

# QFX5100 Guest VM: Data Monitoring Application Setup Guide

Revision 2

## Table of Contents

1	Objectives.....	2
2	Scope.....	2
3	Steps to build the Data Monitoring Application .....	2
3.1	Step 1: Guest VM creation using a JunosV App Engine configuration.....	3
3.2	Step 2: Junos AnalyticsD Configuration .....	3
3.3	Step 3: Installation of Fluentd, Elasticsearch and AnalyticsD plugin .....	4
3.3.1	Fluentd Installation and Configuration .....	5
3.3.2	Elasticsearch Installation and Configuration.....	6
3.3.3	AnalyticsD Plugin Installation.....	7
3.4	Step 4: Launch Fluentd and Elasticsearch on Guest VM.....	11
3.5	Step 5: Configure Kibana on the PC to view the Traffic/Queue statistics .....	11
4	Appendix A – Code Samples .....	13
5	Appendix B – Additional Sample Output .....	17

## 1 Objectives

Using combinations of data acquired from the infrastructure's components and network elements, Junos Network Analytics applies sophisticated data analysis algorithms, control systems techniques, and reporting in order to provide visibility into the aspects of the data center infrastructure's performance and behavior. These aspects can now be directly visualized through a Guest VM application, called the Data Monitoring Application.

This application can be used to process and report the traffic and queue statistics collected from Junos. The analyzed report can then be visualized graphically using third-party visualization tools (like Kibana) running on a desktop or laptop.

The Data Monitoring application uses the Guest VM infrastructure and internally relies on the below third-party tools/libraries:

- Fluentd - Open source data collector designed for processing data streams
- Elasticsearch – Provides distributed RESTful search and analytics
- Kibana - Visualize logs and time-stamped data

This document aims to serve as a step-by-step installation and configuration guide for the Data Monitoring application.

## 2 Scope

This document describes how to leverage QFX5100 Guest VM infrastructure to build applications using third-party software. Juniper Networks is not liable for support of any issues reported with third-party software.

## 3 Steps to build the Data Monitoring Application

The following are the steps that need to be followed for building and using the Data Monitoring application:

1. Guest VM creation using JunosV App Engine configuration
2. Junos AnalyticsD configuration
3. Installation of Fluentd, Elasticsearch and AnalyticsD plugin on the Guest VM created in Step 1
4. Launching Fluentd and Elasticsearch on the Guest VM
5. Configuring Kibana on the PC or laptop to view the Traffic/Queue statistics graphically

### 3.1 Step 1: Guest VM creation using a JunosV App Engine configuration

A Guest VM needs to be created by building a Guest VM package, installing the Guest VM package with the help of the Junos OS CLI, and committing the Guest VM configurations in the Junos configuration mode.

The following document provides a step-by-step guide for creating a Guest VM and can be followed for creating the Guest VM:

[http://www.juniper.net/techpubs/en\\_US/junos13.2/topics/task/installation/qfx-series-guest-vmm.html](http://www.juniper.net/techpubs/en_US/junos13.2/topics/task/installation/qfx-series-guest-vmm.html)

### 3.2 Step 2: Junos AnalyticsD Configuration

For the purpose of reporting the Traffic and Queue statistics from the Junos VM to the Guest VM, CLI options must be configured under the “`set services analytics`” hierarchy in the Junos configuration mode.

The following is the sample configuration that can be used.

```
set groups datamon services analytics export-profiles p1 stream-format json
set groups datamon services analytics export-profiles p1 interface statistics
traffic
set groups datamon services analytics export-profiles p1 interface statistics
queue
set groups datamon services analytics export-profiles p1 system information
set groups datamon services analytics export-profiles p1 system status
traffic
set groups datamon services analytics export-profiles p1 system status queue
set groups datamon services analytics resource-profiles rp1 queue-monitoring
set groups datamon services analytics resource-profiles rp1 traffic-
monitoring
set groups datamon services analytics resource system resource-profile rp1
set groups datamon services analytics resource system polling-interval
traffic-monitoring 10
set groups datamon services analytics resource system polling-interval queue-
monitoring 10
```

```
set groups datamon services analytics collector local file datamon.stat

set groups datamon services analytics collector local file size 10m

set groups datamon services analytics collector address 10.204.38.76 port
24224 transport tcp export-profile p1

set apply-groups datamon
```

**NOTE:** In the above configuration, in “`set groups datamon services analytics collector address 10.204.38.76 port 24224 transport tcp export-profile p1`,” the IP address `10.204.38.76` needs to be replaced with the IP address that will be configured on the management interface of the Guest VM. The port number `24224` should be kept as it is as the same is used by the third party application Fluentd running on the Guest VM.

Please note that the traffic and queue monitoring intervals can also be modified in the configuration depending on your requirements.

- “`set groups datamon services analytics resource system polling-interval traffic-monitoring 10`”
- “`set groups datamon services analytics resource system polling-interval queue-monitoring 10`”

### 3.3 Step 3: Installation of Fluentd, Elasticsearch and AnalyticsD plugin

- Log into the Guest VM and assign the IP address on the external Management interface of the Guest VM. For example, if `eth6` is the external management interface on the Guest VM, then

```
ifconfig eth6 inet <IP-address> netmask <netmask>
```

- Configure the default gateway on the Guest VM

```
route add default gw <gateway>
```

- After the above configurations are done, PING from the Guest VM to the Junos VM should start working. Confirm this by executing the following command on the Guest VM:

```
ping <Junos-VM-IP-address>
```

where `<Junos-VM-IP-address>` is the IP address configured on the Management Interface assigned to the Junos VM.

**NOTE:** Third party applications need to be downloaded through the Internet, so please make sure that Internet connectivity is enabled on the above IP address configured for the Guest VM external management interface.

### 3.3.1 Fluentd Installation and Configuration

1. Install the Fluentd Td-agent:

```
curl -L http://toolbelt.treasuredata.com/sh/install-redhat.sh | sh
```

2. Install the Elasticsearch plugin for Fluentd Td-agent:

```
cd ~
```

```
curl -L https://github.com/ukenu/fluent-plugin-elasticsearch/archive/v0.2.0.zip > ~/master.zip
```

```
yum -y install unzip
```

```
unzip master.zip
```

```
cd fluent-plugin-elasticsearch-0.2.0
```

```
vi fluent-plugin-elasticsearch.gemspec and modify the field s.files to set it to ["lib/fluent/plugin/out_elasticsearch.rb"]
```

```
/usr/lib64/fluent/ruby/bin/fluent-gem build fluent-plugin-elasticsearch.gemspec
```

```
/usr/lib64/fluent/ruby/bin/fluent-gem install fluent-plugin-elasticsearch-0.2.0.gem --no-rdoc --no-ri
```

Execute the following command to confirm that fluent-plugin-elasticsearch has been installed.

```
/usr/lib64/fluent/ruby/bin/fluent-gem list | grep elasticsearch
```

```
cd ~
```

```
rm -rf master.zip
```

3. Obtain the contents of `td-agent.conf` from Appendix A and set up the Td-agent configuration for AnalyticsD:

```
vi ~/td-agent.conf
```

Next, copy and paste the contents of `td-agent.conf` from Appendix A into the above created new file `~/td-agent.conf`.

