Network Working Group
Internet Draft
Expiration Date: October 2005
File name: draft-zinin-rtg-dos-02.txt

Alex Zinin
Alcatel
May 2005

Protecting Internet Routing Infrastructure from Outsider DoS Attacks

draft-zinin-rtg-dos-01.txt

Status of this Memo

Abstract

The mechanism described in this document helps to secure an Internet Service Provider's router infrastructure from outsider attacks, including (but not limited to) Distributed denial of service (DDoS) attacks based on CPU and/or queue exhaustion (e.g., TCP SYN flooding and flooding of invalid MD5-signed routing protocol packets.) The presented approach is based on explicitly marking control packets from trusted sources by different link-layer encapsulation and does not require any modifications to user data exchange protocols, ICMP,

routing protocols or changes to existing hardware in routers, which allows it to be deployed quickly throughout the Internet.

1 Introduction

1.1 Problem Description

The packet authentication mechanisms currently used in Internet routing protocols [OSPF, TCP-MD5] leave a generic threat open for an outside attacker--overloading the control CPUs of the routers with packets that look like they belong to a valid routing protocol adjacency or a peering session, yet are fake and would be discarded because of invalid digest value. Because all IP parameters of valid and faked packets look absolutely identical, it is impossible to reject faked packets earlier in the process. This leads to overloading of internal queues allocated for control traffic (routing and signaling protocols), and hence dropping of legitimate control packets. This, combined with high CPU utilization, results in destruction of routing protocol sessions and finally in denial of service by the network. It is interesting to observe that as security mechanisms in routing protocols become more sophisticated and computationally expensive, it becomes easier for an attacker to mount a CPU-exhaustion-based attack against a router.

Another example of an attack mountable against routers is the simple SYN-flood attack, which could potentially exhaust the router's CPU.

The in-band nature of IP routing and signaling creates a perfect environment for an attacker to put the network itself out of service. The fundamental problems leading to the possibility of a DoS attack on a router are (a) legitimate and forged packets share resources inside the router (such as queues) before the authentication check is performed, and (b) the negative authentication decision is computationally expensive enough to discourage router vendors from performing the check at the line rate. In the latter case, it is important to note that the lack of line-rate processing significantly increases the router's susceptibility to a distributed DoS attack.

1.2 Existing Approaches and Disadvantages

Potential approaches to the problem known to date include:

1. Adding specialized HW elements to the line-card architecture that would allow the line cards to identify packets that need to be authenticated (e.g., OSPF, BGP, RSVP) and perform the MD5 check at the line rate (before the packets are put in any queue), as well as identify TCP SYN packets and limit the rate at which they are sent to the control card.

2. Perform aggressive packet filtering at the edges of the network, on both customer-facing and service provider peering interfaces to make sure that packets destined for the internal routers are not received from outside the network.

3. Use a completely separate set of links for control protocols and customer data, i.e. out-of-band network control.

Below are the disadvantages of these methods (correspondingly):

1. From the service provider's perspective, additional HW increases the cost of the system and requires upgrades of the line-cards of all routers in a service provider's network. From the Internet security perspective, it will take years before a considerable number of service providers upgrade their routing infrastructure, and thus before the threat of DoS attack on the Internet routing system is sufficiently mitigated.

2. Many of today's deployed Internet core routers do not have the ability to perform line-rate access control list (ACL) processing at high speeds, which means that the inter-service provider links will remain insecure. Combined with the fact that not all service providers filter potentially dangerous packets on the customer interfaces, this approach has the same disadvantages from the deployment and Internet security perspective as the first approach.

3. While the out-of-band control scheme is extremely interesting, implementation could require substantial modification to the routing protocols and complete re-architecting of the service provider networks. From the architectural point of view, it would also require a major shift from the assumption that control-plane connectivity implies connectivity of the data plane.

The solution described in this document allows service providers to improve their network without major hardware upgrades , changes to routing protocols or network architecture, and with limited software modifications.

2 Solution

2.1 Overview

The proposed mechanism uses the fact that there are only a limited number of devices in a network that have a legitimate right to send a packet to a router's control plane. This set of devices includes, of

course, other routers in the service provider's network, and network operations center (NOC) machines. The rest of the devices in the Internet, including user hosts, and routers in other service provider networks should not need to send packets to the routers internal to the first provider's network..

