



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials Bootcamp Style (Security 401)"
at <http://www.giac.org/registration/gsec>

Determining Contact Information for Internet Addresses

Andrew Daviel
GSEC Practical Version 1.4b
13 August 2003

Abstract

Very often in computer security work an investigator finds one or more Internet addresses and needs to contact an administrator responsible for them. Often this is a trivial exercise, sometimes not. This paper discusses various methods for identifying organizations responsible for a particular IPV4 Internet address.

Contents

- [Abstract](#)
- [Introduction](#)
- [Address Faking - is an address real ?](#)
- [Local and Unroutable Addresses](#)
- [Broadcast and Network Addresses](#)
- [Multicast Addresses](#)
- [Tracing Multicast Addresses](#)
- [DNS - the Domain Name System](#)
- [Lateral Thinking and Name Resolution](#)
- [Numeric Addresses - ARIN and Friends](#)
- [Proxy Services](#)
- [Contact Identification](#)
- [Telephone Contact](#)
- [Message Composition](#)
- [Replies](#)
- [Considerations for Automated Mail](#)

- [References](#)

Introduction

The procedure for contacting an administrator address can be broken down into a number of steps:

- Extracting an address and timestamp from incident data
- Verifying that the address is genuine
- Determining if the address is a routable external (Internet) address
- Resolving numeric addresses to names (and vice-versa) using the Domain Name System
- Identifying an appropriate Domain Registrar
- Querying the registrar's online database for a domain record.
- Recursively querying netblock databases for a network record.
- Identifying an appropriate security contact from the domain and network records.
- Composing and sending a message

This paper discusses each of these steps.

Address Faking - is an address real ?

Addresses given in computer messages can not always be believed. This may be for one of two reasons:

- the address has the appearance of a system-supplied address, but is in actuality provided by a user process.
- the source address of a network packet is not a real address of the source machine.

The first case is common in unsolicited mail (spam) and online fraud. As most people are by now aware, the "From" and "To" addresses in email are totally falsifiable, and in fact are not used by the system for delivery at all. But they are often the only addresses normally displayed to the user. A fake Internet address in email is an extension of this principle.

Consider the following message headers:

```
Received: from mail.nato.org ([39.24.103.45])
    by relay.parliament.uk (8.11.6/8.11.6) with ESMTTP id gB4FNWU19062
    for <prime.minister@parliament.uk>; Mon, 8 Dec 1941 14:34:34 -0000
Received: from mail.whitehouse.gov (mail.whitehouse.gov [198.137.240.92])
    by mail.nato.org (8.11.6/8.11.6) with ESMTTP id gB4FNWU19062
    for <prime.minister@parliament.uk>; Mon, 8 Dec 1941 10:34:34 -0400
Date: Mon, 8 Dec 1941 10:34:33 -0400
```

From: Franklin D. Roosevelt <president@whitehouse.gov>
To: Winston Churchill <prime.minister@parliament.uk>
Subject: Attack on Pearl Harbor exaggerated

You can see that the message appears to come from the mail server mail.whitehouse.gov. So, does Winston Churchill believe that this mail came from Roosevelt ? No - the only reliable thing in this whole message is that it came from 39.24.103.45 - not from mail.nato.org, which is merely the HELO string that 39.24.103.45 used when it connected to relay.parliament.uk. The rest of the message, including the next Received header, is a fabrication generated or relayed by that computer. It is added by a user process, not by the system it pretends to be.

A slightly different method of faking addresses is used in some online fraud schemes. In this, a long URL is generated which appears to be from a trusted organization. To enhance the page appearance, inline images from the real site are placed in the page along with genuine content such as privacy contacts. The fake URL has a form such as <http://www.paypal.com-secure-0012300001233@some.evil.org/pub/fake.html> It relies on the behaviour of current browsers to display URLs starting at the left in a limited-size window. Here the URL appears to be from the address www.paypal.com but is in fact from some.evil.org. (note that "evil.org" is probably an innocent hosting company, and more analysis is required to find the perpetrator - follow the HTML forms, and the money)

In the above two scenarios network analysis would immediately point to the real address; the fake addresses are there to mislead humans.

In the second case, the source address of the message is faked. Simple network analysis gives the fake address.

Fortunately, this is less common than the first case. The reason is that if the source address is incorrect, the sender cannot easily receive replies, and cannot close a TCP handshake [1] in order to generate an SMTP or HTTP transaction, for example.

However, in some kinds of DoS (denial of service) attack, this is not a consideration. The Slammer SQL [3] worm used UDP and multicast addresses and could have propagated using fake unicast source addresses.

This type of fake-source-address attack is outside the scope of this paper. However, fake addresses may not be entirely inscrutable. For instance, a scanning program such as nmap [2] may intersperse a real address (which is able to collect replies) among fake addresses added merely to frustrate analysis. Or, a machine using a fake source address may be able to monitor return traffic - on the same network segment, for example - and thus able to gather scan data. In these cases, analysis of the source addresses may yield some useful information.

