Port Extension Draft PAR Proposal

PAR Request Date: 19 November 2009

PAR Approval Date:

PAR Signature Page on File: No

Type of Project: Amendment to IEEE Standard

Status: Unapproved PAR, Std 802.1Q-2008

Root Project: 802.1Q-2008

1.1 Project No:

1.2 Type of Document: Standard

1.3 Life Cycle: Full-Use

1.4 Is this document in ballot now? No

2.1 Title: IEEE Standard for Local and Metropolitan Area Networks --- Virtual Bridged Local Area Networks – Amendment
2.1 Amendment/Corrigenda Title
tbd: Bridged Port Extension

3.1 Working Group: Higher Layer LAN Protocols Working Group (C/LM/WG802.1)
Contact Information for Working Group Chair
Tony Jeffree
Email: tony@jeffree.co.uk
Phone: +44-161-973-4278
Contact Information for Working Group Vice-Chair
Paul Congdon
Email: paul.congdon@hp.com
Phone: 916-785-5753

3.2 Sponsoring Society and Committee: IEEE Computer Society/Local and Metropolitan Area Networks (C/LM)
Contact Information for Sponsor Chair
Paul Nikolich

Email: p.nikolich@ieee.org
Phone: 857.205.0050
Contact Information for Standards Representative
None

3.3 Joint Sponsor: None

4.1 Type of Ballot: Individual

4.2 Expected Date of Submission for Initial Sponsor Ballot: 12/2010

4.3 Projected Completion Date for Submittal to RevCom: 06/2011

5.1 Approximate number of people expected to work on this project: 30

5.2 Scope: This amendment specifies protocols, procedures, and managed objects to support Port Extension. A Port Extender attaches to a MAC port of an 802.1Q bridge over an 802 specified full-duplex media and provides additional MAC ports that are logically ports of the 802.1Q bridge to which it is attached (i.e. the "Controlling Bridge"). The protocols, procedures, and managed objects specified in this amendment are expected to specify new behavior in bridges that support port extension as well as the behavior of Port Extenders themselves. In addition, the protocols, procedures, and managed objects specified in this amendment support the cascading of Port Extenders. To the extent technically reasonable, all frame filtering and relay functions remain in the Controlling Bridge. Use of a STag for Multichannel capability as being defined in Remote Port and Reflective Relay Services is envisaged to achieve this objective. A new on-the-wire indication (i.e. a new tag) is envisioned to support remote replication for purposes including frame flooding and group address support.

This project utilizes the work being done in Remote Port and Reflective Relay Services.

5.3 Is the completion of this standard dependent upon the completion of another standard: Yes
If yes, please explain: Remote Port and Reflective Relay Services are intended to be utilized in this amendment.

In addition, it is expected that Port Extenders would take advantage of other standards currently under development, e.g., Priority-based flow control, Enhanced transmission selection, and Congestion Notification. However, completion of this standard is not dependent upon the completion of these other standards.

5.4 Purpose:
The purposes of this project include:

- Reduce the management cost of networks comprised of large number of bridges (such as those commonly found in a data center environments) through significant reduction in both the number of devices to be managed and the management traffic required.

- Decrease total cost of ownership by reducing initial capital expenditure along with management and operational costs.

5.5 Need for the Project: The deployment of thousands of servers in bridged local area networks is becoming common in modern data centers and numbers in the 10's of thousands are starting to appear. This is becoming increasingly feasible and cost effective with the advent of "1U" servers, blade server systems, multicore CPUs, and other high density server technologies. Naturally, this results in a corresponding growth in the number of bridges deployed.

In addition, server virtualization technology has greatly exacerbated this effect. In such a virtualized environment, nearly all physical servers (or blades) contain some form of a bridge (possibly implemented in software as part of the hypervisor or as a physical device such as a PCI adapter card).

With or without server virtualization, this has resulted in a dramatic increase in the capital expenditure costs as well as the operational costs in such networks. In addition, the management overhead of managing the sheer number of bridges, both in terms of points of management and management traffic, has negatively impacted the efficiency with which such environments may be operated. Finally, the embedding of bridges within servers has blurred the traditional bounders of administrative domains complicating the organizational effectiveness of information technology organizations.

At a high level, this project proposes to address these needs by through the definition of a Port Extender device that may be used to replace a large number of these bridges resulting in significantly reduced capital expenditure, points of management, and management traffic and thus reducing total cost of ownership.

5.6 Stakeholders for the Standard: Developers and users of networking for data center environments including networking IC developers, switch and NIC vendors, and users.

