




Reliable and Available Wireless Architecture

Presenter: Pascal Thubert


Authors: P. Thubert, G.Z. Papadopoulos, (+ contributors)

Status Since IETF 114 (-06 to -09)



- Review / help by DetNet chairs, Janos and Lou; added Lou as Contributor
 - Addressed Lou’s Comments (but some left open for discussion next slides)
 - Also Lou’s comments on PAREO TBD with Georgios
- Restructured for new readers
 - Simpler introduction => extracted section 3 on what RAW means
 - Section 5 is called “The RAW Control Loop” as opposed to “The OODA Loop”
- “RAW **and** DetNet” clarifications
 - RAW enhances DetNet to improve the protection against link errors such as transients
 - The RAW methods mostly applicable to wired links as well, e.g., for energy savings

Issues



- Issue: the text may be read as if we are changing the wireless technology

=> New: “Operating at the Layer-3, RAW does not change the wireless technology at the lower layers. OTOH, it can further increase diversity in the spatial, time, code, and frequency domains...”

Issue: is RAW adding new conceptual services:

=> New: “RAW improves the DetNet services by providing elements that are specialized for transporting IP flows over deterministic radios technologies such as listed in [RAW-TECHNOS].”


- Issue: talking about L2-aware *reachability* versus routing.

=> L2 aware, yes, though multihop; still L3 routing executed by forwarding (~like ECMP LB)

- Issue: promiscuous overhearing can be found on wires

=> Unsure which action to take

Issues



- PAREO Issue: can RAW operate without protection or a different protection mechanism?


=> I suspect RAW could steer traffic dynamically within Track with No PAREO at all but dubious. The whole point is end-to-end reliability when any one L2 hop is not reliable => diversity/ Extending PAREO with other protection schemes would be welcome though. New text needed?

Issue: Track and Segment terminology (same as / redundant with TE path?):

=> Track is the term used in the 6TiSCH (one form of deterministic wireless) Architecture (RFC 9030) and RPL drafts to refer to the same concept as RAW => we are not introducing a new term

=> Track refers to a potential that is used statistically, as opposed to an actual (TE) path that was used by a packets. IOW the Track is the aggregation of all the possible paths that the packet or portions of the packet might traverse. Think orbital vs orbit.


=> The Track could be seen as a combinatory aggregation of protection paths, where one such path is selected by the PSE for one or a few packets. We still need the term path to refer to the experience of RAW packets. Q: Is RAW always TE?