Another type of fake address is sometimes used by attack programs. A compromised machine takes over unused or inactive addresses within the local domain. Thus, outgoing packets have a valid real source address, but do not correspond to the resolved name. The compromised machine may be identified from local traffic analysis or from MAC address.

Local and Unroutable Addresses

Within the Internet (IPV4) address space, certain ranges of address have been set aside for

special purposes. Some of these are used within organizations for local routing. It is important to be able to recognize these addresses, both to avoid wasting investigation time and to prevent the sending of spurious complaints to administrative bodies such as ARIN and ICANN.

As described in a number of RFCs (e.g. RFC 1597 [4]), the Internet Assigned Numbers Authority (IANA) has reserved blocks of IPv4 address space for private networks:

```

10.0.0.0      - 10.255.255.255  "24-bit block"  class A      10/8
172.16.0.0   - 172.31.255.255  "20-bit block"  16 x class B 172.16/12
192.168.0.0  - 192.168.255.255 "16-bit" block  255 x class C 192.168/16

```

Address expressions such as 10/8 are a short form of 10.0.0.0/8 or 10.0.0.0/255.0.0.0, i.e. a address block where the top 8 bits are a network address and the bottom 24 bits are a local address. Similarly an address expression 192.168.56/24 refers to the range of addresses from 192.168.56.0 to 192.168.56.255, with a 24-bit network address

The 192.168 addresses are commonly used by small organizations or individuals that use NAT (Network Address Translation). In this scheme, all internal addresses are translated by a gateway machine, which substitutes its own Internet source address on outgoing packets. However, the internal address may appear in email headers, or may appear on the network as a result of a configuration error. These addresses are not supposed to be routed on the Internet, so it is unlikely that they will be seen far from their origin.

The 10/8 addresses have been used by some large ISPs for routing to customers. In this case, the customer has a routable Internet address, but traffic on the ISPs backbone travels over 10/8 addresses. It is unlikely that these addresses would be seen externally on the Internet as source or destination, but they may appear for example in a traceroute toward a customer address. Analysis of the route would be required in order to determine the name of the ISP - this would normally be fairly obvious; only a neophyte sending complaint messages to every address on a traceroute is likely to make this mistake.

Other special use address blocks are defined in RFC 3330 [5]

```

0.0.0.0/8      "This" Network          [RFC1700, page 4]
10.0.0.0/8     Private-Use Networks    [RFC1918]
14.0.0.0/8     Public-Data Networks (X.121) [RFC1700, page 181]
24.0.0.0/8     Cable Television Networks --
39.0.0.0/8     Reserved but subject
                to allocation                               [RFC1797]
127.0.0.0/8    Loopback                [RFC1700, page 5]
128.0.0.0/16   Reserved but subject
                to allocation                               --
169.254.0.0/16 Link Local              --
172.16.0.0/12  Private-Use Networks    [RFC1918]
191.255.0.0/16 Reserved but subject
                to allocation                               --
192.0.0.0/24   Reserved but subject
                to allocation                               --
192.0.2.0/24   Test-Net
192.88.99.0/24 6to4 Relay Anycast      [RFC3068]
192.168.0.0/16 Private-Use Networks    [RFC1918]
198.18.0.0/15 Network Interconnect

```

223.255.255.0/24	Device Benchmark Testing Reserved but subject to allocation	[RFC2544] --
224.0.0.0/4	Multicast	[RFC3171]
240.0.0.0/4	Reserved for Future Use	[RFC1700, page 4]

The "cable television network" addresses are not "special" in the scope of this paper but are used in the normal way by cable modem connections.

The "Public-Data Networks" are listed [6] at IANA. These addresses represent gateways into the X.121 network. They are not much seen in normal work.

Of note are the 169.254 "link local" addresses. These addresses are used for example by Microsoft Windows systems to automatically configure a home network. So they may be seen on a network if used by systems which are unable for some reason to connect with a DHCP server.

Broadcast and Network Addresses

The special address 0.0.0.0 is used to refer to the entire Internet. This address is often seen in firewall and router rules. The corresponding address 255.255.255.255 may be viewed as a broadcast address for the whole Internet (i.e. every system). In practice this can never be seen beyond the source segment.

Within a particular subnet, the highest address is reserved as the broadcast address, while the lowest is used as the network address. Thus for the private network 192.168/16, 192.168.255.255 is the broadcast address and 192.168.0.0 is the network. A broadcast address has the property that packets sent to it will be seen by every host on the subnet, for instance the Unix command "ping -b -c 2 192.168.255.255" will send two ICMP packets to every host. This feature is used by some software for local resource discovery, thus it is often seen on networks. Misconfigured hosts may generate incorrect broadcast addresses; a host with a correct IP address but incorrect broadcast and mask values will appear to function normally for many purposes and may pass unnoticed. Broadcast addresses have also been used in the Smurf [7] DoS attack. Historically many Internet routers passed broadcast addresses; hence the attack could be launched from outside a domain. The use of a fake source address (the victim's) in the ICMP ping would direct all the reply packets at the victim.

