GIAC
CERTIFICATIONS

Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permited without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials: Network, Endpoint, and Cloud (Security 401)"
at http://www.giac.org/registration/gsec

Protecting Network Infrastructure at the Protocol Level

Scope of paper

This paper will briefly discuss attacks and attack prevention methods for network infrastructure protocols. Particular focus will be given to router and routing protocol vulnerabilities such as Routing Information Protocol (RIP), Border Gateway Protocol (BGP), Open Shortest Path First (OSPF), and others.

Routers perform a critical function for each network and if a router is compromised or a route is successfully spoofed, network integrity can be seriously damaged especially if hosts are not using encrypted communications channels. The potential for data manipulation through man-in-the-middle attacks, denial of service, data loss, disruption of network integrity, and packet sniffing is great. Security mechanisms are often available, but are commonly not used because attacks on routing protocols have been rare. Due to the lack of hard data on actual incidents, some approaches outlined in this paper will be theoretical in nature.

Routing is a huge and complex topic; therefore this document will be updated and corrected as I continue my research. Note that I am not a routing engineer and would be glad to accept corrections to any information contained herein.

Commonly known router security issues

Various types of routers have well-known security issues. A collection of some of the commonly known vulnerabilities for network infrastructure equipment vendors such as Cisco, Livingston, Bay and others, can be found at http://www.antionline.com/cgi-bin/anticode/anticode.pl?dir=router-exploits. Most of these vulnerabilities are non routing-protocol level attacks that rely on misconfiguration, bugs in IP packet handling, SNMP insecurities such as default community name strings, weak password or weak password encryption, DOS conditions due to bad IP/UDP packets, etc. These types of attacks are commonly known, and a standard NIDS should be able to be programmed to detect these, at least on an IP based network. IDS are still in the emerging stages as far as non-TCP/IP based routing protocols are concerned. Any of these types of attacks can weaken a network infrastructure and could be used in combination with higher-level protocol-based attacks.

Proper configuration management can resolve many of these common vulnerabilities. This would involve standard procedures such as not using SNMP (or choosing strong passwords/encryption), keeping up to date with vendor patches, proper use of access lists, ingress/egress filtering, firewalls, encrypted management channels and passwords, route filtering, and use of MD5 authentication. However, to understand and implement these security procedures, network engineers must be given the time and training to understand the security implications of their work

© SANS Institute 2000 - 2002 As part of GIAC practical repository. Author retains full rights.

Recent Developments in Infrastructure Defense

A recent development in network defense comes in an IDS called JiNao, which can be found at http://www.anr.mcnc.org/projects/JiNao/JiNao.html. JiNao is funded by DARPA and is currently in development as a joint research project between MCNC and North Carolina State University. JiNao runs on FreeBSD and Linux in on-line mode (using divert sockets) and Solaris in off-line mode, and has been tested in three network testbeds -MCNC, NCSU, and the AF/Rome Laboratory which has a mixture of PC's (operating as routers) and commercial routers. Test results demonstrated various types of successful network infrastructure attacks and also demonstrated that these attacks can be detected with a high degree of accuracy.

At this time, JiNao seems to work mostly with the Open Shortest Path First (OSPF) protocol, but evidently could be expanded to cover other protocols fairly easily. In a nutshell, JiNao features "attack prevention and intrusion detection with highly integrated network management components" (Jou, Gong, et. al, 1999). JiNao functions in some ways like a combination between a network firewall, an intrusion detection system, and a network management system.

Another tool I have found is a modification of *fdget.c*, a program released by Cisco. In 1998, Walt Prue adapted this program "to look at the netflow data records and search for illegitimate default pointing or transit routing from unauthorized source AS's to unauthorized destination AS's." Unfortunately, I was unable to obtain the code itself so further analysis is difficult.

While I've been unable to find much evidence of actual intrusion detection packages for many routing protocols, I image that a high level protocol analysis tool such as the Agilent Advisor (http://onenetworks.comms.agilent.com/) which supports many routing protocols could be customized with filters to detect anomalous behavior.

