aruba
a Hewlett Packard Enterprise company

LACP AND DISTRIBUTED LACP

10:00 GMT | 11:00 CEST | 13:00 GST
August 28th , 2018

Manoj Ramalingam, Aruba ERT

Engineering Resolution Team

AIRHEADS COMMUNITY

Need for LACP (802.3ad)

By Simply adding two links to the switch spanning tree always considers it as a the redundant link and blocks one of the port

Link aggregation solves this as it considers its physical ports as logical link and not physical link

The main goal of a link aggregation protocol is to increase bandwidth. It uses one or more links together to form a link aggregation group.

RSTP/MSTP sees the link aggregation as an interface, and the link aggregation as a whole has a forwarding or blocking status.

LACP

The main goal of a link aggregation protocol is to increase bandwidth. Link Aggregation allows one or more links to be aggregated together to form a Link Aggregation Group, such that a client can treat the Link Aggregation Group as if it were a single link.

The LACP protocol sends periodic packets containing actor and partner information. As long as the partner information remains the same, there is no LACP state change. When the partner information change, LACP recognizes the change and propagates that change throughout the LACP state machines

The ArubaOS switch refers to link aggregations as “trunks”

The switch learns MAC addresses on the link aggregation not on an individual interface.

It supports 8 link per Trunk group in a switch

No hardware limitation as most of the Aruba switch platform supports this feature

LACP is not allowed on ports configured for 802.1X authenticator operation. If you configure port security on a port on which LACP (active or passive) is configured, the switch removes the LACP configuration

How does it work?

It uses LACPDU messages to establish the link aggregations

Ensures links are connected to same peer using system ID and operation key

Ensure comparability of other settings

Manages adding and removing link

LACP work only on Full Duplex ports and ports that are operating at the same data rate

Allows you to assign up to eight physical links to one logical link (trunk) that functions as a single’

When it needs to flood the traffic over a link aggregation, it chooses just one link for the transmission.
This behaviour ensures that the link aggregation does not create a loop

LACPDU

- This protocol requires switches to exchange messages, called LACP data units (DUs), (LACPDU)
- These messages include a system ID, which uniquely identifies the switch.
- Each LACP peer checks the messages that arrive on each link in the aggregation.
- A link aggregation establishes correctly only if the incoming system ID and operational key matches on all the links
- The system ID identifies the switch operational key and the link aggregation group.
- Misconfigured link aggregations LACP would have detected that the system ID for the peer on the two links was different
- This will ensure that all links are connected to the same or other peer and also that all links are connected to the same link aggregation on that peer
- LACPDU will be sent out by the switch port every 1 sec interval
- LACP slow rate will send out LACPDU every 30 sec

Link aggregation should have matching settings for:

Duplex mode (full or half)
Link speed
Media (copper or fiber)
only then It allow compatible links to join the aggregation

LACP-modes

LACP operational modes:

1. Static
2. Dynamic

- Static
  - LACP is specified when trk interface created
  - Uses Active/Active mode—Both sides actively send LACP messages
  - Typically recommended mode
- Dynamic
  - No trk interface is created
  - Instead LACP is enabled on physical interfaces
  - Uses one of two modes:
    - Active/Active
    - Active/Passive—Passive side waits
  - Allows extra links to be on standby
  - Does not support non-default VLAN assignments or various other settings on the aggregation

LACP

Static:

You set up a static mode LACP link aggregation much as you do a manual link aggregation

Physical interfaces that are assigned to that link aggregation then operate in active LACP state.

Both sides of the link send LACPDUs to each other and set up the link

In Link/Duplex setting it do not check for compatibility on links in a manual link aggregation however its not recommended as it would lead to connectivity issues

Dynamic:

Enable LACP on the switch interfaces that you want to function as part of a link aggregation
But you do not actually create a link aggregation (trk) interface

The lacp key option provides the ability to control dynamic trunk configuration. Ports with the same key will be aggregated as a single trunk.

Active—Transmits LACPDUs to advertise that it can create an aggregated links
Passive—Listens for LACPDUs and responds with LACPDUs only after it receives one from an active port

Continued …

As long as one or both sides of each link are in active mode, the link aggregation automatically establishes. (If both sides are in a passive state, then neither will initiate the exchange process.)

When you use dynamic LACP, ArubaOS switches can have standby links beyond the maximum allowed in the aggregation. For example, if the switch supports only eight links, you might be able to add 12.

Links over the maximum are standby links, which cannot be used to forward traffic, However, if a link fails, LACP automatically adds one of the standby links to the link aggregation

```
switch# show lacp
LACP
LACP Trunk Port LACP Admin Oper
Port Enabled Group Status Partner Status Key Key
---- ------- ----- ------ ------- ------ ---- -----
Al Active Dyn1 Up Yes Success 100 100
A2 Active Dyn1 Up Yes Success 100 100
A3 Active Dyn1 Up Yes Success 100 100
A4 Active Dyn1 Up Yes Success 100 100
A5 Active Dyn1 Up Yes Success 100 100
A6 Active Dyn1 Up Yes Success 100 100
A7 Active Dyn1 Up Yes Success 100 100
A8 Active Dyn1 Up Yes Success 100 100
A9 Active Dyn1 Standby Yes Success 100 100
```

A dynamic LACP link aggregation does not have a logical interface in the switch configuration. Therefore, you cannot assign non-default STP, VLAN, and other such settings to it. For this reason, static LACP is typical

Dynamic/static LACP interoperation: A port configured for dynamic LACP can properly interoperate with a port configured for static (TrkX) LACP, but any ports configured as standby LACP links are ignored.

LACP Active and Passive

Active:

The port automatically sends LACP protocol packets

Passive:

The port does not automatically send LACP protocol packets and responds only if it receives LACP protocol packets from the opposite device.

A link having either two active LACP ports or one active port and one passive port can perform dynamic LACP trunking.

A link having two passive LACP ports does not perform LACP trunking because both ports are waiting for an LACP protocol packet from the opposite device.

LACP status

port status:

Up: The port has an active LACP link and is not blocked or in standby mode.

Down: The port is enabled, but an LACP link is not established. This can indicate, for example, a port that is not connected to the network or a speed mismatch between a pair of linked ports.

Disabled: The port cannot carry traffic.

Blocked: LACP, Spanning Tree has blocked the port. (The port is not in LACP standby mode.) This may be caused by a (brief) trunk negotiation or a configuration error, such as differing port speeds on the same link or trying to connect the switch to more trunks than it can support.

Standby:
The port is configured for dynamic LACP trunking to another device, but the maximum number of ports for the dynamic trunk to that device has already been reached on either the switch or the other device. This port will remain in reserve, or "standby" unless LACP detects that another, active link in the trunk has become disabled, blocked, or down. In this case, LACP automatically assigns a standby port, if available, to replace the failed port

LACP Partner: Yes: LACP is enabled on both ends of the link. No, with incorrect config.

LACP Status:

Success: LACP is enabled on the port, detects and synchronizes with a device on the other end of the link, and can move traffic across the link.

Failure: LACP is enabled on a port and detects a device on the other end of the link, but is not able to synchronize with this device, and therefore is not able to send LACP packets across the link.

This can be caused, for example, by an intervening device on the link (such as a hub), a bad hardware connection, or if the LACP operation on the opposite device does not comply with the IEEE 802.3ad standard.

LACP-Load sharing over link aggregation

The link aggregation uses a hash to select a link for each *conversation*:

- Conversation defined by the load-sharing mode
  - IP source and destination address (L3-based) = Default option
  - Source and destination UDP/TCP ports (L4-based) = Option for situations in which few IP addresses are sent on the link
- Link for a conversation stays constant
- Traffic tends to be balanced more evenly as conversations increase

Switch hashes the source and destination address and, based on the hash, assigns the conversation to a link within the aggregation

All subsequent frames from the same workstation to the same server are part of that conversation and assigned to the same physical link, avoiding out-of-order packet delivery

The link aggregation will only select a new link for the conversation if the current one fails.
If you are using LACP, LACP ensures that all packets are delivered in order as the link fails over.

Load sharing - multiple conversations

The hash assigns conversations to links arbitrarily, not based on round-robin.

Conversation of client 1 and sever 1 beings assigned to link 2 in aggregation, when another session begins it could be assigned to same

However, when a link aggregation carries many conversations, the conversations tend to spread out over the links relatively evenly

Our hashing technique will only chose the egress port, If that egress port is fully occupied ,the packets will drop.

All physical links carry precisely the same number of conversations, though, the links might experience significant differences in traffic flow. This is because different conversations will have different bandwidth requirements different levels of burstiness, and so on.

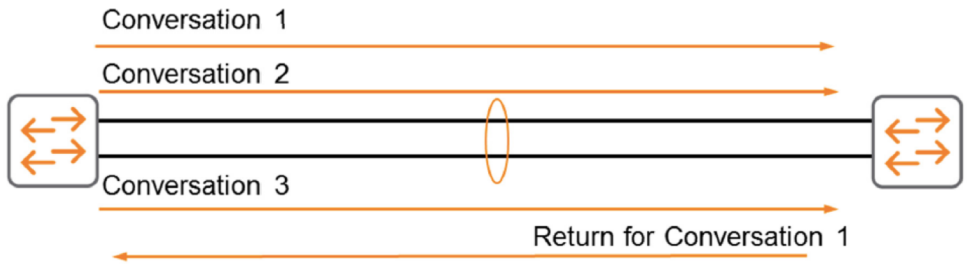

The return traffic is another conversation, and the connected switch might assign it to a different link even when that switch is using the same load-sharing mode.

However, this behaviour does not typically cause issues. Each switch considers the link aggregation as a whole as the logical interface, learns MAC addresses on the link aggregation interface, and accepts traffic on any active link.

