



SRP Multipath Stream Selection (5)

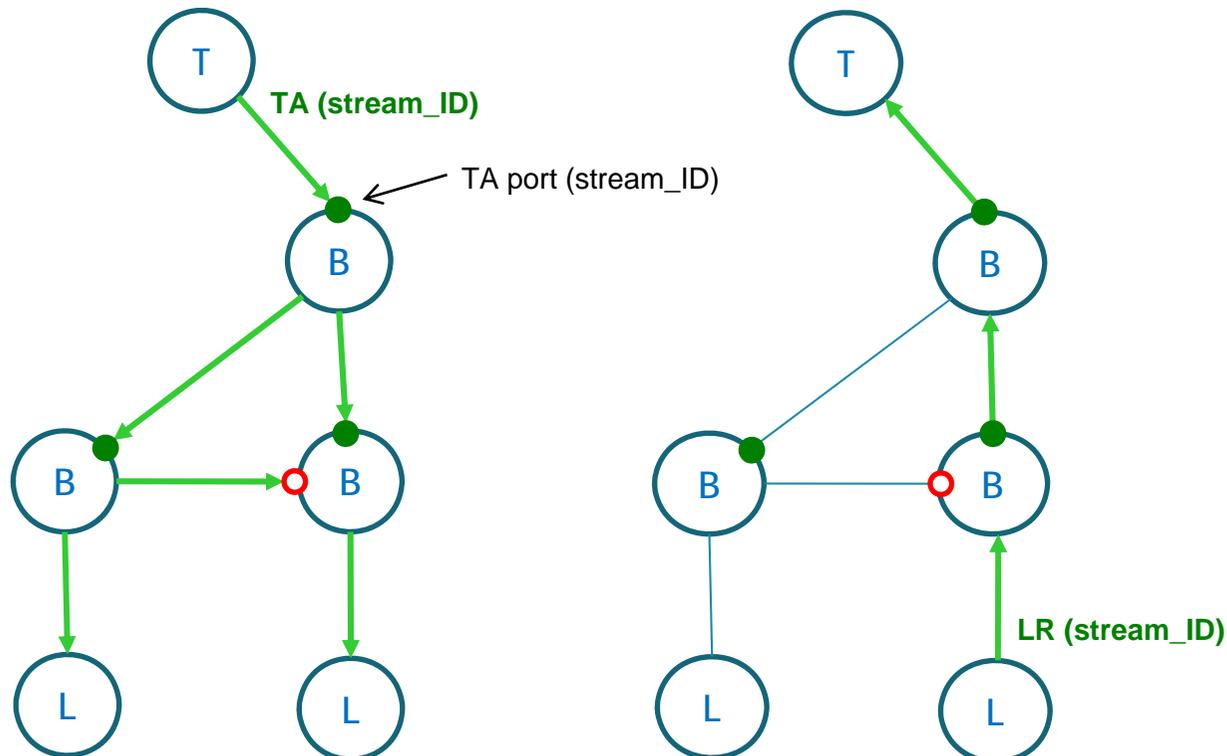
IEEE 802.1 AVB WG
Nov 2011

Presented by: Philippe Klein, Broadcom

Brief background info on MSRP Path Reservation ...



