CONSTRUCTION AND MANAGEMENT OF A SECURE NETWORK IN SPRING-8

M. Ishii, T. Fukui, M. Kodera, T. Ohata, R. Tanaka
SPring-8, Hyogo, Japan

ABSTRACT

SPring-8, a third-generation synchrotron radiation facility, accepts many experiment users coming from outside facilities. Most of users including in-house staffs require a fast, stable and secure network environment to perform experiments.

In SPring-8, the network system consists of three network zones; a Control-LAN for accelerators and beamlines, a BL-USER-LAN for beamline experiment users, and an OA-LAN for the facility public network. A firewall system sets a demilitarized zone (DMZ) between the Control-LAN and the OA-LAN. The DMZ is used for the connection of program development computers and database server computers, and allows bi-directional accesses to the Control-LAN and the OA-LAN. But no direct access between the Control-LAN and OA-LAN is allowed. The firewall system and a Layer 3 switch being installed between the BL-USER-LAN and the OA-LAN interconnect the networks, and block incorrect packets. However, the network system was not secure enough to block worm infection. In the summer of 2003, the computer worm, W32/Blaster worm, explosively went around the world. In SPring-8, a user carried a Blaster worm by a PC, and the infection disturbed the BL-USER-LAN. Some of users brought laptop PCs without security provision and connected the infected PCs to the BL-USER-LAN. The problem was how to protect our internal network along with providing flexible network environment to the experiment users.

In the summer of 2004, we installed a security gateway, an InterSpect from Check Point Software Technologies, to block the spread of worms and attacks inside the BL-USER-LAN. If there is a suspicious or infected computer in the BL-USER-LAN, the InterSpect quarantines and automatically isolates it from the network. In the Control-LAN, most of Windows PCs for measurement instruments are protected by anti virus software, and the server of anti virus software in the DMZ automatically delivers these PCs virus definitions. The virus scanner checks the vulnerability of Windows PCs at the accelerator maintenance period.

This paper describes the design and management of a secure network to block worm infection, reporting on the results of performance measurements of the security gateway.

INTORODUCTION

SPring-8 was opened to the public use in 1997, and currently has 48 beamlines for synchrotron radiation experiments. The total number of experiment users exceeded 9,000 in a year of 2003. A few years ago, network administrator's concern of security was to protect the internal network against attacks coming from outside. Therefore, we installed a firewall system to the boundaries. However, the network system was not secure enough to block worm infections.

Many experiment users and in-house staffs carry laptop PCs, and sometimes connect the PCs to various places in outside. In August 2003, the computer worm, W32/Blster worm [1], explosively went around the world. In SPring-8, a user carried a Blaster worm by a PC, and connected the PC to the internal network. The worm infected the backbone of the inside network and disturbed beamline experiment. A firewall was not useful to prevent worm infection. New quarantine system, a security gateway, was necessary to protect the internal network from computer worms.

In this paper, we describe the network configuration and management policy in SPring-8, selection criteria of a security gateway, network operation experience, and the results of traffic throughput measurements of the gateway.

NETWORK SYSTEM OVERVIEW

The network system in SPring-8 consists of three network zones: a Control-LAN, a BL-USER-LAN and an OA-LAN [2]. Figure 1 shows the network system of SPring-8.

