The paper below is scheduled to appear in
Andrew Herbert & Karen Spärck Jones, editors.
Computer Systems: theory, technology and applications/A tribute to Roger Needham.
Springer-Verlag, 2004. *Monographs in Computer Science.* ISBN 0-387-20170-X

This pre-publication version includes July 2003 corrections of some minor typographical errors.

Slammer: An urgent wake-up call

Jerome H. Saltzer
Massachusetts Institute of Technology
Cambridge, Massachusetts, U.S.A.
1 February 2003

The Slammer worm is an unusually urgent wake-up call, demonstrating as never before the remarkable ease with which an attacker might paralyze the otherwise very robust Internet. Slammer did not quite succeed, because it happened to pick on an occasionally used interface that is not essential to the core operation of the Internet. If Slammer had found a target in a really popular interface, the Internet would have locked up before anyone could do anything about it, and getting things back to even a semblance of normal operation would probably have taken a long time.

How it worked. The basic principle of operation of Slammer was stunningly simple[1]:

1. Discover an Internet port that is enabled in many network-attached computers, and for which a popular listener implementation has a buffer overflow bug that a single, short packet can trigger. IP/UDP ports are thus a target of choice. Slammer exploited a bug in Microsoft SQL Server 2000 and Microsoft Server Desktop Engine 2000, both of which enable the SQL UDP port. This port is used for database queries, and it is vulnerable only on Windows computers that run one of these database packages, so it is by no means universal.
2. Send to that port a packet that overflows the buffer, captures the execution point of the processor, and runs a program contained in the packet.
3. Write that program to go into a tight loop, generating an IP address at random and sending a copy of the same packet to that address, as fast as possible. The smaller the packet, the more packets per second the program can launch. Slammer used packets that were, with headers, 404 bytes long, so a broadband-connected (1 Megabit/second) machine could launch packets at a rate of 300/second, a machine with a 10 Megabits/second path to the Internet could launch packets at a rate of

3,000/second and a high-powered server with an OC-3 (155 Megabits/second) connection might be able to launch as many as 45,000 packets/second.

Forensics: Receipt of this single packet is enough to instantly recruit the target to help propagate the attack to other vulnerable systems. Recruitment modifies no files and leaves few traces, because the worm exists only in volatile memory. If you stop a recruited machine, disconnect it from the Internet, and reboot it, you will find nothing. There may be some counters indicating that there was a lot of outbound network traffic, but no clue why. So one remarkable feature of this kind of worm is the potential difficulty of tracing its source. The only forensic information available is likely to be the payload of the intentionally tiny worm packet.

Exponential attack rate: The second thing that makes this worm significant is how rapidly it increases its aggregate rate of attack. It recruits every vulnerable computer on the Internet as both a prolific propagator and also as an intense source of Internet traffic. The original launcher need merely find one vulnerable machine anywhere in the Internet and send it a single worm packet. This newly-recruited target will immediately begin sending copies of the worm packet to other addresses chosen at random. There are about 4 billion IP addresses, and even though many of them are unassigned, sooner or later one of these worm packets will hit another machine that has the same vulnerability. The worm packet immediately recruits this second machine to help with the attack. The expected time until a worm packet hits yet another vulnerable machine is now half and the volume of attack traffic double. Soon third and fourth machines will be recruited to join the attack; the expected time to find new recruits halves again and the malevolent traffic rate doubles again. This epidemic process proceeds with exponential growth until either a shortage of new, vulnerable targets or bottlenecked network links slow it down; the worm will quickly recruit every vulnerable machine in the Internet[2].

The exponent of growth depends on the average time it takes to recruit the next target machine, which in turn depends on two things: the number of vulnerable targets and the rate of packet generation. If we suppose that the average recruited machine can generate IP addresses and send worm packets at a rate of 1 thousand per second it will hit any one IP address about once every 4 million seconds, or roughly 45 days. At my home, my computer advertises a single IP address, and at the peak I was receiving a worm packet every 80 seconds. Starting with that observation, we can estimate the minimum number of recruits, assuming that the IP address generation mechanism of each worm is independent and memoryless and hits every IP address with equal probability[3]:

observed arrival rate:	$1/80 = 0.0125$ packets/second/IP address
number of IP addresses:	$2^{32} = 4 \cdot 10^9$ IP addresses
aggregate rate:	$.0125 \cdot 4 \cdot 10^9 = 5 \cdot 10^7$ packets/second
assumed rate per recruit:	10^3 packets/second/recruit
number of recruits:	$5 \cdot 10^7 / 10^3 = 50{,}000$

This number is a minimum, because at the peak of the packet storm it is likely that link and router saturation in many parts of the Internet substantially reduced the observed arrival rate. These 50 thousand or more recruits would be launching at least 50 million packets per second into the Internet, and the aggregate extra load on the Internet of these 3200-bit packets probably amounted to something over 150 Gigabits/second, but that is well below the aggregate capacity of the Internet, which is why reported disruptions were localized rather than universal. (Warning: these back-of-the-envelope calculations depend on rough measurements, unconfirmed assumptions and a speculative guess about average rate of packet generation of each recruit. With luck this estimate of the number of recruits may be in the right order of magnitude.)

