



# Maintenance and Interpretation requests to 802.1

July 06, 2012

Presented by  
Aaron Stewart

University of New Hampshire's  
InterOperability Laboratory  
(UNH-IOL)





## Overview

---



- MSRP's MAP definition
- Non-'new' declarations in MSRP
- MAP Context for MSRP
- Use of 'new' MSRP
- MVRP's use of new
- Interpretation of "the Port"
- Do blocking ports transmit MRPDUs
- Use of Flush! in Registrar SM
- MVRP & Static VLAN Reg Entries
- 'taggedness' of MSRPDUs in MST
- Discarding of Applicant and Registrar
- Bridge Port in FailureInformation



## MSRP's MAP definition



MSRP's MAP function is not clearly defined.

5.4.4 indicates that 10.3.1 specifies it

- d) Propagate registration information in accordance with the operation of MAP for the Base Spanning Tree Context, as specified in 10.3.1.
- e) Forward, filter or discard MAC frames carrying any MRP Application address as the destination.

10.3 indicates that 35.2.4 specifies it

function. This subclause specifies the operation of the MAP function for the MMRP application and the MVRP application (11.2.1). The MAP function for the MSRP application is specified in 35.2.4.

35.2.4 indicates that 10.3 is not used

For the Talker and Listener attributes MSRP propagates attributes in a manner different from that described in 10.3 for MMRP and MVRP. In principle, the MAP performs MSRP Attribute Propagation when any of

However, without 10.3, MSRP is lacking some details of attribute propagation, namely, how and when new declarations are propagated.

Discussions with Craig Gunther suggest 35.2.4 is meant to augment the MAP function defined in 10.3 more than replace it, but whether that is the case and the extent to which that is the case is not clear.



- 35.2.4 indicates a number of conditions under which the MSRP Attribute Propagation procedure occurs
  - Receipt of MAD\_Join.indication in which *new* is TRUE is listed
    - a) A MAD\_Join.indication adds a new attribute to MAD (with the *new* parameter, 10.2, set to TRUE);
    - b) A MAD\_Leave.indication is issued by the MAD.
  - However, no mention is made of MAD\_Join.indications in which *new* is FALSE.
- This seems to be an oversight – given that **new** isn’t used by MSRP, MSRP shouldn’t particularly care whether it was declared with a *New* or with a *JoinMT*.
- If ‘Join’ does not cause propagation (as is currently the case in MSRP, undesirable as that may be) then a received Join will cause a Port to Register the attribute but not propagate it.
- The standard should make it clear what action to take when non-new MAD\_Join.indications are received.
  - We presume that non-new joins are propagated



- 35.2.4.5 indicates that
  - “MSRPDU s can carry information about Streams in multiple VLANs, which in an MST environment, can be in different Spanning Tree Instances. ... Therefore there is a single context for MSRP attribute propagation that includes all Bridge Ports. The Declarations are filtered according to the state of the spanning tree, as described in 35.2.4.
- **What does this mean:**
  - “Therefore there is a single context for MSRP attribute propagation that includes all Bridge Ports.”
  - Does this mean the active topology is all Bridge Ports (and hence, MSRPDU s (Domain messages and LeaveAlls) would be transmitted on all ports regardless of whether they are blocking?
- Partial fix: Change “The Declarations are filtered according to the requirements of 35.2.4”
- To: “The Declarations are filtered according to the requirements of 35.2.4 and its subclauses and according to the state of the spanning tree per 35.1.3.1. “



# Use of 'new' MSRP



The only mention of **new** in MSRP is in 35.2.4

For the Talker and Listener attributes MSRP propagates attributes in a manner different from that described in 10.3 for MMRP and MVRP. In principle, the MAP performs MSRP Attribute Propagation when any of the following conditions occur:

- a) A MAD\_Join.indication adds a new attribute to MAD (with the *new* parameter, 10.2, set to TRUE);
- b) A MAD\_Leave.indication is issued by the MAD;
- c) An internal application declaration or withdrawal is made in a station;
- d) When the bandwidth of the underlying media changes (see *bandwidthAvailabilityChanged* notification in 34.3.2);
- e) A port becomes an SRP domain core port (3.177);
- f) If talkerPruning (35.2.4.3.1) is enabled and there is a change in the MMRP (10.9) attributes registered on a port.

