


Link aggregation

In [computer networking](#), **link aggregation** is the combining ([aggregating](#)) of multiple network connections in parallel by any of several methods, in order to increase [throughput](#) beyond what a single connection could sustain, to provide [redundancy](#) in case one of the links should fail, or both. A **link aggregation group (LAG)** is the combined collection of physical ports.



Link Aggregation between a switch and a server

Other umbrella terms used to describe the concept include **trunking**,^[1] **bundling**,^[2] **bonding**,^[1] **channeling**^[3] or **teaming**.

Implementation may follow vendor-independent standards such as [Link Aggregation Control Protocol \(LACP\)](#) for [Ethernet](#), defined in [IEEE 802.1AX](#) or the previous [IEEE 802.3ad](#), but also [proprietary protocols](#).

Motivation

Link aggregation increases the bandwidth and resilience of [Ethernet](#) connections.

Bandwidth requirements do not scale linearly. Ethernet bandwidths historically have increased tenfold each generation: 10 [megabit/s](#), 100 Mbit/s, 1000 Mbit/s, 10,000 Mbit/s. If one started to bump into bandwidth ceilings, then the only option was to move to the next generation, which could be cost prohibitive. An alternative solution, introduced by many of the network manufacturers in the early 1990s, is to use link aggregation to combine two physical Ethernet links into one logical link. Most of these early solutions required manual configuration and identical equipment on both sides of the connection.^[4]

There are three [single points of failure](#) inherent to a typical port-cable-port connection, in either a computer-to-switch or a switch-to-switch configuration: the cable itself or either of the ports the cable is plugged into can fail. Multiple logical connections can be made, but many of the [higher level protocols](#) were not designed to [fail over](#) completely seamlessly. Combining multiple physical connections into one logical connection using link aggregation provides more resilient communications.

Architecture

[Network architects](#) can implement aggregation at any of the lowest three layers of the [OSI model](#). Examples of aggregation at layer 1 ([physical layer](#)) include [power line](#) (e.g. [IEEE 1901](#)) and [wireless](#) (e.g. [IEEE 802.11](#)) network devices that combine multiple frequency bands. OSI layer 2 ([data link layer](#), e.g. [Ethernet frame](#) in LANs or [multi-link PPP](#) in WANs, Ethernet [MAC address](#)) aggregation typically occurs across switch ports, which can be either physical ports or virtual ones managed by an operating system. Aggregation at layer 3 ([network layer](#)) in the OSI model can use [round-robin scheduling](#), hash values computed from fields in the packet header, or a combination of these two methods.

Regardless of the layer on which aggregation occurs, it is possible to balance the network load across all links. However, in order to avoid [out-of-order delivery](#), not all implementations take advantage of this. Most methods provide [failover](#) as well.

Combining can either occur such that multiple interfaces share one logical address (i.e. IP) or one physical address (i.e. MAC address), or it allows each interface to have its own address. The

former requires that both ends of a link use the same aggregation method, but has performance advantages over the latter.

Channel bonding is differentiated from [load balancing](#) in that load balancing divides traffic between network interfaces on per [network socket](#) (layer 4) basis, while channel bonding implies a division of traffic between physical interfaces at a lower level, either per packet (layer 3) or a data link (layer 2) basis.

IEEE link aggregation

Standardization process

By the mid-1990s, most network switch manufacturers had included aggregation capability as a proprietary extension to increase bandwidth between their switches. Each manufacturer developed its own method, which led to compatibility problems. The [IEEE 802.3](#) working group took up a study group to create an interoperable [link layer](#) standard (i.e. encompassing the physical and data-link layers both) in a November 1997 meeting.^[4] The group quickly agreed to include an automatic configuration feature which would add in redundancy as well. This became known as [Link Aggregation Control Protocol](#) (LACP).

802.3ad

As of 2000, most gigabit channel-bonding schemes use the IEEE standard of Link Aggregation which was formerly clause 43 of the [IEEE 802.3](#) standard added in March 2000 by the IEEE 802.3ad task force.^[5] Nearly every network equipment manufacturer quickly adopted this joint standard over their proprietary standards.

