Proposal for 5 Criteria Responses

Ilango Ganga, Intel
John Jaeger, Infinera

HSSG

July 16, 2007

San Francisco, CA

Supporters

Adam Bechtel Yahoo
Mike Bennett - LBNL
Uri Cummings - Fulcrum Microsystems
John Dallesasse - Emcore
Wael Diab - Broadcom
Chris DiMinico - MC Communications
Alan Flatman - LAN Technologies
Howard Frazier - Broadcom
Ilango Ganga - Intel
Garth Gibson - Panasas
Joel Goergen - Force10
Daniel Golding - Tier 1 Research
Larry Green - Ixia
Greg Hankins - Force10
Peter Harrison - Netflix
Robert Hays - Intel
Osamu Ishida - NTT
Kenneth Jackson - Emcore
John Jaeger - Infinera
Alan Judge - Amazon
Scott Kipp - Brocade
Shoukei Kobayashi - NTT
Paul Kolesar - CommScope
David Koenen - HP
Mark Kortekaas - CBS
Glen Kramer - Teknovus
Louis Lee - Equinix
Robert Lingle, Jr. - OFS
Ronald Luijten - IBM Research
Trey Malpass - Malpass Technology
Jeff Maki - Juniper Networks
Arlon Martin - Kotura
David Martin - Nortel
Jim McGrath, Molex
Andy Moorwood - Extreme Networks
Jay Moran - AOL
Shimon Muller - Sun
Gary Nicholl - Cisco
Shinji Nishimura - Hitachi
Mark Nowell - Cisco
Gourgen Oganessyan, Molex
George Oulundsen – OFS
Tom Palkert - Xilinx
Shashi Patel - Foundry Networks
Petar Pepeljugoski - IBM Research
Drew Perkins - Infinera
Jan Peeters Weem, Aprius
Christopher Quesada - Switch & DataVik Saxena - Comcast
Peter Schoenmaker - NTT America
Ted Seely - Sprint
Charles Seitz - Myricom
Henk Steenman - AMS-IX
Andre Szczepanek - Texas Instruments
Geoff Thompson - Nortel
Bruce Tolley - Solarflare
Hidehiro Toyoda - Hitachi
Matt Traverso - Opnext
Bill Trubey - Time Warner Cable
Brad Turner – Juniper Networks
Schelto Van Doorn - Independent
Jason Weil - Cox Communications
Doug Wilson - MSN
Robert Winter - Dell
Bill Woodruff - Aquantia

Broad Market Potential

- **Broad sets of applications**
- **Multiple vendors and numerous users**
- **Balanced cost (LAN versus attached stations)**

- **Bandwidth requirements for computing and core networking applications are growing at different rates, which necessitates the definition of two distinct data rates for the next generation of Ethernet networks in order to address these applications:**
  - Servers, high performance computing clusters, blade servers, storage area networks and network attached storage all currently make use of 1G and 10G Ethernet, with significant growth of 10G projected in '07 and '08. I/O bandwidth projections for server and computing applications indicate that there will be a significant market potential for a 40 Gb/s Ethernet interface.
  - Core networking applications have demonstrated the need for bandwidth beyond existing capabilities and the projected bandwidth requirements for computing applications. Switching, routing, and aggregation in data centers, internet exchanges and service provider peering points, and high bandwidth applications, such as video on demand and high performance computing environments, have demonstrated the need for a 100 Gb/s Ethernet interface.
- **There has been wide attendance and participation in the study group by end users, equipment manufacturers and component suppliers. It is anticipated that there will be sufficient participation to effectively complete the standardization process.**
- **Prior experience scaling IEEE 802.3 and contributions to the study group indicates:**
  - 40 Gb/s Ethernet will provide approximately the same cost balance between the LAN and the attached stations as 10 Gb/s Ethernet.
  - The cost distribution between routers, switches, and the infrastructure remains acceptably balanced for 100 Gb/s Ethernet.
- **Given the topologies of the networks and intended applications, early deployment will be driven by key aggregation & high-bandwidth interconnect points. This is unlike the higher volume end system application typical for 10/100/1000 Mb/s Ethernet, and as such, the initial volumes for 100 Gb/s Ethernet are anticipated to be more modest than the lower speeds. This does not imply a reduction in the need or value of 100 Gb/s Ethernet to address the stated applications.**

