

Network Working Group
Internet-Draft
Expires: August 22, 2003

D. Beard
Nortel Networks
S. Murphy
Network Associates, Inc
Y. Yang
Cisco Systems
February 21, 2003

Generic Threats to Routing Protocols
draft-beard-rpsec-routing-threats-01.txt

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on August 22, 2003.

Copyright Notice

Copyright (C) The Internet Society (2003). All Rights Reserved.

Abstract

Routing protocols are subject to attacks that can harm individual users or the network operations as a whole. The lack of a common set of security requirements has led to the use in existing routing protocol of a variety of different security solutions, which provide various levels of security coverage.

The RPSEC working group intends to deliver in a separate document a set of security requirements for consideration of routing protocol designers. The first step in developing the security requirements is

Internet-Draft Generic Threats to Routing Protocols February 2003

to analyze the threats that face routing protocols. This document describes the threats, including threat sources and capabilities, threat actions, and threat consequences as well as a breakdown of routing functions that might be separately attacked.

Table of Contents

1.	Introduction	3
2.	Routing Functions Overview	4
2.1	Targeted Functions	4
3.	Threat Definitions	6
3.1	Threat Sources	6
3.2	Threat Actions	7
3.3	Threat Consequences	8
3.3.1	Threat Consequence Zone	11
3.3.2	Threat Consequence Periods	11
4.	Generally Identifiable Routing Threats Actions	12
4.1	Deliberate Exposure	12
4.2	Sniffing	12
4.3	Traffic Analysis	13
4.4	Spoofing	13
4.5	Falsification	15
4.5.1	Falsifications by Originators	15
4.5.2	Falsifications by Forwarders	21
4.6	Interference	22
4.7	Overload	23
4.8	Byzantine Failures	23
4.9	Discarding of Control Packets	24
4.10	Network Mapping Threats	25
5.	Multicast Routing Protocol Considerations	26
6.	Security Considerations	28
	References	29
	Authors' Addresses	29
A.	Acknowledgements	31
	Intellectual Property and Copyright Statements	32

Internet-Draft Generic Threats to Routing Protocols February 2003

1. Introduction

The RPSEC working group is tasked to deliver a description of the security requirements for routing protocols. This internet draft discusses an analysis of the threats that face routing protocols, as a precursor to developing a common set of security requirements for routing protocols. Therefore, we intentionally do not address threats to routers (hacking, denial of service flooding attacks, etc.) or to specific routing protocol implementations (bugs, etc.). The security requirements derived from this threat analysis are intended to be guidance to those who are designing routing protocols.

[2.](#) Routing Functions Overview

Routing protocols in general have several common functions:

- o **Transport Subsystem:** The routing protocol transmits messages to its peers using some underlying protocol. For some, as in OSPF, this is IP. For others, this can be a broadcast link layer, as in AODV. Still others may run over TCP. In many cases, the routing protocol is subject to attacks on its underlying protocol.
- o **Neighbor State Maintenance:** Each protocol has a different mechanism for determining its peers in the routing topology. Some protocols have distinct exchange through which they establish peering relationships, e.g., Hello exchanges in OSPF. The peering relationship formation is the first step of topology determination. For protocols that maintain state about their peering relationships, attacks that disrupt the peering relationship can have widespread consequences. For example, if the DR election is disrupted in an OSPF network, an unauthorized router could be chosen as designated router. This might allow unauthorized access to routing information. In BGP, if a router receives a CEASE message, it can break the peering relationship and cause any related topology information to be flushed.
- o **Database Maintenance:** Routing protocols exchange network topology and reachability information. The routers collect this information in routing databases in varying detail. The

maintenance of these databases is a significant portion of the function of a routing protocol. The information in the database must be authentic and authorized; otherwise the function of routing in the overall network is damaged. For example, if an OSPF router sends LSA's with the wrong Advertising Router, the receivers will compute a SPF tree that is incorrect and might not forward the traffic. If a BGP router advertises a NLRI that it is not authorized to advertise, then receivers might forward that NLRI's traffic toward that router and the traffic would not be deliverable. A PIM router might transmit a JOIN message to receive multicast data it would otherwise not receive

[2.1](#) Targeted Functions

Just as a router's functions can be divided into control and data plane (protocol traffic vs. data traffic), so the routing protocol has a control and a data plane. A routing protocol has some message exchanges that are intended only for control of the protocol state. This is the routing protocol control plane. Other message exchanges are intended to distribute the information used to perform the

forwarding function, whether that is to establish a forwarding table in each router or to return a description of the route to use. This is the routing protocol data plane. Each of the routing functions may have both control and data aspects, but there will naturally be an emphasis on one or the other. Neighbor maintenance is likely to be focused on the routing protocol control plane aspects, for example, while database maintenance may have more focus on the routing protocol data plane aspects.

Both the control and the data plane are subject to attack. An attacker who is able to target the routing protocol control plane so as to break a neighbor (e.g., peering, adjacency) relationship can have a strong effect on the behavior of routing in those routers and likely the surrounding neighborhood. An attacker who is able to break a database exchange between two routers can also affect routing behavior. In the routing protocol data plane, an attacker who is able to introduce bogus data can have a strong effect on the behavior of routing in the neighborhood.