802.1AX

The 802.3 maintenance task force report for the 9th revision project in November 2006 noted that certain 802.1 layers (such as [802.1X](#) security) were positioned in the [protocol stack](#) below Link Aggregation which was defined as an [802.3](#) sublayer.^[6] To resolve this discrepancy, the 802.3ax (802.1AX) task force was formed,^[7] resulting in the formal transfer of the protocol to the 802.1 group with the publication of IEEE 802.1AX-2008 on 3 November 2008.^[8]

Link Aggregation Control Protocol

Within the IEEE Ethernet standards, the Link Aggregation Control Protocol (LACP) provides a method to control the bundling of several physical links together to form a single logical link. LACP allows a network device to negotiate an automatic bundling of links by sending LACP packets to their peer, a directly connected device that also implements LACP.

LACP Features and practical examples

1. Maximum number of bundled ports allowed in the port channel: Valid values are usually from 1 to 8.
2. LACP packets are sent with multicast group MAC address `01:80:C2:00:00:02`
3. During LACP detection period
 - LACP packets are transmitted every second
 - Keep-alive mechanism for link member: (default: slow = 30s, fast=1s)
4. Selectable load-balancing mode is available in some implementations^[9]
5. LACP mode :
 - Active: Enables LACP unconditionally.
 - Passive: Enables LACP only when an LACP device is detected. (This is the default state)

Advantages over static configuration

- Failover occurs automatically: When a link has an intermediate failure, for example in a [media converter](#) between the devices, a peer system may not perceive any connectivity problems. With static link aggregation, the peer would continue sending traffic down the link causing the connection to fail.
- Dynamic configuration: The device can confirm that the configuration at the other end can handle link aggregation. With static link aggregation, a cabling or configuration mistake could go undetected and cause undesirable network behavior.^[10]

Practical notes

LACP works by sending frames (LACPDUs) down all links that have the protocol enabled. If it finds a device on the other end of a link that also has LACP enabled, that device will independently send frames along the same links in the opposite direction enabling the two units to detect multiple links between themselves and then combine them into a single logical link. LACP can be configured in one of two modes: active or passive. In active mode, LACPDUs are

sent 1 per second along the configured links. In passive mode, LACPDU's are not sent until one is received from the other side, a speak-when-spoken-to protocol.

Proprietary link aggregation

In addition to the IEEE link aggregation substandards, there are a number of proprietary aggregation schemes including Cisco's [EtherChannel](#) and [Port Aggregation Protocol](#), Juniper's Aggregated Ethernet, AVAYA's [Multi-Link Trunking](#), [Split Multi-Link Trunking](#), [Routed Split Multi-Link Trunking](#) and [Distributed Split Multi-Link Trunking](#), ZTE's Smartgroup, Huawei's Eth-Trunk, and [Connectify's Speedify](#).^[11] Most high-end network devices support some form of link aggregation. Software-based implementations – such as the *BSD *lagg* package, [Linux bonding](#) driver, [Solaris *dladm aggr*](#), etc. – exist for many operating systems.

Linux drivers

The Linux *bonding* driver^[12] provides a method for aggregating multiple [network interface controllers](#) (NICs) into a single logical bonded interface of two or more so-called (*NIC*) *slaves*. The majority of modern [Linux distributions](#) come with a [Linux kernel](#) which has the Linux bonding driver integrated as a [loadable kernel module](#) and the *ifenslave* (if = [network] interface) [user-level](#) control program pre-installed. [Donald Becker](#) programmed the original Linux bonding driver. It came into use with the [Beowulf cluster](#) patches for the [Linux](#) kernel 2.0.

Modes for the Linux bonding driver^[12] (network interface aggregation modes) are supplied as parameters to the kernel bonding module at load time. They may be given as command-line arguments to the *insmod* or *modprobe* commands, but are usually specified in a Linux distribution-specific configuration file. The behavior of the single logical bonded interface depends upon its specified bonding driver mode. The default parameter is `balance-rr`.

