TSN Control & Configuration

Norman Finn
Cisco Systems

Version 2, June 24, 2015

Introduction

- This contribution is in response to the excellent analyses presented in cc-goetz-MRPv2-MSP-v13 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).
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.

There are two network topologies of interest

- Few hops

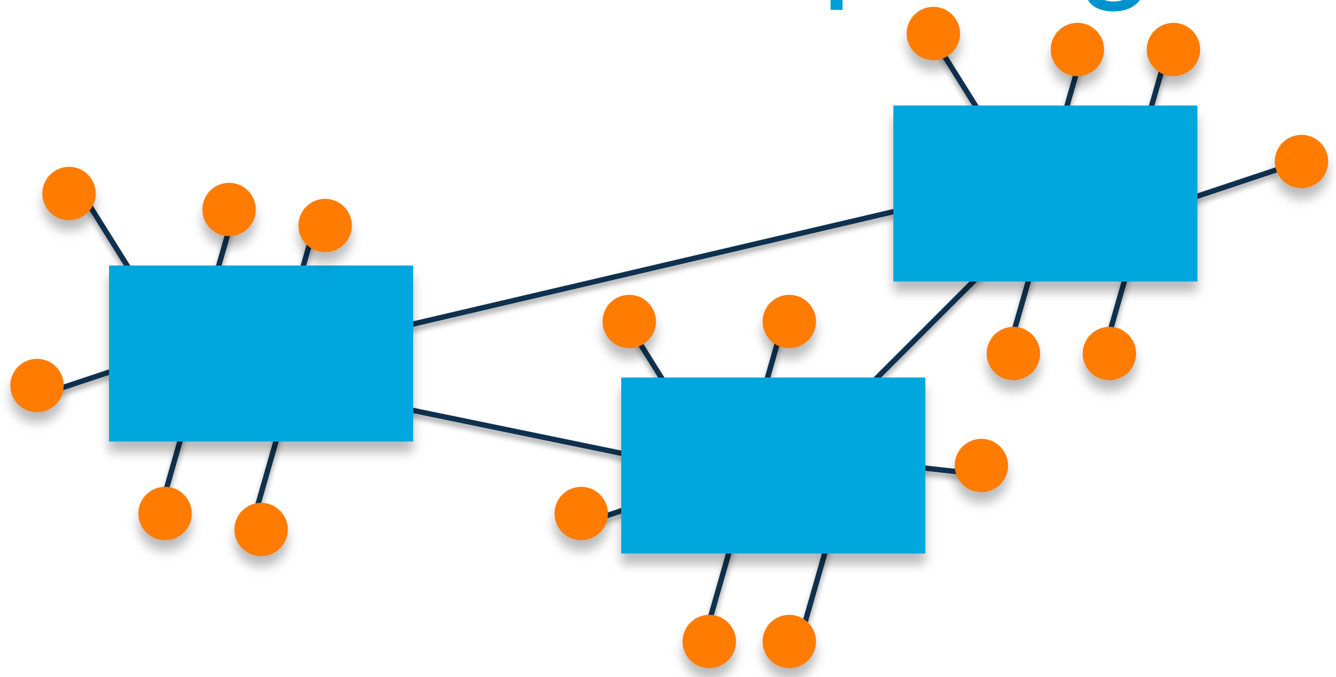

- Many hops

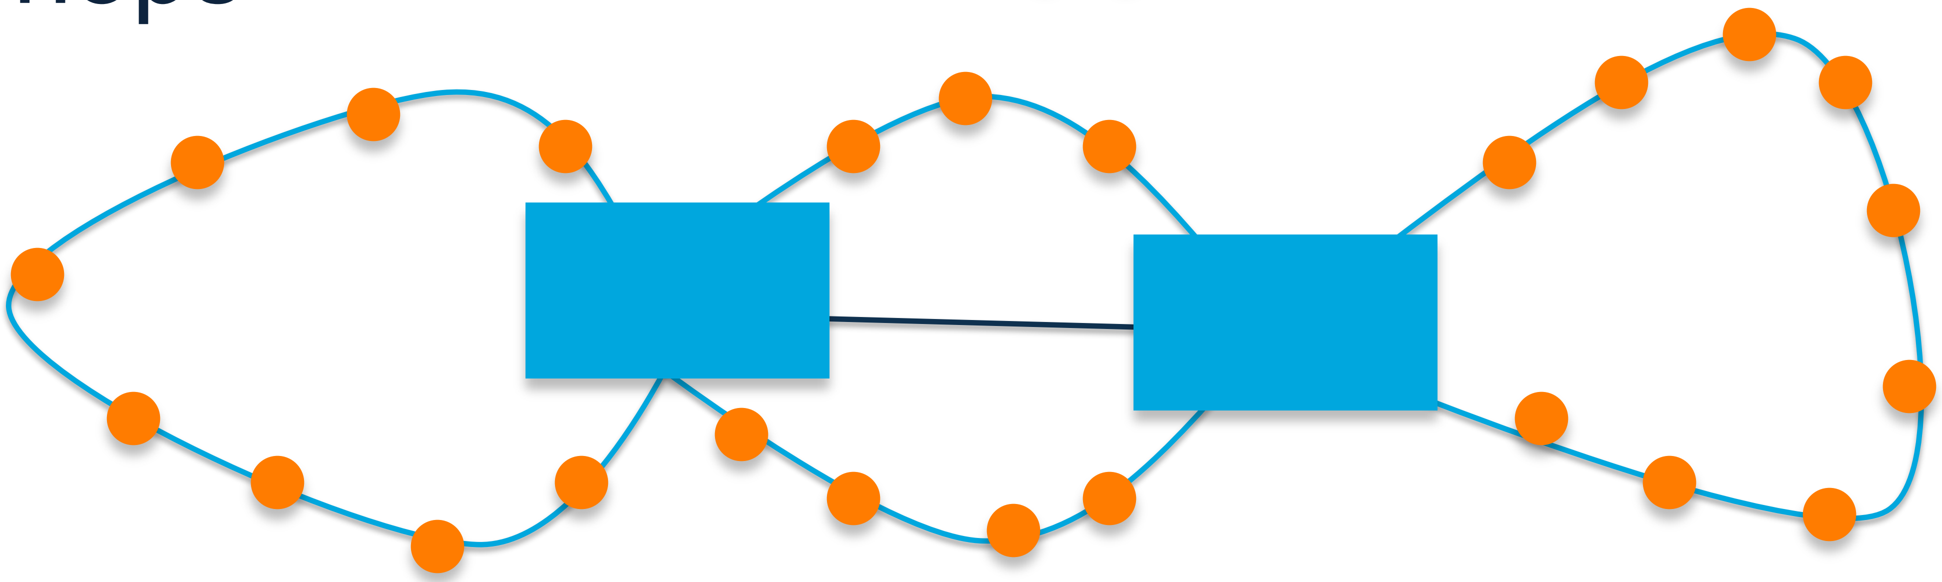

AVB control model

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 flows,
  - Having a third party set up the flows 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”

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

- **If** we assume that the customer **does** want a central controller, and
- **If** we assume a "few hops" network (few switches, directly-connected end systems),

  So that end systems don't need to participate in ISIS;

- **Then**, this model should work just fine.

  But, it's not the only one that will work just fine in this scenario.

Götz et al. 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.

Problems; all-ISIS model, 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 flow. 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.

Problems; all-ISIS model, 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.

Finn, Cummings, et al. model

Finn, Cummings 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 Central Network Controller 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 models to/from the CNC.

