

# ***1:1 oversubscription***

**Kevin Daines**

**World Wide Packets**

**20 September 2004**

# *1:1 oversubscription defined*

- Take a simple device, shown below, equipped with two Gigabit Ethernet ports...



- ...and configure one port as untagged and the other port as tagged and...
- ...Voilà!
  - You have created (the possibility of) *1:1 oversubscription*

# ***Constriction***

- **Using our previous example, if one port is untagged and the other port is tagged, congestion in the form of *constriction*\* can occur**

**\*term borrowed from Hugh Barrass/July 04 CMSG) occurs**

# Tagged constriction

- (1) untagged Gig  $\Rightarrow$  (1) tagged Gig
  - Minimum sized frames
    - 1,488,095  $\Rightarrow$  1,420,454 (adds 4B)
  - Difference of 67,641 min frames
- Solutions
  - Count on sustained rate being less than 95.5%; buffer some line rate bursts; discard excess frames (**widely used**)
  - Source is throttled back to “tagged” frame rate using PAUSE or unspecified egress rate control mechanisms

# Add 2<sup>nd</sup> tag constrictio

- (1) tagged Gig  $\Rightarrow$  (1) double tagged Gig
  - Minimum sized fps:
    - 1,420,454  $\Rightarrow$  1,358,695 (adds 4B)
  - Difference of 61,759 min tagged frames
- Solutions
  - Count on sustained rate being less than 95.7%; buffer some line rate bursts; discard excess frames (**widely used**)
  - Source is throttled back to “double tagged” frame rate using PAUSE or unspecified egress rate control mechanisms

# *Add two tags constrict*

- (1) untagged Gig  $\Rightarrow$  (1) double tagged Gig
  - Minimum sized fps:
    - 1,488,095  $\Rightarrow$  1,358,695 (adds 8B)
  - Difference of **129,400** min frames
  
- Solutions
  - Count on sustained rate being less than **91.3%**; buffer some line rate bursts; discard excess frames (**widely used**)
  - Source is throttled back to “double tagged” frame rate using **PAUSE** or unspecified egress rate control mechanisms

# 10G WAN constriction

- (1) 10 Gig LAN  $\Rightarrow$  (1) 10 Gig WAN
  - Minimum sized fps:
    - 14,880,952  $\Rightarrow$  14,262,857 (pure rate mismatch)
  - Difference of 618,095 min frames
- Solutions
  - Count on sustained rate being less than 95.8%; buffer some line rate bursts; discard excess frames (**widely used**)
  - Source is throttled back to “WAN” frame rate using PAUSE or unspecified egress rate control mechanisms

# MACSec constriction

- (1) Gig  $\Rightarrow$  (1) MACSec Gig
  - Minimum sized fps:
    - 1,488,095  $\Rightarrow$  822,368 (assumes 4B SecTag, 64B ICV)
  - Difference of 665,727 min frames
- Solutions
  - Count on sustained rate being less than 55.3%; buffer some line rate bursts; discard excess frames (*MACSec is being defined in P802.1AE*)
  - Source is throttled back to “MACSec” frame rate using PAUSE or unspecified egress rate control mechanisms



# *PWE3 constriction*

- (1) Gig  $\Rightarrow$  (1) PWE3 Gig
  - Minimum sized fps:
    - 1,488,095  $\Rightarrow$  1,096,491 (assumes 30B overhead)
  - Difference of **391,604** min frames
  
- Solutions
  - Count on sustained rate being less than **73.7%**; buffer some line rate bursts; discard excess frames (**widely used**)
  - Source is throttled back to “PWE3” frame rate using PAUSE or unspecified egress rate control mechanisms

# *Clock tolerance constrictio*

- (1) +100ppm Gig  $\Rightarrow$  (1) -100ppm Gig
  - Minimum sized fps:
    - 1,488,244  $\Rightarrow$  1,487,946 (pure rate mismatch)
  - Difference of **298** min frames
  
- Solutions
  - Count on sustained rate being less than **99.98%**; buffer some line rate bursts; discard excess frames (**widely used**)
  - Source is throttled back to “-100ppm” frame rate using PAUSE or unspecified egress rate control mechanisms

# 802.3 FESG constriction

- (1) legacy Gig  $\Rightarrow$  (1) FESG Gig
  - Minimum sized fps:
    - 1,488,095  $\Rightarrow$  203,583 (*assumes 530B additional header*)
  - Difference of 1,284,512 min frames
- Solutions
  - Count on sustained rate being less than 13.7%; buffer some line rate bursts; discard excess frames (*Study Group phase*)
  - Source is throttled back to “FESG” frame rate using PAUSE or unspecified egress rate control mechanisms

# Summary

- **MACSec and MPLS encapsulation make 802.1Q VLAN tagging rate mismatch seem trivial**
  - **FESG could add 300-500 octets in overhead**
- **Critical question:**
  - **For 1:1 oversubscription/ constriction situations, is today's widely used solution (count on less than line rate traffic, buffer some bursts, discard excess) sufficient?**
  - **Or, should 802.3 CMSG consider specifying an egress rate control mechanism**