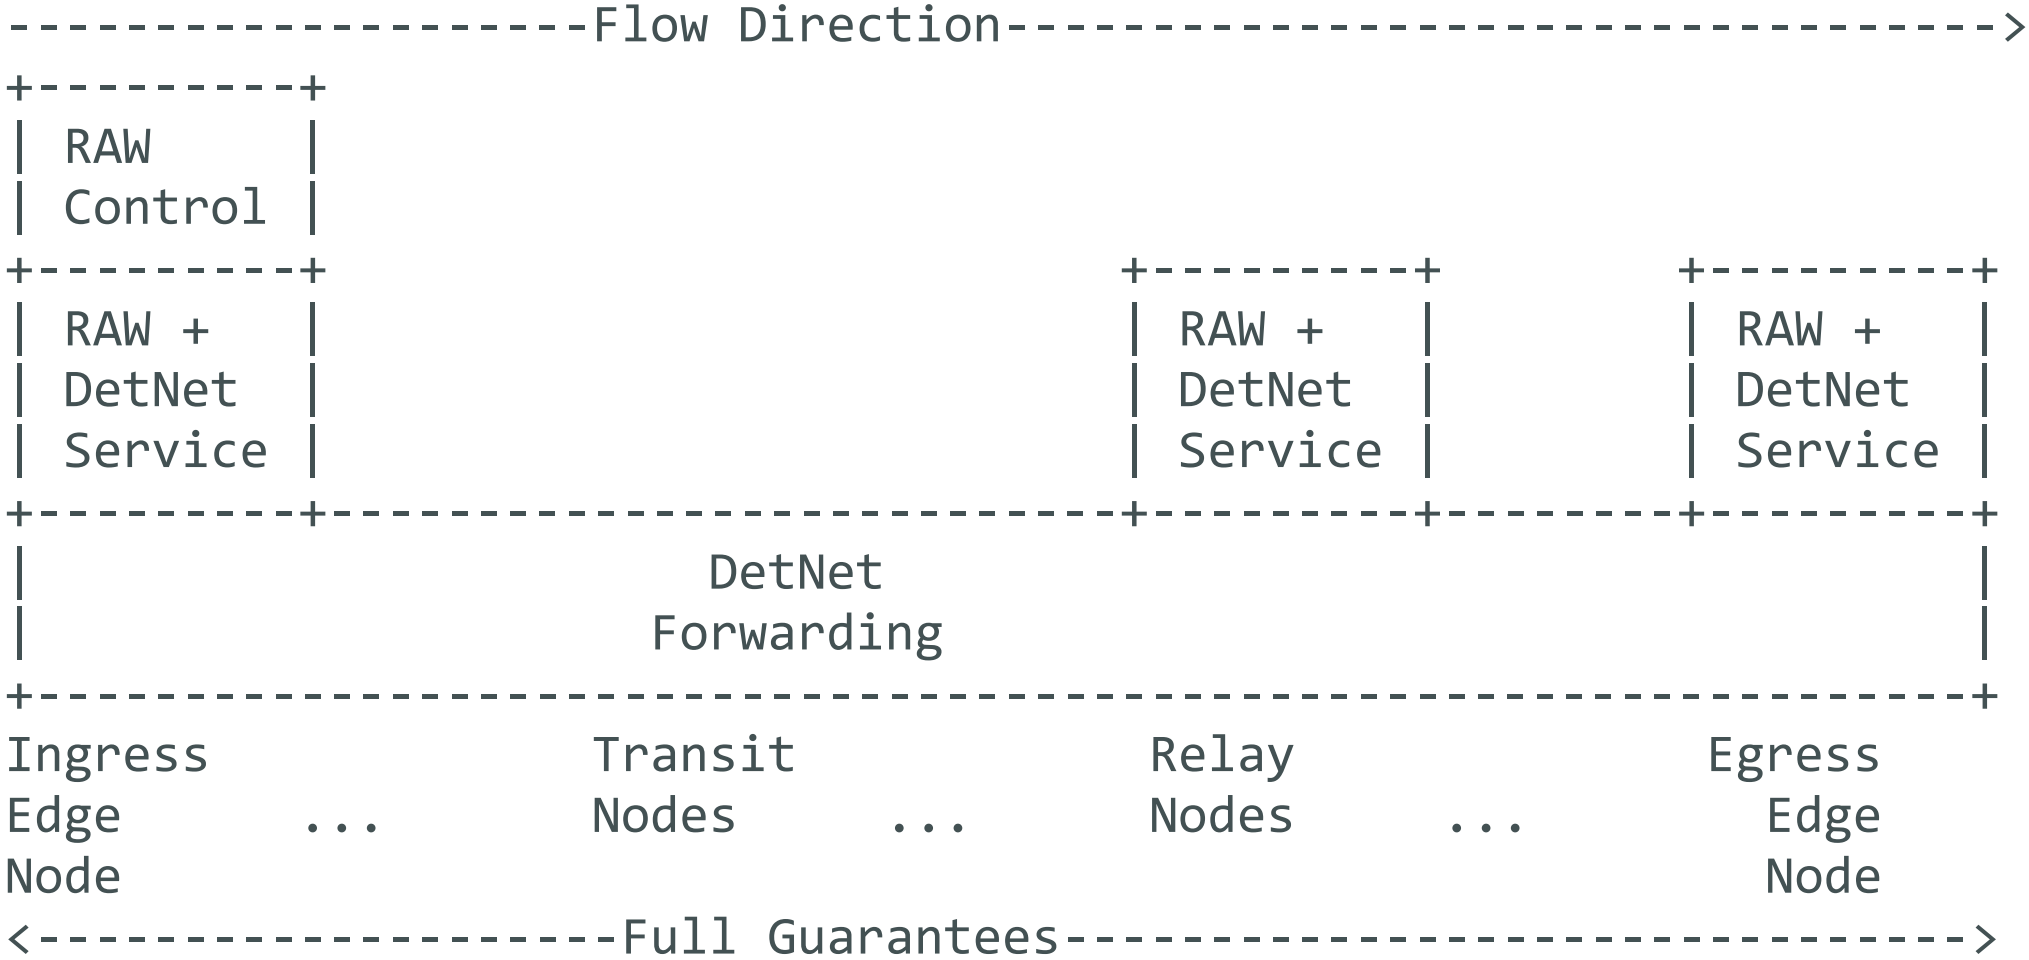

Questions ?

