TSN Control & Configuration

Norman Finn
Cisco Systems

Version 3, July 3, 2015

Introduction

- This contribution is in response to the excellent analyses presented in cc-goetz-MRPv2-MSP-v13, cc-goetz-Next-Steps-v1, and tsn-sexton-feature-priority-request.
- It is an attempt to map the shortest route from the current state of the IEEE 802.1 TSN Task Group to a viable set of TSN standards.

There are four scenarios of interest

1. **Static:** The Talkers, Listeners, and relay systems are configured before power-up. There are no run-time changes to reservations.
2. **Central:** On top of (1), there is a central network controller that learns what was preconfigured, and is responsible for coordinating any changes to those configured reservations with any new reservations. Reservations can be made by Talkers, Listeners, and third-parties (e.g. applications controllers).

   NOTE: I differ strongly from Götz “next steps.” A fully-centralized control model is perfectly within the scope of IEEE 802.1!
3. **Peer-to-peer:** On top of (1), there is no central controller. Talkers and Listeners are responsible for making additional reservations using a peer-to-peer protocol (MSRP or MSRP++).
4. **Mixed:** Some relay systems know about the central controller, and some only know the peer-to-peer protocol.

AVB control model

Current AVB control model: MSRP, MSRP++

- **If** we assume that the customer does not want a central controller,
  - More to come on central controllers – I think there are better control models in that case’
- **If** we assume that we are not using time-scheduled queues,
  - Only a central controller can make the global tradeoffs necessary to use scheduling;
- **If** we assume that the Talkers are capable of describing their streams,
  - Having a third party set up the streams is merely another name for “central controller”;
- **Then**, the current model using MSRP (or improved MSRP) works just fine.
- IMO, we should continue to support this model.

Götz, et al., model: “all ISIS”

ISIS-MRP model: ISIS, MSRP**, “PCEP”

Without going into details, there are three protocols used for controlling the TSN network when a central controller is used:

1. Bandwidth reservations are made along the paths of the streams, as done today, with MSRP (or a newer version of MSRP), whether a central controller is used, or not.
2. A new protocol is used to exchange reservation information, bridge capabilities, schedules, etc. between the central controller and the bridges.
3. Topology information is obtained by the central controller by connecting it directly to the ISIS instance running the bridged network.
4. Stream descriptions and decisions about paths are distributed from the controller or from edge bridges via ISIS.

Finn, Cummings, et al. model

All-Managed Objects (all-MO) model:
extreme version

- When a controller is present, edge relay systems (1, 5) participate in UNI (MSRP++) only to make/accept opaque data registrations. The NetWork Controller (NWC) can see the registrations and order relay system to make others (e.g., replies).
- Topology (LLDP), registrations, schedules, pinned-down paths – **everything** – passes through YANG or MIBs to/from the NWC.

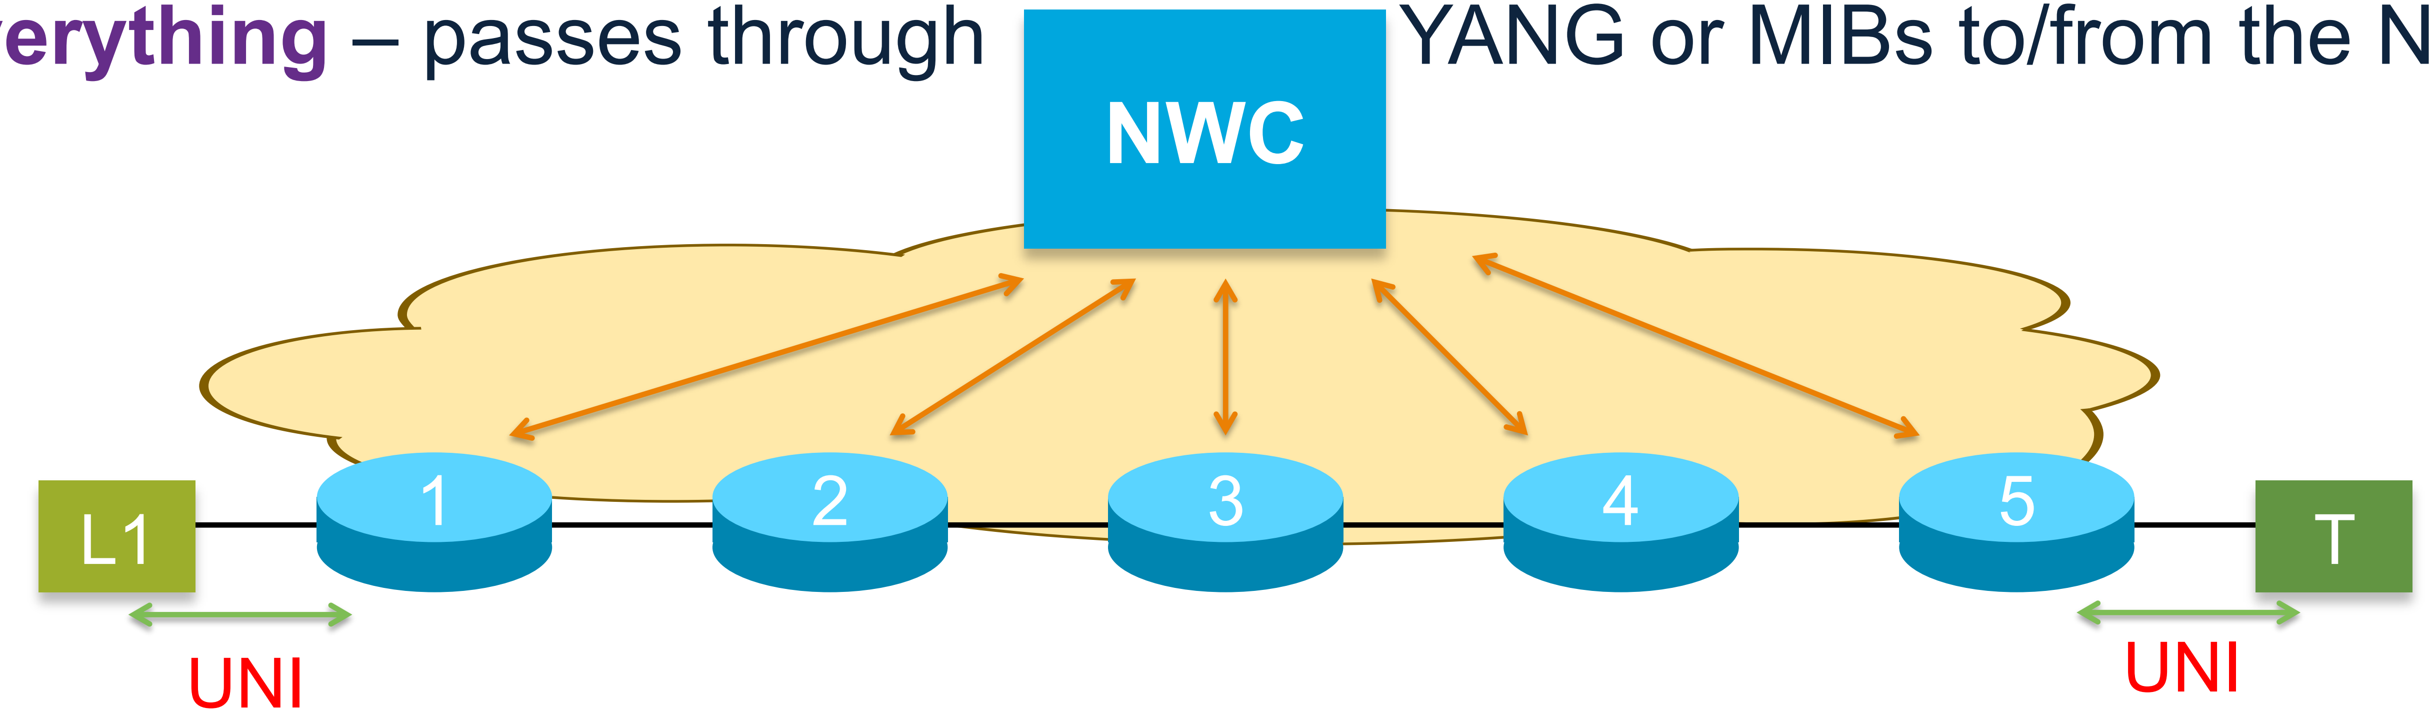

Comparing the models

Comparing the all-MO and ISIS-MRP models

- Let us look at the central controller model.
- In both the all-MO and ISIS-MRP models, the central controller requires exactly the same information from the network, and distributes exactly the same information to the network.
- **There is no difference in the two models in the information that the controller and relay systems must know!**

In the all-MO model:

- There are only **2.5 protocols**:
  - LLDP for topology discovery, which most bridges, routers, and end stations do, already.
  - YANG + carrier (or MIB + carrier), which are necessary, no matter what.
  - A simplified, non-propagating version of MSRP, controlled by YANG(MIB).
- **Any topology control protocol** is satisfactory, including MRP, G.8032, MSTP, RSTP, and SPB. (Or, none, if no redundancy!)
- Control packets are purely best-effort data, unnoticed by all relay systems except for the endpoints (controller and one relay system).
- No relay system maintains any data that is not of direct interest to that relay system.
- Things timed out: One connection to controller, per-port LLDP partners, per-port MRP timers (only if used), topology protocol (if used).

