, you want to run a series of smoke tests in the MI app to ensure the downgrade is successful. Here are some test steps to take:

1.) Log in to MI in a browser and check if the version number on the homepage is the same as the version you just downgraded to (in this example, v4.2).

2.) In the MI homepage go to the Admin menu > Status Monitor > [Application Errors] tab. Check for any new errors based on the timestamp. If you seen anything unusual, please contact Metric Insights Support for help.

3.) Check data source connections.

4.) Ensure you can validate and collect data for metrics and reports.

5.) Check to ensure you can send a test email from Admin > Status Monitor > [Send Test Email] button.

6.) Ensure you can deliver bursts by test sending a burst to yourself.

After the tests pass, the downgrade is complete! For any other questions, please contact Metric Insights Support for help (support@metricinsights.com).

Downgrading Metric Insights to a prior version (AWS)

Metric Insights (MI) does not officially recommend downgrading. However, if you are facing unforeseen issues after upgrading MI to a newer version, here are the steps to take to downgrade to the previous version. Note, the steps below apply to MI running in Amazon Web Services:

[If AMIs are created periodically]

Restore the most recent AMI of a prior version of MI. Then, repoint DNS to this new EC2 instance in Route 53. Ideally, an AMI is created right before an upgrade in case a rollback is necessary.

[If AMIs are not created periodically]

Run the MI installer for the previous version, then restore the most recent backup of MI.

For example:

Metric Insights was running version 4.2 prior to the upgrade to 5.0. Upon upgrading to 5.0, you realize that a downgrade is necessary. *Before attempting to downgrade, please ensure you have the most recent backup available in /var/backups/mi-app-backups (the default location for the standard backup job that runs in MI) or in S3.*

1.) SSH to the MI server and locate the directory where the prior installation package was unpacked (in this example, MI version 4.2). Locate and run the *installer.py* script:

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

Make sure the installation completes successfully, then check the version is 4.2 by running this command:

```
mi-version
```

For any additional questions about running the installer, please see this help doc first: [Install or Update MI via Installation Packages]

2.) Now, restore the most recent backup of version 4.2. Ideally, you'll have created a backup prior to the 5.0 upgrade. The restore command is as follows:

```
mi-app-restore /PATH/TO/version4.2_BACKUP_FILE -vv
```

You can also reference this help doc on restoring backups: [Restore Your Metric Insights Instance]

Make sure the restore completes without errors.

Once the restore is complete, you want to run a series of smoke tests in the MI app to ensure the downgrade is successful. Here are some test steps to take:

1.) Log in to MI in a browser and check if the version number on the homepage is the same as the version you just downgraded to (in this example, v4.2).

2.) In the MI homepage go to the Admin menu > Status Monitor > [Application Errors] tab. Check for any new errors based on the timestamp. If you seen anything unusual, please contact Metric Insights Support for help.

3.) Check data source connections.

4.) Ensure you can validate and collect data for metrics and reports.

5.) Check to ensure you can send a test email from Admin > Status Monitor > [Send Test Email] button.

6.) Ensure you can deliver bursts by test sending a burst to yourself.

After the tests pass, the downgrade is complete! For any other questions, please contact Metric Insights Support for help (support@metricinsights.com).

Where is the Version number?

The version number for your instance is located in lower left corner of your Homepage

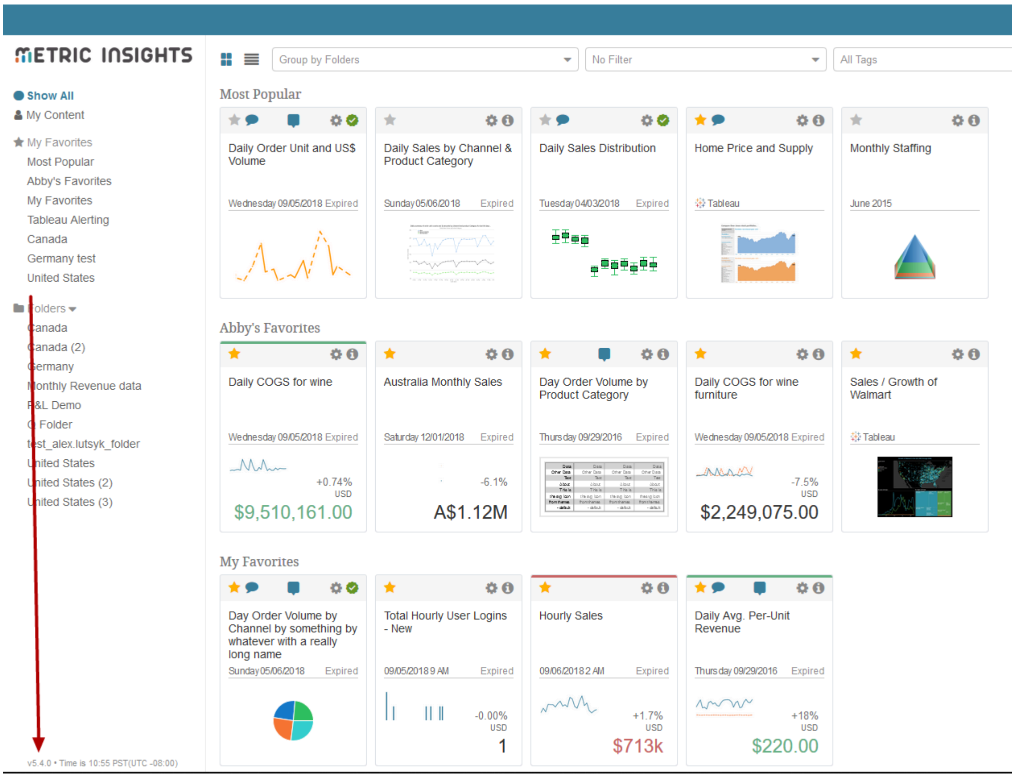

In this example, the current version is '5.4.0'

[v6.x] Getting "Requires: container-selinux >= 2.9" error when installing Metric Insights v6.x on a single server

Issue:

When running the Metric Insights v6.1.1 installer, it attempts to install Docker CE and then stops with the following package dependency error:

```
Error: Package: docker-ce-17.06.0.ce-1.el7.centos.x86_64 (docker-ce-stable)
           Requires: container-selinux >= 2.9
 You could try using --skip-broken to work around the problem
 You could try running: rpm -Va --nofiles --nodigest
```

How do I get the right version of the package container-selinux to proceed with the install? The host operating system is Red Hat 7.7 (RHEL7). The deployment is a single server / simple install type.

Solution:

To resolve the package dependency error, a version of **container-selinux** that is greater than version 2.9 must be available. Note, rather than try to manually install it, we'll attempt to make the package available so that the installer can pull it while installing Docker CE. To do this, first confirm whether container-selinux is installed. Most likely it is not:

```
yum list container-selinux
```

If container-selinux is not installed, check to make sure the appropriate repository is available from which to pull the package from. The repository needed is named

rhel-7-server-extras-rpms.

```
subscription-manager repos --list | grep -B3 'Enabled: 1'
```

If the Extra RPMS repo is not enabled, please enable it:

```
subscription-manager repos --enable=rhel-7-server-extras-rpms
```

If subscription-manager is *not enabled* on the machine, and you don't have the authority to subscribe, you can also install the package directly by running: `sudo rpm -i http://mirror.centos.org/centos/7/extras/x86_64/Packages/container-selinux-2.107-1.el7_6.noarch.rpm`

Once the repository is enabled, or the package installed, *rerun the Metric Insights installer.* It should now pull the right package, or reference the existing package, to continue on with the Docker CE install. Good luck!

Administration

Daily Admin Checklist

Almost all system information can be found on the Admin menu > Status Monitor page (see this doc for more information http://help.metricinsights.com/m/MI_System_Maintenance/l/783129-status-monitor-page).

The following is a list of items you can check to get an overall pulse of the Metric Insights application.

1. Red Tiles:
   Tiles highlighted in red represent issues that need your attention. For example, a red 'Elements with Errors' tile contains a link that will direct you to a list of elements with errors for further troubleshooting. Key tiles to look out for are:
   - Elements with Errors
   - Aborted Data Collections
   - Overdue Triggers
   - Queued Emails
   - Email Server
   - Cron
   - LDAP
   - System Load Average
   - Disk Space Free

2. Pipeline tab:
   This page visualizes the data collection triggers that are running on MI. This is especially useful if you have a chain of dependencies for different triggers. Visualizing this chain allows you to see the relationships between data dependencies, data collection triggers, and notification schedules, i.e. "how does data come into the system and how it goes out."
3. Application Errors tab:
   This page shows all recent application errors. This is especially useful if you or your team encounter issues with the application or the application exhibits any odd behavior. You can also download all system error logs from this page (errors you can't see in the UI) by clicking on the [Get Error Logs] button. You can choose to download the logs to your local machine or send them to the MI Support team for review.

   If sending to MI Support, be sure to submit a ticket to support@metricinsights.com and explain what issues you are encountering.

To get further insight into the system, you can also create reports to run queries against the MI application that will return useful information. Here are some sample queries you can run:

1. How many bursts were delivered for the prior week:

```
SELECT COUNT(DISTINCT nsrld.ns_run_id) as Quantity_of_delivered_Bursts
FROM notification_schedule_distribution AS nsd
JOIN notification_schedule_run_log_detail AS nsrld ON (nsrld.ns_distribution_id=nsd.
notification_schedule_distribution_id)
WHERE nsrld.start_time >=(NOW() - INTERVAL 1 WEEK);
```

2. How many users had failed login attempts for the prior week:

```
SELECT u.user_id AS User_ID, u.username AS User_Name, afl.failed_login_time AS
Failed_Login_Time,
afl.failed_login_attempts AS number_of_failed_attempts
FROM user AS u
JOIN auth_failed_logins AS afl ON (u.user_id=afl.id)
WHERE afl.failed_login_time>=(NOW()-INTERVAL 1 WEEK);
```

3. List of Burst recipients (if *notifications* are enabled):

```
SELECT DISTINCT nsdi.notification_schedule_distribution_id AS Burst_ID, nsd.name AS
Burst_Name, up.user_id AS Burst_Recipient_ID, u.display_name AS Burst_Recipient_Name,
IFNULL(fdei.element_id,nsdi.element_id) AS Element_ID
FROM notification_schedule_distribution AS nsd
JOIN notification_schedule_distribution_item AS nsdi ON (nsdi.
notification_schedule_distribution_id=nsd.notification_schedule_distribution_id
AND( (nsd.content_type='favorites' AND nsdi.favorite_id>0)
OR (nsd.content_type='tiles' AND nsdi.element_id IS NOT NULL)))
LEFT JOIN favorite_dashboard_element_info AS fdei ON (fdei.favorite_id=nsdi.favorite_id)
LEFT JOIN notification_schedule_distribution_group_recipient AS nsdgr ON (nsdgr.
notification_schedule_distribution_id=nsdi.notification_schedule_distribution_id)
LEFT JOIN user_group_member AS ugm ON (ugm.user_group_id=nsdgr.group_id)
LEFT JOIN notification_schedule_distribution_user_recipient AS nsdur ON (nsdur.
notification_schedule_distribution_id=nsdi.notification_schedule_distribution_id)
JOIN user_preference AS up ON (up.user_id=IFNULL(ugm.user_id,nsdur.user_id))
JOIN user AS u ON (u.user_id=up.user_id)
WHERE nsd.enabled_ind='Y'
AND up.email_notification_enabled_ind = 'Y'
AND IFNULL(fdei.element_id,nsdi.element_id) IS NOT NULL;
```

4. List of Alert recipients (if *notifications* are enabled):

```
SELECT DISTINCT uar.alert_rule_id AS Alert_ID, arei.visualization_element_id AS
Element_ID, de.name AS Element_Name,
uar.user_id AS Alert_Recipient_ID, u.display_name AS Alert_Recipient_Name
FROM user_alert_rule AS uar
JOIN alert_rule_element_info AS arei ON (uar.alert_rule_id=arei.alert_rule_id AND uar.
element_id=arei.element_id)
JOIN user_preference AS up ON (up.user_id=uar.user_id)
JOIN user AS u ON (u.user_id=uar.user_id)
JOIN dashboard_element AS de ON (de.element_id=arei.visualization_element_id)
WHERE uar.enabled_ind='Y'
AND up.email_notification_enabled_ind = 'Y';
```

Bulk User Deletion (using Custom Script)

Issue

Is there an automated way to delete a list of users that no longer use Metric Insights?

Resolution

One of the possible ways to set up Bulk User Deletion (e.g. by username) is via Custom Script.

You will need the following components:

- A **Dataset** to extract all the users to be deleted;
- A **Custom Script** to execute the deletion;
- An **API Access** to be able to launch the **Custom Script**.

To proceed, follow the steps below:

1. Set up **API Access** and obtain **API Token**, using our help documentation:

- How to Set up API Access
- How to Test API Access / Get API Token via MI API Test Tool

NB: the generated **API Token** is valid for around 15 minutes (by default), the exact time when the token expires can be seen on the **External Application** Edit Page (see the screenshot below). We recommend to generate the token right before the script execution. The token has to be valid during all the process of execution.

To extend the time the token is valid, go to **Admin** >> **Utilities** >> **System Config** and change the value of **API_TOKEN_LIFE_TIME** variable (consider that the deletion of one user takes around 1 minute on average).

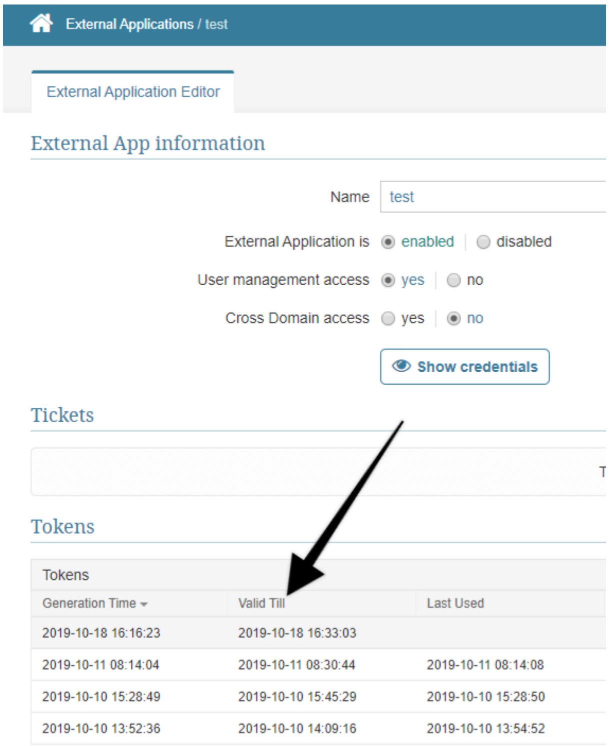

External App information

Name	test
External Application is	● enabled ○ disabled
User management access	● yes ○ no
Cross Domain access	○ yes ● no

Tickets

Tokens

Tokens

Generation Time	Valid Till	Last Used
2019-10-18 16:16:23	2019-10-18 16:33:03	
2019-10-11 08:14:04	2019-10-11 08:30:44	2019-10-11 08:14:08
2019-10-10 15:28:49	2019-10-10 15:45:29	2019-10-10 15:28:50
2019-10-10 13:52:36	2019-10-10 14:09:16	2019-10-10 13:54:52

2. Create a new **Dataset**, that will fetch IDs of all the users to be deleted:

1) Set the system database `Dashboard DB` as **Data Source**;

2) For the SQL query use the following template (*User1* and *User2* to be replaced by any number of usernames in single quotes, divided by comma):

```
SELECT user_id
FROM user
WHERE username in ('User1', 'User2')
```

3) Collect the data for the **Dataset**;

4) Memorize the **Dataset ID** (it is included in the **Dataset** link (see the screenshot below).

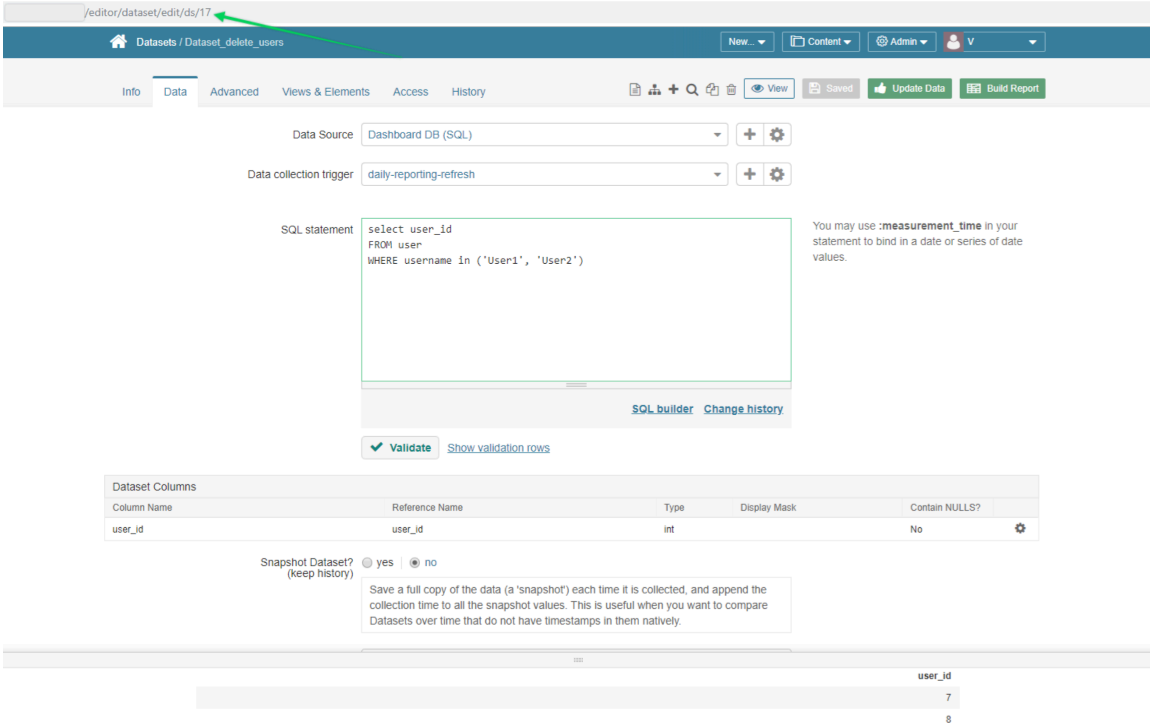

Data Source: Dashboard DB (SQL)

Data collection trigger: daily-reporting-refresh

SQL statement

```
select user_id
FROM user
WHERE username in ('User1', 'User2')
```

You may use **:measurement_time** in your statement to bind in a date or series of date values.