Traffic that is redistributed in the middle of a traffic flow could potentially cause mis-ordered data packets. LACP uses the marker protocol to prevent data packets from being duplicated or reordered due to redistribution. Marker PDUs are sent on each port-channel link

LACP Load balancing Modes

LOAD balancing mode:

L2 Load Balancing:

Hash based on Source and destination MAC when L3 information not available
Hashing will be done on SRC-MAC-ADDR, DEST-MAC-ADDR, and “Physical Port.”

L3 Load balance:

Source and destination IP address when available When Layer 3 information is not available (non-IP or IPv6 traffic), source and destination MAC address
Hashing will be done on SRC-MAC-ADDR, DEST-MAC-ADDR, SRC-IP-ADDR, DEST-IP-ADDR, and “Physical Port.”

In some cases, you might want to set L4-based mode instead.
For example, the link aggregation might support traffic between a limited number of devices or, if network address translation (NAT) is involved, IP addresses. In this case, balancing by TCP or UDP ports could create more conversations.

L4 based:
Source and destination UDP or TCP port when available
When Layer 4 information is not available (non-UDP and TCP traffic), source and destination IP address when available
When Layer 3 information is not available (non-IP traffic), source and destination MAC address

Hashing will be done on SRC-MAC-ADDR, DEST-MAC-ADDR, SRC-IP-ADDR, DEST-IP-ADDR, SRC-TCP/UDP-PORT, DEST-TCP/UDP-PORT, and “Physical Port.”

LACP-Topology

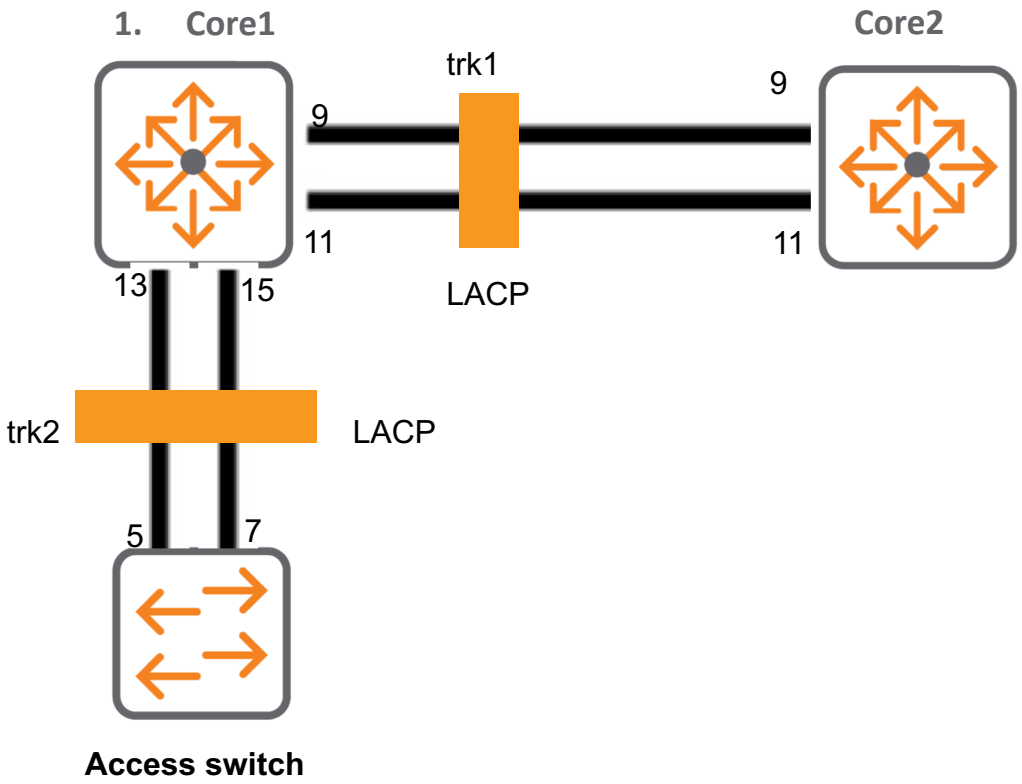

Static LACP configurations:

Core1:

Core1# config terminal

Core1(config)# trunk 9,11 trk1 lacp

Core1(config)# trunk 13,15 trk2 lacp

Core1(config)# vlan 202 tagged trk1

Core1(config)# vlan 19 tagged trk1

Core1(config)# vlan 15 tagged trk1

Core1(config)# vlan 202 tagged trk2

Core1(config)# vlan 10 tagged trk2

Core1(config)# vlan 15 tagged trk2

Configurations:

Core 2:

Core2(config)# trunk 9,11 trk1 lacp

Core2(config)# vlan 202 tagged trk1

Core2(config)# vlan 10 tagged trk1

Core2(config)# vlan 15 tagged trk1

Access-1

Access-1(config)# trunk 5,7 trk2 lacp

Access-1(config)# vlan 202 tagged trk2

Access-1(config)# vlan 10 tagged trk2

Access-1(config)# vlan 15 tagged trk2

Observe the trunks

```
Core2(config)# show trunks

 Load Balancing Method:  L3-based (default)


 Port   | Name                      Type      | Group Type

 ------ + --------------------------------- ---------- + ----- --------

 9      |                          100/1000T  | Trk1  LACP

 11     |                           100/1000T  | Trk1  LACP
```

```
Core1(config)# show trunks

Load Balancing Method:  L3-based (default)


 Port   | Name                      Type      | Group Type

 ------ + --------------------------------- ---------- + ----- --------

 9      |                          100/1000T  | Trk1  LACP

 11     |                           100/1000T  | Trk1  LACP

 13     |                           100/1000T  | Trk2  LACP

 15     |                           100/1000T  | Trk2  LACP
```

Verify LACP

```
Core1(config)# show lacp

                    LACP

      LACP     Trunk    Port          LACP     Admin   Oper
Port  Enabled  Group    Status        Partner  Status  Key     Key
----- -------- -------- ------------- -------- ------- ------- -------
9     Active   Trk1     Up            Yes      Success 0       532
11    Active   Trk1     Up            Yes      Success 0       532
13    Active   Trk2     Up            Yes      Success 0       533
15    Active   Trk2     Up            Yes      Success 0       533
```

```
Core1(config)# show lacp local

      LACP Local Information.

System ID: f40343-0f6260

      LACP              Tx     Rx Timer
Port  Trunk  Mode       Aggregated  Timer  Expired
----- ------ ---------- ----------- ------ --------
9     Trk1   Active     Yes         Slow   No
11    Trk1   Active     Yes         Slow   No
13    Trk2   Active     Yes         Slow   No
15    Trk2   Active     Yes         Slow   No
```

Verify LACP

Core1(config)# show lacp peer

LACP Peer Information.

System ID: f40343-0f6260

Local Port	Local Trunk	System ID	Port	Port Priority	Oper Key	LACP Mode	Tx Timer
9	Trk1	f40343-0f8260	9	0	532	Active	Slow
11	Trk1	f40343-0f8260	11	0	532	Active	Slow
13	Trk2	9cdc71-ffcf40	5	0	533	Active	Slow
15	Trk2	9cdc71-ffcf40	7	0	533	Active	Slow

Core1(config)# show lacp counters

LACP Port Counters.

Port	Trunk	LACP PDUs Tx	LACP PDUs Rx	Marker Req. Tx	Marker Req. Rx	Marker Resp. Tx	Marker Resp. Rx	Marker Error
9	Trk1	1826	1826	0	0	0	0	0
11	Trk1	1826	1826	0	0	0	0	0
13	Trk2	2095	2070	0	0	0	0	0
15	Trk2	2092	2068	0	0	0	0	0

Verify LACP status

```
Core2(config)# show lacp
                        LACP
  LACP     Trunk    Port          LACP     Admin   Oper
  Port     Enabled  Group    Status   Partner  Status   Key    Key
  -----    -------  -------  -------  -------  -------  ------ ------
  9        Active   Trk1     Up       Yes      Success  0      532
  11       Active   Trk1     Up       Yes      Success  0      532

Core2(config)# show lacp peer
LACP Peer Information.
System ID: f40343-0f8260

  Local  Local                Port     Oper    LACP    Tx
  Port   Trunk  System ID     Port  Priority  Key     Mode    Timer
  ------ ------ -------------- ----- --------- ------- -------- -----
  9      Trk1   f40343-0f6260  9     0         532     Active   Slow
  11     Trk1   f40343-0f6260  11    0         532     Active   Slow
```

```
Core2(config)# show trunks

Load Balancing Method:  L3-based (default)

  Port   | Name                      Type       | Group Type
  ------ + ------------------------------ ---------- + ----- --------
  9      |                        100/1000T  | Trk1  LACP
  11     |                        100/1000T  | Trk1  LACP
```

Verify LACP status:

```
Access-1# show lacp
                        LACP
  LACP     Trunk    Port          LACP     Admin   Oper
 Port    Enabled  Group    Status    Partner  Status   Key     Key
 -----  -------  -------  -------  -------  -------  ------  ------
 5      Active   Trk2     Up       Yes      Success  0       533
 7      Active   Trk2     Up       Yes      Success  0       533
```

```
Access-1# show lacp peer
LACP Peer Information.

System ID: 9cdc71-ffcf40

 Local  Local                    Port     Oper    LACP     Tx
 Port   Trunk  System ID         Port  Priority  Key      Mode     Timer
 ------ ------ ----------------- ----- --------- -------- -------- -----
 5      Trk2   f40343-0f6260     13    0         533      Active   Slow
 7      Trk2   f40343-0f6260     15    0         533      Active   Slow
```

Dynamic LACP configuration

Active and Passive configuration:

Core1:

Core1(config)# interface 19,21 lacp passive
Core1(config)# interface 19,21 lacp key 250

Core1(config)# interface 5,6 lacp active
Core1(config)# interface 5,6 lacp key 500

