

Network Working Group
Request for Comments: 3013
BCP: 46
Category: Best Current Practice

T. Killalea
neart.org
November 2000

Recommended Internet Service Provider Security Services and Procedures

Status of this Memo

This document specifies an Internet Best Current Practices for the Internet Community, and requests discussion and suggestions for improvements. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2000). All Rights Reserved.

Abstract

The purpose of this document is to express what the engineering community as represented by the IETF expects of Internet Service Providers (ISPs) with respect to security.

It is not the intent of this document to define a set of requirements that would be appropriate for all ISPs, but rather to raise awareness among ISPs of the community's expectations, and to provide the community with a framework for discussion of security expectations with current and prospective service providers.

Table of Contents

1	Introduction	2
1.1	Conventions Used in this Document.	3
2	Communication.	3
2.1	Contact Information.	3
2.2	Information Sharing.	4
2.3	Secure Channels.	4
2.4	Notification of Vulnerabilities and Reporting Incidents.	4
2.5	ISPs and Computer Security Incident Response Teams (CSIRTs).	5
3	Appropriate Use Policy	5
3.1	Announcement of Policy	6
3.2	Sanctions.	6
3.3	Data Protection.	6
4	Network Infrastructure	6
4.1	Registry Data Maintenance.	6
4.2	Routing Infrastructure	7
4.3	Ingress Filtering on Source Address.	7
4.4	Egress Filtering on Source Address	8
4.5	Route Filtering.	8
4.6	Directed Broadcast	8
5	Systems Infrastructure	9
5.1	System Management.	9
5.2	No Systems on Transit Networks	9
5.3	Open Mail Relay.	9
5.4	Message Submission	9
6	References	10
7	Acknowledgements	12
8	Security Considerations.	12
9	Author's Address	12
10	Full Copyright Statement.	13

1 Introduction

The purpose of this document is to express what the engineering community as represented by the IETF expects of Internet Service Providers (ISPs) with respect to security. This document is addressed to ISPs.

By informing ISPs of what this community hopes and expects of them, the community hopes to encourage ISPs to become proactive in making security not only a priority, but something to which they point with pride when selling their services.

Under no circumstances is it the intention of this document to dictate business practices.

In this document we define ISPs to include organisations in the business of providing Internet connectivity or other Internet services including but not restricted to web hosting services, content providers and e-mail services. We do not include in our definition of an ISP organisations providing those services for their own purposes.

This document is offered as a set of recommendations to ISPs regarding what security and attack management arrangements should be supported, and as advice to users regarding what they should expect from a high quality service provider. It is in no sense normative in its own right. In time it is likely to become dated, and other expectations may arise. However, it does represent a snapshot of the recommendations of a set of professionals in the field at a given point in the development of the Internet and its technology.

1.1 Conventions Used in this Document

The key words "REQUIRED", "MUST", "MUST NOT", "SHOULD", "SHOULD NOT", and "MAY" in this document are to be interpreted as described in "Key words for use in RFCs to Indicate Requirement Levels" [RFC2119].

2 Communication

The community's most significant security-related expectations of ISPs relate to the availability of communication channels for dealing with security incidents.

2.1 Contact Information

ISPs SHOULD adhere to [RFC2142], which defines the mailbox SECURITY for network security issues, ABUSE for issues relating to inappropriate public behaviour and NOC for issues relating to network infrastructure. It also lists additional mailboxes that are defined for receiving queries and reports relating to specific services.

ISPs may consider using common URLs for expanded details on the above (e.g., <http://www.ISP-name-here.net/security/>).

In addition, ISPs have a duty to make sure that their contact information, in Whois, in routing registries [RFC1786] or in any other repository, is complete, accurate and reachable.

2.2 Information Sharing

ISPs SHOULD have clear policies and procedures on the sharing of information about a security incident with their customers, with other ISPs, with Incident Response Teams, with law enforcement or with the press and general public.

ISPs should have processes in place to deal with security incidents that traverse the boundaries between them and other ISPs.

2.3 Secure Channels

ISPs SHOULD be able to conduct such communication over a secure channel. Note, however, that in some jurisdictions secure channels might not be permitted.

2.4 Notification of Vulnerabilities and Reporting of Incidents

ISPs SHOULD be proactive in notifying customers of security vulnerabilities in the services they provide. In addition, as new vulnerabilities in systems and software are discovered they should indicate whether their services are threatened by these risks.