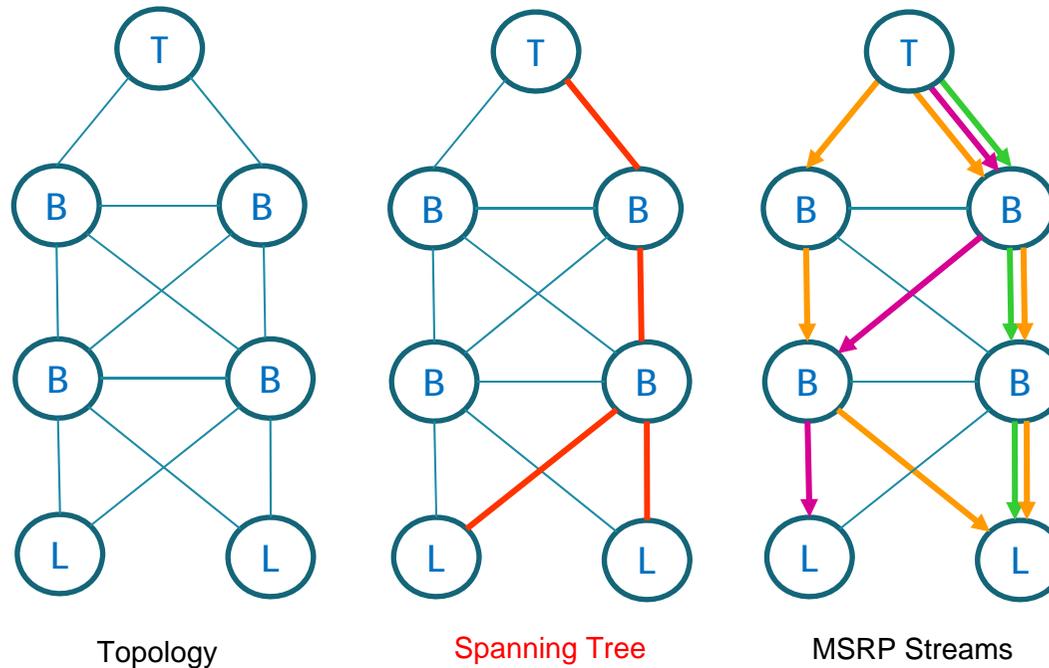
- MSRP Talker Advertise (TA) propagation
- Path creation on MSRP Listener Ready (LR)



The Goals



- Multipath streams to maximize bandwidth offered by the whole network topology
- Stream Redundancy



- Stream follows the spanning tree
- Multi-path streams beyond spanning tree
- Stream Redundancy

Additional SRP Talker Advertise Attributes



- **New SRP Talker Advertise Attributes:**
 - Stream_subID (could either be the Talker ports or the first downstream bridge “splitting” the TA over multiple egress ports)
 - Sequence ID (“easiest” option for loop protection; small counter is good enough)
 - Link Cost (computed from multiple metrics, specific computation on a per profile base) or Link Metrics

Additional SRP Talker Advertise Attributes (cont)



- **New SRP Talker Advertise Attributes:**

- Reference Counters:

- Stream Ref Count: indicates the number of TA with different Stream_subID handled by the bridge for a given stream ID
 - Port Ref Count: indicates the number of TA with different Stream_subIDs handled by a given port of the bridge for a given stream ID
 - Propagates the highest reference counts along the path
 - If a listener requires redundancy without single point of failure, it will select the TAs with a Stream Ref Count equal to 1. If such TAs are not received by the Listener, it indicates a network engineering issue for which the fix is beyond the scope of SRP.
 - The same scheme applies for a single point of failure per “wire”

Bridge's TA Propagation



- TAs are forwarded over bridge's egress port (regardless of their RSTP port state)
 - Loop protection by blocking duplicated TAs based on sequence ID
 - TAs could optionally not be forwarded on ports associated to a given spanning tree, allowing traffic separation between MSRP streams and "IT traffic"
- Configurable bridge selection between same Stream TA based on:
 - Stream SubID & Link Cost (configurable computation) / Metrics (configurable precedence between metrics)
 - Lowest Stream Reference Count

MSRP Stream Forwarding



- MSRP Stream Paths created with the “multipath” MSRP TA/LR scheme presented in the previous slides will be loop free beyond Spanning Trees: However to forward MSRP streams beyond these Spanning Trees without modifying current bridge implementations (which might first filter out ingress datagrams based on bridge’s port states), these streams could be associated to a “Dummy Spanning Tree” for which bridge ports will always be configured to “forwarding” state.

Open Questions



1. What is the effect of a change of the accumulated latency if the path is re-selected "on the fly" (i.e. a link fails and a new spanning tree is established). Does the stream fails on any change ?

Is it possible to configure a threshold so that if a new path is re-selected "on the fly" the full registration does not need to take place again (in case of link failure, the upstream bridge propagates the last stored TA on the ports of the "dummy SP" and the Listener could now receive the TAs thru a different path with a different accumulated latency...) . What happens in such case ?

Open Questions



2. Propagating TAs on every bridge's ports instead of following the spanning tree only increase the overhead traffic of TAs.

Is the increase of traffic an issue compared to the significant traffic a large number of streams creates in the current MSRP version ?

If so could new propagation rules applies, discriminating TA between:

- "flow maintenance" , propagated on registered path

VS

- "advertisement" , propagated on non-registered path

More generally, is a full periodical TA really needed for flow maintenance if no change took place since the previous periodical ?



Thank you