





IEEE 802.3 Congestion Management Study Group

Ottawa, Ontario, CA
27-28 Sept, 2004



Agenda

- Welcome and Introductions
- Appoint/Volunteer Recording Secretary
- Approve meeting minutes
- Goals for this Meeting
- Reflector and Web
- Ground Rules
- IEEE
 - Structure
 - Bylaws and Rules
 - Call for Patents
 - IEEE Standards Process
- Presentations
- Discussions
 - Objectives
 - 5 Criteria
 - PAR (Title, Scope & Purpose)
- Future Meetings
- Motion Madness



Goals for this Meeting

- Hear presentations concerning:
 - Scope of a Congestion Management Project
 - Justification in terms of the 5 Criteria
 - Goals and Objectives for the Project

- Build consensus on:
 - Congestion Management Objectives
 - Responses to the 5 Criteria
 - Project Authorization Request (PAR): Title, Scope, and Purpose



Reflector and Web

- To subscribe to the Congestion Management Study Group reflector send an email to:

listserv@ieee.org

with the following in the body of the message:

*[subscribe stds-802-3-cm <your first name>
<your last name>](#)*

- Congestion Management Study Group web page URL:


http://www.ieee802.org/3/cm_study/



Ground Rules

- 802.3 Rules apply
 - Foundation based upon Robert's Rules of Order
- Anyone in the room may speak
- Anyone in the room may vote
- **RESPECT**... give it, get it
- NO product pitches
- NO corporate pitches
- NO prices!!!
 - This includes costs, ASPs, etc. no matter what the currency
- NO restrictive notices

IEEE Structure





Bylaws and Rules

- Bylaws of the IEEE Standards Association (IEEE-SA):
<http://standards.ieee.org/sa/sa-bylaws.pdf>
