Generation of Protection Messages

Anoop Ghanwani

(anoop@lanterncom.com)

IEEE 802.17 Plenary Meeting

Vancouver, BC, Canada

July 2002

Introduction

- The current specification for the generation of protection messages states that [P802.17/D0.3, Page 112, Line 40]:

  "Protection messages are sent when the local protection state changes, and are repeated based on an exponential back-off timer with start of 1 ms and back-off factor of 8, capped at 1024 ms."

- This presentation describes some problems with this and possible alternatives

The Pitfalls of Exponential Back-off

- It is not reliable
- It has a high overhead and processing requirement at the time immediately following the change in protection status
  - The system has to be designed to handle this level of processing anyway
- If the initial messages are lost for any reason, it be may be a while before everyone knows about the fault
- Introduces unnecessary complexity in the protocol

The only advantage of backing-off to a slower rate is a slightly lower bandwidth overhead

Alternative Solutions

- Combine the protection information with the Type B fairness messages
  - The standard requires that Type B messages be sent every 10 advertisement intervals (approximately every 1 msec)
  - These messages are broadcast and reach all nodes
  - Protection information can easily be piggy-backed on them

OR

- Make the generation of protection messages periodic with a configurable interval from 1-10 msec
  - Primary complaint is the bandwidth overhead

A Word About Bandwidth Overhead

- On a 100 node network running at 10 Gbps, each node sending a 64-byte broadcast message every millisecond requires a total of about 0.5% of the bandwidth

- If we really care about that much bandwidth, we should rather direct our efforts on bettering the fairness algorithm!
  - The gains will be much higher than can be achieved by suppressing protection information

Conclusions

- The generation of protection must be specified so that it is simple and achieves the objectives
  - Timely notification of faults
  - Tolerant to message loss

- Piggy-backing protection information on Type B fairness messages is ideal
  - Type B messages are mandatory and are broadcast very frequently
  - This approach has the least bandwidth overhead

- Alternatively, redefine the protection message generation so that it is periodic with a configurable interval from 1-10 msec