Oracle® Communications Fraud Monitor

User's Guide

Release 3.3.90

E60870-01

April 2015

Oracle Communications Fraud Monitor User's Guide, Release 3.3.90

E60870-01

Contents

3 Installing Fraud Monitor

4 Configuring Fraud Monitor

Glossary

Preface

This guide describes how to install, configure, and use Oracle Communications Fraud Monitor.

The Oracle Communications Session Monitor product family includes the following products:

- Operations Monitor
- Enterprise Operations Monitor
- Fraud Monitor
- Control Plane Monitor

Audience

This guide is intended for system administrators, network administrators, and network operations team who use Oracle Communications Fraud Monitor to monitor calls and detect fraud.

Downloading Oracle Communications Documentation

Oracle Communications Session Monitor documentation and additional Oracle documentation is available from the Oracle Help Center Web Site:

`http://docs.oracle.com`

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at `http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc`.

Access to Oracle Support

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit `http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info` or visit `http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs` if you are hearing impaired.

Document Revision History

The following table lists the revision history for this document:

Version	Date	Description
E60870-01	April 2015	Documentation updates for Session Monitor Release 3.3.90. ▪ Made minor formatting and text changes. ▪ Made graphic screen capture updates. ▪ Added "Downloading Oracle Communications Documentation".

1

Overview of Fraud Monitor

This chapter provides an overview of Oracle Communications Fraud Monitor.

About Fraud Monitor

The Session Monitor architecture consists of the Probe layer, Mediation Engine layer, and the Aggregation Engine layer (see the discussion about Session Monitor architecture in *Session Monitor Installation Guide* for information about the functions performed in each layer).

Fraud Monitor runs on the Aggregation Engine (AE) machine, but relies on the data provided by the Mediation Engines (MEs) to detect fraud. For each established call, the ME that has seen and correlated the call sends to the AE a message when the call is established, then one message every few minutes and finally a message at the end of the call. This allows Fraud Monitor to be aware of the real-time state of all the calls in the system and use this state to apply the different behavioral rules.

Logging In to Fraud Monitor

The Login page allows you to access Fraud Monitor. Enter your user name and password into the indicated fields, then click **Sign in** to proceed to the application.

Figure 1–1 shows the Fraud Monitor Login page.

In the case your user name or password are incorrect, a warning appears below the **Sign in** button and you'll have the opportunity to retry.

Figure 1–1 Fraud Monitor Login Page

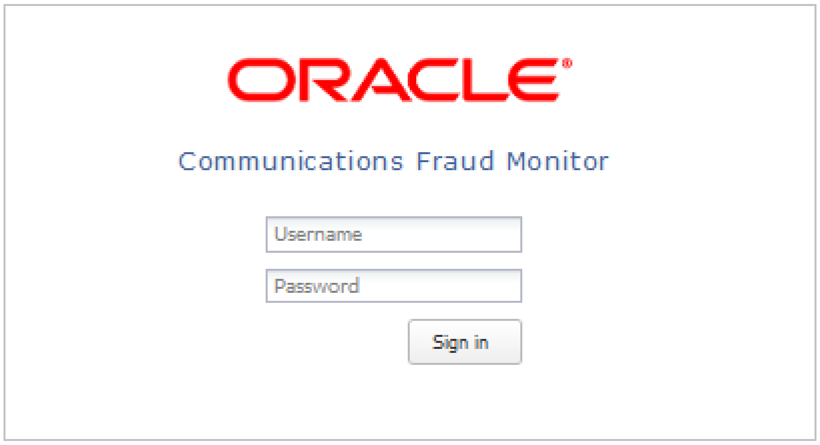

About Using the Fraud Monitor User Interface

The Fraud Monitor user interface has several recurring elements. At the very top, you have the dark bar which gives access to general settings like system information and logout. Below the dark bar on the right is the navigation menu that lets you navigate to the main pages: Overview page, Incidents page, Details page, and the Settings page.

The Overview page shows the current fraud status at a glance. It has big warning indicators as well as a table of the recent potential fraud incidents. From there you can further analyze the latest issues.

The Incidents page displays the full table of recent incidents, with related rules, and from there you can pick an incident for further investigation.

The Details page shows a single user and his incidents. This page contains a short history of the potential fraud that might have occurred.

In the Settings page, you can tweak the rules to improve the fraud pattern matching, manage the users, and configure the ME data sources.

Overview Page

The Overview page is the landing page that appears automatically after you log in. The Overview page displays the status information on processed calls, incidents, users, as well as general information about how Fraud Monitor works. You use this page to check for recent incidents and then navigate to other pages to further investigate the user.

Figure 1–2 shows the Overview page.

