



MATE Live User Guide

Release 6.1
February 2015

Cisco Systems, Inc.
www.cisco.com

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

Text Part Number:

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.

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: www.cisco.com/go/trademarks. 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. (1110R)

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.

© 2015 Cisco Systems, Inc. All rights reserved.



- Overview** 1-1
 - Sample Use Cases 1-2
 - MATE Live Menu 1-2
 - Related Topics 1-2

- Objects, Properties, and Data** 2-1
 - Raw Data 2-2
 - Related Topics 2-3

- Graphs** 3-1
 - Graph Options 3-2
 - Graph Multiple Objects 3-3
 - Graph by Individual Objects 3-3
 - Graph by Multiple Objects 3-4
 - Related Topics 3-6

- Tables and Filters** 4-1
 - Gear Tool Options 4-1
 - Filters 4-2
 - Related Topics 4-3

- Explore** 5-1
 - Example: Explore LSPs 5-2
 - Reports 5-2
 - Explore Table Options 5-3
 - Show More Data 5-3
 - Selections 5-3
 - Graphs 5-4
 - Find Related Objects 5-4
 - Gear Tool Options 5-4
 - Run Report 5-5
 - Related Topics 5-5

Inventory 6-1

Navigation 6-1

Columns 6-1

Related Topics 6-2

Configuring Reports 7-1

Use Cases 7-1

Configuring Reports 7-1

Required Report Tabs 7-1

Grouping Objects 7-3

Time Aggregation 7-4

Running Reports 7-4

Naming Reports 7-5

Applying Filters and Sorting Named Reports 7-5

Report Output 7-6

Gear Tool Options 7-6

Related Topics 7-6

Traffic Reports 8-1

Projecting Traffic 8-2

Project Traffic to Cross Thresholds 8-2

Projecting Traffic by Date 8-3

Example: Linear Trend Data Aggregation 8-4

Traffic Report Tabs 8-4

Required Tabs 8-4

Optional Tabs 8-5

Related Topics 8-5

Health Reports 9-1

Interface Health 9-1

Required Tabs 9-2

LAG Imbalance 9-2

LAG Imbalance and Risk 9-2

LAG Member Utilization 9-3

Additional Utilization, Imbalance, and Risk Data 9-4

Example 9-4

Required Tabs 9-5

LSP Health 9-5

Required Tabs 9-5

Node Health	9-5
Required Tabs	9-5
Related Topics	9-6
Deviation Reports	10-1
Tips for Viewing Graphed Deviations	10-3
Deviation Report Tabs	10-3
Required Tabs	10-3
Related Topics	10-3
Ad Hoc Reports	11-1
Time Intervals	11-2
Percentage Growth	11-3
Ad Hoc Report Tabs	11-4
Required Tabs	11-4
Optional Tabs	11-5
Example: Ad Hoc Data Aggregation	11-5
Example Ad Hoc Report	11-5
Related Topics	11-6
Reports	12-1
Schedule Reports	12-1
Reports Table	12-2
Related Topics	12-2
Report Log	13-1
Report Log Table	13-1
Related Topics	13-2
Map	14-1
Read Network Health	14-2
Traffic Utilization	14-2
Peering Health Panel	14-2
Events Health Panel	14-3
Navigate to Related Explore Information	14-4
Weathermap	14-5
Map Objects	14-6
Control Panel Options	14-6
Navigate to Specific Weathermap	14-7
Set Plot Options	14-8

[Related Topics](#) 14-8



Overview

MATE Live gives immediate and easy access to both current and historical data. Combined, the Explore, Analytics, and Map components offer flexible, interactive means of finding summarized aggregate views or quickly narrowing the search to only relevant details. Explore enables you to interactively navigate to current and historical data, and Analytics allows you to generate reports on this data. Map includes a near-realtime weathermap and health panels from which you can navigate to relevant network data.

This guide includes the following chapters.

- [Objects, Properties, and Data](#)—Lists the objects and properties that can be shipped with MATE Live, and describes the raw data that MATE Live uses.
- Common user interface tools
 - [Graphs](#)—Describes common elements of using the graphs in Explore and in the report output.
 - [Tables and Filters](#)—Describes common elements of using tables, including filtering options.
- [Explore](#)—Describes how to use Explore tables and graphs to navigate to data.
- [Inventory](#)—Describes how to view a hierarchical representation of the current hardware in your network.
- Creating new reports
 - [Configuring Reports](#)—Describes the required steps for all report creation and how aggregation works within reports. It also describes how to name reports, apply filters and sorting to named reports, and export report data.
 - [Traffic Reports](#)—Describes how to create reports to determine current and projected traffic trends, as well as project traffic.
 - [Health Reports](#)—Describes how to create reports to determine the health of network objects and whether LAGs are imbalanced.
 - [Deviation Reports](#)—Describes how to create reports that identify sudden shifts in property values.
 - [Ad Hoc Reports](#)—Describes how to create customized reports.
- Viewing, re-running, and scheduling reports
 - [Reports](#)—Describes how to view and rerun previous reports, as well as how to schedule and email them.
 - [Report Log](#)—Describes how to view and rerun reports.
- [Map](#)—Describes how to use the Map panels and weathermap to determine current network health, as well as how to navigate to more detailed network information.

Sample Use Cases

Following are a few of the MATE Live use cases across diverse groups.

- Operators—Instantly detect backbone congestion, assess currently impacted LSPs, and decide reroutes based on size and path.
- Planners—Develop trending reports for importing into MATE 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.

MATE Live Menu



Note

The availability of menu items depends on the licenses purchased.

Menu	Description
Explore	Enables you to interactively 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.
Inventory	View a hierarchical representation of the current hardware in your network.
Analytics	<ul style="list-style-type: none"> • The Configure Report page provides access to creating multiples types of reports, including health, traffic, deviation, and other user-specified (ad hoc) details. • The Reports page lists those reports that you have named and identified for frequent use. You can rerun reports, access report results, and schedule reports to run at specified dates with results mailed to you or others. • The Report Log page 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 MATE GUI.
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.
Settings	If you are logged in as an administrator, enables you to perform administrative tasks, such as adding users or creating new properties. For more information, see the <i>MATE Live Configuration Guide</i> .

Related Topics

- *MATE Integration and Development Guide*

- *MATE Live Configuration Guide*
- *MATE GUI Visualization Guide*



Objects, Properties, and Data

The Explore and Analytics components enable you to navigate to and report on information about 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.



Note

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

- **Objects**—Network elements discovered and whose information is displayed in MATE Live. Objects include interfaces, interface queues, nodes, RSVP-TE LSPs, and demands. When reporting, these objects can be grouped so that you can aggregate their properties.
- **Properties**—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 objects properties. For information, see the *MATE Live Configuration Guide*.

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

Table 2-1 MATE Live Objects and Properties

Object	Properties		
Demands	<ul style="list-style-type: none"> • Destination • IGP Metric • Maximum Latency • Name 	<ul style="list-style-type: none"> • Path • Path Metric • Service Class 	<ul style="list-style-type: none"> • Source • Tags • Traffic
Interfaces	<ul style="list-style-type: none"> • Admin Status • AS • Capacity • Description • Dropped Packets Out • Error Packets In • Function • IP Address 	<ul style="list-style-type: none"> • Metric • Name • Node • Operating Status • Packets In • Packets Out • Remote AS • Remote Interface 	<ul style="list-style-type: none"> • Remote Node • Tags • Type* • Traffic In • Traffic Out • Util In • Util Out <p>* The Type property appears as a number 1-6. For a description of each of these interface types by number, see the Explore chapter.</p>
Interface Queues	<ul style="list-style-type: none"> • Error Packets In • Error Packets Out • Interface Queue 	<ul style="list-style-type: none"> • Node • Packets In • Packets Out 	<ul style="list-style-type: none"> • Traffic In • Traffic Out
LSPs	<ul style="list-style-type: none"> • Actual Path • Admin Status • Name • Operating Status 	<ul style="list-style-type: none"> • Remote Node • Setup BW • Shortest TE Path 	<ul style="list-style-type: none"> • Source Node • Tags • Traffic
Nodes	<ul style="list-style-type: none"> • CPU0 • CPU1 • Function • IP Address • Mem0 	<ul style="list-style-type: none"> • Mem1 • Model • OS • Name 	<ul style="list-style-type: none"> • 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 MATE Live. When a property value changes, a new raw data point is stored.

For instance, 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.

**Note**

All traffic is displayed in Mbps.

Related Topics

- [Graphs](#)
- [Tables and Filters](#)
- [Explore](#)
- [Configuring Reports](#)
- [Map](#)



Graphs

Access: MATE Live->Explore and report outputs. Click the graph icon associated with the table entry. To graph all objects in a table, click the double-graph icon that is above all the individual graph icons.



