802.1AX -- Link Aggregation:

Editor's Report: January 2019

Version 3

Stephen Haddock

January 17, 2019

802.1AX status

- AX-Rev-d1.0 went to Working Group ballot on December 7, 2018 through January 10, 2019.

- Ballot results:
  - 17 yes; 1 no; 15 abstain.
  - Ballot passes.
  - 103 comments from 4 commenters.

Easy comments

- 61 comments are very straightforward “Proposed Accept” or “Proposed Accept in Principle”. These will be changed to “Accept” or “Accept in Principle” at the end of the Hiroshima meeting without being discussed unless someone requests discussion during the comment resolution sessions:

  1-8, 10-21, 23-31, 39-41, 43-45,
  52-56, 59-62, 74, 77, 79, 80, 85,
  88-101

- On Monday Jan 14 at the Hiroshima meeting we discussed 14 comments:

  Resolved and closed: 9, 32, 33, 38, 47, 50, 51, 58, 102

  Resolved, but keep open until end of Hiroshima meeting: 34, 42, 48

  Still unresolved: 57, 71

discussion topics

- ✓ Solitary Definition: 103
- ✓ ISS Status parameters: 46, 57, 71
- ✓ compareDistributionAlgorithms: 63
- ✓ Renumbering links: 64
- ✓ Counters: 22, 65, 66, 67
- Expose Aggregator in DR-sublayer: 69
- Rename objects: 68, 70, 73
- DRNI System Identifier: 72, 82
- ✓ DRN: 49, 83
- ✓ Clause 8: 75, 76, 78, 86
- ✓ Restructure 9.4: 81, 84
- Sequence numbers: 87

Distributed Relay Number (DRN):

- DRN is a 2 bit value used for the following purposes:
  1. In intermediate states of Gateway and Aggregator selection (i.e. generation of Home/Nbor_Gateway/Aggregator_Mask variables) to indicate whether the selected Gateway/Aggregator is in the Home system, Nbor system, or neither.
  2. As a tie-breaker in the Gateway/Aggregator selection when all related configuration parameters are the same in both Home and Nbor.
  3. In the two MSBs of the operational key when a DR-sublayer is not paired with another DR-sublayer.
- I think we can eliminate DRN as a configured value:
  1. For this purpose, a locally significant value where ‘01’ = Home and ‘10’ = Nbor can be used.
  2. DRNI System Address can be used for the tie-breaker.
  3. Can use the MSB of the administrative Key value for this. Instead of requiring DRN be configured differently in each system for proper operation, require that the administrative key MSB be configured differently.

DRNI System Identifier: Link Aggregation Fundamentals

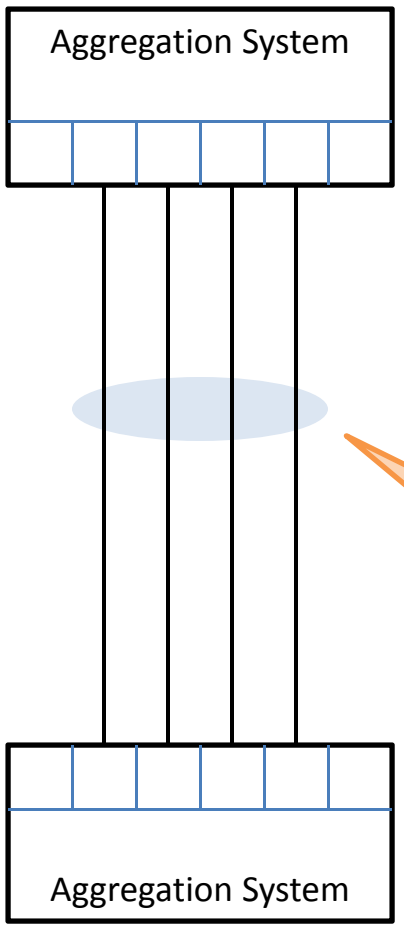

Each Aggregator and AggPort have:

- Actor_Priority
- Actor_System
- Actor_Admin_Key
- Actor_Oper_Key

Actor_Oper_Key typically gets the same value as Actor_Admin_Key, however it can be changed by local processes (e.g. Dynamic Key Allocation specified in 6.7.2 and Annex C). The Actor_Oper_Key is the source of the LACPDU's Actor_Key value.