Dataset Columns

Column Name	Reference Name	Type	Display Mask	Contain NULLS?
user_id	user_id	int		No

Snapshot Dataset? (keep history) ○ yes ● no

Save a full copy of the data (a 'snapshot') each time it is collected, and append the collection time to all the snapshot values. This is useful when you want to compare Datasets over time that do not have timestamps in them natively.

user_id
7
8

3. To create the **Custom Script**:

1) Go to **Admin** >> **Utilities** >> **Custom Scripts**;

2) Create a **New Script**;

3) On the *Info* tab check the **External Application Name** and **Authentication User**, there should be set the ones used in p.1 to generate the **API Token**;

4) On the *Editor* tab paste the following script (replace `datasetIdToSet` with the **Dataset ID** from p.2.4 directly in the 2nd line of the script):

```
customScript.parameters = {
    "datasetId":datasetIdToSet
};

function processData(response){
    if('undefined'!==typeof(response) && response.data && response.data.length>0){
        var data = response.data
            , $t = 0;
        for(var $i in data) if(data.hasOwnProperty($i) && data[$i]['user_id']) {
            customScript.runApiRequest(customScript.homeSite + 'api/user/id/'+data[$i]['user_id'], {"type":"DELETE","async":false});

            var now = new Date().getTime();
            while(new Date().getTime() < now + 60000){ /* do nothing */ }
        }
        customScript.close();
    } else
        customScript.log('no dataset rows');
};

function processError(request,error){
    customScript.result('ajax error');
    customScript.close();
};

customScript.runApiRequest(customScript.homeSite+'api/
dataset_data?dataset='+customScript.parameters.datasetId,{"success":processData,
"error":processError});
customScript.heartBeatTimeout = 30000000;
```

5) Execute the **Script**.

Note 1

We highly recomment to test the solution first: select 1-2 users in the **Dataset** from p.2, as User Deletion process above is irreversible.

Note 2

If you are going to use this **Script** multiple times, you will need to specify the new usernames in the **Dataset** from p.2 and generate each time a new **API token** from p.1.

Upgrades

Getting a "403 Forbidden" Error when upgrading from 3.2 to 3.3 on a VM

ISSUE

Upgrading from 3.2.830 to 3.3.109 on a VM and am getting a "403 Forbidden" error. Note, the VM has outbound access to run **mi-deploy app inc**.

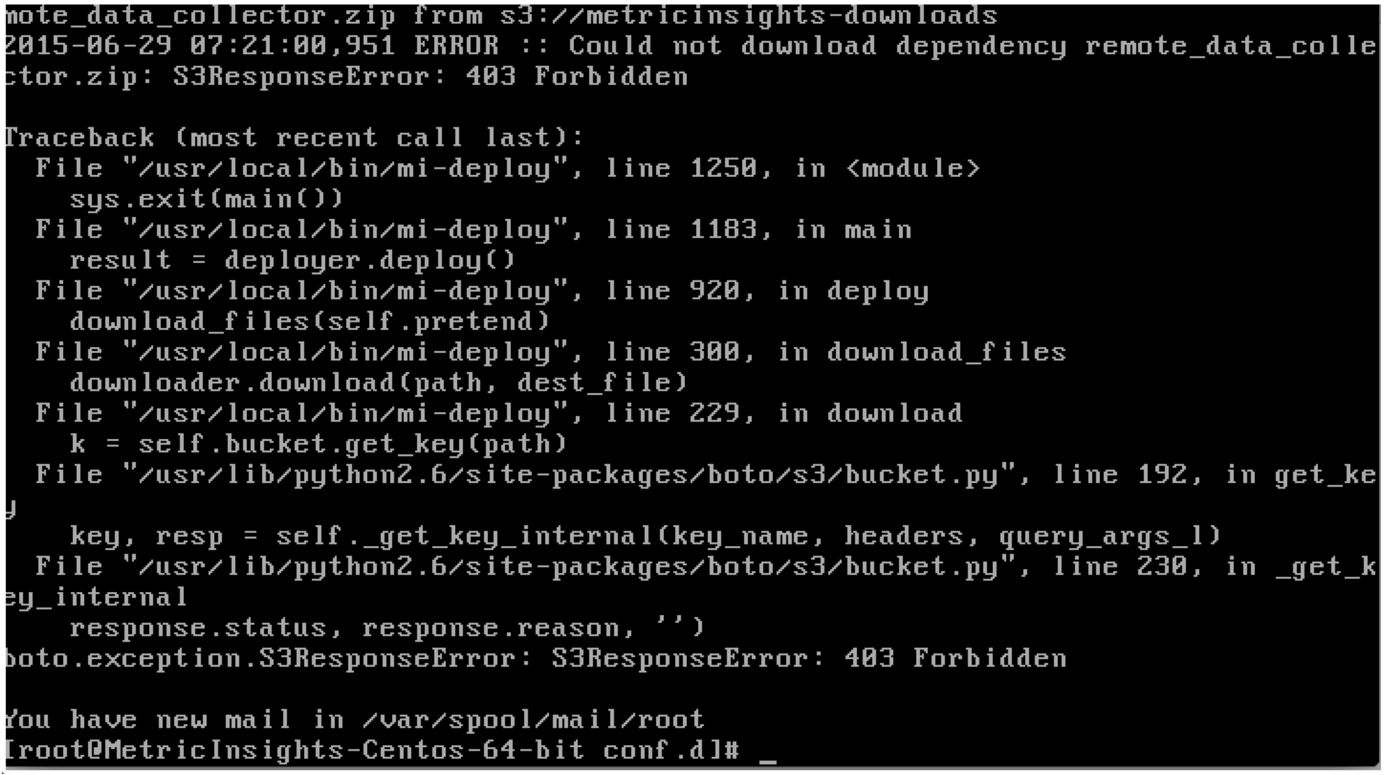

```
mote_data_collector.zip from s3://metricinsights-downloads
2015-06-29 07:21:00,951 ERROR :: Could not download dependency remote_data_colle
ctor.zip: S3ResponseError: 403 Forbidden

Traceback (most recent call last):
  File "/usr/local/bin/mi-deploy", line 1250, in <module>
    sys.exit(main())
  File "/usr/local/bin/mi-deploy", line 1183, in main
    result = deployer.deploy()
  File "/usr/local/bin/mi-deploy", line 920, in deploy
    download_files(self.pretend)
  File "/usr/local/bin/mi-deploy", line 300, in download_files
    downloader.download(path, dest_file)
  File "/usr/local/bin/mi-deploy", line 229, in download
    k = self.bucket.get_key(path)
  File "/usr/lib/python2.6/site-packages/boto/s3/bucket.py", line 192, in get_ke
y
    key, resp = self._get_key_internal(key_name, headers, query_args_l)
  File "/usr/lib/python2.6/site-packages/boto/s3/bucket.py", line 230, in _get_k
ey_internal
    response.status, response.reason, '')
boto.exception.S3ResponseError: S3ResponseError: 403 Forbidden

You have new mail in /var/spool/mail/root
[root@MetricInsights-Centos-64-bit conf.d]# _
```

SOLUTION

This error can occur if the date of the VM is wrong. Please check the date and make sure it's up to date. If not, please correct the date and make sure that the latest tools are installed as well (**mi-deploy comp tools master**). Then, rerun the upgrade; it should install successfully. If the issue persists, please contact Metric Insights Support (support@metricinsights.com)

How do I know the update was successful?

A successful upgrade will not have any errors on the command line after running the upgrade. If you have errors, please contact support. The following steps are what Metric Insights considers a "smoke test" to quickly verify an upgrade was successful.

1. Logon and access your Homepage

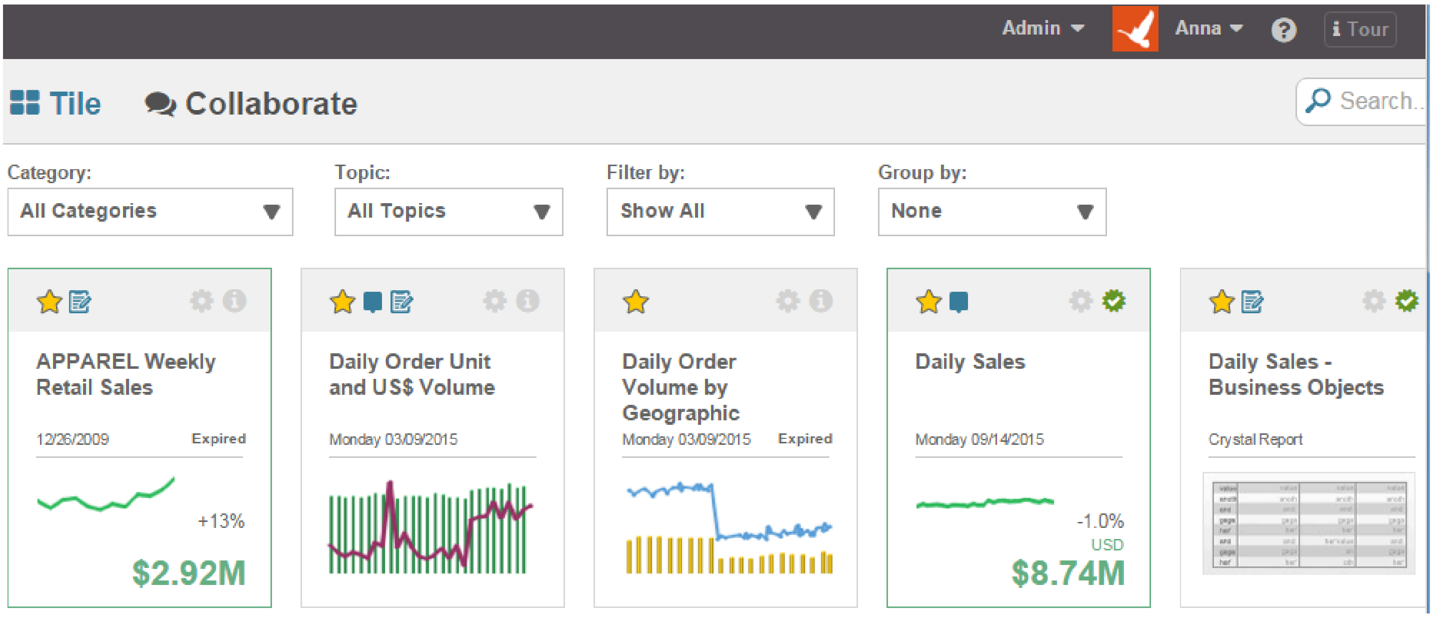

Verify that your tiles appear!

1.1. Verify your version number

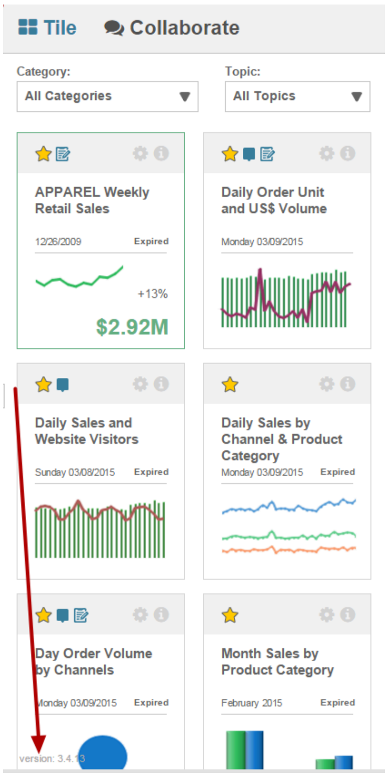

Check that the version number in the bottom left-hand corner of the HomePage is updated to new version

2. Access the Status Monitor

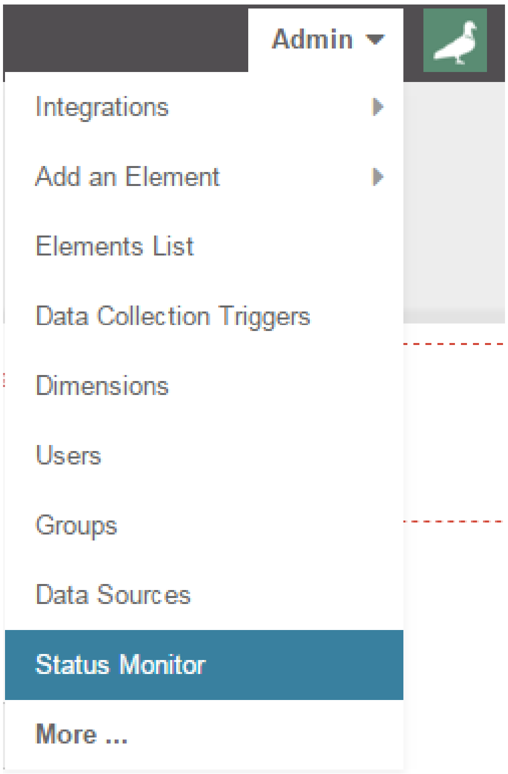

Link from Admin drop-down

2.1. Check Error Messages

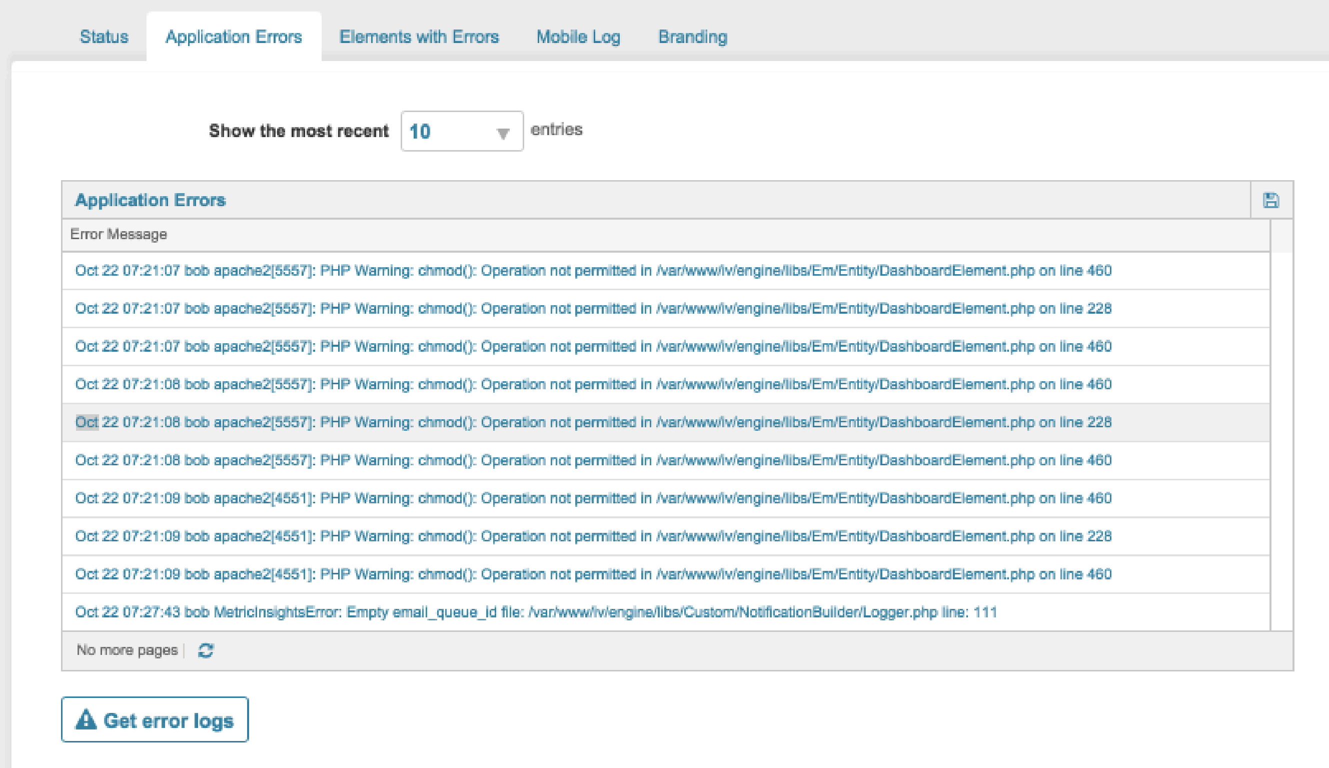

Show the most recent 10 entries

Application Errors

Error Message
Oct 22 07:21:07 bob apache2[5557]: PHP Warning: chmod(): Operation not permitted in /var/www/iv/engine/libs/Em/Entity/DashboardElement.php on line 460
Oct 22 07:21:07 bob apache2[5557]: PHP Warning: chmod(): Operation not permitted in /var/www/iv/engine/libs/Em/Entity/DashboardElement.php on line 228
Oct 22 07:21:07 bob apache2[5557]: PHP Warning: chmod(): Operation not permitted in /var/www/iv/engine/libs/Em/Entity/DashboardElement.php on line 460
Oct 22 07:21:08 bob apache2[5557]: PHP Warning: chmod(): Operation not permitted in /var/www/iv/engine/libs/Em/Entity/DashboardElement.php on line 460
Oct 22 07:21:08 bob apache2[5557]: PHP Warning: chmod(): Operation not permitted in /var/www/iv/engine/libs/Em/Entity/DashboardElement.php on line 228
Oct 22 07:21:08 bob apache2[5557]: PHP Warning: chmod(): Operation not permitted in /var/www/iv/engine/libs/Em/Entity/DashboardElement.php on line 460
Oct 22 07:21:09 bob apache2[4551]: PHP Warning: chmod(): Operation not permitted in /var/www/iv/engine/libs/Em/Entity/DashboardElement.php on line 460
Oct 22 07:21:09 bob apache2[4551]: PHP Warning: chmod(): Operation not permitted in /var/www/iv/engine/libs/Em/Entity/DashboardElement.php on line 228
Oct 22 07:21:09 bob apache2[4551]: PHP Warning: chmod(): Operation not permitted in /var/www/iv/engine/libs/Em/Entity/DashboardElement.php on line 460
Oct 22 07:27:43 bob MetricInsightsError: Empty email_queue_id file: /var/www/iv/engine/libs/Custom/NotificationBuilder/Logger.php line: 111

No more pages |

You can download all the logs using the 'Get error logs' button. Look for new errors.