Core 2:

Core2(config)# interface 19,21 lacp active
Core2(config)# interface 19,21 lacp key 250

Core2(config)# interface 5,6 lacp active
Core2(config)# interface 5,6 lacp key 500

Core1(config)# show lacp

LACP

LACP Port	Trunk Enabled	Port Group	LACP Status	Admin Partner	Oper Status	Key	Key
5	Active	Dyn1	Up	Yes	Success	500	500
6	Active	Dyn1	Up	Yes	Success	500	500
9	Active	Trk1	Up	Yes	Success	0	532
11	Active	Trk1	Up	Yes	Success	0	532
13	Active	Trk2	Up	Yes	Success	0	533
15	Active	Trk2	Up	Yes	Success	0	533
19	Passive	19	Up	No	Success	250	250
21	Passive	21	Up	No	Success	250	250

Core2(config)# show lacp

LACP

LACP Port	Trunk Enabled	Port Group	LACP Status	Admin Partner	Oper Status	Key	Key
5	Active	Dyn1	Up	Yes	Success	500	500
6	Active	Dyn1	Up	Yes	Success	500	500
9	Active	Trk1	Up	Yes	Success	0	532
11	Active	Trk1	Up	Yes	Success	0	532
19	Active	Dyn2	Up	Yes	Success	250	250
21	Active	Dyn2	Up	Yes	Success	250	250

trunk-load-balance

Core1(config)# trunk-load-balance

L3-based Load balance based on IP Layer 3 information in packets.

L4-based Load balance based on Layer 4 information in packets.

Core1(config)# trunk-load-balance

Core1(config)# trunk-load-balance l4-based

Verify the interface that forwards the traffic:

i,e ping was initiated from core-2

Core1(config)# show trunks load-balance interface trk1 mac f40343-0f6260 f40343-0f8260 ip 10.17.172.11 10.17.172.12 inbound-port 9

Traffic in this flow will be forwarded out port 9 based on the configured

Layer 3 load balancing

Core1(config)# show trunks load-balance interface trk1 mac f40343-0f6260 f40343-

0f8260 ip 10.17.172.11 10.17.172.12 inbound-port 11

Traffic in this flow will be forwarded out port 9 based on the configured
Layer 3 load balancing

Troubleshooting & debug commands

CLI Commands:

#Show lacp
#Show lacp counters [<Port-List>]
#show lacp local [<Port-List>]
#Show lacp peer [<Port-List>]
#Show trunks
#show trunks load-balance interface <TRUNK-ID> mac <SRC-MAC-ADDR> <DEST-MAC-ADDR> [ip <SRC-IP-ADDR> <DEST-IP-ADDR> [<SRC-TCP/UDP-PORT> <DEST-TCP/UDP-PORT>]] inbound-port <PORT-NUM> ether-type <ETHER-TYPE> inbound-vlan <VLAN-ID>

Check for RMON logs

#show log –r

Also we can debug
#debug lacp ?
events -Display LACP events
packets -Display all LACP packet

Verify the debug logs using

1. Debug destination session or Debug destination buffer
2. Debug lacp

Debug logs

0003:04:03:10.25 lacp mLACPCtrl:LACP: Port 9 is DOWN

0003:04:03:10.31 lacp mLACPCtrl:LACP: Port 11 is DOWN

0003:04:03:14.49 lacp mLACPCtrl:LACP: Data BPDU Rx. on port 6

0003:04:03:14.55 lacp mLACPCtrl:LACP: Received LACP PDU LENGTH = 124

0003:04:03:14.63 lacp mLACPCtrl:LACP: Data BPDU Rx. on port 19

0003:04:03:14.69 lacp mLACPCtrl:LACP: Received LACP PDU LENGTH = 124

0003:04:03:14.77 lacp mLACPCtrl:LACP: Data BPDU Rx. on port 21

0003:04:03:14.83 lacp mLACPCtrl:LACP: Received LACP PDU LENGTH = 124

0003:04:03:14.91 lacp mLACPCtrl:LACP: LACP data BPDU Tx on port 6

0003:04:03:14.98 lacp mLACPCtrl:LACP: Periodic timer has been set to SLOW_PERIODIC_TIME for port 6

0003:04:03:15.08 lacp mLACPCtrl:LACP: Data BPDU Rx. on port 13

0003:04:03:15.15 lacp mLACPCtrl:LACP: Received LACP PDU LENGTH = 124

0003:04:03:15.23 lacp mLACPCtrl:LACP: Periodic timer has been set to SLOW_PERIODIC_TIME for port 9

0003:04:03:15.33 lacp mLACPCtrl:LACP: Periodic timer has been set to SLOW_PERIODIC_TIME for port 11

0003:04:03:39.64 lacp mLACPCtrl:LACP: Port 11 is UP

0003:04:03:39.80 lacp mLACPCtrl:LACP: Port 9 is UP

0003:04:03:40.16 lacp mLACPCtrl:LACP: LACP data BPDU Tx on port 11

0003:04:03:40.23 lacp mLACPCtrl:LACP: Periodic timer has been set to SLOW_PERIODIC_TIME for port 11

0003:04:03:40.49 lacp mLACPCtrl:LACP: Data BPDU Rx. on port 9

0003:04:03:40.55 lacp mLACPCtrl:Manual LACP: Partner info selected for port

11,partner port 11,partner system-priority 33376,partner port priority 0,partner Key 532

Log -r

Trk2 is Down

LACP Port	Trunk Port Enabled	Trunk Group	LACP Status	Admin Partner	Oper Status	Admin Key	Oper Key
5	Active	Dyn1	Up	Yes	Success	500	500
6	Active	Dyn1	Up	Yes	Success	500	500
9	Active	Trk1	Up	Yes	Success	0	532
11	Active	Trk1	Up	Yes	Success	0	532
13	Active	Trk2	Down	No	Success	0	533
15	Active	Trk2	Down	No	Success	0	533
19	Passive	Dyn3	Up	Yes	Success	250	250
21	Passive	Dyn3	Up	Yes	Success	250	250

---- Reverse event Log listing: Events Since Boot ----

Show log –r
I 01/04/90 04:06:05 00077 ports: port 15 in Trk2 is now off-line
I 01/04/90 04:06:05 00079 ports: trunk Trk2 is now inactive
I 01/04/90 04:06:04 00077 ports: port 13 in Trk2 is now off-line
I 01/04/90 04:04:50 00076 ports: port 21 in Dyn3 is now on-line
I 01/04/90 04:04:50 00076 ports: port 19 in Dyn3 is now on-line
I 01/04/90 04:04:48 00435 ports: port 21 is Blocked by STP
I 01/04/90 04:04:48 00435 ports: port 19 is Blocked by STP
I 01/04/90 04:04:48 00078 ports: trunk Dyn3 is now active
I 01/04/90 04:04:48 00435 ports: port 21 is Blocked by LACP
I 01/04/90 04:04:48 00435 ports: port 19 is Blocked by LACP
I 01/04/90 04:04:28 00077 ports: port 21 in Dyn3 is now off-line
I 01/04/90 04:04:28 00079 ports: trunk Dyn3 is now inactive
I 01/04/90 04:04:28 00077 ports: port 19 in Dyn3 is now off-line
I 01/04/90 04:03:42 00076 ports: port 11 in Trk1 is now on-line

I

LACPDU packet format

```
      8 2018-08-24 15:15:24.393323 HewlettP_0f:62:77      Slow-Protocols       LACP       124 Link Aggregation Control Protocol Version 1.  Actor Port = 9 Partner Port = 9
      9 2018-08-24 15:15:24.393324 HewlettP_0f:62:75      Slow-Protocols       LACP       124 Link Aggregation Control Protocol Version 1.  Actor Port = 11 Partner Port = 11
     10 2018-08-24 15:15:24.417886 HewlettP_0f:82:77      Slow-Protocols       LACP       124 Link Aggregation Control Protocol Version 1.  Actor Port = 9 Partner Port = 9
▶ Frame 8: 124 bytes on wire (992 bits), 124 bytes captured (992 bits) on interface 0
▶ Ethernet II, Src: HewlettP_0f:62:77 (f4:03:43:0f:62:77), Dst: Slow-Protocols (01:80:c2:00:00:02)
▶ Slow Protocols
▼ Link Aggregation Control Protocol
    LACP Version Number: 0x01
    Actor Information: 0x01
    Actor Information Length: 0x14
    Actor System Priority: 25184
    Actor System: HewlettP_0f:62:60 (f4:03:43:0f:62:60)
    Actor Key: 532
    Actor Port Priority: 0
    Actor Port: 9
  ▼ Actor State: 0x3d, LACP Activity, Aggregation, Synchronization, Collecting, Distributing
      .... ...1 = LACP Activity: Active
      .... ..0. = LACP Timeout: Long Timeout
      .... .1.. = Aggregation: Aggregatable
      .... 1... = Synchronization: In Sync
      ...1 .... = Collecting: Enabled
      ..1. .... = Distributing: Enabled
      .0.. .... = Defaulted: No
      0... .... = Expired: No
    [Actor State Flags: **DCSG*A]
    Reserved: 000000
    Partner Information: 0x02
    Partner Information Length: 0x14
    Partner System Priority: 33376
    Partner System: HewlettP_0f:82:60 (f4:03:43:0f:82:60)
    Partner Key: 532
    Partner Port Priority: 0
    Partner Port: 9
  ▼ Partner State: 0x3d, LACP Activity, Aggregation, Synchronization, Collecting, Distributing
      .... ...1 = LACP Activity: Active
      .... ..0. = LACP Timeout: Long Timeout
      .... .1.. = Aggregation: Aggregatable
      .... 1... = Synchronization: In Sync
      ...1 .... = Collecting: Enabled
      ..1. .... = Distributing: Enabled
      .0.. .... = Defaulted: No
      0... .... = Expired: No
    [Partner State Flags: **DCSG*A]
    Reserved: 000000
    Collector Information: 0x03
    Collector Information Length: 0x10
    Collector Max Delay: 10000
    Reserved: 00000000000000000000000000
    Terminator Information: 0x00
    Terminator Length: 0x00
    Reserved: 00000000000000000000000000000000000000000000000000...
```

aruba
a Hewlett Packard
Enterprise company

DT-LACP

AIRHEADS
COMMUNITY

Distributed trunking

Distributed trunking uses a proprietary protocol that allows two or more port trunk links distributed across two switches to create a trunk group.