Figure 1–2 Overview Page

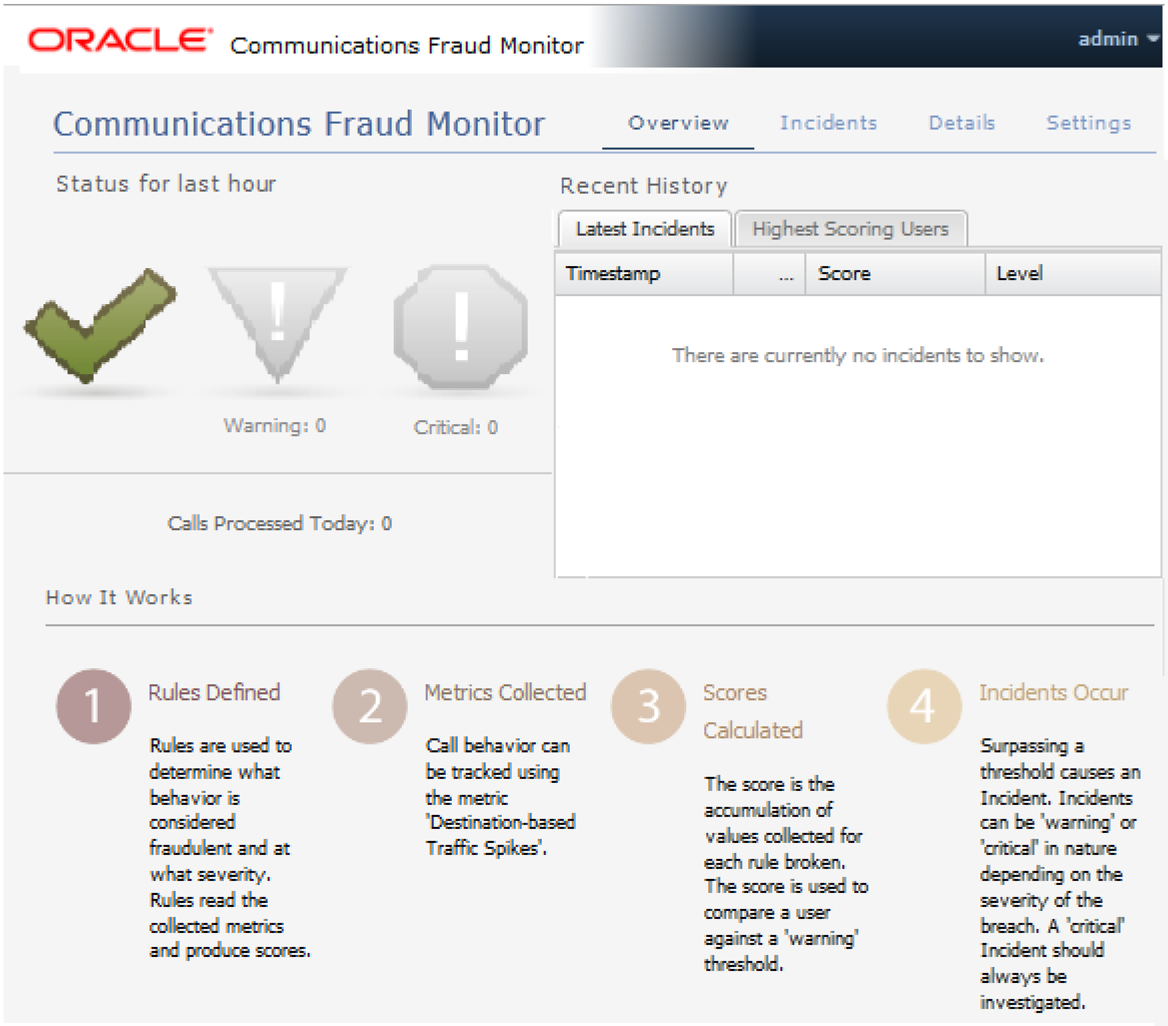

There are currently no incidents to show.

Calls Processed Today: 0

How It Works

1 Rules Defined

Rules are used to determine what behavior is considered fraudulent and at what severity. Rules read the collected metrics and produce scores.

2 Metrics Collected

Call behavior can be tracked using the metric 'Destination-based Traffic Spikes'.

3 Scores Calculated

The score is the accumulation of values collected for each rule broken. The score is used to compare a user against a 'warning' threshold.

4 Incidents Occur

Surpassing a threshold causes an Incident. Incidents can be 'warning' or 'critical' in nature depending on the severity of the breach. A 'critical' Incident should always be investigated.

Viewing the Status for the Last Hour

The Status for the Last Hour section displays the total number of calls processed for the day and the number of incidents detected in the last hour. The incident counts are accompanied by large icons for quickly establishing overall status. If no incidents are detected, a large green tick image is displayed. If any warning incidents are detected, a large orange warning sign image is displayed. If any critical incidents are detected, a large red stop sign image is displayed. It's possible for warning incidents and critical incidents to be detected in the same time frame. When an icon is not being displayed, it remains faded grey in the background.

> **Note:** If the calls processed counter is not increasing, you may not have configured your Mediation Engine correctly. Navigate to the Mediation Engine Configuration tab located on the Settings page.

Viewing the Latest Incidents

On the right-hand side of the page, the Latest Incidents section shows a small list of the most recently detected incidents. A more extensive list can be found on the

Incidents page. Double-clicking on an incident will take you to the Details page for that particular offending user.

Viewing the Highest Scoring Users

On the right-hand side of the page, the Highest Scoring Users section shows a small list of the users with the highest scores. Double-clicking on a user will take you to the Details page for that particular offending user.

How it Works

On the bottom of the page, the How it Works section outlines how Fraud Monitor works in four steps. You might refer back to this anytime you need to be reminded how the components of Fraud Monitor inter-relate.

Incidents Page

The Incidents page lists all the calls which have triggered incidents. This includes warning and critical incidents. See the Settings page to configure incident levels. As incidents are triggered they are added to the Incidents page in real-time.

The latest incidents are on top and incidents are not cleared; they stay forever unless you delete them yourself (see below). The **Suspect** column is the callers number or IP address.

When a line is selected, the panel on the right shows the details of this incident. It displays the caller as well as which rule or rules caused the incident to be triggered.

If you double-click on an incident, you'll go to the Details page for that user. Selecting a row and clicking **Go to User Details** button does the same.

If a user first triggers a *warning* incident and later upgrades that to a *critical* incident, both will be listed.

You can delete an incident by selecting it and clicking **Delete**. This will remove the incident from the list. When user causes multiple incidents of the same level (warning or critical) within 24 hours, a new incident is not triggered. Deleting a row from the incident list will *not* reset that timer. Deleting an incident is useful when, after investigation, you conclude that an incident is not fraudulent.

Times are in the local time zone.

Details Page

The Details page shows information on a particular user. Also, the whitelist for countries of this user can be maintained on this page and all records associated with this user may be deleted.

Figure 1–3 shows the Details page.

The Details page can be reached via the top menu or from the Incidents page.

Typically, in this view one will search a user, if none is selected yet, and then consider the metrics to determine if an incident is justified. In case no user has been chosen, one must be selected using the search field. Then the scores, the calls, and the geographical data for this user are displayed. Each of these topics is explained in the following sections.