Graph a single table entry.




Graph all objects in a table.

Graphs show properties over time, enabling you to 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. [Figure 3-1](#) shows an example of a single Explore graph of an interface's Traffic In and Traffic Out over the last week.

Note that these graphs are available both through the Explore pages and as part of the output generated by most reports.



Figure 3-1 Example Single Graph of Interface Traffic In and Traffic Out Properties



381603

Graph Options

- 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.
- Graph Mode—When graphing multiple objects, change the graph mode. For more information, see the [Graph Multiple Objects](#) section.
 - 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.
- Magnifying glass—Enlarge or decrease the size of the overall graph.
- Show Properties—Select which properties to show in the graph.



381610

- Time Range—Graph properties across a different time range. After changing the time range, you must click Redraw Graphs.
- View values—View values of any point in the graph by hovering your mouse over a data point.
- 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.



Graph 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.




- Individual objects
- Overlay by property
- Stack by property

Graph 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 3-2](#) is an example where the Traffic property for three LSPs is graphed using the Individual Object graph mode.

Figure 3-2 Example of Graphing Individual Objects

One Property
Per Object



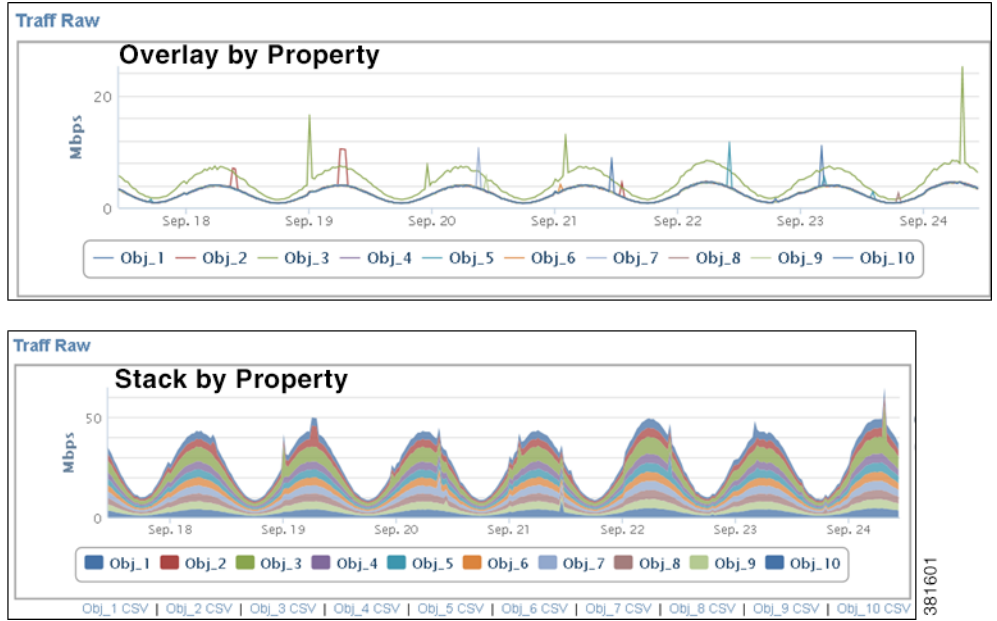
Graph by Multiple Objects

You can graph a single property for multiple objects either by overlaying them or by stacking them. [Figure 3-3](#) shows an example of each using the same data (Traffic property, time frame, and LSPs) as was generated for [Figure 3-2](#).

- **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.

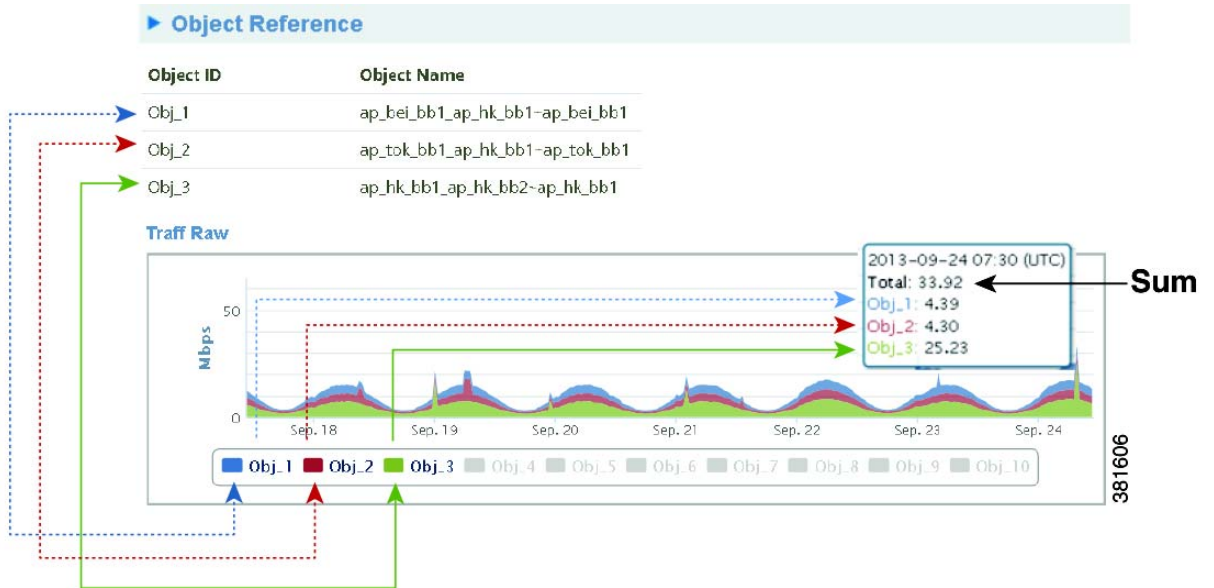
Figure 3-3 Example of Overlay and Stack Graphs

One Property Per Multiple Objects



MATE 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 3-4). 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 3-4 Reading Multi-Object Graphs



Related Topics

- [Objects, Properties, and Data](#)
- [Tables and Filters](#)
- [Explore](#)
- [Configuring Reports](#)
- [Map](#)



Tables and Filters

Access: MATE Live->Explore and report outputs

Tables in the Explore page enable you to view and filter objects based on their current property values, as well as graph property values (raw data) over time and navigate to related objects. Additionally, all reports generate output in a table format.

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

Gear Tool Options



To select any of the following options, click the gear tool.

Common

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

Explore Page

- **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 so that you do not show duplicate traffic statistics. This option is applicable only on the Interfaces page.



Note

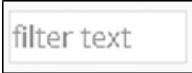

If you select specific interface types from the Explore Interfaces page for use in a report, then the report is only on those interfaces of those selected types. If you select Interfaces as the Object from a Configure Report page, then the report is run on all interfaces regardless of their type.

Report Output

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

Filters

Filters are available in most of the tables.

- Filter text—Filter one or more properties to narrow the list of objects you are analyzing or reporting. 
- Clear filter—To clear all filters in all columns, click the Clear All Filters icon. To clear one of multiple filters, remove the text from the column in which the filter appears. 

- Filter values—Enter a value in one or more <filter text> fields, and press Enter (Return).
 - For non-numeric properties, you can filter using these methods.
 - Partial or full strings
 - Regular expressions
 - isNotNull or isNull

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.

Example regex: In this example, the 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 following regular expression.

```
to=north|to=south$
```

Example regex: In the Description field, to find all descriptions that do not contain the words “test,” “trial” or “alpha”, enter the following regular expression.

```
!~test|trial|alpha
```

If, however, you want to search on descriptions that contain any one of these three fields, enter the following.

```
test|trial|alpha
```

- 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.

Related Topics

- [Objects, Properties, and Data](#)
- [Graphs](#)
- [Explore](#)
- [Inventory](#)
- [Traffic Reports](#)
- [Health Reports](#)
- [Deviation Reports](#)
- [Ad Hoc Reports](#)



Explore

Access: MATE Live->Explore

The Explore component enables you to interactively navigate current network data and retrieve historical object performance data. For example, you can use Explore to do the following.

- Identify which LSPs are currently off their shortest TE paths.
- Determine which LSPs currently are 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.

The Explore page ([Figure 5-1](#)) enables you to view and filter objects based on their current property values, as well as graph property values (raw data) over time and navigate to related objects.

