PROPOSED USE OF GARP FOR DISTRIBUTION OF VLAN MEMBERSHIP INFORMATION

Tony Jeffree
Independent consultant
11A Poplar Grove
Sale, Cheshire
M33 3AX, UK

+44-161-282-3824 (Tel)
+44-973-6534 (Fax)
100271.522@compuserve.com

October 28, 1996

1. Introduction

In the main text of my ballot comments on 802.1P, entitled “Proposed revision of GARP Architecture and State Machines” [1], I alluded to the possibility of using the revised architecture as the basis for describing a further application of the GARP Information Transport (GIT), namely as a vehicle for registering and distributing VLAN membership information. This document describes the rationale for such a mechanism and the basis for achieving it using the machinery available in GIT.

The terminology used in this paper is consistent with the terminology introduced in [1].

2. Rationale

The VLAN environment that we agreed to standardise following the Ottawa interim meeting consists of a single Spanning Tree (the legacy Spanning Tree described in 802.1D), used to support Port-based VLANS. In this environment, each VLAN-aware ingress Port will:

a) Tag any untagged incoming frames (other than BPDUs or GARP PDUs) with a specific VLAN ID, defaulting to VLAN 1 if no management intervention has taken place;
b) Accept any incoming frames that are already tagged, without re-tagging them;
c) Add a tag (as in a)) to any incoming frames that have the “priority only” version of the VLAN header, i.e., VID=0.

Frames are then forwarded by the Forwarding Process in accordance with:

d) Address-related information in the Filtering Database (static, dynamic and Group entries);
e) Knowledge of which VLANs are reachable via the set of possible egress Ports.

At the present time, there is no agreed means of establishing the latter (point e)) information, either statically or dynamically; there is, however, a degree of tacit agreement that it will at least be possible to configure this information by some (management-related?) means.

Clearly, an environment in which manual configuration is required in order to establish, for each egress Port in the Bridged LAN, the set of VLANs reachable from that Port, is not terribly interesting, especially if there is a simple “plug-and-play” solution to be had. Such a solution will potentially allow both Bridges and end stations to declare their VLAN membership(s) and for that information to be used to configure the forwarding behaviour of all Bridges.

3. Proposed solution

The problem (and its solution) is in essence no different to that of handling Group membership, or distributing Port state information, as already specified in P802.1p, and as further discussed in my ballot comments.

The core of the solution is to introduce a third class of GARP Information Declaration (the first two being Groups and Port Filtering Modes); the new class allows declaration of (or registration for) VLAN membership. This new class of declaration is achieved via the following service primitives:

a) ReqJoin <VID>. Allows GAP (or GIP) to declare to GID that it is a member of VLAN <VID>;
b) ReqLeave <VID>. Allows GAP (or GIP) to declare to GID that it is no longer a member of VLAN <VID>;
c) IndJoin <VID>. Allows GID to indicate to GAP (and GIP) that it has registered a member of VLAN <VID>;

d) IndLeave <VID>. Allows GID to indicate to GAP (and GIP) that it has no longer any registered members of VLAN <VID>.

The implications of this new class of declaration are:

e) Instances of the GARP state machine can exist which represent membership of given VLANs. These instances can be subjected to the same set of management controls (Req.X YES/NO...etc) as are described in [1];
f) In addition to representing Group and Port Filtering Mode membership declarations, GARP PDUs will need to be able to represent VLAN membership declarations. This can easily be achieved by using an appropriate (extensible) coding scheme. (Such a scheme would easily allow further declaration types to be invented in the future, if need be);
g) Representing the VLAN membership information in a form usable by the Forwarding Process could be done in the Filtering Database. This would involve a further FDB entry type, the VLAN Registration Entry, closely similar in structure and function to the existing Group Registration Entry. In place of the Group MAC address, this entry type would carry a VLAN ID. These FDB entries would effectively define the filtering (egress) behaviour on all Ports of a Bridge with respect to VLAN membership. <<Note that additional information would need to be held - here or elsewhere - to define whether frames destined for a given VLAN should be tagged or untagged when transmitted through a given Port>>.

3.1 Behaviour of End Stations

VLAN-aware end stations (i.e., end stations capable of issuing VLAN-tagged frames) would issue GARP Joins and Leaves, and participate in GARP protocol activity, as appropriate for the set of VLANs of which they are currently members. As with the current Group membership application of GARP, servers/routers would use incoming VLAN membership information (from other devices on the same LAN segment) to allow them to “prune back to source” any traffic destined for VLANs that currently have no members in the Bridged LAN.

VLAN-unaware end stations would participate in VLANs only as a result of the default Port-based ingress (tagging) and egress (untagging) rule applied by adjacent VLAN aware Bridge Ports.

3.2 Behaviour of Bridges

VLAN-aware Bridges would register and propagate VLAN memberships in a closely similar manner to the current Group memberships.

The GARP state machine representing the default VLAN for a given Port (i.e., the state machine corresponding to the VID that is applied on ingress to untagged frames) behaves as if there is always a member of that VLAN present on the LAN segment attached to the Port. In other words, the Ind.X management control for that VLAN ID is set to YES. This ensures that, even if no end stations register VLAN membership, the membership(s) implied by the default tagging rule on inbound Ports will get propagated across the Bridged LAN.

IndJoin and IndLeave primitives issued by instances of GID are used to update the Port sets in VLAN registration entries in the Filtering Database; hence, the incoming registration/deregistration information received by the Bridge results in automatic configuration of the appropriate egress behaviour on each Port.

The GARP Information Propagation (GIP) rule operates as follows, for all instances of GID associated with all Ports of the Bridge that are Spanning Tree connected:

a) All IndJoin/IndLeave primitives issued by an instance of GID are submitted to GIP;

b) An IndJoin is propagated as a ReqJoin to all other instances of GID;
c) An IndLeave for Group G is propagated as a ReqLeave to a given Port if there are no longer any members of G on the other Ports of the Bridge.

<<Note, that this is a corrected version of the propagation rule as stated for Group propagation in [1]; however, both here and in [1], propagation should only take place between Ports that are Spanning Tree connected - this is not made clear in the description of the rule in [1].>>

4. Some conclusions

The approach described above for distributing/configuring VLAN egress information appears to have a number of advantages, including the following:

a) The manual configuration required in Bridges is limited to configuring the VIDs for default Port-based tagging on ingress, plus any hard wiring that may be deemed desirable by the administrator. In an installation where all end-stations are VLAN aware, this type of configuration may be completely unnecessary. By contrast, the manual configuration that would otherwise be necessary is a potentially difficult and time consuming task;
b) The approach handles the existence of VLAN aware end stations, and how Bridges become aware of the VLANs in which they participate, as well as allowing the Bridges themselves to advertise the VLAN participation of their own Ports;
c) The approach neatly re-uses technology that we have already developed in 802.1p; it can therefore potentially be incorpaorated into 802.1Q with minimal perturbation to its progress.