In the ISIS-MRP model

- There are **4 protocols**:
  - YANG + carrier (or MIB + carrier)
  - ISIS-SPB
  - PCEP
  - MSRP, enhanced with new capabilities.
- Plus, you can’t use any network topology protocol except ISIS.
- Control packets passed via ISIS require processing and software relaying at every hop between controller and target relay system.
- Information passed via ISIS requires every relay system in the network maintain the entire database, with a timer attached to each individual information bundle.
- Things timed out: Per-port LLDP (if used), per-port times per-registration MRP, per-port times per-LSP ISIS (adjacencies and LSPs).

This appears to be an easy decision

- In **both models**, the central controller processes the **same information**.
- **The all-MO model is faster**: relay systems receive and operate upon controller instructions in parallel, instead of in series, and all control data is passed as data, not as hop-by-hop ISIS processing and database propagation.
- **The all-MO model requires somewhat less software**: 2.5 protocols vs. 4. (Or 3, if we say that PCEP → YANG/MIB, or 5, if we don't exclude LLDP.)
- **The all-MO model requires vastly less processing power and memory**: No hop-by-hop processing of stream descriptions or paths, no databases with other systems' data must be maintained, no per-data element timers.

Making it happen

Steps needed to complete all-MO model

1. Define managed objects to support remote control of MRP in P802.1Qcc.

   (See Finn's comments on Qcc D0.4.)

Steps needed to complete ISIS-MRP model

See cc-goetz-Next-Steps-v1 for complete list.

Supporting slides

There are two network topologies of interest

- Few hops

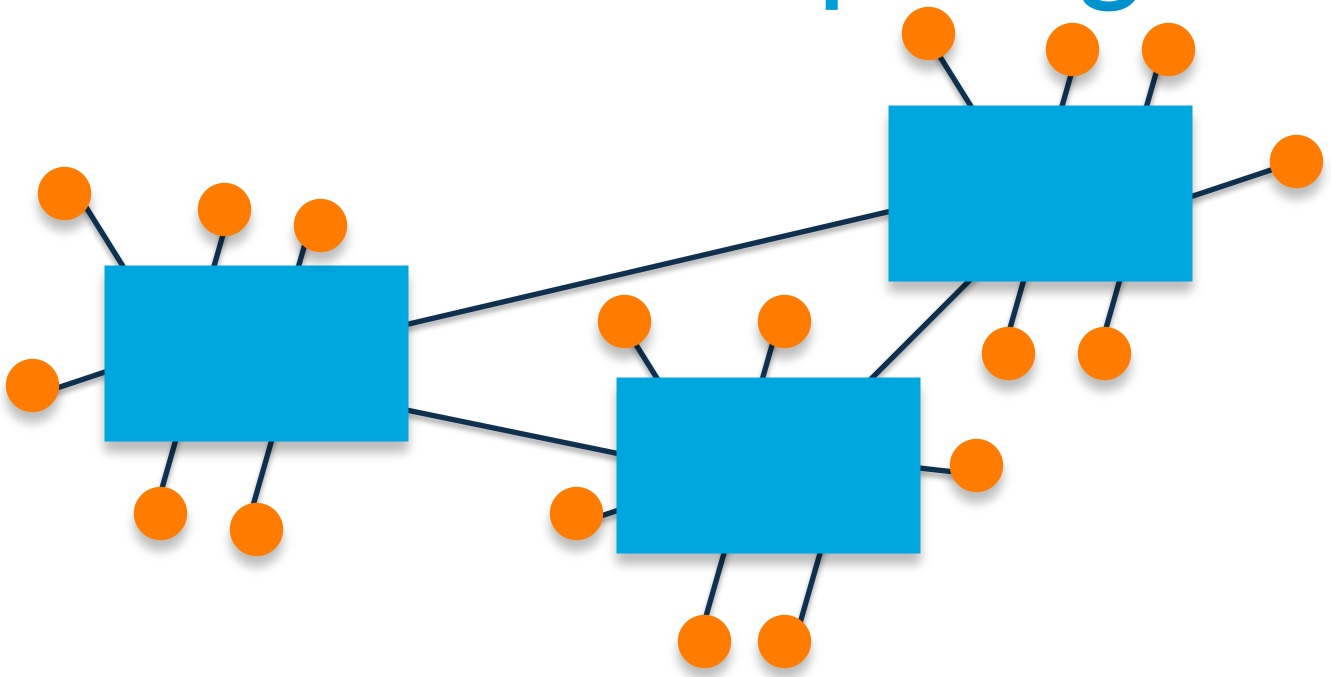

- Many hops

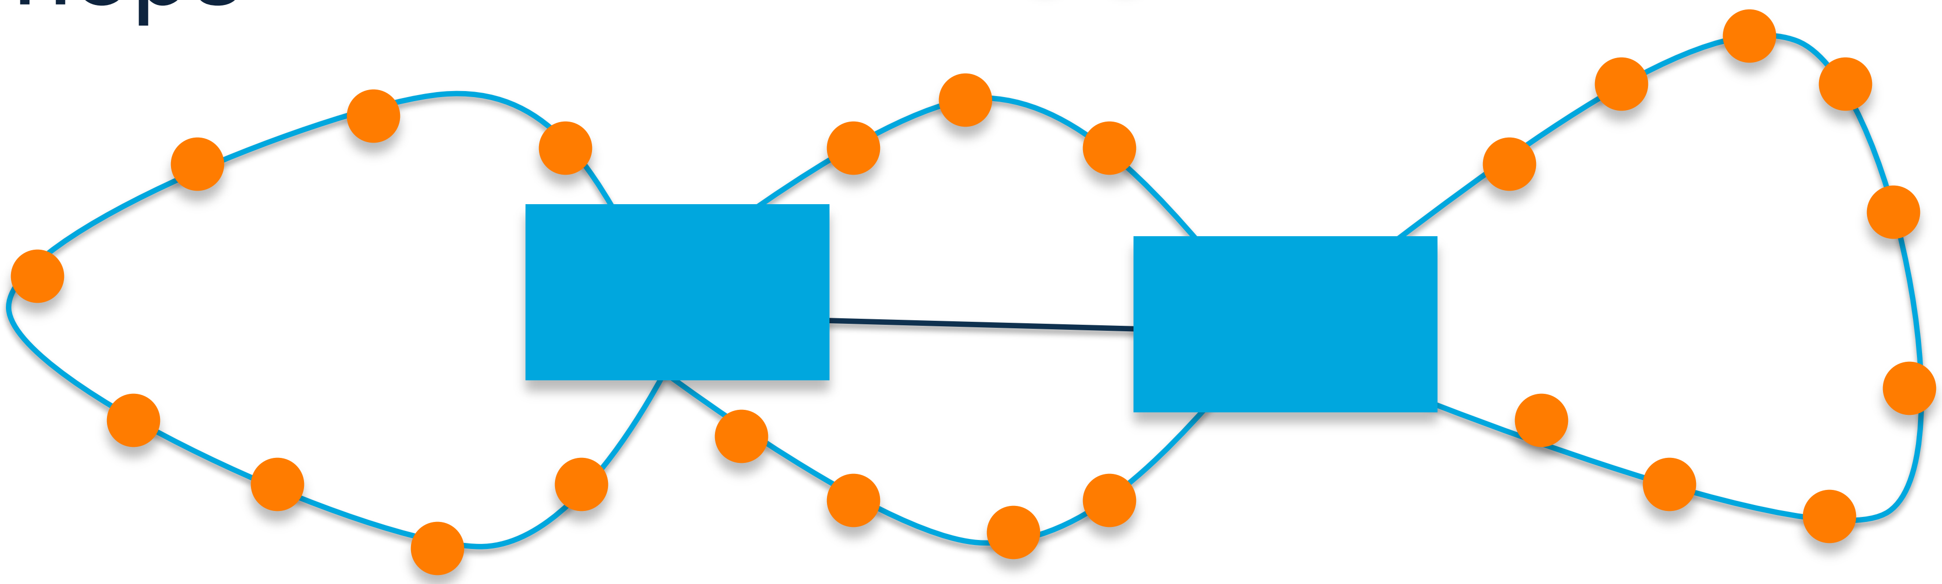

ISIS-MRP model: ISIS, MSRP**, PCEP

- **But, if** we assume a “many hops” network (long chains of two-port end systems), and
- **If** we assume that we are running current ISIS,
  - We’ll talk about the “S2IS” later;
- **Then**, this model has some severe drawbacks, outlined in the next slide.
  - We will talk about an alternative model, a little later.