To form a Link Aggregation Group, all AggPorts (and the selected Aggregator) at one end of the links have the same {Actor_Priority:Actor_System:Actor_Key} value, and all AggPorts (and the selected Aggregator) at the other end of the links have the same {Actor_Priority:Actor_System:Actor_Key} value.

read-write object

read-only object

DRNI System Identifier: DRNI Fundamentals 1

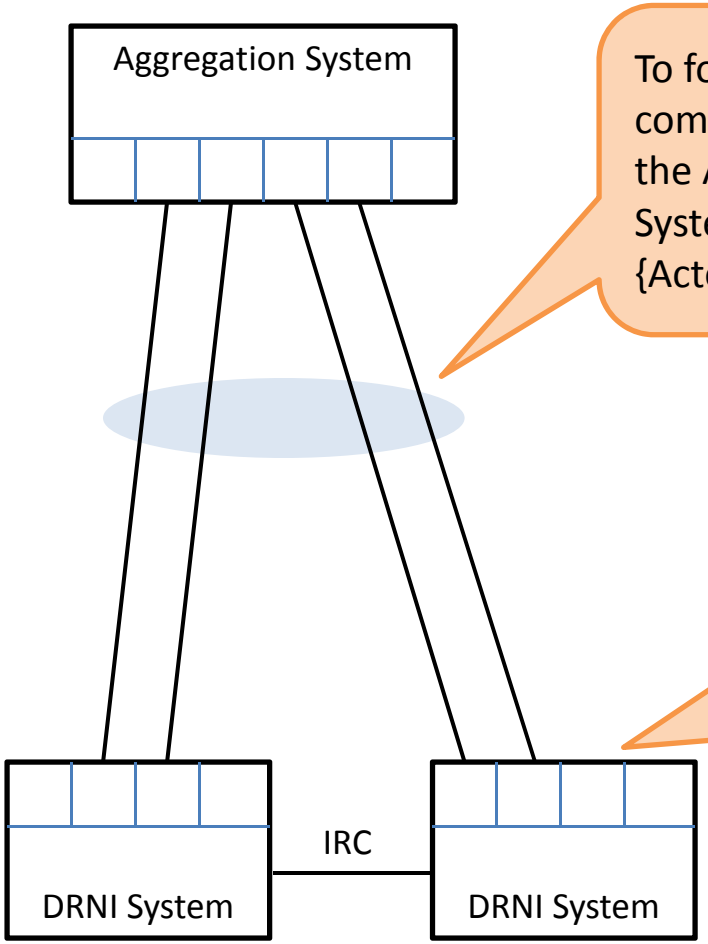

To form a DRNI LAG the DRNI Systems communicate via the IRC, and all AggPorts (and the Aggregator) of the DR-Sublayer in both DRNI Systems have the same {Actor_Priority:Actor_System:Actor_Key} value.

Because the AggPorts forming the DRNI LAG have the same Actor_System value in both DRNI Systems, they use a different Actor_System value than the AggPorts that are not configured to be included in DRNI LAG. (Do not want the AggPorts that are not intended to be in the DRNI LAG using the same Actor_System value in both DRNI Systems!)