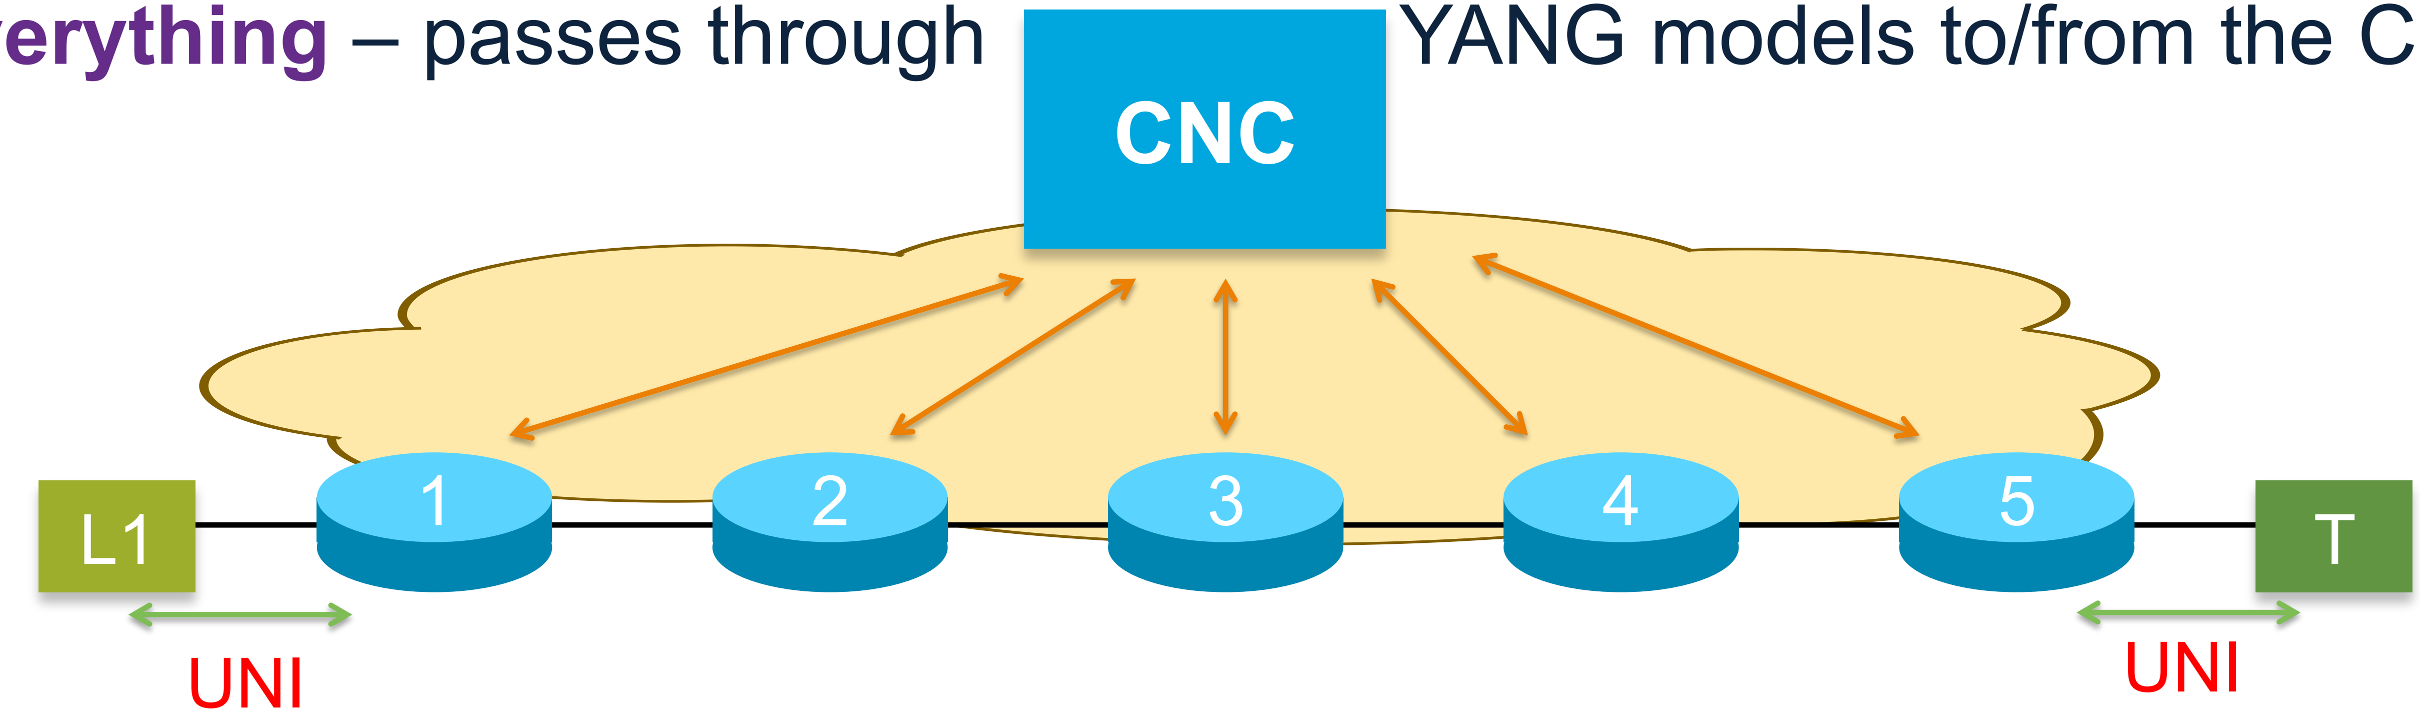

Comparing the models in the many-hops network

Comparing the models: many-hops network

- (The differences are smaller in the few-hops network, so we will not address that case.)
- What must be implemented if **no controller** is used?
- (**Both models are the same**: MSRP*+)

  Note that this precludes scheduled queues and nailed-down paths.

  Why are nailed down paths precluded? Because there are so many issues with what to optimize (non-critical bandwidth, equalized latencies, avoiding fate-sharing for certain flows, when to re-create paths after a topology change, etc.)

What does an end system need to know?

- In a “many hops” network, what does a given **end system** need to know?
- It needs to know **which flows to forward**.
- If using a shaper, with or without CQF, but not scheduled queues, it needs **four integers**:

  Input and output bandwidth on each of its two ports. (Input bandwidth is needed for error checking.)