If MSRP does not intend to make use of **new**, it may be advantageous to explicitly state this, as is done in the case of MMRP (10.12.3).

## 10.12.3 Use of "new" declaration capability

MMRP does not make use of the "new" declaration capability.

If this were done, then **new** would not cause any action (no 'flush' of any filtering database, etc)

Presumably 'new' indications would propagate as 'new' requests.

(per 10.1 last paragraph: "The rules applied to the marking and propagation of newly declared values in this way are common to all MRP Applications; however, the action taken on receipt of an attribute declaration marked as "new" is specific to each MRP Application")



## MVRP's use of new



- 11.2.5, which describes MVRP's use of new, says that when new is received, "any entries in the filtering database for that Port and for the VID corresponding to the attribute value in the MAD\_Join primitive are removed."
- According to 8.8 The Filtering Database, the following entry types may contain information about the relevant port and VID:
  - Dynamic Filtering Entry (Contains: A MAC address, FID, and Port Map for "each outbound Port")
  - Dynamic VLAN Registration Entry (Contains: A VID, and a Port Map for "each outbound Port")
  - MAC Address Registration Entry (Contains: A MAC address, VID, and Port Map for "each outbound Port")
  - Dynamic Reservation Entry (Contains: A VID, a MAC address specification, and a Port Map for "each outbound Port")
  - Static Filtering Entry (Contains: A MAC address, VID, and Port Map for "each outbound Port")
  - Static VLAN Registration Entry (Contains: A VID, and a Port Map for "each outbound Port")
- Every type of entry contains information about some VID (or FID, in the case of Dynamic Filtering Entries), and every port.
- One interpretation of 11.2.5 (which is consistent with 11.2.5) would be, upon receiving a "new" event, to remove all of the listed entries which match the VID which the "new" event was for.
  - This is highly undesirable; in fact, this interpretation puts 11.2.5 in direct conflict with 8.8, because Static Filtering Entries and Static VLAN Registration Entries may only be removed by management. Static filtering information is added to, modified, and removed from the Filtering Database only under explicit management control. It shall not be automatically removed by any ageing mechanism. Management of static filtering information may be carried out by use of the remote management capability provided by
- The only reasonable thing to discard are Dynamic Filtering Entries.
- Beyond the type of entry, "removing an entry" (because they contain a port map), effectively unregisters all ports
  - In the case of Dynamic Filtering Entries (8.8.3), this will only deregister it from one port
- In any case, 11.2.5 should more clearly state what information will be removed when a new indication is received.



## Interpretation of “the Port”



- 10.3 For a given MRP application and MAP Context (10.3.1), and for the set of Ports that are in a Forwarding state as defined by that MAP Context:
  - a) Any MAD\_Join.indication, or any MAD\_Join.request issued by the MRP application, received by MAP from a given Port in the set is propagated as a MAD\_Join.request to the instance(s) of MAD associated with each other Port in the set. If the value of tcDetected (13.23) for the Port and MAP Context is nonzero, then the value of the *new* parameter in the propagated MAD\_Join.request is set TRUE, regardless of the value of this parameter in the indication or request that is being propagated.
- Which of the two previous uses of “Port” is “the Port” referring to?
- This should be stated more clearly.
- It has been suggested that “the Port” refers to “other Port” (the egress port).
- Though, if it refers to the ingress port, it would seem to be more consonant with the purpose of **new**.
  - If a bridge has some attribute registered on a port, and that bridge is connected to another bridge, when the ports begin forwarding, 10.3.d is invoked, causing the newly-forwarding port to transmit a declaration
  - If “the Port” is the egress port, the declaration will not be signaled as a **new** declaration (10.3.d says nothing about propagating as **new**).
  - However, if “the Port” is the ingress port, the neighboring bridge would propagate a received **JoinMt** or **JoinIn** as **New**, thus informing the bridged LAN of the topology change.





## Do blocking ports transmit MRPDUs

- According to 10.3, regarding what to do when a port is removed from the set of forwarding ports

**NOTE**—If a Port is removed from the set, and that Port has declared one or more attributes, then this Port transmits a Leave message (see 10.6) for every attribute that it has declared.

- According to 10.7.6.1, however

Unless stated otherwise in these action definitions, MRPDU transmission as a result of the operation of a state machine in a Bridge occurs only through the Port associated with that state machine, and only if that Port is in the Forwarding state. MRPDUs shall be transmitted using the destination MAC Address and EtherType value defined by the MRP application associated with the state machine.