Round-robin (balance-rr)

Transmit alternate [network packets](#) in sequential order from the first available NIC slave through the last. This mode provides [load balancing](#) and [fault tolerance](#).^[13] This mode can cause congestion control issues due to the packet reordering it can introduce.^[14]

Active-backup (active-backup)

Only one NIC slave in the bond is active. A different slave becomes active if, and only if, the active slave fails. The single logical bonded interface's [MAC address](#) is externally visible on

only one [NIC](#) (port) to simplify forwarding in the [network switch](#). This mode provides fault tolerance.

XOR (balance-xor)

Transmit network packets based on a hash of the packet's source and destination. The default algorithm only considers MAC addresses (*layer2*). Newer versions allow selection of additional policies based on IP addresses (*layer2+3*) and TCP/UDP port numbers (*layer3+4*). This selects the same NIC slave for each destination MAC address, IP address, or IP address and port combination, respectively. Single connections will have guaranteed *in order* packet delivery and will transmit at the speed of a single NIC.^[15] This mode provides load balancing and fault tolerance.

Broadcast (broadcast)

Transmit network packets on all slave network interfaces. This mode provides fault tolerance.

IEEE 802.3ad Dynamic link aggregation (802.3ad, LACP)

Creates aggregation groups that share the same speed and duplex settings. Utilizes all slave network interfaces in the active aggregator group according to the 802.3ad specification. This mode is similar to the XOR mode above and supports the same balancing policies. The link is set up dynamically between two LACP-supporting peers.

Adaptive transmit load balancing (balance-tlb)

Linux bonding driver mode that does not require any special network-switch support. The outgoing network packet traffic is distributed according to the current load (computed relative to the speed) on each network interface slave. Incoming traffic is received by one currently designated slave network interface. If this receiving slave fails, another slave takes over the MAC address of the failed receiving slave.

Adaptive load balancing (balance-alb)

includes *balance-tlb* plus *receive load balancing* (rlb) for IPv4 traffic and does not require any special network switch support. The receive load balancing is achieved by [ARP](#) negotiation. The bonding driver intercepts the ARP Replies sent by the local system on their way out and overwrites the source hardware address with the unique hardware address of one of the NIC slaves in the single logical bonded interface such that different network-peers use different MAC addresses for their network packet traffic.

The Linux Team driver^[16] provides an alternative to bonding driver. The main difference is that Team driver kernel part contains only essential code and the rest of the code (link validation, LACP implementation, decision making, etc.) is run in [userspace](#) as a part of *teamd* daemon.

Usage

Network backbone

Link aggregation offers an inexpensive way to set up a high-speed [backbone network](#) that transfers much more data than any single port or device can deliver. Link aggregation also allows the network's backbone speed to grow incrementally as demand on the network increases, without having to replace everything and deploy new hardware.

Most backbone installations install more cabling or fiber optic pairs than is initially necessary. This is done because labor costs are higher than the cost of the cable, and running extra cable reduces future labor costs if networking needs change. Link aggregation can allow the use of these extra cables to increase backbone speeds for little or no extra cost if ports are available.

Order of frames

When balancing traffic, network administrators often wish to avoid reordering Ethernet frames. For example, TCP suffers additional overhead when dealing with out-of-order packets. This goal is approximated by sending all frames associated with a particular session across the same link. Common implementations use L2 or L3 hashes (i.e. based on the MAC or the IP addresses), ensuring that the same flow is always sent via the same physical link.^{[17][18][19]}

However, this may not provide even distribution across the links in the trunk when only a single or very few pairs of hosts communicate with each other, i.e. when the hashes provide too little variation. It effectively limits the client bandwidth in aggregate.^[18] In the extreme, one link is fully loaded while the others are completely idle and aggregate bandwidth is limited to this single member's maximum bandwidth. For this reason, an even load balancing and full utilization of all trunked links is almost never reached in real-life implementations.

Use on network interface cards

NICs trunked together can also provide network links beyond the throughput of any one single NIC. For example, this allows a central file server to establish an aggregate 2-gigabit connection using two 1-gigabit NICs teamed together. Note the data signaling rate will still be 1Gbit/s, which can be misleading depending on methodologies used to test throughput after link aggregation is employed.

