

Exinda How To Guide: Application Performance Scores (APS)

Exinda ExOS Version 7.4.3

© 2016 Exinda Networks Inc.



Copyright

© 2016 Exinda Networks Inc. All rights reserved. No parts of this work may be reproduced in any form or by any means - graphic, electronic, or mechanical, including photocopying, recording, taping, or information storage and retrieval systems - without the written permission of the publisher.

Products that are referred to in this document may be either trademarks and/or registered trademarks of their respective owners. The publisher and the author make no claim to these trademarks.

While every precaution has been taken in the preparation of this document, the publisher and the author assume no responsibility for errors or omissions, or for damages resulting from the use of information contained in this document or from the use of programs and source code that may accompany it. In no event shall the publisher and the author be liable for any loss of profit or any other commercial damage caused or alleged to have been caused directly or indirectly by this document.

Document Built on Friday, July 22, 2016 at 2:36 PM

Using this guide

Before using this guide, become familiar with the Exinda documentation system.

Documentation conventions

These documentation conventions apply across all of the Exinda documentation sets. All instances of the following may not appear in this documentation

Typographical conventions

- **bold** - Interface element such as buttons or menus. For example: Select the **Enable** checkbox.
- *italics* - Reference to other documents. For example: Refer to the *Exinda Application List*. Also used to identify in the various procedures the response the systems provide after applying an action.
- > - Separates navigation elements. For example: Select **File > Save**.
- `monospace text` - Command line text.
- `<variable>` - Command line arguments.
- `[x]` - An optional CLI keyword or argument.
- `{x}` - A required CLI element.
- | - Separates choices within an optional or required element.

Links

With the exception of the various tables of contents, all links throughout the documentation are **blue**. Most links refer to topics within the documentation, but there may be links that take you to web pages on the Internet. In this documentation we differentiate between these types of links by **underlining** only the external links.

Tips, Notes, Examples, Cautions, etc.

Throughout this manual, the following table styles are used to highlight important information:

- **Tips** include hints and shortcuts. Tips are identified by the light blue icon.

**TIP**

text

- **Notes** provide information that is useful at the points where they are encountered. Notes are identified by the pin and paper icon.

**NOTE**

Text

- **Important** notes provide information that is important at the point where they are encountered. Important notes are identified by the amber triangle.

**IMPORTANT**

Text

- **Cautions** provide warnings of areas of operation that could cause damage to appliances. Cautions are identified by the orange triangle.

**CAUTION**

Text

- **Examples** are presented throughout the manual for deeper understanding of specific concepts. Examples are identified by a pale green background.

EXAMPLE

Text

- **Best Practices** are identified by the "thumbs-up" icon.

**Best Practice:**

It is a best practice to

Table of Contents

Introduction	7
Application Performance Score Reports	8
<i>What to expect</i>	9
If an APS report is not showing data	9
If the thresholds were set by using the baselining feature	9
Evaluating the APS	9
<i>Looking at the results</i>	10
Determining what might be causing a low APS score	10
Determining if this has been a persistent problem	10
Determining if you should be paying attention to the normalized delays or the transaction delays	10
How to have the system notify you if the APS score drops too low	11
Too many APS score lines make the APS chart difficult to read	11
Configure Application Performance Score Objects	12
<i>Creating an Application Performance Score object</i>	16
<i>Manually creating APS thresholds</i>	17
<i>Automatically calculating APS thresholds</i>	19
<i>How to know if baselining is in progress</i>	20
<i>Can I edit APS objects created in the Solution Center?</i>	20
<i>How an Application Performance Score is calculated</i>	21
How are thresholds set?	22
<i>How the Performance Metric thresholds are calculated</i>	22
Application Performance	23
<i>Understanding baselining</i>	24
If the application performance chart is baselining	24
It said it was going to baseline for an hour but it completed in 10 minutes. Was it a valid	24

baseline period?	
Configure Application Performance Metric Objects	25
<i>Creating an Application Performance Metric (APM)</i>	<i>27</i>
How Network Performance Metrics are calculated	29
<i>Round Trip Time</i>	<i>30</i>
<i>Network and Server Delay in a Read Transaction</i>	<i>31</i>
<i>Network and Server Delay in a Write Transaction</i>	<i>32</i>
<i>Delay Normalization</i>	<i>34</i>
<i>Packet loss</i>	<i>35</i>
View a network summary of application groups	36
Monitor the real time application response	37
Monitor the real time TCP health	39
Generate a PDF report of APS results	42

Introduction

The Application Performance Score (APS) object analyzes application metrics such as network delay, server delay, jitter, and others, to determine a measure of the network end-user experience of application performance. You can create an APS and then monitor the behavior of the application as part of the Monitor charts. Also, the application performance solution in the Solution Center automatically configures an APS object and shows the result on the application performance solution page.

This guide describes how to create an APS object, how to interpret the APS monitor information, how the application performance solution uses and displays the APS. It also goes into depth describing how the network metrics are calculated and how these network metrics are used to calculate the APS score.

Application Performance Score Reports

The application performance score (APS) report shows scores that assess the network performance users experience when using business-critical applications. The report is available by going to **Monitor > Service Levels > Application Performance Score (APS)**. The score, ranging between 0 and 10, where 0 is poor and 10 is excellent, indicates whether the app is performing as well as expected or is performing poorly. Multiple applications can be assessed. The scores are graphed over time to show how the scores are changing. The underlying metrics and measures that are used to calculate the scores are shown in the table. You can drill into the details of the APS by clicking on the application name. These charts can answer questions such as:

- Are my important applications performing well from a network perspective for my network users?
- Has this been a persistent problem or is it getting worse?
- If an application is not performing well, what might be causing the problem?



APS Scores							
Name	Score	Transaction Delays (ms)		Jitter (ms)	Loss (%)		RTT (ms)
		Network	Server		Inbound	Outbound	
<input checked="" type="checkbox"/> HTTP	9.52	66.32	125.11	40.18	0.50	0.60	70.25
<input checked="" type="checkbox"/> License DB	4.13	277.85	28.44	19.49	0.40	1.00	265.18
<input checked="" type="checkbox"/> SMTP	9.98	41.86	2.57	1.40	2.10	0.00	250.79

The score includes input from one or more of the following metrics:

- Network delay – the time taken for data to traverse the network (on the wire)
- Server delay – the time taken for a server to respond to the request
- Normalized network delay – the time taken for data to traverse the network, where the delay is measured independent of the transaction size by assuming a normalized packet size of 1024 bytes
- Normalized server delay – the time taken for a server to respond to the request, where the delay is measured independently of the transaction size, by assuming a normalized packet size of 1024 bytes

- Round-trip time – the time taken for
- Jitter – the measure of variability of network delay, defined as one standard deviation of network delay
- Inbound loss – the percentage of packet loss on inbound traffic
- Outbound loss – the percentage of packet loss on outbound traffic

Each metric that contributes to the score has a threshold value set. The threshold may have been set manually or may have been determined automatically by the Exinda appliance observing the traffic for the period of time to determine a baseline threshold values. The table below the chart indicates the current observed values for these metrics and whether that value is considered good or not.