[3.](#) Threat Definitions

Threat is defined in [[SEC-GLOSS](#)] as a potential for violation of security, which exists when there is a circumstance, capability, action, or event that could breach security and cause harm. A threat presents itself when an attacker has the ability to take advantage of an existing security weakness. Threats can be categorized based on various rules, such as threat sources, threat actions, threat consequences, threat consequence zones, and threat consequence periods.

[3.1](#) Threat Sources

Legitimate devices (routers) participate in the routing dialog and computation, intended by the authoritative network administrator, running correct and bug-free code, and using correct and bug-free configuration information. -- By correct and bug-free configuration information, we mean the configurations obey routing protocols and are intended by the authoritative network administrator.

On the other hand, attackers may participate routing, not being authorized, running incorrect codes, or using invalid configurations. In general, attackers can be outsiders or insiders. An insider is an authorized participant in the routing protocol. An outsider is any other host or network. A host is determined to be an outsider or an insider from the point of view of a particular router. Even an authorized protocol speaker can be an outsider to a particular router if the router does not consider the speaker to be a legitimate peer (as could conceivably happen on a multi-access link).

Specifically, threats can be classified into four categories, based on their sources [[DV-SECURITY](#)]:

- o Threat from compromised links: A compromised link is where an attacker can, somehow, access a physical medium and/or have some control over the channel. This threat exists when there is no access control mechanisms applied to physical mediums or channels, or such mechanisms can be circumvented. The attacker may eavesdrop, replay, delay, or drop routing messages, or break routing sessions between authorized routers, without participating in the routing exchange.
- o Threats from compromised devices (e.g. routers): A compromised device (router) is an authorized router with routing software bugs, hardware defects, and / or incorrect/unintended configurations. This threat takes place when there are no mechanisms to verify a device's (router) system integrity, i.e. the router is working correctly as been intended by the

authoritative network administrator, or such mechanisms can be circumvented. The attacker may inappropriately claim authority for some network resources, or violate routing protocols, such as advertising invalid routing information and etc.

- o Threat from unauthorized devices (routers): An unauthorized device

(router) participates in routing exchange and computation, without being authorized (explicitly or implicitly) from the authoritative network administrator. This threat happens when there is no access control mechanism applied to routing sessions/routing exchanges or such mechanism can be circumvented. The attacker may gain knowledge of the network topology through routing exchange, as well as do anything that a compromised router can do.

- o Threat from masquerading devices (routers): A masquerading device (router) illegitimately assumes another router's identity. This threat occurs when there are no (data origin or peer entity) authentication mechanisms, or such mechanisms can be circumvented. The attacker can do anything that an unauthorized router can do.

A device (router) can play multiple roles concurrently. A legitimate OSPF router might be a masquerading RIP router, and a compromised iBGP link might be a compromised OSPF router as well.

3.2 Threat Actions

A threat action is an assault on system security [[SEC-GLOSS](#)], which could be an intentional behavior, or an accidental event.

The actions that might be used to attack routing protocols include:

- o Masquerade: The attacker, whether insider or outsider, may adopt the identity of a legitimate peer. (This is an attack against origin authenticity.)
- o Interception: The attacker gains access to routing information that is considered sensitive. (This is an attack against confidentiality, i.e., privacy.)
- o Falsification: The attacker is able to substitute modified messages for valid routing messages. (This is an attack against integrity.)
- o Misuse: The attacker is able to introduce unauthorized routing information that disrupts routing behavior. (This is an attack against authorized use.)
- o Replay The attacker is able to re-introduce previously transmitted

messages. (This is an attack against freshness.)

These attacks might be used by insider or outsider to accomplish any of the compromises listed below.


[3.3](#) Threat Consequences

A threat consequence is a security violation that results from a threat action [[SEC-GLOSS](#)]. The compromise to the behavior of the routing system can damage a particular network or host or can damage the operation of the network as a whole.

Four types of threat consequences, disclosure, deception, disruption, and usurpation, are identified in [[SEC-GLOSS](#)]. Specifically for threats against routing protocols, these consequences can be described as:

- o Disclosure: Disclosure of routing information happens where a router successfully accesses the information without being authorized. Compromised links can cause disclosure, if routing exchanges lack confidentiality. Compromised devices (routers), unauthorized devices (routers), and masquerading devices (routers) can always cause disclosure, as long as they are successfully involved in the routing exchanges. Please note, although disclosure of routing information can pose a security threat or be part of a later, larger, or higher layer attack, confidentiality is not generally a design goal of routing protocols.
- o Deception: This consequence happens when a legitimate router receives a false routing message and believes it to be true. All attackers (Compromised links, compromised device (routers), unauthorized devices (routers), and masquerading devices (routers) can cause this consequence if the receiving router lacks ability to check routing message integrity, routing message origin authentication or peer router authentication.
- o Disruption: This consequence occurs when a legitimate router's operation is being interrupted or prevented. Subvert links can cause this by replaying, delaying, or dropping routing messages, or breaking routing sessions between legitimate routers. Compromised devices (router), unauthorized devices (routers), and masquerading device (routers) can cause this consequence by sending false routing messages, interfering normal routing exchanges, or flooding unnecessary messages. (DoS is a common threat action causing disruption.)
- o Usurpation: This consequence happens when an attacker gains control over a legitimate router's services/functions. Compromised

Note: an attacker does not have to directly control a router to control its services. For example, in Figure 1, Network 1 is dual-homed through Router A and Router B, and Router A is preferred. However, Router B is compromised and advertises a lower metric. Consequently, devices on the Internet choose the path through Router B to reach Network 1. In this way, Router B steals the data traffic and Router A surrenders its control of the services to Router B. This is depicted in Figure 1.



