

IP Flow Information Export (IPFIX): Applicability and Future Suggestions for Network Security

Elisa Boschi, Tanja Zseby, Lutz Mark, Thomas Hirsch
Fraunhofer FOKUS,
Berlin, Germany
{boschi, zseby, mark, hirsch}@fokus.fraunhofer.de

Abstract

This year, the IP Flow Information Export (IPFIX) protocol will become standard for exporting flow information from routers and probes. Standardized methods for packet selection and the export of per packet information will follow soon from the IETF group on packet sampling (PSAMP). The future availability of network information in a standardized form enables a wide range of critical applications for Internet operation including, accounting, QoS auditing and detection of network attacks. In this paper we present the IPFIX protocol, and discuss its applicability with a special focus on network security. We propose a coupling of IPFIX with AAA functions to improve the detection and defense against network security incidents and for Inter-domain information exchange based on IPFIX utilizing secure transmission channels provided by the AAA architecture.

1. Introduction

The IP Flow Information eXport (IPFIX) protocol defines a format and protocol for the export of flow information from routers or measurement probes. In the past a lot of proprietary solutions were developed for flow information export (e.g. Cisco NetFlow, InMon sFlow, NeTraMet, etc). Now after several years of lively discussions the IETF is about to submit a standard for flow information export, the IPFIX protocol.

Capturing flow information plays an important role for network security, both for detection of security violation, and for subsequent defence. Attack and intrusion detection is one of the five target applications that require flow measurements on which the requirements definition has been based: usage-based accounting, traffic analysis, traffic engineering, QoS monitoring and intrusion detection (Cf. [RFC 3917]).

In this paper we summarize the IPFIX protocol, describe our implementation of it and discuss its applicability for a number of different applications. Particularly, we analyze the IPFIX applicability to security related applications such as anomaly and intrusion detection and discuss the coupling of IPFIX with AAA in the context of inter-domain information exchange.

2. IP Flow Information Export

2.1 IPFIX Summary

Based on the requirements defined in [RFC3917] different existing protocols (NetFlow v9, Diameter, CRANE, IPDR) were evaluated as candidates to provide the basis for a future IPFIX protocol [RFC3955]. The result of the evaluation was that NetFlow v9 provides the best basis for it.

IPFIX is a general data transport protocol that is easily extensible to suit the needs of different applications. The protocol is flexible in both flow key and flow export. The Flow Key defines the properties used to select flows and can be defined depending on the application needs. Flow information is exported using flow data records and the information contained in those records can be defined using template records. A template ID uniquely identifies each template record and provides the binding between template and data records.

As requested by the IESG, IPFIX transport has to fulfil certain reliability and security requirements. Therefore PR-SCTP has been chosen as mandatory basic transport protocol for IPFIX for all compliant implementations. TCP and UDP can be used as optional protocols. Preference to PR-SCTP was given

because it is congestion-aware and reduces bandwidth in case of congestion but still has a much simpler state machine than TCP. This saves resources on lightweight probes and router line cards.

2.2 FOKUS IPFIX Implementation

Fraunhofer FOKUS has developed and uses an Internet measurement application for distributed IP traffic and quality of service measurements. The software is called OpenIMP and consists of a central measurement controller and multiple distributed probes. The probes are remote controlled and send their measurement data to dedicated collectors. The measurement result transfer was implemented via files that were generated on the probes and were sent to the collector via ftp or scp. With the definition of IPFIX there appeared a smart and powerful alternative to the file transfer. So we decided to integrate IPFIX into the measurement application. The implementation was started before the definition of the IPFIX protocol has been finished. This offered us the possibility to send comments to the working group and to influence the further definition of the IPFIX protocol.

The IPFIX protocol was implemented within a separate C-library. This has the advantage that the implementation can be used outside OpenIMP. Using C makes it easy to integrate the code into e.g. C++ or JAVA applications and has the advantage of leading to small binaries. The library supports the IPFIX exporting and collecting processes via providing functions to export and to collect data using IPFIX data types and protocol.

The IPFIX protocol is not complex, which keeps the implementation simple. IPFIX messages can be transferred using SCTP, TCP or UDP as bearer protocol. An IPFIX implementation has to support SCTP-PR whereas support for TCP and UDP is optional. Unfortunately currently there is no major operating system with full support for SCTP-PR (at least the authors are not aware of one). So at present an application has to use TCP to enable the export over ipv6 networks. To test the SCTP code we used Debian Linux with kernel 2.6 that has support for SCTP-PR over ipv4.

