

Yang model for requesting Path Computation

draft-busibel-teas-yang-path-computation-02
IETF 98 – Chicago

Italo Busi (Huawei)
Sergio Belotti (Nokia)
Daniele Ceccarelli (Ericsson)
Victor Lopez, Oscar Gonzales de Dios (Telefonica)
Michael Scharf (Nokia)
Anurag Sharma
Yan Shi (China Unicom)
Ricard Vilalta (CTTC)
Karthik Sethuraman (NEC)

Summary of changes from IETF 97

- Added detailed connectivity matrix dimensioning consideration (section 3.1)
- Added new chapter 5 regarding «Path computation for multiple LSPs»
- Added section for security consideration
- Added Yang model for stateless RPC

Detailed connectivity matrix dimensioning

We have analyzed the set of constraints, and their cardinality, that impact the size of the detailed connectivity matrix

Constraints	Cardinality
Endpoints	$N(N-1)$ unidir
Bandwidth	Technology specific: <ul style="list-style-type: none">• in theory 200 bandwidth values/ranges for ODU/ODUflex [<i>may be reduced in practice</i>]• in practice 4-7 (5 on average) bandwidth ranges for IP
Metrics	8: IGP, TE, hop, MLP, MBP, Delay, Delay Variation, Loss
Bounds	practice 30: 6 metrics (IGP, TE, hop, Delay, Delay Variation, Loss) x 5 ranges
Priority	8 values for setup priority
Local protection	2: true/false flag
Administrative Color	Theoretical: 3×2^{32} (include, exclude-any, exclude-all and 32bits) Practical: few values are used
SRI G	High number [<i>not estimated</i>]

Feedback from the analysis

An approach based only on detailed connectivity matrix is hardly feasible with limited applications

Example: IP Networks


- Endpoints = $N*(N-1)$, Bandwidth = 5, Metrics = 6, Bounds = 20, Priority = 8, Local protection = 2 (no SRLG and no affinities)
- Number of paths: $24.960 * N(N-1) = 300.000$ for $N=4$
- 1K for each path json description: **300 Mbytes** for each domain
- 20% of paths change when a new deployment of traffic occurs:
60 Mbytes of change notifications for each domains traversed by the new e2e LSP

Path Computation for multiple LSPs

Domain A

Domain B

Domain C



3 RPC Requests
>= 1 Path(s)/RPC

9 RPC Requests
>= 1 Path(s)/RPC

3 RPC Requests
>= 1 Path(s)/RPC

VERSUS

1 RPC Request
• 1 ingress port
• 3 egress ports
>= 3 Paths

1 RPC Request
• 3 ingress ports
• 3 egress ports
>= 9 Paths

1 RPC Request
• 3 ingress ports
• 1 egress port
>= 3 Paths

Same number of computed paths but less RPC Requests!

Yang model

- Yang model is provided based on RPC stateless
 - Aligned with the TE-Tunnel YANG model to ensure consistency
 - Working in close relationship with TE-Tunnel YANG model authors to address common open issues
- Stateful Path computation can be achieved with pre-computed tunnels as defined in TE-Tunnel YANG model

GitHub Support

- GitHub Repository
 - <https://github.com/rvilalta/ietf-te-path-computation>
- GitHub support used for
 - Developing and tracking YANG model for stateless RPC
 - Tracking Open Issues, discussions and resolutions

Open Issues

- How to reduce the number of path computation requests in networks with many domains
 - In principle: use the detailed connectivity matrix information together with path computation requests
 - Is this a standardization issue or only an implementation issue, provided that IETF standard provides all the required tools?
- Residual BW [#30]
 - New metric for the minimum unreserved bandwidth over all the links traversed by the computed path
- Topology-id in path constraints [#27]
 - The avoidTopology allows constraining which topologies shall be avoided by path computation, includeTopology list to constraints which topologies shall be considered by path computation.
 - To be agreed the need
- Missing local protection [#24]
 - Use of L flag in the SESSION/ATTRIBUTE object.

Open Issues (2)

- Support of Delay metric [#21]
 - Is it needed to align with draft-ietf-pce-pcep-service-aware-13?
 - Currently te-tunnel uses metric-type TE (min delay) + cost-limit = X msec (if smaller than x msec)
- Multiple metric for path computation [#20]
 - In RFC 5440 a path computation request can include an arbitrary number of METRIC objects.
 - Currently TE-tunnel does not support this
- Optional or mandatory constrains [#19]
 - In PCEP it is possible to specify if a constraint is mandatory, optional, if the path computation must fail if the constraint is not met or to relax the constraint.
 - It should be possible to have the same behavior for path computation RPC and tunnel setup.

Open Issues (3)

- Clarification/discussions with TE-Tunnel authors
 - Representation of IRO and XRO using the explicit route object in TE-tunnel [#29]
 - Usage of Affinities mask [#26]
 - Tiebraker associated behavior [#22]
 - How to know the layer of the tunnel to be setup and/or the path to be computed [#18]
 - Source&Destination reference: not clear the usage of double syntax (ip-address and tp-id) in te-tunnel [#15]
- Capacity units [#28]
 - to be aligned with TE-topology
- How to use a subset of tunnel-params_config grouping for a Path Computation RPC

Next Steps

- Resolve current open issues
 - Continue cooperation with TE Tunnel model authors
- Path computation for multiple LSP
 - Yang has to be updated to support this
- Describe solution for the case where RPC response takes too long time
- Seeking comments and feedbacks from interested WGs to improve document
- Yang solution integration into TE-tunnel draft or in this draft?
- Ready to become WG document?