ISSN (Online) 2394-2320

International Journal of Engineering Research in Computer Science and Engineering (IJERCSE)
Vol 3, Issue 10, October 2016

Redistributing Routing Protocols

[1]Saurabh Kumar, [2]Neha Bhadana, [3]Maneesh Kumar
M .Tech Scholar, Department of Computer Science and Engineering,
Dr. A.P.J Abdul Kalam Technical University, Lucknow (U.P),

Abstract: **-- A distributed routing protocol system that is the use of a routing protocol to advertise routes that are learned by some other means, such as by another routing protocol, static routes, or directly connected routes, is called redistribution. While running a single routing protocol throughout your entire IP internetwork is desirable, multi−protocol routing is common for a number of reasons, such as company mergers, multiple departments managed by multiple network administrators, and multi−vendor environments.**

I. INTRODUCTION

The use of a routing protocol to advertise routes that are learned by some other means, such as by another routing protocol, static routes, or directly connected routes, is called redistribution. While running a single routing protocol throughout your entire IP internetwork is desirable, multi−protocol routing is common for a number of reasons, such as company mergers, multiple departments managed by multiple network administrators, and multi−vendor environments. Running different routing protocols is often part of a network design. In any case, having a multiple protocol environment makes redistribution a necessity. Differences in routing protocol characteristics, such as metrics, administrative distance, classfull and classless capabilities can effect redistribution. Consideration must be given to these differences for redistribution to succeed.

II. PREREQUISITES REQUIREMENTS

There are no specific requirements for this document.

Components Used

The information in this document is based on these software and hardware versions:

- Cisco IOS® Software Release 12.2(10b)
- Cisco 2500 Series Routers

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, make sure that you understand the potential impact of any command.

Conventions

Refer to Cisco Technical Tips Conventions for more information on document conventions.

Metrics

When you redistribute one protocol into another, remember that the metrics of each protocol play an important role in redistribution. Each protocol uses different metrics. For example, the Routing Information Protocol (RIP) metric is based on hop count, but Interior Gateway Routing Protocol (IGRP) and Enhanced Interior Gateway Routing Protocol (EIGRP) use a composite metric based on bandwidth, delay, reliability, load, and maximum transmission unit (MTU), where bandwidth and delay are the only parameters used by default. When routes are redistributed, you must define a metric that is understandable to the receiving protocol. There are two methods to define metrics when redistributing routes. You can define the metric for that specific redistribution only:

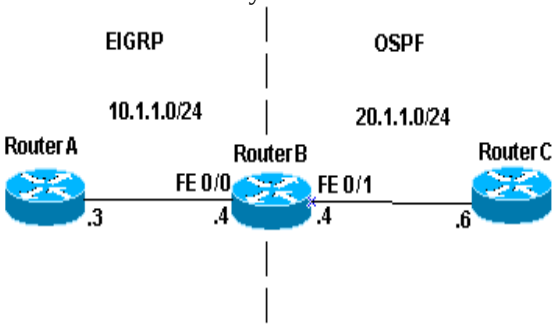

ISSN (Online) 2394-2320

International Journal of Engineering Research in Computer Science and Engineering (IJERCSE)
Vol 3, Issue 10, October 2016

Router rip
Redistribute static metric 1
Redistributeospf 1 metric 1

Or you can use the same metric as a default for all redistribution (Using the default−metric command saves work because it eliminates the need for defining the metric separately for each redistribution.):
Router rip
Redistribute static
Redistributeospf 1
default−metric 1

III. ADMINISTRATIVE DISTANCE

If a router is running more than one routing protocol and learns a route to the same destination using both routing protocols, then which route should be selected as the best route? Each protocol uses its own metric type to determine the best route. Comparing routes with different metric types cannot be done. Administrative distances take care of this problem. Administrative distances are assigned to route sources so that the route from the most preferred source will be chosen as the best path. Refer to Route Selection in Cisco Routers for more information about administrative distances and route selection. Administrative distances help with route selection among different routing protocols, but they can cause problems for redistribution. These problems can be in the form of routing loops, convergence problems, or inefficient routing. See below for a topology and description of a possible problem.