Microsoft Windows

[Microsoft Windows](#) Server 2012 supports link aggregation natively. Previous Windows Server versions relied on manufacturer support of the feature within their [device driver](#) software. [Intel](#), for example, released Advanced Networking Services (ANS) to bond Intel Fast Ethernet and Gigabit cards.^[20]

[Nvidia](#) also supports "teaming" with their Nvidia Network Access Manager/Firewall Tool. [HP](#) also has a teaming tool for HP branded NICs which will allow for non-EtherChanneled NIC teaming or which will support several modes of EtherChannel (port aggregation) including 802.3ad with LACP. In addition, there is a basic layer-3 aggregation (available at least from Windows XP SP3),^[21] that allows servers with multiple IP interfaces on the same network to perform load balancing, and home users with more than one internet connection, to increase connection speed by sharing the load on all interfaces.^[22]

[Broadcom](#) offers advanced functions via Broadcom Advanced Control Suite (BACS), via which the teaming functionality of BASP ("Broadcom Advanced Server Program") is available, offering 802.3ad static LAGs, LACP, and "smart teaming" which doesn't require any configuration on the switches to work. It is possible to configure teaming with BACS with a mix of NICs from different vendors as long as at least one of them is Broadcom and the other NICs have the required capabilities to create teaming.^[23]

Linux and UNIX

[Linux](#), [FreeBSD](#), [NetBSD](#), [OpenBSD](#), [macOS](#), [OpenSolaris](#) and commercial Unix distributions such as [AIX](#) implement Ethernet bonding (trunking) at a higher level, and can hence deal with NICs from different manufacturers or drivers, as long as the NIC is supported by the kernel.^[12]

Virtualization platforms

[Citrix XenServer](#) and [VMware ESX](#) have native support for link-aggregation. XenServer offers both static LAGs as well as LACP. vSphere 5.1 (ESXi) supports both static LAGs and LACP natively with their virtual distributed switch.^[24]

For Microsoft's [Hyper-V](#), bonding or teaming isn't offered from the hyper-visor or OS-level, but the above-mentioned methods for teaming under Windows applies to Hyper-V as well.

Limitations

Single switch

With the modes *balance-rr*, *balance-xor*, *broadcast* and *802.3ad*, all physical ports in the link aggregation group must reside on the same logical switch, which, in most common scenarios, will leave a single point of failure when the physical switch to which all links are connected goes offline. The modes *active-backup*, *balance-tlb*, and *balance-alb* can also be set up with two or more switches. But after failover (like all other modes), in some cases, active sessions may fail (due to ARP problems) and have to be restarted.

However, almost all vendors have proprietary extensions that resolve some of this issue: they aggregate multiple physical switches into one logical switch. The [Split multi-link trunking \(SMLT\)](#) protocol allows multiple Ethernet links to be split across multiple switches in a stack, preventing any single point of failure and additionally allowing all switches to be load balanced across multiple aggregation switches from the single access stack. These devices synchronize state across an [Inter-Switch Trunk \(IST\)](#) such that they appear to the connecting (access) device to be a single device (switch block) and prevent any packet duplication. SMLT provides enhanced resiliency with sub-second failover and sub-second recovery for all speed trunks (10 Mbit/s, 100 Mbit/s, 1,000 Mbit/s, and 10 Gbit/s) while operating transparently to end-devices.

Same link speed

In most implementations, all the ports used in an aggregation consist of the same physical type, such as all copper ports (10/100/1000BASE-T), all multi-mode fiber ports, or all single-mode fiber ports. However, all the IEEE standard requires is that each link be full duplex and all of them have an identical speed (10, 100, 1,000 or 10,000 Mbit/s).

Many switches are PHY independent, meaning that a switch could have a mixture of copper, SX, LX, LX10 or other GBICs. While maintaining the same PHY is the usual approach, it is possible to aggregate a 1000BASE-SX fiber for one link and a 1000BASE-LX (longer, diverse path) for the second link, but the important thing is that the speed will be 1 Gbit/s full duplex for both links. One path may have a slightly longer propagation time but the standard has been engineered so this will not cause an issue.

