



IEEE 802.3
Congestion Management
Study Group

Long Beach, CA
24 May, 2004



Agenda

- Welcome and Introductions
- Appoint/Volunteer Recording Secretary
- 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:
 - Project Authorization Request (PAR) Scope and Purpose
 - Responses to the 5 Criteria
 - Congestion Management Objectives



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



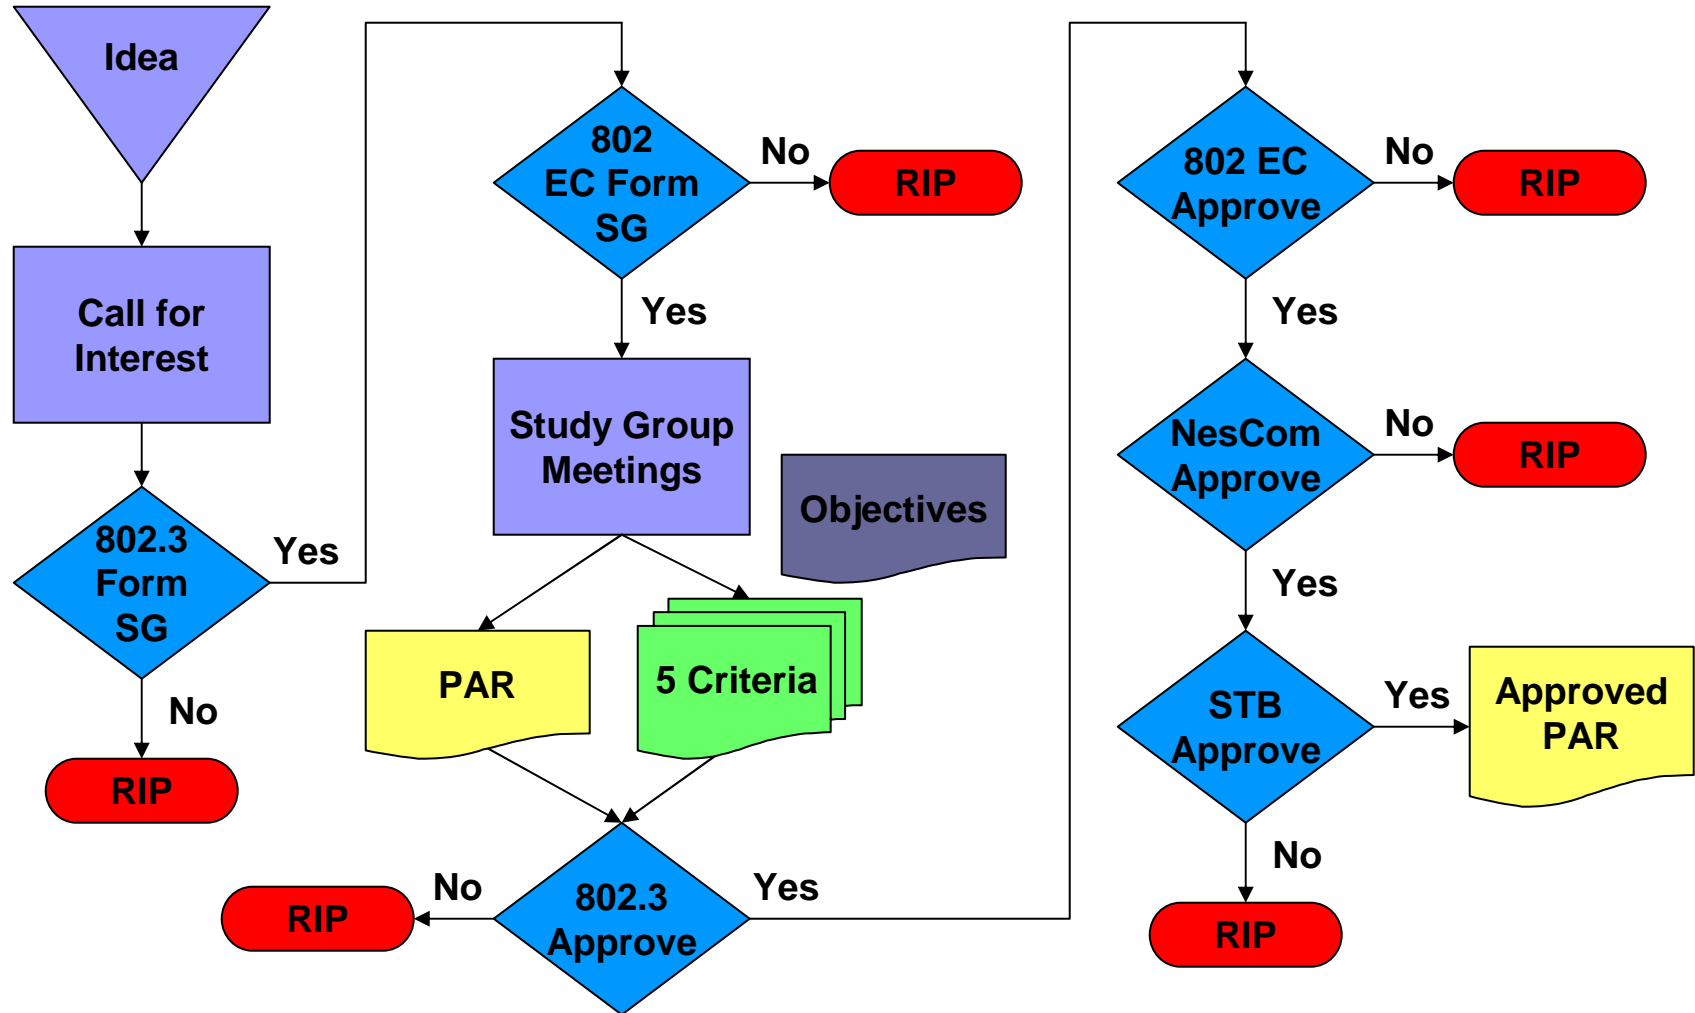
Inappropriate Topics for IEEE SG Meetings

