



EMC® Watch4net®

Version 6.3u2

Flow Listener Administration Guide

P/N: pending

Rev : 46950

EMC²

Copyright ©2009 - 2014 EMC Corporation. All rights reserved. Published in the USA.

Published April, 2014

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

The information in this publication is provided as is. EMC Corporation makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose. Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

EMC², EMC, and the EMC logo are registered trademarks or trademarks of EMC Corporation in the United States and other countries. All other trademarks used herein are the property of their respective owners.

For the most up-to-date regulatory document for your product line, go to EMC Online Support (<https://support.emc.com>).

Contents

Product Overview	iii
Where to get help	v
1 Configuration	1
2 Supported fields	3
3 sFlow support	11
4 Aggregation	13
5 Configuration example	15
1 Event Processing Manager configuration	15
2 <i>Flow Listener</i> configuration	16

Product Overview

The *Flow Listener* is a *Processing Element* that captures and transforms network monitoring packets from flow-shaped data. The *Flow Listener* parses incoming NetFlow or sFlow packets transforming them into events which are then forwarded to the next *Processing Elements*.

Currently, the *Flow Listener* supports versions 5, 7, 8, and 9 of NetFlow, as well as IPFIX, and the basic network monitoring use of sFlow version 5. The *Flow Listener* must be configured to either listen for NetFlow or sFlow data.

Please note that within the scope of the Flow-Listener configuration, NetFlow also encompasses its derivations, such as J-Flow and cFlowd.

When parsing NetFlow data, flows from any of its versions, and derivations, can be received at the same time.

Traffic flows is the term used to encompass all flow data, whether NetFlow, sFlow, or a NetFlow derivation, and is used in other configuration files. For detailed example information on Traffic Flows, see the *APG-Traffic-Flow-Implementation-How-To* PDF.

Note: Subsequent to NetFlow version 9, the NetFlow protocol was generalized from a Cisco-initiated protocol, to an industry-wide backed protocol now named Internet Protocol Flow Information eXport (IPFIX). References to NetFlow version 10 are in fact referring to IPFIX, which is supported by *APG*. For the

purposed of this document, references to NetFlow includes the IPFIX implementation.

Where to get help

EMC support, product, and licensing information can be obtained as follows:

PRODUCT INFORMATION - For documentation, release notes, software updates, or information about EMC products, go to EMC Online Support at:

<https://support.emc.com>

TECHNICAL SUPPORT - Go to EMC Online Support and click Service Center. You will see several options for contacting EMC Technical Support. Note that to open a service request, you must have a valid support agreement. Contact your EMC sales representative for details about obtaining a valid support agreement or with questions about your account.

Configuration

Like all other processing elements, the *Flow Listener* must be run within the Event Processing Manager. To do so, it must be added to the processing configuration file (`<APG-Root>/Event-Processing/Event-Processing-Manager/<Instance-Name>/conf/processing.xml`). Below is a list of parameters that must be defined in order for the *Flow Listener* to receive flows properly. A configuration example is available in the [Configuration example](#) on page 15.

- `port`: Defines the port on which the *Flow Listener* will listen for incoming flows.
- `flow-type`: Defines the family of flows to be parsed (NetFlow or sFlow). Once again, NetFlow includes its derivations, such as J-Flow and cFlowd.
- `keep-flow-timestamp`: Defines if the event timestamp will be set using the flow timestamp or the current timestamp from the *APG* server. Keep in mind the following: Out of order data could be received when using the flow's timestamp. Since *APG* only parses data on a linear timeline, for a particular flow if it receives data with a timestamp before data it has already processed for the flow this data will be dropped. So while it could be useful to use the flow's timestamp, to, for example, record timestamp data from different time zones at their source, you should verify that data will not be sent out of order to *APG*.

To prevent any timestamp mixage, it is recommended to keep it to false.

- `normalize-sample-rate`: Defines if the `BYTES` and `PACKETS` will be multiplied by the `SAMPLING_INTERVAL`. Default value is `false`.
- `nf9-fields`: Defines the configuration file used for the netflow v9 fields. All netflow v9 fields must be defined in that file. Default file is `conf/nf9Fields.properties`.