- If the observed traffic is within the threshold, it is considered good and is coloured green in the APS Scores table.
- If the observed traffic is above the threshold but not above 4 x the threshold, it is considered tolerable and is coloured yellow.
- If the observed traffic is above 4 x the threshold, it is considered poor and is coloured red.
- If there is no colour in the table for a particular metric, it indicates that the metric is not contributing to the calculation of the APS score.

You can use this information to determine which metrics are reporting an issue.

What to expect

If an APS report is not showing data

Either the APS object does not have thresholds set and therefore the score cannot be calculated or there is no traffic for the specified application on the network for the period that is shown on the screen.

If the thresholds were set by using the baselining feature

You should get an application performance score of 9.0 if the same traffic were to be observed.

Evaluating the APS

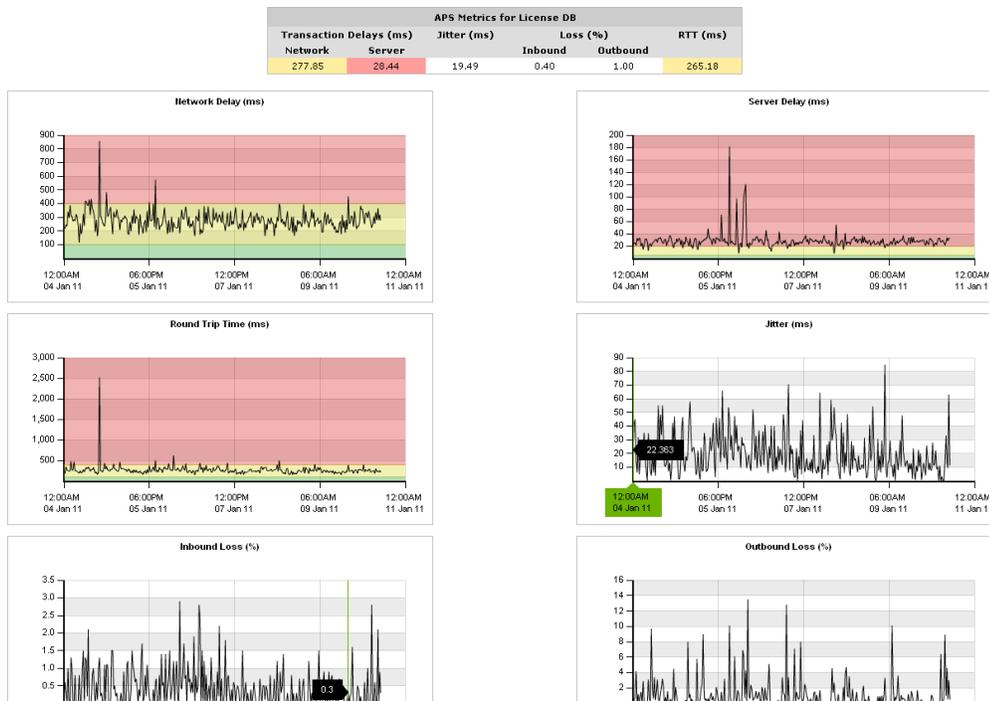
If the thresholds were set by using the baselining feature, a score of 8.5 or higher is considered a good score. The thresholds are automatically set to be slightly higher than the average observed measures so that good would be at the top end of the range of scores.

If you set the thresholds manually such that your thresholds lie at the boundary of what you consider good or not good, then you should consider APS scores greater than 5.0 to be good, as statistically half the time you would expect the observed values to be slightly above your threshold and half the time the observed values would be slightly below the threshold.

Looking at the results

Determining what might be causing a low APS score

Click the APS name in the table. A new screen that shows charts for each underlying metric appears. The background of each chart is coloured to represent value within threshold, within 4 x threshold, and above. If the background of the chart is not coloured, then that metric does not contribute to the calculation of the APS score. You can zoom into these charts by clicking and dragging within a chart to zoom into that x-axis range. All other charts will synchronize their zoom ranges to the specified zoom range. The metrics that are showing poor values could indicate a problem to investigate. For example, if the network delay is good, but the server delay is poor, then you'll know that the network is not the blame and that the server administrator should take a look at the application server.



Determining if this has been a persistent problem

Look at the APS score time line. If the score has been low for an extended period or if it looks like the score is dropping, you'll know that this is a persistent problem that needs addressing.

Determining if you should be paying attention to the normalized delays or the transaction delays

Generally, you should use the transaction delays unless the protocol that is being monitored has large or variable sized packets. The normalized delay measure normalizes the score to reflect a 1024 packet

size allowing you to more easily compare delays when the packets are variable in size.

How to have the system notify you if the APS score drops too low

You can configure the system to send you email if the APS score drops below an APS value that you specify and remains below that value for a specified duration. For example, you can set it to notify you if it drops below 7.0 and stays below 7.0 for 5 minutes.

Too many APS score lines make the APS chart difficult to read

- You can temporarily remove lines from the APS Scores chart by clearing the checkboxes next to the APS name in the table.
- You can zoom into an area of interest by clicking and dragging in the chart to select a smaller time range. This often has the effect of flattening the lines so that it appears less cluttered.

Related Topics

[Configure Application Performance Score Objects \(page 12\)](#)

[How an Application Performance Score is calculated \(page 21\)](#)

Configure Application Performance Score Objects

The application performance score (APS) object is used to assess how network users enjoy the network performance experience of business-critical applications. The score, ranging between 0 and 10, where 0 is poor and 10 is excellent, indicates whether the app is performing as well as expected or is performing poorly. By creating an APS object, you specify an application to monitor. Optionally, you can also specify a network object so that the application is only monitored when observed on that part of the network. You set thresholds on one or more network metrics. Later, traffic for that application is assessed against those thresholds to determine how well the application is performing.

The appropriate thresholds for an application is unique for each network environment. You can manually set the thresholds for the network metrics or you can have the system automatically create threshold values by having the system observe traffic to determine reasonable baseline values. The metrics include network delay, server delay, round trip time, jitter, and network loss. Note that you can manually set the network loss metric, however, it will not be automatically be calculated during the baseline analysis. You can use one or more of these metrics in your APS object. Most applications use transactional protocols. Applications like Citrix XenApp server or Microsoft Remote Desktop use non-transactional protocols that send information between the client and server at arbitrary times. With these types of applications, the standard method of calculating the network delays and server delays does not produce an accurate metric. If the application uses a non-transactional protocol, you must specify that when creating APS object.

For the baselining analysis, traffic is analyzed during the specified period, and a set of metric thresholds is generated. The threshold recommendations target an APS of 8.5. If the application reports an APS below 8.5, the application is performing worse than the baseline. If no traffic is observed during the baselining period, then the appliance will automatically start another the baseline analysis for the next larger time period. Email will be sent for each unsuccessful baseline analysis.



Best Practice:

It is a best practice to start the baseline analysis during a time period when you would expect traffic for the application is typical. This will ensure that the baseline values accurately reflect the typical usage of the application.

This means that if network conditions changes, it is recommended that the thresholds are re-evaluated.



IMPORTANT

APS is not supported for small-packet applications like Citrix and RDP. The metrics are normalized as if the application runs with larger packet sizes, leading to larger values.

You can also set alerts so that you will be notified when the score drops below a certain threshold value. There is an alert trigger delay setting which requires that the score remains below the alert threshold for a specified period of time before triggering the alert. This prevents brief temporary poor scores from appearing like an emergency.

