A Layer 2 Gateway Port Mechanism

François Tallet

Garden Grove 802.1 IEEE Interim, September 2005

Problem definition

Cisco.com

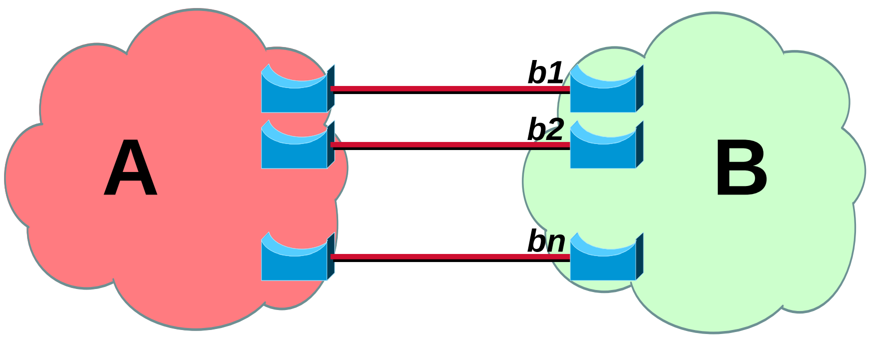

- **A & B are two independent bridged domains (not running any STP the one with the other)**
- **B is redundantly connected to A, via ports *b1*, *b2*… *bn***
- **B wants to elect a unique gateway *bi* to A, while keeping the others ports in standby for redundancy**
- **There must never be a bridging loop, even temporarily**

Run B's STP over A

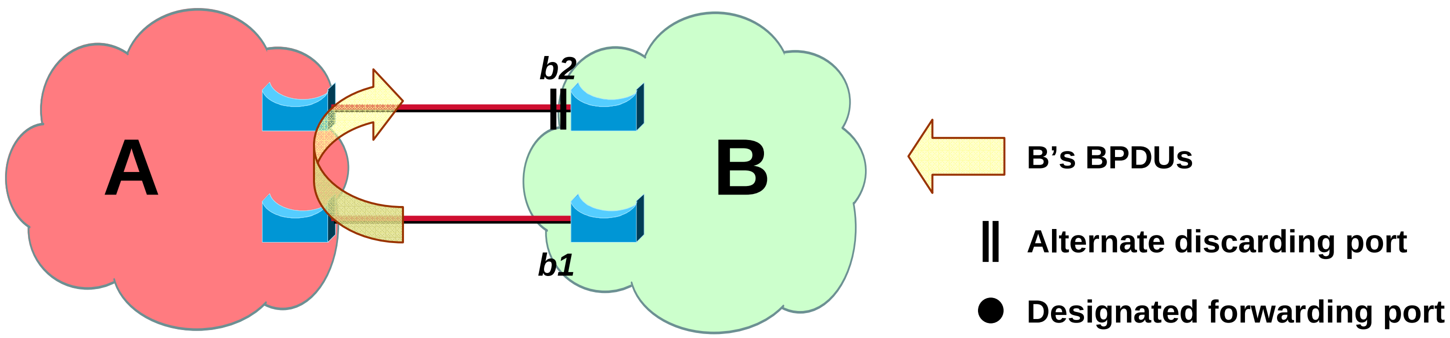

- **A very common solution consists in running B's STP over domain A:**
  B's BPDUs are considered as data traffic by A

Drawbacks:

- **Not necessarily obvious to get *b1* or *b2* to block**
- **B's stability depends on A (what if A drops B's BPDUs?)**
- **B's STP has to be "slower" than A's in order to prevent temporary loops when A reconverges.**
  It is not even enough…