Tools for working with routing protocols

The following section is an incomplete listing of tools that may be used for working with routing protocols. Some of these tools will be mentioned in more depth but a detailed examination is beyond the scope of this paper.

Linux divert sockets, is described as follows:"Divert sockets enable both IP packet interception and injection on the end-hosts as well as on the routers. Interception and injection happen at the IP layer. The intercepted packets are diverted to sockets in the user space, thus they will not be able to reach their destination unless the user space sockets re-inject them. This allows different tricks (e.g., routing and firewall) to be played, outside the operating system kernel, in between the packet interception and

© SANS Institute 2000 - 2002, Author retains full rights.

reinjection." (http://www.anr.mcnc.org/~divert/)

Divert sockets can be found at http://www.anr.mcnc.org/~divert/. Divert sockets were originally implemented on FreeBSD but have been ported to Linux and were used as part of the JiNao IDS.

The Nemesis Packet Injection suite is a powerful network and security utility written by Obecian and can be obtained from http://www.packetninja.net. The latest version at the time of writing is nemesis-1.1 which was released on June 24th, 2000. Nemesis is "a command-line UNIX network packet injection suite" and can be a very powerful tool for testing firewalls, intrusion detection systems, routers, and other elements of a network. It can also be used by attackers and by authorized penetration testers to attempt to circumvent network security at the host and network level. It appears that the intentions of Obecian were to provide a helpful tool to the security community and the networking industry.

The next evolution of Nemesis is a package called *Intravenous*, which has yet to be released as of 11/30/00. Intravenous appears to be carry on the basic functionality of nemesis but within the context of an artificial intelligence engine. Information about Intravenous can be found at the packetninja.net web site.

IRPAS, Internetwork Routing Protocol Attack Suite, written by FX, can be found at http://www.phenoelit.de/irpas/. IRPAS is on it's first generation of code, but the a revision is taking place and shows much promise. IRPAS contains various command line tools that work with Cisco routing equipment at the protocol level. These include *cdp*, which sends Cisco router Discovery Protocol (CDP) messages; *igrp* for injecting Interior Gateway Routing Protocol (IGRP) messages; *irdp* for sending ICMP Router Discovery Protocol messages; *irdresponder*, which responds to IRDP requests with crafted packets; and *ass*, the Autonomous System Scanner, which "works like a TCP port scanner" for Autonomous Systems. The IRPAS website also contains a link to paper on Generic Routing Encapsulation (GRE) vulnerabilities that may allow an outside attacker to bypass NAT and exploit an internal RFC1918 network through a VPN. This paper can be found at http://www.phenoelit.de/irpas/gre.html. More information and possible attack strategies with *irpas* will be included in a separate section of this paper.

FX, the *irpas* developer, sent an example of AS scanning with the new (unreleased) version 2.14 of *ass*, and how the information from *ass* (AS #10 and other data) was used with *igrp* to insert a spoofed route to 222.222.222.0/24. According to FX, IGRP is not used much currently, but the example certainly is interesting. Therefore, at risk of being slightly out of format with the rest of this paper, I will include his test results:

```
test# ./ass -mA -i eth0 -D 192.168.1.10 -b15 -v
ASS [Autonomous System Scanner] $Revision: 2.14 $
      (c) 2k FX <fx@phenoelit.de>
      Phenoelit (http://www.phenoelit.de)
No protocols selected; scanning all
Running scan with:
```