Replace the Td-agent configuration with the above `~/td-agent.conf` with the following command:

```
cp ~/td-agent.conf /etc/td-agent/td-agent.conf
```

4. Make sure `/var/log/fluentd` exists and is owned by the Td-agent:

```
mkdir /var/log/fluentd  
chown td-agent /var/log/fluentd  
chgrp td-agent /var/log/fluent
```

5. Additional details on Fluentd can be found at:

- <http://www.fluentd.org/>
- <https://github.com/fluent/fluentd>

### 3.3.2 Elasticsearch Installation and Configuration

1. Install JVM version 6 or higher, as this is a prerequisite for Elasticsearch:

```
yum install java-1.6.0-openjdk
```

2. Download and install Elasticsearch and its service wrapper:

```
cd ~  
  
curl https://download.elasticsearch.org/elasticsearch/elasticsearch/elasticsearch-1.0.1.tar.gz | tar xvz  
  
mkdir /etc/elasticsearch  
  
mv ~/elasticsearch-1.0.1/* /etc/elasticsearch  
  
rm -rf ~/elasticsearch-1.0.1  
  
yum -y install unzip  
  
curl -L https://github.com/elasticsearch/elasticsearch-servicewrapper/archive/master.zip > master.zip  
  
unzip master.zip  
  
cd elasticsearch-servicewrapper-master/  
  
mv service /etc/elasticsearch/bin
```

```
cd ~
```

```
rm -rf elasticsearch-servicewrapper-master master.zip
```

```
/etc/elasticsearch/bin/service/elasticsearch install
```

3. Obtain the contents of `default-mapping.json` from Appendix A and place it into Elasticsearch:

```
vi ~/default-mapping.json
```

Next, copy and paste the contents of `default-mapping.json` from Appendix A into the above created new file `~/default-mapping.json`.

```
cp ~/default-mapping.json /etc/elasticsearch/.
```

4. Edit the `cluster.name` entry in the Elasticsearch yml file to give Elasticsearch a name:

```
vi /etc/elasticsearch/config/elasticsearch.yml
```

Edit the `cluster.name` entry (for example: `cluster.name: DataMonitoringVM!`).

```
cp ~/elasticsearch.yml /etc/elasticsearch/config/elasticsearch.yml
```

5. Elasticsearch is now ready to use and the following commands can be executed when needed:

```
service elasticsearch start
```

```
service elasticsearch stop
```

```
service elasticsearch restart
```

```
service elasticsearch status
```

6. Additional details on Elasticsearch can be found at:

- <http://www.elasticsearch.org/>
- <https://github.com/elasticsearch/elasticsearch>

### 3.3.3 AnalyticsD Plugin Installation

1. Edit `~/bash_profile` to add a few paths for ruby gems:

```
PATH=$PATH:$HOME/bin:/usr/share/elasticsearch/bin:/usr/lib64/fluent/
ruby/bin/

GEM_HOME="/usr/lib64/fluent/ruby/lib/ruby/gems/1.9.1/"

GEM_PATH="/usr/lib64/fluent/ruby/lib/ruby/gems/1.9.1/"

export PATH

export GEM_PATH

export GEM_HOME
```

2. Export the paths by executing the following command:

```
source ~/.bash_profile
```

3. Generate the file `in_AnalyticsD.rb` and place that in the home directory:

```
mkdir ~/AnalyticsD_utility

cd ~/AnalyticsD_utility
```

Place the files `build-plugin` and `build-wrapper` (provided along with this setup guide) in the current directory. Make sure that these two files are executable:

```
chmod +x build-plugin

chmod +x build-wrapper
```

Obtain the file `in_forward.rb` from Github:

```
curl -L https://github.com/fluent/fluentd/archive/skip-nil-record.zip > skip-
nil-record.zip

unzip skip-nil-record.zip

cp fluentd-skip-nil-record/lib/fluent/plugin/in_forward.rb .
```

Install `csch` if it is not already installed on the Guest VM:

```
yum install csch
```

Generate `in_AnalyticsD.rb` using `in_forward.c` and `build-wrapper`:

```
./build-wrapper in_forward.rb
```

Open the file `in_AnalyticsD.rb` and modify lines 161 and 162:

```

159 record = msg
160 return if record.nil?
161 time = msg[1] <<<<< Remove this statement
162 time = Engine.now if time == 0 <<<<< Remove the condition "if time == 0"
163 Engine.emit(tag, time, record)

```

Thus, the above code snippet in `in_AnalyticsD.rb` would now appear as:

```

159 record = msg
160 return if record.nil?
161 time = Engine.now
162 Engine.emit(tag, time, record)

```

Copy `in_AnalyticsD.rb` to the home directory:

```
cp in_AnalyticsD.rb ~/
```

Clean up the directory `AnalyticsD_utility`:

```
rm -rf AnalyticsD_utility
```

4. Obtain the contents for `fluent-plugin-AnalyticsD-0.1.0.gemspec` from Appendix A and create this file in the home directory:

```
vi ~/fluent-plugin-AnalyticsD-0.1.0.gemspec
```

Copy the contents for this file from Appendix A and paste the same into this new file.

5. Build a directory structure for the Juniper Networks AnalyticsD plugin:

```
mkdir ~/fluent-plugin-AnalyticsD-0.1.0
mkdir ~/fluent-plugin-AnalyticsD-0.1.0/lib
mkdir ~/fluent-plugin-AnalyticsD-0.1.0/lib/fluent
mkdir ~/fluent-plugin-AnalyticsD-0.1.0/lib/fluent/plugin
cp ~/in_AnalyticsD.rb ~/fluent-plugin-AnalyticsD-0.1.0/lib/fluent/plugin/.
cp ~/fluent-plugin-AnalyticsD-0.1.0.gemspec ~/fluent-plugin-AnalyticsD-0.1.0/.
```

6. Build the AnalyticsD gem:

```
cd ~/fluent-plugin-AnalyticsD-0.1.0
fluent-gem build fluent-plugin-AnalyticsD-0.1.0.gemspec
```

"NOTE: a warning like the below might be produced. It is harmless.

WARNING: no homepage specified

Successfully built RubyGem

Name: fluent-plugin-AnalyticsD

Version: 0.1.0

File: fluent-plugin-AnalyticsD-0.1.0.gem

"

7. Install the AnalyticsD gem:

```
fluent-gem install fluent-plugin-AnalyticsD-0.1.0.gem --no-rdoc --no-ri
```

"NOTE: a warning like the below will be produced. It is harmless.

WARNING: Error fetching data: Errno::ETIMEDOUT: Connection timed out - connect(2) (<http://rubygems.org/specs.4.8.gz>)

Successfully installed fluent-plugin-AnalyticsD-0.1.0

1 gem installed

"

8. Execute the following command to view the list of installed gems and verify that the AnalyticsD gem has been successfully installed:

```
fluent-gem list
```

### 3.4 Step 4: Launch Fluentd and Elasticsearch on Guest VM

Execute the following commands on the Guest VM in order to launch Elasticsearch and then Fluentd:

```
service iptables stop
```

```
service elasticsearch stop
```

```
service elasticsearch start
```

```
service td-agent restart
```

After the above commands are executed on the Guest VM, the command “**show analytics collector**” in the Junos OS CLI should confirm that the connection to the Guest VM has been established.

For example:

```
root@qfx5100-c6-sys> show analytics collector
```

Address	Port	Transport	Stream format	State	Sent
10.204.38.76	24224	tcp	json	Established	2

### 3.5 Step 5: Configure Kibana on the PC to view the Traffic/Queue statistics

Download Kibana from the following website and place it on the Desktop:

<https://download.elasticsearch.org/kibana/kibana/kibana-3.0.0.zip>

Unzip the above file. Then a directory structure will get created. In the top level of the directory structure you will find a "config.js" file. Put a line similar to

```
elasticsearch: "http://<IP_ADDRESS>:9200"
```

Here, **IP\_ADDRESS** is the external Management IP address of the Guest VM. Here the port number **9200** is associated to Elasticsearch and should be kept as it is.

