Operations, Administration and Maintenance (OAM) features for RAW

draft-ietf-raw-oam-support-06

Theoleyre, Papadopoulos, Mirsky, **Bernardos**

RAW WG : IETF 116 : March 31, 2023

Updates

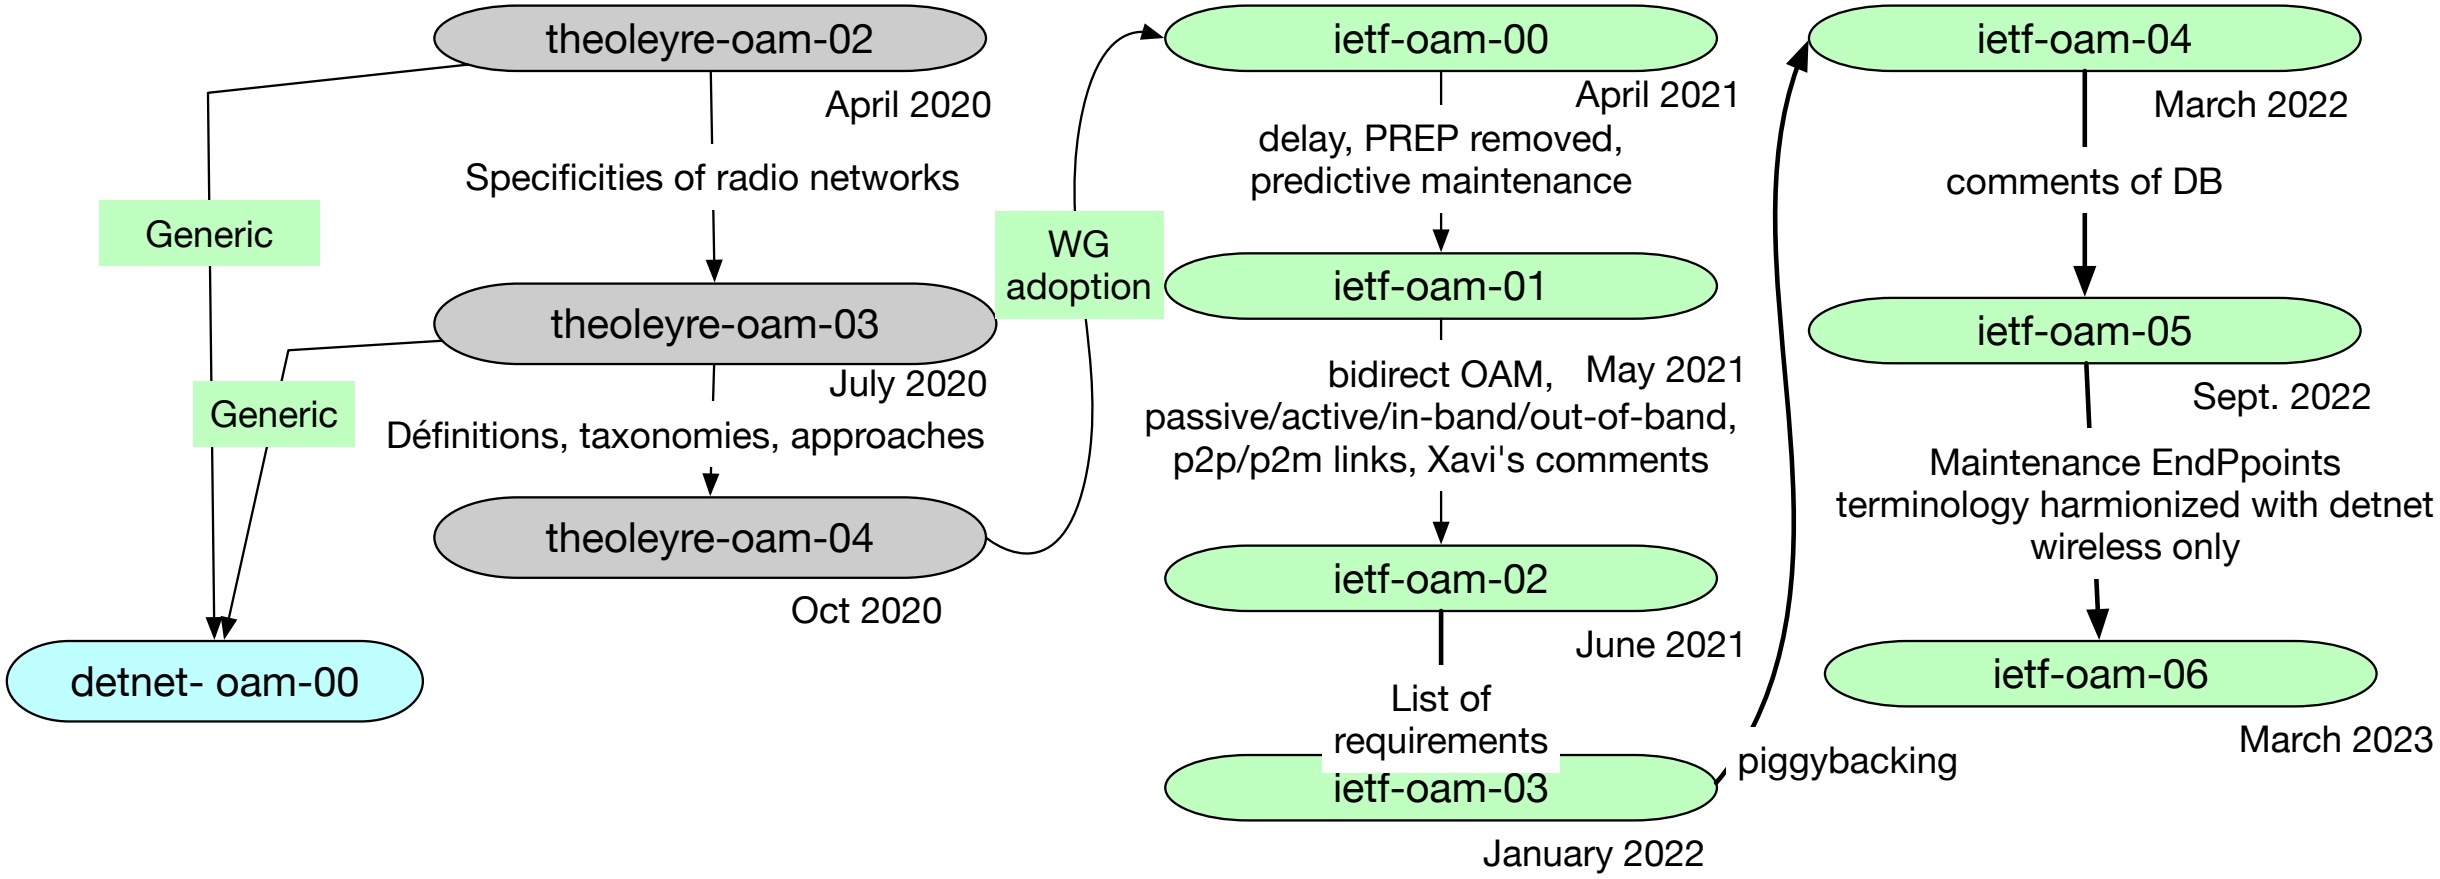

https://github.com/raw-wg/raw-oam/blob/master/README.md

Update - Terminology

- Maintenance EndPoint (MEP)
  - Entity that creates or reacts to OAM messages
- Maintenance Intermediate Point (MIP)
  - Entity that *responds* to OAM messages
- OAM Domain
  - Monitored domain (with MEPs on its edges)

Update - Clarifications

- Interactions between L2 and L3

Last but not least, radio links present volatile characteristics. If the wireless networks use an unlicensed band, packet losses are not anymore temporally and spatially independent. Typically, links may exhibit a very bursty characteristic, where several consecutive packets may be dropped because of, e.g., temporary external interference. Thus, providing availability and reliability on top of the wireless infrastructure requires specific Layer 3 mechanisms to counteract these bursty losses. Besides, Layer 3 has to be *informed* of the physical characteristics to make the right decision, and to avoid exacerbating physical issues (e.g., overloaded link because it became unreliable, overloaded radio channels).

- Distinction RAW / generic (detnet)
  - *« relevant for RAW-specific (aka wireless) segments »*
- PREOF citation
  - I-D.pthubert-raw-architecture
- Anycast
  - *« Anycast transmission typically exploit p2mp links »*

Update - Clarification

- Now two categories for faults:
  1. Fault detection
  2. Fault Identification

- Delay estimation in wireless
  - Queuing, medium access & retransmissions

Update - Clarification

- Revamped the prédiction subsection

5.2. Predictive maintenance