Figure 1–3 Details Page

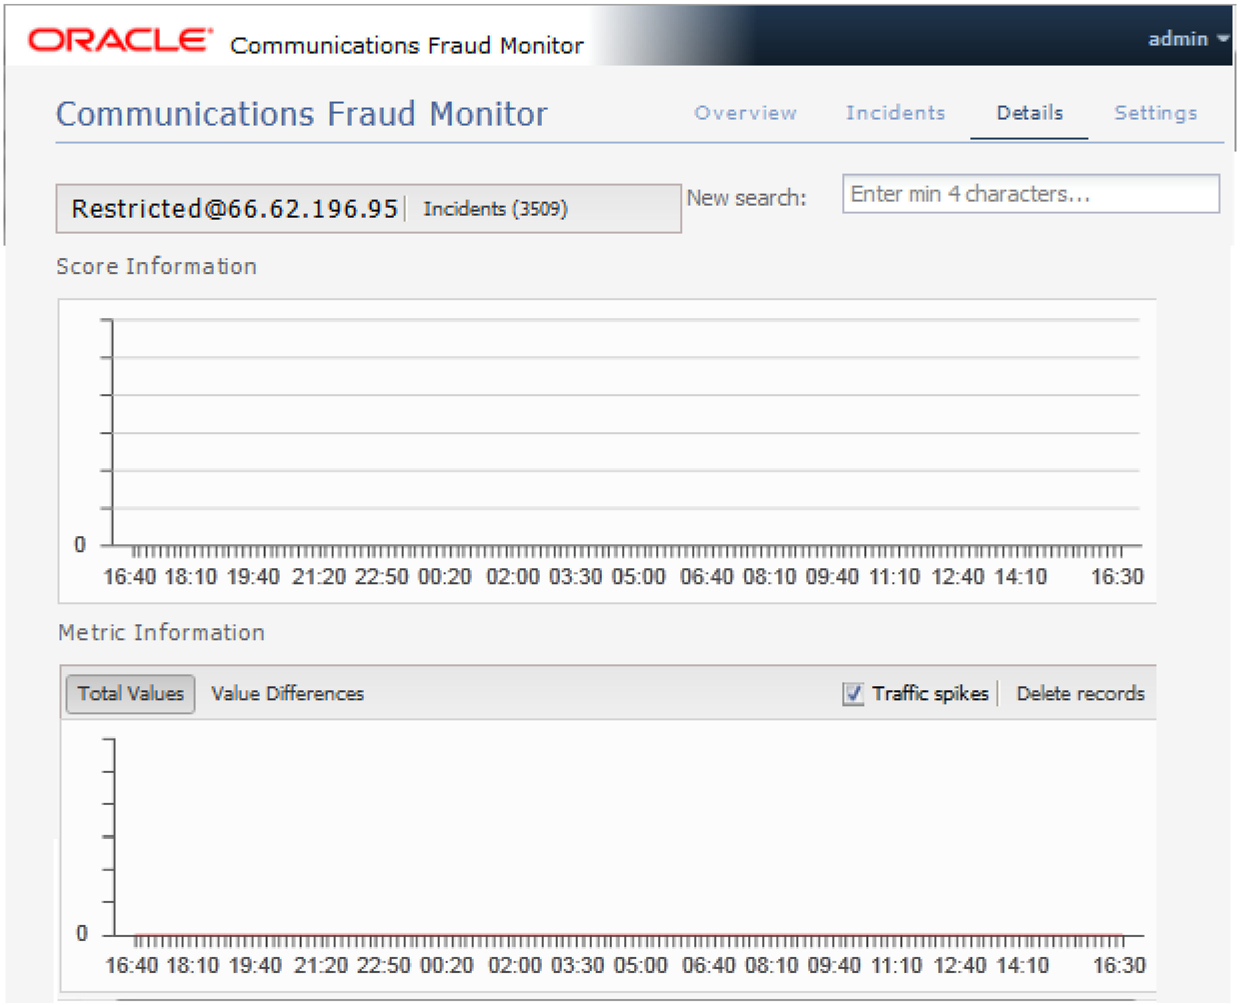

Performing a User Search

The **New search** field allows you to select a user for display by IP or phone number. After four characters matches are shown. Select one of the proposed matches, press return or click **Search** to display the user.

Deleting User-Specific Records

Click **Delete records** to remove all information regarding the user. This includes metric values, scores and incident related information.

Viewing Score Information

The Score Information diagram shows scoring information of all incidents for the last 24 hours, going back from the current time. It is not possible to go further back than 24 hours. For each incident measuring interval, a bar displays the score reached.

Viewing Metric Information

The Metric Information diagram show metrics of selected rules. For each ten minute interval, the values of the last 24 hours are shown in comparison to the average over the last two weeks.

The y-axis displays the number of minutes or calls, while the x-axis specifies the intervals. A red line displays the averages, while the bars show the current data.

By using the check boxes in the top right corner, you can choose what values to display.

- The traffic spikes

The alternative display modes in the top left corner, **Total Values** and **Value Differences**, toggle between displaying the current values as absolute values and as the difference to the average.

User Menu

The User menu is located on the top right corner of the page on the header bar. A drop-down menu appears when you click on your user name.

Figure 1–4 shows the User menu.

Figure 1–4 User Menu

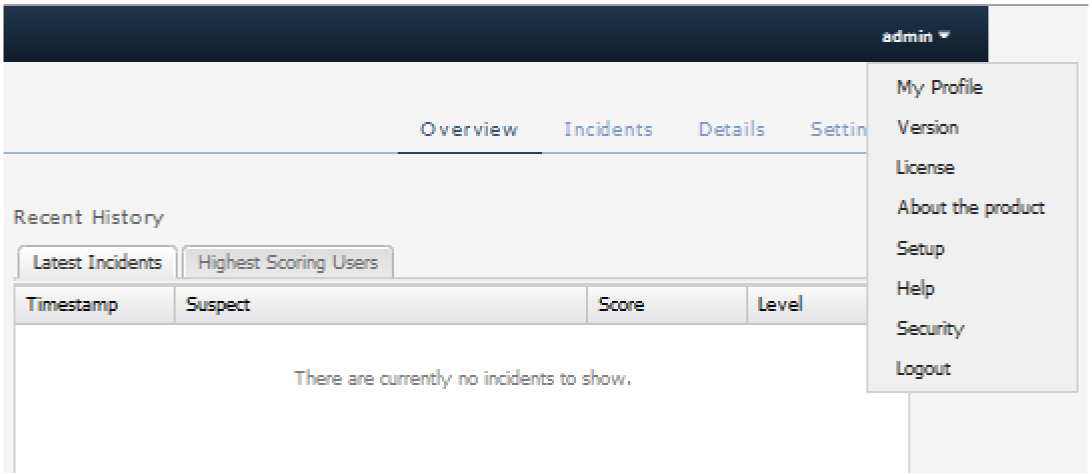

Timestamp	Suspect	Score	Level

There are currently no incidents to show.