Note that "high-speed" network connections as used by some ISPs may use partial class C subnets; hence broadcast addresses such as 24.24.24.252 may be seen.

Domain routers should not normally allow foreign broadcast addresses onto a network, hence any broadcast address should be local. Local network analysis may be necessary to locate a rogue host.

Multicast Addresses

The 224/4 addresses are reserved for multicast use. If broadcast is thought of as one-to-all, and unicast (normal ip traffic) as one-to-one, then multicast is one-to-many. The exact operation of multicast networking is outside the scope of this paper, but essentially multicast traffic is only routed to those machines that have subscribed to it. Most system kernels and TCP stacks are now

multicast-capable, so that an ICMP ping to the multicast address 224.0.0.1 (All Systems on this Subnet) will typically return a significant number of responses.

The one-to-many subscription nature of multicast makes it a natural choice for applications where identical streams of traffic are sent to large numbers of recipients, such as Internet radio, TV, highly popular Web pages etc. However, technical problems (the lack of assured delivery, and historical problems in routing) and political problems (difficulties in assigning responsibility and arranging payment for transit traffic) have so far led to multicast seeing limited use in these areas.

There are a couple of exceptions. Multicast is used by NTP (Network Time Protocol) [8] as the default method of distributing accurate time signals (using address 224.0.1.1). Since NTP is included in Linux and recent Windows systems, external traffic on this address may be seen in cases where a domain does not have a completely managed onsite NTP system.

Multicast is also used by a number of online conferencing systems such as AccessGrid [9] and by a few media broadcast pioneers such as NASA [10]. Multicast is also supported [11] by Windows Media Player in the Windows XP operating system, so that multicast use may be expanded in the future.

Tracing Multicast Addresses

Multicast addresses are not forwarded by normal unicast routing mechanisms; they must be forwarded by a special-purpose multicast router. Historically this was done by the Unix mrouterd daemon; however now this functionality is available in most commercial routers. A multicast domain need not correspond exactly to a unicast domain.

Multicast addresses do not normally indicate the actual source or destinations. However, certain blocks of multicast addresses have been set aside for specific organizations.

The 224.0.0/24 address block is reserved for routing and maintenance protocols, and should not be forwarded beyond the local multicast domain. Hence addresses in this range should only be generated by local machines.

The 239/8 block is reserved for Administratively Scoped [14] addresses - that is, statically assigned by an administrative body. Of these, the 239.255/16 block is reserved for Site-Local use (RFC2365 [12]), and 239.192.000.000-239.251.255.255 is reserved for Organization-Local use. Hence, addresses within these blocks should be assigned and administered within an organization, and should not be forwarded beyond that scope. The presence of unidentified traffic from the Internet on these addresses may indicate a misconfigured multicast router.

The 233/8 block is reserved for GLOP addressing [13]. Each Autonomous System (AS) on the (IPV4) Internet is allocated 254 addresses according to the following scheme:

```
Bits 31-25 : 233
Bits 24-8  : 16-bit AS number
Bits 7-0   : local use
```

The AS may be extracted from the address and looked up in a network database; for instance,

whois.radb.net.

DNS - the Domain Name System

In the early days of the Internet, computer addresses were either given numerically or looked up in a file (the hosts table). In 1987 the current Domain Name System (DNS) was introduced [15][16] which allows distributed lookup of name and address information. The DNS permits the now familiar world of "dot.com" and the Web to function by translating website names to machine addresses. The DNS is currently administered by a number of different private companies, government departments and other organizations which are responsible for different domains. Organizations and ISPs typically run a nameserver for the benefit of their users; these have a copy of the list of root nameservers [17] which is used as a starting point. Searches progress recursively down the tree following assignments to other servers until an address is found. Generally, this system works well - if it does not, the web and email traffic would not be delivered and there would not be an incident to investigate. However, it is possible for a nameserver to be compromised or poisoned, in which case a name may be hijacked to point to a rogue server.

The top-level-domains (TLDs) consist of two-letter domains assigned to countries (using the codes from ISO-3166) plus a number of generic ones. These are [18][19]:

- EDU - Institute of Higher Learning, typically US Universities
- COM - originally US commercial, now general-purpose
- NET - originally for network administration, now somewhat general-purpose
- ORG - originally "other", i.e. not commercial, but now open
- GOV - US government
- MIL - US military
- INT - International organizations, seldom used
- BIZ - Business (since 2000)
- INFO - Informative (since 2000)
- NAME - For individuals (since 2000)
- MUSEUM - Museums
- COOP - Co-operatives
- AERO - Air Transport
- PRO - Professionals (not yet active)

Some country domains have been sold to commercial interests. Examples include TV (Tuvalu),

TO (Tonga), NU (Niue). Also, many country-specific domains include organizations that operate world-wide or wholly outside their nominal country, so that a TLD may be only a "flag of convenience", to borrow a nautical term.