Add New APS Object

APS Name:

Application:

Network Object - Internal:

Network Object - External:

Alert Enable:

Alert Threshold:

Alert Trigger Delay:

Auto Baseline

Auto Baseline Period:

Non-Transactional Protocol

When editing the APS object, you can modify the alert configuration, restart the baselining operation, and modify the threshold values. If you change the network object settings, it is recommended that you re-evaluate the metric thresholds and possibly re-start a baseline.

Edit APS Object			Baseline Info	
APS Name:	<input type="text" value="SalesForce"/>		Status:	Stopped
Application:	<input type="text" value="SalesForce"/>		Average Packet Size (bytes):	757
Network Object - Internal:	<input type="text" value="ALL"/>		Traffic Seen (KB):	6001
Network Object - External:	<input type="text" value="ALL"/>		Start Date:	Mon Feb 10 09:00:00 EST 2014
Alert Enable:	<input type="checkbox"/>		End Date:	Mon Feb 10 10:00:00 EST 2014
Alert Threshold:	<input type="text" value="0.0"/>		Auto Baseline Period:	<input type="text" value="Current Hour"/>
Alert Trigger Delay:	<input type="text" value="5 minutes"/>			
Non-Transactional Protocol:	<input type="checkbox"/>			
Scoring Metrics				
Metric	Config	Baseline		
Normalized Network Delay (ms/kb):	<input type="text" value="48"/>	48		
Normalized Server Delay (ms/kb):	<input type="text" value="22"/>	22		
Network Delay (ms):	<input type="text" value="176"/>	176		
Server Delay (ms):	<input type="text" value="146"/>	146		
Network Jitter (ms):	<input type="text" value="26"/>	26		
Round Trip Time (ms):	<input type="text" value="143"/>	143		
Network Loss (%):	<input type="text"/>	-		
			<input type="button" value="Start Baseline"/>	<input type="button" value="Stop Baseline"/>
			<input type="button" value="Apply Changes"/>	<input type="button" value="Cancel"/>

Related Topics and Tasks

Creating an Application Performance Score object	16
Manually creating APS thresholds	17
Automatically calculating APS thresholds	19
How to know if baselining is in progress	20
Can I edit APS objects created in the Solution Center?	20
How an Application Performance Score is calculated	21

How are thresholds set?	22
How the Performance Metric thresholds are calculated	22

Creating an Application Performance Score object

Use the instructions that follow to create a new APS object. During this set up, you can set a scope for the monitoring process. The scores can focus on specific internal and/or external network objects, or on ALL in one or both categories.

Before you begin...

- If you need to enable alerts, ensure that you have set Email on the **Configuration > System > Setup > Alerts** page. For more information, see the Exinda Web UI help.
- You also need to set up SNMP on the **Configuration > System > Network > SNMP** page. for more information, see the Exinda Web UI help.

To create the object:

1. Go to **Configuration > Objects > Service Levels > Application Performance Score**.
2. Click the **Add New APS Object** button.
3. In the **APS Name** field, type a name for the score.
4. In the **Application** list, select the application traffic to monitor.
5. Open the **Network Object - Internal** drop-down and either select a specific network object or select ALL.
6. Open the **Network Object - External** drop-down and either select a specific network object or select ALL.

NOTE: *By specifying both an internal and external network object, only the application conversations between the specified network objects is tracked.*

7. If you want to be alerted when the application performance score drops below a particular threshold, set the following alert settings:
 - a. Ensure the **Alert Enable** checkbox is selected.
 - b. In the **APS Threshold** field, set a threshold value between 0 and 10.
 - c. In the **Alert Trigger Delay** field, specify how many minutes that the APS score to be below the threshold before the notification is sent.

EXAMPLE

If the alert threshold is set to 7.0 and the alert trigger delay is set to 5 minutes, then the alert needs to be below 7.0 for 5 minutes before the alert is triggered.

8. If you need baselining to start immediately, select the **Auto Baseline** checkbox and select the **Auto Baseline Period**.
9. If the application uses a non-transactional protocol for traffic between the client and server, such as Citrix XenApp Servers or Microsoft Remote Desktop, select the **Non-Transactional Protocol** checkbox.
10. Click **Add New APS Object**.
The object is added to the list of configured APS objects.

Related Topics and Tasks

Manually creating APS thresholds (page 17)

Automatically calculating APS thresholds (page 19)

How to know if baselining is in progress (page 20)

Can I edit APS objects created in the Solution Center? (page 20)

How the Performance Metric thresholds are calculated (page 22)

Manually creating APS thresholds

Metric thresholds can be set manually when initially creating the APS object or upon editing an APS object even if they were automatically determined by the baselining operation. For example, if the baselining operation set all of the thresholds and you really only care about round trip time, normalized server delay, and normalized network delay, then you can remove the threshold settings for the other metrics.

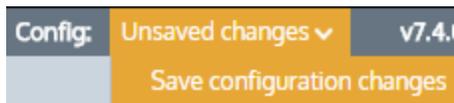
To manually set the threshold values

1. Go to **Configuration > Objects > Service Levels > Application Performance Score**.
2. On the **Add New APS Object** form, uncheck the **Auto Baseline** checkbox. Note if a baseline analysis is running, you'll need to press the **Stop Baseline** button. The threshold values are only editable there is not a baseline running. The metrics will appear on the screen.

Or edit the APS object in the list, then on the **Edit APS Object** form, the **Scoring Metrics** appear at the bottom of the form.

3. Enter or modify the values for the metrics that you are interested in setting thresholds for. Note that any metric that does not have a threshold set will not be analyzed when calculating the APS score.
 - Network delay – the time taken for data to traverse the network (on the wire) in one direction from the client through the Exinda appliance to the server (or in the opposite direction) in ms
 - Server delay – the time taken for a server to respond to the request in ms
 - Normalized network delay – the time taken for data to traverse the network in one direction, where the delay is measured independent of the transaction size by assuming a normalized packet size of 1024 bytes
 - Normalized server delay – the time taken for a server to respond to the request, where the delay is measured independent of the transaction size by assuming a normalized packet size of 1024 bytes
 - Round-trip time – the time taken for a packet to travel from a client through the Exinda appliance to the server and back
 - Jitter – the measure of variability of network delay, defined as one standard deviation of normalized network delay
 - Inbound loss – the percentage of packet loss on inbound traffic
 - Outbound loss – the percentage of packet loss on outbound traffic
4. Click **Apply Changes**.

To save the changes to the configuration file, in the status bar click the **Unsaved changes** menu and select **Save configuration changes**.



Related Topics and Tasks

- Creating an Application Performance Score object (page 16)
- Automatically calculating APS thresholds (page 19)
- How to know if baselining is in progress (page 20)
- Can I edit APS objects created in the Solution Center? (page 20)
- How the Performance Metric thresholds are calculated (page 22)

Automatically calculating APS thresholds

The baselining process can be started when initially creating the APS object or upon editing an APS object. At any time you can restart the baselining process if you would like the system to recalculate the thresholds.

To automatically calculate the thresholds by baselining the traffic:

1. Go to **Configuration > Objects > Service Levels > Application Performance Score**.
2. On the **Add New APS Object** form, ensure that the **Auto Baseline** checkbox is checked and set how long you want the system to observe traffic when calculating the thresholds by using the **Auto Baseline Period** drop-down list.