Start up either Microsoft Internet Explorer or Mozilla Firefox (these are the only browsers supported). Point the browser to a Kibana dashboard to visualize the traffic and queue statistics. For example, if the Kibana directory structure for a user <username> is located at “[C:/Users/<username>/Desktop/kibana-3.0.0milestone4](file:///C:/Users/<username>/Desktop/kibana-3.0.0milestone4)”, the browser can be pointed to the Kibana dashboard *guided.json* by specifying the path as:

<file:///C:/Users/<username>/Desktop/kibana-3.0.0milestone4/index.html#/dashboard/file/guided.json>

### Sample output taken from Kibana

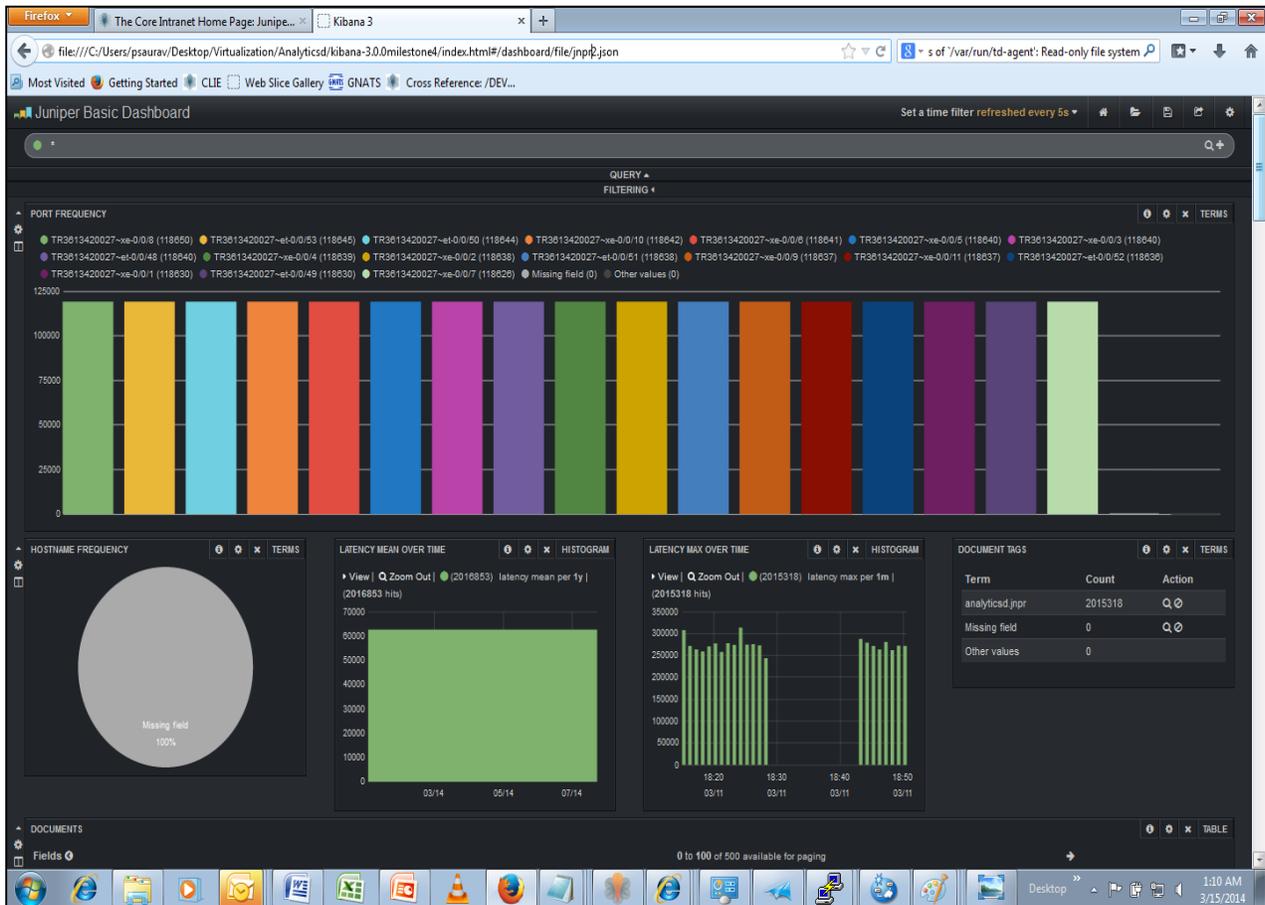


Figure 1. Sample Output from Kibana

For additional sample output, see Appendix B – Additional Sample Output.

Additional details on Kibana can be found at:

- <http://www.elasticsearch.org/overview/kibana/>
- <https://github.com/elasticsearch/kibana>

## 4 Appendix A – Code Samples

### Contents of td-agent.conf

```
####  
  
## Output descriptions:  
  
##  
  
<source>  
  
    type monitor_agent  
  
    bind 0.0.0.0  
  
    port 24220  
  
</source>  
  
<match td.*.*>  
  
    type tdlog  
  
    apikey YOUR_API_KEY  
  
  
    auto_create_table  
  
    buffer_type file  
  
    buffer_path /var/log/td-agent/buffer/td  
  
</match>  
  
  
<match debug.**>  
  
    type file  
  
    path /var/log/fluentd/debug  
  
    flush_interval 1s  
  
</match>  
  
  
####  
  
## Source descriptions:
```

##

```
<source>
  type AnalyticsD
  port 24224
  bind 0.0.0.0
</source>
<match analyticsD.jnpr>
  type elasticsearch
  logstash_format true
  include_tag_key true
  tag_key _jnpr
  flush_interval 1s
</match>
```

```
<source>
  type http
  port 8888
  bind 0.0.0.0
</source>
<match curl.test>
  type elasticsearch
  logstash_format true
  include_tag_key true
  flush_interval 1s
</match>
```

```
<source>
```

```

type debug_agent

bind 127.0.0.1

port 24230

</source>

```

### **Contents of default-mapping.json**

```

{
  "_default_" : {
    "properties" : {
      "port" : { "type":"string","index":"not_analyzed"},
      "hostname" : { "type":"string","index":"not_analyzed"},
      "router-id" : { "type":"string","index":"not_analyzed"},
      "queue-depth" : { "type":"string","index":"not_analyzed"},
      "record-type" : { "type":"string","index":"not_analyzed"},
      "zhtpt" : { "type":"string","index":"not_analyzed"}
    }
  }
}

```

### **Contents of fluent-plugin-AnalyticsD-0.1.0.gemspec**

```

Gem::Specification.new do |gem|

  gem.name          = 'fluent-plugin-AnalyticsD'
  gem.version       = '0.1.0'
  gem.date          = '2014-01-08'
  gem.summary       = "JNPR!"
  gem.description   = "Fluentd Input Plugin for JNPR Devices"
  gem.authors       = ["Jay A. Wilson"]
  gem.email         = 'jayw@juniper.net'
  gem.files         = ["lib/fluent/plugin/in_AnalyticsD.rb"]

```

```
gem.require_paths = ["lib"]  
  
gem.has_rdoc      = false  
  
gem.required_ruby_version = '>= 1.9.2'  
  
gem.add_dependency "fluentd", "~> 0.10.7"  
  
end
```



