




EMC[®] Watch4net[®]

Version 6.4u2

Topology Report Administration Guide

P/N: pending

Rev : 50603

Copyright ©2011 - 2014 EMC Corporation. All rights reserved. Published in the USA.

Published October, 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
Purpose	iii
Intended Audience	iii
Prerequisites	iii
Document Organization	iii
Where to get help	v
1 Product Overview	1
2 Mode	3
1 Simple mode	3
2 Advanced mode	4
3 Drilldown mode	6
4 Topology Service mode	7
5 Custom mode	7
6 Rendering	8
3 Configuration	9
1 Template	9
2 Options	10
2.1 Viewport	10
2.2 Layout	11
3 Highlight	11
3.1 Highlight Attributes	11
4 Bullet Indicator	12
4.1 Bullet Indicator Attributes	12
4.2 Bullet Indicator Ranges	14
5 Image Indicator	15
5.1 Image Indicator Attributes	16
5.2 Image Indicator Ranges	17
6 Sparkline Indicator	18
6.1 Sparkline Indicator Attributes	19
6.2 Sparkline Indicator Ranges	20
7 Tooltips	22
7.1 Tooltip Attributes	22
7.2 Tooltip Contents	22
8 Node	23
8.1 Node Attributes	23
8.2 Nested Elements	24
9 Group	25
10 Edge-Type	26
10.1 Edge-Type Attributes	26
10.2 Edge-Type Ranges	28

11Edge	29
11.1Edge Attributes	30
11.2Nested Elements	31
12Background	32
13Report	32
14Placeholders	33

Preface

Purpose

The *Topology Report* is a category of report designed to show different elements, for example devices, as nodes linked together with arrows (referred to generically as edges). Nodes can have indicators which typically display different metrics about them; edges can also display metrics, this time, typically, about the nodes they connect.

Intended Audience

This guide is intended for advanced *APG* Administrators who will design new *Topology Reports*.

Prerequisites

The *Topology Report* can only be used within the *APG Frontend*, version 5.2 or above.

Document Organization

1. Product Overview	Overview of the report.
2. Mode	The <i>Topology Report</i> modes.
3. Configuration	How to configure the <i>Topology Report</i> .

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.

Product Overview

The topology report allows you to build and display transactions graphically in a report. It can also be used to represent a specific section of a network and to display various data between the different elements comprised within this section. This report is really useful to see graphically how a part of your network is performing or to represent easily a particular transaction. In short, this report allows you to represent different pieces of equipment and services and also the different metrics and KPI's that link them together.

There are two basic components that are used to define the report. First, the **nodes** which represent the different elements of the transaction or the various elements of the topology representing the network section you want to describe. The second type of elements are the **edges** which are used to define the relationships between the nodes. A node can either have none, one or multiple edges linking to other nodes (n-to-n relationship). Edges can be capped with an arrow, pointing from one node to another, representing the flow.

2

Mode

The mode of the *Topology Report* will determine the source of the report. It can be fully customizable, based on the APG topology or based on the subreports of the current report.

To create a new *Topology Report* in the APG Frontend, select **Topology** from the *Report Type* drop-down list while in *Edit Mode*. After, in the *Report details* tab, you must select the appropriate mode.

The screenshot shows the 'Report Details: Topology' tab. At the top, there is a note: 'The simple and advanced mode are based on the APG topology. The drilldown mode is based on the subreports, while the custom mode is completely based on an XML file.' Below this, under the 'Mode' section, there are four radio button options: 'simple', 'advanced' (which is selected), 'drilldown', and 'custom'.

1 Simple mode

The simple mode is based on the APG topology. To use it, at least one topology must be configured in APG. (See *APG Topology CLI documentation for more details on how to create a topology*)

This mode will only display instances of classes from the selected topology which match the current report filter.

For example, if you select the *Router* class which represent the router R1, R2 and R3 and your report filter is *devtype=='Router' & device=='R1'* then only the router R1 will be display on the report.

Note that report filter must englobed the filter of the APG topology to be able to retrieved instances of classes. For example, if the filter *devtype=='Router' & ip='10.128.%'* is selected, then only routers which the IP starts with 10.128 will be displayed. But if the filter *devtype=='Router' & ifaceip='10.128.%'* is selected, then none of the routers will be displayed. This is because

APG topology filter for the routers is `devtype=='Router' & !part` and both filters must matches. So since the `ifaceip` is the interface ip, which contains a property `part` no router instances can be found.

The screenshot shows the 'Report Details: Topology' configuration window. At the top, there are four tabs: 'Filtering & Expansion', 'Report Configuration', 'Report Details: Topology' (which is active), and 'Formula'. Below the tabs, a note states: 'The simple and advanced mode are based on the APG topology. The drilldown mode is based on the subreports, while the custom mode is completely based on an XML file.' Under the 'Mode' section, four radio buttons are present: 'simple' (selected), 'advanced', 'drilldown', and 'custom'. Below this, a 'Simple mode' section is expanded. It contains a 'Topology' dropdown menu set to 'Core Networking'. Below the topology menu are two lists: 'Classes' and 'Relations'. The 'Classes' list has a search bar and five items: 'FileSystem', 'HTTPService', 'Host' (checked), 'IMAPService', and 'Interface' (checked). The 'Relations' list also has a search bar and five items: 'ComposedOf (Host)' (checked), 'HostsServices (Host)', 'PartOf (Interface)', 'RealizedBy (Interface)', and 'SystemPackagedIn (Host)'.

To configure a *Topology Report* in simple mode, first you need to select the topology. After it is selected, the *classes* list will automatically be populated from the APG topology. Select the classes you wish to see on the report. After the classes are selected, the *relations* list will also be automatically populated. The relations represent the edge between the nodes on the reports. So after you choose the classes, you must select the relations. (In the relations, the names in parenthesis represent the class from which the relation comes from).

2 Advanced mode

The advanced mode is an extension of the simple mode. It is also based on the APG topology. You can do everything in the advanced mode that you can do in the simple mode. If you switch from the simple mode to the advanced mode, the advanced mode configuration will be configured to display exactly the same report as the simple mode. But, if you switch from the advanced mode to the simple mode, then some information will be lost.

To configure a *Topology Report* in advanced mode, first you need to select the topology. The *filter* flag indicates whether the nodes will be restricted to the current report filter (like in simple mode). Note that this does not apply to the root nodes. The root nodes are the starting point of the graph. ie. when you add a class with the button *Add classes...* The root nodes are always restricted to the report filter.

For example, if we have the following configuration.




If the report filter is set to *devtype=='Switch'*, and filter flag is not set, the host will be displayed on the graph, even if the *switch* do not match the report filter.

