

Link Aggregation and Ethernet Bonding

Feature Overview and Configuration Guide

Introduction and overview

This guide includes two sample configurations for Link Aggregation Control Protocol (LACP), or dynamic channel groups, as well as a sample configuration for a static channel group. Additionally, it offers a straightforward example of Ethernet bonding on the TQ6702 GEN2-R, a Wi-Fi 6 (802.11ax) Wireless AP Router.

What is link aggregation?

Link aggregation is a key component in resilient network design, since it increases the available bandwidth between network devices and it provides continuity of connectivity if one link is broken between network devices.

By aggregating two or more links together, you can increase the bandwidth between neighboring devices since this is effectively additive, where two links give up to twice the bandwidth of one link. Having more than one link to a neighboring device provides connectivity if one of the links break, where a feature of this resiliency is the speed at which link aggregation reacts to the change of link status in a matter of millisecond.

A link aggregation can only exist between a pair of neighboring switches, where the switch ports that are aggregated on one switch cannot be connected to switch ports that are not aggregated on the other switch. A switch can have multiple link aggregations to different neighbors, or even to the same neighbor if the network is loop protected.

Link aggregation is a broad term that can encompass various types of network links.

What is Ethernet bonding?

In the networking industry, link aggregation and Ethernet bonding are terms often used interchangeably as they both refer to the practice of combining multiple physical network links into a single logical link to increase bandwidth, redundancy, and reliability.

For Allied Telesis products, link aggregation is used to refer to static or dynamic (LACP) aggregation configured on switch ports, where the switch chip manages the link aggregation and the packet hashing and forwarding process is handled in hardware. Link aggregation interfaces have names like 'po1' and 'sa1' depending on the aggregation method in use.

Ethernet bonding is used to refer to static or dynamic (LACP) aggregation configured on router Ethernet ports, where the router's CPU manages the link aggregation and the packet hashing and forwarding process is handled in software. Ethernet bond interfaces have names like 'bond1' and the name is independent of the aggregation method in use by the bond.

Products and software version that apply to this guide

This guide applies to all AlliedWare Plus™ products, running version **5.4.4** or later.

However, support for link aggregation varies between products. For details, see the following documents:

- The [product's Datasheet](#)
- The [AlliedWare Plus Datasheet](#)
- The product's [Command Reference](#)

These documents are available from the above links on our website at alliedtelesis.com.

Feature support may change in later software versions. For the latest information, see the above documents.

Content

Introduction and overview	1
What is link aggregation?	1
What is Ethernet bonding?	2
Products and software version that apply to this guide	2
Static and Dynamic (LACP) Link Aggregation	4
Aggregation criteria	4
Static channel groups	4
Dynamic (LACP) channel groups	4
Load balancing across aggregated links	5
Understanding LACP	5
LACP modes and Data Units	5
Details of the protocol exchange	6
Fine tuning LACP	8
Configuration examples	10
Configuring an LACP channel group	10
Minimal LACP group configuration	13
Configuring a static channel group	13
Configuring Ethernet Bonding on the TQ6702 GEN2-R	14

Static and dynamic (LACP) link aggregation

Channels, either static or dynamic (LACP), increase reliability by distributing the data path over more than one physical link. Channels must be configured on both ends of a link or network loops may result. Ports in a channel group need not be contiguous. A mirror port cannot be a member of either a static or a dynamic channel group.

Aggregation criteria

For individual links to be aggregated into a channel group they must:

- originate on the same device or stack
- terminate on the same device or stack
- be members of the same VLANs
- have the same data rate
- share the same admin port key
- be operating in full duplex mode

The hardware must also be capable and have the capacity to handle the number of links to be aggregated.

Static channel groups

A static channel group, also known as a static aggregator, enables a number of ports to be manually configured to form a single logical connection of higher bandwidth. By using static channel groups you increase connection reliability by distributing the data path over more than one physical link. Static channel groups are best used in simpler environments, usually where neighbor switches are close together, e.g. situated within the same rack, so that you can easily ensure that the correct statically aggregated ports are connected together.

For a static channel group configuration example see the ["This creates the channel group."](#) on [page 13](#) section in this guide.

