


802.1aj Two port MAC Relay status



Two port MAC Relay



- Industry recognises that full 802.1 bridges are sometimes unnecessarily complex
- TPMR (802.1aj) attempts to provide a simpler relay function than a VLAN bridge
- PAR granted December 2004
- Drafts:
 - Initial draft 0.0 May 2005
 - Draft 1.0 July 2005
 - Draft 1.1 August 2005
 - Draft 1.2 November 2005
- This presentation represents the results of the Jan 2006 interim meeting, and a new draft is expected before March 2006

TPMR topics



- Zero configuration option
 - Should work out-of-the-box
- Topologies
- Differences from a bridge
- Link maintenance
- Discovery
- Management
- Forwarding
- MAC types
- Loopback
- Link status propagation

Topologies



- A TPMR can be deployed singly, or in a chain
- A typical application might be as a demarcation device
- A TPMR has exactly two ports
 - Each port can be Ethernet or any MAC or emulated MAC which supports the 802.1 Internal Sublayer Service
 - (implies that the device is not VLAN-aware)
- Protection is not supported in the standard
- Management is from a more intelligent device which proxies the TPMR's managed objects into its own MIB

Differences from a Bridge



- Only two ports
- No MAC address learning
- No VLAN tagging
- No Spanning Tree
 - BPDUs require special treatment (see later slide)
- Mandatory extra features

Link maintenance



- For Ethernet links, 802.3ah EFM OAM may be employed
 - This provides an indication of link up/down
 - Ethernet MAC link down indication is notoriously unreliable
- 802.3ah also provides
 - Link status change information
 - Link statistics including errored seconds etc.
 - Managed object access, which is NOT used in TPMR
- E-LMI (MEF UNI Phase 2) was considered, but is not suited as a link maintenance protocol
 - Intended for CE to retrieve status and service attributes from the network
 - Includes UNI and per-EVC configuration and status information
- Other MACs and emulated MACs can use their own protocol

Discovery




- A mechanism is required to allow discovery of TPMRs, so that the managing device knows what to manage
- Mandatory CFM (802.1ag) is the primary discovery method
 - At least a level 0 MIP is required in TPMR
 - Attached bridge or station can use Linktrace to discover connectivity of attached TPMR chain
 - All TPMRs in a chain can be found, but a method is needed to know when the end of the chain has been reached
 - CFM tells you what kind of device it is (but requires a data point for TPMR)
 - CFM tells you which port number you are connected to

Discovery



- 802.1ag CFM Linktrace is used to find TPMRs to manage



- ❖ Discovery may terminate at UNI or might discover customer equipment
 - ❖ Consider placement of MEP in demarc TPMR?

Discovery



- LLDP (802.1ab) may be used for further probing
 - LLDP support is optional
 - CFM already gives chassis type
 - What additional benefits would this offer?
 - Potential conflict between scope of LLDP and scope of TPMR discovery?
 - Part of the management approach is to hide TPMRs from network management as individual network elements
- Ethernet EFM OAM (802.3ah) could have been chosen for Ethernet links, but is harder to use for chain discovery
 - Hop-by-hop approach

Management



- SNMP over Ethernet, without IP, has been chosen
 - SNMP over IP was rejected because of the desire to avoid IP address management and NMS interaction with individual TPMRs
 - 802.3ah EFM/OAM was rejected because of concerns over scalability to a chain and lack of “Set” capability
 - CORBA was considered too much of a stretch given that nothing else in 802.1 uses it
- There is an Ethertype for SNMP
 - Untagged frames are used
 - Management VLAN option not yet discussed much
- Initial managed object set needs expanding
 - Should include EFM MIB, Interface MIB, extra stuff

Management continued



- Discovery is used first to find what to manage
- Retrieved objects are incorporated into the managing device's MIB
 - Details undecided, but Interface MIB objects are a candidate
- Which ports can be used to manage the device?
 - A management block will be provided to prevent access from the customer port
 - Is this locally provisioned?
 - Is access provided by authentication, to allow a device which is installed the wrong way round to be "recovered" remotely?
 - In a device with different port types, which port then?
 - Suggest that selection of which ports are active for management be left as an exercise for the reader