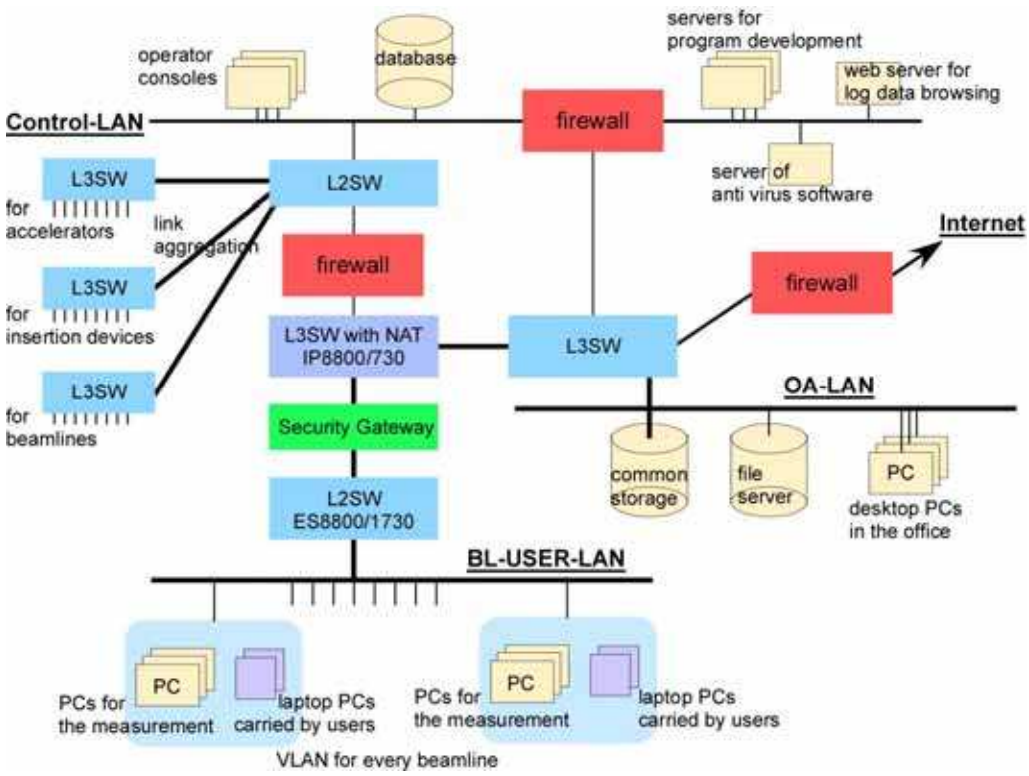

Figure 1: The network system of SPring-8

A Control-LAN is used to control accelerators, insertion devices and beamline components. The backbone is the Gigabit Ethernet. We adopt Link Aggregation (IEEE802.3ad) configuration to keep the network redundancy. A BL-USER-LAN is a network for a measurement system in each beamline. The backbone is the Gigabit Ethernet and uplink connection from each beamline to the backbone is 100Mbps. A Layer3 Switch (L3SW) (NEC [3] IP8800/730) and a Layer2 Switch (L2SW) (NEC ES8800/1730) have redundant power supplies and CPU modules. We introduce VLAN (IEEE802.1q) to build a logically independent and flexible network system. Each beamline is assigned a different IP subnet belonging to a unique VLAN. An OA-LAN is a network for the facility public. An experimental data produced in the beamline is sent to file servers or common data storage systems on the OA-LAN. A demilitarized zone (DMZ) is set by a firewall between the Control-LAN and the OA-LAN. The DMZ is used for the connection of program development computers, a web server for log data browsing and a server of anti virus software, and allows bi-directional accesses to the Control-LAN and the OA-LAN. But no direct access between the Control-LAN and OA-LAN is allowed by the firewall.

When IP packets go out from the BL-USER-LAN to the OA-LAN, an IP subnet of a beamline is rewritten to an IP address for the OA-LAN by a L3SW with a function of IP masquerade (Network Address Translation). Therefore, IP packets directly can't pass through from the OA-LAN to the BL-USER-LAN, which guarantees an access control by unidirectional communication of the network. Moreover, the accesses from one VLAN to other VLAN are restricted by IP filtering of L3SW.

A firewall between the Control-LAN and the BL-USER-LAN passes packets through only predefined IP addresses and limited service ports, and rejects major services such as http, ftp, ping and ssh. The L3SW doesn't change the IP addresses of packets from the BL-USER-LAN to the Control-LAN.

TOWARD A SECURE NETWORK

For Control-LAN

The access from the Control-LAN to the Internet is restricted by a firewall, so it is difficult to keep the antivirus software's virus definition file updated. At worst case, a new type of worm infection in the Control-LAN could cause heavy damage to the operation of SPring-8 and beamline experiment. In the winter of 2004, we installed integrated security software (Symantec [4] Client Security) that can update virus definition file via a proxy server on the DMZ, and kept anti virus software up-to-date. At the same time, we installed network vulnerability assessment tool (Symantec NetRecon).

By port scanning, we can find computers opening vulnerability ports during accelerator maintenance period. We can't install the anti virus software and security patches to an embedded Windows OS of oscilloscope for example on the Control-LAN. In such a case, we refer to the scanning results of other computers and stop unnecessary network services.

For BL-USER-LAN

We had no technical solution to detect computer worms and protect the BL-USER-LAN until the summer of 2004. From September of 2003 to August of 2004, we had troubles eight times in the BL-USER-LAN due to worm attacks. An example of the Worm attacks in the BL-USER-LAN is described as follows.

