Lossless Compression for IP Flow Information Export (IPFIX)

draft-muenz-ipfix-compression-00

Gerhard Münz, Lothar Braun

72nd IETF Meeting, Dublin, 2008

Starting Point

- Paper at E2EMON 2008 [1] on compression of exported packet data

- **Obvious benefits**
  - less bandwidth required between Exporters and Collectors
  - less network I/O performance and smaller buffer sizes at Exporters and Collectors
  - reduced risk of network congestion ➔ reduced losses

- **What about IPFIX?**
  - bandwidth efficiency not an explicit design goal (not mentioned in RFC3917)
  - yet, we do have mechanisms which help reducing the data volume, e.g.
    - reduced size encoding
    - reducing-redundancy with options
    - biflow export
  - and much more, if information loss is acceptable
    - sampling and filtering
    - aggregation

[1] Politopoulos, P., Markatos, E., and S. Ioannidis, “Evaluation of Compression of Remote Network Monitoring Data Streams,” IEEE Workshop on End-to-End Monitoring Techniques and Services E2EMON 2008, April 2008.

Experiments with Lossless Compression

- **DEFLATE (RFC1951)**
  - LZ77 and Huffman coding
  - used in various protocols and applications, e.g. IMAP, SIP, gzip, PNG
- **Applied to Data Sets of different size**
  - 10, 20, 40, 60, and 100 Data Records
  - ZLIB compressed data format (RFC1950)
    - header (2 bytes) and CRC trailer (4 bytes)
- **Measurement data**
  - Flow Records from two networks
    - university network
    - ISP backbone
  - Packet Reports
    - network of a research institute

Compression of Flow Records

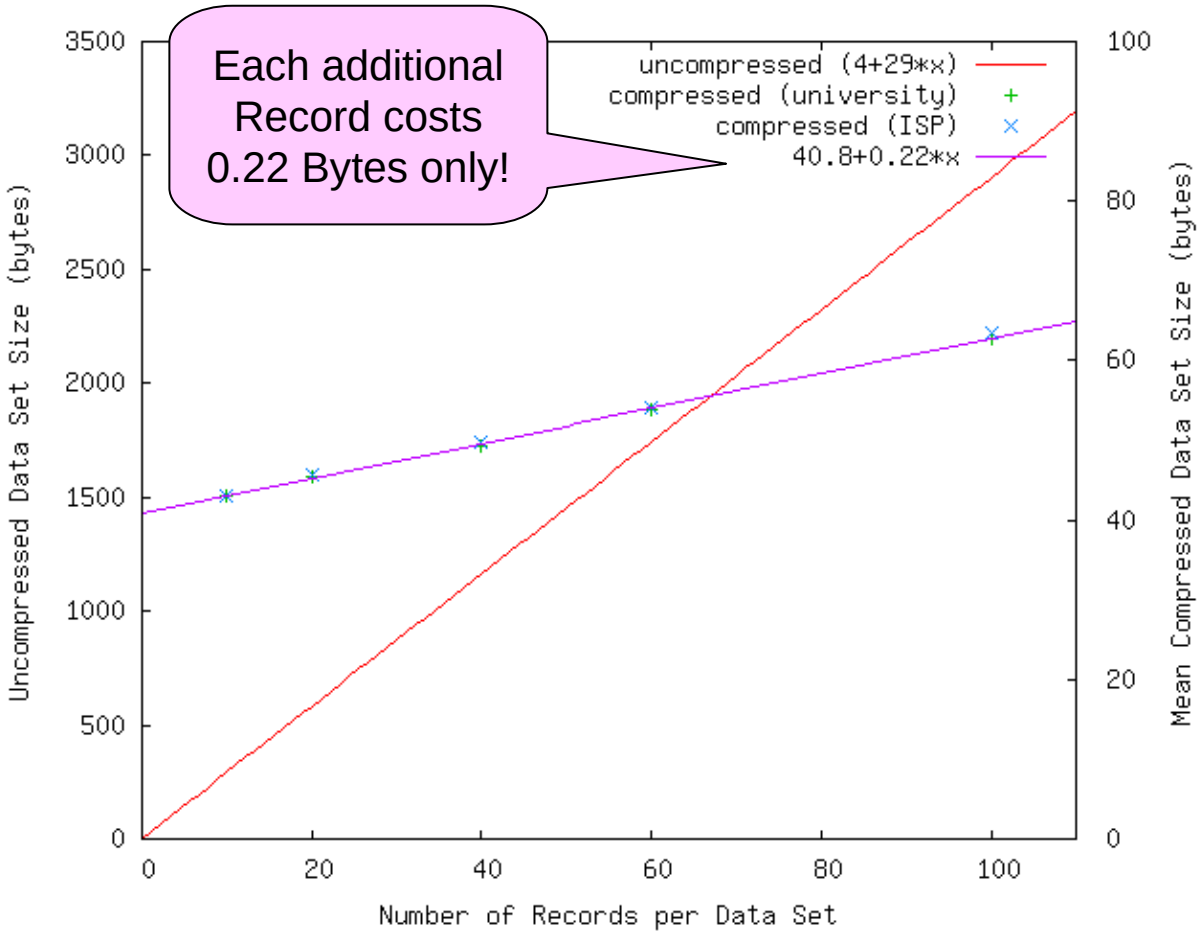

Template:

Field No	Information Element	Length
1	sourceIPv4Address	4
2	destinationIPv4Address	4
3	sourceTransportPort	2
4	destinationTransportPort	2
5	protocolIdentifier	1
6	octetDeltaCount (*)	4
7	packetDeltaCount (*)	4
8	flowStartSeconds	4
9	flowEndSeconds	4