- 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.

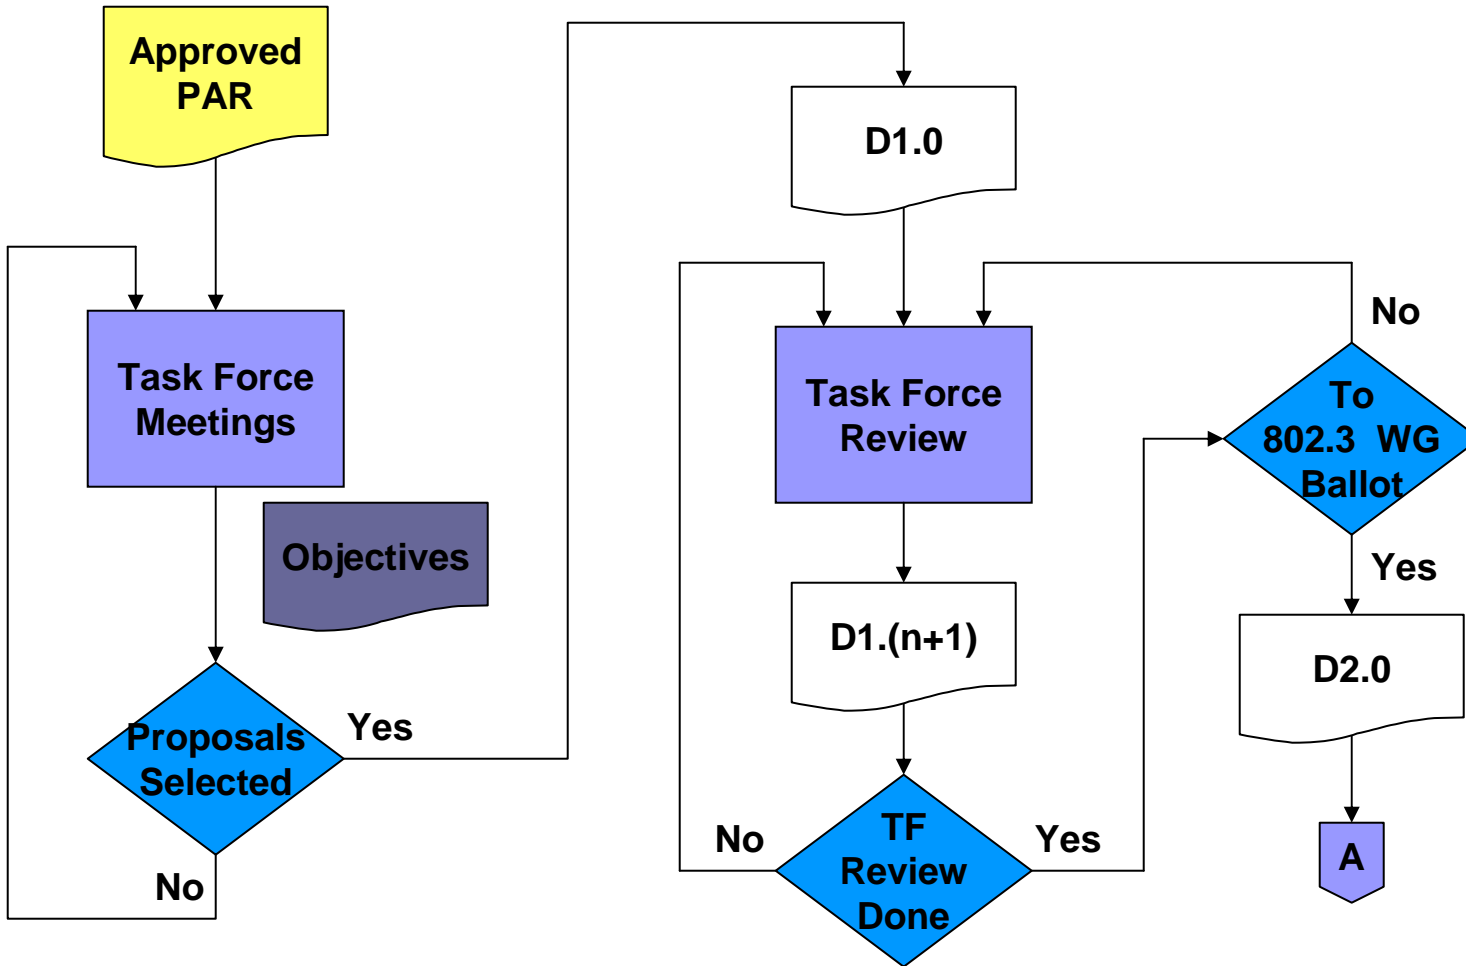
If you have questions, contact the IEEE Patent Committee Administrator at patcom@ieee.org