- None of the action definitions state otherwise
- These statements are contradictory.
- They may be reconciled if 10.3 were to state that the Leave message is transmitted before the port is removed from the forwarding set, however, the current language seems to indicate that Leaves are sent as a response to having already been removed.
- Options are: send all MRPDUs (like BPDUs, LLDP, etc) at all times; send none; or, send some for all or only some of the time after a port is removed from the forwarding set.
  - Receiving a Leave will speed the failover process for a talker (vs receiving nothing and waiting for a LeaveAll timeout).
  - In 10.3, sending a LeaveAll would be faster than sending a Leave, but would trigger additional unnecessary PDUs to be transmitted in a shared environment, making such an option undesirable.
  - If all MRPDUs could be sent when a Port is not forwarding, rules such as 35.1.3.1 “Blocked Declarations” would be violated (ultimately allowing Listener Ready to be propagated on blocked ports)
  - Selectively forwarding some declarations (such as a MAD\_Leave), or only doing so for a limited time (such as implied by the NOTE in 10.3)



# Use of Flush! in Registrar SM

**REQUESTED REVISION:**

STANDARD: 802.1Q-2011  
CLAUSE NUMBER: 10.7.5.2  
CLAUSE TITLE: Flush!

**RATIONALE FOR REVISION:**

The current behavior of the Registrar state table (Table 10-4) does not cause a "Lv" (10.7.6.14) to be sent when in the "IN" state and a Flush! event occurs. This results in the permanent registration of the associated attribute, as the MRP application is never made aware of the Registrar's state change.

**PROPOSED REVISION TEXT:**

Regarding Table 10-4, state "IN", event "Flush!":  
Replace "MT" with  
"Lv  
MT"

**IMPACT ON EXISTING NETWORKS:**

The intention of the informatively named state names "IN" and "MT" is maintained when a "Lv" is signaled, as this will trigger a MAD\_Leave.indication to the MRP application.

The intentions outlined in 10.3.1.1 will similarly be maintained by this change, specifically, quoting from the end of 10.3.1.1:

""

For implementations running over RSTP or MSTP, this gives rise to the risk of information loops when Port roles change; because of the store and forward nature of attribute propagation and the potentially rapid transitions of Port roles (compared to the relatively slow transitions that occurred with STP), these can arise even when there are no data loops.

To prevent such information loops from occurring, the information held by MAD's Registrars for a Port (i.e., information registered on a Port as a result of protocol activity on the LAN to which that Port is connected) is discarded whenever the Port transitions from an Alternate port or Root Port role to become a Designated Port. No such discard is needed for changes in the other direction, i.e., changes from Designated Port to Root Port or Alternate Port.

""

**For reference:**

10.7.5.2 Flush!

A Flush! event signals to the Registrar state machine that there is a need to rapidly deregister information on the Port associated with the state machine as a result of a topology change that has occurred in the network topology that supports the propagation of MRP information. If the network topology is maintained by means of the Spanning Tree state machines, then, for the set of Registrar state machines associated with a given Port and Spanning Tree instance, this event is generated when the Port Role changes from either Root Port or Alternate Port to Designated Port.

When a Flush! event occurs for a given Port and Spanning Tree instance, a leavealltimer! event (10.7.5.22) is also signaled to the LeaveAll state machine for that Port and Spanning Tree instance.

**Table 10-4—Registrar state table**

		STATE		
		IN	LV	MT
EVENT	Begin!	MT	MT	MT
	rNew!	New IN	New Stop leavetimer IN	New IN
	rJoinIn!    rJoinMt!	IN	Stop leavetimer IN	Join IN
	rLv!    rLA!    txLA!    Re-declare!	Start leavetimer LV	-x-	-x-
	Flush!	MT	Lv MT	MT
	leavetimer!	-x-	Lv MT	MT



# MVRP & Static VLAN Reg Entries



## 11.2.1.3 Use of the PVID and VID Set

The initial state of the Permanent Database contains a Static VLAN Registration Entry for the Default PVID, in which the Port Map indicates Registration Fixed on all Ports. This ensures that in the default state, where the value of every PVID of each Port is the Default PVID and where the VID Set of each Port is empty, membership of the Default PVID is propagated across the Bridged Local Area Network to all other MVRP-aware devices. Subsequent management action may change both the Permanent Database and the Filtering Database in order to modify or remove this initial setting, and may change the PVID and/or VID Set value(s) on any Port of the Bridge.