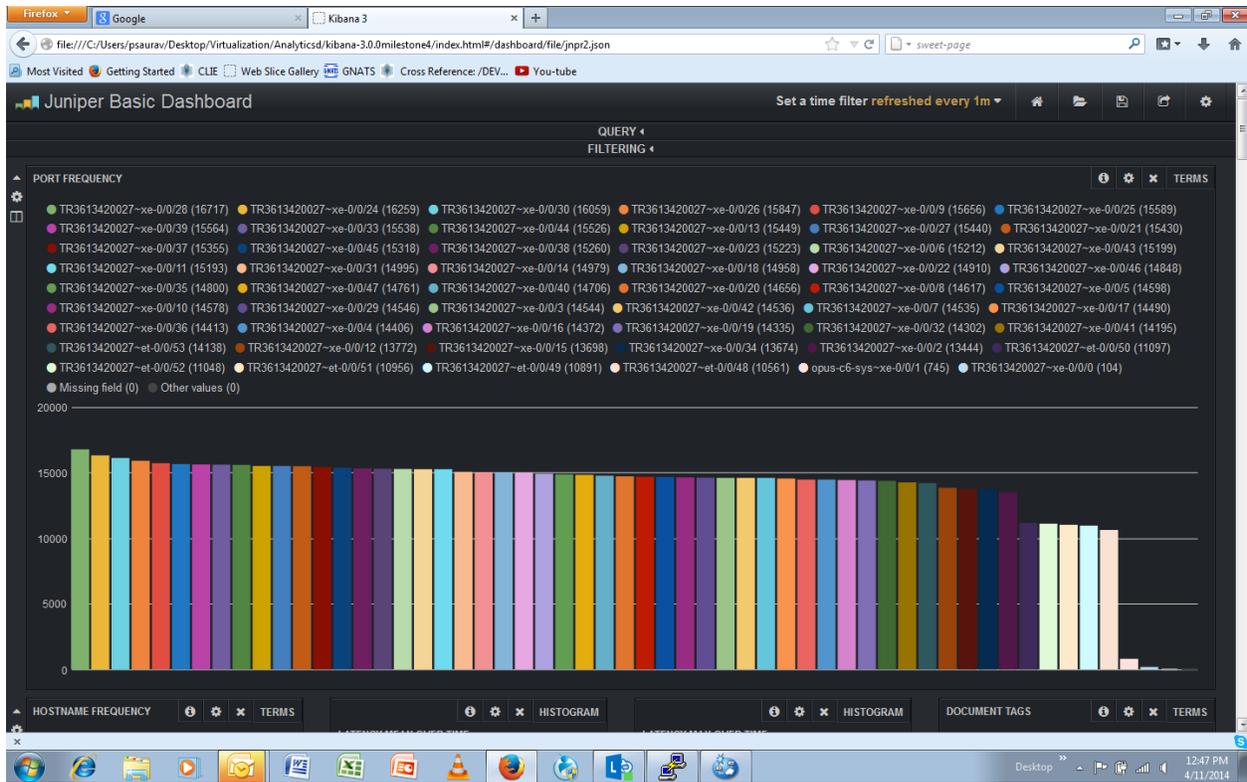


Figure 3. Statistics for Various Ports

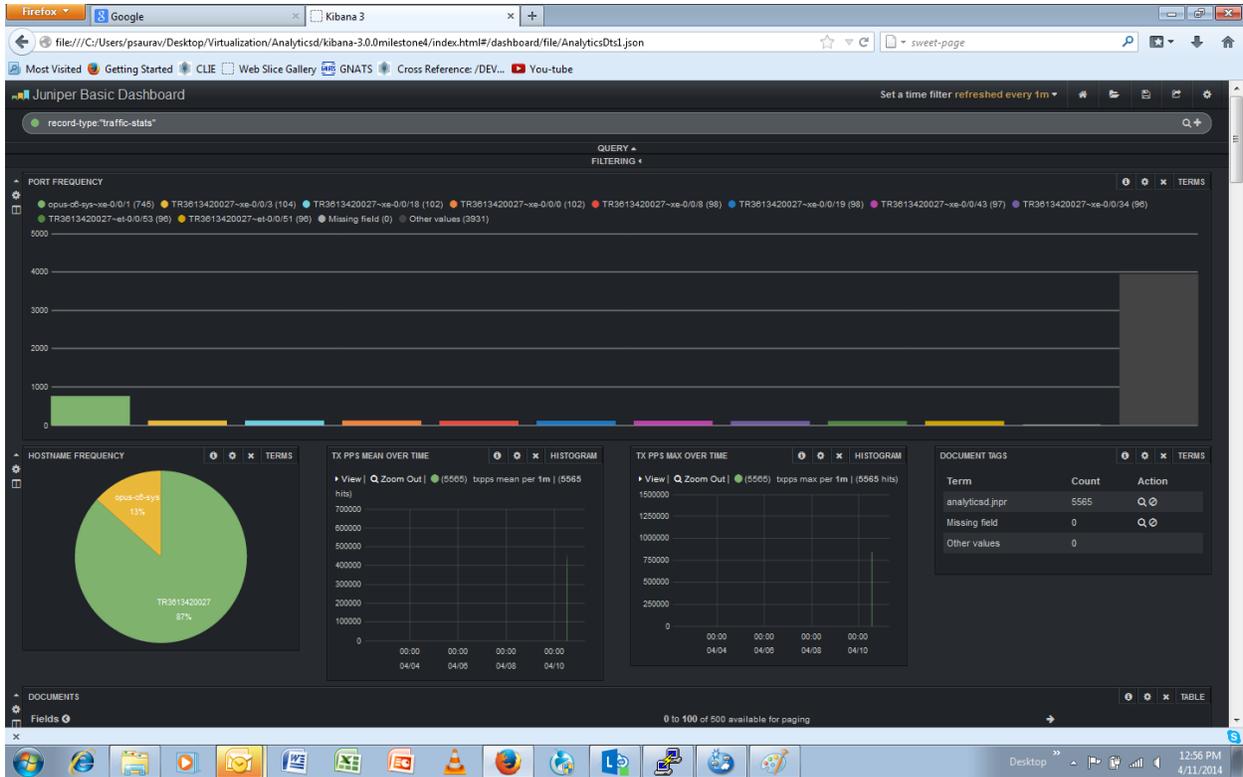


Figure 4. Hostname Frequency and Tx pps Mean/Max Over Time

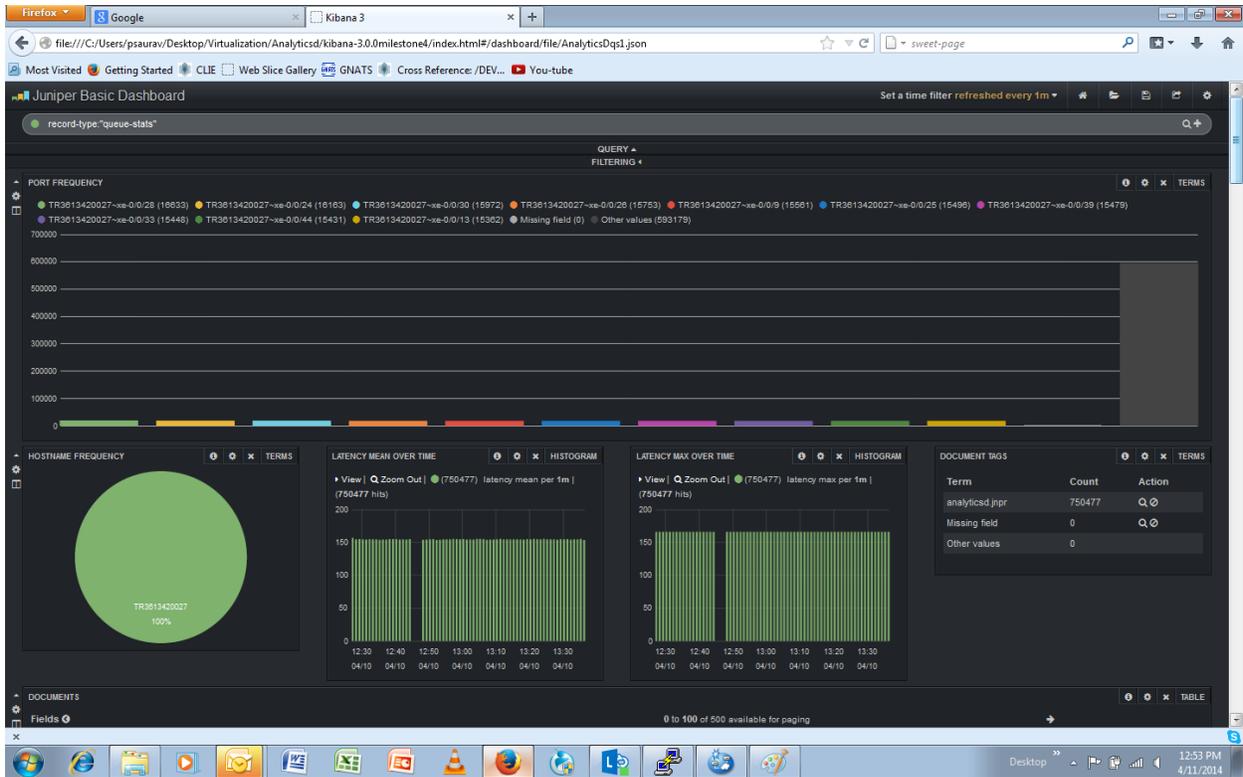


Figure 5. Latency Mean/Max Over Time

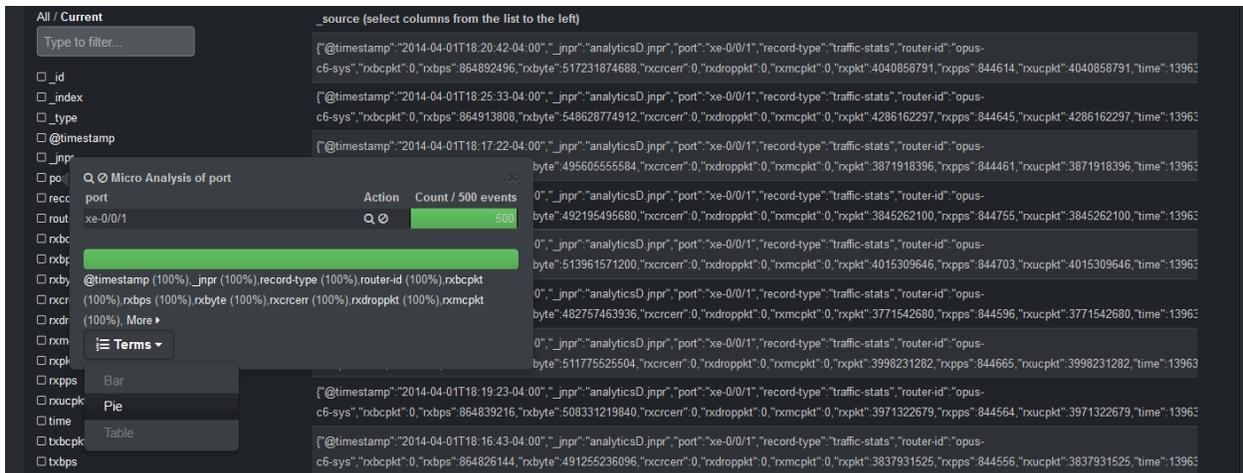


Figure 6. Micro Analysis of a Given Port