Editing the User Profile

You can edit your own profile details by selecting **My Profile** in the User menu. A dialog box appears giving you the option to change your user name, email and password. Fill out the new values for the details you would like to change and click **Finish** to save the changes. Click **Cancel** to exit the window without making changes.

Viewing System Information

You can view the current system information by selecting **System Info** in the User menu.

Viewing License Information

You can view the product license terms and conditions by selecting **License** in the User menu.

Logging Out

You can logout of Fraud Monitor by selecting **Logout** in the User menu. This brings you back to the Login page.

Settings Page

The Settings page lets you configure the rules, manage the users, adapt the notifications and whitelist, as well as perform the Mediation Engine configuration, which sets up the data source needed for a successful operation of Fraud Monitor.

Among the settings, rules is the most important setting. In the Rules section, you can enable and configure patterns that are used to detect fraud and trigger incidents.

2

About Detecting Fraud

This chapter describes some of the common fraud scenarios and fraud detection rules.

How Fraud Monitor Detects Fraud

The Session Monitor probes and Oracle Communications Session Border Controllers with the probe software enabled send monitoring information to the Mediation Engines. The Mediation Engine (ME) then feeds call state information to Fraud Monitor. Fraud Monitor analyzes every incoming call and applies various rules to them. If a rule identifies a call as being suspicious, points can be added to a per-call and per-customer score card. A single rule or a combination of multiple rules may add enough points to trigger a fraud alert. Alerts are on two levels: *warning* and *critical*. Warning level alerts should be investigated while critical level alerts can be considered proven fraud incidents, for example, due to hits on the blacklist which contains known incidents.

Note: A user (also known as a subscriber to distinguish between users of the system and participants in monitored calls) is identified either by his IP address or by the local part of his From SIP URI. If the SIP URI is sip:2125551234@example.com, then the user is shown as 2125551234 in the GUI.

About Fraud Scenarios

The following sections describe some of the common fraud scenarios.

PBX Fraud

Scenario

PBX systems sit on the border between an internal and external network. Users on the internal side (for example, inside an enterprise) may conduct outbound calls and also receive calls. When looking from the outside (visible to Session Monitor or an SBC), the PBX receives calls for a limited set of numbers (for example, the number range of the enterprise) and makes phone calls to almost any number. Depending on the customer, the outbound calls may be directed to a restricted area (for example, mostly local calls).

Detection Method

Whenever possible, multiple metrics should be used to identify fraud. Calls bound to the PBX (as seen from Session Monitor or an SBC) are not subject to fraud in this context but may be part of a fraud scheme (for example, when representing the inbound leg of a forwarded call). In fact, an attacker might bypass the Session Monitor or the SBC monitoring points so that inbound calls are not visible. Nevertheless, fraud may be detected by noticing changes in the outbound call characteristics such as time or destination. Business customers may conduct most of their calls during their business hours (for example, 9AM-5PM local time) while fraudulent users might abuse the system to make long-distance calls outside business hours. Or, an attacker may disable Number Presentation (CLIP) to hide the source of the call. Here too, fraud might be detected by observing a change in the daily distribution of calls as well as the geographical restrictions.

International Revenue Share Fraud

International Revenue Share Fraud (IRSF), Domestic Revenue-Share Fraud (DRSF), and Premium Rate Fraud are closely linked. This scenario is also named Artificial Inflation of Traffic (AIT) by the GSMA. The detection methods for all three scenarios are similar and all covered in this section.

Scenario

An attacker operates a premium number with a revenue share provider in a foreign country. For each call or call minute conducted to this number the attacker receives part of the revenue. The attacker’s goal is to inflate the traffic to this number to increase his revenue. The services provided via this number may range from random announcements to call-through services. To redirect traffic to his number, the attacker may place calls (no connect, just creating a missed call entry) with a spoofed number to victims leading them to call him back. In a more sophisticated scenario, the attacker introduces his premium number into his victims' communication as a call-through service. He may modify VoIP endpoints (PBXes, VoIP enabled routers, and so on.) to carry his number as prefix. A Bluetooth-based attack has been used to replace phone numbers in mobile phones and prefix them with a premium number. This not only increases the revenue for the attacker, but (as above) also allows the attacker to eavesdrop on the phone calls. The most common approach to inflate traffic to the fraudsters phone number is to break into PBX or voicemail systems and call his own number knowing that this costs the PBX or voicemail operator significant amounts of money.

Typically the fraudster can collect revenue from the premium number quicker (for example, each day or each week) than the billing cycle on the originating side (for example, once a month). This allows the fraudster to extract money from the system before the bill hits him on the originating side if he decides to increase the traffic on his own.

Detection Method

The Amount of Traffic to the fraudulent number(s) increases. A hit on the Blacklist may also be triggered.

About Fraud Detection Rules

The metrics described in this section are based on the fraud scenarios above. Multiple rules may be combined to detect a single fraud scenario. Throughout this section the term subscriber relates to either a single IP address or a single phone number.

Traffic Profile

Once a few days of call data for a single subscriber is available a graph with the time of the day on the x-axis may be generated. The y-axis shows the number of calls or call minutes conducted. Once a fraud attack happens the shape of the graph will change.

Blacklist and Whitelist Phone Numbers

A list of specifically allowed and disallowed phone numbers or phone number prefixes can be used to identify fraudulent calls. In case international calls are disallowed by a company policy, an international call may be an indicator of fraud. The customer may add individual numbers or phone number prefixes to a customer-specific blacklist.

A second instance of the blacklist covers numbers and prefixes previously identified as involved in fraudulent behavior. Depending on whether the system observed an exact number hit or a prefix match the scores assigned may differ. A prefix match on its own may not directly trigger a critical alarm but when combined with other metrics (for example, the amount of traffic to the suspicious number) it may generate a critical alarm.