Also, several threat consequences might be caused by a single threat action. In Figure 1, there exist at least two consequences: routers using Router B to reach Network 1 are deceived, while Router A is usurped.

- o Network congestion: more data traffic is forwarded through some portion of the network than would otherwise need to carry the

traffic,

- o Blackhole: large amounts of traffic are directed to be forwarded through one router that cannot handle the increased level of traffic and drops many/most/all packets,

- o Looping: data traffic is forwarded along a route that loops, so that the data is never delivered (resulting in network congestion),
- o Partition: some portion of the network believes that it is partitioned from the rest of the network when it is not,
- o Churn: the forwarding in the network changes (unnecessarily) at a rapid pace, resulting in large variations in the data delivery patterns (and adversely affecting congestion control techniques),
- o Instability: the protocol becomes unstable so that convergence on a global forwarding state is not achieved, and
- o Overload: the protocol messages themselves become a significant portion of the traffic the network carries.

The damage that might result from attacks against a particular host or network address may include:

- o Starvation: data traffic destined for the network or host is forwarded to a part of the network that cannot deliver it,
- o Eavesdrop: data traffic is forwarded through some router or network that would otherwise not see the traffic, affording an opportunity to see the data or at least the data delivery pattern,
- o Cut: some portion of the network believes that it has no route to the host or network when it is in fact connected,
- o Delay: data traffic destined for the network or host is forwarded along a route that is in some way inferior to the route it would otherwise take,
- o Looping: data traffic for the network or host is forwarded along a route that loops, so that the data is never delivered,

It is important to consider all compromises, because some security solutions can protect against one attack but not against others. It might be possible to design a security solution that protected against an attack that eavesdropped on one destination's traffic without protecting against an attack that overwhelmed a router. Or that prevented a starvation attack against one host, but not against a net wide blackhole. The security requirements must be clear as to which compromises are being avoided and which must be addressed by other means (e.g., by administrative means outside the protocol).

[3.3.1](#) Threat Consequence Zone

A threat consequence zone covers an area within which the network operations have been affected by the threat consequences. Possible threat consequence zones can be classified as: a single link or router, multiple routers (within a single routing domain), a single routing domain, multiple routing domains, or the global Internet. The threat consequence zone varies based on the threat action and origin. Similar threat actions that happened at different locations may cause totally different threat consequence zones. For example, when a compromised link breaks the routing session between a distribution router and a stub router, only reach ability from and to the network devices attached on the stub router will be impaired. In other words, the threat consequence zone is a single router. Nonetheless, if the compromised router is located between a customer edge router and its corresponding provider edge router, such an action might cause the whole customer site to lose its connection. In this case, the threat consequence zone might be a single routing domain.

[3.3.2](#) Threat Consequence Periods

Threat consequence period is defined as a portion of time during which the network operations have been impacted by the threat consequences. The threat consequence period is influenced by, but not totally dependent on the duration of the threat action. In some cases, the network operations will get back to normal as soon as the threat action has been stopped. In other cases, however, threat consequences may appear longer than threat action. For example, in the original ARPANET link-state algorithm, some errors in a router

might introduce three instances of an LSA, and all of them would be flooded throughout the network forever, until the entire network was power cycled [[PROTO-VULN](#)].

With appropriate security detection facilities, the network might detect the threat action, implement countermeasures, and resume normal operations even before the threat action has been stopped. In this documentation, we assume such facilities do not exist.

[4.](#) Generally Identifiable Routing Threats Actions

This section addresses generally identifiable and recognized threat action against routing protocols. The threats are not necessarily specific to individual protocols but may be present in one or more of the common routing protocols in use today.

[4.1](#) Deliberate Exposure

Deliberate Exposure is defined as an intentional action that attackers employ to release routing information directly to other routers. This definition presumes that the receiving routers are not authorized to access the routing information. However, an exposure is different from a deliberate exposure. While the deliberate exposure is always a threat action, the exposure is not. Routing protocols are designed to expose routing information. A legitimate router should always expose routing information to its legitimate peers. In some cases, a legitimate router may expose routing information to peering unauthorized/masquerading routers, if it is deceived. However, there is no reason that a legitimate router should keep exposing correct routing information to its peers when those peers have been determined to be unauthorized or masquerading entities.

The consequence of deliberate exposure is the disclosure of routing information.

The threat consequence zone of deliberate exposure depends on the routing information that the attackers have exposed. The more knowledge they have exposed, the bigger the threat consequence zone.

The threat consequence period of deliberate exposure might be longer than the duration of the action itself. The routing information exposed will not be out-dated until there is a topology change of the exposed network.

[4.2 Sniffing](#)

Sniffing is an action whereby attackers monitor and/or record the routing exchanges between authorized routers. Compromised links can sniff the links over which they have control. (Compromised routers, unauthorized routers, and masquerading routers can sniff, but do not need to do this, to access the routing information. They can learn the routing information as long as they are successfully involved in the routing exchanges).