DRNI System Identifier: DRNI Fundamentals 2

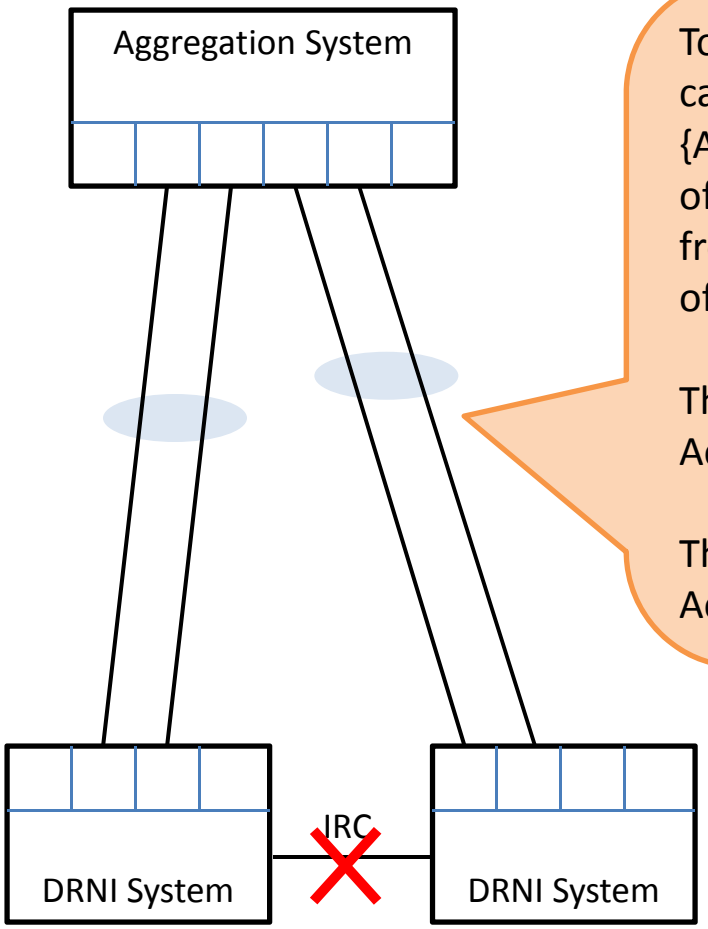

To prevent the formation of a DRNI LAG when the DRNI Systems cannot communicate via the IRC, the {Actor_Priority:Actor_System:Actor_Key} value of the AggPorts (and the Aggregator) in one DRNI System differs from the {Actor_Priority:Actor_System:Actor_Key} value of the AggPorts (and the Aggregator) in the other DRNI System.

This requires being able to dynamically change the Actor_Priority and/or Actor_System and/or Actor_Key value.

The mechanism to change the Actor_Key (i.e. separate Actor_Admin_Key and Actor_Oper_Key objects) already exists.

DRNI System Identifier: How many identifiers?

- DRCP operation needs two “globally unique” identifiers:
  1. Actor_System: Used as the Actor_System value for the DRNI AggPorts and Aggregator in both DRNI Systems when communicating via the IRC, and is otherwise unique within the network. (This is taken from the Aggregator’s Actor_System variable; there is not a duplicate variable in the DR-sublayer).
  2. DRNI_System_Address: Uniquely identifies the DRNI System within the network, and in particular is distinct from the identifier of the paired DRNI System. Presumably such an identifier already exists in the system (and is typically used as the Actor_System value for all AggPorts and Aggregators not configured to be part of a DRNI).

1. DRCP exchanges both ID-1 and ID-2 in DRCPDUs.
   - Actor_System is exchanged to verify both DRNI Systems using the same Actor_System value.
   - DRNI_System_Address is exchanged to be used as a tie-breaker (e.g. in Gateway and Aggregator selection) and to be sure a version 1 DRCP implementation will discard received version 2 DRCPDUs (position of DRNI_System in the version 2 DRCPDU guarantees this).

DRNI System Identifier: How many identifiers?

- DRCP operation needs two "globally unique" identifiers:
  1. Actor_System: Used as the Actor_System value for the DRNI AggPorts and Aggregator in both DRNI Systems when communicating via the IRC, and is otherwise unique within the network. (This is taken from the Aggregator's Actor_System variable; there is not a duplicate variable in the DR-sublayer).
  2. DRNI_System_Address: Uniquely identifies the DRNI System within the network, and in particular is distinct from the identifier of the paired DRNI System. Presumably such an identifier already exists in the system (and is typically used as the Actor_System value for all AggPorts and Aggregators not configured to be part of a DRNI).

1. DRCP exchang identifiers in DRCPDUs.

This is potentially quite confusing. Have "DRNI_System_Address" that is typically used as the Actor_System of the AggPorts that are not part of the DRNI, and is not used as the Actor_System of the AggPorts that are part of the DRNI. This just seems backwards, but we are locked into it since the Actor_System of the DRNI LAG has to be the same in both DRNI Systems, whereas the DRNI_System_Address is by definition different in each DRNI System.

DRNI System Identifier:

What should change when gain/lose IRC connectivity?