Within some country TLDs, common second-level subdomains may be defined. These may sometimes have their own registry. For instance:

- .AC.UK - United Kingdom academic
- .CO.UK - United Kingdom company (commercial)
- .AC.JP - Japan academic
- .CO.JP - Japan commercial
- .COM.XX - generic commercial

The primary tool to determine contact addresses for a resolved address is the whois service [20]; a client is available on most computers. Historically a search for a .com address would start at whois.internic.net (latterly networksolutions.com), but now typically at whois.crsnic.net or a proxy service which will find the appropriate registrar. Many country-specific TLDs (DO, CH) have a whois server at "whois.nic.TLD" (e.g. whois.nic.do), but some do not follow this convention (e.g. AU). Many lists of whois servers exist on the Web [21]. Some whois clients, such as jwhois [23], contain a list of whois servers and will query the most appropriate server automatically. It is necessary to update these tools periodically as the list of servers changes.

Within the general TLDs (.COM, .NET etc.) it is fairly simple to determine the domain for, e.g., C234.TS.ACME.COM (ACME.COM). Within country-specific TLDs it may not be so obvious. For instance, the domain for PPP23.INT.ACME.LADNER.BC.CA would be ACME.LADNER.BC.CA, not BC.CA or LADNER.BC.CA, which are region-specific subdomains.

Lateral Thinking and Name Resolution

In many cases a valid numeric address will not resolve to a name. This may be because no reverse database exists for the subnet, or because an intruder has used an unlisted address. In some cases a name may be discovered by other means.

If a host is running a webserver, often there may be some contact information on the website. For certain web servers, such as Apache, the server may return a contact address and possibly a local name, in an error page. It may also perform a redirect from a numeric address to a named domain.

If a host is running a mailserver such as sendmail, it may respond with its local name in an SMTP response.

These techniques may be useful to contact a small business, department or home user who has only a small number of IP addresses, in cases where the service provider has not enabled reverse DNS.

The first and last host addresses in a subnet (x.x.x.1 and x.x.x.254) are sometimes used for nameservers or routers, hence if the desired numeric address does not resolve it may be possible to resolve one of these, and hence guess a domain name. It is also possible to iteratively query a whole subnet in search of a resolved name, but extensive use of this technique may be interpreted as a network attack.

Where a DNS record exists for a domain or subnet, it may be possible to extract contact information from the SOA record [16], for example (Unix) "host -t SOA some.org" or "host -t SOA 218.173.65.in-addr.arpa" (note reverse field order). Since the SOA does not allow the "@" character, a dot is used, hence "info.sans.org" is "info@sans.org". Note that these contact addresses are intended for network administration, not abuse reporting. Also, experience shows that such addresses are often out-of-date.

In some cases it may be possible to guess the parent organization by executing a traceroute toward the address in question, and looking at the penultimate address. However, the fact that an organization routes a particular address does not mean that it has any responsibility for it.

Numeric Addresses - ARIN and Friends

If a contact address cannot be found from a domain name or webserver, or if it is suspected that the name DNS has been poisoned, then the address may be looked up in a reverse whois database. Since different chunks of IPV4 address space have been allocated to different countries and authoritative bodies [24], there is no single database. The American Registry of Internet Numbers (ARIN) holds records for most addresses assigned to Canada and the USA, It also has referral entries for addresses held in other databases; the phrase used is typically "these address have been further assigned". Major registries are:

- ARIN - North America
- RIPE - Europe
- APNIC - Asia/Pacific
- KRNIC - South Korea
- JPNIC - Japan
- TWNIC - Taiwan
- LACNIC - Latin America

The syntax of queries and form of records varies somewhat between registries and also within registries depending on the age of the record. Typically a query of the form "whois -h whois.arin.net 36.34.132.1" (on some Linux "fwhois 36.34.132.1@whois.arin.net") will return a record from ARIN. For the JPNIC server, a /e modifier will return a record in English, e.g. "whois -h whois.nic.ad.jp 36.34.132.1/e". Sometimes the contact information is listed in the form of a "handle", e.g.

```
n. [Technical Contact]
```

```
XX123JP
```

In this case the handle may be queried to return contact information, e.g. "whois -h whois.nic.ad.jp XX123JP/e".