The previous example will display all the switches (determined by the report filter) linked to all the hosts that are in the topology.

Configuration:

The starting point of the advanced mode configuration is the selection of the root classes. After the root classes are selected you can add relations by clicking on the previously created class.



The **Add relations...** option will open a new window displaying the available relations, you can select one or many relations. After one relation is created, you restart the process by clicking on another class.

The **Edit classes...** option allows you to modify the selection of classes.

The **Delete** option will delete the selected class and the following relations/classes.

This **Hide this class** option allows you to hide a class in the graph. For example:



This example will display the routers linked together. We must hide the interface because in the topology it is the *Interface* that are linked together and not the *Router*. By hiding the interfaces, the resulting graph will only display the routers. Even though the interfaces are hidden, we can still show some metrics about those interfaces.

3 Drilldown mode

The drilldown mode report will render a report based on the subreports of the current one. Through the UI, children can be selected up to a depth of 4, and then configured to be used as groups, nodes or topology instances. These instances will be used to determine the topology elements and relations to retrieve. An exact match will be performed between the tree nodes names and the instances names. Matching elements will generate nodes on the report. Then any relation between those selected instances will generate a corresponding edge in this report.

Groups and nodes will not directly generate elements. They will only appear if they contain some elements generated from the instances definitions. In that case, *groups* will be used to define groups of elements while *nodes* will define groups whose elements won't be displayed in the graph, thus efficiently limiting the amount of nodes to be rendered.

For drawing the nodes, a template is needed which contains the images references. (See [Rendering](#) section on page 8 for more information.) If no template is specified or if the template doesn't contain reference to the subreport, then the subreports icon will be used.

4 Topology Service mode

The topology service mode uses topology service resources that have been defined globally in the web application context. They will appear here, categorized in different groups depending on the way they are defined.

- *Clusters* when multiple resources have been defined as part of a single entity that will provide the same topology,
- *Resources* when resources have been defined as stand-alone topology services,
- *Unknown* when the resource configured in the report has not been found within the defined ones.

Please note that the first one will probably be used most of the time, but the last one will always result in an error and should be fixed by either modifying the selected entry in the list or defining the resource in the web application context. See below for a definition example.

```
<Resource cluster="SRM-Topology" name="topo-svc/srm1"
  type="com.watch4net.apg.v2.gui.report.topomap.configuration.data.TopoServiceResource"
  factory="org.apache.naming.factory.BeanFactory"
  url="https://localhost:48443/Backends/Topology-Service/Default/topology/repository/"
  username="admin" password="changeme" sslValidation="false" />
```

Not all attributes can be modified in this element. `type` and `factory` have to be set with these exact values.


- `cluster` is an optional attribute that identifies the name of the cluster this resource belongs to.
- `name` uniquely identifies this resource. The prefix `topo-svc/` cannot be changed.
- `url` points to the service that will provide the topology data. It needs to be properly escaped. All remaining path elements and parameters can then be provided on the report itself.
- `username` The name of the user, for authentication purposes.
- `password` The password of the user, for authentication purposes.
- `sslValidation` is an optional attribute that should only be used when valid certificates have been setup yet.

WARNING: The password should never be set as plain text but rather obfuscated using the corresponding tool.

5 Custom mode

The custom mode report is configured in an XML file. The location of this XML file must be in the *topomap* directory of the *WebApps-Resources* module when the file is defined as local one. For example: `<APG>/Custom/WebApps-Resources/Default/topomap`.


When using a remote file available through http or https, then the full path needs to be specified, including placeholders for node properties. For instance, one could use `http://my.server/static-xml-service?selected=${device}&type=${devtype}` in order to call a service which accepts a device type and name in order to generate a dedicated map.



To configure this type of report see the [Configuration](#) section.

6 Rendering

After the mode has been selected and configured, you must select the rendering option.



The first option *layout* determines the layout of the report. The *manual* layout exists only for the custom mode.

The *template* option determines how the nodes and edges will be drawn. For the *simple* and *advanced* modes, the **id-ref** of each node is the same as the class. For the *drilldown* mode the **id-ref** is the subreport in edit mode. For example, if we have a subreport named 'Router' and we do an expand on it, then each node of this expansion will have the **id-ref** set to 'Router'. For the custom mode, see the [Template](#) section.

If you select more than one template and the same node is located in multiple templates, only one of them will be valid.

For the edge part of the *simple* and *advanced* mode, its **id-ref** is formatted as *Source.Relation.Destination*. For example, the following configuration:



will have its edge **id-ref** equal to *Switch.Connected.Host*.

If there are hidden classes, they must be also in the edge id. For instance the last example of the advanced mode section would be equal to *Router.ComposedOf.Interface.Connecting.PartOf.Router*.

3

Configuration

This section contains all the configuration options that are available in the templates and *custom* topology report. Please note that although the first version of said XML files is still partially supported, it is not recommended to use them as new features and options have only been added to the second version that is described here.

XSD schemas are available in any installation. They are copied within the *Custom WebApps-Resources topomap* directory, in a dedicated location.

1 Template

This section lists the main XML tags of a *Topology Report* template definition.

Here is an example of a typical *template* element.

```
<template xmlns="http://www.watch4net.com/Frontend/TopologyReport2" xmlns:xsi="http://
www.w3.org/2001/XMLSchema-instance" xsi:schemaLocation="http://www.watch4net.com/
Frontend/TopologyReport2_1/schemas/template.xsd">

  <options> ... </options>

  <bullet-indicator ... </bullet-indicator>
  <sparkline-indicator ... </sparkline-indicator>

  <node ... </node>
  <node ... </node>

  <edge-type ... </edge-type>

  <edge ... </edge>
  <edge ... </edge>

</template>
```

A *template* tag does not accept any attribute. It does, however, accept the following nested elements, in the specified order.

- [Options](#) which configures the global behavior of the maps (see page 10).
- [Highlight](#) which highlights a specific node or group (see page 11).
- [Bullet Indicator](#) which draws a round colored bullet icon (see page 12).
- [Image Indicator](#) which displays an image (see page 15).
- [Sparkline Indicator](#) which renders a small chart called sparkline (see page 18).
- [Tooltips](#) which configures the tooltip contents available on a node, group, or edge (see page 22).
- [Node](#) which defines the look and feel of a node with a specific type (see page 23).
- [Group](#) which defines the look and feel of a group with a specific type (see page 25).
- [Edge-Type](#) which defines a way to determine an edge attributes based on a value (see page 26).
- [Edge](#) which defines the look and feel of an edge with a specific type (see page 29).
- [Background](#) which sets the background of the resulting graph (see page 32).

