PREDICTIVE
UC ANALYTICS

Predictive UC Analytics Web Application First Steps Guide

TeleMate.Net®
SOFTWARE

The Predictive UC Analytics Web Application Interface

The Predictive UC Analytics web interface allows administrators to monitor and analyze UC traffic from anywhere on the network without having to install the admin client. The admin client is needed for initial configuration of the product and its data sources, collection, processing, and maintenance tasks, but many of the advanced analytics features are only available in the web interface. These features include real-time trending, dashboards, monitors, viewpoint maps, and predictive KPI models.

Administrators can also configure additional logins to grant different levels of access to managers and other employees in the organization who need to be able to access that information. To avoid having to manage every privilege for every login, administrators can create custom privilege groups and assign them to logins. Administrators can also configure tenants, granting ownership of a section of the directory organization to tenant administrators. The "Admin View" feature allows administrators to act as other logins to test their security privileges.

Secure Authentication

For all of these authentication methods, we strongly recommend configuring the web server to use HTTPS (secure HTTP). The web server will not allow LDAP accounts to log in at all unless it is secure. Note the URL in the browser the first time you attempt to log in. If it starts with "http://" instead of "https://", please refer to the last section of the First Steps Guide for the admin client.

Predictive UC Analytics supports these authentication methods for authenticating logins:

- **Local:** Predictive UC Analytics manages these user accounts and their passwords locally in its database. You may configure rules for minimum password length, password complexity, and password expiration. If a user forgets their password, an administrator can reset it.
- **LDAP:** Predictive UC Analytics manages these user accounts but not their passwords. The web server forwards these logins and passwords to an LDAP server (usually a Windows Domain Controller) to validate them. You may configure multiple LDAP servers/domains for logins to use, and you may use LDAP security group membership to grant privileges.
- **Alternative Authentication:** Predictive UC Analytics supports a number of alternative authentication methods that require additional professional services to configure and test. The most popular of these methods are Integrated Windows Authentication (IWA) via Kerberos/NTLM, Active Directory Federated Services (ADFS), and Google OAuth.

Note: When you first install the product, the admin client will create the first administrator accounts as local logins with no password. These accounts have full access to the admin client, but they will not be able to log into the web interface until you set an initial password for them in the "User Manager" component of the admin client.

Navigation

After logging in with an admin account, the navigation bar at the top should show from left to right:

- The navigation menu button

- The title of the current page
- An “Admin View” selection for testing other manager accounts.
- Icons for changing manager settings, viewing general info, and viewing the status for tasks running in the background.

If the navigation menu is not already open, click the menu button on the left to open it. The darker menu items represent sub-menus with additional selections beneath them. The pin icon at the very bottom can toggle the menu state between: keep hidden, icons only, or keep open.

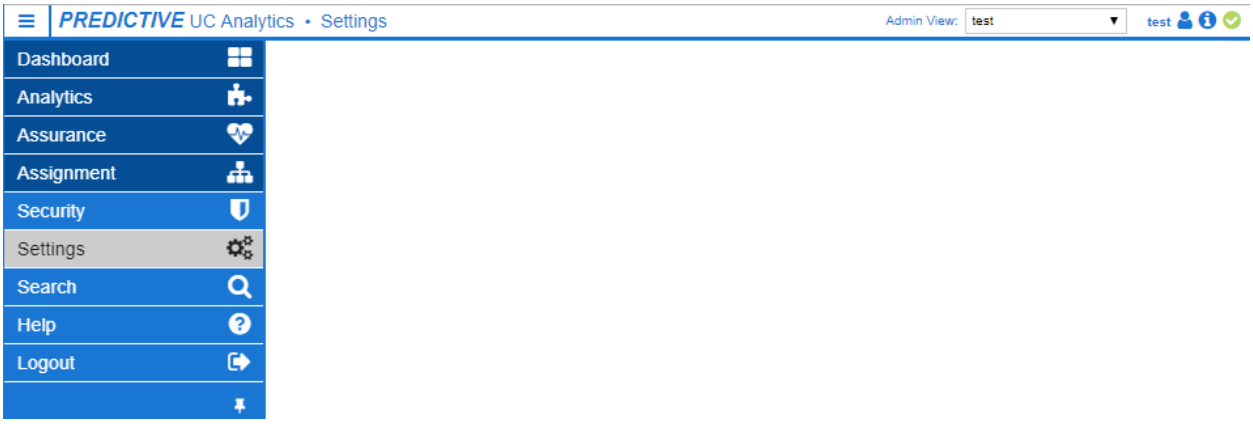

Manager Settings

Click the login name or the person icon next to it in the top-right corner to access manager settings. This will open a dialog box to allow you to change generic options for your login like color scheme, tool-tips, or password. You may only change the password for local logins, not LDAP or alternative logins.

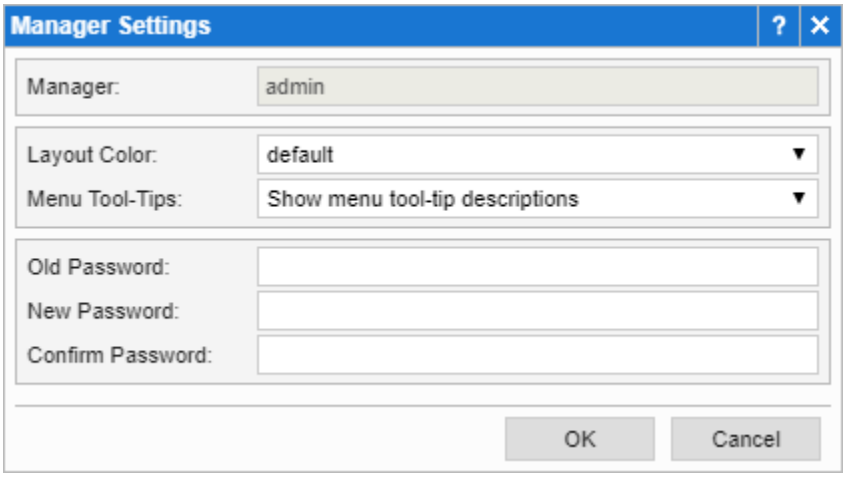

Common Features

The web interface has a few common features and icons that you may see in several different components. These features work the same way regardless of which component they appear in. Explanations for individual components will refer back to this section.

Shareable URL

- Many components will have an icon that that looks like three dots connected with lines. That icon will open a small popup containing a URL and a “Copy” button. The “Copy” button will copy the URL to your clipboard. The URL can be used to share content with other people, or embed content using Iframes in your CRM, wallboard, or intranet sites without having to be authenticated.

Run Report

- Many components will have an icon that that looks like a paper behind a bar chart. That icon will open a dialog to run a report. The report will use the filters from the current entity (ie. Monitors, KPI Models, ViewPoints).

Chart Options

- Many components will have charts with an options icon to allow you to control how the chart looks.

View Details

- Many components will have an icon that that looks like a document with a magnifying glass over it. That icon will open a dialog to display session detail records for the currently selected: directory entity, KPI model, trending chart, etc. Filters and a date range will be pre-configured based on your current selection, and you will be able to adjust the filters in the dialog.

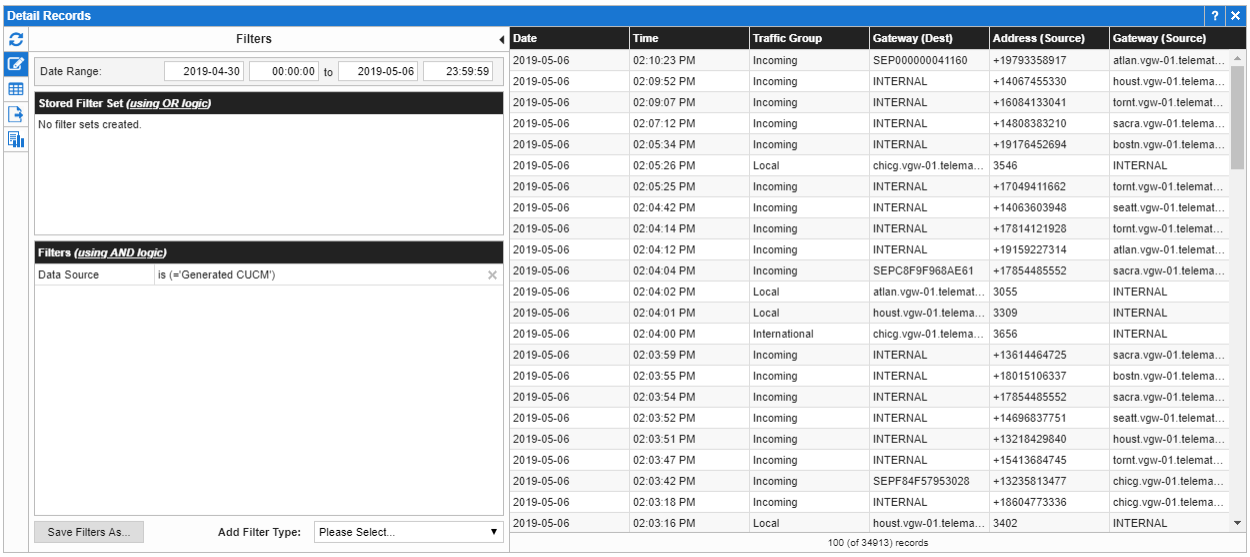

Note: The administrator may apply security restrictions on whether your login can view details, how much detail it can see, and which columns/filters it can see.

The “Filters” icon on the left allows you to modify the filters and the date/time range to apply to the detail records displayed on the right. The “Columns” icon on the left allows you to add, remove, or rearrange the columns displayed on the right. After making changes to filters or columns, click the “Refresh” icon (which looks like a right arrow) to refresh the detail records on the right. This tool provides a great interactive method to test various filter combinations. Once the filters show a specific

data set you need to see often, we recommend that you save them as a stored filter set to reuse later. You may also click the “Run Report” button to run a report with the current filters applied.

Filters

In addition to the filter configuration in the detail dialog above, reports, monitors, viewpoint maps, and KPI models all allow you to configure filters to separate the data you need to see for each specific instance. For example, you may want a monitor that only shows video conferences with poor call quality, a KPI model showing international calls by country, or a report that only shows IM and email to/from addresses ending in “@telemate.net”.

Filters are configured the same way no matter what they are used for, and stored filter sets can be shared between components. We strongly recommend saving combinations of filters that you may need use again later. If you have trouble finding the right combination of filters to apply, we recommend using the view details dialog to test filter changes interactively.

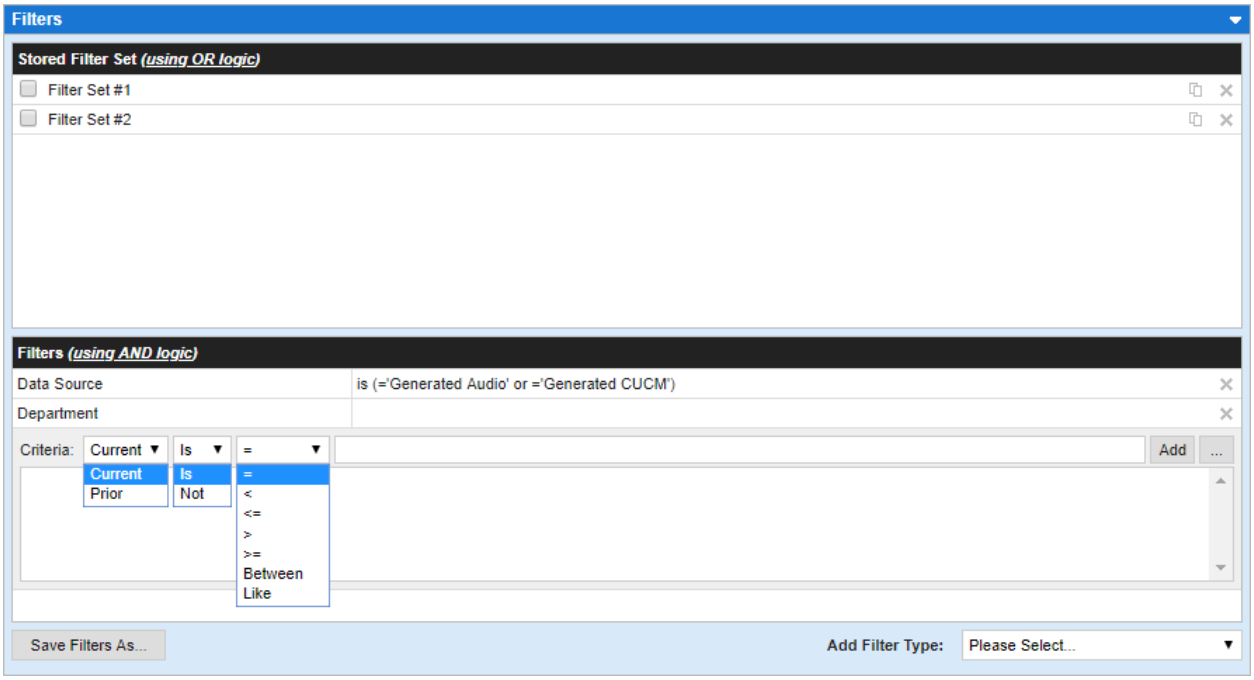

Note: The administrator may apply security restrictions to restrict which filters your login can use. In addition, security filters may exist that will always be applied to your login automatically.

To add a new filter, select the desired filter name from the “Add Filter Type” selection box. When you select one, a new filter definition line will be added to the “Filters” list, and a panel of controls will open below it to allow you to specify the criteria to include or exclude. At any time, you may show or hide the criteria panel by clicking on a filter definition line in that list. In the screenshot above, a “Data Source” filter had been configured previously and the “Department” filter name had just been selected from the “Add Filter Type” selection.

To filter traffic, you need to use the criteria panel to specify exactly what you want to include or exclude. Include filters tell the engine “I only want to include x, y, and z”. Exclude filters tell the engine “I want to include everything except x, y, and z”. For most filter types, the easiest way to do this is to click the small “...” button on the right side of the criteria panel to bring up a list of values to select from:

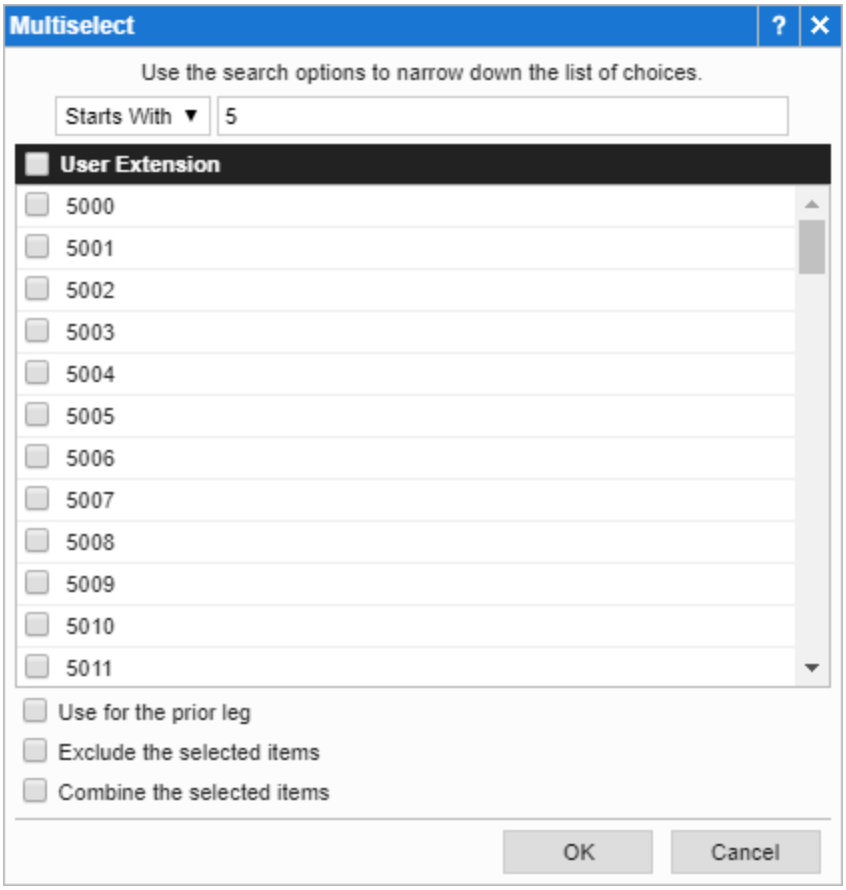

By default this list will only show up to the first 1000 options. If you choose a filter type with thousands of values to choose from, then you will need to use the search options to narrow down the list to find what you are looking for. Select the appropriate values, check the exclude box if you wish to exclude these values from the report, and click OK.

Some filter types do not show a “...” button. These tend to be values not tracked in their own special table in the database, like external phone numbers and email addresses. The number of internal addresses are bound by the size of your organization, but the number of external addresses could easily be in the millions. It is also important to keep in mind that the “...” button works well when you want to select a small number of specific values, perhaps fewer than 25, but not when you want to specify a range of 1000 values.

In those cases, you need to specify the appropriate values from the “Is/Not” selection box, the comparison operator selection box, manually type in the search criteria, and then click the “Add” button to add it to criteria list below. As you add filter criteria, it shows up in the list below the criteria control

panel as well as in the filter definition line above it. The text in the line above may be abbreviated due to limited space, but it attempts to make the selected filter logic clear.

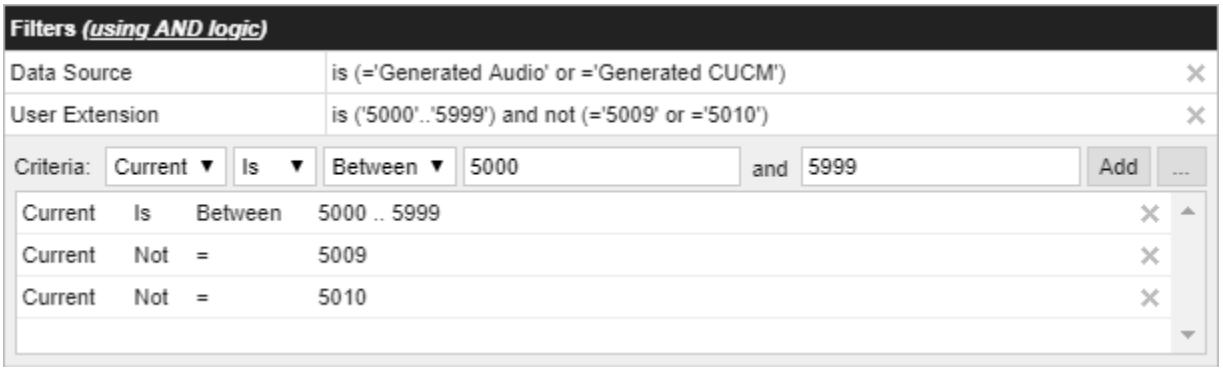

Note: The "LIKE" comparison operator uses SQL wildcards, so '_' will match any single character and '%' matches any number of characters. So "LIKE 5%" would mean "Starts With 5", "Like %5" would mean "Ends With 5" and "LIKE %5%" would mean "Contains 5".

In the example above, traffic will only be included in the report if it matches both the "Data Source" filter AND the "User Extension" filter (note the "using AND logic" reminder above). To match the first filter, the traffic must be generated by the "Demo" or the "UCCX" data source. To match the second, the traffic must be assigned to a "User Extension" in the range between "5000" and "5999" EXCEPT for "5009" and "5010" (which were specifically excluded).

To remove one part of a filter's criteria, for example the "Not = 5009" criteria line above, click the "x" button on the right of that criteria. To remove the entire "User Extension" filter, click the "x" button on the right of that filter's definition line.

As mentioned above, the normal filters always use "AND" logic. Suppose you need to specify filters to include all traffic that is either from the "Demo" data source or has an extension in the range from "5000" to "5999". If you glance above the "Filters" box to the "Stored Filter Sets" box, you should spot a reminder that says "using OR logic". The solution is to add a Data Source filter set to "Is = Demo" and click the "Save Filters As" button. Then clear that filter and add an Extension filter set to "Is Between 5000..5999" and click the "Save Filters As" button. Then clear that filter and check both of the new "Stored Filter Sets".

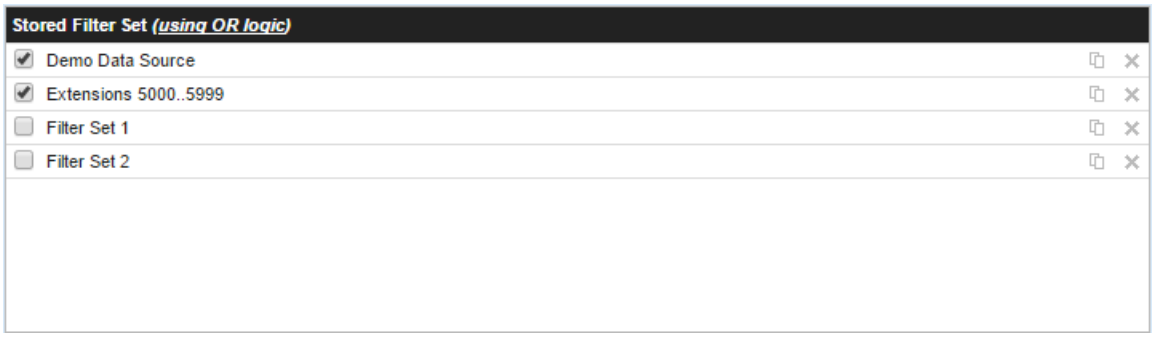

Stored filter sets can be used for much more than providing OR logic. Any time you find yourself needing to use the same filter values multiple times, save them. If you want to combine them with AND logic later on instead of OR logic, click the “Copy” button on the right (which copies their filter criteria down to the “Filters” list below) instead of checking them. When you save a set of filters for a report, you can use that filter set in Monitors and Viewpoint (and vice versa).

Note: Some reports will not provide the option to use or save filter sets. This includes older custom reports and newer reports that provide a limited filter set. Only a report that provides the full set of UC traffic filters can take advantage of stored filter sets.

Manage Destinations (Alarm Notification Delivery)

Most alarm notifications are sent to the same group of people, setting that up multiple times is a pain. In the “Manage Destinations” section, you can create, edit and delete destinations to be used in Predictive UC Analytics.

From the Destinations dialog you can add a destination by clicking the “+” button in the top-right corner of the list. To edit an existing destination, click on its name. To delete a destination, click on the “x” button to the right of it.

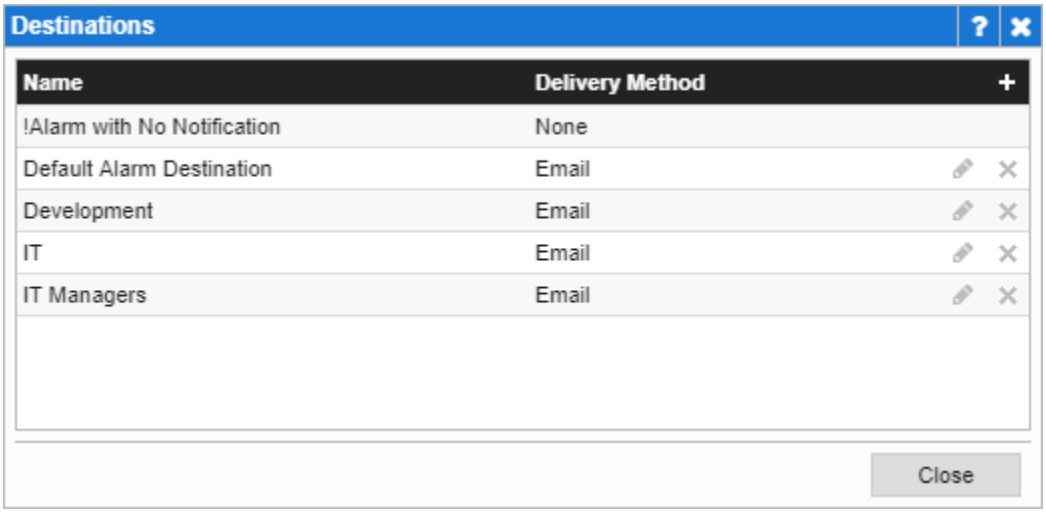

There are a number of alarm delivery methods to choose from. For example, you can send a simple text email to alert a specific person/group, generate an SMTP trap, send an IM message to Cisco Spark, or run a report and have that delivered to someone automatically. The Predictive UC Analytics delivery mechanism allows custom scripts to be written to provide custom delivery methods for you as a professional service.

Note: The “!Alarm with No Notification” destination allows alarms to be created, but no notification will be sent. This is useful when you still want to track alarms in Assurance in areas like UC Devices and Network Devices.

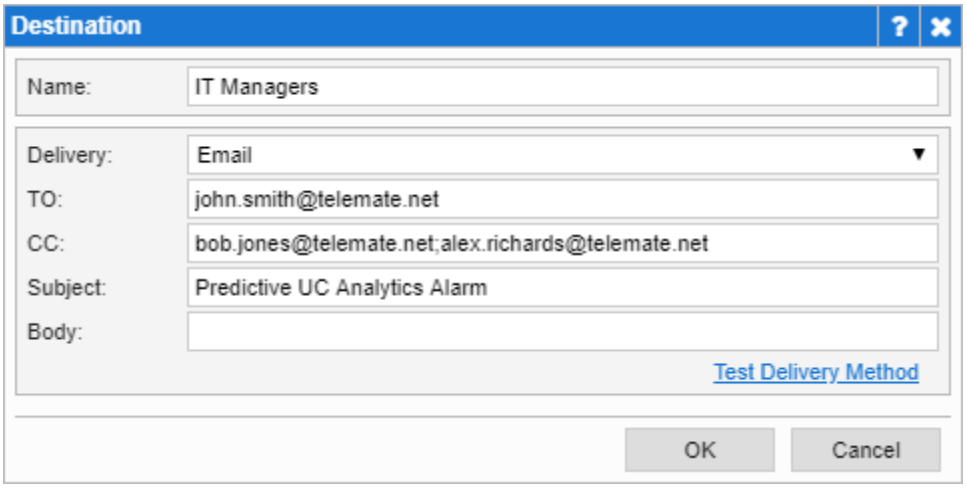

Dashboards

The dashboard component allows you to build custom single-page dashboards that combine display information from several different analytics components. It allows you to add, move, and resize "widgets" that act as windows into KPI models, monitors, viewpoint maps, and real-time contact center statistics. Once completed, a dashboard can be locked to prevent accidental changes, and then it can be shared with specific privilege groups/roles.

We strongly recommend that administrators create dashboards to share with non-administrators instead of allowing non-administrators to create their own. For example, you may create a "Sales Manager" dashboard that looks the same for every sales manager login, but filled with data specific to his or her sales department. A regional sales manager would see the same dashboard filled with data from multiple sales departments.

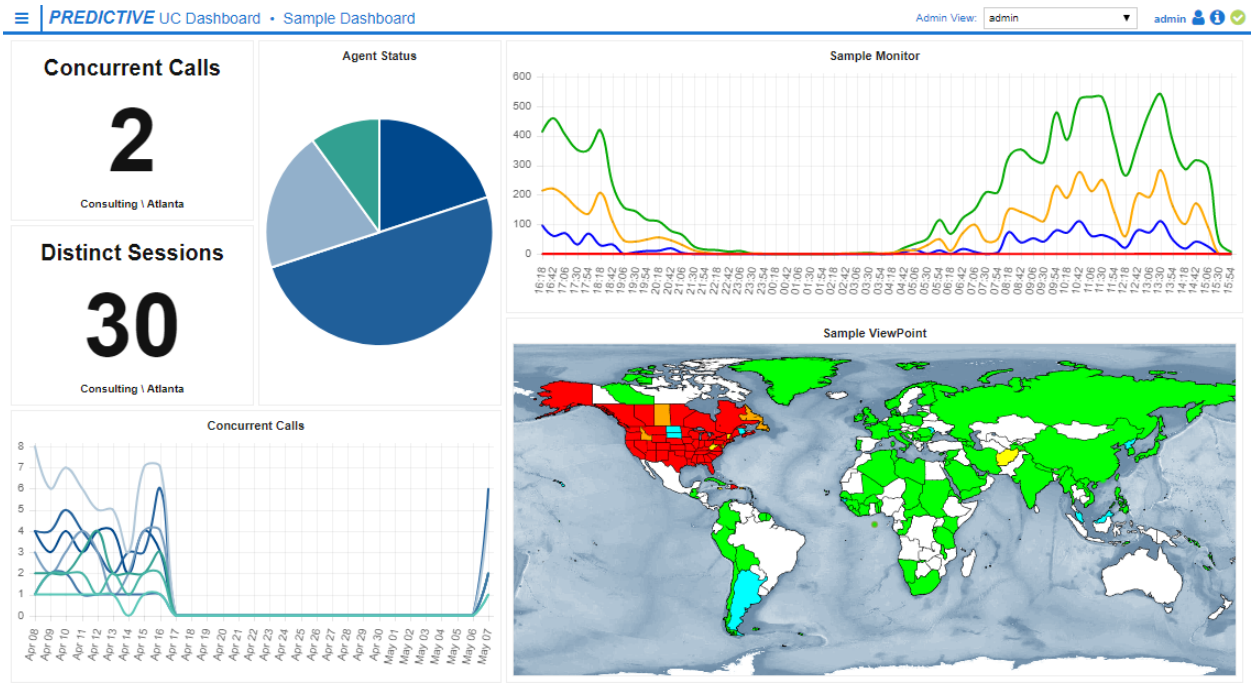

Creating Dashboards

Because you may have several different concerns regarding your data, you can create several different dashboard “pages” that you can switch between quickly. To create a new dashboard, click on the Dashboard category icon in the navigation menu. When hovering over the icon it will become a “+” symbol. To edit the current dashboard page, the gear icon will appear when hovering over the icon for the dashboard in the menu. Clicking either icon will open up the Dashboard properties window.

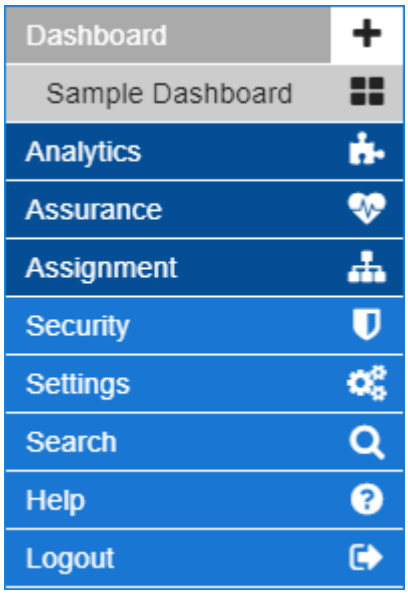

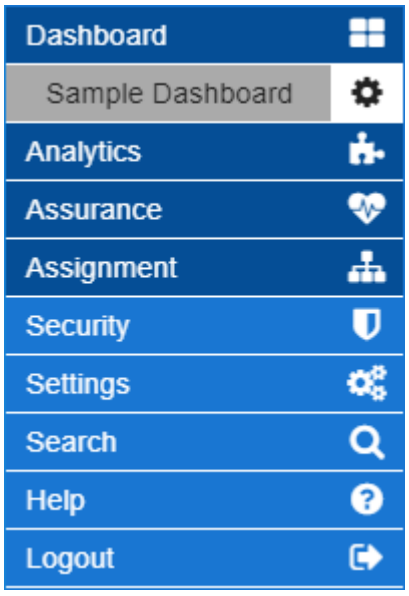

(ex. add icon on Dashboard category) (ex. gear icon on a dashboard page)

Dashboard – Properties

When you create a dashboard, you will be required to specify a name. The icon and description are optional, and will be used in the navigation menu. The description is displayed as a popout when hovering over the dashboard name or icon in the menu. If you are logged in with an admin account with access to view other managers, you will also be allowed to share it with a security group (in which case other members of that group will be able to see it but not modify it).

Dashboards support a number of different types of widgets. Currently, widgets based on Monitors, ViewPoints, KPI Model Scores, KPI Model Views and Contact Center data can be created. You can add a widget by selecting the desired widget type from the "Add Widget" drop down. To edit an existing widget, click on its name. To delete a widget, click on the "x" button to the right of it.

Locking a dashboard will prevent the widgets from being moved and resized on the page. The settings for the dashboard and "widgets" can still be edited via the settings dialog.

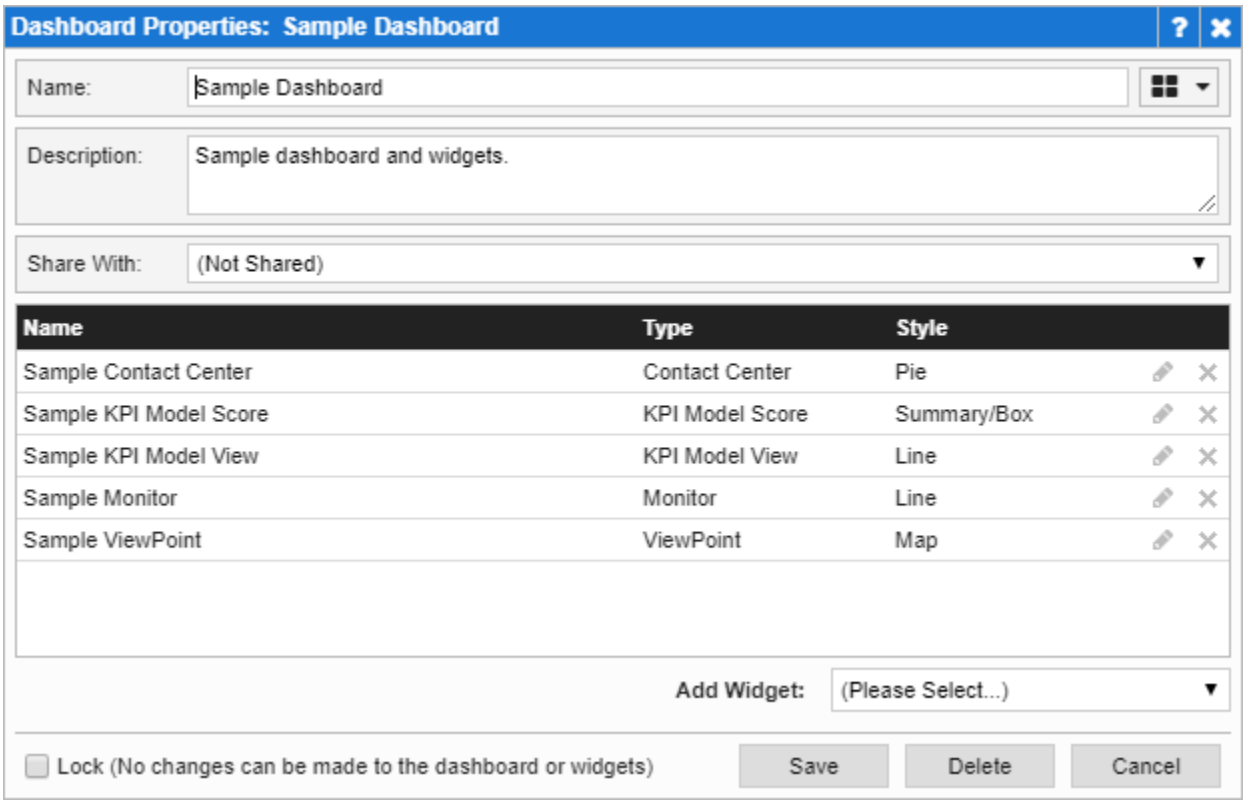

Dashboard Widget – Common

There a number of common properties shared between the different widget types.

1) **Custom Name** – When a custom name is provided it is used as the title for the widget. The custom name can also be used as a description to better identify what is being displayed in the widget.

2) **Title Sizing** – There are three sizing options for widget titles. Default is a static size and will not change if the widget size changes. Fill is dynamic sizing and will adjust based on the size of the widget. Hidden hides the title.

3) **Caption Sizing** – There are three sizing options for widget captions. Default is a static size and will not change if the widget size changes. Fill is dynamic sizing and will adjust based on the size of the widget. Hidden hides the caption.

4) **Chart Style** – Some of the chart styles are not available for every widget. The available styles include Line, Line Filled, Pie, Polar Area, Doughnut, Bar, Bar Stacked, and Summary/Box charts. Summary/Box charts will only display a number representing the value of the selected data.

5) **Legend Style** – There are a number of legend placement styles. Hidden hides the legend. Horizontal displays the legend horizontally across the bottom for the chart. Vertical (Bottom) has the legend displayed vertically below the chart. Vertical (Right) positions a vertical legend to the right of the chart. Vertical (Left) positions a vertical legend to the left of the chart.

Dashboard Widget – Monitor

When you create a widget based on a real-time traffic monitor, first you must select an existing monitor. If you do not have an existing monitor already configured you will need create one in the Monitors component. After selecting your monitor, configure the common widget properties and set the time range slider. The slider narrows the time range displayed in the charts. For real-time monitors the slider allows you to select ranges like “past 1 hour”, “past 2 hours to 5 hours ago”, or past 24 hours.

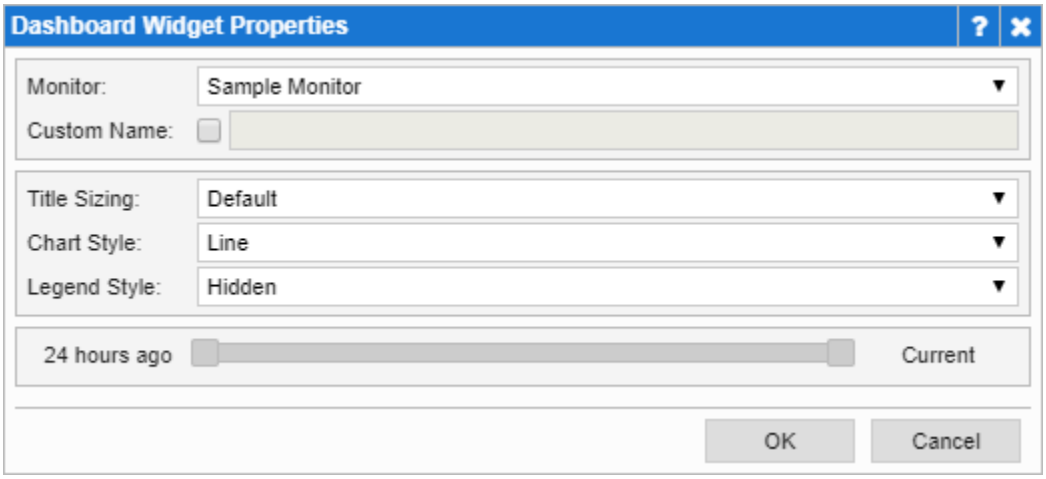

Dashboard Widget – ViewPoint

When you create a widget based on ViewPoint, first you must select an existing ViewPoint view. If you do not have an existing view already configured you will need create one in the ViewPoint component. After selecting your view, configure the common widget properties and set the time range slider. The slider narrows the time range displayed in the charts. For ViewPoint the slider allows you to select ranges like “past 1 hour”, “past 2 hours to 5 hours ago”, or past 24 hours.

Note: Zooming and repositioning the ViewPoint maps require you to hold down the CTRL key while using the mouse. To drag/reposition a map, hold down the CTRL key, click and hold with the mouse button, and then move the mouse. To zoom in or out of a map, hold down the CTRL key and use the mousewheel to control the zoom. The dashboard can not be locked.

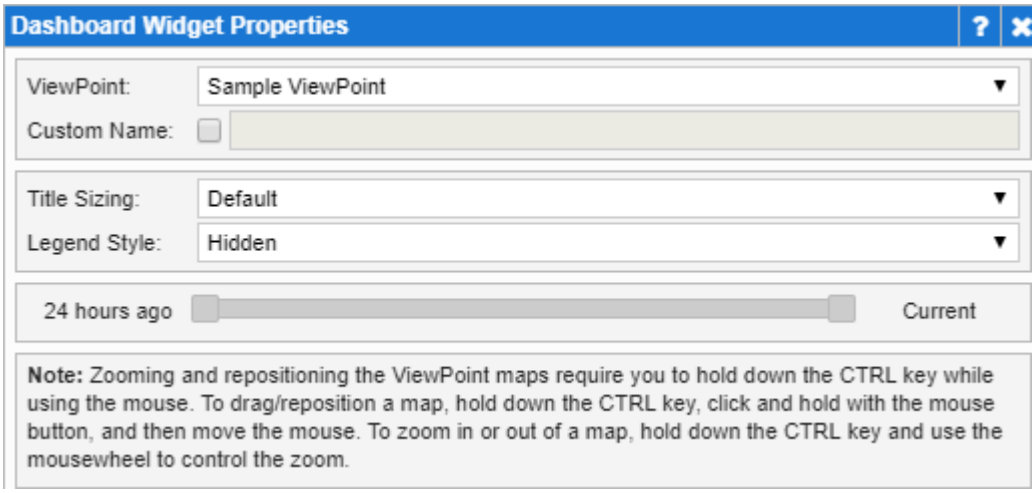

Dashboard Widget – KPI Model Score

Creating a widget based on a KPI Model Score requires you to select the KPI Model and KPI Model View before you can select the KPI Model Score. If you do not have an existing KPI Model already configured you will need create one in the KPI Model component. After selecting your model, view and score, you will be able to select the period type. The lowest period type is determined on what the base period type is set to in the KPI Model. Next you will need to select a client/parent to chart. Since there could be hundreds of client/parent combinations for the model, the list is based on those that have already been checked in the KPI Model component for the selected view.

Note: The KPI Model Score widget is presented in the summary/box style. The value shown is only for a single client and single score.

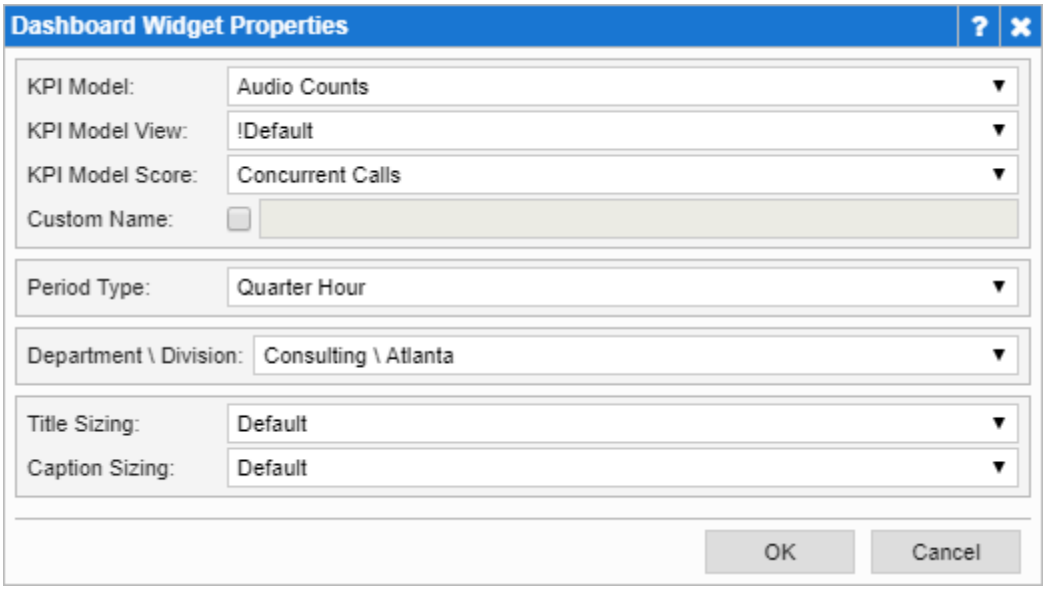

Dashboard Widget – KPI Model View

Creating a widget based on a KPI Model View requires you to select the KPI Model and KPI Model View before you can select the KPI Model Score. If you do not have an existing KPI Model already configured you will need create one in the KPI Model component. After selecting your model, view and score, you will be able to select the period type. The lowest period type is determined on what the base period type is set to in the KPI Model.

The client/parent combinations available for the widget are inherited from the KPI Model View on creation of the widget. If you change the selected clients in the KPI Model View, the clients in the widget will not change. If you need to update the client/parent values after creating the widget you can do so by clicking the link at the bottom of the dialog "Update configured clients from KPI Model View".

Note: The KPI Model View widget is displayed as one of the line or bar chart styles. The data plotted is for a single score.

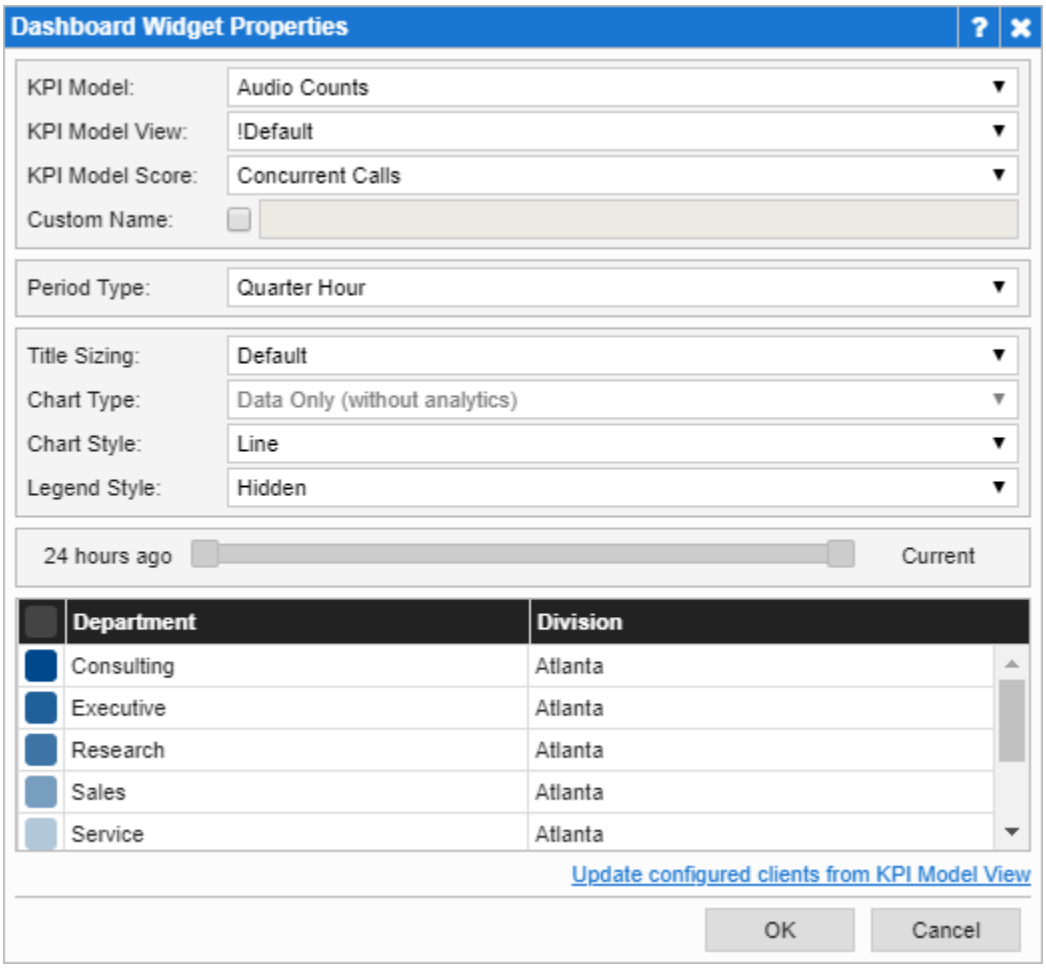

Dashboard Widget – Contact Center

Contact Center widgets do not use existing objects to build the widget. However, it does use the same data that is available in the Contact Center component. The custom name is required for the widget and is used as the title. All data displayed in the widget will come from a single data source and either CSQs or Teams that have been collected from the UCCX. To add CSQs/Teams or Data Fields, click the "+" button in the top-right corner of the list. To edit an existing selection, click on its name. To delete a selection, click on the "x" button to the right of it.

Note: Depending on the chart style selected, the number of allowed CSQs/Teams and Data Fields may be limited. For the Summary/Box style, you are allowed to pick multiple CSQs/Teams and only a single data field. For Pie chart styles, you can either have multiple CSQs/Teams and a single data field or a single CSQ/Team and multiple data fields. For Bar chart styles, you can have multiple CSQs/Teams and multiple data fields. When choosing multiple data fields they must be the same type.

The Summary/Box style has an additional setting for Threshold. If the data field value is greater than or equal to the threshold, the background color of the widget changes to match the selected color. This allows you to set up a visual indicator for when a Contact Center value reaches a certain threshold.

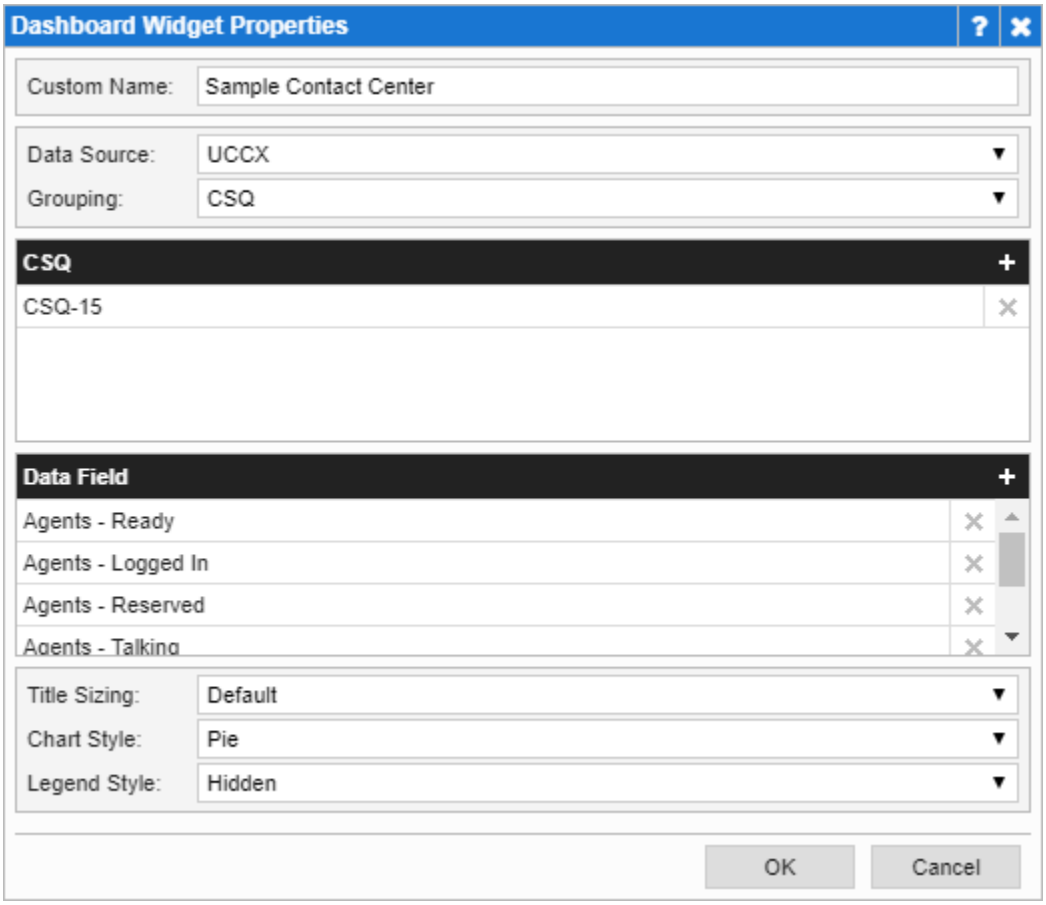

Moving and Resizing Dashboard Widgets

Dashboard Widgets can easily be moved and resized. Moving a widget is as simple as clicking and holding the left mouse button on the widget. You can then drag the widget to the desired location. Resizing the widget can be done by clicking and holding the left button near the edge of the widget. You will see the mouse cursor become a directional arrow when you are in the proper area to initiate a resize.

Note: When moving and resizing, widgets will float to the top of the page to fill up space as it becomes available. Also, widgets can not be moved or resized if the dashboard has been locked. The dashboard can be unlocked in the dashboard properties.

Analytics Features

The Trending Component

The trending component gives you an instant high-level summary of your traffic for the past several days, weeks, months, or quarters to help you quickly spot short-term and long-term trends in your UC traffic without having to run a report. Options are limited to keep it quick and easy to use. For more advanced options, configure a KPI model or a report.

The trending component provides four charts with these five options to control them:

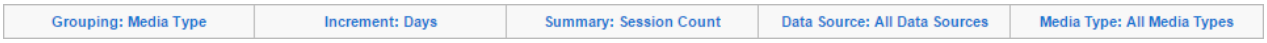

- **Grouping:** Specify whether you want to see traffic grouped by media type (or direction if only one type is available), trend period (to compare recent weeks by day of week or recent years by month or quarter), or data source (to compare data source A against B and/or C).
- **Increment:** Specify whether you want to see recent days, weeks, months, or quarters.
- **Summary:** Specify whether you want to see total session counts, duration, cost, bandwidth, or any other numeric field. Only fields configured as an "accumulator" will be available here.
- **Data Source:** Quick selection to include/exclude traffic by data source. When grouping by data source, you must select data sources here to compare against each other (the charts can only display 5).
- **Media Type:** Quick selection to include/exclude traffic by media type.

Note: In the top-right corner of each chart, you will see "View Details" and "Settings" icons. Refer to the "Common Features" section of this document for the former. The latter allows you to change the chart's view options.

The "Recent Totals" Chart

The recent totals chart shows you the total session count, duration, cost, etc. for recent days, weeks, months, or quarters based on the currently selected options. It also uses statistics calculated from those recent time periods to determine how "complete" the current time period is so it can calculate a "PROJECTED" total for the end of the current time period. For example, if you see 1000 completed calls at 2:15PM today, and the same weekday from prior weeks were 73% complete (on average) by 2:15PM, the "PROJECTED" total for "Today" would be 1370 calls (73% of 1370 is 1000).