2.2. Verify that Cron is running

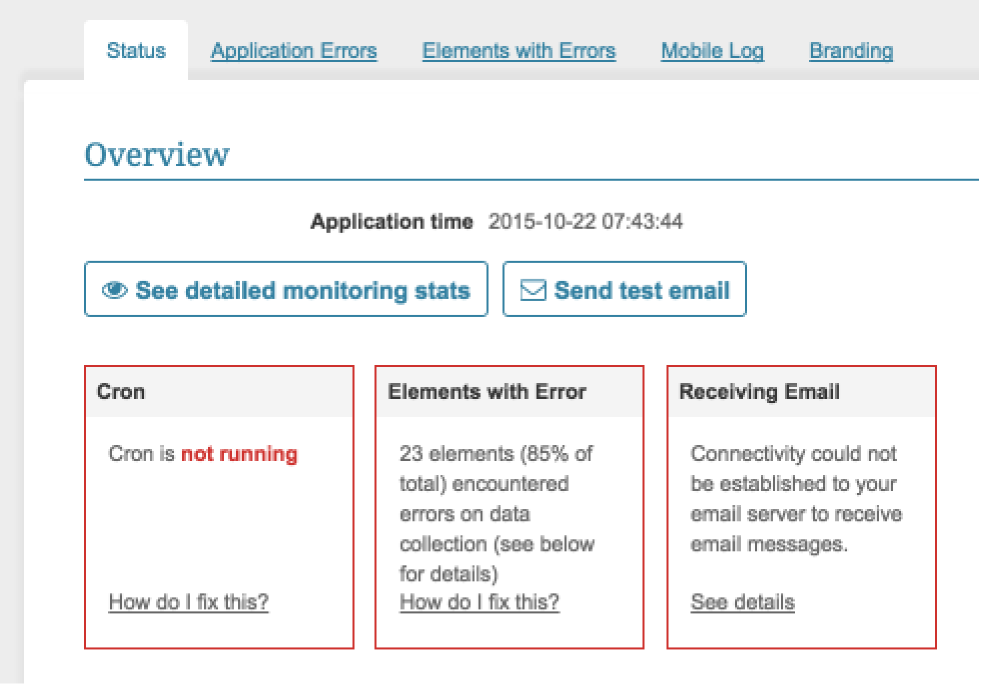

Overview

Application time 2015-10-22 07:43:44

Cron

Cron is **not running**

How do I fix this?

Elements with Error

23 elements (85% of total) encountered errors on data collection (see below for details)

How do I fix this?

Receiving Email

Connectivity could not be established to your email server to receive email messages.

See details

If Cron -- our scheduling service -- is not running you will need to start the service from the command line by running 'service crond start' for Centos or 'service cron start' for Debian versions of Linux.

2.3. Send a test email

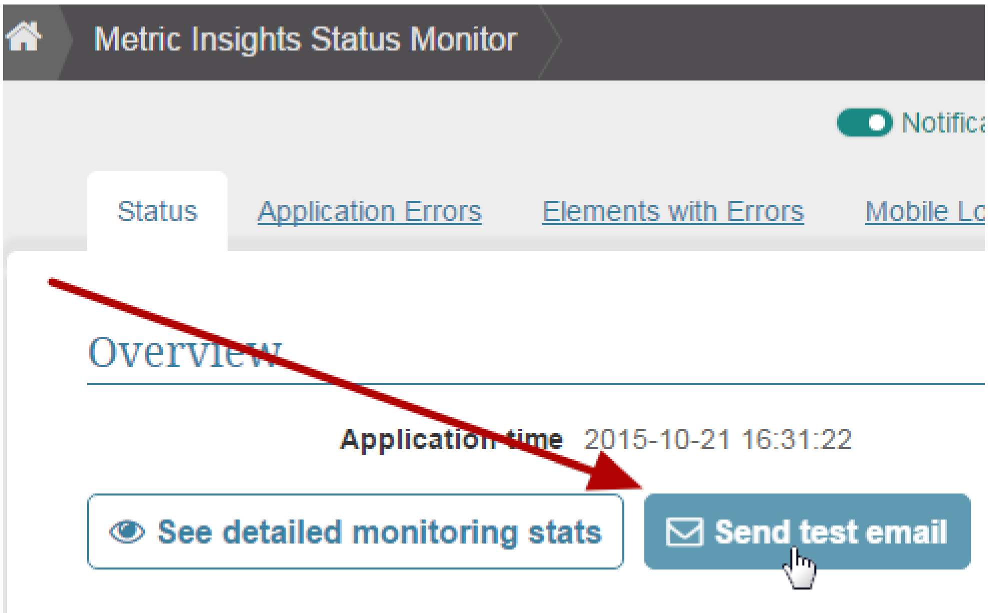

Application time 2015-10-21 16:31:22

Check that email received correctly and in correct format

2.4. Optionally, verify that LDAP is enabled

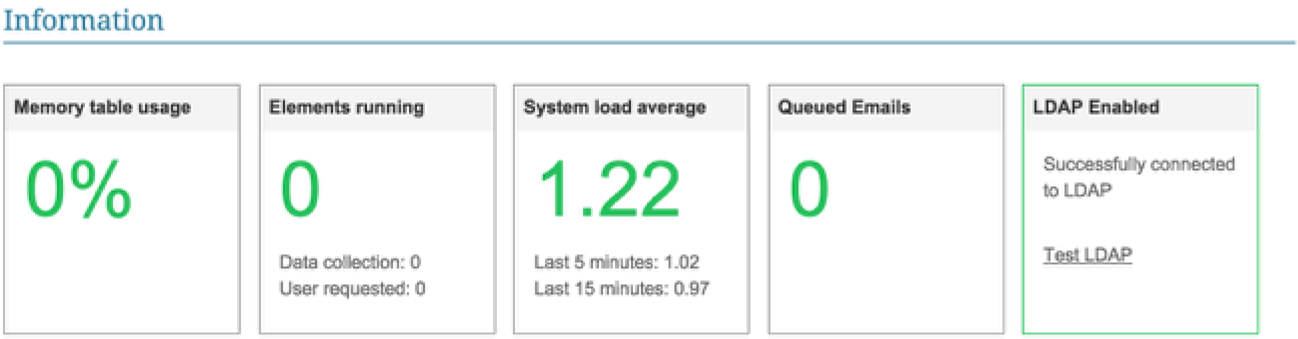

Memory table usage	Elements running	System load average	Queued Emails	LDAP Enabled
0%	0 Data collection: 0 User requested: 0	1.22 Last 5 minutes: 1.02 Last 15 minutes: 0.97	0	Successfully connected to LDAP Test LDAP

3. Access the Metric Editor for an existing element

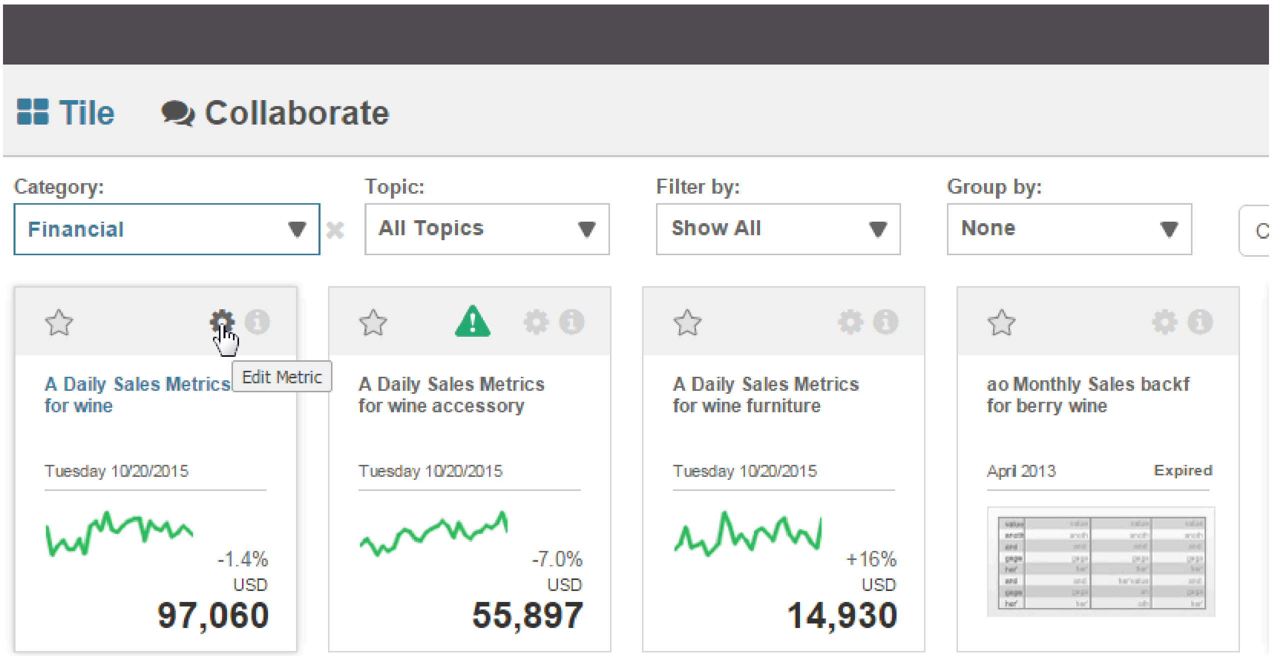
Tile
Collaborate
Edit Metric
A Daily Sales Metrics for wine
Tuesday 10/20/2015
-1.4%
USD
97,060
A Daily Sales Metrics for wine accessory
Tuesday 10/20/2015
-7.0%
USD
55,897
A Daily Sales Metrics for wine furniture
Tuesday 10/20/2015
+16%
USD
14,930
ao Monthly Sales backf for berry wine
April 2013
Expired

Category: Financial

Topic: All Topics

Filter by: Show All

Group by: None

Click the **edit icon** for a metric

3.1. Validate the Element

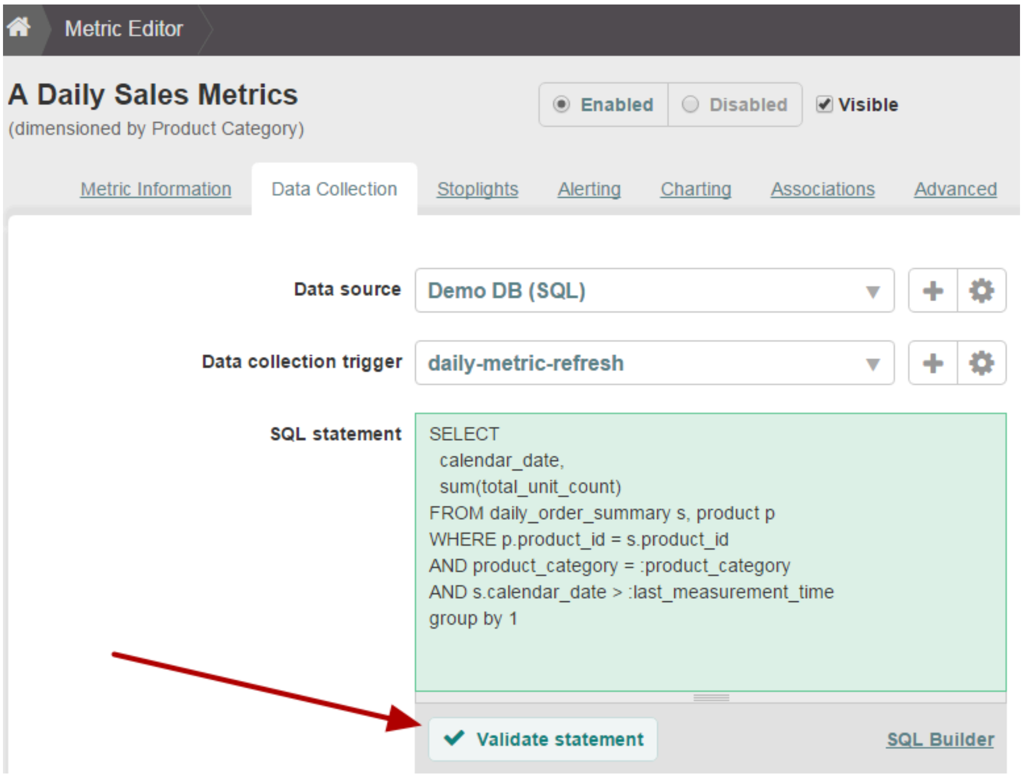

A Daily Sales Metrics

(dimensioned by Product Category)

● Enabled ○ Disabled ☑ Visible

Data source: Demo DB (SQL)

Data collection trigger: daily-metric-refresh

SQL statement:

```
SELECT
  calendar_date,
  sum(total_unit_count)
FROM daily_order_summary s, product p
WHERE p.product_id = s.product_id
AND product_category = :product_category
AND s.calendar_date > :last_measurement_time
group by 1
```

On the *Data Collection tab* and hit "**Validate statement**"

Verify that at least one record is returned.

Note: If you do not have any metrics, you can validate and update a report, or collect an image for an external report.

3.2. Recollect data for the element

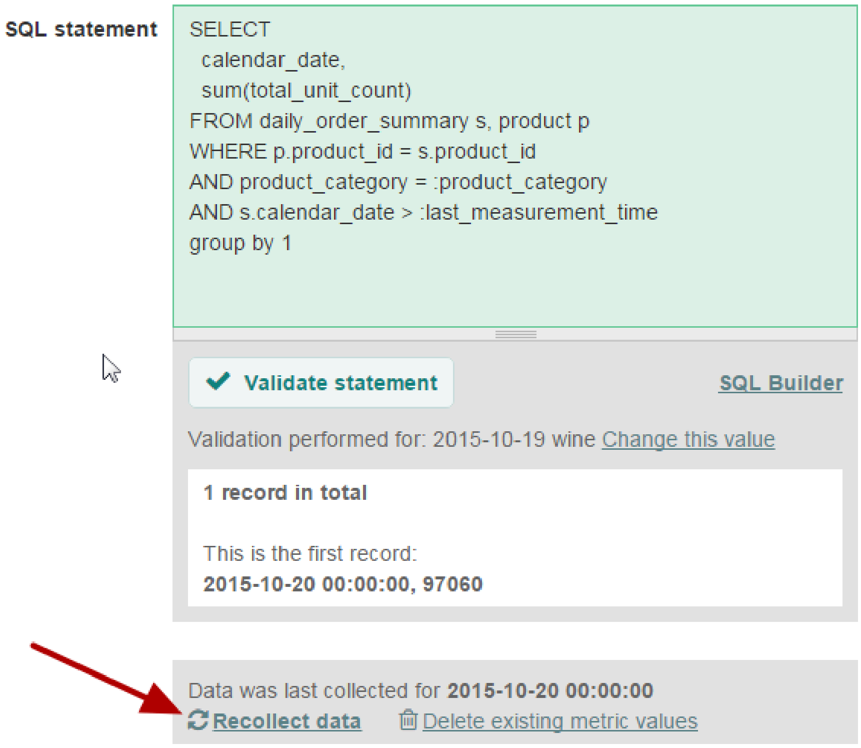

SQL statement

```
SELECT
  calendar_date,
  sum(total_unit_count)
FROM daily_order_summary s, product p
WHERE p.product_id = s.product_id
AND product_category = :product_category
AND s.calendar_date > :last_measurement_time
group by 1
```

Validation performed for: 2015-10-19 wine Change this value

1 record in total

This is the first record:
2015-10-20 00:00:00, 97060

Data was last collected for **2015-10-20 00:00:00**

3.3. Verify that Chart displays correctly

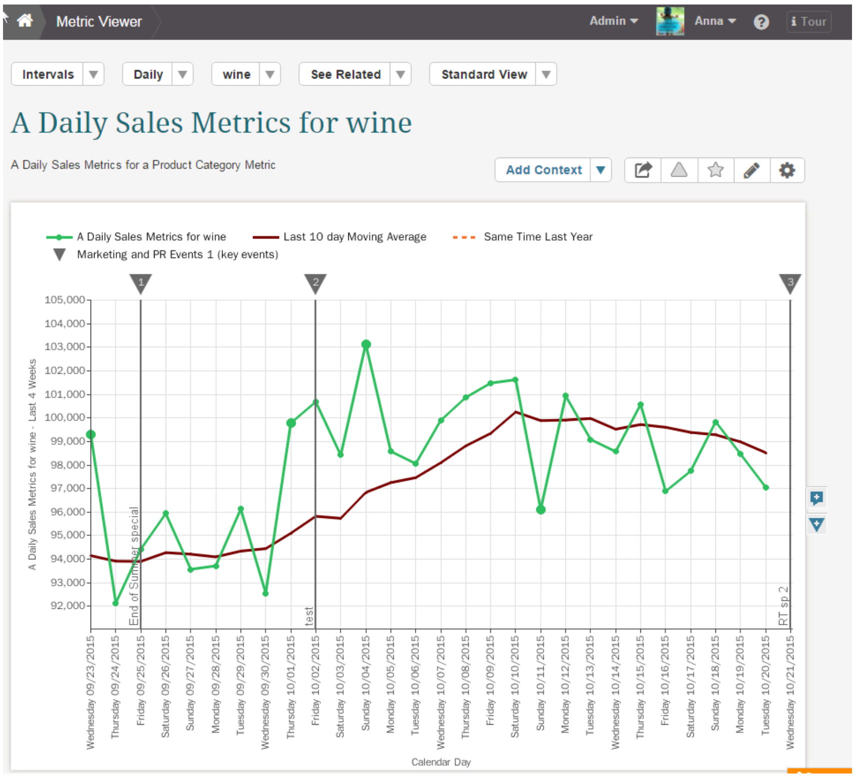

Upgrading-in-place Metric Insights from v5.6.x to v6.2.x

This article covers in detail how to upgrade-in-place Metric Insights from v5.6.x to v6.2.x. This is for non-orchestrated environments of Metric Insights that will run v6 (also known as a *Simple Install*). To ensure the current v5 server is compatible with v6, please see the following help documentation:

- Main Components
- Linux Package Requirements
- System Requirements

In summary, the steps to upgrade entail the following:

1. Create a backup of Metric Insights v5
2. Stop Apache and MySQL services (if MySQL is running locally)
3. Download the Metric Insights v5.6.3 installer and unpack
4. Run the uninstaller script to remove v5 from the server
5. Download the latest Metric Insights v6 installer and unpack
6. Run the installer to deploy v6
7. Restore the v5 backup in v6

1. Create a backup of Metric Insights v5

Create backup on MI v 5.6.x and move it out of the /var/backups/mi-app-backups folder because this folder will be purged

```
$ mi-app-backup -vv
$ mv /var/backups/mi-app-backups/<your-v5-backup.tar.gz> /home/myuser
```

Move or copy any critical files out of the /opt/mi filesystem to ensure a backup copy exists after the uninstall. Place the files or copies in your home directory or in any directory other than /opt/mi. Here's an example of moving the files to /home/myuser/MIv5/

```
$ mv /opt/mi/iv/engine/config/saml.php /home/myuser/MIv5/
$ cp -Rf /opt/mi/iv/data/themes/* /home/myuser/MIv5/
$ cp -Rf /opt/mi/ssl/* /home/myuser/MIv5/
```

2. Stop Apache and MySQL services (if MySQL is running locally)

Stop and disable Apache and MySQL services. Note the service names can differ depending on what flavor of Linux the server is running:

- Debian/Ubuntu: *apache2, mysql/mariadb*
- RedHat/CentOS: *httpd, mysqld/mariadb*

```
$ systemctl stop mariadb
$ systemctl disable mariadb
$ systemctl stop httpd
$ systemctl disable httpd
```

💡 Check ports 80/443/3306 to ensure no lingering Apache and MySQL processes are occupying the ports

```
netstat -tulpn
```

3. Download the Metric Insights v5.6.3 installer and unpack

Contact Metric Insights Support for the v5.6.3 installer. Once you receive the download link, download the install package and copy it to the MI server.

Untar the MI v5.6.3 install package by running the following *tar* command:

```
$ tar xvf MetricInsights-Installer-v5.6.3-Full.tar.gz
```

Move to the v5.6.3 install directory:

```
$ cd MetricInsights-Installer-v5.6.3-Full
```

4. Run the uninstaller script to remove v5 from the server

Use the *uninstaller.py* script to uninstall MI v5.6.x

```
$ ./uninstaller.py --purge --drop-db -vv
```

5. Download the latest Metric Insights v6 installer and unpack

Contact Metric Insights Support for the latest v6 installer. Once you receive the download link, download the install package and copy it to the MI server.

Untar the MI v6 install package by running the following *tar* command:

```
$ tar xvf MetricInsights-Installer-v6.2.1-Full.tar.gz
```

Move to the v6.2.1 install directory:

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

6. Run the installer to deploy v6

Before running the installer, confirm whether the mysql database will run locally, or from a remote database server. Please also identify what timezone the Metric Insights application should be in. Note the timezone values to use below.

If mysql will locally on the same server, run the installer as follows:

```
$ ./installer.py --bind-address 0.0.0.0 --timezone <timezone> --dp-mysql-option
<timezone> -vv
```

If mysql will run from a remote database server, run the installer as shown below. Note, MySQL 8 must already be installed on the remote db server.

```
$ ./installer.py --bind-address 0.0.0.0 --db-hostname <remote-db-server> --db-user
<mysql-root-user> --db-password <mysql-root-password> --timezone <timezone> --dp-mysql-
option <timezone> --components web,dataprocessor,seed,data-analyzer,monitoring -vv
```

7. Restore the v5 backup in v6

Restore your MI v5.6.x backup in MI v6. The backup file can be restored directly from the host server. There is no need to enter the v6 docker containers.

```
$ mi-app-restore /home/myuser/<your-v5-backup.tar.gz> --convert --skip-compatibility-
check --force -vv
```

For more information about Metric Insights v6, please see the following articles:

- v6 Features
- Docker commands to use with v6
- v6 Release Notes

Operating Systems

Getting a "Failed to fetch Backports..." Error when upgrading from Debian 7.4 to latest

This article describes how to properly update from an early release of Debian 7 (wheezy) to the latest.

ISSUE

When trying to update Debian 7 to the latest minor release (7.9) using the command *sudo apt-get update && apt-get upgrade,* the update stops with the following errors:

W: Failed to fetch http://backports.debian.org/debian-backports/dists/wheezy-backports/main/binary-amd64/Packages 404 Not Found [IP: 140.211.15.34 80]

W: Failed to fetch http://backports.debian.org/debian-backports/dists/wheezy-backports/contrib/binary-amd64/Packages 404 Not Found [IP: 140.211.15.34 80]

W: Failed to fetch http://backports.debian.org/debian-backports/dists/wheezy-backports/non-free/binary-amd64/Packages 404 Not Found [IP: 140.211.15.34 80]

Why is this happening?

RESOLUTION

The issue is with the file that lists the sources from which Debian packages can be obtained: */etc/apt/sources.list*

The Backport package for Debian Wheezy is now part of the main archive so your sources.list file must be updated. Open /etc/apt/sources.list in vim editor and you'll find the following source:

deb http://backports.debian.org/debian-backports wheezy-backports main contrib non-free

Comment this out and add the following source instead:

deb http://ftp.us.debian.org/debian/ wheezy-backports main contrib non-free

See the image below to see an example updated sources.list file. Once you've added the new source, save and rerun *sudo apt-get update && apt-get upgrade.* The update to Debian 7.9 should complete successfully!

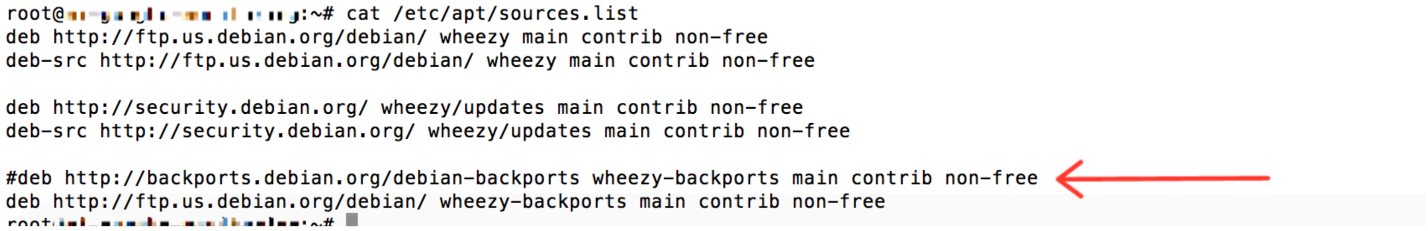

```
root@[redacted]:~# cat /etc/apt/sources.list
deb http://ftp.us.debian.org/debian/ wheezy main contrib non-free
deb-src http://ftp.us.debian.org/debian/ wheezy main contrib non-free

deb http://security.debian.org/ wheezy/updates main contrib non-free
deb-src http://security.debian.org/ wheezy/updates main contrib non-free

#deb http://backports.debian.org/debian-backports wheezy-backports main contrib non-free
deb http://ftp.us.debian.org/debian/ wheezy-backports main contrib non-free
root@[redacted]:~#
```

GPG Error when updating Debian or Metric Insights Tools

This error can come up when updating Debian or Metric Insights Tools.

ISSUE

Getting the following error when updating Debian or updating Metric Insights Tools (mi-deploy comp tools master):

W: GPG error: http://apt.datadoghq.com stable Release: The following signatures couldn't be verified because the public key is not available: NO_PUBKEY 226AE980C7A7DA52

How do I get the public key and finish the update?

RESOLUTION

The public key is actually listed in the error itself. To add the public key, run the following command with the public key specified at the end:

sudo apt-key adv --keyserver keyserver.ubuntu.com --recv-keys <public key>

In the example above, the command to run would look like:

sudo apt-key adv --keyserver keyserver.ubuntu.com --recv-keys 226AE980C7A7DA52

Once the key is added, can you rerun the update to Debian or Metric Insights Tools. Please note that the reason this key is needed is because the apt-get system has a new feature that guarantees the authenticity of servers when pulling updates from Debian sources.

Encryption

How to use Azure Key Vault to encrypt keys

How to...

To add an additional layer of security, we'd like to use Azure Key Vault to help safeguard cryptographic keys and secrets used by Metric Insights. Is this possible?

Solution

Yes, this is definitely possible. Metric Insights recently introduced changes in version 3.3 to leverage Azure Key Vault. In total, there are two keys for Metric Insights: application encryption and file system encryption (for data encryption at rest.)

To enable the use of Azure Key Vault, you will have to make some changes to the Metric Insights configuration. Note that this requires root level access to the Metric Insights server. Here are the steps to enable Azure Key Vault in Metric Insights CentOS systems:

1.) Make sure that the Key Vault is set up in Azure.

Note that there are some powershell commands necessary to set this up. View the following Azure document for more: http://blogs.technet.com/b/kv/archive/2015/06/02/azure-key-vault-step-by-step.aspx.

You will need to create a key for the MI application key, we'll refer to this as **AppKey** in this article. If you'd like to enable file system encrpytion for at rest encryption of data, you will need another key to be used for file system encryption. We'll refer to this as the **FileSystemKey**

2.) Initialize the AppKey in your Azure Key Vault to be the same value as the key stored in /var/www/aes_password. (You can change this value later)

3.) Now, tell Metric Insights to use your new Azure Key vault by adding the following [encryption] section to */etc/mi/insight.conf* in Metric Insights:

```
[encryption]
key_storage = azure
client_id = your-client-id
client_key = 'your-client-secret-key'
auth_url = https://login.windows.net/<your-auth-code>
vault_name = Vault-Name
key_name = AppKey-Name
```

You will need to get this information from your Azure Management Portal. The Client-Id/Client-Key, for example, will be the Oauth2 Application Id and Secret for the application you created for MI in Azure AD. The **auth_url** will be the 'OAuth 2 Authorization Endpoint' for the application. **<TODO: screenshot example from Azure Portal>**

4.) Once insight.conf is saved with the new [encryption] section, verify that you can connect to one of your data sources. You can do this by testing a data source connection or collecting data from an existing metric. Ex: **<SOME URL>**

If this step fails for some reason, you have not configured the [encryption] section of insight.conf correctly.

5.) If step 4 succeeded and you can correctly connect to your data sources, then you can now change the key to your own encryption key. For performance, it's best to pick a key that is somewhere between 16 and 80 bytes, though you can choose larger values for your key if you wish.

To change the application encryption key, you can use the mi-crypt utility that comes with Metric Insights. For example, to update the key to a new random 64 byte value, run:

```
/usr/local/lib/mi/bin/mi-crypt setnewkey default $(openssl rand -hex 64)
```

TODO: NEW SECTION ON FILE SYSTEM ENCRYPTION

This section will explain the steps necessary to setup file encryption:

1.) Get the appropriate *ecryptfs rpm* (send a request to Metric Insights Support) and save to localhost

2.) Install the ecryptfs rpm from localhost: **yum install-local ecryptfs-utils-96-1.el6.x86_64.rpm**

3.) Set up your FileSystemKey in Azure Key Vault. This should be a separate key from your AppKey that was created for application encryption. Due to performance reasons, your FileSystemKey value should not exceed 80 bytes. (You can generate a new key by using `openssl rand -hex <n>` where n is the number of bytes you want). Make sure your FileSytemKey is stored in Azure Key Vault before continuing.
4.) Now tell Metric Insights to use at-rest encryption by setting up an [encryption_fs] section in */etc/mi/insight.conf.* This will be very similar to the AppKey set up above, just with a different key name.

8.) Once insight.conf is saved with your new [encryption_fs] section, you can encrypt the file system with the following command:

```
mi-crypto-mgr enable -p '/usr/local/lib/mi/bin/mi-crypt getkey fs'
```

Updating Keys

Once Azure Key Vault is set, you can update the application key to a new key by running the following command in the Metric Insights instance:

```
/usr/local/lib/mi/bin/mi-crypt setnewkey default <newkey>
```

This will take care of **decrypting** all the secrets stored in the Metric Insights database with the old key, then **re-encrypting** the secrets with the new key.

To update the file system key, use the following command:

```
/usr/local/lib/mi/bin/mi-crypt setnewkey fs <newkey>
```

Authentication

After upgrading to 3.2.830, some LDAP users can't log in

Issue

A client upgraded from 3.2.805 to 3.2.830 recently. Now, some ldap users can log in and others can't.
https://insight.lindenlab.com/home/

For those that can't, as soon as they enter their credentials, they are re-directed back to the login screen.

Client confirms ldap is fine on their side. Users report no issues with ldap for other systems there.

List of LDAP USERS where it:
Works (example): mitchell, fereshteh, steven, sanghavi
Doesn't Work (example): venellyn, matias, dee, garry, fan

Resolution

The issue is a combination of MI defaulting to *secure php session cookies* in 3.2.812, which get set for a specific domain, and Linden using an apache proxy server for MI. They point browsers to https://insight.lindenlab.com which proxies requests over to the MI machine which lives at https://int.insight.lindenlab.com.

This means that if a user does not have a valid PHPSESSID cookie in their browser (they've never logged into MI, or haven't done so recently) MI will try to set a php session cookie for the domain $SERVER[HOST_NAME], which is int.insight.lindenlab.com. This is a *different domain* than what's in the web browser, so the web browser ignores it and refuses to set the PHPSESSID cookie. Without a PHPSESSID cookie, MI can still authenticate, but all further pages think that you haven't yet authenticated, throwing the user back to the login page!

The solution then is to turn off SECURE_SESSION_COOKIES so that MI doesn't force the session cookie to be limited to int.insight.lindenlab.com anymore. PHP will just set a generic PHPSESSID cookie to whatever domain is in the address bar of the URL. This enables a valid PHPSESSID to exist for further MI requests and Linden's proxying works as expected.

To turn off SECURE_SESSION_COOKIES, edit **const.php** and set **SECURE_SESSION_COOKIES = 'N'**.

/var/www/iv/engine/config/const.php

(If you want SECURE_SESSION_COOKIES to be enabled and still want to proxy, you'll need to set a ProxyPassReverseCookieDomain directive in the proxy server's mod_proxy configuration.)

What is the difference between 'default-sp' and 'default-signed-sp' for SAML?

QUESTION

After upgrading to Metric Insights version 4.1, I noticed that there are 2 options for SAML auth sources:

- default-sp
- default-signed-sp

What is the difference between the two and which one do I use?

RESOLUTION

Before v4.1, the Metric Insights application supported only 1 option for *auth source*: **default-sp**. That is used for unsigned requests between SAML Identify Providers (like OKTA, Oracle Access Manager) and the Metric Insights application (Service Provider).

After v4.1, the Metric Insights application added additional support for **default-signed-sp**. This auth source is used for exchanging signed requests (with certificates) between the SAML IdP and Metric Insights (for example, Microsoft ADFS).

Thus, based on the IdP that is used in your organization (signed or unsigned requests) you must choose the appropriate *auth source* to successfully connect Metric Insights to your SAML IdP.

Note: If you do not know which auth source to use, we recommend trying default-sp first.

LDAP SSL certificates

⚠ For successful communication between your *LDAP server* and the *Metric Insights server*, you need to create an **LDAP configuration file** and add it to the **/etc/openldap/** directory on your LDAP server.

