

Cisco WAE Live User Guide

First Published: 2019-12-17

Last Modified: 2022-11-04

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED “AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

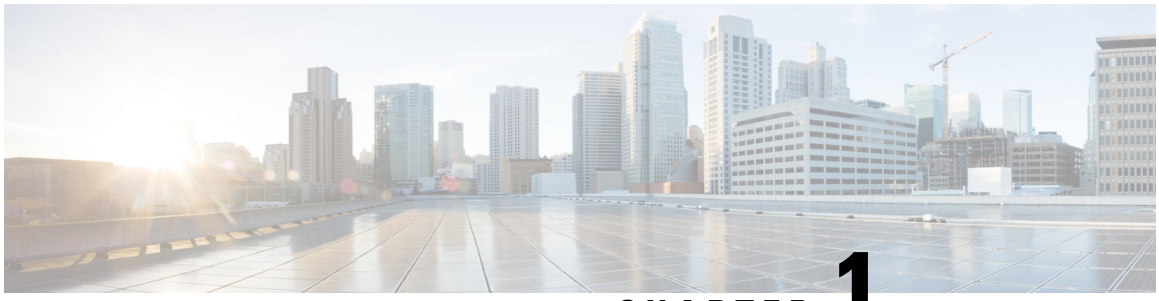
Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

All printed copies and duplicate soft copies of this document are considered uncontrolled. See the current online version for the latest version.

Cisco has more than 200 offices worldwide. Addresses and phone numbers are listed on the Cisco website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: https://www.cisco.com/c/en/us/about/legal/trademarks.html. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)

© 2019-2022 Cisco Systems, Inc. All rights reserved.



CHAPTER 1

Overview

WAE Live is a web-based application that provides visibility into the network infrastructure, state, and traffic utilization.

WAE Live provides immediate and easy access to both current and historical network data. Combined, the Explore, Analytics, and Map tools offer a flexible and interactive means of finding summarized aggregate views or quickly narrowing the search on network data to only relevant details.

WAE Live includes the following components:

- Explore—Lets you view interfaces, interface queues, nodes, LSPs, demands, and reports.
- Inventory—Lets you view your current network hardware in various formats.
- Analytics—Lets you generate reports and analyze data.
- Map—Includes a near-realtime weathermap and health panels from which you can navigate to relevant network data.

WAE Live information is available on a per-network basis if configured by the administrator.

This section contains the following topics:

- Sample Use Cases, on page 1
- WAE Live Menu, on page 2
- Data Insertion, on page 2

Sample Use Cases

You can use WAE Live across diverse groups with different use cases:

- Operators—Instantly detect backbone congestion, assess currently impacted LSPs, and decide reroutes based on size and path.
- Planners—Develop trending reports for importing into WAE Design growth plan tools to determine future needs, such as site-to-site traffic trends per class of service.
- Designers and Engineers—Determine whether LSPs are routed on their shortest paths over time and what events cause them to reroute.
- Peering Coordinators—Monitor priority peers and their growth while verifying that traffic is balanced.
- Business Managers—Determine impact of product launches per market or per site.

- Administrators—Configure and manage the WAE Live application.

WAE Live Menu

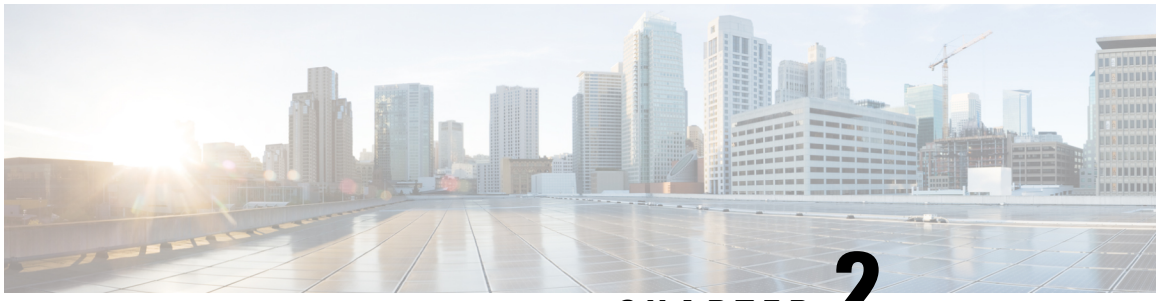
The availability of menu items depends on the licenses purchased and your role (*user* or *admin*).

Menu Item	Description
Explore	Lets you navigate to objects and their properties to view their current or historical state through tables and graphs. You can also use the selections for input when creating reports from Analytics pages. See Exploring Network Data, on page 17.
Inventory	Lets you view a hierarchical representation of the current hardware in your network. See Viewing Inventory, on page 27.
Analytics	• The Configure Report page lets you create multiple reports, including custom (ad hoc), health, traffic, and deviation. See Configuring and Running Reports, on page 33. • The Reports page lists those reports that you have named and identified for frequent use. You can run reports again, access report results, and schedule reports to run at specified dates with results mailed to you or others. • The Report Log lists all reports previously run, including both named and unnamed. You can rerun reports and access report results. This page lists each report's job ID, which is required information for importing traffic reports into the WAE Design GUI. See Using the Report Log, on page 45.
Map	Displays a physical topology (weathermap) in near-realtime to show what is currently happening in the network. Additionally, there are high-level health panels, and you can drill down to see more detailed information. See Viewing Network Health and Weathermaps, on page 69.
Settings	If you are logged in as an administrator, lets you perform administrative tasks, such as configuring the data source and creating new properties.

Data Insertion

The administrator configures the collection of data for use in the WAE Live application. If the data insertion process is halted, a message appears at the top of each page alerting you to the issue.

- If you have a user role, contact your WAE Live or system administrator immediately to fix the issue.
- If you have an admin role, click **Check Status** on the **Settings > Data Source** tab to verify the issue. To troubleshoot the source of the issue, use the ml_insert_ctl tool. For information, see ml_insert_ctl -help .



CHAPTER 2

Understanding Objects, Properties, and Data

The Explore and Analytics tools let you navigate to and report on information about network objects and their properties. You can also navigate to some of these details from the Map page. This information is stored in a time-series optimized datastore that obtains data through the plan files.

This section contains the following topics:

- About Objects and Properties, on page 3
- Raw Data, on page 5

About Objects and Properties

Note Available objects and properties depend on the licenses purchased and how WAE Server is configured. This guide references all possible shipped objects and properties.

WAE Live **objects** are network elements discovered and whose information is displayed in WAE Live. Objects include interfaces, interface queues, nodes, LSPs, and demands. When reporting, these objects can be grouped so that you can aggregate their properties.

WAE Live *properties* are object attributes that are tracked in and retrieved from the datastore. These properties appear as column headings on the Explore page.

If you have an administrator role, you can configure new object properties.

If you are in an Explore or Analytics page and do not see an expected user-defined property, refresh the browser.

Note All traffic is displayed in Mbps.

Table 1: WAE Live Objects and Properties

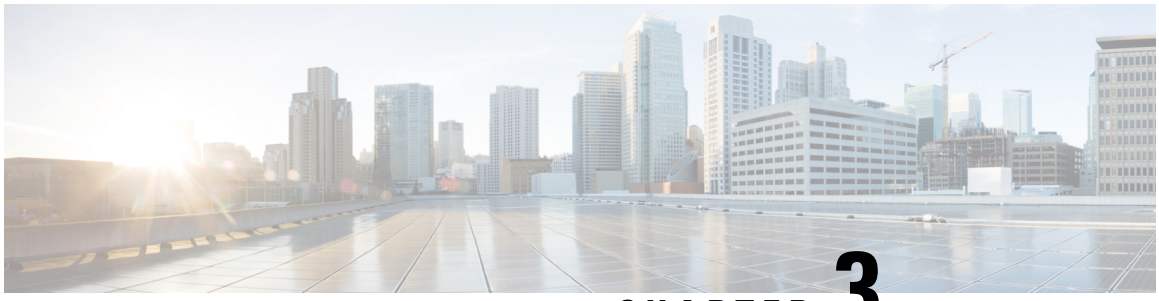
Object	Properties		
Demands	• Destination • IGP Metric • Maximum Latency • Name	• Path Metric • Service Class	• Source • Tags • Traffic
Interfaces	• Admin Status • AS • Capacity • Description • Dropped Packets Out • Error Packets In • Function • IP Address	• Metric • Name • Node • Operating Status • Packets In • Packets Out • Remote AS • Remote Interface	• Remote Node • Tags • Type* • Traffic In • Traffic Out • Util In • Util Out * The Type property appears as a number 1-6. For a description of each of these interface types by number, see Exploring Network Data, on page 17.
Interface Queues	• Error Packets In • Error Packets Out • Interface Queue	• Node • Packets In • Packets Out	• Traffic In • Traffic Out
LSPs	• Actual Path • Admin Status • Name • Operating Status	• Remote Node • SetupBW • Shortest TE Path	• Source Node • Tags • Traffic
Nodes	• CPU0 • CPU1 • Function • IP Address • Mem0	• Mem1 • Model • OS • Name	• SNMP Error • Tags • Vendor

Raw Data

The basis of property values is called *raw data* . Property values within a plan file are stored (in the *datastore*) as raw data points in WAE Live. When a property value changes, a new raw data point is stored.

For example, if a plan file is generated once a day with updated traffic, each week would comprise 7 raw data points. However, if one plan file is generated per hour for updated traffic, each week would comprise 168 raw data points (24 hours * 7 days).

This raw data appears in graphs, showing data over time. The raw data also appears in Explore tables, where each column represents the most recent property value.



CHAPTER 3

Viewing Graphs

You can view graphs in Explore and in the report outputs. Graphs show properties over time, letting you determine when significant events occur. For example, you can graph traffic to identify unexpected shifts, visualize growth trends, and determine whether traffic returns to normal levels after maintenance. Drops in traffic could indicate a routing problem.

This section contains the following topics:

- Graph Icons, on page 7
- View Graph Options, on page 8
- Graphing Multiple Objects, on page 9

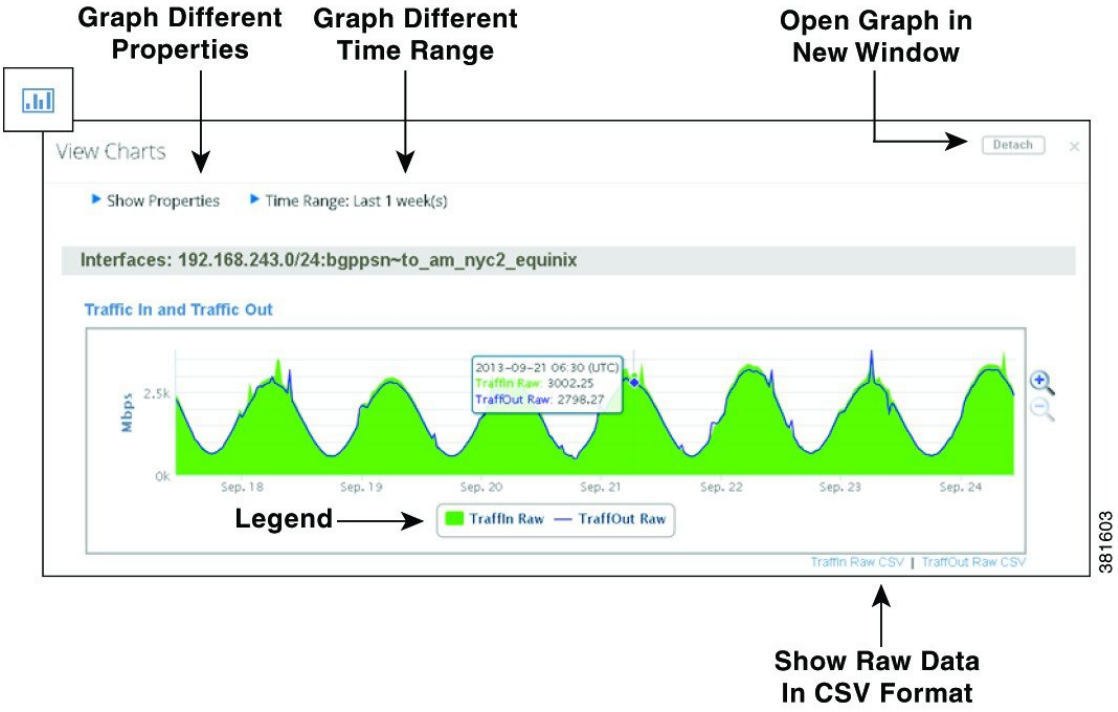
Graph Icons

You can view graphs based on a single table entry or on all objects in a table depending on the icon you choose.

These graph icons are available both through the Explore pages and as part of the output generated by most reports.

Figure 1: Example Single Graph of Interface Traffic In and Traffic Out Properties, on page 8 shows an example of a single Explore graph of an interface's Traffic In and Traffic Out over the last week.

Figure 1: Example Single Graph of Interface Traffic In and Traffic Out Properties



View Graph Options

In the View Charts page, you can choose from any of the following options to generate additional graph views:

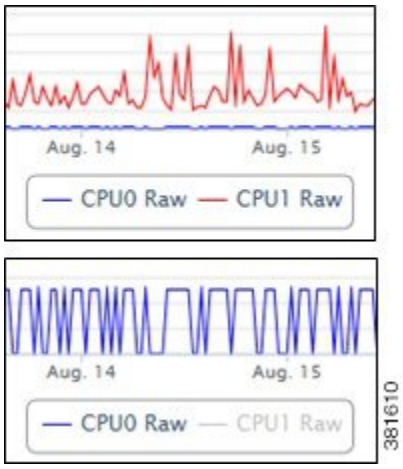
- Show Properties—Choose which properties to show in the graph.
- Time Range—Graph properties across a different time range. After changing the time range, you must click Redraw Graphs.
- Data Gap—A measurement of the time delta between two consecutive data points. This delta determines whether the difference between the two data points should be shown as a gap in the traffic. Enter the value in hours, minutes, and seconds.

Measurement gaps can occur when objects appear and withdraw from a network, when circuits become operational and non-operational, or when demands appear or disappear due to traffic fluctuations. While these are normal conditions, you might find the data easier to read if you adjust the Data Gap value.

Reports include an aggregate measure, such as P95 over one day. When measurement gaps span the aggregation period, the graphs keep the aggregate at a constant rate using the last aggregate calculation. The result is a continuation of a graphed line when possibly there is no data due to network maintenance or failures. You can show these gaps by adjusting the Data Gap value.

- Graph Mode—When graphing multiple objects, change the graph mode. For more information, see Graphing Multiple Objects, on page 9.
  - One property per object.
  - One property across multiple objects in either an overlay or a stacked manner.

- Hide and show graphed properties—Click the property or graphed value (such as a trend) in the legend below the graph to toggle it on and off.



- View values—View values of any point in the graph by hovering your mouse over a data point.



- CSV data—To list the graphed information in CSV format, click the CSV link below the graph. You can then use the browser tools to save or send the data, or you can cut and paste it to a file.
- Magnifying glass—Enlarge or decrease the size of the overall graph.
- Zoom—Zoom into a specific segment of the graph. Click and hold the cursor at the beginning of the range, and then drag it to the end point of the segment you need to enlarge. Zooming helps better read the reported data, such as when a deviation occurs.

Graphing Multiple Objects

You can graph all objects by clicking the **Graph All** icon at the top of the table. Thereafter, there are three available graph modes.

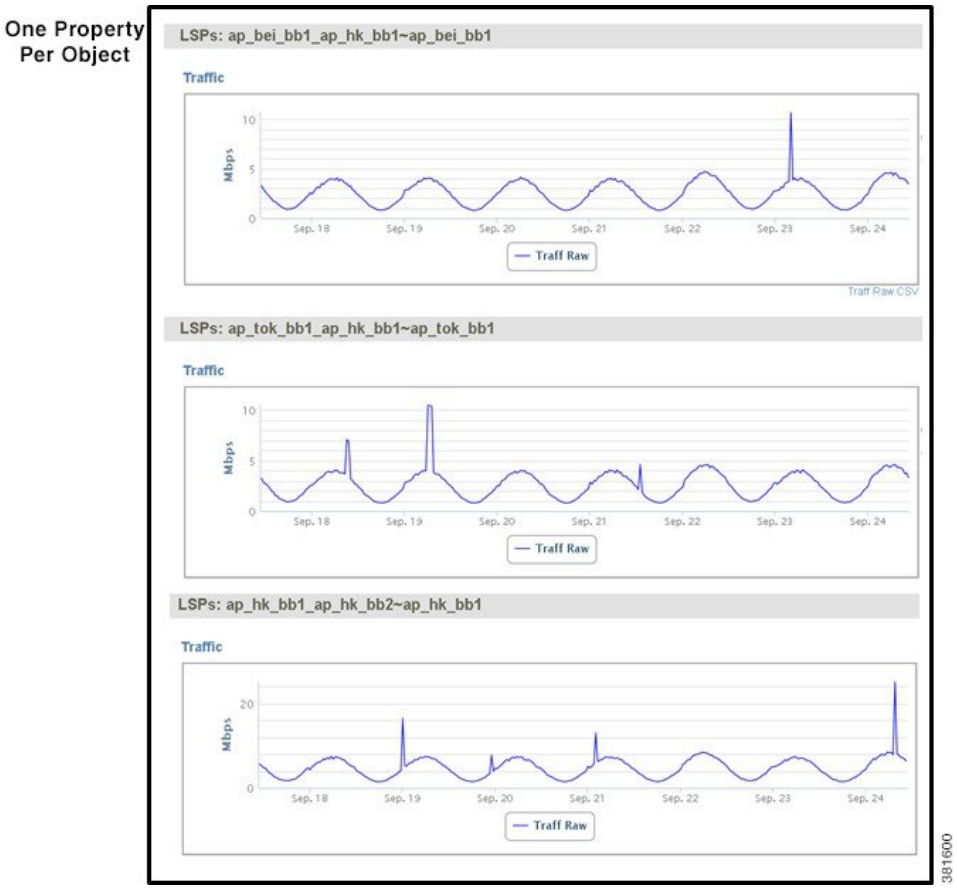
- Graphing by Individual Objects, on page 10
- Graphing by Multiple Objects, on page 10 (**Overlay by property** and **Stack by property**)



Graphing by Individual Objects

Individually graphing each property per object is useful when you want to see all of the data for an object in one place. For example, this mode makes it easy to compare inbound and outbound traffic for a particular interface. Figure 2: Example of Graphing Individual Objects, on page 10 is an example where the Traffic property for three LSPs is graphed using the Individual Object graph mode.

Figure 2: Example of Graphing Individual Objects



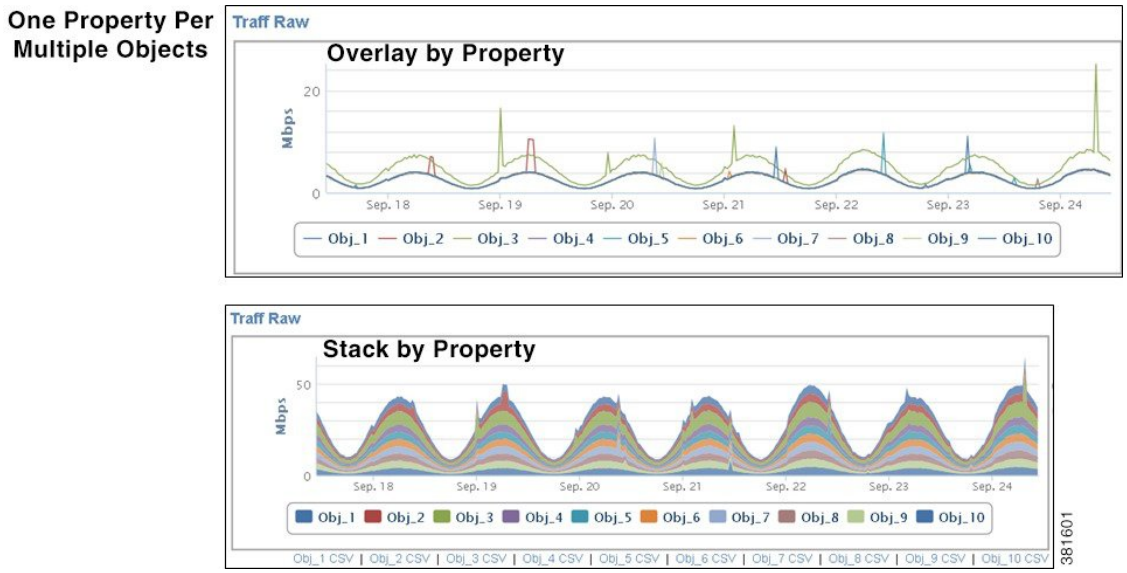
Graphing by Multiple Objects

You can graph a single property for multiple objects either by overlaying them or by stacking them. Figure 3: Example of Overlay and Stack Graphs, on page 11 shows an example of each using the same data (Traffic property, time frame, and LSPs) as was generated for Figure 2: Example of Graphing Individual Objects, on page 10.

- **Overlay by property**—Overlay a single property for multiple objects in one graph. This is useful when you want to compare a single property across multiple objects. For example, if you want to compare inbound traffic data for all interfaces, overlay graphs make it easy to see which interfaces are carrying the most traffic.
- **Stack by property**—Stack a single property for multiple objects in one graph one on top of the other. The size of each colored area is per property and is proportional to the total, which is useful when

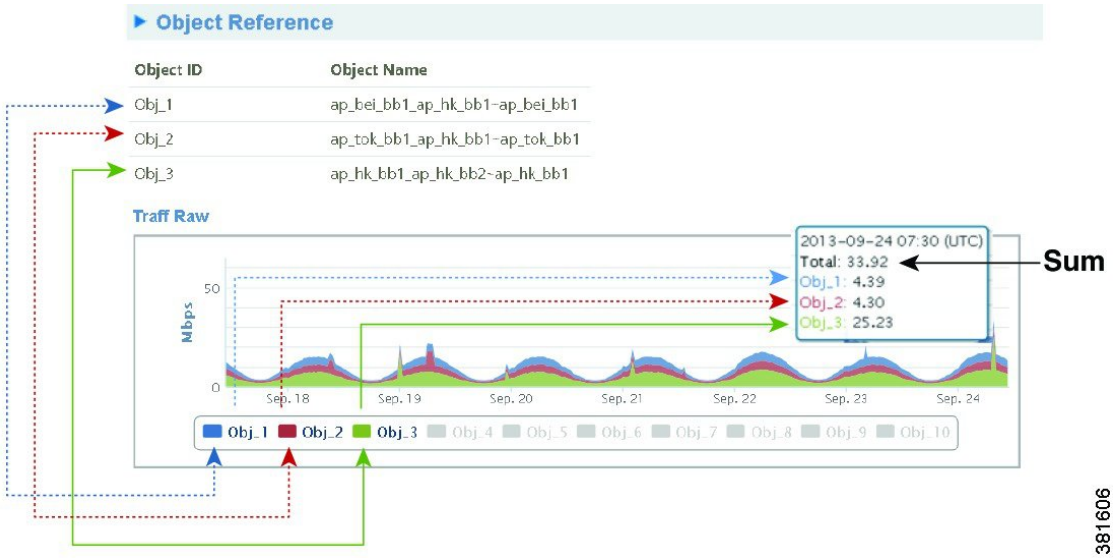
comparing a single property across multiple objects over a given time. For example, you could use stacked graphs to compare how much CPU each node is utilizing at any given time.