Run B's STP over A: transient loop

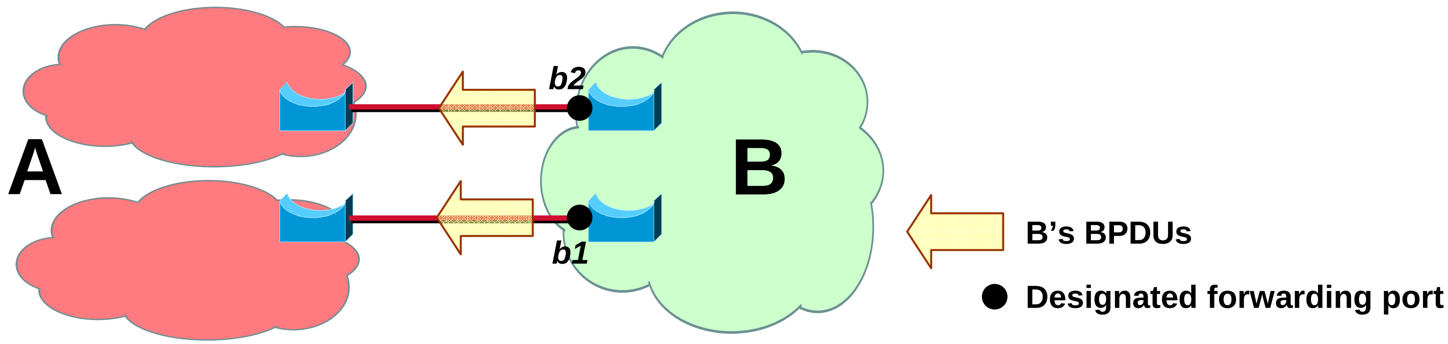

- **If A gets partitioned, both *b1* and *b2* will end up designated forwarding**
- **When A heals, there will be a temporary bridging loop until *b1* receives a BPDU from *b2***
  This may last up to STP B's hello-time (if domain A is intelligent enough to prioritize B's BPDUs.) Here, we would rather want B's STP to be fast…

A hello protocol

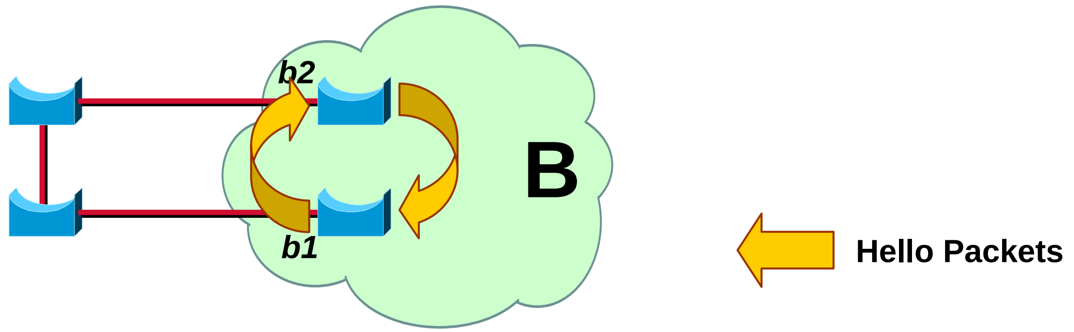

- **Gateway ports start from a discarding state and exchange hellos in order to elect the designated one**

Advantage:

- **Does not rely on region A at all**

Drawback:

- **Suffers from the same potential transient loop issue as the previous solution**

A hello protocol: transient loop

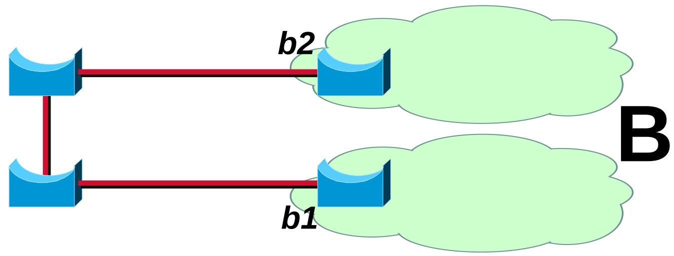

The two parts of B are connected while *b1* and *b2* are forwarding: there is a loop until the hellos from *b1* reach *b2*

- In fact, this is the same solution as “STP B on the top of domain A”. This is: “hello protocol on the top domain B”
- If B is partitioned, the same transient loop can occur when it heals

The hello protocol would need to be understood by each and every bridge of B to prevent this issue