Backup

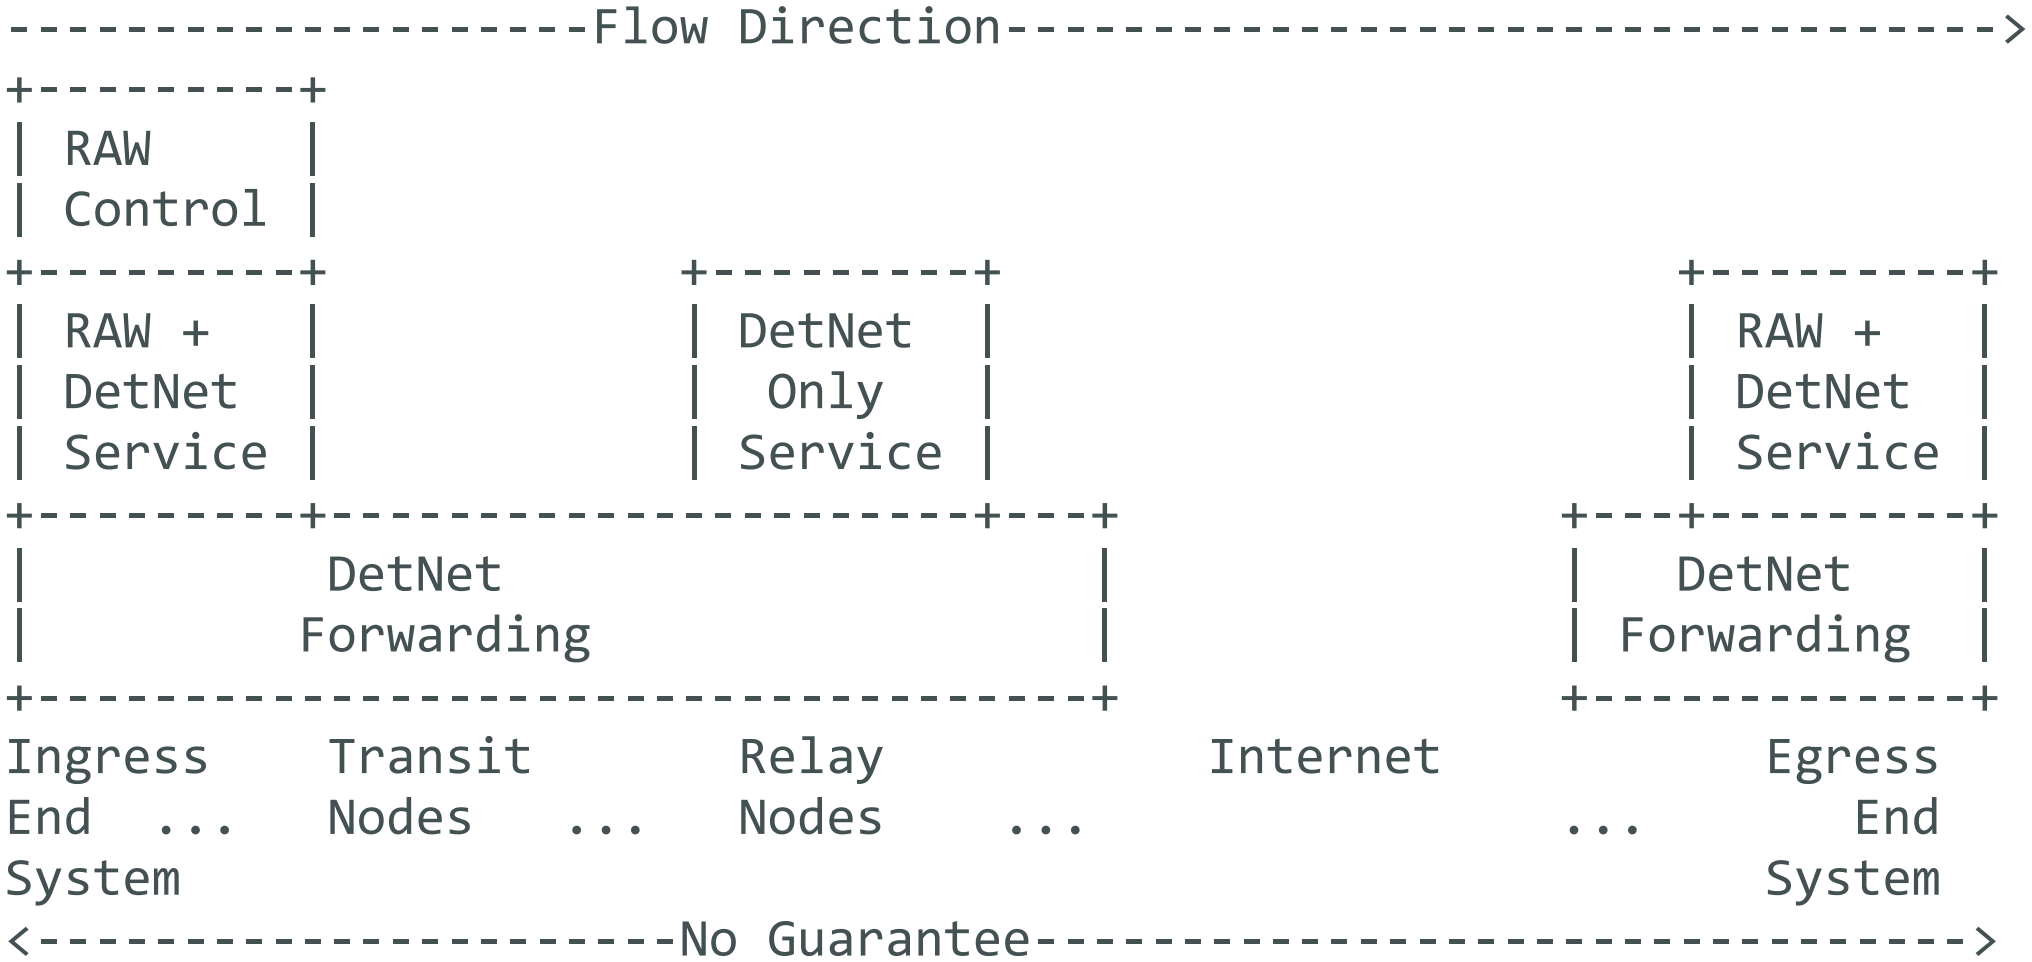

RAW DetNet Services



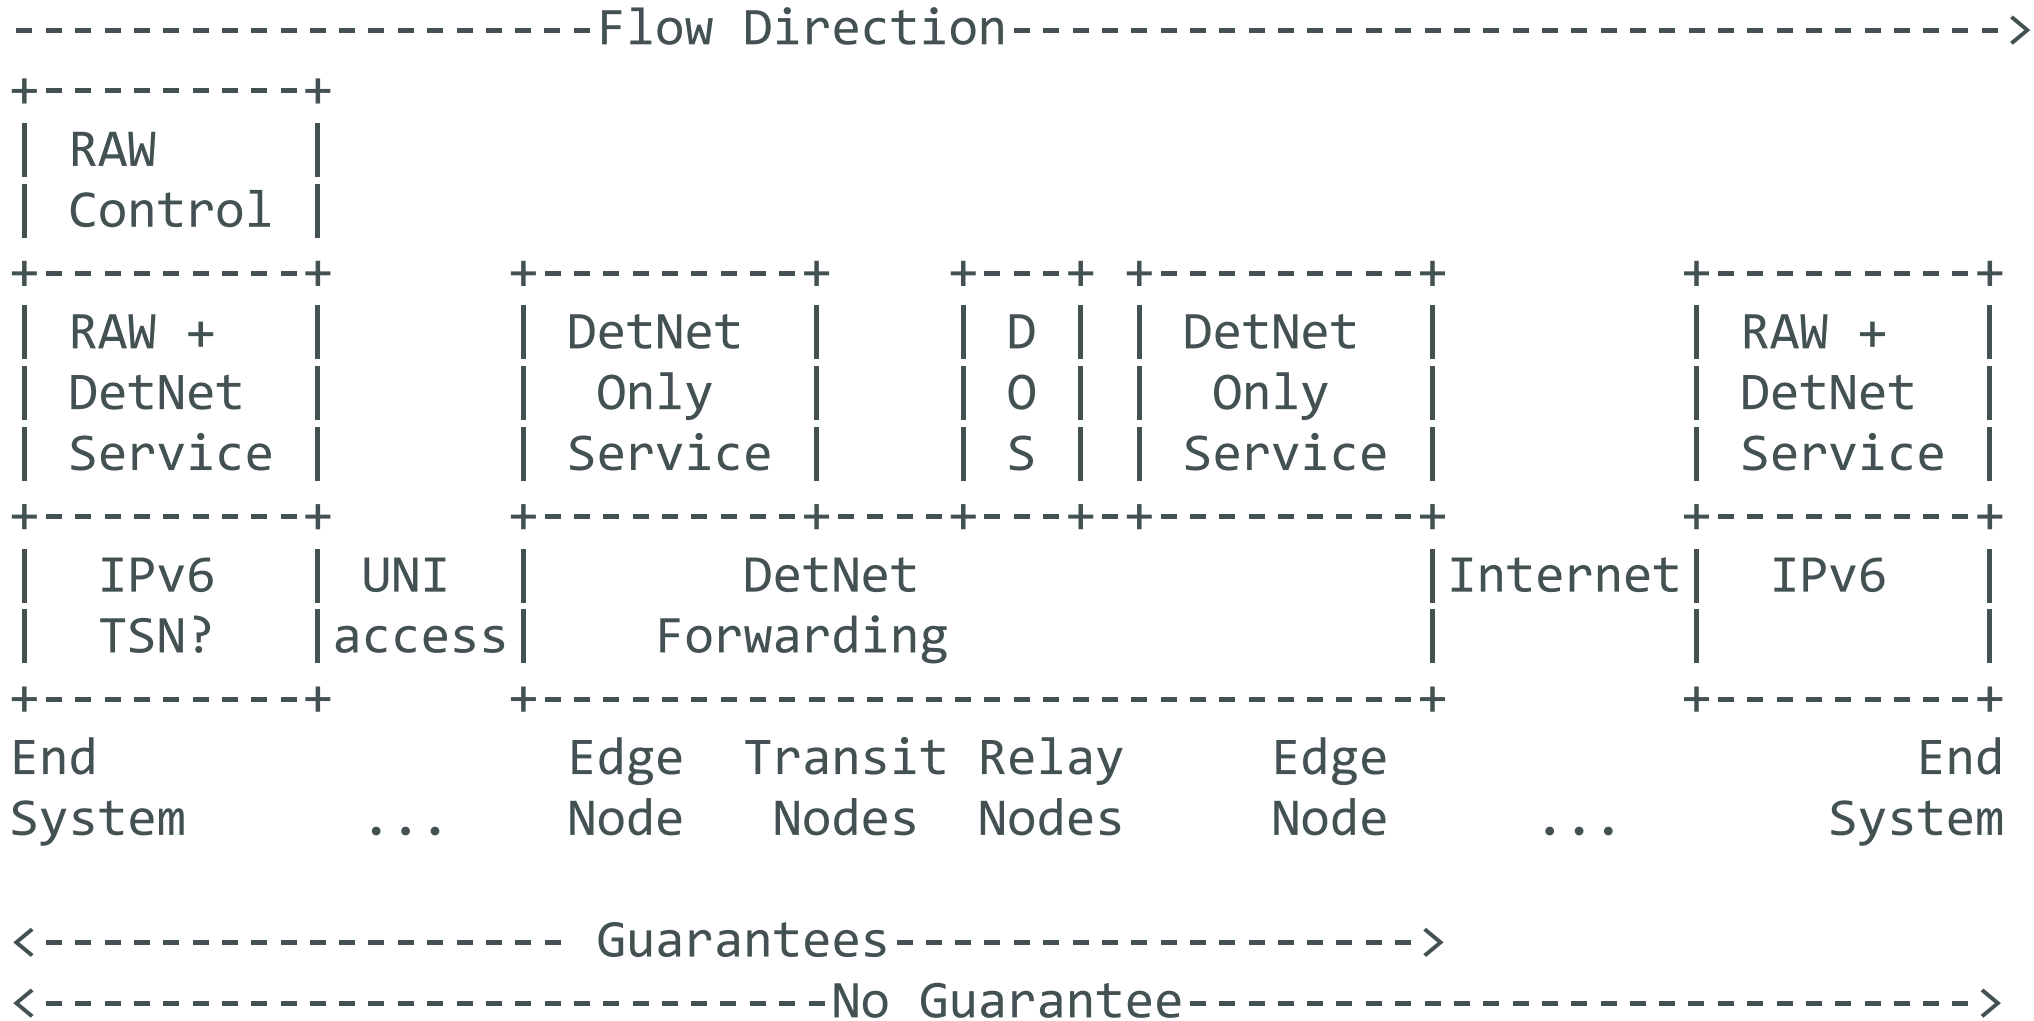

(Strict) RAW over DetNet




Loose RAW



Variation (not illustrated in doc)



Dependencies




RAW improves the reliability of transmissions and the availability of the communication resources, but does not provide scheduling and shaping, so RAW itself does not provide guarantees such as latency for the application payload. Rather, it should be seen as a Dynamic optimization of the use of redundancy to maintain it within certain boundaries.