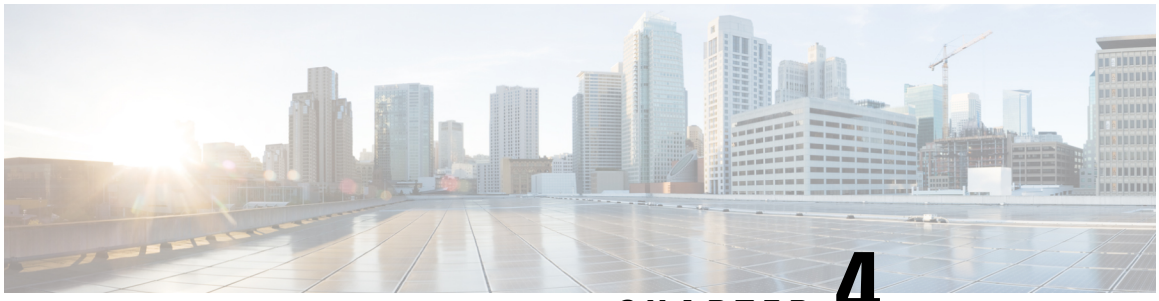
Figure 3: Example of Overlay and Stack Graphs



WAE Live assigns each object an object ID and graph color, and lists this association below each graph, thus identifying which data belongs to which object (Figure 4: Reading Multi-Object Graphs, on page 11). As you move your cursor over the data points, the sum of the property values is listed along with the value of each object. To open or close the Object Reference section, click its arrow.

Figure 4: Reading Multi-Object Graphs





CHAPTER 4

Working with Tables and Filters

Tables in the Explore and report output pages let you view and filter objects based on their current values, as well as graph property values (raw data) over time and navigate to related objects.

Each table lists current information for the selected object type (rows) and properties (columns).

This section contains the following topics:

- Gear Tool Options, on page 13
- Filtering Report Output, on page 14

Gear Tool Options

When working with tables in Explore or in report outputs, you can click the gear tool icon to choose options that operate on report data. These options vary depending on the context and the report.

The **Disable Regex Search** lets you set the filter value for Explore searches. Selecting the check box sets the filter value as a simple string. Else the filter value is set to a regular expression. Regex searches and enabled by default.

Regex searches improves the Explore search response times, especially for complex or malformed regular expressions, or when regular expressions are not intended.

Common Options

The following options are common to all tables:

- Columns—Click to choose which report columns to show in the output. If a column check box is gray, you cannot turn it off. These non-optional properties always appear in the left-most columns of the output.
- Rows per page—Choose the number of objects to show per page in the table.
- Download CSV—Export the table in a tab-delimited format.

Explore Page Options

The following options apply to Explore table options:

- Show Inactive <*Object*> —Show (checked) or hide (unchecked) all objects that have not been detected in the last four hours.

- Show Only Selected—Show only those objects that are selected in the far left column of the table.
- Types—Specify the interface type—such as individual physical interfaces, individual logical interfaces, LAG logical interfaces, LAG physical interfaces, LAG member interfaces—so that you do not show duplicate traffic statistics. This option is applicable only on the Interfaces page.

Note If you choose specific interface types from the Explore Interfaces page for use in a report, the report is only on those interfaces of those selected types. If you choose Interfaces as the Object from a Configure Report page, the report is run on all interfaces regardless of their type.

Report Output Options

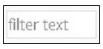
The following options apply to report output table options:

- Edit—Open the Configure Report page to run or configure the report to run again.
- Save As—Name or change the name of the report.
- Permalink—Save the report URL.

Filtering Report Output

Filters are available to narrow report output in most of the tables.

Table 2: Common Filters

Filter	Description
	Filter one or more properties to narrow the list of objects that you are analyzing or reporting on.
	Clears all filters in all columns. To clear one of multiple filters, remove the text from the column in which the filter appears.

You can enter a value in one or more <*filter text*> fields and then press Enter. For non-numeric properties, you can filter text using these methods:

- Entering Partial or Full Strings, on page 14
- Entering Regular Expressions, on page 15
- Filtering Numeric Properties, on page 15

Entering Partial or Full Strings

You can enter partial or full strings as filter text for non-numeric properties.

Example:

If you enter **LON** as the filter in the Interfaces Node column, only interfaces with nodes that include LON in the node name are listed.

Entering Regular Expressions

You can enter regular expressions as filter text in the filter text box or description fields in the tables that support these filtering mechanisms.

Table 3: Using Regular Expression as Filter Text

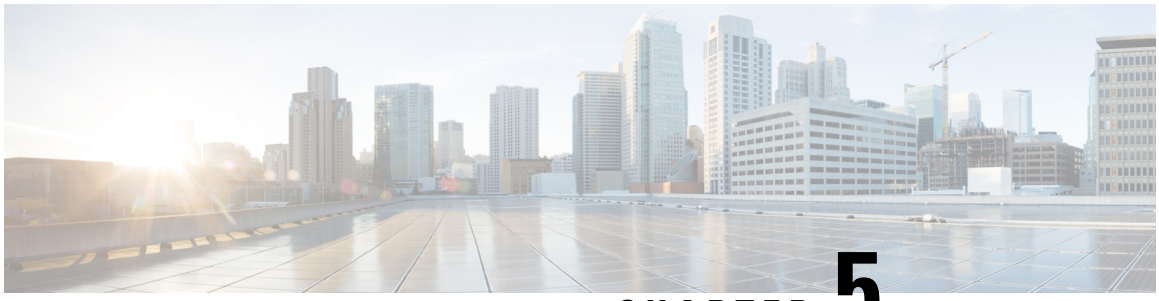
To filter using this criteria	Enter this regular expression
Assume the table descriptions contain "to=" to identify a region into which an interface is going. To search for all descriptions that contain a to= value that starts with "north" (so it will include iterations such as north, northwest, and northeast) or the description contains a to= value of "south" (but not other iterations, such as southwest), enter the regular expression shown in the next column.	**to=north|to=south$**
In the Description field, to find all descriptions that do not contain the words "test," "trial," or "alpha", enter the regular expression shown in the next column.	**!~test|trial|alpha**
If, however, you want to search on descriptions that contain any one of three fields, that is, "test," "trial," or "alpha," enter the regular expression shown in the next column.	**test|trial|alpha**

Filtering Numeric Properties

For numeric properties, you can filter using these methods:

- Complete value
- Operators, such as >, <, and =
- isNotNull or isNull

Example: If you enter **>11894** in the Interfaces Traffic Out column, only interfaces with outbound traffic that is greater than 11894 are listed.



CHAPTER 5

Exploring Network Data

The Explore tool lets you navigate to current network data and retrieve historical object performance data. For example, you can use Explore to:

- Identify which LSPs are currently off their shortest TE paths.
- Determine which LSPs are currently going through a specific interface, or which IP address carries a specific interface and the historical changes on that interface.
- Determine whether IP address or metric changes have occurred on an interface during the last week.

This section contains the following topics:

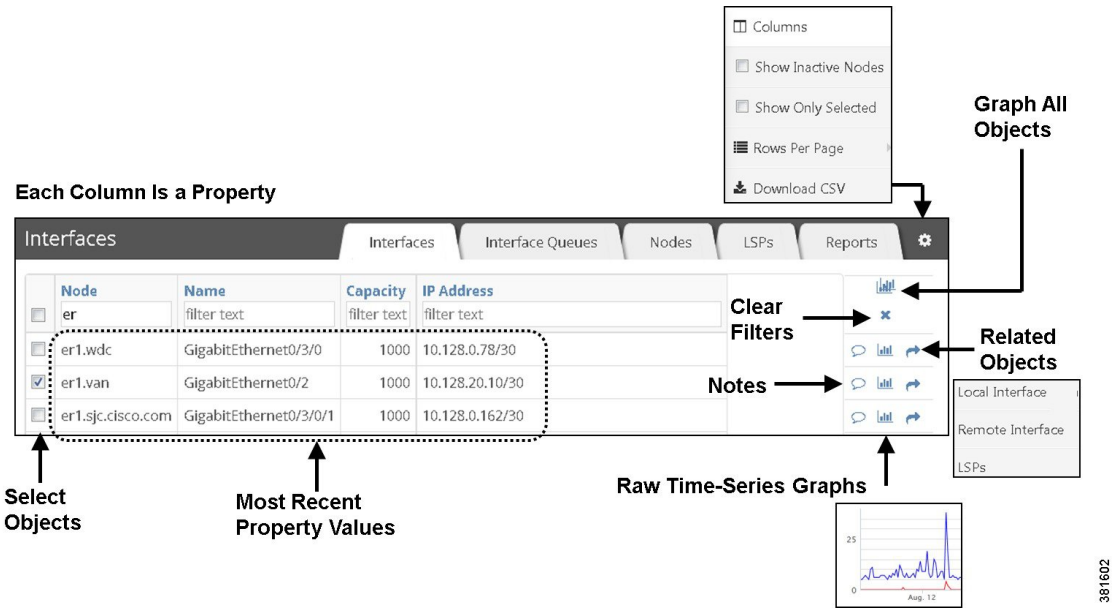
- Navigating with Explore, on page 17
- Working with Tables, on page 18
- Filtering on Interface Descriptions, on page 21
- Exploring LSPs, on page 22
- Reports, on page 23

Navigating with Explore

Each tab in the Explore page represents a table of objects that you can view and filter based on their current values. You can graph those values (raw data) over time and navigate to related objects.

Figure 5: Example Explore Interfaces Table, on page 18 shows the Interfaces tab and its components.

Figure 5: Example Explore Interfaces Table



Working with Tables

In Explore, there are several common tasks that you can perform on table data. This includes the following:

- Choosing Objects to Include in Reports, on page 18
- Clearing Filters, on page 19
- Finding Related Objects, on page 19
- Viewing More Data, on page 20
- Viewing Graphs, on page 21
- Adding Notes to Object Rows, on page 21

Choosing Objects to Include in Reports

In Explore, you can choose the objects that you want to include in reports. You can choose all objects (broad scope) or just a subset of the objects (narrow scope). Each table includes the following check boxes:

- Per-row check boxes—Check the check box next to the desired object. This is a toggle that selects or deselects a single object.

When running reports, these selections are the default filters. For more information on report filters, see Configuring and Running Reports, on page 33.

- Top check box—This is a toggle that selects or deselects all objects of the type you are viewing. For example, you can choose all demands or all nodes.

If you have previously filtered to a subset of objects, this top check box applies only to them.

Clearing Filters

To clear all filters applied on particular objects, click **X** (**Clear All Filters**).



Finding Related Objects

In Explore, you can view objects in one table and then find their related (associated) objects in another table. This lets you jump back and forth between objects. The related object pop-up menu options vary depending on the table you choose and what objects are included in your network model. This can be particularly useful for viewing LSP node hops.

Table 4: Related Objects and Their Locations, on page 19 describes the related objects and where to find them.

Table 4: Related Objects and Their Locations

Table	Related Objects	Where
Interfaces	• Local Node • Remote Node • Remote Interface • LSPs • Individual Logical Interface • Interface to Inventory • Interface Queues	Related Objects Menu
Interface Queues	• Interface	Related Objects Menu

Table	Related Objects	Where
Nodes	• Local Interface • Remote Interface • Node to Inventory • LSPs Source • LSPs Destination • LSPs Actual Path	Related Objects Menu
LSPs	Actual Hops	Actual Path Column
LSPs	• Source Node • Remote Node	Related Objects Menu

Procedure

Step 1 Navigate to the table whose related object you want to find; for example, the LSPs table.

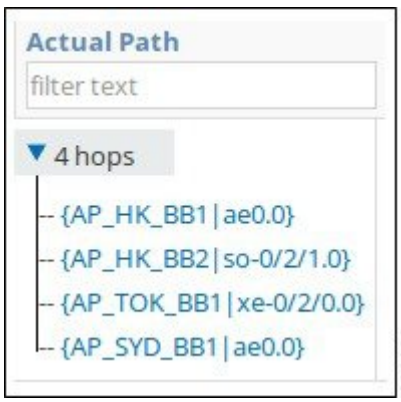
Step 2 In the LSPs table, click the **Related Objects** (arrow).

Step 3 In the Related Objects pop-up menu, choose an object of interest; for example, **Destination Node**. The Explore page changes to list the related information.

Viewing More Data

Table cell information that is blue or cell data that has a small arrow by it can be expanded to show more data. You can view more data in:

LSP Actual Path column—Click the right arrow of an actual path in an LSPs table to see an expanded list of its individual LSP hops. Click any hop to show it in the Nodes table. For an example, see Choosing LSP Hops and Related Nodes, on page 22.



Viewing Graphs

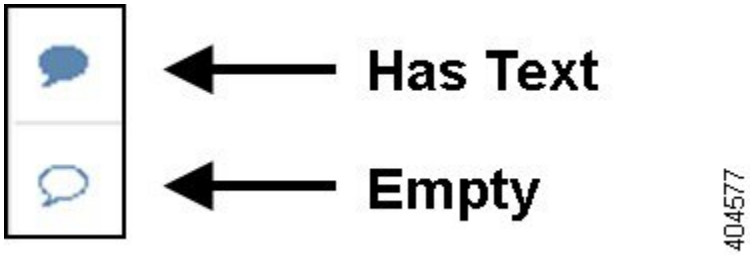
When viewing graphs, you have the following object options:

- Graph all Objects—Graph the data for all objects in the table.
- Graph Single Objects—Graph raw data of an object. From here, you can also list the raw data, which you can then save or send to others.



Adding Notes to Object Rows

You can add notes to the object rows in the Explore tables. All other users can read and add to the notes. This lets multiple WAE Live users share information that others might need to know, such as notes on interfaces that are reaching capacity or notes on when nodes are temporarily down for maintenance.



Notes that have information show as a solid blue to signal there is something to read.

To add a note, read a note, or add content to an existing note, click the **Note** icon.

Filtering on Interface Descriptions

The Interfaces and Interface Queues tabs let you filter text in the Description column on both interface and interface queue object types. This can be useful when your naming scheme for interfaces includes easy-to-filter patterns.

Procedure

Step 1 Choose **WAE Live > Explore** and choose the network (top left) on which you want to explore objects.

If there is only one network configured, the word “default” appears.

Step 2 Click the **Interfaces** or **Interface Queues** tab.

Step 3 In the Description text field, enter the text on which to filter the table rows.

You can use partial or full strings and regular expressions as descriptive text.

After filtering the interface descriptions, you can create reports based on these filtered results on the interfaces or interface queues.

Exploring LSPs

You can explore related LSP objects from the LSPs table as well as drill down on LSP details and then create applicable reports.

Choosing LSP Hops and Related Nodes

From the LSPs tab, you can expand an actual LSP path in the LSPs table to see all of its individual hops in the Actual Path column, Clicking on any hop takes you to the Nodes table where you can see the source node or remote node associated with this LSP.

Procedure

Step 1 Click the **LSPs** tab.

Step 2 In the Actual Path column, click the blue triangle to expand on the available hops.



Step 3 Click the Related Objects (arrow) and choose the source node or remote node that you want to explore for this hop.

This takes you to the applicable source node or remote node for this LSP hop.

Drilling Down on LSP Details

The next procedure assumes that you want to find the three LSPs with the highest setup bandwidth that are currently going through interface ae0.0 on node FRA-BB2 and that exceed 2,000Mbps (2 Gbps) of traffic. It then shows how to find the source node of the LSP with the highest setup bandwidth with these parameters.

Procedure

Step 1 Click the **LSPs** tab.

Step 2 On the LSPs page, enter **FRA-BB2.ae0.0** as the filter in the Actual Path column.

Step 3 In the Traffic column, enter **> 2000** as the filter, and then press **Enter**.



Step 4 Click the **Setup BW** column heading to toggle the list in descending order.

Step 5 Click the **Related Objects** (arrow) of the LSP with the highest setup bandwidth.

Step 6 In the Related Objects pop-up menu, choose **Source Node**.

The resulting table is the source node of the LSP with the highest setup bandwidth that is currently going through interface ae0.0 on node FRA-BB2 and that exceeds 2 Gbps of traffic.

Viewing an LSP Hop Changes Report

The next example report table indicates that in the last LSP Hop Changes report, only 3 out of 36 possible report entries were generated. For the Semi-Annual Node Health report, there were no lines in the report of possible 241 rows once filtering was applied.

Total Lines	Definition ID	Report Name	Run Time	Total Lines After Filter
filter text	filter text	filter text	filter text	filter text
36	44	LSP Hop Changes	13674	3
241	81	Semi-Annual Node Health	7625	0

383329

Reports

The Reports Explore page provides a report summary on the total number of lines in a report and the number of lines when run against a user-specified filter. Explore records only named reports and the most recently run reports.

This high-level report trend can be an indicator of the network's health. For example, if you run a node health report every week and filtered the percentage of time CPUs were over a given threshold, this table could show the number of nodes surpassing this threshold improving or worsening over time. If the number of nodes passing the threshold was increasing, it could be an indicator that the network is running hotter and you need to buy more nodes.

Choose the objects on which you want to run a report, and then click **Run Report**. The Configure Report page opens. For more information, see Configuring and Running Reports, on page 33.

Gear Tool Options

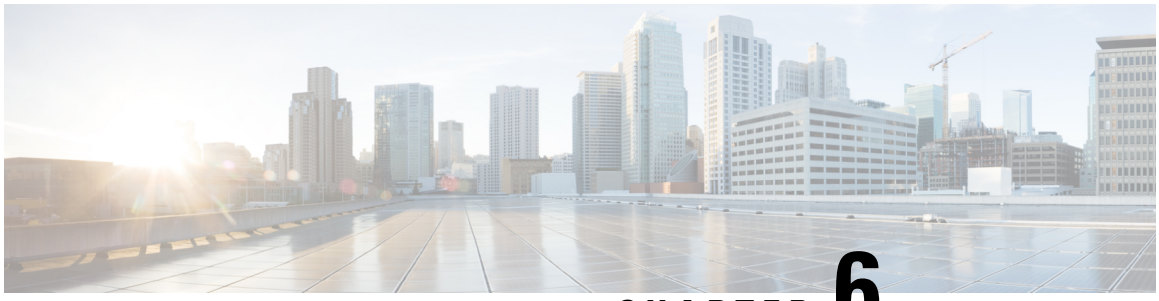
The same set of gear tool options are available across all Explore pages with the exception of the Interface Type. To choose any of the following Explore table options, click the gear tool.

Table 5: Explore Gear Tool Options

Item		Description
Columns		Click to choose which properties appear as column headings. All available properties are listed. If a column (property) check box is gray, then you cannot turn it off. These non-optional properties always appear in the left-most columns of the table.
Types		Specify the interface type so that you do not show duplicate traffic statistics. This option is applicable only on the Interfaces page.
	Individual Physical Interface	An interface that is not in a LAG and does not have an IP address. These interface names do not have a decimal. Example: so-0/2/2, g2/0/0
	Individual logical interface	An interface with an IP address that is not in a LAG. These names might have a decimal. Example: et-0/1/2.0, g2/0/0
	LAG logical interface	An interface with an IP address that contains member interfaces. Examples: ae0.0, Bundle-Ether1, Port-Channel7
	LAG physical interface	The physical representation of an aggregated interface to which a LAG logical interface (such as ae0.0) is assigned. Example: ae0 for Juniper interfaces
	LAG member interface	An interface that is in a LAG. Example: g2/0/0
	Unknown	None of the above interface types.
Disable Regex Search		Select the check box to set the filter value as a simple string. Else the filter value is set to a regular expression.
Show Inactive *<Object>*		Show or hide all objects that have not been detected in the last four hours.
Show Only Selected		Show only those objects that are selected in the far left column of the table.
Rows per page		Choose the number of objects per page to show in the table.
Download CSV		Export the table in a tab-delimited format.



Note If you choose specific interface types from the Explore Interfaces page for use in a report, the report is only on those interfaces of those selected types. If you choose Interfaces as the Object from a Configure Report page, the report is run on all interfaces regardless of their type.



CHAPTER 6

Viewing Inventory

The Inventory pages show inventory data in different views to help you find answers about your hardware inventory.

This section contains the following topics:

- Accessing Inventory Data, on page 27
- Viewing the Inventory Tree, on page 27
- Viewing the Inventory Table, on page 29
- Viewing Inventory Availability, on page 29
- Viewing Hardware Inventory Summary, on page 30
- Columns, on page 30

Accessing Inventory Data

Procedure

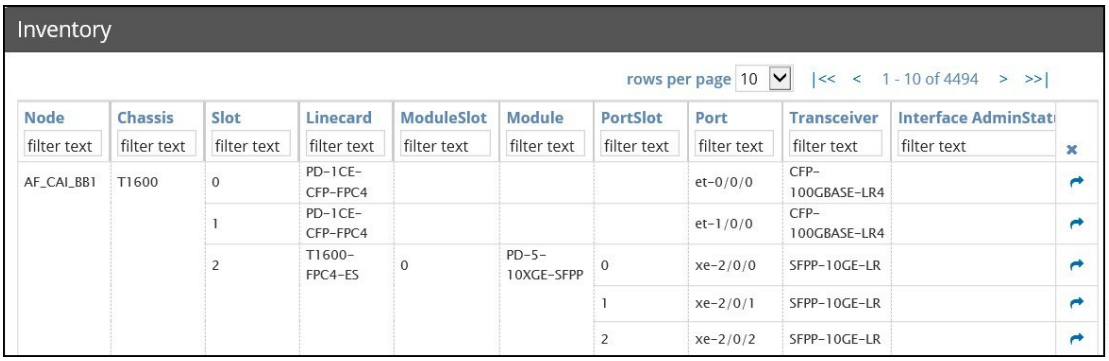
Step 1 Choose **WAE Live > Inventory** and choose the network (top left) on which you want to create the report.

If there is only one network configured, the word “default” appears.

Step 2 On the **Inventory** page, click the Inventory view you want to see; for example, **Tree**. You can also choose Table, Availability, and Summary views.

Viewing the Inventory Tree

The Inventory Tree page provides a hierarchical representation of the current hardware in your network. It also provides information on how the hardware is interconnected within the chassis. This filterable table lets you search for these interconnections. For example, you can find a specific chassis by filtering on its name or you can navigate to connected interfaces.



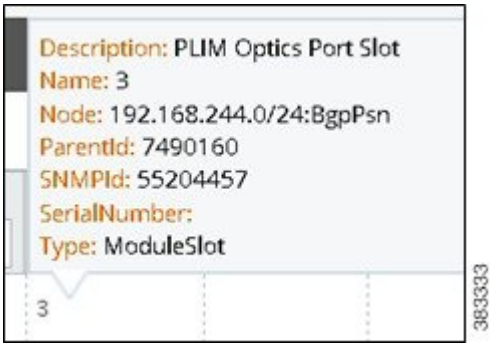
Finding Available Modules or Ports

One use case is to determine if there are any empty slots available to accommodate future growth. For example, if you need to add five10-GB interfaces on a specific router (node), you can determine how many empty module and port slots you have. This can help you plan for future hardware expenses.