- The highlighted sentence suggests that registrations contained in Static VLAN Registration Entries gets propagated via MVRP.
  - However, this statement seems to be intended as a descriptive statement, and not a normative statement.
- Do static, administratively-controlled VLAN registrations get propagated?
  - If so, language should be added to indicate the mechanisms involved with this.
    - For example, 10.3(MAP) does not indicate that it should be run when a static registration is added
  - If not, this statement in 11.2.1.3 is not correct.
- As it is highly likely that 10.7.2 implies that propagation occurs, rewording of this section would suffice.

## 10.7.2 Registrar Administrative Controls

Associated with each instance of the Registrar state machines are *Registrar Administrative Control* parameters. These parameters allow administrative control to be exercised over the registration state of each Attribute value, and hence, via the propagation mechanism provided by MAP, allow control to be exercised over the propagation of declarations.

- a) *Normal Registration*. The Registrar responds to incoming MRP messages as specified by Table 10-4.
- b) *Registration Fixed*. The Registrar ignores all MRP messages, and remains IN (registered).
- c) *Registration Forbidden*. The Registrar ignores all MRP messages, and remains MT (unregistered).

The default value of this parameter is *Normal Registration*.

If the value of this parameter is *Registration Fixed* or *Registration Forbidden*, In and JoinIn messages are sent rather than Empty or JoinEmpty messages.

NOTE—The Registrar Administrative Controls are realized by means of the contents of the Port Map parameters of static entries in the Filtering Database for all MRP applications. In the case of MMRP, the static entries concerned are Static Filtering Entries (8.8.1); in the case of MVRP, the static entries concerned are Static VLAN Registration Entries (8.8.2). The contents of the Port Map parameters in static entries can be modified by means of the management operations defined in Clause 12. In the absence of such control information for a given attribute, the default value “Normal Registration” is assumed.



## 'taggedness' of MSRPDUs in MST



- According to 35.2.4.5, “All MSRPDUs sent and received by MSRP Participants in SST Bridges are transmitted as untagged frames.”
- No analogous statements are made regarding MST environments.
- Are MSRPDUs in an MST environment VLAN-tagged?
- Consensus seems to be that they are not
- Change 35.2.4.5 to indicate that “All MSRPDUs sent and received by MSRP Participants in SST or MST Bridges are transmitted as untagged frames.”



## Discarding of Applicant and Registrar SMs



- Note 11 of Table 10-3 (Applicant SM) states:
  - “In implementations where dynamic creation and discarding of state machines is desirable, the state machine can be discarded when in any of these states, pending a future requirement to declare or register that attribute value”
- A similar NOTE in 10.7.8 (Registrar SM) states:
  - “As with the Applicant, state information is conceptually maintained for all possible values of all Attribute types that are defined for a given application; however, in real implementations of MRP, it is likely that the range of possible Attribute values in some applications will preclude this, and the implementation will limit the state to those Attribute values in which the Participant has an immediate interest.”
- The behavior in the LO state of the Applicant state machine when tx! occurs requires “s” followed by a transition to the VO state. By requiring “s”, the Applicant SM requires all attribute values to be sent in response to (for example) an rLA! while in the VO state.
  - This is clearly not desirable, but it is unclear when to consider the Applicant and Registrar state machines as ‘discarded’.
- Insert a new note 9 before MRP design notes to Table 10-3 (applied to the intersections of STATE columns VO, AO, QO & EVENTS “rLv! || rLA! || Re-declare!”):
  - “This state transition is ignored if responding to rLA! and the Registrar state machine associated with this attribute value is MT.”
- Insert a new note 10 before MRP design notes to Table 10-3 (applied to the intersections of STATE columns VO, AO, QO and EVENTS txLA! and txLAF!):
  - “This state transition is ignored if the Registrar state machine associated with this attribute value is MT.”



- 35.2.2.8.7 indicates that
  - “At the point when a Talker Advertise Declaration is transformed into a Talker Failed Declaration, the Bridge making the transformation adds information that indicates, to the Listeners registering the Talker Failed Declaration, the cause of the failure, and the identity of the Bridge **and Bridge Port** at which the failure occurred.”
- No information is conveyed identifying the Bridge Port.
- Strike reference to Bridge Port