




SonicOS 7.0

# Rules and Policies for Policy Mode

Administration Guide

SONICWALL<sup>®</sup>

# Contents

- Routing ..... 3**
  - About Routing ..... 3
    - About Metrics and Administrative Distance ..... 4
    - Route Advertisement ..... 6
    - ECMP Routing ..... 6
    - Policy-based Routing ..... 6
    - Policy-based TOS Routing ..... 7
    - PBR Metric-based Priority ..... 8
    - Policy-based Routing and IPv6 ..... 8
    - OSPF and RIP Advanced Routing Services ..... 9
    - Drop Tunnel Interface ..... 18
    - App-based Routing ..... 18
  - Rules and Policies > ..... 19
    - Configuring ..... 19
- SonicWall Support ..... 23**
- About This Document ..... 24

# Routing

For SD-WAN routing and route policies, see *Configuring SD-WAN Route Policies*.

## Topics:

- [About Routing](#)
  - [About Metrics and Administrative Distance](#)
  - [Route Advertisement](#)
  - [ECMP Routing](#)
  - [Policy-based Routing](#)
  - [Policy-based TOS Routing](#)
  - [PBR Metric-based Priority](#)
  - [Policy-based Routing and IPv6](#)
  - [OSPF and RIP Advanced Routing Services](#)
  - [Drop Tunnel Interface](#)
  - [App-based Routing](#)
- [Rules and Policies > Route Policy](#)

## About Routing

SonicWall Security Appliances support the following routing protocols:

- [RIPv1 \(Routing Information Protocol\)](#)
- [RIPv2](#)
- [OSPFv2 \(Open Shortest Path First\)](#)
- [OSPFv3](#)
- [PBR \(Policy-Based Routing\)](#)

## Topics:

- [About Metrics and Administrative Distance](#)
- [Route Advertisement](#)
- [ECMP Routing](#)
- [Policy-based Routing](#)
- [Policy-based TOS Routing](#)
- [PBR Metric-based Priority](#)
- [Policy-based Routing and IPv6](#)
- [OSPF and RIP Advanced Routing Services](#)
- [Drop Tunnel Interface](#)
- [App-based Routing](#)

# About Metrics and Administrative Distance

Metrics and administrative distance affect network performance, reliability, and circuit selection.

## Topics:

- [About Metrics](#)
- [About Administrative Distance](#)

## About Metrics

A *metric* is a weighted cost assigned to static and dynamic routes. Metrics determine the best route among several, usually the gateway with the lowest metric. This gateway is usually the default gateway.

Metrics have a value between 1 and 254; see [Metric Value Descriptions](#). Lower metrics are considered better and take precedence over higher costs. SonicOS adheres to Cisco-defined metric values for directly connected interfaces, statically encoded routes, and all dynamic IP routing protocols.

### METRIC VALUE DESCRIPTIONS

| Metric Value | Description   |
|--------------|---------------|
| 1            | Static Route  |
| 5            | EIGRP Summary |
| 20           | External BGP  |
| 90           | EIGRP         |
| 100          | IGRP          |
| 110          | OSPF          |
| 115          | IS-IS         |

| Metric Value | Description    |
|--------------|----------------|
| 120          | RIP            |
| 140          | EGP            |
| 170          | External EIGRP |
| 200          | Internal BGP   |

## About Administrative Distance

*Administrative distance* (admin distance) is a value that influences which source of routes should be used for two identical routes from different sources. The lower the administrative distance value, the more trusted the route.

The admin distance, when set, is only used by the ZebOS components when choosing which routes to:

- Populate into PBR
- Redistribute to other routing protocols when a static route competes with a route received from a particular routing protocol.

The admin distance is not used for prioritizing routes within PBR itself, so unless dynamic routing is in use, the admin distance set for a static route has no effect. When dynamic routing is being used, the admin distance provides a mechanism by which static routes defined in PBR can be compared to otherwise equivalent dynamic routes possibly received from protocols such as OSPF, RIP, or BGP. By default, the admin distance of a PBR static route inserted into the network services module (NSM) is equal to the metric defined for the PBR route. The admin distance of each static route may optionally be set to a different value when a custom value is entered for Admin Distance.

For example, if a simple (destination only) static route (for example, destination = 14.1.1.0/24) is defined with a metric of 10 and the admin distance set to its default of Auto, that route is populated into NSM with an admin distance and metric of 10.

Now assume the same 14.1.1.0/24 route is received from both RIP and OSPF. RIP routes have a default admin distance of 120 and OSPF routes 110, so the static route, with a default admin distance (== the metric) of 10 would be preferred over both routes, and NSM would not populate either the OSPF or RIP route into PBR. If the admin distance of the static route had been set to 115 (keeping the metric at 10), however, then the OSPF route (at 110) would be preferred over the static route, but the RIP route would not. If the OSPF route disappeared, NSM would withdraw the OSPF route and would not populate the RIP route as its 120 AD is greater than the static route's 115 AD.

In either of the above cases, the static route is still preferred in PBR because all non-default routes populated into PBR from NSM are added with a 110 metric, which is greater than the metric of 10 for the static route.

