Chronology of Virus from the MIT Perspective

Jon Rochlis *jon@bitsy.*mi?*edu*

3 PST ?

The first posting mentioning the virus was by Peter Yee (Nasa Ames) at 8:28pm est on Wednesday to the. `tcp-ip` mailing list. Peter stated that **UCB, UCSD, LLNL,** Stanford, and NASA Ames had been attacked, and described the use of sendmail to pull over the virus, including the x* files found in /usr/tmp. The virus was observed to send vax and sun binaries, have **DES** tables built in, and made some use of *.rhosts* and *hosts.equiv* files. A Berkeley'extension was given and Phil Lapsley and Kurt Pires were listed as being knowledgable about the virus.

At 3:10am the first notice of the virus at **MIT** was posted at AMT by Pascal Chesnais (*lacsap@media-lab.mit.edu*). The motd on ***media-lab*** read:

```
--- lacsap Nov 3 1988 03:10am
DO NOT CALL THE GARDEN. IF YOU WANT TO PROTECT YOUR MACHINE TURN OFF SENDMAIL
OR JUST TURN YOUR MACHINE OFF, OR UNPLUG IT FROM THE NETWORK!!!! DO NOT CALL
THE GARDEN"!!!
```

Pascal had spotted the virus earlier but assumed it was just "a local run away program". The group at AMT figured out after midnight that it was a virus and it was coming in via mail. The response was to such down infected machines. The network groups monitoring information shows the media **lab** gateway. first went down at 11:40pm Wednesday, but was back up by 3:00am. Pascal requested that the Network group isolate the building during the Thursday 11:30pm and it remained so Isolated until Friday at 2:30pm.

Pascal now reports that logs on *media-lab* show several scattered `attempts ttloop: peer died: No such file or directory` messages. There were a few every couple of days, several durning the Wednesday afternoon and many starting at 9:48pm. These are caused by opening a telnet connection and immediately closing it: specifically inetd spawns a telnetd, but when telnetd telnetd goes to read from the network, it finds the connection has disappeared. The virus did this in order to determine whether or not to try to infect a target machine.[1] The logs on ***media-lab*** start on October 25th and the following log entries made before the swarm on Wednesday night.

```
Oct 26 15:01:57 media-lab telnetd[23180]: ttloop:  peer died: No such file or
Oct 28 11:26:55 media-lab telnetd[23331]: ttloop:  peer died: No such file or
Oct 28 17:36:51 media-lab telnetd[12614]: ttloop:  peer died: No such file or
Oct 31 16:24:41 media-lab telnetd[18518]: ttloop:  peer died: No such file or
Nov  1 16:08:24 media-lab telnetd[16125]: ttloop:  peer died: No such file or
Nov  1 18:02:43 media-lab telnetd[21889]: ttloop:  peer died: No such file or
Nov  1 18:58:30 media-lab telnetd[24644]: ttloop:  peer died: No such file or
Nov  2 12:23:51 media-lab telnetd[4721]: ttloop:  peer died: No such file or d
Nov  2 15:21:47 media-lab telnetd[13628]: ttloop:  peer died: No such file or
```

[1] The assumption that machines not running a telnetd are not vulernable to attack is quite interesting. It allowed systems like the MIT Project Athena mailhub, *athena.mit.edu,* (on which we perfered to use only *kerberos* authentication), to escape unscathed.

It is not clear whether these represent early testing of the virus, or if they were just truely accidental premute closings of telnet connections. With hindsight we can a telnetd that logged its peer address (even for such error messages) would have been quite useful in tracing the progress and origin of the virus.

At 3:34am est on Thursday, Andy Sudduth from Harvard made his anonymous posting to tcp-Ip. The posting said that a virus might be lose on the Internet and that there were three steps to take to prevent furthur transmission. This included not running fingerd or fixing it not to overwrite the stack when reading its arguments from the net[2], be sure sendmail was compiled without debug, and not to run rexecd.

The posting was make from an Annex terminal server at from Aiken (sp?) Center (?) at Harvard, by teineting the SMTP port of *Iris.brown.edu*. This is obvious since the message was from "foo%bar.apar" and because the last line of the message was "qui\177\177\177", an attempt to get rubout processing out of the brown SMTP server, a common mistake when faking Internet mail.