Layer 2 Gateway Ports

Cisco.com

- **The Layer 2 Gateway Port solution consists in using the STP of domain B as this hello protocol:**
- **We define a Layer 2 Gateway Port (L2GP) as a regular port handled by STP, having the two following additional properties:**
  1. **As long as it is up, an L2GP pretends it is continuously receiving BPDUs, the root bridge ID of which is configured by the user**
  2. **An L2GP with a designated role is always discarding**

L2GP Example

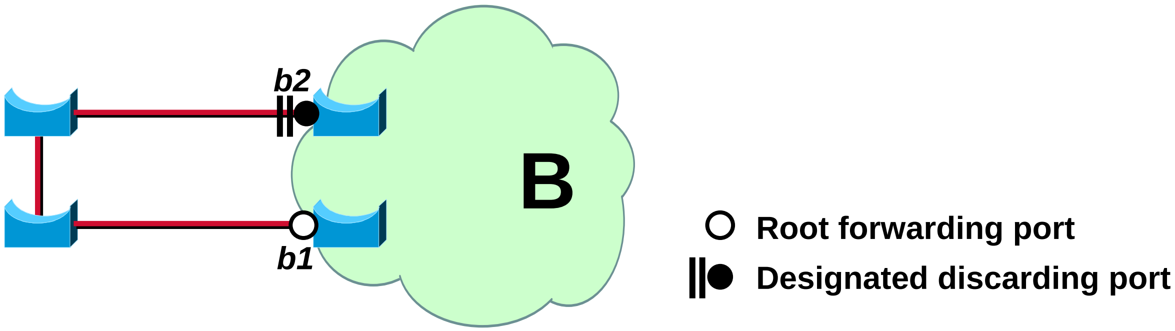

- ***b1*** **and** ***b2*** **are L2GPs configured with a fake bridge ID R1 and R2 respectively**
- **R1 is better than R2, and R2 is better than any bridge in domain B**
- **As a result:**
  - ❑ ***b1*** **is a root port, forwarding (because its bridge “receives” its best BPDU on b1)**
  - ❑ ***b2*** **is a designated port, discarding (rule #2)**

L2GP Example: domain healing with no loop

Cisco.com

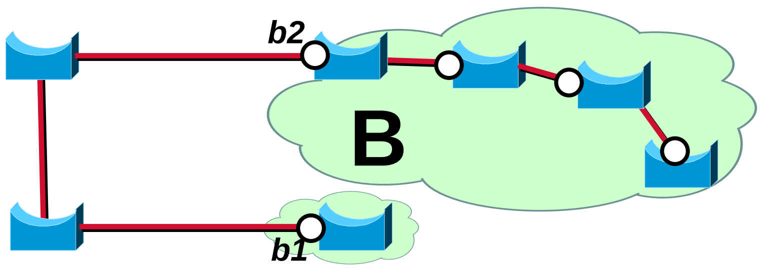

Domain B is split, *b1* and *b2* are root forwarding. A link is then brought up between the two parts.

STP sync

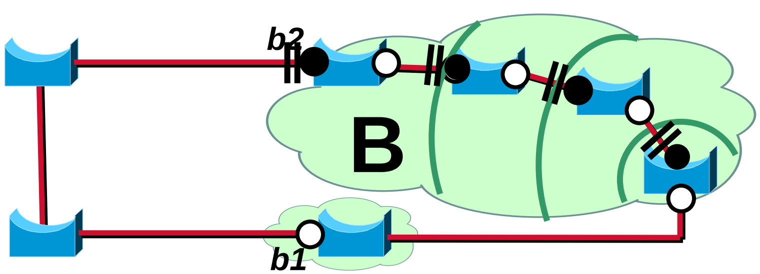

The STP sync mechanism prevents any temporary loop: the designated discarding port "travels toward" *b2*.

