TVR Problem Statement

https://datatracker.ietf.org/doc/draft-taylor-tvr-prb-stmt/

Rick Taylor, *Ori Industries*
Adrian Farrel, *Old Dog*

IETF-115

The Problem *(distilled)*

In a network, links and nodes may disappear and appear with some degree of predictability.

That is, the reachable neighbors in a network may change in a **predictable** or **scheduled** way.

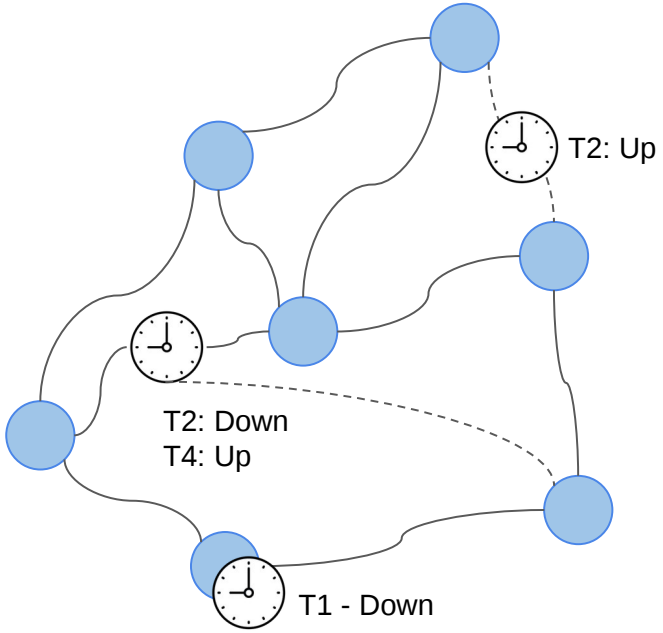

The Problem

- Routing protocols already handle convergence in the face of adjacency **failures**.
    - In general, convergence is fast enough, and fast reroute mechanisms exist to handle transient routing during convergence.
- However, routing does not currently handle the potential connectivity represented by nodes and links that are **scheduled** to appear in the future.
- Nor do they handle planning around predicted outages.

The Problem *(continued…)*

The problems in this space include:

- How does a node or link that is not up have its future presence advertised?
- How does the fact that a node or link is expected to come up (i.e., the future reachability of a neighbor) get advertised along with the predicted up time and up window?
- How does the fact that a node or link is expected to go down (i.e., the future removal of reachability to a neighbor) get advertised along with the predicted down time and down window?

The same two questions apply to advertising known variations in metrics which might not be limited to existing metrics, but might also include local, node-based resources.

Full Text

Routing protocols already handle convergence in the face of adjacency failures. In general, convergence is fast enough, and fast reroute mechanisms exist to handle transient routing during convergence.

However, routing does not currently handle the potential connectivity represented by nodes and links that may appear in the future. Nor do they handle planning around predicted outages.

It is desirable to be able to plan and predict reachability across a network where reachable neighbors will only be present at and for specific known time intervals. This enables less-interrupted and deterministic service delivery without the need to actively re-route at the packet level.

Furthermore, it is not just availability that may vary over time. Any metric associated with a node or link may vary over time. The existing routing protocols can advertise the current state of the metrics for use by the routing system, but it may be beneficial to know about predicted changes to those metrics in order to plan connectivity and service delivery.

The problems in this space include:

- How does a node or link that is not up have its future presence advertised?
- How does the fact that a node or link is expected to come up (i.e., the future reachability of a neighbor) get advertised along with the predicted up time and up window?
- How does the fact that a node or link is expected to go down (i.e., the future removal of reachability to a neighbor) get advertised along with the predicted down time and down window?

The same two questions apply to advertising known variations in metrics which might not be limited to existing metrics, but might also include local, node-based resources.

Questions, Opinions,
and Criticism?