

Network Working Group
Request for Comments: 4593
Category: Informational

A. Barbir
Nortel
S. Murphy
Sparta, Inc.
Y. Yang
Cisco Systems
October 2006

Generic Threats to Routing Protocols

Status of This Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2006).

Abstract

Routing protocols are subject to attacks that can harm individual users or network operations as a whole. This document provides a description and a summary of generic threats that affect routing protocols in general. This work describes threats, including threat sources and capabilities, threat actions, and threat consequences, as well as a breakdown of routing functions that might be attacked separately.

Table of Contents

- 1. Introduction2
- 2. Routing Functions Overview3
- 3. Generic Routing Protocol Threat Model4
 - 3.1. Threat Definitions4
 - 3.1.1. Threat Sources4
 - 3.1.1.1. Adversary Motivations5
 - 3.1.1.2. Adversary Capabilities5
 - 3.1.2. Threat Consequences7
 - 3.1.2.1. Threat Consequence Scope9
 - 3.1.2.2. Threat Consequence Zone10
 - 3.1.2.3. Threat Consequence Periods10
- 4. Generally Identifiable Routing Threat Actions11
 - 4.1. Deliberate Exposure11
 - 4.2. Sniffing11
 - 4.3. Traffic Analysis12
 - 4.4. Spoofing12
 - 4.5. Falsification13
 - 4.5.1. Falsifications by Originators13
 - 4.5.1.1. Overclaiming13
 - 4.5.1.2. Misclaiming16
 - 4.5.2. Falsifications by Forwarders16
 - 4.5.2.1. Misstatement16
 - 4.6. Interference17
 - 4.7. Overload18
- 5. Security Considerations18
- 6. References18
 - 6.1. Normative References18
- Appendix A. Acknowledgments20
- Appendix B. Acronyms20

1. Introduction

Routing protocols are subject to threats and attacks that can harm individual users or the network operations as a whole. The document provides a summary of generic threats that affect routing protocols. In particular, this work identifies generic threats to routing protocols that include threat sources, threat actions, and threat consequences. A breakdown of routing functions that might be separately attacked is provided.

This work should be considered a precursor to developing a common set of security requirements for routing protocols. While it is well known that bad, incomplete, or poor implementations of routing protocols may, in themselves, lead to routing problems or failures or may increase the risk of a network's being attacked successfully, these issues are not considered here. This document only considers

attacks against robust, well-considered implementations of routing protocols, such as those specified in Open Shortest Path First (OSPF) [4], Intermediate System to Intermediate System (IS-IS) [5][8], RIP [6] and BGP [7]. Attacks against implementation-specific weaknesses and vulnerabilities are out of scope for this document.

The document is organized as follows: Section 2 provides a review of routing functions. Section 3 defines threats. In Section 4, a discussion on generally identifiable routing threat actions is provided. Section 5 addresses security considerations.

2. Routing Functions Overview

This section provides an overview of common functions that are shared among various routing protocols. In general, routing protocols share the following functions:

- o Transport Subsystem: The routing protocol transmits messages to its neighbors using some underlying protocol. For example, OSPF uses IP, while other protocols may run over TCP.
- o Neighbor State Maintenance: Neighboring relationship formation is the first step for topology determination. For this reason, routing protocols may need to maintain state information. Each routing protocol may use a different mechanism for determining its neighbors in the routing topology. Some protocols have distinct exchanges through which they establish neighboring relationships, e.g., Hello exchanges in OSPF.
- o Database Maintenance: Routing protocols exchange network topology and reachability information. The routers collect this information in routing databases with varying detail. The maintenance of these databases is a significant portion of the function of a routing protocol.

In a routing protocol, there are message exchanges that are intended for the control of the state of the protocol. For example, neighbor maintenance messages carry such information. On the other hand, there are messages that are used to exchange information that is intended to be used in the forwarding function, for example, messages that are used to maintain the database. These messages affect the data (information) part of the routing protocol.

3. Generic Routing Protocol Threat Model

The model developed in this section can be used to identify threats to any routing protocol.