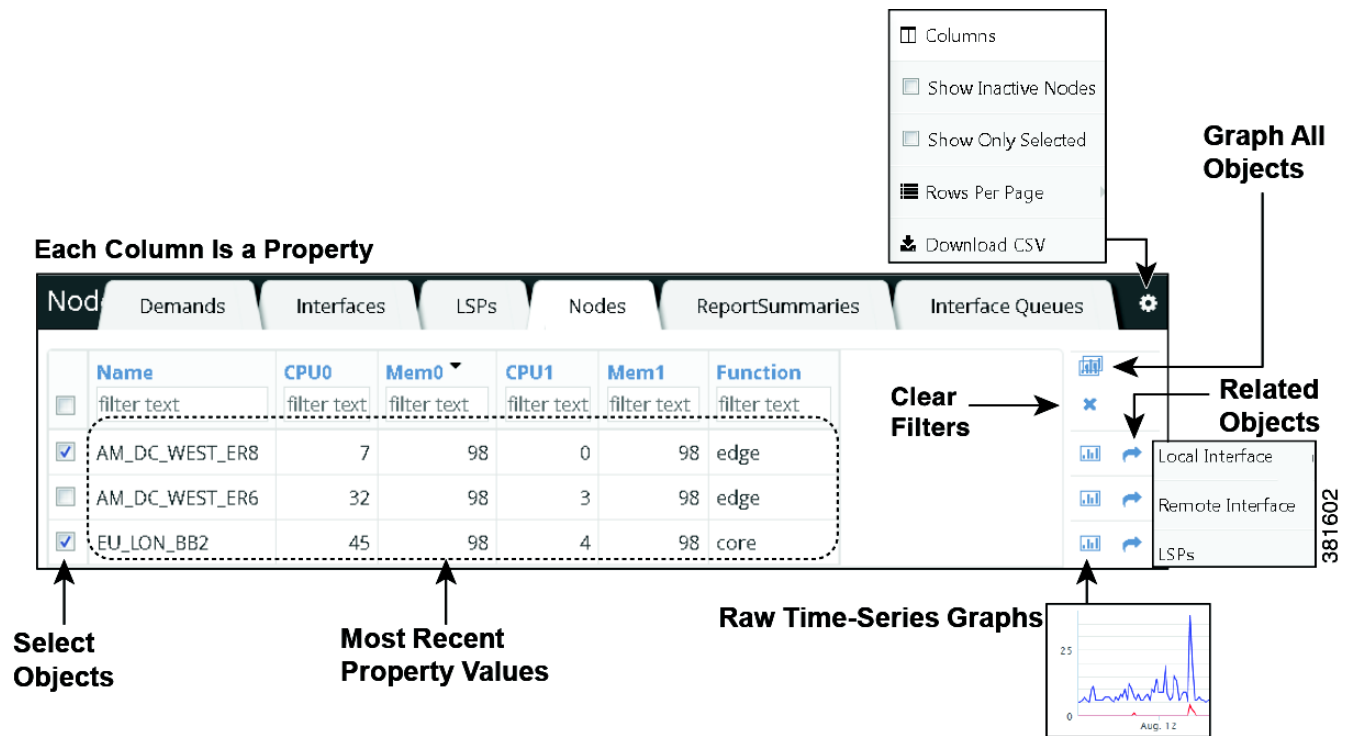
Additionally, the Explore page enables you to view summaries of the named reports over a period of time. Interactions with this report summary is the same as with the objects. See the [Reports](#) section.



Note

For instructions on how to use graphs and tables, see the [Graphs](#) and [Tables and Filters](#) chapter, respectively.

Figure 5-1 Example Explore Interfaces Table



Example: Explore LSPs

This example shows how 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 Gbps of traffic. It then shows how to find the source node of the LSP with the highest setup bandwidth with these parameters.

Actual Path	Traffic
FRA-BB2.ae0.0	>2000

- Step 1** On the Explore LSPs page, enter “FRA-BB2.ae0.0” as the filter in the Actual Path column.
- Step 2** In the Traffic column, enter “>2000” as the filter, and then press Enter (Return).
- Step 3** Click the Setup BW column heading to toggle the list in descending order.
- Step 4** Click the Related Information icon (arrow) of the LSP with the highest setup bandwidth.
- Step 5** Select 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.

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. Only named reports and only the most recently run reports are recorded.

This high-level report trend can be an indicator of the network's health. For instance, 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.

Example: This 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

Explore Table Options

Show More Data

Cell information that is blue or cell data that has a small arrow by it can be expanded to show more data. Following are examples.

- Count column—Click a blue number to see a list of objects included in that count value.
- 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.
- Demand Path column—Click the right arrow of a path in a Demands table to see an expanded list of individual hops. Click any hop to show it in the Interfaces table.

Actual Path
filter text
▼ 4 hops
— {AP_HK_BB1 ae0.0}
— {AP_HK_BB2 so-0/2/1.0}
— {AP_TOK_BB1 xe-0/2/0.0}
— {AP_SYD_BB1 ae0.0}

Selections

- Per-row check boxes—Select 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 information on report filters, see the [Configuring Reports](#) chapter.

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

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

Graphs

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



383330

For more information, see the [Graphs](#) chapter.

Find Related Objects

To find objects that are related to an object of interest, click the related objects (arrow) icon associated with the object of interest, and then select the information you need from the list. The Explore page changes to list the related information.



Example: To find destination node information for a specific LSP, click this icon associated with that LSP and select Destination Node.

Gear Tool Options



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

- Columns—Click to select 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.

If you display the Type column, the property shows as number. Following is a description of each of these interface types and the number associated with them in the Type column.

- (1) 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
- (2) 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
- (3) LAG logical interface—An interface with an IP address that contains member interfaces.
Examples: ae0.0, Bundle-Ether1, Port-Channel7
- (4) 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
- (5) LAG member interface—An interface that is in a LAG.
Example: g2/0/0
- (6) Unknown—None of the above interface types

- 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—Select the number of objects per page to show in the table.
- Download CSV—Export the table in a tab-delimited format.
- Types—Specify the interface type so that you do not show duplicate traffic statistics. This option is applicable only on the Interfaces page.

**Note**

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

Run Report

Select the objects on which you want to run a report, and then click the Run Report button. The Configure Report page appears. For more information, see the [Configuring Reports](#) chapter.

Related Topics

- [Objects, Properties, and Data](#)
- [Graphs](#)
- [Explore](#)
- [Configuring Reports](#)



Inventory

Access: MATE Live->Inventory

The Inventory 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 enables you to search for these interconnections. For instance, you can find a specific chassis by filtering on its name or you can navigate to connected interfaces.

One key use case is for determining if there are any empty slots available to accommodate future growth. For instance, if you need to add five 10-GB interfaces on a specific router (node), you can determine how many, if any, empty module and port slots you have. This can further help you plan for future hardware expenses.

Navigation

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

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

Description: PLIM Optics Port Slot
Name: 3
Node: 192.168.244.0/24:BgpPsn
ParentId: 7490160
SNMPId: 55204457
SerialNumber:
Type: ModuleSlot

3

383333

Columns

- Node—Router name, such as 168.243.0/24:BgpPsn.
- Chassis—A descriptive or exact name of the router type, such as ASR9000.
- Slot—The number of slots within the chassis that contain module slots.
- Module Slots—The number of module slots in which modules can reside.
- Module—The name of the module, such as 10-port 100-Gbps LSR. A module is also sometimes called a *linecard* or *PIC*.

- Port Slot—The number of ports slots in which physical interface ports can reside.
- Port—The port name, such as GigabitEthernet0/1/0/7.
- Transceiver—The name of the transceiver connected to the port, such as DWDM Xenpaks.
- Interface Admin Status—Administrative status (Up or Down) of the interface connected to the port.
- Interface Description—Optional description of the interface connected to the port.

Related Topics

- [Tables and Filters](#)



Configuring Reports

Access: MATE Live->Analytics, click New Report

Access: MATE Live->Explore, select objects, and click Run Report

This chapter describes the following.

- [Configuring Reports](#), including [Required Report Tabs](#) and [Running Reports](#).
 - [Grouping Objects](#) describes the options for grouping objects in the report results.
 - [Time Aggregation](#) describes this feature, which is available in Traffic and Ad Hoc reports.
 - [Naming Reports](#) and [Applying Filters and Sorting Named Reports](#) are applicable to reports that you want to schedule or email.
- [Report Output](#), including how to use the [Gear Tool Options](#).

Use Cases

The Analytics pages enable you to generate traffic, health, ad hoc, and deviation reports. These report tools give you a series of options, such as the following, that enable you to customize the reports.

- On which objects are you creating the report?
- What are the properties on which you are reporting?
- What is the time frame for the report? Do you need to project a trend out farther than that time?
- Are you reporting on individual or grouped objects?
- Are you applying time aggregation operations, and if so, what are they (for example, P95 or average)?

Configuring Reports

Required Report Tabs

- Report Type—The type of report to be run.
- Object—The object type on which the report is generated.

- **Filter**—The subset of objects on which you are reporting. For instance, if the object is an LSP, you can filter to report on all LSP destinations, or you can further refine the report to one or more specific destinations.

You can specify filters from the drop-down lists, or you can base them on selections in the Explore tables, as follows.