To determine if there are available modules or ports, enter “empty” as the column’s filter. If there are no results, there are none available. Returned results indicate which slots are available.

Finding Hardware Details

To find detailed information about any piece of hardware, hover over the top of it. Note that the Parent ID is the SNMP ID of a physical child asset. For example, a slot’s Parent ID is the SNMP ID of the chassis to which it belongs.



To see the Explore data for an interface or a node associated with a port, or to see more details on the node itself, use the Related Info icon (arrow on the far right).

Adding Notes

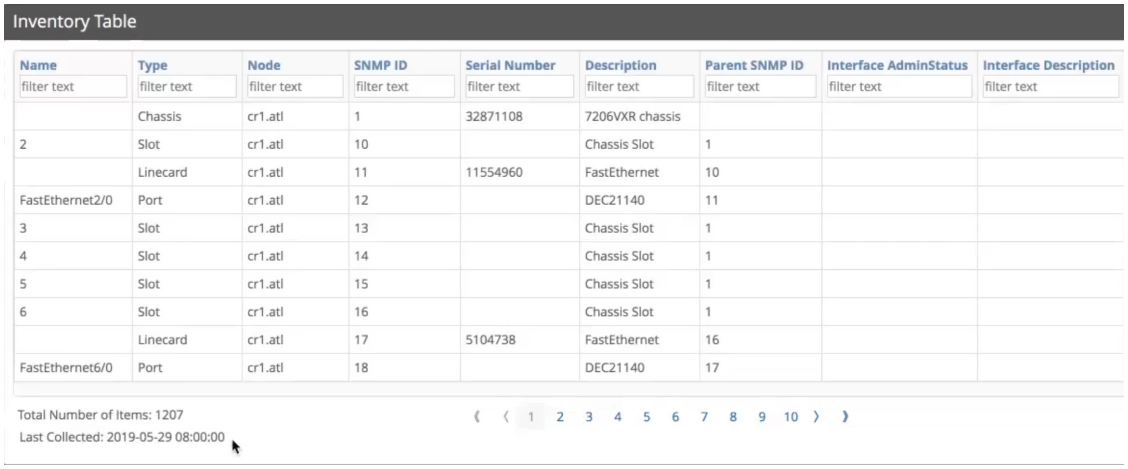
When the **Port** and the **Interface Admin Status** columns are populated, you can add notes by clicking the **[Add]** link in the Notes column. For example, you might add a note that a port is reserved for a specific project or that it is faulty.

All users can see the notes, read the history of the notes, and edit the notes. To edit or read the history, click the note itself to open a dialog box with Edit and Show History links.

Notes cannot be deleted.

Viewing the Inventory Table

The Inventory Table page lists all physical elements in the network. While you can filter on any column to find network elements, the predominate use case is to find serial numbers by using the filtering capabilities in the **Serial Number** field.


Inventory Table

Name	Type	Node	SNMP ID	Serial Number	Description	Parent SNMP ID	Interface AdminStatus	Interface Description
filter text	filter text	filter text	filter text	filter text	filter text	filter text	filter text	filter text
	Chassis	cr1.atl	1	32871108	7206VXR chassis			
2	Slot	cr1.atl	10		Chassis Slot	1		
	Linecard	cr1.atl	11	11554960	FastEthernet	10		
FastEthernet2/0	Port	cr1.atl	12		DEC21140	11		
3	Slot	cr1.atl	13		Chassis Slot	1		
4	Slot	cr1.atl	14		Chassis Slot	1		
5	Slot	cr1.atl	15		Chassis Slot	1		
6	Slot	cr1.atl	16		Chassis Slot	1		
	Linecard	cr1.atl	17	5104738	FastEthernet	16		
FastEthernet6/0	Port	cr1.atl	18		DEC21140	17		

Total Number of Items: 1207
Last Collected: 2019-05-29 08:00:00

« ‹ 1 2 3 4 5 6 7 8 9 10 › »

Last Collected field at the bottom of the table displays the timestamp of the data collected data.

Viewing Inventory Availability

The Inventory Availability page lets you determine how many slots are available per node (router), per line card, and per module. One use case is to identify how to consolidate empty slots and empty ports.

Note Each hardware element that can host is considered to be a “slot.” If a chassis or line card does not have available slots, of if a module does not have available ports, it is not listed in their respective tables.

If a line card is modular, it represents the number of slots it can hold, not ports.

The Modules table shows the number of ports, available ports and ports that are down. A port is considered down if its Operational Status is down.


Chassis

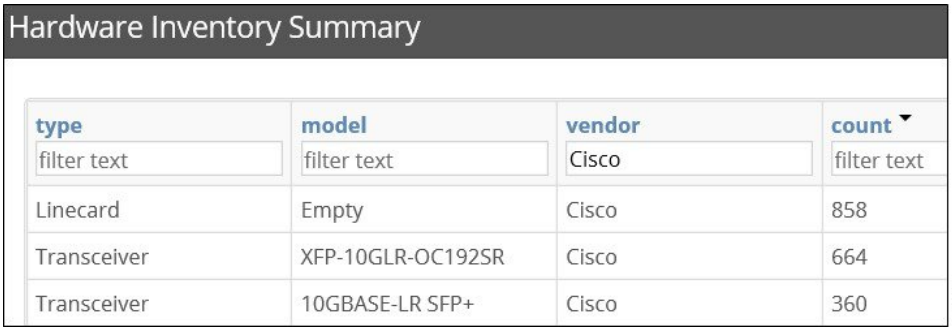
node	model	vendor	serialnumber	number_slots	available_slots
CHI	filter text	filter text	filter text	filter text	filter text
AM_CHI_BB1	CRS-16/S	Cisco	OQAYRLV6	16	10
AM_CHI_ER1	ASR-9010-AC-E	Cisco	769YCK1X	8	7
AM_CHI_ER3	ASR-9010-AC-E	Cisco	SWVEEM3A	8	7

406017

Viewing Hardware Inventory Summary

The Inventory Summary page gives the count of hardware elements in the network on a per-element basis. One use case is if a vendor announces that a line card is being deprecated, you can use this table to determine how many line cards you need to replace. You can understand the impact by filtering to the model in the Availability page to determine how many slots are available, thus indicating where you need to move traffic.

Combined with information on the cost of the hardware, this Summary table helps you estimate the total cost of the network.

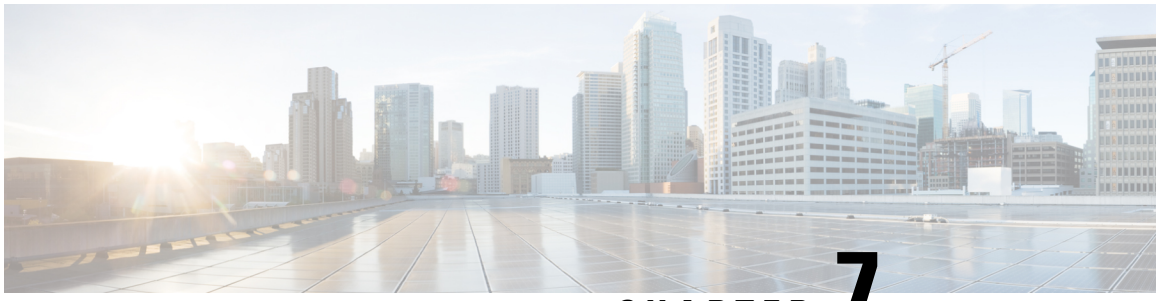


Columns

- Available Slots—Number of slots that are not being used.
- Chassis—A descriptive or exact name of the router type, such as ASR9000.
- Count—Number of network elements of given type.
- Interface Admin Status—Administrative status of the interface connected to the port. Values are Up (active) or Down (inactive).
- Interface Description—Optional description of the interface connected to the port.
- Linecard—A printed circuit board that contains ports or modules (which in turn, contain ports). A line card plugs into a chassis slot.
- Model—A particular unit of hardware that identifies its style or type.
- Module—A self-contained hardware element that contains ports and plugs into a chassis or a line card slot. An example module name is 10-port 100-Gbps LSR. A module is also sometimes called a *PIC* .
- Module Slot—A slot in a line card for holding modules.
- Name—The name of the hardware identified in the Type column. If applicable, such as for chassis and line cards, this is the model.
- Node—Router name, such as 168.243.0/24:BgpPsn.
- Number Slots—Total number of slots an element contains (both used and unused).
- Parent SNMP ID—Hardware has a parent-child relationship based on where the object resides within the router. The chassis has no parent and is considered the root object. Other than the chassis, each object

has one parent and can have one or more child objects. The Parent SNMP ID is the identifier of the parent.

- Port—A socket that accepts a plug. For example, cables are plugged into Ethernet ports. An example port name is GigabitEhternet0/1/0/7.
- Port Slot—A slot in a module for holding physical interface ports.
- Serial Number—A unique series of numbers and letters assigned by the vendor to each hardware element. An example serial number is X8XP16J2.
- Slot—Each hardware element that can host is considered to contain slots. Chassis have slots that contain line cards. Line cards have slots that contain ports or modules. Modules have slots that contain ports.
- SNMP ID—ID of the hardware element as described by the router.
- Transceiver—An electronic device or circuit that transmits and receives analog or digital signals. Transceivers are connected to ports. An example name is DWDM Xenpaks.
- Type—The hardware element category. For example, chassis and module are both types of hardware.
- Vendor—The company that sells the node (router).



CHAPTER 7

Configuring and Running Reports

This chapter describes the analytics report use cases, an example report creation workflow, and how to configure and run WAE Live Analytics reports.

This section contains the following topics:

- Analytics Use Cases, on page 33
- Reports Workflow, on page 33
- Configuring Reports, on page 34
- Choosing a Filter Type, on page 36
- Defining Time Ranges, on page 37
- Ad Hoc and Traffic Report Specific Options, on page 39
- Running Reports, on page 42
- Scheduling Reports, on page 44
- Reports Table Details, on page 45
- Using the Report Log, on page 45

Analytics Use Cases

The Analytics pages let you generate traffic, health, ad hoc, and deviation reports, and visualize their results. These reports let you answer the following questions:

- Which objects do you want to report on?
- What object properties do you want report on?
- What is the time frame for the report? Do you need to project a trend farther out from that time frame?
- Do you want to run a report on individual or grouped objects?
- Do you want to apply time aggregation operations, and if so, what are they (for example, maximum or average)?

Reports Workflow

There are several report configuration options that you can choose from when creating analytics reports for the first time. Configuration options vary depending on the report type.

The next procedure outlines an example ad hoc report configuration process.

Procedure

Step 1 Choose **WAE Live** > **Analytics** and choose the network (top left) on which you want to create the report.

If there is only one network configured, the word “default” appears.

Step 2 On the **Analytics** page, click **New Report**.

Step 3 In the **Configure Report** dialog box, choose the report type (such as **Ad hoc**) and the object type (such as **LSPs**) from the drop-down menus.

The Summary page summarizes your selections so far and presents you with several additional tabs: Filter, Time Range, Group By, Time Interval, and Data Columns.

Step 4 Click the **Filter** tab and the filter type to use. By default, WAE Live assumes you want to choose relevant objects in the Explore view.

To change the default filter, see Choosing a Filter Type, on page 36.

Step 5 Click the **Time Range** tab and specify a time range for the report to run. By default, WAE Live uses the last one week of data when running the ad hoc report.

To change the default time range, see Defining Time Ranges, on page 37.

Step 6 Click the **Group By** tab, and then choose a group or aggregation option. By default, WAE Live uses None (no objects).

To change the grouping or aggregation options, see Grouping Objects, on page 39 and Grouping and Aggregating Objects, on page 41 for more information.

Step 7 Click the **Time Interval** tab and specify the time period over which attributes are aggregated; for example, 1 day. For more information, see Time Aggregation Options, on page 66.

Step 8 Click the **Data Columns** tab and choose the property and time aggregation operation to apply time aggregation combinations. For more information, see Ad Hoc Report Tabs, on page 66.

Step 9 To run the report just once, click **Run**. For more information, see Running the Report without a Name, on page 43.

Step 10 To save and run the report, click **Save & Run**. For more information see, Running the Report with a Saved Name, on page 43.

Step 11 Name the report. For more information, see Naming Reports, on page 43.

After naming and saving a report, it becomes available from the Reports page.

Configuring Reports

You can configure reports on a per-network basis to view network elements of interest and then group these elements to tailor a report for your particular use case.

To create an analytics report, you must configure the following report options:

- Type of report to run from the Report Type drop-down menu

- Type of object to report on from the Object drop-down menu
- Type of filter to use for filtering objects from the Filter tab

You can configure a report to be run just once (ad hoc), or you can save the report and run it again later.

Report Types, Object Types, and Configuration Options

Table 6: Report Types, Object Types, and Report-Specific Configuration Options , on page 35 describes the WAE Live report and object types that you can configure for report creation.

Table 6: Report Types, Object Types, and Report-Specific Configuration Options

Report Type	Definition	Available Object Types	Report-Specific Configuration Options
Ad hoc	A custom report that can be run on any data column object type found in the Explore table. For more information on setting the group by and aggregate by options, see Grouping Objects, on page 39, Grouping and Aggregating Objects, on page 41, and Time Aggregation, on page 42. For more information on setting ad hoc time intervals ad hoc report examples, see Running Ad Hoc Reports, on page 63.	• Interfaces • Interface Queues • Nodes • LSPs • Demands • Reports	• Group By • Aggregate By • Data Columns
Health	A report to determine the health of network objects and whether LAGs are imbalanced. For more information on setting thresholds and health report examples, see Running Health Reports, on page 53.	• Interfaces • LSPs • LAG Imbalance • Nodes	• Utilization Threshold • SetupBW Threshold • LAG Imbalance • CPU Threshold, Memory Threshold
Traffic	A report to determine current and projected traffic trends. For more information on grouping, aggregate, and projection options, and traffic report examples, see Running Traffic Reports, on page 47.	• Interfaces • Demands • LSPs	• Group By, Trends, Projection • Trends, Projection • Trends, Projection
Deviation	A report to determine sudden shifts in property values. For more information on setting parameters and example deviation reports, see Running Deviation Reports, on page 59.	• Interfaces • Interface Queues • Nodes • LSPs • Demands • Reports	• Parameters

Choosing a Filter Type

After deciding the type of report to create and the objects to report on, you must choose a filter type to use when creating your analytics report.

The filter you choose operates on a subset of objects. The complete list of objects and properties that you can filter on is described in Understanding Objects, Properties, and Data, on page 3. There are two filter types:

- Selected objects in the Explore View—this filter is selected by default. This filter is available from the **Configure Report > Filter** tab as well as from the various Explore tables. If there are no objects selected, the Filter page displays (0 Objects).
- Latest value of—this filter represents the latest value of an object. Each object type has different filter values depending on the object type selected. Filter values can include user-defined properties set up by the WAE Live administrator, if any.

  You can specify the filter to use in creating a report from the drop-down list on the Filter tab, or you can base them on selections in the Explore tables.

  Both filters let you have a broad scope report or a more narrow scope report depending on what data you are looking for in the report.

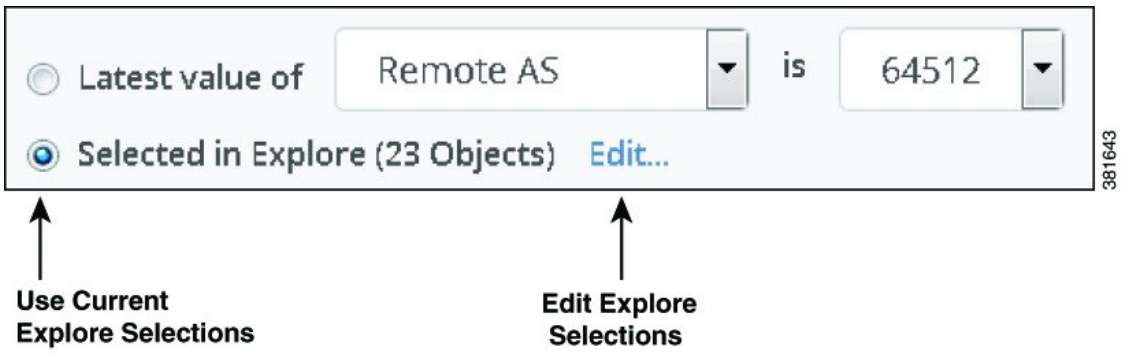
Note If you choose specific interface types from the Explore Interfaces page for use in a report, WAE Live generates the report only on those interfaces of those selected types. If you choose Interfaces as the Object from a Configure Report page, the report is run on all interfaces regardless of their type.

Choosing Objects in Explore

This procedure assumes that you want to narrow your current report scope from the 23 selected LSPs in the Explore LSPs table to just three LSPs.

Procedure

Step 1 In the **Configure Report > Filter** tab, click **Edit**, which takes you to the relevant Explore page.



Step 2 In the LSPs table, choose the three LSPs on which you want to report.

Choosing the Latest Value Filter

When you create reports from the **Analytics > New Report > Configure Report > Filter** tab, the default filter is **Selected in Explore**. You can change the default by clicking the **Latest value of** radio button.

Note The object type properties and value filters that you see in the drop-down menus vary depending on the object you choose.

This procedure shows how to use the **Latest value of** filter and refine a remote autonomous system (Remote AS) to a specific value for reporting purposes.

Procedure

Step 1 In the **Filter** tab, choose the object property to filter on; for example, **Remote AS**.

Step 2 Refine the **Latest value of** list by choosing a specific value from the second list; for example, **64512**.

This lets you run the report on all objects based on the selected property or on all objects matching the refined property value.



Defining Time Ranges

The Time Range tab defines the time frame for when WAE Live generates the report. You can choose a time range using one of these methods:

- Creating a Time Range with the Last Button, on page 38
- Creating a Custom Date Range, on page 38
- Creating a Date Range with the Previous Time Window Button, on page 38

Report time ranges use the time zone that you specify in your user profile.

Creating a Time Range with the Last Button

You can specify that reports use the time range radio button and the time from the Last pull-down menu. For example, you might create a deviation report with a time range that includes all the data collected within the last year.

Procedure

Step 1 In the Configure Report window, click the **Time Range** tab.

Step 2 Click the **Last Radio** button.

Step 3 Enter an integer such as **1** and choose **year(s)** from the drop-down list. You can also choose hour(s), day(s), week(s), and month(s).

Creating a Custom Date Range

You can specify a custom date range for when WAE Live runs reports. You can enter the range manually in *yyyy-mm-dd hh:mm:ss* format, or you can choose dates from the calendar pickers.

Procedure

Step 1 In the Configure Report window, click the **Time Range** tab.

Step 2 Click the **Custom Dates** radio button.

Step 3 In the From field, enter the from date for the time range.

Step 4 In the To field, enter the to date for the time range.

Creating a Date Range with the Previous Time Window Button

You can configure reports with a floating window (**Previous Time Window**). This option reports data based on the time frame when the report was run (the last full window)—such as 7:00 a.m.–7:00 p.m.—instead of the time when the report was defined, such as 6:00 p.m.

The time range for the floating window uses the 24-hour clock and can be defined between 00:00:00 and 23:59:59. You can enter the time range in *hh:mm:ss* format manually (00:00:00), or by choosing the appropriate times from the time pickers.

The time range can transition through two days; for example, you can specify the time window to be From: 22:00:00 (10:00 p.m.) and To: 10:00:00 (10:00 a.m.).

Table 7: Report Window Example

Report Window	Report Run Time	Result
7:00 a.m–7:00 p.m.	6:00 a.m. on January 2	The report covers data from 7:00 a.m.–7:00 p.m. on January 1.
	6:00 p.m. on January 2	The report covers data from 7:00 a.m.–7:00 p.m. on January 1.
	8:00 p.m. on January 2	The report covers data from 7:00 a.m.–7:00 p.m. on January 2.

This procedure shows how to a create report that defines a time window to include data collected from 7:00 a.m. to 7:00 p.m.

Procedure

Step 1 In the Configure Report window, click the **Time Range** tab.

Step 2 Click the **Previous Time Window** radio button.

Step 3 In the From field, enter the from time; for example, 07:00:00 (7:00 a.m.).

Step 4 In the To field, enter the to time; for example,19:00:00 (7:00 p.m.).

Ad Hoc and Traffic Report Specific Options

The ad hoc and traffic reports support the following additional configuration options:

- Grouping Objects
- Time Aggregation

Grouping Objects

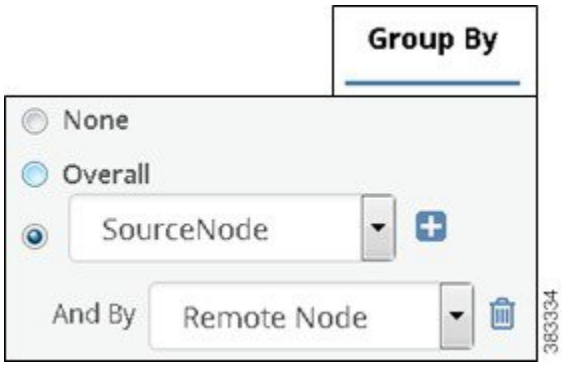
The traffic and ad hoc reports support the ability to group (or aggregate) a set of objects and sum their property values. In the resulting report table, each row represents a group of objects. The columns represent the aggregated values for each group.

The next procedure shows how to create a traffic report that operates on LSPs and is grouped by source nodes and remote nodes.

Procedure

Step 1 In the Configure Report window, click the **Group By** tab.

Step 2 Click the **Select** radio button and choose **Source Node**.



Step 3 Click the **plus button (+)** and choose **Remote Node**.

Group By Options

Table 8: Group By Options, on page 40 describes the Group By tab options and how they affect the report generation process.

Table 8: Group By Options

Group By Option	Definition
None	Does not to group any objects when generating the report. The None radio is selected by default.
Overall	Groups all objects in the report into one row of data. For example, you can group all interfaces to determine their overall traffic. The data in that row is the sum of all interfaces combined.
Group by one property	Group all objects that have the same value for the selected property together. The resulting report shows one row for each grouped property and the data in that row is the aggregate for that property. For example, if you choose LSPs as the object and then group them by Source Node, the resulting report includes one row for each source node. The data in that row is for all LSPs that are sourced from that node.
Group by multiple properties	Further segment a group of objects based on another property selected after clicking the Plus icon. The resulting report shows one row for each set of grouped properties within its parent group. For example, if you group LSPs by their source node and secondly by their remote (destination) node, the report aggregates data for all remote nodes that have the same source node. In effect, this example grouping would create a traffic matrix.