Destination-Based Traffic Spikes

Fraud Monitor can raise an incident if a given destination user receives unusually high traffic, as in an IRSF scenario. For each call, Fraud Monitor monitors the total number of minutes that the destination user has received and compares it to its historical average. If a configurable threshold is exceeded, both the source and destination users accumulate points. This rule can be used to identify possible candidates for blacklisting destination numbers.

3

Installing Fraud Monitor

This chapter describes how to install Oracle Communications Fraud Monitor.

Overview of Fraud Monitor Installation

The following is an outline of the installation procedure for Fraud Monitor:

1. Set up a Mediation Engine machine with the Operations Monitor installed. See *Operations Monitor User's Guide* for configuring the platform devices and correlation.
2. Install Fraud Monitor on an Aggregation Engine machine.
3. Log in to Fraud Monitor and add the Mediation Engine to the list of MEs to which this instance connects.
4. Adjust the default rules configuration, if needed.
5. Configure email notifications and wait for the first incidents to be created.

Figure 3–1 shows an example of Mediation Engine configuration.

Figure 3–1 Mediation Engine Configuration for Fraud Monitor

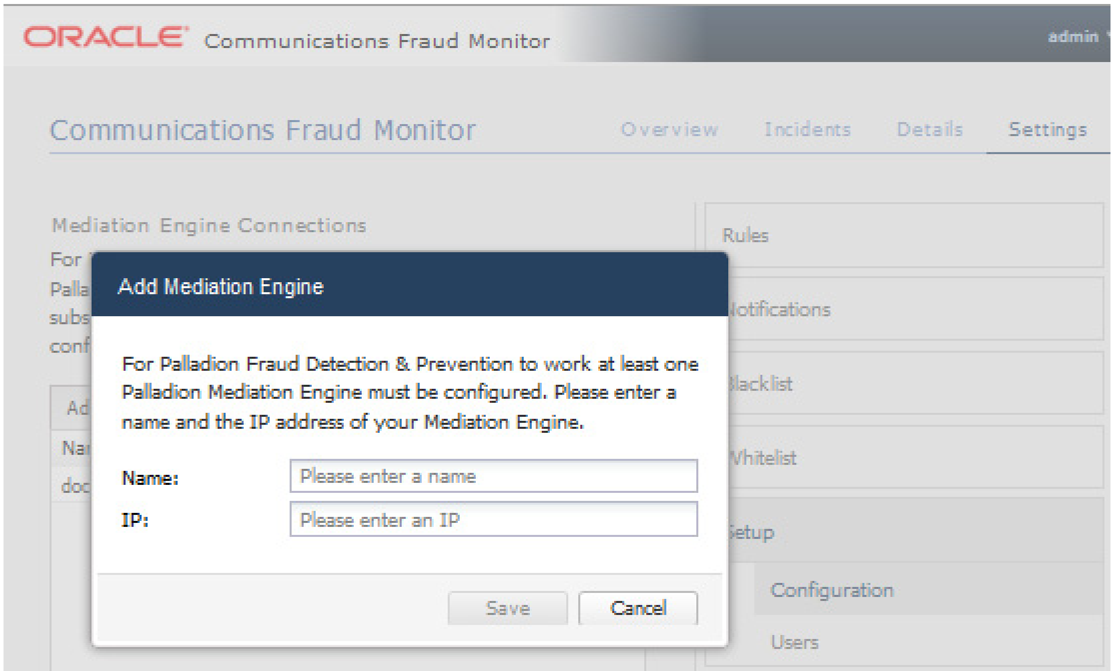

For Palladion Fraud Detection & Prevention to work at least one Palladion Mediation Engine must be configured. Please enter a name and the IP address of your Mediation Engine.

Name: Please enter a name

IP: Please enter an IP

Hardware Requirements

The following minimum requirements must be met to install Fraud Monitor:

- 2.6 GHz Intel Xeon processor, 64-bit with 8 processing threads
- 8 GB RAM
- 70 GB storage on a hardware RAID controller
- 2 Ethernet ports

> **Note:** For production use, Oracle recommends a more thorough sizing exercise completed with your Oracle sales engineer. Higher performance hardware may be required, for example, in cases with:
>
> - High levels of monitored traffic
> - High numbers of concurrent users
> - High volumes of historical information

4

Configuring Fraud Monitor

This chapter describes how to configure Oracle Communications Fraud Monitor.

About Configuring Fraud Detection Rules

The Settings page of the Fraud Monitor user interface lets you configure the rules, manage the users, adapt the notifications and whitelists, as well as perform the Mediation Engine configuration which sets up the data source needed for a successful operation of Fraud Monitor.

The Rules section enables you to configure the patterns that are used to detect fraud and trigger incidents. If the current settings do not trigger any incidents, you may need to change the patterns or raise the points.

> **Note:** Go to the Platform Setup Application and refer to *Session Monitor Installation Guide* for settings (for example, network interfaces, DNS, or SMTP) that affect the server running Fraud Monitor.

Configuring Rules

Fraud Monitor uses configurable rules to find call patterns which are considered fraudulent and classify the severity of the incident with a points system. On the Rules section, you can decide which rules are used, configure them, and restrict their use.

The navigation bar on the right-hand side of the page lists the preconfigured fixed set of rules you can use.

Figure 4–1 shows the navigation bar on the Settings page.

Figure 4–1 Navigation Bar on the Settings Page

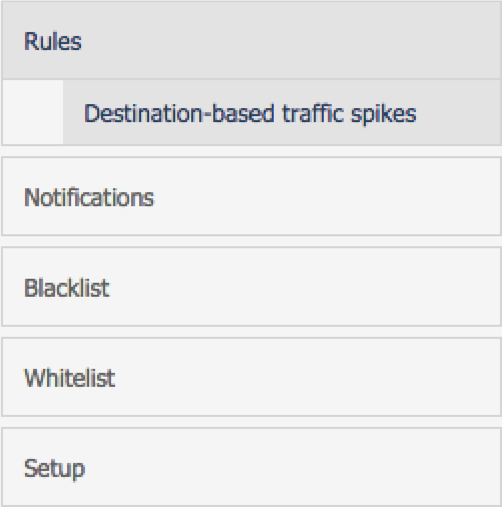

Clicking on a rule opens up its configuration panel in the left column. Use the check boxes next to the rule name in the left column to enable and disable the rule.

