

Configure, Verify, and Debug Link Aggregation Control Protocol (LACP)

Introduction

Link Aggregation Control Protocol (LACP) is a control protocol that automatically detects multiple links between two LACP enabled devices and configures them to use their maximum possible bandwidth by automatically trunking the links together.

Overview of how LACP works

When LACP is enabled on the switch, all ports come under LACP control by default. Ports can be removed from the control of LACP. This will stop LACP frames from being sent from the ports and will be desirable if the devices connected to those ports are not LACP capable.

LACP controlled ports can be in one of two modes, either Active or Passive. A port in Active mode will send LACP Data Units (LACPDU) at regular intervals to seek out partners. A Passive mode port will only begin sending out LACPDUs in response to a received LACPDU. Ports are in Active mode by default.

When we enable LACP, it will send LACPDUs down each 'Active mode' connected link to find any partners that also have LACP enabled. If a reply is received, the device builds up a database of which link goes to which partner device. When LACP detects that two or more links are connected to the same partner device, have the same 'key' (details on keys below), and two or more of this group are not specified to be 'individual', it will aggregate them into one logical link.

Any further physical links added to the same partner system will simply be added to the now existing trunk group (within physical bounds). The same actions will take place if a port is in passive mode and it receives a LACP control packet - it will build a database of connected partners and begin sending control packets out that interface.

Details of the Protocol Exchange

As mentioned, the LACP protocol involves the exchange of LACPDU between the devices at either end of an Ethernet link. In discussions of LACP, the two devices are typically referred to as the "actor" (the device under consideration) and the "partner".

Table 1 shows items of information in the LACPDUs exchanged between actor and partner devices.

Table 1: LACPDUs exchanges between actor and partner

Items (relevant to Actor & Partner)	Description
Port number	The port that the LACPDU was emitted from
System ID	Typically, the MAC address
Port key	
Active/Passing setting	
Timeout setting	This indicates how long a period it waits before timing out the partner/actor
Port's aggregation setting	Whether the port is allowed to become part of an aggregation, or must stay individual
Port's synchronisation setting	If the LACP parameters on the device have recently been updated, but the new settings have not yet been committed to the hardware controlling this port, then the synchronisation status will be 'out of sync'; otherwise it will be 'in sync'
Port's collecting status	Whether or not the receiving of packets is currently enabled on the port
Port's distributing status	Whether or not the transmitting of packets is currently enabled on the port

Of course, the 'partner' items that are sent in the LACPDU are the actor's current view of the partner's status.

After the exchange of a few LACPDUs, the actor and the partner should have agreed on each other's status, and verified that each has correctly recognized the other's status (by checking the values that the other device is sending in the 'partner' fields of the LACPDUs).

Once the actor and partner have come to agreement about each other's settings, the devices can make a decision about whether the ports at each end of the link can be added to an aggregation.

The values exchanged and settings established can be seen for a specific LACP port with the command `SHOW LACP PORT=X` as follows shown in Figure 1 on page 3.

Figure 1: Example output of SHOW LACP PORT=2

```
LACP Port Information
-----

Actor Port ..... 2
  Trunk Group ..... lacp26
  Selected ..... Selected
  Port Priority ..... 32768
  LACP Port Number ..... 2
  Port Key ..... 2
    Admin Key ..... 1
  Mode ..... Active
  Periodic ..... Fast
  Individual ..... No
  Synchronised ..... Yes
  Collecting ..... Yes
  Distributing ..... Yes
  Defaulted ..... No
  Expired ..... No
  Actor Churn ..... No
  Partner Churn ..... No

Partner Information:
  Partner Sys Priority ..... 32768
  Partner System .. 00-00-cd-24-02-32
  Port Key ..... 2
  Port Priority ..... 32768
  Port Number ..... 2
  Mode ..... Active
  Periodic ..... Fast
  Individual ..... No
  Synchronised ..... Yes
  Collecting ..... Yes
  Distributing ..... Yes
  Defaulted ..... No
  Expired ..... No
```

KEYS

As mentioned above, port keys are part of the exchange of LACPDUs between the actor and the partner devices. All LACP controlled ports have two keys, an admin key and a port key. The admin key is a value you can set, but the port key is calculated by the system.