- **THAT’S IT!**

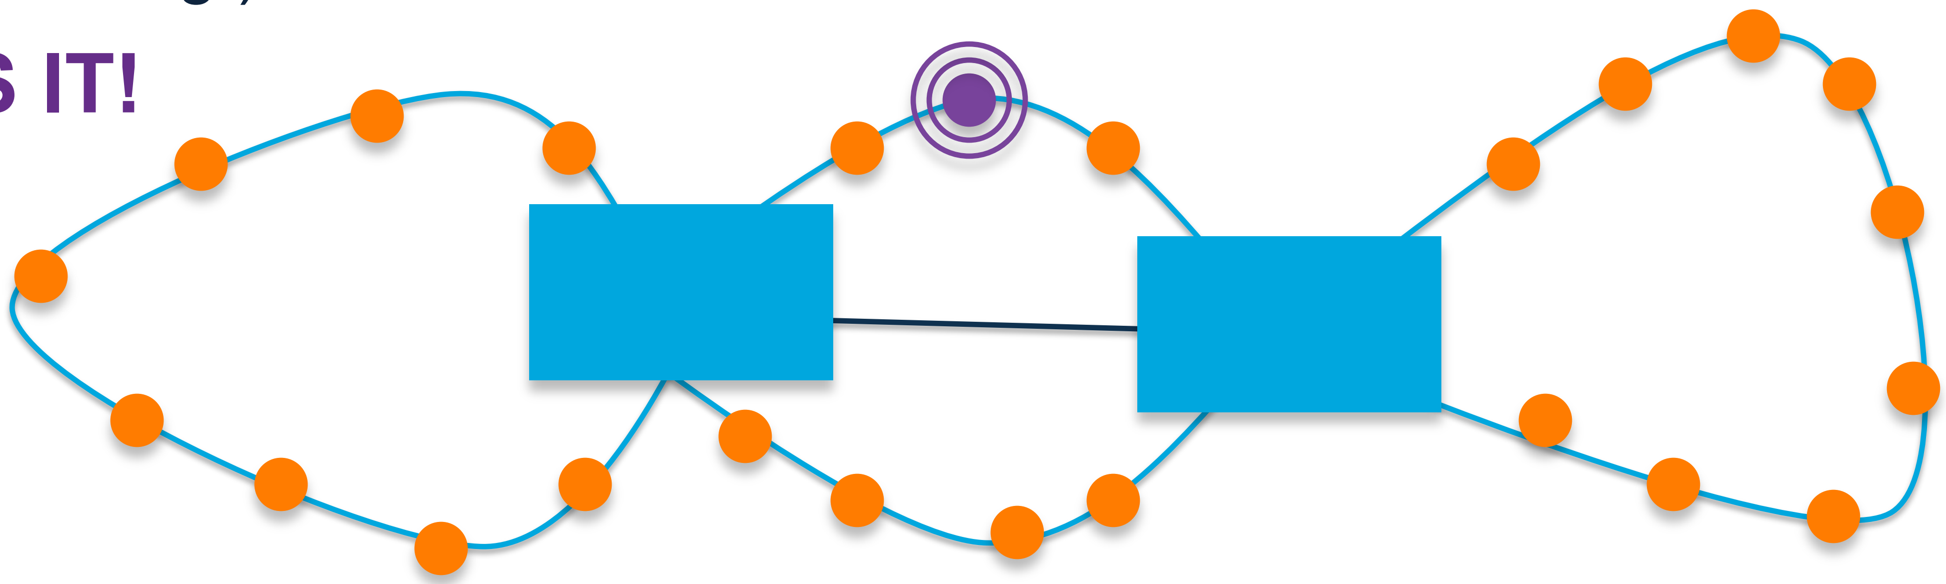

Comparing the models

- Databases maintained by a two-port end system

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

Which flows to pass or not
Four integers

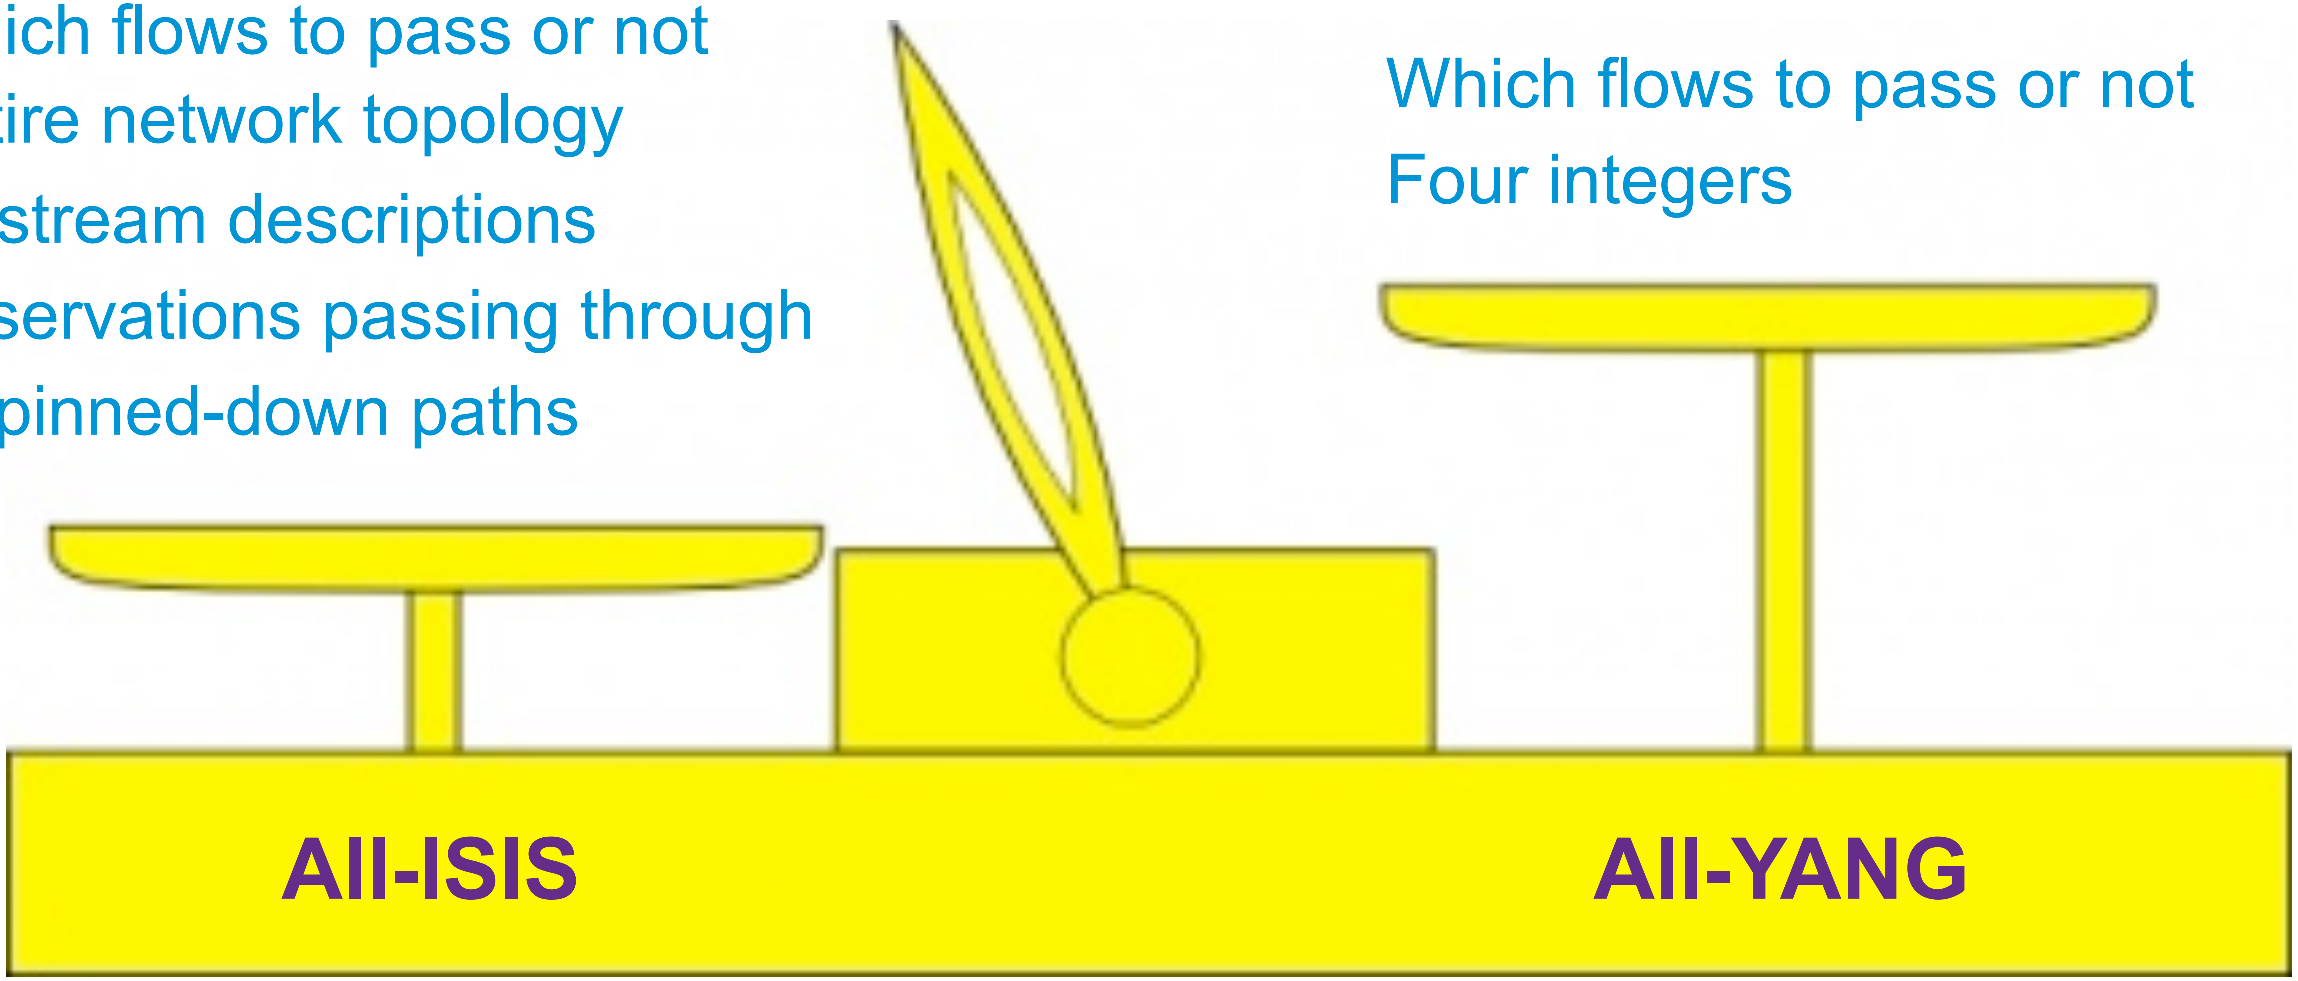

Comparing the models

- 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.