It was ironic that this posting did almost no good. The path it **took** to get to athena was:

```
Received: by ATHENA.MIT.EDU (5.45/4.7) id AA29119; Sat, 5 Nov 88 05:59:13 EST
Received: from RELAY.CS.NET by SRI-NIC.ARPA with TCP; Fri, 4 Nov 88 23:23:24 P
Received: from ca.brown.edu by RELAY.CS.NET id aa05627; 3 Nov 88 3:47 EST
Received: from iris.brown.edu (iris.ARPA) by cs.brown.edu (1.2/1.00)
        id AA12595; Thu, 3 Nov 88 03:47:19 est
Received: from (128.103.1.92) with SMTP via tcp/ip
          by iris.brown.edu on Thu, 3 Nov 88 03:34:46 EST
```

There was a 20 hour delay before the message escaped from re/ay.cs.net and got to *sri-nic.arpa.* Another 6 hours went by before the message was recived by *athena.mit.edu.* Other site have reported similar delays.

At 5:58am Thursday moming Keith Bostic *bostic@okeefe.berkeley.edu*) made the virus bug fix posting. The message went to the `tcp-ip`, `comp.bugs.4bsd.ucb-fixes`, `news.announce`, and `news.sysadmin`. It supplied the compile without debug fix to sendmail (or patch the debug command to a garbage string), as well as the very wise suggestion to rename cc and *ld*, which was effective since the virus needed to compile and link itself.

Gene Spafford ({*spaf@purdue.edu*}) forwarded this to `nntp-managers@ucbvax.berkeley .edu` at 8:06am. Ted Ts'o (*tytso@athena.mit.edu*) forwarded this to an internal Project Athena hackers list (*watchmakers@athena.mit.edu*) at 10:07. He expressed disbelief ("no, it's not April **Ist"),**and thought we at Athena were safe. Though no production Athena servers were infected serveral private workstations and developement machines were, so this proved overly optimistic

[2] this was a level of detail that only the originator of the virus could have know at that point. To our knowledge nobody had yet identified the finger bug, since it *only* affected certain vaxen, and certainly nobody had discovered its mechanism.

During Thursday morning Ray Hirschfeld (*ray@math.mit.edu*) spotted the virus on the MIT math department suns and shut down the math gateway at 10:15am. It remained down until 3:15pm.

Gene Spafford posted a message at 2:50pm Thursday to a large number of people and mailing lists include `nntp-managers` which is how we saw it quickly at MIT. it warned the virus used *rsh* and looked in *hosts.equiv* and *.rhosts* for more hosts to attack.

Around this time the MIT group in E40 (Project Athena and the Network Group), called Milo Medin (*medin@nslpo.nasa.gov*) and found out much of the above. Many of us had not yet seen the messages. He. pointed out that the virus just loved to attack gateways (found via the routing tables) and remarked that it must have not been effective at MIT were we run our own C Gateway code, not Unix. Milo also informed use that DCA had shut down the mailbridges. He pointed us to the group at Berkeley and Peter Yee specifically.

At about 6pm on Thursday, Ron Hoffmann (*hoffmann@bitsy.mit.edu*) observed the virus attempting to log into a standalone router using the Berkeley remote login protocol; the remote login attempt originated from a machine previously believed immune[3]. The virus was running under the userId of nobody, and it appeared that it had to be attacking through the finger service, the only network service running under that userid. At that point, we called the group working at Berkeley; they confirmed our suspicions that virus was spreading through fingerd.

On the surface, it seemed that fingerd was too simple to have a protection bug similar to the one in sendmail; it was a very short program, and the only `exec` it did involved a hard-coded pathname. A check of the modification dates of both */etc/fingerd* and */usr/ucb/finger* showed that both had been untouched, and both were identical to known good copies located on a read-only filesystem.

Berkeley reported that the attack on finger involved "shoving some garbage at it"; clearly some sort of overrun buffer wound up corrupting something.

