

*Help us improve Wikipedia by **supporting it financially**.*

Open Shortest Path First

From Wikipedia, the free encyclopedia

Open Shortest Path First (OSPF) is a dynamic routing protocol for use in Internet Protocol (IP) networks. Specifically, it is a link-state routing protocol and falls into the group of interior gateway protocols, operating within an autonomous system (AS). It is defined as OSPF Version 2 in RFC 2328 (1998) for IPv4^[1]. The updates for IPv6 are specified as OSPF Version 3 in RFC 5340 (2008)^[2].

OSPF is perhaps the most widely-used interior gateway protocol (IGP) in large enterprise networks; IS-IS, another link-state routing protocol, is more common in large service provider networks. The most widely-used *exterior gateway protocol (EGP)* is BGP.

Contents

- 1 Overview
- 2 Neighbor relationships
- 3 Area types
 - 3.1 Backbone area
 - 3.2 Stub area
 - 3.3 Totally stubby area
 - 3.4 Not-so-stubby area
- 4 Path preference
 - 4.1 Traffic engineering
 - 4.2 Other extensions
- 5 OSPF router types
 - 5.1 Area border router
 - 5.2 Autonomous system boundary router
 - 5.3 Internal router
 - 5.4 Backbone router
- 6 Designated router
 - 6.1 Backup designated router
- 7 OSPF Hello Packet
- 8 OSPF in broadcast multiple access topologies
- 9 OSPF in NBMA topologies
- 10 Applications
- 11 Implementations

The Internet Protocol Suite

Application Layer

BGP · DHCP · DNS · FTP · GTP · HTTP ·
 IMAP · IRC · Megaco · MGCP · NNTP ·
 NTP · POP · RIP · RPC · RTP · RTSP · SDP ·
 SIP · SMTP · SNMP · SOAP · SSH · Telnet ·
 TLS/SSL · XMPP · (more)

Transport Layer

TCP · UDP · DCCP · SCTP · RSVP · ECN ·
 (more)

Internet Layer

IP (IPv4, IPv6) · ICMP · ICMPv6 · IGMP ·
 IPsec · (more)

Link Layer

ARP · RARP · NDP · **OSPF** ·
 Tunnels (L2TP) · PPP · Media Access
 Control (Ethernet, MPLS, DSL, ISDN,
 FDDI) · Device Drivers · (more)

- 11.1 Test Equipment
- 12 RFC history
- 13 See also
- 14 References
- 15 Further reading
- 16 External links

Overview

OSPF builds routing tables based solely on the destination IP address found in IP packets. It was designed to support variable-length subnet masking (VLSM, CIDR). OSPF detects changes in the topology, such as link failures, very quickly and converges on a new loop-free routing structure within seconds. For this, each OSPF router collects link-state information to construct the entire network topology of so-called *areas* from which it computes the shortest path tree for each route using a method based on Dijkstra's algorithm. The link-state information is maintained on each router as a link-state database (LSDB) which is a tree-image of the network topology. Identical copies of the LSDB are periodically updated through flooding on all routers in each OSPF-aware area (region of the network included in an OSPF area type - see "Area types" below). By convention, area 0 represents the core or "backbone" region of an OSPF-enabled network, and other OSPF area numbers may be designated to serve other regions of an enterprise (large, business) network - however every additional OSPF area must have a direct or virtual connection to the backbone OSPF area. The backbone area has the identifier 0.0.0.0. Inter-area routing goes via the backbone.

Routers in the same broadcast domain or at each end of a point-to-point telecommunications link form *adjacencies* when they have detected each other. This detection occurs when a router "sees" itself in a hello packet. This is called a two way state and is the most basic relationship. The routers in an Ethernet or frame relay network select a *designated router* (DR) and a *backup designated router* (BDR) which act as a hub to reduce traffic between routers. OSPF uses both unicast and multicast to send "hello packets" and link state updates. Multicast addresses 224.0.0.5 (all SPF/link state routers, also known as AllSPFRouters) and 224.0.0.6 (all Designated Routers, AllDRouters) are reserved for OSPF (RFC 2328). In contrast to the *Routing Information Protocol* (RIP) or the *Border Gateway Protocol* (BGP), OSPF does not use TCP or UDP but uses IP directly, via IP protocol 89. OSPF handles its own error detection and correction, therefore negating the need for TCP or UDP functions.