When security incidents occur that affect components of an ISP's infrastructure the ISP should promptly report to their customers

- who is coordinating response to the incident
- the vulnerability
- how service was affected
- what is being done to respond to the incident
- whether customer data may have been compromised
- what is being done to eliminate the vulnerability
- the expected schedule for response, assuming it can be predicted

Many ISPs have established procedures for notifying customers of outages and service degradation. It is reasonable for the ISP to use these channels for reporting security-related incidents. In such cases, the customer's security point of contact might not be the person notified. Rather, the normal point of contact will receive the report. Customers should be aware of this and make sure to route such notifications appropriately.

2.5 Incident Response and Computer Security Incident Response Teams (CSIRTs)

A Computer Security Incident Response Team (CSIRT) is a team that performs, coordinates, and supports the response to security incidents that involve sites within a defined constituency. The Internet community's expectations of CSIRTs are described in "Expectations for Computer Security Incident Response" [RFC2350].

Whether or not an ISP has a CSIRT, they should have a well-advertised way to receive and handle reported incidents from their customers. In addition, they should clearly document their capability to respond to reported incidents, and should indicate if there is any CSIRT whose constituency would include the customer and to whom incidents could be reported.

Some ISPs have CSIRTs. However it should not be assumed that either the ISP's connectivity customers or a site being attacked by a customer of that ISP can automatically avail themselves of the services of the ISP's CSIRT. ISP CSIRTs are frequently provided as an added-cost service, with the team defining as their constituency only those who specifically subscribe to (and perhaps pay for) Incident Response services.

Thus it's important for ISPs to publish what incident response and security resources they make available to customers, so that the customers can define their incident response escalation chain BEFORE an incident occurs.

Customers should find out whether their ISP has a CSIRT, and if so what the charter, policies and services of that team are. This information is best expressed using the CSIRT template as shown in Appendix D of "Expectations for Computer Security Incident Response" [RFC2350].

3 Appropriate Use Policy

Every ISP SHOULD have an Appropriate Use Policy (AUP).

Whenever an ISP contracts with a customer to provide connectivity to the Internet that contract should be governed by an AUP. The AUP should be reviewed each time the contract is up for renewal, and in addition the ISP should proactively notify customers as policies are updated.

An AUP should clearly identify what customers shall and shall not do on the various components of a system or network, including the type

of traffic allowed on the networks. The AUP should be as explicit as possible to avoid ambiguity or misunderstanding. For example, an AUP might prohibit IP spoofing.

3.1 Announcement of Policy

In addition to communicating their AUP to their customers ISPs should publish their policy in a public place such as their web site so that the community can be aware of what the ISP considers appropriate and can know what action to expect in the event of inappropriate behaviour.

3.2 Sanctions

An AUP should be clear in stating what sanctions will be enforced in the event of inappropriate behaviour.

3.3 Data Protection

Many jurisdictions have Data Protection Legislation. Where such legislation applies, ISPs should consider the personal data they hold and, if necessary, register themselves as Data Controllers and be prepared to only use the data in accordance with the terms of the legislation. Given the global nature of the Internet ISPs that are located where no such legislation exists should at least familiarise themselves with the idea of Data Protection by reading a typical Data Protection Act (e.g., [DPR1998]).

4 Network Infrastructure

ISPs are responsible for managing the network infrastructure of the Internet in such a way that it is

- reasonably resistant to known security vulnerabilities
- not easily hijacked by attackers for use in subsequent attacks

4.1 Registry Data Maintenance

ISPs are commonly responsible for maintaining the data that is stored in global repositories such as the Internet Routing Registry (IRR) and the APNIC, ARIN and RIPE databases. Updates to this data should only be possible using strong authentication.

ISPs should publicly register the address space that they assign to their customers so that there is more specific contact information for the delegated space.

4.2 Routing Infrastructure

An ISP's ability to route traffic to the correct destination may depend on routing policy as configured in routing registries [RFC1786]. If so, and if the registry supports it, they should ensure that the registry information that they maintain can only be updated using strong authentication, and that the authority to make updates is appropriately restricted.

Due care should also be taken in determining in whose routing announcements you place greater trust when a choice of routes are available to a destination. In the past bogus announcements have resulted in traffic being 'black holed', or worse, hijacked.

BGP authentication [RFC2385] SHOULD be used with routing peers.

