

IPFIX Export per SCTP Stream

IPFIX IETF-71 March 10th, 2008

<draft-claise-ipfix-export-per-sctp-stream-03.txt>

Benoit Claise, Paul Aitken, Andrew Johnson, Gerhard Muenz

Potential “Problems” with RFC5101

- **Potential head of line blocking (but also network delay, re-ordering and packet loss) could imply that the data records in a different SCTP stream might arrive before the template records**

RFC5101: collector MAY store

- **Potential head of line blocking (but also network delay, re-ordering and packet loss) could imply that data records are blocked and arrive after the template withdrawal message on another SCTP stream**

RFC5101: wait for 5 seconds

Is 5 seconds enough? Don't know. So we could be losing some data records.

Source of the problems

- **Data records from the same template ID sent across different SCTP streams and/or template definition on one SCTP stream and data records on an other**
- **Or conversely, that templates are not scoped to the streams**

Proposed Solution

- **“Any Data Sets associated with a Template Record MUST be sent on the same stream on which the Template Record was sent”**

Limit the Template ID to a single SCTP stream

Note that we don't change the scope of Template IDs in RFC5101

- **“All IPFIX Messages MUST be sent in order within a stream”**
- **“The Exporter MAY group related Templates and their associated Data Sets within a single stream so that the loss statistics are calculated for the group”**

Advantages

- **No risk that the Collecting Process discards Data Records while waiting for the Template Record to arrive**
- **No need to wait for 5 seconds to send the Template Withdrawal Message**
- **Template ID from a withdrawn Template MAY be reused immediately on the same stream**
- **Main advantage: data loss for each Template Record (or group) can be deduced**
 - From the Sequence Number in the IPFIX Message header (sequence counter of all IPFIX Data Records sent on this PR-SCTP stream)**
- **Backwards compatibility: per-stream exporter can still work with a non-per-stream collector.**

Conclusion

- **This specification offers some advantages**
- **We integrated the new solution discussed on the mailing list**
- **Any feedback?**
- **Working group item?**