ISIS-MRP problems: many hops network

- The end systems must maintain the full topology database.
  - Again, we'll talk about the S2IS possibility, later.
- The end systems must maintain all stream descriptions distributed by ISIS.
  - Knowledge of the ISIS stream descriptions must precede the MSRP** registration phase.
  - Pruning them based on network topology and knowledge of the locations of Talkers and Listeners is conceivable, but requires multiple Dijkstra calculations on the topology in each end system to determine whether it cares about a given stream. This slows down convergence after a topology change. It's probably easier to maintain the whole database.
- Every new or changed stream description anywhere in the network requires every end system in the network to wake up and pay attention to receive two and send two ISIS packets for the new/altered LSP.

ISIS-MRP problems: many hops network

- Consider also the “pinned-down” paths for Seamless Redundancy.
- Passing the path information via ISIS (802.1Qca), again, requires every end system to maintain the entire path database for all streams in the network.
  - Even if an end system knows it’s not on a given stream’s path, it must maintain the path information and pass it on, because the end system may provide the only path through the network from the controller to some relay or end system that requires the information.

S2IS might help, but much more work to do

- IEEE 802.1 has not, so far, even discussed how to make the bridged LAN function properly with the S2IS. In particular, the S2IS cannot know in which direction (left or right) to send a packet without doing the Dijkstra calculation.
- In the ISIS-MRP solution, the S2IS still must maintain the entire stream description database.
- The topology database can be deleted. It may be that a S2IS only need to retain those path descriptions that pass through it, and that it can pass other path descriptions transparently. Passing some information transparently and retaining other information has not been investigated (to my knowledge) but it seems plausible.

Security

- The security model for the all-MO model is that each relay node has a secure relationship (e.g. an IPSec connection) with the NetWork Controller (NWC). That requires one secure connection per relay system, which means o(n) trust relationships.
- All requests are vetted by the NWC. Policy, including security policy, is applied.
- We may or may not have to add a security element for the UNI. MACsec may be sufficient. We may have to add data integrity TLVs to the UNI, in order to avoid MACsec on the physical wire. **We could discuss connecting the Talker or Listener directly to the controller.**
- The security model for MSRP and ISIS is less obvious. Contributions from its proponents are encouraged. Are there widely-shared keys?

Scaling

- Neither model scales super well. Certainly, there is a lot of control traffic near the NetWork Controller. But, this is best-effort data once it leaves the controller, not hop-by-hop ISIS packets to be processed by every system along its path.
- On the other hand, the total number of packets processed by any given relay system, and in particular, by a two-port chained end system, is orders of magnitude lower for the all-MO model than for the ISIS-MRP model.
- The IETF has long recognized that multiple controllers (PCEs) are necessary for scaling to large networks. The same rules apply, here. Whether peered controllers, or a hierarchy of controllers, is better, is To Be Determined. But, we know from experience it can be done. In particular, the IETF PCE WG has a number of drafts on this problem.
- The usual way to scale up multiple ISIS regions involves BGP. No proposal has been made for this, either. Experience with BGP shows clearly that it will be subject to the same scaling issues as its underlying ISIS basis.

Let’s look at an example of the user’s problem:

- We want to do an open source stack for a 2-port chainable end system in the AVnu Alliance.
  - If a chain of two-port end systems is not looped back for redundancy, no topology protocol is required in those systems.
  - If the chain connected at both ends, simple STP, MRP, G.8032, and many other protocols are sufficient for the all-MO model.
- **This author does not believe that the chances of success for this open source effort will be improved by requiring this stack to support the ISIS-MRP model!**
  - The requirements for memory for data base storage, alone, will render the ISIS-MRP solution unacceptable for many potential users.

Comparing the models

- Databases maintained by a two-port end system

Which streams to pass or not
Entire network topology
All stream descriptions
Reservations passing through
All pinned-down paths