If an admin distance of 110 and a metric > 110 are used for the static routes, the metric value passed to NSM would be used by OSPF when it compares the metric of the static route to the OSPF metric (or cost) of any competing OSPF route.

# Route Advertisement

SonicWall Security Appliances use RIPv1 or RIPv2 to advertise its static and dynamic routes to other routers on the network. Changes in the status of VPN tunnels between the Security Appliance and remote VPN gateways are also reflected in the RIPv2 advertisements. Based on your router's capabilities or configuration, choose between:

- RIPv1, which is an earlier version of the protocol, has fewer features, and sends packets through broadcast instead of multicast.
- RIPv2, which is a later version of the protocol, includes subnet information when multicasting the routing table to adjacent routers and route tags for learning routes. RIPv2 packets are backwards compatible and can be accepted by some RIPv1 implementations that provide an option of listening for multicast packets. The RIPv2 Enabled (broadcast) selection, which broadcasts packets instead of multicasting them, is for heterogeneous networks with a mixture of RIPv1 and RIPv2 routers.

## ECMP Routing

SonicOS supports equal-cost multi-path (ECMP) routing, a technique for routing packets along multiple paths of equal cost. The forwarding engine identifies paths by next-hop. When forwarding a packet, the router must decide which next-hop (path) to use. Multi-path routing can be used in conjunction with most routing protocols.

In SonicOS, you can use ECMP routing to specify multiple next hops for a given route's destination. In environments with substantial requirements, there are several reasons for doing this. A router could just use one ISP most of the time, and switch to the other when the first one fails for some reason. Another application of multi-path is to keep a path on standby and enable it only when bandwidth requirements surpass a predefined threshold. SonicOS supports up to four next-hop paths.

Various routing protocols, including Open Shortest Path First (OSPF) and Intermediate System to Intermediate System (ISIS), explicitly allow ECMP routing. Some router implementations also allow equal-cost multi-path usage with RIP and other routing protocols.

## Policy-based Routing

A simple static routing entry specifies how to handle traffic that matches specific criteria, such as destination address, destination mask, gateway to forward traffic, the interface that gateway is located, and the route metric. This method of static routing satisfies most static requirements, but is limited to forwarding based only on destination addressing.

Policy-based Routing (PBR) allows you to create extended static routes to provide more flexible and granular traffic handling capabilities. SonicOS PBR allows for matching based upon source address, source netmask, destination address, destination netmask, service, interface, and metric. This method of routing allows for full control of forwarding based upon a large number of user defined variables.

PBR supports Fully Qualified Domain Name (FQDN). A FQDN cannot be used as the source or destination of the PBR entry, and the PBR entry can be redistributed to advanced routing protocols.

# Policy-based TOS Routing

SonicOS supports policy-based TOS (type of service) routing when defining policy-based routing (PBR) policies by Type of Service (TOS) and TOS mask values. When defined, the TOS and mask values are compared against the associated IP packet's TOS/DSCP field in the IP header when finding a route match.

The TOS value is compared to an 8-bit field in the IP packet header (for information about this header, see RFC 2474, Differentiated Services, and RFC 2168, Explicit Congestion Notification). The TOS value can be used to define services relating to quantitative performance requirements (for example, peak bandwidth) and those based on relative performance (for example, class differentiation).

TOS routing differs from existing SonicOS QoS marking, which does not affect the routing of a packet and cannot forward packets differently based on an inbound packet's TOS field. TOS Routing provides this capability by allowing policy routes to define a TOS Value/TOS Mask pair to be compared to inbound packets for differential forwarding. TOS routing only applies to packets as they enter the Security Appliance.

With TOS routing, it is possible to define multiple policy routes with identical source IP, destination IP, and service values, but differing TOS/TOS mask values. This allows packets with marked TOS fields to be forwarded differently based on the value of the TOS field in the inbound packet.

Any PBR policy routes defined before SonicOS have no values defined for the TOS/TOS mask. Likewise, the default values for TOS/TOS mask fields are zero (no values defined).

Policy routes with a TOS value other than zero are prioritized before all simple destination-only routes, but below any policy routes that define a source or service. When comparing two TOS Policy routes, and assuming both have the same set of source, destination, and service values either defined or not defined, the TOS route with the greater number of TOS mask bits set to 1 is prioritized before TOS routes with fewer TOS mask bits set.

The general prioritization (high to low) of PBR routes is as follows, based on the policy fields defined as anything other than **Any** or zero for TOS:

- Destination, Source, Service, TOS
- Destination, Source, Service
- Destination, Source, TOS
- Destination, Source
- Destination, Service, TOS
- Destination, Service
- Destination, TOS
- Destination
- Source, Service, TOS
- Source, Service
- Source, TOS
- Source
- Service, TOS
- Service
- TOS

# PBR Metric-based Priority

SonicOS supports a metric weighted cost assigned to a route policy for policy-based routing (PBR) that allows the configured metric to take precedence in route prioritization over the route specificity that used by default. Metrics have a value between 0 and 255. Lower metrics are considered better and take precedence over higher ones.