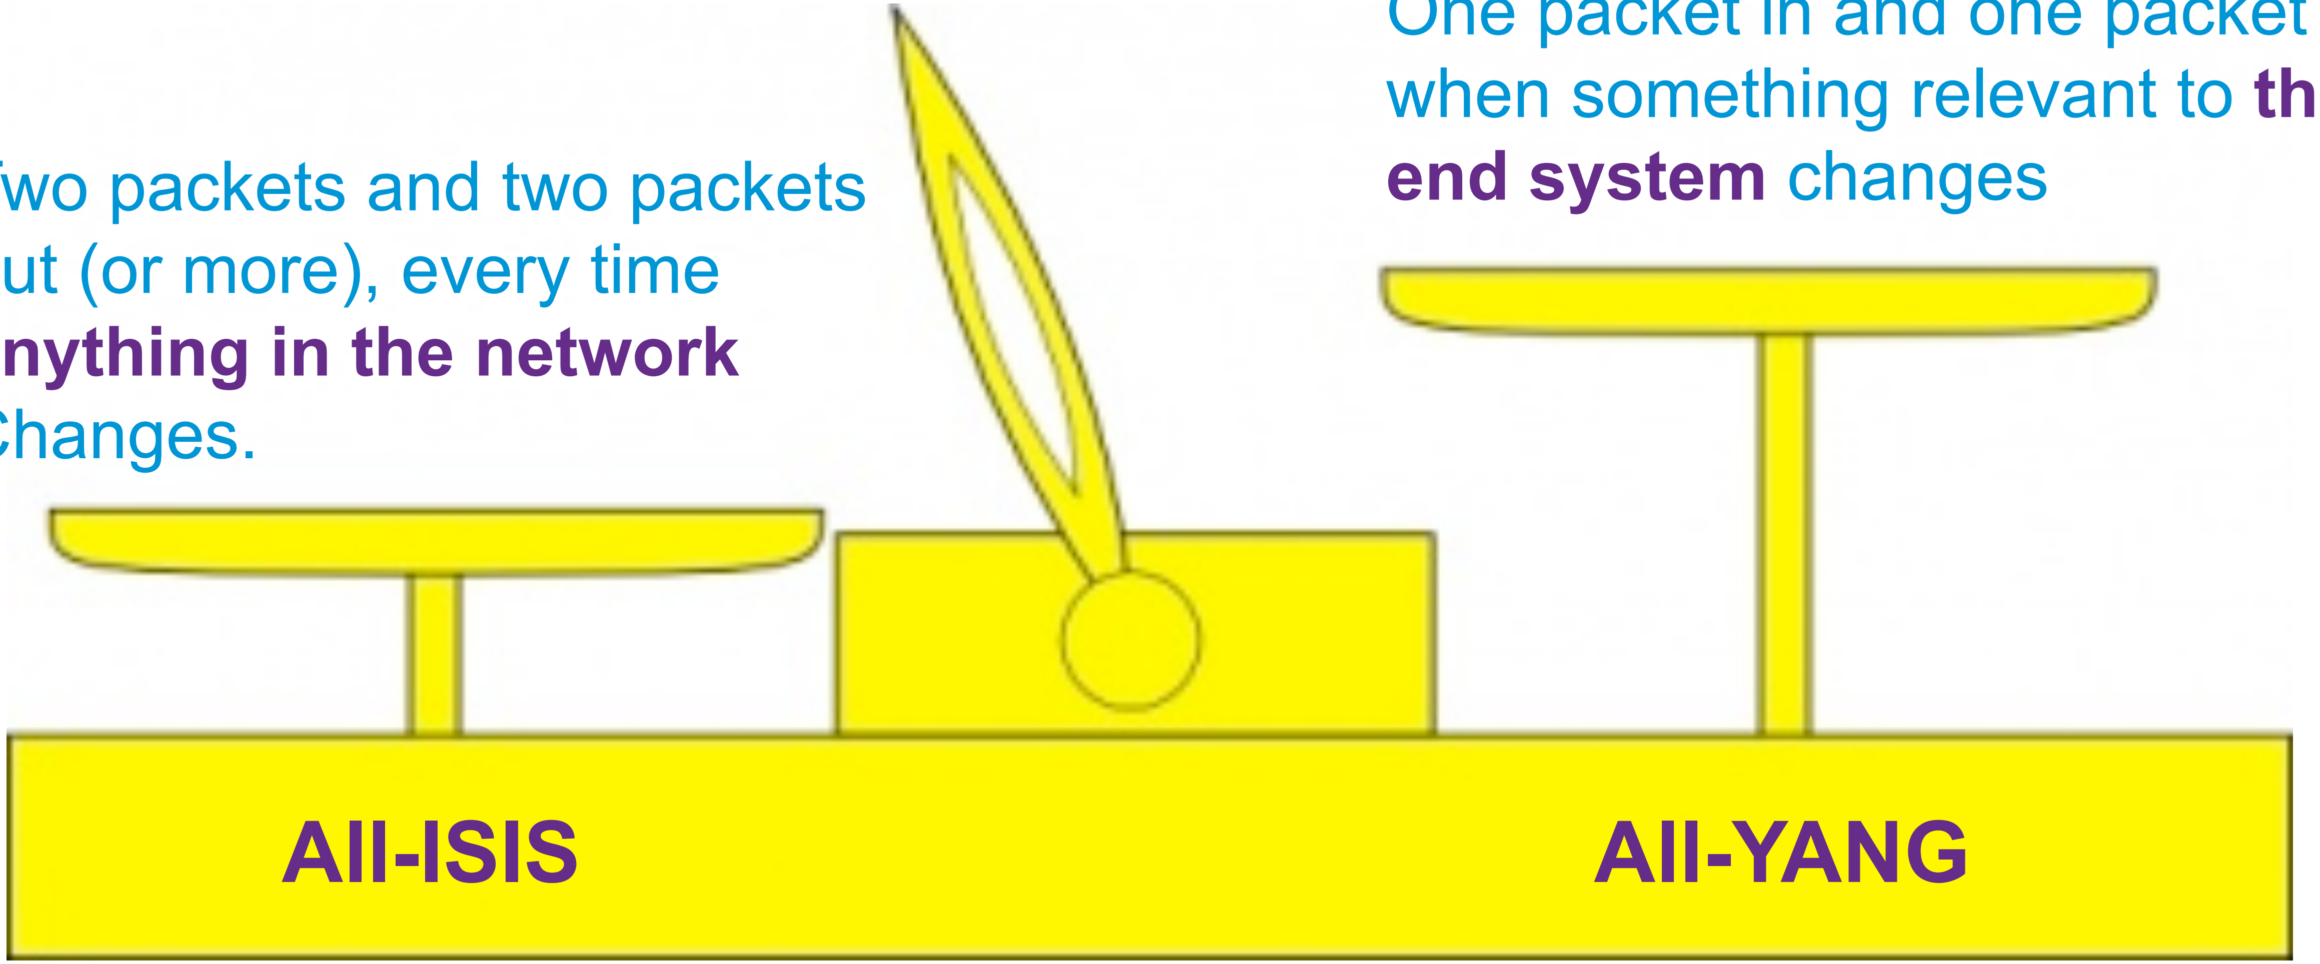

Comparing the models

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

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