Figure 4–2 shows an example of rules configuration.

Every configuration panel has **Add** and **Delete** buttons, which you can use to configure that specific rule. A brief help text is shown above the panel to aid you in the configuration process.

Every rule is assigned a weight. The default is **1.00**. The rule weight can be used to make some rules more important than others.

To restrict the applicability of the rule, click the **Filters** button and enter the caller or callee information in the dialog box.

Figure 4–2 Example of Rules Configuration

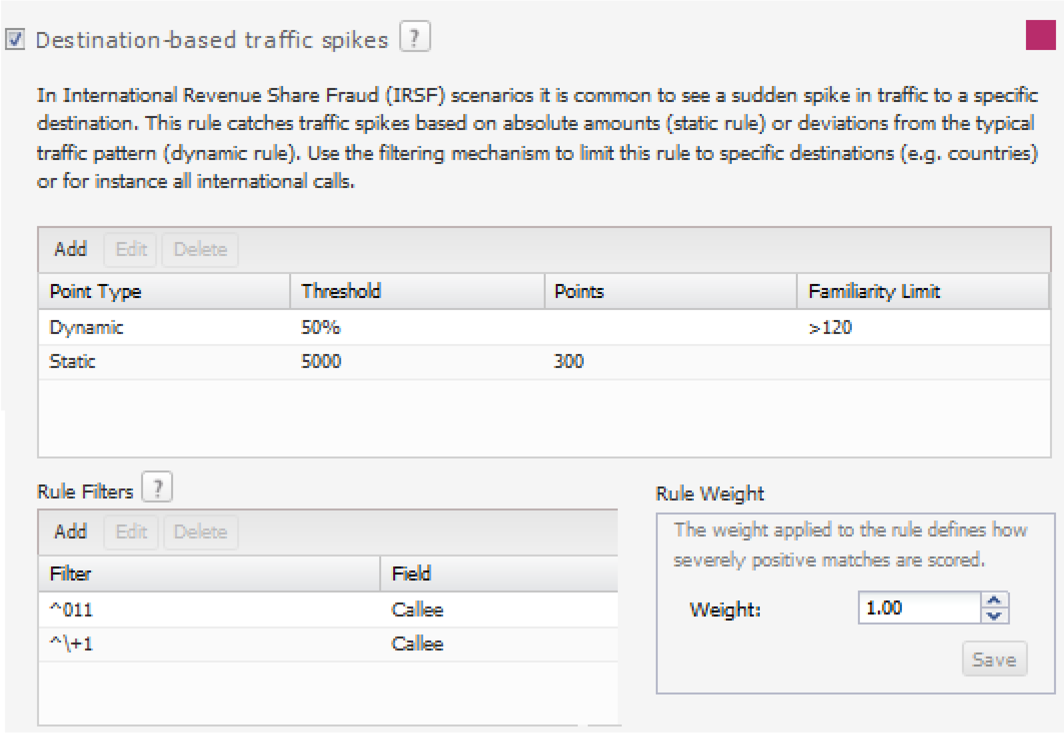

☑ Destination-based traffic spikes

In International Revenue Share Fraud (IRSF) scenarios it is common to see a sudden spike in traffic to a specific destination. This rule catches traffic spikes based on absolute amounts (static rule) or deviations from the typical traffic pattern (dynamic rule). Use the filtering mechanism to limit this rule to specific destinations (e.g. countries) or for instance all international calls.

Point Type	Threshold	Points	Familiarity Limit
Dynamic	50%		>120
Static	5000	300	

Rule Filters

Filter	Field
^011	Callee
^\+1	Callee

Rule Weight

The weight applied to the rule defines how severely positive matches are scored.

Weight: 1.00

Managing User Permissions

On the User section, you can manage several user accounts to work with Fraud Monitor. By default, only the *admin* account exists.

Figure 4–3 shows an example of the users list.

To add another user and enable the user to access Fraud Monitor, click on **Add users...** A window appears where you enter the user name and the email. Then, enter an initial password in the two **Password** fields. Click **Next** to go to the Permission Settings dialog where you specify which actions the new user can perform within the product.

> **Note:** The new user will then be able to connect using the credentials you have chosen. It is recommended that the user change this password at the first connection.

