OPSEC
Internet-Draft
Expires: August 7, 2006

R. Bonica
Juniper Networks
S. Ahmed
Booz Allen Hamilton
February 3, 2006

Network Management Access Security Capabilities
draft-ietf-opsec-nmasc-00

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with Section 6 of BCP 79.

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/1id-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 7, 2006.

Copyright Notice

Copyright (C) The Internet Society (2006).

Abstract

This document describes how network management stations can communicate with the devices that they manage using either the in-band network, an out-of-band network, or a virtual out-of-band network. This document also evaluates each access method in terms of its security capabilities and lists the device capabilities needed to support each method.

Table of Contents

1. Introduction

The Framework for Operational Security Capabilities [4] outlines the proposed effort of the IETF OPSEC working group. This includes producing a series of drafts to codify knowledge gained through operational experience about feature sets that are needed to securely deploy and operate managed network elements providing transit services at the data link and IP layers. Current plans include separate capabilities documents for Packet Filtering; Event Logging; In-Band and Out-of-Band Management; Configuration and Management Interfaces; AAA; and Documentation and Assurance.

This document describes in-band management, out-of-band-management, and a hybrid approach, called virtual out-of-band management.

1.1. Conventions Used In This Document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC2119 [1].

2. Network Management Access Methods

Network management stations can communicate with the devices that they manage using either the in-band network, an out-of-band network, or a virtual out-of-band network. The following sections describe each of the above mentioned network management access methods.

3. In-band Access

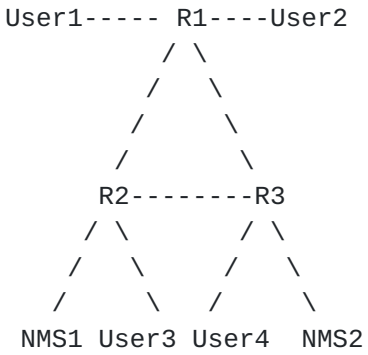

Figure 1: In-Band Access

Figure 1 depicts two network management stations (NMS1 and NMS2) managing three routers (R1-R3). Network management stations use the in-band network to communicate with the routers that they manage. Therefore, network management traffic is intermixed with user traffic on in-band network interfaces.

3.1. Vulnerabilities

RFC 3871 [2] identifies the following security vulnerabilities associated with in-band management:

1. Saturation of customer lines or interfaces can make the device unmanageable unless out-of-band management resources have been reserved.

2. Since public interfaces/channels are used, it is possible for attackers to directly address and reach the device and to attempt management functions.

3. In-band management traffic on public interfaces may be intercepted, however this would typically require a significant compromise in the routing system.

4. Public interfaces used for in-band management may become unavailable due to bugs (e.g., buffer overflows being exploited) while out-of-band interfaces (such as a serial console device) remain available.

Expanding upon the final point, listed above, the in-band network can be misconfigured, such that the managed device becomes isolated with regard to the network management stations. Similarly, DoS attacks against the routers or routing protocols, instability in the routing protocols, or other problems could cause the in-band network to become unavailable. When this happens, operators cannot access the router in order to remedy the configuration error. They become reliant upon physical access to the managed device.

3.2. Required Security Capabilities

RFC 3871 requires the following security capabilities to mitigate the effects of the above mentioned vulnerabilities:

1. Increased priority for management traffic.

2. Use of strong cryptography

3. Selection of cryptographic parameters

4. Use of cryptographic algorithms subject to open review

5. Use of management protocols subject to open review

3.3. Analysis

The cryptographic methods mentioned in Section 3.2 prevent users from accessing management functions and eavesdropping on management traffic. The strength of the cryptographic algorithms deployed should be a determined by cost and perceived threat.

Increasing the priority of management traffic reduces, but does not eliminate, the risk associated with denial of service attacks against the router's management plane. In order to completely eliminate this risk, a network would have to police high priority traffic at each ingress point as well as elevate the priority of management traffic. See Section 5.2 below and also see section 4 of [5].

None of the capabilities mentioned in Section 3.2 address the final vulnerability mentioned in Section 3.1. Failure of the in-band network will render the network unmanageable. This is an inherent weakness of in-band management.