The consequence of sniffing is disclosure of routing information.

The threat consequence zone of sniffing depends on the attacker's

location, the routing protocol type, and, ultimately, what routing information has been recorded. For example, if the compromised link were located in an OSPF totally stubby area, the threat consequence zone should be limited to the whole area. Or, the compromised link could gain knowledge of multiple routing domains, if it sniffs an eBGP session between two providers.

The threat consequence period might be longer than the duration of the action. After the compromised link stops sniffing, its knowledge will not be out-dated until there is a topology change of the disclosed network.

[4.3 Traffic Analysis](#)

Traffic analysis is action whereby attackers gain routing information

by analyzing the characteristics of the data traffic. Compromised links can analyze the data traffic over the links where they have control. (Compromised routers, unauthorized routers, and masquerading routers do not need to do this, although they can, to access the routing information. They learn the routing information by being successfully involved in the routing exchanges).

The consequence of data traffic analysis is the disclosure of routing information. For example, the source and destination IP address of the data traffic, the type, magnitude, and volume of traffic is disclosed.

The threat consequence zone of the traffic analysis depends on the attacker's location and, ultimately, what data traffic has flown through. A compromised link at the network core should be able to gain more information than its counterpart at the edge.

The threat consequence period might be longer than the duration of the traffic analysis. After the attacker stops traffic analysis, its knowledge will not be out-dated until there is a topology change of the disclosed network.

[4.4](#) Spoofing

A spoofing is defined as an action whereby an attacker participates in the routing computation and exchanges with authorized routers by illegitimately assumes a legitimate router's identity. All types of attackers (compromised links, compromised routers unauthorized routers, and masquerading routers) can spoof. When an attacker succeeds to spoof, it plays a role of masquerading router.

The consequences of spoofing are:

- o The disclosure of routing information: The masquerading router will be able to participate in the routing computation and exchanges, and consequently gain access to the routing information.
- o The deception of peer relationship: The authorized routers, which exchange routing messages with the masquerading router, do not realize they are peering with a router that is faking another

router's identity.

Spoofing is special in that it can be used to carry out other threat actions causing other threat consequences. For example, after an attacker spoofs successfully, it can send out unrealistic routing information that might cause disruption of network services. Please note these consequences are directly resulted from other threat actions instead of spoofing, which are also discussed in this documentation. It can be said that spoofing is the means by which one masquerades.

The threat consequence zone covers two different scopes:

The consequence zone of the disclosed routing information depends on what routing information has been exchanged between the attacker and its peers.

The disclosure of routing information: The masquerading router will participate in the routing computation and exchanges, and consequently gain access to the routing information.

There are other consequences caused by a spoofing (masquerading) router. For example, the masquerading router might cause disruption of a network by sending unrealistic routing information. But these consequences are directly resulted from other threat actions instead of spoof.

The threat consequence zone covers two different scopes:

- o The consequence zone of the fake peer relationship will be limited to those routers mistrusting the attacker's identity.
- o The consequence zone of the disclosed routing information depends on the attacker's location, the routing protocol type, and, ultimately, what routing information has been exchanged between the attacker and its deceived peers.

The threat consequence period has two different definitions too:

- o The consequence period of the fake peer relationship is same as

the duration of the spoof. As soon as the attacker stops spoofing,

the fake peer relationship disappears.

- o The consequence period of the disclosed routing information will be longer than the duration of the spoof. After the attacker stops spoofing, its knowledge will not be out-dated until there is a topology change of the disclosed network.

[4.5](#) Falsification

Falsification is defined as an intentional action whereby false routing information is being sent. Routers use routing information to depict network topology, compute routing table, and further forward data traffic. False routing information describes the network in an unrealistic view, whether or not intended by the authoritative network administrator.

To falsify the routing information, an attacker has to be either the originator or a forwarder of the routing information. It cannot be a receiver-only.

[4.5.1](#) Falsifications by Originators

An originator of routing information can launch following falsifications:

[4.5.1.1](#) Overclaiming

An over-claiming is defined as an action that an attacker employs to advertise its ownership of some network resources, while in reality, this ownership does not exist, or the advertisement is not authorized. This is given in Figure 2 and Figure 3 below.




Figure 2




Figure 3

The above figures provide examples. Router A, the attacker, is connected with the Internet through Router B. Router C is authorized to advertise its link to Network 1. In Figure 2, Router A owns a link to the Network 1, but is not authorized to advertise it. In Figure 3, Router A does not own such a link. But in either case, Router A advertises the link to the Internet, through Router B.

Compromised routers, unauthorized routers, and masquerading routers can over-claim network resources.

The consequence of overclaiming includes:

- o Usurpation of the overclaimed network resources. In Figure 2 and 3, it will cause a usurpation of Network 1 when Router B or

other routers on the Internet (not shown in the figures) believe that Router A provides the best path to reach the Network 1. They,

the routers, thereby forward the data traffic, destined to Network 1, to Router A. The best result is the data traffic uses an unauthorized path (Figure 2), and the worst case is the data never reach the destination Network 1 (Figure 3). The ultimate consequence is Router A gains the control over the Network 1's services, by controlling the data traffic.