The OSPF Protocol can operate securely between routers, optionally using a clear-text password or using MD5 to authenticate peers before forming adjacencies and before accepting link-state advertisements (LSA). A natural successor to the Routing Information Protocol (RIP), it was *classless*, or able to use Classless Inter-Domain Routing, from its inception. Multicast extensions to OSPF, the *Multicast Open Shortest Path First* (MOSPF) protocols, have been defined but these are not widely used at present.

Neighbor relationships

As a link state routing protocol, OSPF establishes and maintains neighbor relationships in order to exchange routing updates with other routers. The neighbor relationship table is called an adjacency database in OSPF. Provided that OSPF is configured correctly, OSPF forms neighbor relationships only

with the routers directly connected to it. The routers that it forms a neighbor relationship with must be in the same area as the interface with which it is using to form a neighbor relationship. An interface can only belong to a single area.

Area types

An OSPF network is divided into *areas* that are labeled with 32-bit area identifiers. The area identifiers are commonly, but not always, written in the dot-decimal notation of an IPv4 address. However, they are *not* IP addresses and may duplicate, without conflict, any IPv4 address. The area identifiers for IPv6 implementations of OSPF (OSPFv3) also use 32-bit identifiers written in the same notation. While most OSPF implementations will right-justify an area number written in other than dotted decimal format (e.g., area 1), it is wise to always use dotted-decimal formats. Most implementations expand area 1 to the area identifier 0.0.0.1, but some have been known to expand it as 1.0.0.0.

Areas are logical groupings of hosts and networks, including their routers having interfaces connected to any of the included networks. Each area maintains a separate link state database whose information may be summarized towards the rest of the network by the connecting router. Thus, the topology of an area is unknown outside of the area. This reduces the amount of routing traffic between parts of an autonomous system.

Several "special" area types are defined:

Backbone area

The backbone area (also known as *area 0.0.0.0*) forms the core of an OSPF network. All other areas are connected to it, and inter-area routing happens via routers connected to the backbone area and to their own associated areas. It is the logical and physical structure for the 'OSPF domain' and is attached to all nonzero areas in the OSPF domain. Note that in OSPF the term Autonomous System Border Router (ASBR) is historic, in the sense that many OSPF domains can coexist in the same Internet-visible autonomous system, RFC1996 (ASGuidelines 1996, p. 25) ^[3].

The backbone area is responsible for distributing routing information between nonbackbone areas. The backbone must be contiguous, but it does not need to be physically contiguous; backbone connectivity can be established and maintained through the configuration of virtual links.

All OSPF areas must connect to the backbone area. This connection, however, can be through a virtual link. For example, assume area 0.0.0.1 has a physical connection to area 0.0.0.0. Further assume that area 0.0.0.2 has no direct connection to the backbone, but this area does have a connection to area 0.0.0.1. Area 0.0.0.2 can use a virtual link through the *transit area* 0.0.0.1 to reach the backbone. To be a transit area, an area has to have the transit attribute, so it cannot be stubby in any way.

Stub area

A stub area is an area which does not receive *external* routes except the default route, but does receive *inter-area* routes. This kind of area is useful when, for example, all Internet access goes through autonomous system border routers (ASBRs) in Area 0.0.0.0, but there are on other non-backbone areas allows the routers in the stub area to pick the best route to another area.

Totally stubby area

A *totally stubby area* (TSA), which is a nonstandard but useful extension by Cisco ^[4], is similar to a stub area, however this area does not allow *summary* routes in addition to the *external* routes, that is, *inter-area* (IA) routes are not summarized into totally stubby areas. The only way for traffic to get routed outside of the area is a default route which is the only Type-3 LSA advertised into the area. When there is only one route out of the area, fewer routing decisions have to be made by the route processor, which lowers system resource utilization.

Occasionally, it is said that a TSA can have only one ABR. This is not true. If there are multiple ABRs, as might be required for high availability, routers interior to the TSA will send non-intra-area traffic to the ABR with the lowest intra-area metric (the "closest" ABR).

An area can simultaneously be not-so-stubby and totally stubby. This is done when the practical place to put an ASBR, as, for example, with a newly acquired subsidiary, is on the edge of a totally stubby area. In such a case, the ASBR does send externals into the totally stubby area, and they are available to OSPF speakers within that area. In Cisco's implementation, the external routes can be summarized before injecting them into the totally stubby area. In general, the ASBR should not advertise default into the TSA-NSSA, although this can work with extremely careful design and operation, for the limited special cases in which such an advertisement makes sense.