- Bylaws of the IEEE-SA Standards Board:
<http://standards.ieee.org/guides/bylaws/sb-bylaws.pdf>
- IEEE LAN/MAN Standards Committee (LMSC)
Operating Rules:
<http://www.ieee802.org/rules.pdf>
- IEEE 802.3 Working Group Operating Rules:
<http://www.ieee802.org/3/rules/>



IEEE-SA Standards Board Bylaws on Patents in Standards

6. Patents

IEEE standards may include the known use of patent(s), including patent applications, provided the IEEE receives assurance from the patent holder or applicant with respect to patents essential for compliance with both mandatory and optional portions of the standard. This assurance shall be provided without coercion and prior to approval of the standard (or reaffirmation when a patent becomes known after initial approval of the standard). This assurance shall be a letter that is in the form of either

- a) A general disclaimer to the effect that the patentee will not enforce any of its present or future patent(s) whose use would be required to implement the proposed IEEE standard against any person or entity using the patent(s) to comply with the standard or
- b) A statement that a license will be made available without compensation or under reasonable rates, with reasonable terms and conditions that are demonstrably free of any unfair discrimination

This assurance shall apply, at a minimum, from the date of the standard's approval to the date of the standard's withdrawal and is irrevocable during that period.

Approved by IEEE-SA Standards Board – December 2002



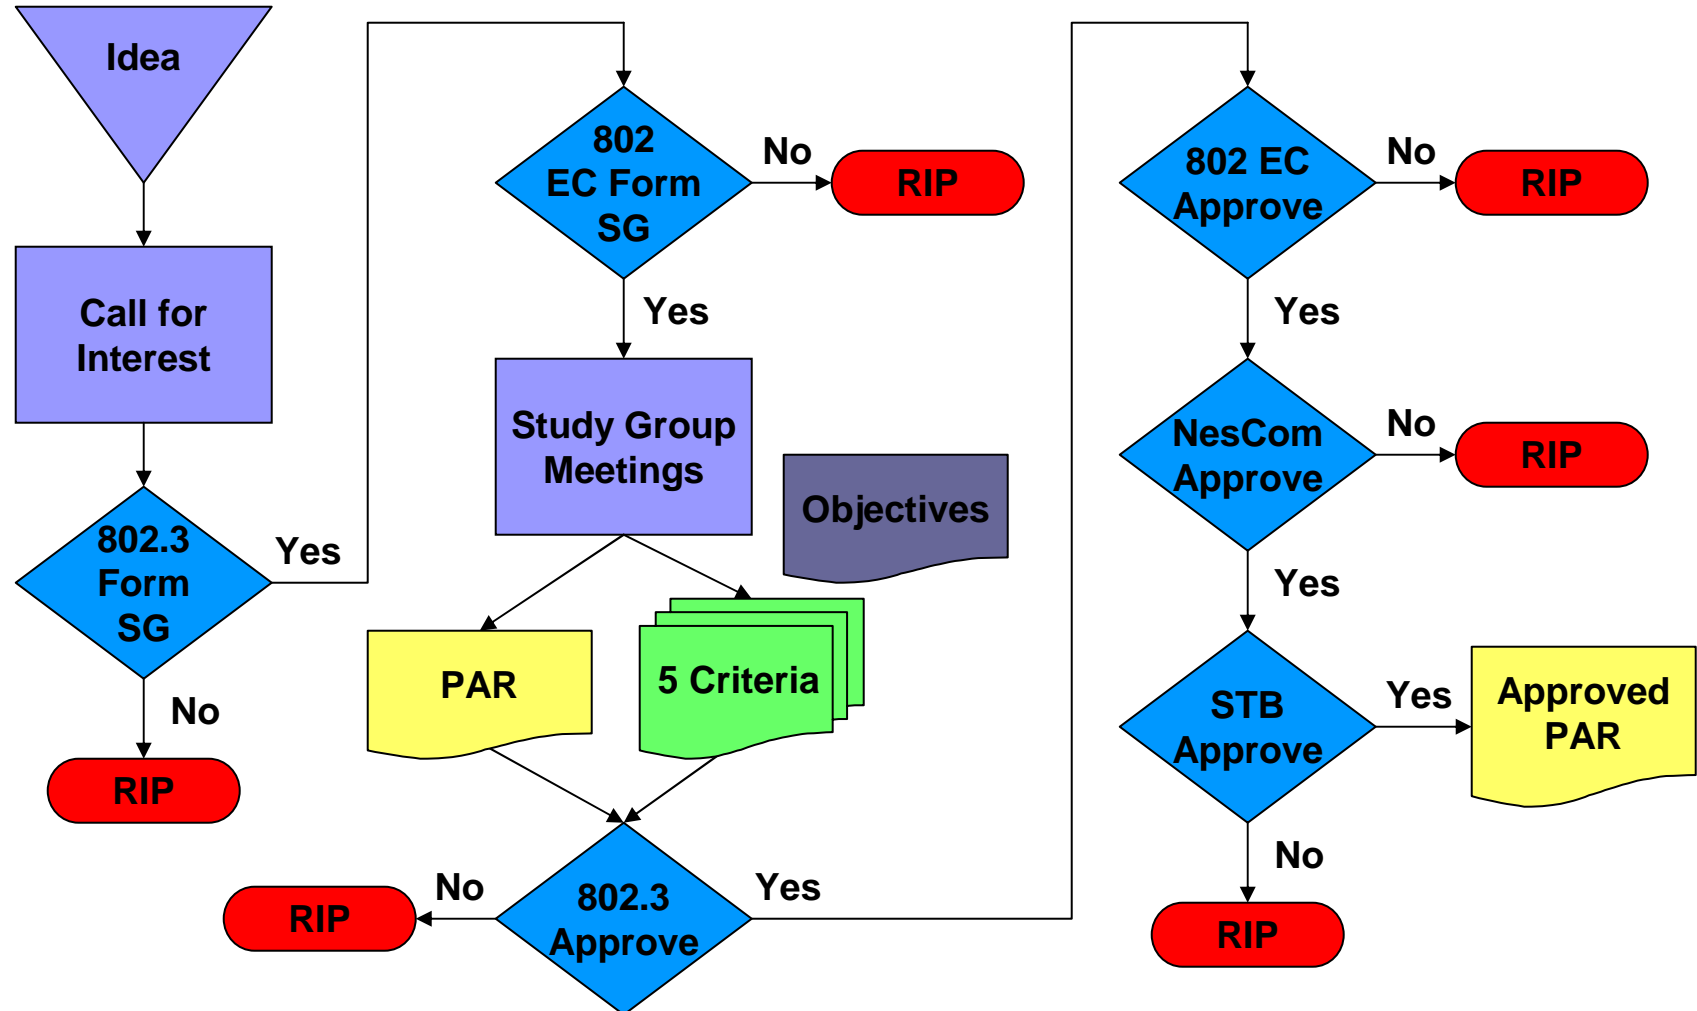
Inappropriate Topics for IEEE SG Meetings