The general prioritization (high to low) of PBR routes is as follows, based on the policy fields defined as anything other than **Any**, or zero for TOS:

- Destination, Source, Service, TOS
- Destination, Source, Service
- Destination, Source, TOS
- Destination, Source
- Destination, Service, TOS
- Destination, Service
- Destination, TOS
- Destination
- Source, Service, TOS
- Source, Service
- Source, TOS
- Source
- Service, TOS
- Service
- TOS

Within these 15 classifications, routes are further prioritized based on the cumulative specificity of the defined route entries. For the source and destination fields, specificity is measured by counting the number of IP addresses represented in the address object. For example, the network address object, `10.0.0.0/24`, would include 256 IP addresses, while the network address object, `10.0.0.0/20`, would represent 4096. The longer /24 (24 bit) network prefix represents fewer host IP addresses and is more specific.

The new metric-weighted option allows the configured metric to take precedence in prioritization over the route specificity. With the option enabled, the precedence used during prioritization is as follows (high to low):

1. Route class (determined by the combination of source, destination, service, and TOS fields with values other than Any or zero)
2. The value of the Metric
3. The cumulative specificity of the source, destination, service, and TOS fields

## Policy-based Routing and IPv6

For complete information on the SonicOS implementation of IPv6, see *IPv6*.

Policy-based Routing is fully supported for IPv6 by selecting IPv6 address objects and gateways for route policies on **POLICY | Rules and Policies** > . You can switch the entries in the table between IPv4 and IPv6.



Routing Information Protocol next generation (RIPng) is an information routing protocol for IPv6 that allows routers to exchange information for computing routes through an IPv6-based network.

For information on route advertisement or for information on setting up Route Policies, see [Route Advertisement](#).

## OSPF and RIP Advanced Routing Services

In addition to Policy-based Routing and RIP advertising, SonicOS offers the option of enabling Advanced Routing Services (ARS). Advanced Routing Services provides full advertising and listening support for the Routing Information Protocol (RIPv1 - RFC1058) and (RIPv2 - RFC2453), and Open Shortest Path First (OSPFv2 – RFC2328). Advanced Routing Service should only be enabled by those environments requiring support for either or both of these dynamic routing protocols.

RIP and OSPF are Interior Gateway Protocols (IGP) that are both widely used by networks of various sizes to automate the process of route distribution. RIP is commonly used within smaller networks, while OSPF is used by larger networks, although network size should not be the only factor used to determine the appropriateness of one protocol over the other – network speed, interoperability requirements, and relative overall complexity, for example, should also be considered. RIPv1 and RIPv2 are both supported by ARS, the largest differences between the two being that RIPv2 supports VLSM (Variable Length Subnet Masks), authentication, and routing updates. [Routing Information Protocol Differences](#) illustrates the major differences between RIPv1, RIPv2, and OSPFv2/OSPFv3:

### ROUTING INFORMATION PROTOCOL DIFFERENCES

|                            | RIPv1                                                 | RIPv2                                                              | OSPFv2/OSPFv3                                                               |
|----------------------------|-------------------------------------------------------|--------------------------------------------------------------------|-----------------------------------------------------------------------------|
| Protocol metrics           | Distance Vector                                       | Distance Vector                                                    | Link State                                                                  |
| Maximum Hops               | 15                                                    | 15                                                                 | Unlimited                                                                   |
| Routing table updates      | Full table broadcast periodically, slower convergence | Full table broadcast or multicast periodically, slower convergence | Link state advertisement multicasts, triggered by changes, fast convergence |
| Subnet Sizes Supported     | Only class-based (a/b/c) subnets support              | Class-based only                                                   | VLSM                                                                        |
| Autonomous system topology | Indivisible and flat                                  | Indivisible and flat                                               | Area-based, allowing for segmentation and aggregation                       |

#### Topics:

- [About Routing Services](#)
- [OSPF Terms](#)

# About Routing Services


## Topics:

- Protocol Type
- Maximum Hops
- Split-Horizon
- Poison Reverse
- Routing Table Updates
- Subnet Sizes Supported
- Autonomous System Topologies

## Protocol Type

Distance Vector protocols such as RIP base routing metrics exclusively on hop counts, while Link state protocols such as OSPF consider the state of the link when determining metrics. For example, OSPF determines interface metrics by dividing its reference bandwidth (100mbits by default) by the interface speed – the faster the link, the lower the cost and the more preferable the path. Consider the example network shown in Example network for determining lowest cost route:

### EXAMPLE NETWORK FOR DETERMINING LOWEST COST ROUTE



In the sample network shown in [Example Network for Determining Lowest Cost Route](#), if Host A wanted to reach Host B, with RIP, the lowest cost route would be from Router A to Router B, across the relatively slow 64kbps link. With OSPF, the cost from Router A to Router B would be 1562, while the cost from Router A to Router C to Router D to Router B would be 364, making it the preferred route.

## Maximum Hops

RIP imposes a hop count of 15 to help prevent routing loops which can occur when bad (for example, stale) routing information is broadcast and propagated through a network either due to misconfiguration, or slow convergence. Consider if the link between Router D and Router E failed in the example in [Maximum Hops](#), and there were no safeguards in place:

- Router A's routing information states that it can reach Network E through Router B or Router C with a metric of 3.
- When the link between Router D and Router E fail, and Router A broadcasts its routing information, Router B and Router C determine that they can reach Network E through Router A with a metric of 4.
- Router B and Router C broadcast this information, and it is received by Router D which then determines it can reach Network E through Router B or Router C with a metric of 5.
- This loop continues until the hop count of 16 (infinity) is reached.