2 Options

Options are used to configure the global behavior of the maps, and especially how the maps react on some specific user interactions.

Here is an example of a typical *options* element for a completely static map.

```
<options>
  <viewport group-expansion="static" group-collapse="static" />
  <layout groups-behavior="static" />
</options>
```

2.1 Viewport

This defines the viewport options that can be set on a report.

Attribute	Description	Required
group-expansion	This defines how the viewport will react when the user expands a group. Possible values are <i>static</i> which does not animate the viewport and <i>fit-bounds</i> which sets current zoom and graph center so that all elements are visible in the viewport.	No; defaults to <i>fit-bounds</i> .
group-collapse	This defines how the viewport will react when the user closes a group. Possible values are <i>static</i> which does not animate the viewport and <i>fit-bounds</i> which sets current zoom and graph center so that all elements are visible in the viewport.	No; defaults to <i>fit-bounds</i> .

2.2 Layout

This defines the layout options that can be set on a report.

Attribute	Description	Required
name	Can be used to specify a core layout for this template. This will only be applied as fallback if the report is configured as using a manual layout.	No.
groups-behavior	This defines layout behaviors for groups, that is whether their layout will be re-computed each time they get expanded or collapsed, or if they will remain at the same position after these operations. Possible values are <code>dynamic</code> which performs a relayout after each group modification, and <code>static</code> which makes sure all groups can be expanded without modify the layout. Note that using a <code>static</code> layout takes more time to compute the initial position of nodes.	No; defaults to <code>dynamic</code> .

3 Highlight

Nodes and groups that have a specific meaning can be highlighted to help the user notice them. For instance, if a map is built based on a specific device, which we would call the *subject of the map*, then this element would be used to identify the corresponding node or group. Here is an example of a typical *highlight* element.

```
<highlight id=".subject-of-the-map" color="#ffc552" width="4" />
```

Such a highlight style can then be referenced within a node or a group through its ID, optionally overriding some attributes.

```
<node>
  <highlight id-ref=".subject-of-the-map" />
</node>
```

3.1 Highlight Attributes

Main Attributes

Attribute	Description	Required
id	Uniquely identifies this element so that it can be referenced elsewhere.	Yes.
color	This defines the color the highlight border will take. Use CSS-like values here, that is something like <code>#ffc552</code> . Not specifying this color will disable the border.	No.

Main Attributes		
Attribute	Description	Required
width	This defines the width of the border; Using 0 here will disable the border.	No.
dash	The highlight border can be plain or dashed. When set, the pattern defined in this attribute will be used to generate dashes. For instance, one could use "3 4" to get some regular dashes.	No.
fill	This defines the color the highlight background will take. Use CSS-like values here, optionally specifying some alpha at the beginning, in other words something like #80ffc552.	No.
labels	When set to <code>false</code> , the highlight won't take the labels into account and will only draw a border around the node icon.	No; defaults to <code>true</code> .

4 Bullet Indicator

An indicator defines visual properties that depend on a value compared its thresholds, in other words how it behaves when the value is within a specific range. This indicator draws a round colored bullet icon at a specific location around the node.

Here is an example of a typical *bullet-indicator* element.

```
<bullet-indicator id=".availability" label="Availability" filter="$NODE(device,␣device)$␣␣&part␣&name='Availability'" position="top-left">
  <range from="0" to="80" color="#70c765" />
  <range from="80" to="90" color="#f5d852" />
  <range from="90" color="#eb6e5a" />
  <range color="#999999" />
</bullet-indicator>
```

Such an indicator can then be referenced within a node or group through its ID, optionally overriding its position.

```
<node>
  <bullet-indicator id-ref=".availability" />
  <bullet-indicator id-ref=".status" position="top-right" />
</node>
```

4.1 Bullet Indicator Attributes

Main Attributes		
Attribute	Description	Required
id	Uniquely identifies this element so that it can be referenced elsewhere.	Yes.

Main Attributes

Attribute	Description	Required
label	The label that describes the element, for instance the name of the selected value.	Yes; unless a formula result is used, in which case it will fall back on the result name.
position	Defines the position around the node at which the indicator will be displayed. Possible values are top-left, top, top-right, left, right, bottom-left, bottom, bottom-right, and center.	No.

Data Computation Attributes

Attribute	Description	Required
filter	Either one of filter or formula-result can be used to specify how the value should be retrieved and computed.	No.
formula-result	Either one of filter or formula-result can be used to specify how the value should be retrieved and computed.	No.
temporal-aggregation	This is the temporal aggregation used when fetching the metric values from the database. This also defines which computed aggregated will be selected within the database. Possible values are average, min, max, sum, last, nbvalues, and lasttimestamp.	No; defaults to average.
spatial-aggregation	This is the spatial aggregation used when fetching the metric values from the database. Possible values are average, min, max, sum, and nbvalues.	No; defaults to average.
rounding	The rounding defines the number of decimals that will be kept. This value must be positive.	No; defaults to 5.
scaling-mode	A scaling mode defines the way the scaling factor will be used to scale the value before it is compared to the thresholds. Possible values are multiply, divide, and unit. The last one performs unit-based scaling.	No.
scaling-factor	A scaling factor is used in conjunction with the scaling mode to modify the resulting value before it is compared to the thresholds.	Only when scaling-mode is specified.

Data Computation Attributes

Attribute	Description	Required
unit	To be used when the unit of the value has to be enforced. It will be displayed after the value.	No.
display-value	When set, the computed value will be displayed within the generated bullet.	No; defaults to false.

4.2 Bullet Indicator Ranges

As each indicator, the bullet one can have one or more threshold-defined ranges that will hold the style attributes with which to decorate the indicator. A range can also define the way its value will be computed, which explains why it looks quite similar to the indicator itself.

Main Attributes

Attribute	Description	Required
from	The range styling attributes will only be applied if the computed value is greater than or equals to this value. If not specified, there will be no check against a minimum value when deciding whether to apply or not this range.	No.
to	The range styling attributes will only be applied if the computed value is strictly lower than this value. If not specified, there will be no check against a maximum value when deciding whether to apply or not this range.	No.
color	This defines the color the bullet will take. Use CSS-like values here, that is something like #f5d852.	Yes.

Data Computation Attributes

Attribute	Description	Required
filter	Either one of <code>filter</code> or <code>formula-result</code> can be used to specify how the value should be retrieved and computed.	No.
formula-result	Either one of <code>filter</code> or <code>formula-result</code> can be used to specify how the value should be retrieved and computed.	No.
temporal-aggregation	This is the temporal aggregation used when fetching the metric values from the database. This also defines which computed aggregated will be selected within the database. Possible values are <code>average</code> , <code>min</code> , <code>max</code> , <code>sum</code> , <code>last</code> , <code>nbvalues</code> , and <code>lasttimestamp</code> .	No; defaults to average.

