Internet Engineering Task Force   Tom Killalea
INTERNET-DRAFT   neart.ie
Valid for six months   December1998

Security Expectations for Internet Service Providers

<draft-ietf-grip-isp-07.txt>

Status of this Memo

Copyright Notice

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

1 Introduction

IMPORTANT NOTE:

AS A RESULT OF DISCUSSIONS IN THE GRIP WORKING GROUP MEETING AT THE 43RD IETF, THIS IS THE LAST VERSION OF THIS DOCUMENT THAT YOU WILL SEE. IT WILL BE REPLACED BY 3 NEW DOCUMENTS IN JANUARY 1999. - tomk.

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

A goal is that customers could have a framework that facilitates the discussion of security with prospective service providers; regrettably, such a discussion rarely takes place today.

Additionally, in informing ISPs of what the community hopes and expects of them, a further goal is 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.

This document is addressed to Internet service purchasing decision-makers (customers) and to ISPs.

It has been argued that vendors begin to care about security only when prompted by customers. I hope that this document will encourage both parties to more readily express how much they care about security, and that discussion between the community and its ISPs will be increased.

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 Incident Response

A Security Incident Response Team (SIRT) 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 SIRTs are described in "Expectations for Computer Security Incident Response" [RFC2350].

Whether or not an ISP has a SIRT, 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.

2.1 ISPs and Security Incident Response Teams (SIRTs)

Some ISPs have SIRTs. 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 SIRT. ISP SIRTs 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 to determine what incident response and security resources are available to you, and define your incident response escalation chain BEFORE an incident occurs.

Customers should find out whether their ISP has a SIRT, and if so what the charter, policies and services of that team are. This information is best expressed using the SIRT template as shown in Appendix D of "Expectations for Computer Security Incident Response" [RFC2350]. If the ISP doesn't have a SIRT they should describe what role if any they will take in incident response, and should indicate if there is any SIRT whose constituency would include the customer and to whom incidents could be reported.

2.2 Assistance with Inbound Security Incidents

When a security incident targeting one of their connectivity customers occurs ISPs should inform the customer of the attack. The ISP may also provide assistance to

- trace the 'apparent' origin of the attack and attempt to determine the veracity of each step in the path (keeping in mind that the source address may be spoofed). In cases where the source address is spoofed the ISP could trace the point at which the bogusly addressed traffic entered the ISP's network.

- obtain contact information for the source of the attack using whois [RFC1834 and RFC1835], the DNS [RFC1034 and RFC1035] or relevant common mailbox names [RFC2142].

- collect and protect evidence of the incident and guard against its destruction or unintentional announcement.

If the event continues then, at the customer's request, ISPs may also assist by logging in order to further diagnose the problem, or by filtering certain types of traffic.

2.3 Assistance with Outbound or Transit Security Incidents

In the case where one of their connectivity customers appears to be the source of a security incident an ISP will frequently be contacted. The ISP may facilitate the administrators at the source and the target of the incident getting in contact with each other, normally by passing contact information for the target to the ISP's customer.

An ISP may also be contacted to assist with incidents that traverse their network but use bogus source addresses, such as SYN flooding attacks [CA-96.21.tcp_syn_flooding]. Assistance in this case might consist of using network traces on a hop by hop basis to identify the point at which the bogusly addressed traffic entered the ISP's network. In tracing such incidents it's frequently necessary to coordinate with adjacent ISPs to form a complete chain of response teams along the path of the attack.

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

2.5 Contact Information

[RFC2142] states that sites (including ISPs) must maintain a mailbox called 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 should 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.

2.6 Communication and Authentication

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

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

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, and ISPs should be forthcoming in announcing to the community when such sanctions are imposed.

4 Protection of the Community

ISPs play a crucial role in helping to improve the security of the Internet. This and following sections describe a number of issues which, should they be addressed by ISPs in a coordinated and timely way, would have a very positive effect on the security of the network, and would make it much more difficult for the perpetrators to cover their tracks.

Later sections cover in some detail issues related to specific services such as ingress filtering and open mail relays. Such issues, if addressed by all the ISPs in a concerted way, could have a very positive effect.

4.1 Cooperation

This section is about protecting the community. This requires that we as a community work together to that end. It's worth observing that many of the most significant successes in securing the Internet could not have been achieved by anyone acting alone.