- Mick's comment #82 (in the comment, but not the suggested remedy) seems to suggest changing the LACPDU Actor_System value, rather than the Actor_Key value, when gain or lose communication over the IRC.
- We originally chose to change Actor_Key because:
  1. Didn't want to change the "address" portion of the system identifier when an IRC or DRNI System failed.
     - This was arguably misguided because it confuses the Actor_System function as an identifier with using it as an address.
  2. Already have the Actor_Admin_Key and Actor_Oper_Key variables that allow dynamically changing the key value.
     - If really prefer to change the Actor_System then can rename the current Actor_System to Actor_Admin_System and add Actor_Oper_System. When DRNI is not supported the Actor_Oper_System would always have the same value as Actor_Admin_System and thus need not be implemented as a separate object.

DRNI System Identifier:
If choose to change Actor_System

Each Aggregator and AggPort have:

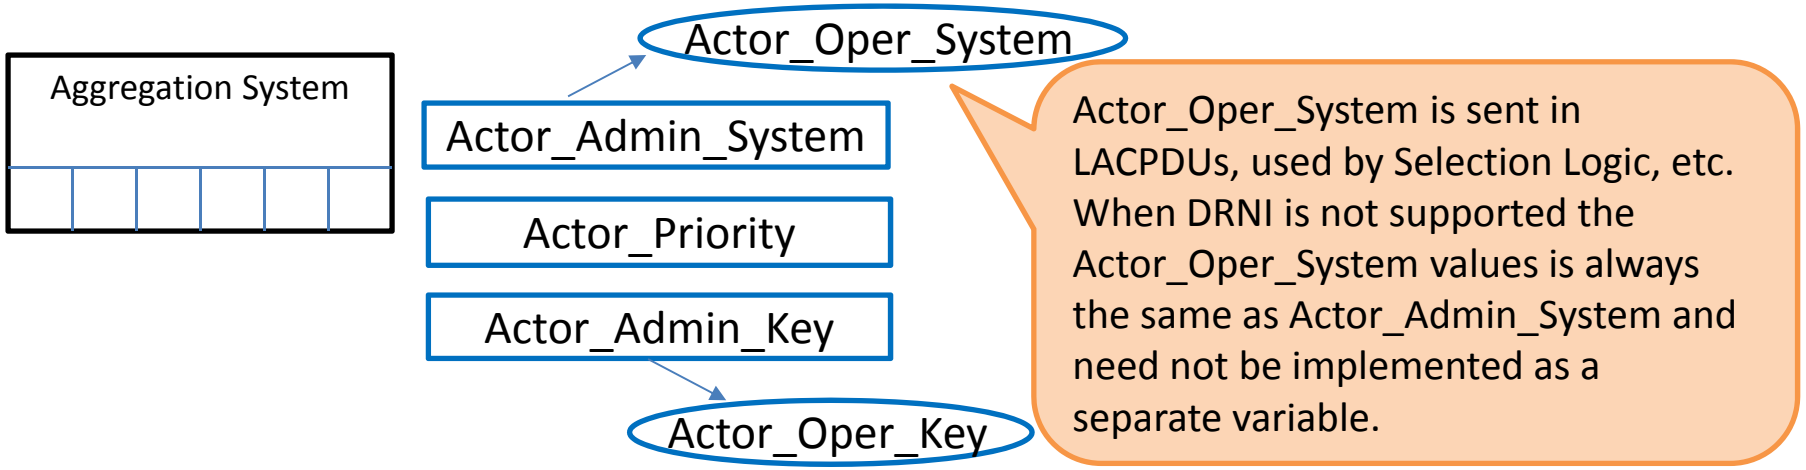

DRNI System Identifier:
If choose to change Actor_System

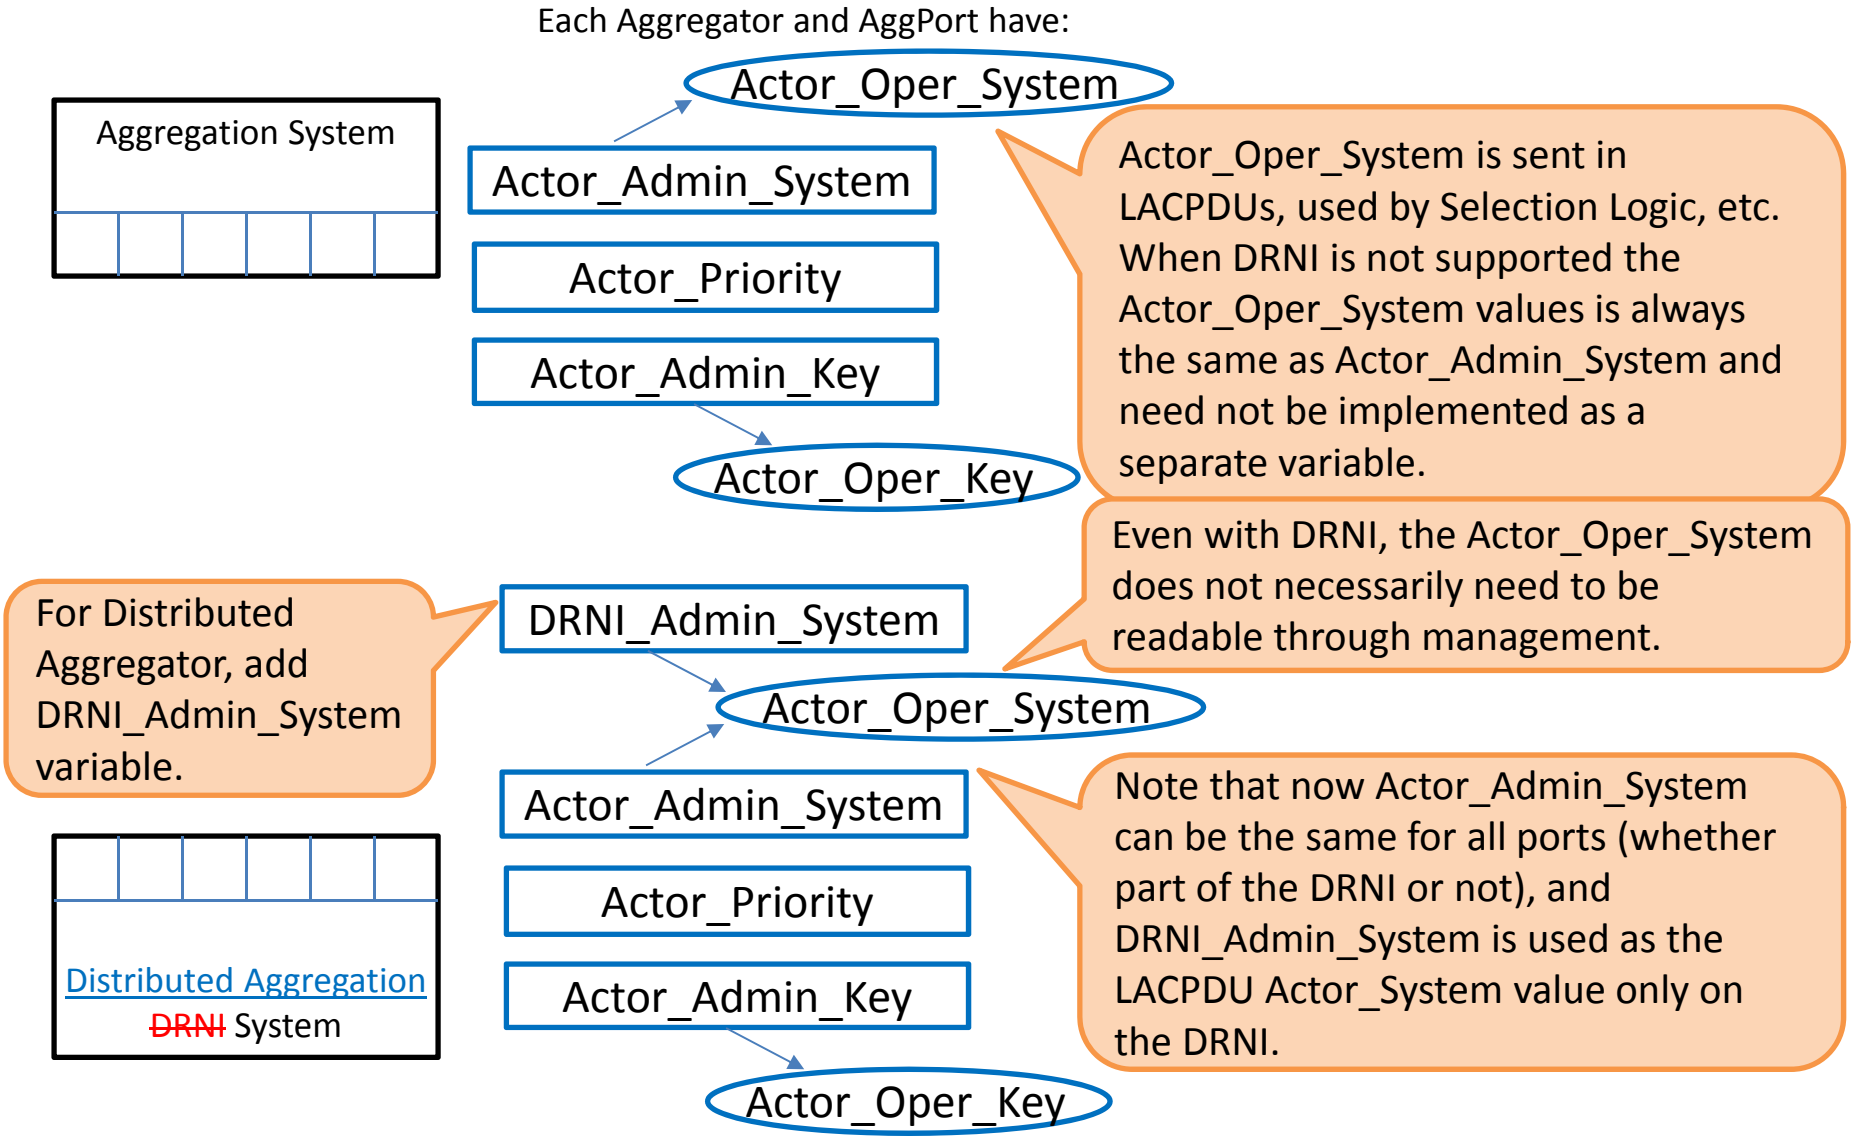

DRNI System Identifier: If choose to change Actor_System

Part of Mick's comment is that we could specify that when the DRNI is formed, the Actor_Oper_System always becomes the highest {Actor_Priority:Actor_Admin_System} value of the paired systems. Then don't need DRNI_Admin_System object at all. Consequence is that we lose the configuration checking provided by verifying that both systems have the same DRNI_Admin_System. Will form a DRNI when any two DRNI Systems get connected via their IRPs. If really want this then should also set Actor_Oper_Key to Actor_Admin_Key of highest priority system.

Can be, but is not necessarily, the same as the Actor_Admin_System value for one of the DRNI Systems.

For Distributed Aggregator, add DRNI_Admin_System variable.

DRNI_Admin_System

Actor_Oper_System

Actor_Admin_System

Actor_Priority

Actor_Admin_Key

Actor_Oper_Key

Note that now Actor_Admin_System can be the same for all ports (whether part of the DRNI or not), and DRNI_Admin_System is used as the LACPDU Actor_System value only on the DRNI.

