

RIFT: Routing in Fat Trees

draft-ietf-rift-rift-20

IETF 119

Brisbane

Tony Przygienda

Jordan Head (presenting)

Pascal Thubert

Bruno Rijsman

Dmitry Afanasiev

Well, we passed IETF last call

- A huge “thank you” to our ADs, past and present:
 - Alvaro Retana
 - Jim Guichard
 - John Scudder
 - Andrew Alston
- And also, to our reviewers:
 - Donald Eastlake
 - Daniam Henriques

What's new in version 20?

- **Structural and Editorial**

- The “Overview” immediate sub-sections are now a top-level sections
- Additions and edits in the “Terminology” section
- Lots of formatting consistency fixes:
 - Backticks
 - **Bold vs. *italics***
 - System ID vs. system ID
 - etc.
- Even more phrasing improvements

What's new in version 20?

- **Figures and Illustrations**

- New ASCII version of Figure 1
- Replaced “via” with “@” in Figure 1
- Removed SVG version of Figure 33
- Edited Figure 8’s title to use “90 degrees” instead of “90o”
- Fixed Figure 8 so E (East) and W (West) weren’t reversed
- More IETF tooling challenges
 - Upload tool’s PDF rendering breaks SVGs

- **Clarified “Northern View” vs. “Southern View”**

- **Northern View:** From the ToF down.
- **Southern View:** From the leaf up.

What's new in version 20?

- **Normative Clarifications**

- **TIDE Generation**

- In the previous version, the language describing the algorithm could potentially allow an implementation to not send all of the TIDEs.
 - HEADERS = At most TIRDEs_PER_PKT headers in TIEDB starting at NEXT_TIDE_ID or higher that SHOULD be filtered by is_tide_entry_filtered and MUST either have a lifetime left >0 or have no content
 - The “*at most*” language could make it legal for an implementation to do the following:
 - Let's say the database contains TIDEs 1, 2, 3, 4, 5 and TIRDEs_PER_PKT is set to 5.
 - NEXT_TIDE_ID = 1
 - HEADERS = 1, 2 (i.e., not all 5)
 - START = 1 (“else” case)
 - END = MAX_TIEID (“then” case)
 - Send 1, 2 as TIDE with range of 1, MAX_TIEID
 - NEXT_TIDE_ID = MAX_TIEID

What's new in version 20?

- **Normative Clarifications**

- **TIDE Generation (continued)**

- **FILTERED_TIEDB**: A filtered view of TIEDB, which retains for consideration only those headers permitted by `is_tide_entry_filtered` and which either have a lifetime left > 0 or have no content.
 - This is NOT a new construct; this is *purely* conceptual. Preceding text in the draft indicates this.
 - **HEADERS** = Exactly `TIRDEs_PER_PKT` headers from `FILTERED_TIEDB` starting at `NEXT_TIDE_ID`, unless fewer than `TIRDEs_PER_PKT` remain, in which case all remaining headers.
 - *“The algorithm will intentionally enter the loop once and send a single TIDE even when the database is empty, otherwise no TIDEs would be sent for in case of empty database and break intended synchronization.”*

What's new in version 20?

- **Normative Clarifications**

- **Attaching Prefixes**

- *“The rule of inheritance **MUST** also be maintained when a new prefix of intermediate length is inserted...”*
 - *“... As the aggregating prefix changes, all the negative routes **MUST** be recomputed, and then again the process may recurse in case of nested negative prefix aggregations.”*

What's new in version 20?

- **Normative Clarifications**

- **TIE Processing**

- *“LSDB is also expected to periodically re-originate the node's own TIEs. Originating at an interval significantly shorter than default_lifetime is **RECOMMENDED** to prevent TIE expiration by other nodes in the network which can lead to instabilities.”*

- **RAIN: RIFT Adjacency Inrush Notification**

- Previously “ECN” (Explicit Congestion Notification)
 - *“A node MAY set in its LIEs the corresponding you_are_sending_too_quickly flag to indicate to the neighbor that it **SHOULD** flood Node TIEs with normal speed and significantly slow down the flooding of any other TIEs.”*

What's new in version 20?