Ethernet aggregation mismatch

Aggregation mismatch refers to not matching the aggregation type on both ends of the link. Some switches do not implement the 802.1AX standard but support static configuration of link

aggregation. Therefore, link aggregation between similarly statically configured switches will work but will fail between a statically configured switch and a device that is configured for LACP.

Examples

Ethernet

On [Ethernet](#) interfaces, channel bonding requires assistance from both the Ethernet [switch](#) and the host computer's [operating system](#), which must "[stripe](#)" the delivery of frames across the network interfaces in the same manner that I/O is striped across disks in a [RAID 0](#) array. For this reason, some discussions of channel bonding also refer to [Redundant Array of Inexpensive Nodes](#) (RAIN) or to "redundant array of independent network interfaces".^[25]

Modems

In analog modems, multiple [dial-up](#) links over [POTS](#) may be bonded. Throughput over such bonded connections can come closer to the aggregate bandwidth of the bonded links than can throughput under routing schemes which simply load-balance outgoing network connections over the links.

DSL

Similarly, multiple [DSL lines](#) can be bonded to give higher bandwidth; in the [United Kingdom](#), [ADSL](#) is sometimes [bonded](#) (<https://evolving.net.uk/bonded-adsl-guide/how-does-bonded-adsl-work/>) to give for example 512kbit/s upload bandwidth and 4 megabit/s download bandwidth, in areas that only have access to 2 megabit/s bandwidth.

DOCSIS

Under the DOCSIS 3.0^[26] and 3.1^[27] specifications for data over cable TV (CATV) systems, multiple channels may be bonded. Under DOCSIS 3.0, up to 32 downstream and 8 upstream channels may be bonded. These are typically 6 or 8 MHz wide. DOCSIS 3.1 defines more complicated arrangements involving aggregation at the level subcarriers and larger notional channels.

Wireless Broadband

Broadband bonding is a type of channel bonding that refers to aggregation of multiple channels at **OSI layers** at level four or above. Channels bonded can be wired links such as a **T-1** or **DSL line**. Additionally, it is possible to bond multiple **cellular links** for an aggregated wireless bonded link.

Previous bonding methodologies resided at lower OSI layers, requiring coordination with **telecommunications companies** for implementation. Broadband bonding, because it is implemented at higher layers, can be done without this coordination.^[28]

Commercial implementations of Broadband Channel Bonding include:

- Wistron AiEdge Corporation's U-Bonding Technology ^[29]
- Mushroom Networks' Broadband Bonding Service ^[30]
- Connectify's Speedify fast bonding VPN - software app for multiple platforms: PC, Mac, iOS and Android ^[31]
- Peplink's SpeedFusion Bonding Technology ^[32]
- Viprinet's Multichannel VPN Bonding Technology ^[33]
- Elsight's Multichannel Secure Data Link ^[34]
- Synopi's Natiplay Internet Bonding Technology ^[35]
- ComBOX Networks multi-wan bonding as a service ^[36]

Wi-Fi

- On **802.11** (Wi-Fi), channel bonding is used in **Super G** technology, referred to as 108Mbit/s. It bonds two channels of standard **802.11g**, which has 54Mbit/s **data signaling rate**.
- On **IEEE 802.11n**, a mode with a channel width of 40 MHz is specified. This is not channel bonding, but a single channel with double the older 20 MHz channel width, thus using two adjacent 20 MHz bands. This allows direct doubling of the PHY data rate from a single 20 MHz channel, but the MAC and user-level throughput also depends on other factors so may not double.