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

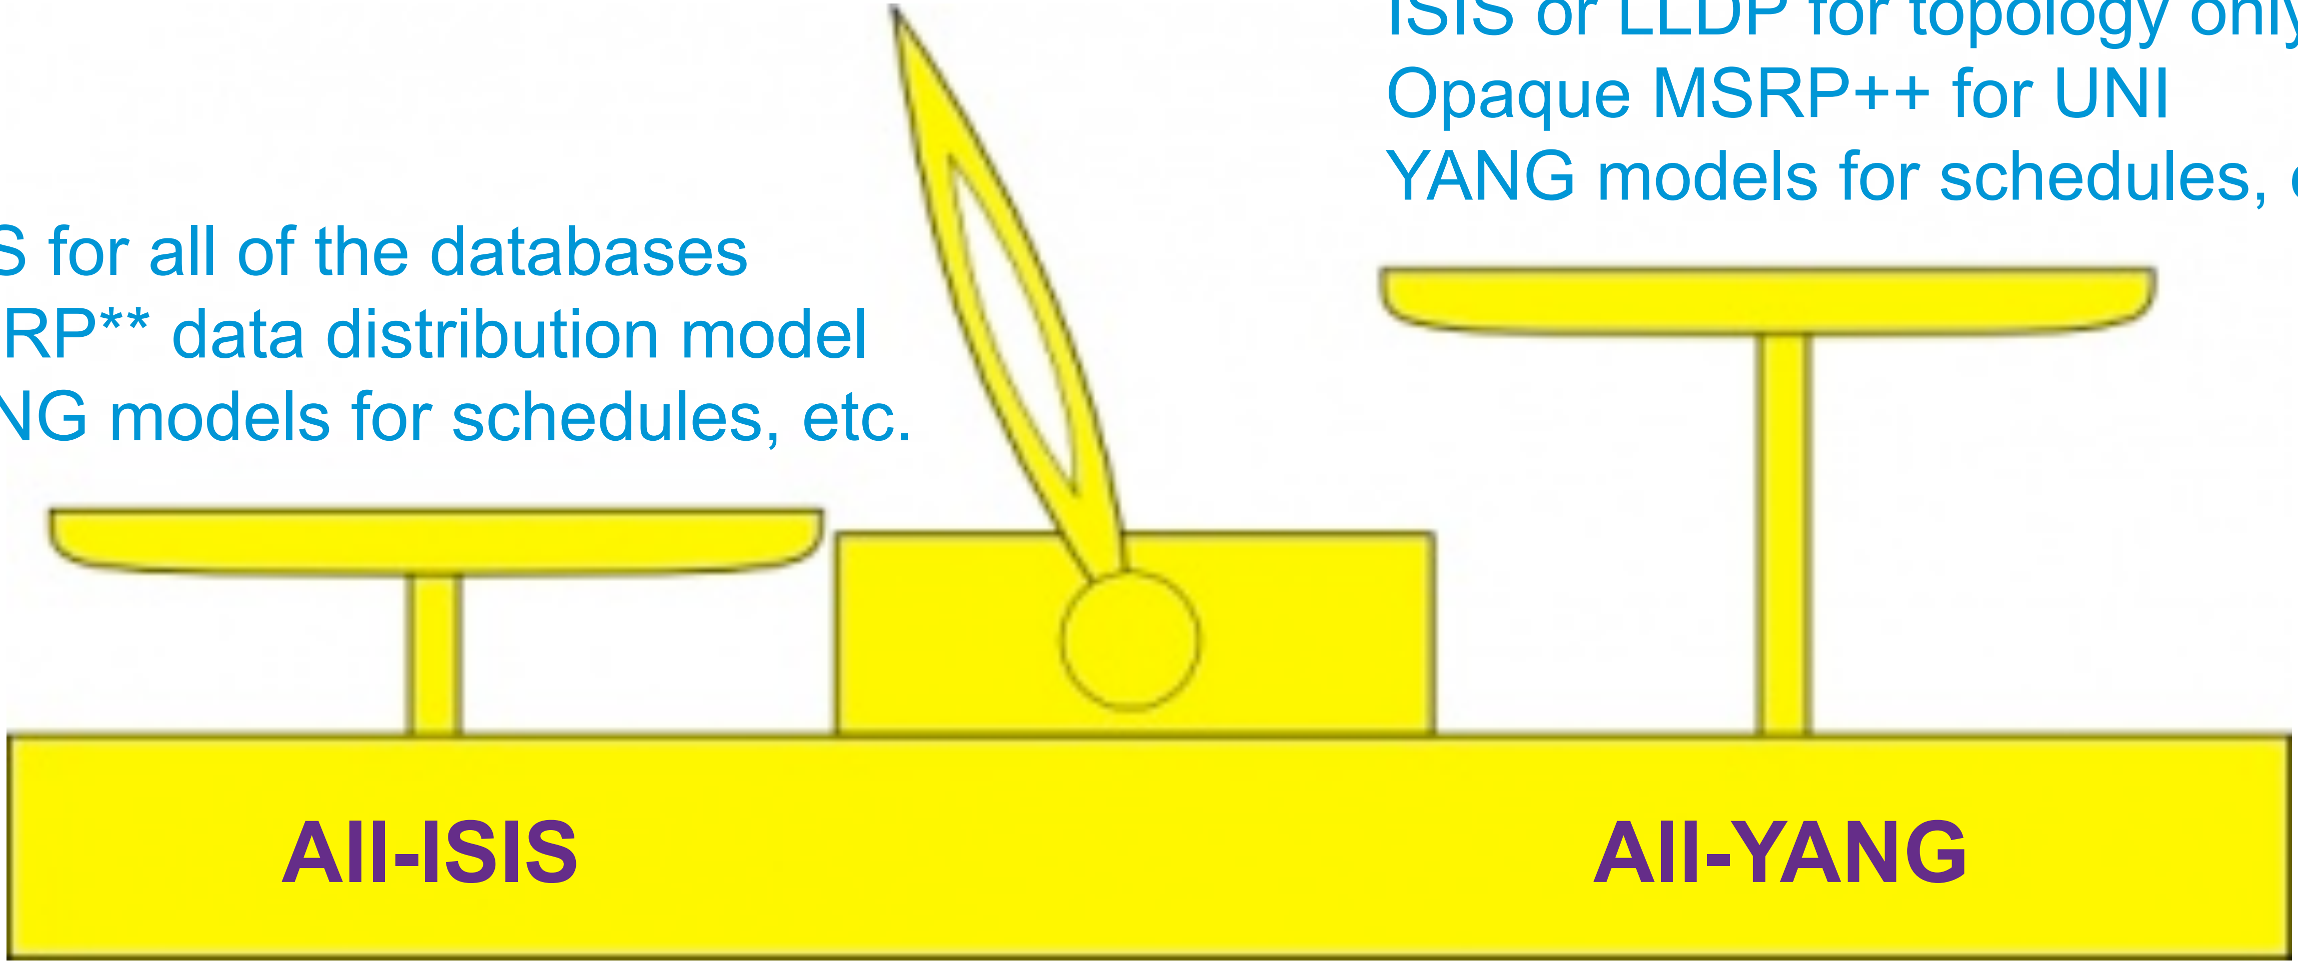

Comparing the models

- What else do we have to standardize?