The main task of the IPFIX exporter is to take the measurement data from one or more metering processes and to send the IPFIX messages to the data collectors. The exporter has to take care that the templates are sent prior to the related data records. For SCTP and TCP the templates have to be resent on a connection reestablishment. For UDP templates have to be resent after a configured timeout. This makes the implementation a bit more complex and requires the exporting process to store all active template definitions.

The IPFIX collector has to maintain a list of sources and per source a list of templates to decode incoming data templates. Because of the template feature of IPFIX the collector does not need any knowledge of the transferred data. All information needed to decode all kind of data is transferred via template records.

2.3 Applicability Scenarios: Accounting and QoS Monitoring

Usage-based accounting is one of the major applications for which the IPFIX protocol has been developed. IPFIX data records provide fine-grained measurement results for highly flexible and detailed resource usage accounting (i.e. the number of transferred packets and bytes per flow). Internet Service Providers (ISPs) can use this information to migrate from single fee, flat-rate billing to more flexible charging mechanisms based on time of day, bandwidth usage, application usage, quality of service, etc.

In order to realize usage-based accounting with IPFIX the flow definition has to be chosen in accordance with the tariff model. If for example the tariff is based on individual end-to-end flows, accounting can be realized with a flow definition determined by source address, destination address, protocol, and port numbers. Another example is a class-dependent tariff (e.g. in a DiffServ network); in this case, flows can be distinguished just by the DiffServ codepoint (DSCP) and source IP address.

QoS monitoring is the passive observation of transmission quality for single flows or traffic aggregates in the network and is one of the target applications of IPFIX. One example of its usefulness is the validation of QoS guarantees in service level agreements (SLAs). IPFIX data records enable ISPs to perform a detailed, time-based, and application-based usage analysis of a network.

3. IPFIX Applicability and Future Suggestions for Detection of and Defense against Network Attacks

The applications described in this section are related to the detection and report of intrusions and anomalous traffic. While describing the applicability of the protocol, we discuss how IPFIX could further support detection of, and reaction to, network attacks.

3.1 *Packet Selection and Packet Information Export*

For some scenarios, the detection of malicious traffic may require further insight into packet content. The PSAMP working group works on the standardization of packet selection methods [ZsMD05] and the export of per packet information [Duff05], [PoMB05]. Recently, the IETF PSAMP [PSAMP] group has decided to also use the IPFIX protocol for the export of per packet information. That means, in future we will get also per packet information from routers in a standardized way.

3.2 *Detection of network incidents and malicious traffic*

IPFIX provides useful input data for basic attack detection functions such as reporting unusually high loads, number of flows, number of packets of a specific type, etc. It can provide details on source and destination addresses, along with the start time of flows, TCP flags, application ports and flow volume. This data can be used to analyze network security incidents and identify attacks like DoS attacks, worm propagation or port scanning. Further data analysis and post-processing functions may be needed to generate the metric of interest for specific attack types.

Already basic IPFIX information allows detecting common attack schemes: A distributed DoS attack generates a large number of flows, often with a high data volume. The *number of newly detected source addresses* is commonly used [TaHo04] as a metric for detecting distributed activities. It correlates strongly with the flow count metric of IPFIX. Also, sudden increases in the occurrence of unusual IP or TCP flags (e.g. “Don’t Fragment”) can be an indicator for malicious traffic [TaAl02, SiPa04]. Based on the IPFIX information, derived metrics can highlight changes and anomalies. The most successful methods for anomaly detection to date are non-parametric change point detection algorithms, such as the cumulative sum (CUSUM) algorithm [WaZS02]. The integration of previous measurement results helps to assess traffic changes over time for detection of traffic anomalies. A combination with results from other observation points allows assessing the propagation of the attack and can help locate the source of an attack.

Detecting security incidents in real-time would require the pre-processing of data already at the measurement device and immediate data export in case a possible incident has been identified. This means that IPFIX reports must be generated upon incident detection events and not only upon flow end or fixed time intervals. IPFIX works in push mode. That means data records are automatically exported without waiting for a request. Placing the responsibility for initiating a data export at the exporting process is quite useful for detection of security incidents. The exporting process can immediately trigger the export of information if suspicious events are observed (e.g. sudden increase of the number of flows).

Security incidents could become a threat to IPFIX processes themselves. If an attack generates a large amount of flows (e.g. by sending packets with spoofed addresses or simulating flow termination), exporting and collecting process may get overloaded by the immense amount of data records that are exported. A flexible deployment of packet or flow sampling methods can prevent the exhaustion of resources.

