Security Expectations for Internet Service Providers

<draft-ietf-grip-isp-expectations-03.txt>

Status of this Memo

   This document is an Internet-Draft and is in full conformance with
   all provisions of Section 10 of RFC2026. Internet Drafts are working
   documents of the Internet Engineering Task Force (IETF), its Areas,
   and its Working Groups. Note that other groups may also distribute
   working documents as Internet Drafts.

   Internet Drafts are draft documents valid for a maximum of six
   months. Internet Drafts 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.

Abstract

   The purpose of this document is to express the general Internet
   community's expectations 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

Internet Draft   Security Expectations for ISPs   1 February 2000

[1] Introduction

   The purpose of this document is to express the general Internet
   community's expectations of Internet Service Providers (ISPs) with
   respect to security. This document is addressed to ISPs.

   By informing ISPs of what the 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.

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 the routing registry [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 or CSIRTs, 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, InterNIC and RIPE databases. Updates to this data should only be possible using strong authentication.

ISPs should 'SWIP' (Shared WhoIs Project) 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 depends on routing policy as configured in the routing registries [RFC1786]. ISPs 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 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 [RFC2267].

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 5.7 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

An SMTP mail server is said to be running as an 'open' mail relay if
it is willing to accept and relay to non-local destinations mail
messages that do not originate locally (i.e., neither the originator
nor the recipient address is local). Such open relays are frequently

used by 'spammers' to inject their Unsolicited Bulk E-mail (UBE)
while hiding their identity. There are only very limited
circumstances in which an administrator can make a justifiable case
for leaving a mail relay on the Internet completely open.

The processes for restricting relaying are well documented. It's
regrettable that some major software vendors ship their Message
Transfer Agents (MTAs) with relaying open by default.

While this is an issue for the whole community, ISPs should be
particularly vigilant in disabling open relaying on mail servers that
they manage because their high-bandwidth connectivity makes them the
preferred injection point for UBE.

ISPs should also strongly encourage their customers to disable open
relaying on their mail servers. Sanctions for running an open mail
relay should be covered in an ISP's AUP.

5.4 Message Submission

To facilitate the enforcement of security policy message submission
should be done through the MAIL SUBMIT port (587) as discussed in
"Message Submission" [RFC2476], rather than through the SMTP port
(25). In addition, message submissions should be authenticated using
the AUTH SMTP service extension as described in the "SMTP Service
Extension for Authentication" [RFC2554]. In this way the SMTP port
(25) can be restricted to local delivery only.

These two measures not only protect the ISP from serving as a UBE
injection point, but also help in tracking accountability for message
submission in the case where a customer sends UBE. Furthermore,
using the Submit port with SMTP AUTH has additional advantages 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.

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.hmso.gov.uk/acts/acts1998/19980029.htm

[RFC1786] Bates, T., Gerich, E., Joncheray, L., Jouanigot, J-M.,
   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", 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", RFC 2196, September
   1997.

[RFC2267] Ferguson, P., and D. Senie, "Network Ingress Filtering:
   Defeating Denial of Service Attacks which employ IP Source
   Address Spoofing", RFC 2267, January 1998.

[RFC2350] Brownlee, N., and E. Guttman, "Expectations for Computer
   Security Incident Response", RFC 2350, June 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.

[RFC2554] Myers, J., "SMTP Service Extension for Authentication",
   RFC 2554, March 1999.

[RFC2644]

Senie, D., "Changing the Default for Directed Broadcasts in
Routers", RFC 2644, August 1999.

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
P.O. Box 81226
Seattle, WA 98108-1226
USA

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

10 Full Copyright Statement

This document expires August 2, 2000.