

Minimizing .1Q

Norman Finn

1.0 What's needed

Assumptions:

- End-stations cannot be trusted to assign themselves to VLANs.
- 100 switches are easier to manage than 10,000 endstations.

Critical issues for interoperability:

1. FDB model. Issues must be resolved to settle interoperability conflicts between models. Different usages make different requirements on which model is required. Stuart Soloway has some suggestions for compromise in this area. We could also chose just one model.
2. PVID-setting. There must be a standard way to specify to which VLAN an untagged frame is assigned. **(done)**
3. Egress list. There must be a standard way to specify the list of VLANs allowed to exit this port untagged. **(done)**
4. Allowed VLANs. There must be a standard way to specify which VLANs are allowed to pass through (or equivalently, are restricted from passing through) this port. And there must exist such a list in the model. Managing GVRP, without using its full functionality, is a clumsy to accomplish this.

Nice to settle, but expensive in time:

5. VLAN ID vs. VLAN name and/or semantics distribution.
6. MIBs for defined controls.
7. MIBs to support gathering VLAN-related statistics.
8. MIBs which account for VLAN's effects on the filtering database.
9. VLAN pruning between switches.
10. Topology discovery.
11. Much-discussed extensions (e.g. multiple spanning trees, per-protocol VLANs, load sharing).

2.0 How to do it

1. FDB Model: in progress.
2. Drop GVRP.
3. Management of allowed VLANs: A breakout session could list the variables.