Routing protocols are subject to threats at various levels. For example, threats can affect the transport subsystem, where the routing protocol can be subject to attacks on its underlying protocol. An attacker may also attack messages that carry control information in a routing protocol to break a neighboring (e.g., peering, adjacency) relationship. This type of attack can impact the network routing behavior in the affected routers and likely the surrounding neighborhood as well. For example, in BGP, if a router receives a CEASE message, it will break its neighboring relationship to its peer and potentially send new routing information to any remaining peers.

An attacker may also attack messages that carry data information in order to break a database exchange between two routers or to affect the database maintenance functionality. For example, the information in the database must be authentic and authorized. An attacker who is able to introduce bogus data can have a strong effect on the behavior of routing in the neighborhood. For example, if an OSPF router sends LSAs with the wrong Advertising Router, the receivers will compute a Shortest Path First (SPF) tree that is incorrect and might not forward the traffic. If a BGP router advertises a Network Layer Reachability Information (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 Protocol Independent Multicast (PIM) router might transmit a JOIN message to receive multicast data it would otherwise not receive.

3.1. Threat Definitions

In [1], a threat is defined 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. Threats can be categorized as threat sources, threat actions, threat consequences, threat consequence zones, and threat consequence periods.

3.1.1. Threat Sources

In the context of deliberate attack, a threat source is defined as a motivated, capable adversary. By modeling the motivations (attack goals) and capabilities of the adversaries who are threat sources, one can better understand what classes of attacks these threats may mount and thus what types of countermeasures will be required to deal with these attacks.

3.1.1.1. Adversary Motivations

We assume that the most common goal of an adversary deliberately attacking routing is to cause inter-domain routing to malfunction. A routing malfunction affects data transmission such that traffic follows a path (sequence of autonomous systems in the case of BGP) other than one that would have been computed by the routing protocol if it were operating properly (i.e., if it were not under attack). As a result of an attack, a route may terminate at a router other than the one that legitimately represents the destination address of the traffic, or it may traverse routers other than those that it would otherwise have traversed. In either case, a routing malfunction may allow an adversary to wiretap traffic passively, or to engage in man-in-the-middle (MITM) active attacks, including discarding traffic (denial of service).

A routing malfunction might be effected for financial gain related to traffic volume (vs. the content of the routed traffic), e.g., to affect settlements among ISPs.

Another possible goal for attacks against routing can be damage to the network infrastructure itself, on a targeted or wide-scale basis. Thus, for example, attacks that cause excessive transmission of UPDATE or other management messages, and attendant router processing, could be motivated by these goals.

Irrespective of the goals noted above, an adversary may or may not be averse to detection and identification. This characteristic of an adversary influences some of the ways in which attacks may be accomplished.

3.1.1.2. Adversary Capabilities

Different adversaries possess varied capabilities.

- o All adversaries are presumed to be capable of directing packets to routers from remote locations and can assert a false IP source address with each packet (IP address spoofing) in an effort to cause the targeted router to accept and process the packet as though it emanated from the indicated source. Spoofing attacks may be employed to trick routers into acting on bogus messages to effect misrouting, or these messages may be used to overwhelm the management processor in a router, to effect DoS. Protection from such adversaries must not rely on the claimed identity in routing packets that the protocol receives.

- o Some adversaries can monitor links over which routing traffic is carried and emit packets that mimic data contained in legitimate routing traffic carried over these links; thus, they can actively participate in message exchanges with the legitimate routers. This increases the opportunities for an adversary to generate bogus routing traffic that may be accepted by a router, to effect misrouting or DoS. Retransmission of previously delivered management traffic (replay attacks) exemplify this capability. As a result, protection from such adversaries ought not to rely on the secrecy of unencrypted data in packet headers or payloads.
- o Some adversaries can effect MITM attacks against routing traffic, e.g., as a result of active wiretapping on a link between two routers. This represents the ultimate wiretapping capability for an adversary. Protection from such adversaries must not rely on the integrity of inter-router links to authenticate traffic, unless cryptographic measures are employed to detect unauthorized modification.
- o Some adversaries can subvert routers, or the management workstations used to control these routers. These Byzantine failures represent the most serious form of attack capability in that they result in emission of bogus traffic by legitimate routers. As a result, protection from such adversaries must not rely on the correct operation of neighbor routers. Protection measures should adopt the principle of least privilege, to minimize the impact of attacks of this sort. To counter Byzantine attacks, routers ought not to trust management traffic (e.g., based on its source) but rather each router should independently authenticate management traffic before acting upon it.