- Don't discuss licensing terms or conditions
- Don't discuss product pricing, territorial restrictions or market share
- Don't discuss ongoing litigation or threatened litigation
- Don't be silent if inappropriate topics are discussed... do formally object.

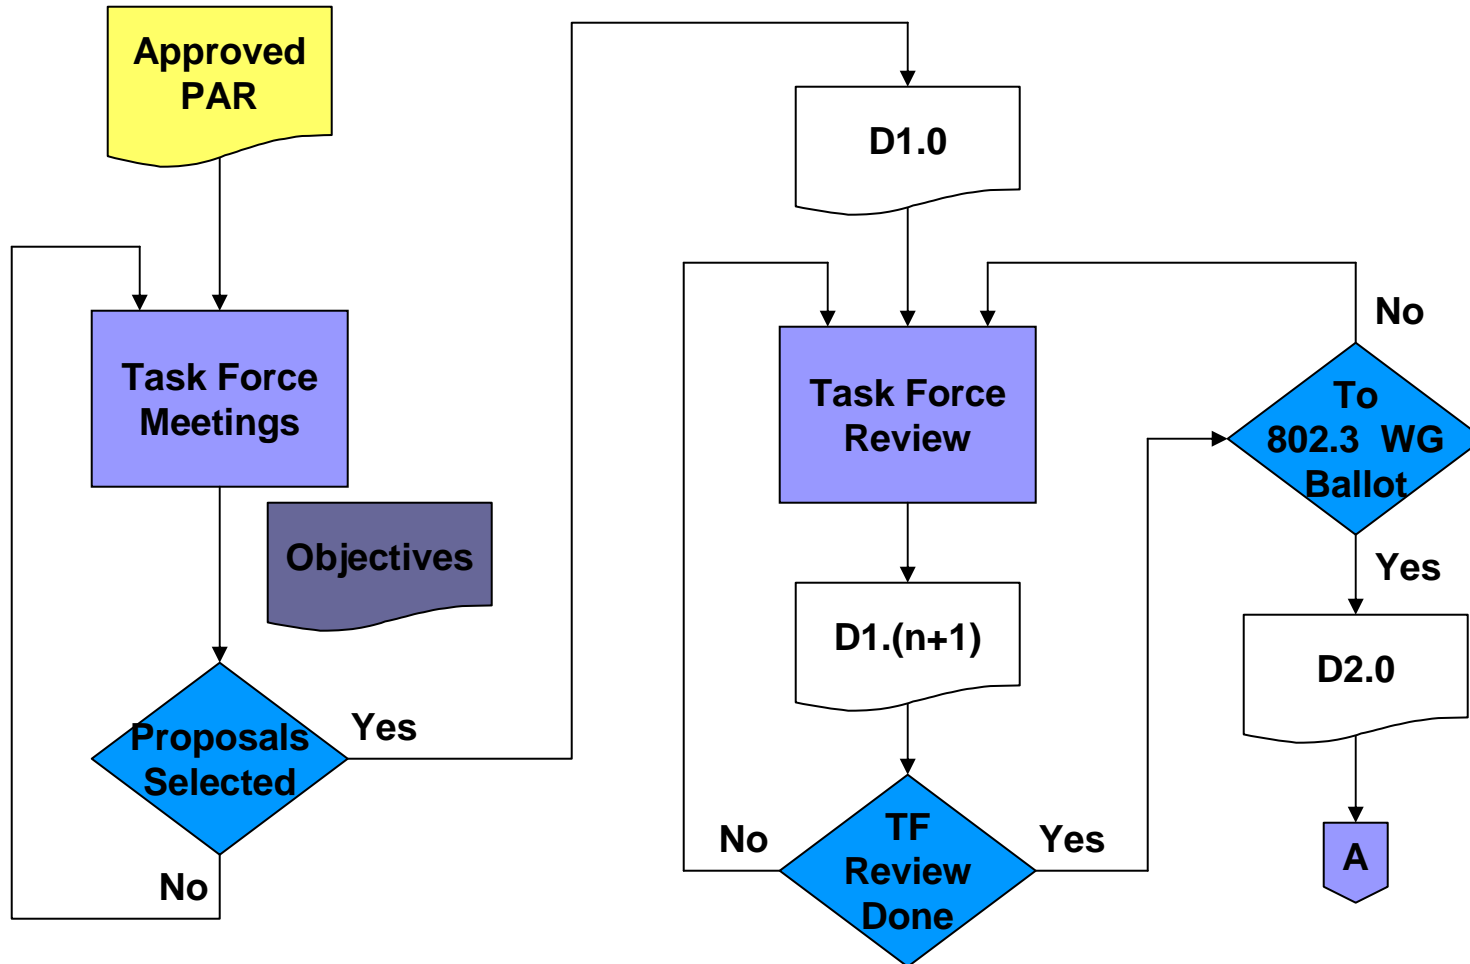
If you have questions, contact the IEEE Patent Committee Administrator at patcom@ieee.org