The grouping selection controls which totals are added to which bars. The example charts below show grouping by Media Type (left) vs. Trend Period (right) while viewing daily (top) vs. monthly (bottom) session count totals for IM and Audio traffic. The top-left chart indicates that traffic fell sharply for both "Today" and "Yesterday". The top-right chart confirms that, but it also shows that traffic also fell off drastically "last week" from Thursday through Sunday and 2 weeks ago on Wednesday. The bottom-left chart shows monthly deviations that make it difficult to see a trend, but the bottom-right chart shows clearly that while most months have seen year-to-year growth, a downward trend started in March and seems to be continuing in April. All of this can be seen very quickly with a few clicks in the trending page.

The "Adoption Rate" Chart

This chart answers the question: "How many users are actively using this service?" When traffic is trending upward or downward, it is often important to see whether it is due to a change in the number of active users versus a change in how active those users have been. The example charts below use the same grouping as the sample charts above to show how the different grouping options affect it.

The "Top 5 Departments" Chart

This chart compares departments against each other to give you totals for the top 5. The totals are over the past few days, weeks, months, or quarters displayed in the "Recent Totals" chart. A separate top 5 pie chart is displayed for each grouping (whether that turns out to be media type, direction, time period, or data source). In the example charts below, you can see that while Sales generates the most overall traffic, Services generates the most Video traffic and Support generates the most App Sharing traffic.

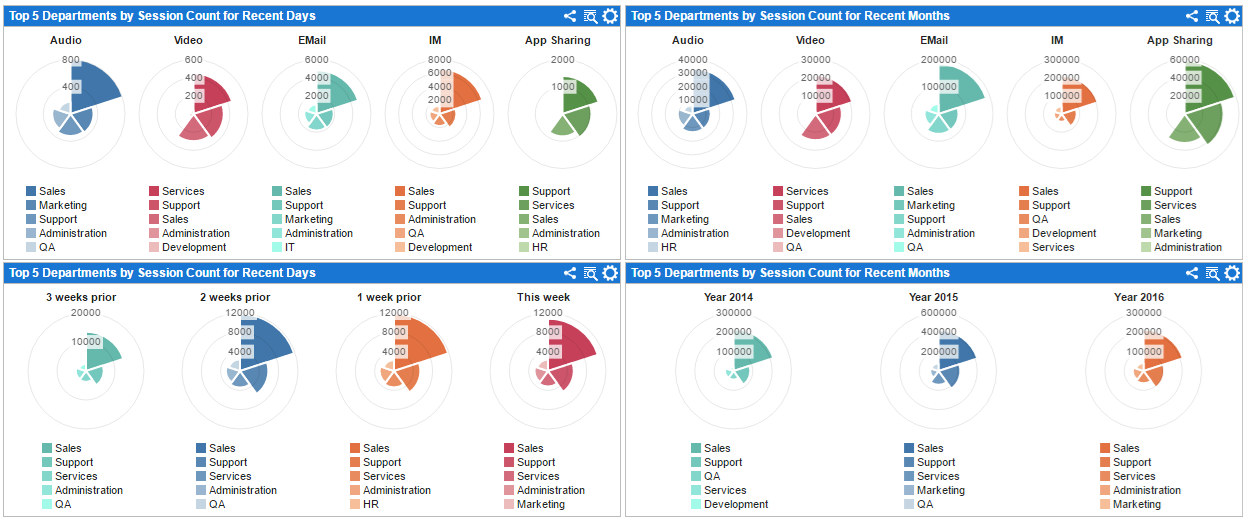

The "Statistical Distribution" Chart

This chart requires a bit more explanation because it contains more densely packed information displayed in a box plot with whiskers, and the grouping options selected have a larger impact on what it displays. The top and bottom of each box are always the 25th and 75th percentile values. The range inside the box is often called the middle 50% or the Inter-Quartile Range (IQR). The horizontal line drawn inside the box is the 50th percentile (or median) value. Outliers are not displayed on this chart, so the whiskers are used to show the min and max values.

When the time increment is set to anything longer than "Days", the box plot always shows a distribution of daily totals. In addition to showing the box plot, this type of distribution will also show the average (or mean) daily totals as dots. The dots are connected with lines connecting to make it easier to spot upward/downward trends over time for a single group. Based on the summary value you choose, these can be daily total session counts, total duration, total cost, total bandwidth, etc. This chart can quickly give insight into what your min, max, and typical days look like as well as how "typical" has changed over time for whatever group selections you choose to look at.

When the time increment is set to "Days", distributions of daily totals for individual days would be meaningless, so the box plots change to show more detailed distributions within each day. What they show depends on the summary value you choose. When the summary value is set to "Session Count", the box plot shows the distribution of session counts throughout the day. For example, if 50% of the day's calls are over at exactly 1:33PM, the median line is drawn there. (This is what the "Recent Totals" chart looks at to determine that today is "X% complete" to calculate "PROJECTED" values.) When the summary value is set to anything else, like "Cost", the box plot shows per-session statistics for that value (e.g. the min, max, and middle 50% session costs for the day).

NOTE: Hover the mouse pointer over any set of boxes to see specific numeric values in the form of "min – 25%/50%/75% - max".

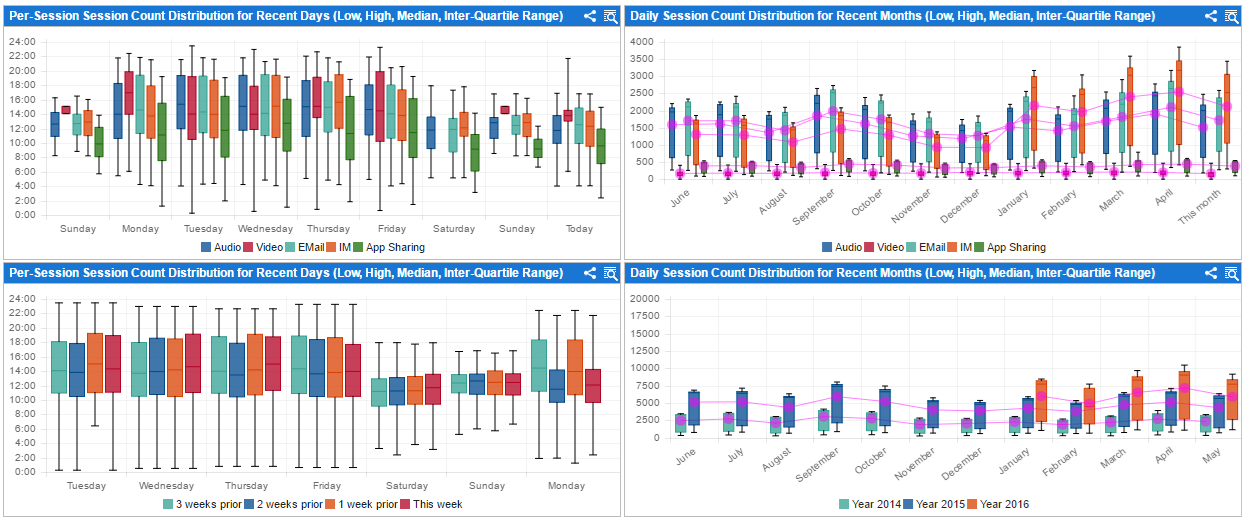

The Contact Center Component

The Contact Center component displays real-time statistics for Cisco Contact Center data sources so you can keep an eye on your agents and CSQs (or Contact Service Queues). This component will only be available if you have purchased a data source license for your Contact Center and the services necessary to configure the Predictive UC Analytics dashboard scripts to collect real-time data from your Contact Center.

The options at the top of the page allows you to filter and hide the different sections:

Data Source Filter	CSQ Filter	View Options

For example, the options allow you to choose a different data source, select a handful of CSQ's to monitor (sometimes you need to monitor a few of them instead of one or all of them), or to hide sections or the web page header to get more out of your screen's real estate.

CSQ Statistics

The information contained in "CSQ Statistics" comes directly from your Contact Center's "Real-Time Snapshot" feature. Depending on how your dashboard collection script is configured, it can either be updated in real-time via the CTI protocol or pulled from the Cisco database periodically (typically every 5 to 30 seconds). It contains some immediate counters, such as how many calls are currently waiting in each CSQ and how many agents are currently available to handle those calls. It also contains some summary information for the current day. You may click on any numeric value in this table to define thresholds for highlighting the value in different colors to make it stand out. For example, if the number of calls waiting jumps to 10 for CSQ 3, you may want to see it highlighted in red so you can take steps to free up agents to handle some of those calls.

Note: As different queues have different volumes, thresholds are saved separately for each queue. The saved thresholds also must be configured per manager.

CSQ Statistics

CSQ	Agents						Calls				Queue Time			Talk Time		Last Update
	Logged In	Not Ready	Ready	Reserved	Talking	Work	Waiting	Handled	Abandoned	Dequeued	Average	Max	Oldest	Average	Max	
CSQ1	5	1	1	1	2	0	0	4	0	0	0:00:29	0:01:07	0:00:00	0:01:06	0:01:49	0:00:02
CSQ2	5	1	1	1	2	0	0	8	0	0	0:01:23	0:03:00	0:00:00	0:02:06	0:03:45	0:00:02
CSQ3	5	1	1	1	2	0	0	1	0	0	0:00:02	0:00:02	0:00:00	0:01:03	0:01:03	0:00:02
CSQ4	4	1	1	1	1	0	1	1	0	0	0:02:52	0:02:52	0:00:00	0:03:46	0:03:46	0:00:02
Totals:	-	-	-	-	-	-	**1**	**14**	**0**	**0**	**0:01:11**	**0:03:00**	**0:00:00**	**0:02:00**	**0:03:46**	-

Clicking the gear icon in the upper-right corner of the "CSQ Statistics" section allows you to hide columns in the text view. An additional option allows you to toggle the display to show the information as a bar chart. As a bar chart the data can be accumulated based on waiting, handled, abandoned and/or dequeued calls.

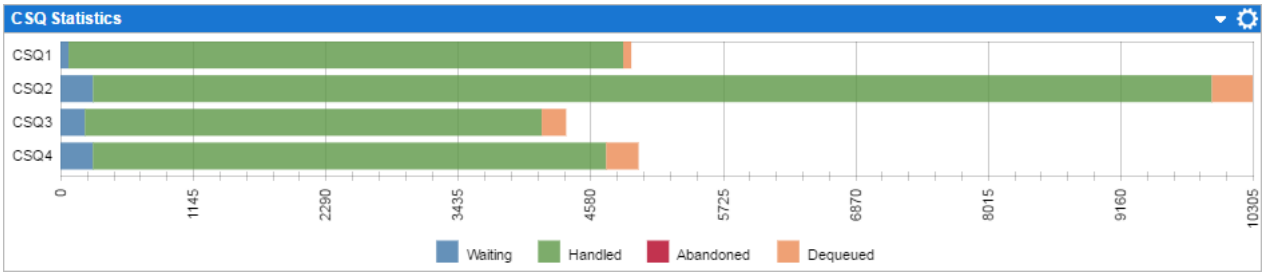

Live Calls by Agent

The "Live Calls by Agent" section is only available if you are collecting live call information via the Cisco Contact Center's CTI protocol. If so, it displays agents grouped by team. Each agent's current state is displayed. If the agent is on the phone, it displays all available call information. This information is updated once per second.

Live Calls by Agent

	Agent	Extension	State	Reason Code	Duration	Active Call							
						Start Time	Direction	ANI or DN	CSQ	Queue Time	Ring Time	Hold Time	Talk Time
Claims													
	Smith, Alaina	4134	Talking		5669:27:27								
	Smith, Amber	4142	Not Ready	6	1:31:52								
	Smith, Anterrica	4139	Talking		5669:20:59								
	Smith, Ashley	4136	Ready		0:00:07								
	Smith, Bradley	4135	Talking		5669:10:45								
	Smith, Cameron	4125	Talking		0:05:02	03:30:55 PM	Incoming	7702675073	CSQ1	0:00:19	0:00:02	0:00:00	0:05:02
	Smith, Charles	4128	Not Ready	2	0:05:01								
	Smith, Clara	4141	Ready		0:00:32								
	Smith, Erin	4130	Not Ready	6	1:32:00								
	Smith, Haley	4122	Talking		0:11:01	03:24:50 PM	Incoming	4042100902	CSQ3	0:00:03	0:00:03	0:00:00	0:11:01
	Smith, Heather	4140	Not Ready	2	0:01:34								
	Smith, Jennifer	4114	Talking		0:05:01	03:30:49 PM	Incoming	7066161881	CSQ1	0:00:23	0:00:08	0:00:00	0:05:01
	Smith, Jessica	4117	Not Ready	1	5669:07:34								
	Smith, Kristen	4129	Not Ready	2	7:10:15								
	Smith, LaTonya	4132	Talking		0:05:40	03:35:45 PM	Outgoing	918037990273	CSQ2	0:02:09	0:00:06	0:00:00	0:00:59
	Smith, Mary	4133	Work		0:00:04								
	Smith, Morgan	4138	Not Ready	2	5669:16:56								
	Smith, Sean	4118	Ready		0:01:38								
	Smith, Sha' Rae	4124	Talking		5669:15:09								
	Smith, Shonterria	4120	Not Ready	2	0:14:52								
	Smith, Tametrius	4126	Talking		5669:08:57								
	Smith, Thwanna	4115	Not Ready	11	0:03:23								

Clicking the gear icon in the upper-right corner of the “Live Calls by Agent” section allows you to hide columns in the text view. An additional option allows you to toggle the display to show the information as a bar chart. As a bar chart the data can be accumulated based on queue time, ring time, hold time and/or talk time.

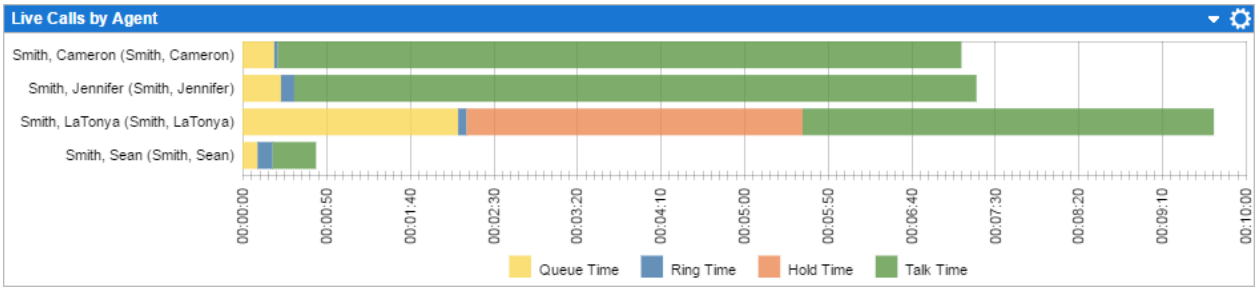

Note: This section will not show calls waiting in a CSQ. See “Live Calls by CSQ” for that information.

Live Calls by CSQ

The “Live Calls by CSQ” section is only available if you are collecting live call information via the Cisco Contact Center’s CTI protocol. If so, it displays all live CSQ calls grouped by CSQ. This information is updated once per second.

Live Calls by CSQ

Start Time	Queue Time	ANI	Agents					
			Name	Extension	Call State	Ring Time	Hold Time	Talk Time
CSQ1								
03:30:49 PM	0:00:23	7066161881	Smith, Jennifer	4114	CONNECT	0:00:08	0:00:00	0:08:01
03:30:55 PM	0:00:19	7702675073	Smith, Cameron	4125	CONNECT	0:00:02	0:00:00	0:08:02
03:37:28 PM	0:00:09	7065954320	Smith, Sean	4118	CONNECT	0:00:09	0:00:00	0:01:39
CSQ2								
03:28:22 PM	0:02:09	8037541706	Smith, LaTonya	4132	CONNECT	0:00:05	0:03:21	0:05:19
03:37:33 PM	0:00:06	4044372610	Unknown	4121	CONNECT	0:00:06	0:00:00	0:00:49
03:37:35 PM	0:00:11	7068826141	Smith, Jessica	4143	CONNECT	0:00:11	0:00:00	0:01:21
CSQ3								
CSQ4								

Note: This section will not show direct in/out calls. See “Live Calls by Agent” for that information.

CSQ Summary

The “CSQ Summary” section is similar to the “CSQ Statistics” section above, but instead of showing the summary counts Cisco tracks, it shows the counts from the “Recently Completed Calls” section, which depends on the time window you specified when configuring the dashboard collection script. It is often more useful to be able to see the number of calls that have been dequeued or abandoned in the past 30-60 minutes than for the entire day.

Summary for last 60 minutes

CSQ	Service Level	Handled			Abandoned			Dequeued			Call Back		
		Count	Average	Max	Count	Average	Max	Count	Average	Max	Count	Average	Max
CSQ1	14.3%	14	0:00:17	0:02:20	0	N/A	N/A	0	N/A	N/A	0	N/A	N/A
CSQ2	23.7%	38	0:00:32	0:03:37	0	N/A	N/A	0	N/A	N/A	0	N/A	N/A
CSQ3	15.4%	13	0:00:18	0:01:35	0	N/A	N/A	0	N/A	N/A	0	N/A	N/A
CSQ4	0.0%	7	0:00:06	0:00:09	0	N/A	N/A	0	N/A	N/A	0	N/A	N/A
Totals:	**18.1%**	**72**	**0:00:24**	**0:03:37**	**0**	**N/A**	**N/A**	**0**	**N/A**	**N/A**	**0**	**N/A**	**N/A**

Agent Summary

The “Agent Summary” section contains a summary of agents assigned to the selected queues that have valid state information for the past 24 hours. These statistics include each agent’s current state and the length of time they have been in that state. You may click on either of those fields to configure thresholds based on an agent’s state duration. Click on the “info” icon next to an agent’s name to see more information about that agent (e.g. CSQ assignments by team, resource group, and skill set).

Agent Summary

	Agent	Extension	State	Reason Code	Duration	Contacts		Ring Time		Hold Time		Talk Time		Work Time	
						Presented	Handled	Average	Max	Average	Max	Average	Max	Average	Max
Claims															
	Smith, Alaina	4134	Talking		5669:50:17	2	2	0:00:02	0:00:02	0:01:15	0:02:30	0:05:47	0:07:28	0:00:09	0:00:15
	Smith, Amber	4142	Not Ready	6	1:54:42	8	8	0:00:03	0:00:07	0:01:32	0:06:07	0:06:11	0:13:10	0:00:13	0:00:15
	Smith, Anterrica	4139	Talking		5669:43:49	8	8	0:00:02	0:00:05	0:02:36	0:06:30	0:07:47	0:19:33	0:00:08	0:00:15
	Smith, Ashley	4136	Work		0:00:04	0	0	0:00:00	0:00:00	0:00:00	0:00:00	0:00:00	0:00:00	0:00:00	0:00:00
	Smith, Bradley	4135	Talking		5669:33:35	1	1	0:00:08	0:00:08	0:03:12	0:03:12	0:14:00	0:14:00	0:00:15	0:00:15
	Smith, Cameron	4125	Talking		0:04:00	0	0	0:00:00	0:00:00	0:00:00	0:00:00	0:00:00	0:00:00	0:00:00	0:00:00
	Smith, Charles	4128	Talking		0:05:56	12	10	0:00:07	0:00:10	0:00:56	0:02:48	0:06:55	0:14:32	0:00:11	0:00:15
	Smith, Clara	4141	Ready		0:14:32	1	1	0:00:06	0:00:06	0:00:50	0:00:50	0:02:09	0:02:09	0:00:15	0:00:15
	Smith, Erin	4130	Not Ready	6	1:54:50	0	0	0:00:00	0:00:00	0:00:00	0:00:00	0:00:00	0:00:00	0:00:00	0:00:00
	Smith, Haley	4122	Ready		0:05:51	4	4	0:00:03	0:00:06	0:00:00	0:00:00	0:09:19	0:13:40	0:00:04	0:00:15
	Smith, Heather	4140	Talking		0:05:22	9	9	0:00:07	0:00:08	0:00:37	0:03:05	0:09:02	0:21:21	0:00:13	0:00:15
	Smith, Jennifer	4114	Not Ready	10	0:03:31	6	6	0:00:02	0:00:05	0:01:35	0:05:56	0:07:53	0:21:33	0:00:10	0:00:15
	Smith, Jessica	4117	Not Ready	1	5669:30:24	6	6	0:00:06	0:00:09	0:00:49	0:02:21	0:04:32	0:12:54	0:00:14	0:00:15
	Smith, Kristen	4129	Not Ready	2	7:33:05	4	4	0:00:04	0:00:07	0:01:55	0:04:15	0:04:29	0:07:03	0:00:15	0:00:15
	Smith, LaTonya	4132	Talking		0:03:05	7	7	0:00:04	0:00:06	0:01:05	0:03:07	0:05:49	0:11:35	0:00:10	0:00:15
	Smith, Mary	4133	Talking		0:00:05	4	3	0:00:05	0:00:10	0:06:17	0:08:28	0:11:38	0:24:38	0:00:05	0:00:09
	Smith, Morgan	4138	Not Ready	2	5669:39:46	6	6	0:00:06	0:00:09	0:00:36	0:01:07	0:07:34	0:18:35	0:00:14	0:00:15
	Smith, Sean	4118	Talking		0:10:20	0	0	0:00:00	0:00:00	0:00:00	0:00:00	0:00:00	0:00:00	0:00:00	0:00:00
	Smith, Sha' Rae	4124	Talking		5669:37:59	7	7	0:00:02	0:00:05	0:02:05	0:06:22	0:06:12	0:12:33	0:00:06	0:00:15
	Smith, Shonterria	4120	Talking		0:05:06	7	7	0:00:06	0:00:08	0:00:34	0:02:17	0:04:20	0:06:20	0:00:13	0:00:15
	Smith, Tametrius	4126	Talking		5669:31:47	5	5	0:00:04	0:00:06	0:01:19	0:05:04	0:04:47	0:10:00	0:00:10	0:00:15
	Smith, Thwanna	4115	Talking		0:01:16	4	4	0:00:03	0:00:06	0:00:44	0:02:40	0:09:44	0:25:21	0:00:15	0:00:15
Totals:		-	-	-	-	**101**	**98**	**0:00:04**	**0:00:10**	**0:01:22**	**0:08:28**	**0:06:55**	**0:25:21**	**0:00:11**	**0:00:15**

Clicking the gear icon in the upper-right corner of the “Agent Summary” section allows you to hide columns in the text view. An additional option allows you to toggle the display to show the information as a bar chart. As a bar chart the data can be accumulated based on calls handled, calls presented or times. The times available include ring time, hold time, talk time and work time. All the times can be selected together so that the bars are stacked.

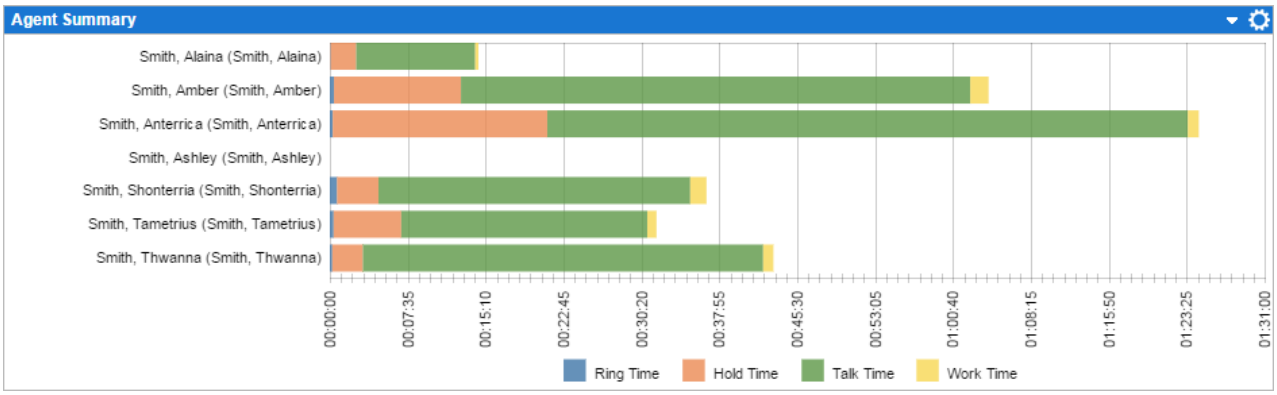

Recent Calls

The "Recent Calls" section shows a listing of recently completed calls that have completed during the time window you specified when configuring the dashboard collection script. For abandoned and dequeued calls, it shows which agent phones rang vs. "No agents available". You may click on any of the column headers in this table to sort by that column.

Calls for last 60 minutes

End Time	Call Type	Duration	ANI or DN	Queue Time	CSQ	Script	Agents					
							Name	Extension	Ring Time	Hold Time	Talk Time	Work Time
10:29:01 AM	Handled	0:03:38	2292467802	0:00:09	CSQ1	Claims	Smith, Jessica	4117	0:00:09	0:00:18	0:02:19	0:00:15
10:28:08 AM	Handled	0:06:56	2299035221	0:00:03	CSQ2	Claims	Smith, Jennifer	4114	0:00:03	0:02:40	0:03:17	0:00:03
10:27:18 AM	Handled	0:08:29	4789872470	0:00:06	CSQ3	Claims	Smith, Charles	4128	0:00:06	0:01:56	0:05:26	0:00:04
10:27:00 AM	Handled	0:02:20	4044039609	0:00:08	CSQ4	Claims	Smith, Heather	4140	0:00:08	0:00:00	0:01:45	0:00:15
10:25:06 AM	Handled	0:07:18	4786253051	0:00:05	CSQ3	Claims	Smith, Amber	4142	0:00:05	0:00:00	0:06:14	0:00:02
10:23:45 AM	Handled	0:02:42	4044039609	0:00:08	CSQ4	Claims	Smith, Heather	4140	0:00:08	0:00:00	0:02:08	0:00:15
10:22:41 AM	Handled	0:12:20	9124279223	0:00:01	CSQ3	Claims	Smith, Kristen	4129	0:00:01	0:04:15	0:07:03	0:00:15
10:22:14 AM	Handled	0:31:07	4782378850	0:00:07	CSQ1	Claims	Smith, Ashley	4113	0:00:07	0:00:41	0:29:28	0:00:15
10:21:05 AM	Handled	0:07:23	9123671353	0:00:04	CSQ1	Claims	Smith, LaTonya	4132	0:00:04	0:00:00	0:05:47	0:00:15
10:19:47 AM	Handled	0:01:38	4047988200	0:00:02	CSQ2	Claims	Smith, Jennifer	4127	0:00:02	0:00:15	0:00:53	0:00:15
10:19:23 AM	Handled	0:10:41	7064835366	0:00:03	CSQ1	Claims	Smith, Haley	4122	0:00:03	0:00:00	0:09:46	0:00:02
10:19:11 AM	Handled	0:06:49	7703379607	0:00:05	CSQ2	Claims	Smith, Sha' Rae	4124	0:00:05	0:01:17	0:04:31	0:00:01
10:16:53 AM	Handled	0:16:21	2294567905	0:00:05	CSQ2	Claims	Smith, Charles	4128	0:00:05	0:01:17	0:14:32	0:00:10
10:16:51 AM	Handled	0:05:43	6787792776	0:00:05	CSQ2	Claims	Smith, Thwanna	4115	0:00:05	0:00:00	0:04:41	0:00:15

Note: This section does not show active calls. See "Live Calls by Agent" or "Live Calls by CSQ" for that information.

The KPI Models Component

A KPI model is like a combination of multiple real-time monitors, real-time dashboards, long-term trending charts, and summary reports all rolled into one compact feature. Each model requires some effort and experimentation up front to configure, but once a KPI model is configured, it runs in the background 24/7 to get you the information you need quickly and effortlessly.

Configuring a new KPI Model requires you to answer these key questions:

- What data is available in my UC reports that needs to be included in (or excluded from) this particular service?

- What entity represents a “client” or “user” of this service? This could be a queue, a gateway, a user, a department, a site, etc.
- Do I want to see these clients broken out by some parent group (e.g. department/user, cluster/gateway, country/state)?
- What Key Performance Indicator (KPI) “scores” do I want to see calculated for each client of this model?
  NOTE: Much of the power in the KPI models feature comes from the flexibility of its scores. Each KPI score can have its own nested set of filters and grouping, and each can be configured to give pass/fail counts, percentages, or accumulated (sum/min/max/avg) values.
- What time period should the scores cover (e.g. hourly, daily, monthly)?

KPI Scores

Before you create a KPI model, first you must choose the KPI Scores. Predictive UC Analytics comes with a number of preconfigured KPI scores. You may use one of those, derive one from an existing score, or create your own.

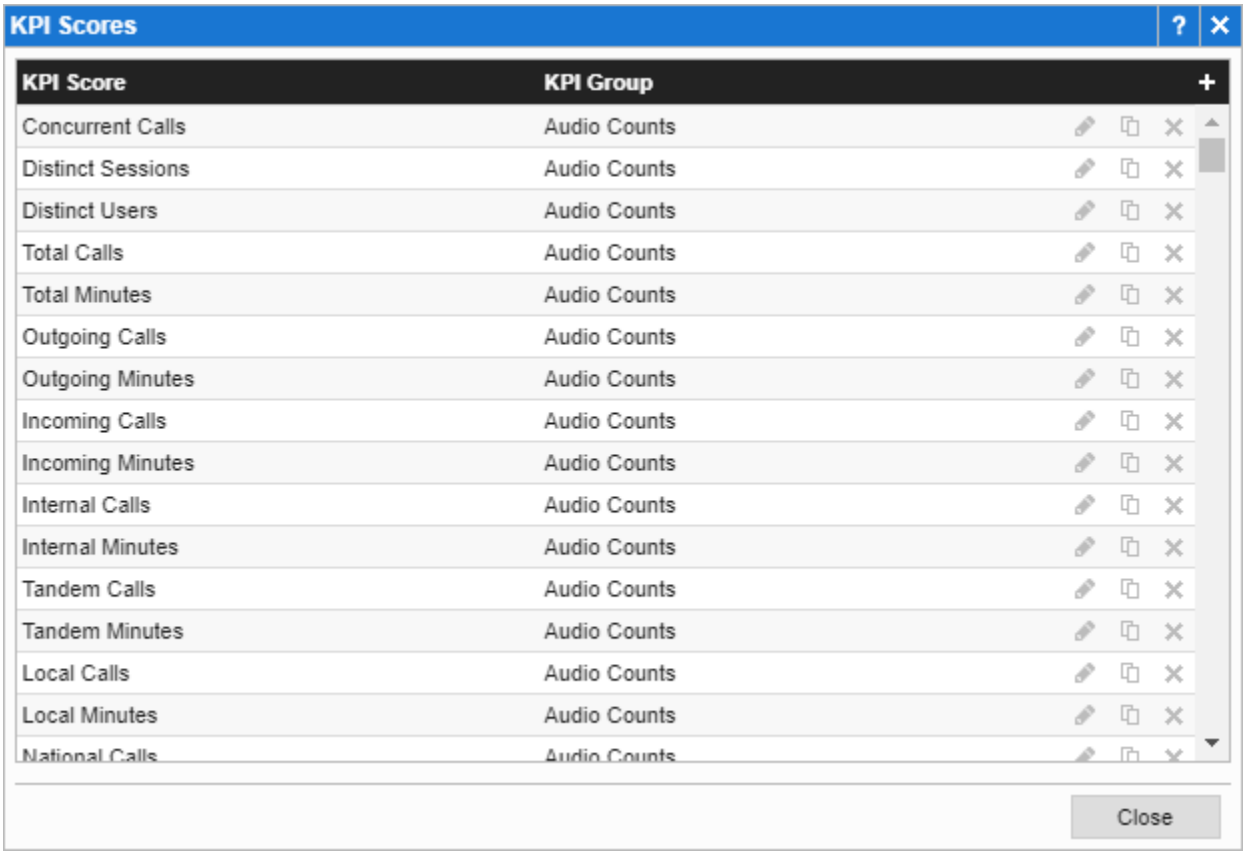

KPI Score Properties – General

When you create a KPI Score, you must first give it a name, group name, and select a score type. Different score types will require additional options to be specified.

- Pass – Calculate a simple count of sessions that pass the filter checks for this score. If no filters are specified, this count will represent the total session count.
- Fail – Calculate a simple count of sessions that do not pass the filter checks for this score. If no filters are specified, this count will always be 0.
- Accumulator – After applying this score's filters, calculate the sum, min, max, or average of any numeric field (e.g. duration, cost, queue time, bandwidth, QoS).
- Grouped Accumulator – After applying this score's filters, calculate a distinct group count or calculate min/max/average of a grouped total of any numeric field.
- Combined – Combines two sub-scores to give you A+B.
- Difference – Combines two sub-scores to give you A-B.
- Ratio – Combines two sub-scores to give you A/B.

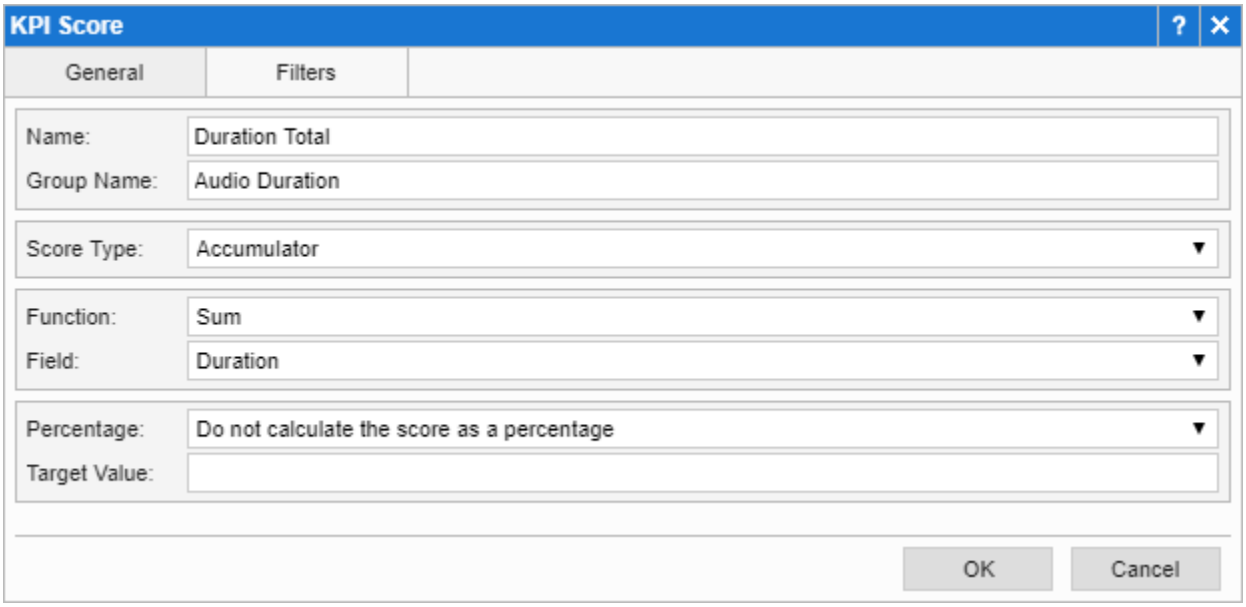

KPI Score Properties – Filters

The Filters tab looks and works exactly the same way the KPI model filters, monitor filters, and report filters work. However, the fact that these filters are nested within the model's global filters gives you a lot of flexibility. For example, you may create one score to give you a percentage of audio sessions with disconnect codes 3, 5, and 10, a second score to give you a minimum QoS score for all video sessions with disconnect code 2, and a third score to give you the average daily cost of all sessions. You can then see these scores grouped by any "client" you wish over any time period you wish. This offers flexibility not found anywhere in the standard reports or monitors.

KPI Models

Predictive UC Analytics comes with a number of preconfigured KPI Models. The preconfigured models are assigned to security groups. You can reuse and customize those models or create your own.

KPI Model Properties – General

When you create a model, first you must specify a name. If you logged in with an admin account, you will also be allowed to assign this model to another login (in which case that login will assume ownership and be allowed to modify it). You will also be allowed to enable the model have KPI Model views shared with security groups. Shared views allow members of a group to see it but not modify it.

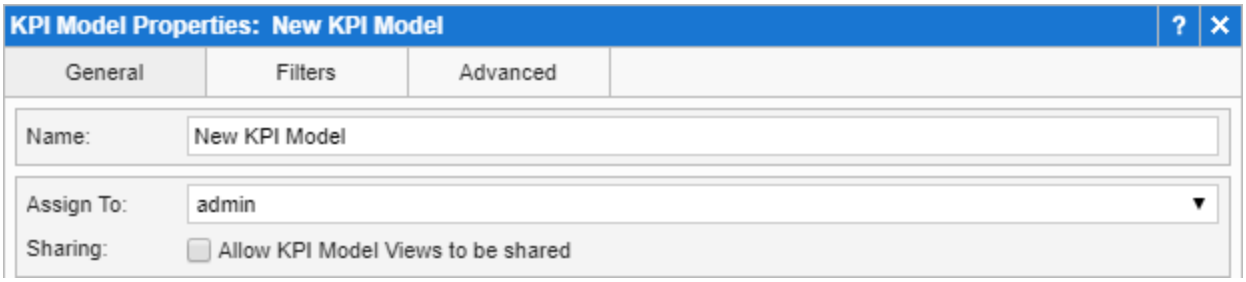

Next you must select the time period type (required), the client column (required), and the parent column (optional) that clients will be grouped by. You may also optionally specify custom names to be used for the client and parent columns of this model. For example, if some of your UC users are employees and some are clients, you may create two model with "User" as the client column, and you may give one a custom name of "Employee" and the other a custom name of "Client".

Before you can save a new model, you must add at least one KPI score to it. To add a score, select the desired score from the "Add KPI Score" selection box. When you select one, the score will be added to the "Scores" list:

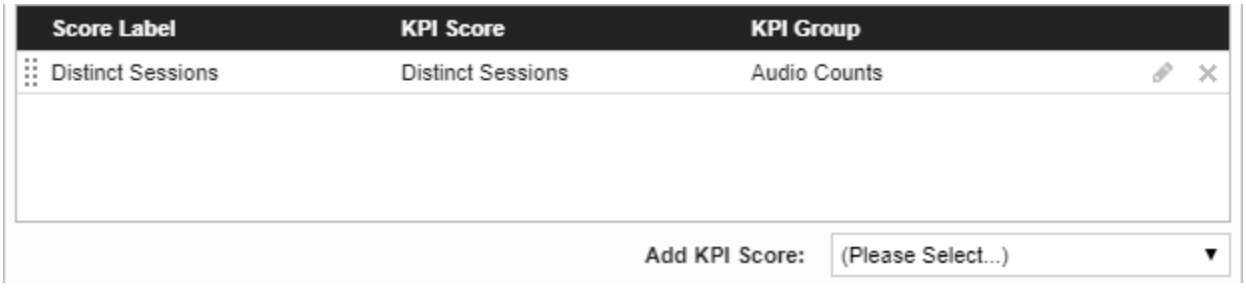

You may add as many scores as you wish. When a model has too many scores to display in a single table in your web browser, you can group scores into different views. If you are an administrator creating a model to share with managers, you can have the same model show one set of scores to manager group

A, another set to manager group B, and so on. After adding a score to the model, you may click on the score's name to modify it. To remove a score from the model, click on the "x" button to the right of it.

KPI Model Properties – Filters

The filters tab allows you to specify what data needs to be included (or excluded) from this model. These filters will impact all scores. Refer to the "Common Features" section of this document for a detailed explanation on how to configure filters. The filters look and work exactly the same way for models, and you can even reuse stored filter sets between reports, monitors, and models.

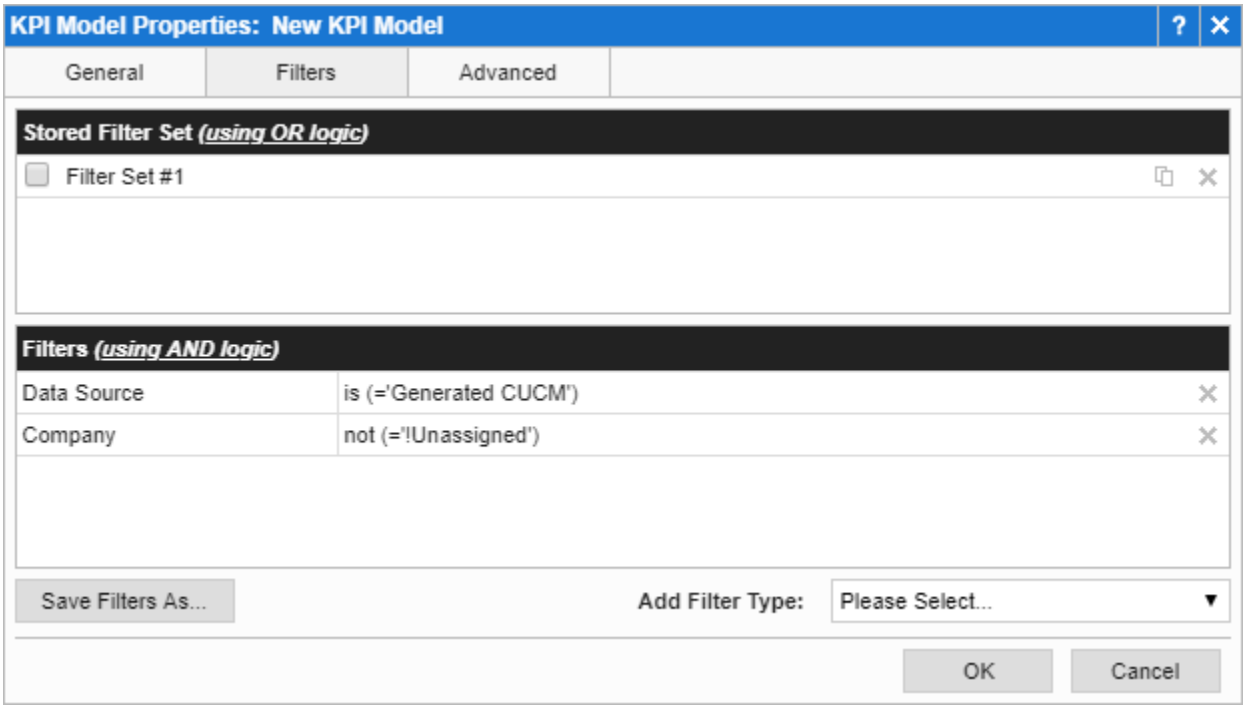

KPI Model Properties – Advanced

The advanced options allow you to select the time zone. If your UC traffic spans multiple time zones, all of it will be converted to this time zone for this model. The other options are related to the "Box" view, which displays all scores for a single client in a very large font in boxes filling up the screen for a wallboard-type display.

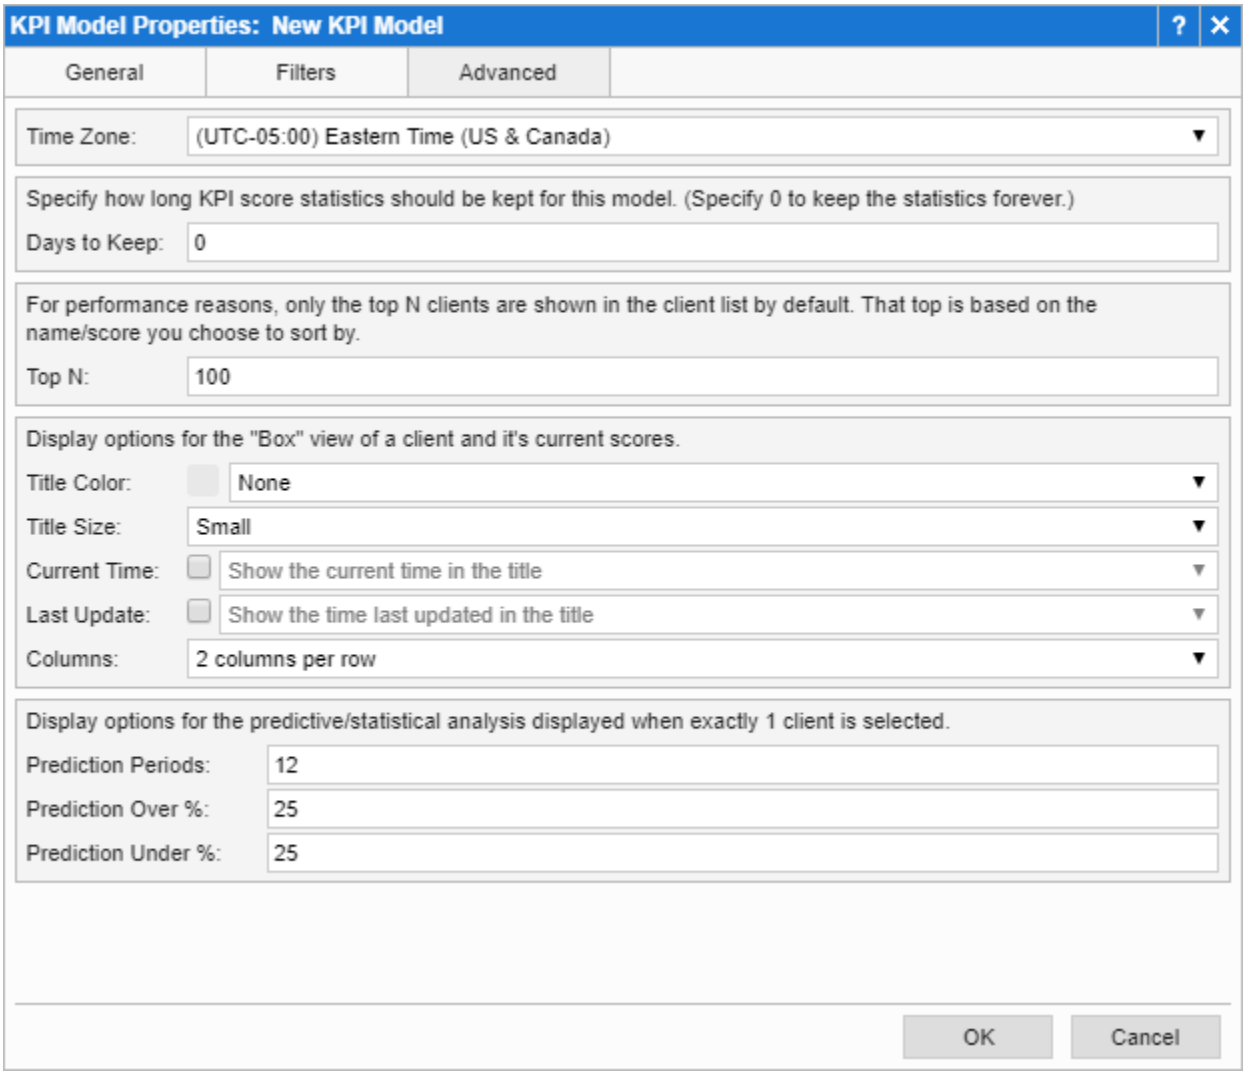

KPI Model Scores

You can modify a score assigned to a KPI model by clicking on its name from the KPI model's General tab. You cannot modify the KPI score itself here, but you can specify a name and thresholds specific to this model.

Note: Color and alarm thresholds are tied to the model's selected time period. If you want to see a score of 15 highlighted in red for an hourly view, that threshold may not make sense for a monthly view.

Score Properties – General

The General tab allows you to override the score name for this model. It also allows you to define thresholds for this score to highlight in different colors, which can help you identify issues at a glance.

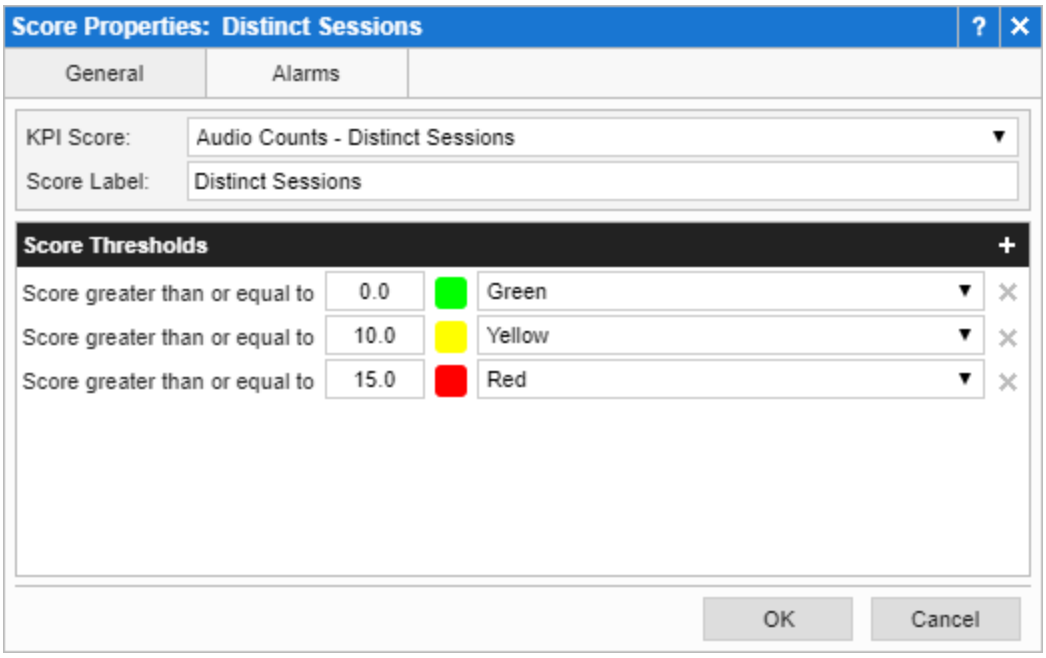

Score Properties – Alarms

This tab is nearly identical to the monitor alarm properties. Each score can have different alarm thresholds as well as different alarm delivery methods.

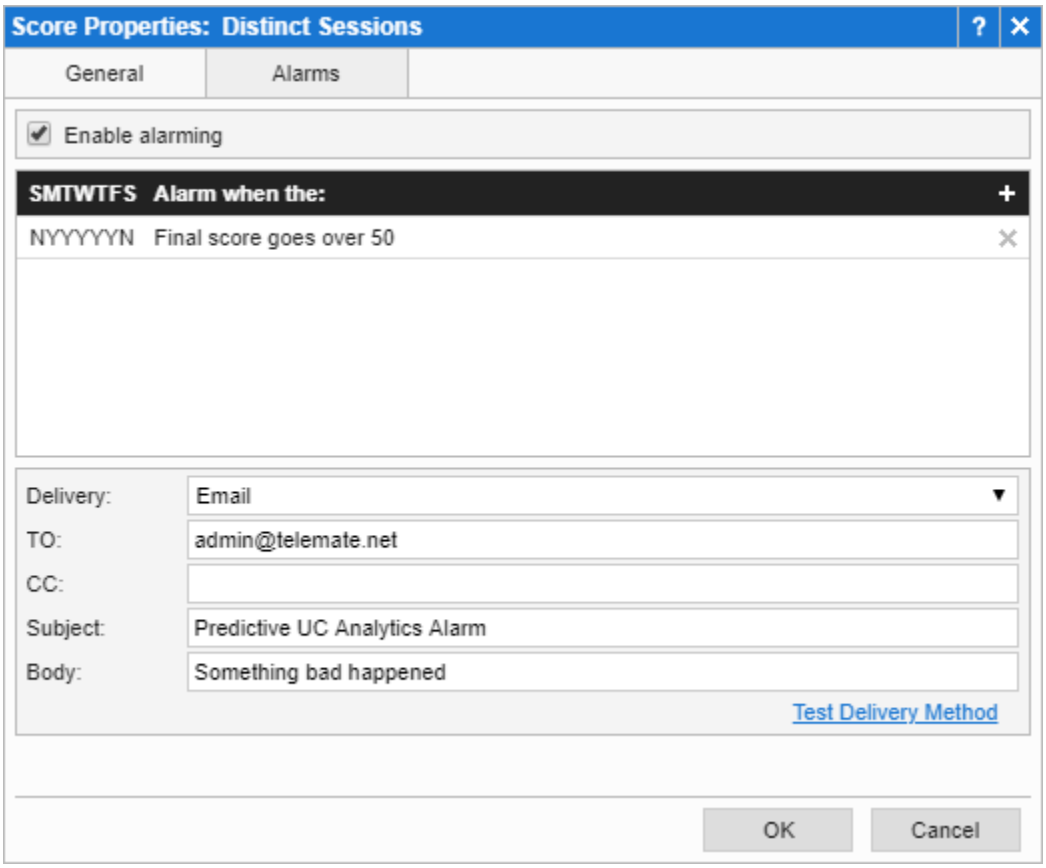

The primary difference between model alarms and monitor alarms is in the timing of when to trigger an alarm. Unlike monitors, which constantly track statistics for all sessions for the past 24 hours and can check an alarm threshold against "the past 20 minutes" at 12:53, and then again at 12:54, the models have fixed time periods and gradually build up the scores over the course of any given hour, day, week, month, etc.

For each alarm threshold, you must choose whether to allow an alarm to be triggered for "partial scores" in the middle of the current period while the score is still changing versus waiting until the period is over to check the threshold against the "final score". Because UC sessions usually are not logged until after they have ended, it is a good idea to specify a delay (in minutes) to wait past the end of a period for active sessions to finish and make their way into the UC Analytics database.

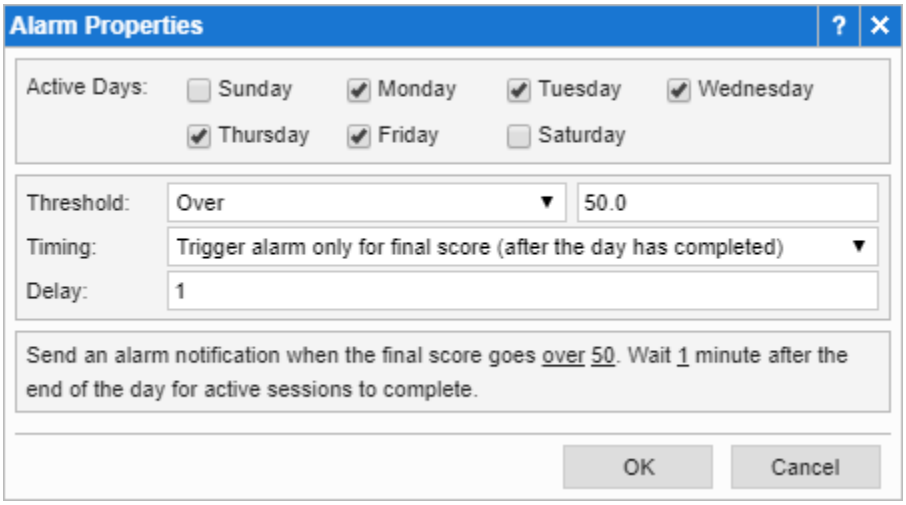

KPI Model Views

While you can easily create 10 KPI models to see 10 different sets of scores, it may be better to create one KPI model with all of those scores and then create a different view for each set of scores. This is not possible when you need one model to show you departments, another to show you users, and another to show you gateways. However, whenever it is possible to add a new view to a model, we recommend that over adding another model with the same "client" and "parent" selections.