Which streams to pass or not
Four integers

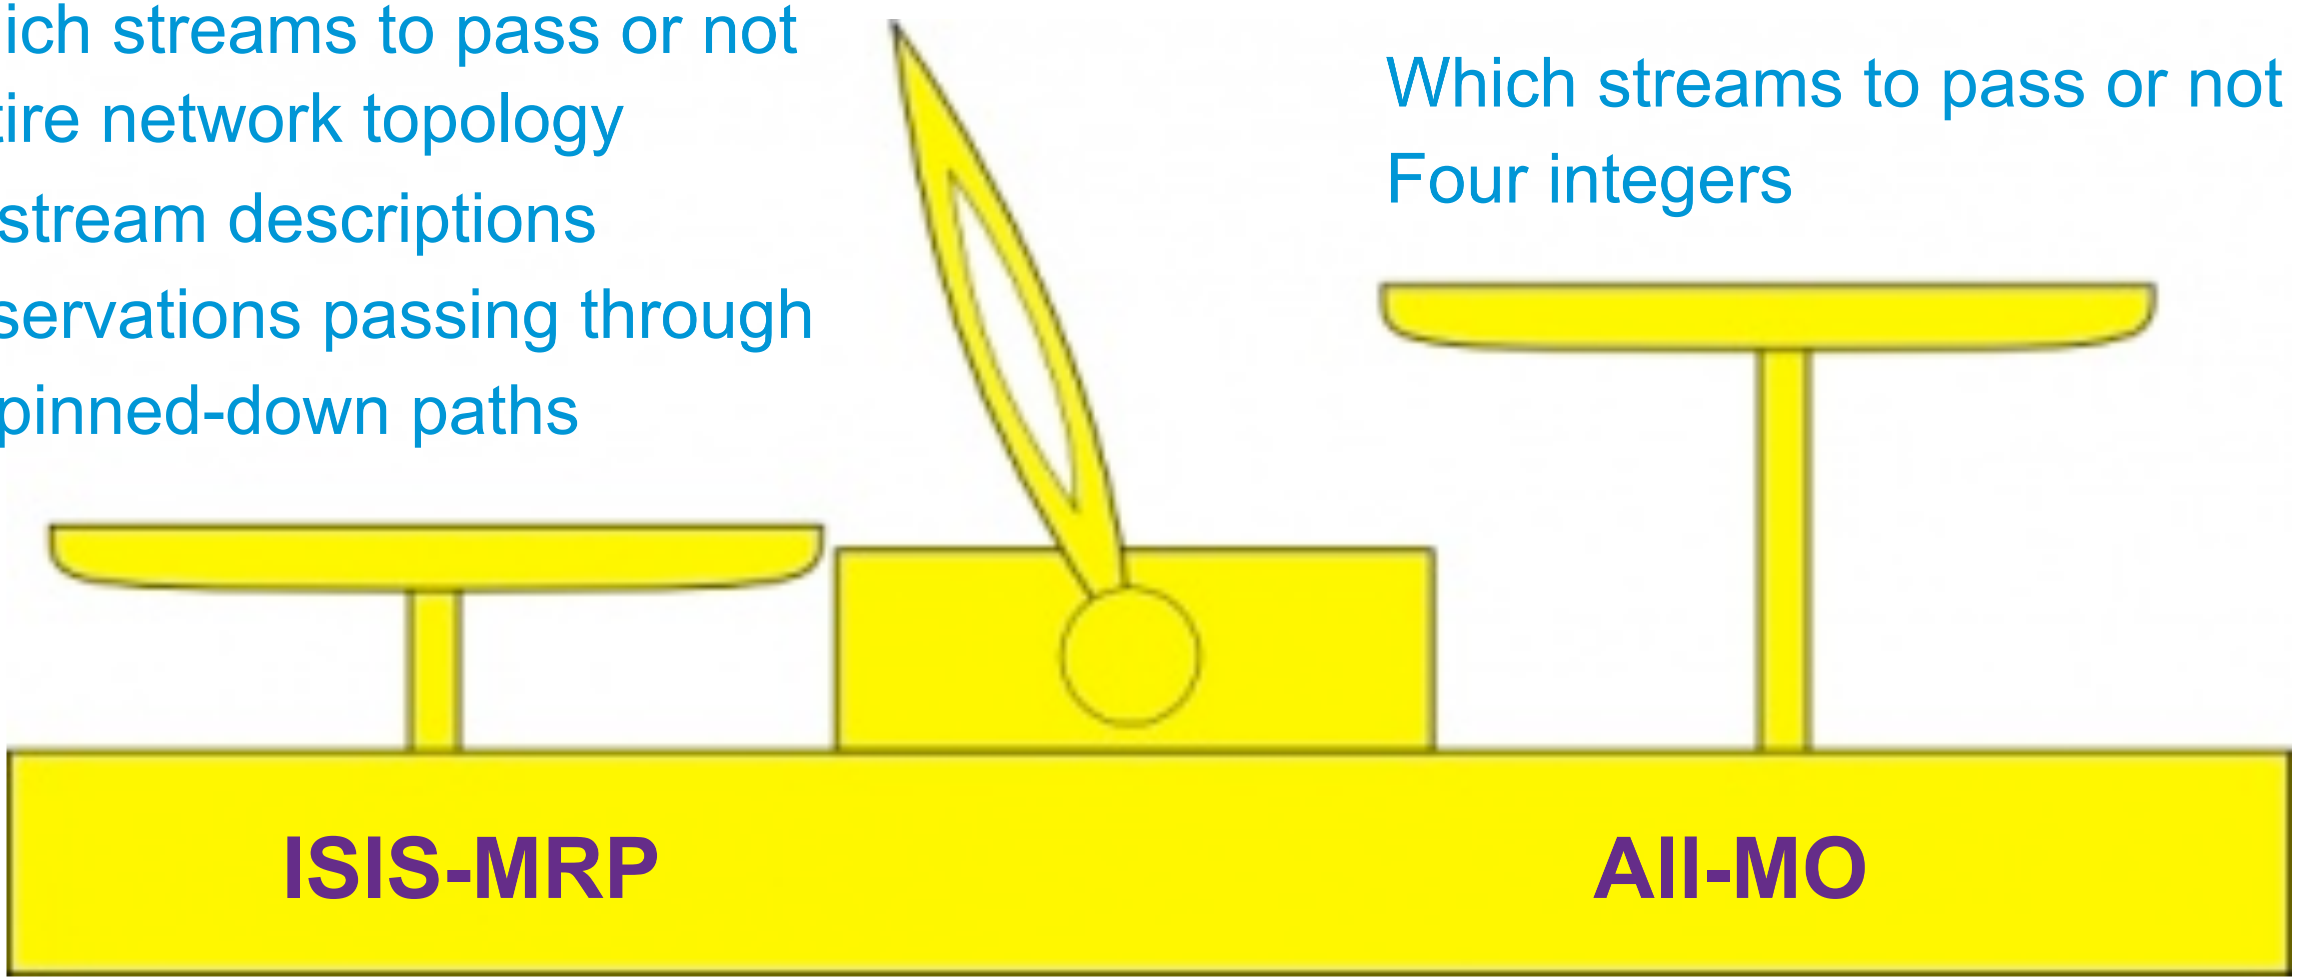

Comparing the models

- Control packets processed by a two-port end system

One packet in and one packet out when something relevant to **this end system** changes

Two packets and two packets out (or more), every time **anything in the network** Changes.