We will assume that any cryptographic countermeasures employed to secure BGP will employ algorithms and modes that are resistant to attack, even by sophisticated adversaries; thus, we will ignore cryptanalytic attacks.

Deliberate attacks are mimicked by failures that are random and unintentional. In particular, a Byzantine failure in a router may occur because the router is faulty in hardware or software or is misconfigured. As described in [3], "A node with a Byzantine failure may corrupt messages, forge messages, delay messages, or send conflicting messages to different nodes". Byzantine routers, whether faulty, misconfigured, or subverted, have the context to provide

believable and very damaging bogus routing information. Byzantine routers may also claim another legitimate peer's identity. Given their status as peers, they may even elude the authentication protections, if those protections can only detect that a source is one of the legitimate peers (e.g., the router uses the same cryptographic key to authenticate all peers).

We therefore characterize threat sources into two groups:

Outsiders: These attackers may reside anywhere in the Internet, have the ability to send IP traffic to the router, may be able to observe the router's replies, and may even control the path for a legitimate peer's traffic. These are not legitimate participants in the routing protocol.

Byzantine: These attackers are faulty, misconfigured, or subverted routers; i.e., legitimate participants in the routing protocol.

3.1.2. Threat Consequences

A threat consequence is a security violation that results from a threat action [1]. To a routing protocol, a security violation is a compromise of some aspect of the correct behavior of the routing system. The compromise can damage the data traffic intended for a particular network or host or can damage the operation of the routing infrastructure of the network as a whole.

There are four types of general threat consequences: disclosure, deception, disruption, and usurpation [1].

- o **Disclosure:** Disclosure of routing information happens when an attacker successfully accesses the information without being authorized. Outsiders who can observe or monitor a link may cause disclosure, if routing exchanges lack confidentiality. Byzantine routers can cause disclosure, as long as they are successfully involved in the routing exchanges. Although inappropriate 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 forged routing message and believes it to be authentic. Both outsiders and Byzantine routers can cause this consequence if the receiving router lacks the ability to check routing message integrity or origin authentication.

3.1.2.1. Threat Consequence Scope

As mentioned above, an attack might damage the data traffic intended for a particular network or host or damage the operation of the routing infrastructure of the network as a whole. Damage that might result from attacks against the network as a whole may include the following:

- 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 unnecessarily re-directed to be forwarded through one router and that router 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.
- o Overcontrol. The routing protocol messages themselves become a significant portion of the traffic the network carries.
- o Clog. A router receives an excessive number of routing protocol messages, causing it to exhaust some resource (e.g., memory, CPU, battery).

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

- 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 consequences, because some security solutions can protect against one consequence but not against others. It might be possible to design a security solution that protects against eavesdropping on one destination's traffic without protecting against churn in the network. Similarly, it is possible to design a security solution that prevents a starvation attack against one host, but not a clogging attack against a router. The security requirements must be clear as to which consequences are being avoided and which consequences must be addressed by other means (e.g., by administrative means outside the protocol).

3.1.2.2. Threat Consequence Zone

A threat consequence zone covers the area within which the network operations have been affected by threat actions. 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 the position of the target of the attack. Similar threat actions that happen at different locations may result in totally different threat consequence zones. For example, when an outsider breaks the routing session between a distribution router and a stub router, only reachability to and from the network devices attached to the stub router will be impaired. In other words, the threat consequence zone is a single router. In another case, if the outsider 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.

3.1.2.3. Threat Consequence Periods

A threat consequence period is defined as the portion of time during which the network operations are 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 persist longer than does the threat action. For example, in the original Advanced Research Projects Agency Network (ARPANET) link-state algorithm, some errors in a router introduced three instances of a Link-State Announcement (LSA). All of them flooded throughout the network continuously, until the entire network was power cycled [2].

4. Generally Identifiable Routing Threat Actions

This section addresses generally identifiable and recognized threat actions against routing protocols. The threat actions 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 occurs when an attacker takes control of a router and intentionally releases routing information to other entities (e.g., the attacker, a web page, mail posting, other routers) that otherwise should not receive the exposed information.

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 outdated 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 to sniff for routing information. Attackers can also sniff data traffic information (however, this is out of scope of the current work).