Distributed Trunking is a link aggregation technique, where two or more links across two switches are aggregated together to form a trunk.

DT provides node-level L2 resiliency in an L2 network, when one of the switches fails.

Both the DT switches are connected by a special interface called ISC port, which exchange information so that both the switches simulate itself as the single switch to the downstream.

DTD (Distributed Trunking Device) is the server/switch which forms a trunk with DTS (Distributed Trunking switch)

Distributed trunking uses the distributed trunking interconnect protocol (DTIP) to transfer DT-specific configuration information for the comparison process and to synchronize MAC and DHCP snooping binding data between the two DT peer switches.

Distributed trunking will use a new proprietary protocol called DTIP (Distributed Trunking Interconnect Protocol) which will run in the Distributed Trunking Switches (DTS)

Sample of DT connection

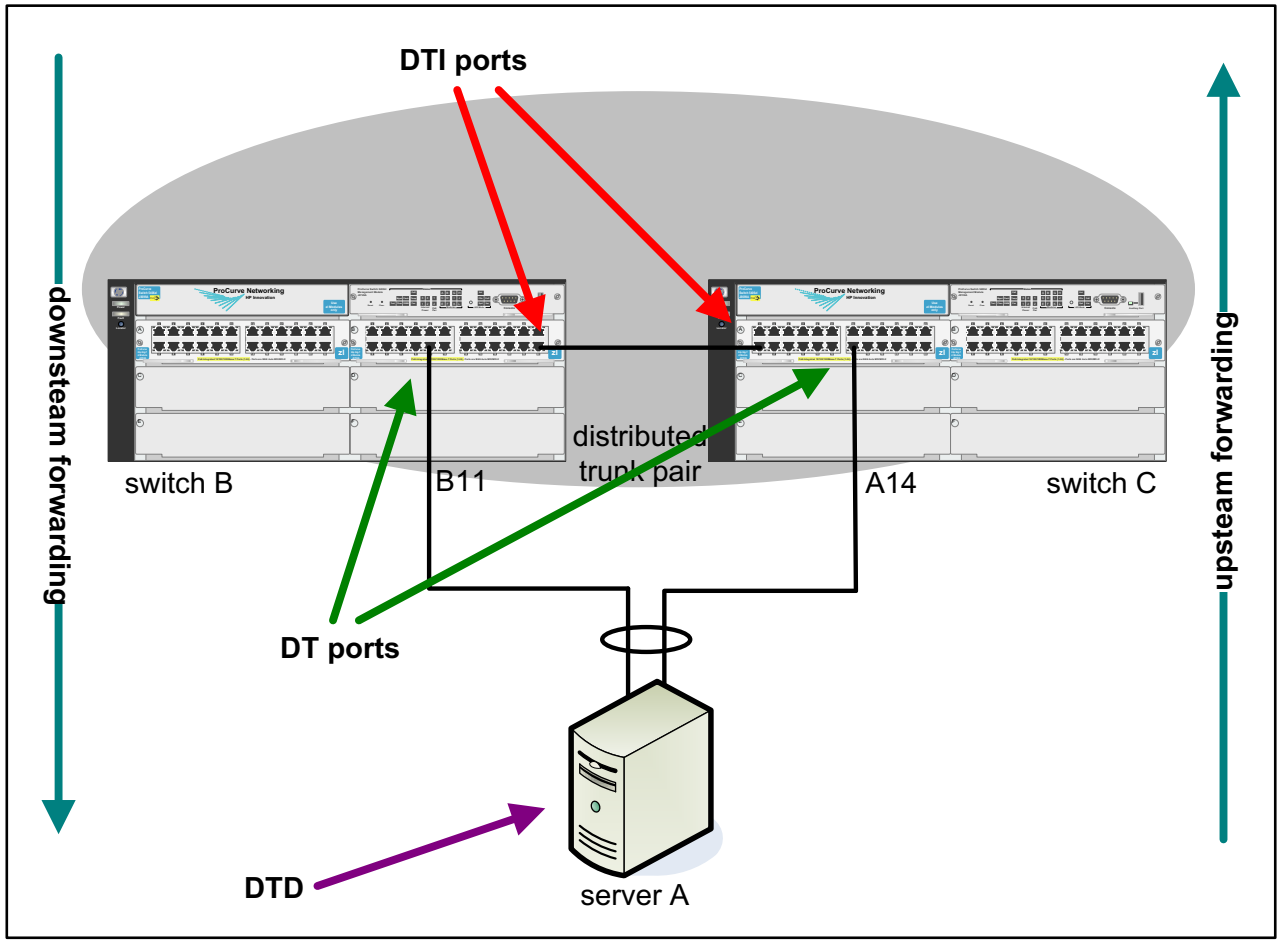

DT-LACP

Each distributed trunk (DT) switch in a DT pair must be configured with a separate ISC link and peer-keepalive link.

Both the distribution switches should Aruba-HPE switches, though downstream network connecting DT link can be a third party server/switch

Each distributed trunk (DT) switch in a DT pair must be configured with a separate ISC link and peer-keepalive link.

The peerkeepalive link is used to transmit keepalive messages when the ISC link is down to determine if the failure is a link-level failure or the complete failure of the remote peer.

STP will not be run on DT links which is connected to the servers

A distributed trunk can span a maximum of two switches.

DT between different type of switches is not supported

Distributed Trunking will be supported with standard LACP running on the DT links towards the server, with proprietary protocol running between DT switches to sync the LACP information across them

ISC(InterSwitch-Connect) port & Peer Keep alive

ISC:

DT-switches are connected by a special interface called the InterSwitch-Connect (ISC) port. This interface exchanges information so that the DT-switches appear as a single switch to a downstream device.

ISC/Distributed Trunking Interconnect Protocol (DTIP) is responsible for **synchronization of DT configuration and dynamic information from the peer**, so that both the aggregation switches behave as if they were a single switch with LACP perspective.

ISC protocol carries messages for many features linked with DT, such as LACP, loop-protect, MAC learns/ageouts, IGMP snooping and DHCP snooping.

An ISC port becomes a member of all VLANs that are configured on the switch. When a new VLAN is configured, the ISC ports become members of that VLAN

Peer-keepalive

The peer-keepalive link is used to transmit keepalive messages when the ISC link is down to determine if the failure is a link-level failure or the complete failure of the remote peer

peer-keepalive interval
Configure the interval between peer-keepalive messages in milliseconds. The default is 1000 milliseconds.

peer-keepalive timeout
Configure the peer-keepalive timeout in seconds. The default is 5 seconds.

Assigning roles to the DT switches

DT switch with the lowest mac address will be DT primary and other DT switch would be secondary, this role comes in to play when ISC link is Down, & which device would forward the traffic

No data or synchronization traffic is sent over the Peer-Keepalive VLAN. Spanning Tree Protocol (STP) cannot run on peer-keepalive links. The DEFAULT VLAN cannot be a peer-keepalive VLAN.

The DT switch which has the lowest MAC gets the Primary role and will be chosen to continue to be the active links. The secondary switch blocks its entire DT links and thereby avoiding a dual active scenario.

Primary DT-switch: The primary DT-switch processes the keep-alive message and continues to forward traffic on all its DT-ports.

Secondary DT-switch: If the secondary DT-switch receives even a single keep-alive pkt. from the primary within timeout time, it disables all the DT-ports

DT-switch failure (Primary/Secondary): If the Primary DT-switch fails, the Secondary DT-switch fails to receive any Keep-alive message during the timeout period and as a result, the Secondary DT-switch assumes that the other DT-peer has failed completely and continues to forward traffic on all its DT-ports.

Keep-alive link failure: When the layer3 connectivity between the DT-switches is lost (keep-alive VLAN goes down), both the DT-switches ignore this condition and continue forwarding as if nothing happened. A log message is recorded in the system log.

So, the Secondary DT-switch would re-enable back all disabled DT-ports after the configured timeout value

Double Failure scenario (Keep-alive link failure followed by ISC link failure): If keep-alive VLAN goes down followed by ISC link failure, the DT-switches have no way to know the state of their peer and hence both the DT-switches forward the traffic on all their DT-ports.

DT requirement

Only one ISC link is supported per switch, with a maximum of 60 DT trunks supported on the switch

All DT linked switches must be running the same software version.

A distributed trunk can span a maximum of two switches.

The limit of 144 manual trunks per switch includes distributed trunks as well.

Each server can have up to four physical links aggregated in a single switch, meaning that there can be a maximum of eight ports (four aggregated links for each DT switch) included in a DT trunk.

Distributed Trunking will come up only if both the switches are configured for DT. The configuration in single switch will never initiate LACP negotiation

Configure bpdu-filter on keep-alive links if MSTP is enabled.

Features not supported include:

- ARP protection is not supported on the distributed trunks.
- QinQ in mixed VLAN mode and distributed trunking are mutually exclusive
- Dynamic IP Lockdown protection is not supported on the distributed trunks.
- Dynamic link aggregation is not supported & should be configured manually
- SVLANs in mixed mode on DT or ISC links
- Meshing
- Multicast routing
- IPv6 routing

DT requirement

DT between different type of switches is not supported

DT is not supported between an HP 5406 switch and an HP 5400R switch.