- A user connected a worm-infected laptop PC to the BL-USER-LAN.
- The worm checked active machines as target for attack by sending ICMP echo packets.
- The L3SW was overloaded by significantly increased ICMP traffic. The CPU utilization of L3SW reached to 100%, afterwards the L3SW couldn't forward packets any more.
- Whole BL-USER-LAN was down. All of users on the BL-USER-LAN couldn't access to any networks.

The L3SW system overload due to the high CPU utilization is related to the NAT architecture of the L3SW. When the L3SW receives a packet, at first it translates IP address of the packet on the CPU. Next, when the L3SW receives the same sessions (same source and destination address), it handles the sessions by using the ASIC (Application Specific Integrated Circuit). Therefore, the CPU is released from the packet processing load. Worm infection sent ICMP echo packets intensively to many different destination IP addresses, whether destination hosts were alive or not, while swept one network segment. As a result, the CPU of the L3SW had to translate all of the received IP addresses, finally reached to its performance limitations. In this case, the worm traffic was no more than 200kbps.

By using VLAN and IP filtering, the BL-USER-LAN established the independency of each beamline network, so the worm infection in the beamline didn't get through other beamlines. The worm infection in the BL-USER-LAN never spread to the Control-LAN, because a firewall rejects the ICMP. On the other hand, the infection might damage the OA-LAN and the Internet. Whenever the throughput was extremely slow, we had to check every time whether the CPU utilization of L3SW reached to 100%. When worm infection occurred, we had to rush to the suspicious beamline, found the computer with worm and removed it from the BL-USER-LAN. The primitive troubleshooting was dependent on the manpower.

In order to introduce new technical solution, we clarified our approach of network security in the BL-USER-LAN as following:

- We want to block the spread of worms inside the BL-USER-LAN, and protect the backbone.
- When worm infection occurs in a beamline, we stop the network service of the infected beamline in order to provide the stable experiment environment to other beamlines.

- All staffs and users should be careful about computer worms and gird their computers from malformed packets by patching and installation of anti virus software.

In the summer of 2004, we installed a security gateway (Check Point [5] InterSpect410) inside the BL-USER-LAN. The InterSpect410 (IS) was set up between the L3SW and the L2SW (NEC ES8800/1730) as shown in figure 1. The IS bridges multi segments to the backbone and is invisible to the IP network. It checks all packets coming from the BL-USER-LAN. Before installation of the IS, total throughput of BL-USER-LAN was about 100Mbps at a maximum. Therefore, we selected the IS, which has a throughput specification of 500Mbps.

The IS selection criteria are listed as follows:

- We don't install software license to the computers of a number of users coming from various facilities.
- A modification has to be minimum for the system installation.
- The system runs by easy management and operation.
- We will be able to integrate management of the firewall and the security gateway in the future.

OPERATION

The IS detects patterns of abnormal network behaviour in accordance with our rules. When the IS detects malicious activity, it prevents the activity by dropping or rejecting the malformed connection. The IS has a unique function "*quarantine*". If a host is in quarantine, it can't communicate with machines of other network zone for 30 minutes. The quarantine quickly isolates the host and eliminates the attack before infection spread out. When a host scans more than 30 inactive ports over a period of 20 seconds, it is automatically quarantined in our guideline. If the suspicious activity still exists at the end of quarantine period, and also the user of the host again tries to communicate, it is once more automatically quarantined. Port scanning is a method of collecting information about open TCP and UDP ports in a network. Once we block the port scanning, we have no way to identify the variety of worms. However, the blocking prevents the pre-emptive attack.