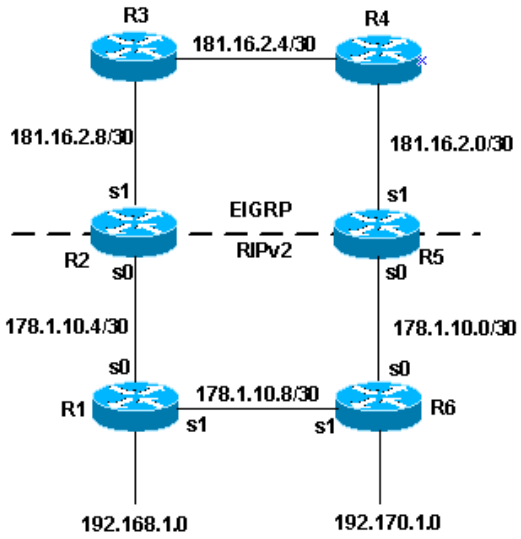

In the above topology, if R1 is running RIP, and R2 and R5 are running both RIP and IGRP and redistributing RIP into IGRP, then there is a potential problem. For example, R2 and R5 are both learning about network 192.168.1.0 from R1 using RIP. This knowledge is redistributed into IGRP. R2 learns about network 192.168.1.0 through R3, and R5 learns about it from R4 using IGRP. IGRP has a lower administrative distance than RIP (100 versus 120); therefore, the IGRP route is what is used in the routing table. Now there is a potential routing loop. Even if split horizon, or any other feature meant to help prevent routing loops comes into play, there is still a convergence problem. If R2 and R5 are also redistributing IGRP into RIP (otherwise known as mutual redistribution) and the network, 192.168.1.0, is not directly connected to R1 (R1 is learning from another router upstream from it), then there is a potential problem that R1 will learn the network from R2 or R5 with a better metric than from the original source.

The mechanics of route redistribution is proprietary on Cisco routers. The rules for redistribution on a Cisco router dictate that the redistributed route be present in the routing table. It is not sufficient that the route be present in the routing topology or database. Routes with a lower Administrative Distance (AD) are always installed in the routing table. For example, if a static route is redistributed into IGRP on R5, and then

ISSN (Online) 2394-2320

International Journal of Engineering Research in Computer Science and Engineering (IJERCSE)
Vol 3, Issue 10, October 2016

IGRP subsequently redistributed into RIP on the same router (R5), the static route is not redistributed into RIP because it never got entered into the IGRP routing table. This is due to the fact that static routes have an AD of 1 and IGRP routes have an AD of 100 and the static route is installed in the routing table. In order to redistribute the static route into IGRP on R5, you need to use the redistribute static command under the router rip command.

The default behavior for RIP, IGRP and EIGRP is to advertise directly connected routes when a network statement under the routing protocol includes the connected interface subnet. There are two methods to get a connected route:

An interface is configured with an IP address and mask, this corresponding subnet is considered a connected route.

• A static route is configured with only an outgoing interface, and not an IP next−hop, this is also considered a connected route.

• Router#conf t
Router(config)#ip route 10.0.77.0 255.255.255.0 ethernet 0/0
Router(config)#end
Router#showip route static
10.0.0.0/24 is subnetted, 1 subnets
S 10.0.77.0 is directly connected, Ethernet0/0

A network command configured under EIGRP, RIP or IGRP that includes (or "covers") either of these types of connected routes includes that subnet for advertisement. For example, if an interface has address 10.0.23.1 and mask 255.255.255.0, the subnet 10.0.23.0/24 is a connected route and will be advertised by these routing protocols when a network statement is configured as follows: router rip | igrp # | eigrp # network 10.0.0.0

This static route, 10.0.77.0/24, is also advertised by these routing protocols, because it is a connected route and it is "covered" by the network statement.