1. ldap.conf file > certificates directory

1. Access your **ldap.conf** file and identify where the SSL certificate should be placed
2. Change directory with `cd /etc/openldap/cacerts`

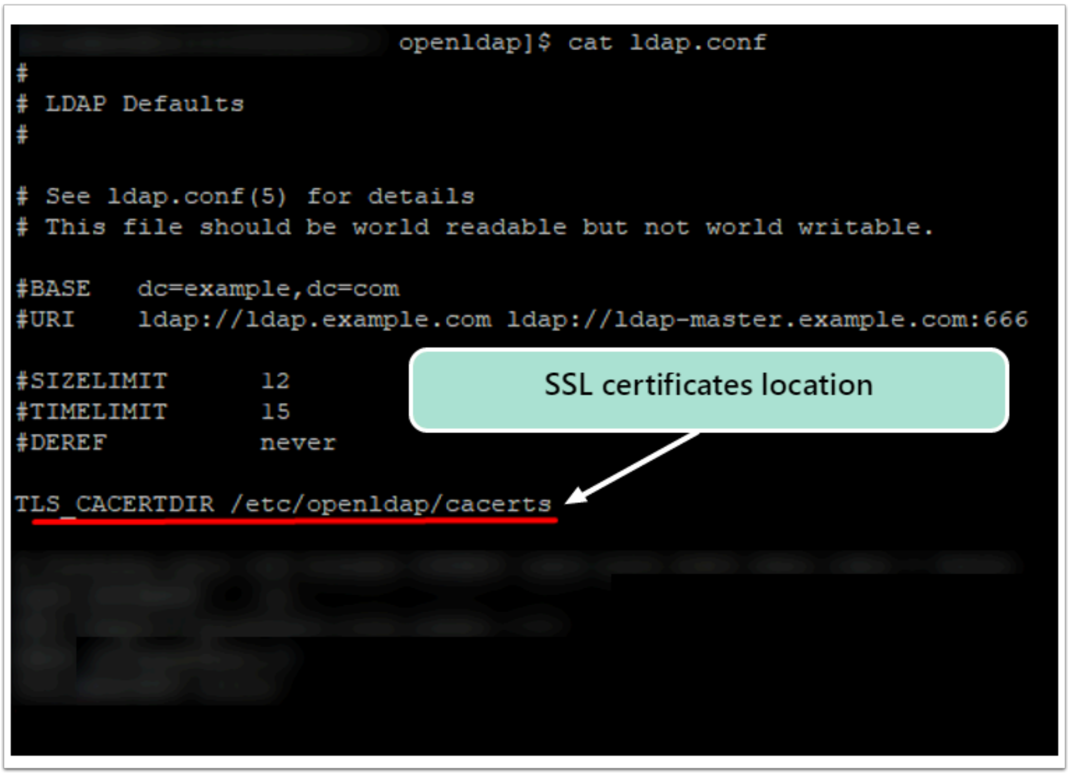

```
[redacted] openldap]$ cat ldap.conf
#
# LDAP Defaults
#

# See ldap.conf(5) for details
# This file should be world readable but not world writable.

#BASE   dc=example,dc=com
#URI    ldap://ldap.example.com ldap://ldap-master.example.com:666

#SIZELIMIT      12
#TIMELIMIT      15
#DEREF          never

TLS_CACERTDIR /etc/openldap/cacerts
[redacted]
```

2. Create a file for your SSL certificate

While in the **/etc/openldap/cacerts** directory, create a new file using `touch GlobalAD.crt`

3. Place your certificate in GlobalAD.crt

1. Open the **GlobalAD.crt** file in Linux editor tool
2. Paste your **SSL certificate** to GlobalAd.crt

4. Set certificate ownership to "apache"

Change ownership of the certificate to **apache:apache**, using the `chown` command

5. If needed, add TLS_REQCERT allow

If you are using self-signed certificates, add `TLS_REQCERT allow` to `/etc/openldap/ldap.conf`

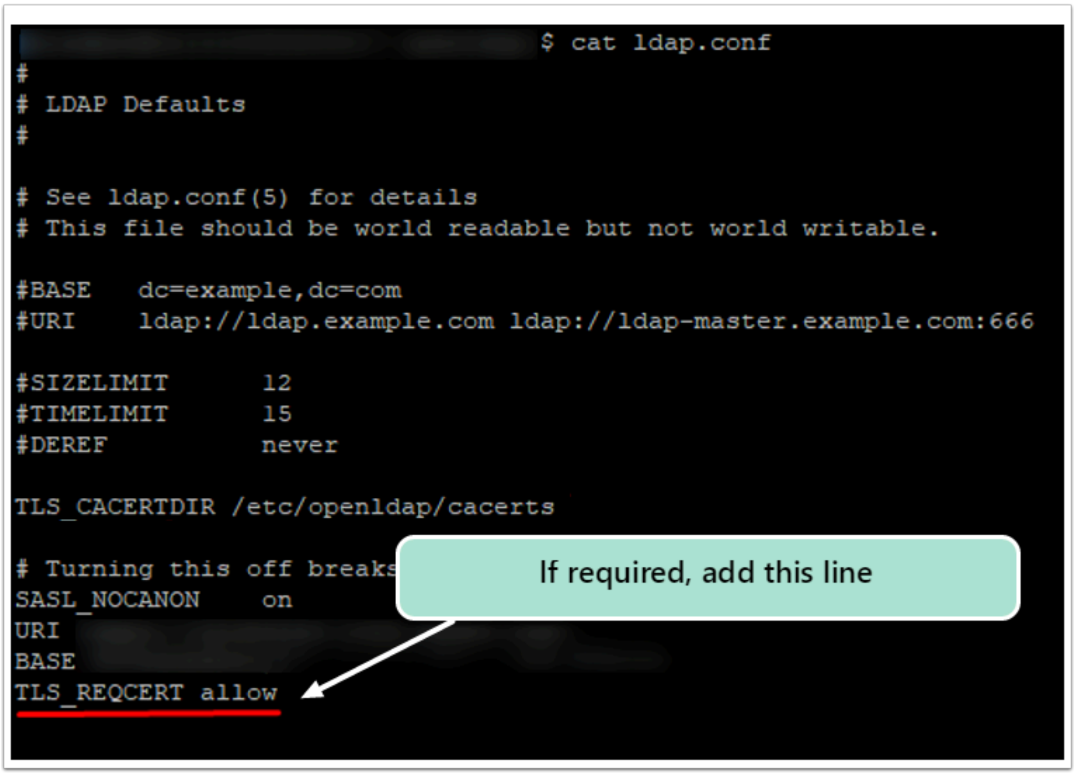

```
[redacted] $ cat ldap.conf
#
# LDAP Defaults
#

# See ldap.conf(5) for details
# This file should be world readable but not world writable.

#BASE   dc=example,dc=com
#URI    ldap://ldap.example.com ldap://ldap-master.example.com:666

#SIZELIMIT      12
#TIMELIMIT      15
#DEREF          never

TLS_CACERTDIR /etc/openldap/cacerts

# Turning this off breaks
SASL_NOCANON    on
URI [redacted]
BASE
TLS_REQCERT allow
```

If required, add this line

How to configure Okta for MI SAML SSL setup

SAML implementation by Okta is one of the Identity Provider options to use in order to set up SAML Single Sign-On (SSO) in Metric Insights. SAML Identity Provider is used to generate a Metadata key required for MI SAML integration.

Configure Okta to generate a Metadata key

ℹ Okta admins can configure Metric Insights app profile in Okta console to generate a Metadata key.

1. Log in as Admin to your Okta console.

2. Create a new Application: go to **Admin** -> **Applications** -> **Add Application**

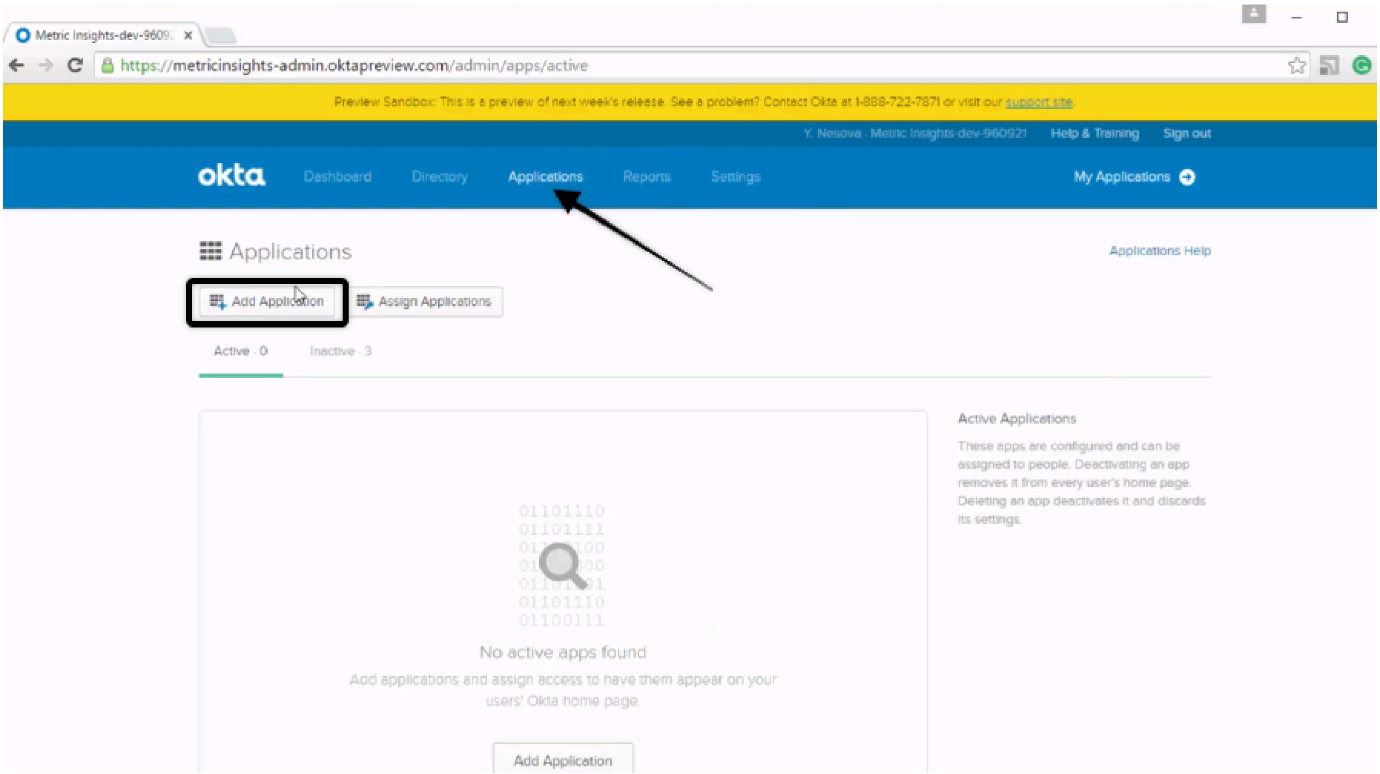

https://metricinsights-admin.oktapreview.com/admin/apps/active

Preview Sandbox: This is a preview of next week's release. See a problem? Contact Okta at 1-888-722-7871 or visit our support site.

Active Applications

These apps are configured and can be assigned to people. Deactivating an app removes it from every user's home page. Deleting an app deactivates it and discards its settings.

No active apps found

Add applications and assign access to have them appear on your users' Okta home page.

3. Click **Create New App** button, in the pop-up choose SAML authentification method and proceed.

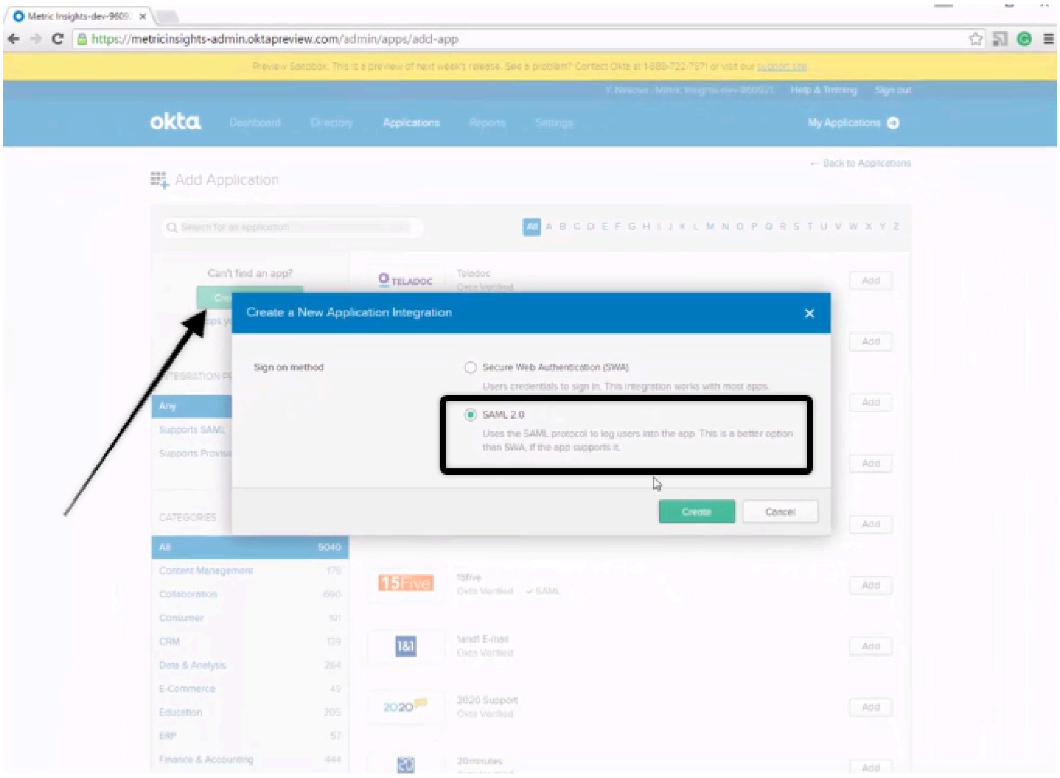

Create a New Application Integration

Sign on method

○ Secure Web Authentication (SWA)
Users credentials to sign in. This integration works with most apps.

● SAML 2.0
Uses the SAML protocol to log users into the app. This is a better option than SWA, if the app supports it.

4. Provide a meaningful app name and move to the **Configure SAML** step.

5. Set up the **General** parameters for the Okta SAML configuration profile:

   5.1. From the MI app Metadata XML within the ***AssertionConsumerService*** section find the ***Location*** parameter.

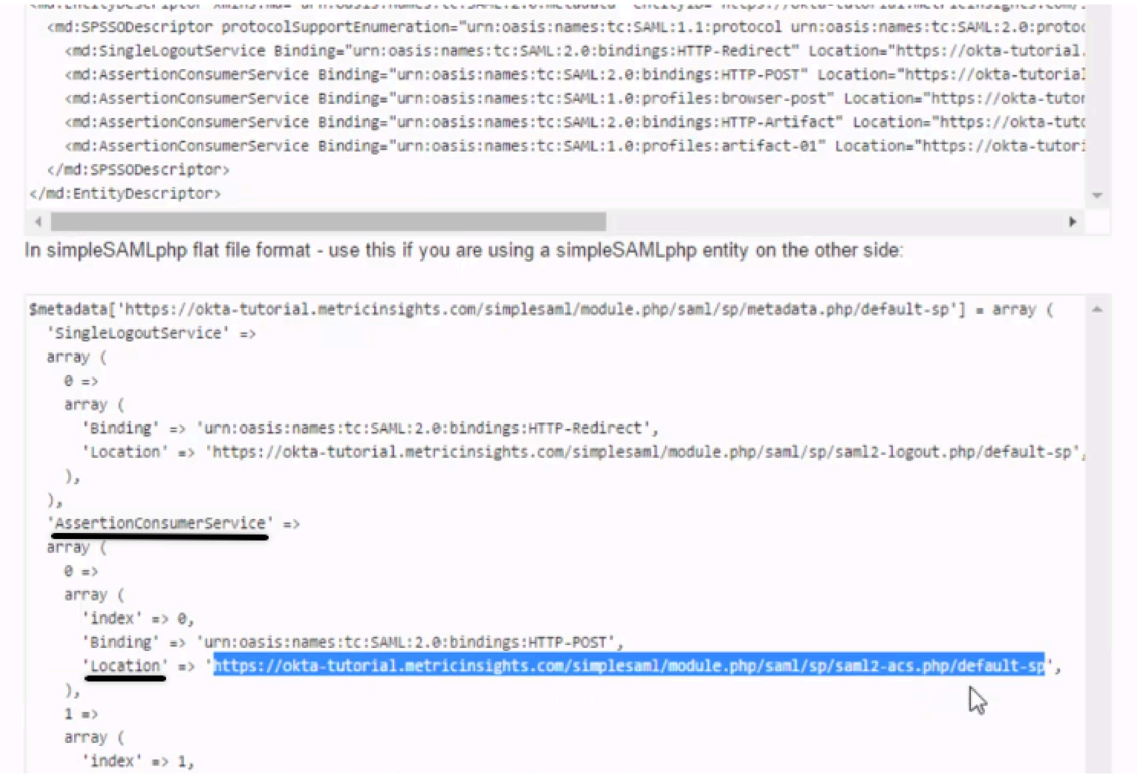

```
  <md:SPSSODescriptor protocolSupportEnumeration="urn:oasis:names:tc:SAML:1.1:protocol urn:oasis:names:tc:SAML:2.0:protoc
    <md:SingleLogoutService Binding="urn:oasis:names:tc:SAML:2.0:bindings:HTTP-Redirect" Location="https://okta-tutorial.
    <md:AssertionConsumerService Binding="urn:oasis:names:tc:SAML:2.0:bindings:HTTP-POST" Location="https://okta-tutorial
    <md:AssertionConsumerService Binding="urn:oasis:names:tc:SAML:1.0:profiles:browser-post" Location="https://okta-tutor
    <md:AssertionConsumerService Binding="urn:oasis:names:tc:SAML:2.0:bindings:HTTP-Artifact" Location="https://okta-tut
    <md:AssertionConsumerService Binding="urn:oasis:names:tc:SAML:1.0:profiles:artifact-01" Location="https://okta-tutori
  </md:SPSSODescriptor>
</md:EntityDescriptor>
```

In simpleSAMLphp flat file format - use this if you are using a simpleSAMLphp entity on the other side:

```
$metadata['https://okta-tutorial.metricinsights.com/simplesaml/module.php/saml/sp/metadata.php/default-sp'] = array (
  'SingleLogoutService' =>
  array (
    0 =>
    array (
      'Binding' => 'urn:oasis:names:tc:SAML:2.0:bindings:HTTP-Redirect',
      'Location' => 'https://okta-tutorial.metricinsights.com/simplesaml/module.php/saml/sp/saml2-logout.php/default-sp',
    ),
  ),
  'AssertionConsumerService' =>
  array (
    0 =>
    array (
      'index' => 0,
      'Binding' => 'urn:oasis:names:tc:SAML:2.0:bindings:HTTP-POST',
      'Location' => 'https://okta-tutorial.metricinsights.com/simplesaml/module.php/saml/sp/saml2-acs.php/default-sp',
    ),
    1 =>
    array (
      'index' => 1,
```

5.2. Set its value for the **Single Sign On URL** parameter in Okta. Make sure that **Use this for Recipient URL and Delivery URL** checkbox is checked.