They key aspect of the proposal is marking of packets from the set of trusted devices in a way that it would either be impossible to spoof by an untrusted device or that would ensure that even if an attacker created such a packet, it would be dropped by the routers already deployed in the Internet today. One option of such marking described in this document is using a different protocol ID in the layer-2 frames when sending IPv4 control packets among the routers. We call this "control IPv4 encapsulation". All Internet routers used today will drop these packets as unrecognized by default. This step makes sure that such a packet marking technique can be relied upon.

The next step is a small modification to the router's local IP processing and encapsulation logic to allow only control-encapsulated IPv4 packets to be sent to the control plane along the normal path. Other packets are considered dangerous and are put on a heavily rate-limited queue. This ensures that outsider attacks do not exhaust resources used for communication with trusted devices. Note that the encapsulation check has O(1) complexity, and can easily be performed at line rate even in legacy routers without major HW or SW modifications. One of the many advantages of this approach is in the fact that no additional packet filtering at the customer or peering interfaces is requires by the service provider, since user data packets always enter the network as dangerous.

Note that the proposed mechanism does not require modification or affect existing Internet user data or network troubleshooting protocols--ICMP will still work they way it works today, so ping, traceroute, TCP path MTU discovery, remain functional. The reason for this is the fact that the proposal only helps routers quickly classify packets as trusted and untrusted, but does not require untrusted packets (e.g., ICMP) to be dropped. Of course, if a router already has a capability to identify ICMP packets and put them on a separate queue, the service provider may decide to configure the router to drop all untrusted packets except for ICMP.

It should also be noted, that this proposal is not an attempt to protect from compromised trusted routers or insider attacks, neither it is an attemp to substitute existing security mechanisms in routing protocols. Instead it helps to protect the routers from outsider (user-level) attacks, such as distributed DoS attacks based on infection of untrusted devices (Internet-connected hosts) with computer viruses turning them into traffic generators targeted against

Internet routers, which is considered to be a bigger threat today. In the situation where a trusted router is compromised, the mechanism still offers additional security by limiting the potential affect of the attack to the boundaries of the trust domain the compromised router participates in through the notion of interface groups. Routers implementing this mechanism and not participating in that domain will not be susceptible to the attack.

Finally, the mechanism allows for gradual deployment across the Internet without a flag day and incremental security gain as it is deployed wider.

2.2 Separating Data and Control Encapsulation

As discussed before, the packet marking technique needs to have the property of default invalidity in order to make sure that no data flowing on the Internet today is considered trusted and is accepted into a service provider's network with such marking if an attacker tried to spoof a packet. Using techniques like DSCP-code marking or IP options does not satisfy this requirement, as it would call for filtering at every customer-facing router in the Internet to make sure that no user data packet is injected with this reserved DSCP value. This is the reason why the author has chosen to use a layer-2 encapsulation technique to achieve this--frames carrying unknown protocols are dropped by todays deployed routers.

This document describes two possible methods for a different layer-2 encapsulation--a separate protocol ID, and a link-local MPLS label. Each has its own advantages and disadvantages discussed below .

Option 1: New Protocol ID

As a protocol ID value is defined for IPv4 and IPv6 for each used media type today (such as Ether_type code), it would be possible to define IPv4-control and IPv6-control protocol IDs.

The advantage of this method is an implicit 100% guarantee that if the protocol ID is selected from an unused space, the packets will be unrecognized. This approach also seems like the "clean" way of doing this.

The disadvantage of this approach is that the control encapsulation protocol ID will need to be defined for each media type used today, which may take a while. Another disadvantage is that in case of an MPLS network, a control packet maybe put on an LSP together with data packets, so the receiving router wouldn't be able to tell the difference. Getting around this problem may require maintaining two

sets of next-hops per route in the data path. See Option 3 below, however.

Option 2: Link-local MPLS Label

This method is more of a hack and relies on the fact that MPLS encapsulation is either defined for or mapped to most of today's used media types. It should be possible to reserve a single label value (or 2 if a separate one for IPv6 is deemed necessary) from the "reserved" range (values 4-15 defined by [MPLS-STACK]), declare it to be link-local, disallow this value from being used for transit MPLS LSPs, and use this as the control encapsulation. Note that since the label would only have significance on the local link, it can be reused on all links. Control messages used for signaling of transit label switched paths (LSPs) can be safely put on top of this label, as there are no order of origin dependencies. Routers that do not support MPLS would not need to have any MPLS code added and could just treat this as a special sequence of octects in the link frame that identifies control encapsulation.