See the Avoiding Problems Due to Redistribution section of this document for tips on how to avoid this problem.

Redistribution Configuration Syntax and Examples IGRP and EIGRP

This output shows an IGRP/EIGRP router redistributing static, Open Shortest Path First (OSPF), RIP, and Intermediate System−to−Intermediate System (IS−IS) routes.

Routerigrp/eigrp 1
Network 131.108.0.0
Redistribute static
Redistributeospf 1
Redistribute rip
Redistribute sis
default−metric 10000 100 255 1 1500

IGRP and EIGRP need five metrics when redistributing other protocols: bandwidth, delay, reliability, load, and MTU, respectively.

Open Shortest Path First (OSPF)

OpenShortestPathFirst(OSPF)isanopenstandard sroutingprotocolthat‘sbeenimplementedbyawidevarietyo fnetworkvendors,includingCisco.ThisworksbyusingtheD ijkstraalgorithm.First,ashortestpathtreeisconstructed,andt hentheroutingtableispopulatedwiththeresultingbestpaths. OSPFconvergesquickly,althoughperhapsnotasquicklyasE IGRP,anditsupportsmultiple,equal-costroutestothesamedestination. But unlike EIGRP, it only supports I routing. OSPF is an IGP protocol. Itisalinkstaterouting protocol. Itissupportedby many operating systems.ItsdefaultADis110,hopcountlimitisunlimited.Itis classlessroutingprotocol,supportsVLSM/CIDR.Bydefault thehighestIPaddressofinterfacewillbeelectedasRouterid.T hebiggestreasonOSPFisthechoiceinlargenetworksisitseffi ciency;insteadofchangingroutingtableviabroadcasttheway RIPdoes,OSPFconfiguredroutersmaintainamapofthenetw ork.Themappingiscal Themappingiscalledthelinkstatedatabase,OSPFrouterskee pthelinkstatedatabaseuptodate.Oncechangeshavebeenmad etolinkstatedatabase,anOSPFrouter‘s link state database is recalculated. Asthenetworksstarttomultiply,thesizeofthelinkstatedatabaseincreases,andacorrespondinghitonrouterperformanceres ults.Areasareconnectedtoeachotherthroughabackbonearea,witheachrouteronlyresponsiblefortheli nkstatedatabaseforthoseareasconnectedtotherou ters.AreaBorderRouters(ABRs)thenconnectone backboneareatoanother.Thebiggestdrawbackof OSPFisitscomplexity;OSPFrequires proper

ISSN (Online) 2394-2320

International Journal of Engineering Research in Computer Science and Engineering (IJERCSE)
Vol 3, Issue 10, October 2016

planningand is moredifficult to configure andadminister.

Features

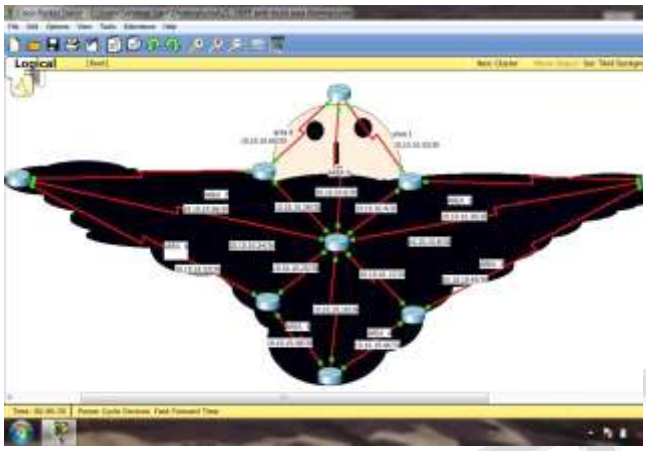

Consists of areas and autonomous systems
Minimizes routing update traffic
Allows scalability
Supports VLSM/CIDR
Has unlimited hop count
Allows multi-vendor deployment (open standard)