4. Out-of-Band Access

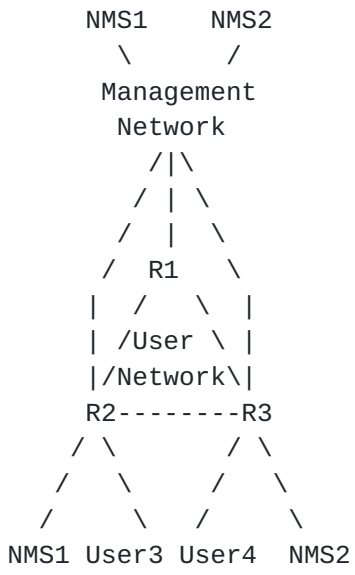

Figure 2: Out-of-Band Access

Figure 2 also depicts two network management stations managing three

routers. In this figure, the network management stations use a dedicated, out-of-band network to communicate with the routers that they manage.

Each router maintains a dedicated management interface. The dedicated management interface is connected to the dedicated, out-of-band management network. Management functions are accessible only through the dedicated management interface. They are not accessible through any other interfaces.

RFC 3871 assumes the following regarding out-of-band management:

- The out-of-band management network is secure

- There is no need for encryption of communication on out-of-band management interfaces

- Security measures are in place to prevent unauthorized physical access

4.1. Vulnerabilities

Although RFC 3871 does not explicitly identify this as a vulnerability, if a router maintains only one dedicated management interface, that interface constitutes a single point of failure. If the dedicated management interface fails, the router will become unmanageable (although it will continue to forward traffic).

Therefore, a router should maintain at least two connections to the management network. Many networks solve this problem by connecting both the dedicated management interface and a terminal server to the out-of-band management network.

4.2. Required Security Capabilities

RFC 3871 states that routers must not forward traffic between dedicated management interfaces and non-management interfaces. The router must never forward a datagram received from a non-management interface through the dedicated management interface. Likewise, the router must never forward a datagram received from the dedicated management interface through a non-management interface.

Operators should refrain from activating dynamic routing protocols on the dedicated management interface. Alternatively, they should rely upon direct or static routes. If static routes are configured, they should be as specific as possible.

4.3. Analysis

Out-of-band management networks isolate network users from communication channels that are dedicated to network management. Therefore, network users cannot access management functions, eavesdrop on management traffic or launch denial of service attacks against the network management plane.

Although the dedicated management interface is somewhat susceptible to misconfiguration, it is less susceptible because its configuration is so simple (i.e., limited to interface definition and a few static routes).

5. Virtual Out-of-Band Access

5.1. Virtual Out-of-Band Access using VPNs

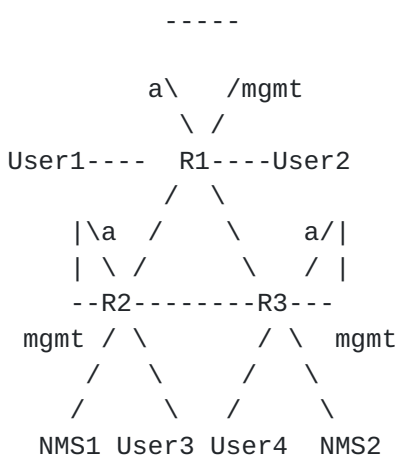

Figure 3: Virtual Out-of-Band Access using VPNs

Figure 3 is identical to Figure 1, except that three loop circuits have been added. Each loop circuit connects an a-end interface to a dedicated management interface. The function of these looping circuits is described below.

In the figure, Routers R1-R3 provide a Layer 3 Virtual Private Network (VPN) [3] service. Although the three routers can support a very large number of Virtual Routing and Forwarding (VRF) instances, for the purpose of example, we will say that they support only two.

The first VRF supports access to the global Internet. This VRF includes interfaces to User1, User2, User3 and User4. It also contains several gateway interfaces to the global Internet (which are

not included in the figure).

The second VRF is dedicated to network management traffic. This VRF includes the interfaces to NMS1 and NMS2, as well as the a-end of each looping interface.

Each router maintains a dedicated management interface that functions exactly as described in Section 4. Management functions are accessible only through the dedicated management interface. They are not accessible any other interfaces.

The dedicated management interface is connected to the a-end of the looping circuit. Therefore, it is accessible only through the management VRF.

Note that the this section describes only one method of constructing a virtual out-of-band management network. An operator could construct a virtual out-of-band management network from in-band pseudowires or an in-band Virtual Private LAN Service. (Layer 3 VPN service is not required.) Likewise, the looping interface need not consume two physical ports. The same results can be achieved with a single, channelized interface or an internal interface.