Select the time period for the baseline based on how popular the application is. For example, if there is a lot of HTTP traffic on the network, the 1 hour period will be long enough to analyze traffic and create an accurate baseline. For an application that is not used very often, use the 1 week baseline period to ensure that enough traffic is analyzed to generate baseline recommendations.

3. Or on the **Edit APS Object** form, set the **Auto Baseline Period** and click **Start Baseline**.



NOTE

The Network Loss metric is not calculated during the baseline analysis.



NOTE

If no traffic matching this APS object is observed during the baseline period, the appliance restarts the baseline analysis for the next larger time period. For example, if no traffic observed during the one hour period, the traffic continues to be analyzed for one day. If no traffic is observed during the one day period, then the traffic is analyzed for a week. If the traffic is analyzed for one week and no traffic has been transferred, the auto baseline analysis stops.

Each time the system unsuccessfully baselines the traffic (that is, when no traffic is observed during the auto baseline period), an email notification is sent to the users configured on the Configuration > System > Network > Email page.

Related Topics and Tasks

Creating an Application Performance Score object (page 16)

Manually creating APS thresholds (page 17)

How to know if baselining is in progress (page 20)

Can I edit APS objects created in the Solution Center? (page 20)

How the Performance Metric thresholds are calculated (page 22)

How to know if baselining is in progress

On the **Application Performance Score** configuration tab, the list of APS objects is shown. If the APS is currently baselining the application traffic, there will be a green checkmark in the **Auto Baseline** column.

Press the **Edit** button for the APS object. The **Baseline Info** section specifies the status (Running or Stopped) and the Start and End Date and time of the baseline period. Note that it also shows the average packet size and the amount of traffic seen.

Related Topics and Tasks

Creating an Application Performance Score object (page 16)

Manually creating APS thresholds (page 17)

Automatically calculating APS thresholds (page 19)

Can I edit APS objects created in the Solution Center? (page 20)

How the Performance Metric thresholds are calculated (page 22)

Can I edit APS objects created in the Solution Center?

Yes. You can modify the thresholds either manually or by re-starting the baseline analysis. You can also enable an alert, and change the internal and/or external network object setting to monitor the application for a subset of your network.



IMPORTANT

Do not change the application being monitored, and do not change the name of the APS object.

Related Topics and Tasks

Automatically calculating APS thresholds (page 19)

How to know if baselining is in progress (page 20)

How the Performance Metric thresholds are calculated (page 22)

How an Application Performance Score is calculated

The Application Performance Score object defines the application traffic that will be monitored, and which application performance metrics should be evaluated providing application performance thresholds to be used in the evaluation.

For each metric, the observed traffic is compared against the threshold and is classified into one of three categories:

- **Good** — The baseline for the application is good, which indicates that the application is performing within the expected levels (below the threshold). The users should be happy with the performance of the application.
- **Tolerated** — The performance of the application is less than expected, but is still performing within a range that users should be able to tolerate (between the threshold and four times the threshold). The performance is not great, but users should tolerate it.
- **Frustrated** — The application is performing poorly (more than four times the threshold). The users will be frustrated with the performance.

The number of good observations for all metrics with a threshold are totaled and given a full weighting; the number of tolerated observations for all metrics with a threshold are totaled and given a half weighting; and all frustrated observations are given a zero weighting. These weighted totals are summed and divided by the total observations.

$$\text{aps} = 10 * ((1 * \text{number of satisfied samples}) + (0.5 * \text{number of tolerated samples}) + (0 * \text{number of frustrated samples})) / \text{total samples}$$

EXAMPLE

For HTTP, a threshold is configured for Network Delay as $T = 300$ msec and a threshold is configured for round-trip time (RTT) as $T = 40$ msec.

In one 10s period, 11 flows are sampled for HTTP with the following results:

- 2 flow samples have a network delay of > 1200 ms (frustrated samples)
- 3 flow samples have a network delay of > 300 ms but < 1200 ms (tolerated samples)
- 6 flow samples have a network delay of < 300 ms (satisfied samples)
- 1 flow sample has a RTT of > 40 ms but < 160 ms (tolerated samples)
- 10 flow samples have a RTT of < 40 ms (satisfied samples)

The APS score is calculated as follows:

$$\text{aps} = 10 * (1 * (6 + 10) + 0.5 * (3 + 1) + 0 * 2) / 22 = 8.1$$

How are thresholds set?

The appropriate thresholds for an application is unique for each network environment. The thresholds can be set manually when configuring an APS object or the Exinda appliance can analyze the traffic for an application for a baseline period and create a recommended set of thresholds.

For more information about creating APS objects and setting the thresholds, read "Configure Application Performance Score Objects" on page 12.

How the Performance Metric thresholds are calculated

The network performance metrics are calculated based on the observed traffic. Each threshold is calculated to be 0.85 of a standard deviation above the average observation for that metric. This ensures that the calculated thresholds target is an APS of 9.0. If the application reports an APS below 9.0, the application is performing worse than the baseline.

Related Topics and Tasks

[Creating an Application Performance Score object \(page 16\)](#)

[Manually creating APS thresholds \(page 17\)](#)

[Automatically calculating APS thresholds \(page 19\)](#)

[How to know if baselining is in progress \(page 20\)](#)

[Can I edit APS objects created in the Solution Center? \(page 20\)](#)

Application Performance

The application performance solution monitors your important applications. On one screen, you can have insight into who is using the application and their perceived performance of the application. You can also see the amount of bandwidth used by the application and the amount of reduction achieved, if applicable. These charts can answer questions such as:

- Is this important application performing well from the perspective for my network users?
- Has this been a persistent problem or is it getting worse?
- If an application is not performing well, what might be causing the problem?
- How much bandwidth does this application use?
- Does it appear to be bandwidth constrained?
- Who are the top users and top hosts of this application?

The application performance chart shows the network user experience of the performance of the application. You should expect the application performance to show a good score (between 8.5 and 10.0). If the score is less than 7.0, you may want to investigate if there is a problem. To see the measures that contributed to the score, click **Show more**.

When first creating an Application Performance monitor, the application performance metric needs a baseline understanding of the observed traffic for the application in your network. The baseline operation starts automatically and observes the traffic for an hour. Once a baseline is set, the monitor compares current traffic against the metrics calculated in the baseline to determine the application performance score.

The inbound and outbound bandwidth charts show how much bandwidth the application is using. You should expect the bandwidth to show spikes instead of raised flat tops. Flat tops in the graph often indicate that the traffic may be limited by policy rules. The chart shows data measured on the WAN-side of the appliance, that is before the accelerated traffic is decompressed for inbound traffic and after acceleration and traffic shaping policies have been applied for outbound traffic. You can choose to overlay the data measured on the LAN-side of the appliance. This shows the amount of reduction achieved due to acceleration and traffic shaping.

The users and hosts bar charts show the WAN-side data volumes consumed by the top users and hosts for the application. Typically, applications are used by multiple users or hosts and the traffic distribution is fairly even amongst the top users or hosts. If one user or host shows considerably more data volume than the other users, it may be reasonable behavior or it may indicate a problem worthy of further investigation. Also, you can choose to show just internal endpoints, that is, hosts and users on the LAN-side of your appliance, or just external endpoints, that is, hosts and users on the WAN-side of your appliance. You can also choose to show just users, just hosts, or both.