Other measures against this sort of situation are also commonly employed by RIP, including:

- [Split-Horizon Routing Table Updates](#)
- [Poison Reverse](#)
- [Routing Table Updates](#)
- [Subnet Sizes Supported](#)
- [Autonomous System Topologies](#)

## Split-Horizon

A preventative mechanism where routing information learned through an interface is not sent back out the same interface. This generally works well on broadcast links, but not on non-broadcast links such as Frame Relay, where a single link can commonly be used to reach two separate autonomous systems.

## Poison Reverse

Also known as route poisoning, an extension of split-horizon where a network is advertised with a metric of 16 (unreachable), helping to ensure that incorrect alternative routes are not propagated.

OSPF does not have to impose a hop count limit because it does not advertise entire routing tables, rather it generally only sends link state updates when changes occur. This is a significant advantage in larger networks in that it converges more quickly, produces less update traffic, and supports an unlimited number of hops.

## Routing Table Updates

As mentioned previously, the practice of sending an entire routing table introduces the problems of slower convergences, higher bandwidth utilization, and increased potential for stale routing information. RIPv1 broadcasts its entire routing table at a prescribed interval (usually every 30 seconds), RIPv2 can either broadcast or multicast, and OSPF multicasts only link state updates whenever a change to the network fabric occurs. OSPF has a further advantage of using designated routers (DR) in forming adjacencies in multiple-access networks (more on these concepts later) so that updates do not have to be sent to the entire network.

## Subnet Sizes Supported

RIPv1 was first implemented when networks were strictly class A, class B, and class C (and later D and E):

|                |                                                                                                                |
|----------------|----------------------------------------------------------------------------------------------------------------|
| <b>Class A</b> | 1.0.0.0 to 126.0.0.0 (0.0.0.0 and 127.0.0.0 are reserved)                                                      |
|                | <ul style="list-style-type: none"><li>Left most bit 0; 7 network bits; 24 host bits</li></ul>                  |
|                | <ul style="list-style-type: none"><li>0nnnnnnnn hhhhhhhh hhhhhhhh hhhhhhhh (8-bit classful netmask)</li></ul>  |
|                | <ul style="list-style-type: none"><li>126 Class A networks, 16,777,214 hosts each</li></ul>                    |
| <b>Class B</b> | 128.0.0.0 to 191.255.0.0                                                                                       |
|                | <ul style="list-style-type: none"><li>Left most bits 10; 14 network bits; 16 host bits</li></ul>               |
|                | <ul style="list-style-type: none"><li>10nnnnnnn nnnnnnnn hhhhhhhh hhhhhhhh (16-bit classful netmask)</li></ul> |
|                | <ul style="list-style-type: none"><li>16,384 Class B networks, 65,532 hosts each</li></ul>                     |
| <b>Class C</b> | 192.0.0.0 to 223.255.255.0                                                                                     |
|                | <ul style="list-style-type: none"><li>Left most bits 110; 21 network bits; 8 host bits</li></ul>               |
|                | <ul style="list-style-type: none"><li>110nnnnn nnnnnnnn nnnnnnnn hhhhhhhh (24-bit classful netmask)</li></ul>  |
|                | <ul style="list-style-type: none"><li>2,097,152 Class Cs networks, 254 hosts each</li></ul>                    |
| <b>Class D</b> | 225.0.0.0 to 239.255.255.255 (multicast)                                                                       |
|                | <ul style="list-style-type: none"><li>Left most bits 1110; 28 multicast address bits</li></ul>                 |
|                | <ul style="list-style-type: none"><li>1110nnnn mmmmmmm mmmmmmm mmmmmmm</li></ul>                               |
| <b>Class E</b> | 240.0.0.0 to 255.255.255.255 (reserved)                                                                        |
|                | <ul style="list-style-type: none"><li>Left most bits 1111; 28 reserved address bits</li></ul>                  |
|                | <ul style="list-style-type: none"><li>1111rrrr rrrrrrrr rrrrrrrr rrrrrrrr</li></ul>                            |

This method of address allocation proved to be very inefficient because it provided no flexibility, neither in the way of segmentation (subnetting) or aggregation (supernetting, or CIDR – classless inter-domain routing) by means of VLSM – variable length subnet masks.

VLSM, supported by RIPv2 and OSPF, allows for classless representation of networks to break larger networks into smaller networks:

For example, take the classful `10.0.0.0/8` network, and assign it a `/24` netmask. This subnetting allocates an additional 16-bits from the host range to the network range ( $24-8=16$ ). To calculate the number of additional networks this subnetting provides, raise 2 to the number of additional bits:  $2^{16}=65,536$ . Thus, rather than having a single network with 16.7 million hosts (usually more than most LAN's require) it is possible to have 65,536 networks, each with 254 usable hosts.

