



Security Threat Response Manager

AQL Flow and Event Query CLI Guide

Release 2009.1

Juniper Networks, Inc.

1194 North Mathilda Avenue

Sunnyvale, CA 94089

USA

408-745-2000

www.juniper.net

Published: 2009-10-12

Copyright Notice

Copyright © 2008 Juniper Networks, Inc. All rights reserved. Juniper Networks and the Juniper Networks logo are registered trademarks of Juniper Networks Inc. in the United States and other countries. All other trademarks, service marks, registered trademarks, or registered service marks in this document are the property of Juniper Networks or their respective owners. All specifications are subject to change without notice. Juniper Networks assumes no responsibility for any inaccuracies in this document or for any obligation to update information in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

FCC Statement

The following information is for FCC compliance of Class A devices: This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to part 15 of the FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. The equipment generates, uses, and can radiate radio-frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case users will be required to correct the interference at their own expense. The following information is for FCC compliance of Class B devices: The equipment described in this manual generates and may radiate radio-frequency energy. If it is not installed in accordance with NetScreen's installation instructions, it may cause interference with radio and television reception. This equipment has been tested and found to comply with the limits for a Class B digital device in accordance with the specifications in part 15 of the FCC rules. These specifications are designed to provide reasonable protection against such interference in a residential installation. However, there is no guarantee that interference will not occur in a particular installation. If this equipment does cause harmful interference to radio or television reception, which can be determined by turning the equipment off and on, the user is encouraged to try to correct the interference by one or more of the following measures: Reorient or relocate the receiving antenna. Increase the separation between the equipment and receiver. Consult the dealer or an experienced radio/TV technician for help. Connect the equipment to an outlet on a circuit different from that to which the receiver is connected.

Caution: Changes or modifications to this product could void the user's warranty and authority to operate this device.

Disclaimer

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR JUNIPER NETWORKS REPRESENTATIVE FOR A COPY.

Configuring DSMs
Release 2009.1

Copyright © 2008, Juniper Networks, Inc.

All rights reserved. Printed in USA.

Revision History

October 2009—Revision 1

The information in this document is current as of the date listed in the revision history.

CONTENTS

ABOUT THIS GUIDE

Conventions	3
Technical Documentation	3
Contacting Customer Support	4

1 USING THE AQL QUERY CLI

About the AQL Query CLI	5
Accessing the AQL Query CLI	6
Using a Select Statement	7
Using Where Clauses	11
Using the Group By Clause	12
Using the Order By Clause	13
Using the Count(*) Clause	13
Using the Distinct Clause	13
Using the Count (Distinct ...) Clause	14
Using the Materialize View Clause	14
Using the Like Clause	15




ABOUT THIS GUIDE

The *AQL Event and Flow Query CLI Guide* provides you with information for using the AQL CLI. This guide assumes you have advanced knowledge of Linux command line functionality.

Conventions

[Table 1](#) lists conventions that are used throughout this guide.

Table 1 Icons

Icon	Type	Description
	Information note	Information that describes important features or instructions.
	Caution	Information that alerts you to potential loss of data or potential damage to an application, system, device, or network.
	Warning	Information that alerts you to potential personal injury.

Technical Documentation

You can access technical documentation, technical notes, and release notes directly from the Juniper Customer Support web site at <https://www.juniper.net/support/>.

Once you access the Juniper Customer Support web site, locate the product and software release for which you require documentation. Your comments are important to us. Please send your e-mail comments about this guide or any of the Juniper Networks documentation to:

techpubs-comments@juniper.net.

Include the following information with your comments:

- Document title
- Page number

**Contacting
Customer Support**

To help you resolve any issues that you may encounter when installing or maintaining STRM, you can contact Customer Support as follows:

- Open a support case using the Case Management link at <http://www.juniper.net/support/>
- Call 1-888-314-JTAC (from the United States, Canada, or Mexico) or 1-408-745-9500 (from elsewhere).

