



EMC® Watch4net®

Version 6.3u2

Topology CLI Administration Guide

P/N: pending

Rev : 40032

EMC²

Copyright ©2010 - 2014 EMC Corporation. All rights reserved. Published in the USA.

Published April, 2014

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

The information in this publication is provided as is. EMC Corporation makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose. Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

EMC², EMC, and the EMC logo are registered trademarks or trademarks of EMC Corporation in the United States and other countries. All other trademarks used herein are the property of their respective owners.

For the most up-to-date regulatory document for your product line, go to EMC Online Support (<https://support.emc.com>).

Contents

Preface	iii
What is the Topology?	iii
Intended Audience	iii
Document Organization	iv
Where to get help	v
1 System Prerequisites	1
2 Tool Overview	3
1 Directory Layout	3
2 Quick Usage Guide	3
3 Configuration Details	5
1 Database Connection Definition	5
2 Topology Group Cleanup	6
3 Topology Refresh	6
3.1 APG Timeseries Import	7
3.2 CSV Data Import	7
4 F.A.Q.	11
1 How do I know the process completed successfully?	11
2 How long will the process take?	11
3 The refresh process is taking a very long time. Why?	11
4 Can I interrupt the refresh process?	12
5 When should I schedule the refresh of the topology?	12
6 What is the import order between CSV and APG data?	12
7 What is the global import order?	12
8 How can I check that the new instances / relations are available?	12

Preface

What is the Topology?

This guide describes the common usage of the topology refresh CLI, which includes the data files format and the CLI configuration. It also explains in this section what the topology is and what it does for the APG users.

The topology is the definition of *Types*, *Instances* for each one, *Relations* between the types and *Topology Links* which are the representation of the relations for the topology instances.

Let's take the example of a router named RTR1. This router has a Processor named PR1. In this simple case, the types are Router and Processor. The instances are RTR1 for the type Router and PR1 for the type Processor.

If we define a relation ComposedOf that links a Router to a Processor, then in the topology, there will be the link RTR1 -ComposedOf-> PR1.

Those relations can then be used in the frontend to easily create or modify reports.

Intended Audience

This guide is meant to be read by any person who will manage and/or schedule the refresh of the topology data through the command line interface.

This process should be scheduled periodically by the server admin so that values and relations are kept up to date. Scheduling this execution through the operating system tools is the admin's responsibility.

Document Organization

1. System Prerequisites	Details the pre-requisites for the topology refresh tool to run as expected.
2. Tool Overview	Gives a general overview of the topology refresh tool layout and files.
3. Configuration Details	Gives a detailed explanation of the configuration file and its options.
4. F.A.Q	Contains answers about the behavior of the refresh process.

Where to get help

EMC support, product, and licensing information can be obtained as follows:

PRODUCT INFORMATION - For documentation, release notes, software updates, or information about EMC products, go to EMC Online Support at:

<https://support.emc.com>

TECHNICAL SUPPORT - Go to EMC Online Support and click Service Center. You will see several options for contacting EMC Technical Support. Note that to open a service request, you must have a valid support agreement. Contact your EMC sales representative for details about obtaining a valid support agreement or with questions about your account.

System Prerequisites

This module can be installed and used on any platform supported by APG. It supports MySQL 5+ and Oracle 11g databases.

It requires access to at least one database that will contain the resulting topology. The database user will require rights to select and insert data, as well as create, alter, and drop tables. In order to use APG timeseries data, the tool will also need to access the corresponding databases, but it can also run using CSV data files only.

- APG's module-manager for the installation
- a proper configuration
- read access to CSV data files
- read access to APG timeseries databases with a generated **data_property_flat**
- full access to a topology database (tables will be automatically created)

2

Tool Overview

1 Directory Layout

Once the installation is complete, the installation directory will contain the following entries.

- `bin` which contains the binaries used to launch the tool
- `conf` which contains the requested configuration file
- `data` which may contain CSV files used to refresh the topology
- `doc` which contains this documentation file
- `lib` which contains the Java binaries

2 Quick Usage Guide

Using this tool is a very simple three steps process.