Dynamic (LACP) channel groups

A LACP channel group, also known as an etherchannel, an LACP aggregator, or a dynamic channel group, enables a number of ports to be dynamically combined to form a single higher bandwidth logical connection. LACP channel groups are best used for complex environments, typically long-distance links, to detect failure between neighbor switches.

For LACP configuration examples see ["Configuration examples"](#) on [page 10](#) and ["Minimal LACP group configuration"](#) on [page 13](#) in this guide.

Load balancing across aggregated links

Link aggregation does not necessarily achieve exact load balancing across the links. The load sharing algorithm is designed to ensure that any given data flow always goes down the same link. It also aims to spread data flows across the links as evenly as possible.

Link aggregation hashes one or more of the source and destination MAC address, IP address and UDP/TCP ports to select a link on which to send a packet. So packet flow between a pair of hosts always takes the same link inside the Link Aggregation Group (LAG). The net effect is that the bandwidth for a given packet stream is restricted to the speed of one link in the LAG.

Understanding LACP

With **static** link aggregation, there is not a great deal to understand—the switch treats the aggregation as effectively a single port, and uses a hashing algorithm to share data across whichever member ports are link-up. Dynamic link aggregation however requires the communication protocol LACP.

LACP definition

LACP is a control protocol that automatically detects multiple links between two LACP enabled devices and enables them to use their maximum possible bandwidth by automatically aggregating the links together.

LACP modes and Data Units

Each aggregation link created by LACP is referred to as a dynamic channel group. Ports in a dynamic channel group can be in one of two modes:

- **LACP active** mode

A port in active mode sends LACP Data Units (LACPDU) at regular intervals to seek out partners.

- **LACP passive** mode

A port in passive mode only begins sending out LACPDUs in response to a received LACPDU.

When ports are added to a dynamic channel group, each connected active mode port sends LACPDUs to find any partner devices that also have LACP enabled. If a port receives a reply, the switch uses the reply to help build a map of connected LACP partners and the links that they share. When LACP detects that two or more links are connected to the same partner, and have the same key (the channel group ID), it aggregates them into one logical link.

When you add any further physical links to the same partner system, the links are added to the already existing aggregation (within the hardware limits on each switch).

Ports in passive mode do not actively seek out LACP partners. However, if a passive mode port receives an LACPDU, then the device it is on includes the new link details in a map of connected LACP partners, and begins sending LACP control packets out of the interface.

Details of the protocol exchange

The LACP protocol involves the exchange of LACPDU between the devices at either end of a link. Each switch refers to itself as the **actor** and refers to the switch that it is negotiating with as its **partner**. The information that the switches exchange is shown in the next table.

ITEMS	DESCRIPTION
Port number	The port that the LACPDU was emitted from.
SystemID	A value used to uniquely identify a switch, so that each switch can work out which of its ports are connected to the same neighbor. The SystemID also acts as a priority value, to enable neighboring switches to decide which switch's configuration takes priority if neighbors share more ports than they can aggregate together.
Port key	Used to identify groups of ports as being available to aggregate together. Ports that are in the same dynamic channel group must all have the same value.
Timeout setting	This indicates the rate that the switch expects to receive LACPDU from its partner.
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 port is allowed to receive packets.
Port's distributing status	Whether or not the port is allowed to transmit packets.

Each LACPDU packet contains an **actor** section, in which the switch sends its own values for those parameters, and a **partner** section, in which it sends its current view of the neighbor's values for these parameters.