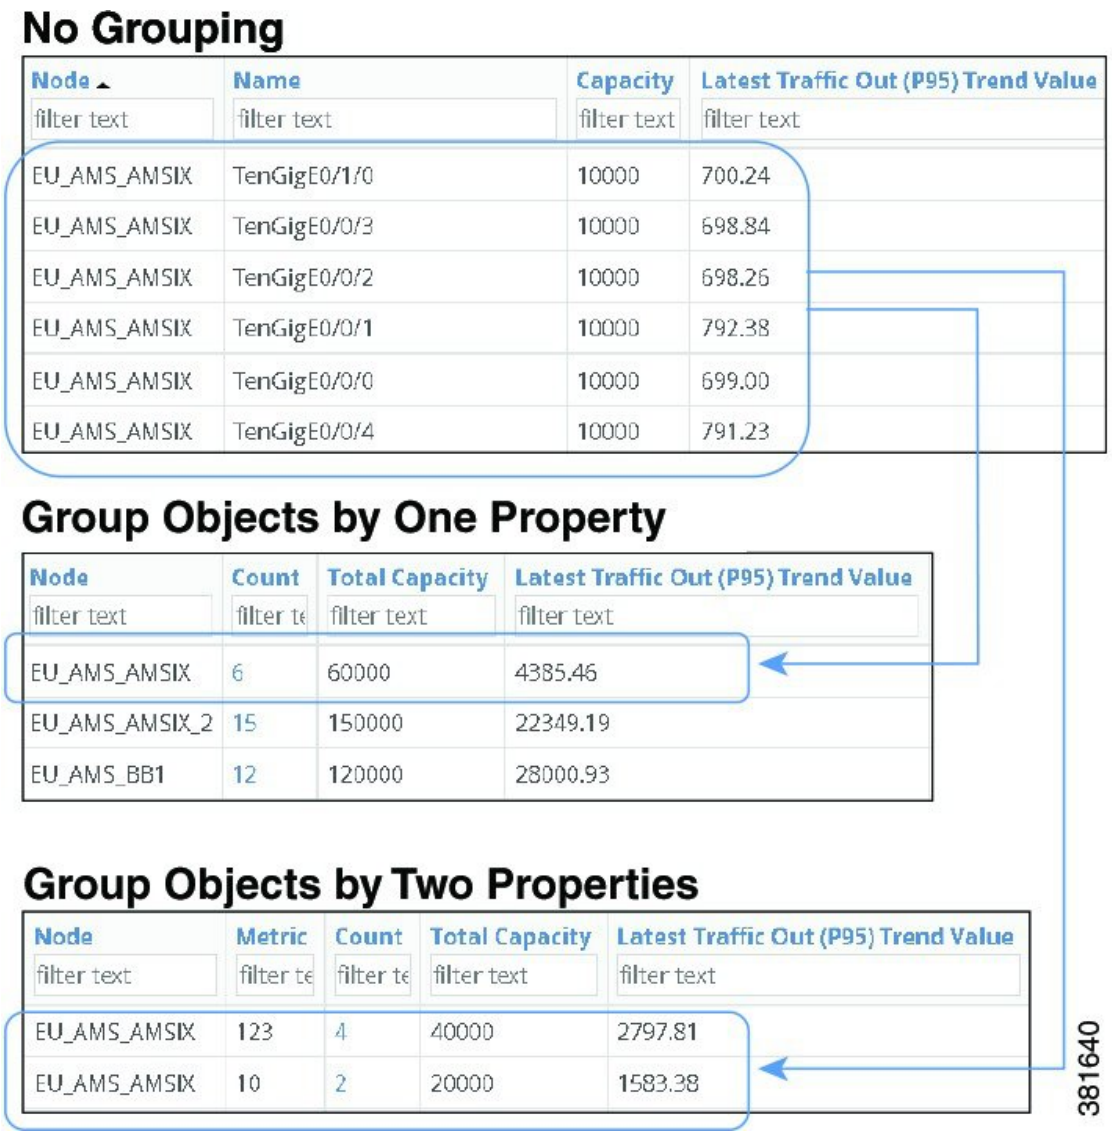
Grouping Examples

Figure 6: Example Effects of Grouping Objects, on page 41 shows a report generated three times:

- One without grouping objects (No Grouping)
- One with objects grouped by nodes (Group Objects by One Property)
- One grouping first by nodes and then by metric (Group Objects by Two Properties)

The Capacity and the Latest Traffic Out (P95) columns adjust according to whether and how the objects are grouped.

Figure 6: Example Effects of Grouping Objects



Grouping and Aggregating Objects

Grouping is the process of choosing what objects to sum together as a group; for example, all interfaces with a specific capacity. Aggregation is a function where multiple values of different objects are grouped together to form a single value; for example, “maximum” or “average.” Note that “average” represents an additional grouping.

Aggregation Options

Table 9: Aggregation Option Values and Their Meanings , on page 42 details the aggregation options and their meanings.

Table 9: Aggregation Option Values and Their Meanings

Aggregation Option	Meaning
Average	Add up all of the object values and find the average value.
Maximum	Look at all of the object values and find the maximum value.
Minimum	Look at all of the object values find the minimum value.
Last	Find the last value in the specified time interval.
Count	Number of times the data was collected during the specified time period. For example, if you collected traffic information once an hour for 24 hours, and the Count was 20, you would know there was a problem during 4 hours.
Change Count	Number of times the value changed. For example, if you choose the Shortest TE Path for LSPs, you can report on the number of times it changed during the last year.
P# Values (P99, P98, P95, P90, and P85)	Find time aggregation computations based on these percentage values. For example, P95 is the value 95% of the time. Many times the top 5% in the chart view is considered as noise, which can be filtered out.

Time Aggregation

The traffic and ad hoc reports support time aggregation features that aggregate multiple property values into a single value for a given time range. For example, you can aggregate maximum raw data values for traffic on a weekly basis.

Time aggregation operations are applied to raw data. For a description of raw data, see Understanding Objects, Properties, and Data, on page 3. For more information and examples of aggregating data in the reports, see Running Traffic Reports, on page 47 and Running Ad Hoc Reports, on page 63.

Running Reports

You can run new reports or existing reports from any of the locations listed in the next sections.

Running New Reports

To run new reports:

- From the **Analytics** > **Reports** page or **Analytics** > **Report Log** page, click **New Report**.
- From the **Explore** page, choose objects and click **Run Report**.

Running Existing Reports

There are two ways to run existing reports (with existing or changed parameters):

- From the **Analytics > Reports** page, choose **View > Run**, which runs immediately with the same parameters. To keep or change the parameters, choose **View > Run** or click the name in the Report Name cell.
- From the **Analytics > Report Log** page, to keep or change the parameters, choose **View > Edit** or click the name in the Report Name cell.

Canceling Running Reports

To cancel reports that are running:

- From the **Analytics > Report Log** page, select the report and click **Cancel** to cancel.

Naming Reports

You can run reports with a saved name or without a name. Additionally, you can change the name of an existing report.

To email the report, you must give the report a unique name when you run it.

Running the Report with a Saved Name

To run the report with a saved name, click the **Save & Run** button after configuring the report. By giving the report a name, you automatically add it to the Analytics Reports page where you can easily identify and schedule it for future use.

Running the Report without a Name

To run the report without giving it a name, click the **Run** button after configuring the report. If you do not save the report with a name, it is given a default name with the format Report Type-Object Type-Definition ID. Reports that use a default name appear only on the Analytics Report Log page.

Changing the Name of an Existing Report

To change the name of an existing report, there are two options:

- From the report output page, choose **Save As** from the menu available through the gear icon.
- Run and save the report again with a new name. This overwrites the existing report name.

Applying Filters and Sorting Named Reports

You can apply filters and sorting to a named report so that it displays according to these filters and sorting each time you open or run the report. Thereafter, you can schedule email alerts to be sent when these conditions are met.

These saved report parameters do not affect the generation of the report. Rather they affect what is displayed in the output (based on the filters that were applied), how the information is displayed (based on how the columns were sorted), and if scheduled appropriately, whether the report is sent via email.

Procedure

Step 1 Run and name the report.

Step 2 On the report output page, apply filters and/or sorting.

- Enter a value in one or more **Filter Text** cells at the top of a table column, and press Enter.
- To sort the results of a property in a specific order, click its column heading one or two times.

Step 3 Save the filters and sorting by clicking the **Save** icon.

Scheduling Reports

You can schedule reports to run once or at regular intervals, and have these results emailed to you and to others. You can also send alerts when the state of your network matches conditions that you define using filters. For example, if want to be notified if interfaces exceed 90% utilization, you can apply and save report filters to meet this criteria, and then schedule an email alert to be sent if this condition is met.

Procedure

Step 1 Choose **WAE Live** > **Analytics** and choose the network on which you want to create the report.

If there is only one network configured, the word “default” appears.

Step 2 If you want to email report results based on specific criteria, set the filters before opening this dialog box. Similarly, if you want emailed reports returned with a specific column sorted, save the report with this column sorted. For information, see Configuring and Running Reports, on page 33.

Step 3 From the Reports page, choose **View** > **Scheduling** or click the Scheduled link associated with the report you want to schedule.

Step 4 To turn on scheduling, choose the options that fit your needs. You can schedule the time of day for the report to run, and you can run reports once on a specific day, daily, weekly, monthly, or on a specified number of weeks or months.

To turn off scheduling, choose **None**.

Step 5 To email scheduled reports, enter one or more email addresses separated by commas.

Step 6 If emailing reports, either specify that they always be sent or only if they contain data (they are non-empty). If you are emailing the report based only on specified criteria (Step 1), choose “non-empty report.”

Step 7 Click **Save**.

Reports Table Details

The Reports table provides the following details and options for your reports:

- Report Name—Name of the report. Click the name to open the Configure Report page with the parameters selected that were used to create the report.
- Time Range—Duration of the report; for example, last month, last year, and so on.
- Last Run—Date and time of the last time the report was run. If the report has been run multiple times, click the icon next to the date to list the available reports. To run one of these reports, choose it and then click **Select**.
- Scheduled—Frequency of how often the report is automatically run. If the report is not scheduled, the column reads None. Click View > Scheduling to schedule the report to run, to email the report, or to stop running the report if already scheduled.
- Owner—Username of the person who created the report. If LDAP is the authentication method, this is the same as the LDAP client username. By filtering on your username, you can view only the reports that you created.
- View options:
  - View—Click to view the report output.
  - View > Run—Run the report with same parameters.
  - View > Edit—Open the Configure Report page with the parameters selected that were used to create the report. If you run and rename the report, it replaces the name of the existing report.
  - View > Scheduling—Open the page that lets you schedule the report to run, to email the report, or to stop running the report if already scheduled.
  - View > Delete—Remove the report from this Reports page. The report remains listed on the Report Log page.

Using the Report Log

The Analytics Report Log lists all previously run reports and information that further defines them, such as their start and end time. You can view report results, run the report again, or edit the report configuration before rerunning.

To view the Report Log:

Procedure

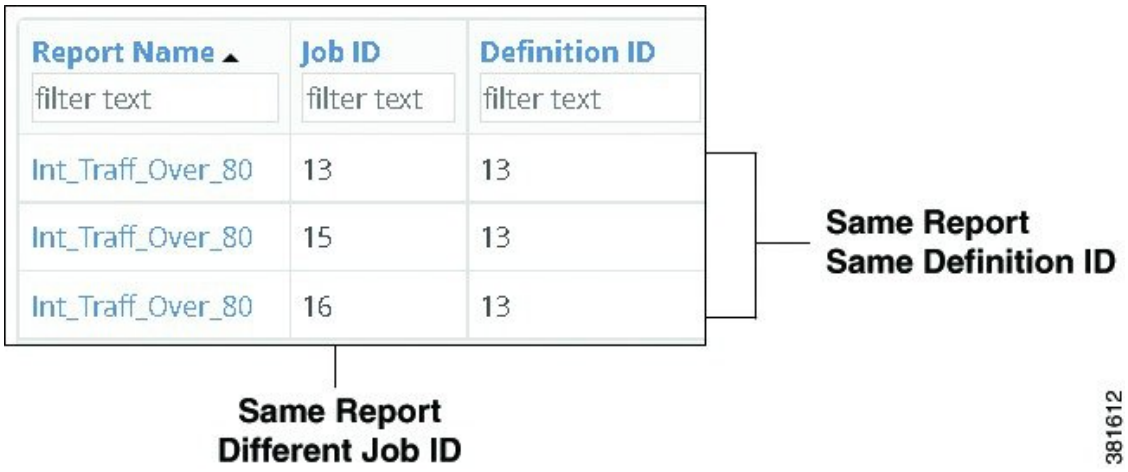
Step 1 Choose **WAE Live > Analytics** and choose the network (top left) on which you want to create the report.

If there is only one network configured, the word “default” appears.

Step 2 On the Analytics page, click **Show Report Log**.

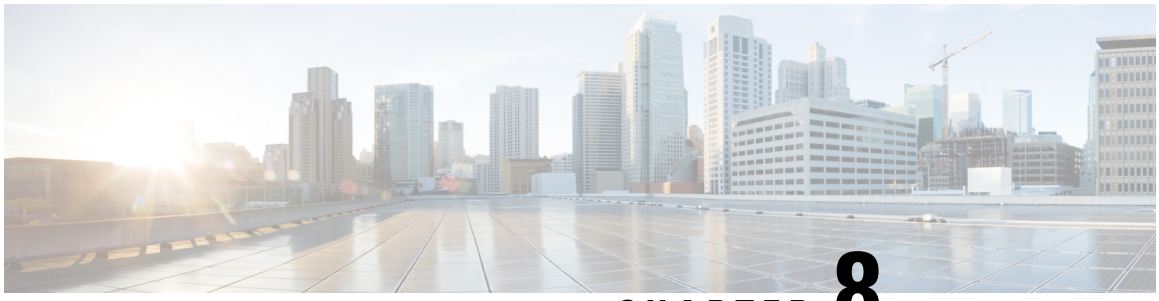
The Report Log provides the following details about your reports:

- Report Name—Name of the report. Click it to open the Configure Report page with the parameters selected that were used to create the report.
- Job ID—Specific instance of a report. A unique Job ID is generated each time a report is run. The Job ID is required information for importing traffic reports into the WAE Design GUI.
- Definition ID—Uniquely identifies the report. If the same report is run multiple times with the same parameters, its Definition ID stays the same.



- Start Time—The time the report was run; for example, 2016-12-11 15:54:26. The time is based on your user profile; for example, UTC-07:00 Pacific Daylight Time.
- End Time—The time the report ended; for example, 2016-12-11 15:54:59.
- Elapsed Time—The time it took for the report to run; for example, 00:00:33.
- % Complete—Percentage of completion for the report.
- Job Status—Status of the report, which can be:
  - Created
  - Pending
  - In progress
  - Complete
  - Cancelled
  - Cancelling
  - Errors
  - Aborted

Step 3 (Optional) Use the View options to run, edit, schedule, download, or delete the report.



CHAPTER 8

Running Traffic Reports

Access: **WAE Live > Analytics**, click **New Report**, and choose Report Type: **Traffic**

Access: **WAE Live > Explore**, choose objects, and click **Run Report**

Traffic reports let you find trends for interface traffic, interface queues traffic, LSP traffic, and demands over a specified time range. These trends are based on time aggregations of the raw traffic data (for example, P95 traffic per week). The resulting report tables provide additional information, such as the number of aggregated objects on which the trending is computed. They also list the percentage at which traffic grows or shrinks.

This section contains the following topics:

- Additional Traffic Report Options, on page 47
- Projecting Traffic, on page 48
- Example: Linear Trend Data Aggregation, on page 49
- Traffic Report Tabs, on page 50

Additional Traffic Report Options

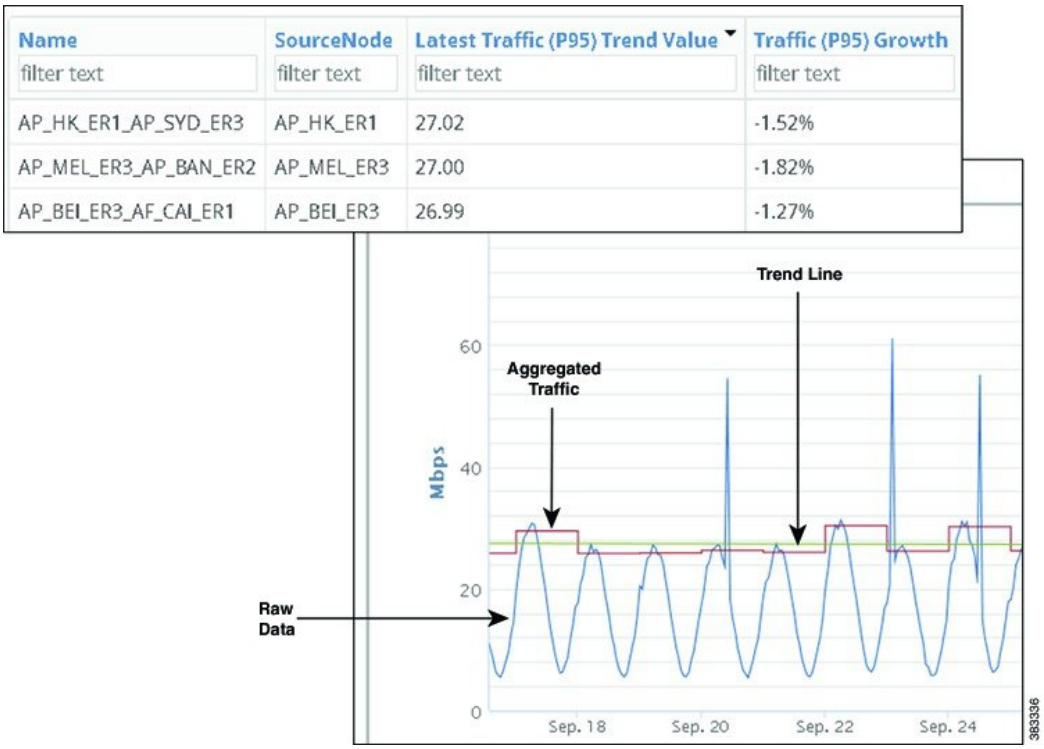
Besides the required Object Type and Filter options, the traffic report presents you with the following configuration options:

- Projection options, which are based on the trends.
  - Carry these historical trends into the future to determine at what date the traffic would be projected to surpass a given threshold.
  - Determine the state of traffic at a specific date.
- Group objects to trend aggregated (summed) traffic data. For information on the effects of grouping objects, see Configuring and Running Reports, on page 33.

Grouping objects is useful for aggregating traffic that can then be imported into WAE Design to create growth plans. For information, see the *Cisco WAE Design Integration and Development Guide*.

To configure reports for a specific network, choose it from the Network menu (top left). If there is only one network configured, the word “default” appears.

Example: This basic LSP traffic report shows the LSP traffic aggregated on a daily P95 basis. The trend values, growth, and trend line are all fairly even, and the traffic spikes upwards with some regularity.



Projecting Traffic

Projecting Traffic to Cross Thresholds

One projection option is to project which objects will exceed a specified utilization in a specified future time range. Once the report is generated, you can graph the historical and projected trend on a per-object or per-group basis over any period of time. These graphs let you visualize the trends and underlying data. They can show a pattern of when thresholds are surpassed, and they can show threshold violations narrowed to a single event.

The resulting graphs display two lines: one showing the trend within the report time period and within the projected time period, and the other showing the threshold. For an example graph, see Figure 7: Example Date Traffic is Projected to Pass Threshold, on page 49.

Example: If you set 50% for two years using a maximum monthly trend, the resulting projected threshold tells you when aggregated monthly maximum traffic is expected to exceed 50% of the interface capacity over the next two years.

Example: Figure 7: Example Date Traffic is Projected to Pass Threshold, on page 49 shows an example traffic report generated to determine which interfaces will exceed 80% of their outbound traffic utilization in the next five weeks based on last month's traffic trend. The trend is based on aggregating daily P95 Traffic Out. Sorting on Traffic Out (P95) Growth shows that the TenGigE0/0/0 interface on the AP_TOK_ER8 node has the highest percentage of traffic utilization growth.

Figure 7: Example Date Traffic is Projected to Pass Threshold

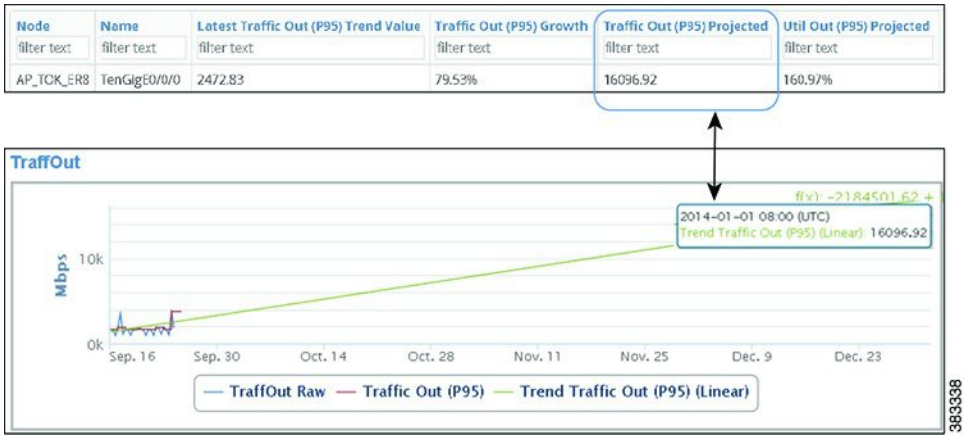


Projecting Traffic by Date

One projection option is to project the amount of traffic for a specified date. This is useful for planning purposes, such as determining when you might need to configure more LSPs or add more capacity.

Example: Figure 8: Example Traffic Projection by Date, on page 49 shows an example report output when projecting traffic by date to January 1, 2014. The report shows that the projected traffic on January 1 will be 16,096.92 Mbps for the same interface shown in Figure 7: Example Date Traffic is Projected to Pass Threshold, on page 49.

Figure 8: Example Traffic Projection by Date



Example: Linear Trend Data Aggregation

If one plan file exists per day, and if the month has 28 days, each interface in the datastore has 28 Traffic Out values for that month. If you run a traffic report on an interface and specify that you want to trend the weekly P95 of Traffic Out over this month, each week raw data is aggregated into a single P95 value. So the first weekly value is an aggregation of seven P95 values. The remaining weekly values are an aggregation of the

seven P95 values for their respective weeks. The result is a time-series aggregation with four data points on which the report is trended. These trends are viewable as graphs.

Traffic Report Tabs

Required Tabs

- For Object, Filter, and Time Range, see Configuring and Running Reports, on page 33.
- Trends—How to compute the trends.