Some queries may return multiple records, for example a search of ARIN for 24.45.0.0 gives:

```
Optimum Online (Cablevision Systems) NETBLK-OOL-3BLK (NET-24-44-0-0-1)
    24.44.0.0 - 24.47.255.255
Optimum Online (Cablevision Systems) OOL-6BMMRNNY2-0821 (NET-24-45-32-0-1)
    24.45.32.0 - 24.45.39.255
```

In these cases the required netblock must be looked up, e.g. "whois -h whois.arin.net OOL-6BMMRNNY2-0821"

Proxy Services

A number of proxy services exist on the Internet which may be used to facilitate contact lookup. These include the Geektools whois proxy (whois.geektools.com) and Sam Spade (www.samspade.org).

These services will automatically recurse public whois records in one operation, and are convenient to use. Occasionally, however, some registrars such as Network Solutions may limit automated access to their whois databases (to discourage address extraction for building mailing lists, and to better control access to their data). So these proxies may not always work. There will also often be slower than direct access, since an extra network traversal is required to obtain the information. However, these services are useful for the occasional user.

Contact Identification

Very often there are many contacts listed for an organization. These may include individuals or departments responsible for website design, sales, network operation etc. A large organization, and in particular an ISP, may have a specific contact address for email abuse, and sometimes a separate address for security-related incidents. In some cases the abuse contact is listed in the whois database. If not, then an attempt should be made to find a security contact - the contacts listed in the whois database are often personal addresses for network support personnel who cannot handle a large volume of mail. The online whois database whois.abuse.net may be used to search for contact addresses. Note that this may list the unverified address "postmaster" if an entry does not exist. As per RFC822, the postmaster address is required for all domains with email service[26], but it may not be actively monitored, and is used for mail service administration not security.

If no specific security contact can be identified, then a contact from the whois record may be used. Often, there are several entries. In the case of a small business with a simple Internet connection, the "administrator" address is typically assigned to the business (hence is often a good choice) while the "technical contact" is assigned to the ISP. In the case of a large business, the "administrator" address may be assigned to management while the "technical contact" is assigned to the IT department (hence that may be the best choice). It is often clear from the domain part of the email address.

RFC 2182[25] defines the role account ABUSE for reporting "inappropriate public behaviour" to

customer relations and the account SECURITY for "security bulletins or queries". However, many more organizations implement the "abuse" account than implement the "security" one, so that in general "abuse" is a better guess for reporting problems. Some organizations use a different name, for instance:

- internet.abuse@shaw.ca
- abuse-nonverbose@qwest.net (shorter autoreplies)
- tosusetnet@aol.com
- netsec@att.net
- complaints@direct.ca

Telephone Contact

Sometimes an email message is not enough, you need to talk to a person (though it may be easier to block an attack at your own upstream than to locate the attacker). The procedure is the same as for finding email contacts, except that records are searched for telephone numbers.

If there is no response from an administrator, or if they are unconcerned, then it may be necessary to contact a parent organization or network provider. In some cases, the technical contact given in the whois record will be for a hosting company or ISP who may be more responsive. If this is not satisfactory, then the netblock database can be queried for the address in question and also for its parent (at ARIN, using the command "whois -h whois.arin.net '<SOME-BLOCK'", or by querying less specific network addresses e.g. 10.1.2.3, 10.1.2.0, 10.1.0.0).

Another technique is to do a traceroute toward the address and query for addresses in reverse order. While the routing organizations may have no responsibility for the end address, they may be able to monitor traffic or perhaps filter an attack if it is in violation of their contracts.

Note that ISPs and other companies will not release customer details without legal process. This varies between jurisdictions but may require a court order or faxed release paper.

As for email messaging, be polite and not accusative on the telephone. It is unlikely that the person at the other end was personally responsible for your trouble, and since many database records are outdated, may have nothing to do with the situation.

Message Composition

Many ISPs have reporting guidelines for email to their abuse department. In general, messages should:

- not contain attachments
- be polite and non-accusatory
- briefly explain the source of the contact address, unless it is a well-known role account

- give the exact time and timezone of the incident
- give the original form of any ip address (numeric or name) as found in data (typically numeric in IDS logs and email headers, named domains in Web pages, email bodies)
- give the exact time that any address was resolved to a number or name
- not contain extraneous information, such as whois records, traceroute logs etc.
- give a valid return address and phone number
- for email (unsolicited mail, viruses), contain all the original mail headers
- for Usenet (news), should contain all the original NNTP headers
- for IDS logs, should contain both source and destination addresses

Many ISPs and other organizations use dynamically allocated addresses from DHCP or PPP (dialup). These may change quite rapidly and be used by a number of different people in the course of a day. Therefore, it is important to record accurate timestamps from incident data. This may conveniently be done using NTP[8] to synchronize system clocks to an accurate source of time such as the US Naval Observatory or a GPS receiver. If an accurate time was not available at the time of the incident, timestamps may be corrected within the limits of clock drift by subtracting the difference between system time and NTP-derived time, provided that the system clock has not been tampered with.