When a control packet for a multi-hop routing session (iBGP or OSPF virtual link) is put on an LSP, an extra label with the reserved value would be added on top of the label stack thus identifying the control packet.

Because service providers generally do not support MPLS on their customer interfaces, and because the label value would be taken from the reserved space, it would be impossible for an Internet user to spoof a control packet using existing Internet infrastructure.

The advantage of this approach is that only a single value for the label would need to be reserved.

The disadvantages are that more modifications of the router microcode are necessary.

Option 3: Combined

It is possible to use the new protocol ID whenever a control packet is not MPLS-encapsulated, and use an extra reserved label whenever it is put on an LSP. See section XXX for more information on support of MPLS networks.

Control-plane software is then modified to make sure that all locally-originated packets that are relevant within the service provider's network only (such as routing protocols, MPLS signaling, telnet, ssh, SNMP, etc.) are control-encapsulated when the outbound

interface is configured as such. Control packets that need to be received by the users (ICMP) are either encapsulated as before (data encapsulation) or also as control. In the latter case, they will be data-encapsulated as soon as they leave the trust domain of the service provider.

2.3 Interface Groups

When deploying this mechanism, the service provider will need to identify a group of interfaces where the control encapsulation should or should not be used. There will most probably be a group of interfaces used for the backbone connection, and another group used for customer connections and peering with other service providers.

The described mechanism uses the notion of an "interface group". There is practically no complexity associated with an interface group--each interface has an interface-group attribute associated with it. Two interfaces are considered to be in one interface group if their interface-group attributes are equal. The service provider is expected to configure the interface group attributes of the interfaces to match the trust communities, as in the following example. Backbone interfaces, interfaces to customer A, interfaces to customer B, interfaces to service provider X, and interfaces to service provider Y, would all be put in separate interface-groups: "backbone", "cust-A", "cust-B", "peer-X", "peer-Y", correspondingly.

As we will see further in the document, when a control-encapsulated packet is forwarded across an interface-group boundary, it become data-encapsulated (untrusted). This is to ensure that if, for example, two service providers are using control encapsulation for their eBGP session, or if an eBGP session between a service provider and a customer is control-encapsulated, forged packets originated by a potentially compromised BGP peer and destined inside of the service provider's network are not considered trusted beyond the border router. In other words, we trust control traffic from a customer or another service provider only as far as it needs to go and no further. Again, once a control-encapsulated packet crosses an interface-group boundary, its encapsulation is changed to data and it will be considered as untrusted by all other routers.

2.4 Modified Local Processing and Packet Encapsulation Procedures

The following new interface parameters used by the modified algorithms are introduced.

InterfaceGroup:
the ID of the group the interface belongs to

IpCtlSendEncap:
defines which encapsulation should be used on the interface to send control packets originated locally by the router or received as control-encapsulated on another interface. Possible values: Data, and Control. Default: Data.

IpCtlRcvEncap:
defines the type of encapsulation that needs to be used in order for the received packet to be allowed for local processing by the RP as trusted. Values: Data, Control, Both . Default: Data.

The router's behavior is modified as follows.

1. A packet addressed to the router itself is considered trusted and is allowed to be locally processed (queued to the control card) if IpCtlRcvEncap of the receiving interface is set to Both, or matches the encapsulation that was used to send the received packet. Otherwise, the packet is put on a "slow" queue (or dropped if the router has the capability to recognize ICMP packets and still allow them to be processed in a rate-limited fashion).

2. The router uses Control encapsulation for an outgoing packet if IpCtlSndEncap of the outbound interface is Control AND the packet:

   a) Has been locally originated by the router itself, OR

   b) Has been received in Control encapsulation AND Interface-Group parameters of the inbound and outbound interfaces are the same (the packet is not leaving its trust domain.)

   Otherwise, (the packet is untrusted or is leaving its trust domain by crossing the interface-group boundary), Data encapsulation is used.

2.5 NOC Support and "Trusted" Interfaces

Hosts on the NOC segments of the service provider's network are an example of trusted devices that are not routers. However, unlike routers, it is unrealistic to expect hosts within the NOC segment to exchange packets using Control encapsulation, as this would require modification to many operating systems. Another specific of a NOC segment is the fact that it majority of cases, it will need to be able to communicate with the rest of the service provider's network using both Data and Control encapsulated packets. The following is an

explanation why.