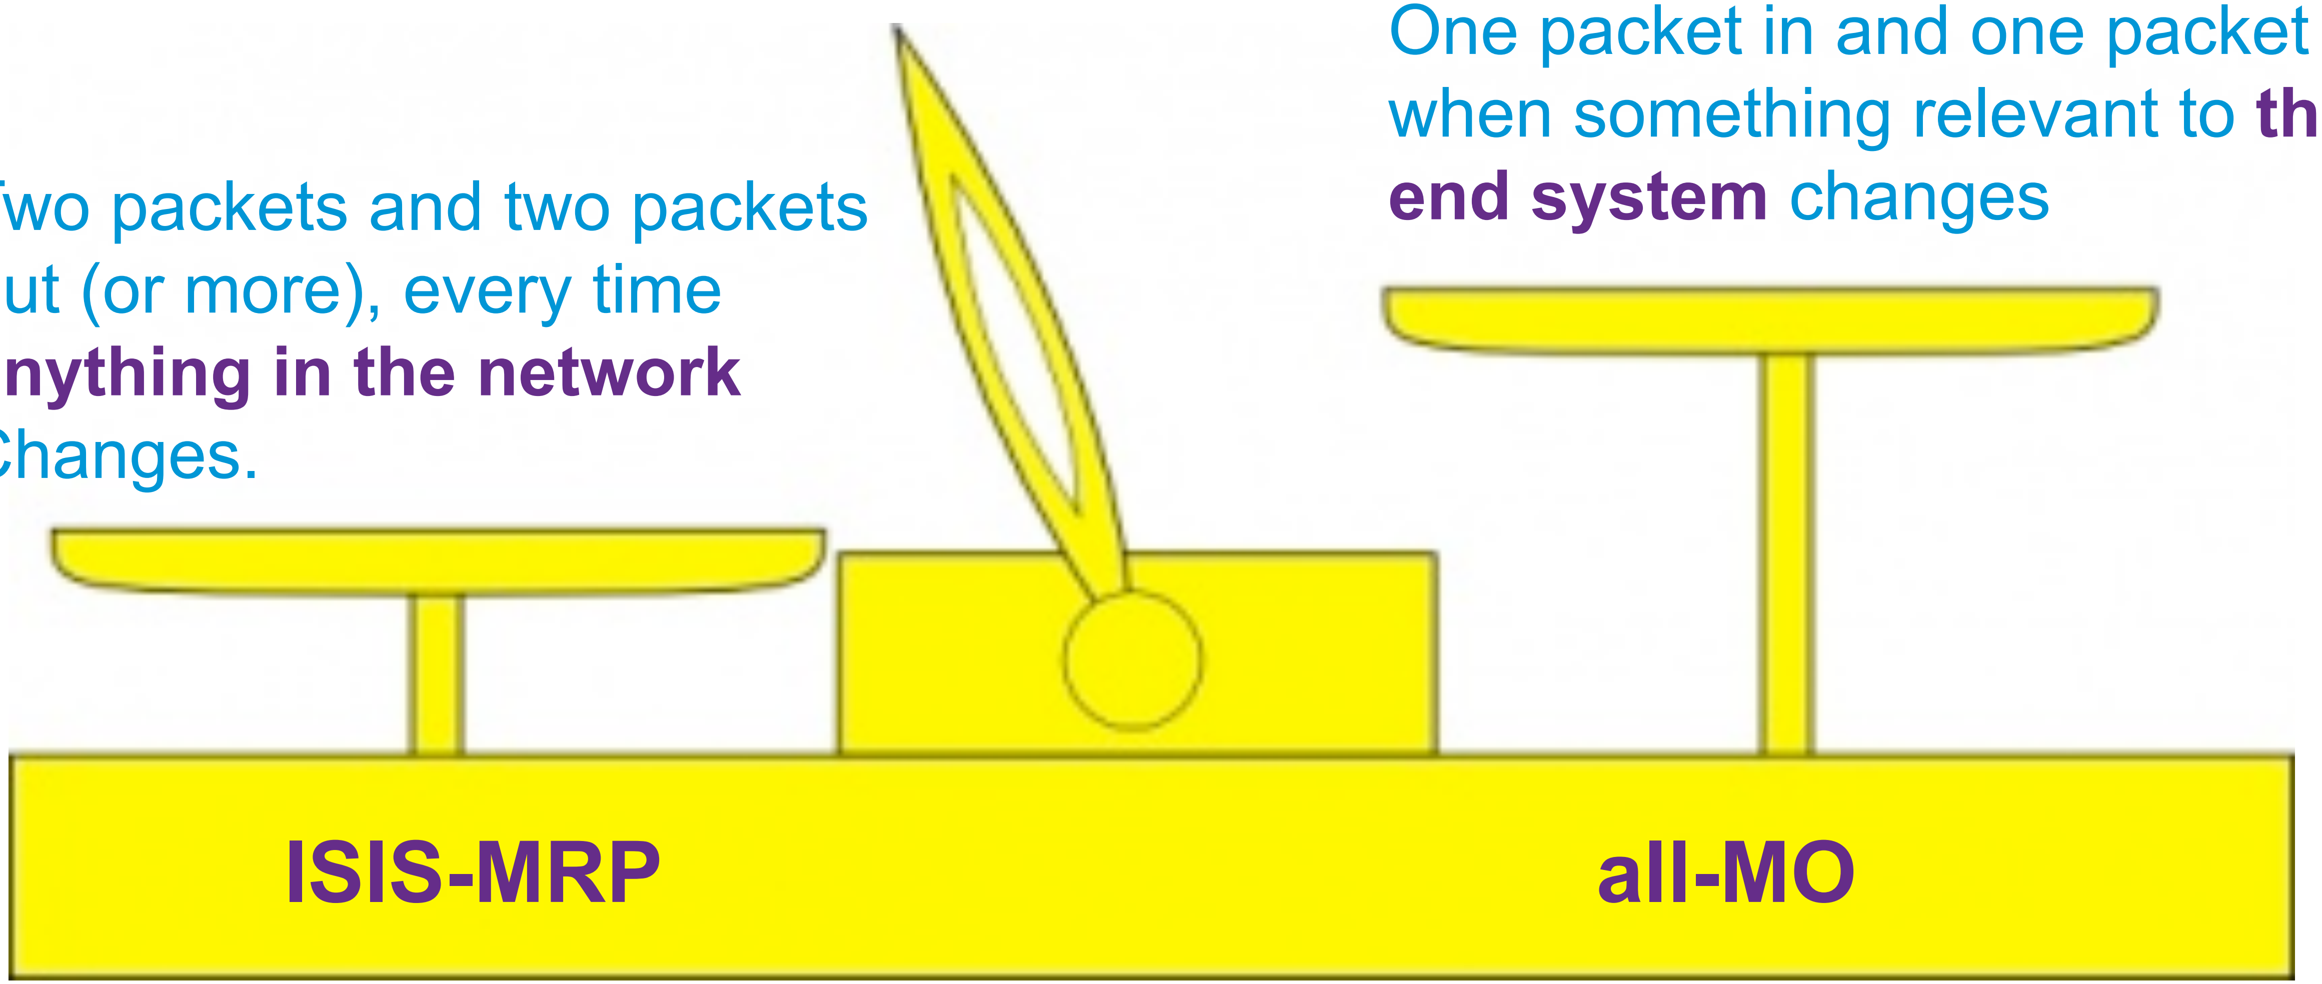

Comparing the models

- What must be implemented if controller **is** used

ISIS for all of the databases
MSRP** data distribution model
YANG/MIB for schedules, etc.

ISIS or LLDP for topology only
Opaque MSRP++ for UNI
YANG/MIB for schedules, etc.