RAW needs to implement self-optimization features. While the network is configured to be fault-tolerant, a reconfiguration may be required to keep on respecting long-term objectives. Obviously, the network keeps on respecting the SLO after a node's crash, but a reconfiguration is required to handle future faults. In other words, the reconfiguration delay MUST be strictly smaller than the inter-fault time.

The network must continuously retrieve the state of the network, to judge about the relevance of a reconfiguration, quantifying:

*the cost of the sub-optimality: resources may not be used optimally (e.g., a better path exists);

*the reconfiguration cost: the controller needs to trigger some reconfigurations. For this transient period, resources may be twice reserved, and control packets have to be transmitted.

Thus, reconfiguration may only be triggered if the gain is significant.

5.2. Predictions

RAW needs to implement self-optimization features. While the network is configured to be fault-tolerant, a reconfiguration may be required to keep on respecting long-term objectives. The network must continuously retrieve the state of the network, to judge about the relevance of a reconfiguration. More precisely, the OAM mechanisms have to provide enough information to predict and quantify:

*the gain of the reconfiguration: what would be the network state after the reconfiguration (e.g., reduction of the bandwidth or energy consumption)?

*the reconfiguration cost: what is the cost (energy, bandwidth) to reconfigure the forwarding and management planes?

Wireless networks exhibit non linear dependencies among links / radio channels / technologies that complexify significantly such predictions.

Next Step

- Draft now stable
- Your comments, suggestions, questions always welcome and greatly appreciated
- WGLC?

Thank you!