By declaring the totally stubby area as NSSA, no external routes from the backbone, except the default route, enter the area being discussed. The externals do reach area 0.0.0.0 via the TSA-NSSA, but no routes other than the default route enter the TSA-NSSA. Routers in the TSA-NSSA send all traffic to the ABR, except to routes advertised by ASBR.

Not-so-stubby area

A *not-so-stubby area* (NSSA) is a type of stub area that can import *autonomous system* (AS) external routes and send them to the backbone, but cannot receive AS external routes from the backbone or other areas. The NSSA is a non-proprietary extension of the existing stub area feature that allows the injection of external routes in a limited fashion into the stub area.

Cisco also implements a proprietary version of a NSSA called a *NSSA totally stubby area*. It takes on the attributes of a TSA, meaning that type 3 and type 4 summary routes are not flooded into this type of area. It is also possible to declare an area both totally stubby and not-so-stubby, which means that the area will receive only the default route from area 0.0.0.0, but can also contain an autonomous system border router (ASBR) that accepts external routing information and injects it into the local area, and from the local area into area 0.0.0.0.

Redistribution into an NSSA area creates special type of LSA known as TYPE 7, which can exist only in an NSSA area. An NSSA ASBR generates this LSA, and an NSSA ABR router translates it into type 5 LSA which gets propagated into the OSPF domain.

Path preference

OSPF uses *path cost* as its basic routing metric, which was defined by the standard not to equate to any standard value such as speed, so the network designer could pick a metric important to the design. In

practice, it is determined by the speed (bandwidth) of the interface addressing the given route, although that tends to need network-specific scaling factors now that links faster than 100 Mbit/s are common. Cisco uses a metric like $10^8/\text{bandwidth}$ (the base value, 10^8 by default, can be adjusted). So, a 100Mbit/s link will have a cost of 1, a 10Mbit/s a cost of 10 and so on. But for links faster than 100Mbit/s, the cost would be <1 .

Metrics, however, are only directly comparable when of the same type. There are four types of metrics, with the most preferred type listed in order below. An intra-area route is always preferred to an inter-area route regardless of metric, and so on for the other types.

1. Intra-area
2. Inter-area
3. External Type 1, which includes both the external path cost and the sum of internal path costs to the ASBR that advertises the route,
4. External Type 2, the value of which is solely that of the external path cost

Traffic engineering

OSPF-TE is an extension to OSPF extending the idea of route preference to include traffic engineering (RFC 3630, ^[5]). The Traffic Engineering extensions to OSPF add dynamic properties to the route calculation algorithm. The properties are:

- Maximum Reservable bandwidth
- Unreserved bandwidth
- Available bandwidth

These fields are distributed between network nodes via the TLV fields of an opaque LSA.

OSPF-TE is commonly used within MPLS and GMPLS networks, as a means to determine the topology over which MPLS paths can be established. MPLS then uses its own path setup and forwarding protocols, once it has the full IP routing map.

Other extensions

RFC3717^[6] documents work in optical routing for IP based on "constraint-based" extensions to OSPF and IS-IS.

OSPF router types

OSPF defines the following router types:

- Area border router (ABR)
- Autonomous system border router (ASBR)
- Internal router (IR)
- Backbone router (BR)

The router types are attributes of an OSPF process. A given physical router may have one or more OSPF processes. For example, a router that is connected to more than one area, and which receives routes from a

BGP process connected to another AS, is both an ABR and an ASBR.

Each router has an identifier, customarily written in the dotted decimal format (e.g.: 1.2.3.4) of an IP address. This ID must be configured to an OSPF instance in order to be considered the ID. If not explicitly configured, the highest logical IP address will assume the router ID role. The router ID does not have to be a part of any routable subnet in the network and often isn't to avoid confusion.

Note: Do not confuse router types with *designated router (DR)*, or *backup designated router (BDR)*, which is an attribute of a router interface, not the router itself.

Area border router

An ABR is a router that connects one or more OSPF areas to the main backbone network. It is considered a member of all areas it is connected to. An ABR keeps multiple copies of the link-state database in memory, one for each area to which that router is connected.