Supported fields

Here's the list of fields the *Flow Listener* supports with the version and type available. Please notes that for NetFlow v9, the fields must be defined in the `nf9-fields` configuration file. Also there is no guarantee that every field will be present in the flows since it depends on the template defined by the NetFlow administrator for the flow. As well, since there can be differently defined sFlow networking monitoring packets that *APG* recognizes, there's also no guarantee that every field will be present in each flow.

Field name	Description	Availability
AAA_USERNAME	AAA username.	NF9 IPFIX
APPLICATION_DESC	Application description.	NF9 IPFIX
APPLICATION_ID	Application classification.	NF9 IPFIX
APPLICATION_MEDIA_BYTE_RATE	Average media bit rate (bps) for all flows during the monitoring interval.	NF9 IPFIX
APPLICATION_MEDIA_BYTES	Number of IP bytes from by media applications received for a specific media stream.	NF9 IPFIX

Supported fields

Rev : 46950

Field name	Description	Availability
APPLICATION_MEDIA_EVENT	Bit 1 is not used. Bit 2 indicates that no media application packets were seen, in other words, a Media Stop Event occurred.	NF9 IPFIX
APPLICATION_MEDIA_PACKET_RATE	Media Packet Rate	NF9 IPFIX
APPLICATION_MEDIA_PACKETS	Number of IP packets produced from media applications received for a specific media stream.	NF9 IPFIX
APPLICATION_NAME	Name associated with a classification.	NF9 IPFIX
BGP_IPV4_NEXT_HOP	IPv4 next hop router address in the BGP domain	NF9 IPFIX
BPG_IPV6_NEXT_HOP	IPv6 next hop router address in the BGP domain	NF9 IPFIX
BYTE_RATE	Average number of packets or bits (as configured) processed by the monitoring system per second during the monitoring interval for all flows.	NF9 IPFIX
BYTES	Number of bytes associated with an IP flow. On NF9/IPFIX, it defines only the incoming bytes.	ALL
DIRECTION	Flow direction: 0 = ingress flow, 1 = egress flow	NF9 IPFIX
DSCP	Differentiated Service Code Point. Other way of accessing the TOS field based on RFC2474	ALL
DSCP_ECN	Explicit Congestion Notification. The 2 LSB contained in the DSCP (TOS) field based on RFC2474	ALL
DST_AS	Destination BGP autonomous system number	NF 5-7-9 IPFIX
DST_MASK	The number of contiguous bits in the destination address subnet mask (slash notation)	NF 5-7-9 IPFIX

Field name	Description	Availability
DST_PORT	The destination port	ALL
DST_PRIORITY	Destination priority	SF5
DST_TOS	Destination type of service	NF9 IPFIX
DST_VLAN	Virtual LAN identifier associated with egress interface	NF9 IPFIX
ENGINE_ID	ID number of the flow switching engine	NF9 IPFIX
ENGINE_TYPE	Type of flow switching engine: 0 = RP, 1 = VIP/Linecard	NF9 IPFIX
EXTRA_PKTS	Packets that exceed the contract	NF8
FIRST_SWITCHED	System uptime at which the first packet of this flow was switched	NF 5-7-9 IPFIX
FLOW_ACTIVE_TIMEOUT	Timeout in seconds for active flows in the cache	NF9 IPFIX
FLOW_INACTIVE_TIMEOUT	Timeout in seconds for inactive flows in the cache	NF9 IPFIX
FLOW_SAMPLER_ID	Identifier shown in “{}show flow-sampler”{}	NF9 IPFIX
FLOW_SAMPLER_MODE	The type of algorithm used for sampling data: 2 = random sampling	NF9 IPFIX
FLOW_SAMPLER_RANDOM_INTERVAL	Packet interval at which to sample	NF9 IPFIX
FLOW_SEQUENCE	Incremental sequence counter of all Export Packets sent from the current Observation Domain by the Exporter	NF9 IPFIX
FLOW_TMST	The field containing the timestamp (will always be present)	ALL
FLows	Number of aggregated flows	NF9 IPFIX
FORWARDING STATUS	Forwarding status	NF9 IPFIX
FRAGMENT_OFFSET	The fragment-offset value from fragmented IP packets	NF9 IPFIX
FW_EVENT	Firewall Event	NF9 IPFIX