Approved by IEEE-SA Standards Board – December 2002


IEEE Standards Process



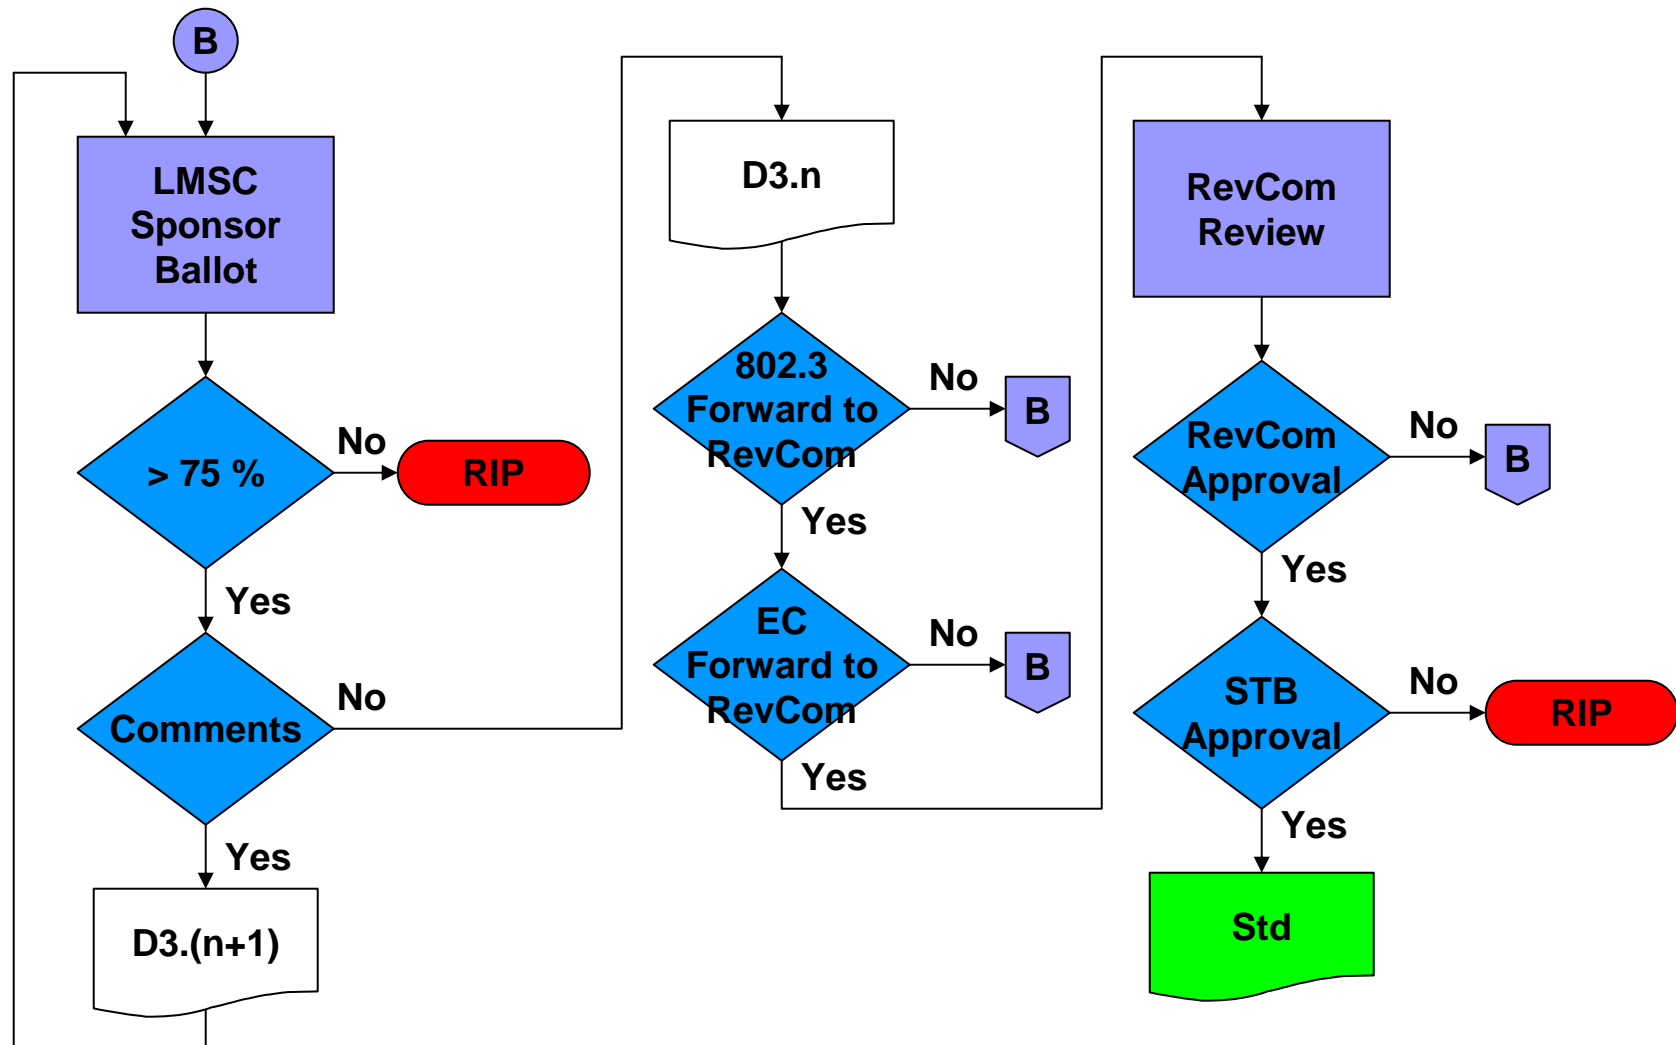
IEEE Standards Process (cont.)