Autonomous system boundary router

An ASBR is a router that is connected to more than one AS and that exchanges routing information with routers in other ASs. ASBRs typically also run a non-IGP routing protocol (e.g., BGP), or use static routes, or both. An ASBR is used to distribute routes received from other ASs throughout its own AS.

Internal router

An IR is a router that only has OSPF neighbor relationships with interfaces in the same area.

Backbone router

Backbone Routers: These are routers that are part of the OSPF backbone. By definition, this includes all area border routers, since those routers pass routing information between areas. However, a backbone router may also be a router that connects only to other backbone (or area border) routers, and is therefore not part of any area (other than Area 0).

Note that: an area border router is always a backbone router, but a backbone router is not necessarily an area border router.

Designated router

A *designated router (DR)* is the router interface elected among all routers on a particular multiaccess network segment, generally assumed to be broadcast multiaccess. Special techniques, often vendor-dependent, may be needed to support the DR function on nonbroadcast multiaccess (NBMA) media. It is usually wise to configure the individual virtual circuits of a NBMA subnet as individual point-to-point lines; the techniques used are implementation-dependent.

Do not confuse the DR with an OSPF router type. A given physical router can have some interfaces that are designated (DR), others that are backup designated (BDR), and others that are non-designated. If no router is DR or BDR on a given subnet, the DR is first elected, and then a second election is held if there is more than one BDR. ^[7]The DR is elected based on the following default criteria:

- If the priority setting on a OSPF router is set to 0, that means it can NEVER become a DR or BDR (Backup Designated Router).
- When a DR fails and the BDR takes over, there is another election to see who becomes the replacement BDR.
- The router sending the Hello packets with the highest priority wins the election.
- If two or more routers tie with the highest priority setting, the router sending the Hello with the highest RID (Router ID) wins. NOTE: a RID is the highest logical (loopback) IP address configured on a router, if no logical/loopback IP address is set then the Router uses the highest IP address configured on its active interfaces. (e.g. 192.168.0.1 would be higher than 10.1.1.2).
- Usually the router with the second highest priority number becomes the BDR.
- The priority values range between 0 - 254, with a higher value increasing its chances of becoming DR or BDR.
- IF a HIGHER priority OSPF router comes online AFTER the election has taken place, it will not become DR or BDR until (at least) the DR and BDR fail.
- If the current DR 'goes down' the current BDR becomes the new DR and a new election takes place to find another BDR. If the new DR then 'goes down' and the original DR is now available, it then becomes DR again, but no change is made to the current BDR.

DR's exist for the purpose of reducing network traffic by providing a source for routing updates, the DR maintains a complete topology table of the network and sends the updates to the other routers via multicast. All routers in an area will form a slave/master relationship with the DR. They will form adjacencies with the DR and BDR only. Every time a router sends an update, it sends it to the DR and BDR on the multicast address 224.0.0.6. The DR will then send the update out to all other routers in the area, to the multicast address 224.0.0.5. This way all the routers do not have to constantly update each other, and can rather get all their updates from a single source. The use of multicasting further reduces the network load. DRs and BDRs are always setup/elected on Broadcast networks (Ethernet). DR's can also be elected on NBMA (Non-Broadcast Multi-Access) networks such as Frame Relay or ATM. DRs or BDRs are not elected on point-to-point links (such as a point-to-point WAN connection) because the two routers on either sides of the link must become fully adjacent and the bandwidth between them cannot be further optimized.

Backup designated router

A *backup designated router* (BDR) is a router that becomes the designated router if the current designated router has a problem or fails. The BDR is the OSPF router with second highest priority at the time of the last election.

OSPF Hello Packet

+	Bits 0–7	8–15	16–18	19–31
0	Version	Type	Packet Length	
32	Router ID			
64	Area ID			
96	Checksum		Authentication Type	
128	Authentication			

160	Authentication		
192	Network Mask		
224	Hello Interval	Options	Router Priority
256	Router Dead Interval		
288	Designated Router		
320	Backup Designated Router		
352	Neighbor ID		
384	...		

OSPF in broadcast multiple access topologies

Neighbor adjacency is formed dynamically using multicast hello packets to 224.0.0.5. A DR and BDR are elected normally, and function normally.

