IEEE P802.3ar
Congestion Management Task Force

Draft edits to PAR, 5 criteria, objectives

26 April 2006

Draft edits to PAR

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

proposed

2.1 Title of Standard

- 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 Rate Limiting

PAR scope

- To specify IEEE 802.3 MAC parameters and minimal augmentation of MAC operation and management parameters of IEEE Std 802.3 to provide rate control and support of IEEE 802 congestion management.

proposed

5.2 Scope of Proposed Standard

- To specify IEEE 802.3 MAC parameters and minimal augmentation of MAC operation and management parameters of IEEE Std 802.3 to provide rate control. and support of IEEE 802 congestion management.

PAR Scope *(contd)*

- Is the completion of this standard is dependent upon the completion of another standard: Yes
- If yes, please explain:
  - This PAR includes work on independent capabilities, and some of the work is not contingent on another project. This PAR is being launched to match the anticipated completion of work to be done within 802.1. The complete capabilities and benefits envisioned in Ethernet networks through Congestion Management capabilities will include both 802.3 and 802.1 work. Serialization of the projects would unnecessarily delay market introduction of the capabilities. If 802.1 work is not completed for the cooperative efforts or the progress would unnecessarily delay either the non-dependent or dependent capabilities also included in this PAR, the work will be split into two PARs.

5.3

- Is the completion of this standard is dependent upon the completion of another standard: Yes No
- If yes, please explain:

- *Note:*
  - *Modified PAR eliminates previous contingency.*

5.4 Purpose of Proposed Standard:

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

no change

5.4 Purpose of Proposed Standard:

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

current

5.5 Need for the project: 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.

proposed

5.5 Need for the project: 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 rate control. This will improve Ethernet in its growing number of applications.

5.6 Stake holders for the project

- Clustering, backplanes, storage, data centers

Draft edits to 5 criteria

Broad market potential

- Broad set(s) of applications
- Multiple vendors, multiple users
- Balanced cost (LAN vs. attached stations)

- 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.
- 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 30 people from at least 16 companies in attendance.
- A standard to support congestion management will respect the balance of cost between LAN and attached stations.

proposed

Broad market potential

- Broad set(s) of applications
- Multiple vendors, multiple users
- Balanced cost (LAN vs. attached stations)

- 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.
- Rate control is an effective technique to reduce buffer requirements when there are known/fixed bottlenecks in the networks.
- 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 and task force meetings, there have been up to 35 people from at least 16 companies in attendance.
- A standard to support congestion management rate limiting will respect the balance of cost between LAN and attached stations.

Compatibility

- 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.

Compatibility

- 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 flow control using PAUSE.
- While this can decrease the frame loss due to oversubscription, the periods of no data transmission result in increased 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. This specification will define a means of decreasing frame loss while permitting increased efficiency in the Ethernet network.
- The specification will be done in a format consistent with the IEEE document requirements

proposed

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 flow control using PAUSE.
- While this can decrease the frame loss due to oversubscription, the periods of no data transmission result in increased delay in the Ethernet link. The use of PAUSE as back pressure can result in congestion spreading and therefore it is rarely used. This is not best solution for addressing known/fixed bottlenecks.
- Congestion management Rate limiting, when used, may will reduce the offered load at the congestion points without spreading congestion. Rate limiting could address bottlenecks due to date rate mismatch as well as mismatch due to protocol overheads. This specification will define a means of decreasing frame loss while permitting increased efficiency in the Ethernet network.
- The specification will be done in a format consistent with the IEEE document requirements

Technical feasibility

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

- 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.

Technical feasibility

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

- 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 (e.g., ifsStretch). Providing common framework, method(s), and parameters will enable interoperability between vendors.

Economic feasibility

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

- Possible solutions investigated for technical feasibility do not add significant complexity to Ethernet devices.
- 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.

Economic feasibility

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

- Possible solutions investigated for technical feasibility do not add significant complexity to Ethernet devices.
- Congestion management Rate limiting 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.

Draft edits to objectives

Objectives

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

Approved by IEEE 802.3 WG on 18-Nov-2004

802.3ar TF Objectives

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

Revised by IEEE 802.3ar TF on 07-March-2006

No vote taken by IEEE 802.3 WG at March 2006 plenary