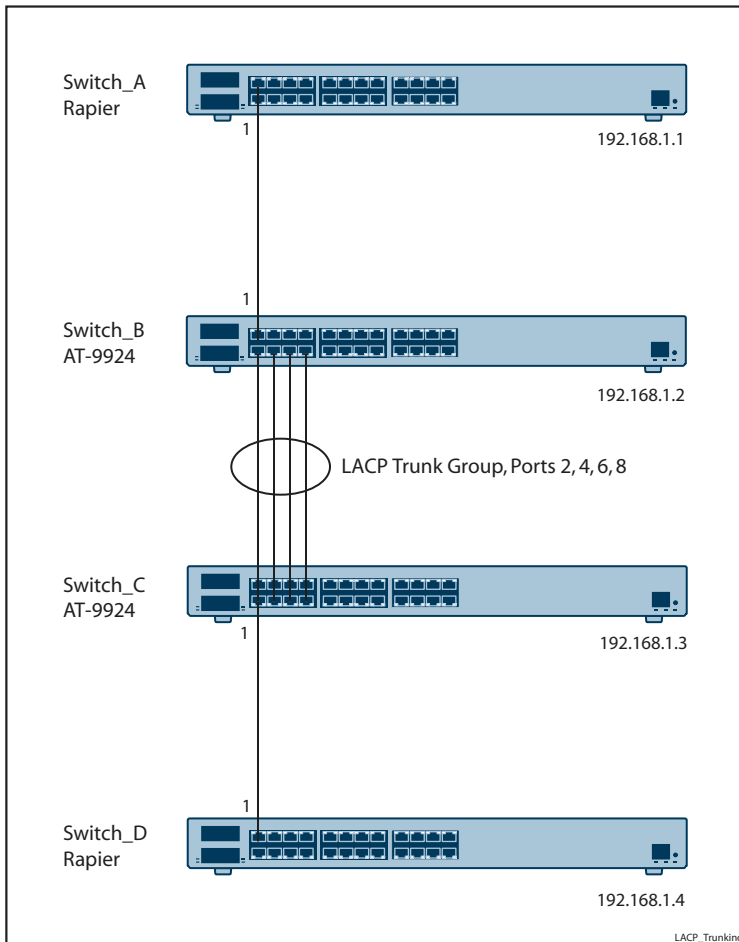
So, the admin key can be set for a specific port, and this will affect the LACP port key generated. This can be used to prevent otherwise aggregatable ports from forming a trunk. By default all ports will be given the same LACP admin port key, which in turn means that all ports that can be aggregated will generate the same LACP port key.

If you want to separate the ports of the switch into separate sets of aggregatable ports, then assign different admin keys to each set of ports.

Configuring LACP

We'll use the network diagram in Figure 2 as an example scenario to discuss configuring, verifying and debugging LACP.

Figure 2: Example network



Configuration of switches

Switch A Rapier

```
set sys name=Switch_A
enable ip
add ip interface-vlan1 ip=192.168.1.1
```

Switch B AT-9924

```
set sys name=Switch_B
enable lacp
```

```
delete lacp port=1

enable ip

add ip interface=vlan 1 ip=192.168.1.2
```

Switch C AT-9924

```
set sys name=Switch_C

enable lacp

delete lacp port=1

enable ip

add ip interface=vlan 1 ip=192.168.1.3
```

Switch D Rapier

```
set sys name=Switch_D

enable ip

add ip interface=vlan1 ip=192.168.1.4
```

LACP is enabled on the two AT-9924 switches B and C, but port 1 is deleted from LACP control, as this port is connected, in each case, to a Rapier switch not running LACP. Ports 2,4,6 & 8 are connected between the LACP enabled AT-9924 switches B and C. LACPDUs are sent between these ports and they form a dynamic trunk group.

The switches provide some SHOW commands via which the operation of the LACP facility can be monitored.

The command SHOW LACP outputs the values of parameters that govern the LACP process.

Figure 3: Example output of SHOW LACP

```
Manager Switch_B> sh lacp

LACP Information
-----
  Status ..... Enabled
  Actor System Priority ..... 32768
  Actor System ..... 00-00-cd-24-02-2b
  LACP Ports ..... 2-24
    Active ..... 2-24
    Passive ..... None
-----
```

LACP shows as enabled with ports 2-24 operating in active mode. These ports will be sending regular LACPDU messages (if connected).