VLSM also allows for route aggregation (CIDR):

For example, if you had 8 class C networks: `192.168.0.0/24` through `192.168.7.0/24`, rather than having to have a separate route statement to each of them, it would be possible to provide a single route to `192.168.0.0/21` which would encompass them all.

This ability, in addition to providing more efficient and flexible allocation of IP address space, also allows routing tables and routing updates to be kept smaller.

## Autonomous System Topologies

An autonomous system (AS) is a collection of routers that are under common administrative control and that share the same routing characteristics. When a group of autonomous systems share routing information, they are commonly referred to as a confederation of autonomous systems. (RFC1930 and RFC975 address these concepts in much greater detail). In simple terms, an AS is a logical distinction that encompasses physical network elements based on the commonness of their configurations.

With regard to RIP and OSPF, RIP autonomous systems cannot be segmented, and all routing information must be advertised (broadcast) through the entire AS. This can become difficult to manage and can result in excessive routing information traffic. OSPF, on the other hand, employs the concept of Areas, and allows for logically, manageable segmentation to control the sharing of information within an AS. An Area ID is an administrative identifier. OSPF areas begin with the backbone area (area 0 or `0.0.0.0`), and all other areas must connect to this backbone area (although there are exceptions). This ability to segment the routing AS helps to ensure that it never becomes too large to manage, or too computationally intensive for the routers to handle.

## OSPF Terms

OSPF is substantially more complicated to configure and maintain than RIP. The following concepts are critical to understanding an OSPF routing environment:

- **Link state** – As it pertains to OSPF, a link is an egress interface on a router, and the state describes characteristics of that interface, such as its cost. Link states are sent in the form of Link State Advertisements (*LSA*) which are contained within Link State Update (*LSU*) packets, one of five types of OSPF packets.
- **Cost** – A quantification of the overhead required to send a packet along a particular link. Cost is calculated by dividing a reference bandwidth (usually 100mbit, or  $10^8$  bit) by an interface's speed. The lower the cost, the more preferable the link. Some common path costs are shown in [Cost Calculation for Different Interfaces](#).


## COST CALCULATION FOR DIFFERENT INTERFACES

| Interface      | Divided by $10^8$ (100mbit) = OSPF Cost |
|----------------|-----------------------------------------|
| Fast Ethernet  | 1                                       |
| Ethernet       | 10                                      |
| T1 (1.544mbit) | 64                                      |
| DSL (1mbit)    | 100                                     |
| DSL (512kbps)  | 200                                     |
| 64kbps         | 1562                                    |
| 56kbps         | 1785                                    |

- **Area** – The network comprising the group of OSPF routers intended to share a common Link State Database. OSPF networks are built around the backbone area (area 0, or 0.0.0.0) and all other areas must connect to the backbone area (unless virtual links are used, which is generally discouraged). Area assignment is interface specific on an OSPF router; in other words, a router with multiple interfaces can have those interfaces configured for the same or different areas.
- **Neighbors** – OSPF routers on a common network segment have the potential to become neighbors by means of sending Hello packets. Hello packets act as a form of advertisement and identification, and if two OSPF routers share a common set of certain characteristics, they become neighbors upon seeing their own router ID in the other router's Hello packet. Hello packets are also used in the DR (Designated Router) and BDR (Backup Designated Router) election process. For two routers to become neighbors, the characteristics that they must have in common are:
  - **Area-ID** – An area ID identifies an OSPF area with a 32-bit value, and is generally represented in an IP address format. OSPF requires at a minimum the backbone area, area 0 (or 0.0.0.0) for operation.
  - **Authentication** – Authentication types can generally be set to none, simple text, or MD5. When using simple text, authentication should be used only for identification, as it is sent in the clear. For security, MD5 should be used.
  - **Timer intervals** – Hello and Dead intervals must be the same. The Hello interval specifies the number of seconds between Hello packets (as a Keepalive function), and the Dead interval specifies the number of seconds after which a router is considered unavailable if a Hello is not received.
  - **Stub area flag** – A Stub area is an area that only requires a single point of egress, and therefore does not require a full list of external link advertisements. The stub area flag on two potential neighbors must be the same to avoid inappropriate link state exchanges. Another factor that affects neighboring is the kind of network. OSPF recognizes three network types:
    - **Broadcast** – For example, Ethernet. In broadcast networks, neighboring can be established with all other routers in the broadcast domain.
    - **Point to Point** – For example, serial links. In point to point (or point to multipoint) networks, neighboring can be established with the router at the other end of the link.
    - **NBMA** (non-broadcast multiple access) – For example, frame relay. In NBMA networks, neighbors must be explicitly declared.