Guarantees such as bounded latency depend on the upper layers (Transport or Application) to provide the payload in volumes and at times that match the contract with the DetNet sublayers and the layers below. Excess of incoming traffic at the DetNet Ingress will cause either dropping, queueing, or reclassification of the packets, and entail loss, latency, or jitter, and moot the guarantees that are provided inside the DetNet Network.


When the traffic from upper layers matches the expectation of the lower layers, RAW still depends on the lower layers to provide the timing and physical resources guarantees that are needed to match the traffic SLA. When the availability of the physical resource varies, RAW will act on the distribution of the traffic to leverage alternates within a finite set of potential resources.

Balázs Varga 's questions (1of3)




- Mix of various OSI Layer functions: The definition of PAREO seems to be very confusing, as it contains a mix of Radio specific and DetNet specific functions. It is confusing as the referred functions work at different layers (e.g., HARQ is part of Radio at L1/L2 vs. PREOF is part of DetNet at L3) and have different "range" (radio acts on radio links vs. PREOF acts across several hops, maybe even end2end). This mix makes it unclear which OSI layer the RAW architecture belongs to. Could you please clarify?
- DetNet leverages lower layers, and RAW will augment that usage to hint about transmission suggestions. Lower Layers do what they like but if the API allows to pass hints, we'll leverage that. In particular, we'll need reliability and timing hints like suggest X retries (min, max), send unicast (one next hop) or multicast (overhearing). The other way around RAW will need hints about L2 conditions like L2 triggers (RSSI, LQI, ETX...) over all the wireless hops. This will be used by both PCE and PSE. Bottom line: to do its job, L3 works on abstractions of L2; in the (dynamic) case of wireless there's more of it.