Trend Fields

WAE Live first calculates the trend at which the property value is moving (its direction and frequency), and then uses this trend to determine the projected traffic. Set the following parameters to determine how this percentage is calculated:

- For interfaces and interface queues, trend the traffic on Traffic Out, Traffic In, or both. For LSPs and demands, only Traffic is available, and it is preselected.
- Interval—The aggregation basis for the trend: P99, P98, P95, P90, P85, Maximum, Average, or Minimum.
- Per—The aggregation interval that is applied to the raw data before calculating the trend. This interval must be less than the report time range.
- Trending Algorithm—Algorithm used to compute the trend. Both the calculation and the graph differ, depending on which option you choose. The graphs are commonly used as follows:
  - Exponential—Used when data values increase or decrease at a constantly increasing rate.
  - Linear—Used with linear data that shows a steady rate of increase or decrease.
  - Logarithmic—Used when data changes by increasing or decreasing quickly and then stabilizing and leveling.
  - Power—Used with data measurements that increase at a specific rate.

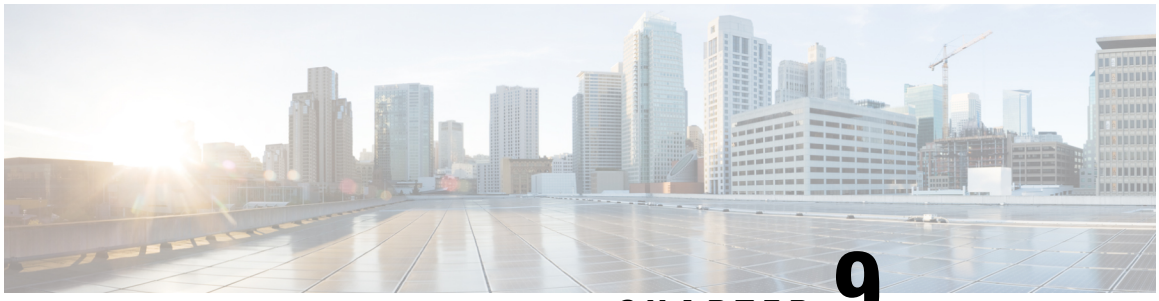
Optional Tabs

- Group By—Aggregate the report results by groups. The Overall option aggregates the results for all objects. For more information, see Configuring and Running Reports, on page 33.

  In the output, you can expand the group to see its objects by clicking the number in the Count column.

  Example: If you choose interfaces as the object and group by nodes, the resulting report includes a row for each node. The data in that row is for all interfaces connected to that node. To see the names of interfaces on that node, click the number in the Count column.
- Projection—Specify whether to project traffic.

- Find <*Object*> Over ___ % or Mbps—Project the traffic trend into the future based on the computed treProjecting Traffic by Date, on page 49nd. WAE Live determines and graphs when the trend will exceed this threshold. See Projecting Traffic to Cross Thresholds, on page 48.
- To Date—Project the traffic amount on a specified date. See .



CHAPTER 9

Running Health Reports

Access: **WAE Live > Analytics**, click **New Report**, and choose Report Type: **Health**

Access: **WAE Live > Explore**, choose objects, and click **Run Report**

Health reports let you monitor and remedy network health by finding network objects that exhibit problematic behavior over time. For example, you can report on specific nodes that surpass 90% CPU utilization.

Each health report uniquely identifies ways in which an object might be considered unhealthy or in the case of LAGs, which LAGs are imbalanced. For example, if you run an interface health report on utilization over 80%, the resulting report tells you the percentage of time each interface has over 80% utilization within the given time range, as well as the number of times the operating status changes. Another example is determining the number of times (count) that a specific event occurs, such as LSP path changes.

After generating the report, you can then graph the properties over the report time range. These graphs allow you to drill down to the specific times at which the event occurred. These graphs could indicate when a specific threshold was violated or they could narrow the threshold violations to a single event.

To configure reports for a specific network, choose it from the Network menu (top left). If there is only one network configured, the word “default” appears.

This section contains the following topics:

- Interface Health Report, on page 53
- Interface Queues Health Report, on page 54
- LAG Imbalance Report, on page 55
- LSP Health Report, on page 58
- Node Health Report, on page 58

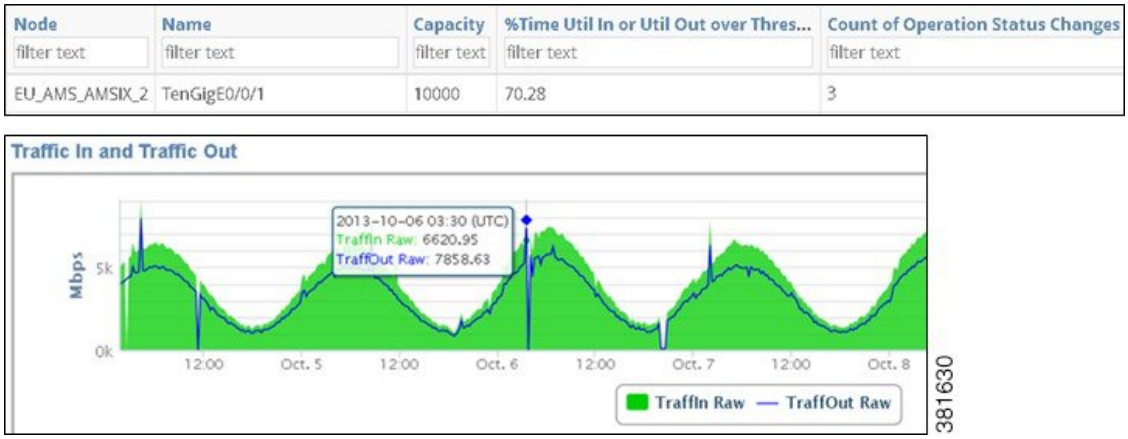
Interface Health Report

An interface health report gives the percentage of time that incoming and outgoing traffic surpassed a user-defined threshold. Additionally, it identifies the number of times the operational status changed.

Example

Figure 9: Example Interface Health Report, on page 54 shows an example health report output for an interface that was above an 80% utilization threshold 70.28% of the time, during which time there were 3 operating status changes. The graph visually shows that one of the peaks in utilization occurred on October 6. Moving the cursor over that point gives the exact date and utilization values.

Figure 9: Example Interface Health Report



Required Tabs

- For Object, Filter, and Time Range, see Configuring and Running Reports, on page 33.
- Utilization Threshold—Threshold interface traffic must surpass to be included in the report.

Interface Queues Health Report

An interface queues health report gives the percentage of time that incoming and outgoing traffic surpassed a user-defined threshold. Additionally, it identifies the number of times the operational status changed.

Example

One use case is to determine if interface queues traffic is exceeding thresholds or trending upward, which might trigger a QoS policy modification or a traffic engineering exercise. An interface queues health report shows threshold definitions, including traffic and utilization based on the overall interface capacity.

An interface queues health report shows the percentage of time an object's property exceeds a threshold, shows a trend line against traffic, and extrapolates traffic according to the observed trend.

Required Tabs

- For Object, Filter, and Time Range, see Configuring and Running Reports, on page 33.
- Utilization Threshold—Threshold interface queues traffic must surpass to be included in the report.

LAG Imbalance Report

A link aggregation group (LAG) is a group of physical interfaces that are bundled into a single logical interface. When traffic is not spread evenly across the members, congestion and dropped packets can occur even though the LAG itself is not congested.

A LAG Imbalance report lets you determine which LAGs are imbalanced and at risk of dropping packets or incurring congestion. The basis for determining whether a LAG is included in the report is whether any of its members surpass a user-specified Util Out threshold. This helps you target potentially problematic LAGs, avoiding those that do not have sufficient traffic to give accurate results.

The LAG Imbalance report is based on two key parameters:

- Include LAGs with members exceeding __%—This value determines which LAGs are included in the report. If a LAG has any member that passes this Util Out threshold (at any time during the report period), then the LAG and all of its members are included in the report.
- High- and Medium-risk threshold—These values set the thresholds that WAE Live uses to determine whether a LAG is imbalanced or at risk. The value must be between 0 and 1.

LAG Imbalance and Risk Report Example

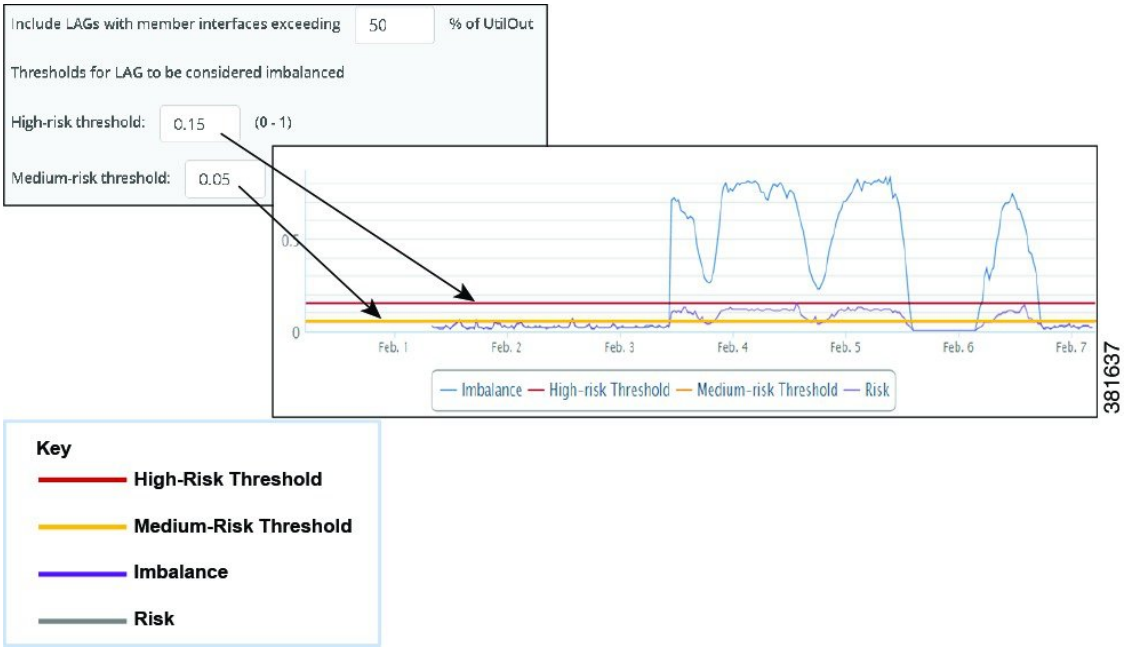
The LAG Imbalance/Risk graph shows the imbalance and risk compared to the high- and medium-risk thresholds, as well as whether and when the LAG passes these thresholds. The difference between risk and imbalance measurements is that risk calculations consider only positive differences, whereas imbalance calculations consider both negative and positive differences. For example, Figure 10: Example LAG Imbalance and LAG Risk Graph, on page 56 shows that a LAG's utilization was significantly imbalanced for several days, while it was only minimally at risk.

- Risk—Highest Util Out of Any Member - Average Util Out of All LAG Members
- Imbalance—Higher of these two values:
  - Risk
  - Average Util Out of All LAG Members - Lowest Util Out of Any Member

Example: This example LAG has four members.

- Util Out values are 0.10, 0.35, 0.55, and 0.40
- Highest Util Out of Any Member = 0.55
- Lowest Util Out of Any Member = 0.10
- Average Util Out of All LAG Members = (0.10 + 0.35 + 0.55 + 0.40)/4 = 0.35
- Risk = 0.55 - 0.35 = 0.20
- Imbalance = 0.25 because it is the higher of the following two values:
  - Risk, which is 0.20
  - 0.35 - 0.10 = 0.25

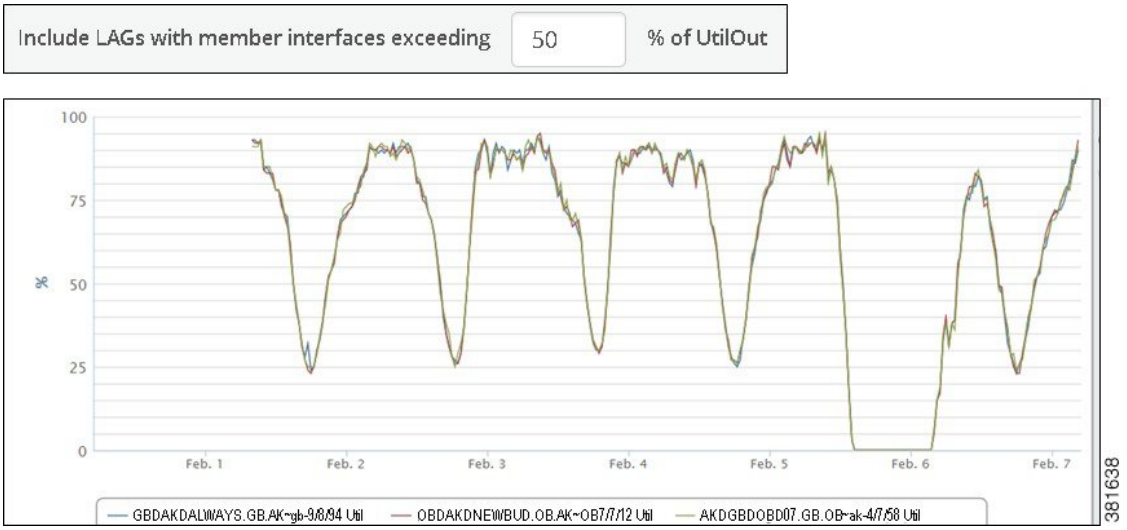
Figure 10: Example LAG Imbalance and LAG Risk Graph



LAG Member Utilization Out Graph Example

The Util Out graph shows the Util Out percentage of all of the LAG's members over the report period (Figure 11: Example LAG Member Util Out Graph, on page 56). If you see that a LAG is imbalanced or at risk, this graph is helpful in determining which of the members is causing these effects.

Figure 11: Example LAG Member Util Out Graph



Additional Utilization, Imbalance, and Risk Data

In addition to these graphs, the report output includes the following utilization, imbalance, and risk values:

- Number of LAG members, which are viewable individually by clicking the blue number in the LAG Members column.
- Peak member utilization.
- Peak aggregated risk.
- Peak aggregated imbalance.
- Most recent aggregated risk.
- Most recent aggregated imbalance.
- Percentage of time the LAG was over the high-risk threshold.
- Percentage of time the LAG was over the medium-risk threshold.

Example

This example compares two scenarios to demonstrate how you could interpret the data.

- LAG A: Nine LAG members are at 60% utilization and one LAG member is at 0%.

  While LAG A is imbalanced, it does not appear to be at high risk. Traffic coming into this LAG will likely not pose a problem since the data suggests that even if one member fails, the other nine members can handle the traffic.

- LAG B: Nine LAG members are at 60% utilization and one LAG member is at 90%.

  LAG B is imbalanced and likely at high risk. One particular member is receiving much more traffic than the other members. If this happens when traffic is coming into LAG B and if that traffic is assigned to the member with 90% utilization, there might not be sufficient bandwidth left to handle the traffic, resulting in dropped packets.

Required Tabs

- For Filter and Time Range, see Configuring and Running Reports, on page 33.
- LAG Imbalance
  - Include LAGs with members exceeding __%—Specify the percentage of Util Out threshold used in determining which LAGs are included in the report. If any LAG member surpasses it, then the LAG is included in the report.
  - High- and Medium-risk threshold—Specify the thresholds used to determine whether a LAG is imbalanced or at risk. The value must be between 0 and 1.

LSP Health Report

An LSP health report identifies the percentage of time that the LSP traffic surpasses the user-specified threshold of setup bandwidth. Additionally, it provides the following information:

- Last setup bandwidth
- Number of actual path changes
- Percentage of time the LSP was not on the shortest TE path

Required Tabs

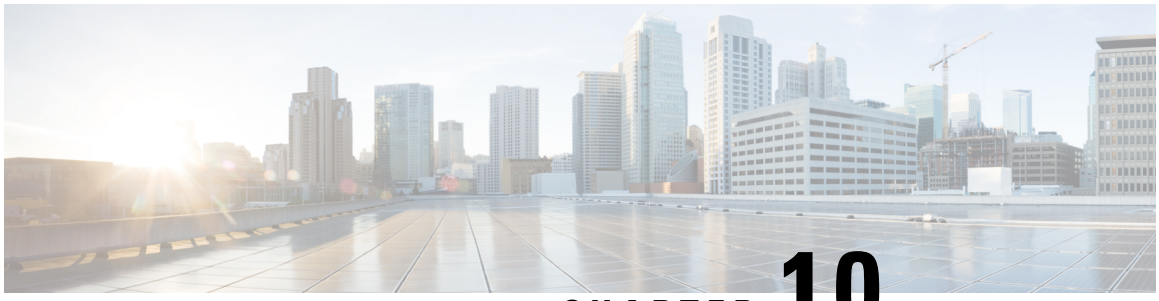
- For Object, Filter, and Time Range, see Configuring and Running Reports, on page 33.
- SetupBW Threshold—Setup bandwidth threshold that LSPs must surpass to be included in the report.

Node Health Report

A node health report identifies the percentage of time that the CPU and memory surpass the user-specified threshold.

Required Tabs

- For Filter and Time Range, see Configuring and Running Reports, on page 33.
- CPU Threshold—CPU threshold that nodes must surpass to be included in the report.
- Memory Threshold—Memory threshold that nodes must surpass to be included in the report.



CHAPTER 10

Running Deviation Reports

Access: **WAE Live > Analytics**, click **New Report**, and choose Report Type: **Deviation**

Access: **WAE Live > Explore**, choose objects, and click **Run Report**

The deviation report detects sudden changes in property values. While the report can be applied to any high-frequency property value, such as nodes' CPUs, the primary use is to detect rapid increases and decreases in LSP or interface traffic. Determining the timestamp can assist you in correlating other network or service events that may have caused a reported surge.

This detection of a deviation is based on how traffic has behaved in the past, which is considered to be the *norm* or the *moving average* . For example, if traffic normally goes up and down within a specific band of Mbps, then this is normal traffic and would not be flagged as a deviation.

This section contains the following topics:

- Deviation Report Parameters, on page 59
- Tips for Viewing Graphed Deviations, on page 61
- Deviation Report Tabs, on page 61

Deviation Report Parameters

The moving average and range of acceptable tolerance for the property value are determined by two configurable fields, each of which creates corresponding lines in the graph.

- Size of Historical Window—Identifies how many past data points to use in the calculation of the moving average.
  - The moving average, which appears in the graph, is how WAE Live expects the property value to change based on past behavior. This is calculated using Bollinger Bands.
  - A data point is a point in time in which you collect data. For example, if you collect every 5 minutes and you want the moving average to be based on the last 24 hours, enter 288.
- Sensitivity—Identifies the tolerance range (standard deviation), which is how far from the moving average the raw data can be before it is flagged as a deviance: the higher the number, the higher the tolerance. This creates two lines that are equal distance on either side of the moving average line. If raw data crosses either the upper or lower band of tolerance, a deviance is recorded. See Figure 12: Deviation Report Configuration and Output, on page 60 for an example.

You can also set a threshold under which property values are not considered for the deviation calculation, thus filtering out insignificant deviation points. If the current data point is larger than the upper band or lower than the lower band, it is compared to the previous data point. If the absolute difference of these two values is lower than the “Ignore Deviations Under” value, the property value is ignored.

Example: Upper band is 1,000 Mbps and lower band is 500 Mbps. The “Ignore Deviations Under” value is 5. The last data point is 999 and the current data point is 1003.

1003 (last data point) - 999 (current data point) = 4

4 < 5 (Ignore Deviations Under)

Therefore, the current data point of 1003 would not be considered as a deviant. Note that had there not been an “Ignore Deviations Under” value set, the 1003 data point would be a deviant since it is larger than the upper band.

In addition to the graph, the report output identifies the following.

- Number of Deviations
- Average Deviation Amplitude Absolute, which is the size difference between the highest and lowest property values
- Average Deviation Amplitude Relative, which is a comparison of the absolute amplitude to the standard deviation

To configure reports for a specific network, choose it from the Network menu (top left). If there is only one network configured, the word “default” appears.

Example: Figure 13: Example Deviation Report, on page 61 shows an example of a report identifying two deviations. Figure 12: Deviation Report Configuration and Output, on page 60 shows the configuration and an enlarged section of its output.

Figure 12: Deviation Report Configuration and Output

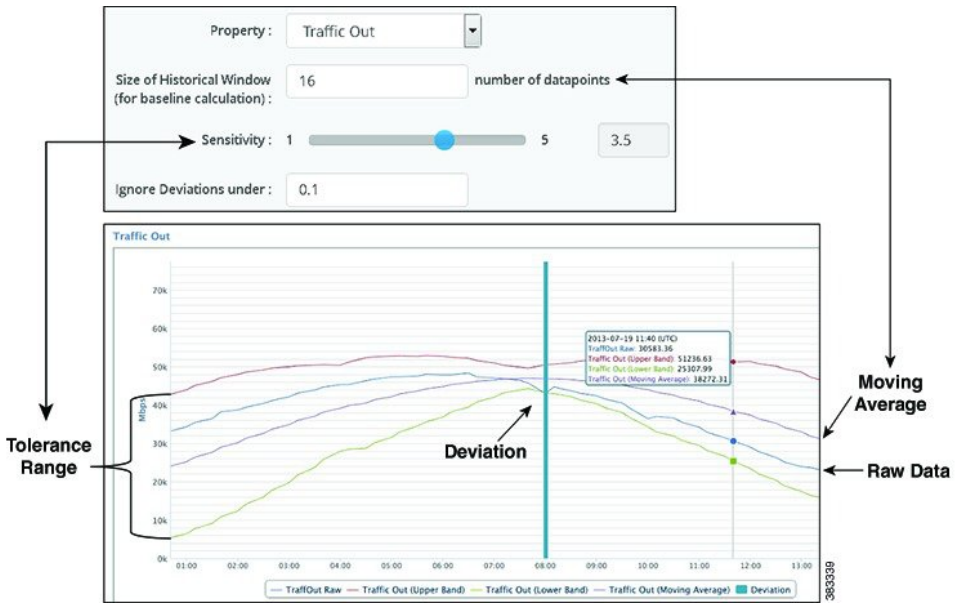
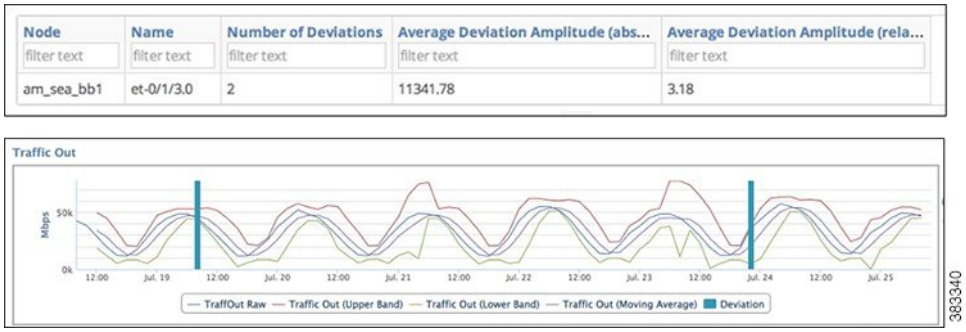


Figure 13: Example Deviation Report



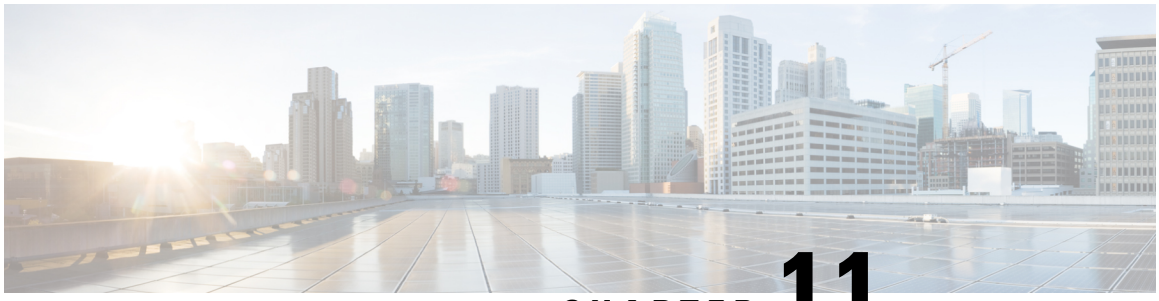
Tips for Viewing Graphed Deviations

- Increasing the “Ignore Deviations Under” value can make it easier to see the deviations.
- If the number of deviations reported in the table are different than what you see in the graph, use the zoom feature to see them. Zoom to a chosen timestamp, and note that timestamp. If you do not see the deviation of interest, zoom to a section of the graph immediately following that timestamp. Repeat this procedure, if needed, to see all of the deviations.
- If the tooltip for the deviation mark shows that the raw data is within the upper and lower boundaries (tolerance range), zoom into that deviation mark as many times as needed to see the raw data cross the boundary.