The next output (collected by a packet sniffing program) shows the contents of an LACPDU for a port that the LACP neighbors have not yet established as part of an aggregated link.

```

Link Aggregation Control Protocol
  Slow Protocols subtype: LACP (0x01)
  LACP Version Number: 0x01
  Actor Information: 0x01
  Actor Information Length: 0x14
  Actor System Priority: 32768
  Actor System: Allied_5f:f5:80 (000a.b75f.f580)
  Actor Key: 2
  Actor Port Priority: 32768
  Actor Port: 5
  Actor State: 0xc5 (Activity, Aggregation, Defaulted, Expired)
    .... ...1 = LACP Activity: Yes
    .... ..0. = LACP Timeout: No
    .... .1.. = Aggregation: Yes
    .... 0... = Synchronization: No
    ...0 .... = Collecting: No
    ..0. .... = Distributing: No
    .1.. .... = Defaulted: Yes
    1... .... = Expired: Yes
  Reserved: 000000
  Partner Information: 0x02
  Partner Information Length: 0x14
  Partner System Priority: 0
  Partner System: 00:00:00_00:00:00 (00:00:00:00:00:00)
  Partner Key: 0
  Partner Port Priority: 0
  Partner Port: 0
  Partner State: 0x02 (Timeout)
    .... ...0 = LACP Activity: No
    .... ..1. = LACP Timeout: Yes
    .... .0.. = Aggregation: No
    .... 0... = Synchronization: No
    ...0 .... = Collecting: No
    ..0. .... = Distributing: No
    .0.. .... = Defaulted: No
    0... .... = Expired: No
  Reserved: 000000

```

After the exchange of a few LACPDUs, the neighbors will have agreed on each other's status. They will also have verified that each has correctly recognized the other's status, by checking the values that their neighbor has sent in the partner fields of the LACPDUs.

Once the neighbors have agreed about each other's settings on a given pair of ports, then those ports are said to be synchronized. The devices can then make a decision about whether they can add the ports at each end of the link to an aggregation.

The values exchanged and settings established can be seen for a specific LACP port with the command **show port etherchannel <port>**.

Fine tuning LACP

There are various settings that you can adjust to fine tune the operation of LACP on a switch and in a network. This section describes the common settings to adjust.

Mode

You can change the mode on links between either passive or active mode. See "[LACP modes and Data Units](#)" on page 5, for more information about the modes.

We recommend that you **always use active mode on links that you want to aggregate**, as there is no downside to using active mode. However, passive mode can have a downside; if both ends of a link are configured for passive mode, then no LACP negotiation will ever occur on that link.

The main purpose for passive mode is to control accidental loops. Passive mode requires one of the neighbor switches to be in active mode.

LACPDU timeout

You can change the rate that your switch expects to receive LACPDUs from its neighbor to either short or long.

If the timeout is set to:

- **long** - the switch expects to receive an LACPDU every 30 seconds. It will time a port out of the aggregation if it does not receive an LACPDU for 90 seconds (this means three consecutive LACPDUs are lost).
- **short** - the switch expects to receive an LACPDU every second. It will time a port out of the aggregation if no LACPDUs are seen for three seconds (this means three consecutive LACPDUs are lost).

The switch indicates its preference using the timeout field in the actor section of its LACPDUs. If that field is set to **1**, then the switch has set the **short** timeout. If the field is set to **0**, then the switch has set the **long** timeout.

Setting the short timeout enables the switch to be more reactive to communication failure on a link. It also does not add much processing overhead to a switch—one packet per second, which is insignificant.

Because of this timeout setting, it is not possible to configure the rate at which the switch sends LACPDUs. A switch must always send LACPDUs at the rate which its neighbor has indicated it expects to receive them.

The interface-mode command **lacp timeout {short|long}** enables the switch to indicate the rate at which it expects to receive LACPDUs from its neighbor.

Port priority

This value deals with the situation where the number of ports that have been configured into a dynamic channel group is larger than the number that the switch hardware can accommodate in a single aggregation. The priority value is used to decide which ports should be included into the aggregation. The higher priority ports (those with a lower priority value) are selected ahead of the lower priority ports.

Excess ports are put into a standby mode, leaving them effectively disabled. However, if a link in their channel group goes down, they will take the place of that link.

The LACP port priority can be specified with the interface-mode command **lacp port-priority**. The default value is 32768 (1 is high).

System priority

LACP neighbors can disagree about port priorities, with one switch configured to leave a certain set of links out of the aggregation, while its neighbor is configured to leave out a different set. This means that LACP requires a mechanism to decide which LACP neighbor's port priority setting is used.

LACP combines the System Priority value with the MAC address of the switch to create a System ID. The switch with the lower System ID becomes the master switch, and the other switch must fall into line with whatever decision the master makes.