1

USING THE AQL QUERY CLI

You can use the AQL Event and Flow Query Command Line Interface (CLI) to access flows and events stored in the Ariel database. This document provides information on accessing and using the AQL query CLI including:

- [About the AQL Query CLI](#)
- [Accessing the AQL Query CLI](#)
- [Using a Select Statement](#)
- [Using Where Clauses](#)
- [Using the Group By Clause](#)
- [Using the Order By Clause](#)
- [Using the Count\(*\) Clause](#)
- [Using the Distinct Clause](#)
- [Using the Count \(Distinct ...\) Clause](#)
- [Using the Materialize View Clause](#)
- [Using the Like Clause](#)

About the AQL Query CLI

The AQL event and flow query CLI allows you to access raw flows and events stored in the Aerial database. The AQL query CLI includes syntax that is a subset of the SQL92 standard and provides support for two tables: events and flows.



Note: *The AQL CLI does not provide support for joining tables.*

The AQL Event and Flow Query CLI functions in the following modes:

- **Interactive mode** - Using a simple shell, you can enter queries interactively and view the results in a standard output. At the query prompt, any valid AQL statement is accepted. If time is not specified (using `-start` and `-end` options), the last minute is assumed as the time range. You can also access previous commands by using your up arrow. This is the default mode.
- **Non-interactive mode** - You can enter the non-interactive mode by adding the `-execute <AQL query>` parameter to the command. The `-execute` command must be followed by a valid AQL query surrounded by double quotes. The non-interactive mode does not include a prompt allowing you to redirect the

output to a file with a regular UNIX pipe syntax. By default, the results are sent to a standard output.

Accessing the AQL Query CLI

To access the AQL query CLI:

- Step 1** Log in to STRM, as root.
Step 2 Enter the following command:

```
/opt/qradar/bin/arielClient
```

The Query prompt appears.

CLI Options [Table 1-1](#) lists the supported CLI options:

Table 1-1 AQL CLI Options

Option	Description
<code>-range <first record> <last record></code>	Limits the number of records sent to the output within the specified range. This is useful for viewing a selection of records generated by an ordered query. For example, if you wish to view the first ten records, you must specify <code>-range 1 10</code> .
<code>-debug</code>	Generates debugging output during execution.
<code>-start <time>, -end <time></code>	Specifies the start and end time of the query. Where <code><time></code> specifies the time. You must specify the time as either a UNIX timestamp or a date using the following format: <code>yyyy/mm/dd-hh:mm:ss</code> . For example: <pre>/opt/qradar/bin/arielClient - start 2007/08/11-01:15:00 -end 2007/08/11-01:17:00</pre>
<code>-exectime <time limit></code>	Specifies the maximum period of time, in seconds, a single query may continue processing.
<code>-execute <AQL query></code>	Allows you to enter non-interactive mode that allows you to process a query that is sent to standard output. If you do not include this option, the command is entered in interactive mode. You must include your query in double quotes.
<code>-f <output format></code>	Allows you to specify the output format for the query results. The table format is an ASCII drawing of a multi-column table while the CSV format provides a comma separated list. Where <code><output format></code> indicates the output format. The options are <code>table</code> or <code>csv</code> .
<code>-remote <host:port></code>	Specifies that you wish to connect to a specific Ariel query host and port.

For example:

If you wish to enter a command in interactive mode:

```
/opt/qradar/bin/arielClient -start 2007/08/11-01:15:00 -end
2007/08/11-01:17:00 -exectime 60
```

```
/opt/qradar/bin/arielClient
```

```
/opt/qradar/bin/arielClient -start 2007/08/11-01:15:00 -end
2007/08/11-01:17:00 -f csv
```

If you wish to enter a command in non-interactive mode:

```
/opt/qradar/bin/arielClient -start 2007/08/11-01:15:00 -end
2007/08/11-01:17:00 -exectime 60 -execute "select * from flows
where sourceIP = '231.12.37.17' and protocol != 'TCP.tcp_ip'"
```

Using a Select Statement

You can use a select statement that includes one or more fields of a flow or event. You can also use an asterisk (*) to denote all columns. All field names are case sensitive, however, the terms `select` and `from` are not case sensitive. The supported fields include:

Table 1-2 Supported Fields

Table	Supported Statement	Normalized Field Name
Flow	anyDestinationFlag	Destination Flags
	anySourceFlag	Source Flags
	application	Application
	applicationId	Application
	bytesIn	Bytes In
	bytesOut	Bytes Out
	destinationAssetName	Destination Asset Name
	destinationASN	Destination ASN
	destinationBytes	Destination Bytes
	destinationByteRatio	Destination Byte Ratio
	destinationDSCP	Destination DSCP
	destinationDscpOnly	Destination DSCP
	destinationFlags	Destination Flags
	destinationIP	Destination IP
	destinationIPSearch	Destination IP
	destinationIfIndex	Destination IF Index
	destinationNetwork	Destination Network
	destinationPackets	Destination Packets

Table 1-2 Supported Fields (continued)

Table	Supported Statement	Normalized Field Name
	destinationPacketRatio	Destination Packet Ratio
	destinationPayload	Destination Payload
	destinationPayloadHex	Destination Payload As Hex
	destinationPort	Destination Port
	destinationPrecedence	Destination Precedence
	destinationPrecedenceOnly	Destination Precedence
	destinationv6	IPv6 Destination
	firstPacketTime	First Packet Time
	flowBias	Flow Bias
	flowDirection	Flow Direction
	flowSource	Flow Source
	flowType	Flow Type
	geographic	Matches Geographic Location
	hasDestinationPayload	Has Destination Payload
	hasSourcePayload	Has Source Payload
	interface	Flow Interface
	lastPacketTime	Last Packet Time
	packetsIn	Packets In
	packetsOut	Packets Out
	protocol	Protocol
	protocolId	Protocol
	remoteHost	Remote Host
	remoteNet	Matches Remote Network
	remoteServices	Matches Remote Service
	sourceASN	Source ASN
	sourceAssetName	Source Asset Name
	sourceByteRatio	Source Byte Ratio
	sourceBytes	Source Bytes
	sourceDSCP	Source DSCP
	sourceDscpOnly	Source DSCP
	sourceFlags	Source Flags
	sourceIP	Source IP
	sourceIPSearch	Source IP
	sourceIfIndex	Source If Index
	sourceNetwork	Source Network
	sourcePacketRatio	Source Packet Ratio

Table 1-2 Supported Fields (continued)

Table	Supported Statement	Normalized Field Name
	sourcePackets	Source Packets
	sourcePacketRatio	Source Packets
	sourcePayload	Source Payload
	sourcePayloadHex	Source Payload As Hex
	sourcePort	Source Port
	sourcePrecedence	Source Precedence
	sourcePrecedenceOnly	Source Precedence
	sourcev6	IPv6 Source
	totalBytes	Total Bytes
	totalPackets	Total Packets
	viewObjectPair	View/Object
Events	category	Low Level Category
	creEventList	Matched Custom Rule
	credibility	Credibility
	destinationAssetName	Destination Asset Name
	destinationIP	Destination IP
	destinationMAC	Destination MAC
	destinationNetwork	Destination Network
	destinationPort	Destination Port
	destinationv6	IPv6 Destination
	device	Log Source
	deviceGroup	Log Source Group
	deviceTime	Log Source Time
	deviceType	Log Source Type
	duration	Duration
	endTime	End Time
	eventCount	Event Count
	eventDirection	Direction
	eventProcessor	Event Processor
	hasOffense	Associated With Offense
	highLevelCategory	High Level Category
	isCREEvent	Is CRE Event
	magnitude	Magnitude
	payload	Payload
	payloadHex	Payload As Hex
	postNatDestinationIP	Post NAT Destination IP

Table 1-2 Supported Fields (continued)

Table	Supported Statement	Normalized Field Name
	postNatDestinationPort	Post NAT Destination Port
	postNatSourceIP	Post NAT Source IP
	postNatSourcePort	Post NAT Source Port
	preNatDestinationIP	Pre NAT Destination IP
	preNatDestinationPort	Pre NAT Destination Port
	preNatSourceIP	Pre NAT Source IP
	preNatSourcePort	Pre NAT Source Port
	protocol	Protocol
	protocolId	Protocol
	qid	Event Name
	relevance	Relevance
	severity	Severity
	sourceAssetName	Source Asset Name
	sourceIP	Source IP
	sourceMAC	Source MAC
	sourceNetwork	Source Network
	sourcePort	Source Port
	sourcev6	IPv6 Source
	startTime	Start Time
	unparsed	Event Is Unparsed
	userName	Username