OSPF in NBMA topologies

As described in RFC 2328 (<ftp://ftp.rfc-editor.org/in-notes/rfc2328.txt>) , has defined the following two official modes for OSPF in NBMA topologies:

- nonbroadcast
- point-to-multipoint

Cisco has defined the following three additional modes for OSPF in NBMA topologies:

- point-to-multipoint nonbroadcast
- broadcast
- point-to-point

Applications

OSPF was the first widely deployed routing protocol that could converge a network in the low seconds, and guarantee loop-free paths. It has a great many features that allow the imposition of policies about the propagation of routes that it may be appropriate to keep local, for load sharing, and for selective route importing more than IS-IS. IS-IS, in contrast, can be tuned for lower overhead in a stable network, the sort more common in ISP than enterprise networks. There are some historical accidents that made IS-IS the preferred IGP for ISPs, but ISP's today may well choose to use the features of the now-efficient implementations of OSPF^[8], after first considering the pros and cons of ISIS in service provider environments^[9].

As mentioned, OSPF can provide better load-sharing on external links than other IGPs. When the default route to an ISP is injected into OSPF from multiple ASBRS as a Type I external route and the same external cost specified, other routers will go to the ASBR with the least path cost from its location. This

can be tuned further by adjusting the external cost.

In contrast, if the default route from different ISPs is injected with different external costs, as a Type II external route, the lower-cost default becomes the primary exit and the higher-cost becomes the backup only.

Implementations

- 6WINDGate (<http://www.6wind.com>) , commercial embedded open-source routing modules from 6WIND including OSPFv2 and OSPFv3
- Vyatta, a commercial open-source router / firewall.
- GNU Zebra, a GPL routing suite for Unix-like systems supporting OSPF
- Quagga, a fork of GNU Zebra for Unix-like systems
- OpenBSD, includes an OSPF implementation
- XORP, a routing suite implementing RFC2328 (OSPFv2) and RFC2740 (OSPFv3) for both IPv4 and IPv6
- BIRD (<http://bird.network.cz>) implements RFC2328 OSPF
- Gated project included an RFC1583 OSPF implementation (UMD OSPF by University of Maryland).
- Routing and Remote Access Service in Windows NT 4.0, Windows 2000, Windows XP and Windows Server 2003 implements OSPFv2.

Test Equipment

Mu Dynamics' Service Analyzer was credited with (<http://labs.mudynamics.com/advisories/MU-200610-01.txt>) discovering a 0-day vulnerability in an open-source OSPF implementation. Other test equipment is also in widespread use for conformance and load or stress testing, such as from Ixia Communications and Agilent Technologies.

RFC history

- 1989, October - First put forward as a proposed standard as RFC 1131.
- 1994, The OSPF NSSA Option, RFC 1587.
- 1994, March - Multicast extensions to OSPF proposed as RFC 1584.
- 1997, July - OSPF version 2, as proposed in RFC 2178
- 1998, April - OSPF version 2, updated in RFC 2328, Internet Standard 54.
- 1999, December - OSPFv3 - OSPF for IPv6, RFC 2740.
- 2003, January - The OSPF NSSA Option updated, RFC 3101
- 2005, October - Prioritized Treatment of Specific OSPF Version 2 Packets and Congestion Avoidance, RFC 4222
- 2006, December - OSPF Version 2 Management Information Base, RFC 4750
- 2007, May - OSPF Version 3 Management Information Base, draft state
- 2008, July - OSPF for IPv6, RFC 5340 (obsoletes RFC 2740)
- 2009, February - OSPF Multipoint Relay (MPR) Extension for Ad Hoc Networks, RFC 5449

See also

- Enhanced Interior Gateway Routing Protocol
- IS-IS
- Mesh networking
- Routing
- Routing Information Protocol
- Route analytics