Bill Sommerfeld (*wesommer@athena.mit.edu*) guessed that this bug might involve overwriting the saved program counter in the stack frame; when he looked at the source for fingerd, he found that the buffer it was using was located on the stack; in addition, the program used the C library `gets` function, which assumes that the buffer it is given is long enough for the line it is about to read. To verify that this was a viable attack, he then went on to write a program which exploited this hole in benign way.[4]

A `risks` digest came out at 6:52pm. It included a message from Cliff Stoll of Harvard (*Stoll@dockmaster.arpa*) which described the spread of the virus on milnet and suggested that milnet

[3] It was running a mailer with debugging turned off

[4] the test virus sent the string "Bozo!" back out the network connection.

sights might want to remove themselves from the network. **Stoll also** made the wonderful statement, "This is bad news." Other messages were from Spafford, Peter Neumann (*neumann@csl.sri.com*), and Matt Bishop (*bishop@bear.dartmouth.edu*). They described the sendmail propagation mechanism.

In the SIPB office Stan Zanarotti (*srz@lcs.mit.edu*) and Ted Ts'o had managed to get a core dump from the virus running on a machine in the MIT Lab for Computer Science (LCS) as well as the vax binary. Stan and Tim Sheppard ***(shep@ptt./cs.mit.edu***had been dealing with the virus from 11am Thursday over in Tech Square. Their first reaction was to shut down the network by powering off **DELNI's.** By 1pm Tim had verified that no fileshad been modified on ***a//spice./sc.mif.edu***and had installed recompiled sendmail. (Tim also reloaded a root partition from tape, just to ensure that he was running trusted software).

Ted and Stan started attacking the virus. Pretty soon they had figured out the xor encoding of the strings and were manually decoding strings. By 9:00pm Ted had written a program to decode ail the strings and we had the list of strings used by the program, except for the built-in dictionary which was encoded in a different fashion (by setting the meta bit of each character).

At the same time they discovered the ip address of ***ernie.berke/ey.edu,***128.32.137.13, and proceeded to take apart thc **send—message** routine to figure out what it was sending to emie, how often, and if a handshake was involved. Stan told Jon Rochlis <jon@bitsy.mit.edu> in the MIT Network Group of the SIPB group's progress. The people in E40 called Berkeley and reported the finding of ernie's address. Nobody seemed to have any idea why that was there.

About this time a camera crew from **WNEV** Channel 7 (the Boston CBS affiliate) showed up at the office of James D. Bruce (*jdb@delphi.mit.edu*), VP for Information Systems. He called Jeff Schiller and headed over to **E40.** Jeff and JIm were interviewed. **The** 60,000 number of hosts was stated along with an estimate of 10% infection of the 2000 hosts at MIT. The infection rate was a pure guess. The virus was the lead story on the 1lpm news, and we were quite suprised that the real worid would pay that much attention. Pieces of the footage shot then were shown on the CBS morning news (but by that point were were too busy to watch).

Sheppard shows up in **E40,** then punts to Tech Square to check his netwatch data for ernie packets. (The machine with the data had been unplugged from the network.)

Serious decompiiing began at midnight. Stan and Ted came to **E40.**

John Kohl had the virus running **by** 5am and obseived many things. They were confirmed by the decompiling which was almost done.

List times of berkeley conversations and ftp exchanges of source code.

Press conference in E40 at noon. 7 camera crews, tons of print media. Total zoo until 3pm.

Bostic asks for our affilations and if we like the idea of posting bug fixes to the virus (we didl).

The Today show comes to the SIPB office Saturday to find out about "hackers".

MIT Cast of Characters

Media Lab
Pascal Chesnais <lacsap@media-lab.media.mit.edu>

VP Information Services
James D. Bruce <jdb@delphi.mit.edu>

Network Group/Athena/SIPB
Jeff Schiller <jis@bitsy.mit.edu>

Athena/SIPB
Mark Eichin <eichin@athena.mit.edu>

LCS/SIPB
Stan Zanarottl <srz@lcs.mit.edu>

Athena/SIPB
Ted Ts'o .ztysto@athena.mit.edu

Apollo/Athena/SIPB
William Sommerfeld <wesommer@athena.mit.edu>

DEC/Athena/SIPB
John Kohl <jtkohl@athena.mit.edu>

Athena/SIPB
Ken Raebum <raeburn@athena.mit.edu>

Network Group/SIPB
Jon Rochlis <jon@bitsy.mit.edu>

Media Lab

Hal Birkeland <hkbirke@athena.mit.edu>

Network Group
Ron Hoffmann <hoffmann@bitsy.mit.edu>

Athena/SIPB
Richard Basch <probe@athena.mit.edu>

LCS
Tim Sheppard <shep@ptt.lcs.mit.edu>