For example:

```
select sourceIP, destinationIP, application from flows where
protocol = 'TCP.tcp_ip'
select category, credibility from events where severity > 8
select * from events where credibility >=9
```

You can also use CIDR-based queries using the select statement. To query by source IP address (sourceIP) or by destination IP address (destinationIP) using a CIDR, use the following format:

```
select <query item> from <flows|events> where
<sourceCIDR|destinationCIDR> = '<CIDR Range>'
```

For example:

```
select * from flows where sourceCIDR = '10.100.100/24'
```

This command returns all flows coming from the 10.100.100 subnet. To capture flows coming from and into the subnet, use the regular OR expression as follows:

```
select * from flows where sourceCIDR = '10.100.100/24' OR
destinationCIDR = '10.100.100/24'
```

You can use the following fields in any logic clause, such as the where or group by clause, to further refine your select statement.

Table 1-3 Supported Fields

Table	Supported Statement	Normalized Field Name
Flow	anyASN	Source or Destination ASN
	anyHost	Source or Destination IP
	anyNetwork.	Source or Destination Network
	destinationTOS	Destination QoS
	icmpType	ICMP Type/Code
	sourceOrDestinationIP	Source or Destination IP
	sourceTOS	Source QoS
Events	anyIP	Any IP
	anyMac	Source or Destination MAC Address
	anyPort.	Any Port
	sourceOrDestinationIP	Source or Destination IP
	sourceOrDestinationNetwork.	Source or Destination Network
	sourceOrDestinationPort	Source or Destination Port

Using Where Clauses

You can restrict your AQL queries using **where** clauses. The supported logical operators in the clause include **and**, **or**, and parentheses. Also, the supported comparison operators include **=**, **<**, **>**, **>=**, **<=**, and **!=**.

For example:

```
select sourceIP, category, credibility from events where
severity > 9 and category = 5013
```

```
select sourceIP, category, credibility from events where
(severity > 9 and category = 5013) or (severity < 5 and
credibility > 8)
```

The **where** clause also supports the **arieltime** variable, which overrides the time settings passed to the AQL CLI. The **arieltime** variable must be used with the **between** keyword to specify the start and end time bounds of the query. All time constraints must be entered as either UNIX timestamps or formatted date/time strings.

You can only use the `arietime` variable once in a single query. Therefore, you can only query a continuous span of time in a single AQL command.

The logical operator for the `arietime` variable and the remainder of the `where` clause should be the `and` operator. We recommend that you use the `arietime` variable as the last constraint of the query and the `and` operator between the `arietime` variable and the rest of the `where` clause.

Using the Group By Clause

You can use the `group by` clause to aggregate your data. Typically, data aggregation is combined with arithmetic functions on remaining columns to provide meaningful results of the aggregation. For example, to enter a query to investigate the IP addresses that sent more than 1 million bytes within all flows in a specific time frame, you must enter:

```
select sourceIP, SUM(sourceBytes) from flows where sourceBytes >
1000000 group by sourceIP
```

The output resembles:

```
-----
| sourceIP          | SUM_sourceBytes |
-----
| 64.124.201.151   | 4282590.0       |
| 10.105.2.10      | 4902509.0       |
| 10.103.70.243    | 2802715.0       |
| 10.103.77.143    | 3313370.0       |
| 10.105.32.29     | 2467183.0       |
| 10.105.96.148    | 8325356.0       |
| 10.103.73.206    | 1629768.0       |
-----
```

However, if you compare this information to a non-aggregate query, the output displays all the IP addresses that are unique:

```
select sourceIP, sourceBytes from flows where sourceBytes >
1000000
```

```
-----
| sourceIP          | sourceBytes      |
-----
| 64.124.201.151   | 1448629          |
| 10.105.2.10      | 2412426          |
| 10.103.70.243    | 1793095          |
| 10.103.77.143    | 1449148          |
| 10.105.32.29     | 1097523          |
| 10.105.96.148    | 4096834          |
| 64.124.201.151   | 2833961          |
| 10.105.2.10      | 2490083          |
| 10.103.73.206    | 1629768          |
| 10.103.70.243    | 1009620          |
-----
```

10.105.32.29	1369660
10.103.77.143	1864222
10.105.96.148	4228522

In addition to the `SUM` operator, the `MIN`, `MAX`, and `AVG` arithmetic aggregation functions are also supported.