<i>Data Computation Attributes</i>		
Attribute	Description	Required
spatial-aggregation	This is the spatial aggregation used when fetching the metric values from the database. Possible values are average, min, max, sum, and nbvalues.	No; defaults to average.
rounding	The rounding defines the number of decimals that will be kept. This value must be positive.	No; defaults to 5.
scaling-mode	A scaling mode defines the way the scaling factor will be used to scale the value before it is compared to the thresholds. Possible values are multiply, divide, and unit. The last one performs unit-based scaling.	No.
scaling-factor	A scaling factor is used in conjunction with the scaling mode to modify the resulting value before it is compared to the thresholds.	Only when scaling-mode is specified.
unit	To be used when the unit of the value has to be enforced. It will be displayed after the value.	No.
display-value	When set, the computed value will be displayed within the generated bullet.	No; defaults to false.

5 Image Indicator

An indicator defines visual properties that depend on a value compared its thresholds, in other words how it behaves when the value is within a specific range. This indicator draws an image at a specific location around the node.

Here is an example of a typical *image-indicator* element.

```
<image-indicator id=".availability" label="Availability" filter="$NODE(device,device)$"
&name='Availability' position="top-left">
  <range from="0" to="80" path="custom/images/bullet_red.png" />
  <range from="80" to="90" path="custom/images/bullet_yellow.png" />
  <range from="90" path="custom/images/bullet_green.png" />
  <range path="custom/images/bullet_gray.png" />
</image-indicator>
```

Such an indicator can then be referenced within a node or group through its ID, optionally overriding its position.

```
<node>
  <image-indicator id-ref=".availability" />
  <image-indicator id-ref=".status" position="top-right" />
</node>
```

5.1 Image Indicator Attributes

Main Attributes

Attribute	Description	Required
id	Uniquely identifies this element so that it can be referenced elsewhere.	Yes.
label	The label that describes the element, for instance the name of the selected value.	Yes; unless a formula result is used, in which case it will fall back on the result name.
position	Defines the position around the node at which the indicator will be displayed. Possible values are top-left, top, top-right, left, right, bottom-left, bottom, bottom-right, and center.	No.

Data Computation Attributes

Attribute	Description	Required
filter	Either one of filter or formula-result can be used to specify how the value should be retrieved and computed.	No.
formula-result	Either one of filter or formula-result can be used to specify how the value should be retrieved and computed.	No.
temporal-aggregation	This is the temporal aggregation used when fetching the metric values from the database. This also defines which computed aggregated will be selected within the database. Possible values are average, min, max, sum, last, nbvalues, and lasttimestamp.	No; defaults to average.
spatial-aggregation	This is the spatial aggregation used when fetching the metric values from the database. Possible values are average, min, max, sum, and nbvalues.	No; defaults to average.
rounding	The rounding defines the number of decimals that will be kept. This value must be positive.	No; defaults to 5.

<i>Data Computation Attributes</i>		
Attribute	Description	Required
scaling-mode	A scaling mode defines the way the scaling factor will be used to scale the value before it is compared to the thresholds. Possible values are multiply, divide, and unit. The last one performs unit-based scaling.	No.
scaling-factor	A scaling factor is used in conjunction with the scaling mode to modify the resulting value before it is compared to the thresholds.	Only when scaling-mode is specified.
unit	To be used when the unit of the value has to be enforced. It will be displayed after the value.	No.
display-value	When set, the computed value will be displayed within the generated image.	No; defaults to false.

5.2 Image Indicator Ranges

As each indicator, the image one can have one or more threshold-defined ranges that will hold the style attributes with which to decorate the indicator. A range can also define the way its value will be computed, which explains why it looks quite similar to the indicator itself.

<i>Main Attributes</i>		
Attribute	Description	Required
from	The range styling attributes will only be applied if the computed value is greater than or equals to this value. If not specified, there will be no check against a minimum value when deciding whether to apply or not this range.	No.
to	The range styling attributes will only be applied if the computed value is strictly lower than this value. If not specified, there will be no check against a maximum value when deciding whether to apply or not this range.	No.
path	The path pointing to the image to display.	Yes.

<i>Data Computation Attributes</i>		
Attribute	Description	Required
filter	Either one of filter or formula-result can be used to specify how the value should be retrieved and computed.	No.

<i>Data Computation Attributes</i>		
Attribute	Description	Required
formula-result	Either one of <code>filter</code> or <code>formula-result</code> can be used to specify how the value should be retrieved and computed.	No.
temporal-aggregation	This is the temporal aggregation used when fetching the metric values from the database. This also defines which computed aggregated will be selected within the database. Possible values are <code>average</code> , <code>min</code> , <code>max</code> , <code>sum</code> , <code>last</code> , <code>nbvalues</code> , and <code>lasttimestamp</code> .	No; defaults to <code>average</code> .
spatial-aggregation	This is the spatial aggregation used when fetching the metric values from the database. Possible values are <code>average</code> , <code>min</code> , <code>max</code> , <code>sum</code> , and <code>nbvalues</code> .	No; defaults to <code>average</code> .
rounding	The rounding defines the number of decimals that will be kept. This value must be positive.	No; defaults to 5.
scaling-mode	A scaling mode defines the way the scaling factor will be used to scale the value before it is compared to the thresholds. Possible values are <code>multiply</code> , <code>divide</code> , and <code>unit</code> . The last one performs unit-based scaling.	No.
scaling-factor	A scaling factor is used in conjunction with the scaling mode to modify the resulting value before it is compared to the thresholds.	Only when <code>scaling-mode</code> is specified.
unit	To be used when the unit of the value has to be enforced. It will be displayed after the value.	No.
display-value	When set, the computed value will be displayed within the generated image.	No; defaults to <code>false</code> .

6 Sparkline Indicator

An indicator defines visual properties that depend on a value compared its thresholds, in other words how it behaves when the value is within a specific range. This indicator draws a small chart called sparkline at a specific location around the node.