Understanding baselining

The following provides some commonly asked questions about application performance baselining.

If the application performance chart is baselining

The application performance score is calculated by analyzing the network performance of the application relative to thresholds set for various network and server metrics. When the application performance object is first set up, the baseline thresholds are set by observing traffic for a period of time. If the application performance chart indicates that it is baselining, it is computing the initial threshold values. Once the baseline period is complete, the application performance chart shows the results. When the Application Performance report is first created, it will automatically start to baseline the traffic. If no traffic is observed for the specified application during the baselining period, the baselining process repeats until traffic is observed and thresholds are calculated.

It said it was going to baseline for an hour but it completed in 10 minutes. Was it a valid baseline period?

Yes, it is a valid period. The Exinda appliance is always observing traffic and it stores its observations for an hour. If near the end of an hour (for instance, at 9:50) you create an application performance page, then the application performance score object will start baselining and can look at the last hour of data collected for the application and can then complete the baseline period in the remaining 10 minutes of the hour.

For more information, see the "*Solution Center – How-to Guide*".

Configure Application Performance Metric Objects

The Application Performance Metric (APM) objects are used to monitor particular application performance metrics. By creating an APM object, you indicate which application to monitor. Optionally, you can also specify a network object so that the application is only monitored when observed on that part of the network. You set a threshold on a single network metric. Later, traffic for that application is assessed against that threshold to determine how well the application is performing. An alert is triggered when the threshold is exceeded for a given length of time.

The following metrics are available:

- bytes lost
- network delay
- server delay
- transaction delay
- normalized network delay
- normalized server delay
- normalized transaction delay
- round trip time
- tcp connections aborted
- tcp connections ignored
- tcp connections refused
- tcp connected started.

Add New APM Object

APM Name:

Metric:

Application:

Network Object - Internal:

Network Object - External:

APM Threshold:

Alert Trigger Delay:

Alert Enable:

Add New APM Object

Cancel



NOTE

APM values are not shown on any report; they are used solely to generate alerts.

Related Topics and Tasks

Creating an Application Performance Score object (page 16)

Creating an Application Performance Metric (APM)

Use the following instructions to create an APM.

Before you begin...

- If you need to enable alerts, ensure that you have set Email on the **Configuration > System > Setup > Alerts** page. For more information, see the Exinda Web UI help.
- You also need to set up SNMP on the **Configuration > System > Network > SNMP** page. for more information, see the Exinda Web UI help.

To create an APM object

1. Go to **Configuration > Objects > Service Levels > Application Performance Metric**.
2. Click the **Add New APM Object** button.
3. Type a name for the APM object.
4. Select the metric that you need to monitor. The following metrics are available:
 - **bytes-lost** — Bytes lost due to retransmissions.
 - **network-delay** — The time taken for data to traverse the network.
 - **server-delay** — The time taken for a server to respond to a request.
 - **transaction-delay** — The total time for a transaction (network delay + server delay)
 - **normalized-network-delay** — The time taken for data to traverse the network where the packet size is normalized to 1024 bytes.
 - **normalized-server-delay** — The normalized measure of the time taken for a server to respond to a transaction request.
 - **normalized-transaction-delay** — The normalized measure of the time taken for a client request to be sent to a server, and the server's reply to be received by the client.
 - **round-trip-time** — The time taken for a packet to travel from a device, cross a network, and return.
 - **tcp-connections-aborted** — The number TCP connections reset after the connection is established. (RST from client or server)

- **tcp-connections-ignored** — The number TCP connections that expire in the SYN-SENT state. No response is received from the server.
 - **tcp-connections- refused** — The number TCP connections that are reset before the connection is established. (RST in SYN-SENT state)
 - **tcp-connections-started** — The number of TCP connections initiated.
5. In the **Application** list, select the application traffic to monitor.
 6. If you want to just monitor the application for a particular internal network object, specify the desired internal network object; otherwise select ALL.
 7. If you want to just monitor the application for a particular external network object, specify the desired external network object; otherwise select ALL.

By specifying both the internal and external network object, only the application conversations between the specified network objects will be tracked.
 8. Select the **Alert Enable** checkbox.
 9. In the **APM Threshold** field, type the threshold that will trigger an alert if the score drops below that value.
 10. In the **Alert Trigger Delay** list, select how long the metric needs to remain below the threshold before the alert is sent.

For example, if the alert is tracking the number of bytes lost, the threshold is set to 100, and the alert trigger delay is set to 5 minutes, then the number of bytes lost needs to be above 100 for 5 minutes before the alert is triggered.
 11. Set the threshold for the APM metric. The units of the threshold is relative to the metric being measured. That is, delays and round trip time are measured in milliseconds, tcp connections and bytes lost are counts.
 12. Click **Add New APM Object**.
The object is added to the list of configured APM objects.

How Network Performance Metrics are calculated

A transaction is defined as a client request followed by a server reply, including both TCP and UDP flows. With each read and write transaction between a client and a server, the following values are measured and used to calculate how long the transaction takes to complete:

- **Round Trip Time** — the time taken for a very small packet to travel across the network and return
- **Network Delay** — the overall time taken for data to cross from a client to a server, or from the server to a client.
- **Server Delay** — the time taken for a server to respond to a request.
- **Total Transaction Delay** — The time taken for data to cross the network from a client to the server and back. Unlike round trip time, this could include large packets and could result in multiple packets being sent to the server, or received from the server.
- **Network Jitter** — measures the variability of the network delay time. This is expressed as a multiple of one standard deviation.
- **Packet Loss** — measures when one or more packets within a transmission are successfully sent, but fail to arrive at the destination.

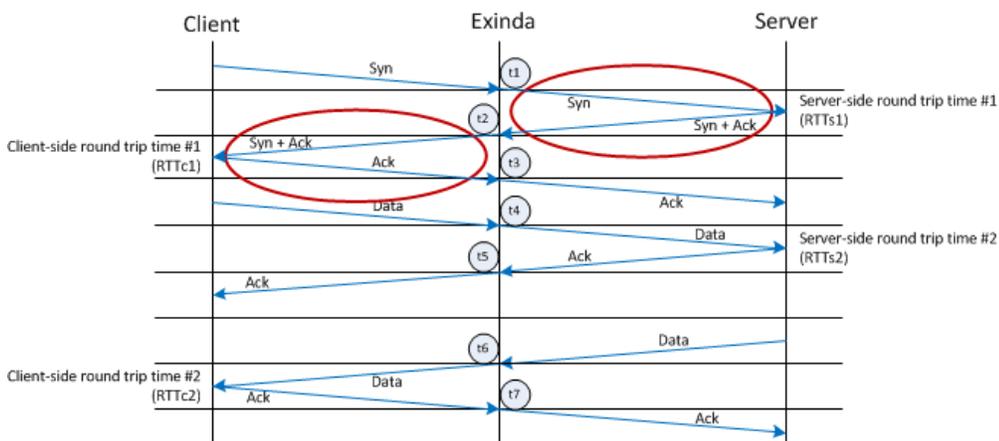
Related Topics

Round Trip Time	30
Network and Server Delay in a Read Transaction	31
Network and Server Delay in a Write Transaction	32
Delay Normalization	34
Packet loss	35

Round Trip Time

Round Trip Time (RTT) is the measure of how long it takes for a very small packet to travel across the network and for an acknowledgment of that packet to be returned. Consider the typical topology where an Exinda appliance is positioned between the client and the server. As each packet is intercepted by the Exinda appliance, it is time stamped with a highly accurate nanosecond resolution clock source. Since the Exinda appliance intercepts the packet after the client sends the packet, the start time is not known and so the RTT is determined by summing the round trip time from the appliance to the server and back (Server RTT), and the round trip time from the appliance to the client and back (Client RTT). As more packets are sent from the client through the Exinda appliance make the round trip, the RTT estimate is updated by averaging new information.