Cooperation may be put on legal ground. For example prior to entering into peering agreements ISPs might specify the steps they will take to cooperatively track security incidents that involve both peers.

4.2 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]).

4.3 Training

It's important that all ISP staff be trained to be security-conscious at all times and to be able to make appropriate use of tools that enhance security. Some issues that they should be particularly aware of include the use of secure channels for confidential information, the risk of attacks that use social engineering, management of data used for authentication, and so on.

4.4 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.

5 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

5.1 Routers

Routers are an excellent platform from which to launch a security attack, as well as being attractive targets of themselves.

Many routers allow an attacker to do dangerous things such as:

- sniff transit traffic
- manipulate routing tables to redirect traffic
- manipulate interface states to disrupt service
- create routing flaps which could potentially cause Denial of Service for large parts of the Internet
- create packets with spoofed addresses, and with any desired flags set
- initiate ICMP packet storms and other Denial of Service attacks
- 'black hole' traffic (e.g., by holding a local route to a null or invalid interface, by holding a local route to an invalid next-hop (one which does not itself have a corresponding route, and does not have a default), or worst yet, by using a dynamic routing protocol to advertise availability of a low-cost route and thus actively drawing traffic toward the black hole)
- launder connections to third-party destinations, facilitated by the router's lack of logging

Such threats are amplified because of the central part in the networking infrastructure that routers occupy, and the large bandwidth frequently available to them.

So access to routers SHOULD be based on one-time passwords or better (e.g., Kerberos) and should be as restricted as possible. Connections to the router should be logged to a different system.

If the router supports different levels of authorisation these levels should be used to restrict privileged access to the router.

Sessions should be encrypted to prevent sessions or data from being stolen and to avoid replay attacks.

Routers should not run the 'small services', which are often enabled by default. These may include bootp, chargen, daytime, discard, echo and finger.

5.2 Switches, Terminal Servers, Modems and other Network Devices

ISPs should be similarly vigilant in their configuration of other network devices. Unfortunately many such devices deployed in the field support only minimal authentication, do authorisation on an all-or-nothing basis and do little or no logging. In the past ISPs have been left with no trail to follow after their switches were reconfigured, their terminal servers were used to launch attacks on third parties or their Uninterruptable Power Supplies were shut down.

Where possible access to such devices should be restricted only to legitimate network administrators.

Network infrastructure devices frequently don't support extensive internal logging because they have no long-term storage, like hosts' hard drives. Many support syslog or SNMP traps however, or at least a short internal event log or debugging mode which can be captured through the console or a through a remote logging session.

Router or switch configurations should always be maintained on a file server, so that they can be restored to previous configuration easily and quickly. These backup configurations should obviously be protected so that they cannot be transferred by unauthorised parties, or overwritten with new bogus configurations.

5.3 Anonymous telnet and other unlogged connections

There are many network devices ranging from low-end routers to printers that accept telnet connections without prompting for a password. Obviously such devices, many of which don't maintain audit trails, are very popular among attackers who wish to cover their tracks.

If ISPs have such devices on their own network they should restrict access to them. In addition, they should encourage their customers to block access to such devices from outside of the customer's network.

The use of telnet to manage network elements is strongly discouraged.

5.4 The Network Operation Centre (NOC) and Network Management

A NOC is a crucial part of an ISP's infrastructure, and should be operated with appropriate regard to security.

A NOC frequently has management control over the configuration information of network devices, and should be vigilant in restricting access to that information. Loading of configuration information into network devices is still frequently done using the TFTP protocol [RFC1350], which not only lacks authentication and uses an insecure channel, but calls for great care in configuration at the server end [CA-91:18.Active.Internet.tftp.Attacks].

A NOC will generally perform a network monitoring function by polling (e.g., with ICMP Echo) a set of network devices periodically. In selecting the set of devices to be polled the crucial role of the devices in 5.2 shouldn't be overlooked.

Beyond simple polling a NOC may also use a network management protocol such as SNMP to communicate with network devices. Usually this will be used to 'get' the value of various variables, such as the number of packets received on a particular interface. However the protocol can be used to 'set' variables, perhaps with serious results (e.g., the device can be reconfigured). In any case, SNMPv1 uses only trivial authentication. Where possible SNMP should be used as a read-only tool to 'get' information from remote devices, and the information gotten should be treated as confidential.

