

**IEEE 802.3
Congestion Management
Study Group July Plenary
Report**

**Portland, OR
12-15 July, 2004**



Bottom Line

- We have some objectives
- We're not done studying the issue
- We don't have 5 Criteria and PAR yet
- We want an extension of the study group



Presentations

■ May

- 7 presentations
- 9 presenters
- 6 companies

■ July

- 6 presentations
- 9 presenters
- 5 companies




Presentation Topics

- Definitions
- Problem Spaces
- Problem Statements
- Layer Stack Discussions
- Market Surveys
- Other standards work
- Preemption
- Simulations, Modeling
- Education
- Requirements

Congestion

Congestion can only happen at a confluence or a constriction

Multiple pipes ...



Fat pipe ...



**Congestion will not necessarily occur, depending on circumstances, but ...
... congestion can never happen in any other case.**

Slide from barrass_1_0704



General Congestion types

- Transitory congestion

- Traffic that can be smoothed over time, without packet drop, because average BW demand is less than capacity and peak demand can be buffered

- Oversubscription

- Traffic that can't be smoothed over time, can result in not being admitted into the network (e.g. admission control), frame drop (e.g. buffer overflow, RED), or back-up into source buffers



Study Group Responses

- Transitory congestion
 - Differentiated Service (priorities)
 - Can provide preferential treatment to higher priority traffic
 - Discussion shows that some in the group believe that .1p addresses this. (Some standard specification of “draining” mechanisms may be helpful.) The MAC Service interface though does not support priority queueing.
 - Many NICs do not implement priorities because they are specified only for bridging.
- Oversubscription
 - Rate limiting can help reduce packet drops in this scenario
 - Can improve latency and latency jitter (not building up buffers, not requiring TCP to resort to retransmit etc.)

Slide from wadekar_2_0704



Problem Space

- Dropped packets within an oversubscribed Ethernet bridged network produce significant latency problems for messaging and storage traffic that cannot be readily resolved at layer 3 or 4.
- PAUSE was originally developed as a means of backpressuring link traffic to an input buffered switch. The use of PAUSE for congestion management has a detrimental affect on latency and it results in simply moving the congestion point away from the actual location.
- A method of limiting network traffic on the oversubscribed link by allocating, rationing, or limiting bandwidth from sources or upstream links is needed.

Slide from wadekar_2_0704



Narrowed Problem Statement

- 802.3 MAC client needs to be able to limit traffic being received from its peer to avoid congestion due to oversubscription

Slide from wadekar_2_0704



Supporting Objectives

Everyone: 17/0/0, 802.3: 8/0/0

- Focus solution to a single link only (hop-to-hop/end-to-end not specified)
- Specify a mechanism to limit the rate of transmitted data using a “pacing” algorithm (not a burst duty cycle)
- Specify the granularity of the rate limiter
- Specify a new MAC Control Opcode and parameter set to support exchange of rate control information
- Do not specify how the MAC Client generates these MA_CONTROL.requests nor how it responds to the reception of MA_CONTROL.indications
- Specify the response to the new MAC Control opcode’s parameter set
- Work with other 802.3 activities on the “long standing inconsistency” between MA_DATA.requests and transmit_frame function call

Slide from wadekar_2_0704



Future Plans

- Meet somewhere in North America during an autumn month to:
 - refine the existing and consider additional objectives
 - develop responses to the 5 criteria
 - generate Title, Scope, and Purpose of the PAR
- Pre-circulate these documents to 802.3
- Prepare a tutorial for the November Plenary meeting
- Request 802.3 to approve the Congestion Management project and recommend same to the EC/NESCOM (P802.3as?)



Question

- Do you feel the study group should extend until November?
 - Y: 16
 - N: 0
 - 17 in attendance



Motion (Extend)

- Move that 802.3 extend the Congestion Management Study Group
- Procedural (50%)
- Moved: Ben Brown
- Second: Brad Booth
- 802.3 voters 46/2/10
- Motion Passes



Thank you!