Figure 4–3 Users List

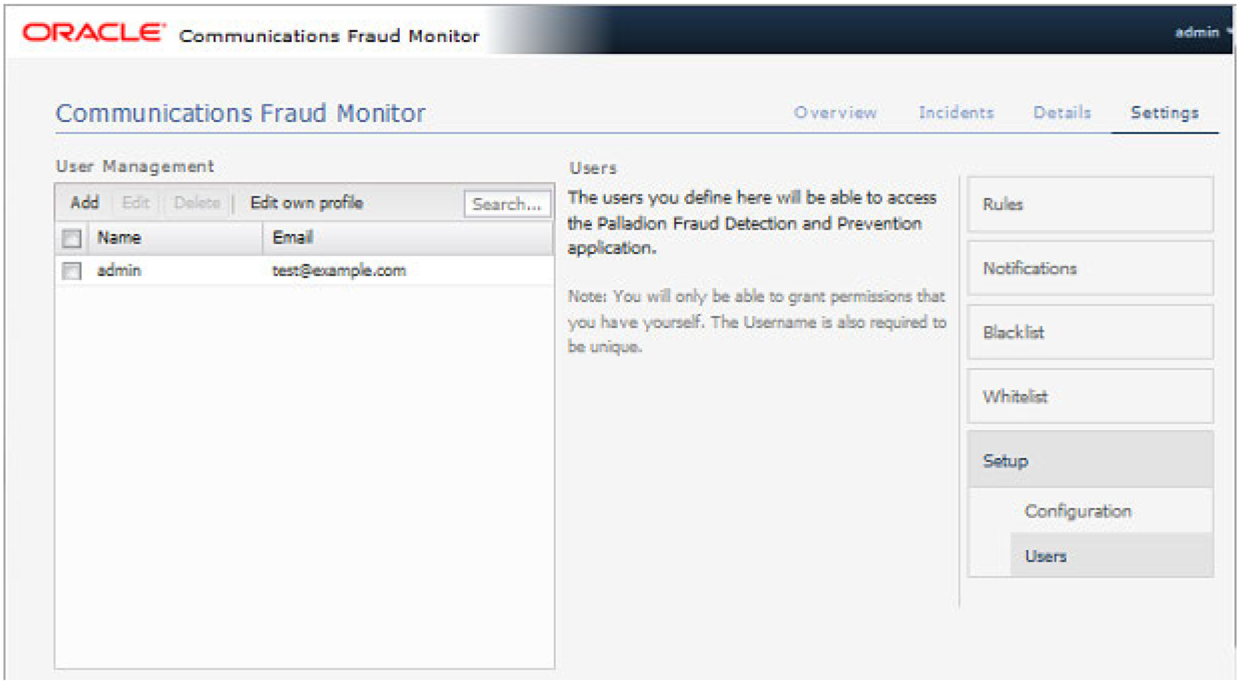

□	Name	Email
□	admin	test@example.com

The users you define here will be able to access the Palladion Fraud Detection and Prevention application.

Note: You will only be able to grant permissions that you have yourself. The Username is also required to be unique.

You are presented with a list of the editing and viewing capabilities a user can have. Each capability is named after a feature of the product. Refer to the corresponding section of this guide for more details. Select the boxes of those that you want to grant. You can use the **Check all** and **Uncheck all** buttons on the top right to check or uncheck all the boxes at once.

Figure 4–4 shows the user permission settings.

Figure 4–4 Granting Capabilities to the New User

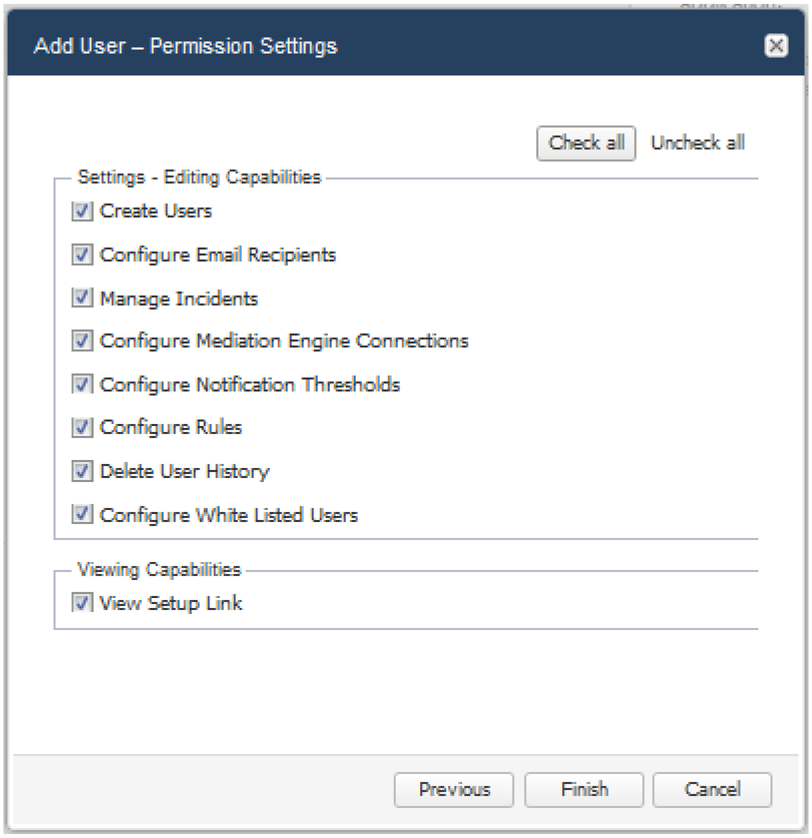

Add User – Permission Settings

Check all Uncheck all

Settings - Editing Capabilities

- ☑ Create Users
- ☑ Configure Email Recipients
- ☑ Manage Incidents
- ☑ Configure Mediation Engine Connections
- ☑ Configure Notification Thresholds
- ☑ Configure Rules
- ☑ Delete User History
- ☑ Configure White Listed Users

Viewing Capabilities

- ☑ View Setup Link

Setting Up Email Notifications

When Fraud Monitor detects an incident, it notifies the users by email.

Figure 4–5 shows an example of the notification settings.

To send email notifications, click on **Add recipient...** In the window that appears, enter the following settings:

- **Name:** A name to identify the new entry in the list of recipients
- **Email:** The email address to which notifications will be sent
- **Incident level:** Select **WARNING + CRITICAL** to receive all notifications, or **CRITICAL** to only receive notification on critical incidents
- **Prefix:** Emails from the system will contain this prefix in the **Subject:** field of the recipient inbox