See also

-
- [Inverse multiplexing](#)

- MC-LAG
- Spanning Tree Protocol

References

1. Guijarro, Manuel; Ruben Gaspar; et al. (2008). "Experience and Lessons learnt from running High Availability Databases on Network Attached Storage" (http://www.iop.org/EJ/article/1742-6596/119/4/042015/jpconf8_119_042015.pdf) (PDF). *Journal of Physics: Conference Series*. Conference Series. IOP Publishing. **119** (4): 042015. Bibcode:2008JPhCS.119d2015G (<https://ui.adsabs.harvard.edu/abs/2008JPhCS.119d2015G>) . doi:10.1088/1742-6596/119/4/042015 (<https://doi.org/10.1088%2F1742-6596%2F119%2F4%2F042015>) . Retrieved 2009-08-17. "Network bonding (also known as port trunking) consists of aggregating multiple network interfaces into a single logical bonded interface that correspond to a single IP address."
2. "IEEE 802.3ad Link Bundling" (http://www.cisco.com/en/US/docs/ios/12_2sb/feature/guide/sbcelacp.html) . Cisco Systems. 2007-02-27. Archived (https://web.archive.org/web/20120419100517/http://www.cisco.com/en/US/docs/ios/12_2sb/feature/guide/sbcelacp.html) from the original on 2012-04-19. Retrieved 2012-03-15.
3. "Cisco Nexus 5000 Series NX-OS Software Configuration Guide - Configuring Port Channels [Cisco Nexus 5000 Series Switches]" (<https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus5000/sw/configuration/guide/cli/CLIConfigurationGuide/EtherChannel.html>) . Cisco. Retrieved 2019-10-25.
4. "IEEE 802 Trunking Tutorial" (http://grouper.ieee.org/groups/802/3/trunk_study/tutorial/index.html) . 1997-11-11. Archived (https://web.archive.org/web/20131207014147/http://grouper.ieee.org/groups/802/3/trunk_study/tutorial/index.html) from the original on 2013-12-07. Retrieved 2013-08-13.
5. "IEEE 802.3ad Link Aggregation Task Force" (<http://www.ieee802.org/3/ad/>) . www.ieee802.org. Archived (<https://web.archive.org/web/20171027055813/http://www.ieee802.org/3/ad/>) from the original on 27 October 2017. Retrieved 9 May 2018.
6. Law, David (2006-11-13). "IEEE 802.3 Maintenance" (http://www.ieee802.org/3/maint/public/maint_open_1106.pdf) (PDF). p. 9. Archived (https://web.archive.org/web/20081007130153/http://www.ieee802.org/3/maint/public/maint_open_1106.pdf) (PDF) from the original on 2008-10-07. Retrieved 2009-08-18. "Proposal to move Link Aggregation to IEEE 802.1 •It is an 802.3 sublayer but it has to go above IEEE Std 802.1x"
7. "IEEE 802.3ax (IEEE P802.1AX) Link Aggregation Project Authorization Request (approved)" (<https://web.archive.org/web/20101116135600/http://standards.ieee.org/board/nes/projects/802-1AX.pdf>) (PDF). 2007-03-22. p. 3. Archived from the original (<http://standards.ieee.org/board/nes/projects/802-1AX.pdf>) (PDF) on 2010-11-16. Retrieved 2018-01-10. "It has been concluded between 802.1 and 802.3 that future development of Link Aggregation would be more appropriate as an 802.1 standard"