Approved by IEEE-SA Standards Board – December 2002

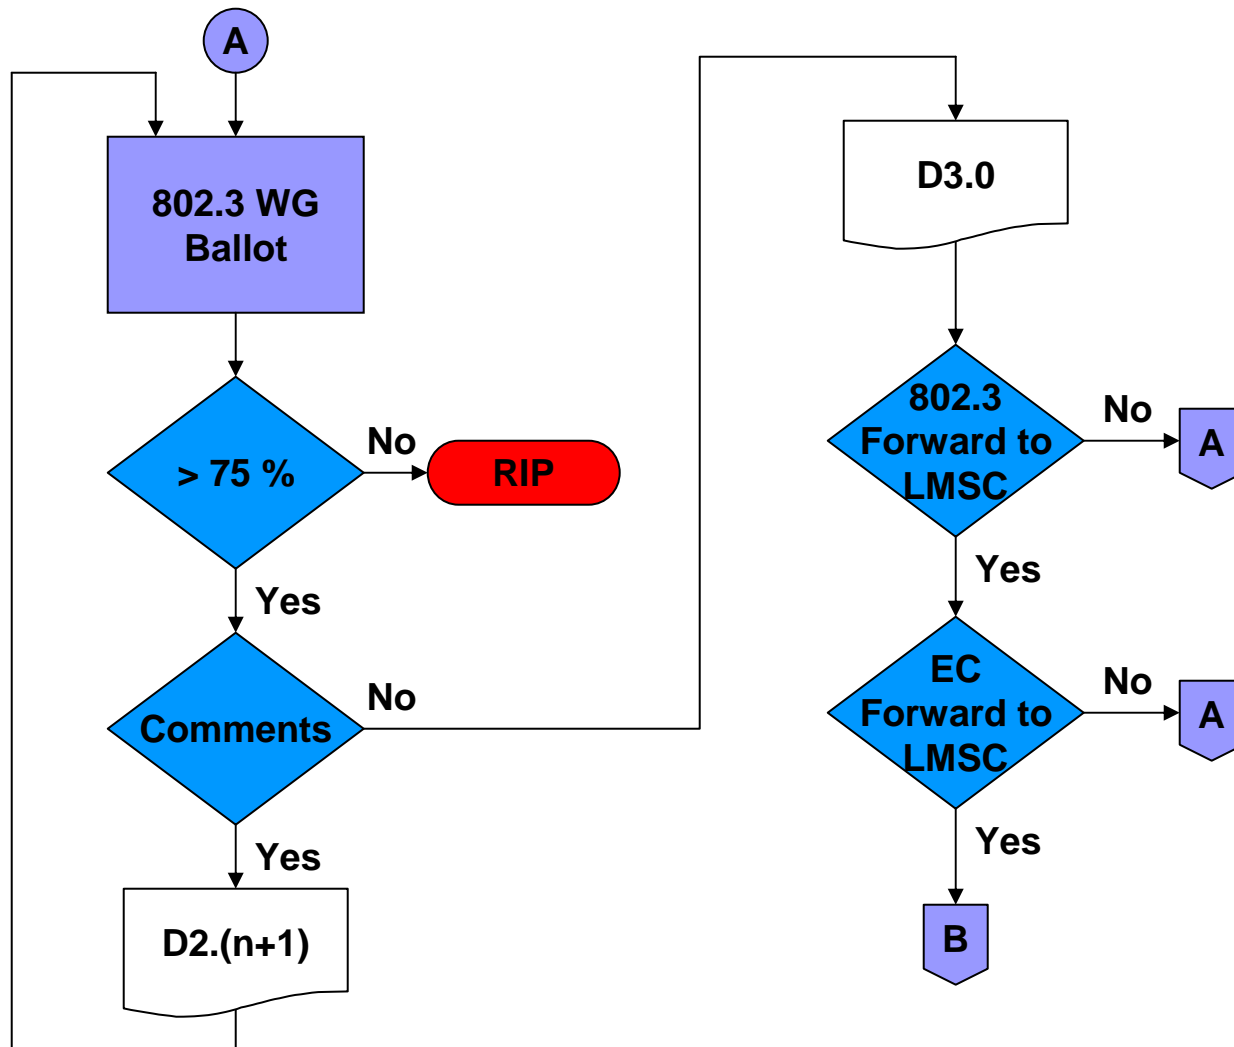
IEEE Standards Process




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.)