© SANS Institute 2000 - 2002 As part of GIAC practical repository. Author retains full rights.

```
        interface eth0
        Autonomous systems 0 to 15
        delay is 1
        in ACTIVE mode

Building target list ...
192.168.1.10 is alive
Scanning ...
Scanning IGRP on 192.168.1.10
Scanning IRDP on 192.168.1.10
Scanning RIPv1 on 192.168.1.10
shutdown ...

>>>>>>>>>>>> Results >>>>>>>>>>>>
192.168.1.10
  IGRP
        #AS 00010       10.0.0.0        (50000,1111111,1476,255,1,0)
  IRDP
        192.168.1.10    (1800,0)
        192.168.9.99    (1800,0)
  RIPv1
        10.0.0.0        (1)


test# ./igrp -i eth0 -f routes.txt -a 10 -S 192.168.1.254 -D 192.168.1.10

routes.txt:
# Format
# destination:delay:bandwith:mtu:reliability:load:hopcount
222.222.222.0:500:1:1500:255:1:0

Cisco#sh ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, * - candidate default
       U - per-user static route

Gateway of last resort is not set

     10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C       10.1.2.0/30 is directly connected, Tunnel0
S       10.0.0.0/8 is directly connected, Tunnel0
C    192.168.9.0/24 is directly connected, Ethernet0
C    192.168.1.0/24 is directly connected, Ethernet0
I    222.222.222.0/24 [100/1600] via 192.168.1.254, 00:00:05, Ethernet0
```

(Thanks for FX for this information).

Rprobe & *srip*, along with an excellent RIP spoofing tutorial (written by humble), can be found at http://www.technotronic.com/horizon/ripar.txt. Rprobe is a utility that will request a copy of a RIP routing table from a routing daemon. Tcpdump or any other sniffer can then be used to obtain the results. Next, srip can be used to send a spoofed RIPv1 or RIPv2 message from any source IP. Srip can insert new routes and deactivate current routes, as long as the attacker/penetration tester knows what parameters to use in the command line. An example of the use of these tools is found in Hacking Exposed, second edition, in the Network Devices section.

© SANS Institute 2000 - 2002, Author retains full rights.

Routed, gated, zebra, mrt, and gasp are some other tools that could be used by an attacker or by the penetration tester to work with routing protocols. Going into detail on all of these tools is beyond the scope of this document.

Routing Information Protocol (RIP)

The Routing Information Protocol (RIP) is a distance vector based routing protocol. All routing decisions are based on the number of hops. An Autonomous System (AS) is the overall administrative entity comprised of hosts, routers, and other network devices. RIP is known as an Interior Gateway Protocol (IGP) since it only works within a specific AS. RIP is not a good choice for large networks because it only supports 15 hops. RIPv1 only communicates routing information relative to itself, whereas RIPv2 can communicate the knowledge of other routers. RIP can work with other routing protocols, and according to Cisco, it is often used in conjunction with OSPF, even though some documents indicate that OSPF is the IGP that should replace RIP. Routing updates delivered via RIP may be redistributed through another routing protocol. If an attacker were to spoof routes with RIP into a network that then redistributed the route through another protocol such as OSPF or BGP without verification, the scope of attack could possibly be extended.

RIP Vulnerabilities and Countermeasures

An auditor (or an attacker) could determine the use of RIP by checking for UDP port 520 through the use of *nmap*. In this example, the port is open without any access lists or any type of filtering-

```
[root@premis]# nmap -sU -p 520 -v router.ip.address.2
interesting ports on  (router.ip.address..2):
Port       State      Service
520/udp    open       route
```

Scans for UDP 520 are listed as number seven in the "Top 10 Target Ports" on the http://www.dshield.org/ web site. This indicates that many people are scanning for RIP, perhaps due to the increased availability of routing attack tools such as those mentioned previously.

