Port Security @ MIX

Concept, live experiences, procedures

Our view

① Absence of L2 routing loops or other Spanning Tree related problems is only a prerequisite for an Exchange Point

② The equation

one port – one mac – one customer

Has to be enforced strictly to ensure 'four-nines' availability of the service

③ Increasing number of LAN extension services and ethernet as dominant technology also in the WAN transport solutions pushes IXes to define policies that ensure strict separation of different ethernet domains

④ Anyway, it has to be simple and scalable…

Port Security technology

All the Switches Vendors have a port security solution available

Basically they all can either:

- Block the traffic coming from the 'offending' mac address(es)
- Disable a port which shows a number of mac addresses higher than a given threshold

Hardcoding mac addresses is not a must as it is also possible to apply PS to dinamically learned MAC addresses.

MIX PS Evolution

Deployment of Port Security at MIX was done together with the adoption of proprietary MAC addresses numbering for the peering devices. We felt this was, above all, a good troubleshooting aid, specifically in LAN extension scenarios.

First implementation was done with 'violation shutdown' policy (in the Foundry terminology, as soon as MAC addresses different from the one assigned from MIX are seen on a customer port it would have been disabled, and the block cleared only through a manual intervention and a 'back to normal' check from both MIX and the customer)

MIX PS Evolution

- After a few months, reviewing this procedure, we felt this was too rude for our customers and we tried to relax it a little by moving towards a 'violation restrict' strategy
- With this new policy, the 'offending' mac address is removed from the Switch CAM and the generated traffic is dropped

This new policy allowed us not to shut down customers' ports in case of L2 addresses leaking.

But…

‘Violation restrict’ drawback

Consider this real life scenario:

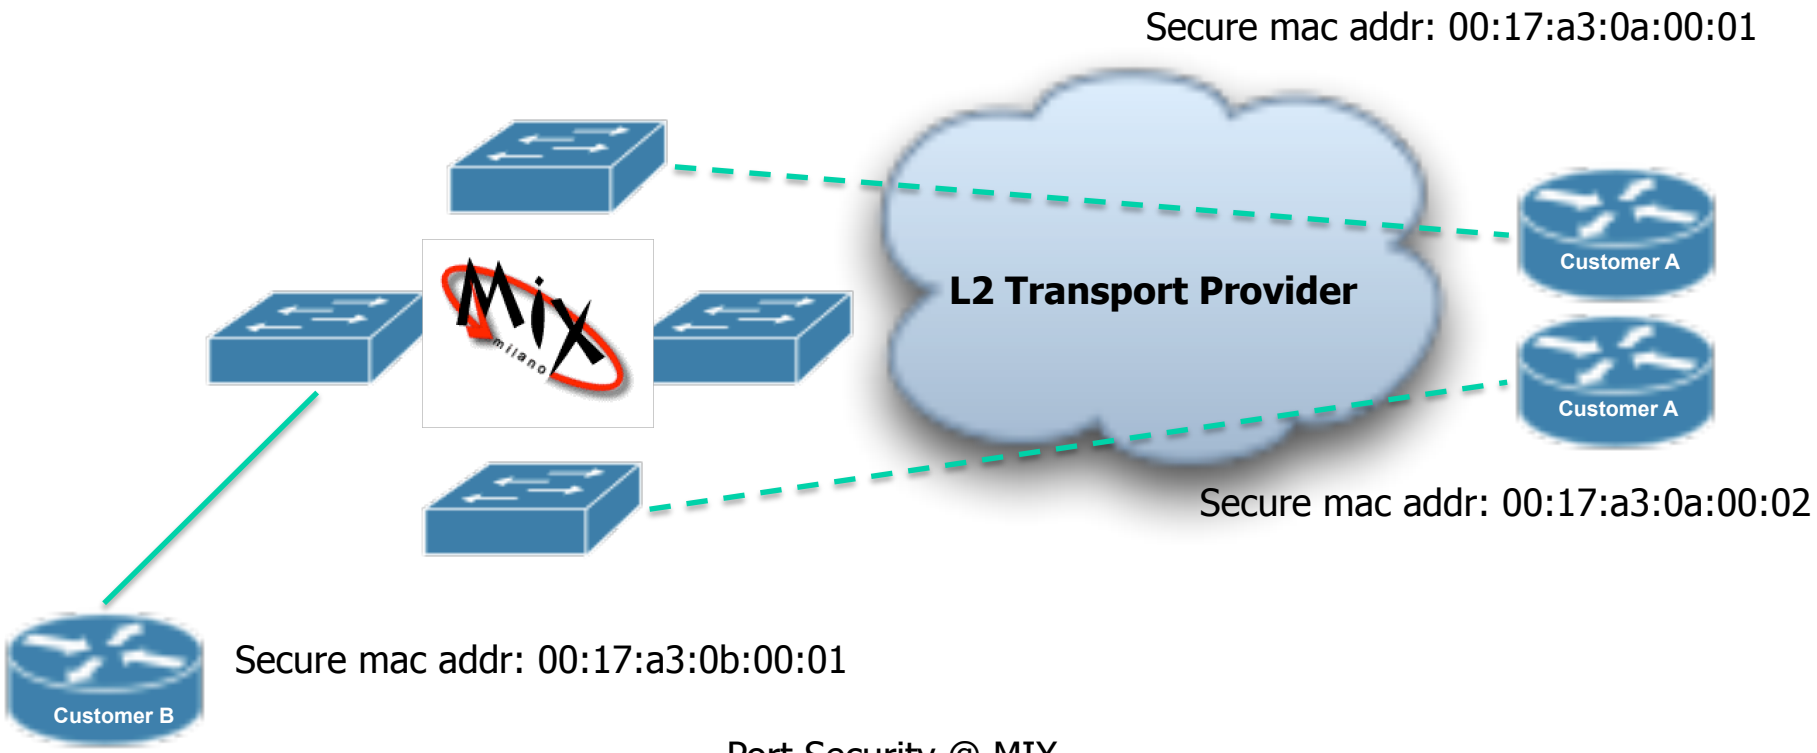

What happened one day…

From our switches perspective, we saw a 'loop' : several macs of the peering LAN were announced back by one of the two remote customer ports