5.1.1. Vulnerabilities

RFC 3871 is silent regarding virtual out-of-band network management. However, because virtual out-of-band management networks rely upon physically in-band channels, they are susceptible to the following vulnerabilities:

1. Saturation of an in-band trunk can make the device unmanageable.

2. Management traffic may be intercepted. However this would typically require a significant compromise in the routing system.

3. Public interfaces used for management may become unavailable due to attacks, bugs, or similar problems (e.g., buffer overflows being exploited).

Expanding upon the final point, listed above, the virtual out-of-band network can be misconfigured, such that the managed device becomes isolated with regard to the network management stations. Similarly, the in-band network may become unavailable due to attacks, bugs, or other problems. When this happens, operators cannot access the router in order to remedy the configuration error. They become reliant upon physical access to the managed device.

5.1.2. Required Security Capabilities

In order to provide a secure management mechanism, the virtual out-of-band management network must effectively separate the management VPN from all user VPNs. Traffic must never cross from the management VPN to a user VPN or vice versa.

Routers must not forward traffic between dedicated management interfaces and non-management interfaces. The router must never forward a datagram received from a non-management interface through the dedicated management interface. Likewise, the router must never forward a datagram received from the dedicated management interface through a non-management interface.

Operators should refrain from activating dynamic routing protocols on the dedicated management interface. Alternatively, they should rely upon direct or static routes. If static routes are configured, they should be as specific as possible.

Operators may also choose to elevate the priority of management traffic so that it will be preserved during periods of trunk congestion.

5.1.3. Analysis

None of the capabilities mentioned in Section 5.1.2 address the final vulnerability mentioned in Section 5.1.1. Failure of the virtual out-of-band network will render the network unmanageable. This is an inherent weakness of virtual management.

5.2. Virtual Out-of-Band Access using CoS

Some significant separation of management traffic may be achieved by assigning all management traffic to a specific Class of Service (CoS) which is separate from the CoS's used for other (particularly user) traffic. Specific link and other resources may then be assigned to the management CoS. Typically this approach may be combined with either in-band management or virtual out-of-band management using VPNs.

5.2.1. Vulnerabilities

Because virtual out-of-band management networks rely upon physically in-band channels, they are susceptible to the same vulnerabilities discussed in Section 5.1.1 above.

5.2.2. Required Security Capabilities

In order to provide a secure management mechanism using a separate class of service to create a virtual out-of-band capability, the network must effectively separate the management CoS from all user CoSs. User traffic must never be permitted to use the management CoS. This requires that ALL PE devices be capable of ensuring that user traffic entering the network via that PE be mapped to a non-management CoS.

5.2.3. Analysis

Failure of the in-band network will render the network unmanageable. This is an inherent weakness of virtual management using CoS.

6. Evaluation

Based on the analysis above, we conclude that out-of-band management is both more secure and more reliable than either of the other options. However, it is typically more expensive than either of the other options.

When out-of-band management does not offer a feasible economic approach, operators must choose between in-band management with cryptographic protection or a virtual out-of-band management network. In either case, the operator must deal with some additional complexity. So, operators should determine which class of threat (DoS, eavesdropping) poses the greatest risk to their network and choose a strategy accordingly.

7. Security Considerations

Security is the subject matter of this entire memo.

8. Acknowledgements

The authors gratefully acknowledge the contributions of:

- Ross Callon for his comments and suggestions.

This listing is intended to acknowledge contributions, not to imply that the individual or organizations approve the content of this document.

Apologies to those who commented on/contributed to the document and

were not listed.

9. References

9.1. Normative References

[1] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[2] Jones, G., "Operational Security Requirements for Large Internet Service Provider (ISP) IP Network Infrastructure", RFC 3871, September 2004.

[3] Rosen, E., "BGP/MPLS IP VPNs", draft-ietf-l3vpn-rfc2547bis-03 (work in progress), October 2004.

9.2. Informative References

[4] Jones, G., "Framework for Operational Security Capabilities for IP Network Infrastructure", draft-ietf-opsec-framework-01 (work in progress), October 2005.

[5] Callon, R. and G. Jones, "Miscellaneous Capabilities for IP Network Infrastructure", draft-callon-misc-cap-00 (work in progress), October 2005.

Authors' Addresses

Ronald P. Bonica
Juniper Networks
2251 Corporate Park Drive
Herndon, VA 20171
US

Email: rbonica@juniper.net

Syed F. Ahmed
Booz Allen Hamilton
8283 Greensboro Drive
McLean, VA 22102
US

Email: ahmed_syed@bah.com

Intellectual Property Statement

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.

Disclaimer of Validity

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.

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.

Acknowledgment

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