- Select the property from the Latest Value list, and then refine it using the second list. This enables you to run the report on all objects based on the selected property or on all objects matching the refined property value.

The screenshot shows a filter configuration interface. At the top, two arrows point to the interface: 'Select Property From List' points to the first dropdown menu, and 'Refine to Specific Value' points to the second dropdown menu. The interface contains a radio button labeled 'Latest value of', followed by a dropdown menu containing 'Remote AS', the text 'is', another dropdown menu containing '64512', and a small vertical label '381634' on the right side.

- Select the properties of only those objects that are currently selected in the associated Explore page. For instance, if you have three LSPs selected in an Explore page, the filter defaults to those three selected LSPs.

Optionally, you can click the Edit link to change these selections, which takes you to the Explore page. For more information, see the [Explore](#) chapter.

The screenshot shows the same filter configuration interface as above, but with a second radio button selected, labeled 'Selected in Explore (23 Objects) Edit...'. Below the interface, two arrows point upwards: 'Use Current Explore Selections' points to the selected radio button, and 'Edit Explore Selections' points to the 'Edit...' link. A small vertical label '381643' is on the right side.

**Note**

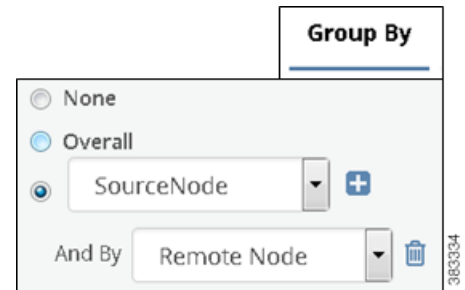
If you select specific interface types from the Explore Interfaces page for use in a report, then the report is only on those interfaces of those selected types. If you select Interfaces as the Object from a Configure Report page, then the report is run on all interfaces regardless of their type.

- **Report Time Range**—The time frame for which the report is generated. Select the time range using one of these methods.
 - Select the range from hours to years and enter a value, for example, 1 month. This option captures a range from the current time backwards.
 - Select a custom date range. You can enter the range in yyyy-mm-dd hh:mm:ss format, or you can select dates from both calendars.

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.

- Overall—Groups all objects in the report into one row of data. For instance, 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 select 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.



Example: Figure 7-1 shows a report generated three times: one without grouping objects, one with objects grouped by nodes, and one grouping first by nodes and then by metric. Notice that both the Capacity and the Latest Traffic Out (P95) columns adjust according to whether and how the objects are grouped.

Figure 7-1 Example Effects of Grouping Objects

No Grouping

Node	Name	Capacity	Latest Traffic Out (P95) Trend Value
filter text	filter text	filter text	filter text
EU_AMS_AMSIX	TenGigE0/1/0	10000	700.24
EU_AMS_AMSIX	TenGigE0/0/3	10000	698.84
EU_AMS_AMSIX	TenGigE0/0/2	10000	698.26
EU_AMS_AMSIX	TenGigE0/0/1	10000	792.38
EU_AMS_AMSIX	TenGigE0/0/0	10000	699.00
EU_AMS_AMSIX	TenGigE0/0/4	10000	791.23

Group Objects by One Property

Node	Count	Total Capacity	Latest Traffic Out (P95) Trend Value
filter text	filter te	filter text	filter text
EU_AMS_AMSIX	6	60000	4385.46
EU_AMS_AMSIX_2	15	150000	22349.19
EU_AMS_BB1	12	120000	28000.93

Group Objects by Two Properties

Node	Metric	Count	Total Capacity	Latest Traffic Out (P95) Trend Value
filter text	filter te	filter te	filter text	filter text
EU_AMS_AMSIX	123	4	40000	2797.81
EU_AMS_AMSIX	10	2	20000	1583.38

381640

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 instance, you can aggregate maximum raw data values for traffic on a weekly basis.

Time aggregation options vary depending on the report and property type. They include P95, average, maximum, minimum, last, and count. (P95 is the value 95% of the time.)

Time aggregation operations are applied to raw data. For a description of raw data, see the [Objects, Properties, and Data](#) chapter. For more information and examples of aggregating data in the reports, see the [Traffic Reports](#) and [Ad Hoc Reports](#) chapters.

Running Reports

You can run a report from any of the following locations.

- Run new reports.
 - From either the Analytics->Reports page or Analytics->Report Logs page, click the New Report button.

- From the Explore page, select objects and click the Run Report button.
- Rerun existing reports (with existing or changed parameters).
 - From the Analytics->Reports page, click View->Run, which runs immediately with the same parameters. To keep or change the parameters, click View->Edit or click the name in the Report Name cell.
 - From the Analytics->Report Log page, to keep or change the parameters, click View->Edit or click the name in the Report Name cell.

**Note**

If you need to email the report, you must give the report a unique name when you run it (see the [Naming Reports](#) section).

Naming Reports

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.

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.

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

- From the report output page, select 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 using the [Reports](#) page.

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.

-
- Step 1** Run and name the report. See the [Naming Reports](#) section.
- 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 (Return). See the [Tables and Filters](#) chapter for more details.
 - 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.



Save

Report Output

The output of all reports is in a table format that can be exported, and most reports enable you to graph the resulting data. For information on how to use tables and graphs, see the [Tables and Filters](#) and [Graphs](#) chapters, respectively.

You can view reports from two primary locations, both accessed from Analytics.

- [Report Log](#)—List of previous reports that you can view or rerun.
- [Reports](#)—List of uniquely named reports that can be scheduled and emailed.

Gear Tool Options



From a report's output page, to select any of the following options, click the gear tool.

- [Edit](#)—Open the Configure Report page to rerun or reconfigure the report to run again.
- [Save As](#)—Name or change the name of the report.
- [Columns](#)—Click to select which report columns show in the output. If a column check box is gray, then you cannot turn it off. These non-optional properties always appear in the left-most columns of the output.
- [Download CSV](#)—Export the table in a tab-delimited format.
- [Permalink](#)—Save the report's URL.
- [Rows per page](#)—Select the number of objects per page to show in the table.

Related Topics

- [Objects, Properties, and Data](#)
- [Explore](#)
- [Creating new reports](#)
 - [Traffic Reports](#)
 - [Health Reports](#)
 - [Deviation Reports](#)
 - [Ad Hoc Reports](#)
- [Viewing, re-running, and scheduling reports](#)
 - [Reports](#)
 - [Report Log](#)
- [Map](#)



Traffic Reports

Access: MATE Live->Analytics, click New Report, and select Traffic Report Type

Access: MATE Live->Explore, select objects, and click Run Report

Traffic reports enable you to trend interface 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 the traffic is growing or shrinking.

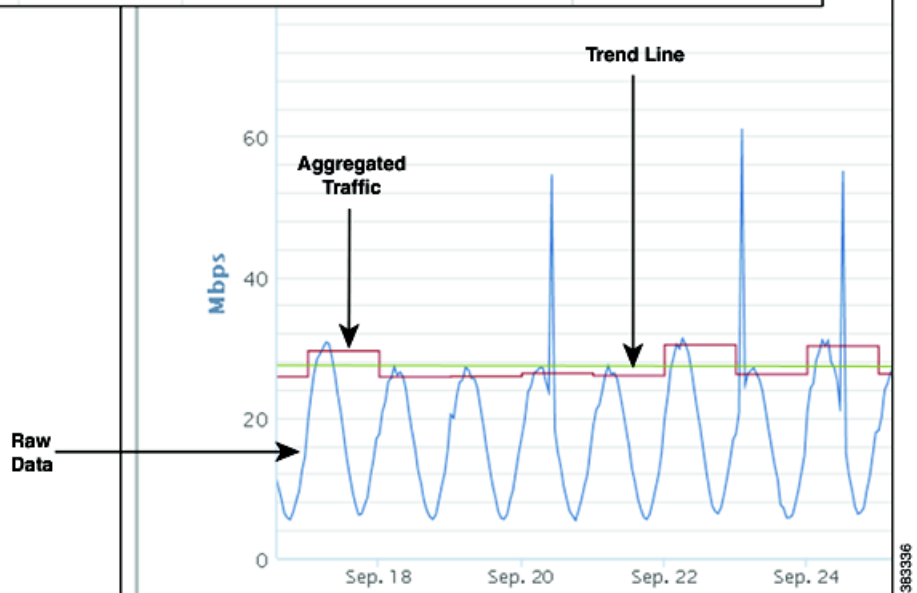
Additional options are available, as follows.

- 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 the [Configuring Reports](#) chapter.

Grouping objects is useful for aggregating traffic that can then be imported into MATE Design for use in creating growth plans. For more information, see the *MATE Integration and Development Guide*.

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 is spiking upwards with some regularity.