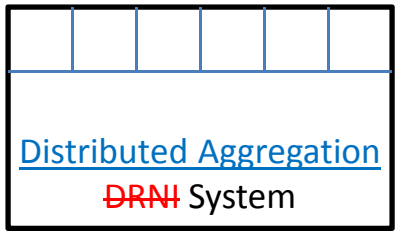

Sequence Numbers

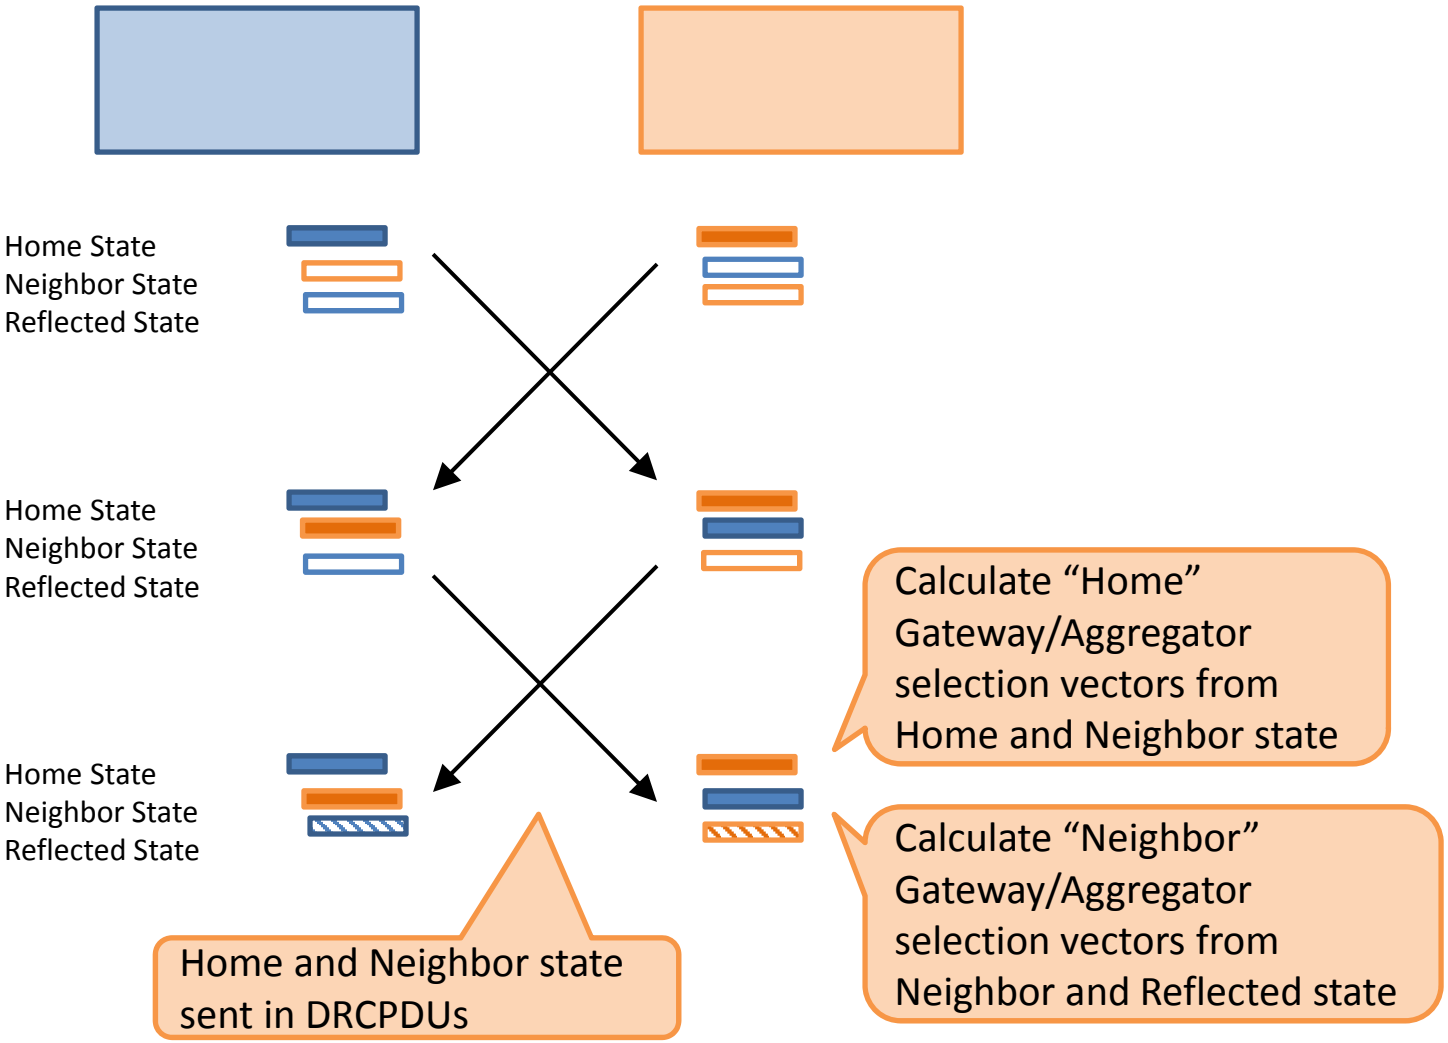

Back up slides

Thank You