UNI/MSRP++ with remote YANG model
LLDP TLV(s) to tie together ISIS and MSRP?
Augmented MSRP for compatibility.

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

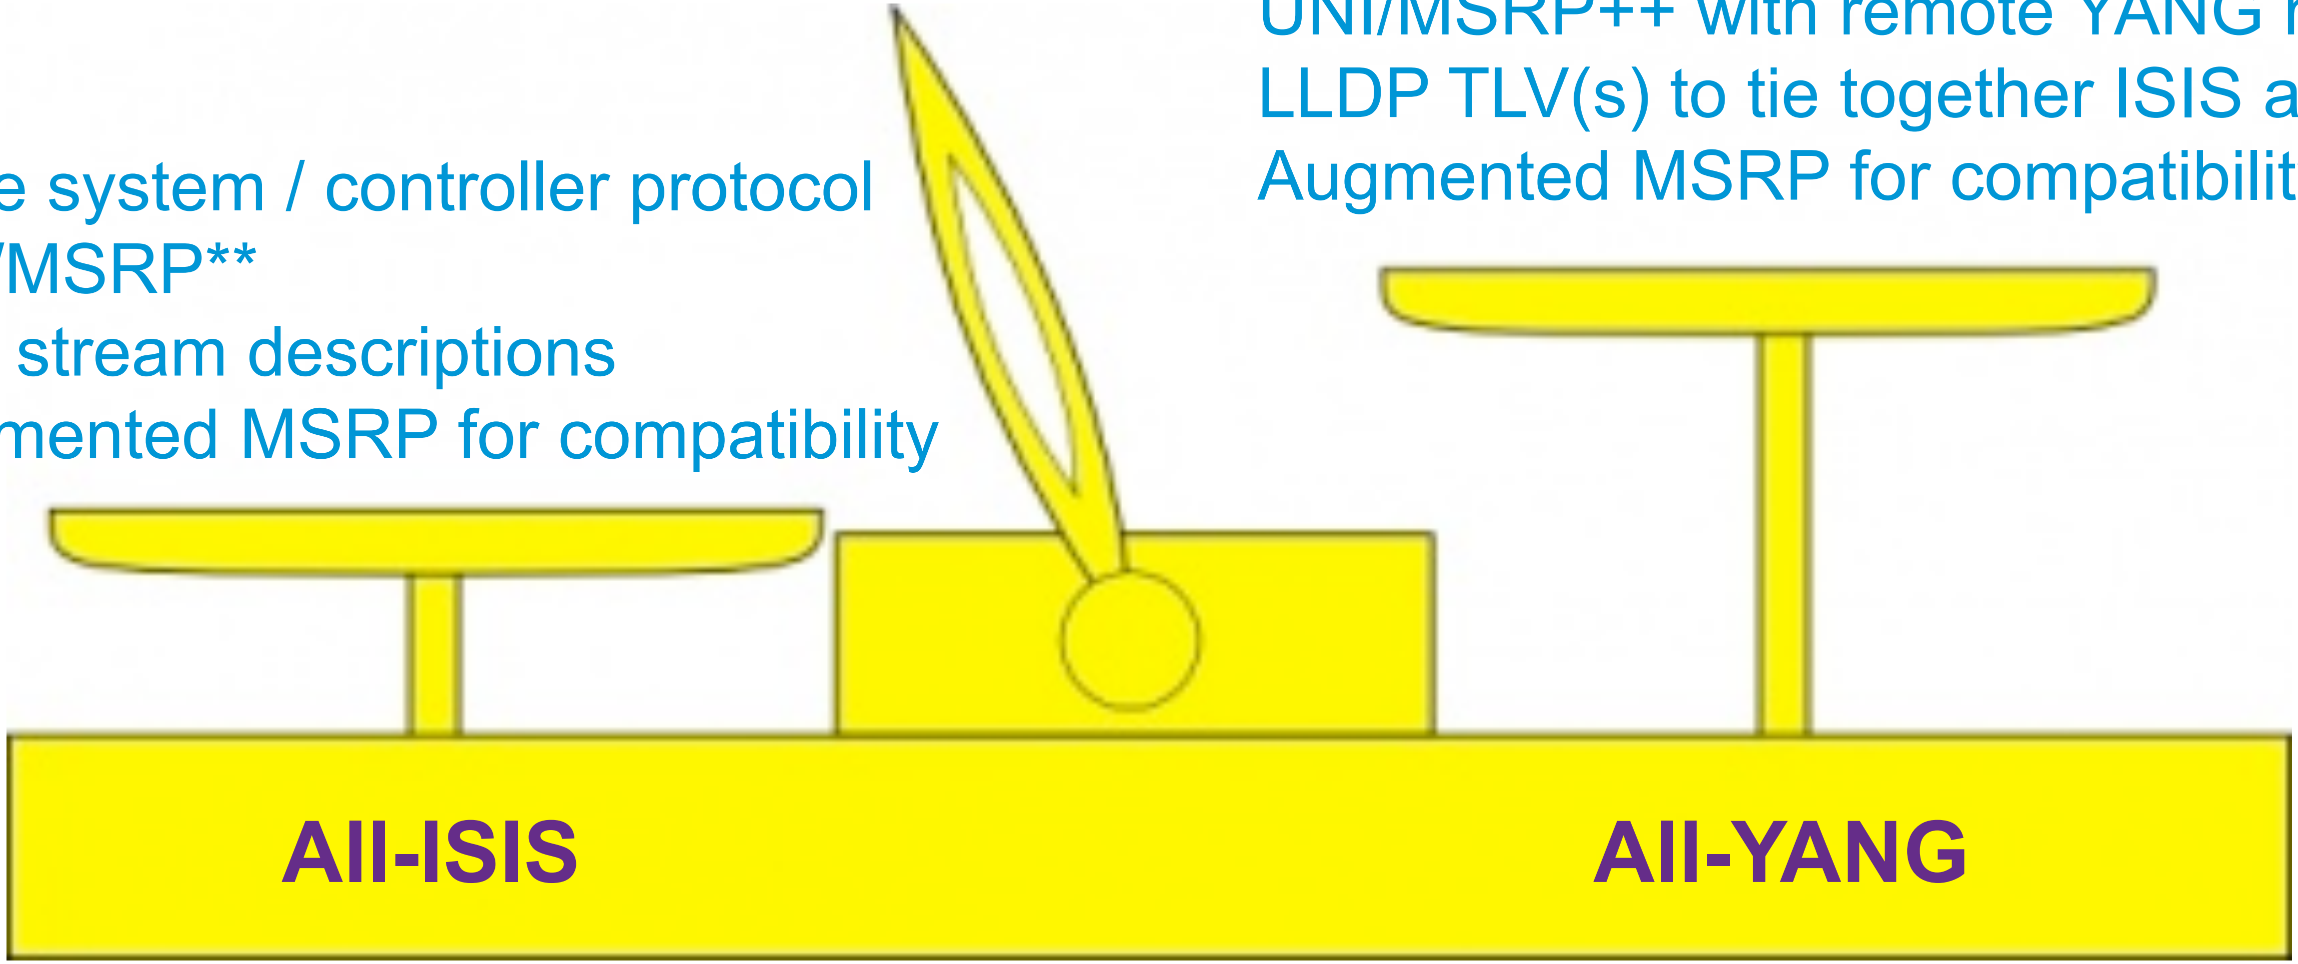

Comparing the models

- Compatibility with non-SPB topology protocols?