A further use of SNMP is in trap reporting, whereby a management station is notified when certain exceptions occur. This information should also be considered confidential, and the NOC should take care that such trap reporting cannot of itself become a Denial of Service attack.

5.5 Physical Security

The physical security of every installation should be given appropriate consideration. This is particularly so for co-located facilities to which people from different organisations and with different security policies have access.

Three types of co-location arrangements are of particular interest:

- customers co-locating equipment at an ISP's facility

- ISPs co-locating equipment at an external facility with authorised 'remote hands'

- ISPs co-locating equipment at an external facility with no authorised physical access

The first case is most likely to directly concern the customer. If an ISP has a co-location facility for the hosting of customer-owned equipment many issues arise surrounding customer access to their co-located equipment.

Ideally every customer would have a fully enclosed locking 'cage', akin to a small room with walls and ceiling of heavy wire mesh fencing, containing the racks in which their equipment is mounted. Customers are allowed access to their own cage under escort by one of the ISP's employees, or with keys that grant access specifically to their cage.

This assignment of separate cages is expensive in terms of space however, so many ISPs compromise by putting all co-located equipment together in a single machine room, and managing the actions of escorted customers very closely. However this may be insufficient to prevent mishaps such as the accidental disconnection of another customer's equipment. If a single machine room is used then the ISP should provide separate locking cabinets for each co-location customer in preference to an open common area.

A customer should always be supervised while in the physical presence of any equipment that is not their own, and should not be allowed to touch, photograph, or examine equipment belonging to another customer.

Also of concern is layer 2 security of co-located equipment. Customer equipment should not be allowed to share a physical network segment with hosts belonging to anyone else, whether another customer or the ISP themselves. It's common for crackers to exploit weak security or unencrypted remote logins on co-located customer-owned equipment to take control of that equipment and put it into promiscuous listening mode on the local network segment, thereby potentially compromising the privacy and security of any other devices on that segment.

When ISPs co-locate network infrastructure components outside of their own premises, such as at peering points or remote POPs, security considerations are extremely important. These locations often play a pivotal role in the network topology, and may be particular targets for attack or vulnerable to accidents. Equipment should ideally be fully enclosed in locking cabinets or cages to limit physical access. If on-site spares are kept, they should likewise be protected from tampering. Whenever possible, security systems and logging card-swipe locks should be employed. Installations should be inspected periodically for the addition of unauthorised equipment which might be used to 'tap' a connection. As with any other facility, hosts should not be attached to transit segments, and hosts should never have unused physical interfaces attached to network segments.

5.6 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.

The internal routing protocol that an ISP uses should be chosen with security in mind. It should not be configured with the assumption that route recalculations are rare and expensive as this would leave the way open for a Denial of Service attack. Routing updates should use the highest level of authentication supported by the internal routing protocol.

If more specific routes to parts of the ISP's allocated address space are heard from external peers then the ISP needs to be judicious in deciding whether to accept the announcement. Only ISPs who have allowed their CIDR address allocations to become fragmented (by allowing customers to take addressess with them when they change providers) have to face this decision.

5.7 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.

5.8 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.

5.9 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 route dampening 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.

5.10 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 SHOULD NOT be configured to allow directed broadcasts onto that medium.

If it is a packet to which the router would respond if received as a unicast, it MAY send a (single) response. If it is not responding (either because it's not appropriate, or because it's been configured not to) it MAY send an ICMP error. It is also appropriate to silently discard such packets. In any case such packets should be counted to detect possible attempts to abuse this feature.

6 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).

In general a 'horses for courses' approach to the provision of systems services should be adopted (i.e., separate systems should be used to deliver each distinct service). Apart from the benefits that accrue in terms of easing systems administration it's widely acknowledged that it's much easier to build secure systems if different services (such as mail, news and web-hosting) are kept on separate systems.

The services discussed in later sections will all benefit from strong security at a lower layer when IPSec is deployed.

6.1 Policy

An ISP's policy with regard to privacy, authentication, accountability, application of security patches, availability and violations reporting should all be of interest to their customers, and should be published in a public place such as the ISP's web site.

A more detailed discussion of Security Policy can be found in the Site Security Handbook [RFC2196].

6.2 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.

If the ISP provides login accounts to customers then the systems that support this service should be isolated from the rest of the ISP's systems networks.