First, if it does not already exist, create a database named `topology`. For instance, on the default MySQL provided with APG, login as root and type the following commands (of course you may want to replace the password by another one).

```
mysql> create database topology;  
mysql> grant all privileges on topology.* to 'apg'@'localhost' identified by 'watch4net';
```

Then edit the configuration file `conf/topology-refresh-config.xml`. In this file, the default APG timeseries database is already configured as well as the `topology` database mentioned above. Check that both usernames and passwords are correct and update them if necessary.

Finally, launch the tool using the `bin/refresh-topology` executable that corresponds to the server operating system.

If you need more information about the configuration and what each of its elements does, then proceed to the next section.

Note: In order to use a specific configuration file instead of the default one, simply provide its absolute path as an argument when launching the executable.

3

Configuration Details

There are three main sections in the single configuration file required by the tool. The first deals with the database connection definition, the second one is about cleanup tasks and the last one configures the way each topology group will be refreshed.

1 Database Connection Definition

A database connection can be defined in a single `database` element.

```
<database id="apg" type="mysql" url="jdbc:mysql://localhost:53306/apg?autoReconnect=true"
username="apg" password="watch4net" />
<database id="topo" type="oracle" url="jdbc:oracle:thin:@localhost:1521:APG" username="apg"
password="watch4net" />
```

The attribute `id` is used to reference this connection in other elements of the configuration. As such, it has to be unique.

The attribute `type` defines the type of database that is targeted. It can be either a `mysql` one either an `oracle` one. It is used to determine how to connect to the defined database.

The `url` is the connection string which defines how to reach the database. It depends on the selected database type and is the same that can be found in *APG's* portal configuration.

The last attributes are the username and password required for the authentication.

2 Topology Group Cleanup

In order to remove unused groups, one can define a `cleanup` element per topology group.

```
<cleanup database="topo" group="Group_Name" />
```

The `database` attribute contains the ID which references one of the databases defined in the above section.

The `group` attribute contains the name of the group that will be removed from this database. Note that if the group does not exist in the specified database, then the cleanup will have no effect.

3 Topology Refresh

This `refresh` element is a bit more complex as it contains sub-elements which specify the data sources to use for the refresh process.

```
<refresh database="topo" group="Core_Networking" remove-previous-data="false">
  <apg-import>
    <database ref="apg-db-1" />
    <database ref="apg-db-2" />
  </apg-import>
  <csv-import>
    <types>data/types.csv</types>
    <relations>data/relations.csv</relations>
  </csv-import>
</refresh>
```

The `database` and `group` attributes are used to target a specific group in a database. The first contains a reference to the defined database while the other one contains the name of the targeted

group.

3.1 APG Timeseries Import

If some APG databases are used as data sources in this refresh process, they have to be set in the `apg-import` element. They consist of database elements that reference a defined database access and accept some optional attributes.

```
<database ref="apg-db-1" />
<database ref="apg-db-2" filter="customer='CustomerA'" types="true" instances="true" topology="true" />
```

In the example above, the tool will recover all instances and topology links for `apg-db-1` without any filter, which are the default settings. It will do almost the same for `apg-db-2`, except that types will also be recovered and all values will be filtered using the specified expression. The filter expression has the same syntax as the ones that can be created in *APG's* frontend.

Note: Although types can be recovered from a timeseries database, relations cannot and have to be specified in a CSV data file. It is also better to have the type definition in a CSV data file too.

3.2 CSV Data Import

The CSV files have to be defined in another element called `csv-import`. There is one element per type of data, and each of them requires a path to the data file which could be either absolute either relative to the tool installation directory.

```
<types>data/types.csv</types>
<instances>data/instances.csv</instances>
<relations>data/relations.csv</relations>
<topology>data/topology.csv</topology>
```

Types Format

Each CSV type entry has to be in the following format:

```
"Type_Name", "Type_Description", "Type_Filter", "Type_Expansion", "Is_Entry_Point?"
```

All attributes except the description and entry point are mandatory. Not providing the entry point flag will default with FALSE. Although not necessary if not required, they should always be escaped. The name has to be unique. The filter restricts the property values this type deals with while the expansion gives the distinct values for it. The entry point information is used to know whether the type should be used at the beginning of a topology drilldown or not.