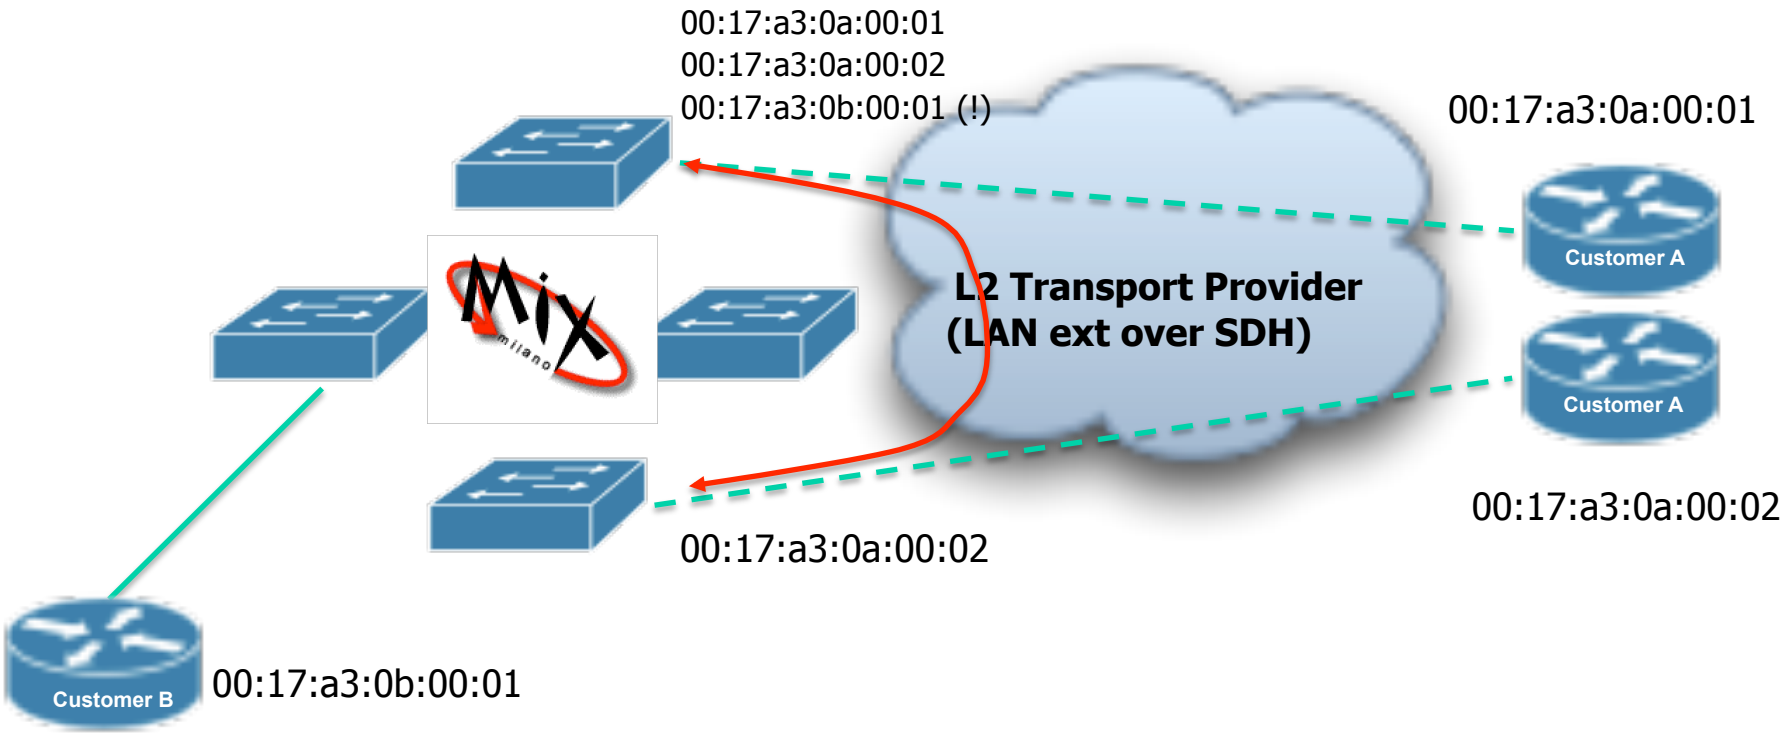

Effect:

Customer B peering sessions with ISPs connected to the switches with the 'looped' port were gone (together with the the ones with 0a:00:02 port of cust A), because PS violation restrict wiped out those MACs from the CAM.