IEEE Standards Process (cont.)



IEEE Standards Process (cont.)





Study Group

- Function is to draft a PAR and 5 Criteria
- Gain approval at WG 802.3, 802 EC, IEEE NesCom and IEEE Stds. Board
- SG only exists for 6 months
 - Extensions can be requested... voted on by 802.3, ratified by EC
- Development of Objectives helps set the goals for the Task Force
- Developing consensus
 - Education helps build consensus
 - Consensus (> 75%) required to move forward



PAR

- Title

- What are we calling this

- Scope

- The focus: Congestion Management over Ethernet Links

- Purpose

- Why do we want to do this



5 Criteria

- **Broad Market Potential**
 - ☐ Broad set(s) of applications
 - ☐ Multiple vendors, multiple users
 - ☐ Balanced cost (LAN vs. attached stations)
- **Compatibility with IEEE Std. 802.3**
 - ☐ Conformance with CSMA/CD MAC, PLS
 - ☐ Conformance with 802.2
 - ☐ Conformance with 802 Functional Requirements
- **Distinct Identity**
 - ☐ Substantially different from other 802.3 specifications
 - ☐ One unique solution for problem
 - ☐ Easy for document reader to select relevant spec



5 Criteria (cont.)

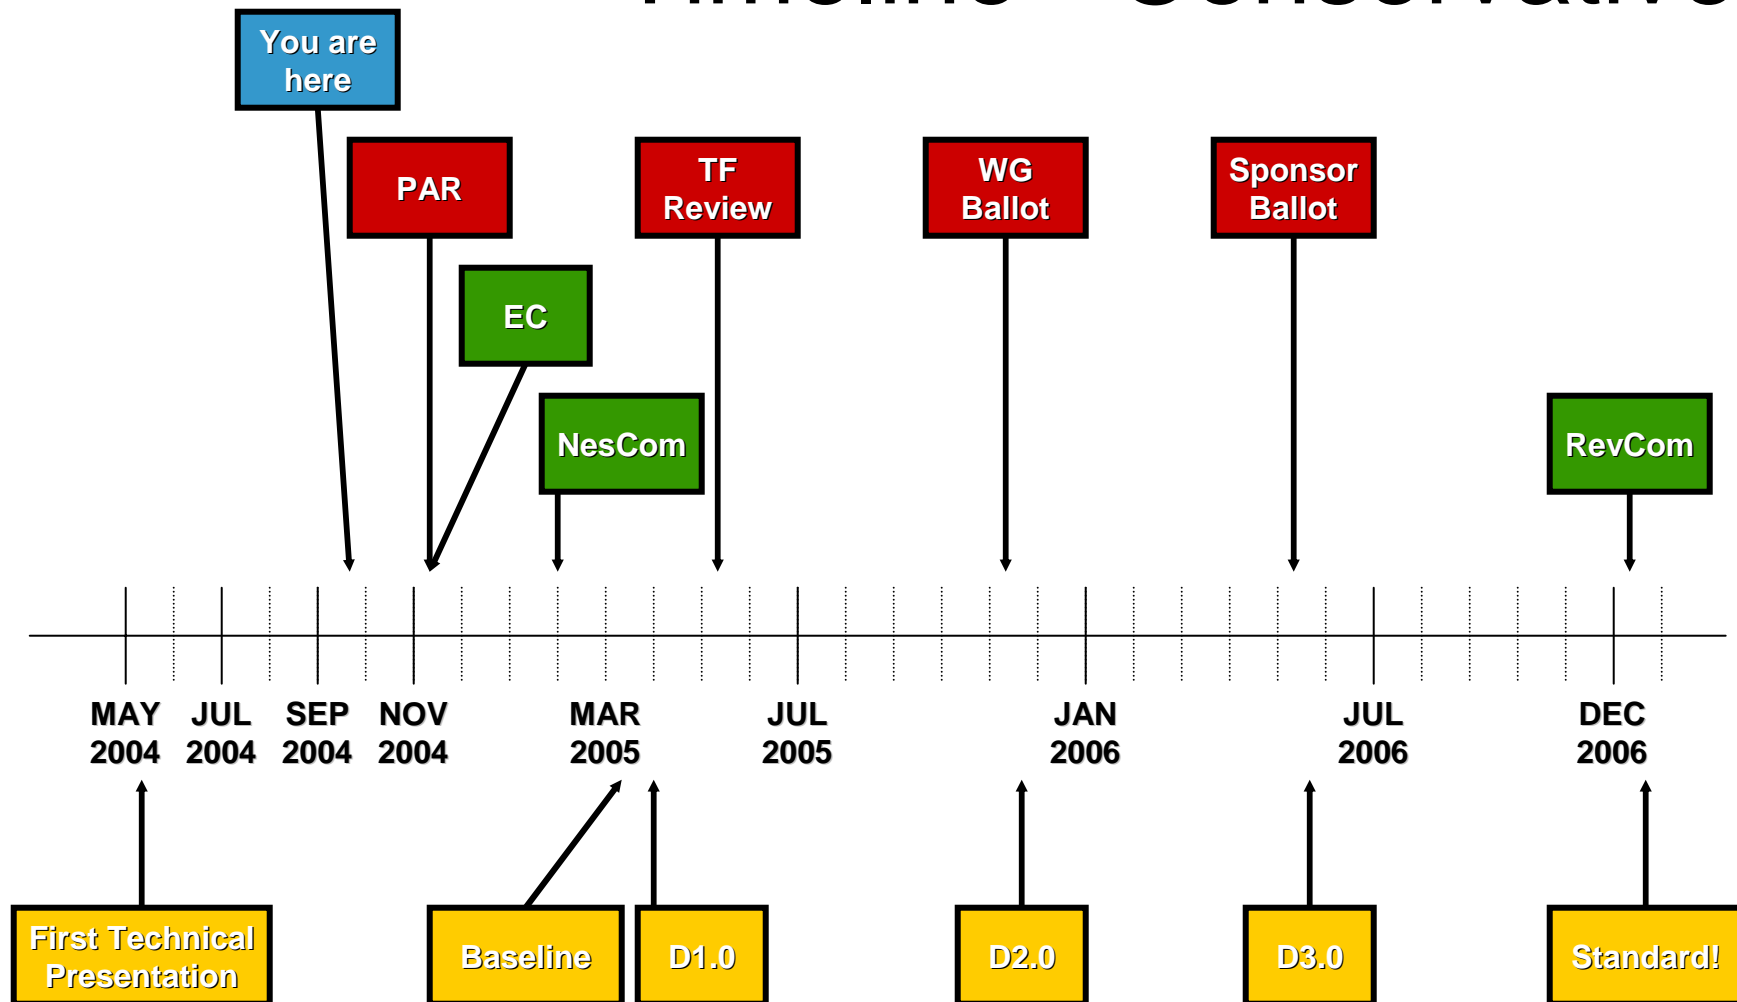
■ Technical Feasibility

- ☐ Demonstrated system feasibility
- ☐ Proven technology, reasonable testing
- ☐ Confidence in reliability


■ Economic Feasibility

- ☐ Cost factors known, reliable data
- ☐ Reasonable cost for performance
- ☐ Total installation costs considered