Name	SourceNode	Latest Traffic (P95) Trend Value	Traffic (P95) Growth
filter text	filter text	filter text	filter text
AP_HK_ER1_AP_SYD_ER3	AP_HK_ER1	27.02	-1.52%
AP_MEL_ER3_AP_BAN_ER2	AP_MEL_ER3	27.00	-1.82%
AP_BEI_ER3_AF_CAI_ER1	AP_BEI_ER3	26.99	-1.27%



Projecting Traffic

Project 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 allow you to 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. For instance, if you set 50% for 2 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.


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 8-1](#).

Example: If you set 50% for 2 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 8-1 shows an example traffic report generated to determine which interfaces will exceed 80% of their outbound traffic utilization in the next 5 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 8-1 Example Date Traffic is Projected to Pass Threshold

Node	Name	Latest Traffic Out (P95) Trend Value	Traffic Out (P95) Growth	Traffic Out (P95) Threshold Time
AP_TOK_ER8	TenGigE0/0/0	2472.83	79.53%	2013-10-12 10:57:26
ASN65200:192.168.200.2	to_EU_LON_ER6	2975.29	68.63%	2013-10-07 09:02:22
AP_HK_BB1	xe-0/0/3.0	4074.16	50.20%	2013-09-29 10:32:23




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 8-2 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 8-1.

Figure 8-2 Example Traffic Projection by Date



Example: Linear Trend Data Aggregation

If one plan file exists per day, and if the month has 28 days, then 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 would be 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. For information on how raw data is collected, see the [Objects, Properties, and Data](#) chapter.

Traffic Report Tabs

Required Tabs

- See the [Configuring Reports](#) chapter.
- Trends—How to compute the trends. See the [Trend Fields](#) section.

Trend Fields

MATE 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, 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: P95, 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. 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 more information, see the [Configuring Reports](#) chapter.

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

Example: If you select 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 trend. MATE Live determines and graphs when the trend will exceed this threshold. See the [Project Traffic to Cross Thresholds](#) section.
 - To Date—Project the traffic amount on a specified date. See the [Projecting Traffic by Date](#) section.

Related Topics

- [Objects, Properties, and Data](#)
- [Graphs](#)
- [Tables and Filters](#)
- [Explore](#)
- Creating new reports
 - [Configuring Reports](#)
 - [Health Reports](#)
 - [Deviation Reports](#)
 - [Ad Hoc Reports](#)
- Viewing, re-running, and scheduling reports
 - [Reports](#)

- [Report Log](#)



Health Reports

Access: MATE Live->Analytics, click New Report, and select Health Report Type

Access: MATE Live->Explore, select objects, and click Run Report

Health reports enable you to 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 instance, if you ran an interface health report on utilization over 80%, the resulting report would tell you the percentage of time each interface had over 80% utilization within the given time range, as well as the number of times the operating status changed. Another example is determining the number of times (Count) that a specific event occurs, such as LSP path changes.

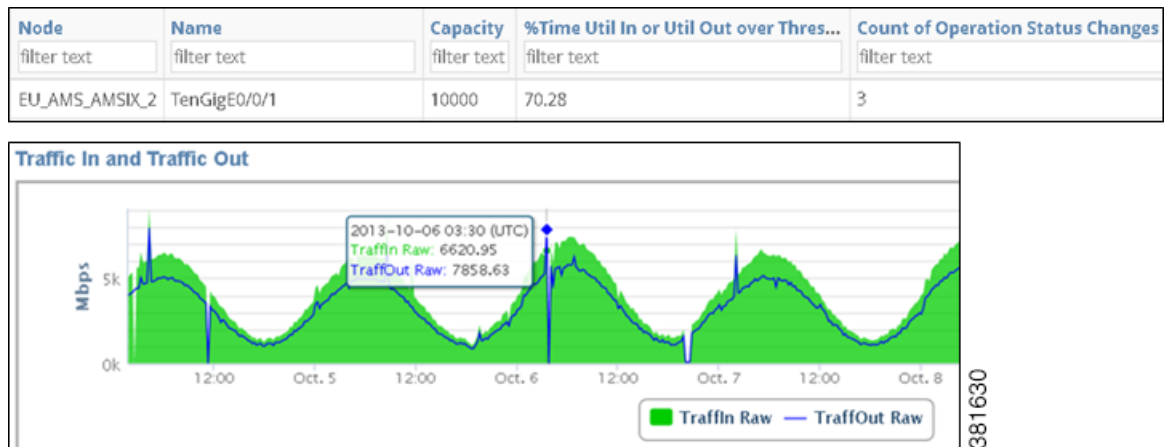
Once the report is generated, you can then graph the properties over the report time range. These graphs allow you to drill down to the specific times at which the event occurred. These graphs could indicate when a specific threshold was violated or they could narrow the threshold violations to a single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

Interface Health

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 9-1](#) 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 9-1 Example Interface Health Report



Required Tabs

- See the [Configuring Reports](#) chapter.
- Util Threshold—Threshold interface traffic must surpass to be included in the report.

LAG Imbalance

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 enables you to 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 MATE Live uses to determine whether a LAG is imbalanced or at risk. The value must be between 0 and 1.

LAG Imbalance and Risk

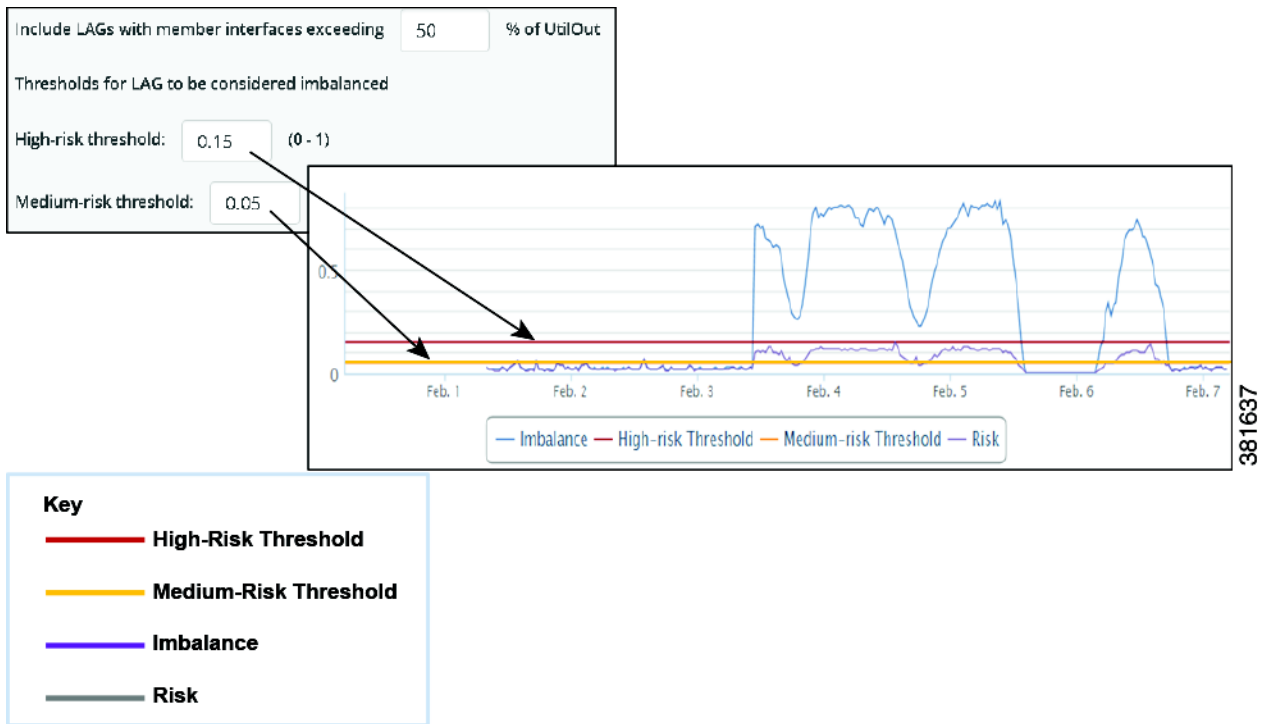
The LAG Imbalance/Risk graph shows the imbalance and risk compared to the high- and medium-risk thresholds, as well as whether and when the LAG passes these thresholds. The difference between risk and imbalance measurements is that risk calculations consider only positive differences, whereas imbalance calculations consider both negative and positive differences. For instance, [Figure 9-2](#) 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
 - $0.35 - 0.10 = 0.25$