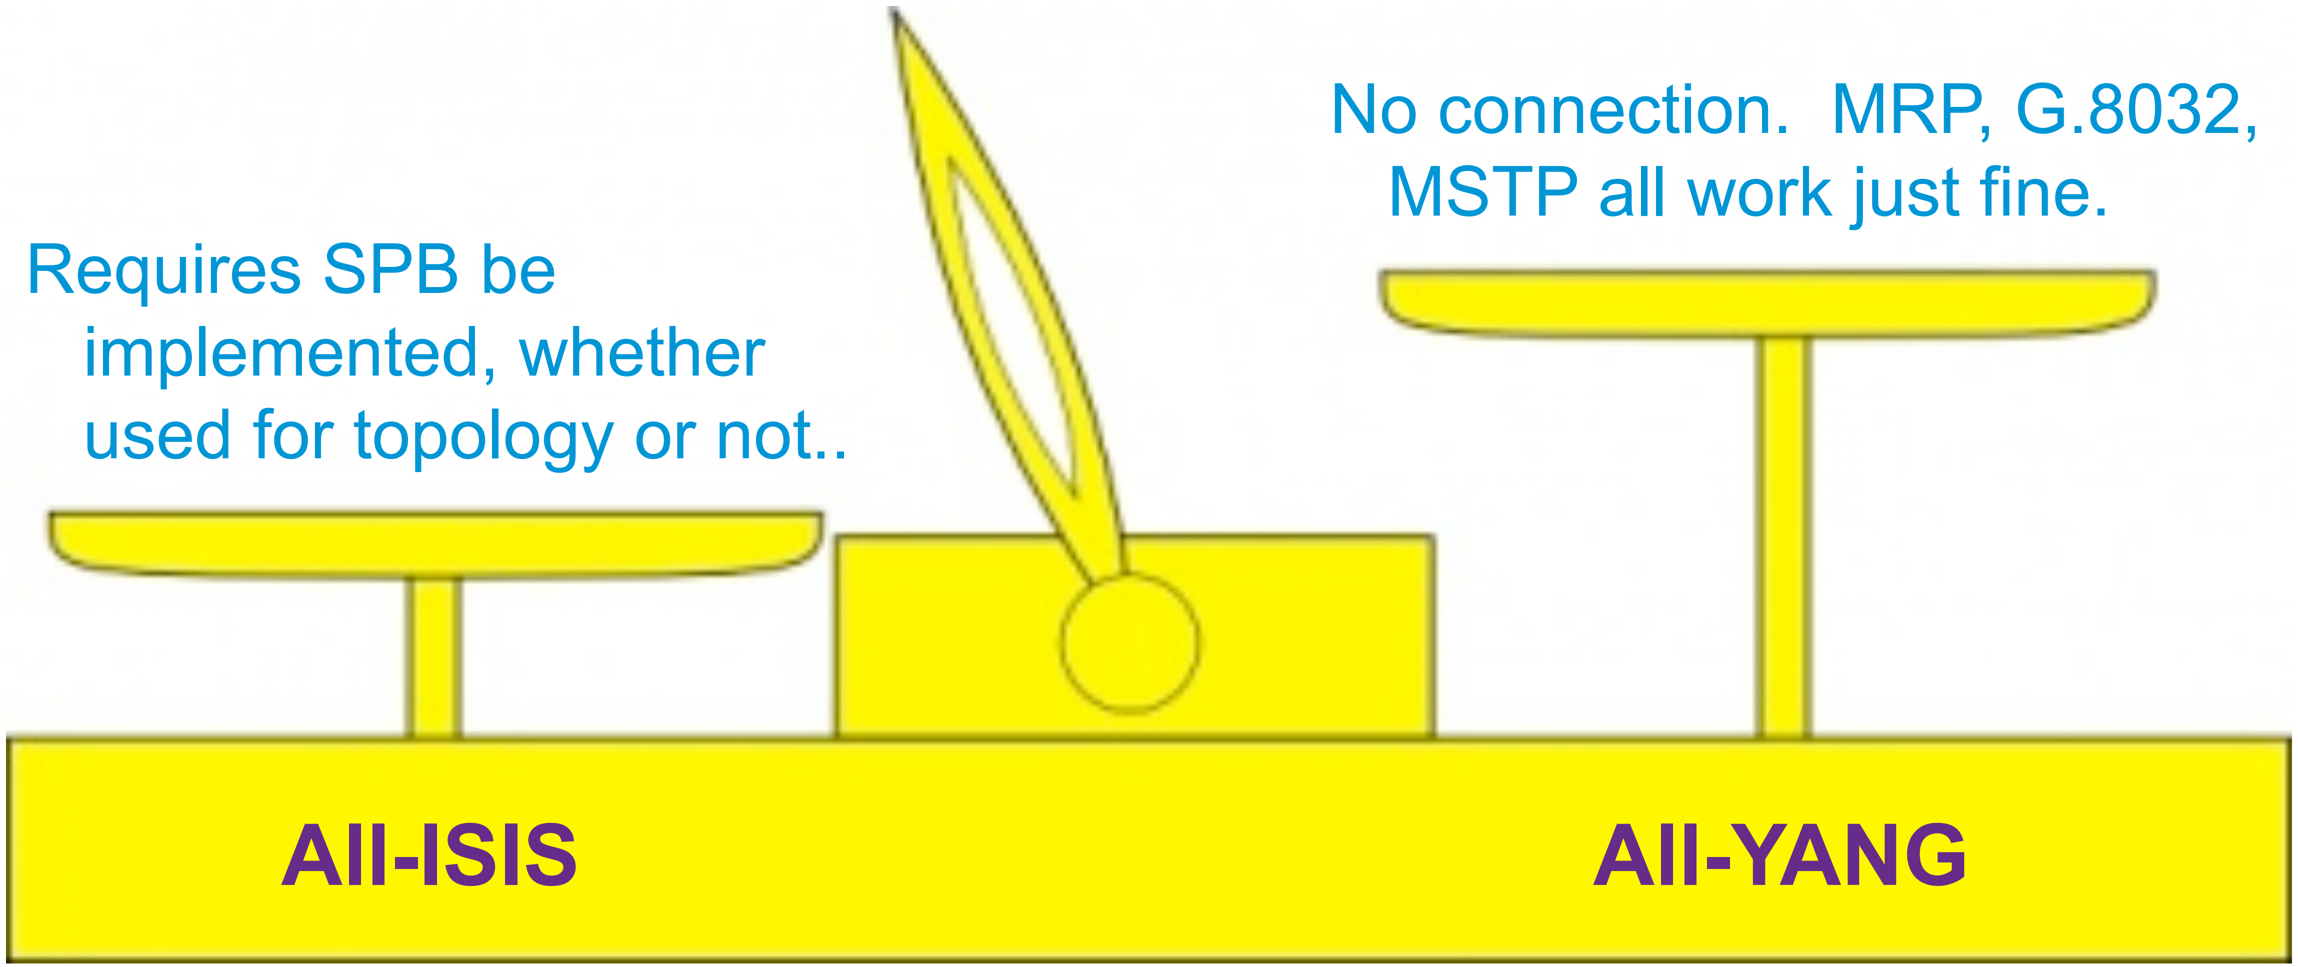

Comparing the models

- 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)