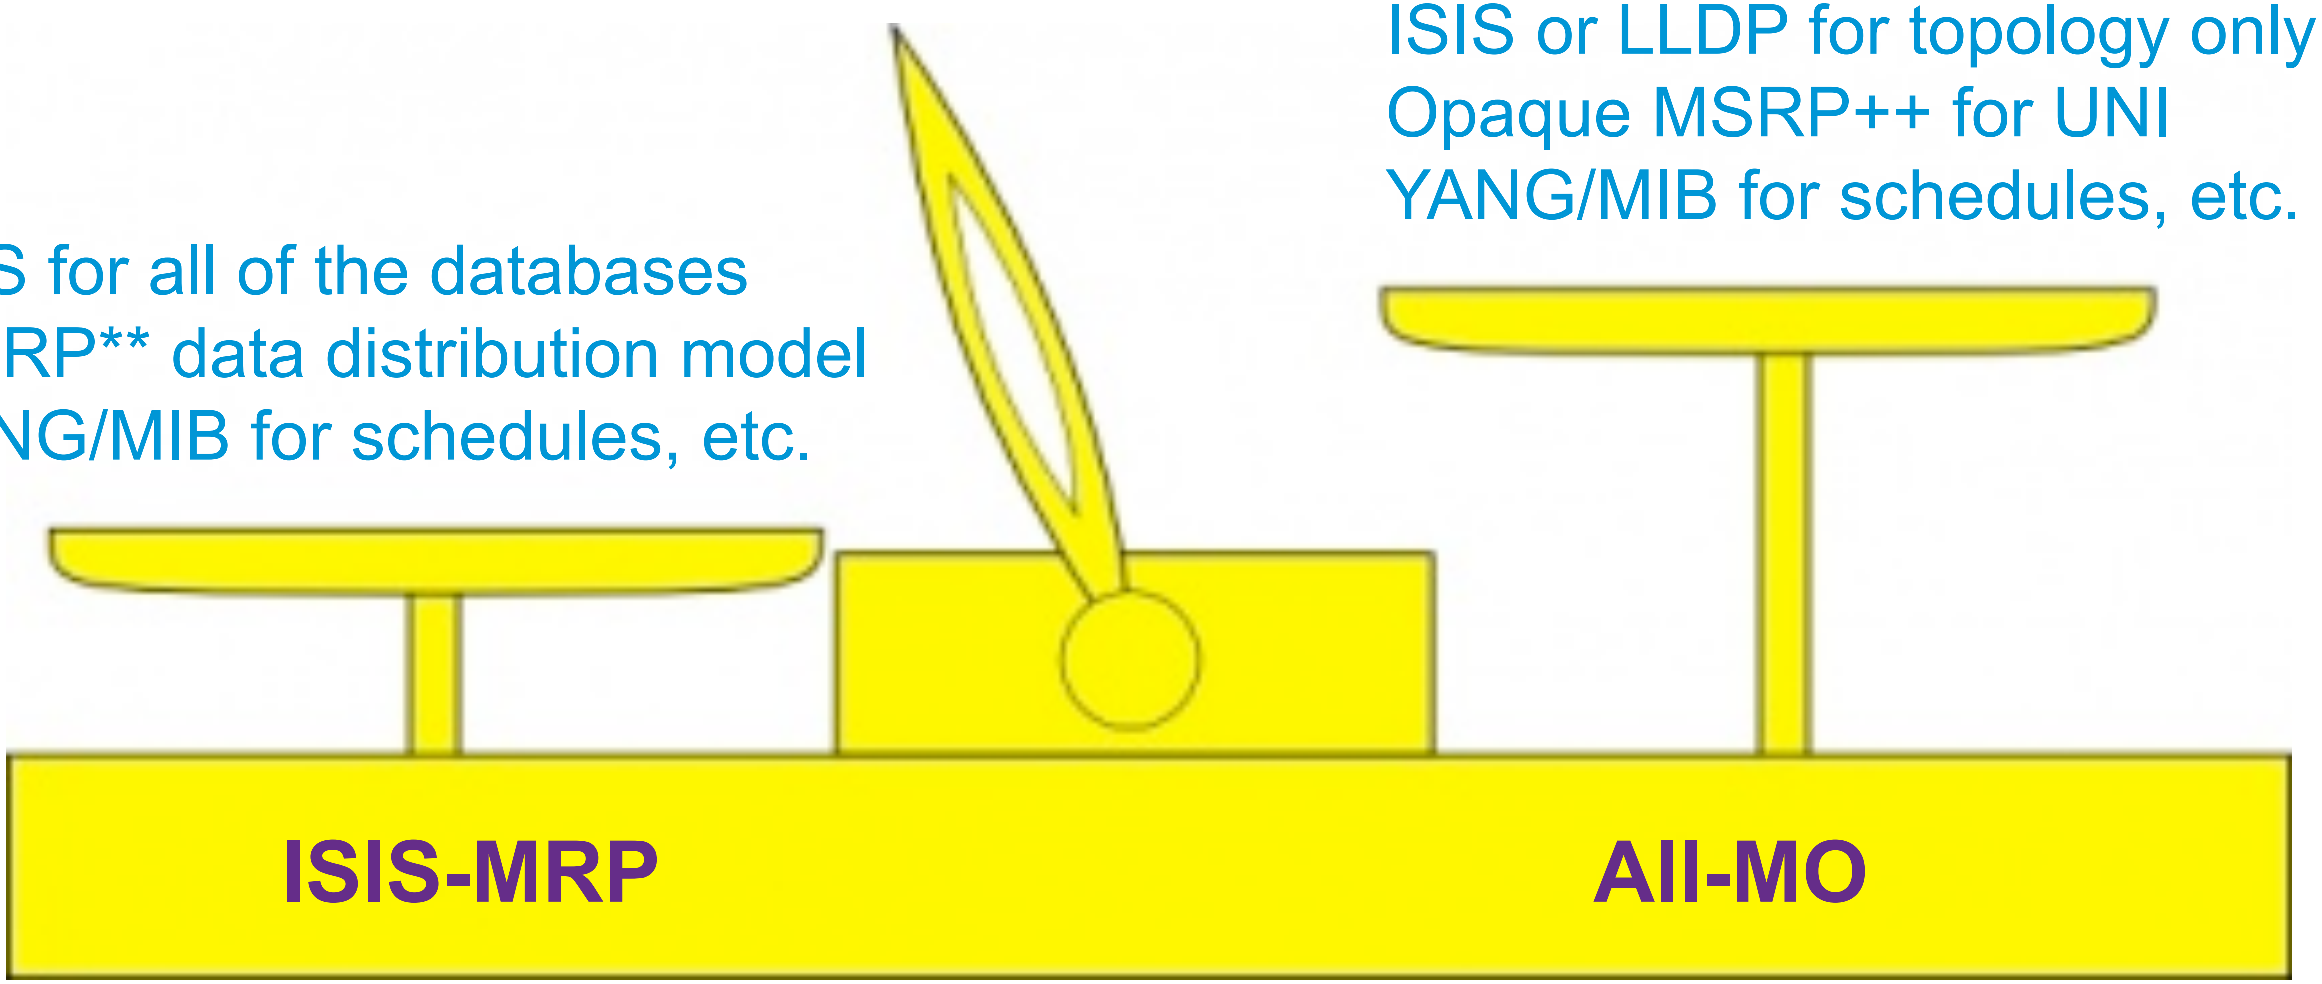

Comparing the models

- What else do we have to standardize?

Edge system / controller protocol
UNI/MSRP**
ISIS stream descriptions
Augmented MSRP for compatibility

UNI/MSRP++ with remote YANG/MIB
Augmented MSRP for compatibility.

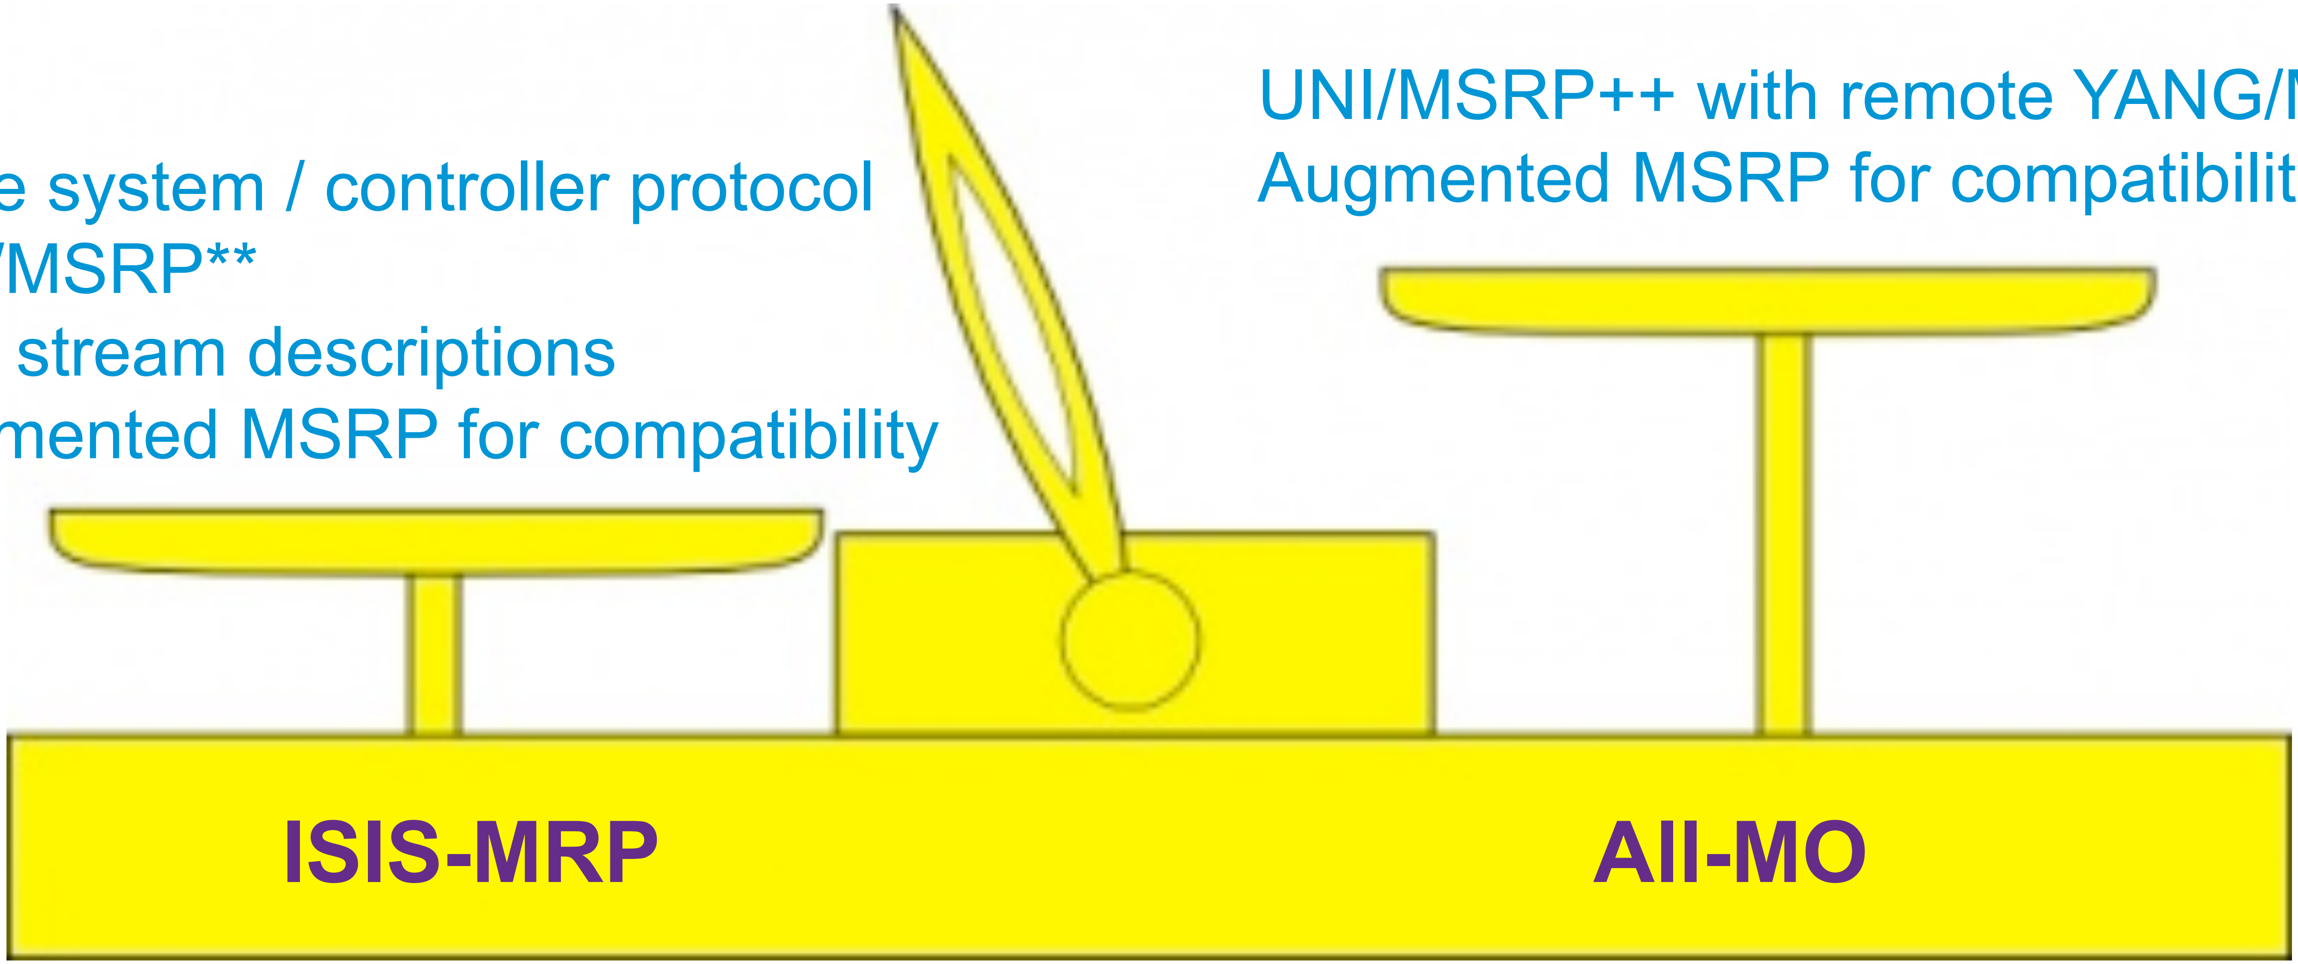

Comparing the models

- Compatibility with non-SPB topology protocols?