Balázs Varga 's questions (2of3)




- The modeling of Radio components from deterministic networking perspective seems to be unclear and different from current work of radio related SDOs. The draft states that the concept is agnostic to the radio technology and agnostic to whether or not radio mesh is applied. Nonetheless, the model applied for the Radio layer seems to be unclear. Please note, e.g., that DetNet Study Item ongoing in 3GPP SA2 models the 5G System as a DetNet router
- The case where the 5G network shows as one virtual switch is opaque to RAW as it is opaque to DetNet. I agree we can improve the discussion on interaction with lower layers in the dependency section following your suggestion. I hope we have more details on the mike...

Balázs Varga 's questions (3of3)



- Related to the Q1, the relationship of RAW and DetNet Layers is unclear. Along the lines of your definition in "Section 3. The RAW Conceptual Model: ... The RAW Nodes are DetNet relays that are capable of additional diversity mechanisms and measurement functions related to the radio interface ..." whereas same section states that " ... the non-RAW subnetwork can be neglected in the RAW computation ...".
- I guess we'll need to clarify. The non-RAW is when RAW is not end to end and latency cannot be guaranteed (loose RAW). Again, I hope there's discussion on the mike.
- See email thread for drawings

Lou's comments (what is RAW?)



1) What does the term "RAW" refer to in the context of the architecture?

Is it a new/standalone set of mechanisms or is it an addition or an extension, or a usage of IETF defined technologies?

I find that reading the architecture, I'm really unsure. The current working is a bit mixed, e.g.,


Reliable and Available Wireless (RAW) provides for high reliability and availability for IP connectivity over a wireless medium.

this sounds like something new/independent

It builds on the DetNet Architecture and discusses specific challenges and technology considerations needed to deliver DetNet service utilizing scheduled wireless segments and other media, e.g., frequency/time-sharing physical media resources with stochastic traffic.

this sounds evolutionary.

Lou's comment; tracks vs TE protection paths



Lou said:

"

In reading the definition of the tracks, the term seems quite aligned/similar to TE protection paths and segments. Keep in mind that DetNet PREOF is just one form of service restoration supported in IETF TE, i.e., the 1+1 form. A track reads to me to be something that can be composed or combine 1:1, 1:N and even 1+N, and have interesting (uncoordinated) protection switching based on actual network/link (channel) state. I suspect we can accomplish the same objectives as Tracks and stay consistent with existing DetNet and TE(AS) terminology.

"

My short answer:

There's too much preconception with the term path as an experience and too much slight overload of the term path in different WGs, sometimes without a clear terminology but rather an intuition that means confusion. We need a term that reflects the statistical expectation and overloading path will entertain / augment the existing confusion. Defining Track allows us to point explicitly on what happens at DetNet and RAW that is really new.

Longer answer:

My view is that the concept of path in the art of IETF is incompatible with that of Track because a path has the classical IP expectation of 1 packet in, the same packet along, and then out. In TE, even in the 1+N case, the packet is replicated first, and then each copy follows a path where the assumption above stands. A copy is an atom that remains in its integrity along the path. The path can thus be defined after the fact as the experience of the (copy of the) packet. And we can define a protection path as a set of such paths for multiple copies, either sent at the same time or deferred.


I believe that in RAW, and even in DetNet, the capability to replicate inside the network already defeats that definition. When a packet is replicated inside the network, which one is the original that follows the path? What is the path for the copy? If a relay fragments a packet into N frags, uses network coding $N \rightarrow P$, $P > N$ between the fragments as redundancy technique, and then distributes the fragments across the feasible next-segments within the Track, the packet is not more an atom, and there's no "path" that the packet experiences. The packet is literally disseminated (as opposed to flooded as an atom) and reconstituted at arrival.

If I'm unclear, maybe a quantic analogy will help. The path is the observation/measurement of how an atomic packet traversed the network (which hole the photon passed through). In the case of network coded fragments, there are superimposed "path" states, each with its own probability. The Track describes the superimposition, not the measurement. It is expressed as the set of statistical possibilities for a future packet, not the experience of one.

