An Information Model for the Monitoring of Network Securit y Functions (NSF)

draft-zhang-i2nsf-info-model-monitoring-03

Liang Xia, Dacheng Zhang
{frank.xialiang|dacheng.zhang}@huawei.com
Yi Wu
anren.wy@Alibaba-inc.com
Rekesh Kumar, Anil Lohiya
{rkkumar|alohiya}@juniper.net
Henk Birkholz
henk.birkholz@sit.fraunhofer.de

IETF 98, March ‘17, Chicago

Mailing List Discussion (recap)

- Does I2NSF need the work of NSF monitoring part? Yes
- Is producing a information model useful? Yes
- If we produce a YANG module, do we still need to publish the information model?
  Not yet decided
- What do you think of the content of the draft? Nobody dislike it, some people say it is a good start, others say it is a key part and very useful 

Purpose of Monitoring

- Ultimately, enable human remediation/intervention
  - Assistance via NSF automation (tier 1)
  - Automation of decision-making (tier 2)
- What has to be monitored?
  - Acquisition of “raw data” that “may be of interest”
  - Boils down to Event and Logs (a loosely defined aggregate of events)
- Publishing data via YANG event notification is one potential acquisition method
  - Why mentioning mothods in the context of an information model?
    The IM handles semantics and should leverage existing models:
  - I-D.ietf-netconf-netconf-event-notifications
  - But! Semantics and orchestration of event streams are out of scope there
  - Other obvious models we should consider? E.g. Sec-Event?

Updates (I)

- Refactoring to include the updated terminology (ext ensive but not complete – first pass)
- Classification of NSF Monitoring Data
  - Circles back to “Purpose of Monitoring”
  - Is a classification useful?
  - Which classes/categories are distinguishable by NSF
- Current Proposal
  - “all”
  - “violates a policy”
  - “impacts operation”

Updates (II)

- Controversial proposal?
  - Feedback ranges from
    - Arbitrary distinction without meaning, to
    - This is exactly how we do things, matching our processes
- Vital question
  - Which NSF / “NSF facing direction” categorizes?
  - How could a classification be represented?
    - Stream semantic
    - Explicit categorization by annotation of payload
    - Implicit categorization by annotating a “basis for categorization”
    - Other approaches?

Logs and Events

- Question to the group
  - Are logs just an aggregate of retained events repre sented in a loosely defined format that requires sp ecialized function to derive “past events” from?
  - Can logs be treated as records of past events or is there more to it?

Next Steps

- Deciding on a way to classify Information Elements in a way that supports the purpose of monitoring
- Clustering and labeling the existing set of Information Elements accordingly
- Assess Event Stream Solutions/Drafts

- Keep on improving…

Thanks!

The authors