If applications such as rdist are used for software distribution and synchronisation then they should be used over a secure channel and with strong authentication, for example over Secure Shell (ssh) [SSH1997].

A system should not be attached to transit segments.

If reusable passwords are permitted then users should be educated about how to choose and care for a password, and proactive password checks, password aging and password guessers should be employed.

6.3 Backup

The importance of backups need not be stressed here. However backups can become the weakest link in a system's security if appropriate care isn't taken of backup media.

If backups are done across the network then a secure channel should be used. If volumes are dumped to staging disks during the backup process then access to the images on those staging disks should be as restricted as possible.

Backups take on additional significance as audit data following a security incident.

Ageing backup media should be destroyed rather than discarded.

The customers of a system or service should be informed of what is and is not backed up. Further, if customers have been informed that certain data is not backed up then it should not be backed up.

6.4 Software Distribution

ISPs frequently engage in application software distribution. The integrity of the software should be assured by distributing with it a checksum that has been produced with a strong digest function such as SHA-1 [SHA].

7 Domain Name Service (DNS)

The DNS is critical to the daily activities of millions of Internet users. Regrettably applications have frequently placed blind trust in the information contained in the DNS, and in the availability of the DNS. However prior to DNSSEC [RFC2065] the DNS protocol lacked security, while widely used implementations of the DNS protocol contain further severe vulnerabilities [VIX1995].

While this section indicates some methods in which the DNS can be made more trustworthy and reliable it cannot be stressed too strongly that name based authentication is inherently insecure.

7.1 DNS Server Administration

In addition to issues raised in section 6 ISPs will need to address the following issues in administering their DNS servers:

- Service Monitoring.
  The service availability (ability to answer queries) should be monitored.

- Clock synchronisation.
  Servers should synchronise their clocks using the NTP protocol [RFC1305] with authentication. At least two NTP servers should be used.

7.2 Authoritative Domain Name Service

An Authoritative Server is one that knows the content of a DNS zone from local knowledge, and thus can answer queries about that zone without needing to query other servers. Customers should consider [RFC2182] when choosing secondary DNS servers.

ISPs commonly operate as secondary (or slave) servers for their customers, and these servers may provide service for thousands of zones. Regardless of the number of zones, administrators of these servers should familiarise themselves with the Operational Criteria for Root Name Servers [RFC2010] as a basis for deciding how to provide highly available service. In particular they should follow these guidelines:

- Recursion should be disabled for queries.

- Zone transfer should be restricted.
  Apart from the significant load presented by zone transfer with resultant exposure to Denial of Service attacks, ISPs should recognise that some of their customers will consider the contents of their zone files to be private.

- Performance Monitoring.
  Key variables such as queries per second and average latency should be monitored.

7.3 Resolution Service

ISPs commonly operate DNS resolution service for their customers. In this scenario customers configure their DNS resolver (client) to resolve queries from the ISP's DNS resolution servers. For resolution servers ISPs should follow these guidelines:

- Recursion must be enabled for queries.
  An implication is that ISPs should not use the same servers for resolution service and authoritative DNS service.

- Zone transfer should be disallowed.
  Even though there may be no zones to transfer, allowing zone transfers would expose the servers to Denial of Service attacks.

- Performance Monitoring.
  Key variables such as queries per second and average latency should be monitored. In addition, the hosts generating the highest number of requests should be periodically reported.

- Name server software.
  A name server package should be run that is not vulnerable to server cache poisoning where malicious or misleading data received from a remote name server is cached and is then made available to resolvers that request the cached data.

8 Email and Mail Services

Email has been the target of some of the most widely reported security attacks, as well as thousands of juvenile hoaxes and pranks.

ISPs have a major role in protecting the community from abuse and in educating their customers in appropriate technologies and in appropriate uses of the technology.

8.1 Mail Server Administration

In configuring mail servers ISPs should follow these guidelines:

- Mail software.
  If possible software that uses a separate receiving/sending agent and a processing agent should be used. A goal is that the receiving/sending agent, which interfaces with remote mail servers, can be run with reduced privilege.

- Restrict remote message queue starting.
  On-demand queue runs (to facilitate customers who receive mail at their own domain and don't have permanent connections) should be restricted, preferably using a strong authentication mechanism. Remote message queue starting is implemented using a variety of mechanisms, one of which is the ETRN SMTP service extension as described in [RFC1985].