To create a new view, click on the "View" drop-down box on the left side of the option bar at the top and select "Add a New View". You will be prompted to name the view and select the scores you want to be visible in this view. You can drag scores up and down to change the order they appear in this view, and you can choose to group certain scores together. Grouped scores can all appear together on one chart. In the screenshot below, the "Grouped KPI Scores" chart would include bars for Incoming Calls, Local Calls, National Calls and Intl Calls.

If the KPI model itself has the "Allow KPI Model Views to be shared" option checked, you may also select a privilege group to share this view with. All logins with access to that group (including tenant logins) will be able to see this view, and each manager will have their own security privileges applied so that they can only see scores applied to departments to which you have granted them access.

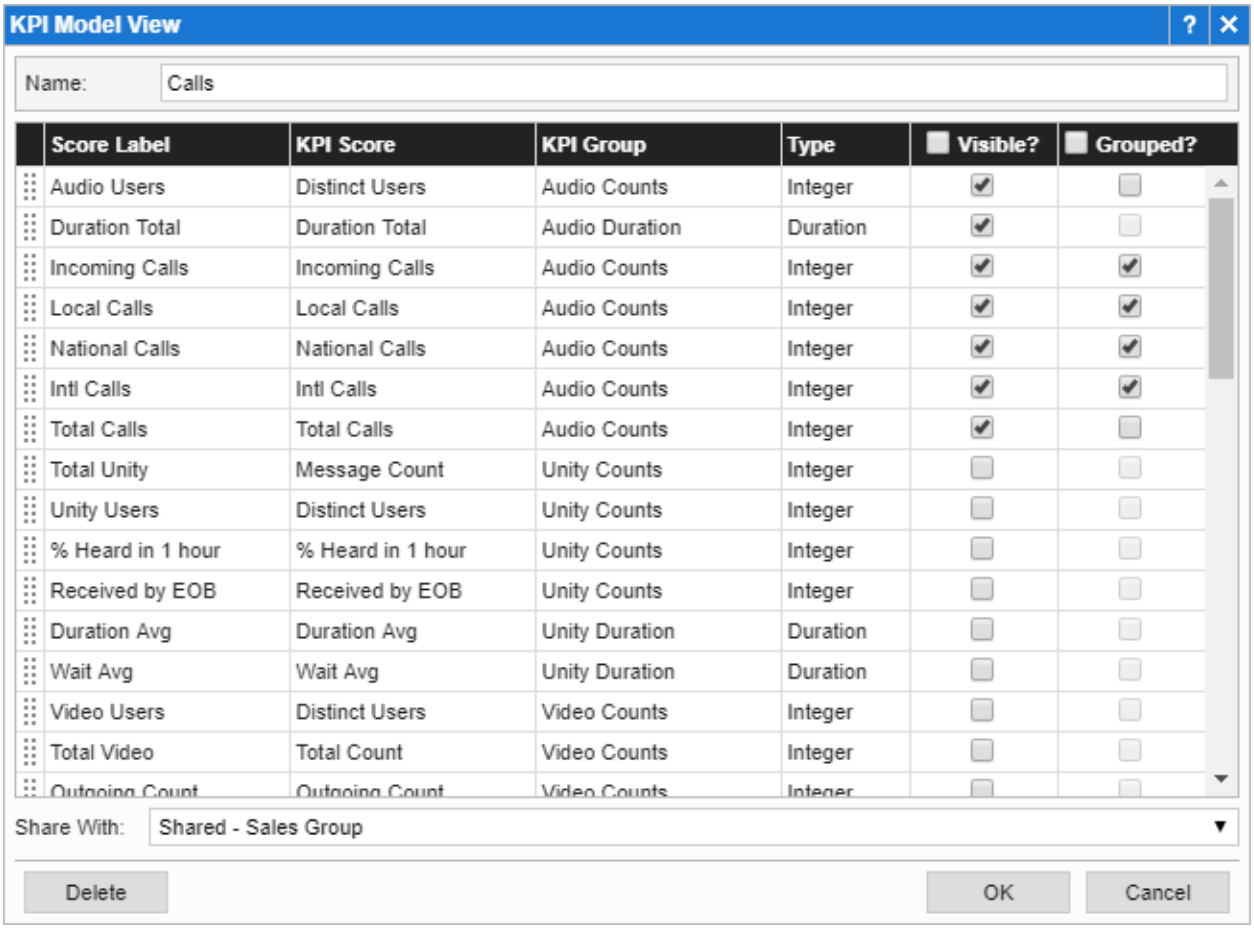

KPI Model Controls

Every time you save a KPI model, the monitoring engine will quickly scan all UC traffic from the past 24 hours to build/rebuild/update the list of detected clients and the set of scores for the current period. If you want the scores to be populated from past data, select the “Rebuild KPI Models” command and specify how far back you want to rebuild the KPI model scores.

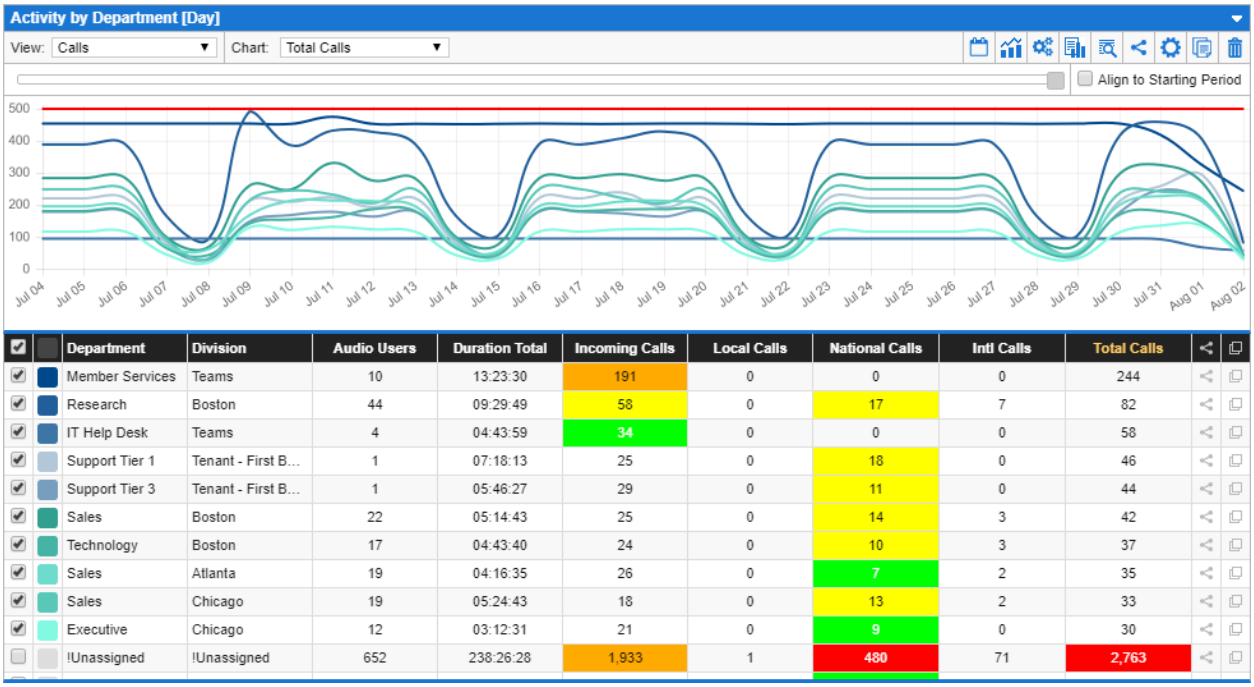

Department	Division	Audio Users	Duration Total	Incoming Calls	Local Calls	National Calls	Intl Calls	Total Calls
Member Services	Teams	10	13:23:30	191	0	0	0	244
Research	Boston	44	09:29:49	58	0	17	7	82
IT Help Desk	Teams	4	04:43:59	34	0	0	0	58
Support Tier 1	Tenant - First B...	1	07:18:13	25	0	18	0	46
Support Tier 3	Tenant - First B...	1	05:46:27	29	0	11	0	44
Sales	Boston	22	05:14:43	25	0	14	3	42
Technology	Boston	17	04:43:40	24	0	10	3	37
Sales	Atlanta	19	04:16:35	26	0	7	2	35
Sales	Chicago	19	05:24:43	18	0	13	2	33
Executive	Chicago	12	03:12:31	21	0	9	0	30
!Unassigned	!Unassigned	652	238:26:28	1,933	1	480	71	2,763

Score Table

The default view for each model shows a simple table with all clients and scores for the "current" time period. You may change the sort order in the table by clicking on any of the column headers at the top of the table. Clicking a column header multiple times will change the sort order between ascending and descending. Clicking the left and right arrows at the bottom-left corner of the table allows you to quickly view scores for prior time periods.

Score Chart

Selecting the checkbox to the left of a client's name will add/remove the client from the chart, which helps you visualize how any score's value trends over time. You may check multiple clients and compare them on the chart, and you may change the color of a client's line on the chart by clicking the colored box next to that client's checkbox. The "Chart" selection at the top next to the "View" selection allows you to change the chart to display a different score. You can also double-click any time period on the chart to view details for that time period.

Score Commands

At the top-right corner of each KPI model, you will see these command buttons:

- **Period Options** – Allows you to change the period type that KPI Model displays. The default period type set in the model properties will affect what period types are available.
- **Display Options** – Allows you to hide the summary table, hide the chart, change the chart style (line, filled line, bar and stacked bar), and display thresholds. The option to display color thresholds is only available for line charts, and only when the current view's "Period Type" matches the KPI model's period type.
- **Current View Properties** – Allows you edit (or delete) the current KPI model view.

- **Run Report** – Allows you to run a report on selected clients for this KPI model.
- **View Details** – Allows you to view recent session detail records for selected clients in this service model. Refer to the "Common Features" section of this document for a more detailed explanation of the detail view.
- **Share** – Allows you to share a link to this KPI model.
- **Properties** – Allows you to change the KPI model definition, filters, and scores.
  Warning: Any changes you make may invalidate color and alarm thresholds configured for existing scores. In addition, if you change the client or parent columns, all data collected for this model will be deleted, and you will need to use "Rebuild KPI Models" to specify how far back to rebuild it.
- **Copy** – Allows you to copy a model to modify in case you need to create a similar model.
- **Delete** – Use to delete this KPI model.

Box View

If you click on any of the "popout" icons in the far right column of the score table, it will open a new browser tab or window to display the "Box View" for that client. The box view was designed to be used as part of a wallboard display, and you can change how it looks in the model's Advanced property tab.

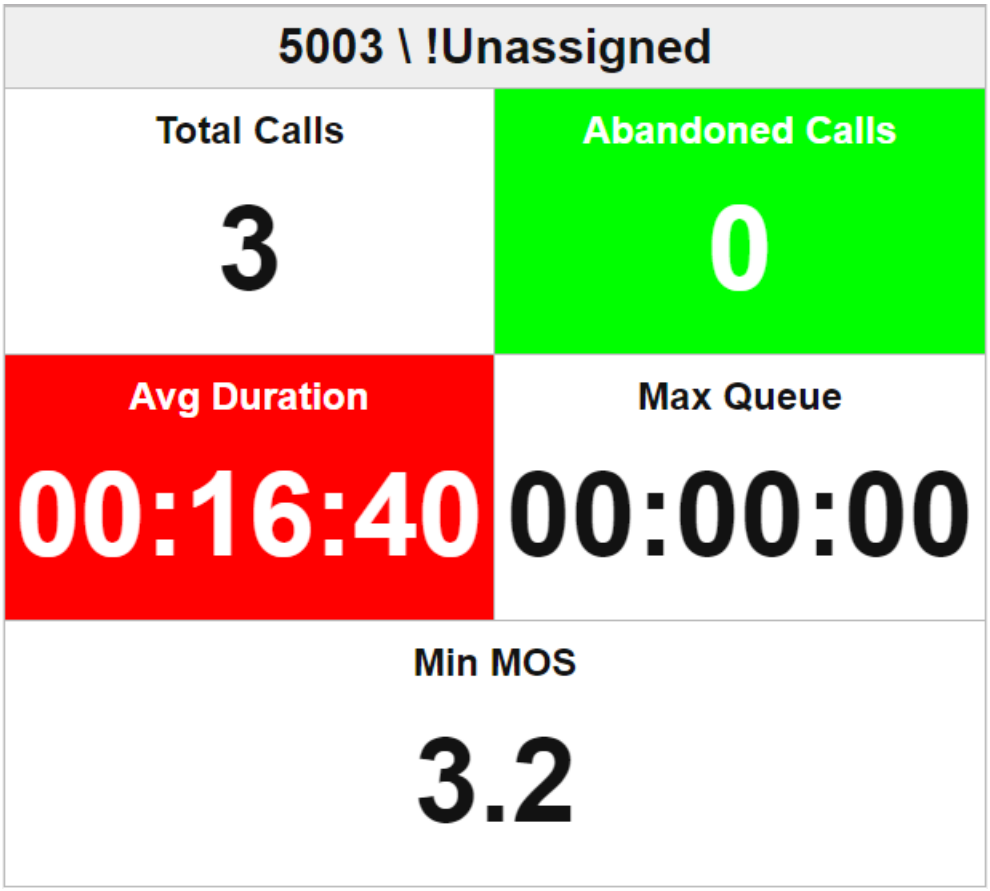

Performance Tips

While the web interface allows you to add as many KPI models and scores as you wish, each modeled score adds a small performance impact to the Predictive UC Analytics server. Most scores and models have an impact so small they are negligible, but the counts for every "group" or "break" get multiplied together (not added together), and some combinations can cause an explosion in the number of summary records generated.

This is easiest to explain by example. Picture a model with 10 simple scores, 10,000 users as clients, and 365 daily summary records. That could create up to 10 x 10,000 x 365 (36.5 million) records. If each user communicates with an average of 25 external contacts per day and you add a grouped accumulator on external address, that one score can add another 10,000 * 365 * 25 (91.25 million) records.

These tips should help you understand what to look out for:

- **KPI models with too many distinct client values** – If you choose a client or parent column with an unbounded number of distinct values (like an external phone number), you could end up with hundreds of thousands, or even millions of distinct clients. The potential for a problem should be clear if you multiply the user count in the example above by 10 or 100.
  Tip: If you really need groupings like this, consider using a KPI report instead of a model, or add filters to the model to limit the number of distinct client values.
- **Grouped Accumulator scores with too many distinct values** – As illustrated in the example above, grouped accumulators are expensive because they add another "group" level. How expensive it is depends on how many distinct group values there are.
  Tip: Use grouped accumulator scores sparingly, especially ones with a very high number of distinct group values (e.g. external address, city).
- **Sharing KPI models** – When you check the "Allow KPI Model Views to be shared" box, the engine adds a hidden break on the department to allow it to enforce security filters for any manager that can see its shared views. This hidden break has little impact on models already breaking by a directory column like department or user, but it can have a significant impact if the model breaks by something else, like external trunk/gateway, country, or state.
  Tip: Avoid checking the "shared" checkbox for every model "just in case". Plan shared models carefully with target managers/stakeholders in mind.
- **Managers creating their own KPI models** – One larger model shared with 10 managers generally performs better than 10 smaller models created by different managers. It is also easier to manage one model assigned to an admin account, and it requires less training for managers.
  Tip: Create shared models and have managers ask admins if they want additional scores (or models) added. Centralized control over models makes it easier to control performance.
- **Rebuilding KPI models** – It often requires a few iterations to get KPI models configured exactly how you want them, and any time you add or change scores or filters, you will need to rebuild them to update them for any day prior to "today". Rebuilding for a large date range can take a while, and slows down other things running on the server while it's rebuilding.

Tip: When testing changes to a model (or multiple), rebuild just a few days back as a litmus test. When you feel confident that changes are complete, start a rebuild for a longer period of time.

The Monitoring Component

Select the monitoring component to manage and view real-time traffic monitors. The monitors allow you to tell the Predictive UC Analytics engine to keep an eye on something specific for you as it processes real-time data. You can then log in periodically to check the monitor component at any time and instantly see all monitored traffic for the past 24 hours represented in a line chart, a pie chart, a small summary table, and a detail listing. Just about everything displayed in the monitor is configurable. You can even set up alerts to have the engine alert you immediately if the monitored traffic exceeds certain thresholds.

The first time you select the monitoring component, the only option available is to create a new monitor.

Monitor Properties – General

When you create a monitor, first you must specify a name. If you are logged in with an admin account, you will also be allowed to assign this monitor to another login (in which case that login will assume ownership and be allowed to modify it) or to share it with a security group (in which case other members of that group will be able to see it but not modify it).

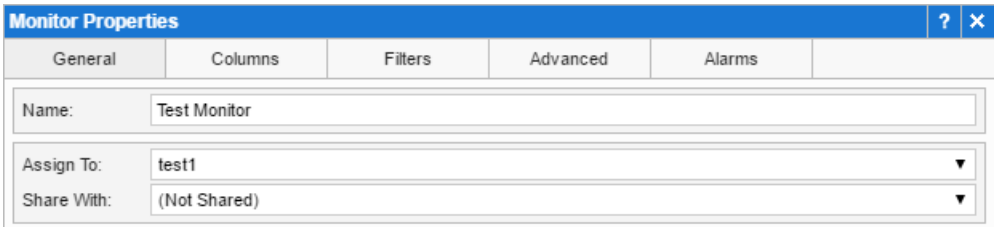

Next you must choose a UC field to monitor, and then you need to specify line criteria. Line criteria is added exactly the same way you add filter criteria (see the reporting section of this document for more detailed information). Note that each line can include a single, set, or range of values or it can exclude a single, set, or range of values. This means that lines can overlap each other if you wish. For example, line 1 could include a range of 5-7 and line 2 could include a range of 6-8. Monitors also have a special comparison operator called “Other”, which includes any traffic that does not belong to any other line.

Note: Line criteria allows you to define what you want to see summarized in the line chart, pie chart, and summary table. However, it does not filter the traffic shown in the detail record listing. Unless you specify filters for the entire monitor, all traffic that your account has access to will show up in the detail record listing.

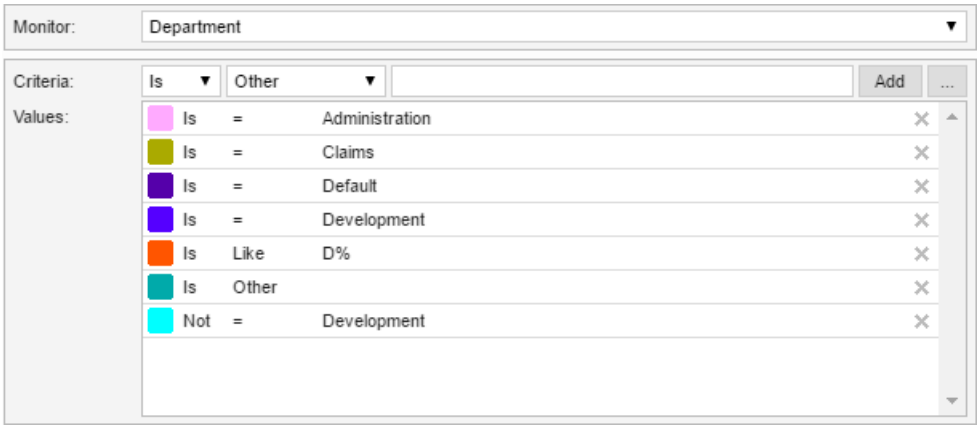

Note: Click on the color next to each line criteria to change the color of that line.

At the bottom of the General tab is a small section to allow you to change the values plotted in the line and pie charts. The “Session Count” and “Concurrent Count” plot functions are special, and will plot the total active session count and max concurrent session count, respectively. If you choose the Sum, Min, Max, or Average plot functions, you will see that function applied to the summary field. For example, you could choose to see the maximum call wait time, the minimum call quality, or perhaps the average duration.

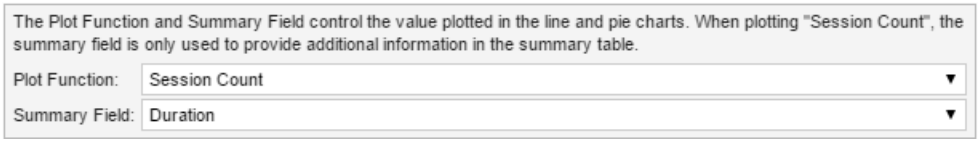

Monitor Properties – Columns

This property tab allows you to choose which columns to display in the detail record listing. To add a column, simply select it from the list at the bottom. To remove a column, click the “x” button on the right. To change the order columns are presented, click and drag them around by the selection icon on the left.

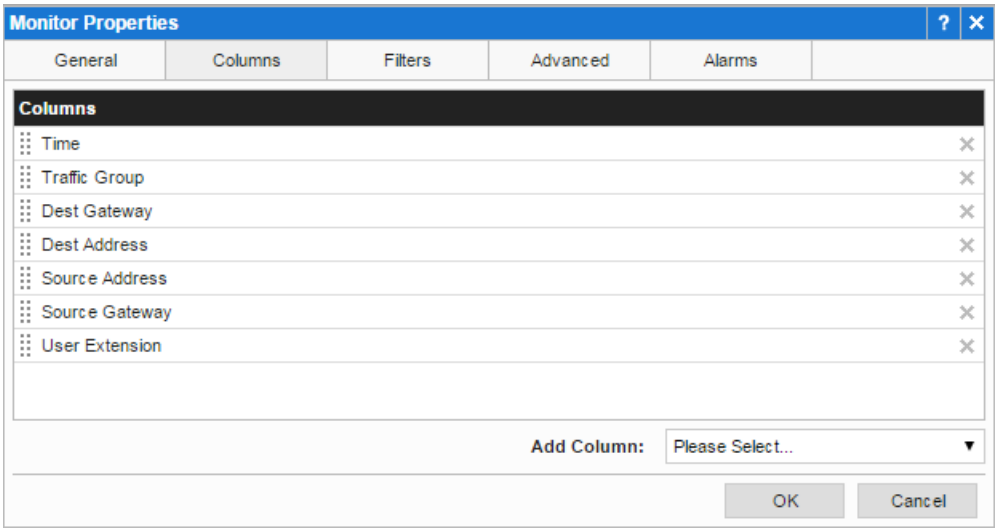

Monitor Properties – Filters

This property tab allows you to specify filters on the monitored data. Refer to the "Common Features" section of this document for a detailed explanation on how to configure filters. The filters look and work exactly the same way for monitors, and you can even reuse stored filter sets between reports and monitors.

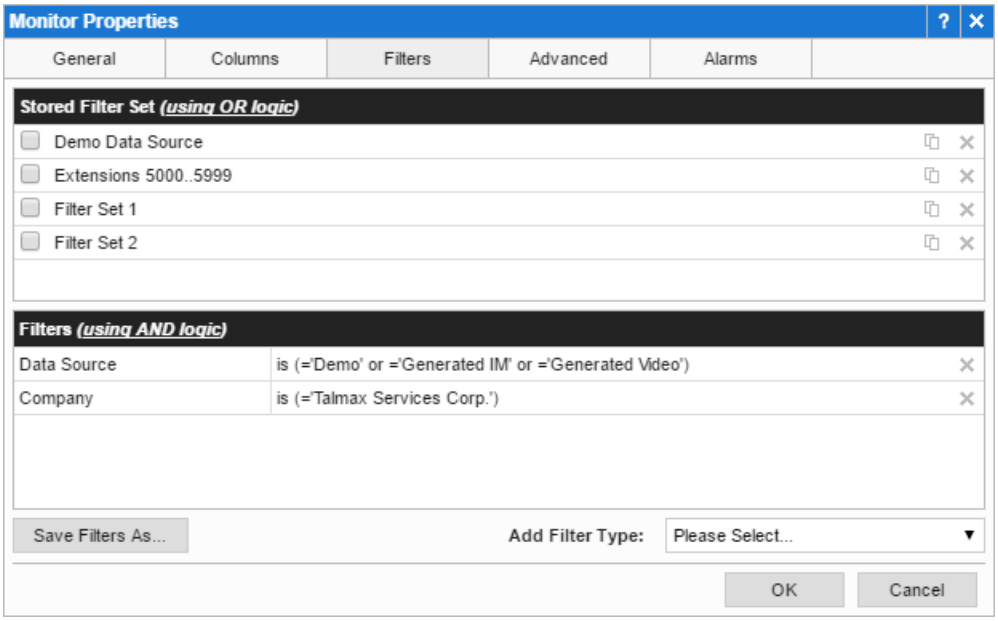

Monitor Properties – Advanced

The advanced tab allows you to select a time zone and specify the maximum number of detail rows to show. It also allows you to switch from a persistent monitor that watches traffic 24/7 in real-time (the default) to a snapshot monitor. A snapshot monitor is a monitor that goes back to a specific date and time and rebuild the monitor as it would have appeared then. Because the time period of a snapshot monitor is fixed, it does not update in real-time and it does not support alarms. Persistent monitors stay running in the background even when you log out. Snapshot monitors are temporary.

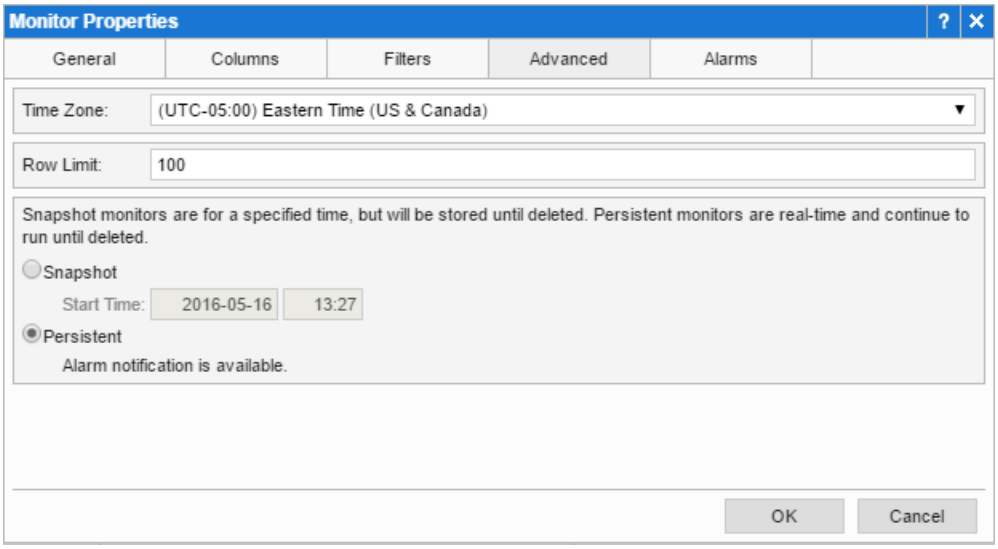

Monitor Properties – Alarms

If the monitor is persistent, you will be able to enable alarms for it. When you enable alarms, you must specify at least one set of alarm criteria and an alarm delivery method. There are a number of alarm delivery methods to choose from. For example, you can send a simple text email to alert a specific person/group, generate an SMTP trap, send an IM message to Cisco Spark, or run a report and have that delivered to someone automatically. As with the report delivery methods, the alarm delivery methods are scriptable so we can create custom delivery methods for you as a professional service.

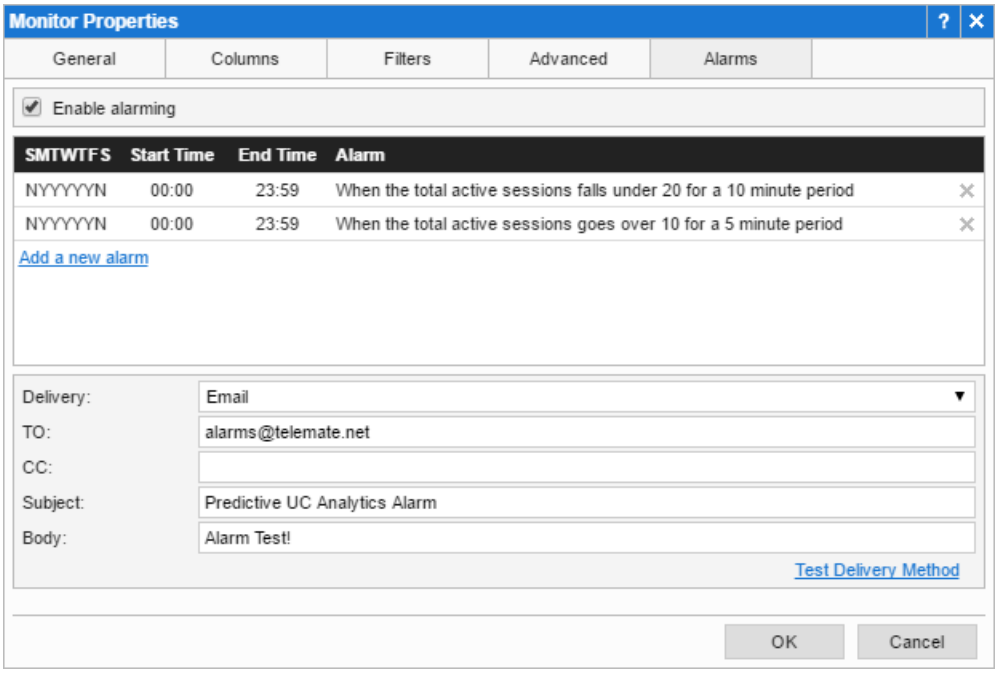

When adding or editing alarm criteria and thresholds, start by choosing the days of the week and times of day you want each alarm threshold to be active. If you need to set a special threshold for a night shift from 8:00PM to 5:00AM, simply set the active time to go from 20:00 to 04:59. After that, you must

specify the threshold values that trigger the alarm. As you change the threshold values, a sentence will be updated below to confirm what conditions will trigger the alarm.

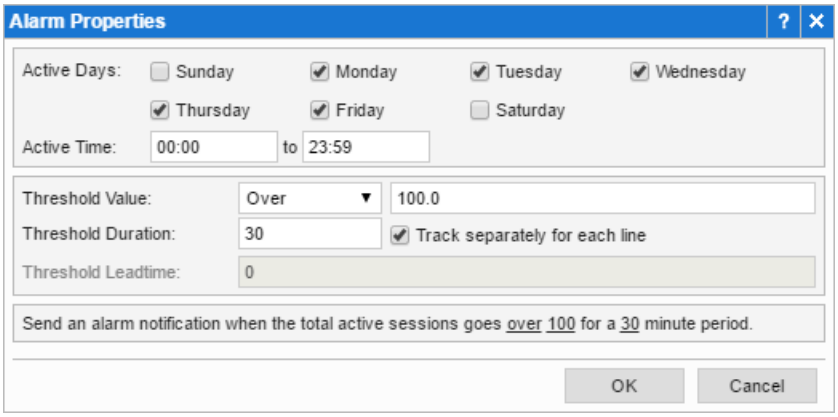

There are three basic types of alarm thresholds: under, over, and single session. The “over” and “under” thresholds are very similar. Put simply, they generate an alarm when any line on the monitor chart goes over or under a certain value for the specified duration. If you set a threshold to “Over 10 within 5 minutes” and the lines are configured to show Session Count or Concurrent Count, an alarm will be generated as soon as a new session is added to the monitor that causes 11 active or concurrent sessions to appear within a 5-minute window. If the lines show an accumulated value like total cost, an alarm will be generated if the total cost for any 5-minute window exceeds $10.00 (costs assigned to sessions with long durations will be pro-rated into 1-minute intervals).

The “under” thresholds work the same way with one exception. UC sessions that can have long durations generally are not logged until the session ends. If the average session duration is 10 minutes and the engine checks the most recent 5-minute window to see if the total is “under” a certain threshold, you will get false alarms. To account for this, you may set a threshold leadtime to tell the engine how many minutes to wait for sessions to finish before generating an “under” alarm. You may also compensate by lowering the threshold value to account for a certain number of active sessions.

The threshold duration also acts like a “snooze button” duration to avoid generating the same alarm repeatedly for overlapping time periods. If you set the threshold to 10 minutes and an alarm is generated for the 12:13 – 12:23 time window, the next allowed alarm would be for the 12:23 – 12:33 time window. When the “Track separately for each line” box is checked, each separate line tracked by the monitor can generate its own alarm every 10 minutes. Otherwise, the entire monitor will wait 10 minutes before sending another alarm.

Note: It is possible to configure overlapping thresholds with conflicting settings. This should be avoided in general, but you may actually want to configure both “under” and “over” alarms for the same time period to ensure that traffic stays within an acceptable range.

Last but not least, if you want to generate an alarm for every single traffic session (e.g. emergency 911 calls), simply configure an “Over 0.0 within 0 minutes” threshold. The 0 duration tells the engine to look at individual sessions instead of totals over time, and the 0.0 value tells it that every session should

generate an alarm. As a side note, if you change the value to 10.0, the engine will still look at individual sessions, but it will only generate an alarm if the “Summary Field” (set in the General tab) exceeds 10.0.

Warning: Be very careful with “single session” alarms, as they can very quickly generate a flood of alarms that can create a mess to clean up. Before you enable one, please double-check the monitor to ensure that it has very few calls appearing on the line chart. Also keep in mind that if a session does not appear on any line in the chart (it may appear in the detail listing but not on any line), it will not trigger an alarm.

Monitor Controls

Once you have created a first monitor and you start seeing data show up in it, the next thing to figure out is what it is showing and what the controls on it are for. Clock-wise from the top-left corner, you should see a line chart, a pie chart, a summary table, and a detail record listing of the most recent traffic sessions.

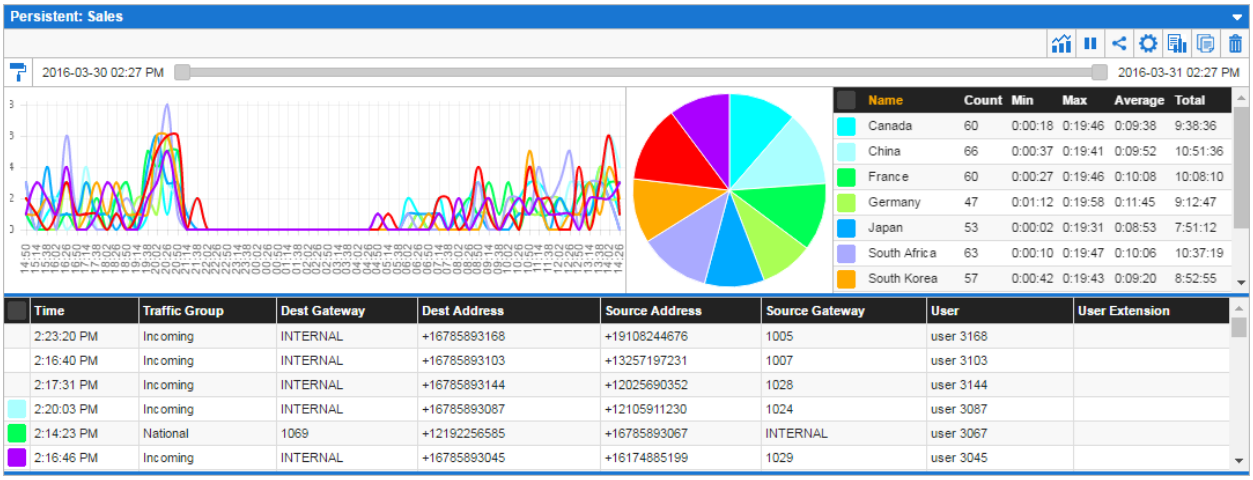

From left-to-right across the top, the controls are:

- **Paint Roller Icon –** Toggles fill mode for the line chart.
- **Time Slider** – Narrows the time range displayed in the line chart, pie chart, and summary table. Want to see the last 1 hour instead of the last 24 hours? Click and drag the block on the left all the way to the right.
- **Chart Icon** – Use to show/hide any of the 4 sections of the monitor. Also contains a selection to copy the display settings from the current monitor to all the other monitors you have access to.
- **Pause Icon** – If you’re trying to look at specific detail records but they keep scrolling out of visible range, click this icon to pause display updates for this monitor temporarily.
- **Properties –** Click on this icon to reopen the property sheet you used to create the monitor to change its properties.
- **Copy –** Click on this icon to create a copy of the current monitor. The property sheet for the new monitor is opened first to allow you to make changes before saving it.

- **Delete –** Click this icon to delete the current monitor.

In addition to those controls, you may resize the charts and summary table vertically and horizontally as needed. The record listing may only be resized vertically. If you have multiple monitors, you can click on their title bars and drag them up or down to view them in the order you want.

Performance Tips

The Predictive UC Analytics web server maintains all persistent monitors for all web users on the server in real-time 24/7. Whether it has to maintain 1 monitor or 1000, it has the same impact on the database server because it only fetches new data once to feed all of the monitors. However, memory and CPU usage on the web server will increase for each new monitor created. If several managers need to see a similar monitor, we recommend creating one monitor and sharing it with a security group. You also may want to carefully limit which managers are allowed to create their own monitors. The administrator can create and either share or assign them to others. An administrator can also use the "Admin View" feature to check and see how many monitors other managers have configured.

Snapshot monitors are not real-time, and they require a separate query to be run to fetch a specific 24-hour period of data. Depending on your call volume, that may run very quickly or it may take a while.

The ViewPoint Component

The ViewPoint component displays a real-time world map of the most recent 24 hours of UC traffic summarized by location. You may configure several views each with different filters, location groupings (e.g. state vs. area code vs. city vs. LATA), size and color criteria, and columns for traffic detail listings. For any view, you may select areas of interest to see a listing of calls going to/from specific locations. **Note:** The term "real-time" depends on how fast your traffic logs are collected and processed, which can vary widely based on the type of data source generating the logs.

Example 1: (calls displayed as different size and color dots by LATA/city)

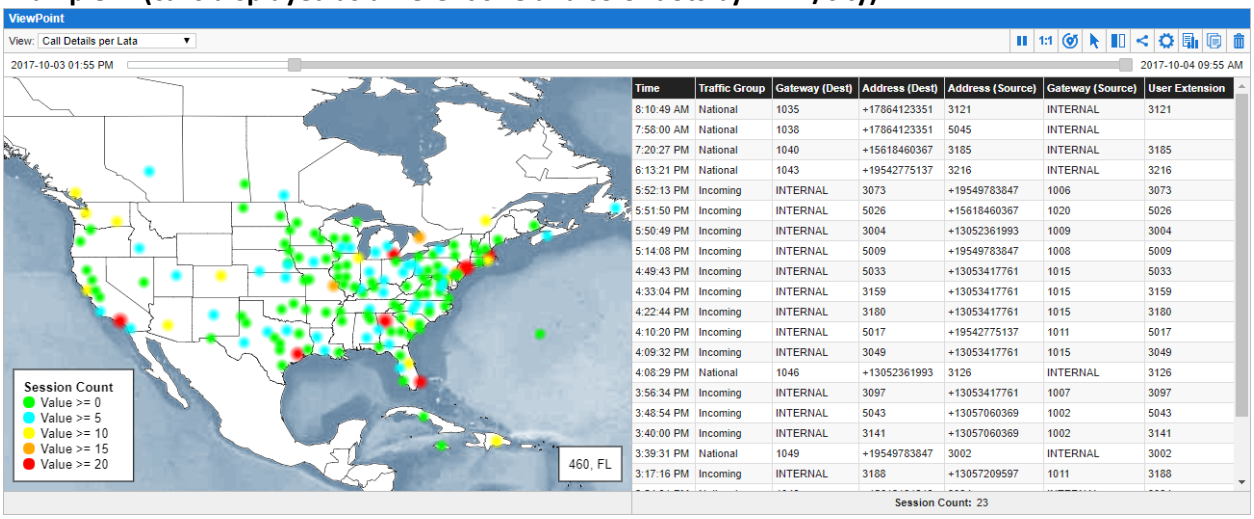

Example 2: (calls displayed as different fills by state/country)

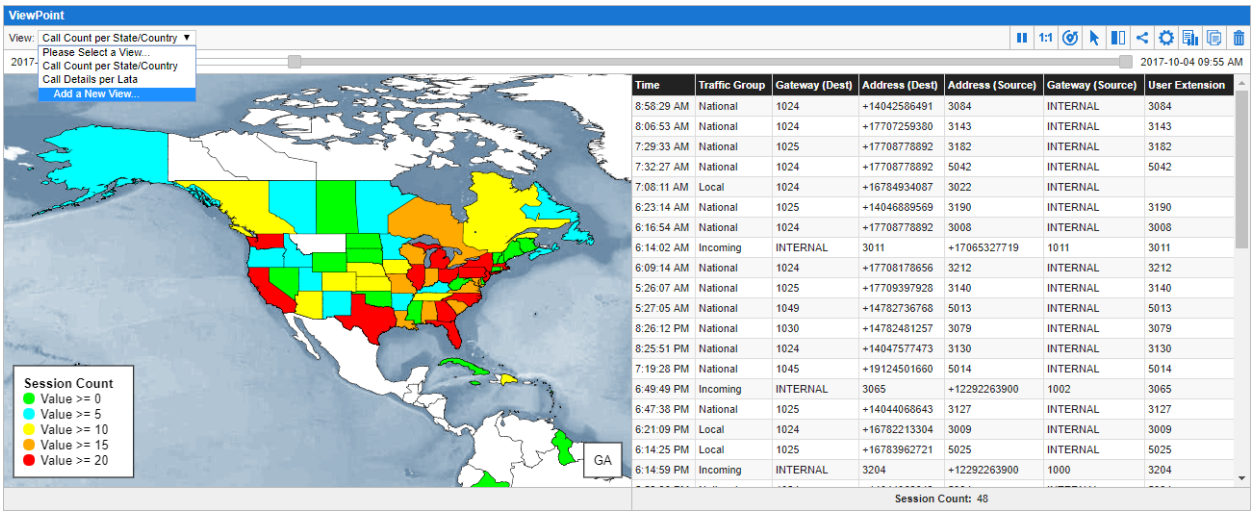

Example 3: (calls displayed as lines for site-to-site)

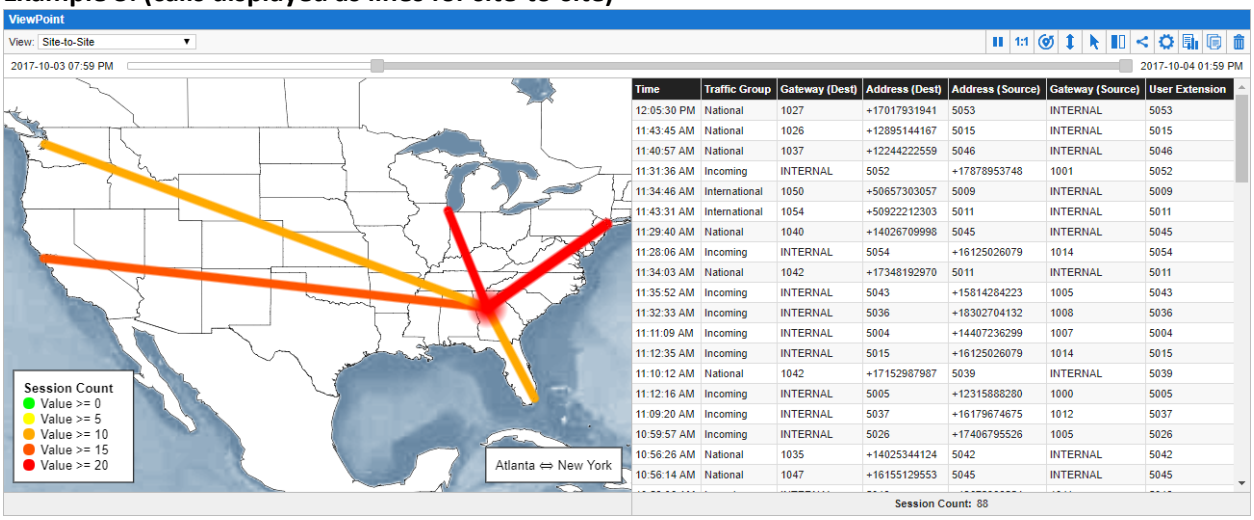

Example 4: (calls displayed as directional lines for site-to-site)

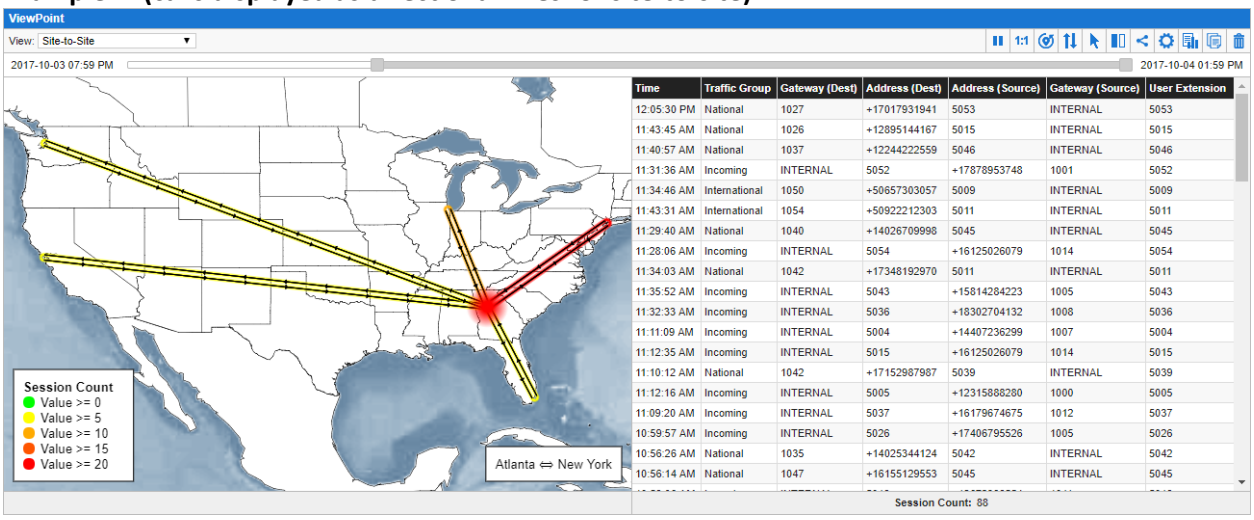

Creating Views

Because you may have several different concerns regarding your data, you can create several different world map views that you can switch between quickly. For example, you may wish to create one view to show you all inbound calls by area code and another to view all outbound calls by state. Or you may wish to create one view to monitor division A and another to monitor division B. To create a new view, click on the "View" drop-down box on the left side of the option bar at the top and select "Add a New View". That will open up the ViewPoint properties window.

ViewPoint Properties: General

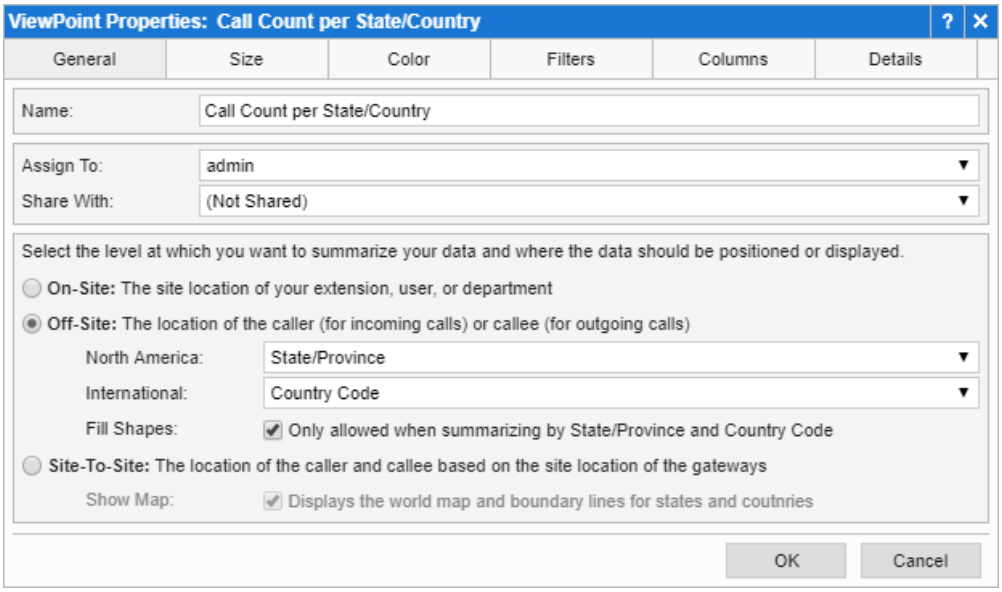

When you create a new view, you must first specify a unique name. After that, choose Off-Site (external address locations), On-Site (internal site locations), or Site-To-Site (network segments modeled using site locations of each internal endpoint). External address locations are determined by the external phone number (e.g. country code and city code), GeoIP location for public IP addresses, etc. Internal site locations can be assigned by data source, gateway group, gateway, or private IP range. They can also be overridden for specific departments, users, or internal resources. With Off-Site you must also specify how you want to summarize calls in North America (by State/Province, Area Code, Lata, or Rate Center) and everywhere else in the world (by Country Code or City Code). If you select State/Province for North America and Country Code for International you can choose to fill in state/country shapes rather than drawing points. However, by choosing state/country fills, you lose the ability to use the point size to represent a summarized value.

ViewPoint Properties: Size

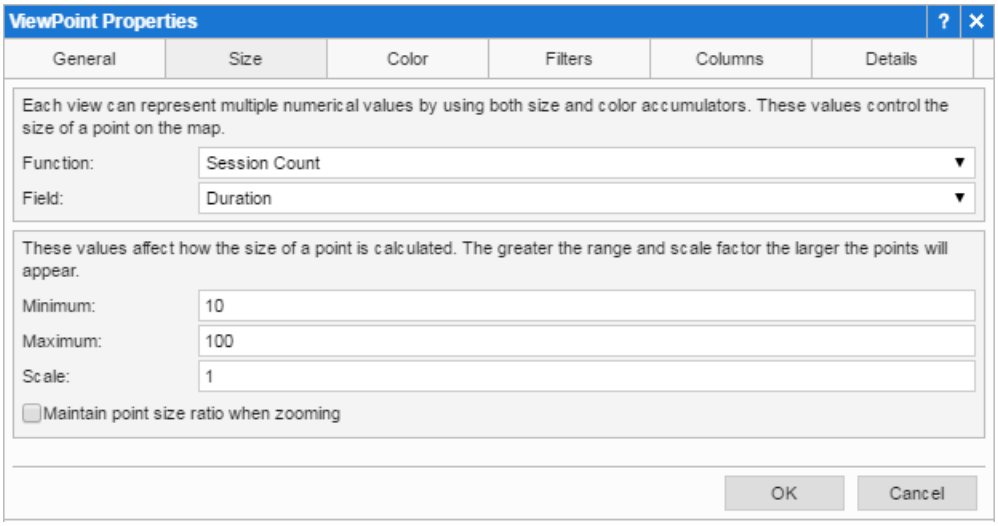

These properties allow you to control the size of the points rendered in this view. For example, you could have the view use total call count, total duration, maximum call cost, or minimum call quality to determine the point size.

The minimum and maximum size define the range of sizes the points use, and the scale is a multiplier that controls how quickly the points go from the minimum to the maximum. The ideal scale value can vary widely based on several factors (call volume, filter selection, summary selection, and accumulator type), so we recommend you try a few different values for each view to see what looks best for each. The “Maintain point size ratio when zooming” checkbox determines whether or not point sizes remain the same size (in pixels) when zooming.

ViewPoint Properties: Color

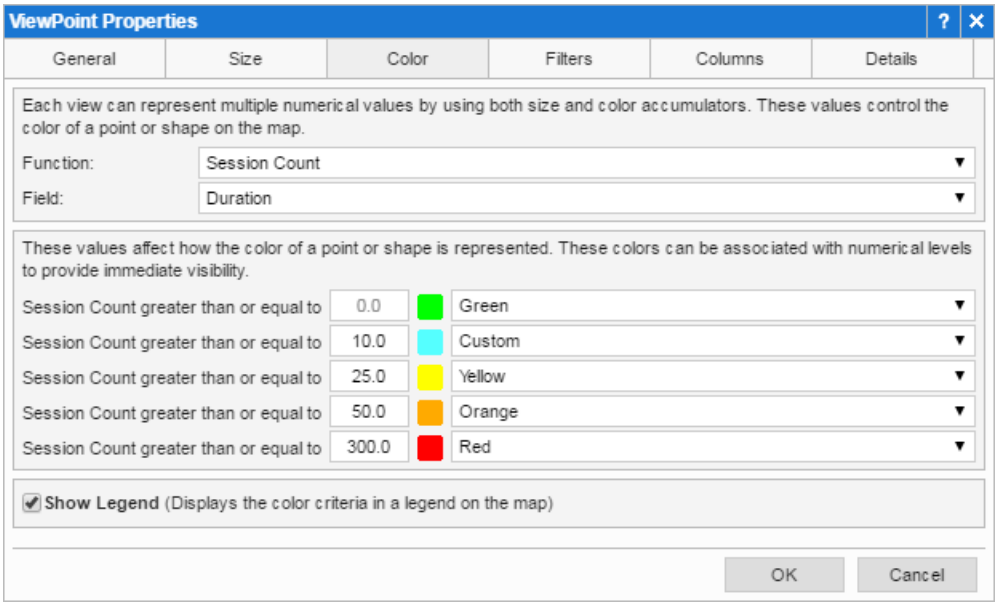

This property tab allows you to control the color of the points or filled regions rendered in this view. Where point sizes grow/shrink smoothly on a logarithmic scale, the colors are fixed based on a set of thresholds you define. For example, you may want to have a region turn red if the total call cost for a region exceeds $100, or perhaps if the maximum or average jitter value for that region exceeds 35.