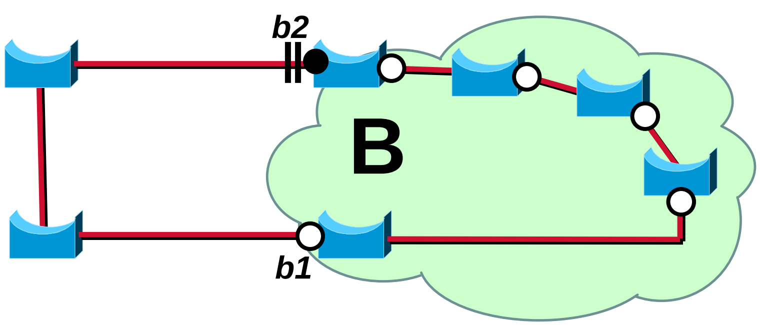

When the designated discarding port has reached *b2*, the domain has converged to its final topology.

GVRP/MVRP for propagating TCs

Cisco.com

- **GVRP/MVRP can still be used at the vlan level to prune unnecessary traffic and advertise topology changes**

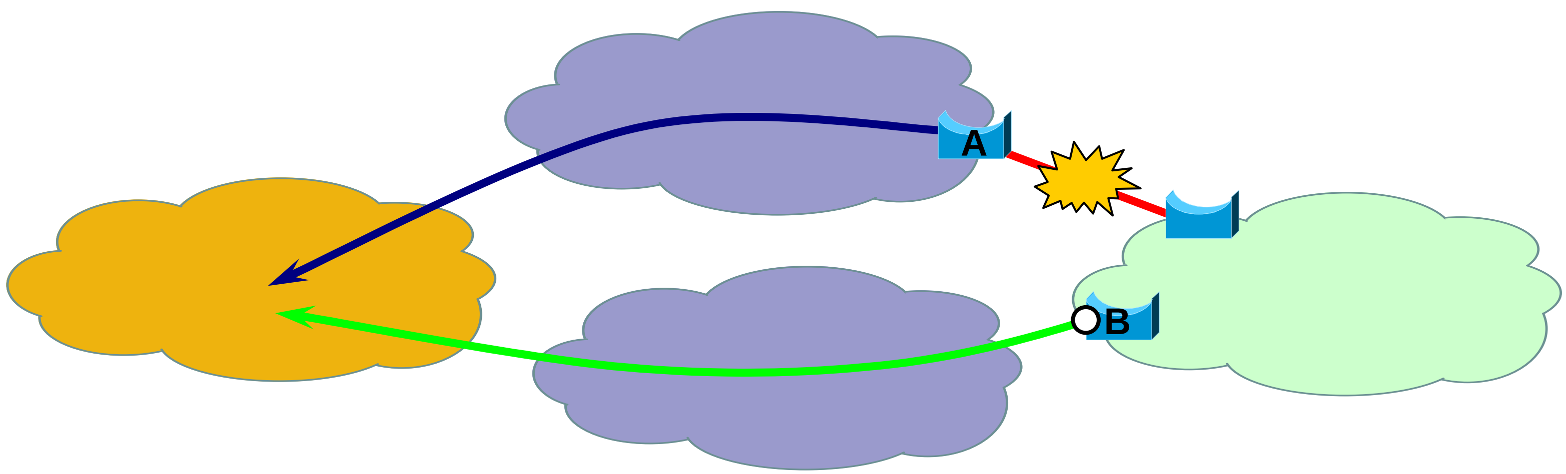

Advantages

Cisco.com

- Simple: very light modification to the STP
- Only the bridges hosting the gateway ports need this particular feature
- No transient loop during reconvergence
- Does not require cooperation from the outside domain
- An easy way to ensure that the domain is never a transit domain for its neighbors

Drawbacks

- The Root Bridge is seen as outside the domain, and practically, is at the edge of the domain (physically).
- Granularity of the instance, not the vlan
- Only tracks the physical state of the Gateway Port (may be enhanced)
- The BPDUs received on a L2GP are ignored. The following misconfiguration creates a loop:

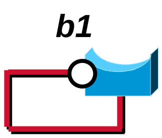

Port b1 is configured as a L2GP, it ignores the BPDUs it receives and goes unconditionally to root forwarding

Possible enhancements

Cisco.com

- **Tracking mechanism: the priority of the fake Root ID could vary depending on external parameters**