Let us now look at some of the detailed information that can be presented about the states of these ports:

Figure 4: Example output of SHOW LACP PORT

```

Manager Switch_B> sh lacp port=2,4

LACP Port Information
-----

Actor Port ..... 2
  Trunk Group ..... lacp26
  Selected ..... Selected
  Port Priority ..... 32768
  LACP Port Number ..... 2
  Port Key ..... 2
  Admin Key ..... 1
  Mode ..... Active
  Periodic ..... Fast
  Individual ..... No
  Synchronised ..... Yes
  Collecting ..... Yes
  Distributing ..... Yes
  Defaulted ..... No
  Expired ..... No
  Actor Churn ..... No
  Partner Churn ..... No

Partner Information:
  Partner Sys Priority ..... 32768
  Partner System .. 00-00-cd-24-02-32
  Port Key ..... 2
  Port Priority ..... 32768
  Port Number ..... 2
  Mode ..... Active
  Periodic ..... Fast
  Individual ..... No
  Synchronised ..... Yes
  Collecting ..... Yes
  Distributing ..... Yes
  Defaulted ..... No
  Expired ..... No

Actor Port ..... 4
  Trunk Group ..... lacp26
  Selected ..... Selected
  Port Priority ..... 32768
  LACP Port Number ..... 4
  Port Key ..... 2
  Admin Key ..... 1
  Mode ..... Active
  Periodic ..... Fast
  Individual ..... No
  Synchronised ..... Yes
  Collecting ..... Yes
  Distributing ..... Yes
  Defaulted ..... No
  Expired ..... No
  Actor Churn ..... No
  Partner Churn ..... No

Partner Information:
  Partner Sys Priority ..... 32768
  Partner System .. 00-00-cd-24-02-32
  Port Key ..... 2
  Port Priority ..... 32768
  Port Number ..... 6
  Mode ..... Active
  Periodic ..... Fast
  Individual ..... No
  Synchronised ..... Yes
  Collecting ..... Yes
  Distributing ..... Yes
  Defaulted ..... No
  Expired ..... No

```

Ports 2 and 4 (which are part of the LACP trunk group of 2,4,6 & 8) are shown above to be in the same trunk group named lacp26. The name of the trunk group to which the port belongs is automatically generated. All LACP trunk group names begin with the letters "LACP" and have the Interface Index appended.

In this particular example the AT-9924 switch has 24 ports, which will have interface index numbers of 1-24. VLAN 1 will have an interface index of 25 and the next available interface index number for our LACP trunk group is 26. (these values can be seen by running the command SH INTERFACE on the switch).

Other information that can be seen includes:

- The mode - Active mode
- The LACPDU periodic time set, which is set to 'fast'

The 'fast' setting causes the port to send LACPDU's every second (the other option is 'slow' which send LACPDU's every 30 seconds)

- Port and Admin keys
- Port priority

The meaning of these values is discussed in the 'Details of the Protocol Exchange' section above.

The information on the port connected at the other end of the link (partner information), can be useful in debugging to check that all settings are correct.

Figure 5: Example output of SHOW LACP PORT COUNT

```
Manager Switch_B> sh lacp port=2,4 count

LACP Port Counters
-----

Port ..... 2
  Received:
    LACP Pkts ..... 910
    Invalid LACP Pkts ..... 0
  Transmitted:
    LACP Pkts ..... 910

Port ..... 4
  Received:
    LACP Pkts ..... 910
    Invalid LACP Pkts ..... 0
  Transmitted:
    LACP Pkts ..... 910
-----
```

Counters are shown for the number of LACPDU's sent and received. Invalid LACPDU's are also counted, which is helpful for debugging purposes.

Figure 6: Example output of SHOW LACP PORT TRUNK

```
Manager Switch_B> sh lacp trunk

LACP Dynamic Trunk Group Information
-----

Trunk group name ..... lacp26:
  Speed ..... 1000 Mbps
  Select ..... source and destination mac address
  Ports in Trunk ..... 2,4,6,8
  LAG ID:
  [(8000,00-00-cd-24-02-2b,0002,00,0000) ,(8000,00-00-cd-24-02-
32,0002,00,0000)]
-----
```

The command SHOW LACP TRUNK used in Figure 6, shows information about the configured trunk group (lacp26) showing ports 2, 4, 6, & 8 as members. One of the parameters shown is the LAG ID as detailed in Table 2 on page 8.

Table 2: The Link Aggregation Group Identifier (LAG ID) shows the following information for the Actor and Partner devices

Actor	Partner
8000	System Priority
00-00-cd-24-02-2b	Switch MAC
0002	Key
00	Port Priority (only shows for 'individual' ports)
0000	Port Number (only shows for 'individual' ports)

Fine Tuning LACP

There are various settings that can be adjusted to fine tune the operation of LACP on the switch and in the Network. Some of these setting are as follows:

Mode

LACP controlled ports can be in one of two modes, either Active or Passive. A port in Active mode will send LACPDU (LACP Data Units) at regular intervals to seek out partners. A Passive mode port will only begin sending out LACPDUs in response to a received LACPDU. Ports are in Active mode by default.