- **Link State Database** – The Link State Database is composed of the LSA's sent and received by neighboring OSPF routers that have created adjacencies within an area. The database, after complete, contains all the link state information for a given area, at which time the Shortest Path First (SPF) algorithm is applied to determine the optimal route to all connected networks based on cost. The SPF algorithm employs the Dijkstra path finding algorithm, which essentially regards all routers as vertices in a graph, and computes the cost between each vertex.
- **Adjacencies** – OSPF routers exchange LSA's with adjacent routers to create the LSDB. Adjacencies are created in different fashions depending on the network type (see Neighbors above). Generally, the network type is broadcast (for example, Ethernet) so adjacencies are formed by the exchanging OSPF packets in a handshake-like fashion (see OSPF Packet types below). To minimize the amount of information exchanged between adjacent routers, segments (broadcast domains) with multiple OSPF routers elect a Designated Router (DR) and a Backup Designated Router (BDR) using Hello packets.
- **DR (Designated Router)** – On multi-access segments, OSPF routers elect a DR and a BDR, and all other routers on the segment create adjacencies with the DR and the BDR. DR election is based on a router's OSPF Priority, which is a configurable value from 0 (not eligible for DR) to 255. The router with the highest priority becomes the DR. In the event of a priority tie, the router with the highest Router ID (based on interface addressing) wins. When a router is the DR, its role is uncontested until it becomes unavailable. LSA's are then exchanged within LSUs across these adjacencies rather than between each possible pairing combination of routers on the segment; see Routing adjacencies: Designated Router (DR). Link state updates are sent by non-DR routers to the multicast address 225.0.0.6, the RFC1583 assigned 'OSPF All Designated Routers' address. They are also flooded by DR routers to the multicast address 225.0.0.5 'OSPF All Routers' for all routers to receive the LSA's.

#### ROUTING ADJACENCIES: DESIGNATED ROUTER (DR)



- **OSPF Packet types** – The five types of OSPF packets are:
  - **Hello** (OSPF type 1) – Sent at a certain interval to establish and maintain relationships with neighboring OSPF routers, and elect Designated Routers. (*Sent during the initialization and the 2-WAY phases on LSDB synchronization*).
  - **Database Description** (OSPF type 2) – Sent between OSPF routers during the creation of an adjacency. *During the Exstart phase of LSDB synchronization*, DD packets establish an ISN (initial sequence number) used to track LSA's, and they establish a master/slave relationship between neighboring OSPF routers. *In the Exchange phase of LSDB synchronization*, they contain short


versions of Link State Advertisements. Because DD exchanges can span multiple packets, they are exchanged in a poll (master) and response (slave) fashion to ensure completeness.

- **Link State Request** (OSPF type 3) – *During the Loading phase of LSDB synchronization*, LSR packets are sent to request database updates from a neighbor. This is the final step in the establishment of an adjacency.
- **Link State Update** (OSPF type 4) – Sent in response to Link State Requests, LSU packets flood adjacencies with Link State Advertisements to achieve LSDB synchronization.
- **Link State Acknowledgment** (OSPF type 5) – To ensure reliability of LSA flooding, all updates are acknowledged.
- **Link State Advertisements** (LSA) – There are 7 types of LSA's:
  - **Type 1** (Router Link Advertisements) - Sent by an OSPF router to describe the links to each area to which it belongs. Type 1 LSA's are only flooded into a router's area.
  - **Type 2** (Network Links Advertisements) – Sent by the DR for an area describing the set of routers within the network. Type 2 LSA's are only flooded into a router's area.
  - **Type 3** (Summary Link Advertisements) – Sent across areas by ABRs (Area Border Routers) to describe the networks within an area. Type 3 LSA's are also used for route aggregation purposes, and are not sent to Totally Stubby Areas.
  - **Type 4** (AS Summary Link Advertisements) – Sent across areas by ABRs to describe networks within a different AS. Type 4 LSA's are not sent to Stub Areas.
  - **Type 5** (AS External Link Advertisements) – Sent by ASBR (Autonomous System Boundary Routers) to describe routes to networks in a different AS. Type 5 LSA's are not sent to Stub Areas. There are two types of External Link Advertisements:
    - **External Type 1** - Type 1 packets add the internal link cost to the external link cost when calculating a link's metric. A Type 1 route is always preferred over a Type 2 route to the same destination.
    - **External Type 2** - Type 2 packets only use the external link cost to determine the metric. Type 2 is generally used when there is only one path to an external AS.
  - **Type 6** (Multicast OSPF or MOSPF) - Called source/destination routing, this is in contrast to most unicast datagram forwarding algorithms (like OSPF) that route based solely on destination. For more information about MOSPF, see RFC1584 – Multicast Extensions to OSPF.
  - **Type 7** (NSSA AS External Link Advertisements) – Sent by ASBRs that are part of an NSSA (see 'Stub Area').
  - **Stub Area** – A stub area is an area that only requires one path, rather than an optimal path. This can be an area with only a single point of egress, or it can be an area where SPF optimization is not necessary. All routers in a stub area must be configured as stub routers, and rather than receiving the full state database, and computing the SPF tree, they receive only summary link information. There are different type of stub area:
    - **Stub area** – The standard stub area receives all LSA's except for LSA type 5 (AS External Link advertisement). This helps to keep the LSDB smaller, and reduces the computational overhead on the router.