8. "IEEE SA - 802.1AX-2008 - IEEE Standard for Local and metropolitan area networks--Link Aggregation" (<http://standards.ieee.org/findstds/standard/802.1AX-2008.html>) . Archived (<https://web.archive.org/web/20130807055327/http://standards.ieee.org/findstds/standard/802.1AX-2008.html>) from the original on 2013-08-07. Retrieved 2013-08-13.
9. "What Is LACP?" (https://support.huawei.com/enterprise/en/doc/EDOC1100086560#EN-US_TOPIC_0169439602) . Retrieved 5 March 2020.
10. "Link aggregation on Dell servers" (<https://web.archive.org/web/20120313024309/http://support.dell.com/support/edocs/network/LAG1855/LAGConsiderationv0.5.pdf>) (PDF). Archived from the original (<http://support.dell.com/support/edocs/network/LAG1855/LAGConsiderationv0.5.pdf>) (PDF) on 13 March 2012.
11. "Connectify commercializes Speedify channel bonding service - FierceWireless" (<http://www.fiercewireless.com/tech/story/connectify-commercializes-speedify-channel-bonding-service/2014-06-15>) . www.fiercewireless.com. Archived (<https://web.archive.org/web/20160628082352/http://www.fiercewireless.com/tech/story/connectify-commercializes-speedify-channel-bonding-service/2014-06-15>) from the original on 28 June 2016. Retrieved 9 May 2018.
12. The Linux Foundation: Bonding (<http://www.linuxfoundation.org/collaborate/workgroups/networking/bonding>) Archived (<https://web.archive.org/web/20101228000337/https://www.linuxfoundation.org/collaborate/workgroups/networking/bonding>) 2010-12-28 at the Wayback Machine
13. "Understanding NIC Bonding with Linux" (<https://www.enterprisenetworkingplanet.com/news/understanding-nic-bonding-with-linux/>) . 2 December 2009.
14. "Linux Ethernet Bonding Driver HOWTO" (<https://www.kernel.org/doc/Documentation/networking/bonding.txt>) . Retrieved 2021-11-27.
15. <https://www.kernel.org/doc/Documentation/networking/bonding.txt>
16. "libteam by jpirko" (<http://www.libteam.org>) . www.libteam.org. Archived (<https://web.archive.org/web/20170914032423/http://libteam.org/>) from the original on 14 September 2017. Retrieved 9 May 2018.
17. "IEEE 802.3ad Link Aggregation (LAG) what it is, and what it is not" (http://grouper.ieee.org/groups/802/3/hssg/public/apr07/frazier_01_0407.pdf) (PDF). Archived (https://web.archive.org/web/20071129153021/http://grouper.ieee.org/groups/802/3/hssg/public/apr07/frazier_01_0407.pdf) (PDF) from the original on 2007-11-29. Retrieved 2007-05-14.
18. Mechanisms for Optimizing Link Aggregation Group (LAG) and Equal-Cost Multipath (ECMP) Component Link Utilization in Networks (<https://datatracker.ietf.org/doc/html/rfc7424>) . doi:10.17487/RFC7424 (<https://doi.org/10.17487%2FRFC7424>) . RFC 7424 (<https://datatracker.ietf.org/doc/html/rfc7424>) .
19. "Outbound Traffic Distribution Across Trunked Links". Procurve 2910al Management and Configuration Guide. Hewlett Packard. February 2009.

20. *Intel Advanced Networking Services* (<http://www.intel.com/support/network/sb/cs-009747.htm>) Archived (<https://web.archive.org/web/20070124201010/http://www.intel.com/support/network/sb/CS-009747.htm>) 2007-01-24 at the *Wayback Machine*
21. *Archiveddocs. "RandomAdapter: Core Services"* ([https://technet.microsoft.com/pl-pl/library/cc784947\(v=ws.10\).aspx](https://technet.microsoft.com/pl-pl/library/cc784947(v=ws.10).aspx)) . *technet.microsoft.com*. Archived ([https://web.archive.org/web/20160425094804/http://technet.microsoft.com/pl-pl/library/cc784947\(v=ws.10\).aspx](https://web.archive.org/web/20160425094804/http://technet.microsoft.com/pl-pl/library/cc784947(v=ws.10).aspx)) from the original on 25 April 2016. Retrieved 9 May 2018.
22. *"Norton Internet Security™ - PC Protection"* (<http://www.pctools.com/guides/registry/detail/951/>) . *www.pctools.com*. Archived (<https://web.archive.org/web/20170401111624/http://www.pctools.com/guides/registry/detail/951/>) from the original on 1 April 2017. Retrieved 9 May 2018.
23. *Broadcom Windows Management Applications* (http://www.broadcom.com/support/ethernet_nic/management_applications.php) Archived (https://web.archive.org/web/20120801043221/http://www.broadcom.com/support/ethernet_nic/management_applications.php) 2012-08-01 at the *Wayback Machine*, visited 8 July 2012
24. *VMware What's New in vSphere 5.1 networking* (<http://www.vmware.com/files/pdf/techpaper/Whats-New-VMware-vSphere-51-Network-Technical-Whitepaper.pdf>) Archived (<https://web.archive.org/web/20140123083838/https://www.vmware.com/files/pdf/techpaper/Whats-New-VMware-vSphere-51-Network-Technical-Whitepaper.pdf>) 2014-01-23 at the *Wayback Machine*, June, 2012. Visited 17 January 2013
25. *Jielin Dong, ed. (2007). Network Dictionary* (https://books.google.com/books?id=On_Hh23IXDUC) . *ITPro collection*. Javvin Technologies Inc. p. 95. ISBN 9781602670006. Retrieved 2013-08-07. "Channel bonding, sometimes also called redundant array of independent network interfaces (RAIN), is an arrangement in which two or more network interfaces on a host computer are combined for redundancy or increased throughput."
26. *DOCSIS 3.0 Physical Interface Specification* (<https://apps.cablelabs.com/specification/CM-SP-PHYv3.0>)
27. *DOCSIS 3.1 Physical Interface Specification* (<https://apps.cablelabs.com/specification/CM-SP-PHYv3.1>)
28. *"Broadband bonding offers high-speed alternative"* (<https://web.archive.org/web/20120707084446/http://engineeringbook.net/articles/broadband-bonding-offers-high-speed-alternative/>) . *engineeringbook.net*. Archived from the original (<http://engineeringbook.net/articles/broadband-bonding-offers-high-speed-alternative/>) on 7 July 2012. Retrieved 5 April 2013.
29. *"Ubonding | Wistron"* (<https://wistronaiedge.com/ubonding>) .
30. *Mushroom Networks' Broadband Bonding Service* (<http://www.mushroomnetworks.com/product/Broadband-Bonding-Service>)
31. *Connectify's Speedify Service* (<http://speedify.com>)