Additionally, technologies such as Active Directory from Microsoft may dynamically update[27] DNS records when mobile or portable devices are moved within a domain. It is therefore helpful if addresses can be resolved at the time of the incident. If this is not possible, for instance in an IDS because of traffic volume, then the time at which the address was resolved should be given.

It is important that timezone information is given. Single timestamps may be given in a familiar form such as in email (RFC822[26]) or HTTP (e.g. "Thu, 14 Aug 2003 00:19:33 GMT"). If there are multiple records with a timestamp in local time, the date and timezone of the first one should be stated explicitly. The Unix Gnu date utility may conveniently be used to convert a variety of formats to either local time or to UTC (Coordinated Universal Time, also known as GMT - Greenwich Mean Time - or Zulu time). For instance: "date -u -d '01/02/03 3:12pm EST'" gives "Thu Jan 2 20:12:00 UTC 2003".

Replies

It is very common to receive no reply from messages, or to receive only an automated response. This does not necessarily mean that your message has been lost or ignored. Large organizations and ISPs receive a substantial amount of mail and may not have the resources to have a human answer every message.

It is sometimes frustrating to receive only a generic reply to what may have taken hours of investigative time. However, in many jurisdictions ISPs and other organizations are constrained by privacy legislation not to reveal any details about customers or employees, so "we have

handled the incident in accordance with our AUP" may be the best that you can expect. On the other hand, you may see personal replies such as "We found the machine and are cleaning it. Thank you for your efforts".

Considerations for Automated Mail

Often, a large number of minor security incidents (e.g. email or network worms) may be encountered. These are usually not serious enough to warrant manual investigation. On the other hand, each source address typically represents a vulnerable machine that may continue to spread infection, and may be used for other unauthorized purposes such as relaying unsolicited mail or storing illicit material. It is desirable therefore that the machine owner be informed of the incident.

Proxy reporting services, such as Dshield [\[22\]](#), exist for this purpose. IDS and personal firewall logs may be processed and mailed to the proxy service, which consolidates reports from a large number of contributors and sends a single report to the security contact of record.

For some organizations, policy may prohibit revealing security incidents to a third party, or the log formats are not in a suitable form for the reporting agent, or the type of incident is outside the scope of the proxy service. In these cases, an in-house automated reporting agent may be used.

When sending automated email, it is important not to send excessive amounts of mail to any one recipient. A hundred email viruses or portscans from one source address in a single day should generate one or two reports, not a hundred. There must be a mechanism to count messages and either suppress duplicates or generate digests.

It is also especially important to report only to appropriate addresses. In some cases, contact information in whois databases has not been updated and the person is no longer associated with the address. (This is more common with numeric whois records.) It is also important to report only routable Internet addresses, i.e. not send mail to ICANN or IETF, and to recurse whois queries properly (not to send mail to ARIN, Network Solutions etc.). This is sometimes a problem in country-specific TLDs - an automated lookup for WWW.ACME.COM.XX may find an entry for COM.XX instead of ACME.COM.XX. One approach is to use the most specific subdomain for which an MX record exists.

The reporting mechanism used at TRIUMF uses a local database of contact addresses, augmented by online search of whois.abuse.net, domain whois databases, and reverse address whois databases. The algorithm used to determine an address is essentially the following:

- If there is a contact listed for the numeric address, use it.
- Resolve the numeric address to a name
- If there is a contact listed for the named address, use it
- If there is a contact listed in whois.abuse.net, use that
- If there is an abuse or security address given in the domain whois record, use that

- Send mail to the nominal "abuse" address for the domain
- If this mail is returned, resend it to a contact from the numeric (netblock) whois database
- If this mail is returned, give up

All contact addresses are checked against the private database; this allows stale addresses from whois to be translated

Typically, automated mail will return a large number of "message undeliverable" replies, while mail to ISPs will commonly return automatic replies. It is common to handle these by means of a mail filter or by using a "donotreply" sending address. It is important that any automated message state clearly that it is an automatically generated message, and that it contain a real email address and phone number, so that the recipient may report contact address errors or start a human-to-human dialog.

It may be useful to insert a tracking number into the subject line and also into the body of the message. Some mail handling systems will quote the subject in a reply, others will quote a part of the message but use a generic subject. However, other systems will send a completely generic response which is difficult to track.

The following is a typical message generated by the TRIUMF system in response to an SQL portscan. Note that the addresses and organization names have been sanitized; a 10/8 address would not have generated a real message.

Date: Mon, 27 Jan 2003 22:49:13 -0800
From: Andrew Daviel <security@triumf.ca>
To: abuse@acme.net
Subject: Scan from 10.71.131.19

Re. port scan from 10.71.131.19 reported 1043736552

This mail was generated automatically to assist in quickly locating malicious mobile code such as Internet worms and trojans.
For more information please see <http://andrew.triumf.ca/reporter.html>

A port scan was detected apparently from 10.71.131.19

Your address abuse@acme.net was obtained from 10.71.131.19@whois.arin.net

This activity may indicate that this computer is infected with the Slammer/Sapphire MS-SQL worm (January 2003)
<http://isc.incidents.org/analysis.html?id=180>

More details about this port and possible infection may be found at
<http://andrew.triumf.ca/cgi-bin/port?udp+1434>
Please forward this information to the machine owner so that they can ensure that no unauthorized or misconfigured programs are running.