Field name	Description	Availability
ICMP_TYPE	Internet Control Message Protocol packet type	NF9 IPFIX
IF_DESC	Interface description	NF9 IPFIX
IF_NAME	Interface name	NF9 IPFIX
IN_DST_MAC	Incoming destination MAC	NF9 IPFIX
IN_PERMANENT_BYTES	Running byte counter for a permanent flow	NF9 IPFIX
IN_PERMANENT_PKTS	Running packet counter for a permanent flow	NF9 IPFIX
IN_SRC_MAC	Incoming source MAC address	NF9 IPFIX SF5
INPUT_SNMP	Input interface index	NF 5-7-9 IPFIX
IP_PRIORITY	IP priority	SF5
IP_PROTOCOL_VERSION	Internet Protocol Version Set to 4 for IPv4 and 6 for IPv6	NF9 IPFIX
IPV4_DST_ADDR	IPv4 destination address	ALL
IPV4_DST_PREFIX	IPv4 destination address prefix	NF9 IPFIX
IPV4_IDENT	The IPv4 identification field	NF9 IPFIX
IPV4_NEXT_HOP	IPv4 next hop router address	NF 5-7-9 IPFIX
IPV4_SRC_ADDR	IPv4 source address	ALL
IPV4_SRC_PREFIX	IPv4 source address prefix	NF9 IPFIX
IPV6_DST_ADDR	IPv6 destination address	NF9 IPFIX SF5
IPV6_DST_MASK	Length of the IPv6 destination mask in contiguous bits	NF9 IPFIX
IPV6_FLOW_LABEL	IPv6 flow label	NF9 IPFIX
IPV6_NEXT_HOP	IPv6 next hop router address	NF9 IPFIX
IPV6_OPTION_HEADERS	Bit-encoded field identifying IPv6 option headers found in the flow	NF9 IPFIX
IPV6_SRC_ADDR	IPv6 source address	NF9 IPFIX SF5
IPV6_SRC_MASK	Length of the IPv6 source mask in contiguous bits	NF9 IPFIX
LAST_SWITCHED	System uptime at which the last packet of this flow was switched	NF 5-7-9 IPFIX
MARKED_TOS	Type of Service of the packets that exceeded the contract	NF8

Field name	Description	Availability
MAX_PKT_LENGTH	Maximum IP packet length	NF9 IPFIX
MAX_TTL	Maximum time to live on incoming packets	NF9 IPFIX
MIN_PKT_LENGTH	Minimum IP packet length	NF9 IPFIX
MIN_TTL	Minimum time to live on incoming packets	NF9 IPFIX
MONITOR_EVENT	Bit 1 indicates that one of the thresholds specified by a react statement for the flow was crossed at least once in the monitoring interval. Bit 2 indicates that there was a loss-of-confidence in measurement.	NF9 IPFIX
MPLS_LABEL_1	MPLS label at position 1 in the stack. 20 bits of MPLS label, 3 EXP (experimental) bits and 1 S (end-ofstack) bit.	NF9 IPFIX SF5
MPLS_LABEL_10	MPLS label at position 10 in the stack	NF9 IPFIX
MPLS_LABEL_2	MPLS label at position 2 in the stack	NF9 IPFIX
MPLS_LABEL_3	MPLS label at position 3 in the stack	NF9 IPFIX
MPLS_LABEL_4	MPLS label at position 4 in the stack	NF9 IPFIX
MPLS_LABEL_5	MPLS label at position 5 in the stack	NF9 IPFIX
MPLS_LABEL_6	MPLS label at position 6 in the stack	NF9 IPFIX
MPLS_LABEL_7	MPLS label at position 7 in the stack	NF9 IPFIX
MPLS_LABEL_8	MPLS label at position 8 in the stack	NF9 IPFIX
MPLS_LABEL_9	MPLS label at position 9 in the stack	NF9 IPFIX
MPLS_TOP_LABEL_IP_ADDR	Forwarding Equivalent Class corresponding to the MPLS Top Label	NF9 IPFIX
MPLS_TOP_LABEL_TYPE	MPLS Top Label Type: 0 = UNKNOWN, 1 = TE-MIDPT, 2 = ATOM, 3 = VPN, 4 = BGP, 5 = LDP	NF9 IPFIX