The ISC link must have a VLAN interface configured for the same VLAN on both DT switches.

VLAN membership for all DT trunk ports should be the same on both DT switches in a DT pair.

IGMP-snooping or DHCP-snooping configuration on a DT VLAN should be the same on both DT switches.

Switch firmware version KB/YA/WB.15.18.xxxx or greater

On DT trunks spanning tree BPDUs are suppressed

For a DT, if IGMP-snooping or DHCP-snooping is enabled on a VLAN that has a DT port as a member port of the VLAN, the same must be configured on the peer DT on the same VLAN.

Loop-protection configuration on a DT VLAN should be the same for both DT switches

ISC link failure

In this scenario, the bridge with the lowest system MAC address acts as the DT primary device; the other device is the DT secondary device.

These roles are used to determine which device forwards traffic when the ISC link is down.

Both the switches detecting the link down including the own switch should do the following things:

Peer-keepalive messages are sent by both the DT switches as soon as the switches detect that the ISC link is down.

Peer-keepalive message transmission (sending and receiving) is suspended until the peer-keepalive hold timer expires

When the hold timer expires, the DT switches begin sending peer-keepalive messages periodically while receiving peer-keepalive messages from the peer switch.

If the DT switch fails to receive any peer-keepalive messages for the timeout period, it continues to forward traffic, assuming that the DT peer switch has completely failed.

Conversely, if the failure is because the ISC link went down and the secondary DT switch receives even one peerkeepalive message from the primary peer, the secondary switch disables all its DT ports

The primary switch always forwards the traffic on its DT ports even if it receives peer-keepalive messages from the secondary DT switch.

In both situations, if the ISC link or the DT switch becomes operational, both the DT peers sync the MAC addresses learned during the failover and continue to forward traffic normally. peer-keepalive timers is halted.

peer-keepalive messages

Distributed trunking uses UDP-based peer-keepalive messages to determine if an ISC link failure is at the link level or the peer has completely failed.

Only peer-keepalive messages are sent over the peer-keepalive VLAN (Layer 3 link.) These messages indicate that the DT switch from which the message originates is up and running. No data or synchronization traffic is sent over the peer-keepalive VLAN.

<parameter> - Configure the various peer-keepalive parameters

o destination < IP-ADDR > - The destination IPv4 address to be used by DT aggregation switches to send peer-keepalive messages to peer DT switch when the ISC is down.
o vlan <VLAN-ID> - The VLAN to which the peer-keepalive interface belongs.
o udp-port <1024-49151> - The source UDP port to be used for transmitting peer-keepalive HELLO messages.
o interval <400-10000 > - The interval between peer-keepalive messages in milliseconds.
o timeout <3-20> - The peer-keepalive timeout in seconds. Default is 5 seconds

ISC link use case

In case of a Distributed trunk (switch to switch model), if the local DT trunk is down for some reason then the traffic will be forwarded on ISC link.

Hence the load balance logic would give the ISC port as the egress port. If the ISC link happens to a trunk then the load balance logic will be applied on the ISC trunk.

Consider the following scenario where trk1 is a DT trunk and trk2 is the ISC trunk, and DT trunk trk1 is down.
If this is exercised on the DT trunk trk1, then the actual load balance will applied on the trk2, ISC trunk,
because the DT trunk is down and the actual traffic will be forwarded on the ISC trunk.

DTIP Frames Overview

DTIP uses SLOW_PROTOCOL_MAC and use LACP type in the ethertype. The sub-type in the LACP pkt will indicate whether the frame is DTIP.

The ISC link connecting two DT aggregation switches runs a protocol that allows the following protocol packets to be exchanged between the two DT aggregation switches

ISC Protocol Packet Type	Value
ISCP_DATA PACKET	0x0
ISCP_NAK PACKET	0x2
ISCP ACK PACKET	0x4
ISCP HELLO PACKET	0x8

The following are shared in the DTIP frame & subtype specified in the DTIP frame.

- DT related configuration for configuration validation.
- MAC learns and age outs
- Local system MAC address
- DT member port status
- LACPDUs received on DT-LACP ports
- Loop-protect management
- DHCP snooping synchronization

ISC protocol header

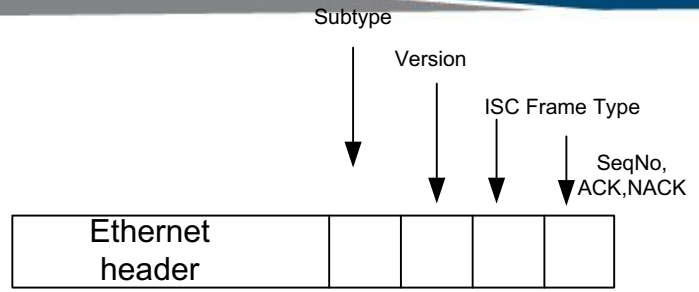

Packet Type (Upper nibble)

Hello packet	ACK	NAK	DATA	28 bit sequence number

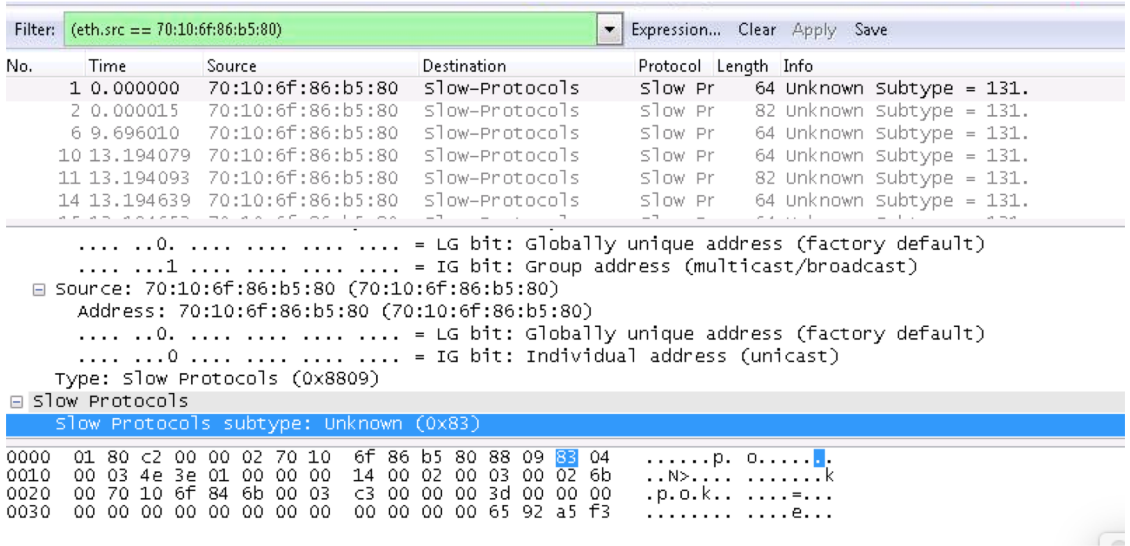

Destination Address: ISC protocol packets carry the Slow_Protocols_Multicast address.

Source Address: The source address is the 48-bit address of the switch's system MAC address from which the ISC protocol messages sent.

Ethernet Type : ISC protocol packet uses ethernet type vale of 0x8809 which identifies the frame as being a Slow Protocol.

Subtype and Version : The slow protocol subtype for ISC protocol is 131

ISC Frame Type: This field contains the values to represent the different ISC packet types.

Seq No : This field carries the sequence number and packet type information. seqNo filed of the packet header tells the packet type.

DT-Topology

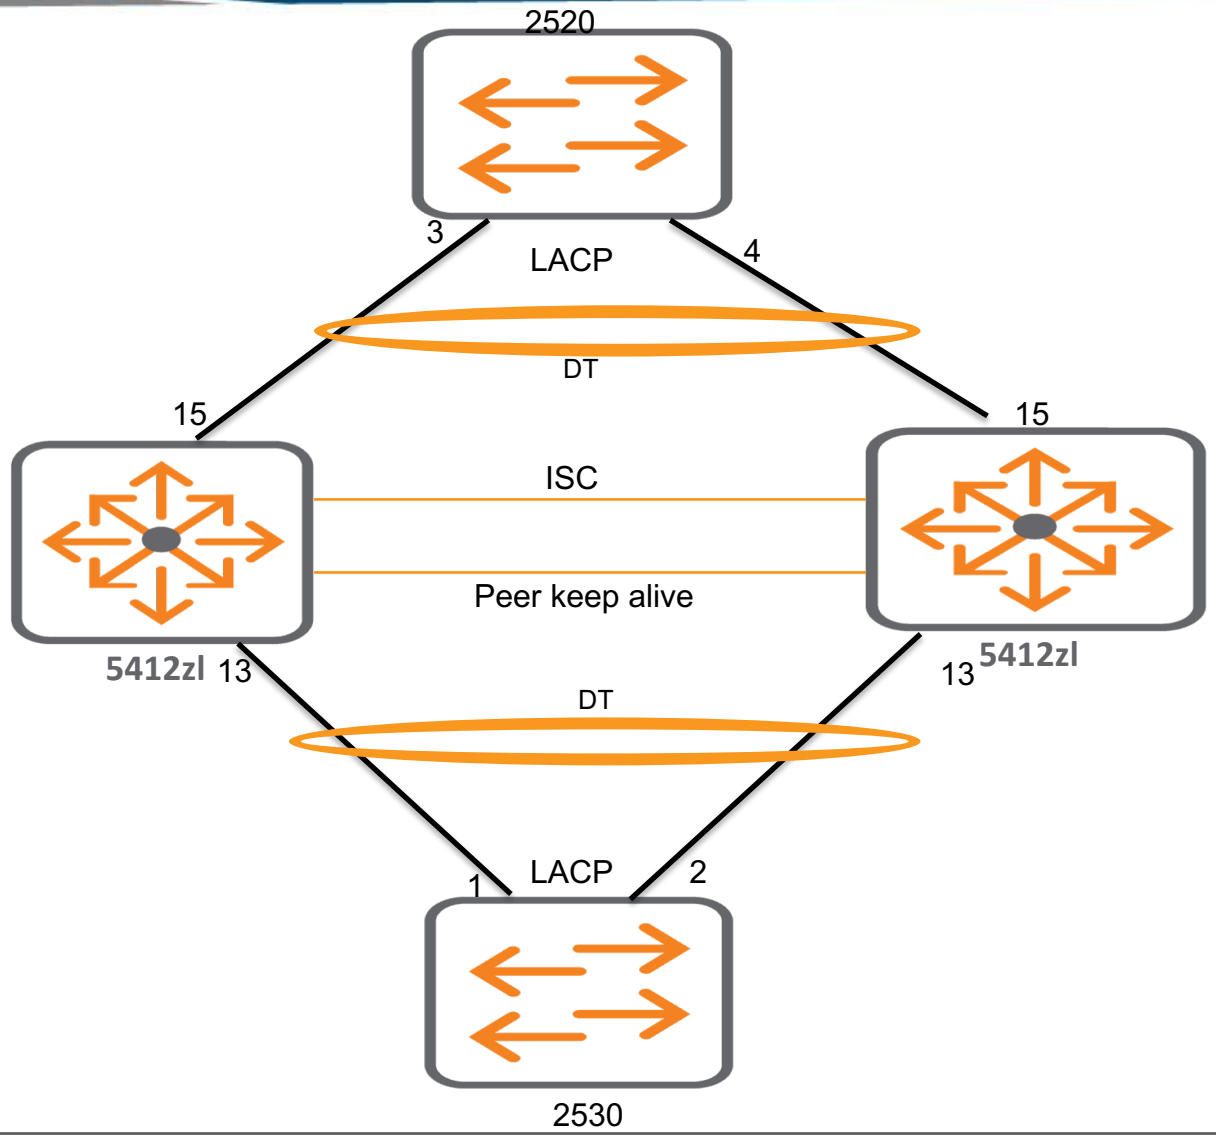
2520
3
LACP
4
DT
15
15
ISC
Peer keep alive
5412zl
13
13
5412zl
DT
1
LACP
2
2530