It is already a common practice to allow incoming telnet, ssh, and snmp packets to routers only if they were originated within a NOC segment (this is usually done by configuring CLI and SNMP-specific filters by access control lists), however, the network administrators are rarely physically located on the NOC premises-many of them work from home and are often mobile. This is why management access to the routers usually requires establishing a secure shell (SSH) session to a server in the NOC, and then from there another SSH session to a router. Of course, the SSH server is usually behind a firewall (let's call it FW1). In our case, this would be a firewall that communicates Data packets with the rest of the service provider's network.

Since all telnet, ssh, snmp, etc. packets going from the NOC segment to the routers in the network need to appear in Control encapsulation, regular data packets exchanged on the NOC segment at some point need to be sent out as Control encapsulated packets. This, of course, introduces a potential security threat (if the hosts on the NOC segment were used to attack the routers, all forged packets would be considered by routers as trusted.) However, it is much less expensive for a service provider to protect its routers from its own NOC segments by installing a firewall (let's call it FW2) that will make sure that only valid packets are sent out as control to the routers in the network.

Note that FW1 and FW2 are only functionally separate, but may physically be the same device.

There are potentially two ways how NOC data packets can be injected as Control into the network: a) FW2's network-facing interface supports Control encapsulation, and b) FW2 has no support of Control encapsulation, but the first-hop router it is connected to performs the "translation". The former case is the most secure, while the latter is the most probable, at least in the beginning. Below is how the router performs the translation function.

The notion of a "trusted interface" is defined by introducing the following parameter:

IpTrustedInterface: When True, identifies a trusted interface. It is expected that only very few interfaces in the service provider's network will be configured as Trusted (for example, interfaces connecting a NOC segment to the rest of the network through a firewall.) Possible values: True, and False. Default: False.

The router's behavior is further modified to accommodate the notion of trusted interface as follows:

1. A packet received on a Trusted interface in any encapsulation is treated as if it was received in Control encapsulation (i.e., is allowed to be locally processed and is sent out using Control encapsulation as long as it stays within the same interface group).

2. All packets (trusted and untrusted) sent out of a Trusted interface are Data-encapsulated.

DISCUSSION: we may want to allow only trusted packets to be sent on a Trusted interface towards NOC. This will make the job of FW2 much easier, but will cut off ICMP messages coming from outside the network or from a different trust domain if the service provider has many.

2.6 ICMP, Ping, and Traceroute

ICMP needs special attention, because its scope of validity is not so well contained as for routing and signaling protocols. Let's consider how this proposals handles ICMP by looking at the following generic combinations for ICMP messages:

1. Originated by and addressed to devices within the same trust domain, for example, an ICMP "Echo Request" message originated by a NOC host and received by a router. Same for an ICMP "Echo Reply". All ICMP messages will be considered trusted. No issues here.

2. Originated by a trusted device (router), addressed to an untrusted one (a ICMP "Destination Unreachable" to an Internet-connected host, for example). The router will inject the packet using Control encapsulation, however, as the packet leaves the service provider's network, it will be sent out using Data encapsulation (see step 2 in the modified router algorithm), as expected by the receiver.

3. Originated by an untrusted host, addressed to a service provider's router. The router will put the packet on the 'untrusted' queue, and it will be processed.

4. Originated by and addressed to an untrusted host. The message will enter and leave the network as untrusted data without touching any router's control plane.

Traceroute from the outside world does not present any problem, because Control-encapsulated ICMP messages sent back to the probing host will be automatically converted to Data as they leave the trust domain.

Traceroute within the network (e.g., from NOC or a router) is not a problem because the messages are exchanged in Control encapsulation. If traceroute crosses multiple trust domains or goes outside the service provider's network, ICMP messages will come back as Data and will go through a FW to NOC or may be received through the slow queue by the router (if traceroute is originated by the router.)

2.6 Routing Protocols

One of the advantages of the described mechanism is that no modification of existing routing protocols is required. Routing protocols still work over IPv4, the only difference is actual layer-2 encapsulation of those packets, which is (in the simplest case) Control for all packets originated by a router.

The routing paradigm remains the same--the messages are sent inbound across the same physical links as data packets. Control and data are only virtually separated, just enough to make a decision on whether a packet should be considered from a trusted source or not.