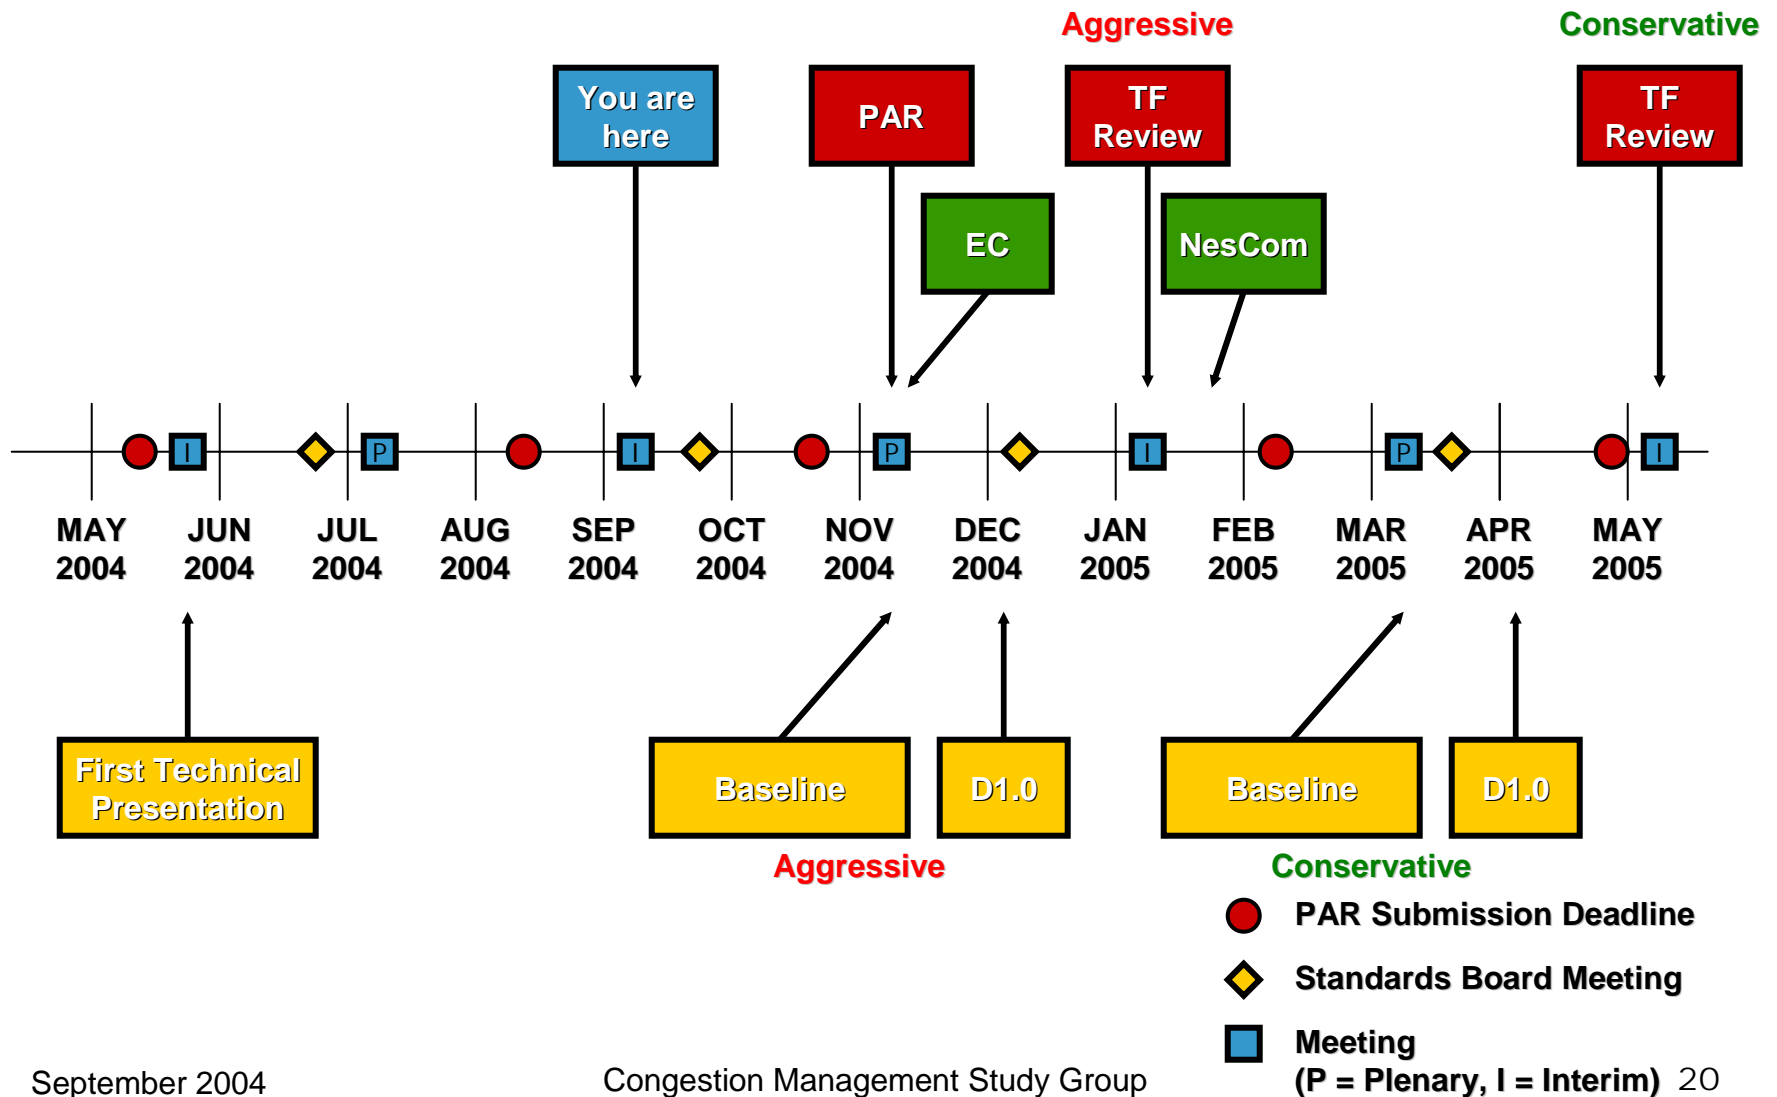
Possible Congestion Management Timeline - Conservative



Possible Congestion Management Timeline - Aggressive



Timeline Detail





Objectives (adopted July, '04)

- 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

Presentations

Monday September 27, 2004				
Presenter	Topic	Length	Start	Finish
	Welcome and Introductions	0:10	8:30 AM	8:40 AM
Brown, Ben	Agenda	0:30	8:40 AM	9:10 AM
Daines, Kevin	1:1 Oversubscription	0:20	9:10 AM	9:30 AM
Booth, Brad	CM Value Proposition: Aligning the objectives	0:30	9:30 AM	10:00 AM
Morning Break		0:20	10:00 AM	10:20 AM
Wadekar, Manoj	Rate Control in Short Range 802.3 interconnects	1:00	10:20 AM	11:20 AM
Thaler, Pat	Congestion Spreading	0:20	11:20 AM	11:40 AM
Lunch		1:20	11:40 AM	1:00 PM
Barass, Hugh	Joint 802.3 & 802.1 work	0:45	1:00 PM	1:45 PM
Group	Discussion of joint 802.3 & 802.1 work	0:30	1:45 PM	2:15 PM
Afternoon Break		0:20	2:15 PM	2:35 PM
Group	Discussion of Objectives, Critters, PAR	2:25	2:35 PM	5:00 PM
Tuesday September 28, 2004				
Presenter	Topic	Length	Start	Finish
Group	Welcome	0:15	8:30 AM	8:45 AM
Group	Discussion of Objectives, Critters, PAR	1:30	8:45 AM	10:15 AM
Morning Break		0:15	10:15 AM	10:30 AM
Group	Discussion of Objectives, Critters, PAR	1:30	10:30 AM	12:00 PM
Lunch		1:00	12:00 PM	1:00 PM
Group	Discussion of Objectives, Critters, PAR	1:45	1:00 PM	2:45 PM
Afternoon Break		0:15	2:45 PM	3:00 PM
Group	Discussion of Tutorial for November Plenary	2:00	3:00 PM	5:00 PM



Questions

- If 802.1 approved a project for Congestion Management, how many in this room would work on it? 18 of 25
- How many would sign on to Hugh's presentation, to be presented to 802.1 on Monday, Oct 4?



