SOLARWINDS WHITEPAPER

Network Routing Protocols – Back to Basics

solarwinds

Routing protocols are used to determine the optimal path for data communication between network nodes. Routers use them to share routing information with other routers to dynamically build global routing tables. The routing protocols are employed when your organization's network grows to the point where static routes are unmanageable. Modern enterprise networks need dynamic routing tables that automatically adjust if there are any traffic or topology changes.

Different Types of Routing Protocols

There are two major classes of routing protocols: Exterior Gateway Protocol (EGP) and Interior Gateway Protocol (IGP). EGP is used to exchange routing information between autonomous systems. For instance, EGP is used in data transfers between ISPs (Internet Service Providers) to ISPs or between autonomous systems to ISPs. Whereas, IGP (Interior Gateway Protocol) is used for exchanging routing information between routers within an autonomous system, like data transfers within your organization's local area network (LAN). IGP can be further classified into two categories: Distance-Vector and Link-State Routing Protocols.

Types of Routing Protocols					
Exterior Gateway Protocols (EGP)	**Interior Gateway Protocols (IGP)**				
Border Gateway Protocol (BGP)	**Distance - Vector**			**Link - State**	
	RIP	**IGRP**	**EIGRP**	**OSPF**	**IS - IS**

In Distance-Vector Routing Protocols, routers communicate with neighboring routers, periodically informing them about network topology changes. Whereas in link-state routing protocol, routers create a roadmap of how they are connected in the network. By calculating the best path from that router to every possible destination in the network, link state routing protocols form the routing table. *RIP* (Routing Information Protocol), *RIPv2*, *IGRP* (Interior Gateway Routing Protocol), and *EIGRP* (Enhanced IGRP) are part of Distance-Vector Routing Protocols. However, *OSPF* (Open Shortest Path First) and *IS-IS* (Intermediate System to Intermediate System) are part of Link-State Routing Protocols.

Open Shortest Path First (OSPF)

OSPF, a link - state routing protocol, is used in large organizations for their Autonomous System (AS) networks. OSPF gathers link state information from available routers and determines the routing table information to forward packets to based on the destination IP address. This occurs by creating a topology map for the network. Any change in the link is immediately detected and the information is forwarded to all other routers, meaning they also have the same routing table information. Unlike RIP, OSPF only multicasts routing information when there's a change in the network. OSPF is used in complex networks that are subdivided to ease network administration and optimize traffic. It quickly calculates the shortest path if topology changes, using minimum network traffic.

Version #	Type	Packet Length
Router ID		
Area ID		
Check Sum	AuthType	
Authentication		
Authentication		

Figure 1: OSPF Packet Format

OSPF allows network admins to assign cost metrics for a particular router so that some paths are given higher preference. OSPF also provides an additional level of routing protection capability and ensures that all routing protocol exchanges are authenticated.

OSPF Message Types

OSPF doesn't send information using UDP. Instead, it builds IP datagrams directly, packaging them using protocol number 89 for the IP protocol field. Different message types of OSPF include:

- ***Hello Packet*** – Sent by routers to set up relationships with neighbors and communicate frequently to keep the connection alive. Hello Packet shares key parameters on how OSPF is to be used within the network.

- ***Database Description*** – The description of the link state database for autonomous systems are transmitted from one router to another.
- ***Link State Request*** – This is requested when a portion of the network needs to be updated with current information. The message specifies exactly which links are requested by the device that wants more current information.
- ***Link State Update*** – This contains the updated information for the requested links. It's sent in response to the LS request.
- ***Link State Acknowledgement*** – This acknowledges the link-state exchange process for link state update message.

OSPF – Pros and Cons

OSPF routing protocol has a complete knowledge of network topology allowing routers to calculate routes based on incoming requests. Additionally, OSPF has no limitations in hop count, it converges faster than RIP, and has better load balancing. A downside with OSPF is that it doesn't scale when there are more routers added to the network. This is because it maintains multiple copies of routing information. An OSPF network with intermittent links can increase traffic every time a router sends information. This lack of scalability in OSPF makes it unsuitable for routing across the Internet.

Routing Information Protocol (RIP)

Routing Information Protocol (RIP), is one of the most commonly used routing protocols for small homogeneous networks. As a distance-vector routing protocol, RIP is used by routers to exchange topology information periodically by sending out routing table details to neighboring routers every 30 seconds. These neighboring routers in turn forward the information to other routers until they reach network convergence. RIP uses the hop count metric with the maximum limit of 15 hops, anything beyond that is unreachable. Because of this, RIP is not suitable for large, complex networks.