RIPv1 is inherently insecure since it has no authentication mechanism and uses the unreliable UDP protocol as a transport. RIPv2 includes an option to set an up to 16-character clear text password string (which could obviously be sniffed) or an MD5 signature. The use of an MD5 signature would obviously make spoofing a much more difficult operation, although evidently RIP packets can be easily spoofed. One likely tool for doing this is the nemesis project's RIP command, *nemesis-rip*. Due to the numerous command line options and prerequisite knowledge, it is unlikely that *nemesis-rip* would be a tool used by script kiddies. Mounting an effective RIP spoof or other attack would still take some degree of knowledge if one were to use *nemesis-rip*. A RIP spoofing

© SANS Institute 2000 - 2002 As part of GIAC practical repository. Author retains full rights.

attack is made easier by tools mentioned in Chapter 10: Network Devices of "Hacking Exposed" second edition. These tools are *rprobe* to obtain a remote networks RIP routing table, standard *tcpdump* (or other sniffer) to view the routing table, *srip* to spoof a RIP packet (v1 or V2), *fragrouter* to redirect routing through our evilhost, and a tool like *dsniff* to collect clear text passwords or other traffic.

Despite the relative ease of spoofing, my research has shown that several very large network providers rely on RIP for some of their routing functions. It is unknown if these network providers have a secure implementation or not. RIP is obviously still in use, but hopefully less people are using RIPv1, and are instead using v2 with it's MD5 security mechanisms in place, or have migrated to OSPF with MD5 authentication.

Border Gateway Protocol (BGP)

BGP is an Exterior Gateway Protocol (EGP) which performs routing between AS's. As of 1998, BGP4 was the most recent standard. There are several types of messages that BGP uses, and the most important one for the sake of this paper is the UPDATE message, which contains routing table update information. A large portion of the global Internet relies upon BGP, and therefore any security problems should be taken seriously. Evidently, the claim by the L0pht several years ago that they could take down the whole Internet in a short time was based around weaknesses in routing protocol security such as BGP.

BGP Vulnerabilities and countermeasures

BGP uses TCP port 179 for communication, therefore an nmap probe of TCP port 179 may indicate the presence of BGP-

```
[root@premis]# nmap -sS -p 179 -v router.ip.address.2
Interesting ports on  (router.ip.address..2):
Port       State      Service
179/tcp    open       bgp
```

-An open BGP port. More vulnerable to attack.

```
[root@premis netw3]# nmap -sS -n -p 179 router.ip.address.6
Interesting ports on  (router.ip.address.6):
Port       State      Service
179/tcp    filtered   bgp
```

A BGP port that is filtered. More resistant to attack.

Since BGP uses TCP for it's transport, this opens up BGP to many of the problems that TCP faces such as SYN flood, sequence number prediction, DOS conditions, and possible advertisement of bad routes (Rauch, Black Hat, Asia 2000). BGP

© SANS Institute 2000 - 2002 As part of GIAC practical repository. Author retains full rights.

does not use it's own sequencing, but relies instead upon TCP sequence numbers. Therefore, if the device in question has a predictable sequence number scheme, there may be an avenue of attack although this is unlikely since the majority of the routers running the Internet are Cisco equipment that do not use predictable sequence numbers.

Some implementations of BGP do not use any authentication by default. Others may use cleartext passwords that are subject to the same problems as RIP. If the authentication scheme is weak, this increases the remote chance that an attacker could send an UPDATE message that would modify routing tables, leading to the types of attacks previously stated.

BGP may propagate spoofed route information in the event that an attacker was able to modify or insert routing packets from a protocol such as RIP that BGP then redistributes. This is more of a flaw in the trust model, and not in the protocol itself. BGP's community configuration may also allow some types of attacks, since it appears that the community name is used in some cases as a trust token that can be obtained. An attack on BGP through its underlying transport protocol (TCP) appears to be difficult, because sessions tend to communicate over a single physical wire between peers. A TCP insertion attack is more likely in an environment where two AS's are connected through a switch. In such a network, an intruder in the same VLAN or with the ability to sniff traffic on the switch (possibly through an ARP spoofing attack using the *dsniff* tools) could intercept traffic, monitor TCP sequence numbers, inject modified packets, and/or hijack connections with a tool such as *hunt*. This type of attack has been demonstrated in a lab environment, but seems unlikely to be something we will see in the wild due to it's complex nature.