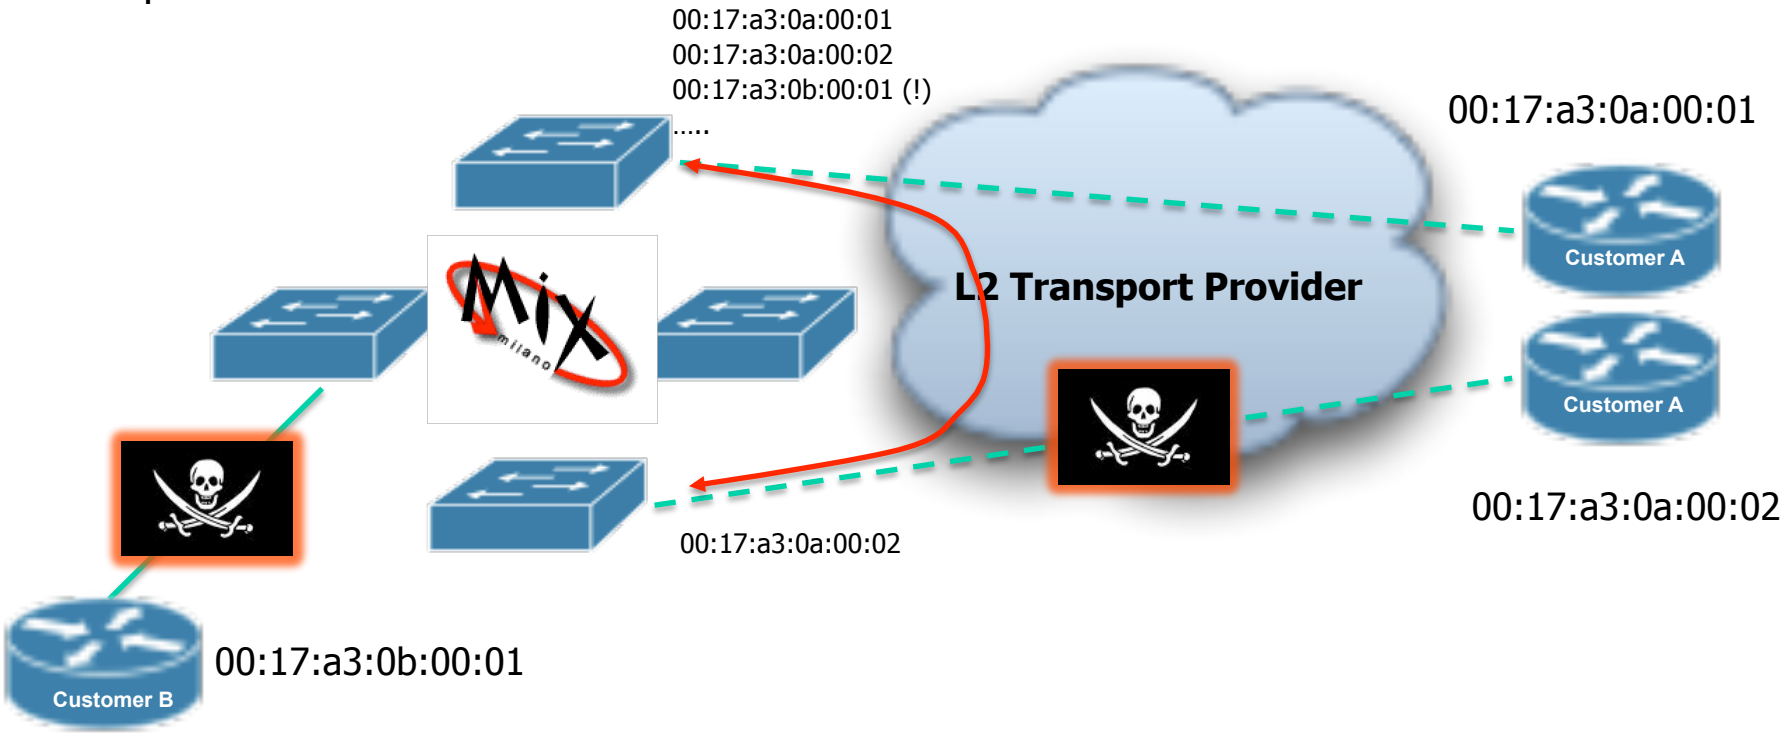

Reconsidering PS…

- ✓ This case forced us to reconsider ‘violation restrict’ policy: under similar circumstances it can create weird behaviour of the peering matrix
- ✓ Hardcoded IEEE MAC OUI was helpful: we were able to recognize at first sight what was happening
- ✓ ‘Violation shutdown’ comes back as the best option. But with some adjustments from the first implementation.
- ✓ We considered also to implement different policies upon different type of customer connections, but without finding the right recipe so far.

Where do we go from here:

- Our new policy will be something like this:
- With new customers, we will always start with a ‘violation shutdown’ policy. Is our own way to ‘quarantine’ newcomers.
- With stable-proven customers, we may return to ‘violation restrict’ approach: we will be more prone to this if they have the peering router colocated with us, a bit more reluctant if they use LAN extension/MAN circuits/campus fibers to reach our switches. In case of problems, we will go back to the ‘violation shutdown’ config.
- In any event, port security alarms are always ‘Critical’ alarms in our NMS system and have to be handled at top priority

Shut down: yes, but how long?

✓ The disable port status in the violation shutdown policy could be kept for a limited amount of time

✓ Basically using a threshold equal to zero, the port stays down forever, waiting for human intervention: we're here now.

✓ An option to mitigate this behavior is to keep the port in shutdown for some time ad then give the customer another chance…

✓ we may try this also: violation shutdown starting with a 2 hrs disable time: we feel less time could sometimes be unnoticed, longer period useless because we would have been looking into it before

Thank You!

- Web http://www.mix-it.net
- General info info@mix-it.net
- Adminstration sg@mix-it.net
- Tech Dept noc@mix-it.net
- Tel +39 02 4091 5701
- Fax +39 02 4091 5693