RIPv1 vs. RIPv2

There are two versions of RIP. RIP version 1 uses classful routing and does not include subnet information while sending out periodic routing table updates. RIPv2 is classless and includes the subnet information supporting

Classless-Inter Domain Routing (CIDR). Unlike RIP version 1, version 2 multicasts the routing updates to the adjacent routers using the address 224.0.0.9. Network convergence happens much faster in RIPv2.

<table>
<tr><td colspan="2">Command</td><td colspan="2">Version</td><td colspan="2">Must be Zero (Unused)</td></tr>
<tr><td colspan="3">Address Family Identifier</td><td colspan="3">Route Tag</td></tr>
<tr><td colspan="6">IP Address</td></tr>
<tr><td colspan="6">Subnet Mask</td></tr>
<tr><td colspan="6">Next Hop</td></tr>
<tr><td colspan="6">Metric</td></tr>
</table>

Figure 2 RIPv2 Packet Format

RIP – Advantages and Disadvantages

Routing Information Protocol has its own advantages in small networks. It's easy to understand, configure, widely used, and is supported by almost all routers. Since its limited to 15 hops, any router beyond that distance is considered as infinity, and hence unreachable. If implemented in a large network, RIP can create a traffic bottleneck by multicasting all the routing tables every 30 seconds, which is bandwidth intensive. RIP has very slow network convergence in large networks. The routing updates take up significant bandwidth leaving behind very limited resources for critical IT processes. RIP doesn't support multiple paths on the same route and is likely to have more routing loops resulting in a loss of transferred data. RIP uses fixed hop count metrics to compare available routes, which cannot be used when routes are selected based on real-time data. This results in an increased delay in delivering packets and overloads network operations due to repeated processes.

Enhanced Interior Gateway Routing Protocol (EIGRP)

EIGRP, a distance vector routing protocol, exchanges routing table information with neighboring routers in an autonomous system. Unlike RIP, EIGRP shares routing table information that is not available in the neighboring routers, thereby reducing unwanted traffic transmitted through routers. EIGRP is an enhanced version of IGRP and uses Diffusing Update Algorithm (DUAL), which reduces the time taken for network convergence and improves operational efficiency. EIGRP was a proprietary protocol from Cisco®, which was later made an Open Standard in 2013.

Version	OPCode	CheckSum
Flag		
Sequence Number		
Acknowledgement Number		
Autonomous System (AS) Number		
Type/Length/Value		

Figure 3 EIGRP Packet Format

EIGRP Packet Types

Different message types in EIGRP include:

- ***Hello Packet*** – The first message type sent when EIGRP process is initiated on the router. *Hello* packet identifies neighbors and forms adjacencies while being multicast every 5 seconds by default (60 seconds on low bandwidth network).
- ***Update Packet*** – Contains route information that is only forwarded when there is a change. They are only sent to the routes that have partial updates. If there's a new neighbor discovered, the packet is then sent to the router as a unicast.
- ***Acknowledgement*** – This is unicast as a response to *Update* packet by acknowledging when they receive an update.
- ***Query*** – This packet is sent to query routes from neighbors. When a router loses a route while sending the multicast, *Query* packet is sent to all neighboring routers to find alternate paths for the router.
- ***Reply*** – These are unicast by routers that know alternate routes for the neighboring routers queried on a network.

EIGRP – Pros and Cons

Speedy network convergence, low CPU utilization, and ease of configuration are some of the advantages of EIGRP. The EIGRP routers store everything as routing table information so they can quickly adapt to alternate routes. The variable length subnet mask reduces time to network convergence and increases scalability. EIGRP also includes MD5 route authentication. Compared to RIP and OSPF, EIGRP has more adaptability and versatility in complex networks. EIGRP combines many features of both link-state and distance-vector. Since EIGRP is mostly deployed in large networks, routers tend to delay sending information during allotted time, which can cause neighboring routers to query the information again, thus increasing traffic.