A legend will be visible in the lower left hand corner of the map for easy reference. It can be hidden by unchecking the Show Legend option.

ViewPoint Properties: Filters

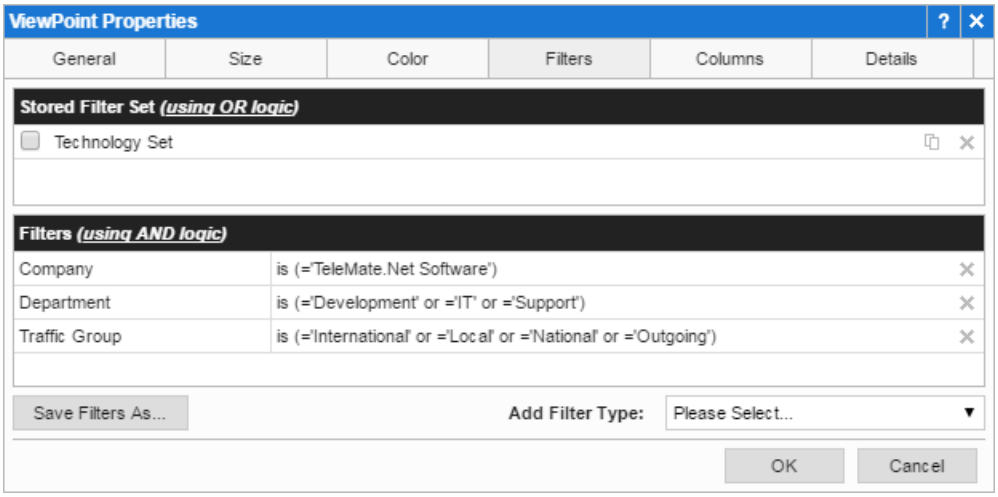

This property tab allows you to specify filters on the traffic included in this view. Refer to the reporting section of this document for a detailed explanation on how to configure filters. The filters look and work exactly the same way for viewpoint, and you can even reuse stored filter sets between reports, monitors, and viewpoint views.

ViewPoint Properties: Columns

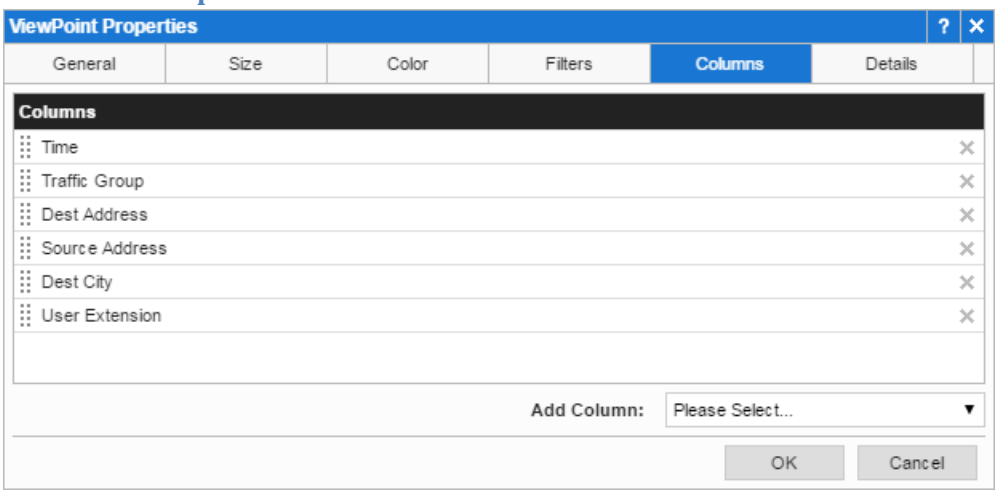

This property tab allows you to choose which columns to display in the detail record listing. It is identical to the "Columns" tab for monitors. To add a column, simply select it from the list at the bottom. To

remove a column, click the “x” button on the right. To change the order columns are presented, click and drag them around by the selection icon on the left.

ViewPoint Properties: Details

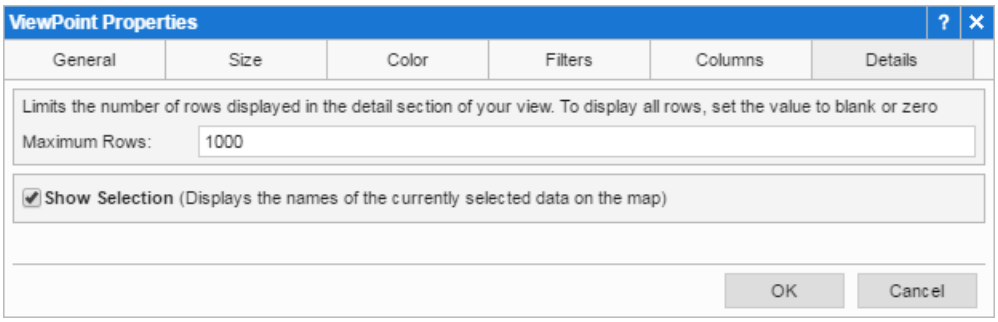

Additional options for the Detail section are located under the Details tab. The Maximum Rows limits the number of rows displayed in a detail listing. If you set the value to zero or leave it blank it will return all rows for selected location and time period. However, returning to many rows could affect the performance and responsiveness of your browser. Additionally, there is an option to display the selected location’s name on the map. The selection’s name is displayed in a box in the lower right corner of the map.

Viewpoint Controls

Once a view is created and selected, you can use the mouse scroll wheel to zoom in and out, you can click-and-drag to pan and see different portions of the map, and you can click on highlighted points/regions to display detailed call info. There is also an option bar at the top with icons you can click for more options:

From left to right, the controls are:

- **View Selection** – Allows you to switch quickly to a different view or create new ones.
- **Time Slider** – Narrows the time range displayed in the map and in the detail listing. Want to see the last 1 hour instead of the last 24 hours? Click and drag the block on the left all the way to the right.
- **Pause/Play –** Click to pause the real-time updates. This is useful when you need to examine specific calls without them moving around. Click again to resume the real-time updates.
- **1:1 –** Restores the zoom to a scale of one-to-one (100%)
- **Point/Line Location Reset –** Restores the default location of the points or lines if they have been manually moved on the map.
- **Cursor –** Toggles the mouse cursor selection mode. The current options are single selection mode (which allows you to select individual points/regions or click and drag to pan the map) and box selection mode (which allows you to click and drag to select a box of points/regions).

- **Layout –** Toggles the orientation of the map and detail section. There are two vertical layouts with the call listing below or above the map and two horizontal layouts with the call listing to the left or right of the map.
- **Properties –** Allows you to edit the properties of the current view.
- **Copy –** Copies the current view's properties to make it easier to create a similar view.
- **Delete –** Allows you to delete the current view when you no longer need it.

In addition to those commands, you may resize the detail section. When the layout orientation has the detail section at the top or bottom the details height can be resized. With detail section on the left or right you can adjust the width of the detail section and the map.

The Reporting Component

Select the reporting component when you need more comprehensive traffic information. Reports are needed to see more detailed traffic information, different kinds of summaries, or to apply more advanced filters and options. Reports are also needed to save traffic information to documents like PDF or Excel and distribute it via email, FTP, etc. Reports can be scheduled to run automatically every night/week/month for you to review.

In addition to the navigation menu, the reporting component has an additional menu on the left side of the page. This menu contains links to completed reports, stored reports, and configurable report templates.

Configuring Reports

When you want to run, save, or schedule a new report, select it from the list of available report templates grouped by report categories. Keep in mind that the reports available to you will depend on the features your product is licensed for and the data you are processing. For example, if you have not configured any expenses and run an expense distribution, you will not even see the "Expenses" category. When you select a report template from the list, it will display a set of options and filters for you to choose from before you run, save, or schedule the report. Note that each report template may have a different set of options and filters.

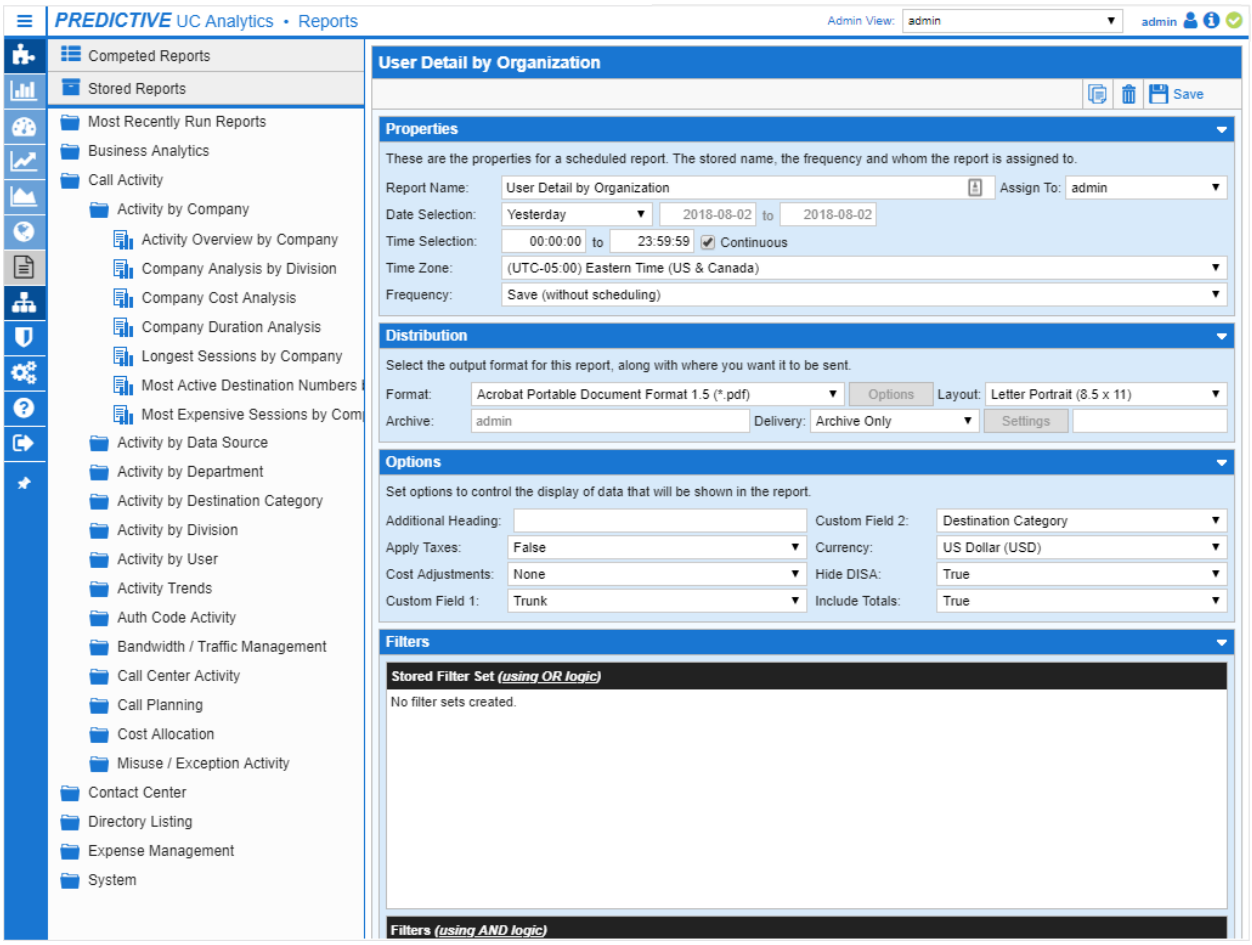

Commands

There are three commands available when configuring a report: copy (only available for stored reports), delete (only available for stored reports), and run/save (which changes depending on the selected frequency).

Properties

The most critical properties when configuring a report are the date range you want to run it for and whether you want to run it immediately, save the currently selected options so you can reuse them later, or schedule it to run automatically. If saving or scheduling it, you must specify the name you want it to show up under in the “Stored Reports” list. If scheduling it, additional options will become available to let you define the frequency.

Note: If you are logged in with an admin account, you can choose to assign a saved/scheduled report to someone else’s account.

Distribution

These options allow you to choose the report’s file format (PDF, Excel, RTF, CSV, TXT, or HTML) and where to deliver it. All completed reports will automatically be stored in an archive, which equates to the user’s “Completed Reports” list. However, completed reports can also be delivered automatically via

email, FTP, Dropbox, Cisco Spark, etc. The Predictive UC Analytics delivery mechanism allows custom scripts to be written to provide custom delivery methods.

Note: If you are logged in with an admin account, you can choose to deliver the report to another account's archive (their completed report list).

Options

These optional formatting options allow you to customize certain aspects of how the final report looks. One example of this would be to remove all charts from a report, which can make the resulting report file significantly smaller, which may be necessary to get around mail server attachment size limits. Each report can have completely different options. For instance, a report with no charts would not have an "Include Charts" option. Some options may be available in many reports, but others may only show up in one report. If you are not certain what an option does, allow the mouse cursor to hover over its name to see a longer description pop up over it. If you are still not certain, try running the report with different options and look for differences.

Filters

Refer to the "Common Features" section of this document for a detailed explanation of how to configure filters.

Stored Reports

Select "Stored Reports" when you want to view, manage, or copy saved or scheduled reports.

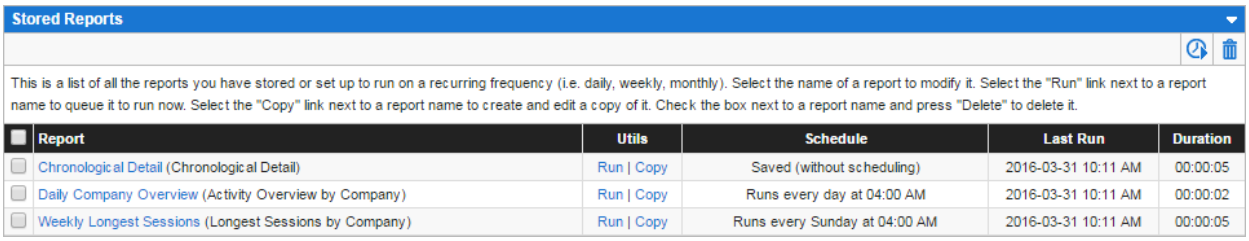

Click the name of a stored report to modify it. Check the box next to it and click the trashcan icon button to delete it. Click the "Run" link to make a copy and queue it to run immediately using all the options you have selected for that stored report. Click the "Copy" link to start modifying a copy of the report, which will be saved under a different name. Note that the "Schedule", "Last Run", "Duration", and "Next Run" columns in the table are for scheduled reports only.

Completed Reports

Select "Completed Reports" when you want to view or manage reports that have recently finished. If you have queued up reports to run that have not completed yet, they will also be shown here in a "Pending Report" list.

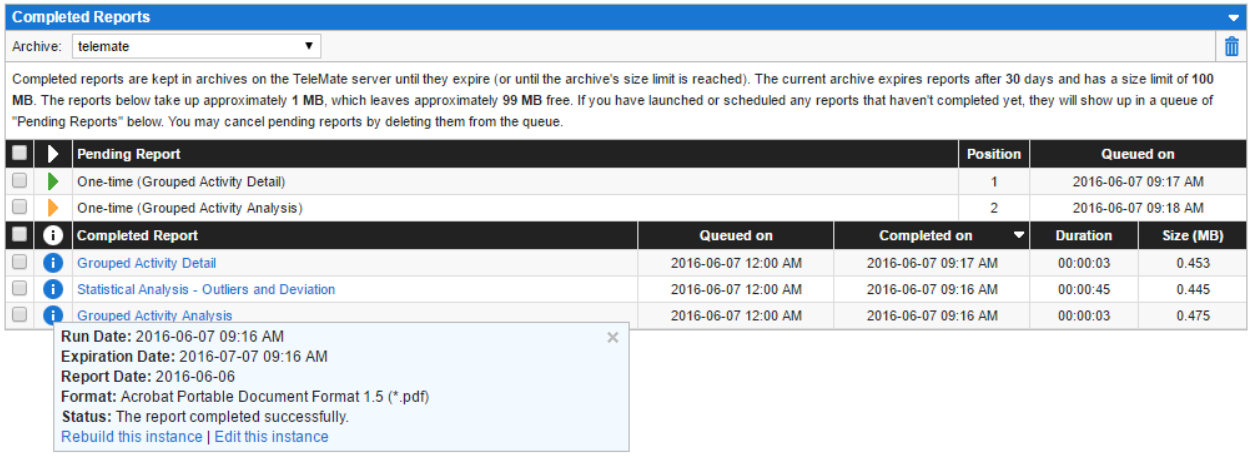

In the screenshot above, one report is in the queue and currently running, a second is in the queue waiting to run, and three completed reports are shown. You may delete pending or completed reports by checking the box next to it and clicking the trash can icon (in the top-right corner).

The icon next to a completed report is very useful. A blue “i” (for info) indicates that it completed without warnings or errors. A yellow exclamation point indicates that it generated a warning. A red “X” indicates that it failed to complete. In all cases for a completed report, you may click on the icon for more information. Note the “Rebuild this instance” and “Edit this instance” links at the bottom in the screenshot above. The first option will delete and rebuild the completed report with the exact same options, and the second option will take you back to the configuration page to allow you to tweak them and run it again.

Click on the name of the completed report to open the report file. This will attempt to open the report in your web browser (e.g. using the Adobe PDF plugin). If you want to download the report file to open it locally, you can right-click on the link instead and select “Save As” (the exact name of the command will depend on which web browser you are using).

Assurance Features

The assurance features monitor and poll devices on your UC network. This component provides a view for IT admins to monitor current health and wellness, and to configure alarms to alert IT admins when they are not monitoring it.

The UC Devices Component

The UC Devices component provides a view into the UC clusters, services, voice gateways, and phones monitored by the assurance feature. The alarms provided in this component focus on “live” call traffic and “current” device or service status. The views provided in this component show live/current information, summary and detailed information for calls for the past week, and statistical predictions based on historical information in prior weeks. The predictions are useful for configuring thresholds for predictive alarms to notify you of statistically abnormal activity.

The analytics engine generates summary statistics and predictions based on logical groups. It will automatically create groups for every new CUCM cluster, location, device pool, trunk, and gateway detected. You may also create your own custom groups by creating custom filters. Once created, the analytics engine will start generating summary statistics and predictions on custom groups as well.

In addition to summary information, this component also provides a great deal of detail on individual call chains, which may aid in troubleshooting certain problems. For each call, the engine builds a chain of live call legs gathered from CUCM SIP trace logs, voice gateway SNMP polls, Expressway syslog events, and in some cases individual phone HTTP status pages. Assuming every device in the path of a call provides data, the engine can build very complex call chains from endpoint to endpoint (e.g. a call originating from cluster A, going through an SME cluster, and then terminating at cluster B, with multiple gateways in-between).

Getting Started

On the left, you should see a tree with root nodes named “Alarms”, “Custom Groups”, and “CUCM Clusters”. Click on the folder icon next to each of these section names to expand or contract the listing of child entries beneath them. Start by expanding “CUCM Clusters”, expanding a specific cluster beneath it, and then selecting “Device Locations” beneath the cluster. This will display a summary table of the current “live” status for every phone device in that cluster grouped by location. It will show a count of phone devices: actively in a call, registered but inactive, and unregistered. If QoS statistics are available for live calls, it will also show counts of live calls impacted by Jitter, Latency, and Packet Loss grouped by location.

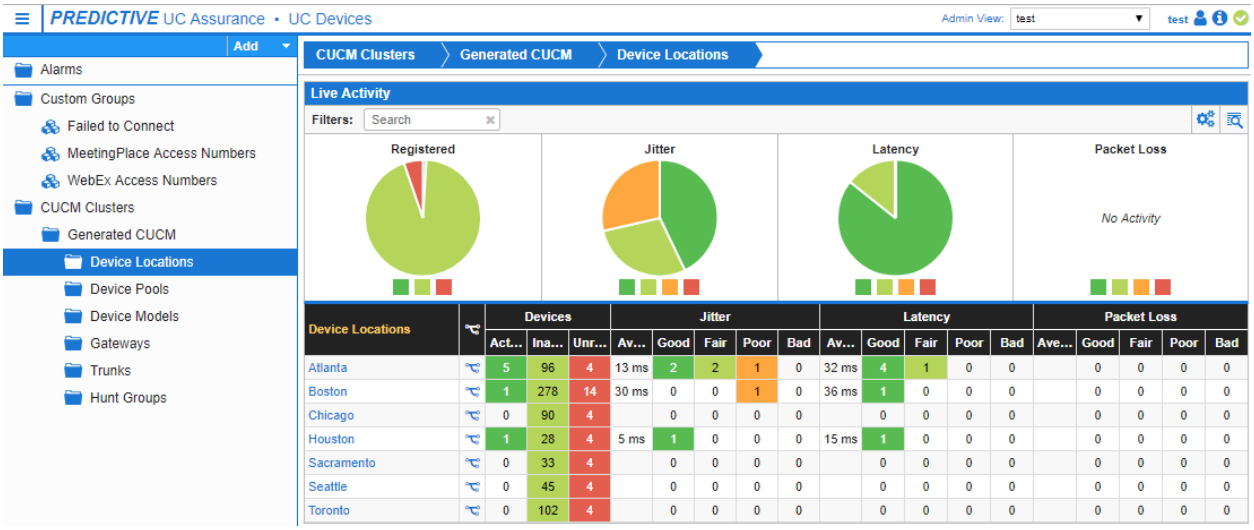

Device Locations	Devices			Jitter					Latency					Packet Loss				
	Act...	Ina...	Unr...	Av...	Good	Fair	Poor	Bad	Av...	Good	Fair	Poor	Bad	Ave...	Good	Fair	Poor	Bad
Atlanta	5	96	4	13 ms	2	2	1	0	32 ms	4	1	0	0		0	0	0	0
Boston	1	278	14	30 ms	0	0	1	0	36 ms	1	0	0	0		0	0	0	0
Chicago	0	90	4		0	0	0	0		0	0	0	0		0	0	0	0
Houston	1	28	4	5 ms	1	0	0	0	15 ms	1	0	0	0		0	0	0	0
Sacramento	0	33	4		0	0	0	0		0	0	0	0		0	0	0	0
Seattle	0	45	4		0	0	0	0		0	0	0	0		0	0	0	0
Toronto	0	102	4		0	0	0	0		0	0	0	0		0	0	0	0

Note: Any time you see a graph icon (🕸) in a group summary table, you may click it to see a network graph of all live call chains that belong to that group.

With the exception of “Hunt Groups”, you will see the same type of summary if you select any of the other groupings beneath the cluster, or any parent grouping including “Custom Groups” and “CUCM Clusters”. When you select a different level, the only difference will be the type of groups displayed in

the table. You may select a group name in the table to navigate down to the next level, or you can simply navigate to it in the tree.

Call/Device Listings for a Specific Group

When you select a specific group, you will see either a listing of phone devices (for a specific Device Location, Pool, or Model) or a listing of call chains (for a specific phone, gateway, trunk, or custom group). The listing of phone devices shows the current state of each phone, such as whether it is registered and whether it is currently in a call. That device-centric view can only show the current device state, so it cannot show recently disconnected calls. However, all other views will show a listing of both live and recently disconnected calls.

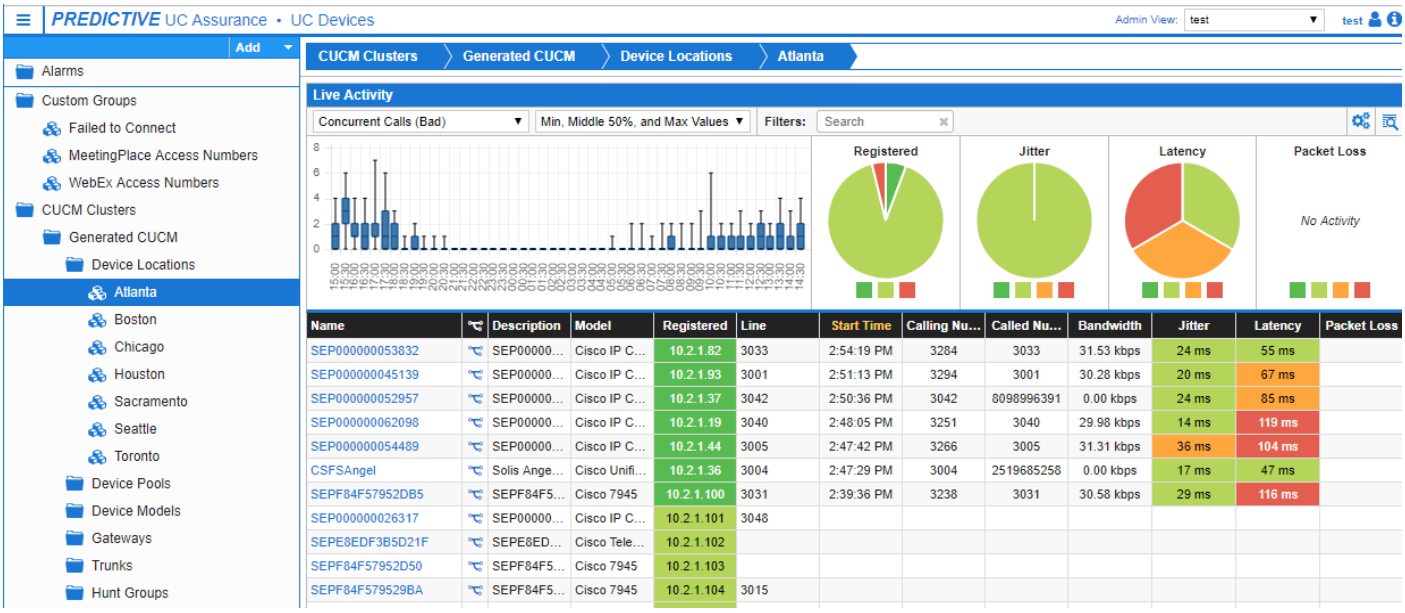

Note: Any time you see a graph icon (📈) in a device/call detail table, you may click on it to see details about that specific call chain.

The table listings can get quite large, so you can sort (ascending or descending) by any column in the table. You can also use the search filter to look for specific names, numbers, or IP addresses in the table. If you have pie charts showing counts of unregistered/idle/active devices, or call counts in different QoS ranges for Jitter, Latency, and Packet Loss, you can also click on slices of the pie charts to filter on the devices/calls they represent.

The line chart shows activity for the past 24 hours summarized into 30-minute buckets for the selected group. For device listings, you can choose unregistered phone device counts. For both call and device listings, you can choose overall Jitter, Latency, Packet Loss, and bandwidth, or choose concurrent call counts that meet various QoS levels. For each statistic, you can choose to see box plots (min and max plus 25th, 50th, and 75th percentiles) or the max, mean, or median values (with predictions). Predictions will not show up until after one full week of data has been collected and aggregated for the group. The predictions will continue to improve as more data is collected. After four weeks, predictive alarms may start to trigger.

Examples of “actual” versus “predicted” max, mean, and median values:

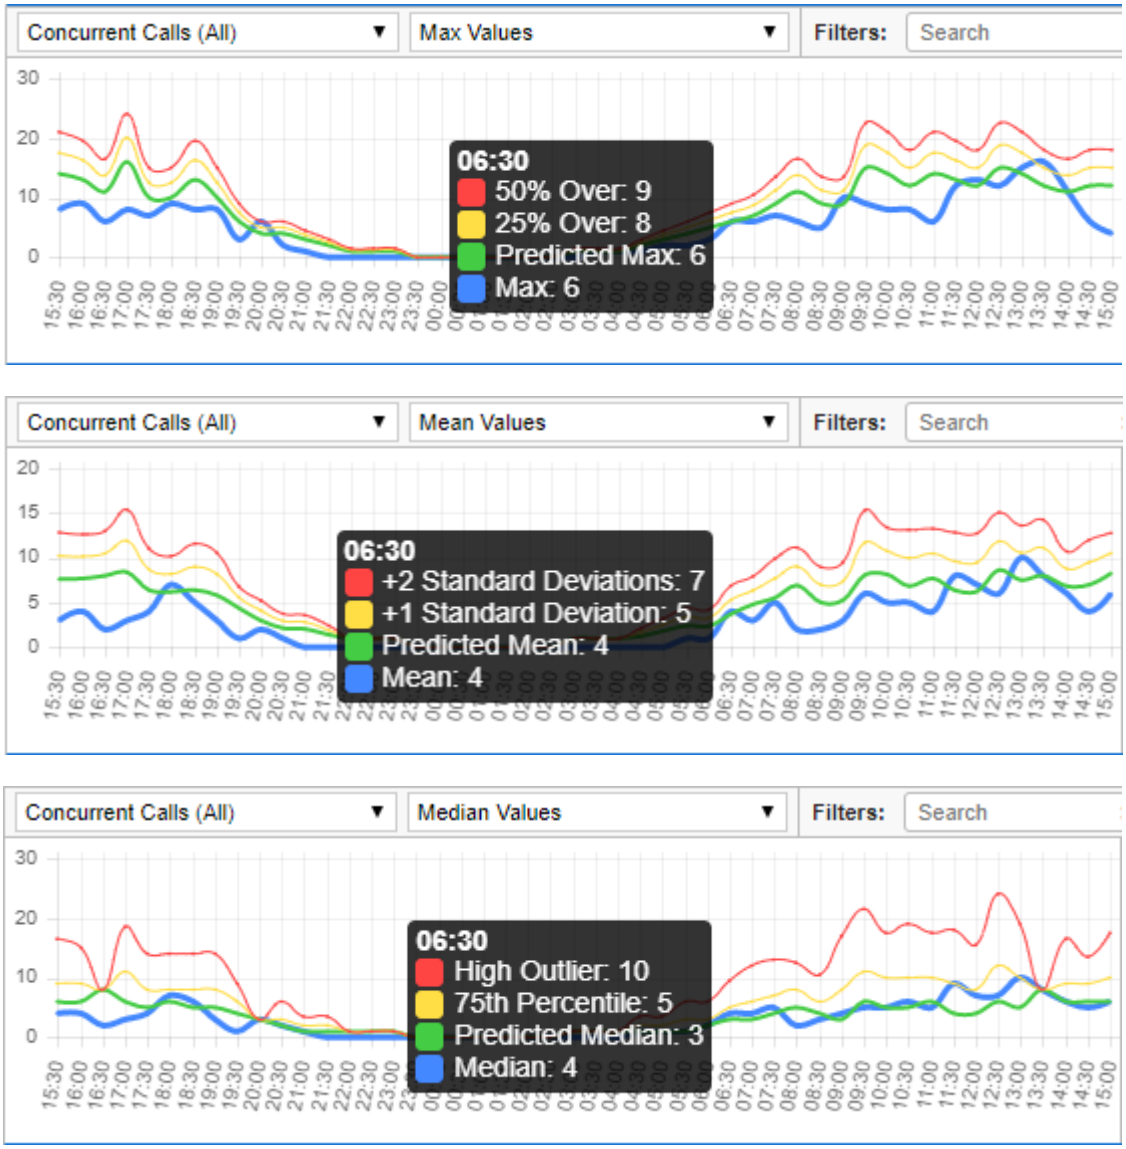

Important: The closer a predicted value gets to zero, the more likely it will be that any percent over will be so close to the predicted value that it will generate false alarms (e.g. when the red lines above get close to the blue lines). To account for that, predictive alarms allow a “minimum offset”.

Individual Call Chain Graphs

When you select the graph icon (📉) for a single call chain, you will see a detailed listing of the call legs included in the chain along with up to three diagrams:

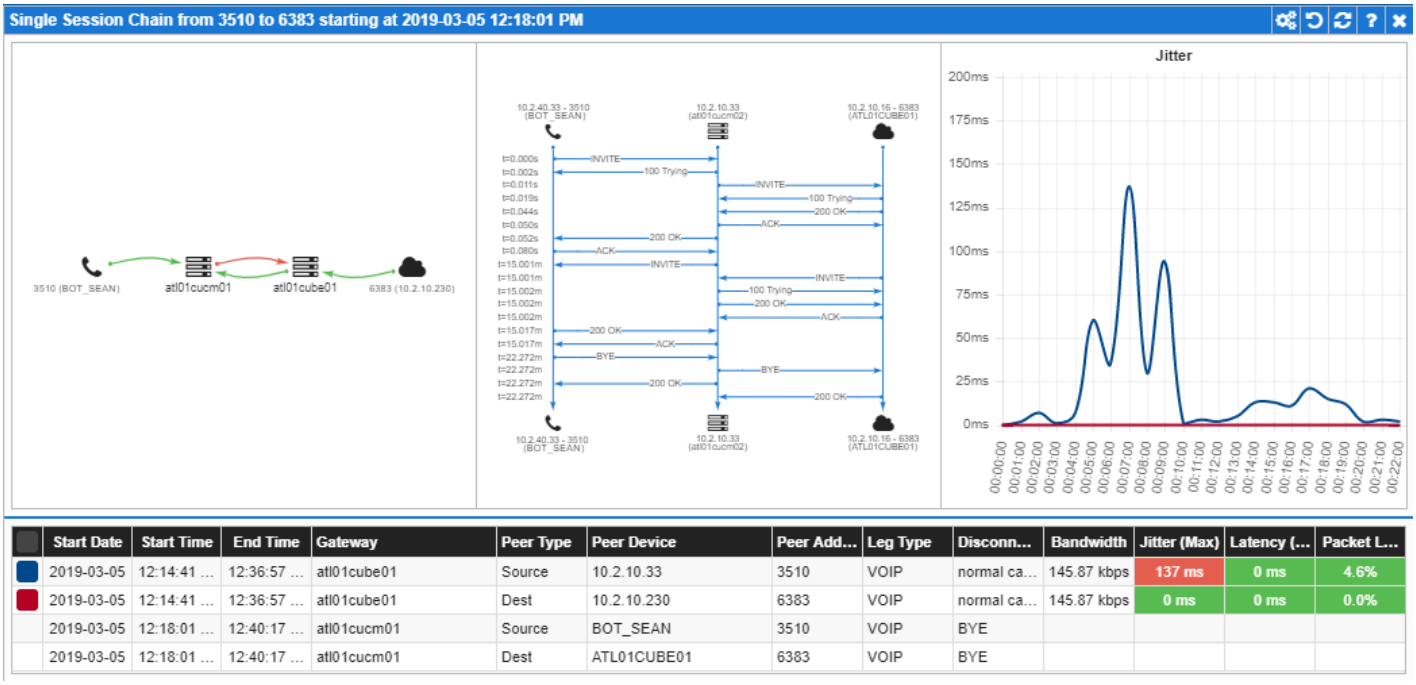

	Start Date	Start Time	End Time	Gateway	Peer Type	Peer Device	Peer Add...	Leg Type	Disconn...	Bandwidth	Jitter (Max)	Latency (...	Packet L...
	2019-03-05	12:14:41 ...	12:36:57 ...	atl01cube01	Source	10.2.10.33	3510	VOIP	normal ca...	145.87 kbps	137 ms	0 ms	4.6%
	2019-03-05	12:14:41 ...	12:36:57 ...	atl01cube01	Dest	10.2.10.230	6383	VOIP	normal ca...	145.87 kbps	0 ms	0 ms	0.0%
	2019-03-05	12:18:01 ...	12:40:17 ...	atl01cucm01	Source	BOT_SEAN	3510	VOIP	BYE				
	2019-03-05	12:18:01 ...	12:40:17 ...	atl01cucm01	Dest	ATL01CUBE01	6383	VOIP	BYE				

Note: A phone icon indicates an internal endpoint, which represents a "phone" device in one of your CUCM clusters. All other endpoints show up as a cloud icon on the network graphs.

The Session Flow Diagram (left) is always available, and shows how the devices involved in the chain connect to each other. Each arrow represents a leg, and each leg can change color based on available QoS statistics. The SIP Ladder Diagram (center) is only available if a SIP trace exists for the chain. The Session Snapshot Diagram (right) is only available if a Voice Gateway or CUBE provided bandwidth and QoS stats for the legs over time. The icons in the top-right corner are:

- Gear (⚙) – Contains view options to show/hide any of the 3 diagrams, and to change the Session Snapshot diagram to show either bandwidth, jitter, latency, or packet loss.
- Reset (↺) – Reset positions and scale in the graphs.
- Refresh (🔄) – Refresh the content to reflect the latest live call info.
- Help (?) – Open the web help page.
- Close (✖) – Close the session chain dialog.

Group Call Graphs

When you select the graph icon (📈) for an entire group, like a Device Location, it will show all current live call chains for the selected group. Because this can show hundreds, or even thousands, of call chains and endpoints, internal and external endpoints will be put into logical groupings to avoid performance and visibility issues inherent in trying to render so many on such a small diagram. It is possible to change the groupings, or to show all endpoints if you wish.

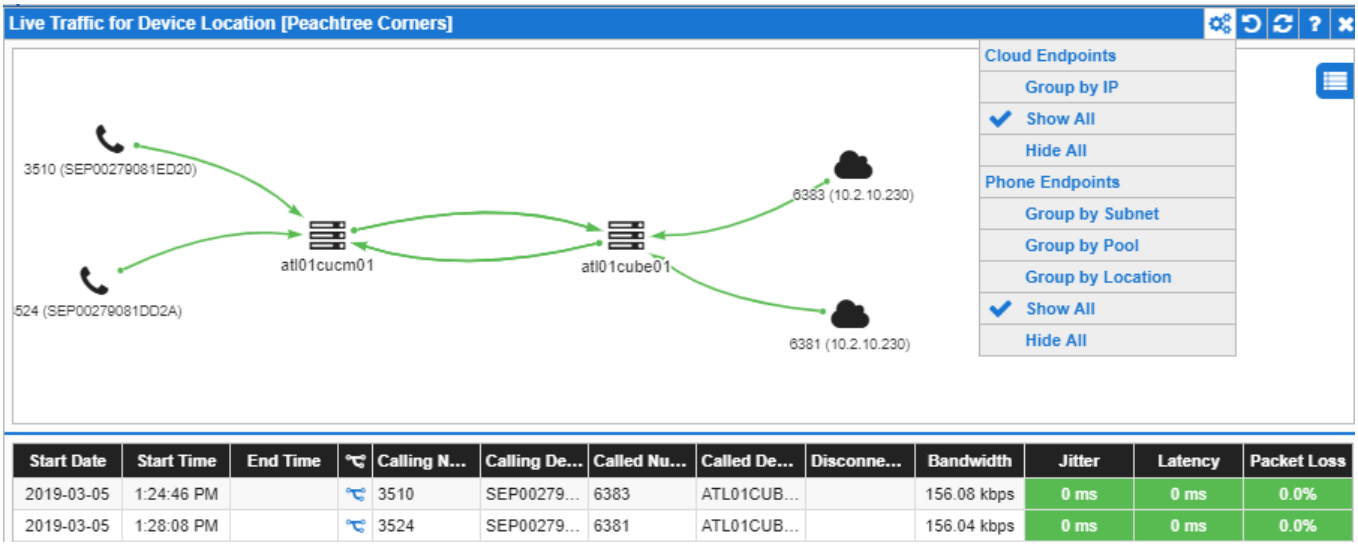

Start Date	Start Time	End Time		Calling N...	Calling De...	Called Nu...	Called De...	Disconne...	Bandwidth	Jitter	Latency	Packet Loss
2019-03-05	1:24:46 PM			3510	SEP00279...	6383	ATL01CUB...		156.08 kbps	0 ms	0 ms	0.0%
2019-03-05	1:28:08 PM			3524	SEP00279...	6381	ATL01CUB...		156.04 kbps	0 ms	0 ms	0.0%

Selecting specific legs in the diagram will filter out call chains in the table below. Note that each call chain also has its own graph icon ( ) to see more details about the specific call chain. The gear icon ( ) allows you to change the grouping options for endpoints.

Custom Groups

Custom groups allow you to apply filters to call chains in ways that would be impossible for the built-in cluster and device type groups. Examples of this would include:

- Isolating calls to and from public WebEx telephone numbers.
- Isolating inter-cluster call chains that include clusters A and B.
- Isolating inter-gateway call chains that include gateways A and B.
- Isolating call chains that had zero “connected” legs.
- Isolating call chains with legs that had particular disconnect codes/reasons.

As with built-in groups, as soon as the group is created, the analytics engine will begin aggregating data for it, and after 4 weeks of data has been collected, predictions will be made and predictive alarms can be triggered.

To add a custom group, hover the mouse cursor over the “Add” link at the top of the tree and select “Add Custom Group”. Enter a unique name for the group, then click the “+” button on the top-right corner of the Filter Group list to add at least one filter group.

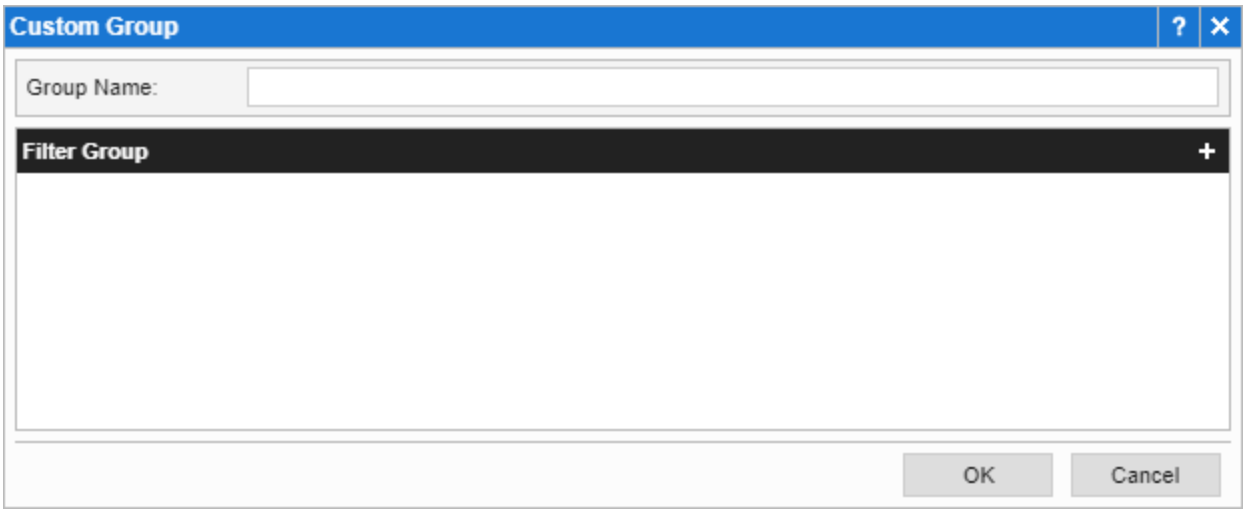

Each filter group should isolate a specific type of call chain. Add multiple filter groups to have your custom group combine multiple different types of call chains in a single group. When adding a filter group, the dialog and controls are very similar to the filter dialog used everywhere else in the product. However, the way they are applied to chains of call legs instead of individual calls makes them different.

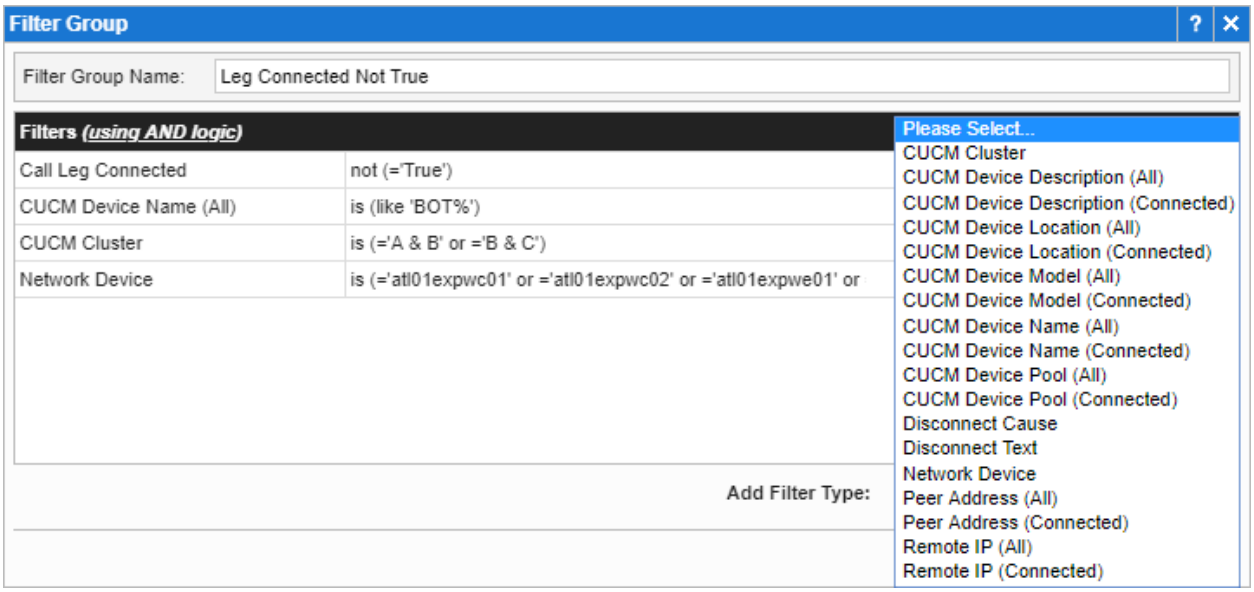

There are a number of important logical distinctions to make with the filter logic concerning “legs” versus “chains”:

- A single “Include” (or “is”) filter will include an entire call chain if at least 1 leg in the chain matches. This works the same way whether looking for a specific value or one of a set (or range) of values.
- The ampersand (&) character has a special meaning in call chain filters. It allows you to change the filter meaning from “one of a set” to “all of a set”, so the “CUCM Cluster” filter in the screenshot above would include inter-cluster calls between clusters A and B or between clusters

B and C. Using “A & B & C” would require a call chain to have 1 leg from each of the 3.
Note: The spaces between names and ampersands are optional.

- Multiple “Include” (or “is”) filter types will include an entire call chain if each include filter applied separately matches the chain. If any fail to match, the call chain will not be included.
- “Exclude” (or “not”) filters override include filters, and will exclude an entire call chain if any leg in the chain matches any exclude filter.
- The “Call Leg Connected” filter is unusual. Set it to “is = True” to only include call chains with at least 1 connected leg, and “not = True” to exclude call chains with at least 1 connected leg (e.g. to keep chains with 0 connected legs).
- Combining “Call Leg Connected” “is = True” with a filter on a specific endpoint will not necessarily give you call chains connected to that endpoint. It will include chains linked to that endpoint but answered by a different endpoint.
- To account for that, each endpoint filter has two types:
    - Endpoint (All) – Include chains linked to this endpoint regardless of whether they are connected or not.
    - Endpoint (Connected) – Include chains that actually connected to this endpoint.
      Note: The definition of “connected” depends on the endpoint. For a ringing phone device, “connected” usually means “answered”. However, a phone system may “connect” inbound calls automatically before routing the call to other devices, so “connected” will not mean “answered” for every call leg.

Hunt Groups

The Hunt Groups section displays the status of every extension and phone device linked to a hunt group. It shows whether or not those devices are registered, and whether or not those devices are currently in a direct call or a hunt group call.

Managing Alarms

The alarms section of the UC Devices component allows you to configure alarm thresholds, assignments, and destinations. It also provides visual cues for triggered alarms, and the ability to view the history of alarms triggered by different groups. You may also schedule maintenance windows to silence alarms during planned down times.

Existing alarm thresholds are listed under “Alarms” in the tree, directly beneath “Scheduled Maintenance”. Select an alarm threshold to see its defined limits, its active days and times, and its assigned alarm destinations:

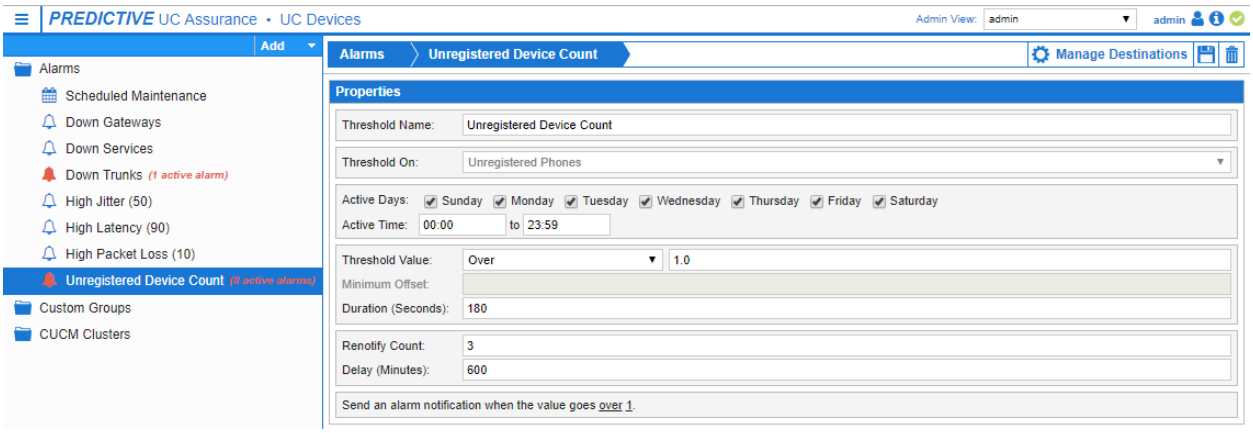

Examples of alarm threshold definitions include:

- Alarm when a specific CUCM service, trunk, or gateway goes down (when listed as “stopped” or “unregistered” in an SXML status request).
- Alarm when the average Jitter of all live calls exceeds 50ms for more than 30 seconds (for selected groups).
- Alarm when there are more than 5 concurrent calls rated “poor or worse” and active for 30 seconds (for selected groups).
- Alarm when the Packet Loss for live calls exceeds the predicted max (based on prior weeks) by 30% for more than 60 seconds (for selected groups).

Pre-Defined Thresholds

Predictive UC Analytics comes installed with some pre-defined alarm thresholds. Because different organizations can have widely different traffic volumes and network performance, most pre-defined thresholds are predictive in nature. Treat these pre-defined thresholds as initial suggestions, and modify or delete them as necessary. These thresholds fall under three basic types:

- Down Services, Gateways, and Trunks
- 30% Over Predicted Max Bandwidth, Jitter, Latency, Packet Loss, Concurrent Calls, Poor Calls, and Unregistered Phones
- 100% Over Median (aka outlier) Bandwidth, Jitter, Latency, Packet Loss, Concurrent Calls, Poor Calls, and Unregistered Phones

Note: You must change or edit the default destination for pre-defined thresholds if you wish alarm notifications to be sent to you.

Creating New Thresholds

Use the “Add” menu to add a new alarm threshold. After giving it a unique name, the most important selection is the “Threshold On” value. That value defines the type of alarm threshold you are creating, and different types of alarms may provide different options and assignment levels. After you save a new alarm threshold for the first time, the “Threshold On” value cannot be changed.

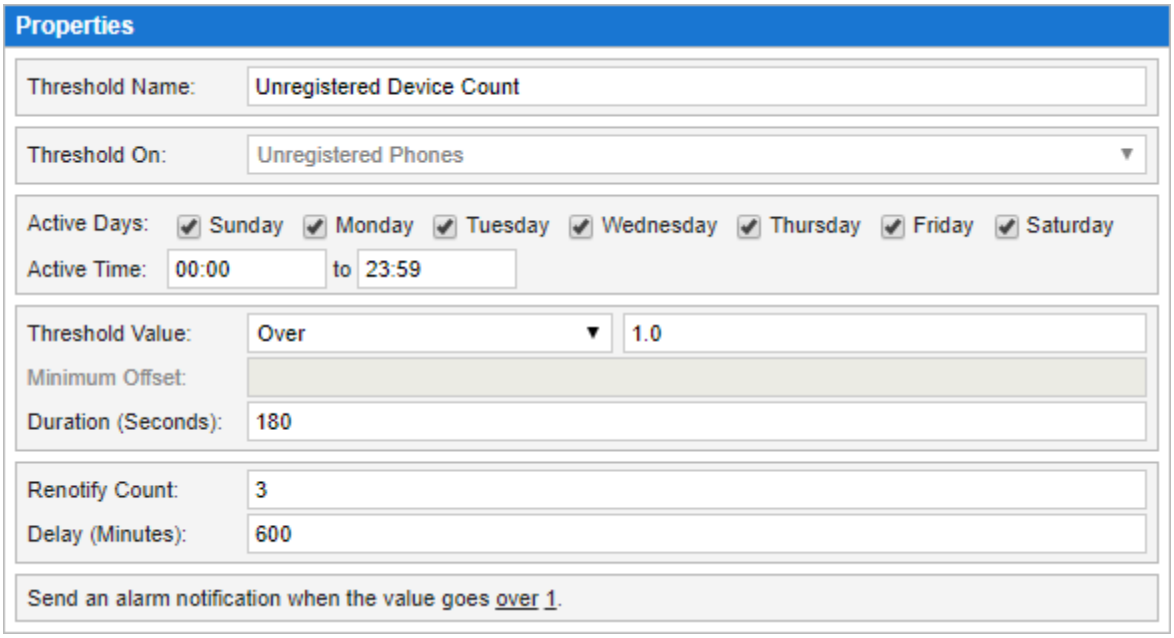

Next check the days of the week you want this threshold to be active, and specify which hours of the day it should be active. In some cases you may need different thresholds for peak and off-peak times.

For “Service Down”, “Gateway Down”, and “Trunk Down” threshold types, you will only be allowed to specify a threshold duration. If one of these goes down and stays down for that duration, the alarm will trigger. For all other threshold types, you must specify a threshold value after selecting the type of value:

- **Over** – Trigger when the threshold type is over some fixed value (e.g. Concurrent Calls Over 10, Jitter over 50)
- **% Over Predicted Max** – Trigger when the current max is some percent over the predicted max. 50% = The red line in the “Max Values” charts for each group.
- **% Over Predicted Median** - Trigger when the current median value is some percent over the predicted median. 100% = The red line in the “Median Values” charts for each group.
- **Deviations Over Predicted Mean** – Trigger when the current mean value is some number of standard deviations over the predicted mean. 2.0 = The red line in the “Mean Values” charts for each group.

Note: We strongly recommend reviewing predicted value charts for various groups you are interested in until you get more comfortable with predictive alarm thresholds.