- **Totally Stubby Area** – A special type of stub area into which LSA types 3 (Summary Links), 4 (AS Summary Links) and 5 are not passed. Only intra-area routes, and a default route are advertised into totally stubby areas.
- **NSSA** (Not So Stubby Area) – Described by RFC3101, NSSA is a hybrid stub area that allows external routes to be flooded within the NSSA area using type 7 LSA's (NSSA AS External Routes), but does not accept type 5 LSA's from other areas. NSSAs are useful when connecting a remote site running a different IGP (such as RIP) to an OSPF site, where the remote site's routes do not need to be distributed back to the main OSPF site. An NSSA ABR (Area Border Router) also has the ability to translate type 7 to type 5 LSA's (this is possible only from the SonicOS CLI).
- **Router Types** – OSPF recognizes 4 types of routers, based on their roles; see [OSPF-Recognized Router Types Example](#).

### OSPF-RECOGNIZED ROUTER TYPES EXAMPLE



- **IR** (Internal Router) - A router whose interfaces are all contained within the same area. An internal router's LSDB only contains information about its own area.
- **ABR** (Area Border Router) – A router with interfaces in multiple areas. An ABR maintains LSDBs for each area to which it is connected, one of which is typically the backbone.
- **Backbone Router** – A router with an interface connected to area 0, the backbone.
- **ASBR** (Autonomous System Boundary Router) – A router with an interface connected to a non-OSPF AS (such as a RIP network) which advertises external routing information from that AS into the OSPF AS.

# Drop Tunnel Interface

A drop tunnel interface prevents traffic from being sent out using an incorrect route when the configured route is down. Traffic sent to a drop tunnel interface does not leave the appliance, but is ostensibly dropped.

A drop tunnel interface should be used in conjunction with a VPN tunnel interface, although a drop tunnel interface can be used standalone. If a static route is bound to a tunnel interface, SonicWall recommends configuring a static route bound to a drop tunnel interface for the same network traffic. That way, if the tunnel interface goes down, the second static route is used and the traffic is effectively dropped. This prevents the data from being forwarded in the clear over another route.

When configuring a route over a VPN tunnel interface, if the tunnel is temporarily down, the corresponding route entry is disabled as well. SonicOS looks up a new route entry for the connections destined for the VPN protected network. In deployments that do not have a backup link for a remote VPN network, no other correct route entry is available. Traffic is sent to a wrong route entry, generally the default route, which causes security issues such as internal data sent without encryption.

For deployments without a backup link, consider configuring the route table as in this example:

```
route n:    local VPN network(source), remote VPN network(destination), VPN TI(egress_if)

route n+1:  local VPN network(source), remote VPN network(destination), Drop If(egress_
if)
```

When the VPN tunnel interface configured as in this example, the traffic matches the drop interface and is not sent out. When the VPN tunnel interface resumes, traffic resumes also.

# App-based Routing

App-based Routing is a kind of PBF (policy-based forwarding) rule that allows traffic to take an alternative path from the next hop specified in the route table and is typically used to specify an egress interface for security or performance reasons.

When an App-based Route entry is created, at the beginning the appliance does not have enough information to identify the application and, therefore, cannot enforce the route entry. As more packets arrive, the appliance determines the application and creates an internal entry in the App-ID cache, which is retained for the session. When a new session is created with the same destination IP address, destination port, and protocol ID, the appliance could identify the application as the same from the initial session and apply the App-based Route. Therefore, a session that is not an exact match and is not the same application, cannot be forwarded based on the App-based Route.

This feature is available only when Gateway AV/Anti-Spyware/Intrusion Prevention/App Control/App Visualization is licensed and App Control is enabled in **POLICY | Rules and Policies >** .

# Rules and Policies >

If you have routers on your interfaces, you configure static routes on the SonicWall appliance on the **POLICY | Rules and Policies >** page. You can create static routing rule policies that create static routing entries that make decisions based upon source address, source netmask, destination address, destination netmask, service, interface, gateway and metric. This feature allows for full control of forwarding based upon a large number of user-defined variables.

## Topics:

- [Configuring](#)

## Configuring

If you have routers on your interfaces, you can configure the SonicWall appliance to route network traffic to specific predefined destinations. Static routes must be defined if the network connected to an interface is segmented into subnets, either for size or practical considerations. For example, a subnet can be created to isolate a section of a company, such as finance, from network traffic on the rest of the LAN, DMZ, or WAN.

When configuring a static route, you can optionally configure a Network Monitor policy rule for the route. When a Network Monitor policy rule is used, the static route is dynamically disabled or enabled, based on the state of the probe for the policy rule. For more information, see [Probe-Enabled Policy-based Routing Configuration](#).

## Topics:

- [Adding Static Routes](#)
- [Probe-Enabled Policy-based Routing Configuration](#)



- a. Select the interface through which these packets are routed from **Interface**.
  - b. Select the address object that acts as a gateway for packets matching these settings from **Gateway**.
  - c. Specify the RIP metric in the **Metric** field.
8. Click **Add** or click to the **Advanced** tab to continue the configuration.
  - a. Optionally select **Disable route when the interface is disconnected**.
  - b. Select **Allow VPN path to take precedence** to allow a matching VPN network to take precedence over the static route when the VPN tunnel is up. This option is not selected by default.
  - c. Enter the ToS hexadecimal value in the **TOS (Hex)** field.
  - d. Enter the ToS Mask hexadecimal value in the **TOS Mask (Hex)** field.
  - e. Enter a value for the **Admin Distance**, or select **Auto** for an automatically created **Admin Distance**.
9. Click **Add** or click to the **Probe** tab to continue the configuration.
  - a. Select a probe type from **Probe**. The default is **None**. If a probe type is selected additional options become available.
  - b. Select **Disable route when probe succeeds**. This option is not selected by default.
  - c. Select **Probe default state is UP**.
10. When you are finished, click **Add**. The route settings are configured for the selected SonicWall appliance (s).