- o Usurpation of the legitimate advertising routers. In Figure 2 and 3, Router C is the legitimate advertiser of Network 1. By overclaiming, Router A also controls (partially or totally) the services/functions provided by the Router C. (This is NOT a disruption, because Router C is operating in a way intended by the authoritative network administrator.)
- o Deception of other routers. In Figure 2 and 3, Router B, or other routers on the Internet, might be deceived to believe the path through Router A is the best.
- o Disruption of data planes on some routers. This might happen on routers that are on the path, which is used by other routers to reach the overclaimed network resources through the attacker. In Figure 2 and 3, when other routers on the Internet are deceived, they will forward the data traffic to Router B, which might be overloaded.

The threat consequence zone varies based on the consequence:

- o Where usurpation is concerned, the consequence zone covers the network resources that are overclaimed by the attacker (Network 1 in Figure 2 and 3), and the routers that are authorized to advertise the network resources but lose the competition against the attacker (Router C in Figure 2 and 3).
- o Where deception is concerned, the consequence zone covers the routers that do not believe the attacker's advertisement and use the attacker to reach the claimed subnets (Router B and other deceived routers on the Internet in Figure 2 and 3).
- o Where disruption is concerned, the consequence zone includes the

routers that are on the path of misdirected data traffic (Router B in Figure 2 and 3).

The threat consequence will cease when the attacker stops overclaiming, and will totally disappear when the routing tables are converged. As a result the consequence period is longer than the duration of the overclaiming.

[4.5.1.2](#) Underclaiming

An underclaiming threat is defined as an action that an attacker illegitimately hides its authorized ownership of some network resources. The attacker could be the only router authorized to claim the network resources, or there might exist some legitimate backup routers. Figures below provide two examples.




Figure 4




Figure 5

Router A, the attacker, owns a link to Network 1 and is authorized to advertise Network 1. Nevertheless, Router A refuses to advertise Network 1. In Figure 4, Network 1 is single-homed with Router A and therefore can only be advertised by Router A. In Figure 5 Network is dual-homed with Router A and B, and both routers are authorized to advertise Network 1 (Router A may or may not provide a preferred path against Router B, the backup router).

Compromised routers, unauthorized routers, and masquerading routers

can underclaim network resources.

The consequence of underclaiming includes:

- o Usurpation of the underclaimed network resources: In Figure 5 when Router A underclaims Network 1, Network 1 is isolated from the rest of the world, and cannot provide services to other devices, though Network 1's own operation is not disrupted. In Figure 4, if the path through Router A is preferred, the underclaiming will force Network 1 to use a sub-optimal path to provide its services. (If the path through Router B is intended to be preferred, the services by Network 1 will not really be hurt even though Router A underclaims).
- o Usurpation of the legitimate backup routers. In Figure 5, Router A's path is preferred but Router A underclaims Network 1, it actually force Router B to serve Network 1. (Again, if Router B's path is intended to be preferred, Router A's underclaim does not really usurp Router B.)
- o Deception of other routers. Routers on the Internet (not shown in Figure 4 or Figure 5) might not be able to reach Network 1 (Figure 5) or have to use a sub-optimal path through Router B when

Router A's path is preferred.

- o Disruption of data planes on some routers. This might happen on routers that are on the sub-optimal paths. In Figure 5, when other routers on the Internet are deceived and use the sub-optimal path through Router B to reach Network 1, they will forward the data traffic to Router C. Router B and C might then become overloaded. (When the path through Router B is intended to be preferred, Router B and C might also be overloaded. However, the disruption in such a case is not a consequence of an underclaim).

Note: Some others type of usurpation might result from an underclaim in routing protocols. Below Figure provides an example.

```
      *---*                      *---*
     /      \                   /      \
    +-----+   +-----+   +-----+   +-----+
```




Figure 6

In Figure 6, Network 2 is attached with the Router B and provides similar services as Network 1. When Router A hides Network 1, devices on the Internet will turn to Network 2 for those services. Although this issue results from an underclaim in routing protocol, this is rather a usurpation issue in related service (application) protocols, and we are not discussing it in detail in this documentation.

The threat consequence zone varies based on the consequence:

- o Where usurpation is concerned, the consequence zone covers the network resources that are underclaimed by the attacker (Network 1 in Figure 4 and 5), and the routers that are intended to be backup with a lower preference (Router B in Figure 5, if Router A's path is preferred).
- o Where deception is concerned, the consequence zone covers the routers that cannot reach the underclaimed network resources or those that have to use sub-optimal paths.
- o Where disruption is concerned, the consequence zone covers the

routers that cannot reach the underclaimed network resources or those that have to use sub-optimal paths.

Like overclaiming, the consequence period is longer than the duration of the underclaiming--the threat consequence will mitigate when the attacker stops underclaiming and will totally disappear when routing tables are converged.

[4.5.1.3](#) Misclaiming

A Misclaiming threat is defined as an attacker action advertising its

authorized ownership of some network resources in a way that is not intended by the authoritative network administrator. An attacker can eulogize or disparage when advertising these network resources. Compromised routers, unauthorized routers, and masquerading routers can misclaim network resources.

The threat consequences of Misclaiming are a combination of consequences from overclaiming and underclaiming. Eulogizing the network resources might cause the same consequences made by overclaiming, while disparaging might trigger the same results from underclaiming.