- Disable VRFY and EXPN.
  No more should be revealed about local users or delivery mechanisms than is necessary.

- Clock synchronisation.
  Servers should synchronise their clocks using the NTP protocol [RFC1305] with authentication. At least two NTP servers should be used.

- Exception Reporting.
  Exceptional conditions such as repeated authentication failures, mail loops and abnormal queue length should be trapped and reported.

- Restrict Access to mail logs.
  Mail logs should only be readable by system administrators.

8.2 Secure Mail

As indicated in 2.6, It's critical that ISPs, and in particular their Security Incident Response personnel, have access to tools that allow them to exchange email securely.

8.3 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.

8.4 Message Submission

To facilitate the enforcement of security policy message submission should be done through the MAIL SUBMIT port (587) as proposed in the work in progress called "Message Submission and Relay", rather than through the SMTP port (25). In addition, message submissions should be authenticated using the AUTH SMTP service extension as described in the work in progess called "SMTP Service Extension for Authentication". 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.

The (undocumented) XTND XMIT POP3 extension which allows clients to send mail through the POP3 session rather than using SMTP may also be considered. It also provides a way to support mobile users at sites where open relaying is disabled, and has the benefit of an authenticated connection and a better audit trail.

8.5 POP and IMAP Services

ISPs who provide POP or IMAP access to mailboxes to their customers should, at a minimum, support the CRAM-MD5 [RFC2195] or APOP [RFC1939] authentication mechanisms. Support for stronger mechanisms should be considered, as should disabling plaintext (user/password) authentication.

9 News Service (NNTP)

As in the case of SMTP, the NNTP protocol [RFC977] used by News suffers from a lack of authentication, so it's trivial to forge news postings. Forgeries can bypass the moderation process, cancel legitimate articles and create havoc for sites that maintain an active file.

The lack of encryption in the protocol and the manner in which many news systems are maintained lead to privacy issues in that it's easy for others to detect what newsgroups and articles you are reading.

9.1 News Server Administration

In configuring news servers ISPs should follow these guidelines:

- News software.
  A news software package should be run that is not vulnerable to maliciously formed news control messages or buffer overflows.

- Disable other services.
  Given news' propensity to consume all available disk space and CPU cycles it's particularly important that news systems do not perform other services.

- Do not interpret batches.
  If incoming batches of articles are supported they should not be fed to a command interpreter.

- Restrict Access to news logs.
  News logs should only be readable by system administrators.

- Authenticate approved headers.
  If possible support for cryptographic authentication of approved messages should be supported, particularly in the case of group control messages.

9.2 Article Submission

As many of the issues relating to open mail relays (8.3) apply to news, ISPs should restrict article submission only to approved customers. Further, the networks from which posting is allowed and the newsgroups to which posting is allowed should be as restricted as possible.

9.3 Control Messages

Control messages attempt to cause the news server to take action beyond filing and passing on the article. Certain control messages, because of the ease with which they can be forged, should be handled with care. While it is up to the ISP to decide whether to take action they must at least propagate control messages even if they do not understand them.

- 'whogets', 'sendsys', 'version' should be ignored by ISPs.

- While 'cancel' messages must be acted on and propagated their sheer volume can sometimes swamp service, and the fact that much of that volume is computer-generated is worrying.

- Systems that require the maintenance of an active file should exercise extreme caution in choosing which if any group control messages (checkgroups, newgroup, rmgroup) will be acted upon.

9.4 Newsfeed Filters

The most obvious form of security problem with news is 'leakage' of articles which are intended to have only restricted circulation. The flooding algorithm is extremely good at finding any path by which articles can leave a subnet with supposedly restrictive boundaries. Substantial administrative effort is required to ensure that local newsgroups remain local [SPE1994].

ISPs who provide customers with the ability to remotely manipulate their inbound filters should use strong authentication for this service.

ISPs should not propagate articles that are excessively crossposted. 10 or more cross-postings is commonly considered to be excessive.

ISPs should impose an upper limit on the article size that they will propagate.

10 Web-hosting Services

Sites frequently choose to out-source the operation and administration of their site to an ISP, and security is often a prominent motivator for doing so. The hosting of such sites and provision of related services is the subject of this section. Further information on the topic can be found in [GAR1997] and [HUG1995].