Compatibility

- **IEEE 802 defines a family of standards. All standards shall be in conformance with the IEEE 802.1 Architecture, Management, and Interworking documents as follows: 802. Overview and Architecture, 802.1D, 802.1Q, and parts of 802.1f. If any variances in conformance emerge, they shall be thoroughly disclosed and reviewed with 802. Each standard in the IEEE 802 family of standards shall include a definition of managed objects that are compatible with systems management standards.**

- As an amendment to IEEE Std 802.3, the proposed project will remain in conformance with the IEEE 802 Overview and Architecture as well as the bridging standards IEEE Std 802.1D and IEEE Std 802.1Q.
- As an amendment to IEEE Std 802.3, the proposed project will follow the existing format and structure of IEEE 802.3 MIB definitions providing a protocol independent specification of managed objects (IEEE Std 802.1F).
- The proposed amendment will conform to the full-duplex operating mode of the IEEE 802.3 MAC.
- As was the case in previous IEEE 802.3 amendments, new physical layers specific to either 40 Gb/s or 100 Gb/s operation will be defined.
- By utilizing the existing IEEE 802.3 MAC protocol, this proposed amendment will maintain maximum compatibility with the installed base of Ethernet nodes.

Distinct Identity

- **Substantially different from other IEEE 802 standards**
- **One unique solution per problem (not two solutions to a problem)**
- **Easy for the document reader to select the relevant specification**

- The proposed amendment is an upgrade path for IEEE 802.3 users, based on the IEEE 802.3 MAC.
- The established benefits of the IEEE 802.3 MAC include:
  - Deterministic, highly efficient full-duplex operation mode
  - Well-characterized and understood operating behavior
  - Broad base of expertise in suppliers and customers
  - Straightforward bridging between networks at different data rates
- The Management Information Base (MIB) for IEEE 802.3 will be extended in a manner consistent with the IEEE 802.3 MIB for 10 / 100 / 1000 / 10000 Mb/s operation.
- The proposed amendment to the existing IEEE 802.3 standard will be formatted as a collection of new clauses, making it easy for the reader to select the relevant specification.
- Bandwidth requirements for computing and networking applications are growing at different rates. These applications have different cost / performance requirements, which necessitates two distinct data rates, 40 Gb/s and 100 Gb/s.

Technical Feasibility

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

- The principle of scaling the IEEE 802.3 MAC to higher speeds has been well established by previous work within IEEE 802.3.
- The principle of building bridging equipment which performs rate adaptation between IEEE 802.3 networks operating at different speeds has been amply demonstrated by the broad set of product offerings that bridge between 10, 100, 1000, and 10000 Mb/s.
- Systems with an aggregate bandwidth of greater than or equal to 100 Gb/s have been demonstrated and deployed in operational networks.
- The proposed project will build on the array of Ethernet component and system design experience, and the broad knowledge base of Ethernet network operation.
  - The experience gained in the development and deployment of 10 Gb/s technology is applicable to the development of specifications for components at higher speeds. For example, parallel transmission techniques allow reuse of 10 Gb/s technology and testing.
  - Component vendors have presented data on the feasibility of the necessary components for higher speed solutions. Proposals, which either leverage existing technologies or employ new technologies, have been provided.
- The reliability of Ethernet components and systems can be projected in the target environments with a high degree of confidence. Presentations demonstrating this have been provided.

Economic Feasibility

- **Known cost factors, reliable data**
- **Reasonable cost for performance**
- **Consideration of installation costs**

- The cost factors for Ethernet components and systems are well known. The proposed project may introduce new cost factors which can be quantified.
- Presentations indicate that for the server market and computing applications the optimized rate to provide the best balance of performance and cost is 40 Gb/s. For the network aggregation market and core networking applications, the optimized rate offering the best balance of performance and cost is 100 Gb/s.
- In consideration of installation costs, the project is expected to use proven and familiar media, including optical fiber, backplanes, and copper cabling technology.
- Network design, installation and maintenance costs are minimized by preserving network architecture, management, and software.