The consequence zone and period are also similar to those of overclaiming or underclaiming.

[4.5.2](#) Falsifications by Forwarders

When a legitimate router forwards routing information, it must or must not modify the routing information, depending on the routing information and the routing protocol type. For example, in RIP, the forwarder must modify the routing information by increasing the hop count by 1. On the other hand, the forwarder must not modify the type 1 LSA in OSPF. In general, forwarders in distance vector routing protocols are authorized to and must modify the routing information, while most forwarders in link state routing protocols are not authorized to and must not modify most routing information.

As a forwarder authorized to modify routing message, an attacker does not forward necessary routing information to other authorized routers. Unauthorized aggregation (summarization) is special type of understatements.

[4.5.2.1](#) Misstatement

This is defined as an action whereby the attacker describes route attributes in a wrong way. For example, in RIP, the attacker

increases the path cost by two hops instead of one. Another example is, in BGP, the attacker deletes some AS numbers from the AS PATH.

When forwarding routing information that should not be modified, an

attacker can launch the following falsifications:

- o Deletion: Attacker deletes valid data in the routing message.
- o Insertion: Attacker inserts false data in the routing message.
- o Substitution: Attacker replaces valid data in the routing message with false data.
- o Replaying: Attacker replays out-dated data in the routing message.

All types of attackers (Compromised links, compromised routers, unauthorized routers, and masquerading routers) can falsify the routing information when they forward the routing messages.

The threat consequences of these falsifications by forwarders are similar to those caused by originators: Usurpation of some network resources and related routers; deception of routers using false paths; and disruption of data planes of routers on the false paths. The threat consequence area and period are also similar.

[4.6](#) Interference

Interference is defined as a threat action where attackers inhibit exchanges on legitimate routers. Attackers can do this by adding noise, not forwarding packets, replaying out-dated packets, delaying responses, denial of receipts, and breaking synchronization.

Compromised links can interfere with the routing exchanges over the links where they have control. Compromised, unauthorized and masquerading routers can slowdown their routing exchanges or create flapping routing sessions of the legitimate peering routers.

The consequence of interference is the disruption of routing operations.

The consequence zone of interference varies based on the source of the threats:

- o When a compromised link launches the action, the threat consequence zone covers routers that are using the link to exchange the routing information. Routers behind might be disrupted too.

- o When compromised routers, unauthorized routers, or masquerading routers are the attackers, the threat consequence zone covers routers with which the attackers are exchanging routing information, and router behind.
- o The threat consequences might disappear as soon as the interference is stopped, or might not totally disappear until the networks are converged. Therefore, the consequence period is equal or longer than the duration of the interference.

[4.7](#) Overload

Overload is defined as a threat action whereby attackers place excess burden on legitimate routers. Attackers can overload data plane or control plane. Because data plane is involved in routing exchanges, overload of data plane will also influence the routing operations.

The consequence of overload is the disruption of routing operations. The consequence zone varies based on several factors:

- o When compromised links launch an overload action against the control plane, the consequence zone covers routers that are using the links to exchange the routing information, and routers behind.
- o When compromised links launch an overload action against the data plane, the consequence zone covers routers that are physically connected by the links, and routers behind.
- o When Compromised routers, unauthorized routers, or masquerading routers launch an overload action against the control plane, the threat consequence zone covers routers with which the attackers are exchanging routing, and routers behind.
- o When Compromised routers, unauthorized routers, or masquerading routers launch an overload action against the data plane, the threat consequence zone covers of routers with which the attackers have physical connections, and routers behind.

The threat consequences might disappear as soon as the overload is stopped, or not disappear until networks are converged.

[4.8](#) Byzantine Failures

When a host or network behaves in a way contrary to the protocol specification or in a way that is not authorized, the behavior is

called a "Byzantine failure"[[BYZANTINE](#)]. These failures can include

timing error (producing messages at intervals contrary to the specification), protocol errors (producing messages at variance with the specification, e.g., responding with the incorrect message type), or data error (producing messages that carry faulty data).

Byzantine attacks may be seen where any intermediate node or group of nodes can intentionally create routing loops, misrouting packets on non-optimal paths, or selectively dropping packets (black hole). Another way to state the problem is that Byzantine failures occur when a processor returns incorrect or malicious data. Under such an attack, only the source and destination nodes are assumed to be trusted. Detecting a Byzantine error is harder than the fail-stop model in the sense that at least one other processor must do the same computation to confirm the results. What isn't clear is just how much validation is required to determine whether a Byzantine failure has occurred

[4.9](#) Discarding of Control Packets

Similar to Byzantine threats discussed above, uncontrolled discarding of control packets lies in the same plane. That is, discarding of control packets will have the same consequence as an incorrect routing control packet propagated in the network by a compromised router. In distance vector protocols the consequences may not be as dire because of the protocol behavior, i.e. the routing update, is exchanged only with the neighbor. However in the case of link state routing protocols, the threat associated to discarding of control packet can become a serious issue, as the routing updates are flooded in the network. Exploitation of this threat was discussed by S.F. Wu B. Vetter and F. Wang from the perspective of an insider attacks in a Link State Routing environment. It is worth considering this threat in more detail.