Forwarding



- General idea is to be transparent to protocols the TPMR does not implement
 - But some protocols are filtered out by the MAC, e.g. Pause
 - Transparent to BPDUs
 - Transparent to LACP (despite the layering violation)
 - Treatment of other reserved addresses needs defining
- No modification of user data frames (e.g., tagging)
- Multiple queues are optional
 - Extract priority from Q-tag and 802.1ae LinkSec tag
 - Only for integrity-protected frames – unencrypted
 - Recognise L2 control protocols and place in fastest queue
 - Typically BPDUs
- Otherwise like 802.1d/Q
 - Note that MRP (802.1ak) needs special handling in a Q-bridge

Loopback

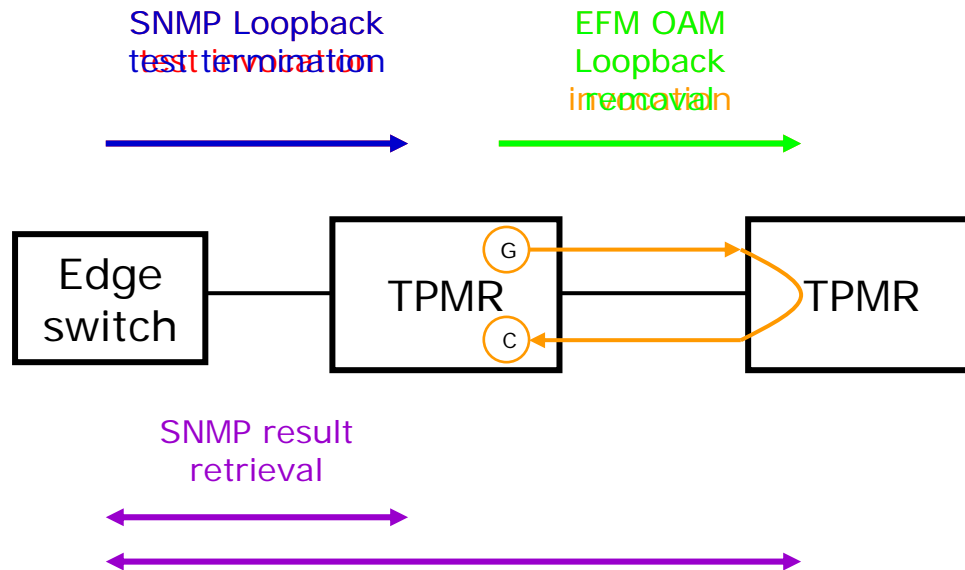


- Optional per-link loopback on Ethernet with 802.3ah EFM/OAM
 - Invoked by SNMP to previous hop
- Multi-hop CFM-based (802.1ag) loopback
 - Uses a special loopback frame
 - Can contain arbitrary data inside a TLV
 - Non-intrusive, in that user data continues
 - Stateless (no invocation or termination command required)
 - Limitation: TPMR is not VLAN-aware, so uses untagged CFM only
- Stateful per-VLAN loopback is not supported
 - CFM group rejected this idea as not sufficiently useful at resolving data-driven errors – still controversial
 - Could be provided using an EFM/OAM extension invoked from the previous hop using SNMP



Optional per-link loopback

- Per-link loopback on Ethernet with 802.3ah EFM/OAM
 - Invoked by SNMP to previous hop
 - Beware that EFM OAM loopback discards returned frames
 - Frame generator (G) and checker (C) needed to support this





Link status propagation

- The draft proposes classic link-loss forwarding where one port is disabled when failure is detected on the other port
- This has problems regarding management reachability in the case of failure of the customer-side port
- A proposal is needed for the standard
 - It needs a signalling-based method
 - Could Y.1731 AIS be used as part of this?

Open issues



- Interworking of Discovery with LLDP
- Managed object set
- Management VLAN option
- Details of forwarding behaviour for reserved layer-2 control protocol addresses
- Link status propagation

Turning
metro visions
into value



Thank you

JMessenger@advaoptical.com