- **Normative Clarifications**

- **Level Determination Procedure (ZTP)**

- **Old:** *“A node that lost all adjacencies with HAL value MUST hold down computation of new DERIVED_LEVEL for a short period of time...”*
 - **New:** *“A node that lost all adjacencies with HAL value MUST hold down computation of new DERIVED_LEVEL **for at least one second...**”*

What's new in version 20?

- **Further Clarification**

- **Balancing Bandwidth (BAD Computation)**

- **Context:**

- *“On a node, L, use Node TIEs to compute from each non-overloaded northbound neighbor N to compute 3 values:”*

- **Old**

- **L_N_u**: as sum of the bandwidth available to N
 - **N_u**: as sum of the uplink bandwidth available on N
 - **T_N_u**: as sum of $L_N_u * OVERSUBSCRIPTION_CONSTANT + N_u$

- **New**

- **L_N_u**: sum of the bandwidth available from L to N (to account for parallel links)
 - **N_u**: sum of the uplink bandwidth available on N
 - **T_N_u**: $L_N_u * OVERSUBSCRIPTION_CONSTANT + N_u$

What's new in version 20?

- **Further Clarification**

- **Balancing Bandwidth (BAD Computation)**

- **Context:**

- *“On a node, L, use Node TIEs to compute from each non-overloaded northbound neighbor N to compute 3 values:”*

- **Old**

- **L_N_u**: as sum of the bandwidth available to N
 - **N_u**: as sum of the uplink bandwidth available on N
 - **T_N_u**: as sum of $L_N_u * OVERSUBSCRIPTION_CONSTANT + N_u$


- **New**

- **L_N_u**: sum of the bandwidth available from L to N (to account for parallel links)
 - **N_u**: sum of the uplink bandwidth available on N
 - **T_N_u**: $L_N_u * OVERSUBSCRIPTION_CONSTANT + N_u$

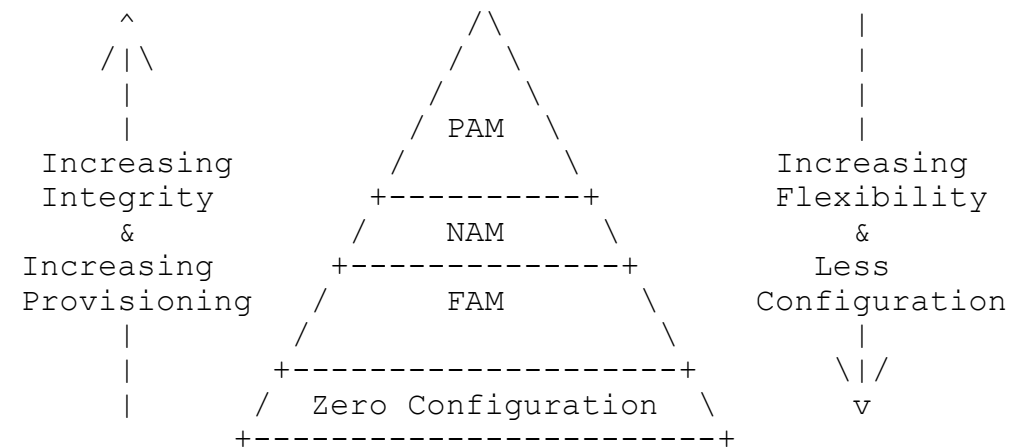
What's new in version 20?

- Removed the Level Association Model from the Security Model

Old



New



What's new in version 20?

- **References**

- Moved the Applicability draft reference to informative
- Removed references to RFC6830 (LISP) in favor of RFCs 9300 and 9301
- Added the DayOne book to the informative references
- Added a reference RFC4086 (Random Number Generation)
- Replaced valley-free routing (VFR) reference with an IEEE reference
- And more

What's new in version 20?

- **IANA Considerations**

- Tony had lots of discussion with IANA and clarified some language to conform to their procedures.
- Schema Versions bumped from 7.0 to 8.0
- Added Expert Review criteria

What's next?

- -21 will be published once the upload tool is open again, it addresses:
 - Comments from IANA on last review.
 - Re-adding Alankar Sharma as co-author.
- That's it. Next stop IESG.

Questions?