
From NetFlow to IPFIX

the evolution of IP flow information export

presented by

Carsten Schmoll - Fraunhofer FOKUS - Berlin, DE

Elisa Boschi - Hitachi Europe - Zurich, CH
Brian Trammell - CERT/NetSA - Pittsburgh, PA, US

FRNOG 11 - Paris, Fri - November 30, 2007

What is IPFIX?


- Emerging IETF standard for flexible export of IP flow data from routers or other metering processes: *IP Flow Information eXport*
- Defines:
 - a rich, easily extensible information model,
 - a template-driven data representation,
 - and a unidirectional protocol for export of IP flow data over a variety of transport protocols.
- Does not define specific requirements for flow assembly, flow key selection, aggregation.

History and Motivation

- IETF IPFIX working group started in 2001 to define a standard flow export protocol.
- Cisco NetFlow V9 selected as basis for this new protocol.
 - Evolution of previous NetFlow versions.
 - Adding templates for flexible data definition.
- Developed protocol from this basis to meet defined requirements.
 - IPFIX information model provides a superset of NetFlow V9 information model
 - the two are not directly interoperable without binary message translation.

Data Representation

- Templates in the message stream describe the data sets.
- Allows flexible and efficient (binary) representation of flows on the wire.



Information Model

- The information model supports reporting a wide variety of information elements (IEs):
 - “Five-tuple” (IPv4, IPv6) and standard packet/byte counters
 - Packet treatment: e.g., routed next hop and AS
 - Detailed counters: e.g., sum of squares, flag counters
 - Timestamps down to nanosecond resolution
 - All ICMP, TCP, UDP header fields
 - Layer 2, VLAN, MPLS, and other sub-IP information
- New IEs registered with IANA.
- Enterprise-specific IEs for private extensions.

Current Status

- **It's taken longer than we'd thought, but we're nearly done ...**
- Core IPFIX protocol documents completed in 2006, (probably) to be published as RFCs in 2007.
- Working group continuing to define extensions to the protocol and applications of the protocol:
 - Bidirectional flow export
 - Redundancy reduction for increased export efficiency
 - Flow storage and File-based interoperability
 - MIB and XML-based configuration for IPFIX devices
- Multiple implementations tracking the draft standards available now.

Bidirectional Flow Export

- Bidirectional flow (biflow) metering and analysis is applicable to several use cases:
 - data reduction
 - separation of “answered” traffic from unanswered
 - use in full reconstruction of TCP sessions
- The IPFIX protocol has no direct support for bidirectional flow export in a single-record
- This extension allows “reversal” of any element within the Information Model for bidirectional flow export.
- Approved as RFC.

Reducing Redundancy

- Technique for bandwidth-saving information export
 - Separates the export of flow records such that attributes common to several flow records are sent only once.
 - Links common flow properties to specific properties with a unique identifier.
 - similar to foreign keys in relational databases (1:n and n:m)
- Approved as RFC.

Flow Storage

- Many analysis tools interoperate not via direct communication, but via file exchange.
 - exchange available via a variety of transport methods (HTTP, FTP, SSH+SCP, SMTP+MIME, ...)
 - files support a variety of useful operations (compression, encryption, replay, etc.)
 - files are a natural unit of grouping related flow data (e.g. a single security incident or query result).
- Existing de-facto standard for flow storage:
NetFlow PDU files
 - Not extensible for data fields not in NetFlow.

Flow Storage: IPFIX as basis

- IPFIX defines a template-driven data representation and a rich, easily extensible information model, so:
- Ideal basis for a flow storage format
 - Extensible and self-describing, unlike V5 PDU files
 - Adequate semantic flexibility for flow data without overhead of e.g. XML.
 - Additional applicability to IPFIX (or NetFlow V9) collection infrastructures.

IPFIX Files

- An IPFIX file is any serialized stream of IPFIX Messages.
 - Alternately, a “file transport” for IPFIX.
- Provides a set of extensions:
 - File contents
 - Error detection and recovery
 - Extended type information for enterprise-specific information elements.
- To be published as an RFC in 2008.

IPFIX Implementations (1)

- **YAF (Yet Another Flowmeter)**
 - takes packets from the wire or libpcap dumpfiles.
 - writes IPFIX Files or exports IPFIX Messages.
 - supports bidirectional flow export.
- **SiLK (System for Internet Level Knowledge)**
 - large-scale flow storage and command-line analysis suite.
 - supports NetFlow V5 and IPFIX flow collection.
 - can analyze IPFIX Files directly, as well.
- **libfixbuf: an IPFIX library in C**
 - Used by YAF and SiLK
 - Available from <http://tools.netsa.cert.org/>

IPFIX Implementations (2)

- **OpenIMP**
 - provides metering processes, export/collection, and analysis tools.
 - specifically focused on active and passive quality of service measurement.
 - available from <http://www.ip-measurement.org/openimp/>
- **libipfix: another IPFIX library in C**
 - supports Reducing Redundancy extension
 - supports IPFIX File and mysql storage
 - used by OpenIMP
 - available from <http://meteor.fokus.fraunhofer.de/libipfix/>

IPFIX Implementations (3)

- **Versatile Monitoring Toolkit (VERMONT)**
 - provides metering processes, export/collection, and monitoring tools.
 - implements IPFIX and related PSAMP (packet sampling) protocol.
 - available from <http://vermont.berlios.de/>
- **ntop**
 - web-based traffic measurement application
 - acts as IPFIX collecting process
 - available from <http://www.ntop.org/>
- **Cisco IOS**
 - is going to support it sooner or later for sure

FIN

- IPFIX is an emerging standard for flexible flow export, representation, and storage.
 - For those who want to follow the progress:
 - <http://tools.ietf.org/wg/ipfix> - WG tools page
 - <http://www.ietf.org/html.charters/ipfix-charter.html>
- Implementations available & tested
 - IPFIX interoperability events in July '05, March '06, and November '06 so far.

Questions?

- ask now
- or later:

carsten.schmoll@fokus.fraunhofer.de

bht@cert.org

elisa.boschi@hitachi-eu.com

Additional Slides

Comparison to sFlow

- sFlow is a packet sampling protocol
 - Intended for many of the same applications as NetFlow and IPFIX.
 - Use of packet sampling instead of flow assembly reduces state overhead on measurement device.
 - Analogous to PSAMP, which extends IPFIX for export of sampled packet data.
- Both provide flexible export, but...
 - sFlow provides message types for flexibility,
 - IPFIX provides templates and information elements:
 - IPFIX allows definition of novel message types on the fly.