4.3 Ingress Filtering on Source Address

The direction of such filtering is from the edge site (customer) to the Internet.

Attackers frequently cover their tracks by using forged source addresses. To divert attention from their own site the source address they choose will generally be from an innocent remote site or indeed from those addresses that are allocated for private Internets [RFC1918]. In addition, forged source addresses are frequently used in spoof-based attacks in order to exploit a trust relationship between hosts.

To reduce the incidence of attacks that rely on forged source addresses ISPs should do the following. At the boundary router with each of their customers they should proactively filter all traffic coming from the customer that has a source address of something other than the addresses that have been assigned to that customer. For a more detailed discussion of this topic see [RFC2827].

There are (rare) circumstances where ingress filtering is not currently possible, for example on large aggregation routers that cannot take the additional load of applying packet filters. In addition, such filtering can cause difficulty for mobile users. Hence, while the use of this technique to prevent spoofing is strongly encouraged, it may not always be feasible.

In these rare cases where ingress filtering at the interface between the customer and the ISP is not possible, the customer should be encouraged to implement ingress filtering within their networks. In general filtering should be done as close to the actual hosts as possible.

4.4 Egress Filtering on Source Address

The direction of such filtering is from the Internet to the edge site (customer).

There are many applications in widespread use on the Internet today that grant trust to other hosts based only on ip address (e.g., the Berkeley 'r' commands). These are susceptible to IP spoofing, as described in [CA-95.01.IP.spoofing]. In addition, there are vulnerabilities that depend on the misuse of supposedly local addresses, such as 'land' as described in [CA-97.28.Teardrop_Land].

To reduce the exposure of their customers to attacks that rely on forged source addresses ISPs should do the following. At the boundary router with each of their customers they should proactively filter all traffic going to the customer that has a source address of any of the addresses that have been assigned to that customer.

The circumstances described in 4.3 in which ingress filtering isn't feasible apply similarly to egress filtering.

4.5 Route Filtering

Excessive routing updates can be leveraged by an attacker as a base load on which to build a Denial of Service attack. At the very least they will result in performance degradation.

ISPs should filter the routing announcements they hear, for example to ignore routes to addresses allocated for private Internets, to avoid bogus routes and to implement "BGP Route Flap Dampening" [RFC2439] and aggregation policy.

ISPs should implement techniques that reduce the risk of putting excessive load on routing in other parts of the network. These include 'nailed up' routes, aggressive aggregation and route dampening, all of which lower the impact on others when your internal routing changes in a way that isn't relevant to them.

4.6 Directed Broadcast

The IP protocol allows for directed broadcast, the sending of a packet across the network to be broadcast on to a specific subnet. Very few practical uses for this feature exist, but several different security attacks (primarily Denial of Service attacks making use of the packet multiplication effect of the broadcast) use it. Therefore, routers connected to a broadcast medium MUST NOT be configured to allow directed broadcasts onto that medium [RFC2644].

5 Systems Infrastructure

The way an ISP manages their systems is crucial to the security and reliability of their network. A breach of their systems may minimally lead to degraded performance or functionality, but could lead to loss of data or the risk of traffic being eavesdropped (thus leading to 'man-in-the-middle' attacks).

It's widely accepted that it's easier to build secure systems if different services (such as mail, news and web-hosting) are kept on separate systems.

5.1 System Management

All systems that perform critical ISP functions such as mail, news and web-hosting, should be restricted such that access to them is only available to the administrators of those services. That access should be granted only following strong authentication, and should take place over an encrypted link. Only the ports on which those services listen should be reachable from outside of the ISP's systems networks.

ISPs should stay up to date for more secure methods of providing services as they become available (e.g., IMAP/POP AUTHorize Extension for Simple Challenge/Response, [RFC2195]).

5.2 No Systems on Transit Networks

Systems should not be attached to transit network segments.

5.3 Open Mail Relay

ISPs should take active steps to prevent their mail infrastructure from being used by 'spammers' to inject Unsolicited Bulk E-mail (UBE) while hiding the sender's identity [RFC2505]. While not all preventive steps are appropriate for every site, the most effective site-appropriate methods should be used.

ISPs should also strongly encourage their customers to take the necessary steps to prevent this activity on their own systems.

5.4 Message Submission

Message submissions should be authenticated using the AUTH SMTP service extension as described in the "SMTP Service Extension for Authentication" [RFC2554].