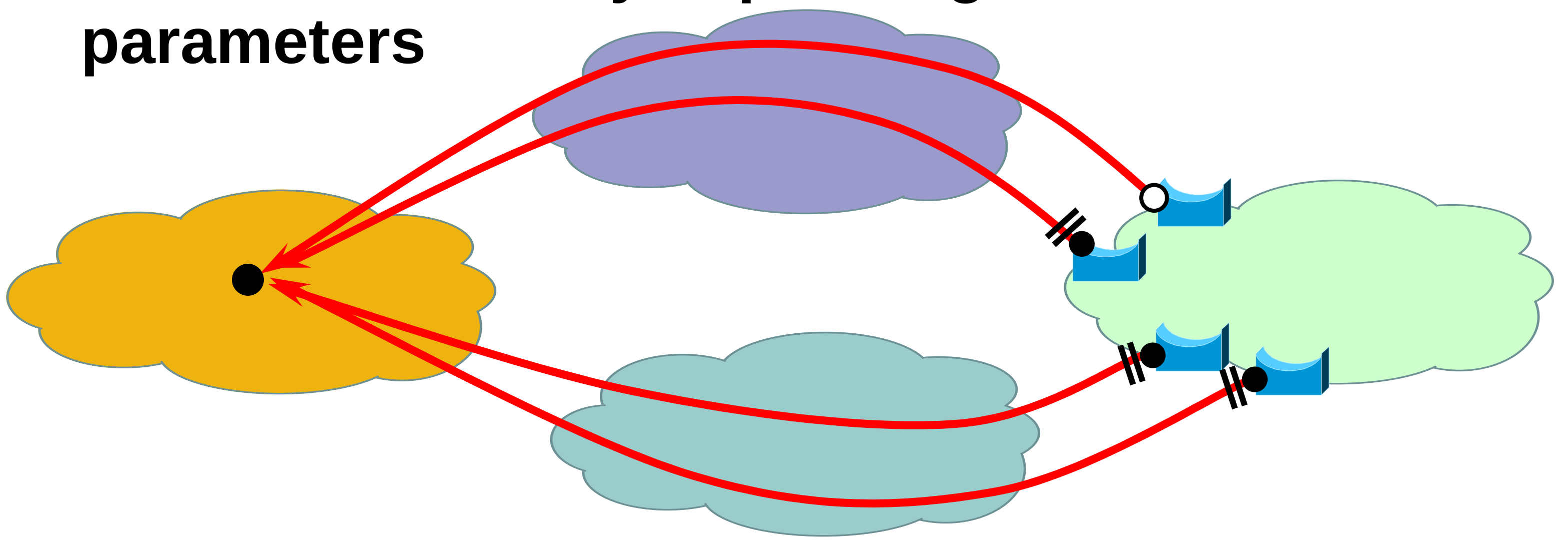

- **For instance some kind of polling (CFM?) could be achieved through the L2GPs to determine their relative priority**

Possible enhancements

Cisco.com

- **Rule #1 could be optional. The BPDUs from domain A would select the forwarding L2GP:**
  - ❑ **Interesting to just prevent the domain from being a transit area**
  - ❑ **Useful for stability purposes, even in an enterprise network**

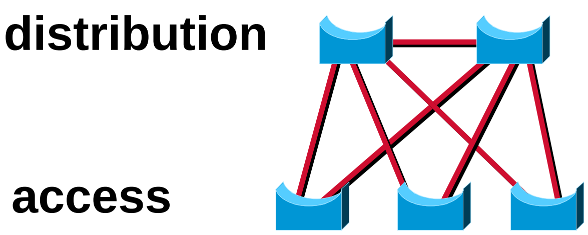

If the uplinks of the access bridges are configured as L2GPs, the access will never provide redundancy for the distribution.

Possible enhancements

Cisco.com

- **A L2GP can stay in discarding when it receives its own information (protect against misconfiguration)**

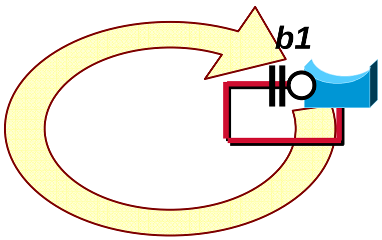

CISCO SYSTEMS