

Interspersion of Express Traffic

Technical Proposal for a Wire Compliant Solution

Thursday, May 16, 2012

Geoff Thompson
GraCaSI S.A.

GraCaSI

STANDARDS ADVISORS

Geoffrey O. Thompson

PRINCIPAL

**158 PASEO COURT
MOUNTAIN VIEW, CA 94043-5286
USA**

PHONE: +1.540.227.0059

E-MAIL: <THOMPSON@IEEE.ORG>

Design Rules:

GraCaSI

Thompson Proposal Design Rules for” "Interspersing Express Traffic" (Proposed for consensus w/i DMLT SG)

1. No padding in any fragments produced by our process
2. All packets on the wire must be "legal" Ethernet packets
 - a. Max Length not an issue here
 - b. Segmentation packets will have locally valid CRC
 - c. Segmentation packets will not have padding added
3. All packets entering the Bridge Relay Sub-Layer must be "legal" Ethernet packets.

1 & 2c Result in poorer performance (i.e. higher worst case latency)
therefore are open for discussion.

**THIS AREA DESERVES SIGNIFICANT DISCUSSION
HIGH IMPACT ON BEST POSSIBLE PERFORMANCE**

Proposal (1):

GraCaSI

Method to "distinguish traffic" at transmit end will be two transmit bridge ports

- One for Normal traffic
- One for Distinguished/Express traffic

Method to "distinguish traffic" at receiver

- All "normal" traffic will have:
 - VLAN method TAG and
 - IET EtherType in that TAG

Express Traffic requires EtherType (No LLC)

- EtherType \neq IET type value

Prefer to have TAG added at transmit Bridge Relay

Proposal (2):

GraCaSI

All "Express Traffic" will:

- NOT have IET EtherType

All "Non-Express Traffic" (unsegmented) will:

- Have TAG with IET-1 EtherType

All "Non-Express Traffic" segments
(i.e. "following segments")

- Will have TAG with IET-2 EtherType

IET-1 TAG added at transmit Bridge Relay

IET-2 TAG inserted at segmentation sub-layer

Proposal (3):

GraCaSI

RECEIVE SIDE

- Only 1 receive MAC/Bridge Port necessary, could be 2 for symmetry.
- Major design decision (& cross layer issue)
 - Where to do reassembly of Segmented Packets?

See next slide

Proposal (4):

GraCaSI

Reassembly of Received Segmented Packets

A. Reassemble below receiving MAC

- Makes process transparent at receiving MAC and its service interface
- Makes reassembled packet subject to fragmentation "down the line"
- Breaks layer model by making this a packet store & fwd layer

B. Reassemble above the receiving MAC

- In a new (Dot 3?) shim layer between MAC and Bridge Layer
- Makes process transparent at Bridge Layer interface
- Makes reassembled packet subject to fragmentation "down the line"
- Breaks layer model due to asymmetry

C. Reassemble in receiving Bridge Layer (or not)

- Doesn't have to be a port specific process
- More suitable for a software process
- EnTAG/DeTAG already a Bridge process
- Segmentation packets can be forwarded "as-is" through other bridges
 - Improves Express Traffic performance through rest of network
- Reassembly takes place (last resort) at bridge with end station port (or possibly even w/i "bridge process" in end station)
- Minimizes new port hardware required by std & implementation

Proposal (5):

GraCaSI

TAG Requirements:

- Express Traffic from Bridge to MAC Transmit Port
 - NO TAGGING (No length or processing overhead to slow it)
 - Any EtherType value except IET
- Normal Unsegmented Traffic from Bridge to MAC Transmit Port
 - VLAN Type TAG
 - Identifying EtherType/Subtype (IET-1)
 - Length (for managing segmentation/reassembly)
 - If “Length” = Physical Length then it hasn't been broken up
 - Other information (TBD, should reserve some space)
- Chop Operation (i.e. Interspersion Request = Segment Operation)
 - Terminate IET-1 segment with valid CRC, Length > 46
 - VLAN Type TAG inserted in front of new segment
 - Identifying EtherType/Subtype (IET-2)
 - Local CRC flag (Redundant, implied w/ IET-2 EtherType)
 - (Says this CRC is not the CRC of the reassembled packet)
 - Length (for managing segmentation/reassembly)
 - Countdown value from IET-1 Length Value
 - Value must be > 46 (See slide 3)
 - Other information (TBD, should reserve some space)

Features (1):

GraCaSI

- DML/E Traffic is “Distinguished” by virtue of not having a VLAN Tag of known Type value.
- All tagging/untagging overhead penalties are pushed over to the best effort traffic stream.
- Any Type value except IET is OK for Express Traffic stream.

Features (2):

GraCaSI

- No changes to the MAC
- Packets on wire are all legal Ethernet
- Minimal changes to packet analyzers
 - 1 or 2 new types (1 w/ 2 sub-types)
 - New TAG format w/ small # of simple fields
- No changes whatsoever to Express Traffic packets

Features (3):

GraCaSI

- No changes to the MAC
- Packets on wire are all legal Ethernet
- Minimal changes to packet analyzers
 - 1 or 2 new types (1 w/ 2 sub-types)
 - New TAG format w/ small # of simple fields
- No changes whatsoever to Express Traffic packets

Features (4):

GraCaSI

- No changes to the MAC
- Packets on wire are all legal Ethernet
- Minimal changes to packet analyzers
 - 1 or 2 new types (1 w/ 2 sub-types)
 - New TAG format w/ small # of simple fields
- No changes whatsoever to Express Traffic packets

Features (5):

GraCaSI

- Insertion of standard TPMR is completely benign
- Insertion of normal bridge into network only breaks Express Traffic aspect
- Reach of segmentation is programmable w/i bridge set (w/ bridge layer reassembly)
- No changes whatsoever to Express Traffic packets

Features (6):

GraCaSI

- Minimizes time to segment for express transmit
 - Calculate by LENGTH field contents
 - No waiting to see length of normal packet
 - No tagging/reframing delays for express traffic
 - All overhead pushed to lower priority frames

Further Work:

GraCaSI

- Presentation to date is a half baked scheme
- Needs more work (TAG format, State machines)
- Fairly complete as single hop solution
 - Allowing padding in segments would improve latency.
- Allowing multi-hop and multi-site fragmentation will add complexity.
- A multi-hop/keep things segmented across the network should improve overall network performance (i.e. more ATM like)

GraCaSI

Discussion

GraCaSI

STANDARDS ADVISORS

Geoffrey O. Thompson

PRINCIPAL

**158 PASEO COURT
MOUNTAIN VIEW, CA 94043-5286
USA**

PHONE: +1.540.227.0059

E-MAIL: <THOMPSON@IEEE.ORG>