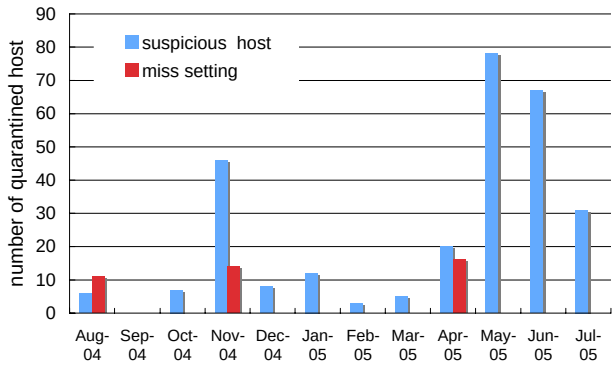

Figure 2: The monthly number of host quarantined by the InterSpect

We accumulated a logging data of the IS from Aug-04 to Jul-05. Figure 2 shows the monthly number of quarantined hosts. "*Suspicious*" is a type of a Sweep Port Scan (SPS) or a peer-to-peer (P2P) connection such as a file sharing service for example Kazza and Gnutella. For past a year, the number of hosts quarantined by a SPS was 265 and by P2P was 15. "*Miss setting*" is caused by wrong detection of worm pattern string. The detection can block not only the spread of worms but normal connections. For example, a pattern string for detection of Sasser worm [6] exploiting the vulnerability of Microsoft's Local Security Authority Service (LSASS) is just a string, "\\lsarpc$". The communication management of domain authentication and active directory also uses LSASS. Therefore, normal connections are miss-quarantined.

From Sep-04 to mid Oct-04, the number of quarantined hosts was zero, because we did not block the SPS. But as the result, Trojan horse occurred. From Jan-05 to Mar-05, SPring-8 storage ring was in maintenance. Consequently, there were fewer suspicious hosts than usual because activity of experiment users was low during this period.

The BL-USER-LAN has been working well for a year being protected by the security gateway.

PERFORMANCE

Set up of measurement

We measured the actual throughput of the IS. Figure 3 shows a set up of performance measurement. Connections of network switches and PCs were Gigabit Ethernet (GE). Two client PCs were set on the BL-USER-LAN, and two server PCs were on the OA-LAN. The IS used the transparent mode and located between the L3SW and the L2SW. Table 1 shows the specifications of PCs. We used netperf2.4.0 [7] as a throughput measurement tool for many different types of network. TCP_STREAM was used for this measurement. We also used PSPaser1.0.2 [8] as a bandwidth control tool to derive maximum throughput. It ran on client PCs (client01 and client02). Tuning parameters such as socket buffer and tick were set to the system defaults.

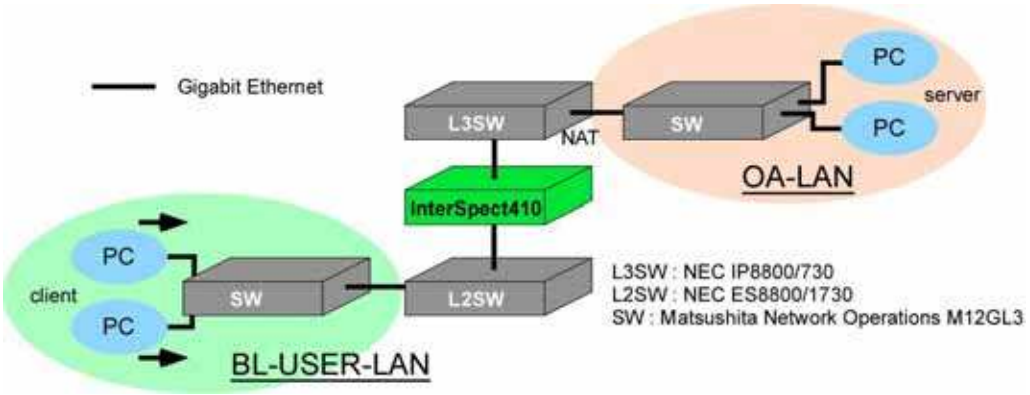

Figure 3: set up of performance measurement

Table 1: Specifications of computers

	Client	Server
CPU	Pentium 4 3.8GHz x 1	Pentium M 1.8GHz x 1
Memory	1GB DDR2	1GB SO-DIMM
Kernel	Linux 2.6.11	Linux 2.6.10
NIC	Broadcom NetXtreme BCM5751 GE	Intel 82541 GE

Measurement

Figure 4-(a) and (b) show the traffic performance when the IS was bypassed. Figure 4-(c) and (d) show the traffic performance with the IS. The horizontal axes of the figures are packet rate (GPR) of client01 generated by PSPacer. A vertical axis is the throughput of client01 or client02. Figure 4-(a) indicates the client01 throughput is proportional with respect to the GPR. In figure 4-(b), the throughput of clinet02 was fixed to 400Mbps by PSPacer. When total throughput of client01 and client02 reached to 900Mbps, the throughput of client01 was slightly below the GPR. The result shows the maximum bare throughput of the set up is about 900Mbps. The IS guarantees a throughput up to 500Mbps as can be seen in figures 4-(c) and (d). When the total throughput went over 500Mbps, the IS was unsteady.