Time is maintained by NTP and should be accurate.

Scan report generated from Snort portscan log (www.snort.org)
Source: 10.71.131.19 (10.71.131.19)

Destination port: 1434 () UDP Count: 51 (or more)
Jan 28 06:14:54.99 GMT 10.71.131.19.4935 > 192.0.148.128.1434: udp 376
Jan 28 06:15:38.82 GMT 10.71.131.19.4935 > 192.0.163.80.1434: udp 376
Jan 28 06:16:22.67 GMT 10.71.131.19.4935 > 192.0.178.32.1434: udp 376
Jan 28 06:17:06.56 GMT 10.71.131.19.4935 > 192.0.193.240.1434: udp 376
Jan 28 06:17:50.45 GMT 10.71.131.19.4935 > 192.0.208.192.1434: udp 376
...
Jan 28 06:49:04.18 GMT 10.71.131.19.4935 > 192.0.85.179.1434: udp 376
etc.

--

TRIUMF is a high-energy physics research facility located in Vancouver, Canada
Andrew Daviel Tel. +1-604-222-7376 <http://www.triumf.ca>

References

RFCs are generally available from <http://www.ietf.org/rfc/rfc###.txt>, where ### is the RFC number.

1. Postel, J. TCP protocol. [RFC 761](#). January 1980.
- 1(a). Cerf, V., and Kahn, R. A Protocol for Packet Network Intercommunication. IEEE Transactions on Communications Vol. COM-22, No. 5, pp 637-648. May 1974.
2. Fyodor. The Art of Port Scanning. Phrack Magazine Volume 7, Issue 51 September 01, 1997. URL: www.insecure.org/nmap/p51-11.txt (13 August 2003).
- 2.(a) Fyodor. Nmap. URL: www.insecure.org/nmap (13 August 2003).
3. InterNetStormCenter. Port 1434 MS-SQL Worm. January 2003. URL: isc.incidents.org/analysis.html?id=180 (13 August 2003).
4. Rekhter, Y. et al. Address Allocation for Private Internets. [RFC 1597](#). March 1994.
5. IANA. Special-Use IPv4 Addresses. [RFC 3330](#). September 2002
6. IANA. Public Data Network Numbers. June 2001. URL: www.iana.org/assignments/public-data-network-numbers (13 August 2003).
7. CERT. Advisory CA-1998-01 Smurf IP Denial-of-Service Attacks. January 1998. URL: www.cert.org/advisories/CA-1998-01.html (13 August 2003).
8. Mills, David. NTP - Network Time Protocol (version2) Specification and Implementation. [RFC 1119](#). September 1989. URL: www.ietf.org/rfc/rfc1119.ps (13 August 2003).
9. AccessGrid Project. Argonne National Laboratory. URL: www.accessgrid.org (13 August 2003).
10. Brutzman, Don and Macedonia, Mike. MBONE, the Multicast BackBONE. 1995. URL: www-mice.cs.ucl.ac.uk/multimedia/projects/mice/mbone_review.html (13 August 2003).
11. Renous, Raphael. MSMQ3.0 for Windows XP White Paper. Microsoft. April 2001. URL: www.microsoft.com/msmq/MSMQ3.0_whitepaper_draft.doc (13 August 2003).

12. Meyer, D. Administratively Scoped IP Multicast. [RFC 2365](#). July 1998
13. Meyer, D. and Lothberg, P. GLOP addressing in 233/8. [RFC 2770](#). February 2000
14. Meyer, D. Administratively Scoped IP Multicast. [RFC 2365](#). July 1998
15. Mockapetris, P. Domain Names - Concepts and Facilities. [RFC 1034](#). November 1987
16. Mockapetris, P. Domain Implementation and Specification. [RFC 1035](#) November 1987.
17. InterNIC. Root Name Servers. Nov 5, 2002. URL: <ftp://ftp.internic.net/domain/named.cache> (13 August 2003).
18. Postel, J. Domain Name System Structure and Delegation. [RFC 1591](#). March 1994
19. InterNIC. FAQs on New Top-Level Domains. 24-Jan-2002. URL: www.internic.net/faqs/new-tlds.html (13 August 2003).
20. Williamson, S. Transition and Modernization of the Internet Registration Service. [RFC 1400](#). March 1993.
21. Power, Matt. List of Internet whois servers. March 2000. URL: <ftp://rtfm.mit.edu/pub/whois/whois-servers.list> (13 August 2003).
22. Euclidian Consulting. Dshield - Distributed Intrusion Detection System. URL: www.dshield.org (13 August 2003).
23. Oberg, Jonas. jwhois, Whois client. Free Software Foundation. Dec 1999. URL: www.gnu.org/software/jwhois (13 August 2003).
24. Gerich,E. Guidelines for Management of IP Address Space. [RFC 1466](#). May 1993.
25. Crocker,David. Mailbox Names for Common Services, Roles and Functions. [RFC 2142](#). May 1997.
26. Crocker,David. Standard for ARPA Internet Text Messages. [RFC 822](#). August 1982.
27. Microsoft. Configuring Dynamic Update and Secure Dynamic Update. URL: www.microsoft.com/technet/prodtechnol/windowsserver2003/proddocs/deployguide/dnsbc_dhc_oold.asp (13 August 2003).