How to Choose a Suitable Routing Protocol for Your Network

Distance vector routing Protocols like RIP and EIGRP are ideal for small networks that are simple and non-hierarchical. Enterprises use link state routing protocols like OSPF and IS-IS for their large and hierarchical networks, while distributed networks will use BGP to establish routing information between autonomous systems. For instance, network administrators using OSPF will have advanced knowledge about complex networks which helps them in troubleshooting routing related issues.

Network admins choose routing protocols based on convergence time: the time taken for all the routers to collect the status of current topological information about the network. If you have three routers in the network and one of the links that connects the network has failed, the information about the status should be immediately available in all the routers by the process of convergence. The slower they converge, the harder they become for network admins to troubleshoot. Focusing on easing network route management, network configuration and troubleshooting become very important to admins managing large enterprise networks.

How SolarWinds® Can Help

Advanced network monitoring tools have the ability to monitor network route information and provide real-time views on issues that might affect the network. **SolarWinds Network Performance Monitor (NPM)** is a network fault, availability, and performance management software that enables you to monitor and alert on large, complex network routes including support for major routing protocols (RIP v2, OSPF v2 & v3, EIGRP, BGP). As your network size and complexity grow and change it becomes increasingly difficult to monitor and troubleshoot routing related performance issues. Manually retrieving route information using CLI is time consuming and error prone.

SolarWinds NPM gives you the ability to monitor network route information and receive alerts when issues arise, enabling you to reduce your time to resolution by providing a combined view of real-time network route information alongside device information. You can now view routing tables, changes in default routes, BGP transitions and flapping routes. It also provides visual alert indicators and audible alarms, and includes more than sixteen built-in network alert delivery methods and responses, including email, pages, SNMP traps, text-to-speech, syslog messages, and the launching of an external application.

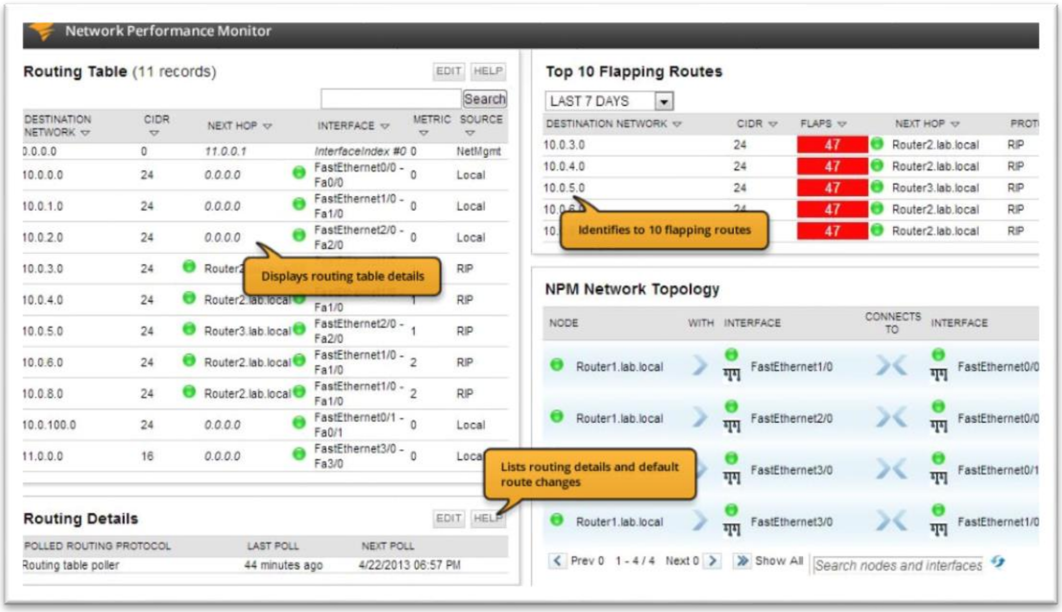

LEARN MORE »   DOWNLOAD FREE TRIAL

About SolarWinds

SolarWinds (NYSE: SWI) provides powerful and affordable IT management software to customers worldwide. Focused exclusively on IT pros, we strive to eliminate the complexity in IT management software that many have been forced to accept from traditional enterprise software vendors. SolarWinds delivers on this commitment with unexpected simplicity through products that are easy to find, buy, use, and maintain, while providing the power to address any IT management problem on any scale. Our solutions are rooted in our deep connection to our user base, which interacts in our online community, thwack®, to solve problems, share technology and best practices, and directly participate in our product development process. Learn more at http://www.solarwinds.com.