5.3. For the **Audience URI** parameter in Okta set the **Entity ID** link from the **Federation** tab on MI simpleSAML installation page.

5.4. **Default Relay State** parameter is optional, but you can specify the link to your MI app here (https://*<serverIP>*or*<DNS.com>*)

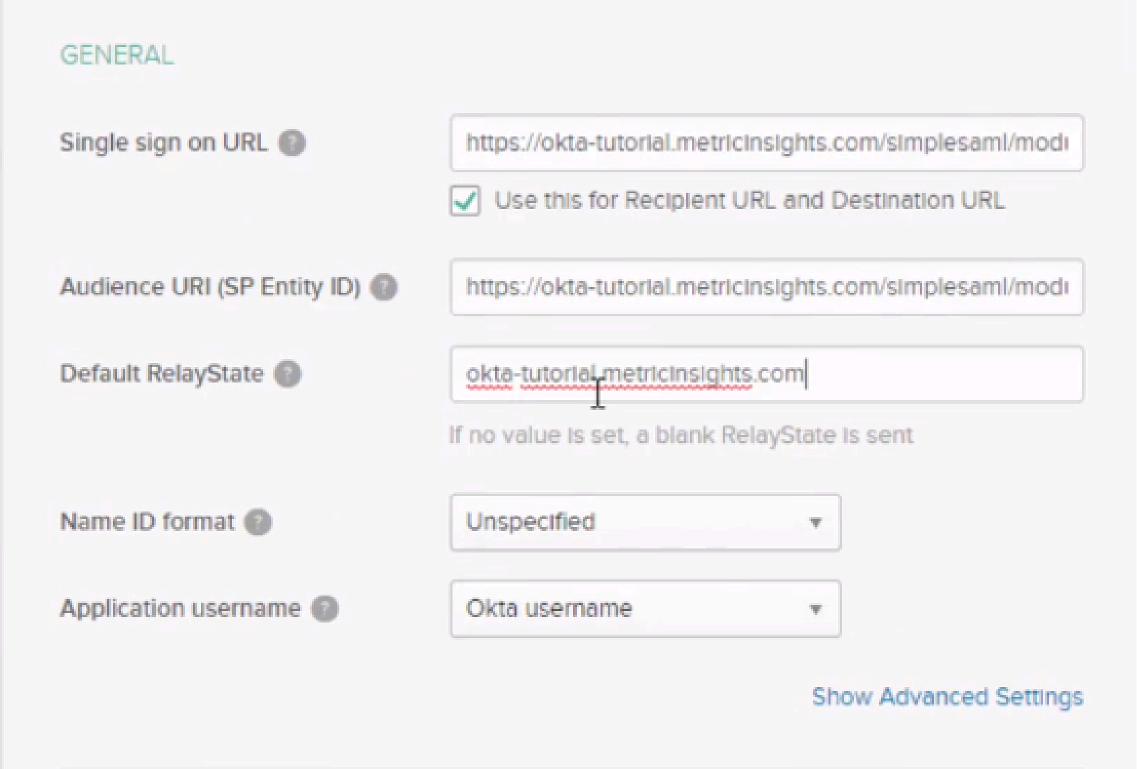

GENERAL

Single sign on URL: https://okta-tutorial.metricinsights.com/simplesaml/modı

☑ Use this for Recipient URL and Destination URL

Audience URI (SP Entity ID): https://okta-tutorial.metricinsights.com/simplesaml/modı

Default RelayState: okta-tutorial.metricinsights.com

If no value is set, a blank RelayState is sent

Name ID format: Unspecified

Application username: Okta username

5.5. Specify the following **Attribute Statements**:

- *firstName*
- *lastName*
- *uid (user.login)*
- *email*

5.6. Proceed to the next step within Okta SAML configuration process.

6. On the Feedback step set the following parameter values:

- For the ***Are you a customer or partner*** choose *I'm an Okta customer adding an internal app*
- *For the **App type*** pick *This is an internal app that we have created*

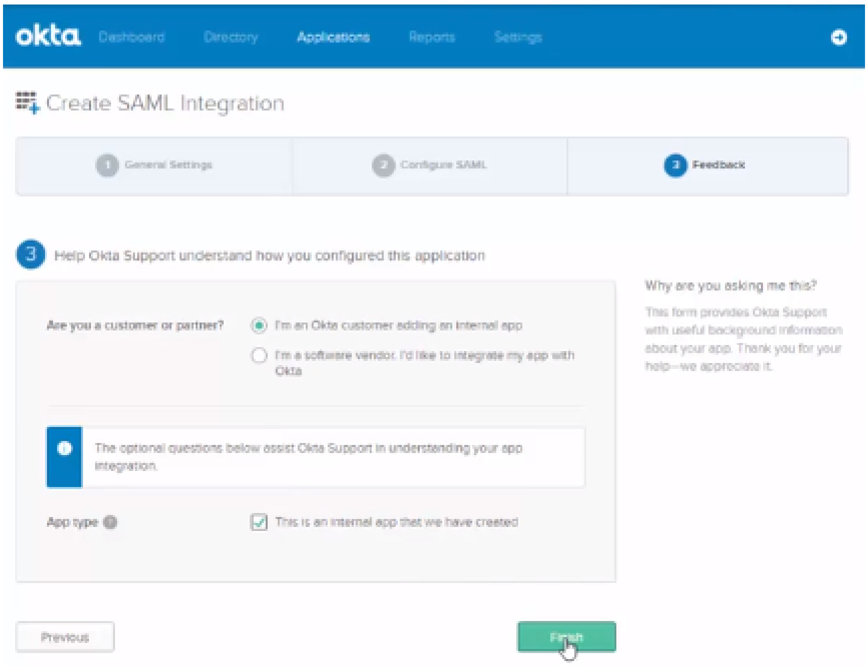

Create SAML Integration

3 Help Okta Support understand how you configured this application

Are you a customer or partner? ● I'm an Okta customer adding an internal app ○ I'm a software vendor. I'd like to integrate my app with Okta

The optional questions below assist Okta Support in understanding your app integration.

App type ☑ This is an internal app that we have created

Why are you asking me this?

This form provides Okta Support with useful background information about your app. Thank you for your help—we appreciate it.

7. Finish the Okta SAML configuration process.

Get the IdP Metadata key in Okta app

Okta admins can get the link to the required **.xml file** on the SAML integration profile page in Okta.

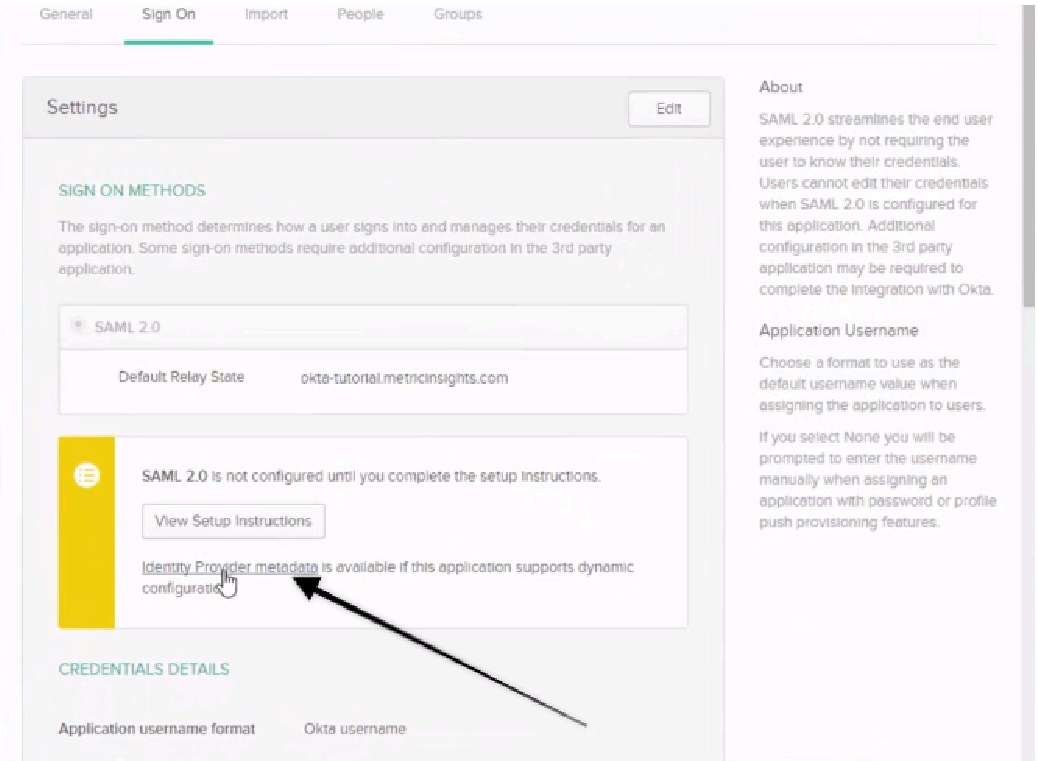

Settings

SIGN ON METHODS

The sign-on method determines how a user signs into and manages their credentials for an application. Some sign-on methods require additional configuration in the 3rd party application.

● SAML 2.0

Default Relay State okta-tutorial.metricinsights.com

SAML 2.0 is not configured until you complete the setup instructions.

Identity Provider metadata is available if this application supports dynamic configuration.

CREDENTIALS DETAILS

Application username format Okta username

About

SAML 2.0 streamlines the end user experience by not requiring the user to know their credentials. Users cannot edit their credentials when SAML 2.0 is configured for this application. Additional configuration in the 3rd party application may be required to complete the integration with Okta.

Application Username

Choose a format to use as the default username value when assigning the application to users.

If you select None you will be prompted to enter the username manually when assigning an application with password or profile push provisioning features.

⚠ Check that the Metadata key file (the **.xml file** donloaded from Okta) has the *firstName*, *lastName* and *Email* attributes defined. If they are not specified, they are to be set directly in the final MI SAML configuration file **saml.php**.

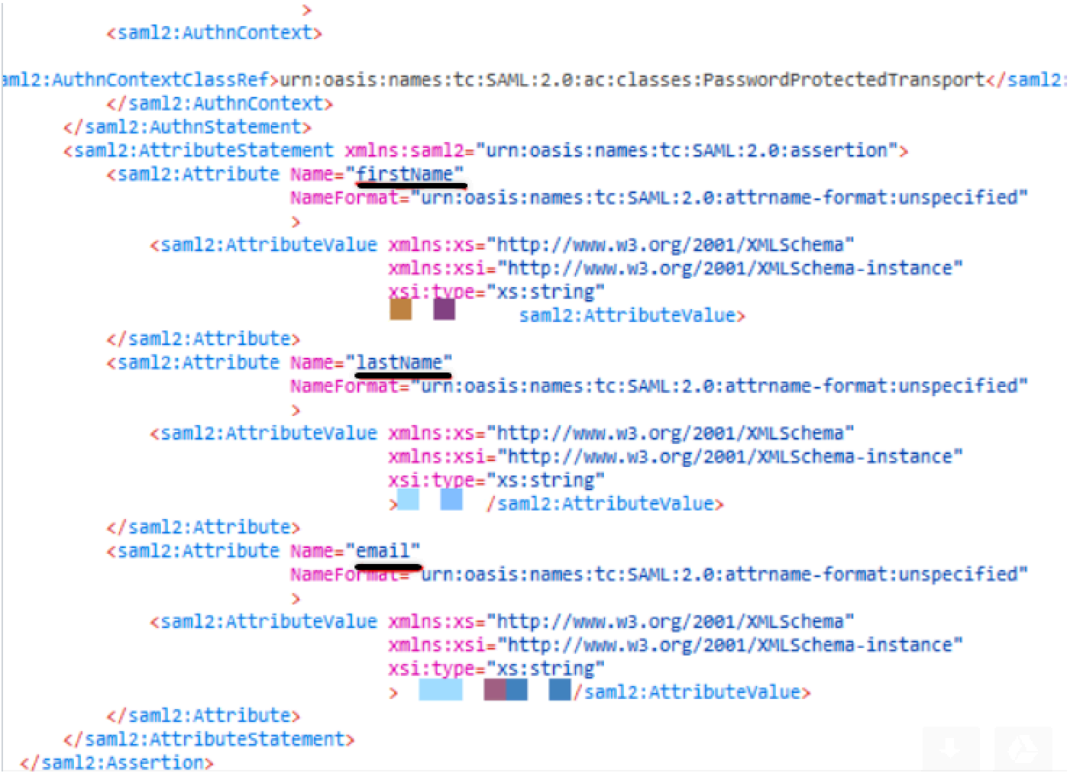

```
                            >
         <saml2:AuthnContext>

aml2:AuthnContextClassRef>urn:oasis:names:tc:SAML:2.0:ac:classes:PasswordProtectedTransport</saml2:
         </saml2:AuthnContext>
     </saml2:AuthnStatement>
     <saml2:AttributeStatement xmlns:saml2="urn:oasis:names:tc:SAML:2.0:assertion">
         <saml2:Attribute Name="firstName"
                          NameFormat="urn:oasis:names:tc:SAML:2.0:attrname-format:unspecified"
                          >
             <saml2:AttributeValue xmlns:xs="http://www.w3.org/2001/XMLSchema"
                                   xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
                                   xsi:type="xs:string"
                                   [redacted]     saml2:AttributeValue>
         </saml2:Attribute>
         <saml2:Attribute Name="lastName"
                          NameFormat="urn:oasis:names:tc:SAML:2.0:attrname-format:unspecified"
                          >
             <saml2:AttributeValue xmlns:xs="http://www.w3.org/2001/XMLSchema"
                                   xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
                                   xsi:type="xs:string"
                                   >[redacted] /saml2:AttributeValue>
         </saml2:Attribute>
         <saml2:Attribute Name="email"
                          NameFormat="urn:oasis:names:tc:SAML:2.0:attrname-format:unspecified"
                          >
             <saml2:AttributeValue xmlns:xs="http://www.w3.org/2001/XMLSchema"
                                   xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
                                   xsi:type="xs:string"
                                   > [redacted]/saml2:AttributeValue>
         </saml2:Attribute>
     </saml2:AttributeStatement>
 </saml2:Assertion>
```

Database

Fine Tuning MySQL Parameters

Adjusting MySQL parameters in /etc/mysql/conf.d

For Metric Insights v 6+ Simple Installation:

- put the `01-metricinsights.cnf` file on the host under `/opt/mi/config/mysql/external_config/`
- restart mysql container with `mi-control restart mysql`

For Metric Insights v5.x

Metric Insights deploys an *01-metricinsights.cnf* file in */etc/mysql/conf.d* with standard mysql parameters. If your dba feels the need to fine tune these parameters* after running mysqltuner for example, **do not** edit 01-insightsettings.cnf.

Each new deploy of MI will overwrite 01-metricinsights.cnf with our default parameters. However, MI will load additional cnf files found in the directory alphabetically/sequentially if there's more than one. In this case 01-insightsettings gets loaded first.

To apply your own parameter settings, create an 02-metricinsights.cnf file (if there already isn't a second *.cnf file). It can also be *yourcompanyname.cnf* because that will get loaded after 01.

MySQL paramenters include:

```
key_buffer_size =
join_buffer_size =
sort_buffer_size =
read_buffer_size =
read_rnd_buffer_size =
innodb_buffer_pool_size =
innodb_log_file_size =
```

NOTE: The example configuration files below are for servers with 12/16/32GB RAM and should be considered baseline specs. Additional fine tuning may be necessary if performance is still an issue (consult with your dba).

If using the 32gb_settings.cnf file, please see the section below on *resizing the innodb_log_file parameter SAFELY.*

(Do not use a file that is provisioned for more RAM than your server has.)

12gb_settings.cnf

16gb_settings.cnf

32gb_settings.cnf

64gb_settings.cnf

Resizing the Innodb Log Files Safely

Simply changing the *innodb_log_file_size* parameter and restarting mysql will not work. If you do, **mysql will actually fail to restart**, producing an error log.

To properly adjust the parameter:

1.) Make sure *innodb_fast_shutdown* is set to *1* (it is by default from Metric Insights). To check, log on to mysql in the MI Instance (or MI Instance db server) and enter:

show variables like '%innodb_fast_shutdown%';

It should be set to 1. If not, set it to 1 by entering:

set global innodb_fast_shutdown=1;

2.) Next, shut down mysql

Debian - *service mysql stop*

Centos - *service mysqld stop*

3.) Remove (but don't delete) the innodb log files *ib_logfile0* and *ib_logfile1* from */var/lib/mysql/*

4.) Modify* the *innodb_log_file_size* parameter in your second cnf file, e.g., *02-insightsettings.cnf* or *yourcompanyname.cnf*

5.) Start mysql

Debian - *service mysql start*

Centos - *service mysqld start*

NOTE: if you do not see an *innodb_log_file_size* parameter, simply enter it in the cnf file (without it explicitely stated, the log file size is set to 5MB by default). For any additional questions, please contact support@metricinsights.com.

Ports used by Metric Insights to connect to different databases

When you establish connection to a database, the Port number will be set by default, based on your choice of JDBC Driver. You can change it if necessary per the following article:

http://help.metricinsights.com/m/Connecting_to_Data_Sources/l/104147-create-new-sql-data-source

Note about HTTP and HTTPS: HTTPS URLs begin with "https://" and use port 443 by default, whereas HTTP URLs begin with "http://" and use port 80 by default.