As mentioned above, the closer a predicted value gets to zero, the more likely it will be that any percent over will be so close to the predicted value that it will trigger false alarms. To account for that, predictive threshold values allow you to specify a “minimum offset”. For example, consider a predicted max value of 2 concurrent calls at 10:00PM on Wednesday. In that case, 1 additional concurrent call would be 50% over the predicted max, which may periodically trigger an alarm you don’t want to see. However, if you specify a minimum offset of 3, that would prevent the alarm from triggering until 2+3 concurrent calls

were reached, which would make false alarms due to a few employees working later than usual much less likely.

The threshold duration helps avoid triggering alarms for live call and QoS counts that only go over defined thresholds for a very brief time period. It waits until the problem persists for a period of time, and then triggers the alarm.

The renotify count allows you to put a limit on the number of alarm notifications to receive per 24-hour period. If an alarm threshold sends you an email, and you set this value to 3, you will not see more than 3 emails per 24-hour period (per group triggering the alarm).

The renotify delay allows you to specify the amount of time to wait between sending additional alarm notifications for any group triggering an alarm threshold. If an alarm threshold sends you an email, and you set this value to 15, there will be at least a 15-minute delay in between each email to notify you that a group is still triggering this alarm threshold.

Assigning Thresholds to Groups

After you have configured an alarm threshold, you must decide which groups (at which levels) to assign it to. You can usually enable alarms at any level in the assignment tree, which in most cases will also enable it for levels below that. Optionally, you may also specify different alarm destinations to deliver each group to, or disable alarms at different levels.

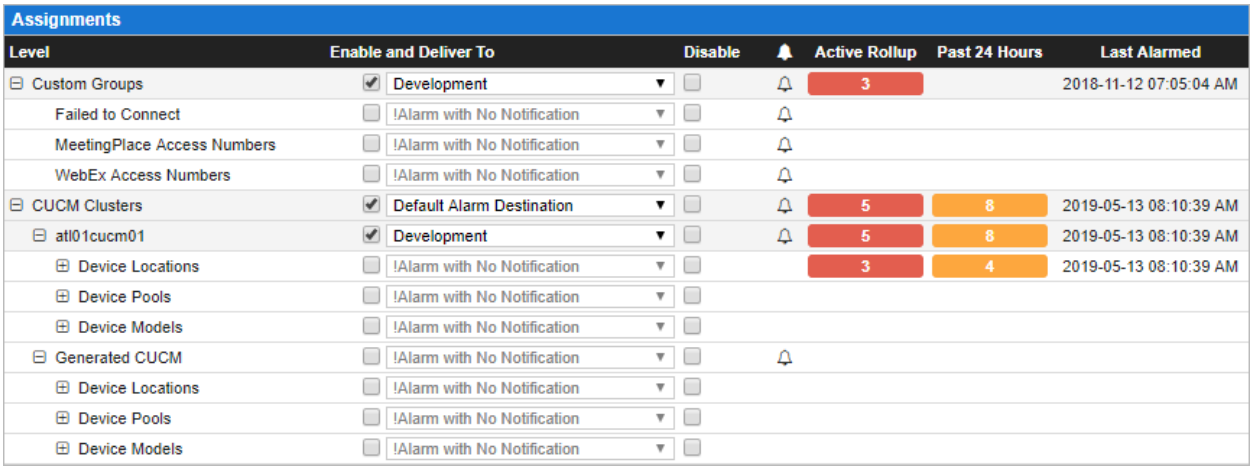

Note: You may specify different alarm destinations for different levels, in which case a single alarm can generate multiple notifications. This may be useful if you want alarms for all clusters to go to a global administrator and alarms for individual clusters to go to more local admins for each.

Checking the “Enable” checkbox for the “CUCM Clusters” group, or for a specific cluster, will apply the alarm threshold values to a cluster-wide group (e.g. all phones/calls in the cluster). If you want the threshold values applied to specific locations within a cluster, you must check the “Enable” checkbox for the “Device Locations” level, or for specific locations beneath it. The same holds true for device pool and model groups.

Checking the “Disable” checkbox at any given level disables this alarm threshold for the specified level and everything beneath it. The primary use for the “Disable” checkbox would be to apply an alarm threshold to all but a few groups. For example, instead of checking the “Enable” box for 48 of 50 location checkboxes, you can check one “Enable” box above them in the tree, and check the “Disable” box for the two you want to be exempt.

To the right of the “Disable” checkbox is an alarm state icon. The alarm state icon can show:

- - This threshold has not been assigned to this level (or any of its parent levels).
- - This threshold has been actively disabled for this level.
- - This threshold is enabled for this level, but the alarm is not triggered.
- - This threshold is currently triggered at this level.

The “Active Rollup” column shows the count of currently active triggered alarms at or below the current level. The “Past 24 Hours” column shows a similar count, but for all alarms triggered within the past 24 hours. The “Last Alarmed” column shows the most recent time any group at or below the current level triggered this alarm threshold. Clicking on any of these columns will display more details regarding when and where the alarm threshold was triggered:

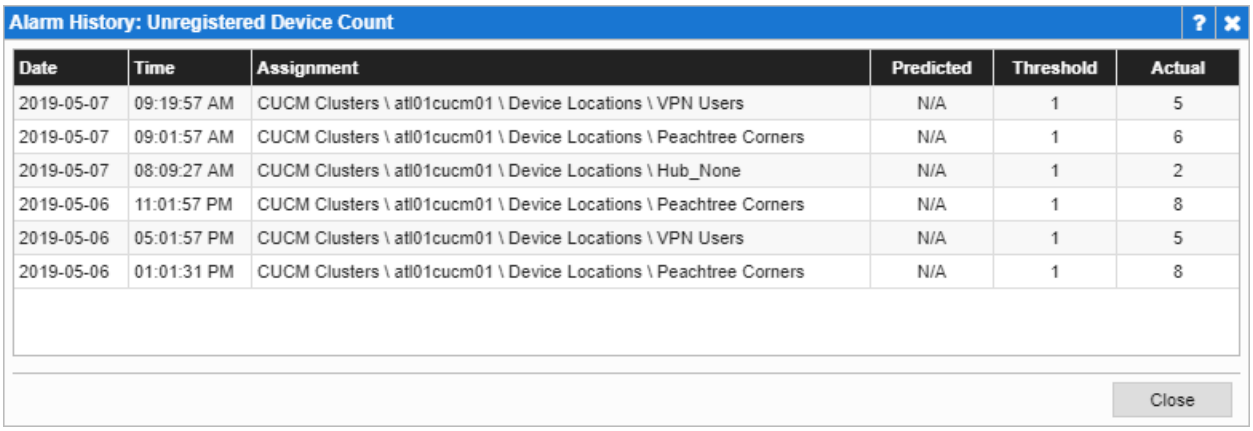

Date	Time	Assignment	Predicted	Threshold	Actual
2019-05-07	09:19:57 AM	CUCM Clusters \ atl01cucm01 \ Device Locations \ VPN Users	N/A	1	5
2019-05-07	09:01:57 AM	CUCM Clusters \ atl01cucm01 \ Device Locations \ Peachtree Corners	N/A	1	6
2019-05-07	08:09:27 AM	CUCM Clusters \ atl01cucm01 \ Device Locations \ Hub_None	N/A	1	2
2019-05-06	11:01:57 PM	CUCM Clusters \ atl01cucm01 \ Device Locations \ Peachtree Corners	N/A	1	8
2019-05-06	05:01:57 PM	CUCM Clusters \ atl01cucm01 \ Device Locations \ VPN Users	N/A	1	5
2019-05-06	01:01:31 PM	CUCM Clusters \ atl01cucm01 \ Device Locations \ Peachtree Corners	N/A	1	8

Scheduled Maintenance

Use the Scheduled Maintenance section of alarms to define windows of time to silence alarm notifications. You can define maintenance windows globally or for specific clusters, and you can configure them in advance. The engine will still track alarm thresholds during a maintenance window, but it will avoid sending notifications. Alarm notifications will resume as soon as a maintenance window ends, so if alarm thresholds are still triggered at that time, the notifications will be sent.

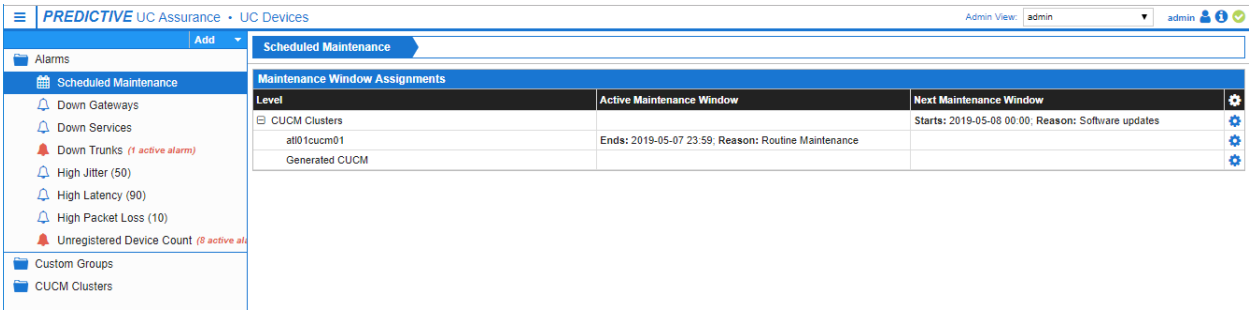

To edit global maintenance windows click on the gear icon in the right column of the "CUCM Clusters" row under "Maintenance Window Assignments". To edit maintenance windows for a specific cluster, select the gear icon to the right of that cluster. From the Maintenance Windows dialog, click the "+" icon in the top-right corner of the list to add a new window. To edit an existing window, click on the pencil icon for that row. To copy a window, click on the copy icon for that row. To delete a window, click on the "x" button on its row.

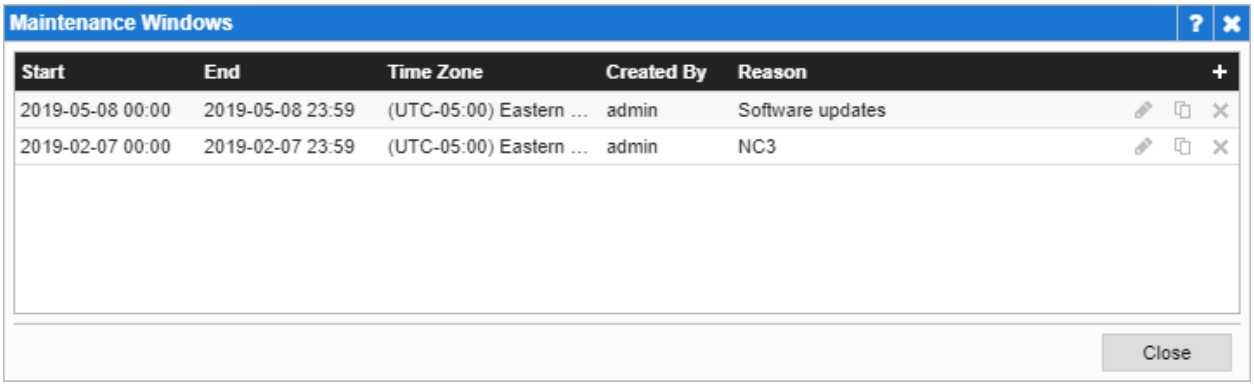

When adding or editing a maintenance window, simply specify the start date and time, end date and time, time zone, and a descriptive reason for the maintenance window. Maintenance windows can overlap, but you cannot create two maintenance windows for the exact same time period.

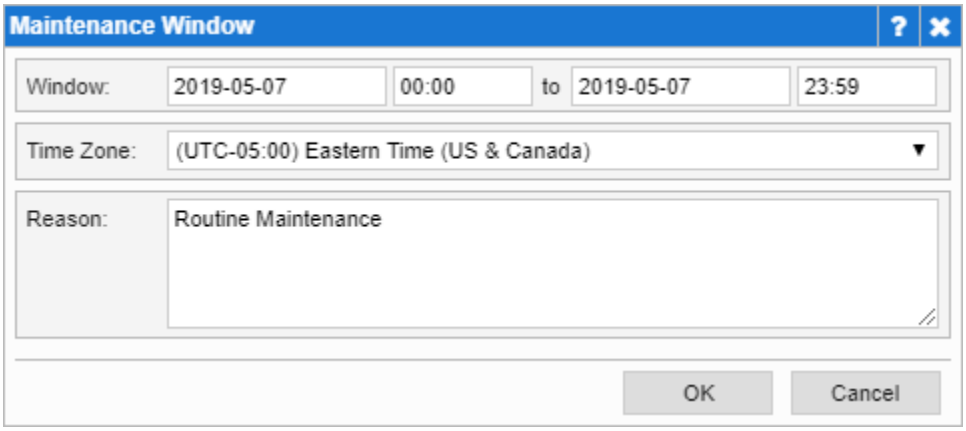

The Network Devices Component

Where the UC Devices component focuses primarily on live calls and UC device and service states, the Network Devices component focuses primarily on network interfaces and ports of your primary and secondary UC servers. The Collector Pro can be configured to use SNMP to periodically poll your UC devices to gather basic information such as uptime and the configuration and status of all network interfaces. It can also be configured to periodically test ICMP (ping) and/or TCP ports to make sure they are up.

Device Listing

The main view for Network Devices shows all polled devices grouped by location and by device type. Because location names polled by SNMP are unreliable, the class C subnet of the device's IP address will be used as a default location. The engine attempts to auto-detect a device type using the SNMP "sysObjectID" value.

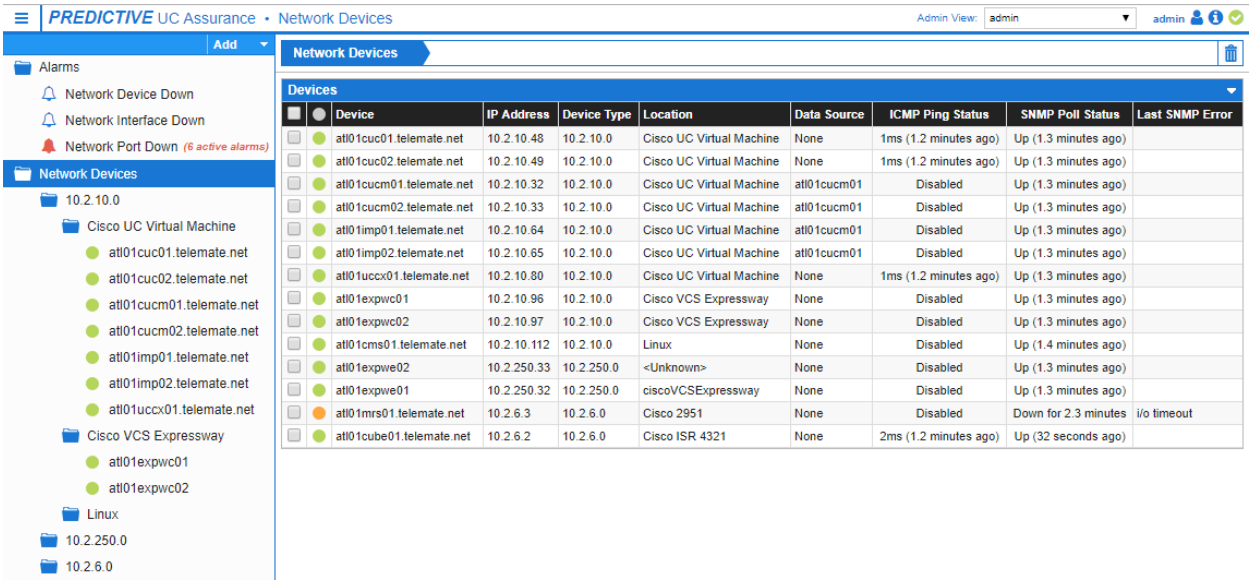

Click on the folder icon next to a location or device type name in the tree to expand or collapse that node in the tree. Click on the name itself to see only devices specific to that group.

Location

Select a specific location in the tree to limit the device list to a single location. You should notice a new "move" icon (⇅) near the top-right corner of the page.

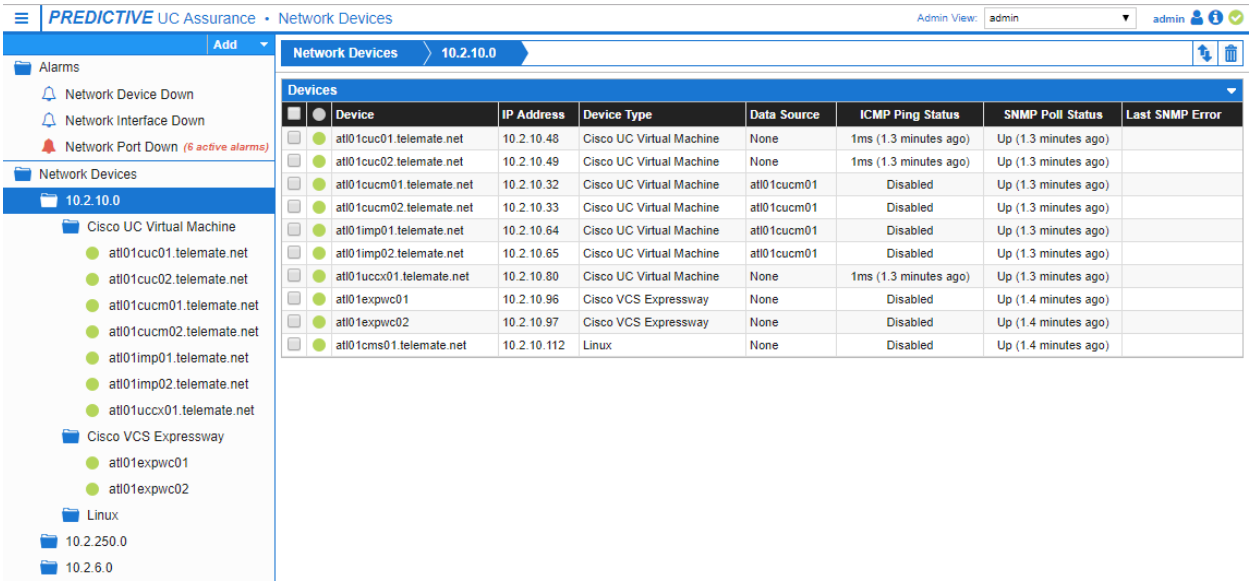

To move a few devices to a new location, select the checkboxes next to the devices you wish to move, and then click the “move” icon. To rename the location, select the checkbox at the top of the table to select all, and then click the “move” icon. Enter any location name you wish and click OK.

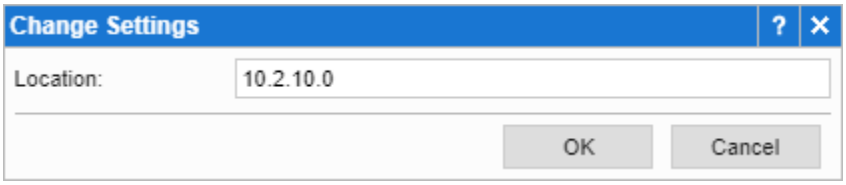

Device Type

Select a specific device type under a location in the tree to limit the device list to a single location and type. Everything will look and behave pretty much the same, but the “move” icon will also allow you to change or rename the device type.

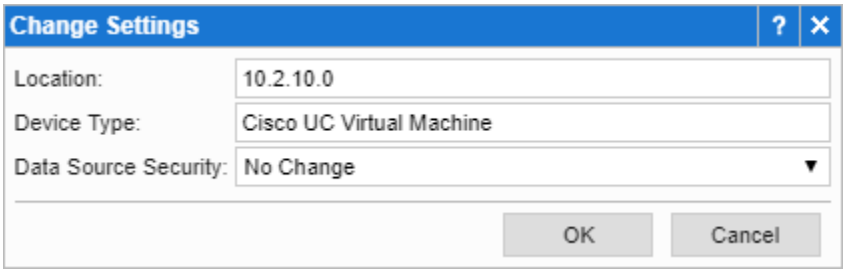

Note: You can also assign network devices to specific data sources. The only reason to do this would be to apply data source security. For example, if you have 4 clusters and an IT manager login with access to 1 cluster, assigning network devices to data sources would prevent that IT manager from seeing devices that belong to other clusters.

Device

Select a specific device in the tree to view that device’s current and recent status.

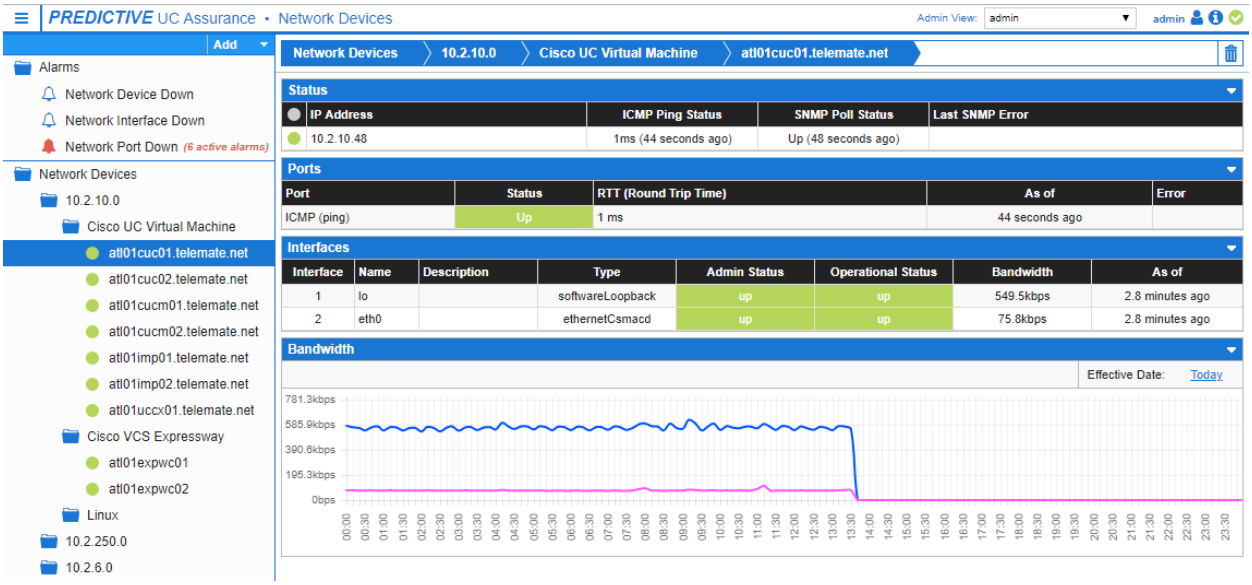

The Status section of this view simply shows whether the device itself appears to be up or down. This depends entirely on whether recent port tests or SNMP poll attempts have succeeded or failed. If any of them have succeeded, the device state is “up”. If none have succeeded, the device state is “down”.

Status

IP Address	ICMP Ping Status	SNMP Poll Status	Last SNMP Error
10.2.10.48	1ms (44 seconds ago)	Up (48 seconds ago)	

The Ports section of this view will only appear if ICMP or TCP port tests have been configured for this device. If so, the most recent status of each port configured will be displayed.

Ports

Port	Status	RTT (Round Trip Time)	As of	Error
ICMP (ping)	Up	1 ms	44 seconds ago	

The Interfaces section of this view will only appear if SNMP polling has been configured for this device. If so, the most recent status of each network interface will be displayed.

Interfaces

Interface	Name	Description	Type	Admin Status	Operational Status	Bandwidth	As of
1	lo		softwareLoopback	up	up	549.5kbps	2.8 minutes ago
2	eth0		ethernetCsmacd	up	up	75.8kbps	2.8 minutes ago

The Bandwidth section of this view will only appear if SNMP polling has been configured for this device. If so, the difference in bandwidth counts will be calculated from one poll to the next to build a chart of bandwidth usage over time for each network interface. Note the “Effective Date” selection in the top-right corner, as you may change it from today’s date to look at bandwidth from prior days.

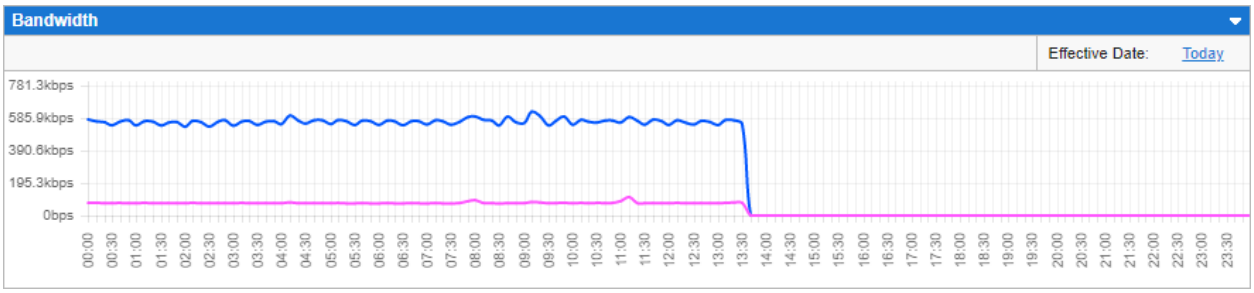

Managing Alarms

The alarms section of the Network Devices component allows you to configure alarm thresholds, assignments, and destinations. It also provides visual cues for triggered alarms, and the ability to view the history of alarms triggered by different groups. It looks and works almost exactly like the alarms section for UC Devices, so please read that description above for a more detailed explanation. The only differences are in the threshold types allowed and the assignment levels available.

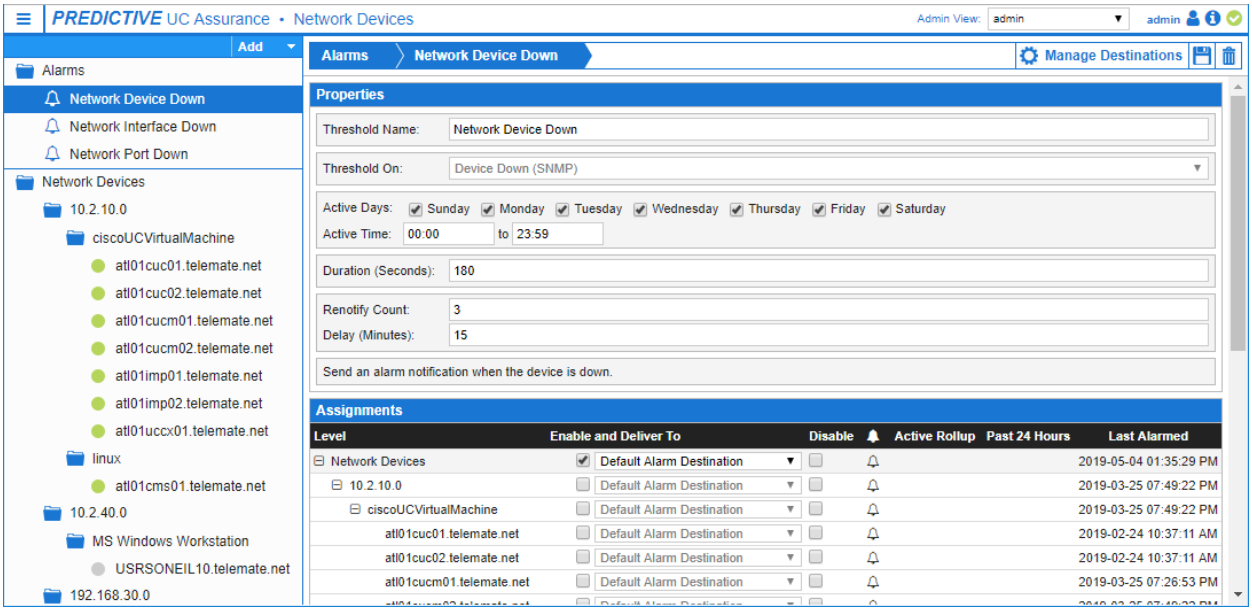

Pre-Defined Thresholds

Predictive UC Analytics comes installed with some pre-defined alarm thresholds. Treat these pre-defined thresholds as initial suggestions, and modify or delete them as necessary. For network devices, only two pre-defined thresholds are configured:

- Network Device Down – This threshold will trigger an alarm if any monitored device goes down for 180 seconds.
- Network Port Down – This threshold will trigger an alarm if any monitored network port goes down for 180 seconds.

Note: You must change or edit the default destination for pre-defined thresholds if you wish alarm notifications to be sent to you.

Creating New Thresholds

Use the “Add” menu to add a new alarm threshold. After giving it a unique name, the most important selection is the “Threshold On” value. That value defines the type of alarm threshold you are creating, and different types of alarms may provide different options and assignment levels. After you save a new alarm threshold for the first time, the “Threshold On” value cannot be changed.

Assigning Thresholds to Groups

Again, this is virtually identical to how it works for the UC Devices section. One difference is that the “Interface Down (SNMP)” threshold type will not allow assignments at parent levels. You must select each specific interface to monitor because many network devices have several virtual network interfaces that may not make sense to monitor.

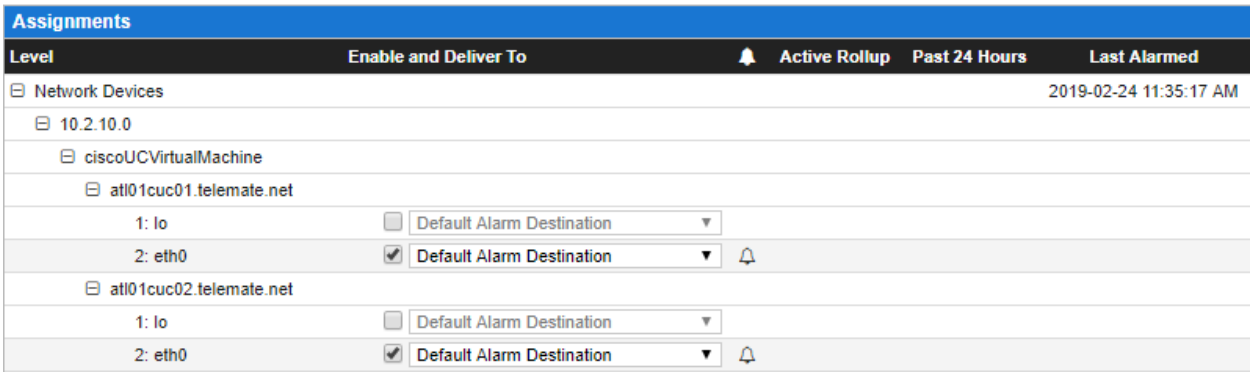

Assignment Features

The Data Sources Component

The Data Sources component allows you to view and manage gateways (or trunks on more traditional phone systems) and gateway groups for your data sources. If you are also collecting provisioning tables from a CUCM data source, you will also have the option to view device pools, locations and additional device properties. The left-hand side of this section provides controls for navigating through the tree.

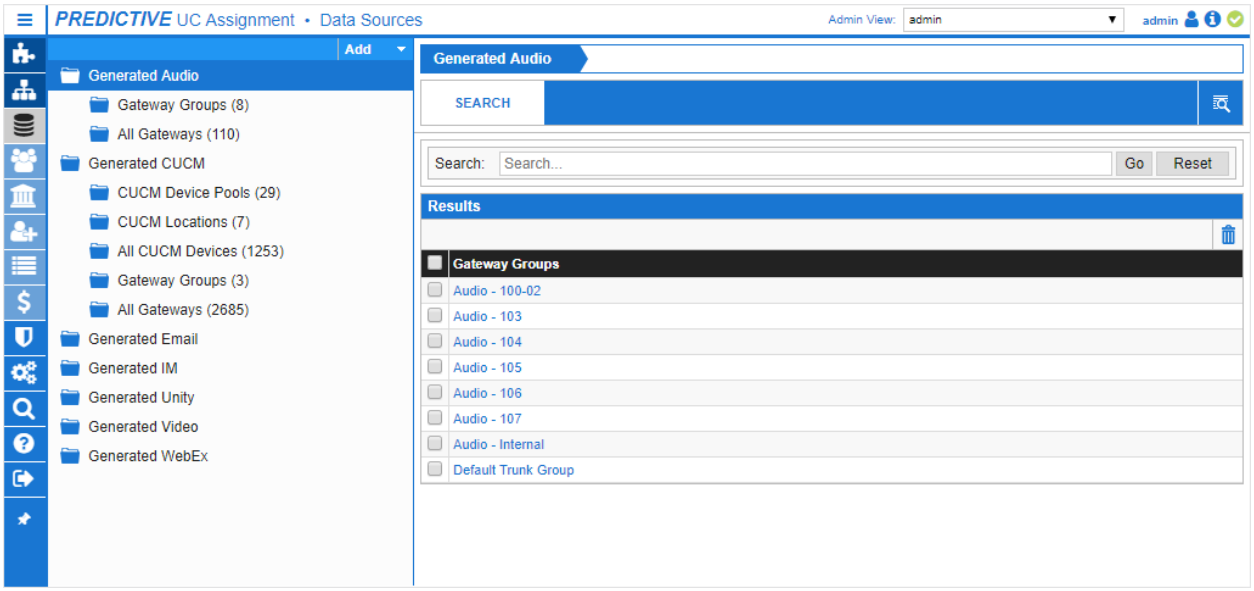

Note: In the data sources tree, the option to add a new gateway group will only appear if the currently selected item a valid "parent" to assign the gateway group to.

Once you select an entity at any level, a searchable list of its children will appear on the right-hand side of this section. Entities at the lowest level of the tree will display settings instead of searchable listing. Gateway groups have the option to display the listing and settings.

Controls

While some levels of the tree have special options, some general options are the same on every level. No matter where you are in the tree, you will see a "You are here" section at the top.

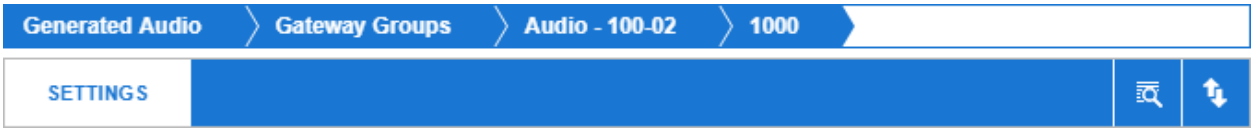

The line of text shows the hierarchy of levels and their names. The names show where the currently selected entity is in the tree. The first name is the highest level and the last name is the lowest and the currently selected entity. Each name is a link and can be used to quickly jump to that level. The image above shows the "1000" gateway in the "Audio - 100-02" gateway group of the "Generated Audio" data source. On the right of the image, you will see buttons for making changes to the current entity:

1) **View Details** – Select to view recent session detail records for the selected directory entity. Refer to the "Common Features" section of this document for a more detailed explanation of the detail view.

2) **Transfer (Gateways only)** – Select to transfer gateways from its current group to a new gateway group.

3) **Delete (Gateway Groups only) –** Select to delete the currently selected gateway group. You must transfer any gateways currently assigned to the group to another group before you can delete the gateway group.

Gateway Group Settings

When you select a gateway group, it will have a "Settings" section for general properties. The properties in this section are name, comment, costing method, and site location. Site locations can be used to determine the physical location of activity. Setting a site location at this level will override the one set for the data source.

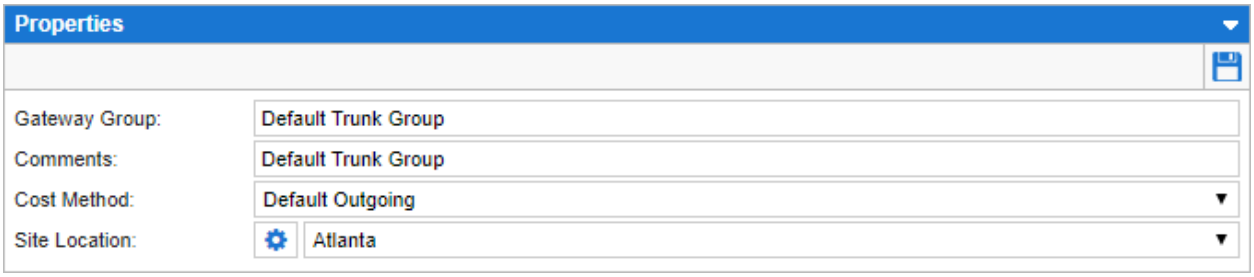

CUCM Device Settings

When you select a CUCM device, it will have a "Settings" section for general properties. The properties in this section are view only.

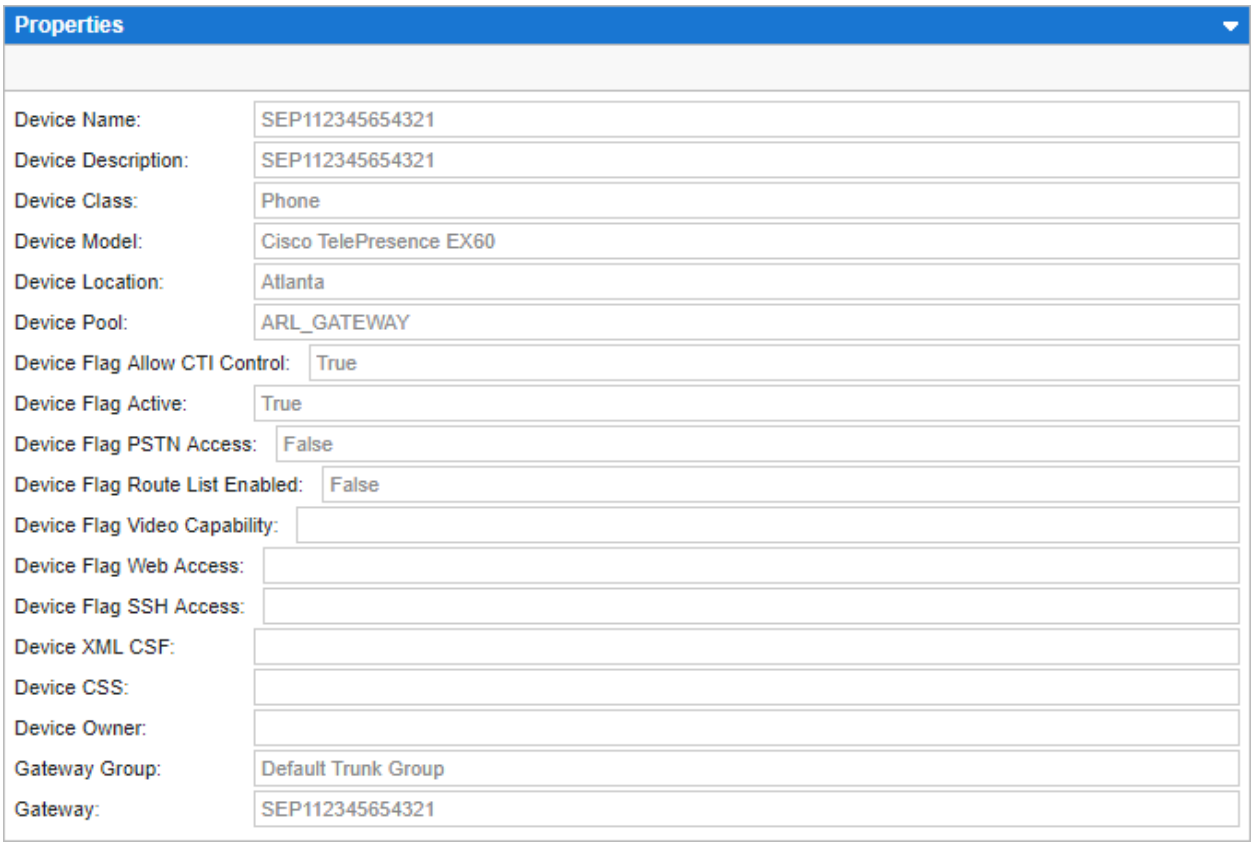

The Directory Component

The Directory component allows you to view and manage companies, divisions, departments, users, and resources/addresses in your organization. The left-hand side of this section provides controls for navigating through your directory tree. Select the "Organization" tab when it is easier to navigate by clicking through a tree of companies, divisions, and departments. When it is easier to type in a few characters of what you are looking for, select the "Find" tab. We recommend that you get familiar with both navigation methods, as the "easier" method changes based on what you need to accomplish.

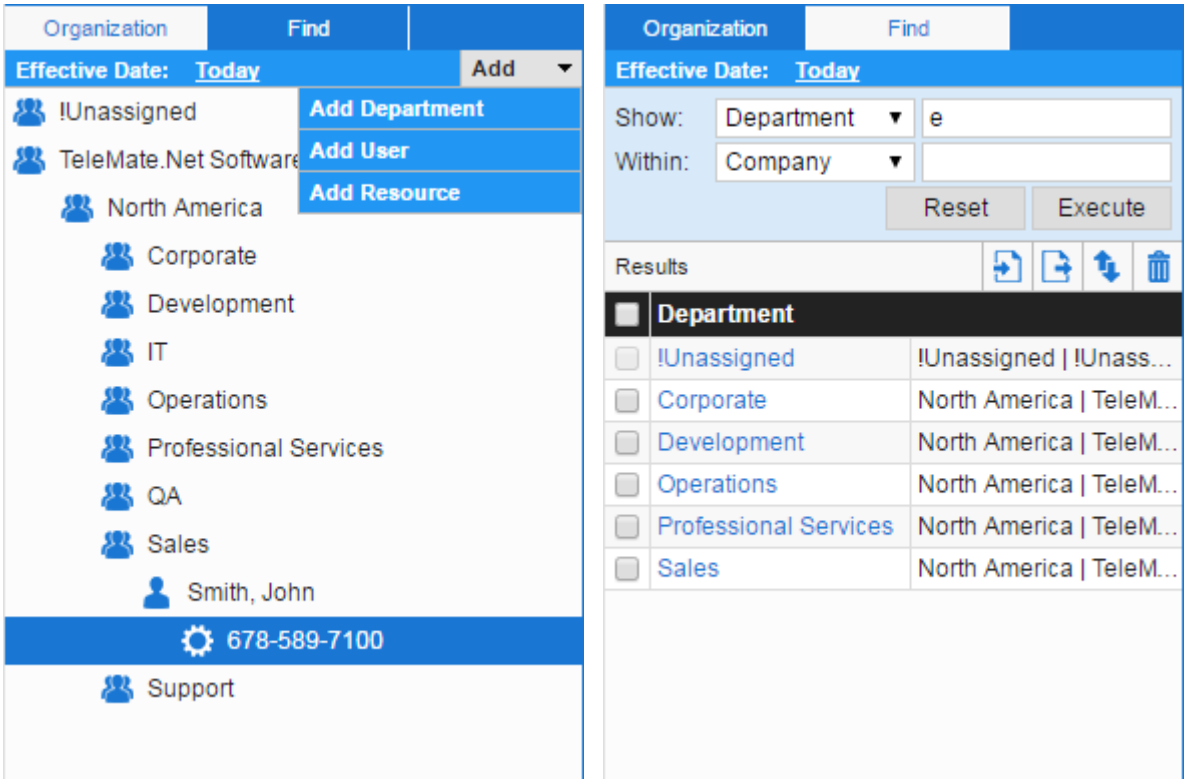

Note: In the Organization tree, the option to add a new entity at a specific level will only appear if the currently selected item provides a valid "parent" to assign it to. For example, if you select a company, you cannot add a department because it is not clear which division it should be assigned to.

Regardless of how you choose to navigate, once you select a directory entity at any level, its properties and children will appear on the right-hand side of this section. Note that entities at different levels of the organization will have different options and properties. For example, the user level allows you to specify a job title and the resource (extension) level allows you to specify call plans.

Common to All Levels

While some levels of the tree have special options, some general options are the same on every level. No matter where you are in the tree, you will see a "You are here" section at the top.

The line of text shows the hierarchy of levels and their names. The names show where the currently selected directory entity is in the directory tree. The first name is the highest level and the last name is the lowest and the currently selected directory entity. Each name is a link you can use to jump to that directory level. The image above shows the "Sales" department in the "Atlanta" division of the "TeleMate" company. On the right of the image, you will see a number of buttons for making changes to the current entity:

6) **View Details** – Select to view recent session detail records for the selected directory entity. Refer to the "Common Features" section of this document for a more detailed explanation of the detail view.

7) **History** – Select to view/modify the list of historical changes (e.g. transfers) for this entity. For example, if "John Smith" was transferred from the "West Coast Sales" department to the "East Coast Sales" department on June 1st, the history listing for "John Smith" will show that transfer. If the transfer was a mistake and you need to change or delete it, you should do that here. **Note:** If historical tracking is disabled, each entity will only have one history record to represent where it belongs (for all dates). If you change it, you change it for all dates.

8) **Transfer –** Select to transfer any directory entity from its current parent to a new parent. When historical tracking is disabled, this simply replaces its only history record. When tracking is enabled, it creates a new history record, which requires you to specify an effective date.

9) **Create Invoice Batch** – Refer to the "Invoice Batches" section below for an explanation.

10) **Delete –** Select to delete the currently selected entity. This is a recursive deletion that can have unintended consequences when historical tracking is enabled, so be very careful when deleting anything in the directory. The only time deletion is completely safe is when you recently created a directory entity by mistake and no calls or expenses are assigned to it. Otherwise, we recommend you transfer entities to "!Unassigned" instead of completely deleting them.

Reminder: Background Tasks

Operations like transfer and delete may take a while because they may affect many records in the database, and they may require integrity checks or locks to clear before they complete, so they are run as background tasks to keep the web browser from getting stuck for long periods of time and/or timing out. Check the status of background tasks by clicking on the "info" icon in the top-right hand corner of the page. These are examples of in progress, successful and failed deletion tasks:

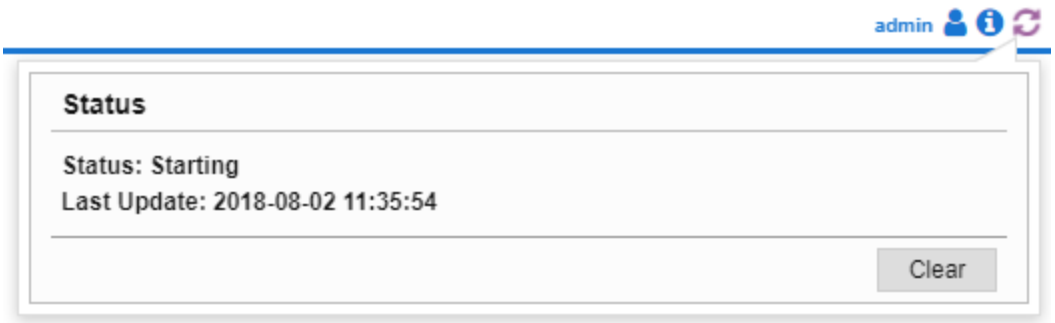

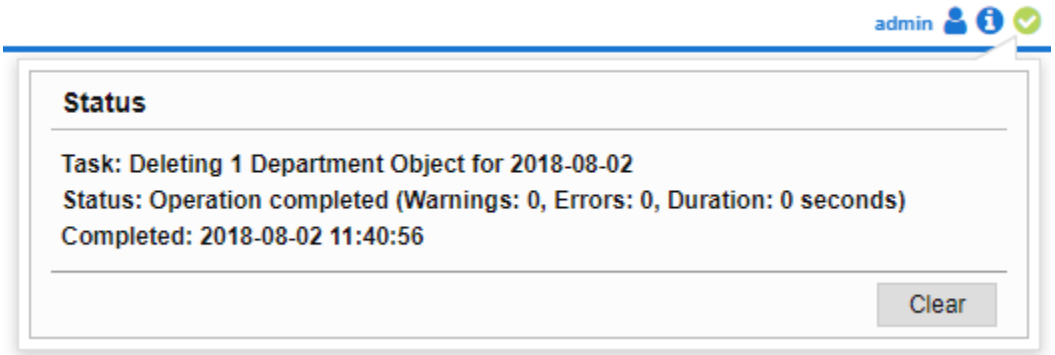

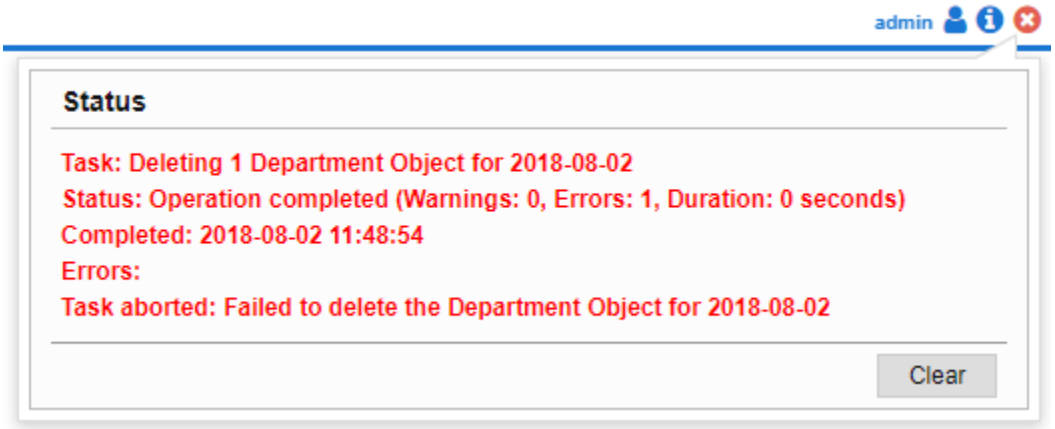

Properties

When you select a directory entity at any level, it will have a “Properties” section for general properties. The most common properties in this section are name, description, and email address (for email breakout report distribution), but some levels will have different properties:

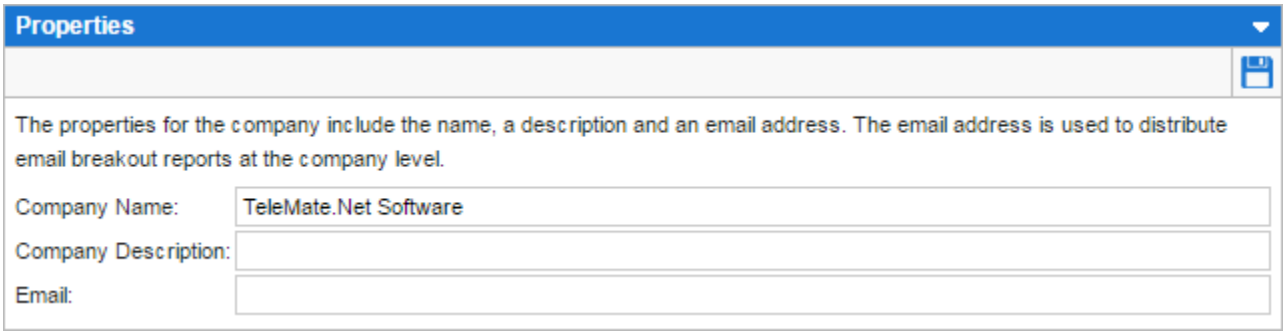

Children

When you select a directory entity at any level, except the very bottom of the tree, it will have a "Children" section. This section allows you to manage all children, grand-children, etc. beneath the selected entity:

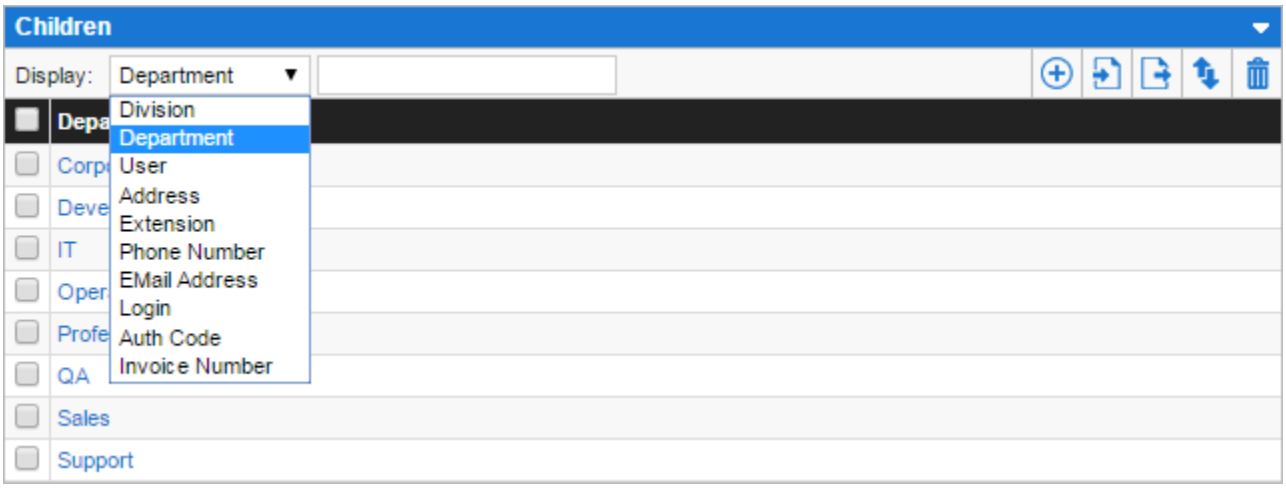

The "Display" option defaults to show the children directly beneath the current entity, but you can change it to show any level of the tree below the current one. The text box next to the display selection allows you to search the selected level. So if you select "Phone Number" and type "351", it will show you all internal phone number resources beneath the current entity that contain "351".

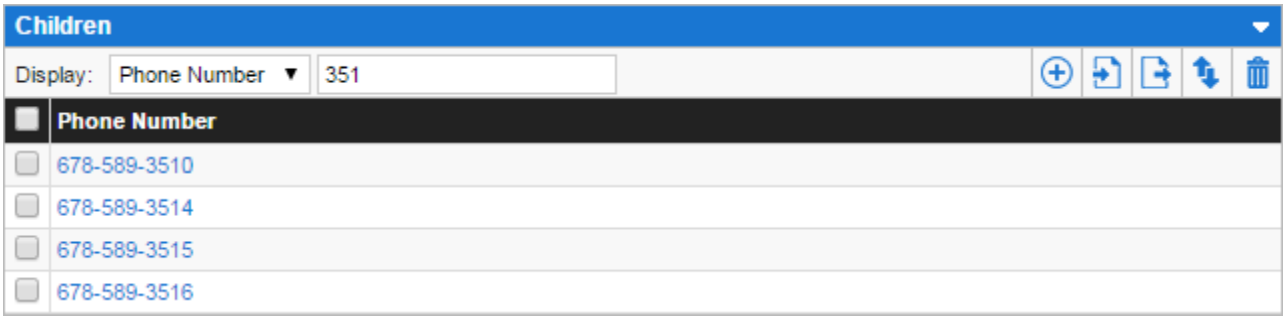

The buttons available on the right allow you to:

1) **Add –** Create a new child entity (direct children only, not grand-children).