The following diagram illustrates how the round trip time is calculated:



Server RTT:

- $RTTs1 = t2 - t1$
- $RTTs2 = t5 - t4$

Client RTT:

- $RTTc1 = t3 - t2$
- $RTTc2 = t7 - t6$

$$\text{Average Server RTT} = (RTTs1 + RTTs2) / 2$$

$$\text{Average Client RTT} = (RTTc1 + RTTc2) / 2$$

Average Total RTT = avRTTs + avRTTc

Network and Server Delay in a Read Transaction (page 31)

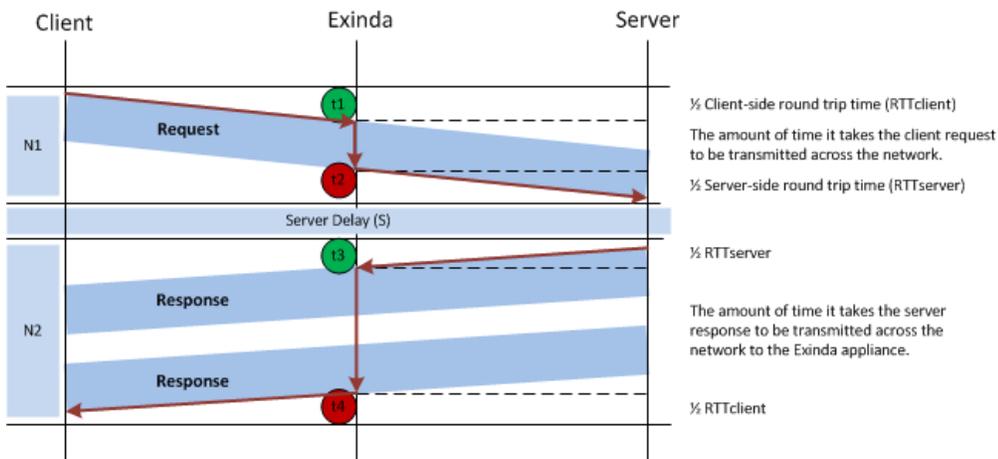
Network and Server Delay in a Write Transaction (page 32)

Delay Normalization (page 34)

Packet loss (page 35)

Network and Server Delay in a Read Transaction

When a client computer requests information from the server, the request and response are tracked to determine how long it takes for the client to send the request, and the server to send the requested data back to the client. The diagram below shows the flow of information between the client, the Exinda appliance, and the server, and identifies the points in the transaction where time stamps are acquired.



Network Delay for Read Request (N1)

1. The client sends a request to the server.
2. When the request passes through the Exinda, the time stamp is noted as the beginning of the request (t1).
3. When the end of the request passes through the Exinda, the time stamp is noted (t2).
 $t2 - t1 =$ The amount of time it takes the client request to pass through the Exinda appliance.
4. The server receives the complete client request.

Server Delay for Read Request (S)

- After the server receives a request from the client, the server takes some time to process the request. This is the Server delay (S).

Network Delay for Read Response (N2)

- The server's response to the client request is sent, and may be sent in a number of packets.
- When the first response passes through the Exinda, the time stamp is noted (t3).
- When the end of the last response passes through the Exinda, the time stamp is noted (t4).
 $t4 - t3 =$ The amount of time it takes the data requested by the client to pass through the Exinda appliance.
- The client receives the data requested from the server.

Total Time for Read Transaction

The total transaction time for a Read transaction is calculated as $\text{Transaction time} = N1 + S + N2$ where:

- $N1 = \frac{1}{2} \text{RTT}_{\text{client}} + (t2 - t1) + \frac{1}{2} \text{RTT}_{\text{server}}$
- $S = (t3 - t2) - \text{RTT}_{\text{server}}$
- $N2 = \frac{1}{2} \text{RTT}_{\text{server}} + (t4 - t3) + \frac{1}{2} \text{RTT}_{\text{client}}$

Related Topics

Round Trip Time (page 30)

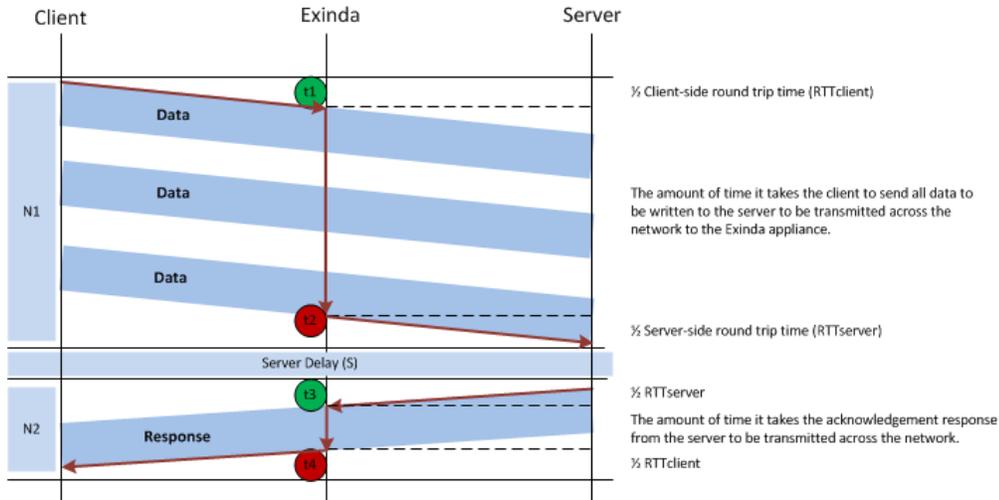
Network and Server Delay in a Write Transaction (page 32)

Delay Normalization (page 34)

Packet loss (page 35)

Network and Server Delay in a Write Transaction

When a client computer sends information to be written to the server, the request and response are tracked to determine how long it takes for the client to send the data to the server, and the server to send an acknowledgment of receiving the data back to the client. The diagram below shows the flow of information between the client, the Exinda appliance, and the server, and identifies the points in the transaction where time stamps are acquired.



Network Delay for Write Request (N1)

1. The client sends data to be written on the server, and may be sent in a number of packets.
2. When the first data packet starts passing through the Exinda, the time stamp is noted as the beginning of the packet (t_1).
3. When the end of the last data packet passes through the Exinda, the time stamp is noted (t_2).
 $t_2 - t_1 =$ The amount of time it takes the client to send data through the Exinda appliance.

4. The server receives all the data from the client.

Server Delay for Write Request (S)

5. There is a very small delay between receiving the data from the client and the acknowledgement that is sent from the Server back to the client. This is the Server delay (S).