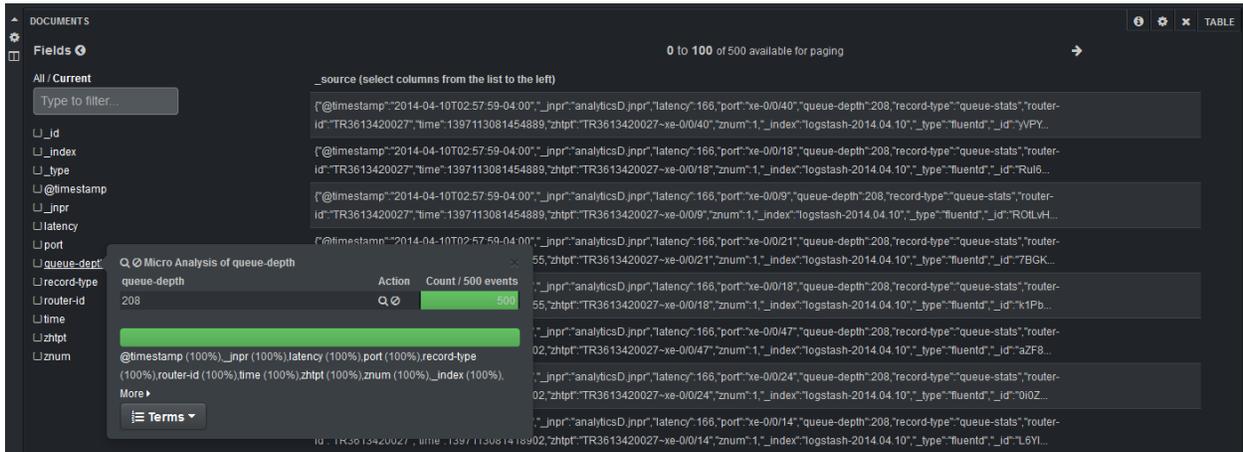


Figure 7. Micro Analysis of Queue Depth

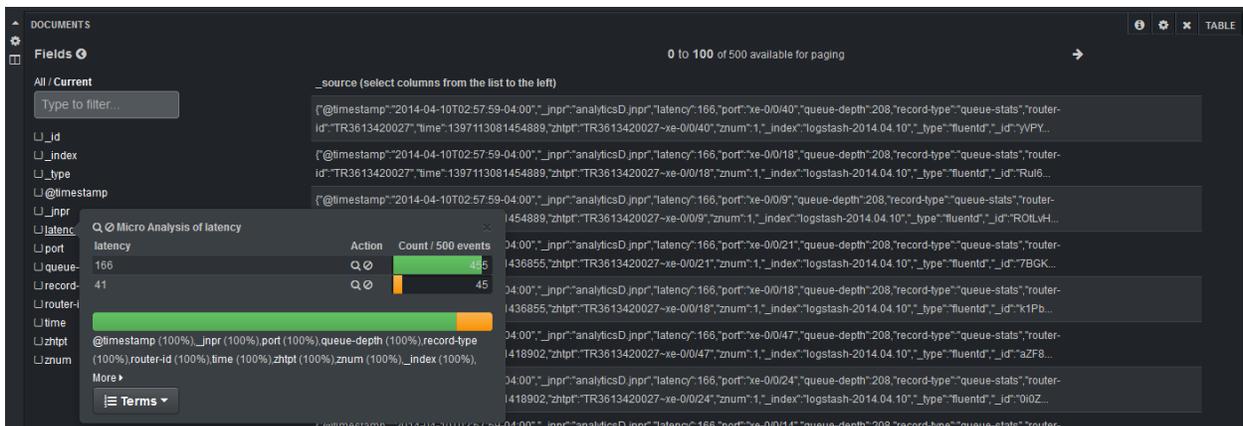


Figure 8. Micro Analysis of Latency



Figure 9. Queue Records

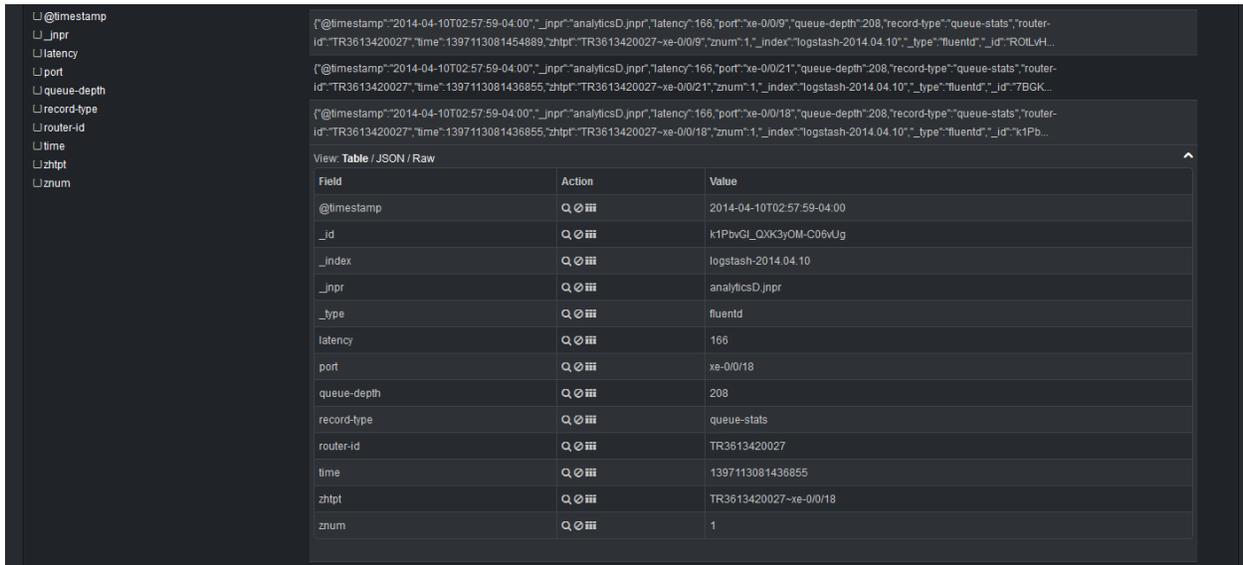


Figure 10. Queue Record Expanded

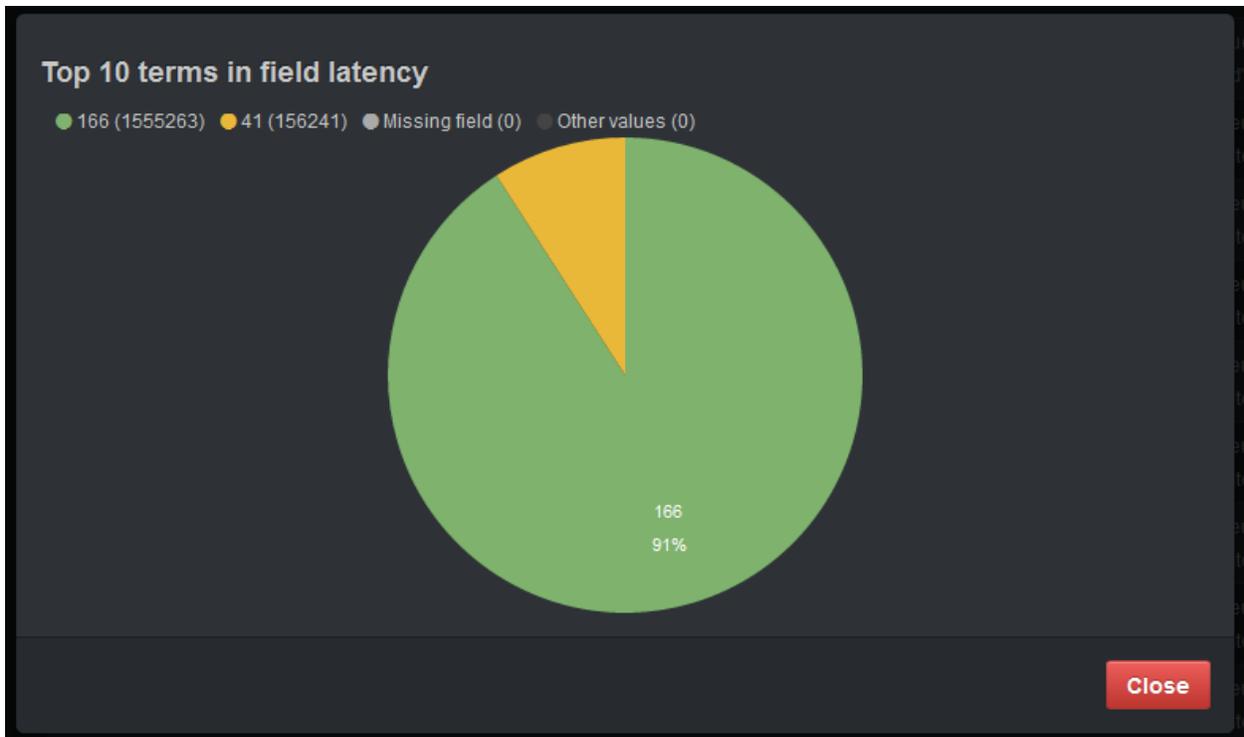


Figure 11. Top Ten Latency in a Pie Chart