LACPDU Period

The PERIODIC parameter specifies the speed the LACP port will request that the partner send LACPDU update messages. A port in fast mode will receive one LACPDU per second. A port in slow mode will receive one LACPDU per thirty seconds. The default is Fast.

Key Values

As mentioned earlier the admin key can be set for a specific port, which affects the LACP port key generated. This can be used to prevent otherwise aggregated ports from forming a trunk. By default, all ports will be given the same LACP admin port key, which in turn means that all ports that can be aggregated will generate the same LACP port key.

Port Priority

The LACP port priority can be specified. This value is used to decide which ports should be selected when being added to a trunk group (if the hardware can handle less links than physically present between the two devices). Excess ports will be put into a 'standby' mode, which is effectively disabled, but they will take the place of a link in a trunk that goes down. The default value is 500 (one is high).

The **Mode**, **LACPDU Period**, **Key Value** and **Port Priority** are set with either of the following commands:

```
ADD LACP PORT=[{port-list|ALL}] [ADMINKEY=key] [PRIORITY=priority]
[MODE={ACTIVE|PASSIVE}] [PERIODIC={FAST|SLOW}]
```

```
SET LACP PORT=[{port-list|ALL}] [ADMINKEY=key] [PRIORITY=priority]
[MODE={ACTIVE|PASSIVE}] [PERIODIC={FAST|SLOW}]
```


Ports Removed

Ports can be removed from the control of LACP. This will stop LACP frames from being sent across the link and will be desirable if the device on the other end is not LACP capable. Ports are removed from LACP control with the following command:

```
DELETE LACP PORT={port-list|ALL}
```

Device Priority

The PRIORITY parameter specifies the value that this device's priority will take. This specifies which device can make changes to its configuration (the "master" device has the lowest System ID: priority:SystemMacAddress). The default is 32768. Set with the following command:

```
SET LACP PRIORITY=priority
```

Debugging LACP

Using the SHOW commands is the first step in debugging an LACP problem to confirm that settings are correct on all devices in a network running LACP.

LACP debugging is also available with the following command:

```
ENABLE LACP DEBUG=[MSG|PACKET|STATE|ALL NON_DOC |TRACE|DEV|PERSISTENT]
```

Table 3: Options of information on the screen output

Message	Decode of LACPDUs being exchanged
Packet	All LACP incoming and outgoing packets
State	Internal state machine changes
All	Message, Packet and State information

An example output is shown in Figure 7 on page 10.

Figure 7: Example output of ENABLE LACP DEBUG

```
Manager Switch_B> ena lacp debug=all

Info (1122003): Operation successful.

Manager Mustang_1>
T lacpApiPortManagedByLacpGet
M Rxing on Port 4

Dest MAC = 01-80-c2-00-00-02, Src MAC = 00-00-cd-24-02-32
Eth type = 34825, subtype = 1, version = 1

TLV type (Actor) = 1, TLV info Length = 20
Priority = 32768, ID = 00-00-cd-24-02-32
Key = 2, Port Priority = 32768, Port = 6
Activity = 1, Timeout = 1
Aggregation = 1, Synchronisation = 1
Collecting = 1, Distributing = 1
Defaulted = 0, Expired = 0

TLV type (Partner) = 2, TLV info Length = 20
Priority = 32768, ID = 00-00-cd-24-02-2b
Key = 2, Port Priority = 32768, Port = 4
Activity = 1, Timeout = 1
Aggregation = 1, Synchronisation = 1
Collecting = 1, Distributing = 1
Defaulted = 0, Expired = 0

TLV type (Collector) = 3, TLV info Length = 16
Collector delay = 0

TLV type (Terminator) = 0, TLV info Length = 0

P
LACP Port4 Rx: 0180c200 00020000 cd240232 88090101 01148000 0000cd24
02320002 80000006 3f000000 02148000 0000cd24 022b0002 80000004 3f000000
03100000
00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000
00000000
00000000 00000000 00000000 00000000 00000000 00000000 00000000
S Port 4 Receive Machine nextState: Current
S Port 4 Receive Machine nextState: No Change
S Port 8 Periodic Tx Machine nextState: Periodic Tx
S Port 8 Periodic Tx Machine nextState: Fast Periodic
S Port 8 Periodic Tx Machine nextState: No Change
M Txing on Port 8

Dest MAC = 01-80-c2-00-00-02, Src MAC = 00-00-cd-24-02-2b
Eth type = 34825, subtype = 1, version = 1
```

The 'message' information (highlighted above) is very useful in debugging to confirm that the values sent between the Actor and Partner devices is consistent with values seen in the show commands for specific ports.

Packet and state debug information can be seen following the message debug.