Network Delay for Write Response (N2)

6. The server's acknowledgment response to the client that the data has been received is sent.
7. When the response passes through the Exinda, the time stamp is noted (t_3).
8. When the end of the response passes through the Exinda, the time stamp is noted (t_4).
 $t_4 - t_3 =$ The amount of time it takes the server response to pass through the Exinda appliance.

9. The client receives the response from the server.

Total Time for Write Transaction

The total transaction time for a Write transaction is calculated as Transaction time = $N1 + S + N2$ where:

- $N1 = \frac{1}{2} RTT_{client} + (t2 - t1) + \frac{1}{2} RTT_{server}$
- $S = (t3 - t2) - RTT_{server}$
- $N2 = \frac{1}{2} RTT_{server} + (t4 - t3) + \frac{1}{2} RTT_{client}$

Related Topics

Round Trip Time (page 30)

Network and Server Delay in a Read Transaction (page 31)

Delay Normalization (page 34)

Packet loss (page 35)

Delay Normalization

To create accurate comparisons of the network delay experienced by a transaction, the appliance must analyze packets of the same size (normalized). All other factors being equal, the transaction delays should increase with the amount of data transferred or the transaction size.

To make the APS score independent of transaction size, the transaction delay metrics are normalized using a constant of 1024 bytes. The normalized network delay is calculated as follows:

$$\text{Normalized Network Delay} = \text{Total Network delay} * 1024 / \text{transaction bytes}$$


NOTE

Due to the nature of normalization, protocols that naturally use very small sized packets (like RDP and Citrix) will produce APM values higher than the real ones given such that the results are being multiplied by 1024. The number of bytes used to normalize the calculation of the network delay during a transaction can be configured through the CLI, as well as disabling normalization altogether.

Related Topics

Round Trip Time (page 30)

Network and Server Delay in a Read Transaction (page 31)

Network and Server Delay in a Write Transaction (page 32)

Packet loss (page 35)

Packet loss

Packet loss occurs when one or more packets within a transmission are successfully sent, but fail to arrive at the destination. Packet loss can be caused by a variety of factors including network congestion, faulty network components such as hardware or drivers, or corrupted packets within the transmission. If the transmission experiences packet loss, it may cause the following:

- Jitter in video conferences
- Gaps in audio during VoIP communications
- Performance issues when streaming media

To recover from packet loss, data must be retransmitted to the destination to complete requests successfully. The amount of data retransmitted per flow is used to calculate the Network Efficiency metric.

$$\text{Efficiency} = 100\% * (\text{transferred} - \text{retransmitted}) / \text{transferred}$$

and

$$\text{Network Loss} = 100 - \text{Efficiency}$$


NOTE

Network loss, not efficiency, is used when calculating APS.

Related Topics

[Round Trip Time \(page 30\)](#)

[Network and Server Delay in a Read Transaction \(page 31\)](#)

[Network and Server Delay in a Write Transaction \(page 32\)](#)

[Delay Normalization \(page 34\)](#)

View a network summary of application groups

Each table shows the top Application Groups together with the number of packets, number of flows data transferred and throughput statistics.

1. With your browser, open the Exinda Web UI.

`https://UI_IP_address.`

2. Type the **User Name** and **Password**.

3. Click **Login**.

The Exinda Web UI appears.

4. Ensure you are in **Advanced** mode.

5. Click **Monitor > Application Groups** and switch to the **Groups** tab.

6. To expose Round trip time, Normalized Delays, Transaction Delays, and Efficiency statistics for each Application Group, click **Show Details**.

Top 30 Inbound Application Groups					
Name	Packets	Data (MB)	Throughput (kbps)		Flows
			Average	Max	
[+] Show Details					
Web	96259	116.599	76.11	8118.82	645
Social Networking	27639	36.786	181.54	4231.32	81
Other	11319	1.801	0.78	7.70	626
Exinda	3239	1.741	4.89	40.82	48
Streaming	737	0.970	264.86	366.70	2
Software Updates	691	0.898	105.13	183.42	3
Mail	1056	0.506	3.95	106.94	25
File Services	711	0.173	3.46	37.98	19
Thin Client	491	0.113	1.90	6.61	3
Voice	410	0.052	0.42	4.64	43
Interactive	142	0.034	1.54	5.05	17

Top 30 Outbound Application Groups													
Name	Packets	Data (MB)	Throughput (kbps)		Flows	RTT (ms)	Normalized Delays (ms/kb)			Transaction Delays (ms)			Efficiency (%)
			Average	Max			Network	Server	Total	Network	Server	Total	
[+] Hide Details													
Web	53524	7.024	4.69	156.41	631	51	4280	309	4589	3942	717	4659	100.00
Social Networking	13556	1.267	6.25	79.05	81	81	290	32	322	1097	68	1165	100.00
Other	1949	0.362	0.92	9.58	159	48	511	341	852	1641	285	1926	100.00
Mail	802	0.362	2.96	4.28	22	163	5595	35	5630	24270	177	24447	98.84
File Services	753	0.167	3.10	23.26	22	29	9	6108	6117	5	1412	1417	100.00
Thin Client	581	0.065	1.09	2.14	3	92	11202	7284	18486	612	379	991	100.00
Voice	381	0.039	0.31	3.36	43	191	46216	1096018	1142234	24492	19440	43932	99.62
Streaming	371	0.027	7.31	10.35	2	18	10	503	513	13	698	711	100.00
Software Updates	399	0.026	3.00	4.77	3	0	1	13	14	30	9	39	100.00
Interactive	69	0.008	1.60	2.66	3	38	181	0	181	33	0	33	100.00

7. To view the data for individual applications within a group, click the application group name.

Monitor the real time application response

The APM values are available as a real time display. The real time display shows the APM values by application for the selected time period. As well as the APM values, the number of flows and the number of transactions are shown.

Display the report in the Exinda Web UI

1. With your browser, open the Exinda Web UI.

`https://UI_IP_address.`

2. Type the **User Name** and **Password**.

3. Click **Login**.

The Exinda Web UI appears.

4. Ensure you are in **Advanced** mode.

5. Click **Monitor > Real Time** and switch to the **Application Response** tab.

The following report opens.

Application Response									
Application Name	RTT (ms)	Normalized Network (ms/kb)	Normalized Server (ms/kb)	Normalized Delay Total (ms/kb)	Network (ms)	Server (ms)	Transaction Delay (ms)	Transaction Count	Flows
HTTP	3074.50	1.94	0.98	2.92	73.19	2.16	75.36	38	79
FTP	9.81	0.00	0.00	0.00	0.00	0.00	0.00	0	6
mDNS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0	4
ICMPV6	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0	1
HTTPS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0	4
DNS	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0	202
SMTP	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0	2

6. To change how often the table is refreshed, select an **Auto-Refresh Rate** from the list.

Display the report in the Exinda CLI

1. With your browser, open the Exinda Web UI.

`https://UI_IP_address.`

2. Type the **User Name** and **Password**.

3. Click **Login**.

The Exinda Web UI appears.

4. Ensure you are in **Advanced** mode.

5. Click **Configuration > System > Tools > Console**.

6. Type the appliance username and password at the prompts. Do *one* of the following:

- To enter privileged EXEC (enable) mode, at the prompt type the following command:

```
hostname > enable
```

The hostname # prompt appears.