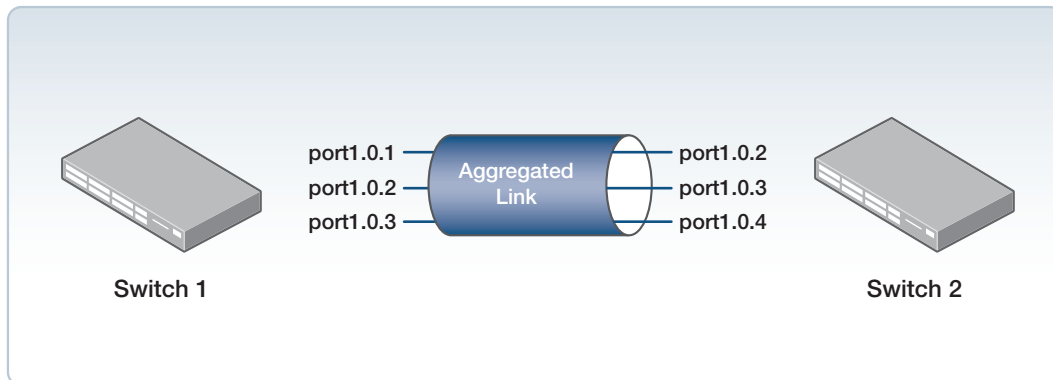
The system priority of a switch can be set in global configuration mode, using the command **lacp system-priority**.

Note: AlliedWare Plus supports IEEE 802.3ad link aggregation and uses the Link Aggregation Control Protocol (LACP). LACP does not interoperate with devices that use Port Aggregation Protocol (PAgP).

Configuration examples

Configuring an LACP channel group

The following example shows how to configure three links between two Allied Telesis managed Layer 3 switches. The three links are assigned the same administrative key (I), so that they aggregate to form a single channel (I). They are viewed by the Spanning Tree Protocol as one interface.



Switch 1 Step 1: Enter Global Configuration mode

```
awplus# configure terminal
```

Step 2: Set the system priority of the switch

```
awplus(config)# lacp system-priority 2000
```

This priority is used to determine which switch in the system is responsible for resolving conflicts in the choice of aggregation groups. A lower numerical value has a higher priority. Switch 1 has a higher priority than Switch 2 in this configuration.

Step 3: Configure port 1.0.1

```
awplus(config)# interface port1.0.1
```

Step 4: Add the interface to channel group 1

```
awplus(config-if)# channel-group 1 mode active
```

Step 5: Exit Interface Configuration mode

```
awplus(config-if)# exit
```

Step 6: Configure port 1.0.2

```
awplus(config)# interface port1.0.2
```

Step 7: Add the interface to channel group 1

Enable link aggregation so that it may be selected for aggregation by the local system.

```
awplus(config-if)# channel-group 1 mode active
```

Step 8: Exit Interface Configuration mode

```
awplus(config-if)# exit
```

Step 9: Configure port 1.0.3

```
awplus(config)# interface port1.0.3
```

Step 10: Add the interface to channel group1

Add this interface to a channel group **1** and enable link aggregation so that it may be selected for aggregation by the local system.

```
awplus(config-if)# channel-group 1 mode active
```

Step 11: Choose the dynamic aggregator logical interface 'po1' for channel-group 1.

```
awplus(config-if)# interface po1
```

Switch 2 Step 1: Enter Global Configuration mode

```
awplus# configure terminal
```

Step 2: Set the system priority of this switch

```
awplus(config)# lacp system-priority 3000
```

This priority is used to determine which switch in the system is responsible for resolving conflicts in the choice of aggregation groups. A lower numerical value has a higher priority. Switch 2 has a lower priority than Switch 1 in this configuration.