HTTP is not encrypted and is vulnerable to man-in-the-middle and eavesdropping attacks, which can let attackers gain access to website accounts and sensitive information, and modify webpages to inject malware or advertisements. HTTPS is designed to withstand such attacks and is considered secure against them (with the exception of older, deprecated versions of SSL).

The list of default Ports for SQL data sources

Database	Default port
MySQL Connector/J	3306
PostgreSQL	5432
Netezza	5480
Oracle	1521
Microsoft SQL Server (JTDS)	1433
Sybase (JTDS)	7100
Hadoop Hive	10000
Teradata	1025
Vertica	5433
HSQLDB	9001
Apache Hive2	10000

RDS Parameters for MySQL

The following RDS parameters for MySQL are minimum recommendations only. This is not a *one-size fits all*, but is a good starting point if deploying Metric Insights in your own Amazon cloud.

Note: the *metricinsights* parameters below are for a **db.m3.large**

RDS Parameters for MySQL 5.6

MySQL Parameter Comparison: Default vs. Metric Insights

Parameter	default.mysql5.6	metricinsights
character_set_client	<engine-default>	utf8
character_set_connection	<engine-default>	utf8
character_set_results	<engine-default>	utf8
character_set_server	<engine-default>	utf8
collation_connection	<engine-default>	utf8_unicode_ci
collation_server	<engine-default>	utf8_unicode_ci
innodb_buffer_pool_instances	<engine-default>	4
innodb_buffer_pool_size	{DBInstanceClassMemory*3/4}	2147483648
innodb_flush_log_at_trx_commit	<engine-default>	2
innodb_log_buffer_size	8388608	67108864
innodb_sort_buffer_size	<engine-default>	8000000
join_buffer_size	<engine-default>	16777216
key_buffer_size	16777216	67108864
log_bin_trust_function_creators	<engine-default>	1
max_heap_table_size	<engine-default>	4294967296
query_cache_size	<engine-default>	67108864

query_cache_type	<engine-default>	1
read_buffer_size	262144	2097152
read_rnd_buffer_size	524288	4194304
slow_query_log	<engine-default>	1
sort_buffer_size	<engine-default>	2097152
table_definition_cache	<engine-default>	2048
table_open_cache	<engine-default>	2048
tmp_table_size	<engine-default>	1073741824

Encrypting database data at rest in Metric Insights

QUESTION

How the Metric Instance instance and the database elements are encrypted at rest? We are looking specifically at the backend data stores themselves.

ANSWER

MySQL does not provide an encryption mechanism by default for the physical db files in /var/lib/mysql. There is an option for it but the decryption key must lie in plaintext on the same server (low level security).

If encryption at rest is needed, you could use RDS encryption settings (unyour AWS cloud) but that's if you've moved the MI database off of EC2 to RDS to begin with.

Please reference to this doc: http://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Overview.Encryption.html

You could also potentially use AES block encryption using the command *mi-crypto-mgr*. However, this requires *ecryptfs* and will use a key on disk to encrypt the db files. This method is untested and most likely does not work on RedHat7/Centos7 because ecryptfs was deprecated.

Logging

How to move the log files from /var/log to a new location

If you need to move the log files from /var/log directory to a new location then follow these steps:

1. Stop syslog.
   Run command in terminal to stop syslog (e.g. '*systemctl stop rsyslog.service*')
2. Edit the following file: /opt/mi/config/external/etc/rsyslog.d/metricinsights.conf
   Change all occurrences of '/var/log/' to your desired destination.
3. Edit the following file: /app/mi/config/external/etc/logrotate.d/metricinsights
   Change all occurrences of '/var/log/' to match the destination set in step #2 above.
4. Run the following command to move your log files and create symlinks to them:
   for x in log debug error; do mv /var/log/mi.$x /my/new/location/mi.$x ln -sf /my/new/location/mi.$x /var/log/mi.$x; done
5. Restart syslog.
   Run command in terminal to stop syslog (e.g. *'systemctl restart rsyslog.service'* or *'service rsyslog restart'*)
6. Update the LOGGING_SYSLOG_FILE_NAME config variable to reflect the new location of your log files.

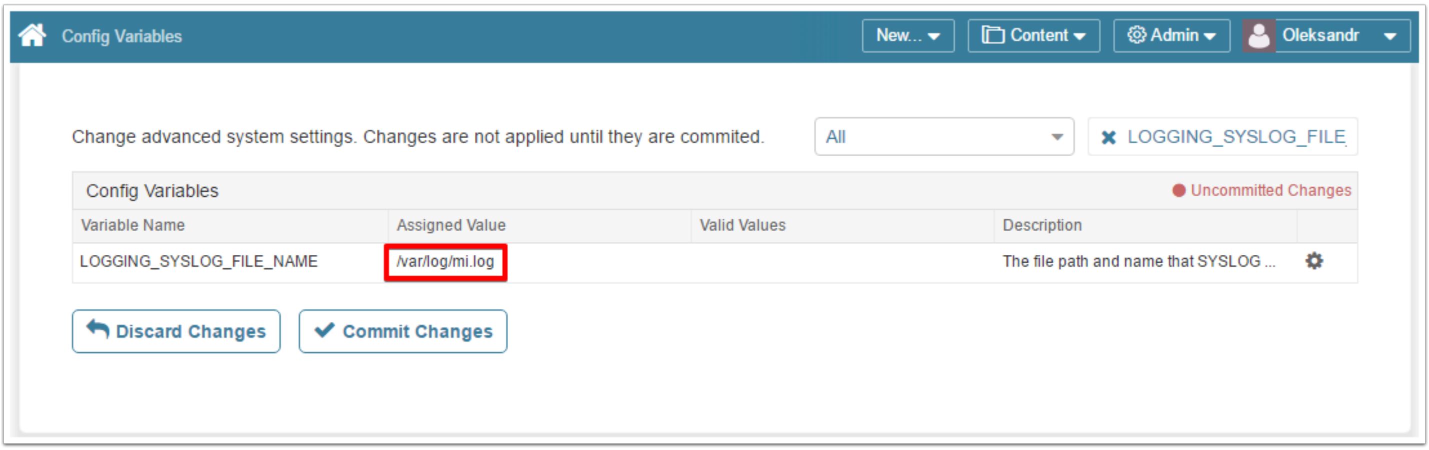

Change advanced system settings. Changes are not applied until they are commited.

All

LOGGING_SYSLOG_FILE

Config Variables — ● Uncommitted Changes

Variable Name	Assigned Value	Valid Values	Description
LOGGING_SYSLOG_FILE_NAME	/var/log/mi.log		The file path and name that SYSLOG ...

SQL Queries for Database statistics

MySQL statements for various tasks

The queries below are used to retrieve data from the Dashboard database. Therefore, Dashboard DB (SQL) must be selected as a Data Source when creating Datasets/Elements from Dataset/Element Editors in Metric Insights.

1. Upgrade all Regular Users to the Power User role

USE CASE:

- The provided query will **upgrade all Regular Users to Power Users in selected Groups**

```
UPDATE user
SET is_power_user_ind='Y'
WHERE user_id IN
(SELECT ugm.user_id
FROM (SELECT * FROM user) AS u
JOIN user_group_member AS ugm ON(u.user_id=ugm.user_id)
WHERE (u.is_power_user_ind='N' AND u.is_administrator_ind='N') AND ugm.user_group_id
IN(<YourTargetedGroupID>));
```

1.1. Get the group ID for the above query

```
SELECT ugm.user_group_id, ugm.user_id, u.username
FROM user_group_member AS ugm
JOIN user AS u ON(u.user_id=ugm.user_id)
WHERE (u.is_power_user_ind='N' AND u.is_administrator_ind='N')
ORDER BY ugm.user_group_id;
```

2. Elements with specific Filter Values

USE CASE:

- The provided query will retrieve **a list of External Reports that are filtered to include ONLY specific values**.

```
SELECT it.element_id, it.name AS filter_name, it.value
FROM (
    SELECT de.element_id, ef.name, efv.value
    FROM dashboard_element de
        INNER JOIN plugin_connection_profile pcp ON pcp.plugin_connection_profile_id =
de.plugin_connection_profile_id
        INNER JOIN external_report_reference err ON err.plugin_connection_profile_id =
pcp.plugin_connection_profile_id AND err.external_report_reference_id = de.
external_report_external_id
        INNER JOIN external_filter ef ON ef.external_report_reference_id = err.
external_report_reference_id
        INNER JOIN external_filter_usage efu ON efu.external_filter_id = ef.
external_filter_id AND efu.item_type = 'element' AND efu.item_id = de.element_id
        INNER JOIN external_filter_value efv ON efv.external_filter_id = ef.
external_filter_id
    WHERE de.type = 'external report'
        AND de.data_fetch_method = 'plugin'
        AND de.external_report_auto_update_image_ind = 'Y'
        AND ef.value_source = 'manual'
        AND efu.value_set = 'all'

    UNION ALL

    SELECT de.element_id, ef.name, efv.value
    FROM dashboard_element de
        INNER JOIN plugin_connection_profile pcp ON pcp.plugin_connection_profile_id =
de.plugin_connection_profile_id
        INNER JOIN external_report_reference err ON err.plugin_connection_profile_id =
pcp.plugin_connection_profile_id AND err.external_report_reference_id = de.
external_report_external_id
        INNER JOIN external_filter ef ON ef.external_report_reference_id = err.
external_report_reference_id
        INNER JOIN external_filter_usage efu ON efu.external_filter_id = ef.
external_filter_id AND efu.item_type = 'element' AND efu.item_id = de.element_id
        INNER JOIN external_filter_usage_value efuv ON efuv.external_filter_usage_id =
efu.external_filter_usage_id
        INNER JOIN external_filter_value efv ON efv.external_filter_value_id = efuv.
external_filter_value_id
    WHERE de.type = 'external report'
        AND de.data_fetch_method = 'plugin'
        AND de.external_report_auto_update_image_ind = 'Y'
        AND ef.value_source = 'manual'
        AND efu.value_set = 'selected'

    UNION ALL

    SELECT de.element_id, ef.name, sv.value_display_name AS value
    FROM dashboard_element de
        INNER JOIN plugin_connection_profile pcp ON pcp.plugin_connection_profile_id =
de.plugin_connection_profile_id
        INNER JOIN external_report_reference err ON err.plugin_connection_profile_id =
pcp.plugin_connection_profile_id AND err.external_report_reference_id = de.
external_report_external_id
        INNER JOIN external_filter ef ON ef.external_report_reference_id = err.
external_report_reference_id
        INNER JOIN external_filter_usage efu ON efu.external_filter_id = ef.
external_filter_id AND efu.item_type = 'element' AND efu.item_id = de.element_id
        INNER JOIN segment_value sv ON sv.segment_id = ef.source_segment_id
    WHERE de.type = 'external report'
        AND de.data_fetch_method = 'plugin'
        AND de.external_report_auto_update_image_ind = 'Y'
        AND ef.value_source = 'segment'
        AND efu.value_set = 'all'

    UNION ALL

    SELECT de.element_id, ef.name, sv.value_display_name AS value
    FROM dashboard_element de
        INNER JOIN plugin_connection_profile pcp ON pcp.plugin_connection_profile_id =
de.plugin_connection_profile_id
        INNER JOIN external_report_reference err ON err.plugin_connection_profile_id =
pcp.plugin_connection_profile_id AND err.external_report_reference_id = de.
external_report_external_id
        INNER JOIN external_filter ef ON ef.external_report_reference_id = err.
external_report_reference_id
        INNER JOIN external_filter_usage efu ON efu.external_filter_id = ef.
external_filter_id AND efu.item_type = 'element' AND efu.item_id = de.element_id
        INNER JOIN external_filter_usage_value efuv ON efuv.external_filter_usage_id =
efu.external_filter_usage_id
        INNER JOIN segment_value sv ON sv.segment_value_id = efuv.segment_value_id
    WHERE de.type = 'external report'
        AND de.data_fetch_method = 'plugin'
        AND de.external_report_auto_update_image_ind = 'Y'
        AND ef.value_source = 'segment'
        AND efu.value_set = 'selected'
) it
WHERE value like '%19%';
```

3. User stats (general)

3.1. Usage activity

3.1.1. Hourly usage activity

USE CASE:

- Retrieve a **count of elements viewed per hour**

```
SELECT DATE_FORMAT(view_time, "%Y-%m-%d %H:00:00") "Date", COUNT(distinct element_id)
"Elements Viewed"
FROM dashboard_element_view_log_detail dash
WHERE view_time > :last_measurement_time
GROUP BY 1
```

3.1.2. Daily usage activity

USE CASE:

- Retrieve a **count of elements viewed per day**

```
SELECT DATE(view_time) "Date", COUNT(distinct element_id) "Elements Viewed"
FROM dashboard_element_view_log_detail dash
WHERE  view_time > :last_measurement_time
GROUP BY 1
```

3.2. Login stats

3.2.1. Last login time per User

USE CASE:

- Retrieve the **last login time for User (user_id)**
- This will show the time when a User last entered their login and password
- Recommendation: use to build a Metric

```
SELECT user_id,
IFNULL(last_login_time, 'not available') AS Last_login_date
FROM user
GROUP BY 1
```

3.2.2. List of Users since last login time

USE CASE:

- Retrieve a **list of Users since last login time**
- Show User information, last login time, and the count of days since that time until now
- Recommendation: use to build a Report

```
SELECT username, first_name, last_name,
IFNULL(last_login_time, 'Not Available') AS Last_login_date,
DATEDIFF(CURDATE(),last_login_time) AS Days_since_last_login
FROM user
```

4. User Engagement (Objects and Elements)

4.1. Homepage elements

4.1.1. Available elements on the Homepage (per User)

USE CASE:

- Retrieve **available elements on the Homepage per User**

```
select d.user_id as user_id, u.username as username, count(element_id) as
number_of_available_elements_on_HP
from user_dashboard_element_instance d
join user u on u.user_id=d.user_id
where in_dashboard_ind_flag = 'Y'
group by 1;
```

4.1.2. Available elements on the Homepage (per User by User types)

USE CASE:

- Retrieve **available elements on the Homepage per User (by User type)**

```
select case is_administrator_ind and is_power_user_ind
when is_administrator_ind = 'Y' then 'Administrator'
when is_power_user_ind ='Y' then 'Power User'
else 'Regular User'
end as usertype, d.user_id as user_id,
u.username,
count(element_id) as number_of_available_elements_on_HP
from user_dashboard_element_instance d
join user u on u.user_id=d.user_id
where in_dashboard_ind_flag = 'Y'
group by 2;
```

4.1.3. All available elements on the Homepage (by count of all Users)

USE CASE:

- Retrieve the **count of all available elements on the Homepage (by count of all Users)**

```
select (select count(element_id) from user_dashboard_element_instance where
in_dashboard_ind_flag = 'Y')/count(user_id)
from user_dashboard_element_instance
```

4.1.4. Homepage search and the number of returned Tiles (per query)

DESCRIPTION

- Table **homepage_search** is used to track searches performed by Users on the Homepage
- Table **homepage_search** contains the following: homepage search id, user id, search text, the time of search and the number of Tiles returned

USE CASE:

- Retrieve data on **Homepage searches and the count of Tiles that were returned by search queries**

```
SELECT *
FROM homepage_search
ORDER by homepage_search_id DESC
limit 100;
```

4.2. Favorites

4.2.1. List of Favorite elements

USE CASE:

- Retrieve a **list of Favorite elements**
- Fetch a list of Favorites for each user
- Recommendation: build a Report

```
SELECT u.user_id AS User_ID, u.username AS User_Name,  fdei.element_id AS Element_ID,
de.name AS Element_Name,  f.display_name AS Fovirites_Name
FROM favorite_dashboard_element_info AS fdei
JOIN dashboard_element AS de ON(de.element_id=fdei.element_id)
JOIN user AS u ON(u.user_id=fdei.user_id)
JOIN favorite AS f ON(f.favorite_id=fdei.favorite_id)
GROUP BY fdei. favorite_dashboard_element_id
ORDER BY u.username
```

4.3. Viewing and usage statistics

4.3.1. Average number of viewed Tiles (per period)

USE CASE:

- Retrieve the **average number of Tiles per period**

```
SELECT user_id, ifnull(count(DISTINCT
element_id)/DATEDIFF(MAX(view_time),MIN(view_time)), 'N/A') as
AVG_number_of_tiles_per_period
FROM dashboard_element_view_log_detail
GROUP BY 1;
```

4.3.2. Average number of viewed Tiles (per month)

USE CASE:

- Retrieve the **average number of Tiles per month**

```
SELECT user_id, cnt/vt as a
FROM (
SELECT user_id, CONCAT(YEAR(view_time),'-',MONTH(view_time)) AS vt,  count(DISTINCT
element_id) AS cnt FROM dashboard_element_view_log_detail
GROUP BY 1
) AS t
GROUP BY 1;
```

4.3.3. Most viewed elements per month

USE CASE:

- Retrieve **most viewed elements per month (Top 10)**
- Displays which elements were viewed the most within the last 30 days
- Recommendation: build a Report

```
SELECT devld.element_id AS "Element ID",
de.name AS "Element Name",
count(distinct devld.user_id) AS "Number of Distinct Portal Users",
count(devld.user_id) AS "Total Number of Views"
FROM dashboard_element_view_log_detail AS devld
JOIN dashboard_element de ON (devld.element_id = de.element_id)
WHERE DATE(devld.view_time)>(curdate()- interval 1 month)
GROUP BY 1
ORDER BY 4 desc Limit 10
```

4.3.4. List of last viewed elements

USE CASE:

- Retrieve the **list of last viewed elements**
- This Report will fetch data about the last viewed elements, including User information and viewing time

```
SELECT u.user_id AS User_id, u.username AS User_name, de.name AS Viewed_element_name,
dc.category AS Element_category_name,
IFNULL(devl.last_view_time, 'Not Available') AS Last_viewed_on
FROM dashboard_element_view_log AS devl
JOIN user AS u ON(u.user_id=devl.user_id)
JOIN dashboard_element AS de ON(de.element_id=devl.element_id)
JOIN dashboard_category AS dc ON(dc.category_id=de.category_id)
ORDER BY last_view_time DESC
```