Proposed new objectives from Brad

- Specify a mechanism to support the ~~exchange~~ **communication** of congestion ~~control~~ information
- Specify a mechanism to limit the rate of transmitted data **on an Ethernet link**
 - ☐ ~~Support point to point links only~~
 - ☐ ~~Support full duplex operation only~~
- Preserve the MAC/PLS service interfaces
- ~~Preserve the 802.3/Ethernet frame format at the MAC Client service interface~~

Red text is inserted Strikethrough text is deleted



More new objectives from Pat

- Minimize throughput reduction in non-congested flows



Objectives (adopted September, '04, replaces those adopted in July)

- Specify a mechanism to support the communication of congestion information
- Specify a mechanism to limit the rate of transmitted data on an Ethernet link
- Preserve the MAC/PLS service interfaces
- Minimize throughput reduction in non-congested flows



Broad Market Potential

Broad set(s) of applications

Multiple vendors, multiple users

Balanced cost (LAN vs. attached stations)

Red text is deleted

Underscore text is inserted

- Ethernet **links** networks are being used in an increasing number of application spaces (clustering, backplanes, storage, data centers, etc.) that are sensitive to frame **discard and latency** delay, delay variation and loss. **Frame discard occurs when the MAC Client is oversubscribed in comparison to the capabilities of today's Ethernet full duplex links. Available solutions trade off latency for reduced frame discard. Study Group** presentations have shown that Ethernet networks can experience higher throughput, lower **latency** delay, and **fewer frame discards** lower frame loss by performing congestion management **within 802.3**. **This** which will improve Ethernet in its growing number of applications spaces.
- During the discussion of the WG 802.3 motion to initiate this study group, 23 people from 16 companies indicated that they plan to participate in the standardization effort for congestion management. This level of commitment indicates that a standard will be developed by a large group of vendors and users. During the study group meetings, there have been up to **22** 30 people from at least 16 companies in attendance.
- A standard to support congestion management will **maintain** respect the balance of cost between LAN and attached stations.



Compatibility with IEEE Std 802.3

Conformance with CSMA/CD MAC, PLS

Conformance with 802.2

Conformance with 802

- The proposed standard will conform to the 802.3 MAC, and therefore will be consistent with 802.1d, 802.1Q, and relevant portions of 802.1f.
- As was the case in previous 802.3 standards, additional MAC Control sublayer functionality and MAC Control frame opcodes may be defined.
- The proposed standard will conform to the 802.3 MAC Client Interface, which supports 802.2 LLC.
- The proposed standard will conform to the 802.1 Architecture, Management and Internetworking.
- The proposed standard will define a set of systems management objects, which are compatible with OSI and SNMP system management standards.
- The proposed standard will conform to the requirements of IEEE Std 802-2001.



Distinct Identity

Substantially different from other 802 & 802.3 specs

One unique solution for problem

Easy for document reader to select relevant spec

- The current 802.3 standard specifies a means of **XON/XOFF** flow control using PAUSE. While this can decrease the frame loss due to oversubscription, the periods of no data transmission result in increased **latency delay** in the Ethernet link. The use of PAUSE as back pressure can result in congestion spreading and therefore it is rarely used.
- Congestion management, when used, may reduce the offered load at the congestion points without spreading congestion **will be performed to prevent oversubscription of the MAC Client without halting data transmission.** This specification will define a means of decreasing frame loss while permitting **decreased latency** increased efficiency in the Ethernet network **link**.
- The specification will be done in a format consistent with the IEEE document requirements thus making it easy for implementers to understand and to design.

Red text is deleted

Underscore text is inserted



Technical Feasibility

Demonstrated system feasibility
Proven technology, reasonable testing
Confidence in reliability

Red text is deleted
Underscore text is inserted

- Mechanisms for congestion management using congestion indication are known in the industry for some protocols and standards. Simulations of similar protocols show there are alternatives that can be feasibly implemented to accomplish the objectives within IEEE 802.
- The inclusion of congestion indication in layer 2 devices was anticipated in RFC 3168 “The Addition of Explicit Congestion Notification (ECN) to IP”.
- Rate control is commonly implemented in Ethernet devices.
- Ethernet supports a link level PAUSE mechanism using MAC Control frames, today. The means of exchanging congestion management information will use comparable technology.
- The testing for the generation of and response to any new MAC Control frame opcodes will be similar to the testing currently available for the MAC Control frame’s PAUSE opcode. Any such testing would rely on upper bounds on propagation delays for the media and the sublayers within an endstation and would need to be well defined throughout the document as they are today for PAUSE.
- This standard will maintain the reliability of the store-and-forward implementation of the MAC Control sublayer.
- Simulation efforts will evaluate the interaction with common protocols running above the MAC Client interface.



Economic Feasibility

Cost factors known, reliable data
Reasonable cost for performance
Total installation costs considered

Red text is deleted
Underscore text is inserted

- Possible solutions investigated for technical feasibility do not add significant complexity to Ethernet devices **the MAC or the MAC Control. Cost for the support of additional MAC Control opcodes is negligible.**
- **Simulations illustrate how congestion management reduces frame loss and consequently the costs of buffering frames in Ethernet switches and end stations. Simulations have shown reduced latency and increased throughput, which improves the overall performance of Ethernet.**
- Congestion management standardization will increase the broad market potential of Ethernet which will increase deployment and further reduce cost.
- System design, installation and maintenance costs are minimized by utilizing Ethernet system architecture, management, and software.



PAR Title

Information technology --

Telecommunications and information
exchange between systems -- Local and
metropolitan area networks -- specific
requirements Part 3: Carrier Sense
Multiple Access with Collision Detection
(CSMA/CD) Access Method and Physical
Layer Specifications Amendment:
Enhancements for Congestion
Management



PAR Scope

Red text is deleted
Underscore text is inserted

To specify IEEE 802.3 MAC parameters and minimal augmentation of MAC operation and management parameters **additions to and appropriate modifications** of IEEE Std 802.3 **for** to provide rate control and support of IEEE 802 congestion management.



PAR Purpose (14)

This project will enable accelerated deployment of Ethernet into emerging limited-topology applications that require improved delay, delay variation and frame loss characteristics.