For instance, let's see what the definition of the `Router` and `Interface` types would look like:

```
"Router", , "devtype='Router'", "device", "TRUE"  
"Interface ", , "parttype=' Interface '", "device, part", "FALSE"
```

Please note that when omitting the description, the field can be left empty but the commas have to remain here.

Instances Format

Each CSV instance entry has to be in the following format:

```
"Instances_Type_Name", "Instance_Name_1", "Instance_Name_2", . . .
```

There has to be exactly as many names as the there are elements in the type expansion, and they have to be in the exact same order.

Let's continue with the previous example and create an instance of `Router` and two instances of `Interface`:

```
"Router", "RTR1"  
"Interface ", "RTR1", "FastEthernet0"  
"Interface ", "RTR1", "FastEthernet1"
```

Although the names give a hint about the fact the interfaces belong to the router, there is still no relation that defines it.

Relations Format

Each CSV relation entry has to be in the following format:

```
"Relation_Name","Relation_Description","Source_Type_Name","Destination_Type_Name" [, "Reverse_Relation_Name"]
```

The relation name, source and destination are mandatory. The name has to be unique. The description can be left empty. The last field is optional and can be completely omitted. It simply gives information about the opposite relation that links the destination to the source.

In order to create a composition relation, we would have:

```
"ComposedOf",,"Router","Interface","PartOf"
"PartOf",,"Interface","Router","ComposedOf"
"ComposedOf",,"Router","Processor"
```

Topology Links

When types, instances and relations have been defined, links between the instances can be created. Each CSV topology link entry has to be in the following format:

```
"SrcType:SrcInstName1,SrcInstName2,...","RelationName","DstTypeName:DstInstName1,DstInstName2,..."
```

Given that format, it is possible to guess instances and relations from a topology data file.

As a consequence, two optional attributes are available on a `topology` element of the configuration:

- `guessRelations` which can be set to `true` in order to create relations from this file too
- `guessInstances` which can be set to `true` in order to create instances from this file too

Now, let's define links between the instances used in the above examples:

```
"Router:RTR1","ComposedOf","Interface:RTR1,FastEthernet0"
"Router:RTR1","ComposedOf","Interface:RTR1,FastEthernet1"
"Interface :RTR1,FastEthernet0","PartOf","Router:RTR1"
"Interface :RTR1,FastEthernet1","PartOf","Router:RTR1"
```


4

F.A.Q.

1 How do I know the process completed successfully?

Every step is written in the standard output. It also gives information about the number of types, instances, relations and topology links that have been created or updated. Any error will appear in the standard error, with detailed message and possibly a stacktrace too.

2 How long will the process take?

In most cases, less than a minute. But of course, it also depends on the server load and the data to process. On a standard APG installation, it should take around 20 seconds.

3 The refresh process is taking a very long time. Why?

Check that the `data_property_flat` table exists. If it is not available, the topology refresh will not be efficient at all. This accelerator table is mandatory!

4 Can I interrupt the refresh process?

Yes. Everything will still be working, except that some topology data might not be available.

5 When should I schedule the refresh of the topology?

During an idle time, and if APG import is used, then a few minutes after the (re)generation of the `data_property_flat` table so that it could benefit its new values.

6 What is the import order between CSV and APG data?

It's always APG first, and then CSV so that CSV data can override data coming from APG's databases.

7 What is the global import order?

1. cleanup
2. APG types
3. CSV types
4. APG instances
5. CSV instances
6. CSV relations
7. APG topology
8. CSV topology

8 How can I check that the new instances / relations are available?

It is possible through the Report Wizard of the frontend, in the topology selection. The instances will appear when the right group and type are selected.