Deviation Report Tabs

Required Tabs

- For Filter and Time Range, see Configuring and Running Reports, on page 33.
- Parameters:
  - Size of Historical Window—Enter the number of past data points (collections) to use in the calculation of the moving average.
  - Sensitivity—Slide the scale to a number from 1 to 5 (including up to one decimal) to specify the tolerance range, which is how far from the moving average the raw data can be before it is flagged as a deviance.
  - Ignore Deviations Under—(Optional) Enter a number (up to one decimal) to set the threshold under which property values are not considered for the deviation calculation.



CHAPTER 11

Running Ad Hoc Reports

Access: **WAE Live > Analytics**, click **New Report**, and choose Report Type: **Ad hoc**

Access: **WAE Live > Explore**, choose objects, and click **Run Report**

Ad hoc reports offer a flexible means of customizing reports to meet specific requirements, making them useful for a variety of use cases. For example, you can average the sum of traffic exiting grouped nodes, or determine how many times an LSP left its shortest TE path.

The basic ad hoc report creates a table of data in which each table value is the result of the time aggregation operation applied to the property over the entire report time range. Each row is an object. Each column is defined by a combination of properties and time aggregations (both of which are selected in the Data Columns fields).

This section contains the following topics:

- Time Aggregation Options, on page 63
- Aggregating by Time Intervals, on page 64
- Graphing Percentage Growth, on page 65
- Ad Hoc Report Tabs, on page 66
- Example: Ad Hoc Data Aggregation, on page 67
- Example Ad Hoc Report, on page 67

Time Aggregation Options

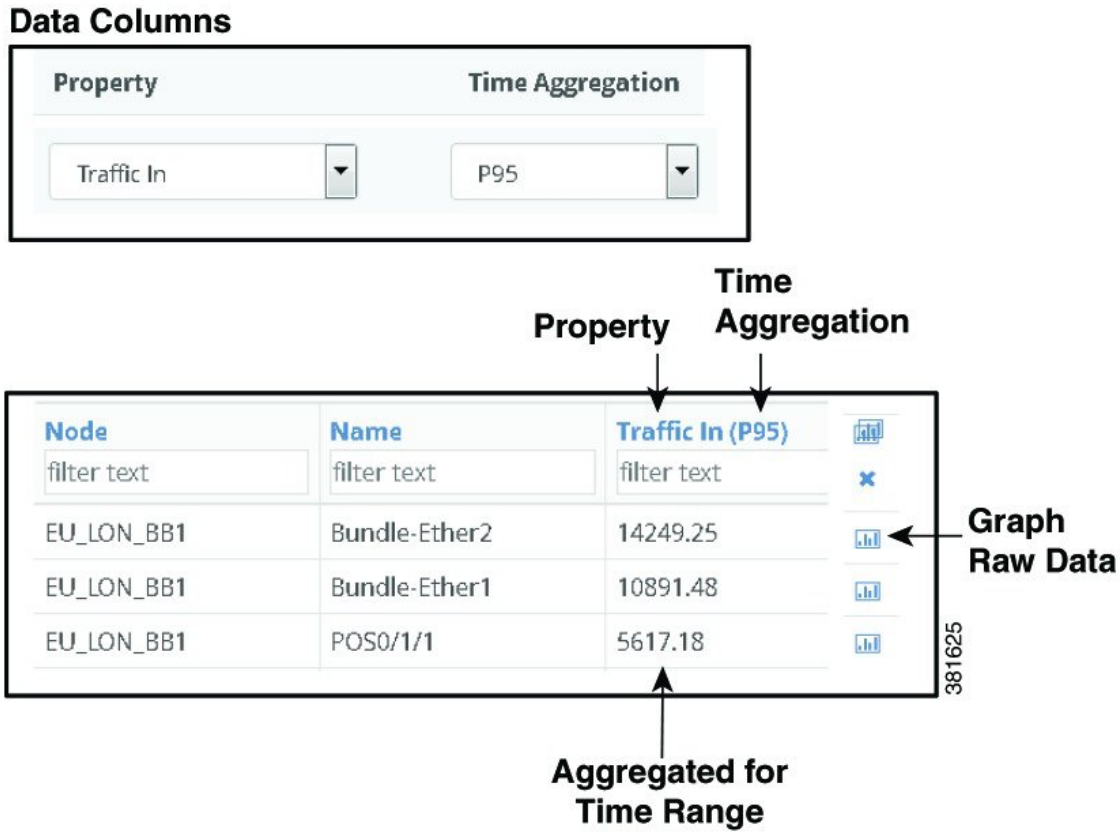
Time aggregation operations vary, depending on the selected object:

- Change Count—Number of times the value changed. For example, if you choose Shortest TE Path for LSPs, you can report on the number of times it changed during the last year.
- Count—Number of times the data was collected during the specified time period. For example, if you collected traffic information once an hour for 24 hours, and the Count was 20, you would know there was a problem during 4 hours.
- P99, P98, P95, P90, P85, Maximum, Minimum, Average—Time aggregation computations based on these values.
- Last—Last value in the specified time interval.
- Percentage Growth—Computation of the percentage of growth.

To configure reports for a specific network, choose it from the Network menu (top left). If there is only one network configured, the word “default” appears.

Example: Figure 14: Example Ad Hoc Report, on page 64 shows the output for aggregating P95 Traffic In for interfaces connected to the EU_LON_BB1 node. Interface Bundle-Ether2 (row 1) has 14,249.25 Mb of P95 traffic for the last three months.

Figure 14: Example Ad Hoc Report



Aggregating by Time Intervals

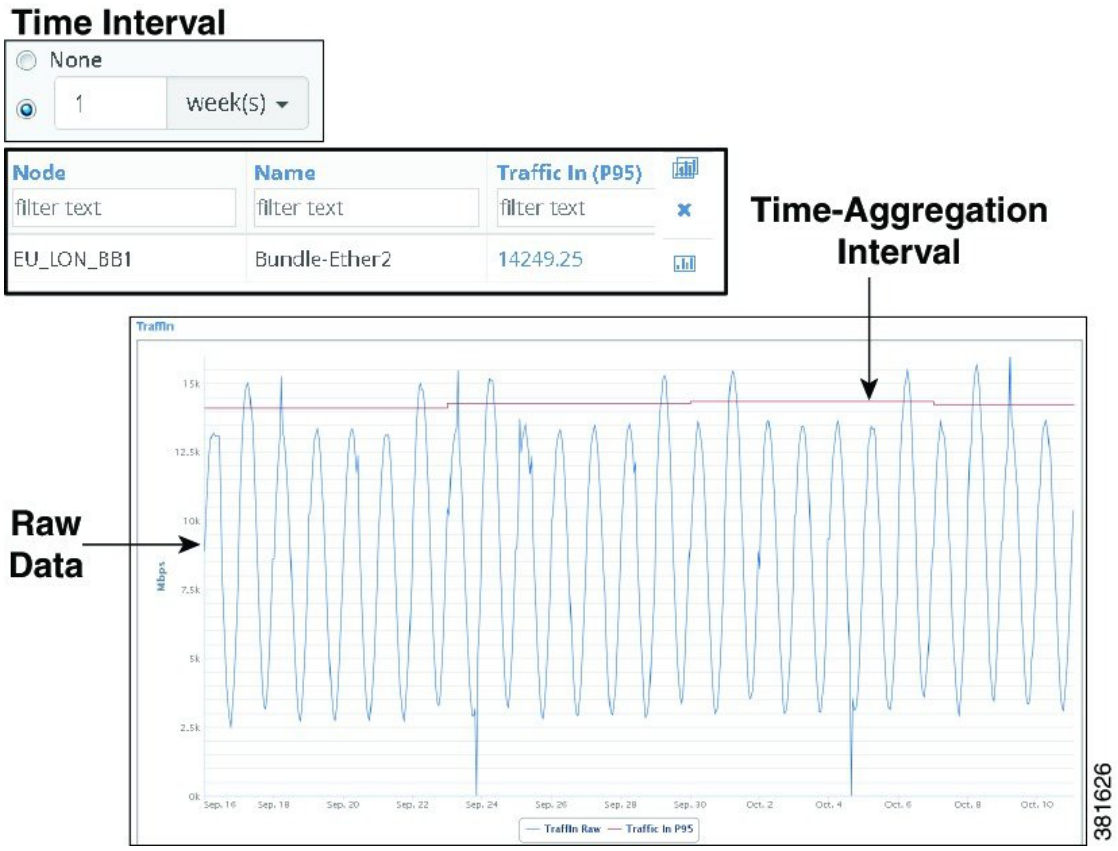
Each combination of property and time aggregation value can be aggregated by time intervals ranging from hours to years. This broadens the possibilities for how you view the data. For example, you can report on the average number of dropped packets for the last six months (Time Range), and aggregate those averages on a weekly basis (Time Interval). Sorting the Dropped Packets Out column would then show the most problematic interfaces. Another example would be to create yearly reports on a quarterly basis.

These intervals are viewable as graphs. You can graph the time-aggregation value by clicking the aggregation value, or you can graph both the raw data and the time-aggregation interval by clicking the graph icon.

Example: Figure 15: Example Ad Hoc Report with Graphed Time Aggregations Intervals, on page 65 extends the report created in Figure 14: Example Ad Hoc Report, on page 64. Here, each property and time aggregation is further aggregated into a one-week time interval. The difference is that the graph now shows the P95

aggregation applied to one-week increments across three months. By looking at these increments you can see that while the raw data fluctuations, the traffic is very steady on a per-week basis.

Figure 15: Example Ad Hoc Report with Graphed Time Aggregations Intervals



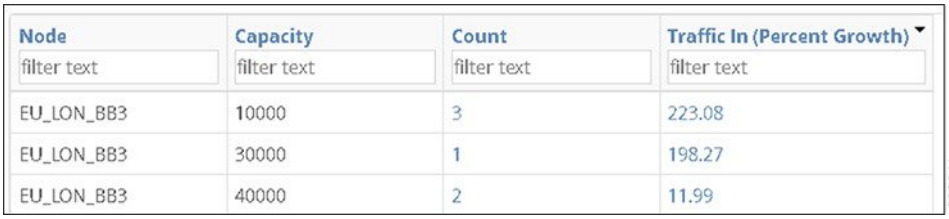
Graphing Percentage Growth

Ad hoc reports let you determine and graph the percentage of growth for any high-frequency property over a specified time range. The resulting data is based on time aggregations of the raw traffic data (for example, average traffic per day). The report tables provide additional information, such as the number of aggregated objects on which the trending is computed. They also list the percentage at which the traffic is growing or shrinking.

One potential use for the percentage growth time aggregation would be to identify how fast your traffic is growing so you can plan for increasing capacity in the future, or whether it is declining, which could be an indication of a diminishing market or service.

Example: Figure 16: Example Ad Hoc Report of Percent Growth, on page 66 shows the output of an ad hoc report that was grouped by node and then by capacity. The time range was 6 months, and the growth computation was based on the daily maximum traffic. Sorting on Traffic In (Percent Growth) shows that the EU_LON_BB3 node with 10,000 Mbps capacity had the largest percentage growth in the last 6 months at 223.08 percent.

Figure 16: Example Ad Hoc Report of Percent Growth



Node	Capacity	Count	Traffic In (Percent Growth)
filter text	filter text	filter text	filter text
EU_LON_BB3	10000	3	223.08
EU_LON_BB3	30000	1	198.27
EU_LON_BB3	40000	2	11.99

383341

Ad Hoc Report Tabs

Required Tabs

- For Object, Filter, and Time Range, see Configuring and Running Reports, on page 33.
- Data Columns—Apply the time aggregated operation to the associated property. The selected property and time aggregation combination defines the output columns for the report. Defining each Data Column is a two-step process.

  - Add or delete properties, as needed.

  – To add a property, click the Plus icon.

  – To delete a property, click the Trash Can icon.

  For each selected property, identify how you want the data time aggregated for the report time range. See Time Aggregation Options, on page 63. If you specify Percent Growth, you must further identify how you want this growth calculated. See Percent Growth Options, on page 66.

Time Aggregation Options

Time aggregation operations vary, depending on the selected object as follows.

- Change Count—Number of times the value changed.
- Count—Number of times the data was collected during the specified time period.
- P99, P98, P95, P90, P85, Maximum, Minimum, Average—Time aggregation computations based on these values.
- Last—Last value in the specified time interval.
- Percentage Growth—Computation of the percentage of growth.

Percent Growth Options

WAE Live first calculates the trend at which the property value is moving (its direction and frequency), and then uses this trend to determine the percentage of growth. Set the following parameters to determine how this percentage is calculated.

- Interval— The aggregation basis for the percentage: P99, P98, P95, P90, P85, Maximum, Average, or Minimum.

- Per—The aggregation interval that is applied to the raw data before calculating the percentage of growth. This interval must be less than the report time range.
- Trending Algorithm—Algorithm used to compute the percentage of growth. Both the calculation and the graph differ, depending on which option you choose. Following is a list of how these graphs are commonly used.
  - Exponential—Commonly used when data values increase or decrease at a constantly increasing rate.
  - Linear—Commonly used with linear data that often shows a steady rate of increase or decrease.
  - Logarithmic—Commonly used when data changes by increasing or decreasing quickly and then stabilizing and leveling.
  - Power—Commonly used with data measurements that increase at a specific rate.

Optional Tabs

- Group By—Aggregate the report results by groups. The Overall option aggregates the results for all objects. For information, see Configuring and Running Reports, on page 33.

  In the output, you can expand the group to see its objects by clicking the number in the Count column.
- Time Interval—Apply the time aggregation operations defined in the Data Columns by time intervals. For more information, see Aggregating by Time Intervals, on page 64.

  Example: If you choose Traffic as an LSP property in the Data Columns section and P95 as the Time Aggregation operation, and if you choose 1 week as the Time Interval, the resulting graph shows P95 traffic aggregated on a weekly basis (one data point per week, which is the P95 value of Traffic over that week).

Example: Ad Hoc Data Aggregation

If one plan file is generated per hour, then each day is comprised of 24 raw data points. If you run an ad hoc report for maximum LSP traffic for the last 12 weeks, and if you use weekly intervals, the resulting report first aggregates the maximum LSP traffic for those 12 weeks per object (provided those objects are not grouped). The number of raw data points included in this aggregation would be 169,344 (84 days * 24 raw data points per day). These aggregated values appear in the table columns. Additionally, the intervals would be aggregated per week for a total of 168 raw data points per each of the 12 weeks (168 = 7 days * 24 raw data points per day). These intervals are viewable as graphs. For information on how raw data is collected, see Understanding Objects, Properties, and Data, on page 3.

Example Ad Hoc Report

Figure 17: Example Ad Hoc Report, on page 68 shows an example ad hoc report generated to determine the aggregated weekly maximum outgoing traffic per node over the past year. It also includes a calculation for the percentage of growth The totals in the resulting report are for three months, and the graphs show the traffic aggregated in weekly increments. Clicking the number in the Count column shows the names of the interfaces.

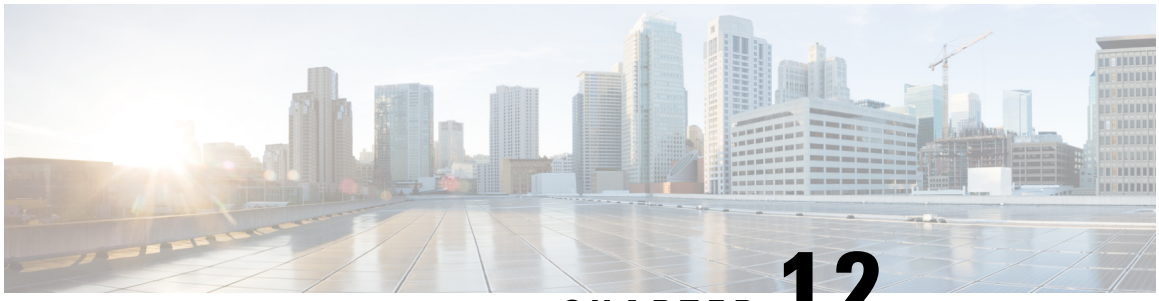
The table is filtered by EU nodes and sorted on the Traffic Out (Maximum) column to show which European nodes have the highest outbound traffic using the given aggregation. For the European nodes with the highest traffic, the table also shows that these nodes have little fluctuation in their growth, but that all but one has decreasing traffic. The node showing the greatest percentage of loss is EU_MAD_BB2 at -23.25 percent loss over the last three months. Yet, you can also determine that while it is showing a percentage growth loss over time, the node has traffic spikes.

Report Configuration	Parameters
Report Type	Ad Hoc
Object	Interfaces
Filter	Node
Time Range	Last three months
Group By	Node
Time Interval	1 week
Data columns	Traffic Out, Maximum
	Traffic Out, Percentage Growth based on a maximum interval of 1 day using a linear trend

Figure 17: Example Ad Hoc Report



Node	Count	Traffic Out (Maximum)	Traffic Out (Percent Growth)
EU	filter text	filter text	filter text
EU_LON_BB1	27	226203.00	2.12
EU_PAR_BB1	26	209134.62	-2.95
EU_PAR_BB3	16	129778.63	-6.47
EU_AMS_BB2	19	85742.99	-1.28
EU_LON_BB2	13	85254.13	-6.34
EU_AMS_BB1	12	48909.92	-7.01
EU_MAD_BB1	7	48035.67	-3.69
EU_MAD_BB2	8	41286.45	-23.25

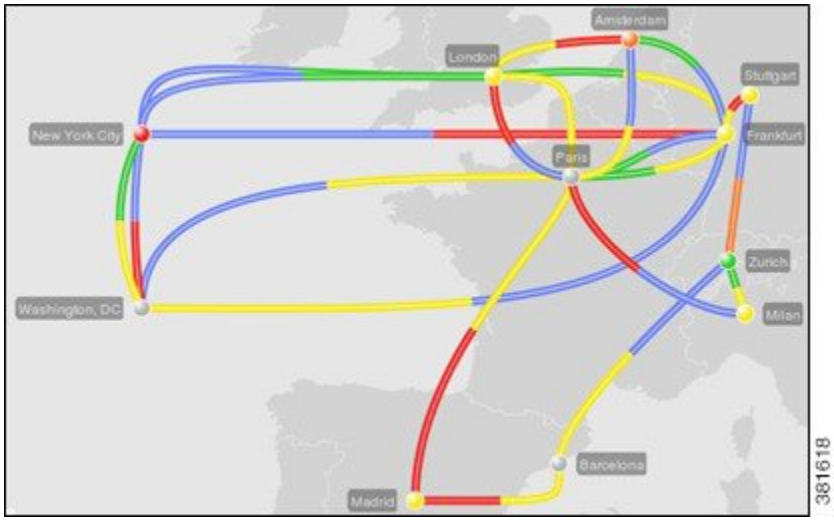


CHAPTER 12

Viewing Network Health and Weathermaps

Access: **WAE Live > Map**

The Map page provides a high-level view of the current network health through a near-realtime weathermap and health panels of critical peering traffic and network issues. Because the information is current and highly visual, you can see at a glance whether and where there are problems in the network, such as congestion. Because the Map and Explore capabilities are tightly coupled, you can immediately drill down to detailed information to troubleshoot network issues.



WAE Live weathermaps are generated by the network topology and traffic information that WAE Server discovers. This information is stored in a series of plan files called *snapshots* that are navigable by date.

The Peering and Events panels offer immediate information on network health. The Peering panel lets you monitor top peers and verify traffic is balanced across peers and sites. The Events panel lists the predominate types of network concerns, such as interfaces down, so you can determine how many there are total and how many have changed since the last snapshot.

To view a weathermap for a specific network, choose it from the Network menu (top left). If there is only one network configured, the word “default” appears.

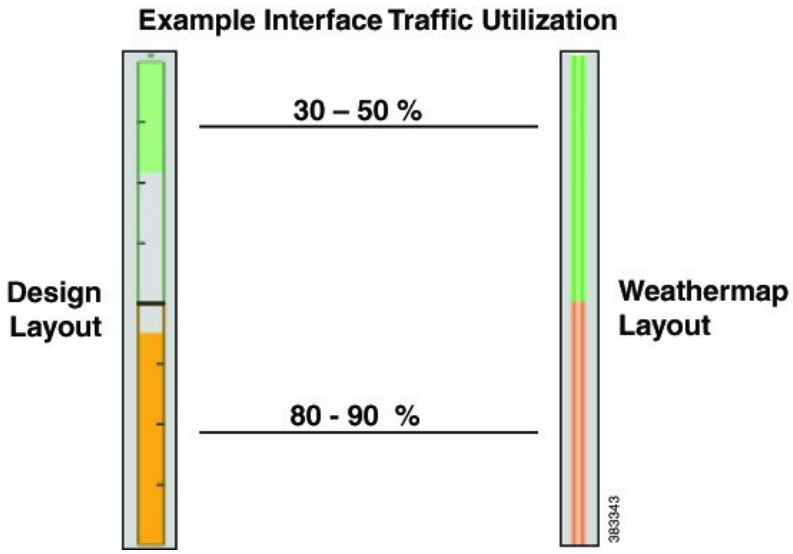
Note The visualization of the weathermap is based on how your administrator laid out the template. For detailed descriptions of the objects, backgrounds, and traffic utilization colors, see the Cisco WAE Network Visualization Guide.