The table provides a detailed view of the latency terms. It lists the term, its count, and an action icon (magnifying glass and trash can) for each entry. The data is consistent with the pie chart above.

Term	Count	Action
166	682337	🔍 🗑️
41	68140	🔍 🗑️
Missing field	0	🔍 🗑️
Other values	0	

Figure 12. Top Ten Latency in a Table

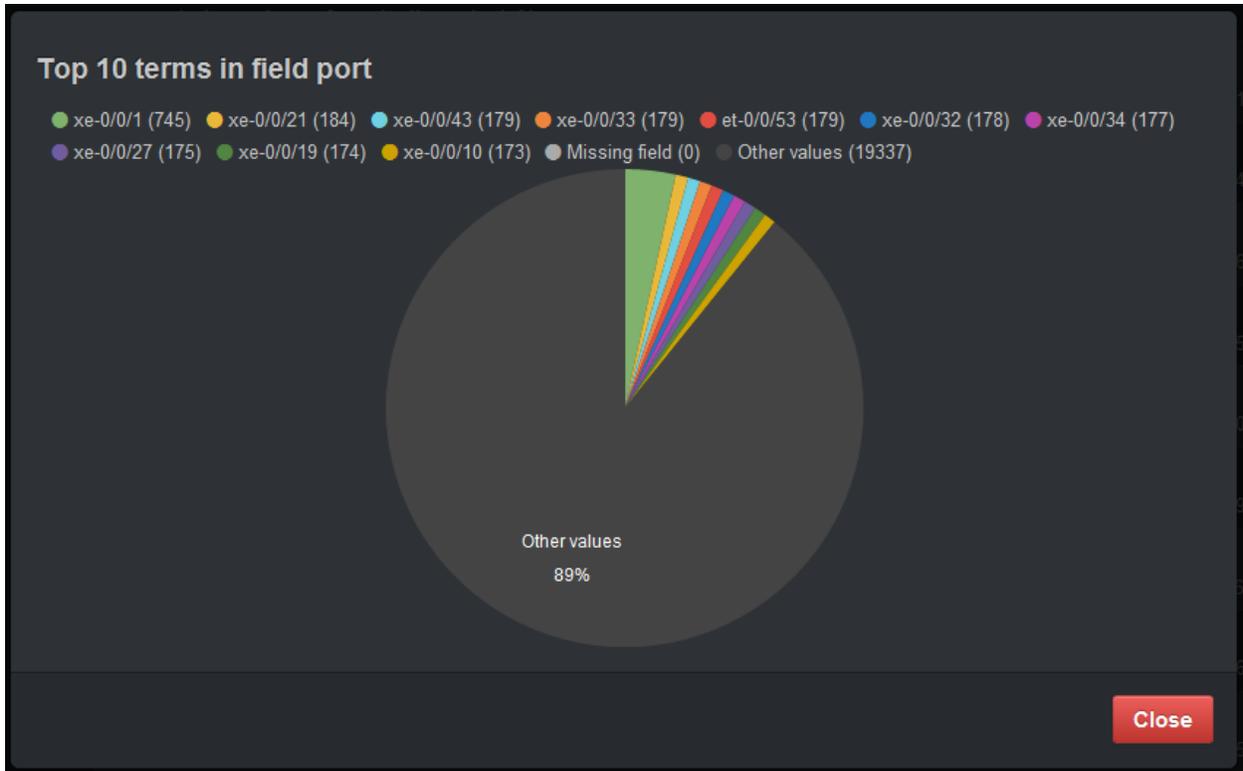
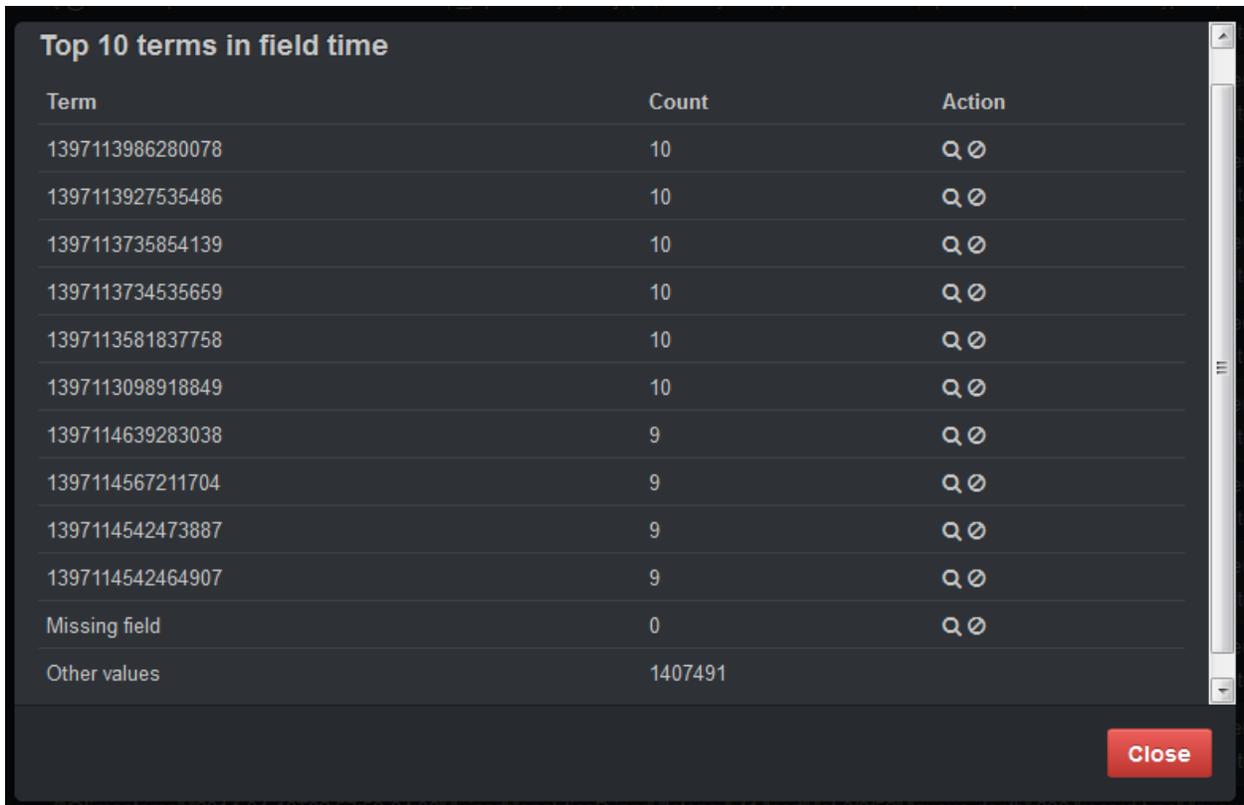


