



Two PAR Proposal

Dan Dove

Dove Networking Solutions

for

ProCurve Networking by HP

Preface

- This presentation presumes that sufficient material will be presented in July 2007 to complete the justification for a project to create a 40G Ethernet standard.
- This presentation provides a basis for the HSSG to advance a two-PAR proposal should we determine that a single-PAR approach is not preferred.
- Either of these PARs should be able to stand-alone as well.
- The content of these PARs is subject to amendment by the HSSG prior to advancement to 802.3.
- Slides 3 – 12 constitute the 100G PAR, Objectives, and 5 Criteria
- Slides 13 – 22 constitute the 40G PAR, Objectives, and 5 Criteria

HSSG “PAR A” (Working Draft)

- Title -
 - IEEE Standard for 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: Media Access Control Parameters, Physical Layers and Management Parameters for 100 Gb/s Operation

HSSG “PAR A” (Working Draft)

- Scope –
 - Define 802.3 Media Access Control (MAC) parameters, physical layer specifications, and management parameters for the transfer of 802.3 format frames at 100 Gb/s.

HSSG “PAR A” (Working Draft)

- Purpose

- The purpose of this project is to extend the 802.3 protocol to an operating speed of 100 Gb/s in order to provide a significant increase in bandwidth while maintaining maximum compatibility with the installed base of 802.3 interfaces, previous investment in research and development, and principles of network operation and management. The project is to provide for the interconnection of equipment satisfying the distance requirements of the intended applications.

HSSG “PAR A” (Working Draft)

- Need –
 - The project is necessary to provide a solution for applications that have been demonstrated to need bandwidth beyond the existing capabilities. These include data center, internet exchanges, high performance computing and video-on-demand delivery. 100Gb/s Ethernet will provide the next generation of high speed interconnect for these services.

“PAR A” Objectives

- Support full-duplex operation only
- Preserve the 802.3 / Ethernet frame format at the MAC Client service interface
- Preserve minimum and maximum FrameSize of current 802.3 Std
- Support a BER better than or equal to 10^{-12} at the MAC /PLS service interface.
- Support a speed of 100 Gb/s at the MAC/PLS interface
 - Provide Physical Layer specifications which support 100 Gb/s operation over:
 - at least 40km on SMF.
 - at least 10km on SMF.
 - at least 100 meters on OM3 MMF.
 - at least 10m over a copper cable assembly.

“PAR A” 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:**
 - 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:**
 - 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.**

“PAR A” 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 100 Gb/s operation will be defined.

“PAR A” 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.
- By utilizing the existing IEEE 802.3 MAC protocol, this proposed amendment will maintain maximum compatibility with the installed base of Ethernet nodes.
- 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.

“PAR A” 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.

“PAR A” 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 from component and equipment suppliers and their customers indicate:
 - 100 Gb/s Ethernet will be optimized for core networking applications that are performance driven, and will offer better value than alternate approaches or technologies.
- In consideration of installation costs, the project is expected to use proven and familiar media, including deployed optical fiber, backplanes, and copper cabling technology.
- Network design, installation and maintenance costs are minimized by preserving network architecture, management, and software.

HSSG “PAR B” (Working Draft)

- Title -
 - IEEE Standard for 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: Media Access Control Parameters, Physical Layers and Management Parameters for 40 Gb/s Operation

HSSG “PAR B” (Working Draft)

- Scope –
 - Define 802.3 Media Access Control (MAC) parameters, physical layer specifications, and management parameters for the transfer of 802.3 format frames at 40 Gb/s.

HSSG “PAR B” (Working Draft)

- Purpose

- The purpose of this project is to extend the 802.3 protocol to operating speeds of 40 Gb/s in order to provide a significant increase in bandwidth while maintaining maximum compatibility with the installed base of 802.3 interfaces, previous investment in research and development, and principles of network operation and management. The project is to provide for the interconnection of equipment satisfying the distance requirements of the intended applications.

HSSG “PAR B” (Working Draft)

- Need –
 - The project is necessary to provide a solution for applications that have been demonstrated to need bandwidth beyond the existing capabilities. These include high speed server interconnect.

“PAR B” Objectives

- Support full-duplex operation only
- Preserve the 802.3 / Ethernet frame format at the MAC Client service interface
- Preserve minimum and maximum FrameSize of current 802.3 Std
- Support a BER better than or equal to 10^{-12} at the MAC /PLS service interface.
- Support a speed of 40 Gb/s at the MAC/PLS service interface
 - Provide Physical Layer specifications which support 40 Gb/s operation over:
 - at least 100 m on OM3 MMF
 - at least 10 m over a copper cable assembly
 - at least 1 m over a backplane

“PAR B” 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 computing applications indicate that there will be a significant market potential for a 40 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.
- **Given the topologies of the networks and intended applications, early deployment will be driven by ultra-high performance servers and then mature to higher volume end system application typical for 10/100/1000 Mb/s Ethernet. As such, the initial volumes for 40 Gb/s Ethernet are anticipated to be more modest than the lower speeds but will grow with time and serve a large market need. This does not imply a reduction in the need or value of 40 Gb/s Ethernet to address the stated applications.**

“PAR B” 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 40 Gb/s operation will be defined.

“PAR B” 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.
- By utilizing the existing IEEE 802.3 MAC protocol, this proposed amendment will maintain maximum compatibility with the installed base of Ethernet nodes.
- 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.

“PAR B” 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.
- Projections indicate server bandwidth demand will exceed 10 Gb/s within a few years and center around the 40 Gb/s rate a number of years before 100 Gb/s is applicable for server connections.
- 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.

“PAR B” 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 from component and equipment suppliers and their customers indicate:
 - 40 Gb/s Ethernet will be optimized for computing applications that are cost sensitive, and will offer better value than alternate approaches or technologies.
- In consideration of installation costs, the project is expected to use proven and familiar media, including deployed optical fiber, backplanes, and copper cabling technology.
- Network design, installation and maintenance costs are minimized by preserving network architecture, management, and software.

Process Proposal

It is essential that a process for resolving our forward direction be considered within the HSSG in an open and transparent way. I am proposing that we determine our forward direction with public debate using the following process.

1. Motion to proceed with two PARs or Motion to proceed with a dual-PAR or Motion to proceed with any individual PAR.
2. Debate that Motion and vote.
3. If vote is successful
 - review, amend, and forward the PAR(s).Else (until four PAR options are exhausted (40G only, 100G only, two PAR, dual-PAR)
 - Motion to proceed with the alternative PAR(s).
4. If that vote is successful
 - review, amend and forward the alternative PAR(s).Else (if all four options are exhausted)
 - We have determined that no project should come from the HSSG.

*** Success should not be determined by whether we (the IEEE 802.3 HSSG) have forwarded a PAR to 802.3, but whether we are serving the industry with our actions.