Distribute Trunking configuration

Switch1:

Create a ISC link

SwitchDT-1(config)# vlan 204
SwitchDT-1(vlan-204)# tagged A11
SwitchDT-1(config)# switch-interconnect A11

Create a keepalive link

SwitchDT-1(config)# vlan 100
SwitchDT-1(vlan-100)# untagged A12

SwitchDT-1(vlan-100)# ip address 10.10.10.8/24

SwitchDT-1(config)# distributed-trunking peer-keepalive vlan 100
SwitchDT-1(config)# distributed-trunking peer-keepalive destination 10.10.10.10

Create LACP Trunks

SwitchDT-1(config)# trunk A13 trk1 dt-lacp
SwitchDT-1(config)# trunk A15 trk2 dt-trunk

SwitchDT-1(config)# vlan 204

SwitchDT-1(vlan-204)# tagged trk1
SwitchDT-1(vlan-204)# tagged trk2

SwitchDT-1(config)# ip default-gateway 10.17.204.254
SwitchDT-1(config)# vlan 204
SwitchDT-1(vlan-204)# ip address 10.17.204.50/24

switchDT-1(config)# write memory

CLI configurations

Switch2:

Create ISC link:

```
switchDT-2# config terminal
switchDT-2(config)# vlan 204
switchDT-2(vlan-204)# tagged A11
switchDT-2(vlan-204)# exit
switchDT-2(config)# switch-interconnect A11
```

Peer keep alive

```
switchDT-2(config)# vlan 100
switchDT-2(vlan-100)# untagged A12
switchDT-2(vlan-100)# ip address 10.10.10.9/24
switchDT-2(vlan-100)# exit

switchDT-2(config)# distributed-trunking peer-keepalive vlan 100
switchDT-2(config)# distributed-trunking peer-keepalive destination 10.10.10.8
```

LACP trunks

```
switchDT-2(config)# trunk A13 trk1 dt-lacp
switchDT-2(config)# trunk A15 trk2 dt-trunk
switchDT-2(config)# vlan 204
switchDT-2(vlan-204)# tagged trk1
switchDT-2(vlan-204)# tagged trk2
switchDT-2(config)#

switchDT-2(config)# ip default-gateway 10.17.204.254
switchDT-2(config)# vlan 204
switchDT-2(vlan-204)# ip address 10.17.204.204/24
```

CLI configurations

Switch-3

```
switch3# config terminal

switch3(config)# ip default-gateway 10.17.204.254

switch3(config)# vlan 204

switch3(vlan-204)# ip address 10.17.204.60/24


switch3(config)# trunk 1-2 trk3 lacp

switch3(config)# vlan 204
switch3(vlan-204)# tagged trk3
switch3(vlan-204)# exit


switch3(config)# write memory
```

Command line

Swicth4:

Switch4# config terminal

Switch4(config)# ip default-gateway 10.17.204.254

Switch4(config)# vlan 204

Switch4(vlan-204)# ip address 10.17.204.70/24

Switch4(config)# trunk 3-4 trk4 lacp

Switch4(config)# vlan 204

Switch4(vlan-204)# tagged trk4

Switch4(vlan-204)#

Switch4(vlan-204)# exit

Switch4(config)#

Switch4(config)# write mem

Verify Trunks

SwitchDT-1# show trunks. →view the status of trunks on the switch

Load Balancing Method: L3-based (default)

Port	Name	Type	Group	Type
A13		100/1000T	Trk1	dt-lacp
A15		100/1000T	Trk2	dt-trunk

SwitchDT-1# show lacp. → LACP status

LACP

Port	LACP Enabled	Trunk Group	Port Status	LACP Partner	Admin Status	Oper Key	Oper Key
A13	Active	Trk1	Up	Yes	Success	0	962

Verify DT status

```
SwitchDT-1# show distributed-trunking status

 Distributed Trunking Status

  Switch Interconnect (ISC)          : Up
  ISC Protocol State                 : In Sync
  DT System ID                       : 98f2b3-968f00
  Oper Role Priority                 : 32768
  Peer Oper Role Priority            : 32768
  Switch Role                        : Secondary


  SwitchDT-1# show distributed-trunking config

  Distributed Trunking Information

   Switch Interconnect (ISC) : A11
   Admin Role Priority       : 32768
   System ID                 : 98f2b3-968f00
   DT trunk                  : Trk2
   DT lacp                   : Trk1
```

Compare parameters on the Local and Peer switch

SwitchDT-1# show distributed-trunking consistency-parameters global. -> “Display global peer consistency details.”

```
                        Local         Peer

                        -----         ----
Image Version                 KB.16.05.0004   KB.16.05.0004

IP Routing                  Disabled        Disabled

Peer-keepalive interval       1000          1000

PIM-DM Support              Disabled        Disabled

PIM-SM Support              Disabled        Disabled

IGMP enabled VLANs on Local                   :
IGMP enabled VLANs on Peer                    :
PIM-DM-DT Enabled VLANs on Local              :
PIM-DM-DT Enabled VLANs on Peer               :
PIM-SM-DT Enabled VLANs on Local              :
PIM-SM-DT Enabled VLANs on Peer               :
DHCP-snooping Enabled on Local              : No
DHCP-snooping Enabled on Peer               : No
DHCP-snooping Enabled VLANs on Local          :
DHCP-snooping Enabled VLANs on Peer           :
DHCP-snooping Max-Binding Configured on Local   : No
DHCP-snooping Max-Binding Configured on Peer    : No
DHCPv6-snooping Enabled on Local            : No
DHCPv6-snooping Enabled on Peer             : No
DHCPv6-snooping Enabled VLANs on Local        :
DHCPv6-snooping Enabled VLANs on Peer         :
DHCPv6-snooping Max-Binding Configured on Local : No
DHCPv6-snooping Max-Binding Configured on Peer  : No
Loop-protect enabled VLANs on Local           :
Loop-protect enabled VLANs on Peer            :
MLD enabled VLANs on Local                    :
MLD enabled VLANs on Peer                     :
```

Check Local and Peer switch config

SwitchDT-1# show distributed-trunking consistency-parameters trunk trk1. →Display peer consistency details for the specified trunks

Allowed VLANs on Local : 1,204

Allowed VLANs on Peer : 1,204

Name	Local Value	Peer Value
Loop Protect	Disabled	Disabled

SwitchDT-1# show distributed-trunking consistency-parameters trunk trk2

Allowed VLANs on Local : 1,204
Allowed VLANs on Peer : 1,204

Name	Local Value	Peer Value
Loop Protect	Disabled	Disabled

SwitchDT-1# show distributed-trunking consistency-parameters trunk trk3

Allowed VLANs on Local :

Allowed VLANs on Peer :