Applying access lists to filter port 179, using MD5 authentication, using a secure transport medium for BGP communications, and performing route filtering (see http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/12cgcr/np1_c/1cprt1/1cbgp.htm#40309) are some suggested precautions that can be taken, in addition to the standard router security settings such as egress/ingress filtering, etc. Secure BGP (SBGP) is being proposed, but does not appear to have any degree of implementation yet.

Open Shortest Path First (OSPF)

OSPF is a dynamic link-state routing protocol that keeps a map of the entire network and uses the information in this map to determine the shortest path between network points. An area is a grouping of hosts, routers and other network devices that connect to each other. Each area has it's own link-state database.

OSPF communicates by flooding a network with Link State Advertisements (LSA) which describes the status of one or more network links. Each router participating in that network receives these LSA messages. Other routers are found and communications maintained through the use of Hello packets that are generated every ten

© SANS Institute 2000 - 2002 As part of GIAC practical repository. Author retains full rights.

seconds and sent to 224.0.0.5. An OSPF hello packet header, sniffed with *iptraf*, appears as follows:

```
OSPF hlo (a=3479025376 r=192.168.19.35) (64 bytes) from 192.168.253.67
to 224.0.0.5 on eth0
```

A border router, 192.168.253.67, has sent a helo packet to multicast (224.0.0.5) which tells other routers and hosts that it knows how to contact area a (a= 3479025376) from 192.168.19.35.

Once a router receives a Hello packet, it begins a process to synchronize its database with other routers.

An LSA header is composed of the following elements: LS age, option, LS type, Link state ID, Advertising Router ID, LS sequence number, LS checksum, and length.

OSPF Vulnerabilities and Countermeasures

OSPF uses protocol type 89, therefore the presence of OSPF can be determined through an *nmap* protocol scan unless a network is configured to not respond to these types of queries through proper configuration of access lists.

root@premis security]# nmap -sO –router.ip.address.252
Interesting protocols on (router.ip.address.252):
Protocol State Name
89 open ospfigp

OSPF is inherently more secure than RIP, featuring several built in security mechanisms. However, various elements of an LSA can be modified by intercepting and re-injecting OSPF packets. The JiNao team developed a Linux implementation of FreeBSD's divert sockets and used this in their tests.

OSPF can be configured to use no authentication, text-based password authentication, or MD5. If an attacker gained the correct level of access, they could use a tool such as dsniff to monitor OSPF packets and obtain the cleartext password. Alternatively, an attacker could be running divert sockets or possibly any one of the various ARP spoofing tools that redirect traffic. Numerous tools exist to create a variety of dangerous scenarios.

The JiNao team developed and implemented four OSPF attacks. These are basically DOS attacks but may have other applications if other elements of the packets are changed. In brief:

Max Age attack	The maximum age of a LSA is one hour (3600). Attacker sends LSA packets with maxage set. The original router that sent this LSA then contests the

© SANS Institute 2000 - 2002 As part of GIAC practical repository. Author retains full rights.

	sudden change in age by generating a refresh message in a process called "fight-back". Attacker continually interjects packets with the maxage value for a given routing entity which causes network confusion and may contribute to a DOS condition.
Sequence++ attack	Attacker continually injects a larger LSA sequence number, which indicates to the network that it has a fresher route. The original router contests this in the "fight back" process by sending it's own LSA with an even newer sequence number than the attackers sequence number. This creates an unstable network and could contribute to a DOS condition.
Max Sequence attack	The maximum sequence number 0x7FFFFFFF is injected by an attacker. Attackers router then appears to be the freshest route. This creates the same "fight-back" condition from the original router – IN THEORY. In practice, they found that in some cases, the MaxSeq LSA is not purged and remains in the link state database for one hour, giving an attacker control for that time period.
Bogus LSA attack	Refers to a bug in an implementation of gated. This attack crashed gated and required that all gated processes be stopped and restarted to purge the bad LSA, thereby causing a DOS condition. This attack may not affect hardware routers and is most likely fixed in more recent versions of gated.