Intellectual Property

6.1.a. Has the IEEE-SA policy on intellectual property been presented to those responsible for preparing/submitting this PAR prior to the PAR submittal to the IEEE-SA Standards Board? Yes
If yes, state date: tbd
6.1.b. Is the Sponsor aware of any copyright permissions needed for this project? No
6.1.c. Is the Sponsor aware of possible registration activity related to this project? Yes, an Ethertype is needed for the new tag used for remote replication.

7.1 Are there other standards or projects with a similar scope? No

7.2 Future Adoptions
Is there potential for this standard (in part or in whole) to be adopted by another national, regional, or international organization? No

7.3 Will this project result in any health, safety, security, or environmental guidance that affects or applies to human health or safety? No

7.4 Additional Explanatory Notes: (Item Number and Explanation) none

Five Criteria

1. **Broad Market Potential**
    a. **Broad sets of applicability**
    Data centers containing hundreds or thousands of deployed bridges are common. These include data centers that have deployed high density server solutions including "1U" servers, server blade racks, etc. Deployments such as these are expected to significantly benefit from the technologies proposed. Additionally, data centers that have deployed server virtualization technology are expected to enjoy even greater benefits.
    b. **Multiple vendors and numerous users**
    There has been interest expressed by multiple vendors in this technology. In addition, many vendors have announced products supporting similar technology in a proprietary fashion. This technology is applicable to bridge, NIC, server, and software vendors. Given the wide deployment of networks that would benefit from this technology, numerous users may clearly be expected.
    c. **Balanced costs (LAN versus attached stations)**
    This technology has been expressly designed for balanced costs. It is deployable with no change to existing attached stations (that is, the technology interoperates with existing NIC cards). The design of the Port Extender function has been carefully considered to keep costs constrained. This has been a high priority since it is expected that Port Extenders may well outnumber bridges in typical deployments and are likely to be integrated in with attached stations.

2. **Compatibility**
    The combination of Port Extenders and their Controlling Bridge result in an 802.1Q bridge, thus compatibility with external devices is assured. In particular, such a combination will fully interoperate with neighbor bridges (whether embedded in stations or external), as well as existing NIC cards. Finally, this technology will assume full benefit of other Data Center Bridging technologies under development including Priority-based flow control, Enhanced Transmission Selection, and Congestion Notification.

3. **Distinct Identity**
    a. **Substantially different from other IEEE 802 standards**
    IEEE Std 802.1Q is the authoritative specification for Bridges. No other IEEE 802 standard addresses port extension by bridges.
    b. **One unique solution per problem (not two solutions to a problem)**
    The need to provide port extension has not been anticipated by any other standard. Consequently, this is the only solution to this problem. Importantly, this proposal address the needs produced by

both external and embedded bridge devices along with server virtualization with a common solution thereby eliminating the need for an additional solution in the future.

c. **Easy for the document reader to select the relevant specification**
IEEE Std 802.1Q is the natural reference for port extension of 802.1Q bridges.

4. **Technical Feasibility**
   a. **Demonstrated system feasibility**
   Similar techniques have been deployed as proprietary enhancements to 802.1Q bridging and are supported by multiple vendors. In additions, roughly analogous techniques have been deployed in Fibre Channel that have been widely adopted. These deployments have shown that the technology proposed is feasible.
   b. **Proven technology, reasonable testing**
   This technology has been proven on an operational basis in data centers using proprietary implementations. The resulting behavior remains that of an 802.1Q bridge thus existing testing methodologies remain applicable. The on-the-wire indication of ingress / egress port numbers is intuitively reasonable to test and has been shown to be such in the existing proprietary implementations.
   c. **Confidence in reliability**
   The overall behavior is that of an 802.1Q bridge; the reliability of such has been firmly established. Furthermore, the simplicity of the Port Extenders compared to that of the bridges they replace, along with the associated reductions in management complexity, is expected to yield increase in reliably over that achievable today.
   d. **Coexistence of 802 wireless standards specifying devices for unlicensed operation**
   Not applicable.

5. **Economic Feasibility**
   a. **Known cost factors, reliable data**
   Port Extenders are expected to cost significantly less than existing bridges. This is both intuitive and supported by experience in existing deployments of this technology. In addition, the resultant reduction in management complexity brings significant cost advantages.
   b. **Reasonable cost for performance**
   The proposed technology reduces overall system cost while maintaining existing performance (both in raw bandwidth and feature / functionality) for a wide variety of deployments thus cost for performance is benefited.

c. Consideration of installation costs

Due to the simplicity of the Port Extender device, initial capital expenditure and initial configuration costs are expected to be significantly reduced.