Name	Local Value	Peer Value
Loop Protect	Disabled	Disabled

View peer-keepalive link parameters & ISC status

SwitchDT-1# show distributed-trunking peer-keepalive

Distributed Trunking peer-keepalive parameters

Destination : 10.10.10.10

VLAN : 100

UDP Port : 1024

Interval(ms) : 1000

Timeout(sec) : 5

SwitchDT-1# show distributed-trunking statistics peer-keepalive

DT peer-keepalive Status : Up
DT peer-keepalive Statistics

Tx Count : 6014

Rx Count : 1042

SwitchDT-1# show distributed-trunking statistics switch-interconnect

Switch Interconnect Port : A11

Switch Interconnect Status : Up

Switch Interconnect Statistics

Protocol Packets Sent : 17797

Protocol Packets Received : 17547

MAC Learns Sent : 378

MAC Learns Received : 474

MAC Age-Outs Sent : 549

MAC Age-Outs Received : 784

Status of DT Local and Peer

```
switchDT-2# show distributed-trunking status
 Distributed Trunking Status
  Switch Interconnect (ISC)          : Up
  ISC Protocol State                 : Out Of Sync
  DT System ID                       : 98f2b3-968f00
  Oper Role Priority                 : 32768
  Peer Oper Role Priority            : Unknown
  Switch Role                        : Unknown

switchDT-2# show logging DT
----  Event Log listing: Events Since Boot  ----
I 08/26/18 14:48:27 02012 mtm: AM1: A non-multicast client: Non-Mcast client DT,
            is registered with client ID:  1
I 08/26/18 14:48:35 03242 DT: AM1: Distributed-trunking keep-alive receive
            socket bind is successful. UDP port 1024 is used.
I 08/26/18 14:48:35 03242 DT: AM1: Distributed-trunking keep-alive receive
            socket bind is successful. UDP port 1024 is used.
I 08/26/18 14:48:35 03305 DT: AM1: ISC port  is deleted from keepalive VLAN 100
I 08/26/18 14:49:46 03226 DT: AM1: Switch-interconnect link is up
I 08/26/18 14:49:58 05201 DT: AM1: Switch-interconnect link is In Sync.
```

```
switchDT-2# show distributed-trunking status
 Distributed Trunking Status
  Switch Interconnect (ISC)          : Down
  ISC Protocol State                 : Out Of Sync
  DT System ID                       : 98f2b3-968f00
  Oper Role Priority                 : 32768
  Peer Oper Role Priority            : 32768
  Switch Role                        : Primary

switchDT-2# show logging DT
----  Event Log listing: Events Since Boot  ----
I 08/26/18 14:48:27 02012 mtm: AM1: A non-multicast client: Non-Mcast client DT, is registered with client ID:  1
I 08/26/18 14:48:35 03242 DT: AM1: Distributed-trunking keep-alive receive socket bind is successful. UDP port 1024 is used.
I 08/26/18 14:48:35 03242 DT: AM1: Distributed-trunking keep-alive receive    socket bind is successful. UDP port 1024 is used.
I 08/26/18 14:48:35 03305 DT: AM1: ISC port  is deleted from keepalive VLAN 100
I 08/26/18 14:49:46 03226 DT: AM1: Switch-interconnect link is up
I 08/26/18 14:49:58 05201 DT: AM1: Switch-interconnect link is In Sync.
W 08/26/18 14:51:43 03225 DT: AM1: Switch-interconnect link is down
I 08/26/18 14:51:43 05202 DT: AM1: Switch-interconnect link is Out of Sync.
----  Bottom of Log : Events Listed = 8  ----
```

ISC failure

Logs when secondary DTSs receives keep alive from Primary

I 08/25/18 19:07:32 02012 mtm: AM1: A non-multicast client: Non-Mcast client DT, is registered with client ID: 1

I 08/25/18 19:07:38 03242 DT: AM1: Distributed-trunking keep-alive receive socket bind is successful. UDP port 1024 is used.
I 08/25/18 19:07:38 03242 DT: AM1: Distributed-trunking keep-alive receive socket bind is successful. UDP port 1024 is used.
0002:02:17:50.94 DT eDTKeepAlive:Keep Alive packet sent, payload size = 50
I 08/25/18 19:07:38 03305 DT: AM1: ISC port is deleted from keepalive VLAN 100
I 08/25/18 19:08:11 03226 DT: AM1: Switch-interconnect link is up
I 08/25/18 19:08:25 05201 DT: AM1: Switch-interconnect link is In Sync.

W 08/25/18 19:18:09 03225 DT: AM1: Switch-interconnect link is down
I 08/25/18 19:18:09 05202 DT: AM1: Switch-interconnect link is Out of Sync.
I 08/25/18 19:18:12 03231 DT: AM1: Received keep-alive message from peer DT switch.
I 08/25/18 19:18:12 03233 DT: AM1: Switch role is Secondary, disabling all DT trunks.
I 08/25/18 19:18:29 03226 DT: AM1: Switch-interconnect link is up
I 08/25/18 19:18:39 05201 DT: AM1: Switch-interconnect link is In Sync.
I 08/25/18 19:18:40 03235 DT: AM1: Enabling all DT trunks.
W 08/27/18 20:47:04 03225 DT: AM1: Switch-interconnect link is down
I 08/27/18 20:47:04 05202 DT: AM1: Switch-interconnect link is Out of Sync.
0002:02:17:51.98 DT eDTKeepAlive:Keep Alive packet sent, payload size = 50
[24;1H[2K-- MORE --, next page: Space, next line: Enter, quit: Control-C[24;1H[24;1H[2K[24;1H[1;24r[24;1HI 08/27/18 20:47:07 03226 DT: AM1: Switch-interconnect link is up

I 08/27/18 20:47:07 03231 DT: AM1: Received keep-alive message from peer DTswitch.

I 08/27/18 20:47:07 03233 DT: AM1: Switch role is Secondary, disabling all DT trunks.

Logs when secondary doesn't received any packet from Primary

I 08/27/18 20:54:55 03226 DT: AM1: Switch-interconnect link is up

I 08/27/18 20:55:05 05201 DT: AM1: Switch-interconnect link is In Sync.

I 08/27/18 20:55:07 03235 DT: AM1: Enabling all DT trunks.

W 08/27/18 20:58:25 03225 DT: AM1: Switch-interconnect link is down

I 08/27/18 20:58:25 05202 DT: AM1: Switch-interconnect link is Out of Sync.

I 08/27/18 20:58:29 03231 DT: AM1: Received keep-alive message from peer DT switch.

I 08/27/18 20:58:29 03233 DT: AM1: Switch role is Secondary, disabling all DT trunks.

I 08/27/18 20:59:18 03226 DT: AM1: Switch-interconnect link is up

I 08/27/18 20:59:33 03234 DT: AM1: Failed to receive keep-alive messages from peer DT switch.

I 08/27/18 20:59:33 03235 DT: AM1: Enabling all DT trunks.

ISC failure

SwitchDT-1 –Secondary

0002:01:43:53.17 DT eDTKeepAlive:Keep Alive packet sent, payload size = 50

0002:01:43:54.18 DT eDTKeepAlive:Keep Alive packet sent, payload size = 50

0002:01:43:54.25 DT eDTKeepAlive:Keep Alive packet received, payload size = 50

0002:01:43:55.18 DT eDTKeepAlive:Keep Alive packet sent, payload size = 50

0002:01:43:55.25 DT eDTKeepAlive:Keep Alive packet received, payload size = 50

0002:01:43:56.18 DT eDTKeepAlive:Keep Alive packet sent, payload size = 50

0002:01:43:56.25 DT eDTKeepAlive:Keep Alive packet received, payload size = 50

MAC learn:

0002:01:47:50.32 DT mDTCtrl:MAC learn data packet sent with seq no 1

0002:01:47:50.44 DT mDTCtrl:LACP partner data packet sent with seq no 2

0002:01:47:50.54 DT mDTCtrl:DT port data packet sent with seq no 3

0002:01:47:50.66 DT mDTCtrl:DT port data packet sent with seq no 4

0002:01:47:51.32 DT mDTIscpRcv:ACK sent for seq no 1

0002:01:47:51.41 DT mDTIscpRcv:ACK sent for seq no 2

0002:01:47:51.51 DT mDTIscpRcv:ACK sent for seq no 3

0002:01:47:51.61 DT mDTIscpRcv:ACK sent for seq no 4

Mac-learning on DTSs

Switch-1

HswitchDT-2# show mac-address

Status and Counters - Port Address Table

MAC Address	Port	VLAN
00fd45-b21360	Trk2	1
00fd45-b2137c	Trk2	1
00fd45-b2137d	Trk2	1
94f128-0d03c0	Trk1	1
94f128-0d03fe	Trk1	1
94f128-0d03ff	Trk1	1
98f2b3-96a700	A11	1
98f2b3-96b7f5	A11	1
98f2b3-96a700	A12	100
98f2b3-96b7f4	A12	100
98f2b3-96a700	A11	204
f40343-344cf6	A11	204

Switch -2

SwitchDT-1#show mac-address

Status and Counters - Port Address Table

MAC Address	Port	VLAN
00fd45-b21360	Trk2	1
00fd45-b2137c	Trk2	1
00fd45-b2137d	Trk2	1
94f128-0d03c0	Trk1	1
94f128-0d03fe	Trk1	1
94f128-0d03ff	Trk1	1
98f2b3-968f00	A11	1
98f2b3-969ff5	A11	1
98f2b3-968f00	A12	100
98f2b3-969ff4	A12	100
98f2b3-968f00	A11	204
f40343-344cf6	A20	204