In a test lab environment, theses attacks were successfully used to force OSPF to change routes by changing the link cost, thereby redirecting all network traffic through a specific host/router of choice. Evidently, these types of attacks have not yet been seen in the wild, and may never be since there are so many other easier-to-exploit security holes in an average network.

These attacks, and others, could possibly be delivered by nemesis-ospf. However, due to it's complexity, nemesis-ospf is hardly a tool that script kiddies will use. The number of options is truly staggering and seems to require a detailed knowledge of OSPF that many attackers and network administrators will not have. It is likely that only skilled network engineers and those who work with WAN equipment would know enough to really put nemesis-ospf to use. Some reports indicate that nemesis-ospf does not always work properly, so this tool may be of limited value.

OSPF authentication requires a key, which evidently needs to be passed back and forth each time a router authenticates itself to another router and attempts to pass OSPF messages. Router hello packets are configured by default to pass between routers every 10 seconds, which gives an attacker many opportunities to sniff the key. If an attacker was able to sniff a network and obtain the key, OSPF packets could possibly be forged, especially if the packets were redirected instead of just blindly spoofed. Such an attack is

© SANS Institute 2000 - 2002 As part of GIAC practical repository. Author retains full rights.

probably difficult and unlikely, especially since many other security flaws will most likely exist in most networks.

It is suggested not to run routing on hosts that don't need dynamic routing. Most hosts can function just fine with static routes. Dynamic routing protocols could open up hosts to attack. For instance, several years ago the *gated* software was found to have an authentication problem in some settings where it accepted all 1's in the authentication header. Evidently, this has been fixed since then. While evidence shows that most intruders that we are aware of at this time do not scan for routing protocols other than RIP, the possibility of targeted router attacks and attack databases does hopefully encourage network operators to secure their infrastructure.

CDP & IRDP attacks using IRPAS

It is my opinion that IRPAS and tools like it may be future agents to cause some degree of chaos in the Internet, including successful system penetrations and network compromise. It is hoped that the existence of such tools will also encourage people to take infrastructure security more seriously. The *cdp* program can be used within a local network segment to perform a denial of service attack on some Cisco routers, causing the routers to reboot and/or crash when the devices are flooded with garbage characters. It can also be used to spoof, which could open the door to other dangerous applications. Please see the examples posted on the IRPAS web page at http://www.phenoelit.de/irpas/docu.html.

One possible attack scenario would use the *cdp* tool to take a router out of service, then the *irdp* and *irdresponder* tools to send notification of a new route with a higher numeric preference value. If the targeted routers could not communicate with the router that had been crashed by a DOS attack, the new route with a higher preference value would then be used instead. If an attack of this nature were to succeed, an attacker could then insert their system in the traffic path relatively easily.

This type of attack could also be used to affect certain hosts that are configured to use IRDP. Windows 98 is configured to use IRDP by default. Windows NT must be manually configured to support an IRDP environment, and will broadcast three ICMP Router Solicitation messages at boot time. A vulnerability in IRDP implementations was found in various Windows and Sun machines by the L0pht, who released their security advisory August 11, 1999. http://www.l0pht.com/advisories/rdp.txt

The router solicitation message does not appear to have any type of authentication other than some very basic criteria that evidently are met in the *irdp* and *irdresponder* tools. These criteria are (from RFC1256): - "IP Source Address is either 0 or the address of a neighbor (i.e., an address that matches one of the router's own addresses on the arrival interface, under the subnet mask associated with that address.) - ICMP Checksum is valid. - ICMP Code is 0. - ICMP length (derived from the IP length) is 8 or more octets". In today's day and age, a non-authenticated protocol is a dangerous thing.

© SANS Institute 2000 - 2002 As part of GIAC practical repository. Author retains full rights.