Figure 13. Top Ten Ports in a Pie Chart



Term	Count	Action
1397113986280078	10	🔍 🗑️
1397113927535486	10	🔍 🗑️
1397113735854139	10	🔍 🗑️
1397113734535659	10	🔍 🗑️
1397113581837758	10	🔍 🗑️
1397113098918849	10	🔍 🗑️
1397114639283038	9	🔍 🗑️
1397114567211704	9	🔍 🗑️
1397114542473887	9	🔍 🗑️
1397114542464907	9	🔍 🗑️
Missing field	0	🔍 🗑️
Other values	1407491	

Close

Figure 14. Top Ten Timestamps in a Table



Term	Count	Action
0	9914	Q ∅
624	384	Q ∅
280	272	Q ∅
272	266	Q ∅
616	237	Q ∅
632	205	Q ∅
288	169	Q ∅
296	151	Q ∅
264	115	Q ∅
640	84	Q ∅
Missing field	0	Q ∅
Other values	10784	

Figure 15. Top Ten Tx bps in a Table

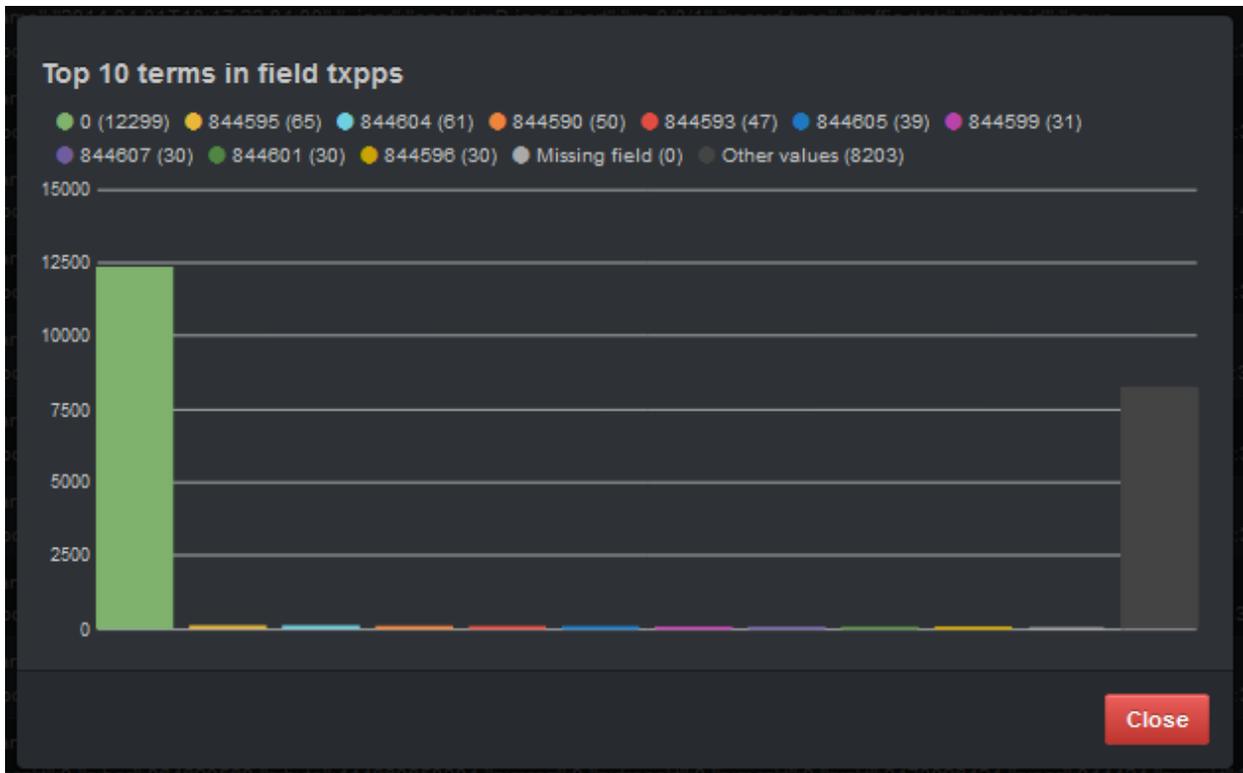


Figure 16. Top Ten Tx pps in a Bar Graph

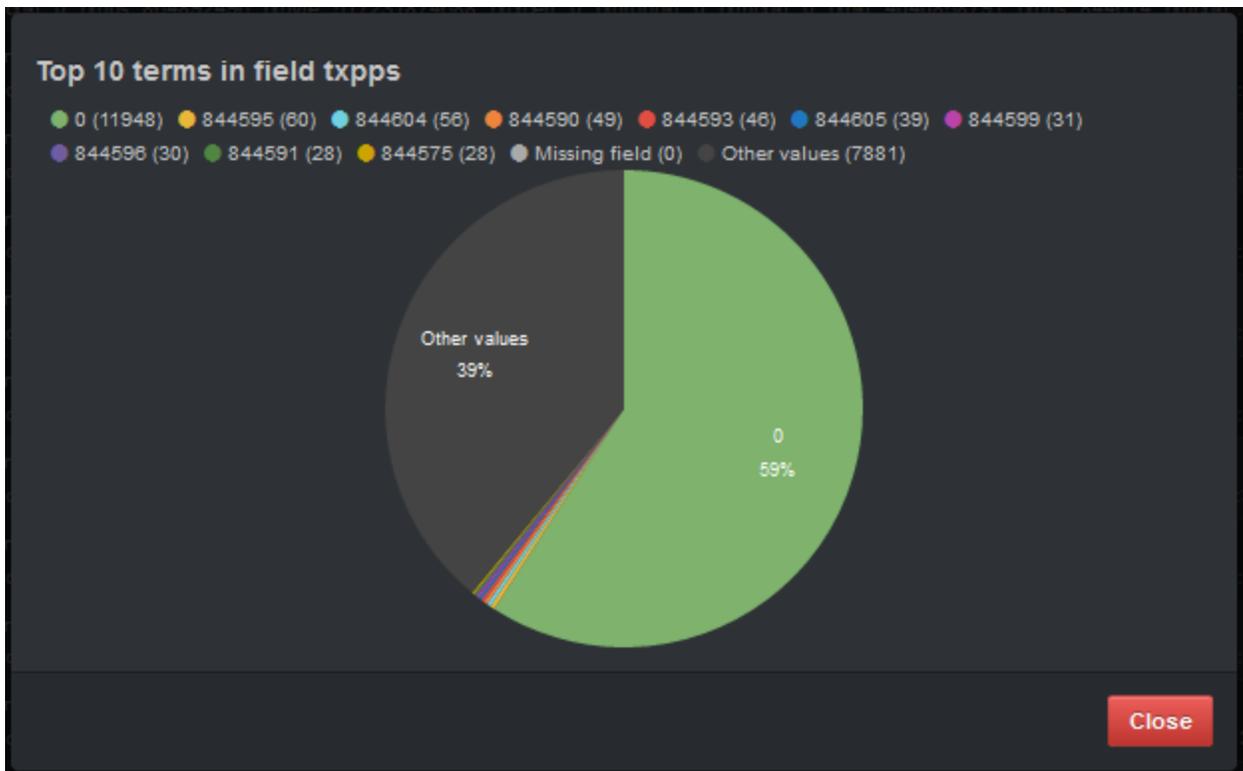
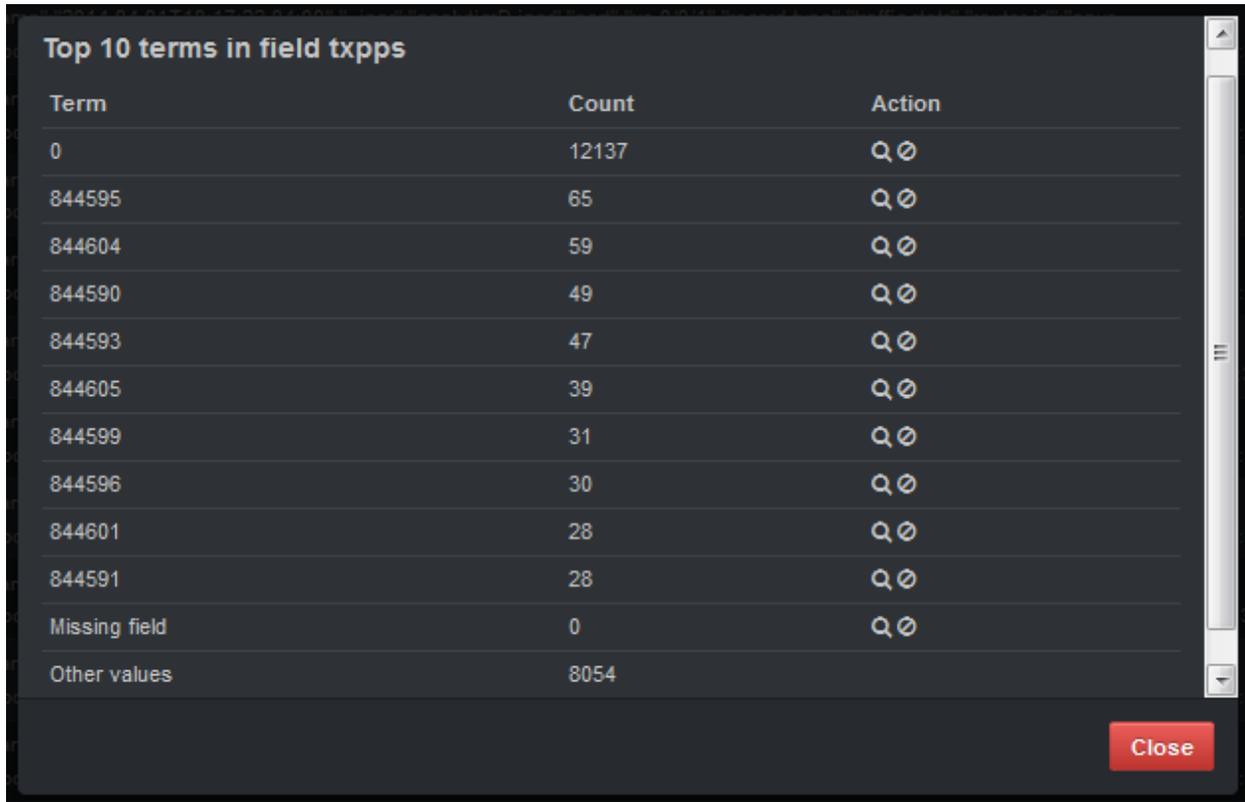


Figure 17. Top Ten Tx pps in a Pie Chart



Term	Count	Action
0	12137	🔍
844595	65	🔍
844604	59	🔍
844590	49	🔍
844593	47	🔍
844605	39	🔍
844599	31	🔍
844596	30	🔍
844601	28	🔍
844591	28	🔍
Missing field	0	🔍
Other values	8054	

Figure 18. Top Ten Tx pps in a Table

The screenshot shows the 'Micro Analysis of rxbps' table in the Data Monitoring Application. The table has columns for 'Action' and 'Value'. The data rows include:

Action	Value
Q Q !!!	2014-04-01T18:17:22-04:00
Q Q !!!	3WFJHH2SZONjuZGMFe1qA
Q Q !!!	logatash-2014.04.01
Q Q !!!	analyticsD.jnpr
Q Q !!!	fluentId
Q Q !!!	xe-0/0/1
Q Q !!!	traffic-stats
Q Q !!!	opus-c6-sys
Q Q !!!	0
Q Q !!!	864729184
Q Q !!!	49560555584
Q Q !!!	0
Q Q !!!	0

Figure 19. Micro Analysis of Rx bps

The screenshot shows a list of traffic records in the Data Monitoring Application. The records are displayed as JSON-like strings with fields like @timestamp, \_jnpr, port, record-type, router-id, opus-c6-sys, rxbcpkt, rxbps, rxbyte, rxrcrcr, rxdroppkt, rxmcpkt, and rxpkt. The records are sorted by time, showing various timestamps and associated data points.