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 the routing information that has been recorded. For example, if the outsider is sniffing a link that is in an OSPF totally stubby area, the threat consequence zone should be limited to the whole area. An attacker that is sniffing a

link in an External Border Gateway Protocol (EBGP) session can gain knowledge of multiple routing domains.

The threat consequence period might be longer than the duration of the action. If an attacker stops sniffing a link, their acquired knowledge will not be out-dated until there is a topology change of the affected network.

4.3. Traffic Analysis

Traffic analysis is an action whereby attackers gain routing information by analyzing the characteristics of the data traffic on a subverted link. Traffic analysis threats can affect any data that is sent over a communication link. This threat is not peculiar to routing protocols and is included here for completeness.

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

The threat consequence zone of the traffic analysis depends on the attacker's location and what data traffic has passed through. An attacker at the network core should be able to gather more information than its counterpart at the edge and would therefore have to be able to analyze traffic patterns in a wider area.

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 outdated until there is a topology change of the disclosed network.

4.4. Spoofing

Spoofing occurs when an illegitimate device assumes the identity of a legitimate one. Spoofing in and of itself is often not the true attack. Spoofing is special in that it can be used to carry out other threat actions causing other threat consequences. An attacker can use spoofing as a means for launching other types of attacks. For example, if an attacker succeeds in spoofing the identity of a router, the attacker can send out unrealistic routing information that might cause the disruption of network services.

There are a few cases where spoofing can be an attack in and of itself. For example, messages from an attacker that spoof the identity of a legitimate router may cause a neighbor relationship to form and deny the formation of the relationship with the legitimate router.

The consequences of spoofing are as follows:

- o The disclosure of routing information. The spoofing router will be able to gain access to the routing information.
- o The deception of peer relationship. The authorized routers, which exchange routing messages with the spoofing router, do not realize that they are neighboring with a router that is faking another router's identity.

The threat consequence zone is as follows:

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

Note: This section focuses on addressing spoofing as a threat on its own. However, spoofing creates conditions for other threats actions. The other threat actions are considered falsifications and are treated in the next section.

4.5. Falsification

Falsification is an action whereby an attacker sends false routing information. 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. False routing information describes the network in an unrealistic fashion, whether or not intended by the authoritative network administrator.

4.5.1. Falsifications by Originators

An originator of routing information can launch the falsifications that are described in the next sections.

4.5.1.1. Overclaiming

Overclaiming occurs when a Byzantine router or outsider advertises its control of some network resources, while in reality it does not, or if the advertisement is not authorized. This is given in Figures 2 and 3.

data never reaches the destination Network 1, as in Figure 3. The ultimate consequence is that Router A gains control over Network 1's services, by controlling the data traffic.

- o Usurpation of the legitimate advertising routers. In Figures 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, as Router C is operating in a way intended by the authoritative network administrator.)
- o Deception of other routers. In Figures 2 and 3, Router B, or other routers on the Internet, might be deceived into believing that the path through Router A is the best.
- o Disruption of data planes on some routers. This might happen to routers that are on the path that is used by other routers to reach the overclaimed network resources through the attacker. In Figures 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 Figures 2 and 3), and the routers that are authorized to advertise the network resources but lose the competition against the attacker (Router C in Figures 2 and 3).
- o Where deception is concerned, the consequence zone covers the routers that do believe the attacker's advertisement and use the attacker to reach the claimed networks (Router B and other deceived routers on the Internet in Figures 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 Figures 2 and 3 and other routers in the Internet on the path of the misdirected traffic).

The threat consequence will not cease when the attacker stops overclaiming and will totally disappear only when the routing tables are converged. As a result, the consequence period is longer than the duration of the overclaiming.

4.5.1.2. Misclaiming

A misclaiming threat is defined as an action whereby an attacker is advertising some network resources that it is authorized to control, but in a way that is not intended by the authoritative network administrator. For example, it may be advertising inappropriate link costs in an OSPF LSA. An attacker can eulogize or disparage when advertising these network resources. Byzantine routers can misclaim network resources.

The threat consequences of misclaiming are similar to the consequences of overclaiming.