Advantages

Routes calculated with OSPFarealwaysloop free.
OSPF can scale much more easily than RIP.
Reconfiguration for network topology changes is faster.
To decrease routing over head
To speed up convergence
To confine network instability to single areas of the network

(i)Minimum routing updates.
(ii) Priorities on all the CISC Or outers theprioritыis 1.
(iii) The routershavinghighestIPaddressbecome BRD (BorderDestinationRouter).

Router ID

The Router ID (RID) is an I address used to identify the router.
CiscochoosestheRouterIDbyusingthehighestIPaddressofallconfiguredloopbackinterfaces.Ifnoloopbackinterfacesareconfiguredwithaddresses,OSPF will choose the highest IP address of all active physical interfaces.

OSPF Areas

AnOSPFareaisagroupingofcontiguousnetworksandrouters.Allroutersinthe same area share a common AreaID.

Broadcast (multi-access)

Networks such as Ethernet allow multiple devices to connect to (ORACCESS) the same network, as well as provide abroad cast ability inwhichasinglepacketisdeliveredtoallnodesonthenetwork. InOSPF, aDRandaBDRmustbeelected for each broadcast multi-access network.

Non-broad castmulti-access

Non-Broadcast Multi-Access (NBMA) networks are types such as Frame Relay, X.25, and Asynchronous Transfer Mode (ATM). These networks allow for multi-access, but have nobroadcastabilitylikeEthernet.So, NBMAnetworksrequirespecialOSPFconfiguration to function properly and neighbor relationships must be defined.

Point-to-point

Point-to pointreferstoatypeofnetworktopologyconsistingofadirectconnectionbetweentworoutersthatprovidesasinglecommunicationpath.Thepoint-topointconnectioncanbephysical,asinaserialcabledirectlyconnectingtworouters,orit can be logical.

Point-to-multipoint

Point-to-multipointreferstoatypeofnetworktopologyconsistingofaseriesof connections between single interface on one router and multiple destination routers.

Steps to apply OSPF

Syntax:

Router (config)#router ospf<ospf process id>
Router (config-router)#network <network address><wild card mask> area<area

ISSN (Online) 2394-2320

International Journal of Engineering Research in Computer Science and Engineering (IJERCSE)
Vol 3, Issue 10, October 2016

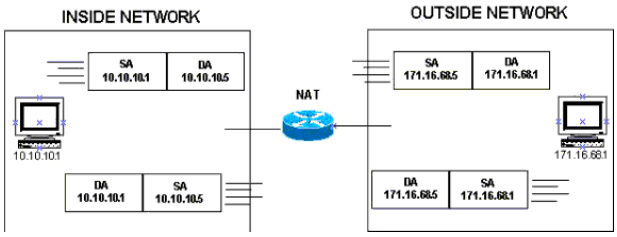

Fig. : Configuration example of RIP routing

(RIP (Routing Information Protocol)

RoutingInformationProtocolisatruedistance-vectorroutingprotocol.ItisanIGB(InterGatewayProtocol). Itsendsthecompleteroutingtableouttoallactiveinterfacesevery30secondstoitsimmediateneighbor Itsendsthecompleteroutingtableouttoallactiveinterfacesevery30secondstoitsimmediateneighbor.Thisisslowconvergencemeansthatoneroutersendsarequesttootheraboutitsroutеornetworkgetnetworkswhicharenotassignedtoitafterallthreeroutershavesamenetworks,thisprocessisrepeatedtosendandreceiverequestsoitiscalledslowconvergenceRIPonly useshopcounttodeterminethebestwaytoremotenetwork,butithasamaximumallowablehopcountof0-15 by default, meaning that 16 is deemed unreachable.
RIPversion1usesonlyclassfullrouting,whichmeansthatall devicesinthenetworkmustusethesamesubnetmask.RIPversion2providessomethingcalledprefixrouting,anddoessend subnetmaskinformationwiththerouteupdates.Thisiscalled classless routing.