Figure 20. Traffic Records

The screenshot shows a table view of traffic records. On the left, there is a list of checkboxes for various traffic-related fields. The main table has columns for Field, Action, and Value. The fields listed include timestamps, IDs, indices, JNPR IDs, types, ports, record types, router IDs, and various traffic metrics like ndropkt, ndbps, ndbyte, nroerr, ndroppkt, nrcmpkt, ncpkt, npps, nrcpmt, time, and various 'b' (byte) and 'p' (packet) metrics.

Field	Action	Value
@timestamp		2014-04-01T18:19:23-04:00
_id		N431eXcHRBmZdHhIdz2qQ
_index		logstash-2014.04.01
_jnpr		analyticD_jnpr
_type		fluentd
port		xe-0/0/1
record-type		trafficstats
router-id		opus-05-5ys
ndropkt		0
ndbps		804839216
ndbyte		508331219840
nroerr		0
ndroppkt		0
nrcmpkt		0
ncpkt		3971322679
npps		844564
nrcpmt		3971322679
time		1395390763898047
tbcpkt		0
tbbps		0
tbbyte		0
trroerr		0
tdroppkt		0
trcpkt		0
trpkt		0
trpps		0

Figure 21. Various Fields of a Traffic Record

---

This document describes how to leverage QFX5100 Guest VM infrastructure to build applications using third-party software. Juniper Networks is not liable for support of any issues reported with third-party software.

Juniper Networks, and Junos are registered trademarks of Juniper Networks, Inc. in the United States and other countries. The Juniper Networks Logo and the Junos logo are trademarks of Juniper Networks, Inc. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners.

Juniper Networks assumes no responsibilities for any inaccuracies that might appear in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.