## Probe-Enabled Policy-based Routing Configuration

You can optionally configure a Network Monitor policy for the route. When a Network Monitor policy is used, the static route is dynamically disabled or enabled, based on the state of the probe for the policy.

Policy-based Routing is fully supported for IPv6 by selecting IPv6 address objects and gateways for route policies on the **POLICY | Rules and Policies >** page. IPv6 address objects are listed in the **Source**, **Destination**, and **Gateway** columns of the **Route Policies** table. Configuring routing policies for IPv6 is nearly identical to IPv4.

### *To configure a policy-based route:*

1. Navigate to the **POLICY | Rules and Policies >** page.
2. Click **+Add** (in the bottom left corner). The **Adding Rule** dialog displays.

3. Click the **Probe** view and select the appropriate Probe Network Monitor object or select **Create a new Network Monitor Object...** to dynamically create a new object.

**NOTE:** Typical configurations do not have **Disable route when probe succeeds** checked because typically a route is disabled when a probe to the route's destination fails. This option is provided to give you added flexibility for defining routes and probes.

4. Select the **Probe default state is UP** to have the route consider the probe to be successful (such as in the UP state) when the attached Network Monitor policy is in the UNKNOWN state. This is useful to control the probe-based behavior when a unit of a High Availability pair transitions from IDLE to ACTIVE, because this transition sets all Network Monitor policy states to UNKNOWN.
5. Click **Add** to apply the configuration.

# SonicWall Support

Technical support is available to customers who have purchased SonicWall products with a valid maintenance contract.

The Support Portal provides self-help tools you can use to solve problems quickly and independently, 24 hours a day, 365 days a year. To access the Support Portal, go to <https://www.sonicwall.com/support>.

The Support Portal enables you to:

- View knowledge base articles and technical documentation
- View and participate in the Community forum discussions at <https://community.sonicwall.com/technology-and-support>.
- View video tutorials
- Access <https://mysonicwall.com>
- Learn about SonicWall Professional Services
- Review SonicWall Support services and warranty information
- Register for training and certification
- Request technical support or customer service

To contact SonicWall Support, visit <https://www.sonicwall.com/support/contact-support>.

# About This Document

SonicOS Rules and Policies for Policy Mode Administration Guide

Updated - May 2024

Software Version - 7.0

232-005343-00 Rev B

Copyright © 2024 SonicWall Inc. All rights reserved.

The information in this document is provided in connection with [[[Undefined variable Company\_Information. the ]]] and/or its affiliates' products. No license, express or implied, by estoppel or otherwise, to any intellectual property right is granted by this document or in connection with the sale of products. EXCEPT AS SET FORTH IN THE TERMS AND CONDITIONS AS SPECIFIED IN THE LICENSE AGREEMENT FOR THIS PRODUCT, [[[UNDEFINED VARIABLE COMPANY\_ INFORMATION. THE ]]] AND/OR ITS AFFILIATES ASSUME NO LIABILITY WHATSOEVER AND DISCLAIMS ANY EXPRESS, IMPLIED OR STATUTORY WARRANTY RELATING TO ITS PRODUCTS INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT. IN NO EVENT SHALL [[[UNDEFINED VARIABLE COMPANY\_ INFORMATION. THE ]]] AND/OR ITS AFFILIATES BE LIABLE FOR ANY DIRECT, INDIRECT, CONSEQUENTIAL, PUNITIVE, SPECIAL OR INCIDENTAL DAMAGES (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF PROFITS, BUSINESS INTERRUPTION OR LOSS OF INFORMATION) ARISING OUT OF THE USE OR INABILITY TO USE THIS DOCUMENT, EVEN IF [[[UNDEFINED VARIABLE COMPANY\_ INFORMATION. THE ]]] AND/OR ITS AFFILIATES HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. [[[Undefined variable Company\_Information. the ]]] and/or its affiliates make no representations or warranties with respect to the accuracy or completeness of the contents of this document and reserves the right to make changes to specifications and product descriptions at any time without notice. and/or its affiliates do not make any commitment to update the information contained in this document.

For more information, visit <https://www.sonicwall.com/legal>.

## End User Product Agreement

To view the [[[Undefined variable Company\_Information. the ]]] End User Product Agreement, go to:

<https://www.sonicwall.com/legal/end-user-product-agreements/>.

## Open Source Code

SonicWall Inc. is able to provide a machine-readable copy of open source code with restrictive licenses such as GPL, LGPL, AGPL when applicable per license requirements. To obtain a complete machine-readable copy, send your written requests, along with certified check or money order in the amount of USD 25.00 payable to "SonicWall Inc.", to:

General Public License Source Code Request  
Attn: Jennifer Anderson  
1033 McCarthy Blvd  
Milpitas, CA 95035