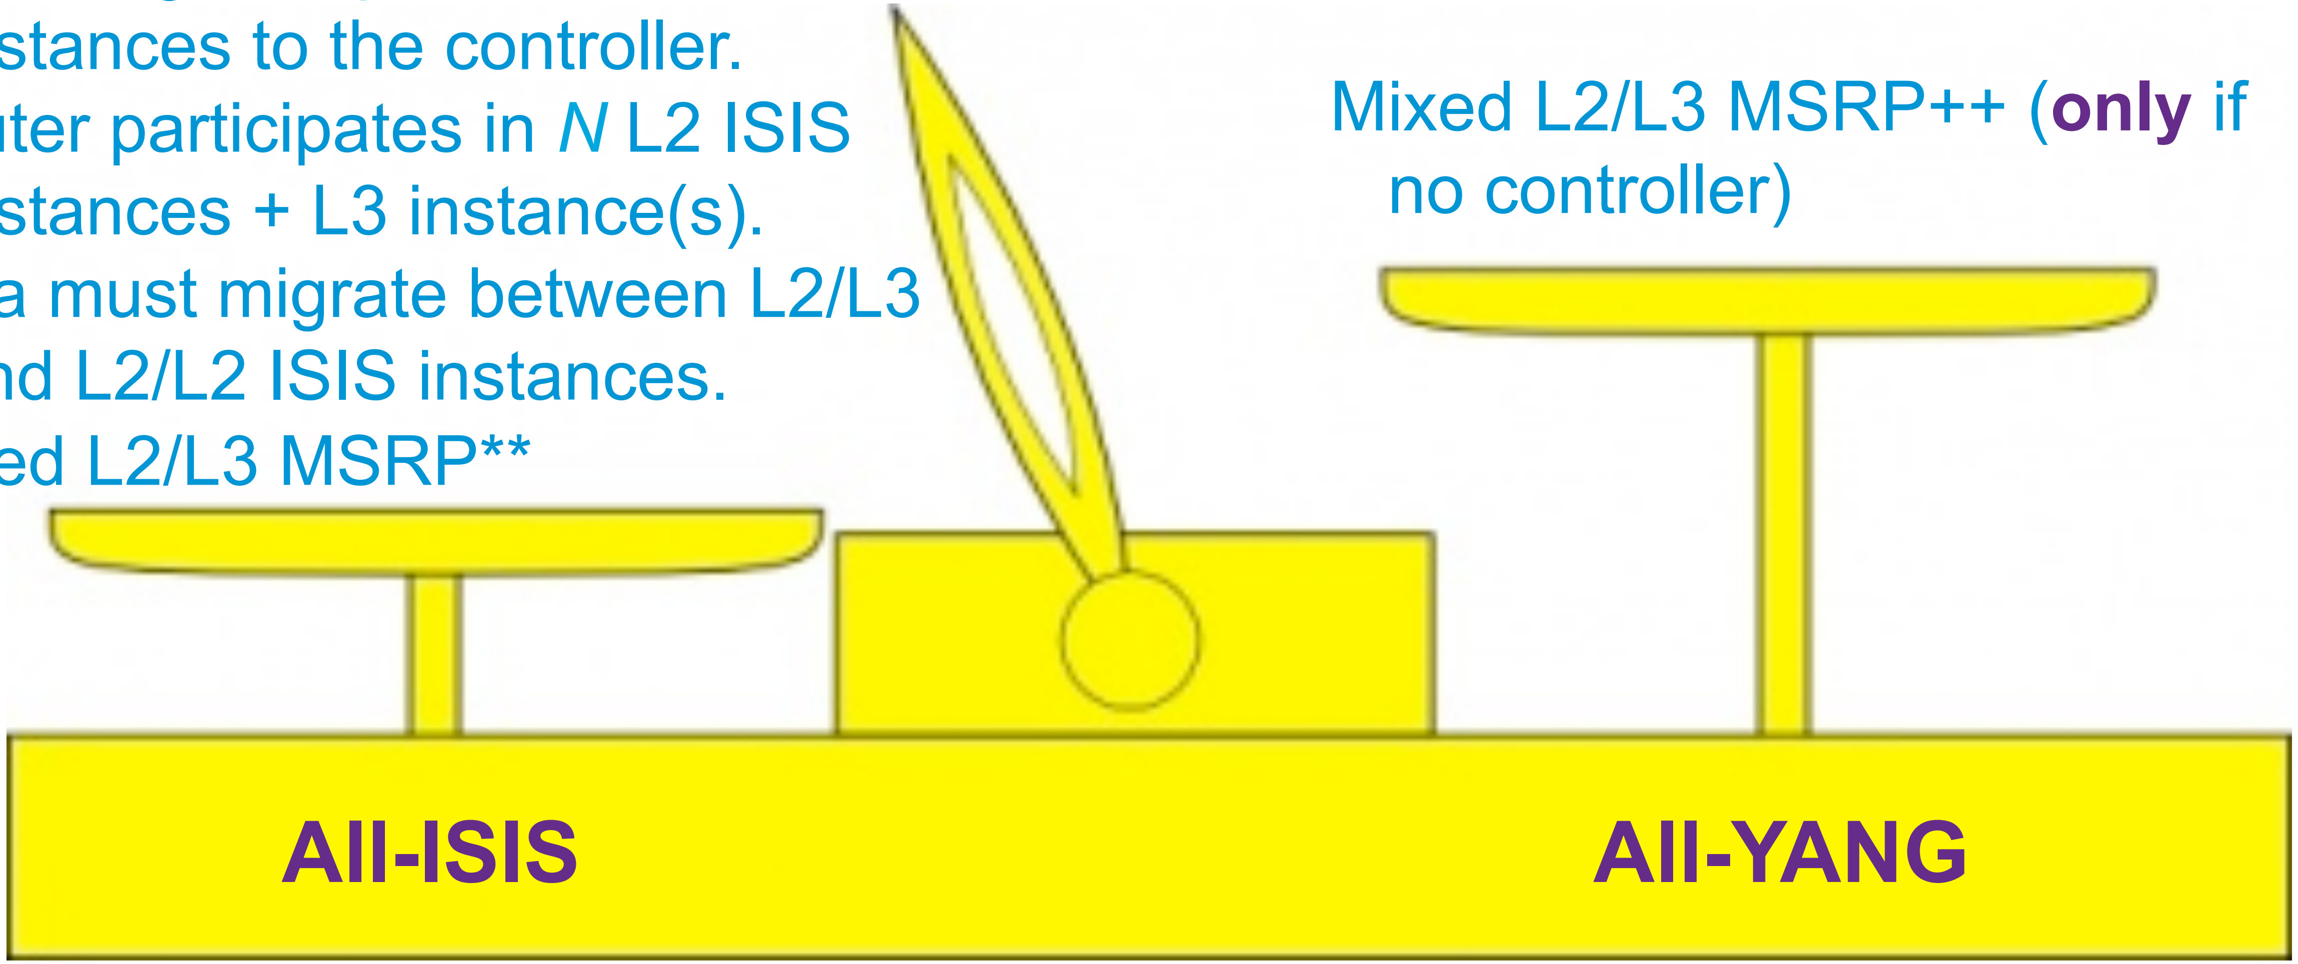

Simple 2-port Intermediate System (S2IS)

If we do the S2IS, we still have more work

- 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 all-ISIS solution, the S2IS still must maintain the entire stream description and multiple path databases; only the topology database is deleted. So, the difference between the models is still quite large.

Summary

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 Finn/Cummings 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 model.
- This author does not believe that the ISIS-based solution will achieve market acceptance, and will merely prompt other Standards Development Organizations to develop their own, much simpler, solutions.

Thank you.