Because the IS-IS routing protocol encapsulates its PDUs in L2 frames, as opposed to IP packets, it is not susceptible to the outsider attacks, and hence no modification to IS-IS encapsulation is required. If IS-IS-in-IP is used, the routers need to make sure that the IP packets is Control-encapsulated. Note that the fact the IS-IS routing protocol is not susceptible to outsider attacks does not mean that ISP running IS-IS should not be worried about those attacks. There's a whole set of potential CPU-based attacks which an outsider could mount, and this set is constantly growing.

2.7 Multicast

There are two aspects of IP mutlicast we're interested in from the routing security point of view: routing protocols, and (S,G) state.

From the routing protocols perspective, service provider's routers are protected by the presented mechanism as with unicast.

The link between data and control plane required to maintain the (S,G) state is part of the multicast architecture and may be considered by some as an issue (it is definitely safer to decouple control and data planes of the network as much as possible). Presented security mechanism does not affect it in anyway. The service provider will have to make an informed decision whether to deploy multicast in its network or not keeping in mind the possibility of some router implementations not being able to keep up with large amounts of (S,G) state.

2.8 MPLS Networks

When the mechanism is deployed in an MPLS network, it is possible for any IP packet (including a control one) that is sent over multiple hops to be put on an LSP due to an LSR using either LDP-derived FECs or IGP-shortcut FECs. Because layer-2 encapsulation is not preserved when an IP packet is put on an LSP, it will be impossible for the receiving router to tell the difference between data and control packets if just a different link-layer protocol ID was used to mark trusted packets.

To solve this problem, the LSR putting the control packet on an LSP, adds an extra inner label with the reserved value described before to the label stack.

If penultimate hop popping (PHP) is used in the network, the tail-end LSR may not even notice the fact that the packet has traveled on a LSP if the MPLS-label approach is used for encapsulation, because the LSR will receive the packet with only one--reserved--label.

If the PHP mechanism is not used, the receiving LSR, after popping the outer label, will need to recognize the reserved value of the inner label and treat the packet as Control-encapsulated.

3 Deployment Considerations

The following subsections discuss how the described mechanism would be deployed in a service provider's network. Note that we consider the final setup, after all transitional steps. The transition scenarios are described in a separate subsection

3.1 Backbone-only Routers

Routers where all interfaces are connected to internal links will most often have all of them configured to be in the same interface group. It is possible of course, to have multiple Control trust domains within a single service provider's network if for example, BGP AS confederations are used. In this case, each member-AS would be a separate trust domain and some BGP speakers would have more than one interface group. One consideration related to running a network with multiple trust domains is the fact that control message that are not naturally scoped to a single trust domain (such as ICMP) will be encapsulated as Data once they leave the trust domain they have been originated in. This means that Control encapsulation-aware firewalls connecting the NOC segment need to also receive and process Data-encapsulated ICMP.

Receiving and sending encapsulation of control packets would be set

to Control on all interfaces.

3.2 Customer-facing Routers

Customer facing routers will have more than one interface group.

One group will be configured for all backbone links. In this group receive and send encapsulation will be configured as Control.

For each customer, all interfaces providing connections to it will be configured as a separate interface group. The type of encapsulation is expected to be Data for a long time, before customer routers start supporting Control encapsulation. With Data encapsulation, the router is allowed to send Data-encapsulated control packets to the control plane CPU. Other packets, supposedly both valid data and potentially forged packets, are forwarded onwards to the network using Data encapsulation, so other routers in the network won't allow these packets to the control plane in case of an attack.

When Control encapsulation is supported by the customer routers, the service provider will configure send and receive control packet encapsulation on those links to be Control. This will prevent DoS attacks on the customer-facing router on those links.

3.3 Peer-facing Routers

Peer-facing routers will be configured similar to the customer-facing routers. If the peering routers do not support Control encapsulation, the routers are configured to allow Data-encapsulated packets to be received by the control CPU. Potential attacks against the border router could be prevented by the BGP TTL hack (though implementing Control encapsulation seems easier.) service provider's internal routers will not be susceptible to the attacks originated in other service providers, because forged packets will be sent as Data and won't be allowed to the routers' control plane CPUs. When Control encapsulation is supported, the border router will be protected from the DoS attack on the links to those service providers supporting this technique.