PAR Purpose (14a)

Ethernet networks are being used in an increasing number of application spaces (clustering, backplanes, storage, data centers, etc.) that are sensitive to frame delay, delay variation and loss.


Study Group presentations have shown that Ethernet networks can experience higher throughput, lower delay, and lower frame loss by performing congestion management. This will improve Ethernet in its growing number of applications.



PAR Purpose (14) from Brad 9/22 / Gary 9/23

To improve the performance of an 802.3 link in terms of latency and frame discard in the presence of congestion.

Congestion due to oversubscription results in frame discards. This project will enable an exchange of congestion control information, which will **define / support** a means of decreasing frame discard and permit decreased latency.



PAR Purpose (14) from Jose 9/24

To improve the performance of an 802.3 link in terms of latency and frame discard in order to avoid congestion.


In case that the congestion appears, this project will enable an exchange of congestion control information, which will **define / support** a means of decreasing frame flow and permit decreased latency.



PAR Purpose (14a) from Brad 9/22

Ethernet as the transport for new applications has increased the need for, and the market benefit from, improved reliability of the frame delivery and improved latency performance.


Addition of these capabilities will accelerate Ethernet deployment into new billion dollar markets. An IEEE 802.3 standard will improve interoperability of equipment for these new Ethernet markets.



PAR Purpose (14a) from Ben 9/22

Ethernet as the next transport for ~~new~~ a range of applications has increased the need for, and the market benefit from, improved reliability of the frame delivery and improved latency performance.


Addition of these capabilities will accelerate Ethernet deployment into these new Ethernet ~~billion dollar~~ markets. An IEEE 802.3 standard will improve interoperability of equipment for ~~these new Ethernet~~ markets.



PAR Purpose (14a) from Brad / Gary 9/23 / Manoj 9/24

Ethernet is becoming the **transport / physical interconnect** of choice for a broader range of applications. These applications have increased the **market / need** for improved frame delivery and latency performance.

Addition of these capabilities in an IEEE 802.3 standard will accelerate Ethernet deployment and will improve interoperability of equipment in these new markets.



PAR Purpose (14a) from Matt 9/23

A new range of applications has demonstrated the need for improved reliability and latency over Ethernet links and networks. Enhancements to 802.3 to address these needs will accelerate the use of Ethernet as a transport for these applications.

An 802.3 standard is required to provide interoperability among multiple vendors.



Motion (Objectives)

- Move that the Congestion Management Study Group adopt the Objectives on slide 24 of agenda_2_0904.pdf replacing those adopted in July
- Move: H. Barrass Second: B. Booth
- Technical: 75%
- All in the room
- Y: 23 N: 0 A:3
- Motion: Passes
- Total in Room: 26
- Companies in the room: 16



Motion (Objectives)

- Move that the Congestion Management Study Group adopt the additional Objective “Minimize throughput reduction in non-congested flows”
- Move: Pat Thaler Second: George Cravens
- Technical: 75%
- All in the room
- Y: 21 N: 0 A: 4
- Motion: Passes



Motion (Broad Market Potential)

- Move that the Congestion Management Study Group adopt the response to the Broad Market Potential Criteria as presented on slide 27 of this presentation, to be included in critters_0904.pdf
- Move: Asif Hazarika Second: George Cravens
- Technical: 75%
- All in the room
- Y: 20 N: 0 A: 1
- Motion: Passes



Motion (Compatibility)

- Move that the Congestion Management Study Group adopt the response to the Compatibility Criteria as presented on slide 28 of this presentation, to be included in critters_0904.pdf
- Move: Hugh Barrass Second: Manoj Wadekar
- Technical: 75%
- All in the room
- Y: 17 N: 0 A: 3
- Motion passes



Motion (Distinct Identity)

- Move that the Congestion Management Study Group adopt the response to the Distinct Identity Criteria as presented on slide 29 of this presentation, to be included in critters_0904.pdf
- Move: Gopal Hegde Second: George Cravens
- Technical: 75%
- All in the room
- Y: 19 N: 0 A: 2
- Motion passes




Motion (Technical Feasibility)

- Move that the Congestion Management Study Group adopt the response to the Technical Feasibility Criteria as presented on slide 30 of this presentation, to be included in critters_0904.pdf
- Move: Tanmay Gupta Second: Matt Squire
- Technical: 75%
- All in the room
- Y: 18 N: 0 A: 3
- Motion passes



Motion (Economic Feasibility)

- Move that the Congestion Management Study Group adopt the response to the Economic Feasibility Criteria as presented on slide 31 of this presentation, to be included in critters_0904.pdf
- Move: David Martin Second: Kevin Daines
- Technical: 75%
- All in the room
- Y: 19 N: 0 A: 2
- Motion passes



Motion (PAR)

- Move that the Congestion Management Study Group adopt the Title, Scope, Purpose (14) and Purpose (14a) of the PAR as presented on slides 32-35 of this presentation, to be included on the PAR form
- Move: Gopal Hegde Second: George Cravens
- Technical: 75%
- All in the room
- Y: 16 N: 0 A: 2
- Motion passes



Motion (Forward)

- Move that the Congestion Management Study Group forward the PAR, 5 Criteria, and Objectives to EC and 802.3 for consideration at the November Plenary
- Move: Manoj Wadekar Second: George Cravens
- Technical: 75%
- All in the room
- Y: 16 N: 0 A: 2
- Motion Passes

Future Meetings

- Nov 2004 Plenary
 - ☐ Week of the 14th
 - ☐ San Antonio, TX
 - ☐ St. Anthony/Wyndham
- How many plan to attend the CMSG?
 - ☐ ?
- January 2005 Interim:
 - ☐ Vancouver, IEEE host, Jan 24-28: Yes – 6
 - ☐ Sacramento, Intel host, Jan 17-21: Yes – 8
 - ☐ Sacramento, CMSG only with 802.1, Jan 10-14: Yes – 14





Adjourn

Thank you!