Andrew Daviel
GSEC Practical Version 1.4b
13 August 2003

Upcoming Training

Click Here to
{Get CERTIFIED!}



San Diego Fall 2017 - SEC401: Security Essentials Bootcamp Style	San Diego, CA	Oct 30, 2017 - Nov 04, 2017	vLive
SANS Seattle 2017	Seattle, WA	Oct 30, 2017 - Nov 04, 2017	Live Event
SANS San Diego 2017	San Diego, CA	Oct 30, 2017 - Nov 04, 2017	Live Event
SANS Gulf Region 2017	Dubai, United Arab Emirates	Nov 04, 2017 - Nov 16, 2017	Live Event
Community SANS Colorado Springs SEC401~	Colorado Springs, CO	Nov 06, 2017 - Nov 11, 2017	Community SANS
SANS Miami 2017	Miami, FL	Nov 06, 2017 - Nov 11, 2017	Live Event
Community SANS Vancouver SEC401^	Vancouver, BC	Nov 06, 2017 - Nov 11, 2017	Community SANS
SANS Sydney 2017	Sydney, Australia	Nov 13, 2017 - Nov 25, 2017	Live Event
SANS Paris November 2017	Paris, France	Nov 13, 2017 - Nov 18, 2017	Live Event
Community SANS St. Louis SEC401	St Louis, MO	Nov 27, 2017 - Dec 02, 2017	Community SANS
Community SANS Portland SEC401	Portland, OR	Nov 27, 2017 - Dec 02, 2017	Community SANS
SANS San Francisco Winter 2017	San Francisco, CA	Nov 27, 2017 - Dec 02, 2017	Live Event
SANS London November 2017	London, United Kingdom	Nov 27, 2017 - Dec 02, 2017	Live Event
SANS Khobar 2017	Khobar, Saudi Arabia	Dec 02, 2017 - Dec 07, 2017	Live Event
SANS Munich December 2017	Munich, Germany	Dec 04, 2017 - Dec 09, 2017	Live Event
Community SANS Ottawa SEC401	Ottawa, ON	Dec 04, 2017 - Dec 09, 2017	Community SANS
SANS Austin Winter 2017	Austin, TX	Dec 04, 2017 - Dec 09, 2017	Live Event
SANS vLive - SEC401: Security Essentials Bootcamp Style	SEC401 - 201712,	Dec 11, 2017 - Jan 24, 2018	vLive
SANS Bangalore 2017	Bangalore, India	Dec 11, 2017 - Dec 16, 2017	Live Event
SANS Cyber Defense Initiative 2017	Washington, DC	Dec 12, 2017 - Dec 19, 2017	Live Event
SANS Cyber Defense Initiative 2017 - SEC401: Security Essentials Bootcamp Style	Washington, DC	Dec 14, 2017 - Dec 19, 2017	vLive
Community SANS Nashville SEC401^	Nashville, TN	Jan 08, 2018 - Jan 13, 2018	Community SANS
SANS Security East 2018	New Orleans, LA	Jan 08, 2018 - Jan 13, 2018	Live Event
Community SANS Hawaii SEC401	Honolulu, HI	Jan 08, 2018 - Jan 13, 2018	Community SANS
Mentor Session - SEC401	Memphis, TN	Jan 09, 2018 - Mar 13, 2018	Mentor
Northern VA Winter - Reston 2018	Reston, VA	Jan 15, 2018 - Jan 20, 2018	Live Event
SANS Amsterdam January 2018	Amsterdam, Netherlands	Jan 15, 2018 - Jan 20, 2018	Live Event
Mentor Session - SEC401	Minneapolis, MN	Jan 16, 2018 - Feb 27, 2018	Mentor
Las Vegas 2018 - SEC401: Security Essentials Bootcamp Style	Las Vegas, NV	Jan 28, 2018 - Feb 02, 2018	vLive
SANS Las Vegas 2018	Las Vegas, NV	Jan 28, 2018 - Feb 02, 2018	Live Event
SANS Miami 2018	Miami, FL	Jan 29, 2018 - Feb 03, 2018	Live Event