If the compromised (bad) router partitions the network, i.e. the router is the only path between two good routers, then the bad router can avoid forwarding the routing information on to the network on the other side.




Figure 7

In this scenario, the network is partitioned and either side may not receive correct updates and the update packets may be dropped. Clearly if F is positioned such that the network is not partitioned, then the correctness of the protocol in such circumstances depends on the mechanism of transmitting routing updates. In the case of a typical LSRP like OSPF, reliable flooding is used that guarantees that the updates are received by each and every router in the network. Hence even when a set of bad routers partition a network, if there exists at least one good path between all the routers then this threat can be deterred by designing a robust transmitting mechanism for control updates.

[4.10](#) Network Mapping Threats

Based on a simple set of inputs, computers can generate graphical and quantitative representations of informal knowledge networks within an organization. If there were no preventive measures in place, network map knowledge obtained by unauthorized access to intelligence can be costly and expensive threats. Motivation for snooping can range from curiosity to voyeur tendencies. The threat with router plane data snooping is the fact that it looks to historical information to be an indication of what will happen in the future. The principal threat

aspect is that the snooped data can be used to develop a network topology. When unauthorized attackers develop a model, they attempt to create one that will be relevant for all situations going forward. Although these models may not be exact for every situation, they can be applied with a reasonable amount of certainty without introducing any biases based on past information.

[5](#). Multicast Routing Protocol Considerations

Based on a simple set of inputs, computers can generate graphical and quantitative representations of informal knowledge networks within an organization. If there were no preventive measures in place, network map knowledge obtained by unauthorized access to intelligence can be costly and expensive threats. Motivation for snooping can range from curiosity to voyeur tendencies. The threat with router plane data snooping is the fact that it looks to historical information to be an indication of what will happen in the future. The principal threat aspect is that the snooped data can be used to develop a network topology. When unauthorized attackers develop a model, they attempt to create one that will be relevant for all situations going forward. Although these models may not be exact for every situation, they can be applied with a reasonable amount of certainty without introducing any biases based on past information.

In general, multicast routing updates can be fabricated, modified, replayed, deleted, and snooped. For example, unauthorized nodes can simply participate in the multicast routing protocol dialog when no access control mechanisms are defined for the protocol. Non-routing devices can masquerade as an authorized router and inject spurious routing updates, perhaps using source routing attacks or TCP session hijacking attacks. Communication links can be compromised by an intruder to facilitate the manipulation of routing messages. Individual routers can be attacked and compromised to run modified

software, or use a modified configuration.

Multicast communication may be specifically targeted by security threats, due to its potential for communicating with large numbers of receivers simultaneously. An attacker may attempt to use multicast sessions in order to spread specific data to recipients, or may use multicast traffic patterns to overload links as a denial-of-service (DOS) attack.

In some architecture such as PIM-DM, even routers which are not actively participating in the multicast tree must maintain state information on active groups within the routing domain.

Multicast routing protocols are at least as susceptible as unicast routing protocols to security threats. In general, multicast routing updates can be fabricated, modified, replayed, deleted, and snooped. For example, unauthorized nodes can simply participate in the multicast routing protocol dialog when no access control mechanisms are defined for the protocol. Non-routing devices can masquerade as an authorized router and inject spurious routing updates, perhaps using source routing attacks or TCP session hijacking attacks. Communication links can be compromised by an intruder to facilitate

the manipulation of routing messages. Individual routers can be attacked and compromised to run modified software, or use a modified configuration.

Just as with unicast routing, the key vulnerabilities of multicast routing lie in the introduction of misleading routing information, through non-existent (black hole) or incorrect routes, or in intercepting the routing information for malicious purposes. Incorrect routing information can form the basis for DOS attacks, while intercepting routing information (particularly group membership information) can reveal compromising topological information.

Denial-of-service attacks may come either from senders or receivers in the multicast model. That is, if uncontrolled, senders may create large numbers of multicast groups, thus potentially creating a processing burden on multicast routers throughout the domain. Receivers, if uncontrolled, may join large numbers of multicast groups, thus causing the establishment of paths from the senders in each group to the receiver, as well as causing the flow of packets

for each of the groups to converge on the receiver.

[6.](#) Security Considerations

This entire informational draft RFC is security related. Specifically it addresses security of routing protocols as associated with threats to those protocols. In a larger context, this work builds upon the recognition of the IETF community that signaling and control/management planes of networked devices need strengthening. Routing protocols can be considered part of that signaling and control plane. However, to date, routing protocols have largely remained unprotected and open to malicious attacks. This document discusses inter and intra domain routing protocol threats as we know them today and lays the foundation for a future draft which fully discusses security

requirements for routing protocols.