Unicast traffic and Multicast flooding

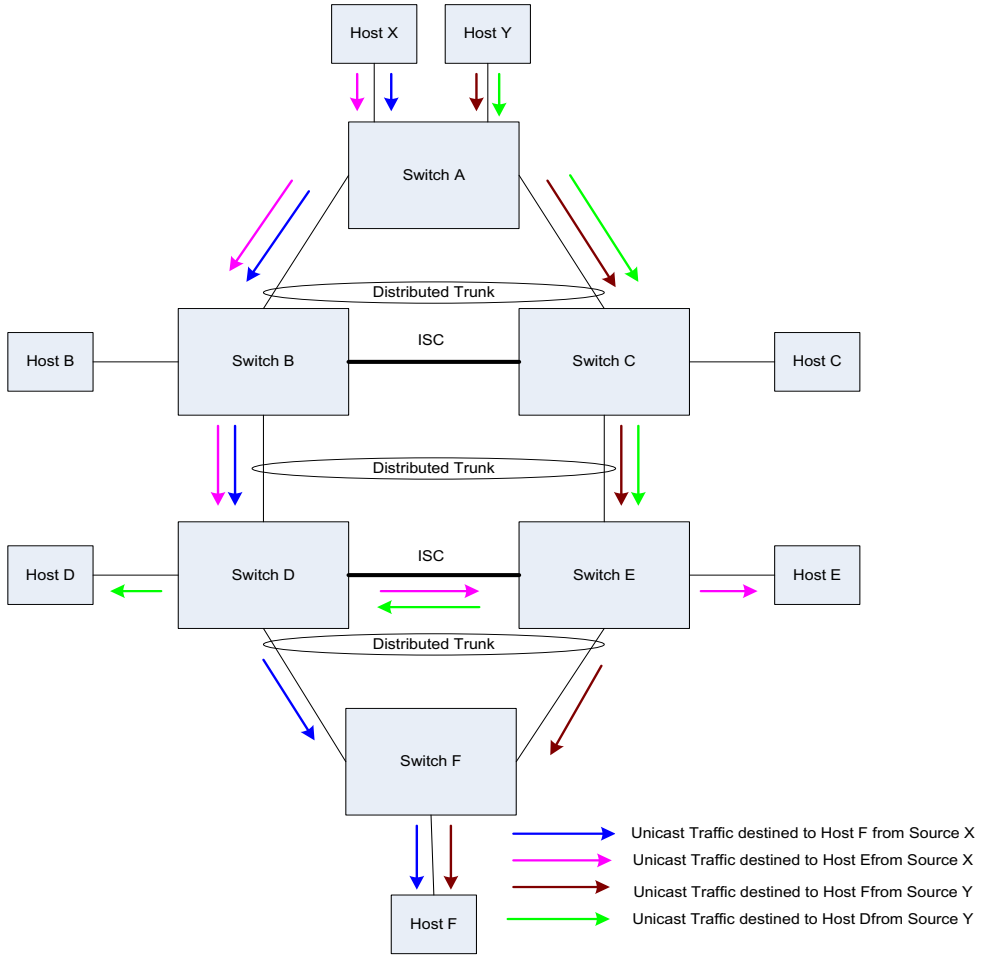

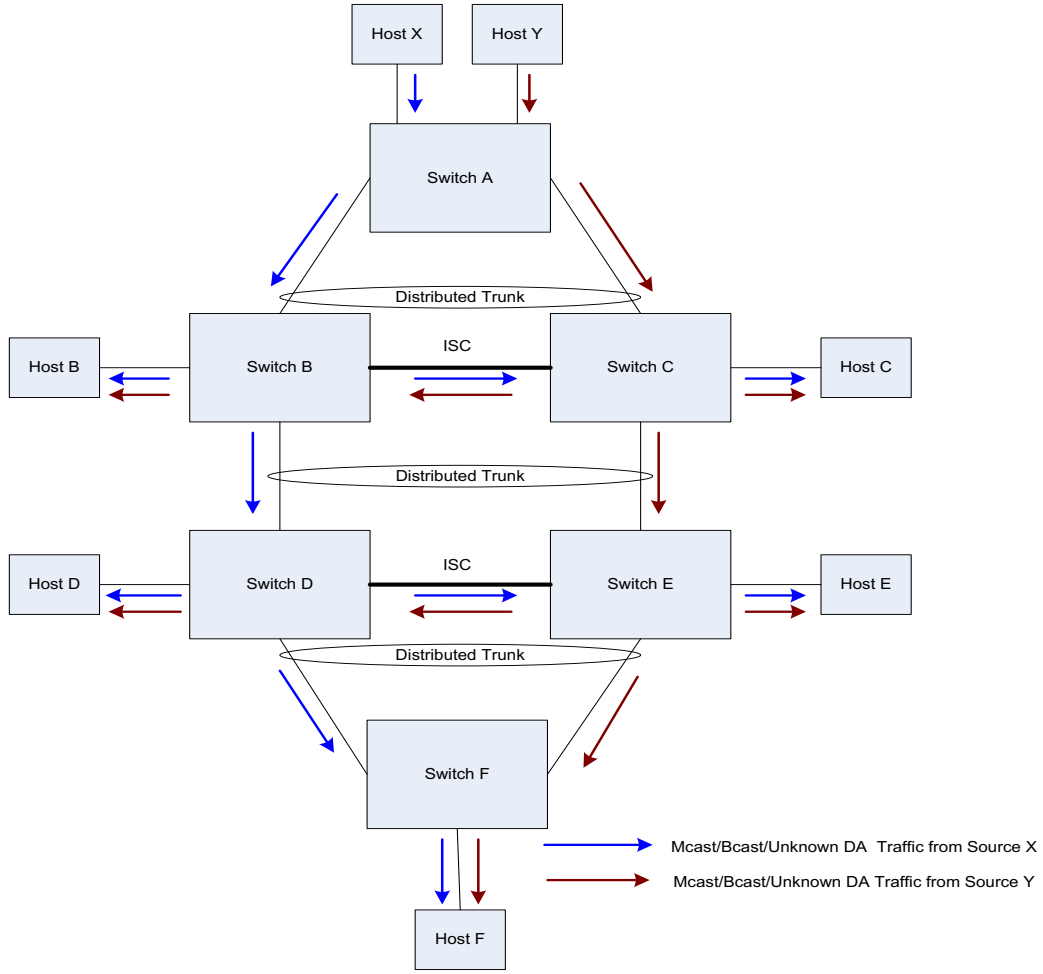

ISCP packet format

```
    17 2018-08-25 22:04:24.227280 HewlettP_84:0c:80 Slow-Protocols Slow Protocols  82 Subtype = 131
    18 2018-08-25 22:04:24.227922 HewlettP_86:b5:80 Slow-Protocols Slow Protocols  64 Subtype = 131
    19 2018-08-25 22:04:24.227937 HewlettP_86:b5:80 Slow-Protocols Slow Protocols  64 Subtype = 131
    20 2018-08-25 22:04:24.670434 HewlettP_84:0c:ff LLDP_Multicast LLDP           287 NoS = 70:10:6f:
    21 2018-08-25 22:04:30.923682 HewlettP_86:b5:80 Slow-Protocols Slow Protocols  64 Subtype = 131
    22 2018-08-25 22:04:34.167376 HewlettP_84:0c:80 Slow-Protocols Slow Protocols  64 Subtype = 131
    23 2018-08-25 22:04:40.923925 HewlettP_86:b5:80 Slow-Protocols Slow Protocols  64 Subtype = 131

▶ Frame 17: 82 bytes on wire (656 bits), 82 bytes captured (656 bits)
▼ Ethernet II, Src: HewlettP_84:0c:80 (70:10:6f:84:0c:80), Dst: Slow-Protocols (01:80:c2:00:00:02)
  ▼ Destination: Slow-Protocols (01:80:c2:00:00:02)
      Address: Slow-Protocols (01:80:c2:00:00:02)
      .... ..0. .... .... .... .... = LG bit: Globally unique address (factory default)
      .... ...1 .... .... .... .... = IG bit: Group address (multicast/broadcast)
  ▼ Source: HewlettP_84:0c:80 (70:10:6f:84:0c:80)
      Address: HewlettP_84:0c:80 (70:10:6f:84:0c:80)
      .... ..0. .... .... .... .... = LG bit: Globally unique address (factory default)
      .... ...0 .... .... .... .... = IG bit: Individual address (unicast)
    Type: Slow Protocols (0x8809)
▼ Slow Protocols
    Slow Protocols subtype: Unknown (0x83)
▼ Data (67 bytes)
    Data: 0400031b230200000035000200040000000000000000000...
    [Length: 67]
```

Keep alive packets

```
    46 2018-08-25 23:03:46 100.0.0.2   100.0.0.1   UDP  96 1024 → 1024 Len=50
    47 2018-08-25 23:03:47 100.0.0.1   100.0.0.2   UDP  96 1024 → 1024 Len=50
    48 2018-08-25 23:03:47 100.0.0.2   100.0.0.1   UDP  96 1024 → 1024 Len=50
    49 2018-08-25 23:03:48 100.0.0.1   100.0.0.2   UDP  96 1024 → 1024 Len=50
    50 2018-08-25 23:03:48 100.0.0.2   100.0.0.1   UDP  96 1024 → 1024 Len=50
```

```
▶ Frame 47: 96 bytes on wire (768 bits), 96 bytes captured (768 bits)
▶ Ethernet II, Src: HewlettP_84:0c:80 (70:10:6f:84:0c:80), Dst: HewlettP_86:b5:80 (70:10:6f:86:b5:80)
▶ Internet Protocol Version 4, Src: 100.0.0.1, Dst: 100.0.0.2
▼ User Datagram Protocol, Src Port: 1024, Dst Port: 1024
    Source Port: 1024
    Destination Port: 1024
    Length: 58
    Checksum: 0x315b [unverified]
    [Checksum Status: Unverified]
    [Stream index: 0]
▼ Data (50 bytes)
    Data: 70106f840c800000000000000000002000000000000000...
    [Length: 50]
```

Troubleshooting commands

#show lacp distributed	→ distributed trunks and LACP status
#show distributed- trunk consistency- parameters global	→ compare for the configuration mismatch between DTSs
# show switch-interconnect	→ status of ISC link
# show distributed-trunking status	→view the status of Distributed
#show distributed-trunking config	→view the status of the configuration of Distributed
#show distributed-trunking peer-keepalive	→ view the peer-keepalive link parameters
#show log DT	→view the log events specifically for Distributed Trunking

Debug DT packets on the switch

#debug distributed-trunking keepalive
#debug distributed-trunking iscp
Debug destination session or Debug destination buffer

QUESTIONS?

aruba
a Hewlett Packard
Enterprise company

THANK YOU!

aruba
a Hewlett Packard
Enterprise company