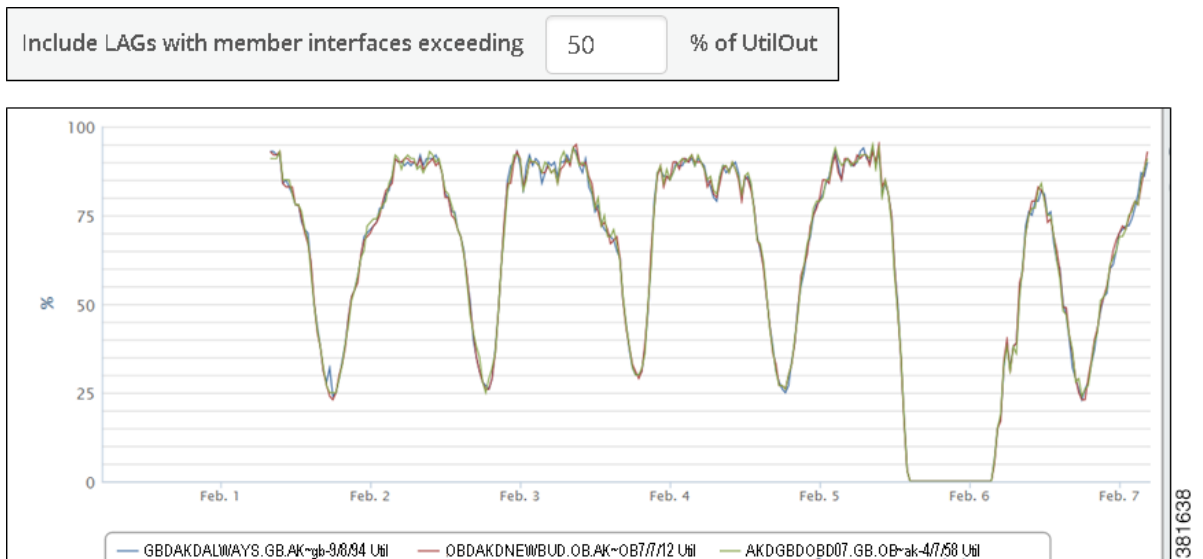
Figure 9-2 Example LAG Imbalance and LAG Risk Graph



LAG Member Utilization

The Util Out graph shows the Util Out percentage of all of the LAG’s members over the report period (Figure 9-3). 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 9-3 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 on 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

- See the [Configuring Reports](#) chapter.
- 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

An LSP health report identifies the percentage of time that the LSP traffic surpasses the user-specified threshold of setup bandwidth (set in the Setup BW Threshold tab). Additionally, it gives 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

- See the [Configuring Reports](#) chapter.
- Setup BW Threshold—Setup bandwidth threshold that LSPs must surpass to be included in the report.

Node Health

A node health report identifies the percentage of time that the CPU and memory surpass a threshold that you set in the CPU Threshold and Memory Threshold tabs, respectively.

Required Tabs

- See the [Configuring Reports](#) chapter.
- 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.

Related Topics

- [Objects, Properties, and Data](#)
- [Graphs](#)
- [Tables and Filters](#)
- [Explore](#)
- [Creating new reports](#)
 - [Configuring Reports](#)
 - [Traffic Reports](#)
 - [Deviation Reports](#)
 - [Ad Hoc Reports](#)
- [Viewing, re-running, and scheduling reports](#)
 - [Reports](#)
 - [Report Log](#)



Deviation Reports

Access: MATE Live->Analytics, click New Report, and select Deviation Report Type

Access: MATE Live->Explore, select 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 instance, 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.

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 MATE 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. So, for example, if you collect every 5 minutes and you want the moving average to be based on the last 24 hours, you would 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 10-1](#) for an example.

You can also set a threshold under which property values are not considered for the deviation calculation, thus filtering out non-significant 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. For example, if the “Ignore Deviations Under” entry is 5, and if at a given collection the upper band is 1,000 Mbps and the lower band is 500 Mbps, data would be dropped if it were larger than 1,005 or lower than 495 Mbps.

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

Example: Figure 10-2 shows an example of a report identifying two deviations. Figure 10-1 shows the configuration and an enlarged section of its output.

Figure 10-1 Deviation Report Configuration and Output

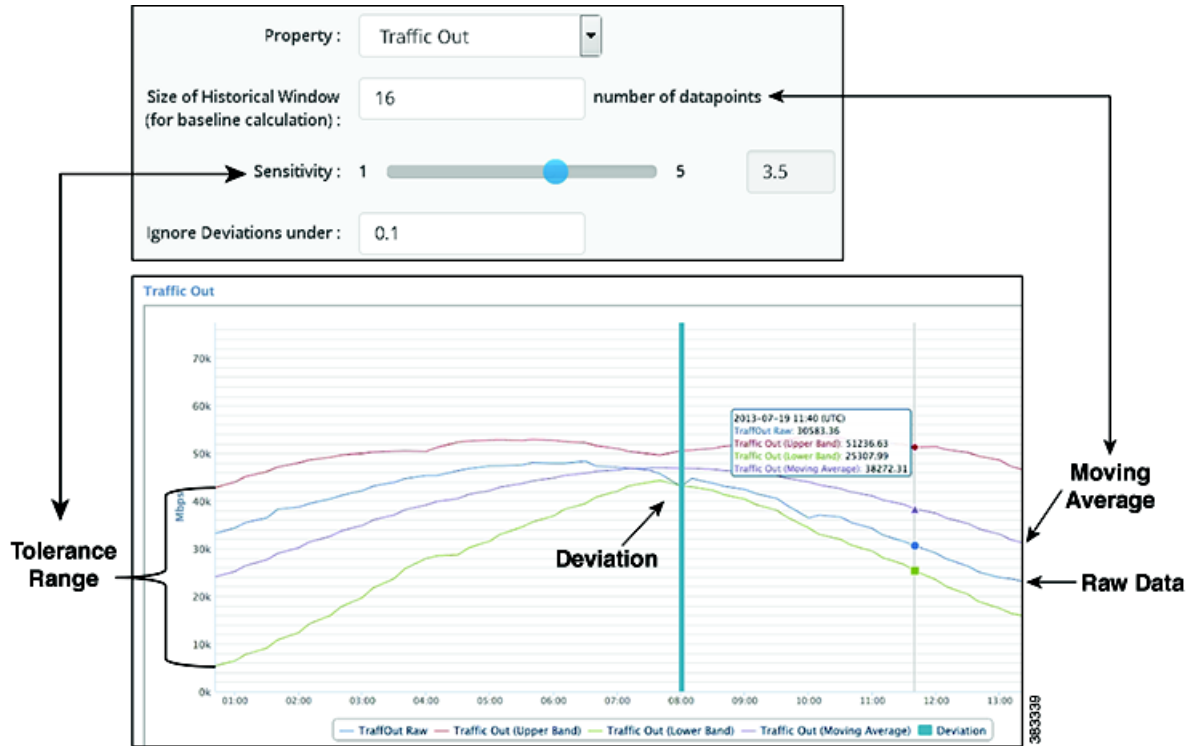
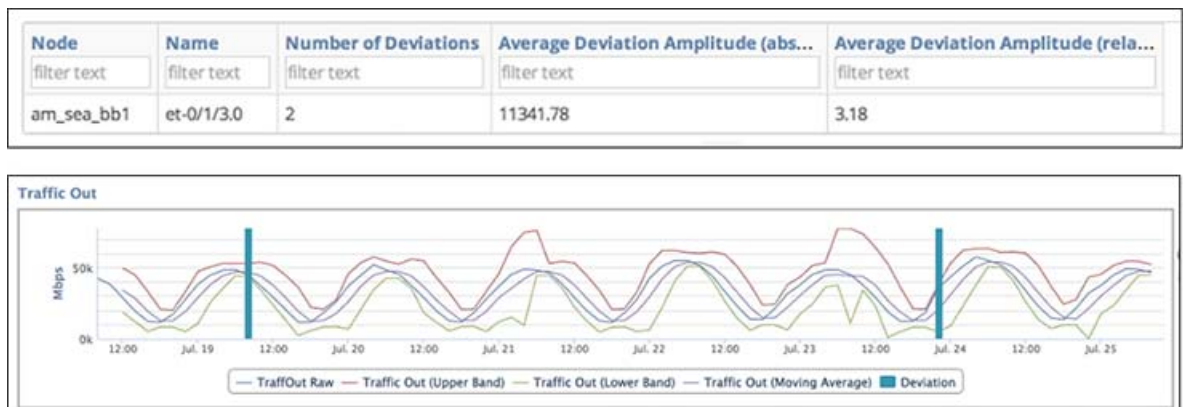


Figure 10-2 Example Deviation Report



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

- See the [Configuring Reports](#) chapter.
- 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.