2) **Import –** Use to re-import a file you have modified after exporting it. The file format must be the same as the format generated by the Export button (see below).
   Note: The import is actually global and not specific to any directory entity. The button only appears here because it goes with the Export button, which is specific to this entity.

3) **Export –** Quickly export the properties of all children beneath the currently selected entity to a csv file. You can modify those properties in Excel or some other csv editor.

4) **Transfer –** Executes a batch transfer for every child selected in the list.

5) **Delete –** Execute a batch deletion for every child selected in the list.

Note: There are a few special cases when some of these buttons are not applicable, in which case they may be hidden, grayed out, or give an error message. For example, you may not transfer invoices, and you may not transfer any level of the built-in "!Unassigned" bin.

Company Level

The company level has no additional properties. Since it is the highest level in the organizational tree, companies have no parent, which means they cannot be transferred and do not have history.

Division Level

The division level has no additional properties.

Department Level

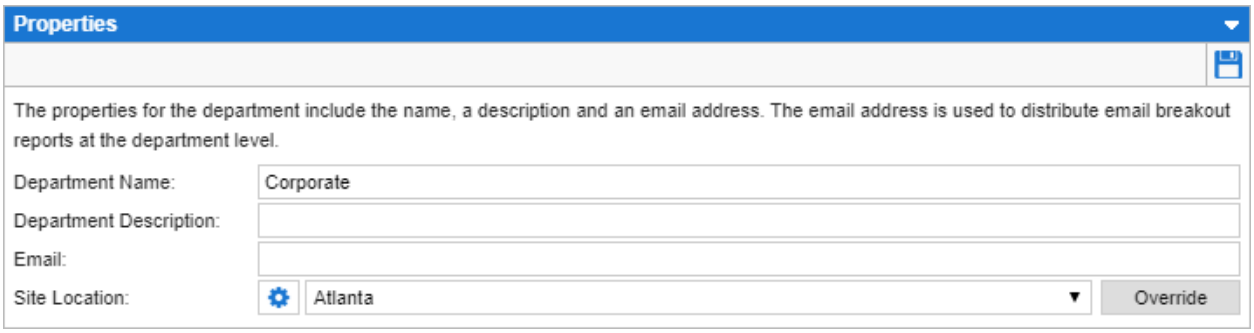

In addition to its common properties, a department can have a site location. Site locations can be used to determine the physical location of an Extension. When you assign a department to a site location, all users and extensions in that department will be assigned to that site location unless you explicitly override it for individual users or extensions. The gear icon is a link that allows you add, edit or delete

site locations. The Override button clears any site locations set at the user or extension level, so that they will use the location set at the department level.

User Level

The user level has more properties than any other level:

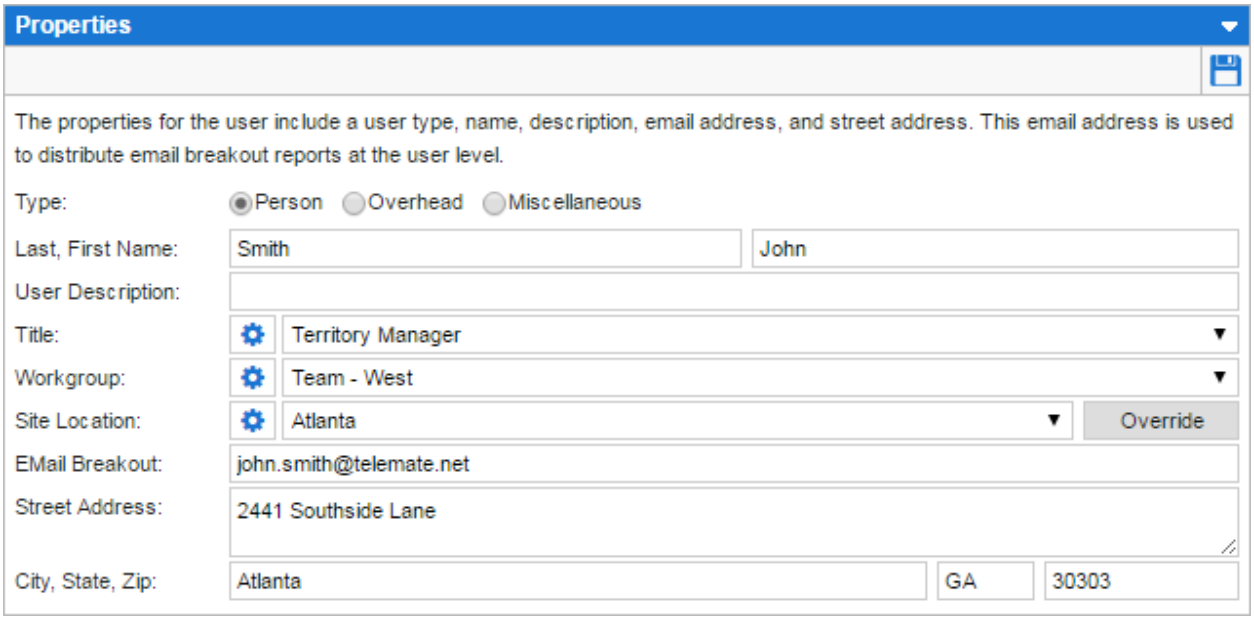

A user can be categorized as one of three types: person, overhead, or miscellaneous. The type selection is purely informational, and most of our customers only use it as a reminder when they create an overhead user under a department with the same name to assign shared resources (e.g. fax machines) and expenses (e.g. the cost of the phone line for that fax machine). The street address properties are also informational, but they show up on all of our bill reports.

You can also select a title and/or a workgroup for each user. While these selections are primarily informational, they are also available as filters on almost every report that includes directory information, so you can use these selections as an alternate way to group users together outside of the department they belong to. If you need to create or edit titles or workgroups, click on the gear icon link located next to the title or workgroup.

You can also select a site location. Similar to site locations in departments, they can be used to determine the physical location of an Extension. When a site location is set at the user level all extensions assigned to that user will have that site location unless it was explicitly set at the extension level to something else. If you need to create or edit a site location, click on the gear icon link located next to the site location. The Override button clears any site locations set at the extension level, so that they will use the location set at the user level.

User-Defined Data

Our customers often wish to import or enter additional information for each user in the directory, but different customers have different needs. As a result, we reserve 10 fields that you can define to use for anything you wish. To use one of these fields, you must first enable it from the server-side client in the UC Columns section.

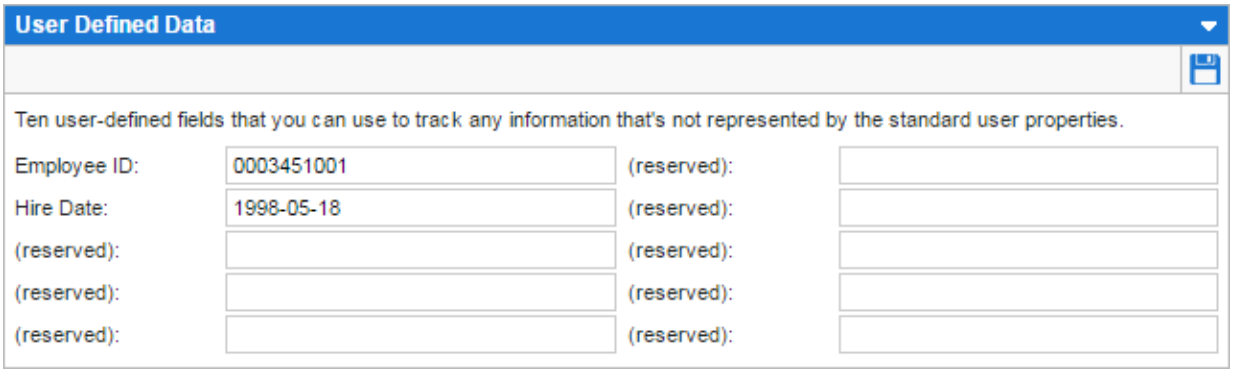

Devices

You may also add custom "devices" (or phone numbers) for a user, such as their home phone, cell phone, etc. This is purely informational, and is not used by any of our standard reports (though custom reports can be written to use this information). The devices feature is disabled in the web interface by default. To enable it, click on the "Settings" link in the navigation menu and look for the "Features" option in the "General Settings" section. In "Features", enable "Additional User Devices" and click the OK button.

Expenses

If you have an Enterprise license for Predictive UC Analytics, you can also manage expense code assignments for the selected user here. Refer to the documentation on the Expenses section for more information.

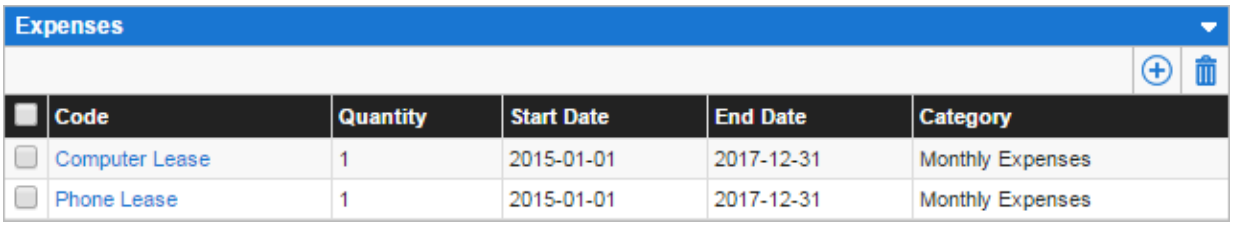

Address Level

The term "address" is an aggregate that includes any unique identifier that can be used to assign traffic sessions to a user in the Predictive UC Analytics directory. These can be extensions, internal phone numbers, authorization codes (sometimes called DISA, pin, authentication, or access codes), calling card numbers, email addresses, VoIP or network logins/addresses, etc. If we use it when processing UC logs to identify "ownership" of a session from an accounting viewpoint, we call it an "address". You may also see addresses referred to as "resources" in some places.

Options

Aside from authorization codes (which are special because they override call ownership for other resource types and are not tied to a physical location), all resource types can be assigned a Site Location and a Publish Type. These work exactly the same way that the Title and Workgroup selections work at the user level, and they can be used for just about anything you need to use them for. For example, you may wish to create special publish types for hunt pilots, customer service lines, fax machines, floor extensions that any employee can pick up, or dedicated extensions.

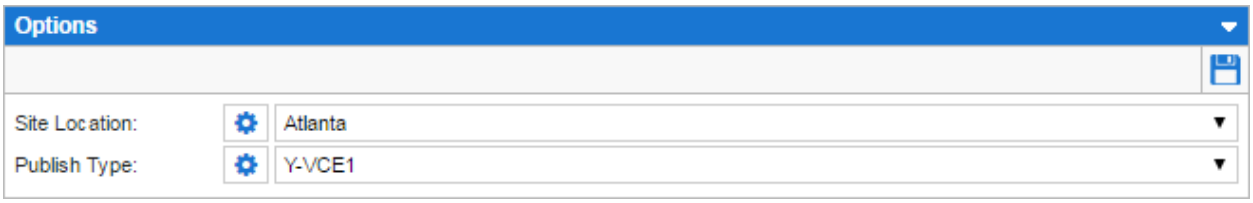

Call Plans

In addition, addresses can be assigned to call plans to help you track when a particular address is going over (or under) certain daily, weekly, or monthly targets. Call plans can be configured by clicking on the gear icon in the top-right corner Call Plan Assignments section. By running the reports in the "Call Planning" category, you can watch for over-use (e.g. cell phones costing you extra by exceeding the number of minutes included in their plan) or under-use (e.g. call center agents not answering enough calls).

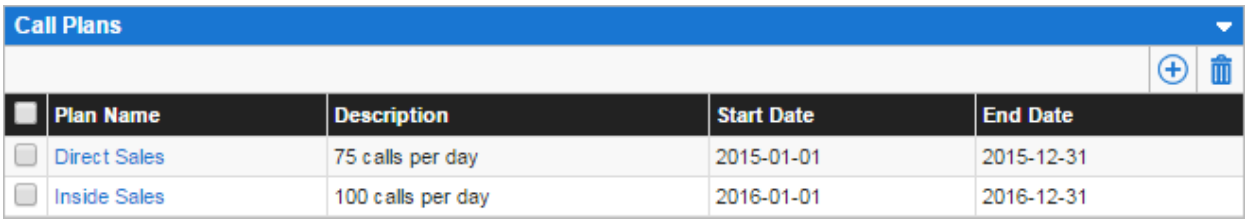

The Expenses Component

The Expenses component can only be accessed if you have an Enterprise license for Predictive UC Analytics and it has been enabled in the "Features" section of "Settings". It allows you to allocate and distribute specific costs to the users responsible for them. For example, you may wish to allocate a one-time charge of $50.00 to a specific user to cover the installation of a phone line on a specific date, and then a recurring charge of $20/month while that phone is in use. You can also create a very large expense and distribute the cost across several users/departments (and across several months/years). You can then run monthly distributions and reports to calculate budget expenses for departments or divisions, bills for clients, etc.

The Expenses feature also ties into our invoicing features by allowing you to create expense codes that assign different percentages of call costs for different call types. This allows you to account for call costs as well as to calculate fees based on a percentage of call costs (e.g. local/federal USF charges). You may also specify which taxes to apply to each expense code.

Information is presented as a tree of expense categories, expense codes, assigned users, distribution periods and applied expense distributions.

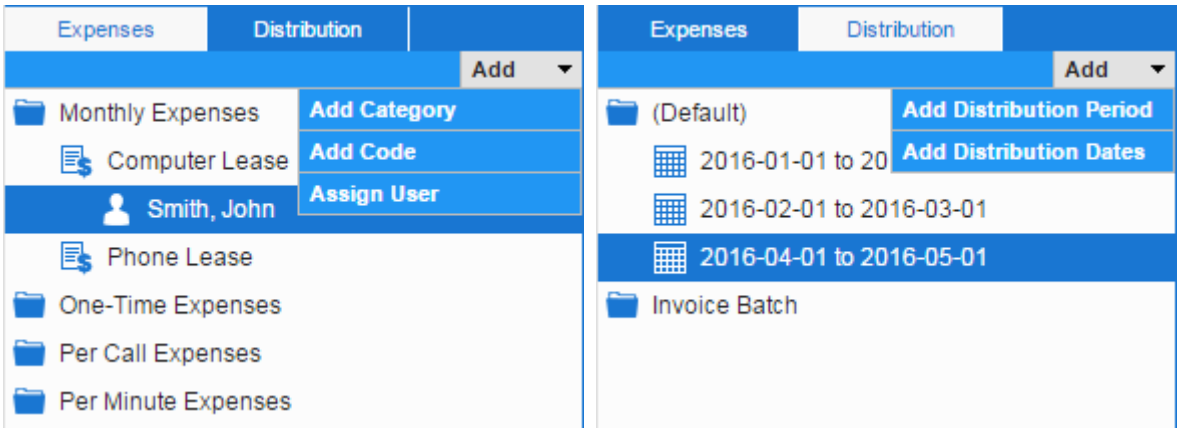

Categories

The category level of the tree allows you to manage a specific expense category. You can modify the selected category's properties as well as add/edit/delete expense codes assigned to this category.

Click the trashcan icon in the title bar of the expense category to delete the selected category along with all codes, assigned users and distributions associated with the category.

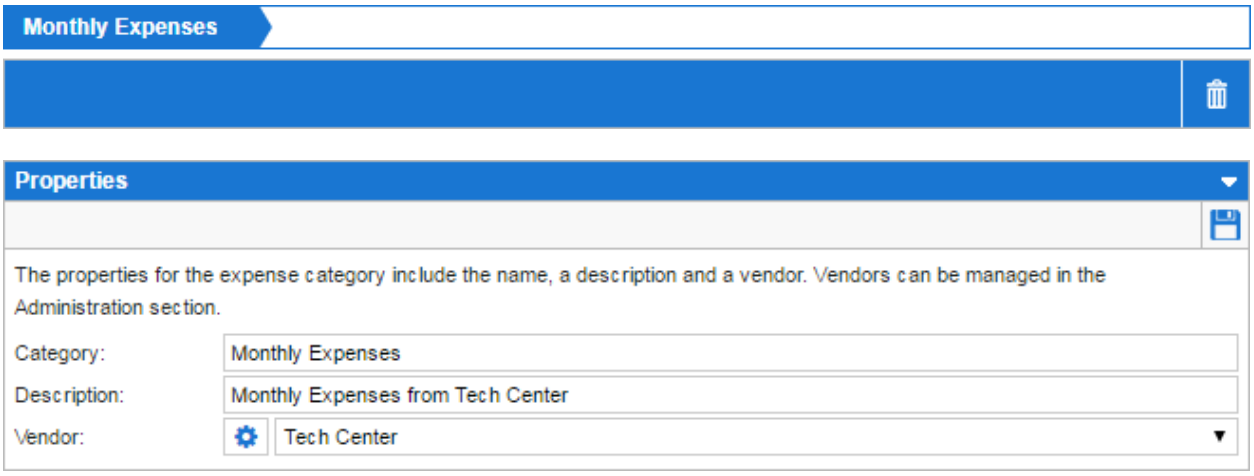

Unless you execute a search, the Children section will show all of the codes for the selected category. To assign new codes, select the "Add" button in the upper left corner.

Codes

The code level of the tree allows you to manage a specific expense code. You can modify the selected code's properties as well as manage users assigned to this expense code.

Click the trashcan icon in the title bar of the expense code to delete the code and all assigned users and distributions associated with the code.

The Expense Code properties are straightforward. You can also assign a general ledger number. If you need to create or edit a general ledger number, click on the gear icon link located next to the general ledger number.

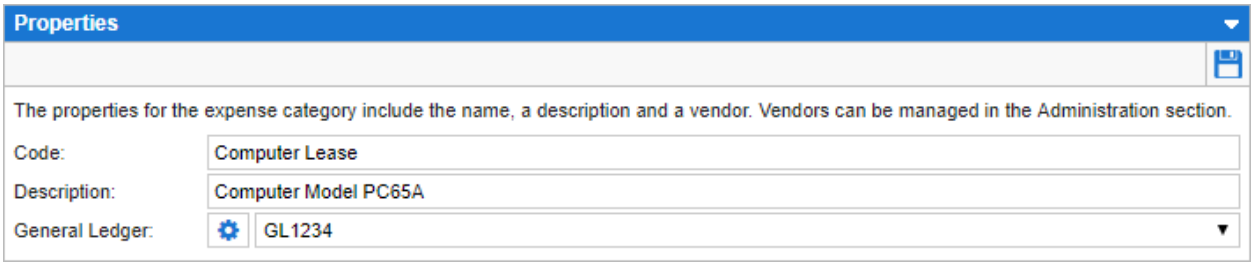

There are several distribution types that affect how the expensed cost is distributed to assigned users:

- **Period, Once** - Distributes a one-time cost to any user you choose on any date you choose. A good example of this would be a fixed charge for installing a phone line. If John Smith has 3 phone lines installed on June 5, you can assign the “Phone Line Installation Fee” expense code to John for June 5 and set the assigned quantity to 3. When you distribute costs for June and run bill/expense reports, the appropriate cost will be applied.
- **Period, Recurring** - Distributes a recurring expense to all assigned users from the specified start date to the specified end date (pro-rating as necessary). Any monthly service fee would be a good example of this. In the John Smith example above, you could assign 3 recurring service line fees of $20 each starting June 5. June would be pro-rated for 26 out of 30 days for a total of $52. Subsequent months would be $60 until you choose to terminate the assignment of those lines (which will be pro-rated if you terminate mid-month).
- **Assigned Users, Call Records, and Call Duration** – All of these distribution types take a single amount across multiple users/departments, and often across multiple months. For example, let’s say you spend $200,000 to upgrade your phone system and you want to expense the cost across all departments over the course of 5 years. Create the expense code and specify the expense amount and date range. Select “Assigned Users” to spread the amount evenly across all assigned users, “Call Records” to weight the distribution based on how many calls each user makes, or select “Call Duration” to weight the distribution based on how long each user spends on the phone.

- **Invoice Calls** –This distribution type allows you to add up the costs of individual calls for each user and distribute a percentage of that call cost. Further, it allows you to include/exclude calls by call type. This can be used to apply markups, but it is also useful for calculating things like local/federal USF fees, which may not apply to all call types. It is also useful if you need to apply different taxes to different call types (e.g. some states charge sales tax on intrastate calls while other states do not).

The date range selection is really only used for the distribution types that spread a single cost across multiple users and periods. While it can be used to terminate a recurring expense code, it is better to terminate individual user assignments to that expense code.

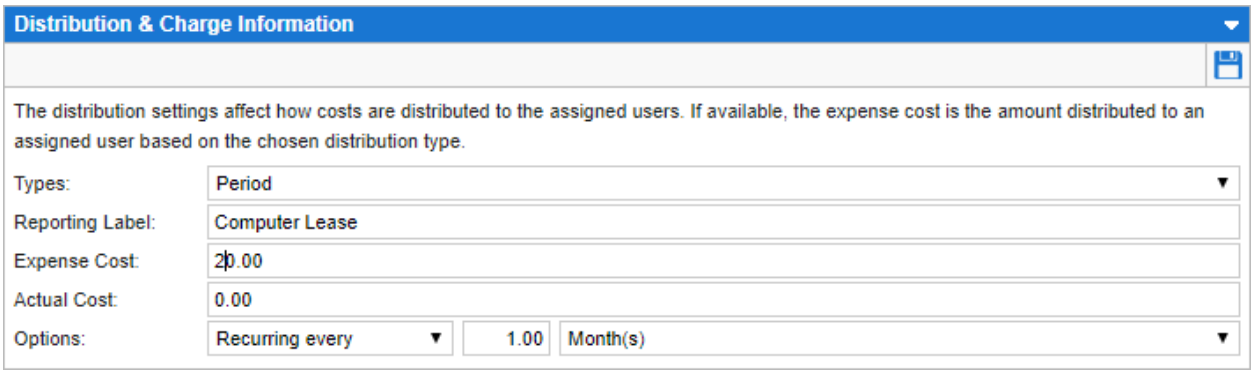

Taxes allow you to apply the selected taxes to the distributed cost of the expense in certain reports.

Note: The tax assignments are currently only used in our invoice reports and custom reports. The standard Expense reports do not currently use these tax assignments.

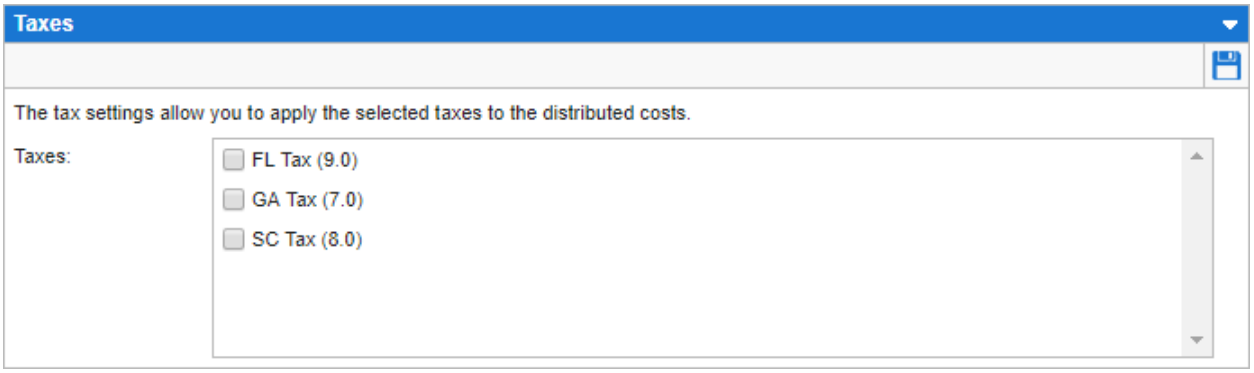

Unless you execute a search, the listing below will show all of the users assigned to the selected code. To assign new users, select the "Assign" or “Add” button in the upper left corner. To manage an assigned user, select the user name. To unassign a user, select its checkbox and click the "Delete" button.

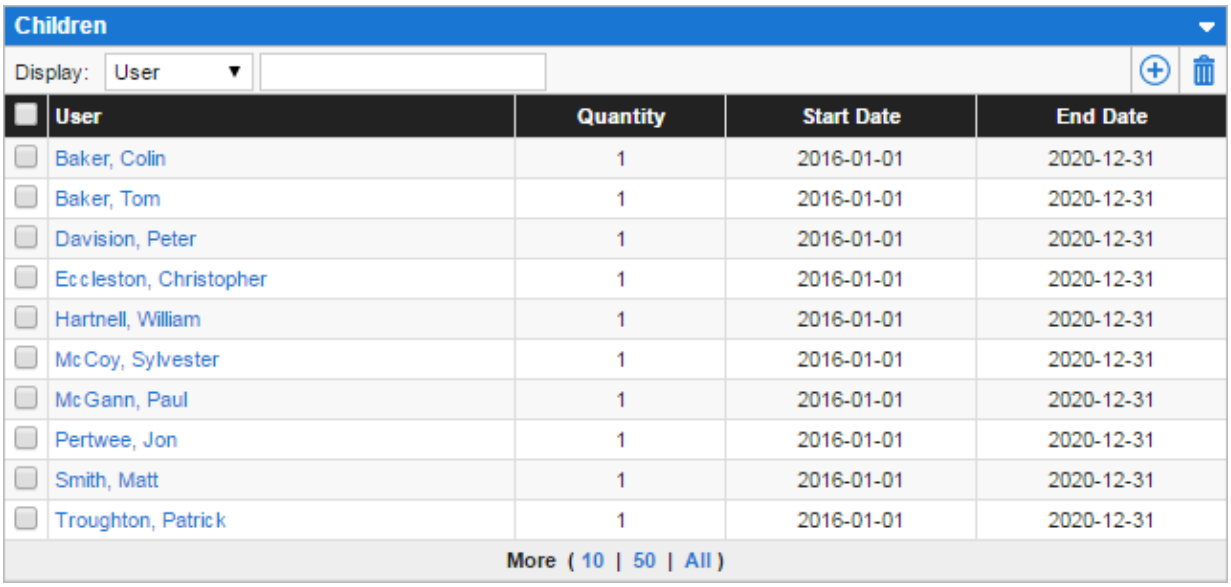

Assigned User

The user assignment level allows you to manage the assignment of a user to an expense code. You can modify the selected assignment's properties as well as manage distribution or billing runs applied to this assignment.

Click the trashcan icon in the title bar of the assigned user to delete the assignment and all distributions associated with the assignment.

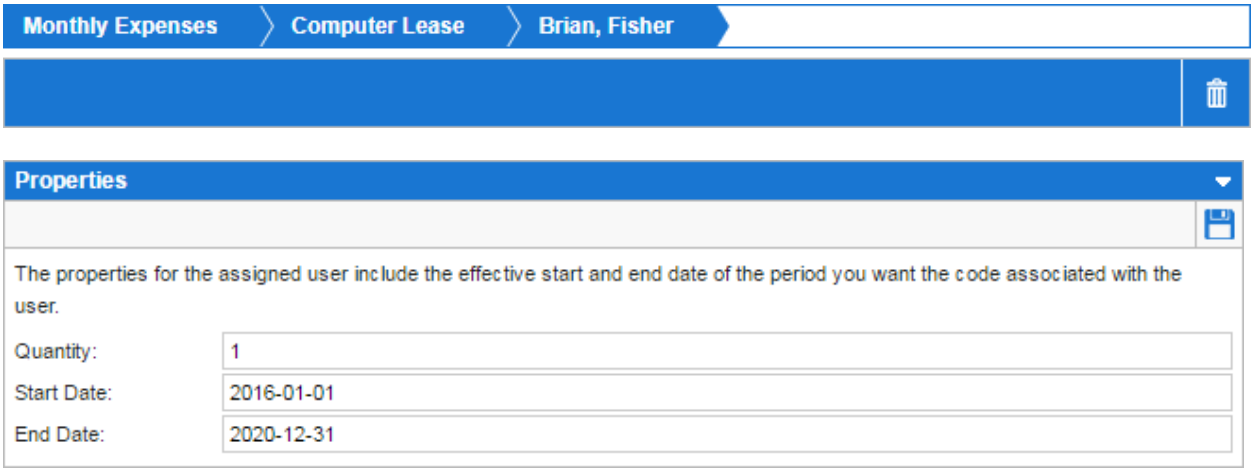

Unless you execute a search, the listing below shows the distribution for the assigned user.

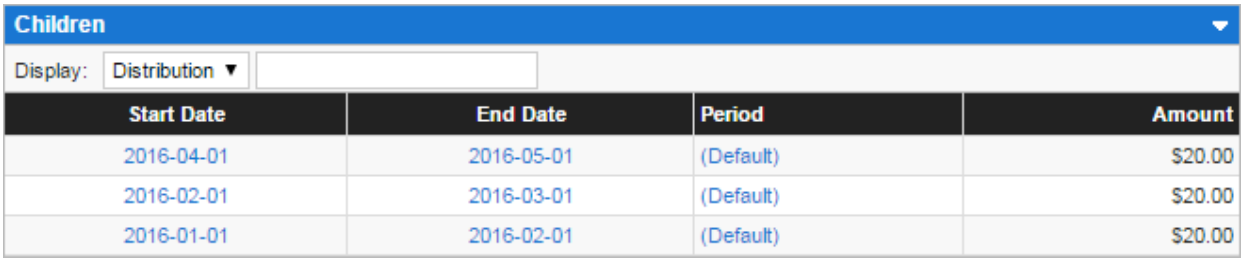

Distribution Period

This page allows you to manage all of your expense periods and distributions. Most of our customers only have one expense period (e.g. "(Default)"), which gets created automatically when they run either expense distribution reports or invoice batch reports. The primary reason to have more than one period would be if you had different users that needed to be on different billing cycles (or perhaps just have bills generated on different dates). For example, your US division may need to generate bills on the 1st of every month but your UK division may need to wait until the 8th.

Note: We recommend you use distribution/invoice reports to manage your periods and distributions. Most of our customers only check this information to trouble-shoot something that looks wrong in a report or to delete older distributions that they wish to clean out of the system.

Important: Unlike date ranges specified when running a report, distribution date ranges appear to overlap. They do not actually overlap, as the end date is not included in the distribution. We are just following a common convention used by accounting software when displaying distribution ranges.

Click the trashcan icon in the title bar of the distribution period to delete the period and all distributions associated with the period.

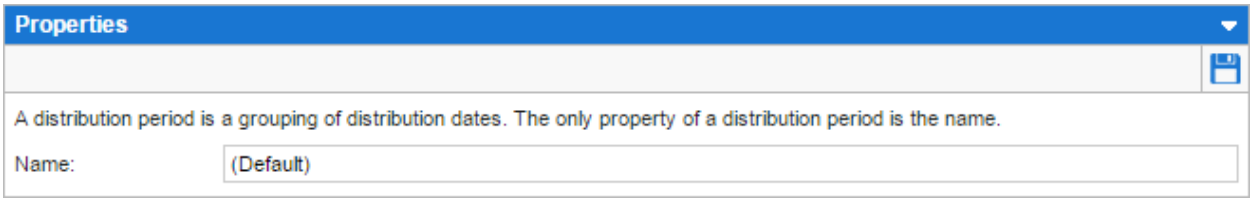

Unless you execute a search, the listing below shows the distribution dates for the selected billing period. Here you may add/edit/delete the distributions.

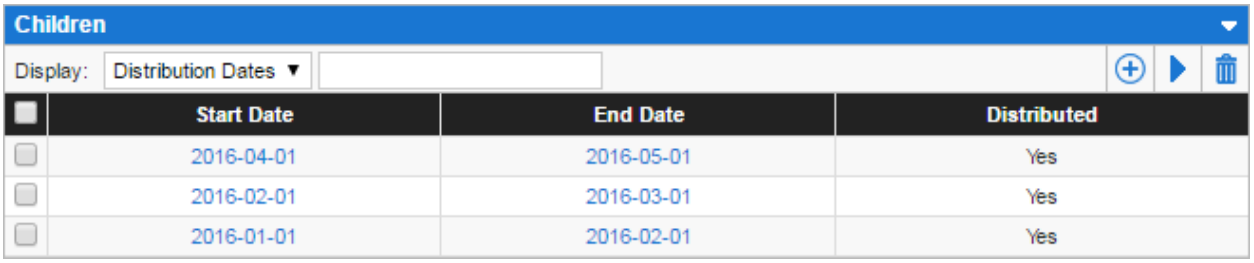

Distributions

This page allows you to manage the selected expense distribution. Any changes to a distribution's date range may cause inconsistencies in the data for that distribution. If you must change a distribution's date range, you should re-run the distribution to fix the distributed amounts by clicking the button with the green arrow.

Note: We recommend that you use distribution/invoice reports to manage your periods and distributions. Most of our customers only check this information to trouble-shoot something that looks wrong in a report or to delete older distributions that they wish to clean out of the system.

To run generate a distribution click the green arrow in the title bar. If you want to delete the distribution click the trashcan icon.

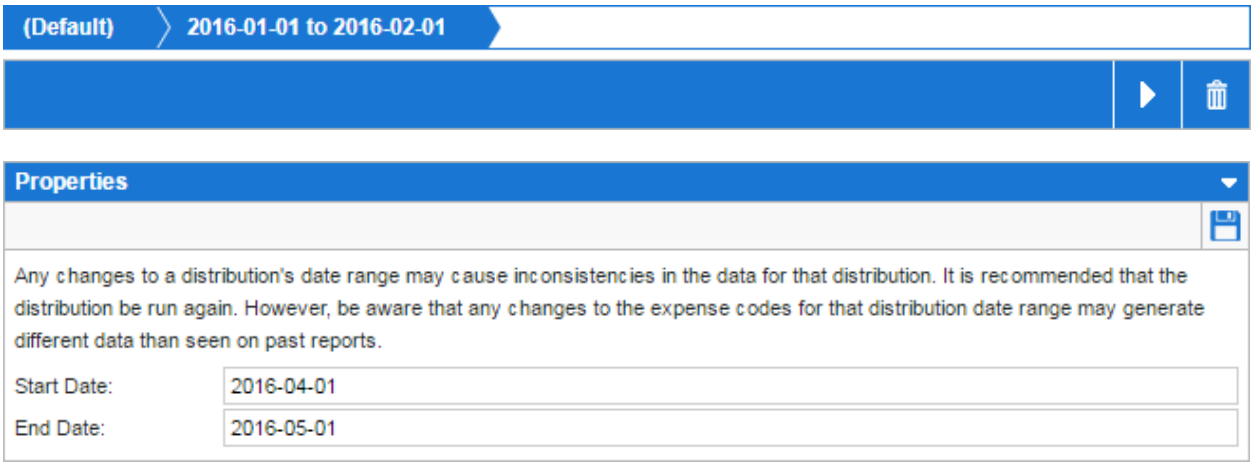

Unless you execute a search, the listing below shows the distribution costs for the selected distribution range.

Children

Display: Distribution Detail ▼

Amount	Quantity	User	Expense Code	Expense Category
$20.00	1.00000	Baker, Colin	Computer Lease	Monthly Expenses
$20.00	1.00000	Baker, Tom	Computer Lease	Monthly Expenses
$20.00	1.00000	Davision, Peter	Computer Lease	Monthly Expenses
$20.00	1.00000	Eccleston, Christopher	Computer Lease	Monthly Expenses
$20.00	1.00000	Hartnell, William	Computer Lease	Monthly Expenses
$20.00	1.00000	McCoy, Sylvester	Computer Lease	Monthly Expenses
$20.00	1.00000	McGann, Paul	Computer Lease	Monthly Expenses
$20.00	1.00000	Pertwee, Jon	Computer Lease	Monthly Expenses
$20.00	1.00000	Smith, John	Phone Lease	Monthly Expenses
$20.00	1.00000	Smith, Matt	Computer Lease	Monthly Expenses

More (10 | 50 | All)

The Imports Component

If your company directory is maintained in an external application or database, it is usually better to import it into the Predictive UC Analytics directory than to manually enter the information twice. Directory imports can even be scheduled to handle moves/adds/changes automatically after the initial import. Predictive UC Analytics supports two types of imports. You can import from a delimited file or

from a server with Lightweight Directory Access Protocol (LDAP) support, like Microsoft's Active Directory or Novell's eDirectory.

Template Listing

This listing allows you to manage, delete, run or schedule directory imports.

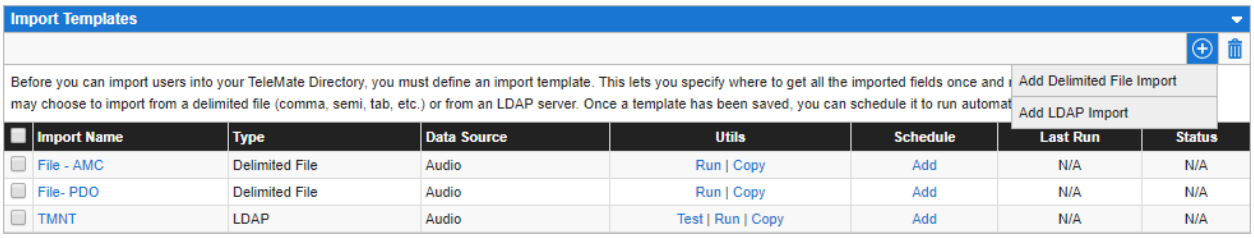

Hover the mouse over the “Add” button in the upper right corner of the listing and select a template type to create a new one. Click on an existing template’s name to modify it. Use the checkbox next to a template name and click the trashcan icon button to delete an import template. To run an import manually, select the "Run" link in the Utils column. Select the “Test” link before running an LDAP import to verify that it will pick up the LDAP attributes correctly. To schedule an import, or to modify its schedule frequency, select the “Add”, “Edit”, or “Clear” links in the Schedule column.

When you schedule an import, you will be asked to specify a frequency and run time for the import. You may also choose to be notified by email when the import has completed.

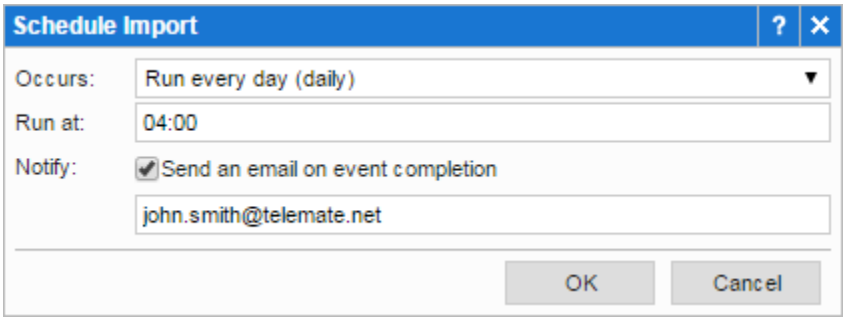

Delimited File Import

Predictive UC Analytics can import from a delimited file by specifying which fields in the file map to which organization properties.

Step 1: File Properties

The first step requires you to specify a delimited text file to import. This file **must** be visible to the server, either on one of the server’s local drives or on a network share (the Predictive UC Analytics services must have read access to that share). If the file is not on the server but you can access it from your web browser, select the “Upload File” button to upload the file to the server. Before selecting the “Upload File” button, you **must** specify the path on the server to upload it to. If you’re not sure where to put it, you can use an environment variable to specify a temporary folder (e.g. %TEMP%\myimport.txt). If you want to schedule this import, you need to specify a path on the server that some external process will push updated directory files to periodically.

Warning: If you have multiple data sources configured, the Data Source selection here is very important! You may **not** import a single file for multiple data sources. You must split them up into separate files and configure a separate import for each.

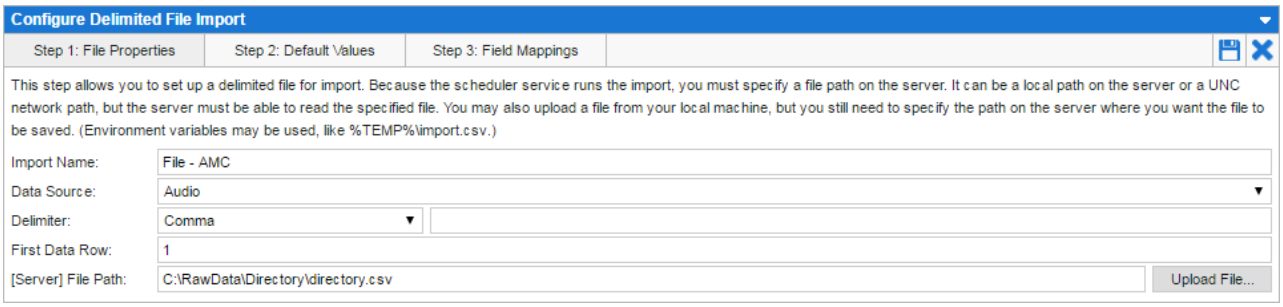

Step 2: Default Values

The second step allows you to choose a default value for the Company, Division and Department fields. You can either type in the name or use the "Choose…" button to pick a value from a list of Companies/Divisions/Departments that have already been created. If the delimited file does not contain one of them or the field is blank, the import will use the default value.

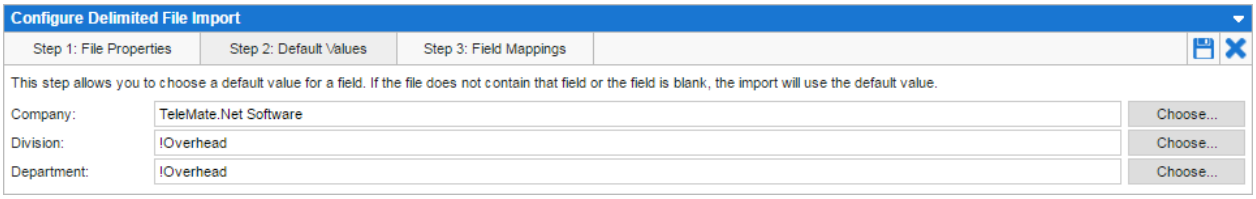

Step 3: Field Mappings

The third step shows a few sample records parsed from the delimited file specified in Step 1. If the sample records do not look right for any reason, double-check the file and the settings you specified for it in Step 1. Once everything looks okay, use the selection boxes to specify where each field from the file should be pulled into Predictive UC Analytics, and click the Save button (in the top-right corner).

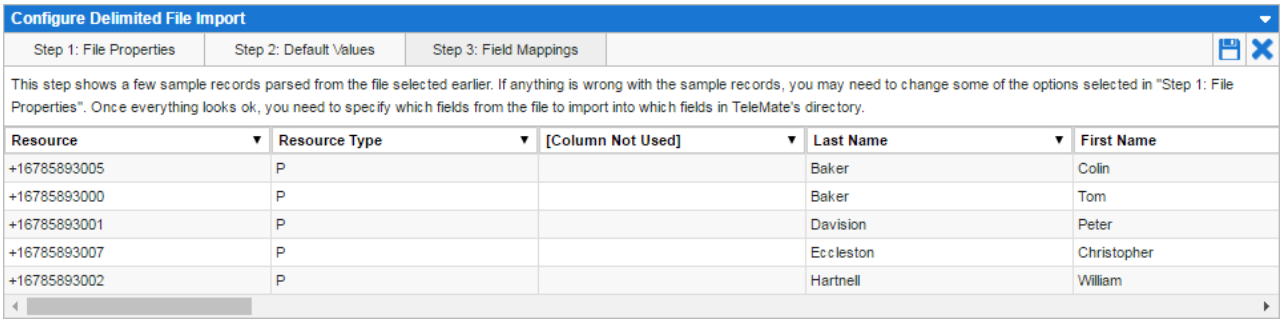

LDAP Import

Predictive UC Analytics can connect directly to your LDAP server to import directory information. Unfortunately, the mapping between your LDAP directory and the Predictive UC Analytics directory is rarely a perfect fit, which tends to make the configuration complex in all but the simplest cases. Fortunately, Predictive UC Analytics provides some very flexible and powerful options to "clean up" the

mapping. It can even be used to clean up some of the data it finds in LDAP, which almost always contains data entry errors that are hard to find until the data is forced into a structure that makes errors stand out clearly. TeleMate also provides professional services to help you through the complexity to get the most out of your LDAP integration.

The process of configuring an LDAP import generally includes:

1) If you have multiple data sources to import from a single LDAP server, you must first find a way to identify the data source for each extension you want to import. Extension 3001 on phone system A is **not** the same as extension 3001 on phone system B. You may or may not have overlapping extensions, but the import cannot make assumptions, so you must tell it which data source extension 3001 belongs to.

2) Next you must decide how you want to map LDAP Organizational Units, Groups, and/or attributes to the Predictive UC Analytics Company/Division/Department hierarchy. This depends on how you have organized your LDAP directory and how you want users to appear in the Predictive UC Analytics directory. It ends up being at least a little different for every customer. Some customers want to use the location attribute for division and the department attribute for department. Others want to use groups or organizational units to decide where to put everyone. Others want to use extension ranges or phone number prefixes to decide.

3) Most customers want to be able to use LDAP to "terminate" users in Predictive UC Analytics. The import cannot delete users because they historical data linked to them, so this is generally done by creating a copy of the main LDAP import. The copy will check for some sign in LDAP that a user has been deactivated (e.g. a special OU for terminated users or the "enabled" flag for ActiveDirectory accounts). This alternate import will transfer users out of their current department to a "disabled" area in the Predictive UC Analytics directory.

4) Once the basic structure of the LDAP import is laid out, it is time to worry about the details. You need to specify how to pick up each LDAP attribute, run a "Test" to spot-check examples of users being imported, and tweak the details of how those attributes are picked up.

5) Finally, run the import to see what shows up in the Predictive UC Analytics directory. The first run usually highlights a number of data entry problems that you may not be aware of in your LDAP directory. At this point, it is very common for customers to run through a few passes where they delete the entire organization in Predictive UC Analytics, fix some things in LDAP, perhaps make a few changes in the import definition, and run it again.

Step 1: Connection Settings

The first step requires you to supply all the information Predictive UC Analytics needs to connect to the LDAP server. You must also choose a "default" data source for any extensions imported. (This data source selection can be overridden using custom scripting checks on LDAP attributes in step 4.)

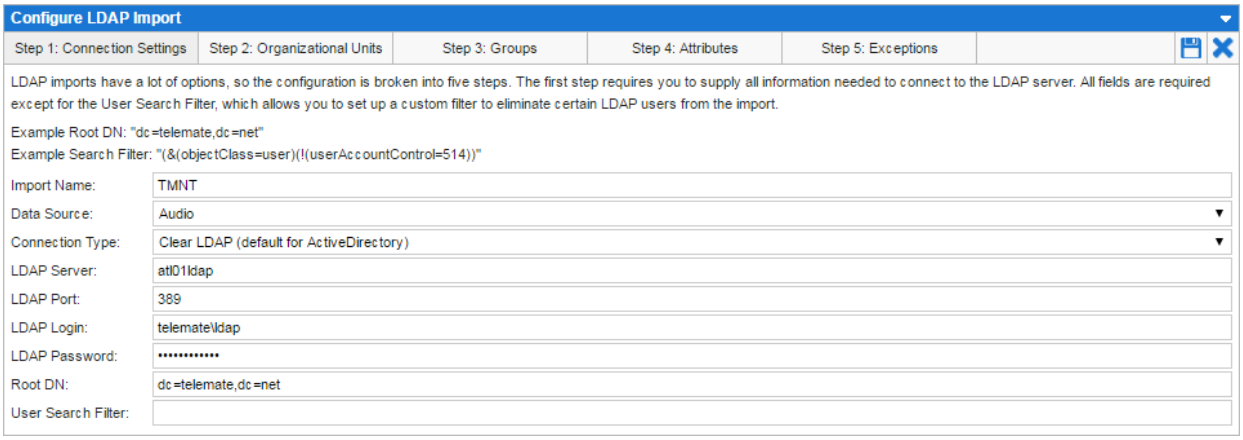

- **Import Name –** A unique name to identify the Import in Predictive UC Analytics.
- **Data Source –** The data source that any imported extensions or DISAs will be associated with.
- **Connection Type –** The type of LDAP connection used, Clear or Secure.
- **LDAP Server –** The IP or hostname of the LDAP server. The value can be multiple values separated by a comma. It allows you to provide a backup server in case the other LDAP server goes down. It is suggested that you use the full host name like "dc01.telemate.net".
- **LDAP Port –** The port number specifies which TCP port is used to connect to the server. If the LDAP server is not using its default port you should set it here. The default for the Clear LDAP connection type is 389 and Secure LDAP is 636.
- **LDAP Login –** This login should have read access to the RootDN. Some examples are:
    - Active Directory: telemate\joe.smith
    - Active Directory: cn=Joe Smith,ou=Development,ou=Telemate.Net Software,dc=telemate,dc=net
    - eDirectory: cn=admin,o=test;
    - Open Directory: uid=netspective,cn=users,dc=qa,dc=xserve,dc=com
- **LDAP Password –** The password to authenticate the login.
- **Root DN –** Specify the Distinguished Name (DN) of the “root” object to use for this import. Most customers just use the root node of your LDAP tree (e.g. your entire domain or organization), but in some cases it is better to create a few distinct imports that each check a different branch of the tree.
    - Active Directory: dc=telemate,dc=net
    - eDirectory: o=test
    - Open Directory: dc=xserve,dc=com
- **User Search Filter –** The search filter is an advanced option to help you weed out LDAP objects you want to avoid importing. Its syntax, and how you use it, depends on the type of LDAP server you have.

Note: All fields are required except for the User Search Filter. Also, if you are experiencing slow connection times check your server names to ensure that they are valid. Failed DNS resolution can slow down the connection times.

Step 2: Organizational Units

The second step allows you to map LDAP Organizational Units to Predictive UC Analytics Departments.

Important: The LDAP import will **only** import users it finds in Organizational Units and Groups you tell it to import. If you do not select anything in steps 2 or 3, the import will run, but it will not do anything.

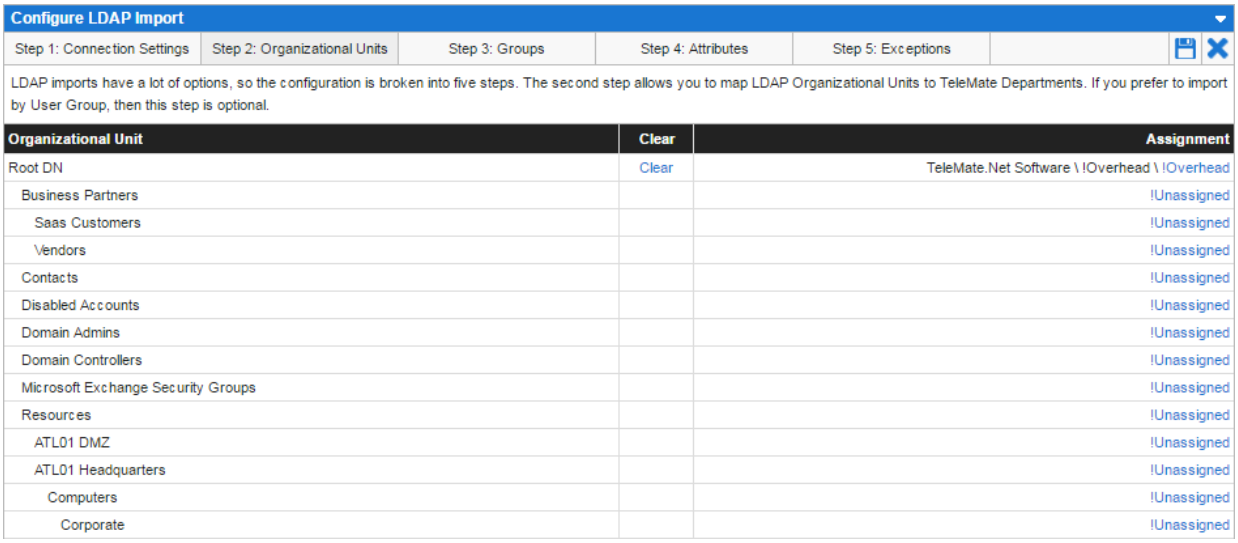

Find the Organizational Unit (OU) that you would like to import and click on the "!Unassigned" link to the right of it to select a Predictive UC Analytics Company/Division/Department to map it to. If you have not created the appropriate Company/Division/Department in Predictive UC Analytics yet, you must switch to the Organization page to create it. (We recommend that you open the Organization page in a new browser tab or window so you can switch back and forth easily.) When you click the "!Unassigned" link, the LDAP Assignment dialog will open with a list of the top 100 departments. The list can be narrowed down by typing part of the name of a department into the "Contains" field. When you find the department you want, select the department and click "OK". After you have assigned an OU to a department, you will see a "Clear" link you can use to remove the association.

It is important to keep in mind that OU's are organized in a tree. When you map an OU to a department, all users beneath that OU (and its children) will be mapped to that department. However, the department you link an OU to is just the "default" department for that OU. The default department for any OU can be overridden by: departments mapped to child OU's, LDAP Group mappings, LDAP attribute assignments, etc.

Step 3: Groups

The instructions for step 3 are exactly the same as the instructions for step 2, so read the explanation above first.