Field name	Description	Availability
MUL_DST_BYTES	IP multicast outgoing byte counter for packets associated with the flow	NF9 IPFIX
MUL_DST_PKTS	IP multicast outgoing packet counter for packets associated with the flow	NF9 IPFIX
MUL_IGMP_TYPE	Internet Group Management Protocol packet type	NF9 IPFIX
NF7_FLAGS	NetFlow 7 flags	NF7
OUT_BYTES	Number of bytes of the outgoing packets	NF9-IPFIX
OUT_DST_MAC	Outgoing destination MAC address	NF9 IPFIX SF5
OUT_PKTS	Number of the outgoing packets	NF9 IPFIX
OUT_SRC_MAC	Outgoing source MAC	NF9 IPFIX
OUTPUT_SNMP	Output interface index	NF 5-7-9 IPFIX
PACKETS_DROPPED	IP packet drops by any intermediate system in any of the monitored flows.	NF9 IPFIX
PACKETS	Number of packets associated with an IP flow. On NF9/IPFIX, it defines only the incoming packets.	NF 5-7-9 IPFIX
PROTOCOL	The IP protocol	ALL
ROUTER_SC	NetFlow 7 Catalyst 5000 field	NF7-8
ROUTER_SRC	The IP adress of the router sending the flows (will always be present)	ALL
SAMPLER_NAME	Name of the flow sampler	NF9 IPFIX
SAMPLING_ALGORITHM	The type of algorithm used for sampling: 1 = Deterministic Sampling, 2 = Random Sampling	NF 5 - 9 IPFIX
SAMPLING_INTERVAL	The rate of packet sampling (100 = 1 out of 100 packets)	ALL
SOURCE_ID	A 32-bit value that identifies the Exporter Observation Domain.	NF9 IPFIX
SRC_AS	Source BGP autonomous system number	NF 5-7-9 IPFIX

Field name	Description	Availability
SRC_MASK	The number of contiguous bits in the source address subnet mask (slash notation)	NF 5-7-9 IPFIX
SRC_PORT	The source port	ALL
SRC_PRIORITY	Source priority	SF5
SRC_VLAN	Virtual LAN identifier associated with ingress interface	NF9 IPFIX
TCP_FLAGS	Cumulative of the TCP flags for the aggregated flows	ALL
TIMESTAMP_INTERVAL	Timestamp Interval	NF9 IPFIX
TOS	Type of service. On NF9/IPFIX, it defines only the the TOS for the incoming interface. Will return the IP Precedence based on RFC791/1349	ALL
TOS_REST	The other 4 bits contained in the TOS field	ALL
TOTAL_BYTES_EXP	Number of bytes exported by the Observation Domain	NF9 IPFIX
TOTAL_FLOWS_EXP	Number of flows exported by the Observation Domain	NF9 IPFIX
TOTAL_PKTS_EXP	Number of packets exported by the Observation Domain	NF9 IPFIX
TRANSPORT_EVENT_PACKET_LOSS	Number of loss events (number of contiguous sets of lost packets).	NF9 IPFIX
TRANSPORT_PACKETS_EXPECTED	Number of packets expected.	NF9 IPFIX
TRANSPORT_PACKETS_LOST	Number of packets lost.	NF9 IPFIX
TRANSPORT_PACKETS_LOST_RATE	Rate of packets lost, in percent.	NF9 IPFIX
TRANSPORT_ROUND_TRIP_TIME	Average of all round-trip-times.	NF9 IPFIX
TRANSPORT_RTP_JITTER_MAX	Maximum threshold for the deviation of the difference in packet spacing at the receiver compared to the sender for a pair of packets.	NF9 IPFIX