References

[SEC-GLOSS] R.Shirey, Internet Security Glossary, [RFC 2828](#), May 2000

[DV-SECURITY] B.R.Smith, S.Murthy, and J.J. Garcia-Luna-Aceves,

Securing Distance-Vector Routing Protocols, Symposium on Network and Distributed System Security 1997, Feb. 1997

[PROTO-VULN] E.Rosen, Vulnerabilities of Network Control Protocols: An Example, Computer Communication Review, Jul. 1981

[BYZANTINE] R.Pearlman, Network Layer Protocols with Byzantine Robustness, August 1988

[OSPF-SIG] S. Murphy, M. Badger, and B. Wellington, OSPF with Digital Signatures, [RFC2154](#), June 1997

[OSPFv2] J.Moy, OSPF Version 2, [RFC 2328](#), April 1998

[SENSOR-IDS] V.Mittal and G.Vigna, Sensor-Based Intrusion Detection for Intra-Domain Distance-Vector Routing, Proceedings of the ACM Conference on Computer and Communication Security (CCS'02), Washington, DC, November 2002

[DOS-IDS] S.Cheung et. al., Protecting Routing Infrastructures from Denial of Service using co-operative intrusion detection, In Proceedings of the 1995 IEEE Symposium on Security and Privacy

[DIST-MONITOR] K.A. Bradley et. al., A distributed Network Monitoring approach

[ATTACK-LS] S.F. Wu B. Vetter, and F. Wang. An Experimental Study of Insider Attacks in a Link State Routing Protocol, In 5th IEEE International Conference on Network Protocols, Atlanta, GA, 1997.

[IGMP] B. Cain, S. Deering, I. Kouvelas, B. Fenner, and A. Thyagarajan, Internet Group Management Protocol, Version 2, [RFC 3376](#), October 2002

[PIM-SM] D. Estrin, D. Farinacci, A. Helmy, D. Thaler, S. Deering, M. Handley, V. Jacobson, C. Liu, P. Sharma, and L. Wei, Protocol Independent Multicast-Sparse Mode (PIM-SM): Protocol Specification, [RFC 2362](#), June 1998

[THREATS] - A. Ballardie and J. Crowcroft, Multicast-Specific Security Threats and Counter-Measures;; In Proceedings "Symposium on Network and Distributed System Security", February 1995, pp.2-16.
(<http://cs.ucl.ac.uk/darpa/IDMR/mcast-sec-isoc.ps.Z>)

Authors' Addresses

Dennis Beard
Nortel Networks
3500 Carling Avenue
Nepean, Ontario K2H 8E9
Canada

Phone:
EMail: beardd@nortelnetworks.com

Sandy Murphy
Network Associates, Inc
3060 Washington Rd.
Glenwood, MD 21738
USA

Phone: 443-259-2303
EMail: Sandra_murphy@nai.com

Yi Yang
Cisco Systems
7025 Kit Creek Road
RTP, NC 27709
USA

Phone:
EMail: yiya@cisco.com

Internet-Draft

Generic Threats to Routing Protocols

February 2003

[Appendix A](#). Acknowledgements

This draft would not have been possible save for the excellent efforts and team work characteristics of those listed here.

Ayman Musharbash - Nortel Networks

Paul Knight - Nortel Networks

Elwyn Davies - Nortel Networks

Ameya Dilip Pandit - Graduate student - University of Missouri

Senthilkumar Ayyasamy - Graduate student - University of Missouri

Internet-Draft Generic Threats to Routing Protocols February 2003

Intellectual Property Statement

The IETF takes no position regarding the validity or scope of any intellectual property or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; neither does it represent that it has made any effort to identify any such rights. Information on the IETF's procedures with respect to rights in standards-track and standards-related documentation can be found in [BCP-11](#). Copies of claims of rights made available for publication and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementors or users of this specification can be obtained from the IETF Secretariat.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights which may cover technology that may be required to practice this standard. Please address the information to the IETF Executive Director.

Full Copyright Statement

Copyright (C) The Internet Society (2003). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other

Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assignees.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION

Beard, et al.

Expires August 22, 2003

[Page 32]

Internet-Draft

Generic Threats to Routing Protocols

February 2003

HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

Acronyms

AODV - Ad-hoc On-demand Distance Vector routing protocol

AS - Autonomous system. Set of routers under a single technical administration. Each AS normally uses a single interior gateway protocol (IGP) and metrics to propagate routing information within the set of routers. Also called routing domain.

AS-Path - In BGP, the route to a destination. The path consists of the AS numbers of all routers a packet must go through to reach a destination.

BGP - Border Gateway Protocol. Exterior gateway protocol used to exchange routing information among routers in different autonomous systems.

eBGP - External BGP. BGP configuration in which sessions are established between routers in different ASs.

iBGP - Internal BGP. BGP configuration in which sessions are established between routers in the same ASs.

LSRP - Link-State Routing Protocol

LSA - Link-State Announcement

M-OSPF - Multicast Open Shortest Path First

NLRI - Network layer reachability information. Information that is carried in BGP packets and is used by MBGP.

OSPF - Open Shortest Path First. A link-state IGP that makes routing decisions based on the shortest-path-first (SPF) algorithm (also referred to as the Dijkstra algorithm).

PIM (and PIM DM) - Protocol Independent Multicast. A protocol-independent multicast routing protocol. PIM Sparse Mode routes to multicast groups that might span wide-area and interdomain internets. PIM Dense Mode is a flood-and-prune protocol.

RIP - Routing Information Protocol. Distance-vector interior gateway protocol that makes routing decisions based on hop count.

SPF - Shortest-path first, an algorithm used by IS-IS and OSPF to make routing decisions based on the state of network links. Also called the Dijkstra algorithm.

TCP - Transmission Control Protocol. Works in conjunction with Internet Protocol (IP) to send data over the Internet. Divides a message into packets and tracks the packets from point of origin to destination.

