Ring Configuration and Provisioning

Anoop Ghanwani

(anoop@lanterncom.com)

Constantinos Bassias

(cbassias@lanterncom.com)

Motivation

- There are several configurable parameters that have ring-wide significance
  - How do we ensure these parameters in sync across all nodes?
  - What happens if they are not in sync or if there is a misconfiguration?

- Correct operation cannot be ensured using the methods for bandwidth provisioning currently specified
  - The MAC Data Path Clause supports only a single shaper per ringlet for each of the provisioned classes
  - Yet the Topology Clause suggests that bandwidth can be spatially provisioned

- This presentation describes these issues and proposes solutions to them
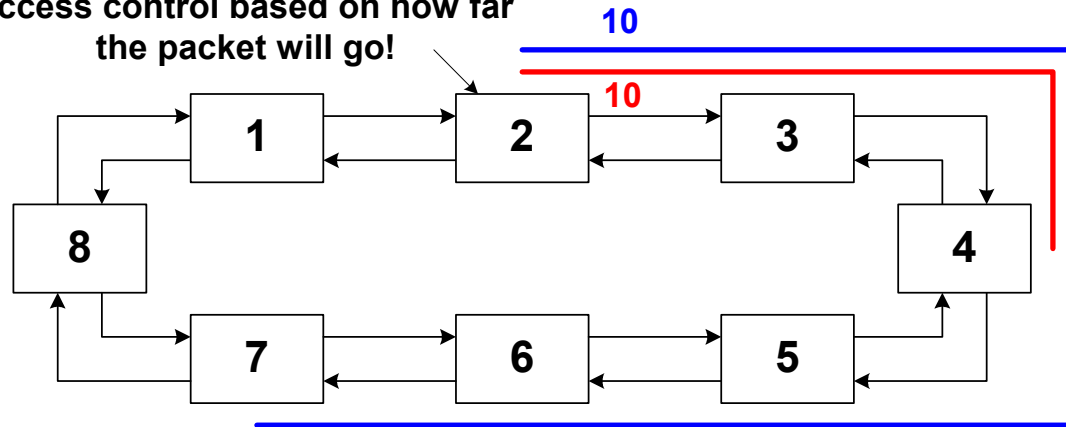
Parameters with Ring-wide Significance

- Parameters used by the fairness algorithm
  - Aging interval must be the same at all nodes
  - Other parameters such as AGECOEF, RAMPCOEF, etc. need to be in sync for correct operation

- Provisioned bandwidth
  - Total provisioned bandwidth must not exceed the ringlet capacity

- Wrap capability
  - Set only if configured for wrapping
  - Wrapping is used only when all stations on the ring support it and are configured for it

- Maximum frame size
  - Use the minimum value of MTU among all nodes

Checking the Ring Configuration

- Advertise parameters using TLVs in the topology messages
  - Most of these are defined today, but others need to be defined
- Define the behavior for the following scenarios
  - Stations are not in sync
  - A station is using a value that is out of the allowable range
- Actions when out of sync
  - Use the smallest/largest value on the ring as appropriate for correct operation
- Actions on misconfiguration
  - Continue to operate the ring, but flag an alarm
  - Stop sourcing any traffic that would be affected by the misconfigured parameter, flag an alarm, but continue to operate

Spatial Bandwidth Provisioning

For ClassA, Node 2 sets shaperA for 20. It does not perform access control based on how far the packet will go!
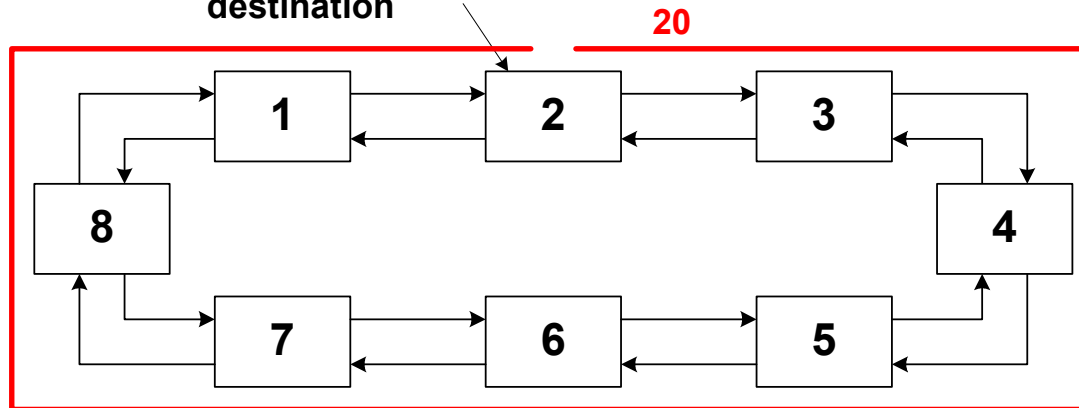

- Provisioned bandwidth of 10 units is required from 2—4 and 2—7
- Provisioning would take place as follows:
  - ShaperA at Node 2 is set to limit access to 20 units
  - 20 units are provisioned on each of the segments from 2—4
  - 10 units are provisioned on each of the segments from 4—7
  - Nothing is provisioned on the segments from 7—2
- ShaperA is not spatially aware
  - A ClassA frame sourced by Node 2 destined for Node 8 will be admitted as long as there are credits in the shaper
  - This will adversely affect the performance of other nodes' traffic

Bandwidth Provisioning

- Because access control for ClassA and ClassB-CIR is not spatially aware, bandwidth provisioning also cannot be spatially aware

- In other words, each node *must* uniformly provision bandwidth used by ClassA and ClassB-CIR all the way around a ringlet

Uniform Bandwidth Provisioning

For Class A, Node 2 does limits total traffic to 20 units regardless of destination

- Provisioned bandwidth of 10 units is required from 2—4 and 2—7
- Provisioning would take place as follows:
  - ShaperA at Node 2 is set to limit access to 20 units
  - 20 units are provisioned on each of the segments from 2—2
- Unused provisioned bandwidth can be reclaimed for ClassA1 and ClassB

Ring-wide Consistency Check for Provisioned Bandwidth

- Advertise the provisioned bandwidths for a given node
  - Provide the breakdown by Class – A0, A1, CIR B
  - A0 is required for the downstream shaper (ShaperD)

- Every node performs a consistency check to ensure that the total provisioned bandwidth does not exceed the capacity of the ringlet
  - If the check fails, an alarm is generated

Summary

- P802.17 should use the topology messages to detect misconfiguration of parameters with ring-wide significance

- P802.17 should also define the behavior for scenarios where the parameters are out-of-sync or where there are misconfigurations

- For bandwidth provisioning, we must only allow for uniform per-node provisioning around a ringlet

- Spatial bandwidth provisioning is possible if we have per-destination node shaping for ClassA traffic as described in Annex H, but the MAC Data Path clause needs to be clear on this issue