Related Topics

- [Objects, Properties, and Data](#)
- [Graphs](#)
- [Tables and Filters](#)
- [Explore](#)
- Creating new reports
 - [Configuring Reports](#)
 - [Traffic Reports](#)
 - [Health Reports](#)
 - [Ad Hoc Reports](#)
- Viewing, re-running, and scheduling reports
 - [Reports](#)
 - [Report Log](#)



Ad Hoc Reports

Access: MATE Live->Analytics, click New Report, and select Ad Hoc Report Type

Access: MATE Live->Explore, select 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.

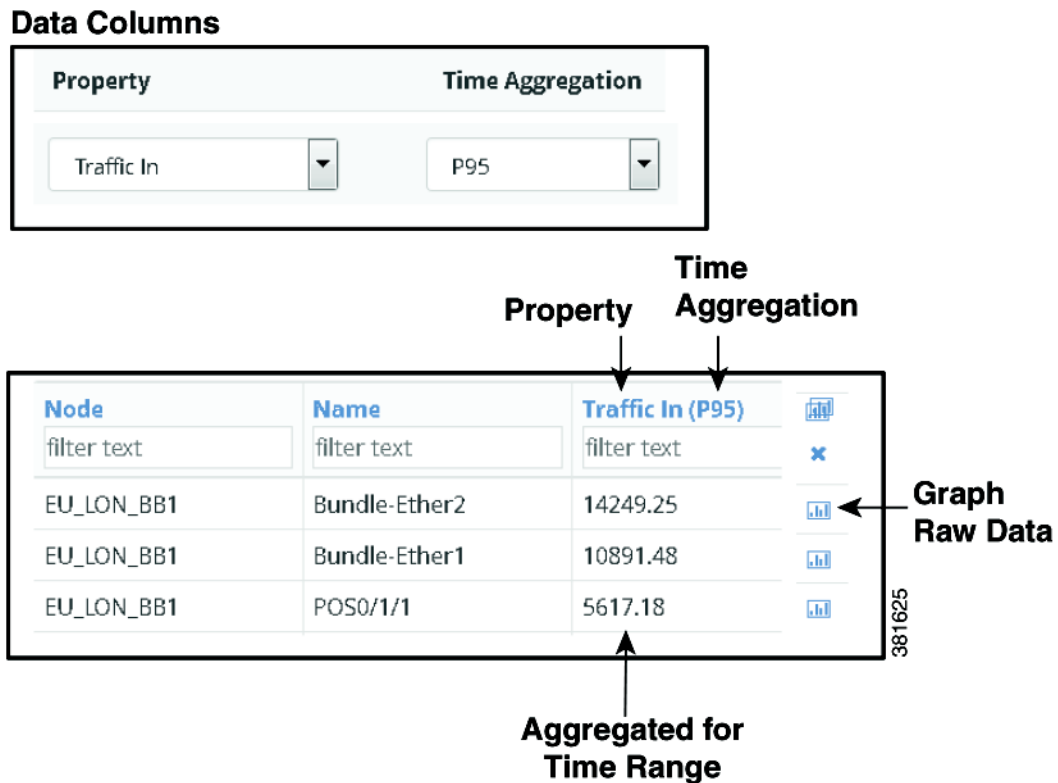
The basic ad hoc report creates a table of data in which each table value is the result of the time aggregation operation applied to the property over the entire report time range. Each row is an object. Each column is defined by a combination of properties and time aggregations (both of which are selected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

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

- **Change Count**—Number of times the value changed. For example, if you select 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.
- **P95, 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.

Example: [Figure 11-1](#) 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 11-1 Example Ad Hoc Report



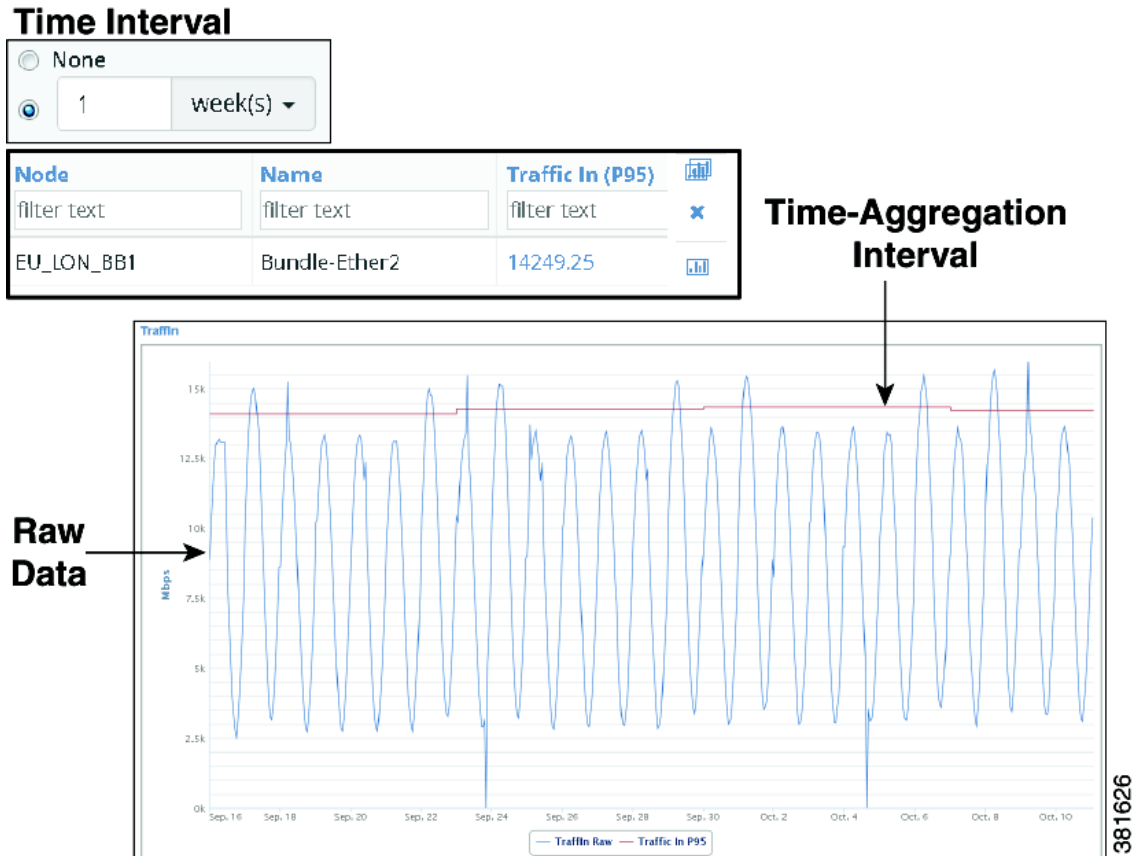
Time Intervals

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 11-2 extends the report created in Figure 11-1. 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 fluctuates, the traffic is very steady on a per-week basis.

Figure 11-2 Example Ad Hoc Report with Graphed Time Aggregations Intervals



Percentage Growth

Ad hoc reports enable you to 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 11-3 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 11-3 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

363341

Ad Hoc Report Tabs

Required Tabs

- For Filter and Time Range, see the [Configuring Reports](#) chapter.
- 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 the [Time Aggregation Options](#) section. If you specify Percent Growth, you must further identify how you want this growth calculated. See the [Percent Growth Options](#) section.

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.
- P95, 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

MATE 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: P95, 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 the [Configuring Reports](#) chapter.

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 the [Time Intervals](#) section.

Example: If you select Traffic as an LSP property in the Data Columns section and P95 as the Time Aggregation operation, and if you select 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 the [Objects, Properties, and Data](#) chapter.

Example Ad Hoc Report

[Figure 11-4](#) 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 11-4 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



Related Topics

- [Objects, Properties, and Data](#)
- [Graphs](#)

- [Tables and Filters](#)
- [Explore](#)
- Creating new reports
 - [Configuring Reports](#)
 - [Traffic Reports](#)
 - [Health Reports](#)
 - [Deviation Reports](#)
- Viewing, re-running, and scheduling reports
 - [Reports](#)
 - [Report Log](#)



Reports

Access: MATE Live->Analytics

Access: MATE Live->Analytics->Report Log, click Show Reports

The Reports page is the page from which most report management occurs. It lists all reports that have been assigned names.

From this page you can view when the report was last ran, rerun the report, change its name, schedule it to run at specific intervals or dates, and have the report automatically emailed.

Schedule 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 you 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.

-
- Step 1** 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 the [Configuring Reports](#) chapter.
 - Step 2** From the Reports page, either select View->Scheduling or click the Scheduled link associated with the report you want to schedule. The Schedule dialog box appears.
 - Step 3** To turn on scheduling, select 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, select None.
 - Step 4** To email scheduled reports, enter one or more email addresses separated by commas.
 - Step 5** 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), select “non-empty report.”
 - Step 6** Click Save.