Features

RIPversion1andversion2, with the ability to configure individual network cards with separate versions Calculations used to avoid routing loops and speed recovery of the network when ever topology changes occur.
Route filters; you can con figure RIP to except information from only certain networks, and also choose which routes will be shared with RIP routers. Peer filters, which allow control over which router announcement sare accepted. Default administrative distance is 120.
Simplepasswordauthenticationsupport.Buttherearesignificantdrawbacks, whichmakesRIPa poor, if notunusablesolution for large networks.
Forexample,
themaximumhopcountusedforRIProutersis15, makingnetwork16hops away (or more) unreachable where RIP is concerned

IGRP (Interior Gateway Protocol)

Interior Gateway Routing Protocol (IGRP)is a Cisco-proprietary distance vectorroutingprotocol.TouseIGRP,allyourroutersmustbe Ciscorouters.IGRPhasamaximumhopcountof255withadefaultof100.IGRPusesbandwidthanddelayofthelinebydefaultasametricfordeterminingthebestroutetoaninternetwork. Reliability,load,andmaximumtransmissionunit(MTU)can alsobeused,althoughtheyarenotusedbydefault.

EIGRP (Enhanced Interior Routing Protocol)

EIGR Pies sometime preferred to ashy bridrouting protocol because it has characteristics of both distance-vectorandlink-stateprotocols.Itsendstraditionaldistancevectorupdatescontaininginformationaboutnetworksplusthecostofreachingthemfromtheperspectiveoftheadvertingrouter.EIGRPhasamaximumhopcountof255

Features

Powerful features that make EIGRP a real standout from IGRP. Support for IP,IPX, and AppleTalk via protocol-dependent modules Considered classless (sameas RIPv2and OSP
Support for VLSM/CIDR.
Support for summaries and discontinuous networks.
Efficient neighbor discovery.
Communication via Reliable Transport Protocol (RTP).
Best path selection via Diffusing Update Algorithm (DUAL).

ISSN (Online) 2394-2320

International Journal of Engineering Research in Computer Science and Engineering (IJERCSE)
Vol 3, Issue 10, October 2016

Cisco calls EIGR Pa distance vector routing protocol, or sometimes an advanced distance vector or even a hybrid routing protocol.
EIGRP supports different Network layer protocols through the use of protocol-dependent modules (PDMs). Each EIGRPPDM will maintain as eparateseries of table containing the routing information that applies to a specific protocol. It means that there will be IP/EIGRP tables, IPX/EIGRP tables, and AppleTalk/EIGRP tables.

Neighbor Discovery

BeforeEIGRProutersarewillingtoexchangeroute switheachother,theymustbecomeneighbors.Tomaintainth eneighbourshiprelationship,EIGRProutersmustalsocontin uereceivingHellosfromtheirneighbors.EIGRProutersthatb elongtodifferentautonomoussystems(ASes)don'tautomati callyshareroutinginformation and automatically share routing information and they don't become neighbors. There are three conditions that must be met forneigh borship establishment: Hello or ACK received
AS numbers match
Identical metrics (K values)

Avoiding Problems Due to Redistribution

In the section on administrative distance you saw how redistribution can potentially cause problems such as below optimal routing, routing loops, or slow convergence. Avoiding these types of problems is really quite simple never announce the information originally received from routing process X back into routing process X. Example 1