This section contains the following topics:

- Assessing Traffic Utilization, on page 70
- Weathermap, on page 74

Assessing Traffic Utilization

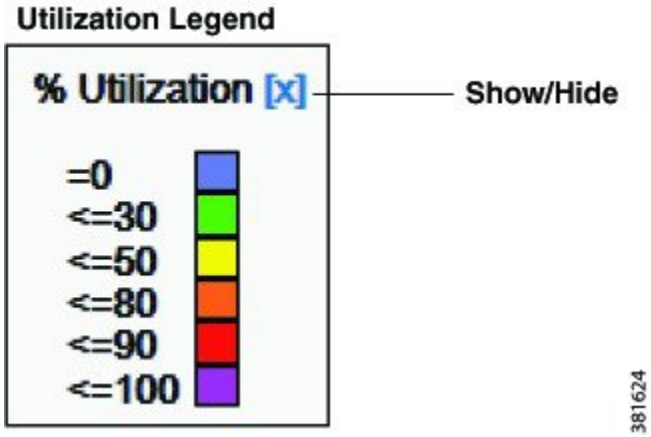
WAE Live offers multiple ways to quickly assess network health. The most immediate way is to look at the color of the interfaces, nodes, and sites in the weathermap and to look at the color of the peer icons in the Peering health panel. These colors represent the outbound traffic utilization on the interface. If parallel circuits[1] have been grouped, these utilizations represent the average of the grouped interfaces.



Default utilization colors depends on how your administrator set up the visualization of the plan file. This visualization might be either in a Design or Weathermap color scheme. Default colors are blue, green, yellow, orange, red, and purple (Design) or magenta (Weathermap), and sequentially each of these colors represents increased utilization.

Note these colors are configurable in the WAE Design GUI where the administrator creates the template for the WAE Live plan files. The utilization legend at the bottom of the weathermap identifies the template's utilization colors and their values.

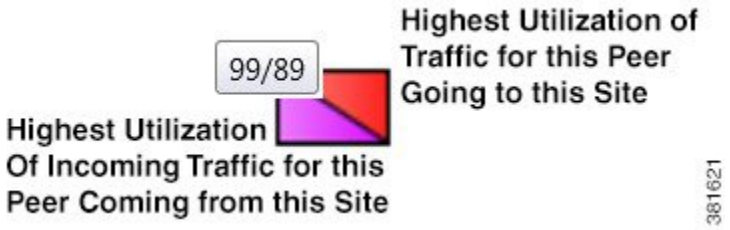
[1] Each circuit connects a pair of interfaces on different nodes. Circuits are parallel if they have the same source node and same destination node as each other.



Peering Health Panel

The Peering health panel is based on BGP information. It shows the highest traffic utilization per peer and per site. This panel lets you monitor critical peers and their growth so that you can quickly detect and overcome congestion issues.

The far left column lists the peers, while each top column lists the associated sites. The icons represent the traffic flowing between the two. While the color gives you a visual indication of traffic, hovering the mouse over each icon lists the numeric values.



To navigate to graphs for all interfaces connected to that peer, click the icon. Graphs appear for only those interfaces connecting this peer on its associated site. To go to the associated Explore page for these interfaces, click the All Interfaces link at the top of the graphs.

Note The administrator can configure this panel name to something other than “Peering.”

Events Health Panel

The Events health panel lists the number of current events and the number of changes since the last plan file was received.

- Event—The event being tracked. The default set is as follows, though these can be configured differently.
  - Circuits down
  - Nodes down

- Interfaces errors/drops
- External (peering) interfaces over 80%
- Internal (non-peering) interfaces over 80%

- Count—The number of objects of that event type.

Example: If there are 10 nodes total that are down, the Nodes Down Count is 10.

- Last Delta—How many objects have changed their state since the last snapshot was taken.

Example: If 4 nodes came up since the last snapshot, and 3 nodes went down, the Last Delta column would show +4/-3. Note, if the Count was previously 10, it would now be 11.



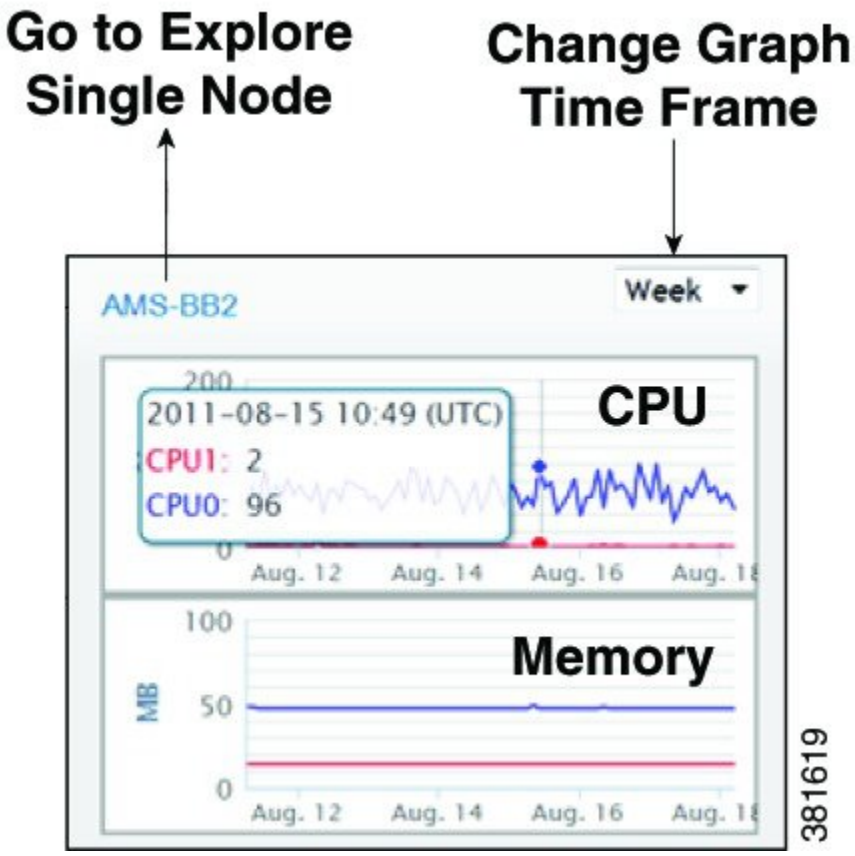
Note The administrator can configure this panel name to something other than "Events" and can configure a different set of events to appear in this panel.

Navigating to Related Explore Information

From a weathermap, you can navigate to related Explore information for nodes and interfaces. Note that regardless of the plan file date, the resulting Explore information is the most recent. For more information on Explore capabilities, see Exploring Network Data, on page 17 .

Graphing Node CPU and Memory Usage

To graph node CPU and memory usage, click its parent site. From here you can graph usage for this particular node, or you can go to the Explore tables for one or more nodes in the site.



Graph Node CPU and Memory Usage	View Node Tables
Click a node in the site view. • The top pane graphs both CPU0 (blue) and CPU1 (red). • The lower pane graphs both Memory0 (blue) and Memory1 (red). • To view the numeric values, move the cursor over the graph, and a tooltip lists the values.	• To view an Explore table for a single node, open its graph and click its name link. • To view an Explore table containing data for all nodes in a site, click the All Nodes link in the site view.

Graphing Interface Traffic In/Traffic Out

To graph the traffic coming into an interface and leaving it, click the interface. The height of the interface is in relationship to the capacity so you can see how close the interface is to being filled. If the weathermap is displaying an interface group, all interfaces within it appear in the graph.

From here you can go to the Explore tables for one or more interfaces in the group by clicking on name of the interface.

Graph Interface Traffic In/Traffic Out	View Interface Tables
Click an interface or set of grouped interfaces. (Interfaces with a common source and destination are sometimes grouped in the template to improve visualization.) A single interface and grouped interfaces look the same in the weathermap. • If you click an interface group, multiple graphs appear. • The Traffic In (ingress) is green, and the Traffic Out (Egress) is blue. • To view the numeric values, move the cursor over the graph, and a tooltip lists the values.	• To view an Explore table for a single interface, open its graph and click its name link. • To view an Explore table for all interfaces in a group, open its graph and click the All Interfaces link. • To view an Explore table for all interfaces in a site, click the site to open it. Then click the All Interfaces link in the site view.

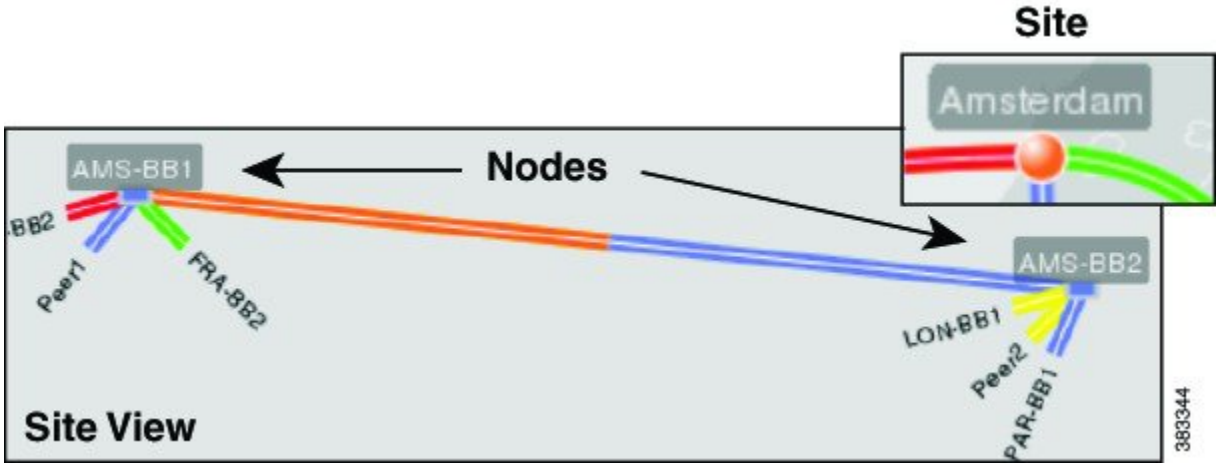


Weathermap

By default, the weathermap shows the current topology. The control panel at the top lets you navigate to different snapshots (plan files from different times and dates), view different QoS service classes and layouts, view specific plots, and change the manner in which circuits are displayed (Plot Options). From the weathermap and from its sites, you can navigate to more specific information about interfaces and nodes as described in Navigating to Related Explore Information, on page 72.

Map Objects

The objects appearing in on the Map page represent the network topology. How they appear can differ depending on how the layout was set up to visualize the network.



- A node (router) appears as either a disk or a rectangle. Nodes can reside inside and outside of sites.
- A site is a WAE construct to make visualization of networks easier to read. Sites appear as either a square (which is a peer or non-peer site), a circle (non-peer site), or cloud (peer site).

  Unless a site is empty, there is a hierarchy of the sites and nodes that a site contains. A site can be both a parent and a child site. A site that contains another site is called a *parent* site. The contained nodes and sites are called *children* . Often the child nodes and sites are geographically co-located. For example, a site might be a PoP where the routers reside.

  - In the weathermap, the parent site shows all egress inter-site interfaces of all nodes contained within it, no matter how deeply the nodes are nested. Similarly, this is true for each child site view within it.
  - In a site view, the “All Nodes” and “All Interfaces” links show a count for (and filter to) all nodes and interfaces directly within that site. That is, it does not show nodes and interfaces for nested sites.
  - To open a site view, choose a site from the Site drop-down list in the control panel, or double-click a site in the weathermap.

- Interface—Each pair of interfaces connects two different nodes. Their colors represent the outbound traffic utilization. For more information, see Assessing Traffic Utilization, on page 70.

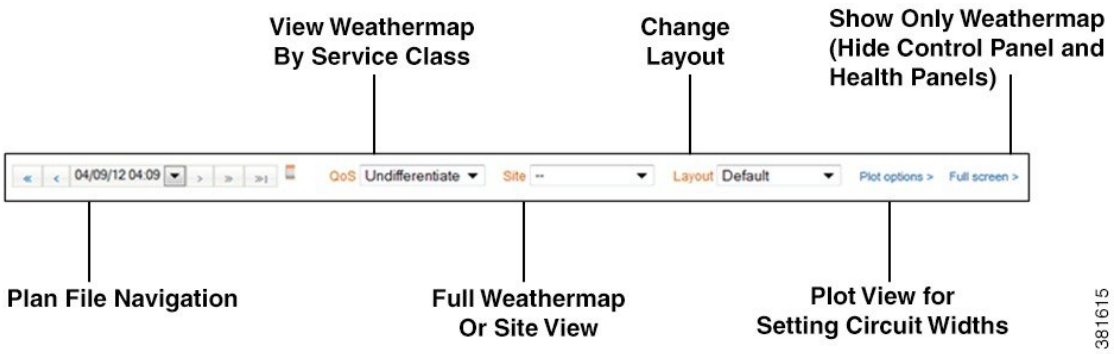
Control Panel Options

- Plan File Navigation—Click to save the file to a location of your choice. See Navigating to a Specific Weathermap, on page 76.
- QoS—Displays the type of traffic to view in the weathermap as differentiated by the QoS model. Note that if you navigate to an Explore page while displaying QoS traffic, the Explore page shows only undifferentiated traffic.
  - Undifferentiated traffic—Aggregate traffic on an interface.
  - Service class—Administratively defined classification of traffic that is not discovered.
  - Queue—Interface traffic belonging to the selected queue. Queues and their traffic are discoverable by WAE Collector.

- – They can be manually created and configured during the snapshot process to be associated with a set of user-defined QoS parameters that specify how these queues are prioritized and what percentage of traffic they carry.
- – Queue names are preceded by a “Q:” in the list.

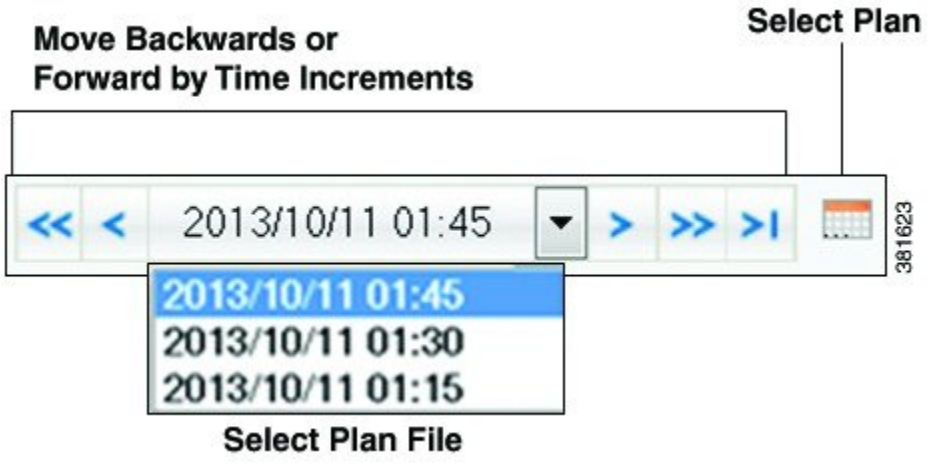
- Site—Displays a list of sites available for viewing, as well as the entire weathermap. If you click on a site, the weathermap shows the nodes and intrasite circuits within the site.
- Layout—The manner in which plan files are laid out to visualize the network is called a *layout* . This option displays a list of available weathermap layouts. Example layouts might include schematic layouts, geographic layouts, or separate layouts for different continents.

  The default layout is configurable by the administrator. Upon first opening WAE Live, this is the layout that you see. However, the last layout viewed before closing the browser is the one that appears the next time you open WAE Live provided you are using the same browser.
- Plot Options—Let you modify visual circuit attributes of the weathermap. See Set Plot Options, on page 77.
- Full Screen—Lets you view an enlarged weathermap without the control panel and without the health panels.



Navigating to a Specific Weathermap

By default, the weathermap shows the most recent snapshot (plan file) taken. You can, however, navigate to plan files created at different times and on different days.



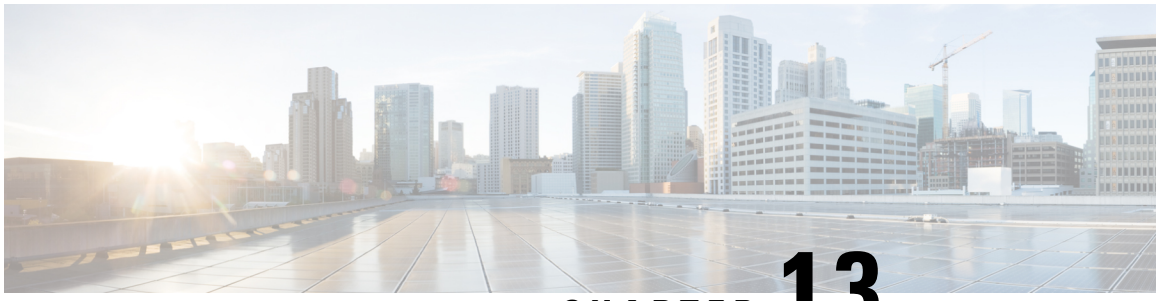
Plan File Navigation Options

- Single forward and back arrows—Move forwards or backwards by one increment in time.
- Double left arrow—View a plan file from yesterday.
- Double right arrow—View a plan file for the next day (assuming you are not on the current day).
- Double arrow with a bar—View the most recent plan file.
- Drop-down list—Choose a different plan file based on its date and time from the date drop-down list.
- Calendar icon—Show a calendar wherein you can navigate to the desired month, and choose the desired date.

Set Plot Options

To modify the visual attributes of the circuits, click **Plot Options**. Options are retained for future sessions. To restore values to the system defaults, click **Reset to Defaults**.

The circuit width is based on its capacity (the higher the capacity, the wider the circuit). Using these options, you can set the minimum width and the capacity it represents, as well as how fast the width grows as the capacity grows.



CHAPTER 13

Performing Administrative Tasks

This chapter is for administrators who configure and manage the WAE Live application.

To configure WAE Live, it is assumed that you already have:

- An admin role—You know that you have admin privileges when you see the admin icon in the upper-right corner of the WAE Live Home page after logging in.
- Installed the Cisco WAE Live server, including the mld data store. For information on installing the WAE Live mld data store, see the *Cisco WAE Server Installation Guide*.
- Collected WAE Live data from plan files, SNMP, or other collection methods, such as a snapshot.
- The appropriate license.

This chapter assumes that you are familiar with the CLI commands to configure the WAE Live server and to troubleshoot operations on your WAE server installation.

Note $CARIDEN_ROOT is the directory in which the software is installed. On Linux, the default $CARIDEN_ROOT is /opt/cariden. $CARIDEN_HOME is the directory in which the WAE Design, and WAE Live executables and binaries are installed. On Linux, the default is /opt/cariden/software/Live/current .

This section contains the following topics:

- About Network Data, on page 79
- About Advanced Configuration Settings, on page 80
- Troubleshooting, on page 80
- Multi-Network Management, on page 80
- Data Source Settings, on page 83
- Database Management, on page 85
- Explore and Analytics Settings, on page 86
- Map Settings, on page 91

About Network Data

WAE Live collects data from WAE Server, which can include the creation and addition of a simulated demand traffic matrix. The data is stored in *plan files* . WAE Live uses two internal mechanisms, a time-series optimized

data store and an internal Map archive, to store this data and make it available for use in the application. The Map component accesses the Map archive and shows the visualization of these plan files in a weathermap layout. The Explore and Analytics components access the data store.

Note Cisco recommends using only one network in WAE Live to pick plan files.

If a newly discovered node cannot be assigned to an existing site, WAE Live creates an extra node site and places the unassigned node there.

About Advanced Configuration Settings

Access: **WAE Live > Settings > General** tab

The General tab includes an advanced configuration (Advanced Config) link. Clicking this link displays data configuration details (config.xml) for your WAE Live server. While you can edit this file, be advised that any changes you make could adversely impact your WAE Live setup. We recommend that you consult with your support representative for assistance before changing any properties or parameters in this file.

Troubleshooting

Relevant log files are located in $CARIDEN_ROOT/logs/wae-live and $CARIDEN_ROOT/logs/wae-web-server.

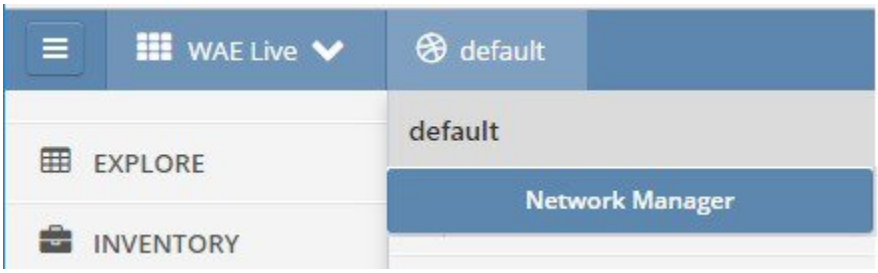
- Data collection is the process of getting the plan files from WAE Server. WAE Live then inserts this data to internal locations for use in the WAE Live application. If the data insertion process is halted, a message appears at the top of each page alerting you to this issue. To troubleshoot the source of the issue, do one of the following:
  - Navigate to **WAE Live > Settings > Data Source** table and click **Check Status** button to check the status.
  - Use the ml_insert_ctl tool. For information, see ml_insert_ctl -help.
- If you need to contact a support representative, use the mate_tech_support tool to create a tar file of support information. The output is stored in the /tmp/MATE_TS directory by default. To change the directory in which the results are stored, use the -tar-path option.

Example: This creates a tar file of support information and stores the output in the /troubleshooting directory:

```
mate_tech_support -tar-path /troubleshooting
```

Multi-Network Management

Access: **WAE Live > Network** icon (**default**) > **Network Manager**



WAE Live supports multiple networks. If only one network exists, the Network menu shows “default” as the name of the default network. If there are multiple networks, the name of the network that is in use shows.

To configure collections for a specific network, you must first create the network. After creating the network, WAE Live automatically creates a map archive. These map archives are stored in $CARIDEN_ROOT/data/mldata/archives as <network_name>-archive . $CARIDEN_ROOT is the installation directory, which is /opt/cariden by default.

After you create a network, all users have access to its data through the Explore, Analytics, and Map components.

Adding a Local Network

If you have not already run collections for a network, any newly added networks have no data to display. For information on running collections, see the *Cisco WAE User Guide*.

When you add a network, it is assumed that you want to add local network, which is your current WAE server.

Procedure

Step 1 As an admin, click the default Network icon at the top of the WAE Live Settings page. Initially, this says “default”; choose the Network Manager.

On a fresh installation with no networks added, there are empty sections for Active Networks, Inactive Networks, and Remote Network Links. If you have already added networks, this page displays all active, inactive, and network links.

Step 2 In the upper-right corner of the Network Manager page, click **Add**.

By default, the Add Network dialog box automatically uses the Local Network option.

Step 3 Enter a unique network name to identify the network from which you are collecting; for example, testnetwork1.

Step 4 (Optional) Enter a description of the network or other useful notes about the network.

Step 5 Set the status, and then click **Save**.

After successfully creating a network, the network name becomes available from the default network drop-down list and appears on the Network Manager page.