An important point to keep in mind here is the fact that trust domains of the service providers are not merged when they peer with each other. Links used to peer with other service providers are put in a separate interface group from the backbone interface group. This means that even if routers of another service provider are compro-mised and forged packets are sent as Control to us, they would first be translated to Data encapsulation by that service provider's border router, but even if they are not for some reason (or if the service provider's border router is compromised), our border router will

"translate" any forged control packets into Data as they cross the boundary between the peering and the backbone interface group.

3.4 Internet Exchgange Points and Help from LAN Switches

In the case where a LAN segment is used to connect devices under different administrative control (as in the case of an Internet Exchange point, for example), it is possible that some connected devices may not be trusted enough for others to agree to receive control-encapsulated packets from them. In order to avoid line-rate source-based filtering, LAN switches may be equipped with a small bit of additional functionality controlling whether inbound control-encapsulated packets are allowed on a specific port. Because this check is done on a per-port basis and the switch does not need to look further than the layer-2 frame, this check can easily be performed at the line rate without performance degradation. The LAN switch administrator would then have to ensure that control frames are allowed only from trusted devices.

From the security perspective this essentially means that the service provider, normally only marginally trusting its IX peers, would need to trust the IX administrator's decision on whether the remote device is trusted or not, in other words, when a service provider agrees to accept control frames on an IX-connected interface, it essentially agrees to trust the security of that IX. However, though this means that IX-connected routers are less secure from each other, even if a trusted IX router is compromised, the effect of the attack is limited to the border router by the notion of the interface groups. Besides, IX-connected routers still have the same level of security against user-level attacks, which is thought of as a bigger threat than an attack from a compromised or untrusted IX-connected device.

3.5 NOC

As describe before, NOC segments can be connected to a service provider's network either through a Control-encapsulation-aware FW, or through a regular FW connected to a router implementing Trusted interfaces.

3.5 Transition Scenarios

To be completed.

Note that no flag day is required and gradual deployment gives incremental security increase.

4 Security Considerations

The described proposal does not claim to provide complete protection of routers against all types of attacks. Instead, it raises the bar by attempting to prevent attacks mounted by outsiders that have no access to the SP's network except for basic IP connectivity. These types of attacks are considered to be the immediate threat on the Internet routing system and the proposal attempts to protect against it without requiring expensive hardware upgrades. By virtually separating control and data packets, the level of security in IP networks is raised to the one normally found in ATM or Frame Relay networks, where routing and signaling are virtually out-of-band. This level of security is considered by many to be just enough to feel comfortable.

Insider attacks, based on the physical access to the SP's equipment or on compromising a trusted device (such as a router or a NOC-attached host) are not prevented by this mechanism.

The described proposal relies on the notion of a trust domain, which implies that if a router is configured to accept Control-encapsulated packets on an interface, the administrator administrator has full control of the devices attached to the segment and capable of sending Control-encapsulated packets (in reality, any connected device should be assumed to be capable of doing so), and those devices are authorized to send them. In other words, physical security needs to be insured by the SP. This practically means that no devices that with high probability can be compromised by an outside attacker (such as servers, or hosts) should be allowed on the segments used for router connections. Point-to-point links used between routers encourage this requirement by their very nature, while LAN segments require more attention to ensure no unauthorized devices have access to them. Fortunately, this is already the best current practice that the service providers follow. In the situations where a device connected to an otherwise trusted segment is considered to be highly susceptible to being compromised, some help from the LAN switch used to implement the segment is required. See Section XXX for more information.

Finally, because the described mechanism does not prevent from insider attacks, it should not be considered as a substitute for existing or future authentication mechanisms in routing protocols or other security measures used in the service provider networks (e.g., SSH). Instead, they should be considered complimentary to each other and used together. In fact, the more elaborate and computationally expensive routing protocol-specific mechanisms become, the easier it will be for an outside attacker to bring a router to its knees, and the more important it will be to separate control and data encapsulation in the Internet.

4. Intellectual Property Considerations

5. Acknowledgements

Thanks to Ben Crosby, Steve Buchko, Randy Bush, John Heasley, Radia Perlman, and Tony Li for an early review of this work.

6. References

[OSPF] J. Moy. OSPF version 2. Technical Report RFC 2328, Internet Engineering Task Force, 1998.

[TCP-MD5] A. Heffernan. Protection of BGP Sessions via the TCP MD5 Signature Option. RFC 2385. 1998.

[MPLS-STACK] RFC 3032. MPLS Label Stack Encoding.

7. Authors' Addresses

Alex Zinin
Alcatel
E-mail: zinin@psg.com

Full Copyright Statement