Here is an example of a typical *sparkline-indicator* element.

```
<sparkline-indicator id=".availability" label="Availability" filter="$NODE(device,
device)$&part,&!part,&name='Availability'" position="right" lower-bound="0"
upper-bound="100">
  <range from="0" to="80" color="#70c765" />
  <range from="80" to="90" color="#f5d852" />
  <range from="90" color="#eb6e5a" />
```

```
<range color="#99999" />
</sparkline-indicator>
```

Such an indicator can then be referenced within a node or group through its ID, optionally overriding its position.

```
<node>
  <sparkline-indicator id-ref=".availability" />
  <sparkline-indicator id-ref=".status" position="bottom-right" />
</node>
```

6.1 Sparkline Indicator Attributes

Main Attributes

Attribute	Description	Required
id	Uniquely identifies this element so that it can be referenced elsewhere.	Yes.
label	The label that describes the element, for instance the name of the selected value.	Yes; unless a formula result is used, in which case it will fall back on the result name.
position	Defines the position around the node at which the indicator will be displayed. Possible values are top-left, top, top-right, left, right, bottom-left, bottom, bottom-right, and center.	No.
number-of-values	This attribute defines how many points the sparkline will have. A minimum of 2 is required in order to draw a line.	No; defaults to 42.
lower-bound	Lowest boundary (value) of the sparkline drawing.	No.
upper-bound	Highest boundary (value) of the sparkline drawing.	No.

Data Computation Attributes

Attribute	Description	Required
filter	Either one of filter or formula-result can be used to specify how the value should be retrieved and computed.	No.

Data Computation Attributes

Attribute	Description	Required
formula-result	Either one of <code>filter</code> or <code>formula-result</code> can be used to specify how the value should be retrieved and computed.	No.
temporal-aggregation	This is the temporal aggregation used when fetching the metric values from the database. This also defines which computed aggregated will be selected within the database. Possible values are <code>average</code> , <code>min</code> , <code>max</code> , <code>sum</code> , <code>last</code> , <code>nbvalues</code> , and <code>lasttimestamp</code> .	No; defaults to <code>average</code> .
spatial-aggregation	This is the spatial aggregation used when fetching the metric values from the database. Possible values are <code>average</code> , <code>min</code> , <code>max</code> , <code>sum</code> , and <code>nbvalues</code> .	No; defaults to <code>average</code> .
rounding	The rounding defines the number of decimals that will be kept. This value must be positive.	No; defaults to 5.
scaling-mode	A scaling mode defines the way the scaling factor will be used to scale the value before it is compared to the thresholds. Possible values are <code>multiply</code> , <code>divide</code> , and <code>unit</code> . The last one performs unit-based scaling.	No.
scaling-factor	A scaling factor is used in conjunction with the scaling mode to modify the resulting value before it is compared to the thresholds.	Only when <code>scaling-mode</code> is specified.
unit	To be used when the unit of the value has to be enforced. It will be displayed after the value.	No.

6.2 Sparkline Indicator Ranges

As each indicator, the sparkline one can have one or more threshold-defined ranges that will hold the style attributes with which to decorate the indicator. A range can also define the way its value will be computed, which explains why it looks quite similar to the indicator itself.

Main Attributes

Attribute	Description	Required
from	The range styling attributes will only be applied if the computed value is greater than or equals to this value. If not specified, there will be no check against a minimum value when deciding whether to apply or not this range.	No.

Main Attributes

Attribute	Description	Required
to	The range styling attributes will only be applied if the computed value is strictly lower than this value. If not specified, there will be no check against a maximum value when deciding whether to apply or not this range.	No.
color	This defines the color the sparkline line will take. Use CSS-like values here, that is something like #f5d852.	Yes.

Data Computation Attributes

Attribute	Description	Required
filter	Either one of <code>filter</code> or <code>formula-result</code> can be used to specify how the value should be retrieved and computed.	No.
formula-result	Either one of <code>filter</code> or <code>formula-result</code> can be used to specify how the value should be retrieved and computed.	No.
temporal-aggregation	This is the temporal aggregation used when fetching the metric values from the database. This also defines which computed aggregated will be selected within the database. Possible values are <code>average</code> , <code>min</code> , <code>max</code> , <code>sum</code> , <code>last</code> , <code>nbvalues</code> , and <code>lasttimestamp</code> .	No; defaults to <code>average</code> .
spatial-aggregation	This is the spatial aggregation used when fetching the metric values from the database. Possible values are <code>average</code> , <code>min</code> , <code>max</code> , <code>sum</code> , and <code>nbvalues</code> .	No; defaults to <code>average</code> .
rounding	The rounding defines the number of decimals that will be kept. This value must be positive.	No; defaults to 5.
scaling-mode	A scaling mode defines the way the scaling factor will be used to scale the value before it is compared to the thresholds. Possible values are <code>multiply</code> , <code>divide</code> , and <code>unit</code> . The last one performs unit-based scaling.	No.
scaling-factor	A scaling factor is used in conjunction with the scaling mode to modify the resulting value before it is compared to the thresholds.	Only when <code>scaling-mode</code> is specified.
unit	To be used when the unit of the value has to be enforced. It will be displayed after the value.	No.

7 Tooltips

Tooltips can be assigned to either nodes or edges. They replace the default ones that usually show the node name or edge path and can be enhanced with various indicators and properties. Here is an example of a typical *tooltip* element.

```
<tooltip id=".node-tooltip">
  <bullet-indicator id-ref=".availability" />
  <sparkline-indicator id=".availability" label="Availability" filter="$NODE(device,
device)$&lt;&lt;part&lt;&lt;name='Availability' lower-bound="0" upper-bound="100">
    <range from="0" to="80" color="#70c765" />
    <range from="80" to="90" color="#f5d852" />
    <range from="90" color="#eb6e5a" />
    <range color="#999999" />
  </sparkline-indicator>
  <properties filter="!parttype">
    <property name="devtype"/>
    <property name="device"/>
    <property name="ip"/>
    <property name="part"/>
  </properties>
</tooltip>
```

It could then be referenced within another element using its ID.

```
<node>
  <tooltip id-ref=".node-tooltip" />
</node>
```

7.1 Tooltip Attributes

Main Attributes

Attribute	Description	Required
id	Uniquely identifies this element so that it can be referenced elsewhere.	No.
title	When set, this will enforce the tooltip title text. If not set, the tooltip title will default to either the node name or the edge path.	No.
label	This defines a short description of the node, edge, or tooltip contents.	No.

7.2 Tooltip Contents

The tooltip can be configured to display some specific indicators. These indicators are the same ones that can also be added around each node, and are defined likewise.

- [Bullet Indicators](#) as described on page 12,
- [Image Indicators](#) as described on page 15,
- [Sparkline Indicators](#) as described on page 18,

- and [Properties](#) as described on page 25.

8 Node

In a template, *node* tags are used to configure the look and feel of some elements of the topology report, typically devices, interfaces, and so on... Nodes that represent elements that are in the *APG* database can also have related metrics assigned to them. Otherwise, for example, if the node represents something outside the scope of *APG*, it can be assigned a label describing what it is, and its relationship to the other nodes.