10.1 Webhosting Server Administration

In addition to issues raised in section 6 ISPs will need to address the following issues in administering their web-hosting servers:

- Service Monitoring.
  The service availability (ability to answer HTTP requests) should be monitored.

- Clock synchronisation.
  Servers should synchronise their clocks using the NTP protocol [RFC1305] with authentication. At least two NTP servers should be used.

- DNS.
  DNS lookups should not be performed on web clients when they connect because they expose the web servers to DNS-based Denial of Service attacks, and they adversely affect performance.

- Process User and Group.
  The web daemon should be run as a user and group that is set up specifically for that purpose, and that user/group should have minimal privilege. This user should be different from the maintainers of the web content.

- DocumentRoot.
  Everything below this directory should be subject to the strictest scrutiny. If possible chroot should be used to change the HTTP daemon's root directory.

- UserDir.

Users other than administrators should not be permitted on the server. If users have accounts then the 'UserDir' directive, if permitted, should not access their private accounts. In particular, scripts should not be permitted to be run from their accounts.

- Partitioning of Virtual Sites.
  A single server that hosts multiple sites (virtual domains) SHOULD be set up such that all data, programs and logs for the different sites are partitioned from each other such that no access to the configuration or data of each other's sites is possible. In addition, it should not be possible to access the data or programs of one customer's site using a URL that has the name of another customer's site in it's host part.

- Access Control.
  Restricted access to certain parts of a site should be facilitated using a strong authentication mechanism such as a certificate or a one-time password device. An alternative is to use well-chosen passwords in conjunction with SSL which at least avoids passwords being passed across the network in plaintext.

- Security Patches and Service Packs.
  The stakes in running a web server are particularly high, so administrators should be particularly vigilant in applying security patches and Service Packs as they are released.

10.2 Server Side Programs

Server side programs such as those that use the Common Gateway Interface (CGI) or other server side interfaces are important to the flexibility of the web as a communications medium. However that flexibility introduces security risks and a weak program might threaten all of the virtual hosts on the server that runs it. An ISP's policy with regard to what programs it will allow is a good indicator of security policy in general.

ISPs should consider the guidelines on server side programs and CGIs:

- Security Policy.
  ISPs should give their customers clear guidelines about how to write secure programs for their hosting environment, and give specific indications about what programming practices will result in a program being rejected.

- Program Installation.
  Customers should not be allowed to install their own programs. All programs and scripts should be submitted to the ISP first to be checked for conformance with security policy. The programs

SHOULD be installed such that only the server administrators have permission to modify them.

- Process User and Group.
  Programs should be run as a user and group that is set up specifically for that purpose, and that user/group should have minimal privilege (many sites use 'nobody').

- Display by Browsers.
  Programs SHOULD never be allowed to be viewed by browsers. One implication of that is that they SHOULD NOT be put under the DocumentRoot.

- Partitioning of Virtual Sites.
  Programs SHOULD NOT be accessible through the site of another customer on the same server, or to the webmaster of that other customer.

- User Input.
  Expressions SHOULD NOT be evaluated based on user input except when used with the equivalent of Perl's tainting features.

- Processing Limit.
  All programs SHOULD have a limit set on real and CPU time, and on the amount of disk space that they can consume.

- Paths.
  All paths SHOULD be full or starting at DocumentRoot, and the PATH variable should be set by the server administrator.

10.3 Data and Databases

Data that is written by server-side programs should be considered confidential. To prevent them being read by browsers their permission should be such that they're not readable by the web daemon process.

If access to a back-end database is provided then programs that facilitate such access should have the least privilege that is absolutely necessary.

Data that relates to state management (cookies) that is stored on the server should be considered confidential and should not be accessible from browsers.

10.4 Logs and Statistics Reporting

The logs generated by the web daemon process can be useful from the security viewpoint in providing an audit trail of site activity, however their more common use is for billing and for market and site analysis.

These logs should be considered highly confidential.

- The only manipulation of them done by the ISP should be that which is necessary to generate billing information and periodically rotate them.

- They should be stored outside of DocumentRoot to prevent access by a browser to them.

- Access to them, whether in raw or summarised format, should be provided to the customer over a secure channel.