At present, the throughput of each beamline is less than 1 Mbps. One beamline is 320Mbps at a maximum. We can manage the current BL-USER-LAN with the recorded throughput of a 500Mbps. However, if users require higher throughput in the future, the IS will become a bottleneck. We may consider installation of an InterSpect of the higher specification as an option.

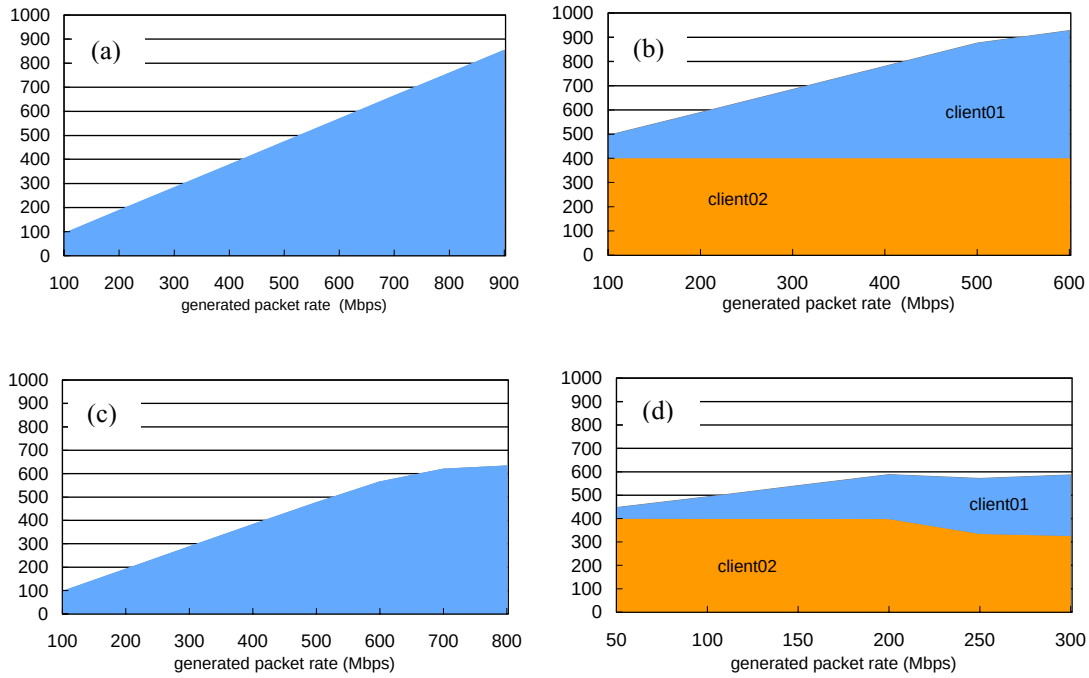

Figure 4: Traffic performance: (a) shows one session without the IS, (b) shows multi sessions without the IS, (c) shows one session with the IS, (d) shows multi sessions with the IS. Blue (orange) area shows throughput of the client01 (client02), respectively.

SUMMARY

We have a mission to keep providing the stable and secure network to the experiment users in SPring-8. Experiment users and in-house staffs brought computer worms being carried by laptop PCs, and worm infection disturbed the BL-USER-LAN operation in the past. As a first step to construct a secure network, we installed an integrated security software system and network vulnerability assessment tool in the Control-LAN. Second, we introduced a security gateway in the BL-USER-LAN as a quarantine system. Blocking of Port Scanning was effective to prevent the spread of worms. A simple detection algorithm of worm pattern strings worked to find suspicious hosts, but sometimes stopped normal connections. The traffic throughput performance of the security gateway was measured, and the actual throughput was 500Mbps. The secure network system has been working effectively to support synchrotron radiation experiments in SPring-8.

REFERENCES

[1] "CERT Advisory CA-2003-20 W32/Blaster worm", http://www.cert.org/advisories/CA-2003-20.html/
[2] M. Ishii *et al*., "Upgrade of SPring-8 Beamline Network with VLAN Technology over Gigabit Ethernet", ICALEPCS'01, http://www.slac.stanford.edu/econf/C011127/TUAP056.pdf
[3] NEC Corporation, http://www.nec.com/
[4] Symantec Corporation, http://www.symantec.com/
[5] Check Point Software Technologies, Inc., http://checkpoint.com/
[6] W32/Sasser, http://www.cert.org/current/archive/2004/09/23/archive.html
[7] Netperf, http://www.netperf.org/
[8] PSPacer, http://www.gridmpi.org/pspacer-1.0/index.en.jsp/