Here is an example of some typical *node* elements.

```
<node id="Host" image="custom/images/topomap/Host.svg">
  <bullet-indicator id-ref=".availability" />
  <image-indicator id-ref=".status" />
</node>
<node id="Fabric" image="custom/images/topomap/Fabric.svg" />
<node id="StorageArray" image="custom/images/topomap/Array.svg">
  <sparkline-indicator id-ref=".usage" />
</node>
<node id="Adapter" image="custom/images/topomap/NetworkAdapter.svg" />
<node id="Interface" image="custom/images/topomap/Interface.svg" scale="0.4" />
<node id="Port" image="custom/images/topomap/Port.svg" scale="0.4" />
```

When multiple templates are used, a node can extend the definition from a previous template as follows.

```
<node id="Host" />
  <sparkline-indicator id-ref=".health" />
</node>
<node id="StorageArray" image="custom/images/topomap/MyArray.svg" />
<node id="Interface" id-ref="Interface" scale="0.5" />
<node id="Port" id-ref="Interface" scale="0.5" />
```

8.1 Node Attributes

Attribute	Description	Required
id	Uniquely identifies this element so that it can be referenced elsewhere.	Yes.
label	The label that describes the element, for instance the name of the selected value.	No.
display-label	Used to control the display of the node label.	No; defaults to true.
title	The title is displayed hover the element, and further describes it.	No.

Attribute	Description	Required
image	This image will be used to visually represent the node in the topology report. It is highly recommended to set a different image on every type of node in order to easily differentiate them.	No; defaults to a generic image.
href	If this attribute is set, the node gets a drill down action. This reference can be absolute, using a complete <code>http</code> address, or can point to a node in the tree, either by ID, or by name.	No.
href-delimiter	It is possible to drill down to a sub-report nested multiple levels down under the parent topology report. For example, suppose we have the following APG tree, and the href attribute is <code>root-report2/report3</code> and the href-delimiter attribute is <code>/</code> . Then when clicking on the performing a drill down on the node, report3 will be the selected report.	No.
use-report-filter	If set to <code>true</code> , the report filter will be merged within the filter of this node. It gets usefull when a specific topology report selects elements with a limited filter, and we would like to be sure indicators will select the right subset of metrics.	No; defaults to <code>false</code> .
filter	This filter can be used by indicators or edges when computing values based on this node.	No.
x	Defines the initial position of the node on the map. Essential when defining static nodes in an XML report with a manual layout.	No.
y	Defines the initial position of the node on the map. Essential when defining static nodes in an XML report with a manual layout.	No.
scale	Scales up or down the node representation on the map. For instance, one could want to have smaller nodes for interfaces, and bigger nodes for large groups of devices. The value must be between 0 (exclusive) and 10 (inclusive).	No; defaults to 1.

8.2 Nested Elements

A node can have multiple nested elements that define additional properties and visual indicators to be displayed around it. Different types of elements can be added: the tooltip contents (see [Tooltips](#) on page 22), a highlight (see [Highlight](#) on page 11), some indicators (see [Bullet Indicator](#) on page 12, [Image Indicator](#) on page 15, and [Sparkline Indicator](#) on page 18) and properties. Properties are described below; they are composed of a *properties* element con-

taining one *property* element for each property to be retrieved and displayed.

Properties

It is usually a good idea to display properties at the bottom or on the right side of the node, due to the space they require. Also, note that there are only nine positions around and on a node, so only nine of these elements can be defined within a single node.

Properties Attributes

Attribute	Description	Required
filter	This filter will be used when fetching the properties from the database. Note that the node's filter is always used when retrieving properties.	No.
position	Location, around the node, where the properties will be displayed. Possible values are top-left, top, top-right, left, right, bottom-left, bottom, bottom-right, and center.	No; defaults to bottom

Property Attributes

Attribute	Description	Required
name	Defines the name of a property to be retrieved.	Yes.
alias	Optional alias for the property name that replaces the default one. If no alias is provided, and the property is documented, then the alias from the documentation will be used; else it will be the name.	No.

9 Group

A group defines a set of nodes that have been aggregated together. It can also have its own style and indicators, which makes it a specific type of node. The *group* element extends the *node* itself, so all attributes and nested elements described respectively on page 23 and page 24 apply to the group as well.

Here is an example of some (not so) typical *group* elements. Note the *id-ref* sharing the same ID: this means the group inherits from either a group defined in a parent template, or from a node defined in the same template. This allows a group to share common attributes with a node of the same type without having to redefine them.

```
<group id="Host" id-ref="Host" layout="HIERARCHICAL={&quot;orientation&quot;:&quot;left&quot;}&quot;" />
<group id="Fabric" id-ref="Fabric" expanded="true" />
<group id="Switch" id-ref="Switch" layout="ORGANIC" />
<group id="StorageArray" id-ref="StorageArray" layout="HIERARCHICAL={&quot;orientation&quot;:&quot;right&quot;}&quot;" />
```

Styling Attributes

Attribute	Description	Required
expanded	By default, a group appears collapsed (like a node) and does not display its children. The user can manually perform the expansion operation though. Set this to <code>true</code> in order to have it opened when drawing the graph.	No; defaults to <code>false</code> .
color	The color is used as background in order to visually identify an expanded group.	No; defaults to <code>#fafafa</code> .
layout	The layout is used to assign an inner layout to a group. Beware that this forces an independant layout of the group elements, which results in suboptimal positioning of external edges.	No.

10 Edge-Type

An edge type defines a type of link between two element in the graph. Like the *node*, an *edge* can have some indicators which are used to change its visual aspects depending on different thresholds and values. These indicators are called *edge-type*.

Here is an example of some typical *edge-type* elements that could be used between a host and its interfaces.

```
<edge-type id=".utilization" label="LinkUtilization" filter="$SOURCE(device,device)$_&
&$DESTINATION(part,part)$_&$parttype='Interface'&$name='CurrentUtilization'">
  <range to="20" width="1" color="#70c765" />
  <range from="20" to="50" width="2" color="#f5d852" />
  <range from="80" width="4" color="#eb6e5a" />
</edge-type>
```

Such an edge type can then be referenced within an edge through its ID.

```
<edge>
  <edge-type id-ref=".utilization" />
</edge>
```

10.1 Edge-Type Attributes

Main Attributes

Attribute	Description	Required
id	Uniquely identifies this element so that it can be referenced elsewhere.	Yes.

Main Attributes

Attribute	Description	Required
label	The label that describes the element, for instance the name of the selected value.	No.
display-value	Used to control the display of the computed value (with its label) within the edge tooltip.	No; defaults to true.