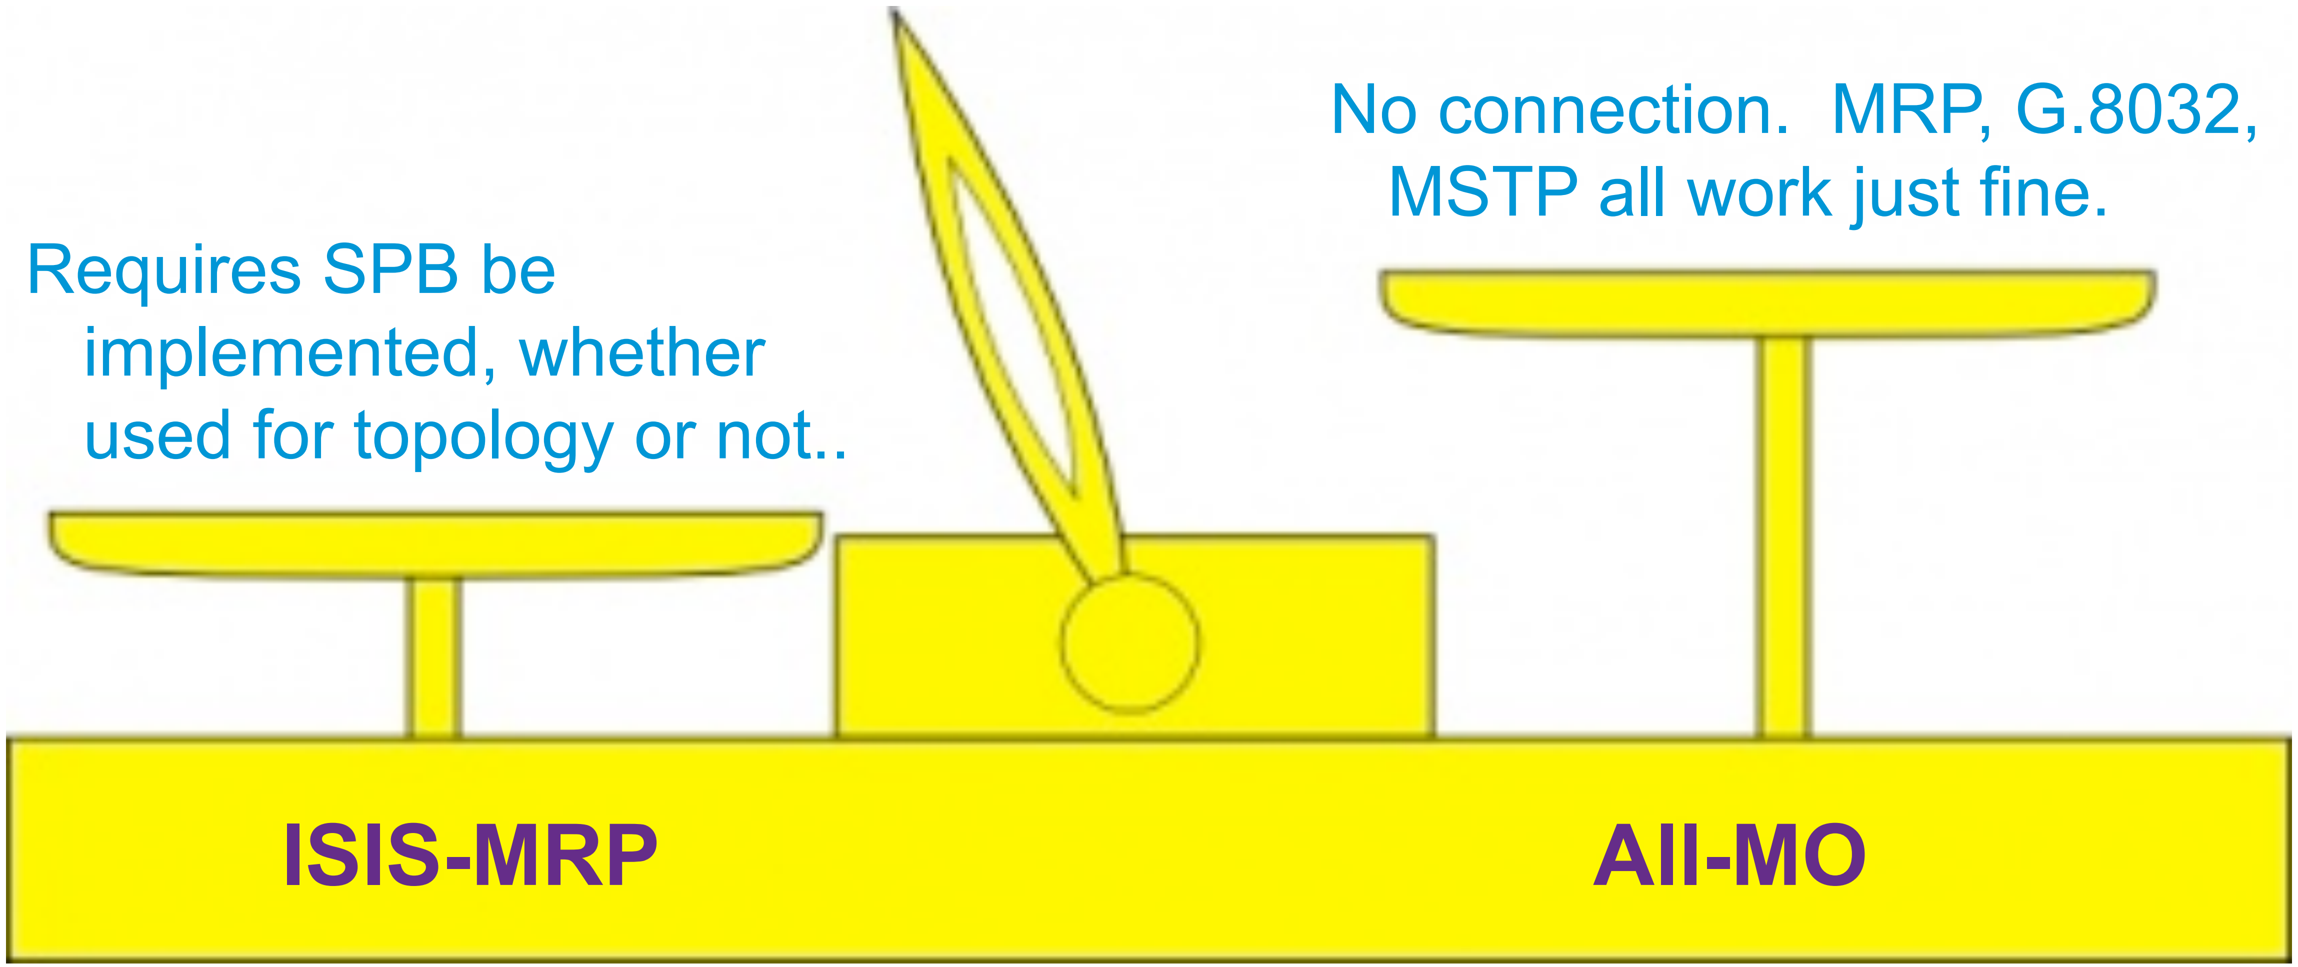

Comparing the models

- What new things are required in a mixed L2/L3 network?