10.5 Push and Streaming Services

ISPs frequently provide their customers with the ability to deliver content using protocols other than HTTP. Where such add-on services are provided, both the customer and the ISP should be aware of the security implications of providing such services.

10.6 Commerce

Many ISPs set up the means whereby their customers can sell goods and services through their web-hosted sites. Though a server that can exchange information with a browser over SSL is sometimes described as a 'secure server' this term can be misleading, and ISPs that host commerce applications should consider the following:

- Encrypted Transactions.
  Transactions should never be stored on the server in unencrypted form. Public key cryptography may be used such that only the customer can decrypt the transactions. However even when transactions are passed directly to a financial institution and to the customer some part of the transaction will have to be stored by the ISP for audit trail purposes.

- Transaction Transfer.
  If transactions are not processed immediately but instead are transferred to the customer in batches then that transfer should occur over a secure channel such as SSL and only after strong authentication has taken place. Transaction files should be carefully rotated so that every transaction occurs exactly once.

- Backups.
  If transactions are written to backup media then the physical security of the backup media should be assured.

10.7 Content Loading and Distributed Authoring

The loading of content onto the ISP's server should happen over a secure channel.

If server support for Distributed Authoring tools is enabled, then this should be administered with great care to ensure that strong authentication takes place and that access is given only to the customer's virtual site, and only to that customer's content maintainer.

10.8 Search Engines and other tools

ISPs frequently install tools such as search engines, link checkers and so on for use by their customers. Many such tools create a very great processing overhead when run and so running them on-demand should not be allowed to avoid Denial of Service attacks.

Search engines should be configured so that their searches are restricted to those parts of a site that are available to all.

The output of link checkers should be considered confidential, and should only be available to the content maintainer of the customer's site.

11 References

[CA-91:18.Active.Internet.tftp.Attacks] "Active Internet tftp Attacks", ftp://info.cert.org/pub/cert_advisories/

[CA-95.01.IP.spoofing] "IP Spoofing Attacks and Hijacked Terminal Connections", ftp://info.cert.org/pub/cert_advisories/

[CA-96.21.tcp_syn_flooding] "TCP SYN Flooding and IP Spoofing Attacks", 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

[GAR1997] Garfinkel, S., "Web Security and Commerce", O'Reilly and Associates, Sebastopol, CA, 1997.

[HUG1995] Hughes Jr., L., "Actually Useful Internet Security Techniques", New Riders Publishing, Indianapolis, IN, 1995.

[RFC977] Kantor, B and P. Lapsley, "Network News Transfer Protocol", RFC 977, February 1986.

[RFC1350] Sollins, K. R., "The TFTP Protocol (revision 2)", STD 33, RFC 1350, July 1992.

[RFC1034] Mockapetris, P. V., "Domain names - concepts and facilities", STD 13, RFC 1034, November 1987.

[RFC1035] Mockapetris, P. V., "Domain names - implementation and specification", STD 13, RFC 1035, November 1987.

[RFC1305] Mills, D., "Network Time Protocol (Version 3) Specification, Implementation", RFC 1305, March 1992.

[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.

[RFC1939] Myers, J., and M. Rose, "Post Office Protocol - Version 3", RFC 1939, May 1996.

[RFC1985] De Winter, J. "SMTP Service Extension for Remote Message Queue Starting", RFC 1985, August 1996.

[RFC2010] Manning, B., and P. Vixie, "Operational Criteria for Root Name Servers", RFC 2010, October 1996.

[RFC2065] Eastlake 3rd, D., and C. Kaufman, "Domain Name System Security Extensions", RFC 2065, January 1997.

[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.

[SHA] NIST, FIPS PUB 180-1: Secure Hash Standard, April 1995.

[SPE1994] Spencer, H., "News Article Format and Transmission", ftp://ftp.zoo.toronto.edu/pub/news.txt.Z

[SSH1997] SSH (secure Shell) Remote Login Program, http://www.cs.hut.fi/ssh/

[VIX1995] Vixie, P., "DNS and BIND Security Issues", ftp://ftp.vix.com/pri/vixie/bindsec.psf, 1995.

12 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.

13 Security Considerations

This entire document discusses security issues.

14 Author's Address

Tom Killalea
1516 2nd Ave
Seattle, WA 98101
USA

Phone: +1 206 694-2196
E-Mail: tomk@neart.ie

15 Full Copyright Statement

This document expires June 20, 1999.