References

1. ^ Moy, J. (April 1998). "OSPF Version 2". Internet Engineering Task Force. <http://www.ietf.org/rfc/rfc2328.txt>. Retrieved on 2007-09-28.
2. ^ Coltun, R.; D. Ferguson, J Moy, A. Lindem (July 2008). "OSPF for IPv6". Internet Engineering Task Force. <http://www.ietf.org/rfc/rfc5340.txt>. Retrieved on 2008-07-23.
3. ^ Hawkinson, J; T. Bates (March 1996). "Guidelines for creation, selection, and registration of an Autonomous System". Internet Engineering Task Force. <ftp://ftp.rfc-editor.org/in-notes/rfc1930.txt>. Retrieved on 2007-09-28.
4. ^ What Are OSPF Areas and Virtual Links? (<http://www.cisco.com/warp/public/104/8.html>) ,Cisco Document ID: 13703,December 2005
5. ^ Katz, D; D. Yeung (September 2003). [Traffic Engineering (TE) Extensions to OSPF Version 2 "<http://www.ietf.org/rfc/rfc3630.txt>"]. Internet Engineering Task Force. Traffic Engineering (TE) Extensions to OSPF Version 2. Retrieved on 2007-09-28.
6. ^ Rajagopalan, B; J. Luciani & D. Awduche (March 2004). [IP over Optical Networks: A Framework "<http://www.ietf.org/rfc/rfc3717.txt>"]. Internet Engineering Task Force. IP over Optical Networks: A Framework. Retrieved on 2007-09-28.
7. ^ RFC2328 page 75
8. ^ Berkowitz, Howard (1999), "OSPF Goodies for ISPs", *North American Network Operators Group NANOG 17*, Montreal, OSPFforISPs
9. ^ Katz, Dave (2000), "OSPF and IS-IS: A Comparative Anatomy", *North American Network Operators Group NANOG 19*, Albuquerque, OSPFvsISIS

Further reading

- Andrew Colton, *OSPF for Cisco Routers* (<http://www.amazon.com/OSPF-Cisco-Routers-CCIE-Library/dp/0972286217/>)
- Jeff Doyle, Jennifer Carroll, *Routing TCP/IP* (<http://www.ciscopress.com/bookstore/product.asp?isbn=1587052024>) , Volume 1, 2nd Edition
- John T. Moy, *OSPF: Anatomy of an Internet Routing Protocol*.
- William R. Parkhurst, *Cisco OSPF Command and Configuration Handbook*. ISBN 978-1-58705-071-8
- "Configuring OSPF Authentication". <http://www.netcordia.com/resources/tech-tip-25.asp>. Retrieved on 2009-03-19.

External links

- IETF OSPF Working Group (<http://www.ietf.org/html.charters/ospf-charter.html>)
- OSPF Basics (<http://www.setup32.com/network-administration/networking/know-ospf.php>)
- Design and Implementation of OpenOSPF (Paper) (<http://www.networx.ch>)

- /OpenOSPF%20-%20Paper.pdf)
- Design and Implementation of OpenOSPF (Presentation) (<http://www.networx.ch/OpenOSPF%20-%20Presentation.pdf>)
- Cisco OSPF (http://www.cisco.com/en/US/tech/tk365/tk480/tsd_technology_support_sub-protocol_home.html)
- Cisco OSPF Areas and Virtual Links (http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a0080094aaa.shtml)
- OSPF Tutorial (<http://www.geocities.com/Heartland/4394/work/ospf.html>)
- Summary of OSPF v2 (<http://www.freesoft.org/CIE/Topics/89.htm>)
- IS-IS and OSPF difference discussion (<http://www.join.uni-muenster.de/Dokumente/drafts/draft-bhatia-manral-diff-isis-ospf-01.txt>) (Vishwas Manral, Manav bhatia and Yasuhiro Ohara)
- Data Communication Lectures of Manfred Lindner - Part OSPF Fundamentals (http://www.ict.tuwien.ac.at/skripten/datenkomm/infobase/L41-OSPF_Fundamentals_v4-5.pdf)
- Data Communication Lectures of Manfred Lindner - Part OSPF Areas (http://www.ict.tuwien.ac.at/skripten/datenkomm/infobase/L42-OSPF_Advanced_v4-4.pdf)
- TCP/IP Guide - OSPF (http://www.tcpipguide.com/free/t_OpenShortestPathFirstOSPF.htm)

Retrieved from "http://en.wikipedia.org/wiki/Open_Shortest_Path_First"

Categories: [Internet protocols](#) | [Internet standards](#) | [Routing protocols](#) | [Link protocols](#)

- This page was last modified on 21 April 2009, at 17:23 (UTC).
 - All text is available under the terms of the GNU Free Documentation License. (See **Copyrights** for details.)
- Wikipedia® is a registered trademark of the Wikimedia Foundation, Inc., a U.S. registered 501(c)(3) tax-deductible nonprofit charity.