Extending multiple ISIS instances to the controller.
Router participates in *N* L2 ISIS instances + L3 instance(s).
Data must migrate between L2/L3 and L2/L2 ISIS instances.
Mixed L2/L3 MSRP**

Mixed L2/L3 MSRP++ (**only** if no controller)

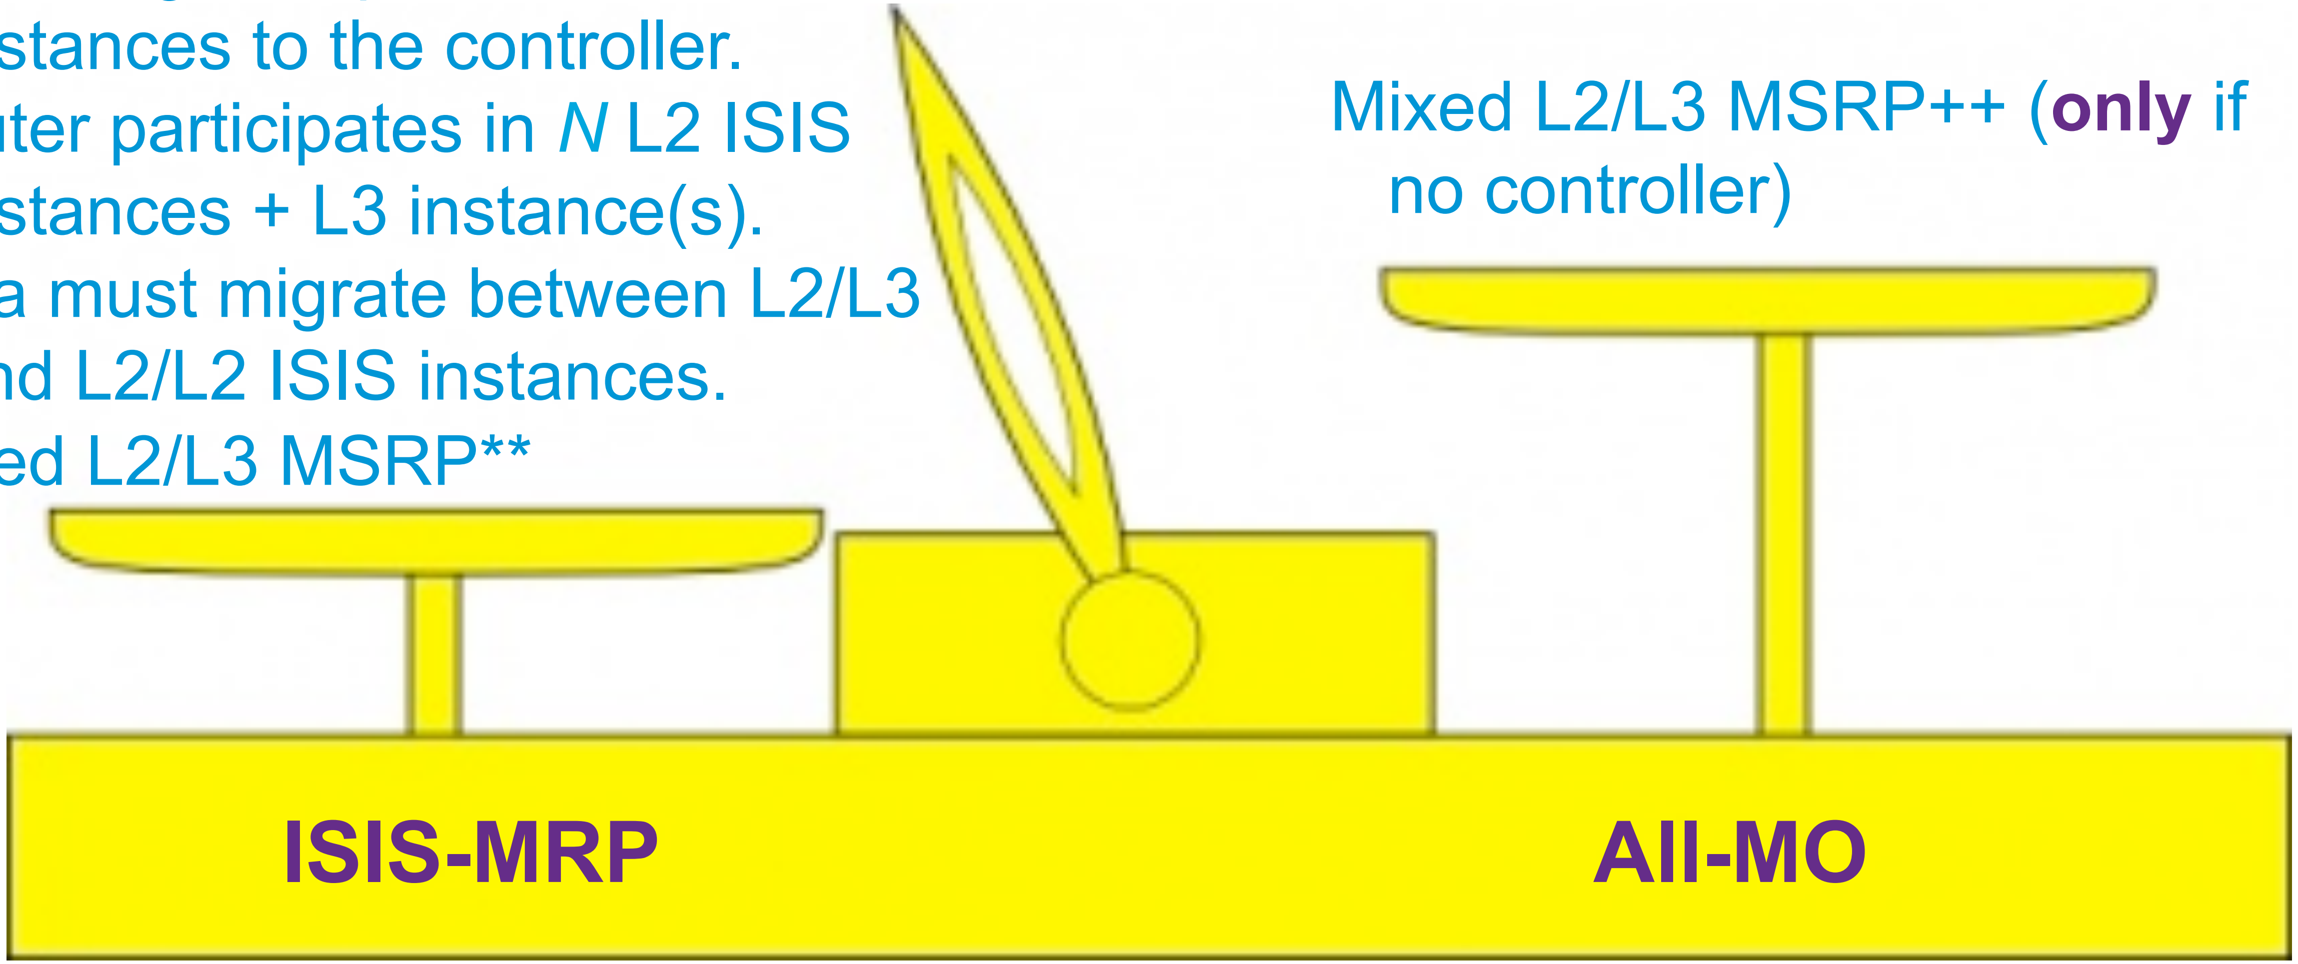

Can the use of VLANs help?

- It may be that dividing a network up into geographically-based VLANs might help reduce the sizes of the databases that must be kept.
- At present, computing VLAN forwarding requires each node to have the full topology; so the S2IS won't work; the two-port end stations still have to run full SPB.
- Furthermore, I believe that the proposal for distributing stream descriptions uses ISIS, not MxRP. ISIS floods all information everywhere, regardless of VLAN. Flooding information along VLANs gets one into the trap of requiring full connectivity to create anything, and allow topology failures to cause essential information to be deleted, lengthening the time it takes to restore service.
- This idea needs more explanation from its proponents.

Summary

- **If we accept as a given**, that we need to support the scenario where a central controller is present only during configuration time, or perhaps during startup time, and it then disappears …
- **Then** the extreme all-MO model, as presented in this deck, requires nothing else to be completely dynamic, except:
  1. A requirement to run LLDP for topology discovery.
  2. The ability to run the MSRP++ UNI interface remotely, via a YANG (or MIB) model.
- This author believes very strongly that it is absolutely essential to standardize the trivial amount of additional work required by the all-MO model.
- If 802.1 insists that only one model be standardized, it should be the all-MO model, not the ISIS-MRP model.
- Until the discussion has progressed further, this author has no strong opinion on whether the ISIS-MRP model should be pursued any further.

Thank you.