Data Computation Attributes

Attribute	Description	Required
filter	Either one of <code>filter</code> or <code>formula-result</code> can be used to specify how the value should be retrieved and computed.	No.
formula-result	Either one of <code>filter</code> or <code>formula-result</code> can be used to specify how the value should be retrieved and computed.	No.
temporal-aggregation	This is the temporal aggregation used when fetching the metric values from the database. This also defines which computed aggregated will be selected within the database. Possible values are <code>average</code> , <code>min</code> , <code>max</code> , <code>sum</code> , <code>last</code> , <code>nbvalues</code> , and <code>lasttimestamp</code> .	No; defaults to average.
spatial-aggregation	This is the spatial aggregation used when fetching the metric values from the database. Possible values are <code>average</code> , <code>min</code> , <code>max</code> , <code>sum</code> , and <code>nbvalues</code> .	No; defaults to average.
rounding	The rounding defines the number of decimals that will be kept. This value must be positive.	No; defaults to 5.
scaling-mode	A scaling mode defines the way the scaling factor will be used to scale the value before it is compared to the thresholds. Possible values are <code>multiply</code> , <code>divide</code> , and <code>unit</code> . The last one performs unit-based scaling.	No.
scaling-factor	A scaling factor is used in conjunction with the scaling mode to modify the resulting value before it is compared to the thresholds.	Only when <code>scaling-mode</code> is specified.
unit	To be used when the unit of the value has to be enforced. It will be displayed after the value.	No.

10.2 Edge-Type Ranges

This element defines how a specific edge is displayed when the value is within a specific range. Ranges are checked in the order they are defined, and the first one that fits, given the computed value, is selected and applied.

Main Attributes

Attribute	Description	Required
id	Uniquely identifies this element so that it can be referenced elsewhere.	Yes.
label	The label that describes the element, for instance the name of the selected value.	No.
display-value	Used to control the display of the computed value (with its label) within the edge tooltip.	No; defaults to true.
from	The range styling attributes will only be applied if the computed value is greater than or equals to this value. If not specified, there will be no check against a minimum value when deciding whether to apply or not this range.	No.
to	The range styling attributes will only be applied if the computed value is strictly lower than this value. If not specified, there will be no check against a maximum value when deciding whether to apply or not this range.	No.
color	This defines the color the edge line will take. Use CSS-like values here, that is something like #f5d852.	No; defaults to black.
width	This defines the width the edge line will have.	No; defaults to 1.
dash	An edge line can be plain or dashed. When set, the pattern defined in this attribute will be used to generate dashes. For instance, one could use "3 4" to get some regular dashes.	No.

Data Computation Attributes

Attribute	Description	Required
filter	Either one of <code>filter</code> or <code>formula-result</code> can be used to specify how the value should be retrieved and computed.	No.
formula-result	Either one of <code>filter</code> or <code>formula-result</code> can be used to specify how the value should be retrieved and computed.	No.

<i>Data Computation Attributes</i>		
Attribute	Description	Required
temporal-aggregation	This is the temporal aggregation used when fetching the metric values from the database. This also defines which computed aggregated will be selected within the database. Possible values are average, min, max, sum, last, nbvalues, and lasttimestamp.	No; defaults to average.
spatial-aggregation	This is the spatial aggregation used when fetching the metric values from the database. Possible values are average, min, max, sum, and nbvalues.	No; defaults to average.
rounding	The rounding defines the number of decimals that will be kept. This value must be positive.	No; defaults to 5.
scaling-mode	A scaling mode defines the way the scaling factor will be used to scale the value before it is compared to the thresholds. Possible values are multiply, divide, and unit. The last one performs unit-based scaling.	No.
scaling-factor	A scaling factor is used in conjunction with the scaling mode to modify the resulting value before it is compared to the thresholds.	Only when scaling-mode is specified.
unit	To be used when the unit of the value has to be enforced. It will be displayed after the value.	No.

11 Edge

An edge defines a link between two element whose visual properties depend on a value compared to the defined thresholds. It can either have its own style or rely on one or more edge-type styles.

Here is an example of some typical *edge* elements that could be used between two interfaces.

```
<edge id="Interface.ConnectedTo.Interface">
  <edge-type label="LinkUtilization" filter="($SOURCE(device,device)$&part,$SOURCE(part,part)$&DESTINATION(device,device)$&DESTINATION(part,part)$&parttype='Interface'&name='CurrentUtilization')">
    <range to="20" width="1" color="#70c765" />
    <range from="20" to="50" width="2" color="#f5d852" />
    <range from="80" width="4" color="#eb6e5a" />
  </edge-type>
</edge>
```

Such an edge can then be referenced from another edge through its ID.

```
<edge id-ref="Interface.ConnectedTo.Interface">
  <edge-type>
    <range dash="3_4" />
  </edge-type>
</edge>
```

11.1 Edge Attributes

Note that only a value computed on an edge can be displayed alongside the edge line.

Main Attributes

Attribute	Description	Required
id	Uniquely identifies this element so that it can be referenced elsewhere.	Yes.
label	The label that describes the edge value.	No.
display-label	Used to control the display of the previously defined label alongside the edge line.	No; defaults to false.
display-value	Used to control the display of the computed value alongside the edge line.	No; defaults to false.
href	If this attribute is set, the edge gets linked to a corresponding node in the tree, which can then be used in filters. This reference can point to a node in the tree, either by ID, or by name.	No.
href-delimiter	It is possible to select a sub-report nested multiple levels down under the parent topology report. For instance, suppose we have the following APG tree, and the href attribute is root-report2/report3 and the href-delimiter attribute is /. Then when replacing the placeholder, report3 will be the selected report.	No.

Data Computation Attributes

Attribute	Description	Required
filter	Either one of filter or formula-result can be used to specify how the value should be retrieved and computed.	No.
formula-result	Either one of filter or formula-result can be used to specify how the value should be retrieved and computed.	No.

<i>Data Computation Attributes</i>		
Attribute	Description	Required
temporal-aggregation	This is the temporal aggregation used when fetching the metric values from the database. This also defines which computed aggregated will be selected within the database. Possible values are average, min, max, sum, last, nbvalues, and lasttimestamp.	No; defaults to average.
spatial-aggregation	This is the spatial aggregation used when fetching the metric values from the database. Possible values are average, min, max, sum, and nbvalues.	No; defaults to average.
rounding	The rounding defines the number of decimals that will be kept. This value must be positive.	No; defaults to 5.
scaling-mode	A scaling mode defines the way the scaling factor will be used to scale the value before it is compared to the thresholds. Possible values are multiply, divide, and unit. The last one performs unit-based scaling.	No.
scaling-factor	A scaling factor is used in conjunction with the scaling mode to modify the resulting value before it is compared to the thresholds.	Only when scaling-mode is specified.
unit	To be used when the unit of the value has to be enforced. It will be displayed after the value.	No.