(*) reduced size encoding

➔ Mean compression ratio from 6.8 (10 records) to 46 (100 records)

Compression of Packet Reports containing Packet Sections

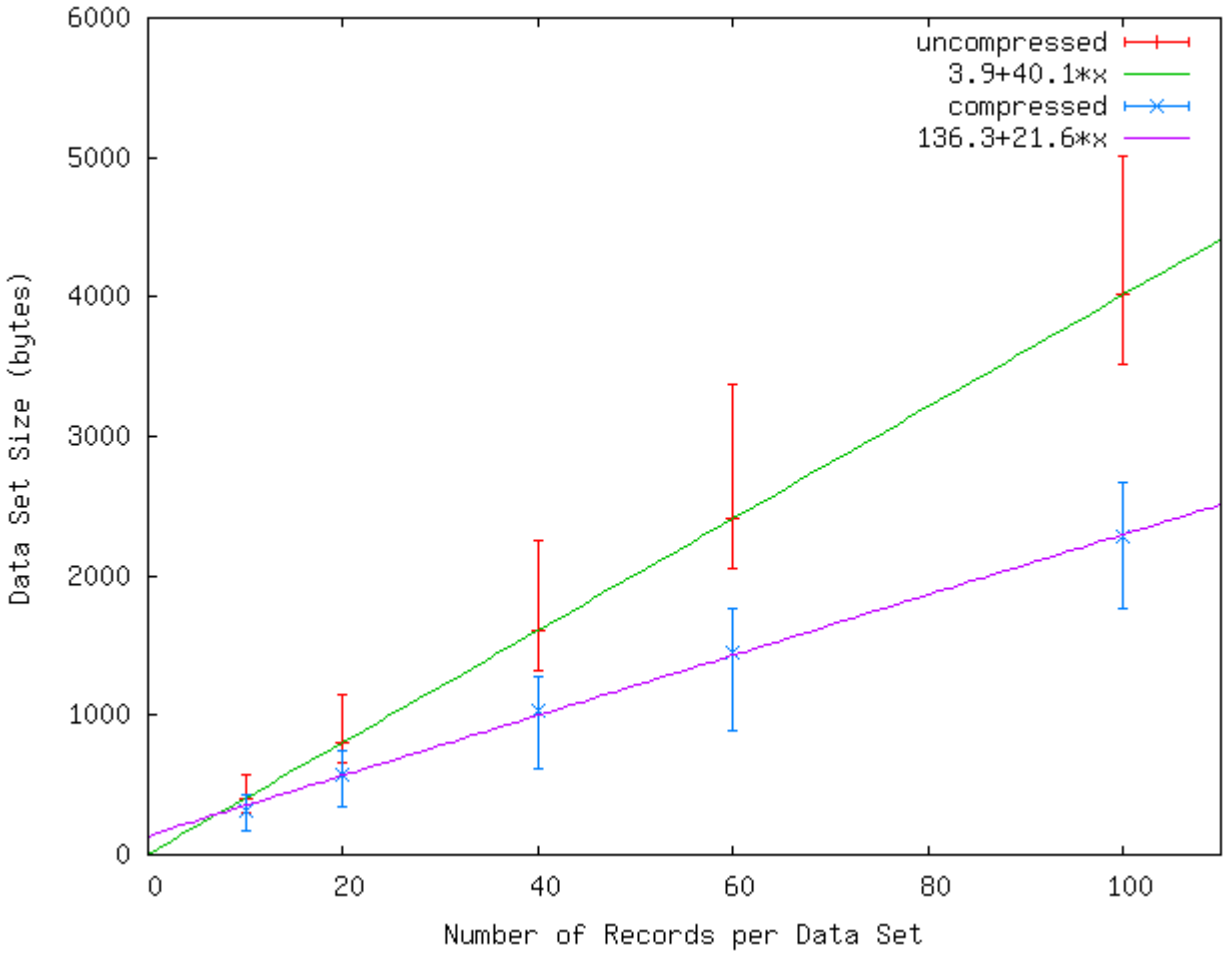

Template:

Field No	Information Element	Length
1	observationTimeSeconds	4
2	ipHeaderPacketSection (*)	variable

(*) IP and transport layer

➔ Mean compression ratio from 1.27 (10 records) to 1.76 (100 records)

Realization of IPFIX Compression

- **Compressed IPSec tunnel using IPComp (RFC3173)**
  - compression at IP datagram level
- **TLS/DTLS compression (RFC3749)**
  - compression at level of SCTP messages, UDP datagrams, TCP stream
- **IPFIX extension I: Compressed IPFIX Message**

- **IPFIX extension II: Compressed Data Sets**
  - Deflate Template Set (same structure as normal Template Set)

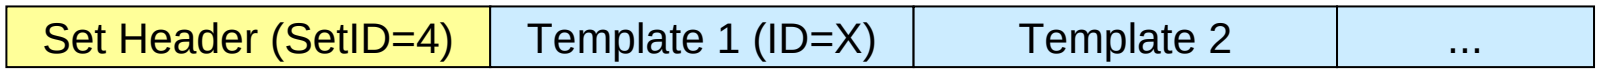

  - corresponding Data Set containing compressed Data Records

Conclusion

- **Significant compressibility of Data Records**
  - very high for typical Flow Records
  - much lower for Packet Reports containing packet sections
- **Possible realizations based on network and transport layer**
  - IPSec tunnel
  - TLS/DTLS (mandatory according to RFC5101!)
- **We can think of more flexible IPFIX-specific solutions**
- **Not sure, if this is an IPFIX standardization issue**

Thank you.