Field name	Description	Availability
TRANSPORT_RTP_JITTER_MEAN	Mean deviation of the difference in packet spacing at the receiver compared to the sender for a pair of packets.	NF9 IPFIX
TRANSPORT_RTP_JITTER_MIN	Minimum threshold for the deviation of the difference in packet spacing at the receiver compared to the sender for a pair of packets.	NF9 IPFIX
TRANSPORT_RTP_SSRC	The configured maximum number of SSRCS that can be monitored within same flow (as defined by the protocol, source/destination address, source/destination port). The range is from 1 to 50.	NF9 IPFIX
VENDOR_43	Vendor proprietary	NF9 IPFIX
VENDOR_51	Vendor proprietary	NF9 IPFIX
VENDOR_65	Vendor proprietary	NF9 IPFIX
VENDOR_66	Vendor proprietary	NF9 IPFIX
VENDOR_67	Vendor proprietary	NF9 IPFIX
VENDOR_68	Vendor proprietary	NF9 IPFIX
VENDOR_69	Vendor proprietary	NF9 IPFIX
VENDOR_87	Vendor proprietary	NF9 IPFIX
XLATE_DST_ADDR_IPV4	Xlate destination IPv4 address.	NF9 IPFIX
XLATE_DST_PORT	Xlate destination port.	NF9 IPFIX
XLATE_SRC_ADDR_IPV4	Xlate source IPv4 address.	NF9 IPFIX
XLATE_SRC_PORT	Xlate source port.	NF9 IPFIX

3

sFlow support

To this date, only a fragment of sFlow version 5 is supported. More support will be added if necessary. The currently supported sample formats are the ones under 1000 which are:

1. Raw packet header
2. Ethernet frame data (supported also on VLAN)
3. IPv4 data
4. IPv6 data

Most of the received samples should be Raw packet headers since that's the suggested configuration from the sFlow RFC. A Raw packet header will simply be a copy of the Ethernet header and the network layer header from the sampled data. The most common of the network layer headers are supported:

1. IPv4 header
2. IPv6 header
3. MPLS header

Every other headers type or sample format will be dropped. The currently supported counter flows are:

1. Generic Interface Counters
2. Ethernet Interface Counters
3. VLAN Counters
4. Processor Information

In order to help differentiate sFlow counters with regular flows, a field is added only in the counters named SFLOW_COUNTER with a value of 1.

4

Aggregation

Another processing element, the Events Aggregator, can be used to generate aggregates from multiple flows on custom properties and outputting either events or time-series. This practice can be recommended if the incoming flow of events make the generation of top level reports too slow. The configuration example in section [5](#) shows how to chain an Events Aggregator to the *Flow Listener*. Note that this step is *optional*, and the lines referring to this component can be omitted if no aggregations are needed.

5

Configuration example

1 Event Processing Manager configuration

```
<?xml version="1.0" encoding="UTF-8"?>
<processing-manager xmlns="http://www.watch4net.com/Events/DefaultProcessingManager"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://www.watch4net.com/Events/DefaultProcessingManager_
DefaultProcessingManager.xsd">

  <!-- databases -->
  <processing-element name="DW" config="APG-Datastore-Writer/Default/conf/ds-netflow-archive.
xml" />

  <!-- Events Aggregator (optional) -->
  <processing-element name="E2TC" config="Events-Aggregator/Default/conf/flow-aggregator-ts.
xml" data="DW" />

  <!-- collectors -->
```

```
<processing-element name="FLOW" config="Flow-Listener/Default/conf/flowListener.xml" data="E2TC" />
</processing-manager>
```

2 *Flow Listener* configuration

```
<port>9996</port>
<socket-buffer-size>128</socket-buffer-size>
<bind-address>192.168.0.25</bind-address>
<flow-type>netflow</flow-type>
<keep-flow-timestamp>false</keep-flow-timestamp>
```

The above configuration specify that the flow listener will listen on port 9996, uses a socket buffer size of 128 MB (optional/default) and bind on the local address 192.168.0.25 (optional) for any version of NetFlow packets.