Figure 4–5 Notification Settings for an Email Recipient

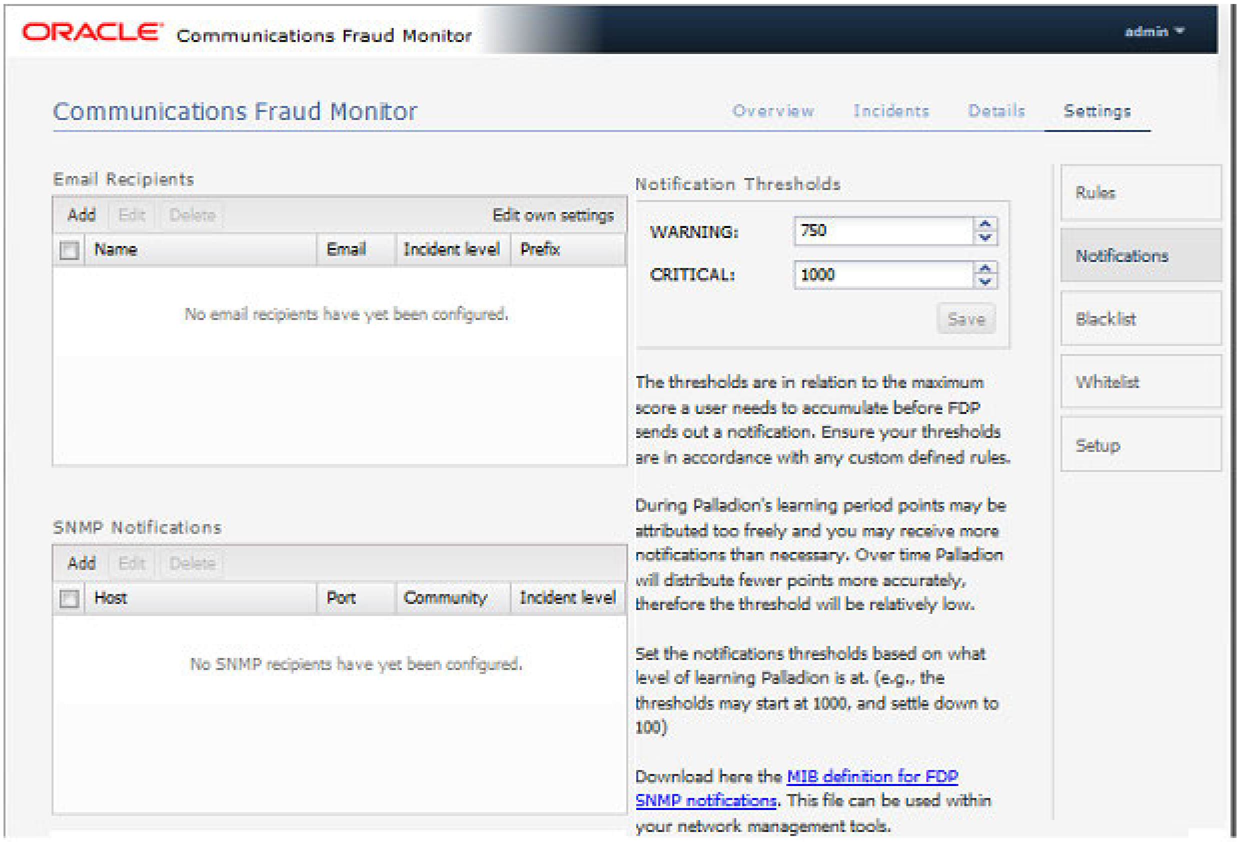

Email Recipients

Add Edit Delete Edit own settings

□	Name	Email	Incident level	Prefix

No email recipients have yet been configured.

SNMP Notifications

Add Edit Delete

□	Host	Port	Community	Incident level

No SNMP recipients have yet been configured.

Notification Thresholds

WARNING: 750

CRITICAL: 1000

Save

The thresholds are in relation to the maximum score a user needs to accumulate before FDP sends out a notification. Ensure your thresholds are in accordance with any custom defined rules.

During Palladion's learning period points may be attributed too freely and you may receive more notifications than necessary. Over time Palladion will distribute fewer points more accurately, therefore the threshold will be relatively low.

Set the notifications thresholds based on what level of learning Palladion is at. (e.g., the thresholds may start at 1000, and settle down to 100)

Download here the MIB definition for FDP SNMP notifications. This file can be used within your network management tools.

Adjusting the Notification Levels

To receive more or less notifications, you can adjust the two levels, warning and critical, in number of Incident points. The rules specified in the Rules page assign points to each user of the network. If the number of points for a user exceeds the threshold warning (1000 by default), an email is sent to all recipients of level **WARNING**. If it exceeds the level critical (1500 by default), the notification is sent to all recipients.

This is a global sensitivity adjustment. You can choose the amount of points each single rule attributes in the Rules section.

Specifying Blacklist Phone Numbers

The Blacklist contains phone numbers, IP addresses, and SIP User Agents which have been verified in fraudulent activity. You can enable and disable the Blacklist feature for specific data types in the configuration menu.

The Blacklist information provided by Oracle is in the international format. You can append a prefix to international numbers or provide a regular expression to transform the number.

The Global Blacklist is read-only and can be uploaded using the **Update** menu. You can also add and remove individual entries in the Custom Blacklist area.

Specifying Whitelist Phone Numbers

You can add and remove whitelist entries. Both IP addresses and phone numbers are possible. After adding or removing white-list entries, click **Save**. The new rules will go into effect immediately.

Phone numbers or IP addresses matching a whitelist entry are not used for point calculation. This filtering is done before any processing by any rule.

Calls which match a whitelist entry can still raise incidents. For example, if you block a certain caller IP address a call can still trigger an incident if the callee phone number is on the blacklist.

Both the phone number and the IP address of the caller and of the callee are tested against the list. The comparison is against the complete value, so there are no regular expressions or substring comparisons. If there are alphanumeric letters in the number, these will be treated as case sensitive.

Some rules only check against the caller's IP address or phone number. Filtering based on values you would expect in the callee won't significantly effect these rules.

Specifying Mediation Engine Configuration

The Mediation Engine Configuration section lists the Mediation Engines which are used as a source for the call data. Fraud Monitor won't have any data to analyze unless it's connected to a Mediation Engine.

On first use, Fraud Monitor will ask for the name, IP number, and port number of the Mediation Engine (see "Overview of Fraud Monitor Installation"). That information may be changed under this section if needed.

Name is free-form and merely a reminder for the system administrator. IP number is the IP address to which Fraud Monitor will try to connect. Port is 12000 by default, and configured in the Mediation Engine.

After adding or changing a connection, Fraud Monitor will test the connection. Any errors will display in a dialog box.

After changing an existing entry the entry will have a small red marking to show it has unsaved changes. To save the data, click **Save** on the lower right. Adding a new instance does not require saving.

It is possible to configure multiple Mediation Engines. All their call-data will be used for all rules. Fraud Monitor does not distinguish which Mediation Engine a call comes from.

Glossary

IRSF

International Revenue Share Fraud is a fraud scenario in which the fraudulent user earns money by directing lots of traffic to his or her own revenue share numbers.

Probe

A machine which filters and processes network traffic. It doesn't calculate the statistics.

PBX

Private Branch Exchange is a telephone exchange (often SIP based) that connects a business to a VoIP carrier (Communication Service Provider).

RTP

Real-time Transport Protocol. Used for transporting media. Defined in **RFC 3550**. For more information, see the IETF Tools website at:

`http://tools.ietf.org/html/rfc3550.html`

VLAN

Virtual Local Area Network is a technique to separate a network into distinct, isolated broadcast domains.

See `https://en.wikipedia.org/wiki/Virtual_LAN`.