Using the Order By Clause

You can add a single `order by` clause to the end of your AQL CLI query. Only one field can be used in the `order by` clause. Also, sorting can be switched between ascending or descending by appending the `asc` or `desc` keyword to the `order by` clause, respectively. By default, the query returns results in descending order.

For example:

```
select sourceBytes, sourceIP from flows where sourceBytes >
1000000 order by sourceBytes
```

Or, if you wish to display results in ascending order:

```
select sourceBytes, sourceIP from flows where sourceBytes >
1000000 order by sourceBytes asc
```

Combining the `group by` and the `order by` clauses in a single query is useful for creating data, such as, TopN lists to determine the most abnormal events or the most bandwidth intensive IP addresses. For example, the following query displays the top traffic intensive IP address in a descending order:

```
select sourceIP, sum(sourceBytes) from flows group by sourceIP
order by sum(sourceBytes) desc
```

Using the Count(*) Clause

You can use the `count(*)` clause to count the number of records matching your query. For example, if you wish to count all events with credibility equal to or greater than 9:

```
select count(*) from events where credibility >= 9
```

Using the Distinct Clause

You can use the `distinct` clause to select unique rows based on a column or a group of columns. This clause is similar to the `group by` clause, however, the `distinct` clause ensure ANSI SQL compatibility. For example:

```
select distinct sourceIP, sourcePort from flows where
sourceBytes > 1000000
```

Using the Count (Distinct ...) Clause

You can use the standard SQL `Count(Distinct ...)` clause to obtain unique counts. Using the AQL CLI, you can only use one field. For example, if you wish to view all the IP addresses that are connected to a specific IP address over time:

```
select count(distinct sourceIP) from flows where destinationIP =
'192.168.61.71'
```

Or, if you wish to view the number of unique source IP addresses communicating with a particular destination IP address:

```
select destinationIP, count(distinct sourceIP) from flows group
by destinationIP
```



Note: Using this clause may require additional system resources. Therefore, depending on the query, the amount of time to return results may vary.

Using the Materialize View Clause

The `materialize view` clause allows you to produce query results as a static view and run subsequent queries against the view. You can also specify the period of time that the `materialized view` is accessible.

The syntax for the `materialized view` includes:

```
materialize view <time> NameOfView as select <statement>
```

Where:

- `<time>` specifies the time you wish the `materialized view` to be accessible.
- `<statement>` specifies a valid select statement.

For example, if you wish to create a `materialized view` containing flows with more than 1,000,000 source bytes, enter the following:

```
materialize view LargeSourceBytesFlows as select * from flows
where sourceBytes >1000000
```

To select from this view, enter the select statement as you would a valid table:

```
select * from LargeSourceBytesFlows
```

You can also use an aggregation statement on a materialized view:

```
select sourceIP, sum(sourceBytes) from LargeSourceBytesFlows
group by sourceIP
```



Note: You cannot create a `materialized view` statement based on a previously created materialized view.

If you wish to create a **materialized view** to select from a record set with ambiguous column names, you can define aliases for all computed columns. For example:

```
materialize view MyView as select sourceIP, sum(sourceBytes) as
srcBytesSum from flows group by sourceIP
```

Then you can refer to the alias in a subsequent query against **MyView**:

```
select * from MyView orderBy srcBytesSum
```

Using the Like Clause

You can search text fields using the standard **like** clause. You can also use the two wild card options supported by the AQL Query CLI: **%** and **_**. The percentage (**%**) wild card option matches zero or more characters while the **_** wild card option only matches one character.

For example:

If you wish to match names such as Joe, Joanne, Joseph, or any other name beginning with Jo, enter the following clause:

```
select * from events where userName like 'jo%'
```

If you wish to match names beginning with Jo that are three characters long, such as, Joe or Jon, enter the following clause:

```
select * from events where userName like 'jo_'
```

You can enter the wild card option at any point in the command. For example:

```
select * from flows where sourcePayload like '%xyz'
```

```
select * from events where payload like '%xyz%'
```

```
select * from events where payload like '_yz'
```