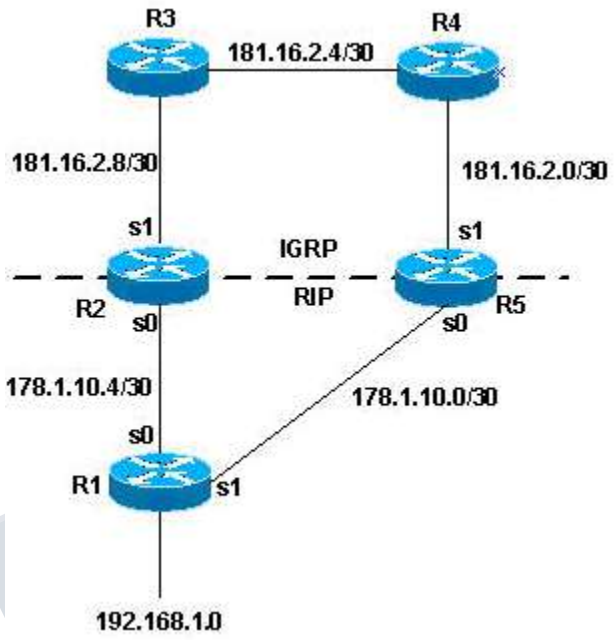

In the previous topology, R2 and R5 are doing mutual redistribution. RIP is being redistributed into IGRP and IGRP is being redistributing into RIP, as this configuration shows.

R2
routerigrp 7
Network 181.16.0.0
Redistribute rip metric 1 1 1 1 1
Router rip
Network 178.1.0.0
Redistributeigrp 7 metric 2
R5:
Routerigrp 7
Network 181.16.0.0
Redistribute rip metric 1 1 1 1 1
Router rip
Network 178.1.0.0
Redistributeigrp 7 metric 2

With the previous configuration you have the potential for any the the problems previously described. In order to avoid them, you can filter routing updates as follows:
R2:
Routerigrp 7

ISSN (Online) 2394-2320

International Journal of Engineering Research in Computer Science and Engineering (IJERCSE)
Vol 3, Issue 10, October 2016

Network 181.16.0.0
Redistribute rip
Metric 1 1 1 1 1
distribute−list 1 in s1
Router rip
Network 178.1.0.0
Redistributeigrp 7 metric 2
access−list 1 deny 192.168.1.0
access−list 1 permit any
R5:
Routerigrp 7
Network 181.16.0.0
Redistribute rip metric 1 1 1 1 1
distribute−list 1 in s1
Router rip
Network 178.1.0.0
redistributeigrp 7 metric 2
access−list 1 deny 192.168.1.0
access−list 1 permit any

The distribute lists added to the configurations, as shown above, filter any IGRP updates that come into the serial 1 interface of the routers. If the routes in the updates are permitted by access list 1, the router accepts them in the update; otherwise it does not. In this example, the routers are being told that they should not learn network 192.168.1.0 through the IGRP updates they receive on their serial 1 interface. Therefore, the only knowledge these routers have for network 192.168.1.0 is through RIP from R1. Also keep in mind that in this case it is not necessary to use the same filter strategy for the RIP process because RIP has a higher administrative distance than IGRP. If routes that originate in the IGRP domain were fed back to R2 and R5 through RIP, the IGRP routes would still take precedence.

REFERENCE

1) Leeds City College, CCNA – Routing Protocols andConcepts, http://www.leedscitycollege.ac.uk/courses/index.php/3260/ccna-routing-protocols-and-concepts/ (Nov. 16, 2010). [*dead link*]

2) Jump up^ Diane Teare, *Implementing Cisco IP Routing (ROUTE)*: *Foundation Learning Guide* 22 (Cisco Press, Mar. 2013)

3) Jump up^ Teare, *Implementing Cisco IP Routing (ROUTE)* at 22–25

4) Jump up^ Cisco Systems, Inc., *Configuring a Gateway of Last Resort Using IP Commands*, http://www.cisco.com/c/en/us/support/docs/ip/routing-information-protocol-rip/16448-default.html(last updated Aug. 10, 2005).

5) Jump up^ Teare, *Implementing Cisco IP Routing (ROUTE)* at 28–29

6) Jump up^ *Id*. at 26–28.