Reports Table

- **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.
- **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, select it and then click Select.
- **Scheduled**—Frequency of how often the report is automatically run. Click the link to schedule the report to run, to email the report, or to stop running the report if already scheduled. See the [Schedule Reports](#) section.
- **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 enables you to schedule the report to run, to email the report, or to stop running the report if already scheduled. See the [Schedule Reports](#) section.
 - **View->Delete**—Remove the report from this Reports page. The report remains listed on the Report Log page.

Related Topics

- [Tables and Filters](#)
- [Creating new reports](#)
 - [Configuring Reports](#)
 - [Traffic Reports](#)
 - [Health Reports](#)
 - [Deviation Reports](#)
 - [Ad Hoc Reports](#)
- [Report Log](#)



Report Log

Access: MATE Live->Analytics, click Show Report Log

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

The reports appear by default names unless you renamed them.

Like other Explore pages, you can filter and sort the Report Log table, as described in the [Tables and Filters](#) chapter.

Report Log Table

- 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 MATE 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.

Report Name ▲	Job ID	Definition ID
filter text	filter text	filter text
Int_Traff_Over_80	13	13
Int_Traff_Over_80	15	13
Int_Traff_Over_80	16	13

**Same Report
Different Job ID**

**Same Report
Same Definition ID**

381612

- View options
 - View—Click to view the report output.
 - View->Edit—Open the Configure Report page with the parameters selected that were used to create the report.

- View->Download CSV—Export the table in a tab-delimited format.

Related Topics

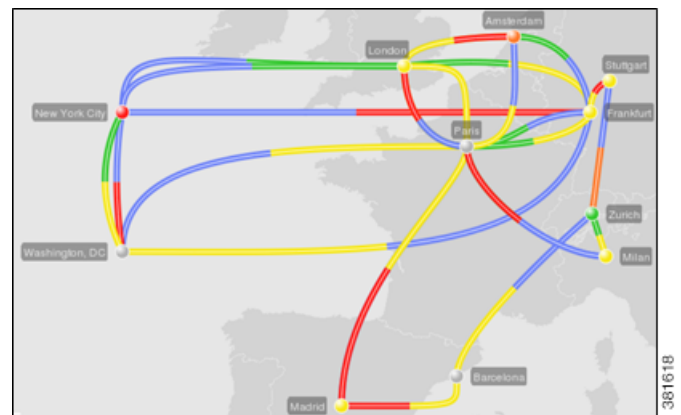
- [Tables and Filters](#)
- [Creating new reports](#)
 - [Configuring Reports](#)
 - [Traffic Reports](#)
 - [Health Reports](#)
 - [Deviation Reports](#)
 - [Ad Hoc Reports](#)
- [Reports](#)



Map

Access: MATE 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. Since 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.



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

Two panels, Peering and Events, offer immediate information on network health. The Peering panel enables you to 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, enabling you to determine how many there are total and how many have changed since the last snapshot was taken.



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 *MATE Visualization Guide*.

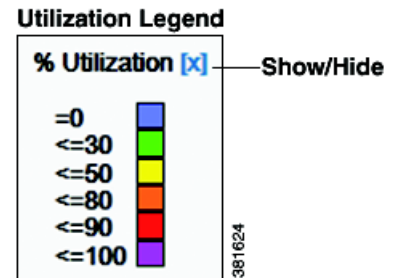
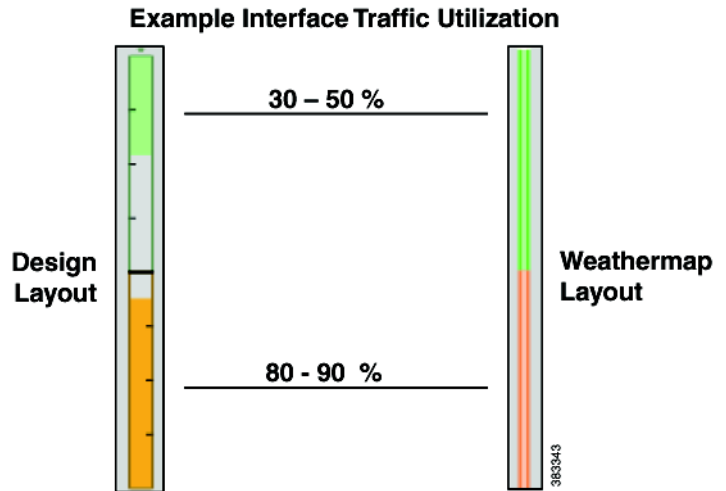
Read Network Health

Traffic Utilization

MATE 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 circuits¹ 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 MATE GUI where the administrator creates the template for the MATE Live plan files. The utilization legend at the bottom of the weathermap identifies the template's utilization colors and their values.



Peering Health Panel

The Peering health panel (Figure 14-1) is based on BGP information. It shows the highest traffic utilization per peer and per site. This panel enables you to monitor critical peers and their growth so that you can quickly detect and overcome congestion issues.



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.

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.” For information, see the *MATE Live Configuration Guide*.

Figure 14-1 Peering Health Panel



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.

Events	Count	Last Δ
Nodes Down	11	+4/-3

**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. For information, see the *MATE Live Configuration Guide*.

Navigate 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 the [Explore](#) chapter.

Node CPU and Memory

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
<p>Click a node in the site view.</p> <ul style="list-style-type: none"> 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. 	<ul style="list-style-type: none"> 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.

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
<p>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.</p> <ul style="list-style-type: none"> • 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. 	<ul style="list-style-type: none"> • 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 (Figure 14-2) at the top enables you to 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 the [Navigate to Related Explore Information](#) section.

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 both inside and outside of sites.
- A site is a MATE 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 instance, 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, follow one of these two steps.
 - Select a site from the Site drop-down list in the control panel.
 - 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 [Traffic Utilization](#) section.

Control Panel Options

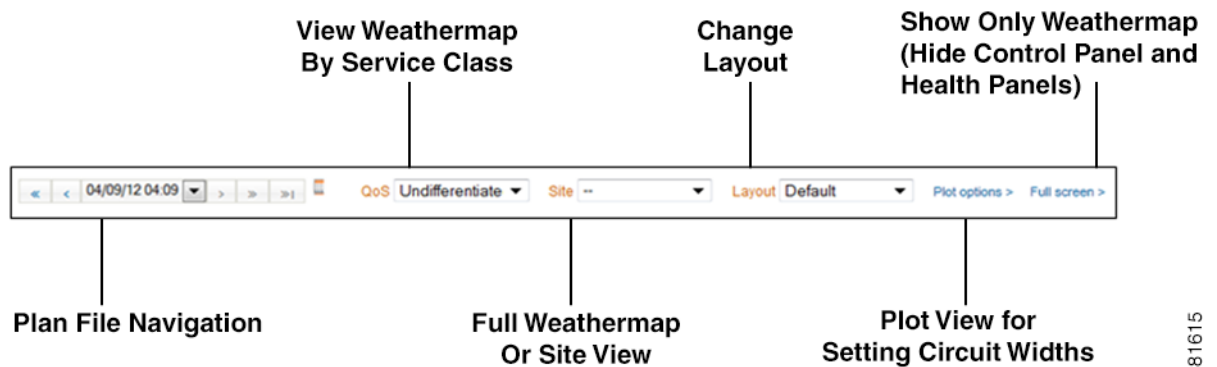
- Plan File Navigation—Click to save the file to a location of your choice. See the [Navigate to Specific Weathermap](#) section.
- QoS—Displays the type of traffic to view in the weathermap as differentiated by the QoS model. Note that if you navigate to an Explore page while displaying QoS traffic, 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 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 MATE 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 MATE Live provided you are using the same browser.

- **Plot Options**—Enables you to modify visual circuit attributes of the weathermap. See the [Set Plot Options](#) section.
- **Full Screen**—Enables you to view an enlarged weathermap without the control panel and without the health panels.

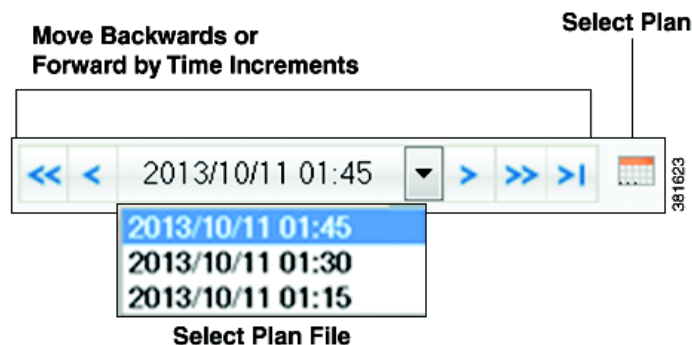
Figure 14-2 Weathermap Control Panel



381615

Navigate to 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.



381623

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—Select 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 select the desired date.

Set Plot Options

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

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.

Related Topics

- [Objects, Properties, and Data](#)
- [Explore](#)