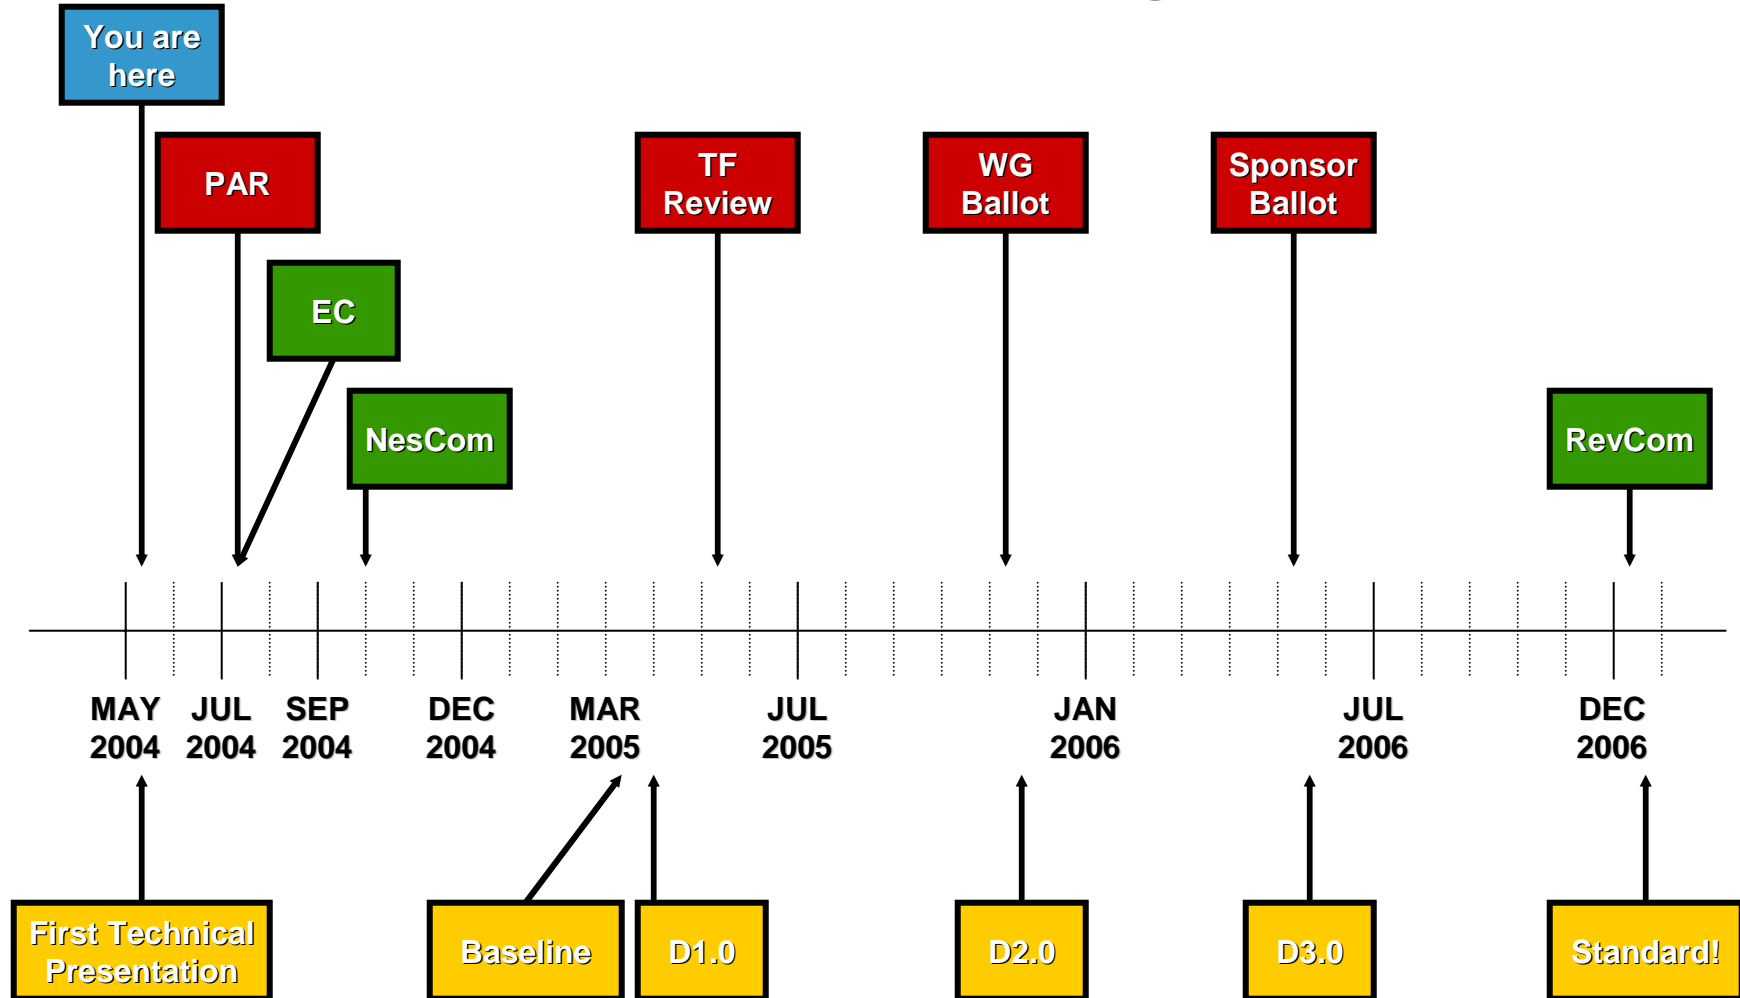
■ Technical Feasibility

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

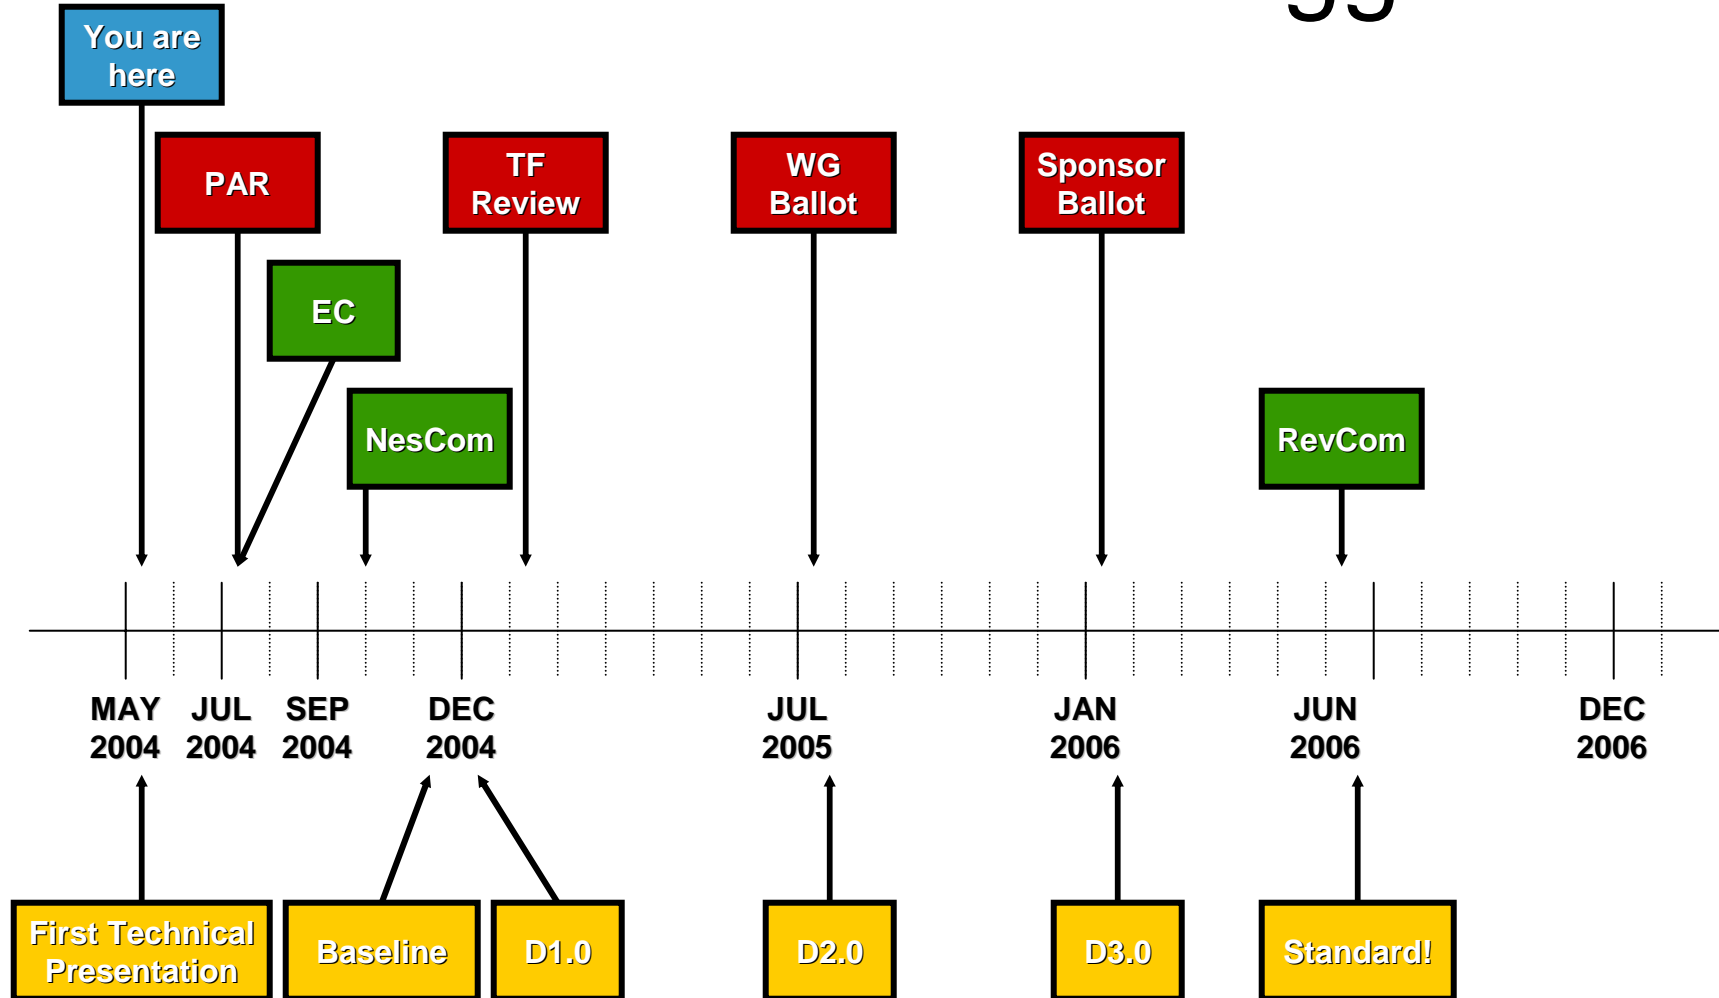
■ Economic Feasibility

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

Possible Congestion Management Timeline - Conservative

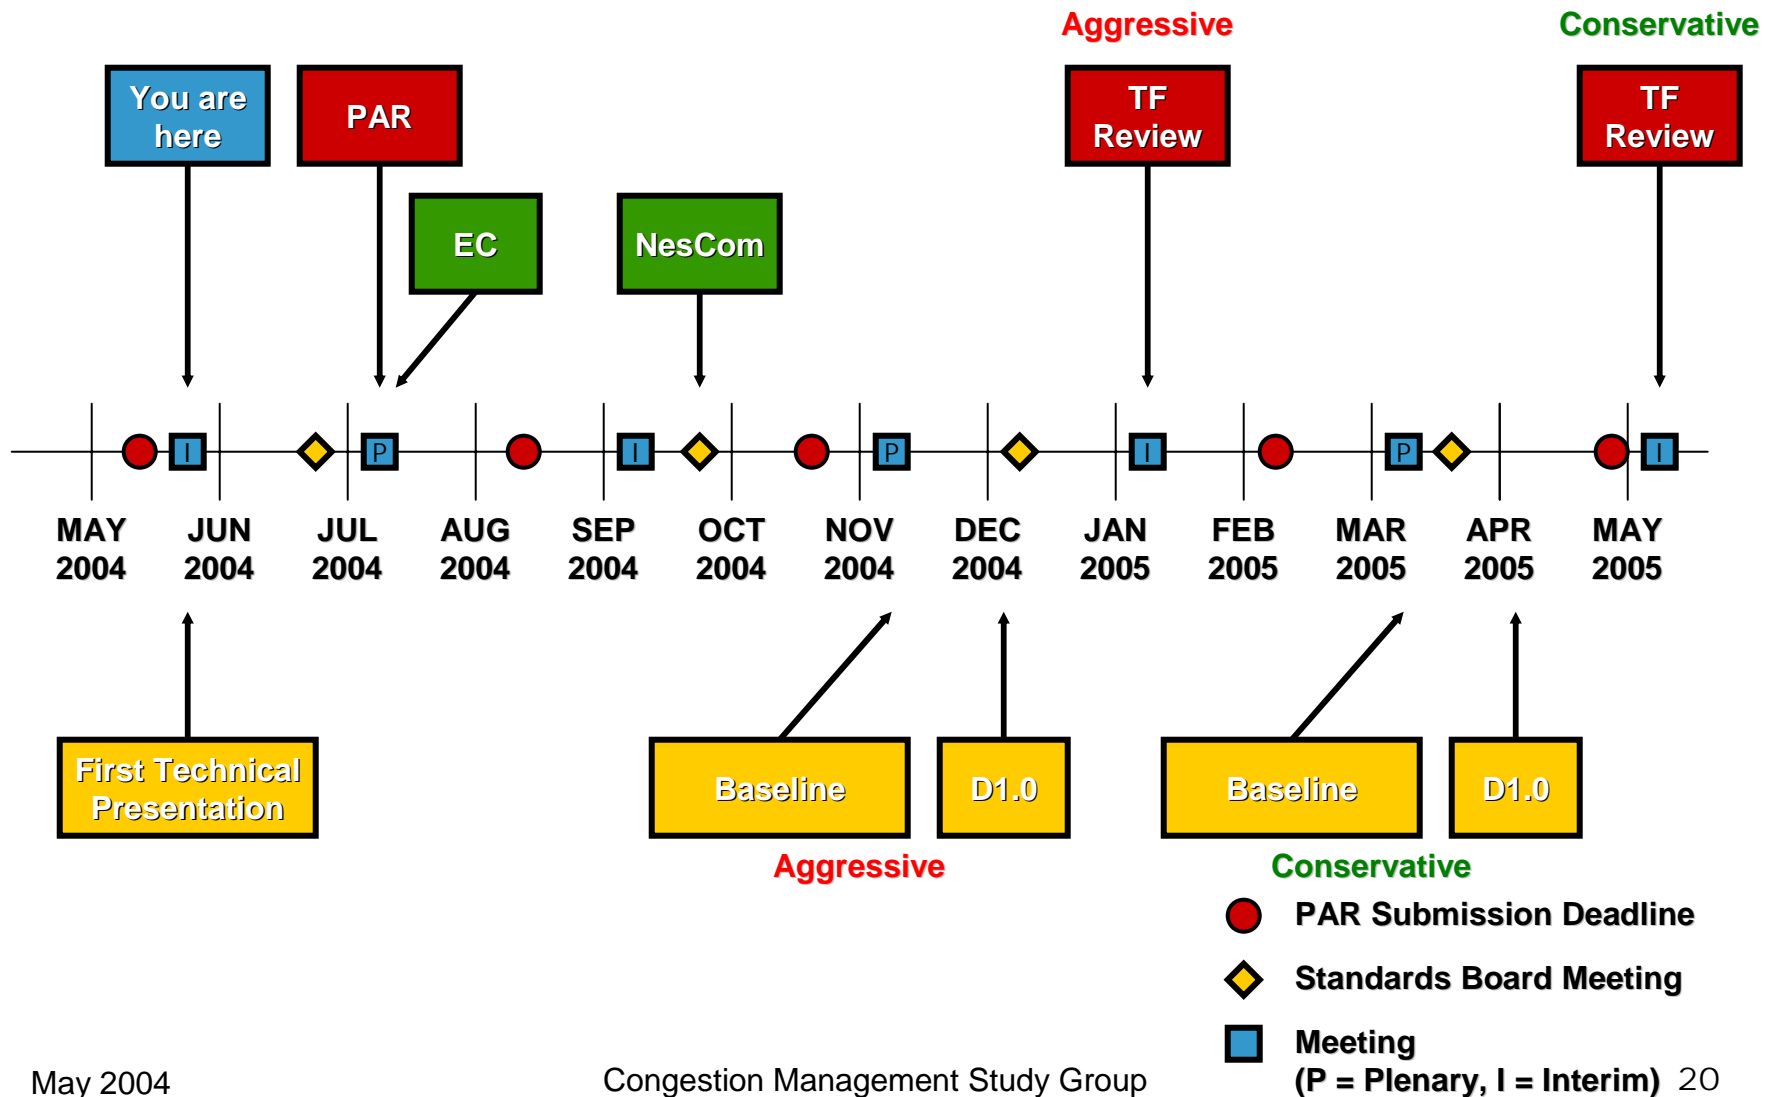


Possible Congestion Management Timeline - Aggressive





Timeline Detail



May 2004

Congestion Management Study Group

20

Presentations

Monday May 24, 2004				
Presenter	Topic	Length	Start	Finish
	Welcome and Introductions	0:10	8:30 AM	8:40 AM
Brown, Ben	Agenda	0:20	8:40 AM	9:00 AM
Grow, Bob	The Structure for Congestion Management	0:20	9:00 AM	9:20 AM
Hegde, Gopal	IEEE Congestion Management	0:50	9:20 AM	10:10 AM