With 50 thousand vulnerable ports scattered through a space of 4 billion addresses, the chance that any single packet hits a vulnerable port is one in 120 thousand. If the first recruit sends one thousand packets per second, the expected time to hit a vulnerable port would be about two minutes. In four minutes there would be four recruits. In six minutes, eight recruits. In half an hour, nearly all of the 50 thousand vulnerable machines would probably be participating[4].

Extrapolation: The real problem appears if we redo that analysis for a port to which five million vulnerable computers listen: the time scale drops by two orders of magnitude. With that many listeners, a second recruit would receive the worm and join the attack within one second, two more one second later, etc. In less than 30 seconds, most of the 5 million machines would be participating, each launching traffic onto the Internet at the fastest rate they (or their Internet connection) can sustain. This level of attack, about two orders of magnitude greater than the intensity of Slammer, would almost certainly paralyze every corner of the Internet. And it could take quite a while to untangle, because the overload of every router and link would hamper communication among people who are trying to resolve the problem. In particular, it could be very difficult for owners of vulnerable machines to learn about and download any necessary patches.

Prior art: Slammer used a port that is not widely enabled, yet its recruitment rate, which determines its exponential growth rate, was at least one and perhaps two orders of magnitude faster than that reported for the previous generation of fast-propagating worms, Code Red and Nimda[5]. Those worms attacked much more widely-enabled ports, but they took longer to propagate because they used complex multipacket protocols that took much longer to set up. Interestingly, Slammer did not use any of several propagation enhancement techniques suggested by Staniford, et al. Instead, the Slammer attack demonstrates the power of brute force. By choosing a UDP port, infection can be accomplished by a single packet, so there is no need for a time-consuming protocol interchange. And the smaller the packet size, the faster a recruit can then launch packets to discover other vulnerable ports.

Another risk: The worm also revealed a risk of what in the Internet are called class A or CIDR /8 networks. At the time that my computer, which advertises a single IP address, was receiving one Slammer worm packet every 80 seconds, a class C network (which advertises 256 addresses) would have been receiving three packets per second, a class B

network (which advertises 65 thousand addresses) would have been receiving 750 packets/second, and a class A network (which advertises 16 million IP addresses) would have been receiving 200,000 packets/second, with a data rate of about 640 Megabits/second. In confirmation, incoming traffic to the M. I. T. class A network border routers peaked at a measured rate of around 500 Megabits/second with the 155 Megabits/second link to the public Internet saturated.[6] Being the home of 16 million IP addresses has its hazards.

Lessons: From this incident we can draw some important lessons for different Internet participants: For users, the perennial but often-ignored advice to disable unused Internet ports does more than help a single computer resist attack, it helps protect the entire Internet. For vendors, shipping an operating system that by default activates a listener for a feature that the user does not explicitly request is hazardous to the health of the Internet. For implementers, the importance of diligent care in network listener implementations, especially on widely activated UDP ports, just ratcheted up another notch or two.

Acknowledgement: This note benefited greatly from review by and ensuing discussion with Hari Balakrishnan.

Notes.

1. This description of the operation of Slammer is based on a preliminary report found at

Internet Worm W32/SQL/Slammer.worm
McAfee Security Virus Information Library
`http://vil.nai.com/vil/content/v_99992.htm`
(URL verified 30 January 2003)

2. The initial rate of spread up to the point that Internet bandwidth limitations begin to cap it can be described by a well-known formula called the logistic equation, applicable to population growth and epidemics. An analysis of the application of the logistic equation to Internet worm recruitment rate can be found in

Stuart Staniford, Vern Paxson, and Nicholas Weaver
How to own the Internet in your spare time
Proceedings of the 11th USENIX Security Symposium, San Francisco, August 5-9, 2002
`http://www.icir.org/vern/papers/cdc-usenix-sec02/`
(URL verified 30 January 2003)

3. An early report by disassemblers of Slammer indicates that its pseudo-random number generator was defective, and that the equal probability assumption did not apply, at least during the initial propagation of the worm. See

David Moore, et al.
The spread of the Sapphire/Slammer worm

http://www.caida.org/outreach/papers/2003/sapphire/sapphire.html
(URL verified 1 February 2003)

4. These estimates both of the speed of onset and the Slammer worm's relatively mild effect on the Internet as a whole are confirmed by published measurements that show packet loss rates averaged across many servers increasing from near zero to a peak of a little under 20% in less than 30 minutes.

Matrix Net Systems Event Advisories, Slammer Worm Attack
Weekly summary, January 24 through January 30, 2003
http://www.matrixnetsystems.com/ea/index.jsp
(URL verified 1 February 2003)

In their paper cited above, Moore, et al., report observing Slammer to have an initial recruitment rate of 7/minute, about 15 times as fast as our calculation. This observation suggests that there were actually many more vulnerable hosts than estimated here. The alternative explanation, that the generation rate of the average recruit was far higher than 1 thousand packets/second, seems unlikely, though a somewhat higher generation rate may have contributed part of the difference.

5. The above-cited paper by Staniford et al., reported that Code Red had an initial recruitment rate of about 2 recruits/hour. Our lower-bound estimate for Slammer of 0.5 recruits/minute is 15 times greater, and the measurement of Slammer by Moore et al. of 7/minute is 200 times greater.

6. The M. I. T. router traffic statistics were reported in an e-mail message:

From: James D. Bruce (Director of Information Systems)
To: the MIT community,
Date: 28 January 2003 09:51:51 EST
Subject: Weekend Network Outage