Note: The sum of probabilities is typically more than 1 due to redundancy methods, and the PSE dynamically DECIDEs the new values for those statistics to place them on segments that appear to work at this time based on ORIENTATION by the PCE (D and O in OODA).

In terms of IETF inheritance, the term Track is already defined in the art of a deterministic wireless RFC (RFC 9030) and a method to program Tracks in a wireless network is being defined at ROLL (draft-ietf-roll-dao-projection). The definition of Track at RAW, ROLL, and 6TiSCH are consistent, though ROLL can only build Tracks that are DODAGs (it cannot build bidir North/South segments).

Lou's comments: PCE-only?




4) Is RAW limited to PCE-based centralized solutions?

DetNet introduced the term Controller Plane to cover all types of control supported by the IETF TE architecture, i.e., fully centralized, fully distributed, or any hybrid combination of centralized/distributed control. The Architecture reads as supporting only one combination of - PCE for paths, PSEs for tracks (aka protection segments) . PSEs read as also doing the actual protection switching, but this is outside the scope of this comment.

Hereto, I see no reason for the architecture to limit the scope of the Controller Plan solutions that could be standardized as part of RAW. (Yes PCE-based approaches are likely the first to be standardized, but that's not an architecture level decision.) “

> I agree with Lou. But how to compute Tracks and provide Orientation in a distributed fashion is beyond me.

Lou's comments: Wireless only?




2) Are RAW solutions limited to IPv6 and a limited set of wireless technologies? I think the framework says/implies no, but the architecture is less inclusive. e.g., RAW provides DetNet elements that are specialized for IPv6 flows [IPv6] over selected deterministic radios technologies [RAW-TECHNOS].

I would expect that at least at the Architecture level that there would be no exclusion of IETF networking technologies (including v4 and MPLS) and any SDO-defined wireless technology. I certainly understand needing to focus and prioritize work on specific technologies, but that is practical choice not a limitation that should be codified in the architecture.

Somewhat related, but of less importance, it's probably worth mentioning that RAW forwards/switches at the IP, not link layer.

> I agree with Lou. But how to compute Tracks and provide Orientation in a distributed fashion is beyond me.

Lou's comments (rewrite abstract)



OLD

Reliable and Available Wireless (RAW) provides for high reliability and availability for IP connectivity over a wireless medium. The wireless medium presents significant challenges to achieve deterministic properties such as low packet error rate, bounded consecutive losses, and bounded latency. This document defines the RAW Architecture following an OODA loop that involves OAM, PCE, PSE and PAREO functions. It builds on the DetNet Architecture and discusses specific challenges and technology considerations needed to deliver DetNet service utilizing scheduled wireless segments and other media, e.g., frequency/time-sharing physical media resources with stochastic traffic

NEW

Reliable and Available Wireless (RAW) provides for high reliability and availability for IP connectivity across any combination of wired and wireless network segments. The RAW Architecture extends the DetNet Architecture and other standard IETF concepts and mechanisms to adapt to the specific challenges of the wireless medium. This document defines an architecture element for the RAW data plane, in the form of an OODA loop, that optimizes the use of constrained spectrum and energy while maintaining the expected connectivity properties. It also introduces a new Control plane Function to prepare alternate paths to go around local failures. The loop involves OAM, PCE, and PREOF extensions, and a new component called the Path Selection Engine (PSE).


The new text does 2 things:

- Use Lou's proposal with the exception of the explicit reference to TE
- Present the PSE as a data plane component to the RAW architecture, more may come.

Along the same line, should we change the title to something more focused on the PSE or the OODA loop?

Discussion: whether to mention TE(AS)

Next Steps



- An Ad hoc team met, and also the draft was discussed at the last interim
- Lou still has some questions / issues open that we need to sort out
- How far are we from WGLC?
- Note:

The RAW architecture is normative ref. to draft-ietf-roll-dao-projection, which is closing WGLC.

The draft provides a way for the controller to set up RAW Tracks using new signaling in RPL.