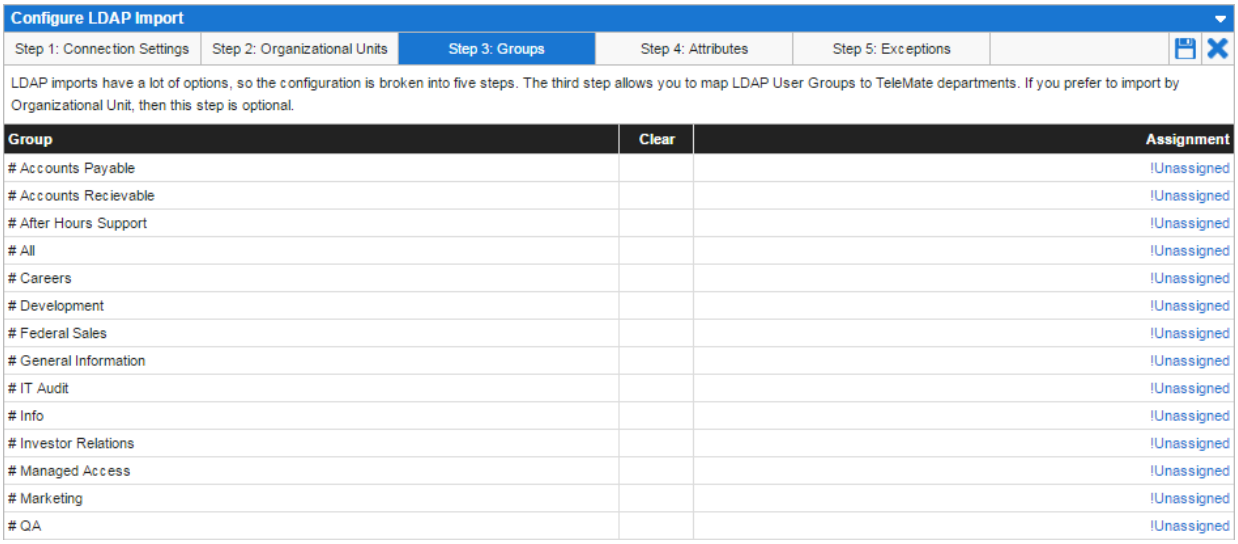

The only difference that needs to be pointed out here is that LDAP Group associations take precedence over LDAP OU associations. For example, imagine you have a "John Smith" in a "Sales" OU that belongs to a "Sales Engineer" Group. You can map both to different departments in Predictive UC Analytics, but "John Smith" can only belong to one department. In this case the import will choose the "Sales Engineer" department.

Step 4: Attributes

The fourth step allows you to map LDAP User attributes to Predictive UC Analytics directory fields. In the simplest cases, this is a simple mapping where you click on the name of each field you care about, and then select an LDAP attribute to assign to it. Because the Data Source, Company, Division, and Department fields all have default values you specified in steps 1-3, the Last Name field is the only field that is "required" to import an LDAP User. However, you can use LDAP attributes to pick up several additional fields, and you can use them to override all of those default values.

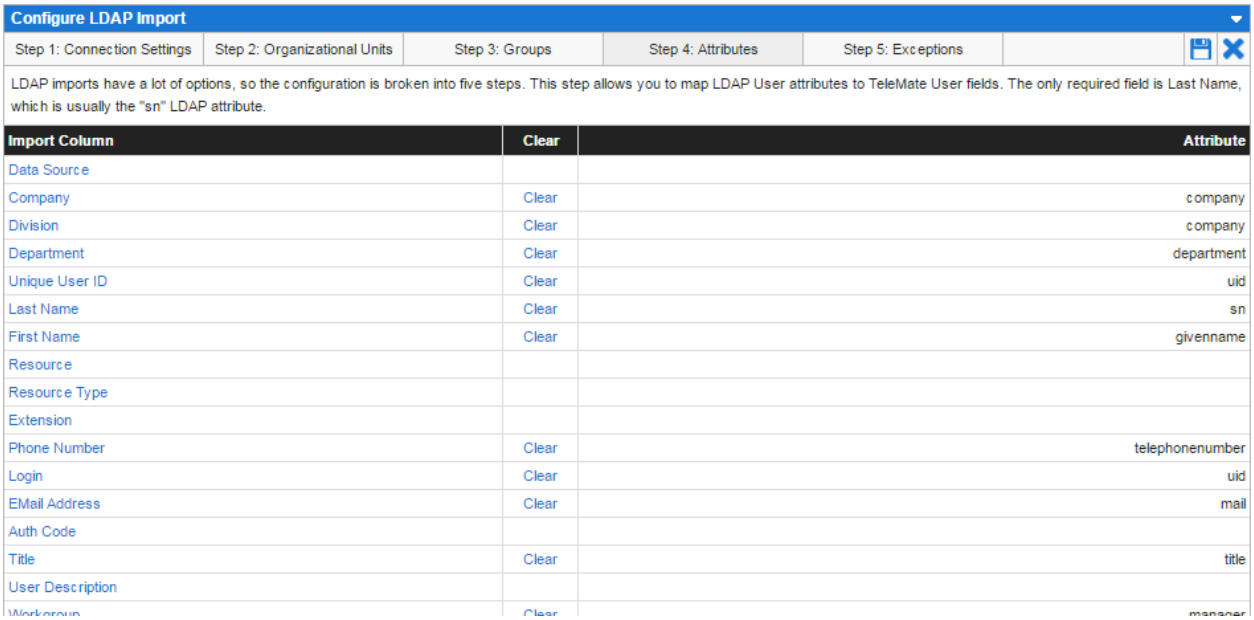

For most of our customers, that simple mapping is not enough. The data in the LDAP directory often is too "messy" for a clean mapping. For example, you may have a number like "1(678)589-7100" for users in one location and need to grab the last 5 digits (stripping the dash) to get the extension logged by your phone system. For users in another location, the numbers may look like "+1.678.589.7100 ext 7105" and need to grab all digits after "ext". Your LDAP directory may contain department or division names with typographical errors, or perhaps department numbers that you need converted to names.

To help you overcome these issues, Predictive UC Analytics allows custom scripting code to be specified for each and every import column. This custom scripting allows you to do virtually anything, but as it requires programming skills, the majority of our customers require assistance from our professional services group to configure it.

Step 5: Exceptions

The fifth step is optional, and we recommend that you avoid using it if you can because the users in this list must be maintained manually. This step is a crutch designed specifically for users that end up in the wrong department no matter what you do in steps 1-4. It may be due to users that currently work in one location/department but technically belong to another, or that belong to multiple LDAP Groups you've mapped to different departments, or something else entirely.

Advanced: Custom Scripting Examples

Warning: Unless you have some knowledge of basic programming/scripting, do not attempt to change these without assistance from technical support/services. This section is primarily for support/services to refer to when assisting you with cleaning up some of the data found in your LDAP attributes. The scripting language is called Ruby. Documentation for the Ruby scripting language can be found at: http://www.ruby-lang.org/. The most useful variables you can access from the script are:

- **value** - An array of values for the selected LDAP attribute.
- **user** - A hash table of attributes for the LDAP user object.

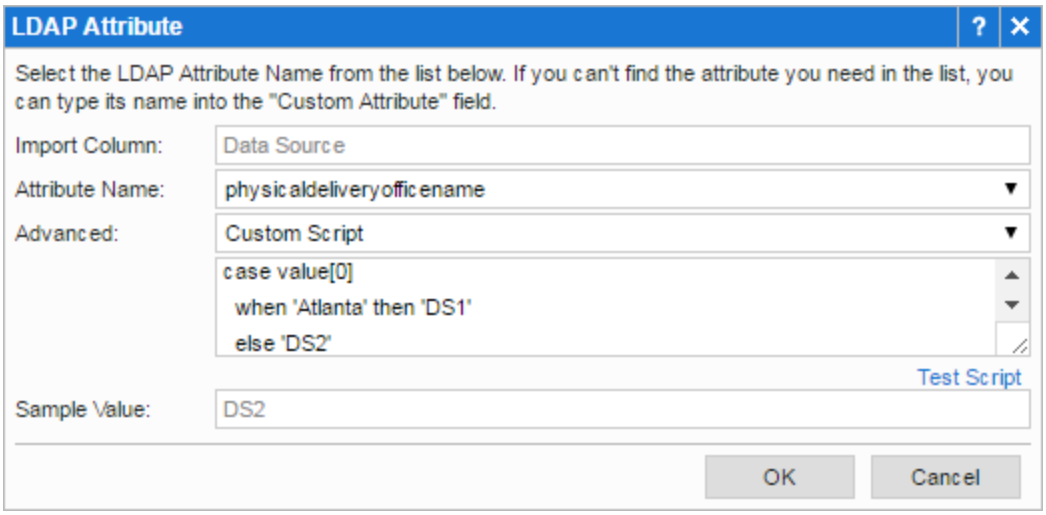

It is easiest to learn to use scripts by example, especially when you can get it to work by copying/pasting existing scripts and changing the names. For instance, let's say you need to select the data source for an extension based on the first digit in the extension. You could select your LDAP attribute containing the extension for the "Data Source" column, and then enter a script like this:

```
case value[0][0,1]
  when '4','5' then 'Data Source 1'
  when '6' then 'Data Source 2'
  else nil # Do not override the default data source
end
```

In this example, "value[0]" gets the first value of the selected attribute (LDAP allows multiple values per attribute), and the "[0,1]" gets a substring starting at position 0 for length 1 (i.e. get the first character of the first value). The "case" statement makes it easy to map one set of values to another. In the case above, any extension starting with a '4' or '5' gets mapped to 'Data Source 1'. Any extension starting with a '6' gets mapped to 'Data Source 2'. All other extensions return nil (which means no value). If no Data Source is selected, the import will use the default Data Source you selected in step 1. Here are some other examples:

```
# You can use complex regular expressions to check your values
# This script singles out extension ranges 4100-4399 and 4400-4799
case value[0]
  when /^4[1-3]\d\d$/ then 'Data Source 1'
  when /^4[4-7]\d\d$/ then 'Data Source 2'
  else nil # Do not specify a data source
end
```

```
# You can also use if-then-else logic, variables, loops, data conversion, and math calculations.
# This script singles out extensions 4100-4399, 4400-4799, and the Atlanta location.
ext = user['extension'][0].to_i
```

```
loc = user['l'][0]
if ext >= 4100 and ext <= 4399
 'Data Source 1'
elsif ext >= 4400 and ext <= 4799
 'Data Source 2'
elsif loc == 'Atlanta'
 'Data Source 3'
else
 nil # Do not specify a data source
end

# If you have multiple values in one LDAP attribute, you can scan them all
ret = nil # This will be the value we return to the import
value.each do |v| # Loop through all values in the "value" array
 if v[0,1] == '4' or v[0,1] == '5'
  ret = 'Data Source 1'
 elsif v[0,1] == '6'
  ret = 'Data Source 2'
 end
end
ret # Do not forget to return the selected value to the import
```

The Invoices Component

The Predictive UC Analytics invoicing features are an extension of the Expenses feature, which means they can only be accessed if you purchased an Enterprise license. The invoicing features are disabled in the web interface by default. To enable them, click on the “Settings” link in the navigation menu and look for the “Features” option in the “General Settings” section. In “Features”, enable “Invoices” and click the OK button.

Once enabled, you will see new invoicing options in several places in the Directory component. The first thing you should notice is a new “Invoices” link in the navigation menu. From the “Invoices” component you can create, review, or delete invoices by batch.

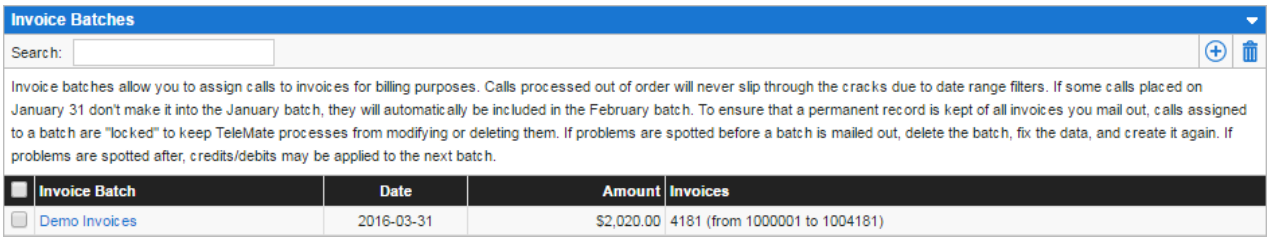

Generating Batches

Keep in mind that we recommend against manually creating or deleting invoices (or batches). The recommended steps to get started with invoice batches are:

1) Go to the Reporting component and select the “Invoice Batch” report (or your customized version of it, as most customers want their invoice report tailored to fit their company).

2) Specify a “Batch Level”, which determines whether invoice batches should be generated for a set of companies, divisions, or departments. To keep the list of batches easy to manage, always choose the highest level (e.g. choose company unless you have different divisions or departments that have different billing cycles).

3) Specify a “Batch Type”. For the first few billing cycles, we recommend you use the “Test Batch” first to spot-check invoices for potential problems (e.g. calls processed incorrectly, expense code assignments incorrect). When all visible problems have been fixed, switch to “Final Batch”. In the future, if you need to reprint bills from prior batches, select “Review Existing Batches”.

4) After running a final batch, double check it. If problems are detected before invoices are sent out, we recommend that you delete the batch, fix the problem, and re-run the batch. As your list of invoice batches grows over the years, this will help keep it clean and easy to manage.

5) In most cases when problems are detected by individual clients/customers after bills are sent out, we recommend applying adjustments (credits/debits) to those users in the directory so that corrections will show up in their next invoice. This allows you to preserve a record of the original invoices in your database.

6) If you have a special case where an individual client needs a bill to be fixed and reprinted, it may be best to manually delete a specific invoice (preserving the rest of the batch that you sent out to your other clients) and then manually run a special batch just for that one client. To print the new invoice, you would then run the “Invoice Batch” report in “Review Existing Batches” mode and filter on the custom batch.

7) If/when you feel comfortable automating invoice batch generation (e.g. everything runs smoothly for a few billing cycles), you can schedule the Invoice Batch report to generate invoices automatically. We still recommend spot-checking them.

Managing Invoices

In addition to being able to review invoices by the batch that generated them, the Organization tree in the Directory component also allows you to search for them. You can either search by invoice number in the “Find” tab, or you can select any entity in your organization and change the “Children” section to display Invoices assigned under the current entity. When you look at the invoices assigned to a specific user, you will also see a special link named “!Unassigned Calls” that represents the next invoice (which has not been generated yet). Click that link if you wish to view/modify recent adjustments and calls waiting to be assigned to that user’s next invoice.

There is also a “Create Invoice Batch” button at the top of the page (next to the delete button) that allows you to manually create a batch of invoices for the currently selected entity. We generally recommend that our customers use the “Invoice Batch” report instead, but it is not required.

Regardless of how you find an invoice, when select one you will see all adjustments, expense amounts, and call records assigned to that invoice. You may not modify anything assigned to an invoice because those records are "locked" to ensure that you always have a record of the bills you sent out to clients or customers. You may delete an individual invoice to unlock those records, but we recommend you try to avoid doing that. Instead, we recommend that you add adjustments (credits and debits) and select specific calls to credit back to that user, which you can do directly from the invoice page.

Reminder: The adjustments you create (and calls you credit) from an invoice will not be visible on the current page after you save them. This is because they can not be applied to an invoice because all of its records are "locked". New adjustments must wait to be applied to the next invoice. To view/modify recent adjustments or calls, you must select the special "!Unassigned Calls" link for the current user.

Applying Adjustments

There are two methods to apply an adjustment to a user. The first method is to simply click the "Add" button in the Adjustment section of an invoice (or the "!Unassigned Calls" page):

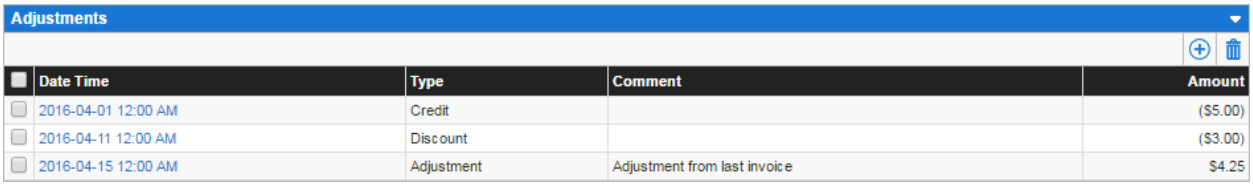

Clicking the Add button will bring up the "Invoice Adjustment" dialog. This dialog requires you to specify a date, an adjustment amount (positive for debits, negative for credits), an expense code (which is necessary to track which general ledger to apply the amount to and which taxes to apply to it), and a comment (a description which typically shows up on the user's invoice).

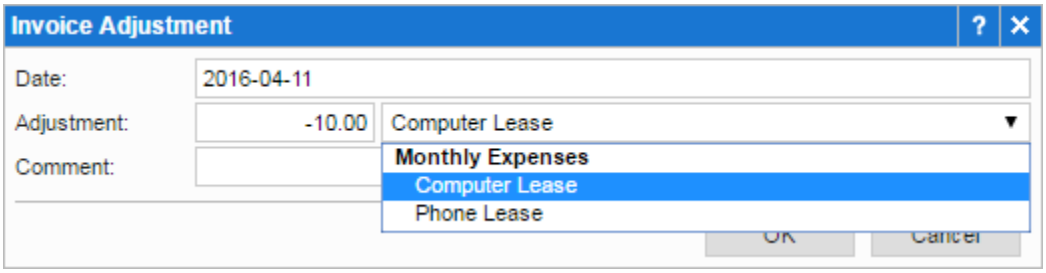

The second method is to credit specific calls. This will create a copy of the call record with a negative amount, which will show up on (and be applied correctly to) the next invoice. To credit specific calls, simply select them in the invoice call listing and select the "Credit" button in the top-right corner of the call detail section.

Note: When looking at "!Unassigned Calls", which represents the "next" invoice, the "Credit" button is replaced with a "Delete" button, which essentially credits the call before your customer even sees it.

Detail

Search: Dialed Number

Date Time	Resource	Dialed Number	Location	Traffic Type	Duration	Cost
2016-01-01 07:09 AM	+16785893001	ANI +19139472838	WCLEVELAND, KS	Incoming	00:18:33	$0.00
2016-01-01 03:29 PM	+16785893001	+12148612025	DALLAS, TX	Interstate	00:01:28	$0.01
2016-01-01 07:26 PM	+16785893001	+232232794569	Mobile Service, SI	International	00:02:51	$0.02
2016-01-01 07:33 PM	+16785893001	ANI +13372411354	JEANERETTE, LA	Incoming	00:18:20	$0.00
2016-01-01 07:55 PM	+16785893001	ANI +12156079497	PHILA, PA	Incoming	00:02:38	$0.00
2016-01-01 08:05 PM	+16785893001	ANI +19098250629	COLTON, CA	Incoming	00:06:48	$0.00
2016-01-02 05:36 AM	+16785893001	ANI +17203509607	DENVER, CO	Incoming	00:03:48	$0.00
2016-01-02 07:14 AM	+16785893001	+17707956601	ATLANTA NW, GA	Interstate	00:19:01	$0.01
2016-01-02 01:04 PM	+16785893001	ANI +13093439660	GALESBURG, IL	Incoming	00:04:44	$0.00
2016-01-04 06:57 AM	+16785893001	ANI +16238482966	PHOENIX, AZ	Incoming	00:09:05	$0.00

More (10 | 50 | 100)

The Costing Component

The costing component allows you to view and manage costing methods and tariffs. These are used to assign costs to session activity when it is processed by Predictive UC Analytics processing engine. The left-hand side of this section provides controls for navigating through a listing of configured costing methods and tariffs.

Important: The entire costing component can be enabled or disabled. It also contains many individual advanced features that can be enabled or disabled. If a costing feature you need is hidden, refer to "Features" in the "General Settings" section above.

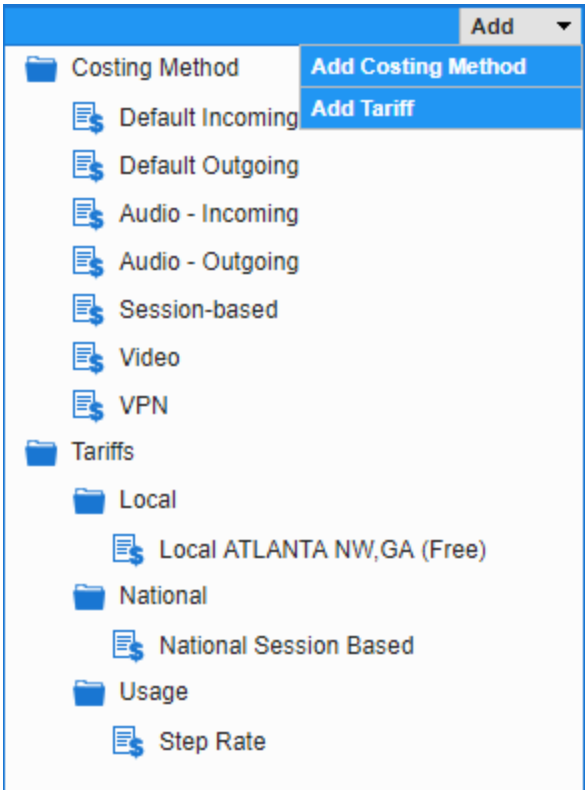

Once you select a costing method or tariff, the properties for it will appear on the right-hand side of this section.

Any changes made to cost settings, including both cost methods and tariffs, will not take effect until the next day. A notification of this will be show in the upper right hand corner as a reminder. The notification will also allow you to "Rebuild Processed Data" for a specified date arrange if you want to have the costing changes to be applied to historical data.

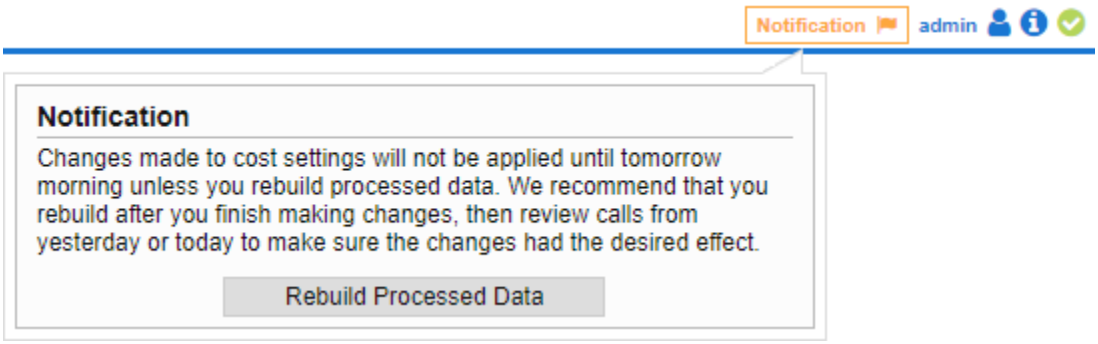

Common to All Levels

Whether you select a cost method or a tariff, at the top you will see a "You are here" section.

The line of text in the image above shows the current selection is the "National Session Based" tariff in the "National" tariff type of the "Tariff" costing section. On the right of the image, you will see a number of buttons for making changes to the current entity:

1) **Lock/Unlock** – Lock the selected cost method or tariff to prevent it from being modified or deleted. You can unlock it again later if you need to modify it.

2) **Import Destination and Rates –** This option is only available for a tariffs that are not usage-based. See the "Importing Destinations and Rates" section of this document for more detailed information.

3) **Copy –** Click on this icon to create a copy of the current cost method or tariff. It will start creating a new one with the same properties, allowing you to modify and save it or cancel.

4) **Delete –** Select to delete the current cost method or tariff. You will not be allowed to delete it if it is currently in use. For a tariff, you will need to find the cost methods using it and either delete them or assign a different tariff. For a cost method, you will need to find the data sources, gateway groups, departments, or tenants using it and assign a different cost method.

Costing Methods

There are two types of costing methods supported in the web interface, standard and VPN costing methods. The general properties are the same for both, but they each have different additional properties. If enabled, a currency and applicable taxes may be selected here. Taxes selected here do not affect processing, but they can be applied in “Call Activity” reports that have an “Apply Taxes” option.

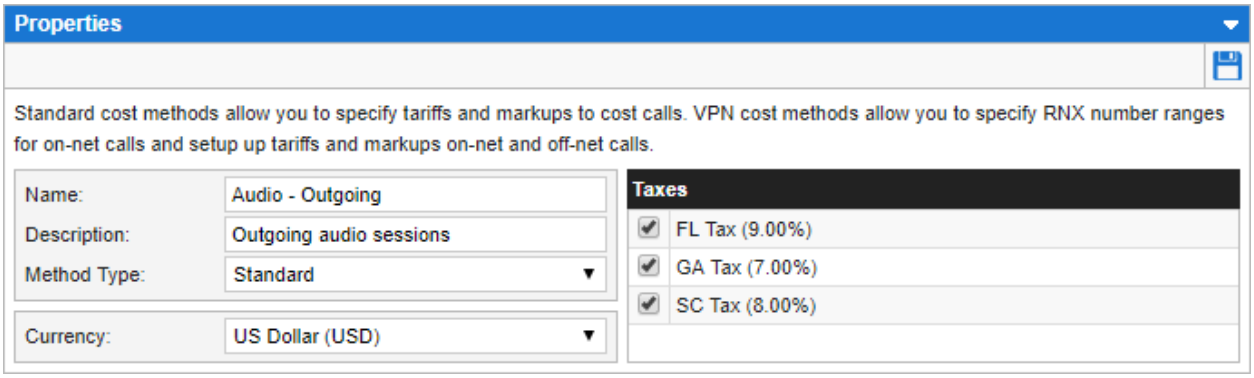

Settings for a Standard Cost Method

The tariffs section of a standard costing method is where you can assign tariffs to incoming, local, national and international traffic. If enabled, special cost adjustments like percent markups, minimum costs, and surcharges can also be set for each of the traffic types at this level.

The advanced options can be used to override dialing prefixes and our provided NANP local dialing areas to determine whether a call is local, national, or international. One example would be to count Alaska and Hawaii as international to avoid giving them the same national rate used for the other 48 states.

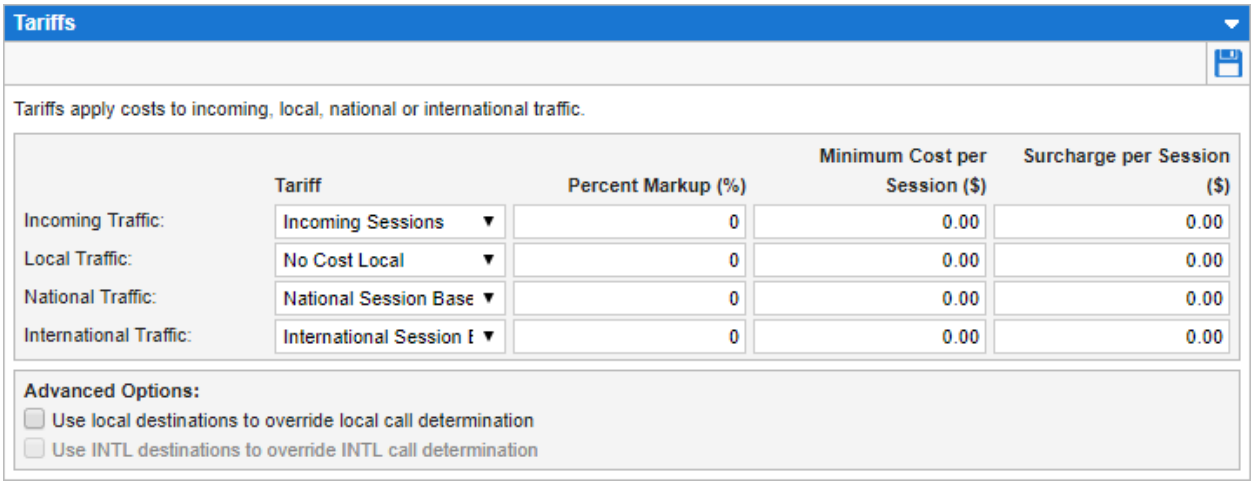

Note: The advanced options will be disabled if the appropriate tariff selected has a “default” destination. This is because a default destination would override all traffic, forcing all traffic to be either local or international.

Settings for a VPN Cost Method

When you select a VPN cost method, you will have access to additional properties to configure RNX ranges (phone numbers considered “on-net”), On-Net, and Off-Net settings for the cost selected costing method.

In the RNX settings section you have buttons available on the right that allow you to either add a new RNX entity or delete entities that have been selected. Clicking on the RNX range of an existing range will allow you to edit the properties of the range.

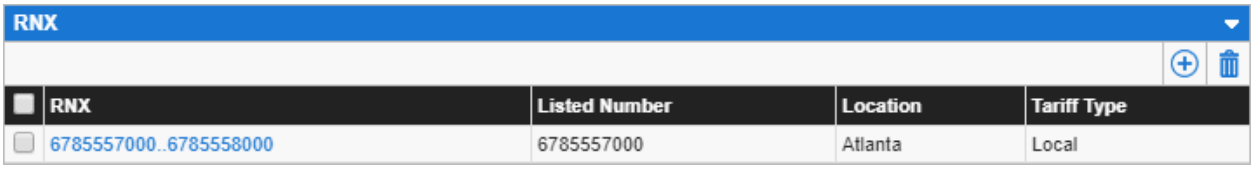

The options for the On-Net and Off-Net sections of a VPN costing method are the same. You can assign tariffs to incoming, national and international traffic. Percent markups, minimum costs, and surcharges can also be added to each of the traffic types at this level.

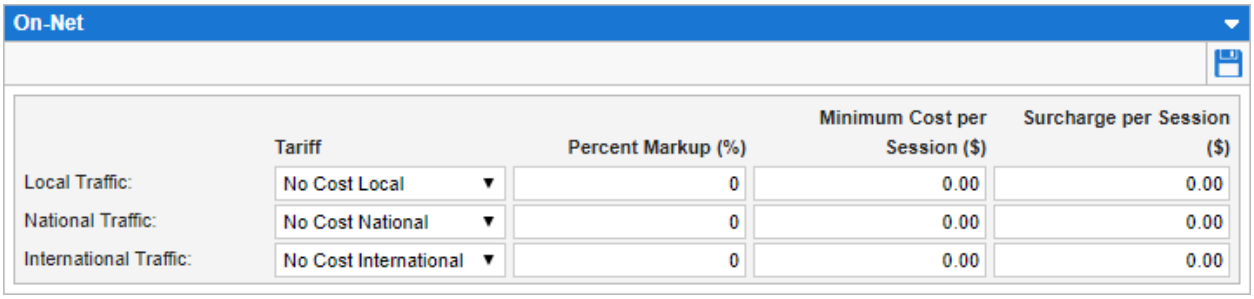

Tariffs

Tariffs are divided into 5 main tariff types: incoming, local, national, international, and usage-based. Usage-based tariffs are special, and apply rates based on monthly usage tiers. All other tariffs are destination-based, and apply rates based on destinations (or rate centers) configured using dialing patterns.

Destination-Based Tariffs

When you add a new destination-based tariff (incoming, local, national, or international), the only required option is its name, which must be unique. If you have currency, holidays, and evening/weekend rate features enabled, you will see the additional properties displayed in the screenshot below.

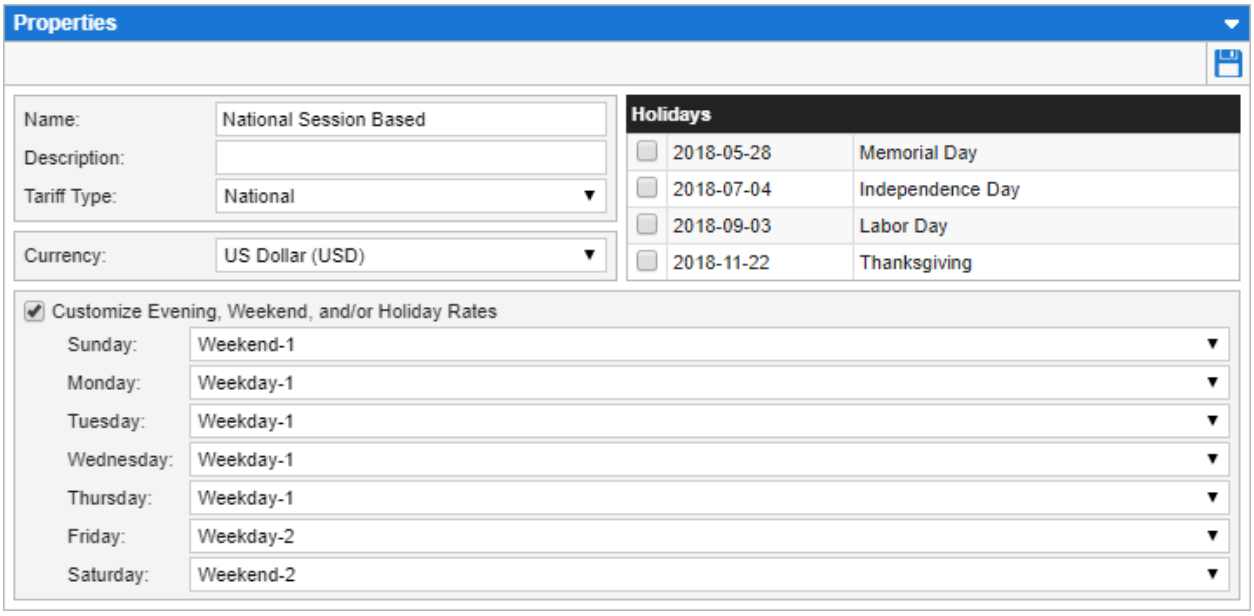

Rates themselves must be assigned to destinations.

If you want to apply the same rate for all traffic matching the current tariff type (incoming, local, national, or international), simply select the "!Default" destination and enter the rate. If you enabled evening/weekend rates, you may need to enter different rates for different day types and start times.

If most destinations use the same rate but there are a few exceptions, select the "!Default" destination and enter the default rate first. Then add the exception destinations, specify the dialing patterns that identify them, and set their rates separately. A good example of this would be having the same long distance rates for all US states except for Alaska and Hawaii.

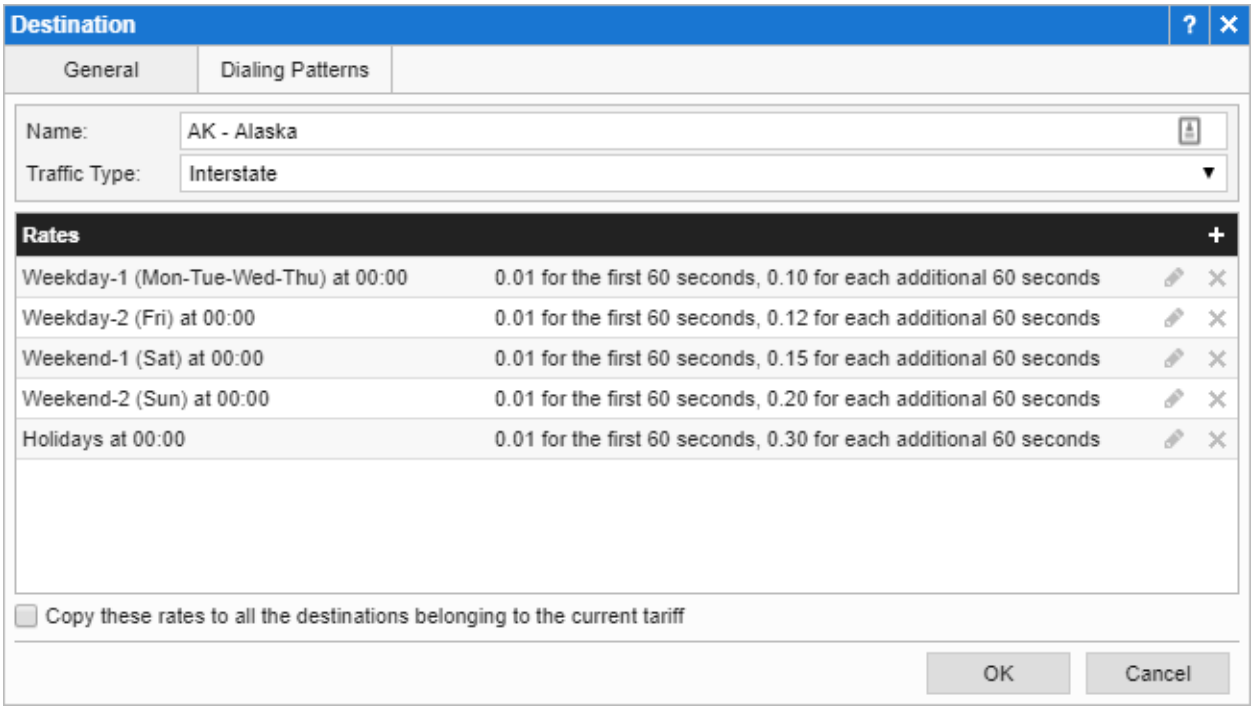

If this tariff has evening/weekend rates, you will see a table of rates for different day types that you can add, edit, or delete. If not, the table will be replaced by a section asking you to enter initial and additional rates and time periods. Set the additional rate to zero if you want a fixed cost per call.

If this is the "!Default" destination, you should not add any dialing patterns to it. If not, click the "Dialing Patterns" tab to switch to it, and click the "+" button in the top-right corner of the dialing patterns list. To edit an existing dialing pattern, click on its name. To delete a dialing pattern, click on the "x" button to the right of it.

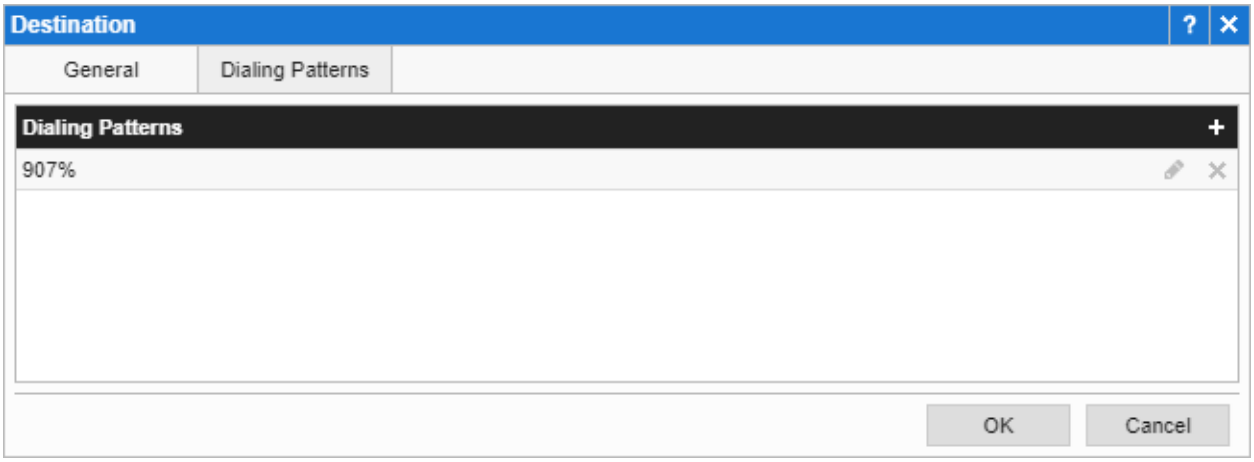

If you have many different destinations with many different dialing patterns, it can be very time-consuming to add them all one at a time. To make that easier, we allow you to copy and paste lists of dialing prefixes into the dialing pattern dialog from other documents open in Notepad, Excel, etc. The dialing prefixes can be separated by spaces, commas, on their own lines, etc.

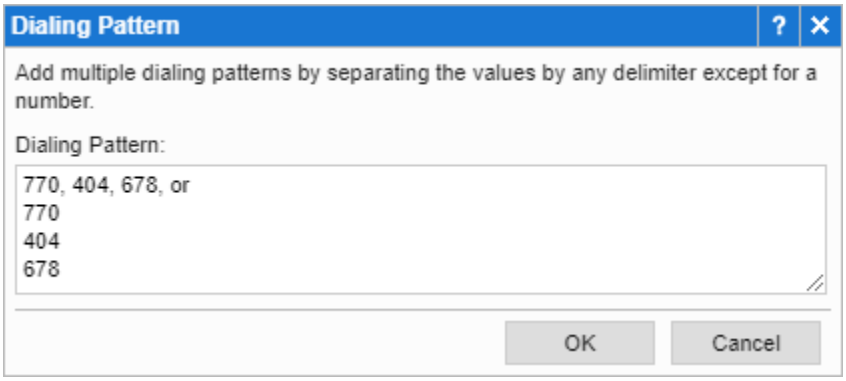

Predictive UC Analytics can also use its own built-in location tables to auto-populate destinations and prefixes for you. For local NANP tariffs, you can select an area code and exchange, and it will populate all prefixes for its Local Dialing Area. For national NANP tariffs, you can have it populate all prefixes for the 50 US states. For international tariffs, you can have it populate all known countries by country code (with an option to split mobile prefixes so they can be assigned different rates).

Importing Destinations and Rates

If you can get a list of destination names, rates, and dialing prefixes in a file from your service provider, we strongly recommend that you use our tariff import feature rather than copying it all by hand (which is more error-prone). This import can be very simple or very complex depending on your needs. Contact support if you have any problems. Depending on how your provider gave you the rates, you may need to manipulate the file a bit before importing it. The import requires a CSV file. Most providers can provide rates in an Excel spreadsheet, which is easy to save to a CSV file.

The image above shows that we are importing destination and rates into the “National Session Based” tariff. On the right of the image you will see two buttons, “Run Import” and “Cancel”. “Run Import” will import the chosen CSV file based on the options chosen. “Cancel” will return you to the tariff’s property page.

To start configuring an import, click the “Select File…” button. Once you have chosen your CSV file, click “Load File”. Loading the file will upload the file to the server so it can analyze the file and assist you with configuring the import.

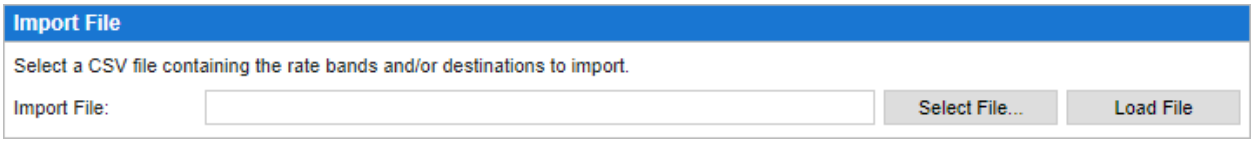

Next, you will need to select a Traffic Type to assign to any new destinations created during the import. If you enabled evening/weekend rates for the current tariff, you will also need to specify the day type

and start time to use for this pass. For evening/weekend rates, you will need to run multiple import passes (once for each day type and time the rate changes).

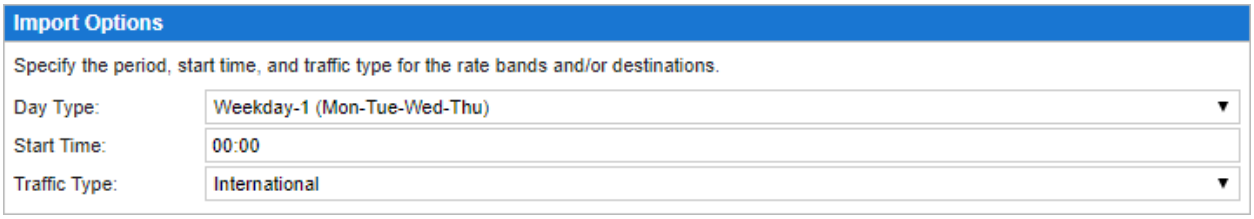

Once the file is loaded, you will be presented with two side-by-side tables. The table on the left shows you the columns the import needs, and initially they will be empty or set to some default value. The table on the right shows the content from the original columns of the CSV file.

To configure each import column, simply select the column header above it. After you configure an import column, look at the updated values below it to make sure it looks correct. If not, select the column header again and fix it.

Note: Some problems may require you to make changes to the spreadsheet, save it to a CSV file again, and then load it again.

Import Columns and Values

Importing			From Selected File		
Destination Name	Dialing Pattern	Initial Rate	A	B	C
Afghanistan	01193%	1.55 for the first 60 seconds	Afghanistan	93	1.54696
Albania	011355%	1.48 for the first 60 seconds	Albania	355	1.47688
Algeria	011213%	0.85 for the first 60 seconds	Algeria	213	0.85072
Andorra	011376%	0.72 for the first 60 seconds	American Samoa	1684	0.88096
Angola	011244%	1.98 for the first 60 seconds	Andorra	376	0.72368
Antarctica (Casey)	011672%	1.83 for the first 60 seconds	Angola	244	1.97608
Antarctica (Scott)	011672%	1.83 for the first 60 seconds	Anguilla	1264	0.9364
Argentina	01154%	0.78 for the first 60 seconds	Antarctica (Casey)	672	1.83064

Configuring the “Destination Name” column only involves choosing a column from the CSV file you uploaded. This import field is required.

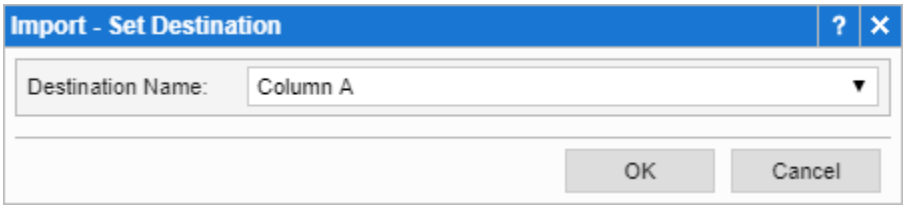

The “Dialing Pattern” column is optional, but we strongly recommend that you configure it if your CSV file does contain phone number prefixes. In addition to choosing a column from the CSV file, you have additional options to add an optional dial prefix or to filter the import by country code. The dial prefix is generally used to add a prefix like 011 or 00 to country codes listed in the CSV. The filter is generally

used to split a rate file containing both NANP and non-NANP rates into two separate tariffs (e.g. one national and one international).

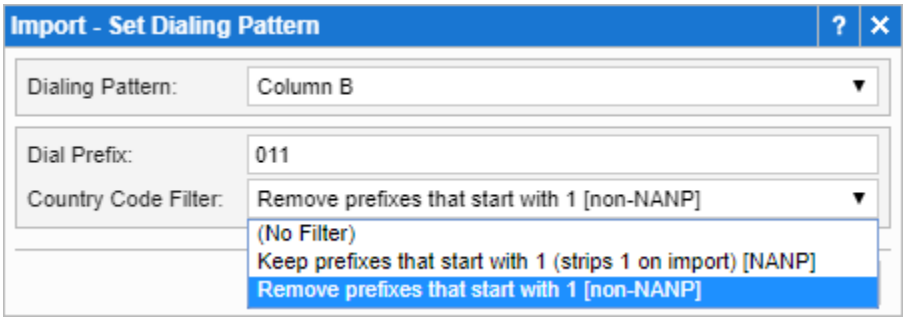

The “Initial Rate” and “Additional Rate” columns have the exact same options. For both period and cost, you can either choose a column from the CSV file or enter a fixed value. The only difference between initial rate and additional rate is that the initial rate has a checkbox that allows you to copy the settings you select over to the additional rate (because initial and additional rates are often the same).

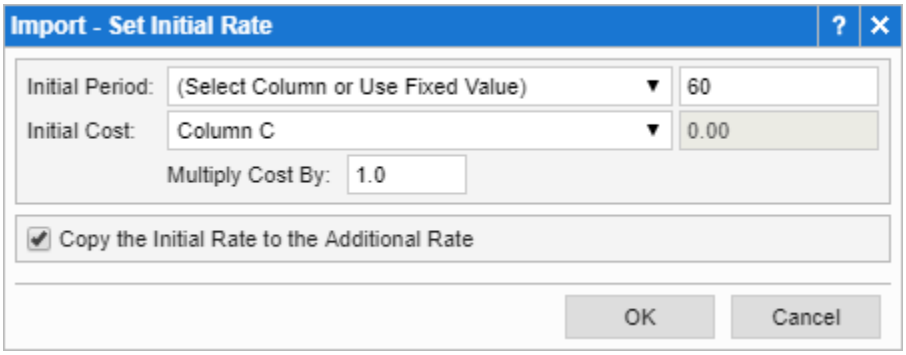

Settings for Usage-Based Tariffs

Usage-based tariffs have the same basic properties as standard tariffs. In addition to the basic properties, usage-based tariffs require you to define a base unit period. Some common base unit periods would be:

- Per session (e.g. per call, email, text message, etc.)
- Per minute (set to 60 seconds)
- Per KB (set to 1024 bytes)
- Per MB (set to 1048576 bytes)

Note: Each session will be rounded up to the next base unit. If a video call has a base unit of 1MB, a video call that used 10.5MB would be rounded up to 11MB.

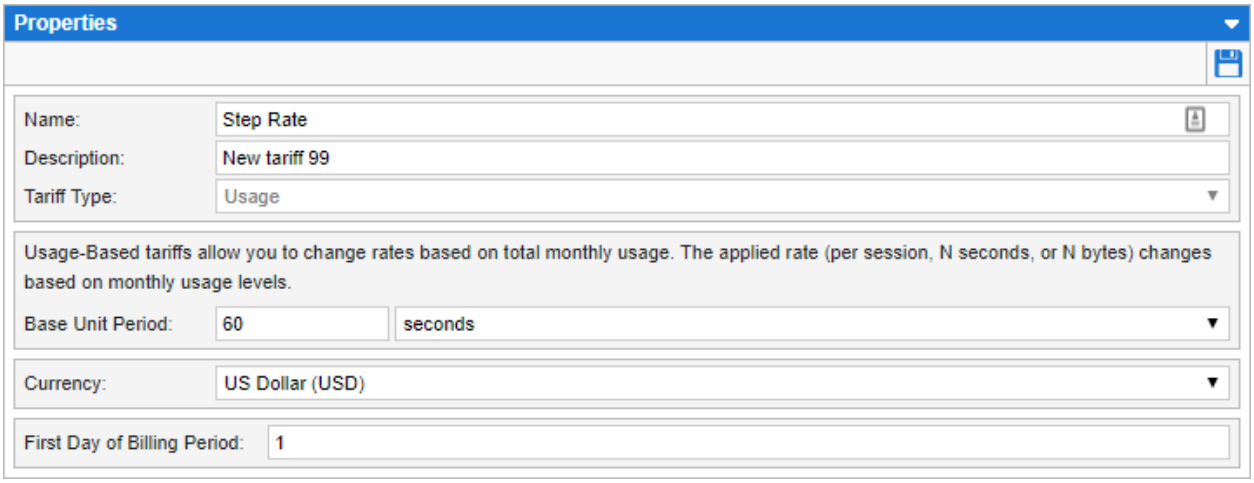

In the Usage Rates section you have buttons available on the right that allow you to either add a new monthly usage rate or delete ones that have been selected. Clicking on the monthly usage rate will allow you to edit the properties of the rate.

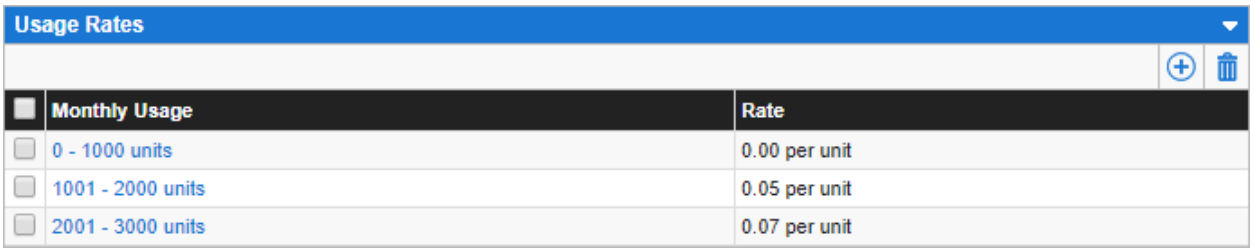

When adding a new usage rate tier, enter the upper bound for the number of units and the amount to charge per unit. In the example above, the first 1000 units in each month would be free (e.g. included in the provided service), the next 1000 units in each month would be rated at $0.05 per unit, and the next 1000 units would be rated at $0.07 units per month.

The Security Manager Component

The Security Manager component is only visible to accounts with admin privileges. Select it to configure privilege groups, manager accounts, and (if tenant features are enabled) tenants.

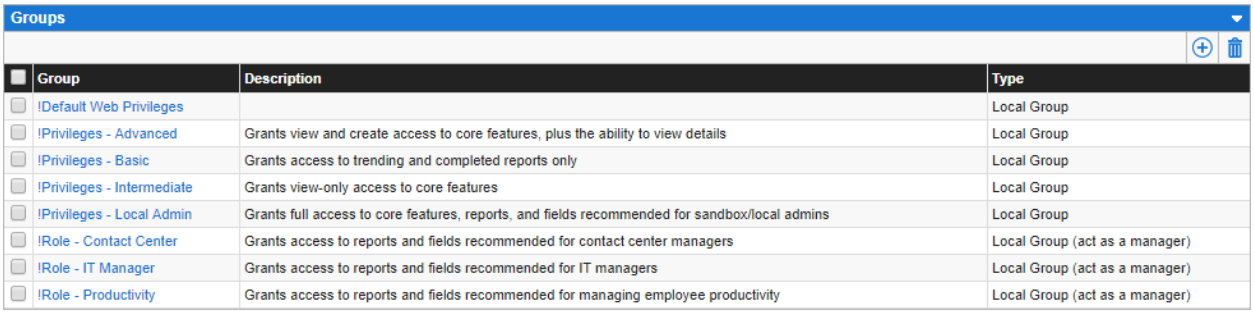

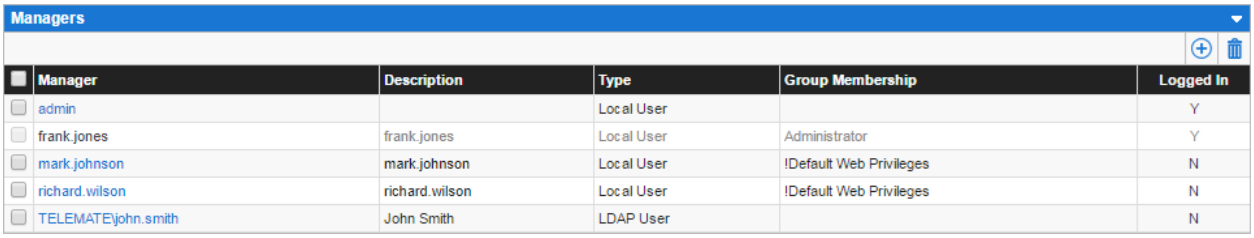

Managers

Manager	Description	Type	Group Membership	Logged In
admin		Local User		Y
frank.jones	frank.jones	Local User	Administrator	Y
mark.johnson	mark.johnson	Local User	!Default Web Privileges	N
richard.wilson	richard.wilson	Local User	!Default Web Privileges	N
TELEMATE\john.smith	John Smith	LDAP User		N

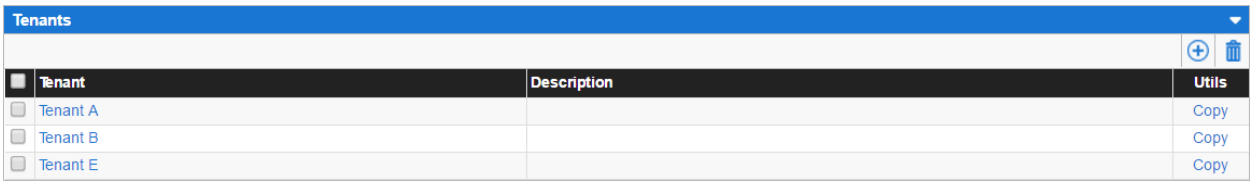

Tenants

Tenant	Description	Utils
Tenant A		Copy
Tenant B		Copy
Tenant E		Copy

Note: To enable tenant features, select the “Settings” link from the navigation menu. Then select “Features” in the General Settings section, and enable “Tenants”.