Step 3: Configure port 1.0.2

```
awplus(config)# interface port1.0.2
```

Step 4: Add the interface to channel group1

```
awplus(config-if)# channel-group 1 mode active
```

Step 5: Exit Interface Configuration mode

```
awplus(config-if)# exit
```

Step 6: Configure port 1.0.3

```
awplus(config)# interface port1.0.3
```

Step 7: Add the interface to channel group1

Enable link aggregation so that it may be selected for aggregation by the local system.

```
awplus(config-if)# channel-group 1 mode active
```

Step 8: Exit Interface Configuration mode

```
awplus(config-if)# exit
```

Step 9: Configure port 1.0.4

```
awplus(config)# interface port1.0.4
```

Step 10: Add the interface to channel group1

Add this interface to a channel group **1** and enable link aggregation so that it may be selected for aggregation by the local system.

```
awplus(config-if)# channel-group 1 mode active
```

Minimal LACP group configuration

The following minimal LACP group configuration example creates LACP channel group 2 and enables link aggregation on switch ports **1.0.1** and **1.0.2** within this channel group. Note that all aggregated ports must belong to the same VLAN.

Step 1: Enter Global Configuration mode

```
awplus# configure terminal
```

Step 2: Enter Interface Configuration for the switch ports to aggregate into the channel group

```
awplus(config)# interface port1.0.1-port1.0.2
```

Step 3: Assign the switch ports to channel group 2 in active mode

This creates the channel group.

```
awplus(config-if)# interface port1.0.2
```

Configuring a static channel group

The following example creates a static channel group and adds switch ports **1.0.1** and **1.0.2**.

Step 1: Enter Global Configuration mode

```
awplus# configure terminal
```

Step 2: Enter Interface Configuration to configure port 1.0.1

```
awplus(config)# interface port1.0.1
```

Step 3: Add port1.0.1 to static-channel-group 2

```
awplus(config-if)# static-channel-group 2
```

Step 4: Exit the Interface Configuration mode and return to the Global Configuration mode

```
awplus(config-if)# exit
```

Step 5: Configure port 1.0.2

```
awplus(config)# interface port1.0.2
```

Step 6: Add port1.0.2 to static-channel-group 2

```
awplus(config-if)# static-channel-group 2
```

Configuring Ethernet Bonding on the TQ6702 GEN2-R

The Ethernet bonding feature allows Ethernet ports on supported products to be combined to create higher bandwidth logical interfaces. Bonded interfaces can be set up to function in either a statically configured mode or an LACP aggregator mode. These modes are compatible with the static-channel and dynamic LACP link aggregation links supported on Allied Telesis switch products.

By default, Ethernet bonds operate in LACP mode but do not support any further LACP fine-tuning options. The mode can be changed using the “aggregator-type (static | lacp)” command.

Ethernet bonds are otherwise logically equivalent to underlying Ethernet ports and support the same configuration commands that Ethernet ports do. Ethernet bonds can also be assigned to bridge groups, or act as parent interfaces for dot1q or ppp encapsulated sub-interfaces.

Example

To configure the Ethernet ports on a wireless router to act as a bond:

Step 1: Enter Global Configuration mode

```
awplus# configure terminal
```

Step 2: Create a new bond interface with ID 1

```
awplus(config)# bond 1
```

Step 3: Enter Interface mode for ethernet1 and ethernet2 and assign to the bond

```
awplus(config)# interface eth1
awplus(config-if)# bond-group 1
awplus(config-if)# exit
awplus(config)# interface eth2
awplus(config-if)# bond-group 1
awplus(config-if)# exit
```

Step 4: (Optional) Change the aggregation mode of the bond interface to static

```
awplus(config)# interface bond1
awplus(config-if)# aggregator-type static
```

Step 5: Exit Interface Configuration mode

```
awplus(config-if)# exit
```

Once configured, the bonded Ethernet ports can be cabled to the appropriately configured aggregator/bond ports on other Allied Telesis products.

C613-22069-00 REV C



NETWORK SMARTER

North America Headquarters | 19800 North Creek Parkway | Suite 100 | Bothell | WA 98011 | USA | T: +1 800 424 4284 | F: +1 425 481 3895

Asia-Pacific Headquarters | 11 Tai Seng Link | Singapore | 534182 | T: +65 6383 3832 | F: +65 6383 3830

EMEA & CSA Operations | Incheonweg 7 | 1437 EK Rozenburg | The Netherlands | T: +31 20 7950020 | F: +31 20 7950021

alliedtelesis.com

© 2023 Allied Telesis, Inc. All rights reserved. Information in this document is subject to change without notice. All company names, logos, and product designs that are trademarks or registered trademarks are the property of their respective owners.