- To enter configuration (config) mode, at the prompt type the following commands:

```
hostname # configure terminal
```

The hostname (config)# prompt appears.

8. To display real time APM data from the CLI, use the following command:

```
(config) # show realtime apm applications
```

The following results are displayed:

```
ex-240 (config) # show realtime apm applications
```

Application	RTT (ms)	Network (ms)	Server (ms)	Transaction (ms)	Transactions	Flows
ExindaWM	956.04	77706.24	206863.37	226125.26	48	4
Unclassified	459.74	35040.99	15000.30	37512.24	8	44
Replify	292.75	2660.00	0.00	2655.70	4	1
HTTP	256.16	202.86	147.08	338.41	10	9
HTTPS	217.45	97.34	26.83	124.18	10	6
CIFS	108.53	186.69	89.73	231.30	2	2
SSH	71.48	386.28	0.00	336.24	2	1
ExindaCom	0.00	0.00	0.00	0.00	0	16
mDNS	0.00	0.00	0.00	0.00	0	3
ICMP	0.00	0.00	0.00	0.00	0	7
ssdp	0.00	0.00	0.00	0.00	0	1
IGMP	0.00	0.00	0.00	0.00	0	15
NTP	0.00	0.00	0.00	0.00	0	2
sLm	0.00	0.00	0.00	0.00	0	1

```
ex-240 (config) # █
```

Monitor the real time TCP health

The Real Time Host Health report shows the Retransmitted Bytes, Aborted Connections, Refused Connections, Ignored Connections, and Flow Count for each internal and external host monitored by the Exinda appliance.



Version Info:

A new internal mechanism was implemented in the ExOS 7.4.2 firmware that prevents Exinda Appliances from being affected by DDoS attacks. As a consequence, the “Ignored Connections” historical report is no longer available, but the Real Time TCP Health report still includes the number of ignored connections.

Display the report in the Exinda Web UI

1. With your browser, open the Exinda Web UI.

`https://UI_IP_address.`

2. Type the **User Name** and **Password**.
3. Click **Login**.
The Exinda Web UI appears.
4. Ensure you are in **Advanced** mode.
5. Click **Monitor > Real Time** and switch to the **Host Health** tab.
The report appears:

Health					
Internal IP	Retransmitted (bytes)	Aborted	Refused	Ignored	Flows
192.168.80.66	0	0	0	0	1
172.16.1.240	0	1	0	0	13
172.14.1.10	0	0	0	0	90
172.16.0.160	0	0	0	0	2
172.16.0.115	0	0	0	0	4
192.168.100.252	0	0	0	0	1
192.168.0.207	0	0	0	0	1
172.16.1.149	0	0	0	0	4
172.16.1.242	0	0	0	0	1
172.16.0.63	0	0	0	0	2
192.168.90.180	0	0	0	0	1
192.168.0.35	0	0	0	0	1
192.168.0.178	0	0	0	4	5
192.168.0.83	0	0	0	0	1
0.0.0.0	0	0	0	0	1
172.16.0.179	0	0	0	0	1
192.168.0.145	0	0	0	0	1
192.168.0.209	0	0	0	0	1
192.168.0.114	0	0	0	0	1
192.168.20.115	0	0	0	0	1
172.16.0.252	0	0	0	0	1
172.16.0.114	0	0	0	0	2
192.168.60.12	0	0	0	0	1
172.16.0.119	0	1	0	0	1
192.168.0.171	0	0	0	0	78
192.168.0.54	0	0	0	0	1

Health					
External IP	Retransmitted (bytes)	Aborted	Refused	Ignored	Flows
189.47.235.215	0	0	0	0	1
182.237.13.217	0	0	0	0	1
114.27.0.33	0	0	0	0	1
209.162.180.188	0	0	0	0	1
46.146.120.14	0	0	0	0	1
77.124.175.104	0	0	0	0	1
190.50.180.129	0	0	0	0	1
203.2.192.124	0	0	0	0	1
95.25.240.132	0	0	0	0	1
81.182.22.28	0	0	0	0	1
92.112.162.185	0	0	0	0	1
116.48.3.56	0	0	0	0	1
78.37.161.202	0	0	0	0	1
114.143.5.78	0	0	0	0	1
183.98.3.27	0	0	0	0	1
98.82.49.248	0	0	0	0	1
87.242.31.140	0	0	0	0	1
239.255.255.250	0	0	0	0	1
190.44.86.103	0	0	0	0	1
85.228.237.61	0	0	0	0	1
98.226.208.222	0	0	0	0	1
239.255.255.100	0	0	0	0	1
211.206.58.76	0	0	0	0	1
122.106.153.166	0	0	0	0	1
70.55.149.222	0	0	0	0	2
96.244.84.2	0	0	0	0	1

- To change how often the table is refreshed, select an **Auto-Refresh Rate** from the list.

Display the report in the Exinda CLI

- With your browser, open the Exinda Web UI.

`https://UI_IP_address.`

- Type the **User Name** and **Password**.
- Click **Login**.
- The Exinda Web UI appears.*
- Ensure you are in **Advanced** mode.
- Click **Configuration > System > Tools > Console**.
- Type the appliance username and password at the prompts. Do *one* of the following:
 - To enter privileged EXEC (enable) mode, at the prompt type the following command:

`hostname > enable`

The hostname # prompt appears.

- To enter configuration (config) mode, at the prompt type the following commands:

`hostname # configure terminal`

The hostname (config)# prompt appears.

3. To display realtime TCP health from the CLI, use the following command:

```
(config) # show realtime apm hosts
```

The following results are displayed:

```
ex-240 (config) # show realtime apm hosts

Internal
Host          Retransmissions Aborted Refused Ignored Flows
-----
172.16.1.240  0           0      0      0      13
192.168.0.176 0           0      0      0      1
172.16.0.213  0           0      0      0      1
192.168.50.147 0          0      0      0      1
192.168.0.179 0           0      0      0      1
172.16.0.63   0           0      2      0      3
172.16.1.242  0           0      0      0      1
192.168.40.96 0           0      0      0      1
192.168.0.178 0           0      0      0      6
0.0.0.0       0           0      0      0      1
192.168.0.209 0           0      0      0      1
192.168.50.143 0          0      0      0      1
172.16.0.252  0           0      0      0      1
172.16.0.108  0           0      0      0      3
172.16.1.149  0           0      0      0      3
172.16.0.67   0           0      0      0      5
172.16.0.190  0           1      0      0      4
192.168.0.118 0           0      0      0      1
192.168.0.145 0           0      0      0      1
192.168.0.207 0           0      0      0      1
```

Generate a PDF report of APS results

Create a report that contains the APS, TCP health, and TCP efficiency for a specified period of time.

1. With your browser, open the Exinda Web UI.

`https://UI_IP_address.`

2. Type the **User Name** and **Password**.
3. Click **Login**.
The Exinda Web UI appears.
4. Ensure you are in **Advanced** mode.
5. Click **Report** and switch to the **PDF Reports** tab.
6. Click **Add New PDF Report**.
7. In the Report Selection area select **APS, TCP Health**, and **TCP Efficiency**.
8. In the Report Details area, type a name for the report.
9. Specify how often the report will be generated.
10. Click **Add New Report**.
11. To generate the report, locate the report in the list and click **PDF**.