To create a privilege group, manager account, or tenant click the “+” icon located in the top-right corner of that section. To delete one, check the box next to it and click the trash icon for that section. To modify one, click its name. We have provided a few built-in privilege groups to help you get started. The groups that start with “!Privileges” grant different levels of access to basic web interface features. The groups that start with “!Roles” grant access to sets of report templates and filters that make sense for managers in those roles.

If you have configured LDAP Sources (refer to the System Settings section above), you may also specify an LDAP User when creating a new Manager account or an LDAP Group or Organizational Unit (OU) when creating a new privilege group. If you create a privilege group tied to LDAP, you can avoid having to create individual manager accounts. Any user that logs in with valid LDAP credentials belonging to one of those LDAP Groups or OU’s will have a manager account created automatically the first time they log in. Every time that user logs in, his LDAP OU and group membership will be checked to ensure that he or she still has access.

Configuring a Privilege Group

Configuring a basic privilege group is simple. When you configure a Local (non-LDAP) Group, the only required value you need to specify is a group name. When you configure an LDAP Group or Organizational Unit, you must select an LDAP Source (which is essentially selecting a Windows domain if you are using ActiveDirectory), and then you must pick from the list of Groups and OU’s discovered from that LDAP server. If the list appears to be out of date, click the Refresh button to the right of the list.

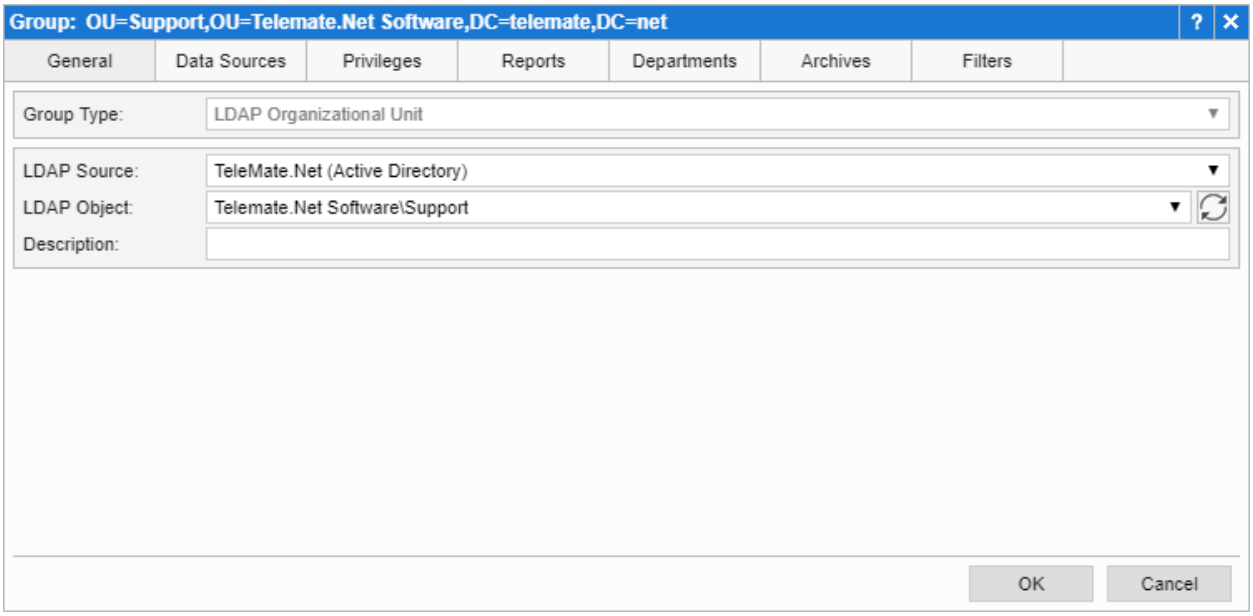

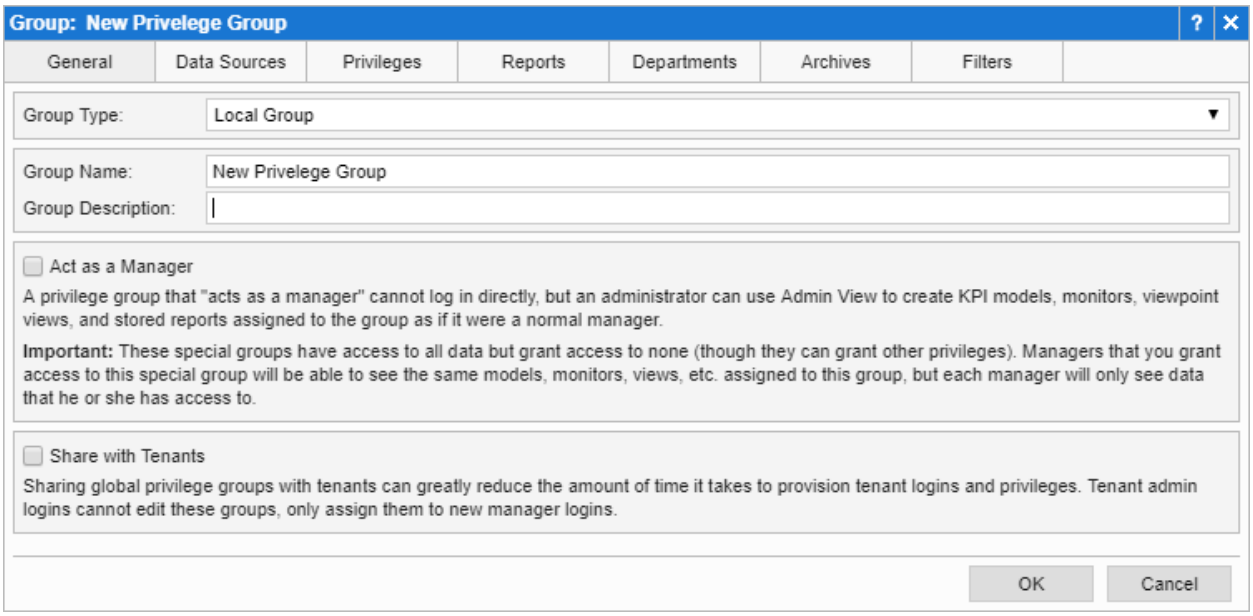

All of the other tabs provide lists of checkboxes. If you want this privilege group to grant access to specific data sources, privileges, reports, departments, archives, or filters, then select the appropriate tab and check the appropriate box.

Local privilege groups can also “Act as a manager”. They cannot actually log in, but admins can create and assign shared KPI models, monitors, viewpoint views, etc. to them. The built-in “!Role” privilege groups all have pre-configured KPI models, monitors, and viewpoint views to help provide new admins with a better out-of-the-box experience. Rather than having to create them from scratch, admins can simply modify the defaults.

Configuring a Manager Account

Configuring a manager account is very similar to configuring a privilege group, but it has a few additional options to set up. Those additional options will depend on whether you set the “Manager Type” to “Local User”, “LDAP User”, or “Alternative Authentication User”. When a new manager account is created, a report archive is created with the same name for this manager to send completed reports to. Archives have a default size limit (in MB) and expiration time (in days), and you may override these when creating a new manager account.

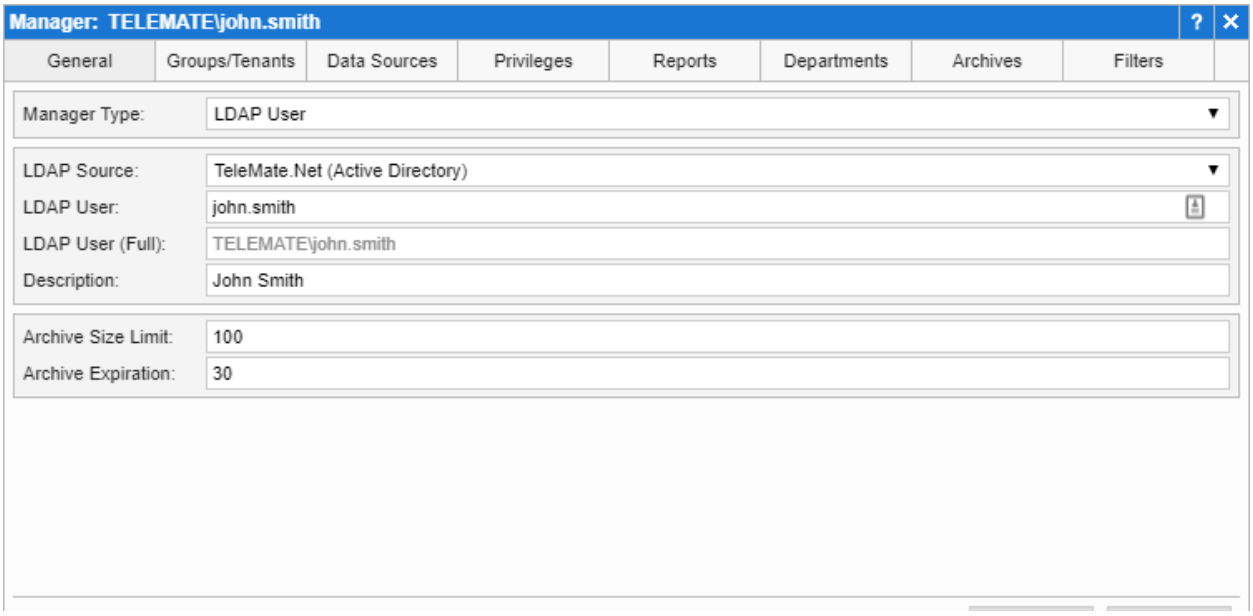

Note: Accounts created in the admin client have special privileges, and their privileges may only be changed in the admin client. However, if you want to change a local user to an LDAP user (or vice versa), you must do it here.

Until you are comfortable with provisioning new manager accounts, we recommend you follow these steps:

1) Click on the Groups/Tenants tab and select one “!Privileges” group and one “!Role” group.

2) At first, save the new manager and then edit it again (to make it easier to see access privileges granted by the selected groups).

3) Click on the Data Sources tab and check any data sources that have traffic you want this manager to see.

4) Click on the Privileges tab and verify the privileges granted by the assigned “!Privileges” group. If you want to make changes, consider whether it would be better to modify the group’s privileges, to create a new group, or to assign additional privileges to this manager directly. **Note:** Privileges granted by a group cannot be removed here. To remove them, you must either modify the group or remove the group from this manager.

5) Click on the Reports tab and verify the report templates granted by the assigned "!Roles" group.

6) Click on the Departments tab and check any companies, divisions, or departments that have traffic you want this manager to see.

7) Most managers only have access to their own personal report archive. However, you may click on the Archives tab if you want users to have access to each other's archives (or to a shared archive you create).

8) Click on the Filters tab and verify the filters/columns granted by the assigned "!Roles" group.

9) Use the "Admin View" feature to temporarily switch to the new manager login and test its privileges. Make sure that it has access to everything it needs (and no more than that). If you need to make changes, switch back to your admin login, make the changes, and test it again.

Important: If a manager doesn't have access to at least one data source and one department, that manager won't be able to see any data. While you can create privilege groups that grant access to sets of these, it makes little sense if each manager is the head of a different department (and should only be allowed to see his or her own).

Configuring a Tenant

A tenant is similar to a privilege group, but instead of granting privileges to accounts that belong to it, it limits privileges. Each tenant effectively defines a "sandbox" for a new set of logins to play in. After creating a tenant, you can use "Tenant View" to switch to that tenant and create administrator accounts within it. No tenant login can see anything outside their sandbox, and tenants cannot overlap each other. To do this, you must:

- Assign each tenant to a branch in your directory. Depending on the size of the tenant, you may wish to assign it to a specific company, division, or department. If you specify a branch that already exists in your directory, the tenant will assume ownership of it. Otherwise, the company, division, or department you specify will be created when you create the tenant.
- Assign each tenant a costing method for each media type. The costing methods will be assigned to the Tenant Level. For more information on costing methods see the section on "The Costing Component".
- Assign any internal addresses (e.g. extensions, phone numbers, logins, email addresses) that belong to a new tenant to its branch in the directory. If you specify an existing branch in the directory, these resources may already be assigned. Otherwise, you may use the "Address Ranges" tab to assign addresses to a new tenant. You may also use a directory import.
  Note: A special $Unassigned directory user will be created beneath each tenant. Address ranges defined for a tenant will automatically transfer matching addresses to this user (unless they are already assigned to a user under that tenant).

- Assign any privileges you want to be available under this tenant. For example, you may only want to allow logins created under this tenant to be able to manage their directory, choose from three report templates to run, and choose from three filters to apply to those reports.

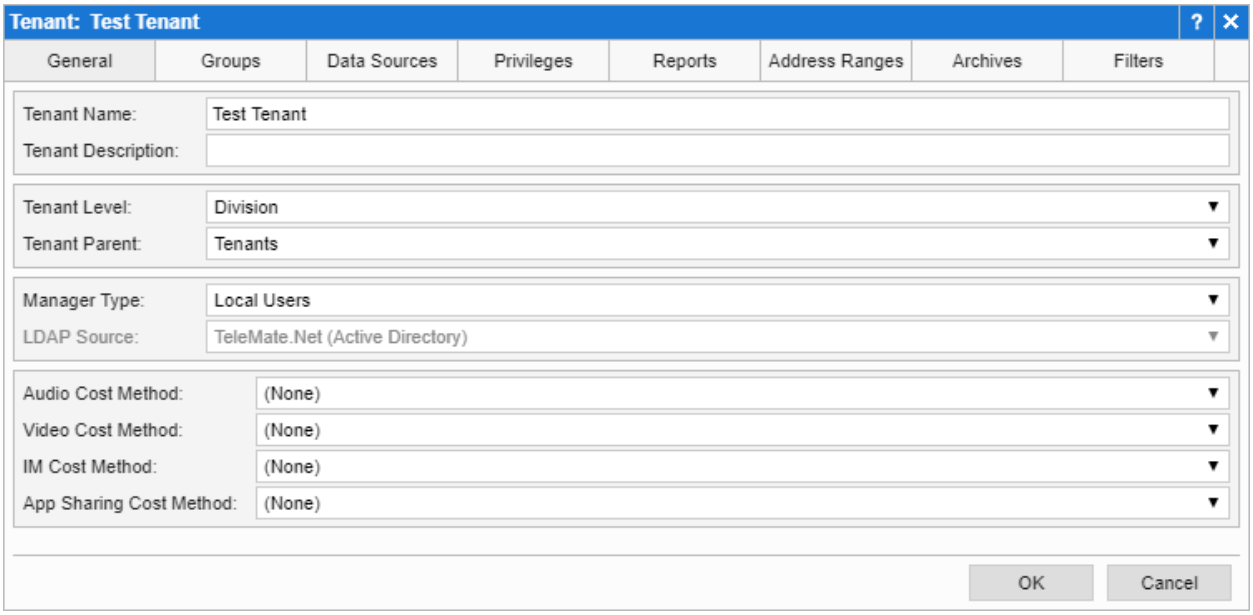

After you create a tenant, you should use the "Tenant View" selection to switch to that tenant. The first thing you should notice is that the new tenant has no privilege groups or logins. When you create the first manager login under that tenant, you should notice that the available privileges are restricted to those you chose to allow the tenant to have. At this point, you may choose to create a tenant "administrator" login to delegate the management of additional logins under this tenant.

Next, we recommend that you check the directory and run some test reports to verify that the tenant can access the data it needs to access (and nothing else). You may allow tenant administrators to manage the directory within their sandbox, but they may never be allowed to move anything in or out of the sandbox, so their addresses must be provisioned correctly.

In addition to having their own branch in the directory and their own set of logins, tenants also have their own distinct set of global settings. If you use "Tenant View" and click on the "Settings" link, you will see that each tenant has their own set of: Titles, Workgroups, Site Locations, Publish Types, Call Plans, Vendors and General Ledger Numbers. If the tenant has access to Expenses, they will also only be able to see Expense Categories and Codes created under that tenant. This helps ensure that no tenant is able to see potentially sensitive information configured by another tenant.

Admin View

While the "Admin View" feature is not technically a part of the Security Manager component, it is similar in that it is only available to accounts with special admin privileges. The admin view allows you to see what any other manager account would see when they log in. This is a great tool for testing changes to security privileges. It is also very useful for being able to see (and modify) any saved reports, monitors, viewpoint views, etc. a manager has set up.

The example screenshot below shows the “admin” manager account (see the login name in the top-right corner) viewing the “telemate” manager account. The component bar shows that this manager only has access to Reports, Viewpoint, and Monitors. Selecting the “Monitors” component makes it clear what monitor privileges he has. The lack of a “Create Monitor” link shows that he lacks the “Create new monitors” privilege. In that case, the only reason he can see the “Monitors” component at all is that an admin account has created a monitor called “Sample Monitor” and shared it with a security group that he belongs to. The lack of “Properties” and “Delete” icons shows that the admin account did not actually assign this monitor to “test1”.

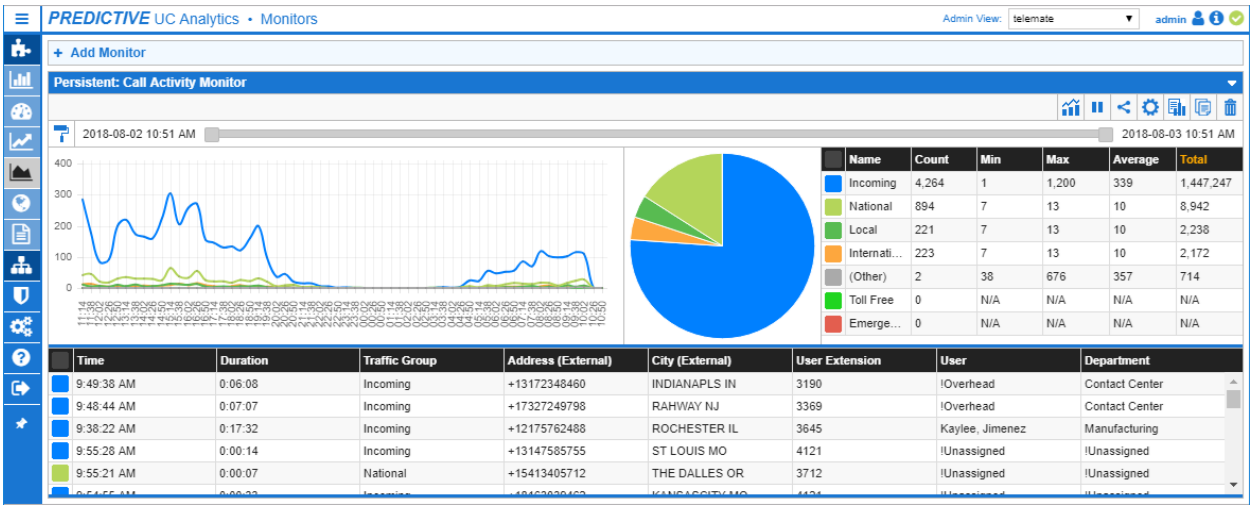

Note: When using admin view, you will be restricted to that manager’s privileges for almost everything. However, you will always be able to reassign (or share) reports, monitors, and views created by other accounts. This is very useful when deleting manager accounts that have automated reports or alarms.

Tenant View

Like the “Admin View” feature, the “Tenant View” feature allows the manager to see what any other manager account within a Tenant sandbox would see when they log in.

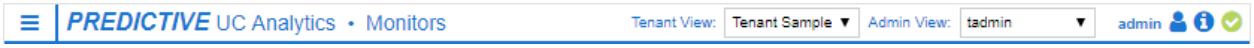

System Settings

Select the “Settings” link from the navigation menu to access system settings. The system settings allow you to configure system-wide options, like mail server settings and custom text strings, as well as, component-specific settings.

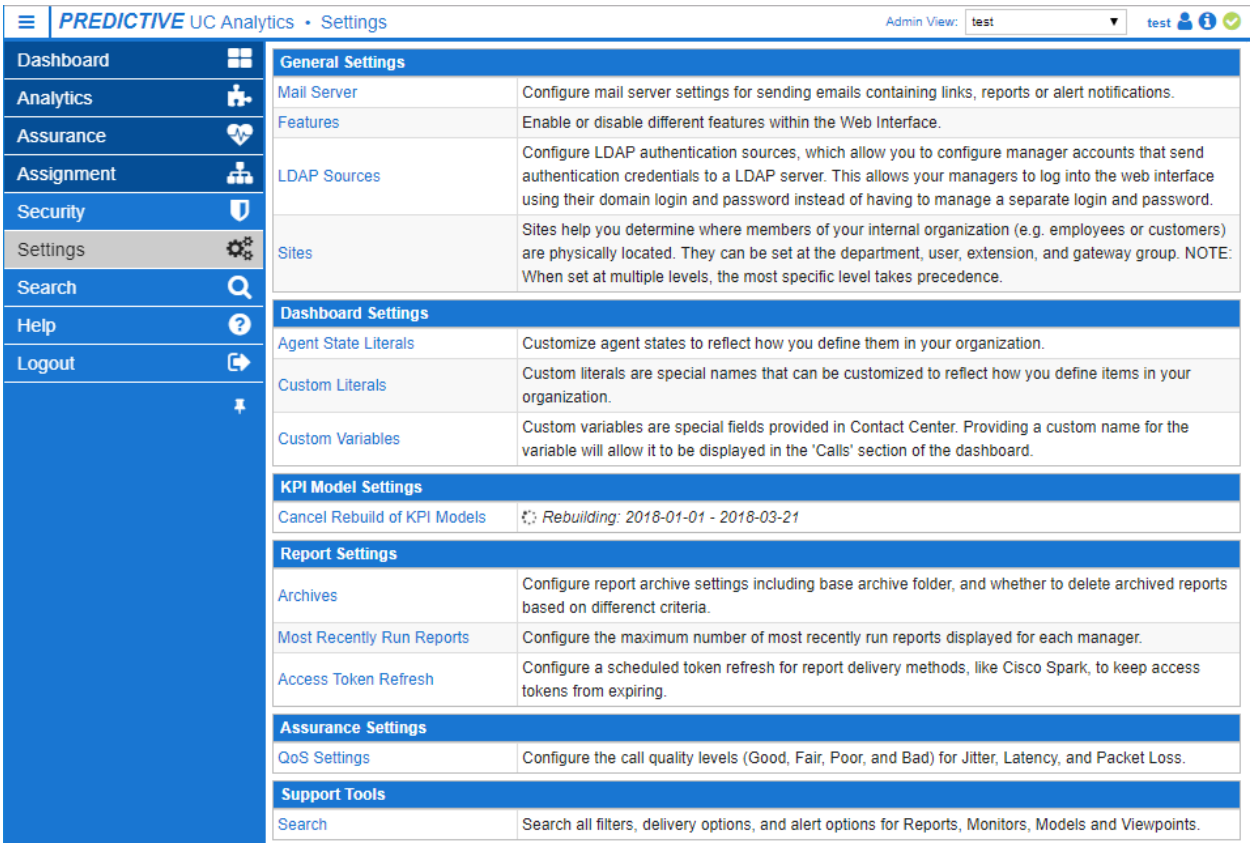

General Settings

The Mail Server Settings dialog allows you to configure and test the SMTP server you wish to use to email reports and alerts.

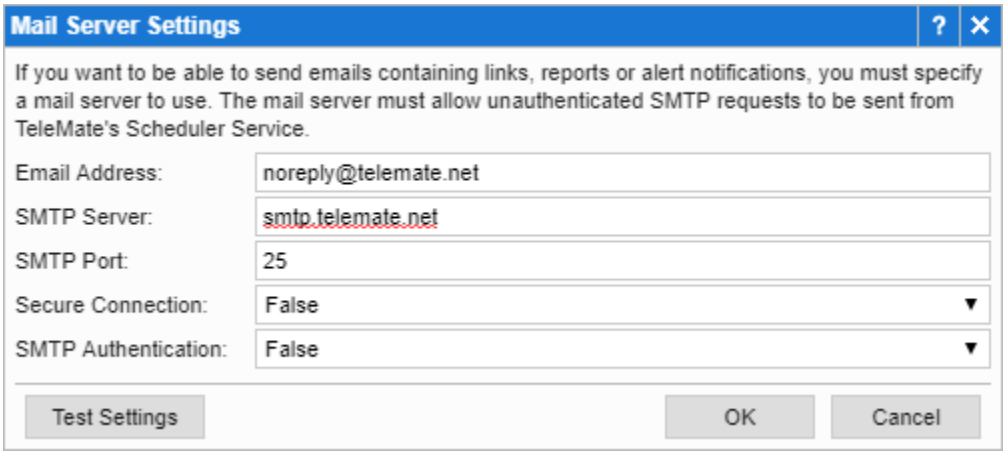

The Features dialog allows you to disable advanced features you do not plan to use in the web interface. This can streamline and simplify navigation.

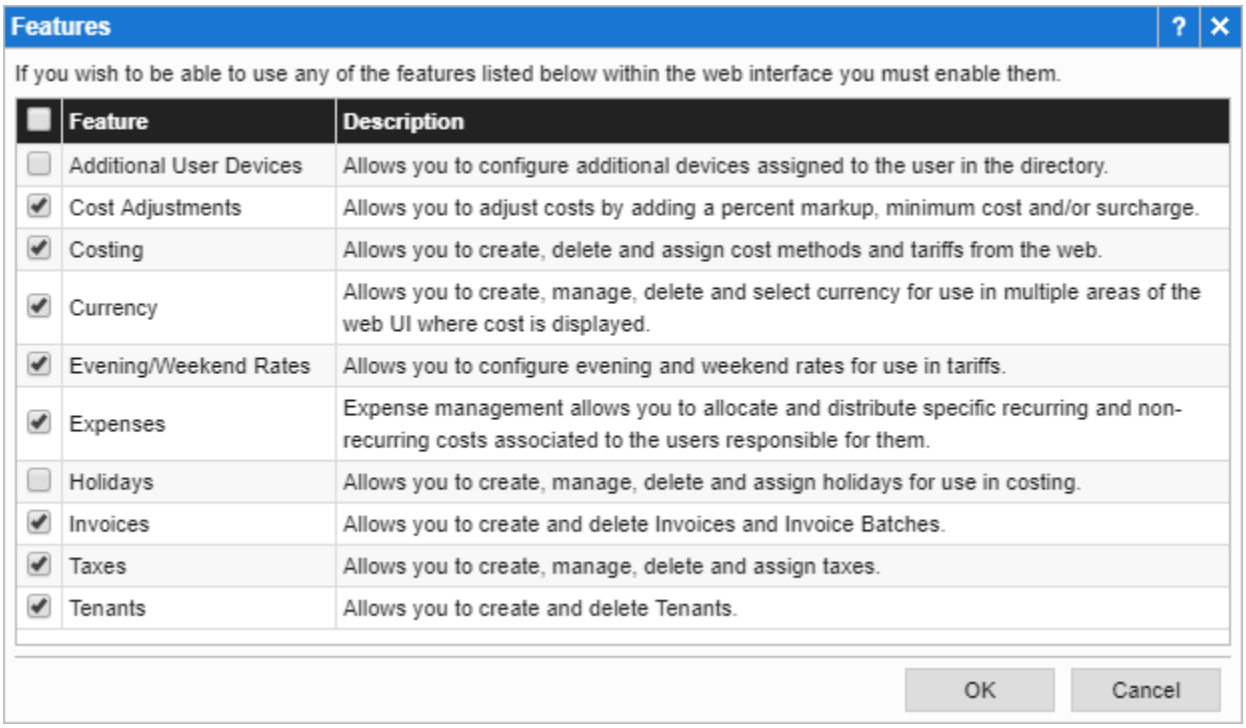

The LDAP Sources dialog allows you to configure LDAP (or Active Directory) authentication sources. You must configure these before you can configure manager accounts that use their domain login and password, or for granting access by LDAP security group or organizational unit.

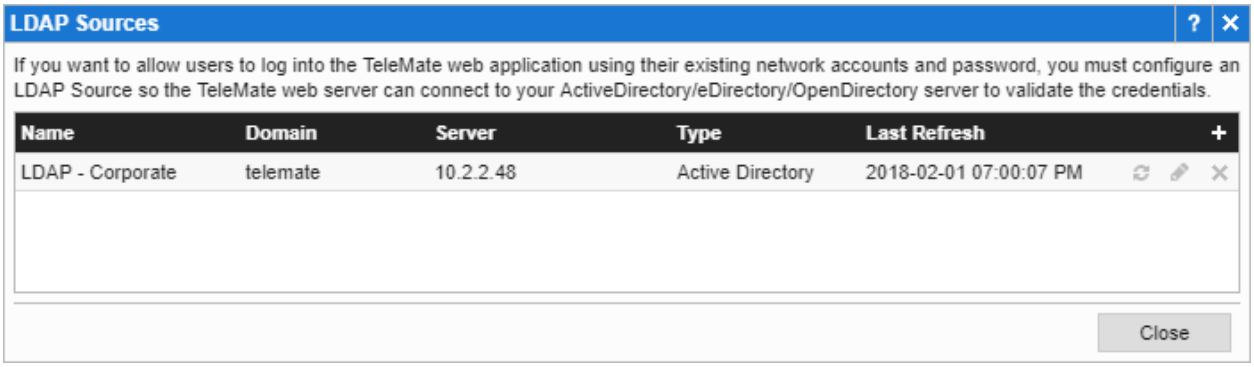

The "Domain" Name: You must enter a unique domain name for each LDAP source. For Active Directory, you should specify the short Windows domain name (e.g. the "telemate" in "telemate\login1"). For eDirectory, specify "o=myorg" to have a "login1" attempt validate "cn=login1,o=myorg".

The "Server" Name: This value accepts hostnames or IP addresses separated by commas. We strongly recommend that you provide at least two, one primary and one secondary. When using host names, we recommend that you use the fully qualified domain name (e.g. dc1.telemate.net instead of dc1).

Performance Issues: If you experience slow authentication, check your LDAP Source configuration to make sure the host names or IP addresses are valid, and when using host names, make sure the DNS

servers resolve them quickly. Failed DNS resolution and failed connection attempts can greatly slow down authentication.

Sample Active Directory LDAP Source:

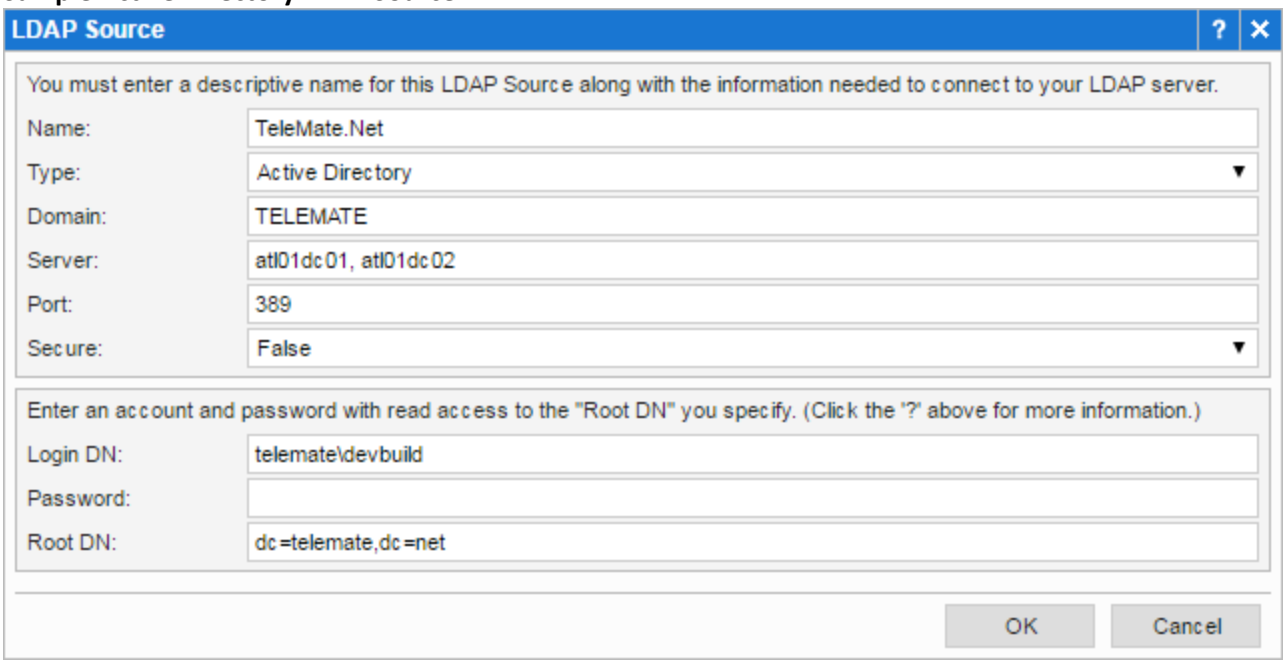

Sample eDirectory LDAP Source:

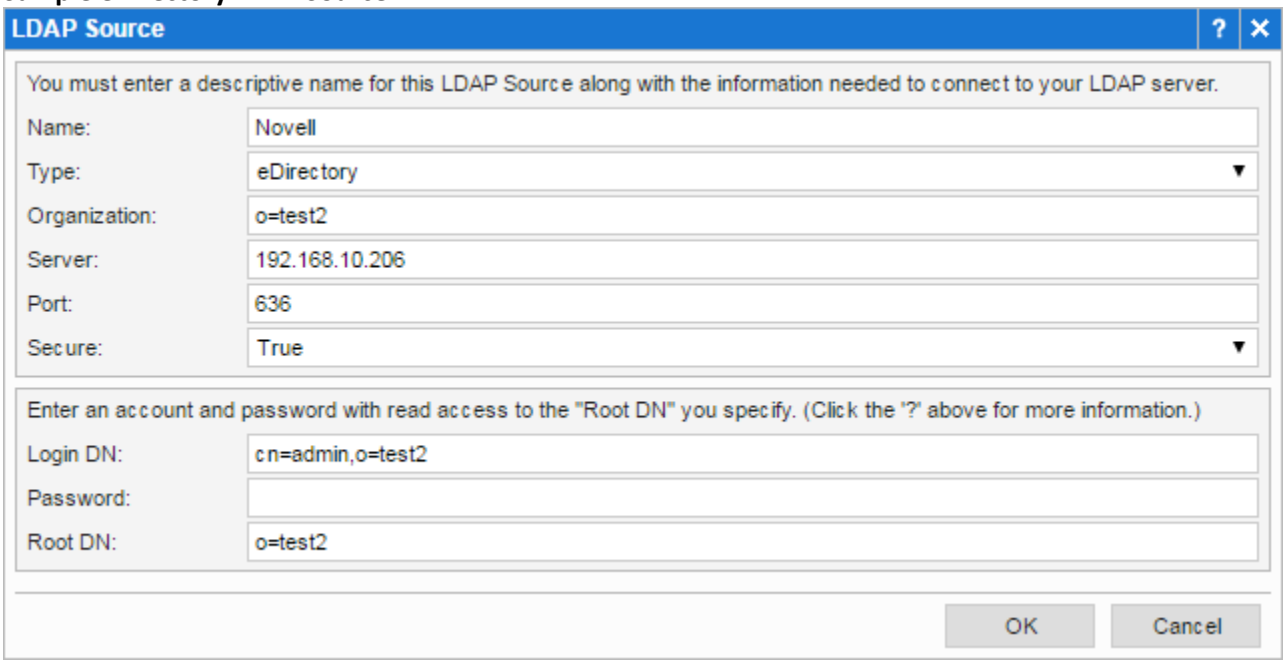

The Site Locations dialog allows you to add/edit/delete internal site locations. These site locations define time zones and physical locations for internal endpoints. Each data source has a default site location for its internal endpoints, but that may be overridden for specific CUCM locations, IP address ranges, gateway groups, gateways, departments, users, or extensions.

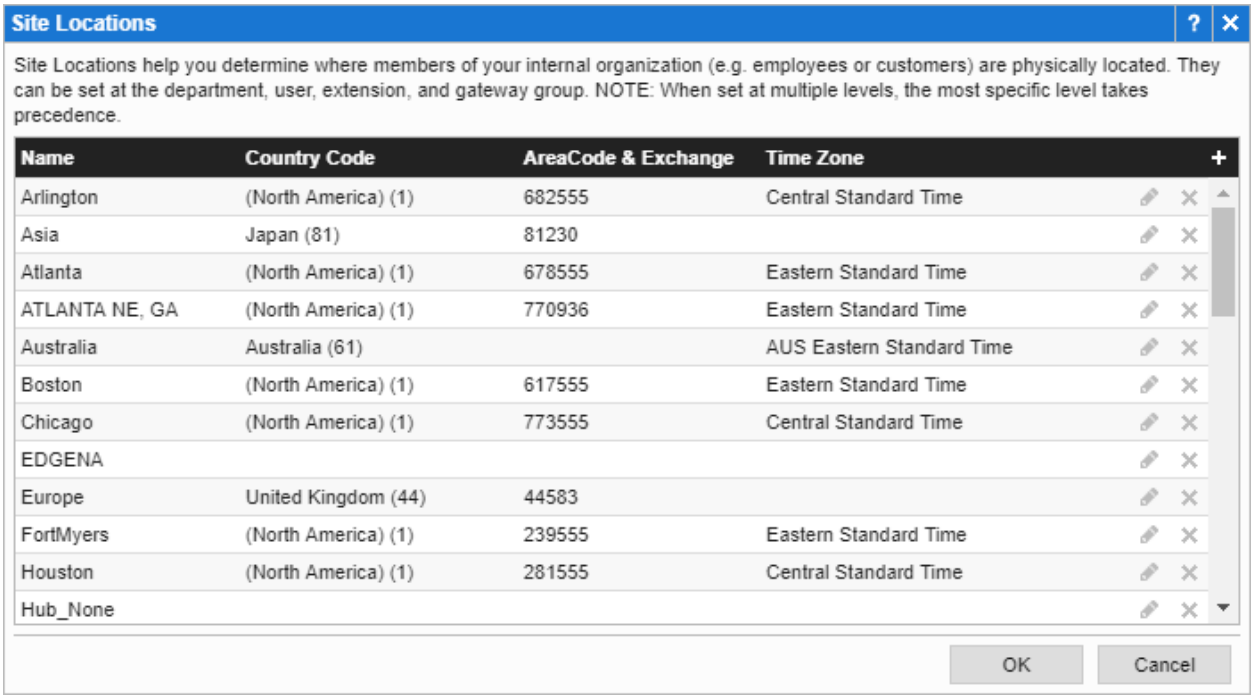

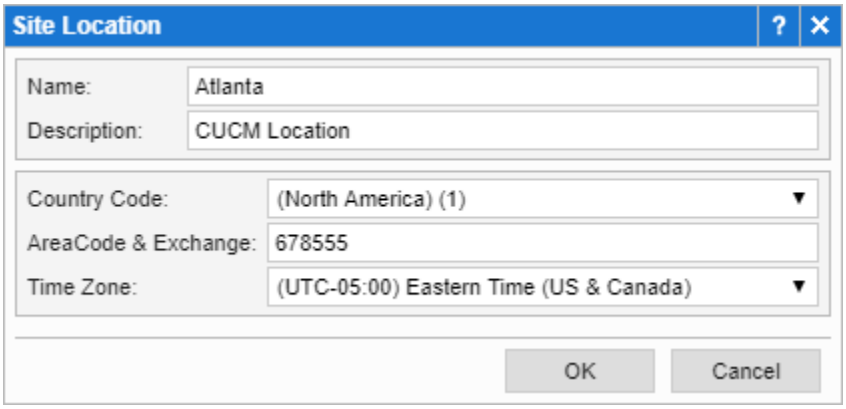

Contact Center Settings

The contact center settings allow you to configure the global options for the Contact Center.

The “Agent State Literals” allows you to configure custom values for Agent States which are the statuses reported by the Contact Center. The “Custom Literals” section allows you to configure special names to reflect how you define items in your organization. You may also specify custom names for the 10 custom variables supplied by Contact Center scripts. You must specify a custom name for the variable before it can be displayed in the 'Calls' section of the contact center dashboard.

KPI Model Settings

The “Rebuild All KPI Models” command rebuilds scores for all KPI models for all managers in a single pass. This allows you to rebuild all of them all at once, which saves time and minimizes the performance impact on the server. If you want to rebuild the models, select the “Rebuild All KPI Models” specify how far back you want to rebuild the KPI model scores.

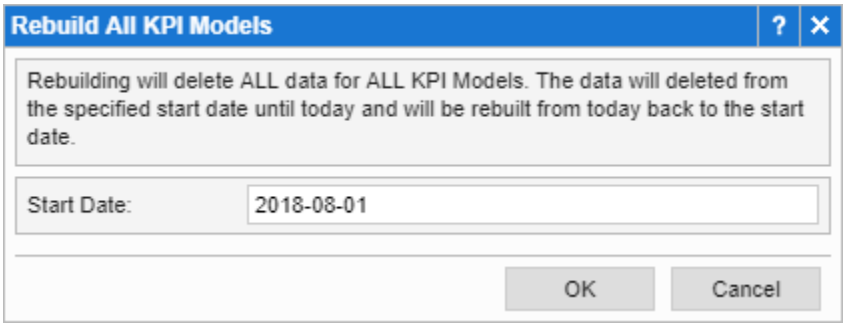

Report Settings

The report archive settings allow you to configure the global report archive options. Each web login gets its own archive – which corresponds to a directory beneath the specified root folder on the Predictive UC Analytics server – where completed reports are saved for a number of days (which you specify when you provision a user in the Security Manager). There are also options to purge archived reports. When "Delete Expired" is enabled, archived reports will be deleted before new reports are run. When "Free Space" is enabled, the oldest archived reports will be deleted to free up space for new archived reports if the archive size limit has been reached.

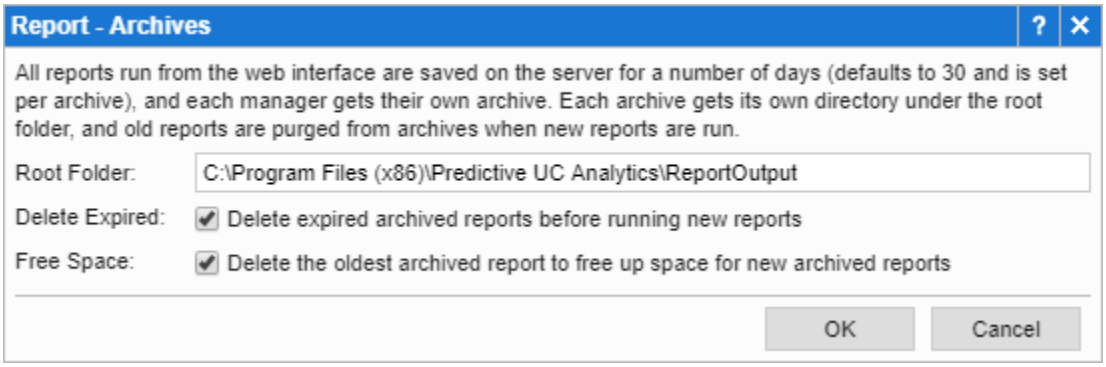

The web interface also keeps track of each user's most recently run reports, including the filter and destination options selected for those reports. You can change the number of most recently run reports tracked by Predictive UC Analytics here.

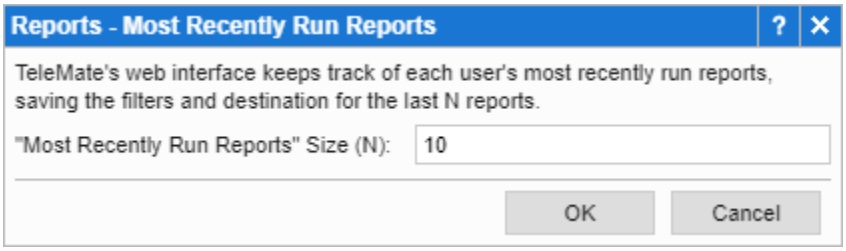

The Access Token Refresh settings allow you to schedule a nightly refresh of OAuth tokens that expire over time if you do not refresh them. This is only needed if you deliver reports or alarms to a web service that has a refresh token, like Cisco Spark/Teams.

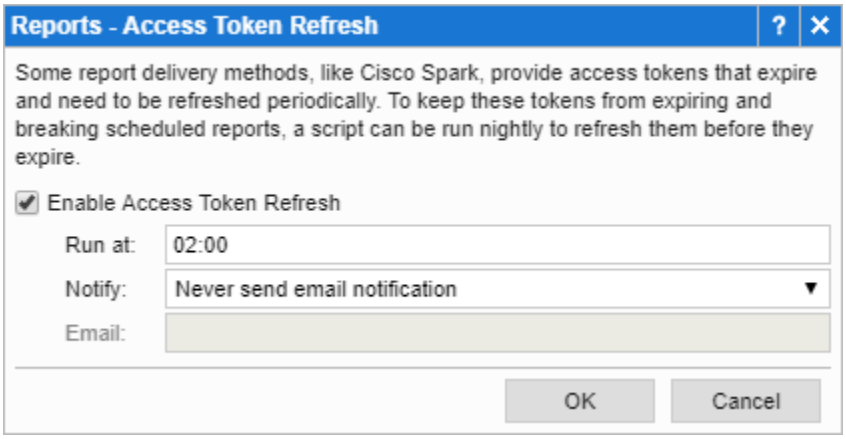

Cost Settings

The Cost Settings section contains options for configuring currencies, taxes, and holidays. They will only show up if you have enabled them in the "Features" section.

Predictive UC Analytics comes with some currencies pre-configured. To create a currency, click the "+" button in the top-right corner of the Currency list. To edit an existing currency, click on its name. To delete a currency, click on the "x" button to the right of it.

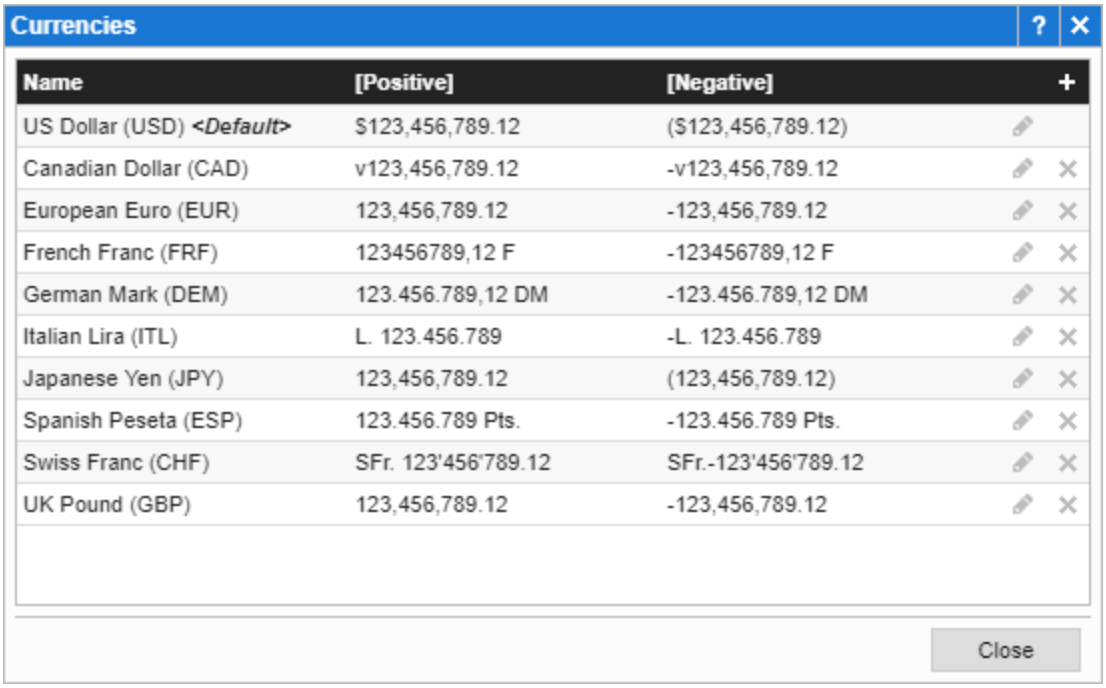

Note: You may specify a conversion rate for every currency except the default. When you run a report in a different currency, its conversion rate will be applied to convert all costs from the default currency to the selected currency.

Assurance Settings

The Assurance Settings section will only show up if you licensed the assurance feature set. It allows you to change the good/fair/poor/bad thresholds for Jitter, Latency, and Packet Loss for the assurance displays and alerts.

Support Tools

This section provides a search to help administrators more quickly find things configured by other user accounts. A good example would be that various user accounts have configured reports or alarms to go to jim.smith@mycompany.com, Jim is leaving the company, and you need to track down everything sending emails to him.

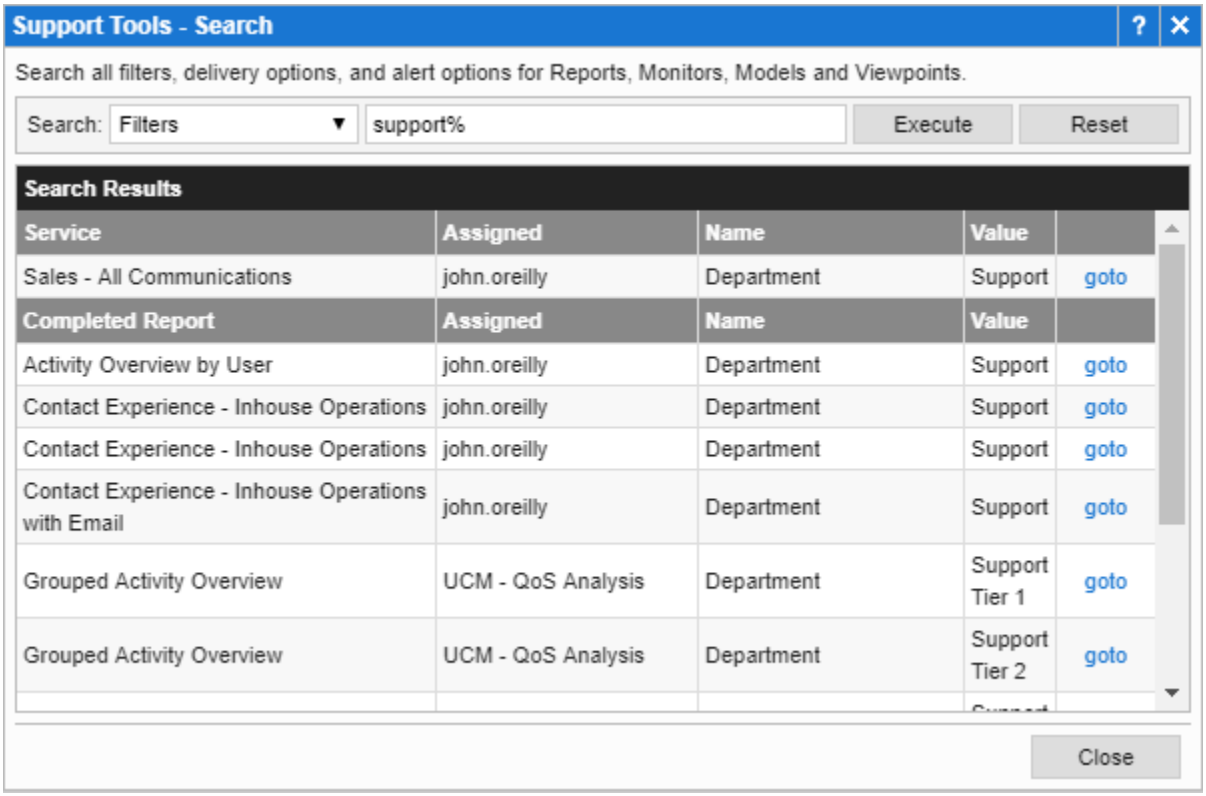

Help

Click to open a web-based version of this First Steps Guide.

Search

The navigation menu contains a link for “quick search”. Clicking the link will give you the option for a date selection (i.e. today, yesterday, last week, etc.) and an address value. This will launch a detail records search with filters set from your choices. Refer to the “View Details” in the “Common Features” section of this document for a more detailed explanation of the detail view.

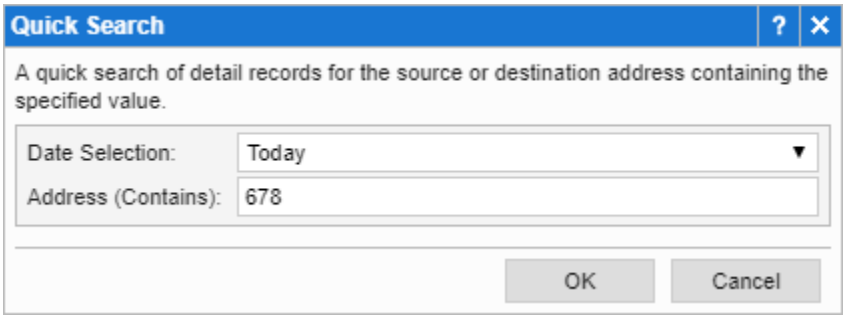

Logout

Click to log out and/or change to a different login.