Morning Break		0:20	10:10 AM	10:30 AM
Thatcher, Jonathan	Congestion Management -- Managing the Layer Stack	1:10	10:30 AM	11:40 AM
Martin, David	A Survey of Standards Efforts on Traffic and Congestion Management in Ethernet Networks	0:20	11:40 AM	12:00 PM
Lunch		1:00	12:00 PM	1:00 PM
Merchant, Shashank	Congestion Management in a Bladed System	0:15	1:00 PM	1:15 PM
Lynskey, Eric	Preemption Simulations	0:20	1:15 PM	1:35 PM
Wadekar, Manoj	Proposal for 802.3 Enhancements for Congestion Management	1:00	1:35 PM	2:35 PM
Afternoon Break		0:25	2:35 PM	3:00 PM
Group	Discussion of Objectives	2:00	3:00 PM	5:00 PM
Tuesday May 25, 2004				
Presenter	Topic	Length	Start	Finish
Group	Welcome	0:15	8:30 AM	8:45 AM
Group	Discussion of Scope, Purpose, 5 Criteria	1:15	8:45 AM	10:00 AM
Morning Break		0:15	10:00 AM	10:15 AM
Group	Motion Madness - do we precirculate?	1:30	10:15 AM	11:45 AM



Task Force Objectives

- Provide a mechanism for rate limiting
 - Straw poll: 22 in favor / 3 against
- Support for full duplex
 - 23 in favor / 1 against
- Optimize solutions for short, 10 Gbps links
 - Deferred



Study Group Objectives

- Evaluate 802.3x with finer granularity
 - 19 in favor / 1 against
- Address latency, latency variation and frame loss
 - 21 in favor / 1 against
- Evaluate rate limiting
 - Feed forward/back
 - Static/dynamic
 - 21 in favor / 3 against
- Consider preemption
 - 5 in favor / 9 against



IEEE 802.3 CMSG Objectives

- No changes to PHYs
 - 19 in favor / 1 against
- Not intended to work with existing PAUSE
 - Deferred
- This is not an end-to-end flow control
 - 21 in favor / 2 against
- No new methods for traffic classification
 - 16 in favor / 2 against
- No reordering of packets within a class
 - 21 in favor / 1 against



IEEE 802.3 CMSG Objectives

- Support up to 100 m of media (copper or optical)
 - Deferred
- Support 100 Mb/s, 1 Gb/s, and 10Gb/s
 - Deferred
- Be consistent with IEEE 802.3 and IEEE 802.1 layer architecture
 - 21 in favor / 0 against
- Provide predictable, consistent network-wide operation
 - Deferred
- Be consistent with slow protocols (e.g. OAM)
 - 15 in favor / 0 against



IEEE 802.3 CMSG Objectives (cont.)