4.3.5. Dimension and time period changes for Metrics (at last View)

USE CASE:

- Retrieve **a Dimension and time period changes**
- This Report shows what Dimension and time period were selected by User during the last Metric View

```
SELECT u.username AS User, u.user_id AS User_ID, de.name AS Metric_Name, de.element_id
AS Metric_ID, sv.value_display_name AS Seen_Dimension_Values, uco.last_updated_time AS
View_time, REPLACE(SUBSTRING_INDEX(SUBSTRING_INDEX(uco.
overlay_state,'"interval_unit":"',-1),'}',1),'","interval_value":',' ') AS
Last_Viewed_Time_interval
FROM user_chart_overlay AS uco
JOIN user AS u ON(u.user_id=uco.user_id)
JOIN dashboard_element AS de ON(uco.element_id=de.element_id)
JOIN segment_value AS sv ON(uco.segment_value_id=sv.segment_value_id)
ORDER BY u.username, uco.last_updated_time DESC, sv.value_display_name
```

4.3.6. Default View type for Metrics

USE CASE:

- Retrieve the **default View type for Metrics**
- This Report will show what type of Metric View was selected (Standart, Stoplight, Target, Projection, etc.)

```
SELECT de.element_id AS Element_id, de.name AS Element_name, devld.view_time AS
Viewed_time, u.username AS User
, SUBSTRING(uco.overlay_state,10,POSITION('"' IN REPLACE(uco.
overlay_state,'{"view":"',''))-1) AS View_name
 FROM dashboard_element AS de
 JOIN dashboard_element_view_log_detail AS devld ON(devld.element_id=de.element_id)
 JOIN user_chart_overlay AS uco ON(de.element_id=uco.element_id)
 JOIN user AS u ON(u.user_id=devld.user_id)
 WHERE uco.overlay_state LIKE '{"view":"%'
 GROUP BY u.user_id
```

4.3.7. Most used Datasets

USE CASE:

- Retrieve **most used Datasets (Top 10)**
- This Report will contain the most used Datasets and the count of elements sourced from them

```
SELECT de.dataset_id As Dataset_ID, d.name AS Dataset_Name, count(distinct de.
element_id) AS Number_of_element
FROM dashboard_element AS de
JOIN dataset AS d ON(de.dataset_id=d.dataset_id)
GROUP BY de.dataset_id
ORDER BY 3 Desc LIMIT 10
```

5. User Engagement (Notifications)

5.1. Favorites and Shared Folders

5.1.1. All elements in favorite Folders for a Digest with enabled Notifications

DESCRIPTION:

- Table **favorite_dashboard_element_info** contains all elements in Favorite Folders including Shared Folders.

USE CASE:

- Retrieve **all elements in favorite Folders for a Digest with enabled Notifications**

```
SELECT DISTINCT fdei.user_id, fdei.element_id
FROM favorite_dashboard_element_info AS fdei
JOIN favorite AS f ON (f.favorite_id=fdei.favorite_id)
JOIN user_preference AS up ON (up.user_id=fdei.user_id)
WHERE f.include_in_favorites_digest_ind='Y'
AND up.email_notification_enabled_ind = 'Y';
```

5.2. Alerts

5.2.1. All User Alert Subscriptions for elements with enabled Notifications

DESCRIPTION:

- Table **user_alert_rule** is used for storing information about user alert subscriptions since 4.0.
- Table **alert_rule_element_info** contain elements from alert rule scope with visualizations.

USE CASE:

- Retrieve **all User Alert Subscriptions for elements with enabled Notifications**

```
SELECT DISTINCT uar.user_id, arei.visualization_element_id
FROM user_alert_rule AS uar
JOIN alert_rule_element_info AS arei ON (uar.alert_rule_id=arei.alert_rule_id AND uar.
element_id=arei.element_id)
JOIN user_preference AS up ON (up.user_id=uar.user_id)
WHERE uar.enabled_ind='Y'
AND up.email_notification_enabled_ind = 'Y';
```

5.3. Bursts

5.3.1. Burst Recipients with enabled Notifications

DESCRIPTION:

- Table **user_alert_rule** is used for storing information about user alert subscriptions since 4.0.
- Table **alert_rule_element_info** contains elements from alert rule scope with visualizations.

USE CASE:

- Retrieve **Burst Recipients with enabled Notifications**

```
SELECT dlarv.notification_schedule_distribution_id as Burst_id, nsd.name AS
Burst_name,  nsd.enabled_ind AS Burst_Enabled, u.user_id AS Subscribed_User_id, u.
username AS Subscribed_user_name,  u.email AS User_mail
FROM distribution_list_all_recipients_view AS dlarv
JOIN notification_schedule_distribution AS nsd ON dlarv.
notification_schedule_distribution_id=nsd.notification_schedule_distribution_id
JOIN user AS u ON u.user_id=dlarv.user_id
WHERE u.enabled_ind='Y'
ORDER BY dlarv.notification_schedule_distribution_id, nsd.name;
```

⚠ The contents of **"distribution_list_all_recipients_view"** (from the above query) are provided below:

```
CREATE ALGORITHM=UNDEFINED VIEW `distribution_list_all_recipients_view` AS
    SELECT
        `nsdur`.`user_id` AS `user_id`,
        `nsdur`.`notification_schedule_distribution_id` AS
`notification_schedule_distribution_id`
    FROM `notification_schedule_distribution` as `nsd`
        JOIN `notification_schedule_distribution_user_recipient` as `nsdur` ON `nsdur`.
`notification_schedule_distribution_id` = `nsd`.`notification_schedule_distribution_id`
    WHERE `nsd`.`recipient_scope` = 'selected'
    UNION
    SELECT
        `f`.`user_id` AS `user_id`,
        `nsd`.`notification_schedule_distribution_id` AS
`notification_schedule_distribution_id`
    FROM `notification_schedule_distribution` as `nsd`
        JOIN `notification_schedule_distribution_item` as `nsdi` ON `nsdi`.
`notification_schedule_distribution_id` = `nsd`.`notification_schedule_distribution_id`
AND `nsd`.`content_type` = 'favorites'
        JOIN `favorite` as `f` ON (`f`.`source_shared_favorite_id` = `nsdi`.
`favorite_id` OR `f`.`copied_from_favorite_id` = `nsdi`.`favorite_id`)
    WHERE `nsd`.`recipient_scope` = 'all'
    UNION
    SELECT
        `nsd`.`created_by_user_id` AS `user_id`,
        `nsd`.`notification_schedule_distribution_id` AS
`notification_schedule_distribution_id`
    FROM `notification_schedule_distribution` as `nsd`
    WHERE `nsd`.`send_owner_distribution_ind` = 'Y';
```

6. User Engagement (Notes, Annotations, Events)

6.1. List of User Comments

USE CASE:

- Retrieve the **list of User comments**

- This query will return all comments for Notes, Annotations, and Events

```
SELECT uc.element_id As ElementID, de.name AS Element_Name, de.type AS Element_Type, sv.
value_display_name As Dimension_value, uc.scope AS First_comment_type, un.text AS
First_comment, u_un.username AS First_comment_user, un.created_time As
First_comment_time, uc.text AS Second_comment, u_uc.username AS Second_comment_user, uc.
last_updated_time AS Second_comment_time
FROM user_comment AS uc
JOIN dashboard_element AS de ON(de.element_id=uc.element_id)
LEFT JOIN segment_value AS sv ON(sv.segment_value_id=uc.segment_value_id)
JOIN user_note AS un ON(uc.user_note_id=un.user_note_id)
JOIN user AS u_un ON(un.user_id=u_un.user_id)
JOIN user AS u_uc ON(uc.user_id=u_uc.user_id)
WHERE uc.scope='note'

UNION

SELECT uc.element_id, de.name, de.type, sv.value_display_name,  uc.scope, ua.
annotation_text, u_ua.username, ua.annotation_time, uc.text, u_uc.username, uc.
last_updated_time
FROM user_comment AS uc
JOIN dashboard_element AS de ON(de.element_id=uc.element_id)
LEFT JOIN segment_value AS sv ON(sv.segment_value_id=uc.segment_value_id)
JOIN user_annotation AS ua ON(uc.user_annotation_id=ua.user_annotation_id)
JOIN user AS u_ua ON(ua.user_id=u_ua.user_id)
JOIN user AS u_uc ON(uc.user_id=u_uc.user_id)
WHERE uc.scope='annotation'

UNION

SELECT uc.element_id, de.name, de.type, sv.value_display_name, uc.scope, ne.name, ne.
last_updated_by, ne.last_notable_event_activity_time, uc.text, u_uc.username, uc.
last_updated_time
FROM user_comment AS uc
JOIN dashboard_element AS de ON(de.element_id=uc.element_id)
LEFT JOIN segment_value AS sv ON(sv.segment_value_id=uc.segment_value_id)
JOIN notable_event AS ne ON(uc.notable_event_id=ne.notable_event_id)
JOIN user AS u_uc ON(uc.user_id=u_uc.user_id)
WHERE uc.scope='event'

ORDER BY 2
```

7. System performance

7.1. System load

7.1.1. Count of parallel processes per day for the last month

USE CASE:

- Retrieve the **daily count (maximum and average) of parallel processes for the last month**

```
SELECT DATE(collection_time), MAX(total_process_count), AVG(total_process_count)
FROM mysql_processlist_log
WHERE collection_time>NOW() - INTERVAL 30 DAY
GROUP BY 1;
```

7.1.2. System load average per hour

USE CASE:

- Retrieve the **system load average value per hour**

```
SELECT DATE_FORMAT(collection_time, "%Y-%m-%d %H:00:00"),
AVG(total_process_count)
FROM mysql_processlist_log
WHERE collection_time>NOW() - INTERVAL 30 DAY
GROUP BY 1
```

7.1.3. Data collection issues

7.1.4. List of errors upon data collection

USE CASE:

- Retrieve the **list of errors upon data collection**
- This basic query will return all elements with errors

```
SELECT *
FROM (
    SELECT 'trigger' AS caller,  element_id, segment_value_id, start_time, success_ind,
error_message
    FROM update_trigger_event_run_log_detail
    WHERE success_ind = 'N'

    UNION ALL

    SELECT 'editor' AS caller,  element_id, segment_value_id, start_time, success_ind,
error_message
    FROM editor_data_collection_detail
    WHERE success_ind = 'N'
) it
WHERE error_message != 'No rows are returned'
```

7.1.5. Elements with data collection exceeding 60 minutes

USE CASE:

- This basic query will return all elements with long-running data collection (>= 60 min)

```
SELECT *
FROM (
    SELECT 'trigger' AS caller,  element_id, segment_value_id, TIMESTAMPDIFF(MINUTE,
start_time, finish_time) AS _mins
    FROM update_trigger_event_run_log_detail
    WHERE success_ind = 'Y'
        AND TIMESTAMPDIFF(MINUTE, start_time, finish_time) >= 60
    UNION ALL
    SELECT 'editor' AS caller,  element_id, segment_value_id, TIMESTAMPDIFF(MINUTE,
start_time, finish_time) AS _mins
    FROM editor_data_collection_detail
    WHERE success_ind = 'Y'
        AND TIMESTAMPDIFF(MINUTE, start_time, finish_time) >= 60
) it
ORDER BY _mins DESC
```

7.1.6. List of overdue Trigger runs

USE CASE:

- Retrieve the **list of Triggers with overdue runs**
- This basic query will return a list of overdue Trigger ids, the start time and reasons for overdue runs

```
SELECT ute.update_trigger_event_id _id, ute.name, rl.run_id, rl.run_start_time, IF (rl.
event_aborted_ind = 'Y', 'Aborted', 'Timed out') _reason
FROM update_trigger_event ute
INNER JOIN update_trigger_event_run_log rl ON rl.update_trigger_event_id = ute.
update_trigger_event_id
WHERE rl.run_timed_out_ind = 'Y' OR rl.event_aborted_ind = 'Y'
ORDER BY rl.run_start_time DESC
```

Network

SSH Port Tunneling with Putty

This instruction will be useful if you want to reach instances that are behind firewall outside of network by using the Putty SSH Client and SSH Port Tunneling.

Here is an instruction on how you will be able to do it: https://howto.ccs.neu.edu/howto/windows/ssh-port-tunneling-with-putty/

For example to access the Tableau server at: https://10.80.7.94 you need to:

1) In Putty

Under Add new forwarded port:, enter the following information:

Source port: [port on local machine]   (Source port: 8443)

Destination: [hostname of ccis machine]:[port on ccis machine]   (Destination: 10.80.7.94:443)

Click Add button.

2) Under Add new forwarded port:, enter the following information:Then connect to MI via this Putty.

3) Then access Tableau via url in web browser: https://localhost:8443

Also this technique is useful to reach MI instances that are behind a firewall - https://subdomain.metricinsights.com, Destination will look like subdomain.metricinsights.com:443.

How to set up a Static IP for a Virtual Machine (CentOS)

The following lists steps to set up a Static IP for a Metric Insights Virtual Machine (CentOS). Note, you will have to get the IP from the customer's network administrator.

ISSUE

We've deployed the Metric Insights OVA in our virtual player, but company policy dictates that we not use a Bridged connection. We must use NAT instead and our network administrator has provided us with an IP for the VM. How do we set this up?

SOLUTION

To set up a Static IP, please follow the steps below in the VM. Ensure you have the static IP, gateway, and netmask to properly configure the VM:

1.) From the command line, make a backup of the network card file **ifcg-eth0** or **ifcg-enp0s3** in /etc/sysconfig/network-scripts/ then proceed with making changes to either file (the file repesent depends on the CentOS release):

# cp /etc/sysconfig/network-scripts/ifcfg-eth0 /etc/sysconfig/network-scripts/ifcfg-eth0.bak

# vim /etc/sysconfig/network-scripts/ifcfg-eth0

2.) Make the following changes to ifcg-eth0 in the vim editor:

```
DEVICE="eth0"
NM_CONTROLLED="yes"
ONBOOT=yes
HWADDR=20:89:84:c8:12:8a
TYPE=Ethernet
BOOTPROTO=static
NAME="System eth0"
UUID=5fb06bd0-0bb0-7ffb-45f1-d6edd65f3e03
IPADDR=192.168.90.1
NETMASK=255.255.255.0
GATEWAY=192.168.90.254
```

Only the parameters BOOTPROTO, IPADDR, NETMASK, and GATEWAY have to be adjusted (or added if missing). You do not have to edit the other lines.

3.) Next, edit /etc/sysconfig/**network**:

# vim /etc/sysconfig/network

Add the following entries:

```
NETWORKING=yes
HOSTNAME=company_name.metricinsights.com
GATEWAY=192.168.90.254
```

4.) After configuring the IP, we now need to configure DNS. This can be done in /etc/**resolv.conf**:

# vim /etc/resolv.conf

Add or edit the following **nameserver** entries:

```
nameserver        8.8.8.8
nameserver        8.8.4.4
```

You can add two or more nameserver entries. The VM will try the second one in case the first nameserver is unreachable.

5.) Add the IP and hostname in /etc/**hosts**:

# vim /etc/hosts

```
127.0.0.1   localhost localhost.localdomain localhost4 localhost4.localdomain4
192.168.90.1   company_name.metricinsights.com
```

Note, the new hostname will only be applied after a system reboot. **6.)** To apply the network changes, enter the following command:

# services network restart

7.) After the network service restart you can confirm the static IP has been applied to the VM by running:

# ifconfig

You should see an output like this:

```
root@company_name.metricinsights.com:~# ifconfig
eth0      Link encap:Ethernet  HWaddr 20:89:84:c8:12:8a
          inet addr:192.168.90.1  Bcast:192.168.0.255  Mask:255.255.255.0
          inet6 addr: 2001:db8::c0ca:1eaf/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:200197 errors:0 dropped:67 overruns:0 frame:0
          TX packets:69689 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:64103748 (64.1 MB)  TX bytes:14106191 (14.1 MB)


lo        Link encap:Local Loopback
          inet addr:127.0.0.1  Mask:255.0.0.0
          inet6 addr: ::1/128 Scope:Host
          UP LOOPBACK RUNNING  MTU:65536  Metric:1
          RX packets:10365 errors:0 dropped:0 overruns:0 frame:0
          TX packets:10365 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0
          RX bytes:875114 (875.1 KB)  TX bytes:875114 (875.1 KB)
```

You should be able to reach the Metric Insights UI now by entering the IP in your web browser. Also, you can check for Internet access from the VM by pinging an external website like google.com. Note, if the network administrator has restricted internet access from the NAT network, you will be unable to ping the internet.

"Device eth0 does not seem to be present, delaying initialization" ERROR on VM reboot

This error indicates that the VM can't identify the network adapter.

ISSUE

After rebooting the Metric Insights VM, I am getting a ***"Device eth0 does not seem to be present, delaying initialization"*** error. The VM then defaults to 127.0.0.1 (localhost). The network adapter itself is functioning correctly (other servers are connected to it). I can also confirm that the **ifcfg-eth0** file is missing in /etc/sysconfig/network-scripts. What is this happening?

RESOLUTION

The reboot probably left the **ifcfg-eth0** file in a state of limbo, possibly because a new MAC address was assigned to the network adapter. For example, if you try to recreate the file, you'll get a prompt that the file exists even though it is not visible.

To fix this, delete the networking interface rules from the command line so that they can be regenerated after another reboot:

```
Delete networking interface rules:
# rm /etc/udev/rules.d/70-persistent-net.rules

Reboot the VM:
# reboot
```

On reboot, you should see a new MAC address being generated and an IP being assigned to the VM. Confirm the IP by hitting it from a browser; it should take you to the Metric Insights login page.