32. *Peplink's SpeedFusion Bonding Technology* (<http://www.peplink.com/technology/speedfusion-bonding-technology/>)
33. *Viprinet's Multichannel VPN Bonding Technology* (<http://www.viprinet.com/en/technology/viprinet-principle>)
34. *Elsight Multichannel Secure Data Link* (<http://el-sight.com/product/multichannel/>)
35. *Synopi's Natiplly Internet Bonding Technology* (<http://www.synopi.com/natiplly/>)
36. *comBOX multi-wan services Broadband bonding* (<https://www.combox-networks.com>)

General

- "Chapter 5.4: Link Aggregation Control Protocol (LACP)". *IEEE Std 802.1AX-2008 IEEE Standard for Local and Metropolitan Area Networks – Link Aggregation* (<http://standards.ieee.org/getieee802/download/802.1AX-2008.pdf>) (PDF). IEEE Standards Association. 2008-11-03. p. 30. doi:10.1109/IEEESTD.2008.4668665 (<https://doi.org/10.1109%2FIEEESTD.2008.4668665>) . ISBN 978-0-7381-5794-8.
- Tech Tips - Bonding Modes (<http://techtips.blogfa.com/post-9.aspx>)

External links

- IEEE P802.3ad Link Aggregation Task Force (<http://grouper.ieee.org/groups/802/3/ad/index.html>)
- Mikrotik link Aggregation / Bonding Guide (<http://wiki.mikrotik.com/wiki/Manual:Interface/Bonding>)
- Configuring a Shared Ethernet Adapter (SEA) - IBM (https://web.archive.org/web/20190701201158/https://www.ibm.com/support/knowledgecenter/POWER5/iphc3_p5/iphc3_vios_configuring_sea.htm?origURL=api%2Fredirect%2Fpowersys%2Fv3r1m5%2Findex.jsp)
- Managing VLANs on mission-critical shared Ethernet adapters - IBM (<http://www.ibm.com/developerworks/aix/library/au-managevlans/index.html?ca=drs>)
- Network overview by Rami Rosen (section about bonding) (<https://web.archive.org/web/20121019131903/http://www.linuxfoundation.org/collaborate/workgroups/networking/networkoverview>)

Retrieved from

["https://en.wikipedia.org/w/index.php?title=Link_aggregation&oldid=1100957605"](https://en.wikipedia.org/w/index.php?title=Link_aggregation&oldid=1100957605)

Last edited 1 month ago by Kvng

WIKIPEDIA