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

4.5.2. Falsifications by Forwarders

In each routing protocol, routers that forward routing protocol messages are expected to leave some fields unmodified and to modify other fields in certain circumscribed ways. The fields to be modified, the possible new contents of those fields and their computation from the original fields, the fields that must remain unmodified, etc. are all detailed in the protocol specification. They may vary depending on the function of the router or its network environment. For example, in RIP, the forwarder must modify the routing information by increasing the hop count by 1. On the other hand, a forwarder must not modify any field of the type 1 LSA in OSPF except the age field. 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 messages, an attacker might also falsify by not forwarding routing information to other authorized routers as required.

4.5.2.1. Misstatement

This is defined as an action whereby the attacker modifies route attributes in an incorrect manner. For example, in RIP, the attacker might increase the path cost by two hops instead of one. In BGP, the attacker might delete some AS numbers from the AS PATH.

Where forwarding routing information 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.

A forwarder can also falsify data by replaying out-dated data in the routing message as current data.

All types of attackers, outsiders and Byzantine 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 zone and period are also similar.

4.6. Interference

Interference is a threat action whereby an attacker inhibits the exchanges by legitimate routers. The attacker can do this by adding noise, by not forwarding packets, by replaying out-dated packets, by inserting or corrupting messages, by delaying responses, by denial of receipts, or by breaking synchronization.

Byzantine routers can slow down their routing exchanges or induce flapping in the routing sessions of legitimate neighboring routers.

The consequence of interference is the disruption of routing operations.

The consequence zone of interference depends on the severity of the interference. If the interference results in consequences at the neighbor maintenance level, then there may be changes in the database, resulting in network-wide consequences.

The threat consequences might disappear as soon as the interference is stopped or might not totally disappear until the networks have converged. Therefore, the consequence period is equal to 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. For example, it is possible for an attacker to trigger a router to create an excessive amount of state that other routers within the network are not able to handle. In a similar fashion, it is possible for an attacker to overload database routing exchanges and thus to influence the routing operations.

5. Security Considerations

This entire document is security related. Specifically, the document 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 that are currently known and lays the foundation for other documents that will discuss security requirements for routing protocols. This document is protocol independent.

6. References

6.1. Normative References

- [1] Shirey, R., "Internet Security Glossary", RFC 2828, May 2000.
- [2] Rosen, E., "Vulnerabilities of network control protocols: An example", RFC 789, July 1981.
- [3] Perlman, R., "Network Layer Protocols with Byzantine Robustness", PhD thesis, MIT LCS TR-429, October 1988.
- [4] Moy, J., "OSPF Version 2", STD 54, RFC 2328, April 1998.
- [5] Callon, R., "Use of OSI IS-IS for routing in TCP/IP and dual environments", RFC 1195, December 1990.
- [6] Malkin, G., "RIP Version 2", STD 56, RFC 2453, November 1998.
- [7] Rekhter, Y., Li, T., and S. Hares, "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, January 2006.

- [8] ISO 10589, "Intermediate System to Intermediate System intra-domain routing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode network service (ISO 8473)", ISO/IEC 10589:2002.

Appendix A. Acknowledgments

This document would not have been possible save for the excellent efforts and teamwork characteristics of those listed here.

- o Dennis Beard, Nortel
- o Ayman Musharbash, Nortel
- o Jean-Jacques Puig, int-evry, France
- o Paul Knight, Nortel
- o Elwyn Davies, Nortel
- o Ameya Dilip Pandit, Graduate student, University of Missouri
- o Senthilkumar Ayyasamy, Graduate student, University of Missouri
- o Stephen Kent, BBN
- o Tim Gage, Cisco Systems
- o James Ng, Cisco Systems
- o Alvaro Retana, Cisco Systems

Appendix B. Acronyms

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.

LSA - Link-State Announcement

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).

Authors' Addresses

Abbie Barbir
Nortel
3500 Carling Avenue
Nepean, Ontario K2H 8E9
Canada

E-Mail: abbieb@nortel.com

Sandy Murphy
Sparta, Inc.
7110 Samuel Morse Drive
Columbia, MD
USA

Phone: 443-430-8000
E-Mail: sandy@sparta.com

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

E-Mail: yiya@cisco.com

Full Copyright Statement

Copyright (C) The Internet Society (2006).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights 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; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat 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 implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgement

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).