SMTP AUTH is preferred over IP address-based submission restrictions in that it gives the ISP's customers the flexibility of being able to submit mail even when not connected through the ISP's network (for example, while at work), is more resistant to spoofing, and can be upgraded to newer authentication mechanisms as they become available.

In addition, to facilitate the enforcement of security policy, it is strongly recommended that messages be submitted using the MAIL SUBMIT port (587) as discussed in "Message Submission" [RFC2476], rather than through the SMTP port (25). In this way the SMTP port (25) can be restricted to local delivery only.

The reason for this is to be able to differentiate between inbound local delivery and relay (i.e., allow customers to send email via the ISP's SMTP service to arbitrary receivers on the Internet). Non-authenticated SMTP should only be allowed for local delivery.

As more and more mail clients support both SMTP AUTH and the message submission port (either explicitly or by configuring the SMTP port), ISPs may find it useful to require that customers submit messages using both the submission port and SMTP AUTH; permitting only inbound mail on port 25.

These measures (SMTP AUTH and the submission port) not only protect the ISP from serving as a UBE injection point via third-party relay, but also help in tracking accountability for message submission in the case where a customer sends UBE.

6 References

- [CA-95.01.IP.spoofing] "IP Spoofing Attacks and Hijacked Terminal Connections",
ftp://info.cert.org/pub/cert_advisories/
- [CA-97.28.Teardrop_Land] "IP Denial-of-Service Attacks",
ftp://info.cert.org/pub/cert_advisories/
- [DPR1998] The UK "Data Protection Act 1998 (c. 29)",
<http://www.hmsso.gov.uk/acts/acts1998/19980029.htm>
- [RFC1786] Bates, T., Gerich, E., Joncheray, L., Jouanigot, J., Karrenberg, D., Terpstra, M. and J. Yu, "Representation of IP Routing Policies in a Routing Registry (ripe-81++)", RFC 1786, March 1995.

- [RFC1834] Gargano, J. and K. Weiss, "Whois and Network Information Lookup Service", RFC 1834, August 1995.
- [RFC1835] Deutsch, P., Schoultz, R., Faltstrom, P. and C. Weider, "Architecture of the WHOIS++ service", RFC 1835, August 1995.
- [RFC1918] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G. J. and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2142] Crocker, D., "Mailbox Names for Common Services, Roles and Functions", RFC 2142, May 1997.
- [RFC2195] Klensin, J., Catoe, R. and P. Krumviede, "IMAP/POP AUTHorize Extension for Simple Challenge/Response", RFC 2195, September 1997.
- [RFC2196] Fraser, B., "Site Security Handbook", FYI 8, RFC 2196, September 1997.
- [RFC2350] Brownlee, N. and E. Guttman, "Expectations for Computer Security Incident Response", BCP 21, RFC 2350, June 1998.
- [RFC2385] Heffernan, A., "Protection of BGP Sessions via the TCP MD5 Signature Option", RFC 2385, August 1998.
- [RFC2439] Chandra R., Govindan R. and C. Villamizar, "BGP Route Flap Damping", RFC 2439, November 1998.
- [RFC2476] Gellens R. and J. Klensin, "Message Submission", RFC 2476, December 1998.
- [RFC2505] Lindberg, G., "Anti-Spam Recommendations for SMTP MTAs", BCP 30, RFC 2505, February 1999.

- [RFC2554] Myers, J., "SMTP Service Extension for Authentication", RFC 2554, March 1999.
- [RFC2644] Senie, D., "Changing the Default for Directed Broadcasts in Routers", BCP 34, RFC 2644, August 1999.
- [RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, RFC 2827, May 2000.

7 Acknowledgements

I gratefully acknowledge the constructive comments received from Nevil Brownlee, Randy Bush, Bill Cheswick, Barbara Y. Fraser, Randall Gellens, Erik Guttman, Larry J. Hughes Jr., Klaus-Peter Kossakowski, Michael A. Patton, Don Stikvoort and Bill Woodcock.

8 Security Considerations

This entire document discusses security issues.

9 Author's Address

Tom Killalea
Lisi/n na Bro/n
Be/al A/tha na Muice
Co. Mhaigh Eo
IRELAND

Phone: +1 206 266-2196
EMail: tomk@neart.org

10 Full Copyright Statement

Copyright (C) The Internet Society (2000). 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 assigns.

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 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.