3.3 *Sharing Information with Neighbor Domains*

For inter-domain measurements it is required to exchange *result* data, and eventually to allow remote configuration, across multiple administrative domains. Result data can be both measurements of QoS metrics on an end-to-end path, or monitoring information for troubleshooting, or information regarding attacks (either notifications of anomalous traffic or specific measurements to get further insight in case suspicious behavior was observed). Although ISPs can control and monitor their own network, they have minimal or no information at all about the characteristics and performance of other networks, nor the means of requesting and acquiring it. IPFIX provides the standard format and protocol for this information exchange.

3.4 Sharing Information with AAA Functions

One approach to do this is to use existing AAA components which provide a secure data transfer between domains by using the DIAMETER protocol and also provide functions for authentication and authorization which are useful to control access to data [RFC3334]. Furthermore, AAA servers usually keep accounting and auditing requirements, which can be used to directly derive measurement demands. For anomaly and intrusion detection the strong relation to AAA components can provide further benefits. Potential attackers can be identified and stopped from injecting traffic into the network. This is especially powerful, if AAA components from different administrative domains work together.

The combination of IPFIX and AAA functions can be beneficial also for attack detection. Such an interoperation enables further means of attack detection, advanced defense strategies and secure inter-domain cooperation. A AAA system has secure channels to neighbor AAA servers and can inform neighbors about incidents or suspicious observations. Through this system an ISP could also react to an attack by for example requesting the denial of access for potential attackers. A further benefit would be if AAA functions could invoke further measurements.

4. Conclusions

IPFIX is the upcoming standard for IP flow information export. The protocol is well suited to support applications such as QoS measurement, accounting, and anomaly and intrusion detection. In particular, for security applications, IPFIX can be used to exchange information with neighbors not only about incidents or anomalous traffic or to receive information previously requested to track attackers or classify the attacks. The joint use of IPFIX and AAA functions can add further benefits and be useful to track and stop attackers.

Some implementations of the protocol already exist, one of them coming from the authors of this paper, and interoperability events have been planned in the next months.

5. References

- [RFC3917] J. Quittek, T. Zseby, B. Claise, S. Zander, Requirements for IP Flow Information Export, October 2004
- [Clai05] B. Claise (Editor), IPFIX Protocol Specification, Internet Draft <draft-ietf-ipfix-protocol-14.txt>, work in progress, May 2005
- [PoMB05] G. Pohl, L. Mark, E. Boschi Export of per-packet information using IPFIX, Internet Draft <draft-pohl-perpktinfo-03.txt>, work in progress, February 2005
- [QuBr05] J. Quittek, S. Bryant, J. Meyer, Information Model for IP Flow Information Export, Internet Draft <draft-ietf-ipfix-info-06>, work in progress, February 2005
- [PSAMP] <http://www.ietf.org/html.charters/psamp-charter.html>
- [Duff05] Nick Duffield (Ed.), A Framework for Packet Selection and Reporting, Internet Draft draft-ietf-psamp-framework-08, work in progress, January 2005
- [RFC3334] T. Zseby, S. Zander, G. Carle, Policy-Based Accounting, Request for Comments: 3334, October 2002
- [TaHo04] Application and Analyses of Cumulative Sum to Detect Highly Distributed Denial of Service Attacks using Different Attack Traffic Patterns", H. Takada and U. Hofmann, IST INTERMON newsletter issue 7, Feb 2004
- [ZsMD05] Tanja Zseby, Maurizio Molina, Nick Duffield, Saverio Niccolini, Fredric Raspall: Sampling and Filtering Techniques for IP Packet Selection, Internet Draft <draft-ietf-psamp-sample-tech-06.txt>, work in progress, February 2005
- [ZsBB05] Tanja Zseby, Elisa Boschi, Nevil Brownlee, Benoit Claise, IPFIX Applicability, Internet Draft <draft-ietf-ipfix-as-05.txt>, work in progress, May 2005
- [TaAl02] An Empirical Analysis of NATE Network Analysis of Anomalous Traffic Events. Carol Taylor, Jim Alves-Foss, CS Department, U Idaho
- [SiPa04] V. Siris & F. Papagalou; Application of Anomaly Detection Algorithms for Detecting SYN Flooding Attacks; Proceedings of IEEE Global Telecommunications Conference (Globecom 2004), Dallas, USA, 29 November - 3 December 2004.
- [WaZS02] H Wang, D Zhang, K G Shin, Detecting SYN flooding attacks, in Proc IEEE INFOCOM, 2002