- To define 802.3 congestion control support that, at a minimum, will do nothing to degrade the operation of existing upper layer protocols and flow/congestion control mechanisms, but has the explicit goal of facilitating the improved operation of some existing and emerging protocols, over 802.3 full-duplex link technology.
 - Deferred



Broad Market Potential

Broad set(s) of applications

Multiple vendors, multiple users

Balanced cost (LAN vs. attached stations)

- **Ethernet links have begun carrying an ever widening variation of traffic types as they get used in an increasing number of application spaces including cluster interconnect, backplanes, data centers, etc. Today, all traffic is treated equally on these links, which impacts throughput, latency, and frame discard for traffic over layer 2 networks. This limitation is a barrier to broader acceptance of Ethernet in these application spaces.**
- **Presentations have been made to indicate that IEEE 802.3 can positively impact throughput, latency, and frame discard by combining traffic differentiation with rate/flow control enhancements. Presentations also showed that a substantial market potential might be at risk without these features.**
- **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. (Review these numbers at the May meeting for additional support.)**
- **A standard to support congestion management will maintain 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 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 Interworking.
- 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 specification does not explicitly cover congestion management for differentiated traffic types. The PAUSE opcode for the MAC Control frame is specific to all traffic on the link.**
- **The standard will define the means to differentiate traffic types and controls (e.g. one or more additional opcodes for the MAC Control frame) to support congestion management of differentiated traffic types.**
- **The specification will be done in a format consistent with the IEEE document requirements thus making it easy for implementers to understand and design to.**
- **The proposed specification will support congestion management of differentiated traffic types similar to other networking technologies (Fibre Channel, IB, PCI Express) but does so using the IEEE 802.3 MAC, which breaks down the barrier to broader acceptance of Ethernet as a valid competitor to these technologies.**



Technical Feasibility

Demonstrated system feasibility
Proven technology, reasonable testing
Confidence in reliability

- **Ethernet supports different traffic types, today.**
- **Ethernet supports a link level PAUSE mechanism using MAC Control frames, today.**
- **Ethernet supports various means of differentiating traffic types (EtherType, VLAN Priority), today.**
- **The testing for the generation of and response to any new MAC Control frame opcodes is 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 project will not modify PMA/PMD, the MAC, nor the bridge and therefore will not be introducing any significant impact on system reliability.**



Economic Feasibility

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

- **The component costs will benefit from cost reduction associated with Moore's Law. Further integration of functionality will reduce cost.**
- **Costs for the support of additional MAC Control functions are negligible when compared to a MAC chip.**
- **Congestion management standardization will increase deployment and diversity of supply base to further reduce cost.**
- **Ethernet IP re-use will lower implementation 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

- The scope of this project is to specify additions to and appropriate modifications of IEEE Std 802.3 for congestion management. This includes but is not limited to improved latency for high priority traffic and reduced need for frame discard as a stimulus for flow control in Ethernet networks.




PAR Purpose (14)

- The purpose of this project is to improve Ethernet performance for high priority traffic in the presence of transitory congestion or congestion caused by oversubscription of network resources. This differentiation of service will enable easier convergence of the different types of traffic that Ethernet carries.




PAR Purpose (14a) (1 of 3)

- Ethernet configurations are shrinking. Servers used to be pedestals interconnected with traditional enterprise Ethernet links. A significant portion of servers are now implemented as rack mount devices again using traditional enterprise Ethernet links. Today, servers are migrating into chassis with the Ethernet connections now moving onto the backplane. At the same time these physical form factor migrations have been going on, Ethernet has expanded to handle new traffic types. IP telephony over Ethernet, video over Ethernet and Ethernet becoming an interconnect for telecommunications servers and telecommunications interface blades are driving demand for improved performance capabilities of Ethernet in these high density environments.



PAR Purpose (14a) (2 of 3)

- Ethernet as the transport for IP telephony and similar traffic types has increased the need for and the market benefit from granting certain traffic classes preferred latency performance. This differentiation of service can provide both shorter latency in the presence of congestion and less latency jitter for high priority traffic.
- Ethernet as the transport for storage has some contrasting requirements. Storage and similar applications benefit significantly from improved reliability of the frame delivery. Network congestion results in frame drops within the network, and improved flow control or rate management will result in fewer frame drops.



PAR Purpose (14a) (3 of 3)

- Addition of these capabilities will accelerate Ethernet deployment into new billion dollar markets. While proprietary solutions for these problems are in the market, an IEEE 802.3 standard will improve interoperability of equipment for these newer Ethernet markets.



Motion

- **Move that the CMSG adopt the 5 Criteria as contained within this document.**

- Move:

- Second:

- All voters

- Yes:

- No:

- Abstain:

- 802.3 Voters

- Yes:

- No:

- Abstain:

- 75% Required: Motion



Motion

- **Move that the CMSG forward the PAR (based on the Title, Scope, and Purpose slides in this document), 5 Criteria, and Objectives to 802.3 for consideration at the July Plenary.**

- Move:

- Second:

- All voters

- Yes:

- No:

- Abstain:

- 802.3 Voters

- Yes:

- No:

- Abstain:

- 75% Required: Motion



Action Items

- Presentation Follow-ons

- Bob Grow
- Gopal Hegde
- Jonathan Thatcher
- David Martin
- Shashank Merchant
- Eric Lynskey
- Manoj Wadekar

- Add names to the support list – work together

- Education gains consensus – use the reflector

- Additional review of Objectives, 5 Criteria, PAR



Question

- Do you feel the study group should extend until November?
 - Y:24
 - N:2
 - 30 in attendance

Future Meetings

- July 2004 Plenary:
 - Week of the 11th
 - Portland, OR
 - Embassy Suites
- How many plan to attend the CMSG?
 - 22
- Sep 2004 Interim
 - ???





Adjourn

Thank you!