Summary

Computer networks such as the Internet are very dependent upon routing protocols for proper operation. Routing protocol attacks have not been explored as thoroughly as IP based attacks, but this is obviously changing since tools such as *nemesis* and *irpas* are starting to appear. Other tools for infrastructure protection such as the *JiNao* IDS are also starting to appear but are yet to be widely deployed. Infrastructure vulnerabilities due to misconfiguration or protocol weaknesses can severely affect network security on all levels, therefore it is vitally important for network engineers to receive the time and the training to properly implement security measures when designing or maintaining networks.

References:

Thanks to the following individuals:

Batz, FX, Sebastien Barbereau, Feiyi Wang of the MCNC, J. Oquendo

Antionline collection. Commonly known router exploits. URL:
http://www.antionline.com/cgi-bin/anticode/anticode.pl?dir=router-exploits

Cisco Systems. "Improving Security on Cisco Routers". URL:
http://www.cisco.com/warp/public/707/21.html

Convery, Sean (CCIE #4232) and Trudel, Bernie (CCIE #1884).
"SAFE: A Security Blueprint for Enterprise Networks". URL:
http://www.cisco.com/warp/public/cc/so/cuso/epso/sqfr/safe_wp.htm

Frank Jou, Y. "Scalable Intrusion Detection for the Emerging Network Infrastructure". URL:
http://www.anr.mcnc.org/projects/JiNao/JiNao.html

Prue, Walt. "Re: Some abuse detection hacks". NANOG list. (Mon, 9 Mar 1998) URL:
http://www.cctec.com/maillists/nanog/historical/9803/msg00035.html

"Divert Sockets for Linux". URL:
http://www.anr.mcnc.org/~divert/

Obecian. "The nemesis packet injection tool-suite". URL:
http://www.packetninja.net/nemesis

FX. "IRPAS – Internetwork Routing Protocol Attack Suite". URL:
http://www.phenoelit.de/irpas/

Humble. "Spoofing RIP (Routing Information Protocol)". URL:
http://www.technotronic.com/horizon/ripar.txt

Cisco Press. "Routing Information Protocol". (8 Dec 1999). URL:
http://www.cisco.com/cpress/cc/td/cpress/fund/ith2nd/it2444.htm

Rekhter, Y. "A Border Gateway Protocol 4 (BGP-4)". Request for Comments 1771. (Mar 1995). URL:

© SANS Institute 2000 - 2002 As part of GIAC practical repository. Author retains full rights.

http://www.isi.edu/in-notes/rcf1771.txt

Moy, J. "OSPF Version 2". Request for Comments 1583. (March 1994). URL:
http://www.isi.edu/in-notes/rfc1583.txt

Cisco Press. "Designing & Implementing an OSPF Network". (2 Aug 2000). URL:
http://www.cisco.com/cpress/cc/td/cpress/design/ospf/on0407.htm - xtocid1636554

Cisco Press. "RIP and OSPF redistribution". (12 May 2000). URL:
http://www.cisco.com/univercd/cc/td/doc/cisintwk/ics/cs001.htm

Grefer, Roland. Re: "Anyone know what IP protocol #54 is?". SANS Institute Global Incident Analysis Center, Detects Analyzed 11/10/00. (10 Nov 2000). URL:
http://www.sans.org/y2k/111000.htm

IANA. "Protocol numbers". URL:
http://www.isi.edu/in-notes/iana/assignments/protocol-numbers

Ahmad, Dave & Rauch, Jeremey. "Routers, Switches & more: The glue that binds them all together" Black Hat Briefings 200, Las Vegas USA. (26 July 2000). URL:
http://www.blackhat.com/html/bh-multi-media-archives.html

Batz. "Security Issues Affecting Internet Transit Points and Backbone Providers". Black Hat Briefings 1999, Las Vegas. (7-8 July 1999). URL:
http://www.blackhat.com/html/bh-multi-media-archives.html

Oquendo, J. "Theories in DoS". URL:
http://www.antioffline.com/TID/

Silicosis. "L0pht security advisory". (11 August 1999). URL:
http://www.l0pht.com/advisories/rdp.txt

© SANS Institute 2000 - 2002 Author retains full rights.