Adding a Remote Network Link

The network can reside on a separate WAE Live server. A remote network link lets you access another WAE Live server without having to remember its name.

Procedure

Step 1 As an admin, click the default Network icon at the top of the WAE Live Settings page. Initially, this says “default”; choose the Network Manager.

Step 2 Click **Add**.

Step 3 Click the **Remote Network Link** radio button.

Step 4 Enter a hostname.

Step 5 Enter a network name.

After adding the remote network link, it appears under the Remote Network Links section on the Network Manager page.

Setting Networks to Active/Offline

For local networks only, you can set their state to either Active or Offline.

- Active—Data can be collected for this network, and the network is visible to all users.
- Offline—Data cannot be collected for this network, and the network is not visible to users.

If you configure collection for an active network, set it to offline, and then later change it back to active, the collection restarts automatically.

Viewing or Editing Networks

Procedure

Step 1 As an admin, choose **Network Manager** from the list of available networks in the upper-left network link. A list of active and inactive networks opens.

Step 2 Click the network name that you want to edit.

Step 3 Enter your changes, and click **Update**.

Deleting Networks

If you delete a network and then add it back at a later time, you must reconfigure its collection.

Procedure

Step 1 As an admin, choose **Network Manager** from the list of available networks in the upper-left network link. A list of active and inactive networks opens.

Step 2 Click the network name that you want to delete.

Step 3 Click **Delete**, and then click **OK** to confirm the deletion.

What to do next

Data Source Settings

Access: **WAE Live > Settings > Data Source** tab

The Data Source tab lets you configure how often to collect network data, specify the network source for the data, and check the status of the latest data collected.

All WAE Live settings and application data, including the Data Source tab, are applicable on a per-network basis. Before configuring or making changes to any of these pages, choose the appropriate network from the Network menu (top left).

Multi-network collection is only supported in the following cases:

- The Collection Source is the 7.1 Remote Archive.
- The Collection Source is None.

About Data Collection

Data collection is the process of getting plan files; WAE Live then inserts this data to internal locations for use in the WAE Live application.

The WAE Live > Settings > Data Source tab identifies the source and status of data collection.

By default, WAE Live checks for a plan file every 30 seconds. The Insertplan.Source.Pollfrequency property in the Advanced Config feature sets this default. If none is defined, the default applies.

Data Collection Sources

Data is collected from WAE 7.1 Remote Archive and the plan is used in WAE Live.

Collecting Data

Procedure

Step 1 At the top of the WAE Live Settings page, click the default Network icon. Initially, this says “default”; choose the Network Manager.

Step 2 Navigate to **WAE Live > Settings > Data Source**.

Step 3 For the collection source, click the **7.1 Remote Archive** radio button.

Step 4 Enter the hostname.

Step 5 Enter the port details.

Step 6 Enter a username. This is the administrative username for the server.

Step 7 Enter a password. This is the password for the administrative user.

Step 8 Click **Save**.

Checking the Data Collection Status

Procedure

Step 1 Choose **WAE Live > Settings > Data Source**.

Step 2 At the top of the Data Source tab, click **Check Status**.

The Data Collection Status dialog box shows the last collected filename, the time the plan file was collected, and any source archives for the plan files.

Disabling Data Collection

Procedure

Step 1 Choose **WAE Live > Settings > Data Source**.

Step 2 At the bottom of the Data Source tab, click **None** and then **Save**.

This disables data collection and stops the scheduler for the current WAE Live network.

Inventory Configurations

This page must be configured to collect the plan files containing topology and statistics for the same network used in the inventory collection process. Inventory collection for use in WAE Live must be configured in WAE Server.

Data is visible on WAE Live when 7.1 Remote Archive which is the remote WAE network specified in the Collecting Data, on page 84 section has the inventory data collected.

Database Management

Access: **WAE Live > Settings > Database Management** tab

WAE Live allows purging of stale data from the datastore. Purging only removes data from datastore and reserves the disk space for future use. This action does not reduce OS data file footprint.

Note Purging affects only the datastore and does not delete data from the archives. To remove data from the archives, you must manually delete any files that are no longer needed.

You can set the following criteria for purging of data:

- Timestamp Limit
- Disk Space Limit
- Disk Space Reserve Ratio
- Resource Usage Time

Note The purging of data is triggered by the combination of **Resource Usage Time** and **Time Stamp Limit**. Therefore, it is important to set **Resource Usage Time** along with **Time Stamp Limit**.

The following topic explains how to set purge criteria:

Setting Purge Criteria

Set the following criteria to purge stale data from datastore. You can choose to set one or multiple criteria for purging.

Procedure

Step 1 Enter the **Time Stamp Limit**. Select unit (days, weeks, months, years) from the dropdown list.

Step 2 Enter the **Disk Space Limit**. Select unit (MB, GB, TB) from the drop-down list.

Step 3 Enter the **Disk Space Reserve Ratio** in %. Disk Space Reserve Ratio is used to limit the disk space used by datastore. This value indicates the Reserved Ratio that datastore file can grow against Disk Size Limit.

Setting the value to 10 means 10% of disk space is reserved and the datastore files can grow up to `(Disk Size Limit) * (100-10)/100`.

Step 4 Enter the **Resource Usage Time** in milliseconds. It is the time duration for which the purge job is allowed to run. For example, if you enter the value as 1800000, it means that the purge job is allowed to run for 30 minutes.

Explore and Analytics Settings

Access: **WAE Live > Settings > Explore and Analytics** tab

WAE Live supports multiple interface types that you can view on the Explore page and are used when creating reports from the Analytics pages. The Settings > Explore and Analytics tab lets you specify the default interface type used. You can override the default for the Explore pages, but not for the Analytics pages.

WAE Live acts on a set of objects (network elements) that are discovered or created prior to WAE Live accessing the plan file. For example, interfaces might be discovered by WAE Collector while demands thereafter can be added to the plan file. Each object has a set of default properties (attributes) that are tracked in and retrieved from the data store. For example, an LSP's Shortest TE Path, a node's CPU, or an interface's Traffic In are all properties that are collected and stored by default. In addition to these default objects, from the Settings > Explore and Analytics tab you can create user-defined properties to fit your specific needs.



Note These settings are configured for all networks; that is, they are not configurable on a per-network basis.

Default Interface Types

You can specify the default interface types to show in Explore and use in Analytics (reporting):

- Individual logical interface—An interface with an IP address that is not in a LAG. These names might have a decimal.

  Example: et-0/1/2.0, g2/0/0

- Individual physical interface—An interface that is not in a LAG and does not have an IP address. These interface names do not have a decimal.

  Example: so-0/2/2, g2/0/0

- LAG logical interface—An interface with an IP address that contains member interfaces.

  Examples: ae0.0, Bundle-Ether1, Port-Channel7

- LAG physical interface—The physical representation of an aggregated interface to which a LAG logical interface (such as ae0.0) is assigned.

  Example: ae0 for Juniper interfaces

- LAG member interface—An interface that is in a LAG.

  Example: g2/0/0

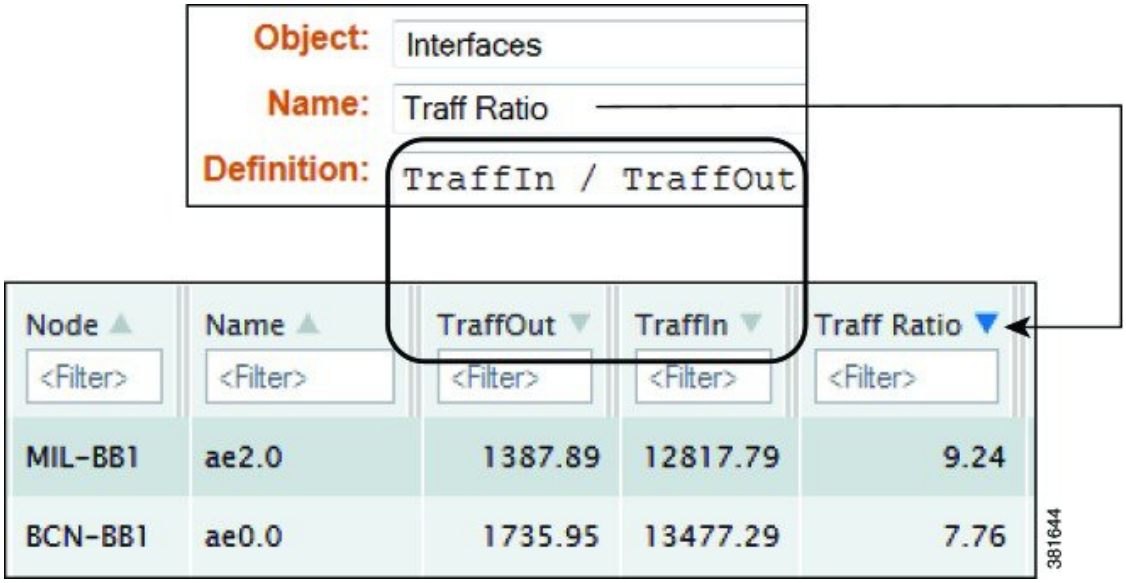
- Unknown—None of the preceding interface types.

User-Defined Properties

You can define your own properties to use when creating reports. User-defined properties derive their values from other properties. For example, rather than tracking incoming traffic (Traffic In) and outgoing traffic (Traffic Out), you might need to track the ratio of incoming to outgoing traffic (Figure 18: Example of a User-Defined Property Definition Using Simple Arithmetic Expression, on page 87).

Note For two-word properties, use one word in the definition. For example, for the Remote Node property, use "remotenode." For Traffic In and Traffic Out properties, use "traffin" and "traffout" in the definition.

Figure 18: Example of a User-Defined Property Definition Using Simple Arithmetic Expression



As with default properties, these user-defined properties are available in the Explore and Analytics pages.

Each user-defined property consists of an object, a user-specified name, and a user-created definition.

Adding a User-Defined Property

Procedure

Step 1 Choose **WAE Live > Settings > Explore and Analytics**.

Step 2 Click **Add Property**.

Step 3 From the Add New Property dialog box, choose the object type to which the newly created property applies (interfaces, interface queues, nodes, and so on).

Step 4 Enter a property name. Once the property is saved, this name appears as a selectable property for the given object in Explore and Analytics pages.

Step 5 Enter a definition value.

Step 6 Click **Preview** to view the results that the object and definition return.

Step 7 Click **Save**.

Acceptable Definition Field Values

The definition field accepts alphanumeric characters and regular expressions. Each definition is an expression using one or more default properties. The following rules apply:

- Properties used in the definition must be default properties; that is, you cannot use user-defined properties in defining other user-defined properties.
- Property names in the Definition field are not case-sensitive.
- A user-defined property can be an integer, a double (floating point number), string, or boolean.
- You can use any combination of the following in the definition. You can group any of these with parenthesis, which are executed in the order in which they appear.
  - Arithmetic expressions: +, -, *, /, and % (Figure 18: Example of a User-Defined Property Definition Using Simple Arithmetic Expression, on page 87)
  - Relational operators: >, <, =, <=, and >=
  - Logical operators: AND, OR, and NOT
  - Regular Expression Functions, on page 88
  - CASE Statements, on page 89
  - Edit (pencil icon)—Edit the user-defined property.
  - Delete (trash can icon)—Delete a user-defined property. There is no undo.

Regular Expression Functions

In the Definition field for your user-defined property, you can use regular expressions to filter results using regexp_extract and regexp functions.

regexp_extract

regexp_extract Syntax
regexp_extract(formatted_result, identifier1, pattern1, identifier2, pattern2, ...)

This searches a string with a list of regular expressions to produce a formatted string result compiled from back references. Back references are the results of capturing groups from the regular expressions and referencing them in order by $N, where N is a number that is greater than or equal to 1. This N identifies the captured string returned by a back reference in the order it appears in the matching pattern. At a minimum, only three arguments are needed for regexp_extract. You can use any number of pairs of attributes and expressions.

Example: In this example, both captured groups ($1 and $2) come from the regular expression for Description. The expression for Node does not have any capturing groups. It is used only as an additional filter. If any of the regular expressions fail to match, the overall regexp_extract fails to match.

"regexp_extract("$1 ($2)", Description, ";interconnect to ([^\(]*) \(([^\)]*)\)", Name, "^(?!.*lo)")

regexp

boolean regexp Syntax
regexp(identifier1, pattern1)

This regexp syntax returns true, false, or null.

Example: If Node starts with 'palt', then sum the Traffic In and Traffic Out values. If the Node starts with something else, or is null, no value is displayed.

case when regexp(node, '.*PALT.*') then traffin + traffout else null end

Example: If your node names include site names, this regexp syntax is useful for creating site properties for interfaces.

regexp_extract('$1', node, '([^-]+)-(.*)')

CASE Statements

The CASE statement allows conditional values based on a set of expressions. All result-expressions must be of the same type. That is, they can all be numeric or all be string, but they cannot be a mixture of both string and numeric.

CASE Statement Syntax
CASE WHEN <boolean-expression> THEN <result-expression>
[WHEN ...]
[ELSE <result-expression>]
END

Example: This example creates a user-defined property if a node name matches a remote node and if neither name is null.

case

when regexp_extract('$1', node, '(....).*\.(..)\..*') isnull then null

when regexp_extract('$1', remotenode, '(....).*\.(..)\..*') isnull then null

when regexp_extract('$1', node, '(....).*\.(..)\..*') = regexp_extract('$1', remotenode, '(....).*\.(..)\..*') then true else false end

Expression Syntax

Two syntaxes are used for the supported expressions: EBNF (Extended Backus–Naur Form) and a simpler query. Note that you can use one or multiple words for isNull and isNotNull.

- “isNull” can be represented by “is null”
- “isNotNull” can be represented by “is not null”

Simple Query

The following table shows the syntax for a simple query.

Simple Syntax
expression := value { +, -, *, /, %, IN, is [not] null, case } [expression]
when-expression := expression [{AND, OR, NOT } when-expression]
case := CASE WHEN when-expression THEN expression [ELSE expression] END
value := { object-attribute, number, string, regexp(), regexp_extract() }

EBNF

The following table shows the grammar in EBNF format.

SELECT Rule	Syntax
expression	case | logical
logical :=	relational {(OR, AND, SEMICOLON, COMMA) relational }
relational :=	additive {(<, <=, > , >=, =, ==, !=) additive }
additive :=	multiplicative {(+, -, ||) multiplicative }
multiplicative :=	unary {(*, /, %) unary }
unary :=	(!, not, ~, +, -) factor | factor isnull | factor
factor :=	NUMERIC | STRING | NULL | TRUE | FALSE | ID | regexp | regexp_extract | in_expr | '(' expression ')'
regexp :=	REGEXP '(' literal COMMA STRING ')' |
regexp_extract :=	REGEXP_EXTRACT '(' STRING COMMA reg_arg_list ')' opt-as
reg_arg_list :=	literal COMMA STRING |

SELECT Rule	Syntax
in_expr :=	IN '(' string_list ')'
string_list :=	STRING COMMA string_list STRING
case :=	CASE when when_list opt_else END
when_list :=	when when_list |
when :=	WHEN expression THEN expression empty
opt_else :=	ELSE expression | empty
opt_as :=	AS ID |
literal :=	ID | STRING

Map Settings

Access: **WAE Live > Settings > Map** tab

The Settings > Map tab lets you configure the WAE Live Map panels and set a default layout for the weathermap. These configurations apply to all networks and are not configurable for a specific network.

Map Archives and Templates

WAE Live automatically creates the template (template.pln) file from the most recently collected plan file if no template exists. These templates and plan files are stored in map archives for use in the Map component. If visual elements have not been applied, the plan file is visualized in the default, unprocessed state. To create the visual elements of a template file and save it to a map archive, use the WAE Design GUI.

About Advanced Configuration Settings

Access: **WAE Live > Settings > Map** tab > **Advanced Config** (in lower-left corner)

As is the case with the WAE server, the Map tab includes an advanced configuration (Advanced Config) link. Clicking this link displays configuration details for the Map component. While you can edit this file, be advised that any changes you make could adversely impact your Map setup. We recommend that you consult with your support representative for as**sistance before changing any parameters.**

Default Map Archive Setting

Assuming that you are using the default map archive path setting and the default installation directory /opt/cariden , the property that defines the path to the map file is:

```
"Map.ArchivePath": "/opt/cariden/data/mldata/archives/default-archive"
```

Overriding the Default Map Archive Setting

You can override the default map archive setting. The Map.ArchivePath property controls the location and map archive name. It uses the following format:

```
"Map.ArchivePath": "<map_archive_path>/
<network_name>
-archive"
```

For example, if you want to change the name of the default archive to reflect the “Acme” network’s weathermap, change this property to:

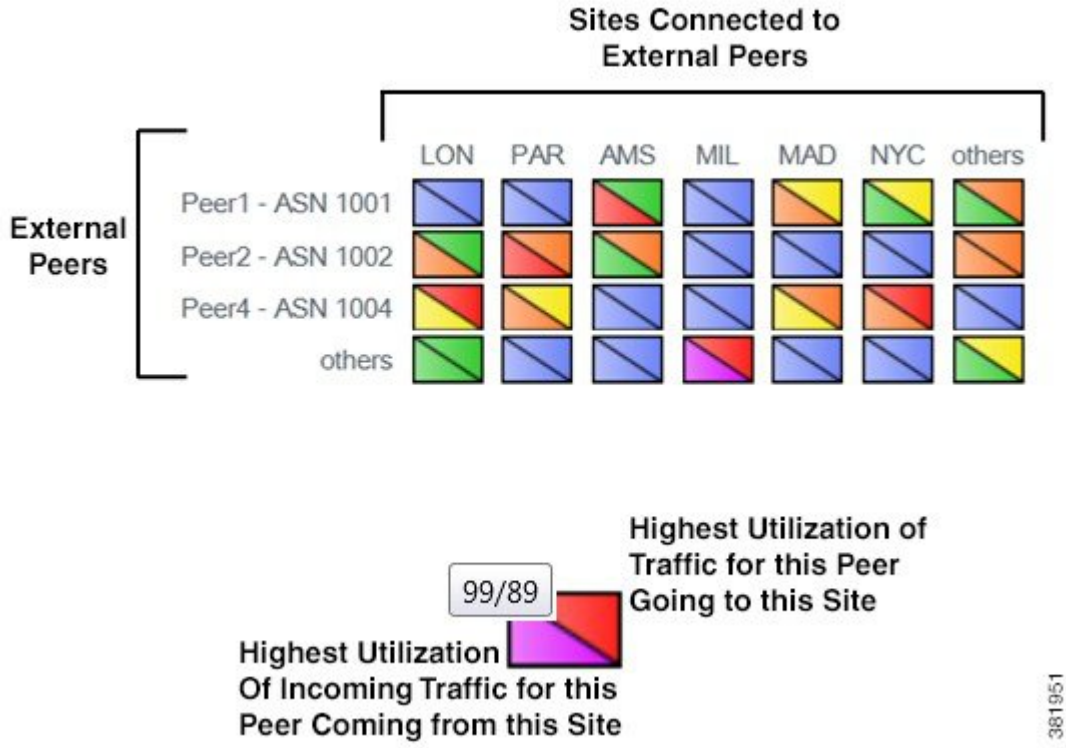
```
"Map.ArchivePath": "/opt/cariden/data/mldata/archives/acme-archive"
```

Map archives are automatically created when you make a change to the Map.ArchivePath property.

Configuring the Traffic Panel

The Traffic Widget lets you configure the Peering panel on the Map page. This panel shows the highest traffic utilization per peer and per site. The far left column lists the peers, while each top column lists the associated sites. The icons in the panel represent the traffic flowing between the two.

For this panel to have data, you must configure the collection of BGP.

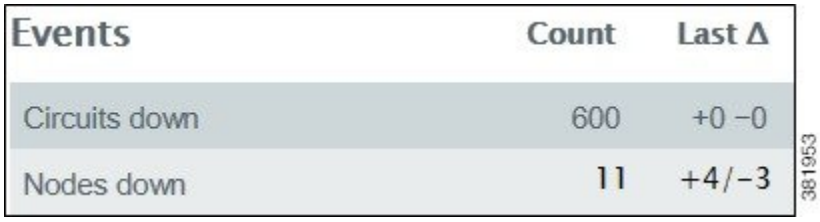


Traffic Widget Fields

- Enabled—If checked, the panel shows in the Map page.
- Title—The name of the panel to appear on the Map page.
- Rows—Configure the ASNs (external peers) used in the panel. Their names appear on the left side of the panel. Each ASN represents a row in the panel.
  - Top ___ ASNs by Traffic Utilization—The ASNs with the highest traffic utilization. Enter a whole number from 1 to 20.
  - Except—To exclude specific ASNs that would be captured by the highest utilization, enter one or more ASNs separated by commas.
  - Only These—To use only a specific set of ASNs, regardless of their traffic utilization, enter one or more ASNs separated by commas.
- Columns—Configure the sites used in the panel. Their names appear across the top of the panel. Each site represents a column in the panel.
  - Top ___ Sites by Traffic Utilization—The sites with the highest traffic utilization. Enter a whole number from 1 to 20.
  - Only These—To use only a specific set of sites, regardless of their traffic utilization, enter one or more site names separated by commas.

Configuring the Events Panel

The Events Widget lets you configure which events appear on the Map page. This panel tracks a set of current events and the number of changes since the last plan file was received.



Configuration of this panel is a combination of this UI and the default_timeplot_summary_format.txt file located in the $CARIDEN_HOME/etc/matelive directory.

For an event to appear, it must be configured in both locations. Combining the examples in this section would result in two events in the panel: Circuits down and Nodes down.

- default_timeplot_summary_format.txt file—Defines the metrics used for the events. Each metric is listed in the “Queries” section. For information on how to configure this file, see the comments at the end of the file.

Example: This is an excerpt from the default file.

```
----------Queries----------
# Circuit Down
CircuitsDown SELECT Old_Name AS Details
```

```
CircuitsDown FROM Joined_Circuits
CircuitsDown WHERE ((Old_Name IS NOT NULL) AND (Name IS NULL)) OR ((Old_Active = 'T') AND
(Active = 'F'));
#Node Down
NodesDown SELECT Old_Name AS Details
NodesDown FROM Joined_Nodes
NodesDown WHERE ((Old_Name IS NOT NULL) AND (Name IS NULL)) OR ((Old_Active = 'T') AND
(Active = 'F'));
```

- Events Widget UI—Lists the metrics and their display names. Each metric must match an event metric (query) that appears in the default_timeplot_summary_format.txt file. Example:

Metric	Display Name
CircuitsDown	Circuits down
NodesDown	Nodes down

381952

Events Widget Fields

- Enabled—If checked, the panel shows in the Map page.
- Title—The name of the panel to appear on the Map page.
- Add Metric—Click to add a new event for display on the Map page.
  - Metric—The name of the event metric that appears in the default_timeplot_summary_format.txt file.
  - Display Name—The event name that appears in the panel on the Map page.
- Edit (pencil icon)—Edit the event metric name or the display name.
- Delete (trash can icon)—Delete a metric from the UI so that it no longer appears in the Events panel. There is no undo.