11.2 Nested Elements

Usually, an edge only has one or more *edge-type* nested elements. Each type will be checked separately, meaning for each type, a range can apply. This proves useful to combine multiple different style attributes coming from different types, and thus from different values. Each style will be applied one after the other, eventually overriding any previously set attribute.

One can also configure the contents of the edge tooltips (see [Tooltips](#) on page 22) through a nested *tooltip* element. If provided, this element has to appear just before the first *edge-type* child element, or at the end if no such child element exists.

When used in a report, and not a template, edges need to be given source and destination as additional elements. These elements have to be the first child elements of the edge.

```
<edge id-ref="Interface.ConnectedTo.Interface">
  <from node="host1-interface1" />
  <to node="host2-interface2" />
</edge>
```

Attribute	Description	Required
node	Contains the ID of a node defined in this report.	Yes.
position	The position on the node from (or to) which the edge line will start (or arrive). Possible values are top-left, top, top-right, left, right, bottom-left, bottom, bottom-right, and center.	No.

12 Background

It is possible to define a specific background for the *Topology Report*. Note that this only proves useful when using manual positioning of elements. Else they are more than likely to be positioned outside of said background.

Here is an example of a typical *background* element.

```
<background x="100" y="200" height="900" width="1400" image="custom/images/topomap/my-background-image.jpg"/>
```

Attribute	Description	Required
x	X-coordinate of the location at which the background should be displayed.	No; defaults to 0.
y	Y-coordinate of the location at which the background should be displayed.	No; defaults to 0.
width	Expected width, in pixels, of the background image.	Yes.
height	Expected height, in pixels, of the background image.	Yes.
image	Path to the image file to use as background.	Yes.

13 Report

This section lists the main XML tags of a custom XML-based *Topology Report* definition.

Here is an example of a typical *topology-report* element. The referenced elements obviously come from an external template that will need to be provided when configuring the *Topology Report*.

```
<topology-report xmlns="http://www.watch4net.com/Frontend/TopologyReport2" xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" xsi:schemaLocation="http://www.watch4net.com/Frontend/TopologyReport2../schemas/topology-report.xsd">
  <node id="H1" id-ref="Host" label="H1" indicators-filter="devtype=='Host'&device=='H1'" />
```

```

<node id="RTR1" id-ref="Router" label="RTR1" indicators-filter="device=='RTR1'&
    &_devtype=='Router'" />

    <group id="group" label="RTR2">
        <node id="RTR2" id-ref="Router" label="RTR2" indicators-filter="device=='RTR2'&
            &_devtype=='Router'" />
    </group>

    <edge id="H-R" id-ref="Host.ConnectedTo.Router" label="ConnectedTo">
        <from node="H1" />
        <to node="RTR1" />
        <edge-type>
            <range color="#880044" />
        </edge-type>
    </edge>

    <edge id="R-H" label="ConnectedTo">
        <from node="H1" />
        <to node="RTR1" />
        <edge-type id-ref="et0" />
    </edge>

</topology-report>

```

A *topology-report* tag does not accept any attribute. It does, however, accept the following nested elements, in the specified order.

- **Node** which defines the look and feel of a node with a specific type (see page 23).
- **Group** which defines the look and feel of a group with a specific type (see page 25).
- **Edge** which defines the look and feel of an edge with a specific type (see page 29).
- **Background** which sets the background of the resulting graph (see page 32).

14 Placeholders

It is possible to put some placeholders on the edge's filter and all the href attributes. (for the [Node](#) and [Edge](#))

Filters:

Note: the source node is the one specified by the *from* attribute of the edge and the destination node is the one specified by the *to* attribute of the edge.

- **\$SOURCE\$**: The *source* placeholder can be set on any edge or edge-type *filter*. It will be replaced by the source node *filter*.
For example, if the source node has the filter *devtype='Router' & device='RTRA'* and the edge's filter is *name='Availability' & \$SOURCE\$* and the edge's from attribute is the node 'RTRA', then its filter will be transformed to *name='Availability' & (devtype='Router' & device='RTRA')*.
- **\$DESTINATION\$**: The *destination* placeholder has the same behavior as the source placeholder except that it will take the destination node as replacement.
- **\$SOURCE(prop1, value1)\$, \$DESTINATION(prop1, value1)\$**: This special placeholder will be replaced by a value of a property of the source or destination node.

For example, let's take the node's filter `devtype='Router' & device='RTRA'`, and the edge filter `$SOURCE(toDevice, device)$`. The filter will be transformed to `toDevice='RTRA'` because the device property of the source node is equal to 'RTRA', and the `toDevice` is just a literal.

- **\$NODE(prop1, value1)\$**: The *node* is only available on the indicators. It will be replaced by a value of a property of the parent node.
- **#href**: The *href* is available on all filters. It will be replaced by the filter of the href node.

For *advanced mode* with hidden nodes only:

- **\$NODE1\$, \$NODE2\$, etc:** The *NODE1* placeholder has the same behavior as the source placeholder except that it will take the first hidden node as replacement. The *NODE2* will take the second hidden node.
- **\$NODE1(prop1, value1)\$, \$NODE2(prop1, value1)\$, etc:** Same as `$SOURCE(prop1, value1)$` but it will take the property value of the first, second, etc. hidden node.

Href:

- **\$NODE\$**: This placeholder is only available for the [Node](#). It will be replaced by the name of the current node.
- **\$NODE(prop)\$**: This placeholder is only available for the [Node](#). It will be replaced by the property *prop* of the inherited filter.

For example, suppose we have the following filter `devtype='Router' & device='RTRA'` and we set the *href* attribute to `$NODE(devtype)$/$NODE(device)$`. Then the href attribute will be transformed to `Router/RTR4`.

- **\$SOURCE(prop)\$**: This placeholder is only available for the [Edge](#). It will be replaced by the property *prop* of the inherited filter of the source node.
- **\$DESTINATION(prop)\$**: This placeholder is only available for the [Edge](#). It has the same behavior of the precedent (`$SOURCE$`) placeholder.

For *advanced mode* with hidden nodes only:

- **\$NODE1(prop)\$, \$NODE2(prop)\$, etc:** This placeholder is only available for the [Node](#). It will be replaced by the property *prop* of the inherited filter of the first, second, etc. hidden node.