

NSPIXP's

Akira Kato

The logo for WIDE, featuring the word "WIDE" in a bold, black, sans-serif font. A red dot is positioned to the right of the letter "E", centered vertically with the middle of the letter.

WIDE Project
kato@wide.ad.jp

NSPIXP-2

☆ The first production quality IX in Japan

- Colocated in a carrier housing
- Based on Digital's FDDI switches
 - Primary and Backup in 1996 Sep
- Expanded to Four FDDI switches in 1997 Mar
 - 2*Pri + 2*Bak
- Introduction of GigE
 - SSR8600 in 1999
 - Bigiron8000 in 2000
 - Bigiron15000 in 2001
- Retirement
 - FDDI switches: 2000
 - SSR8600: 2001

Current NSPIXP-2

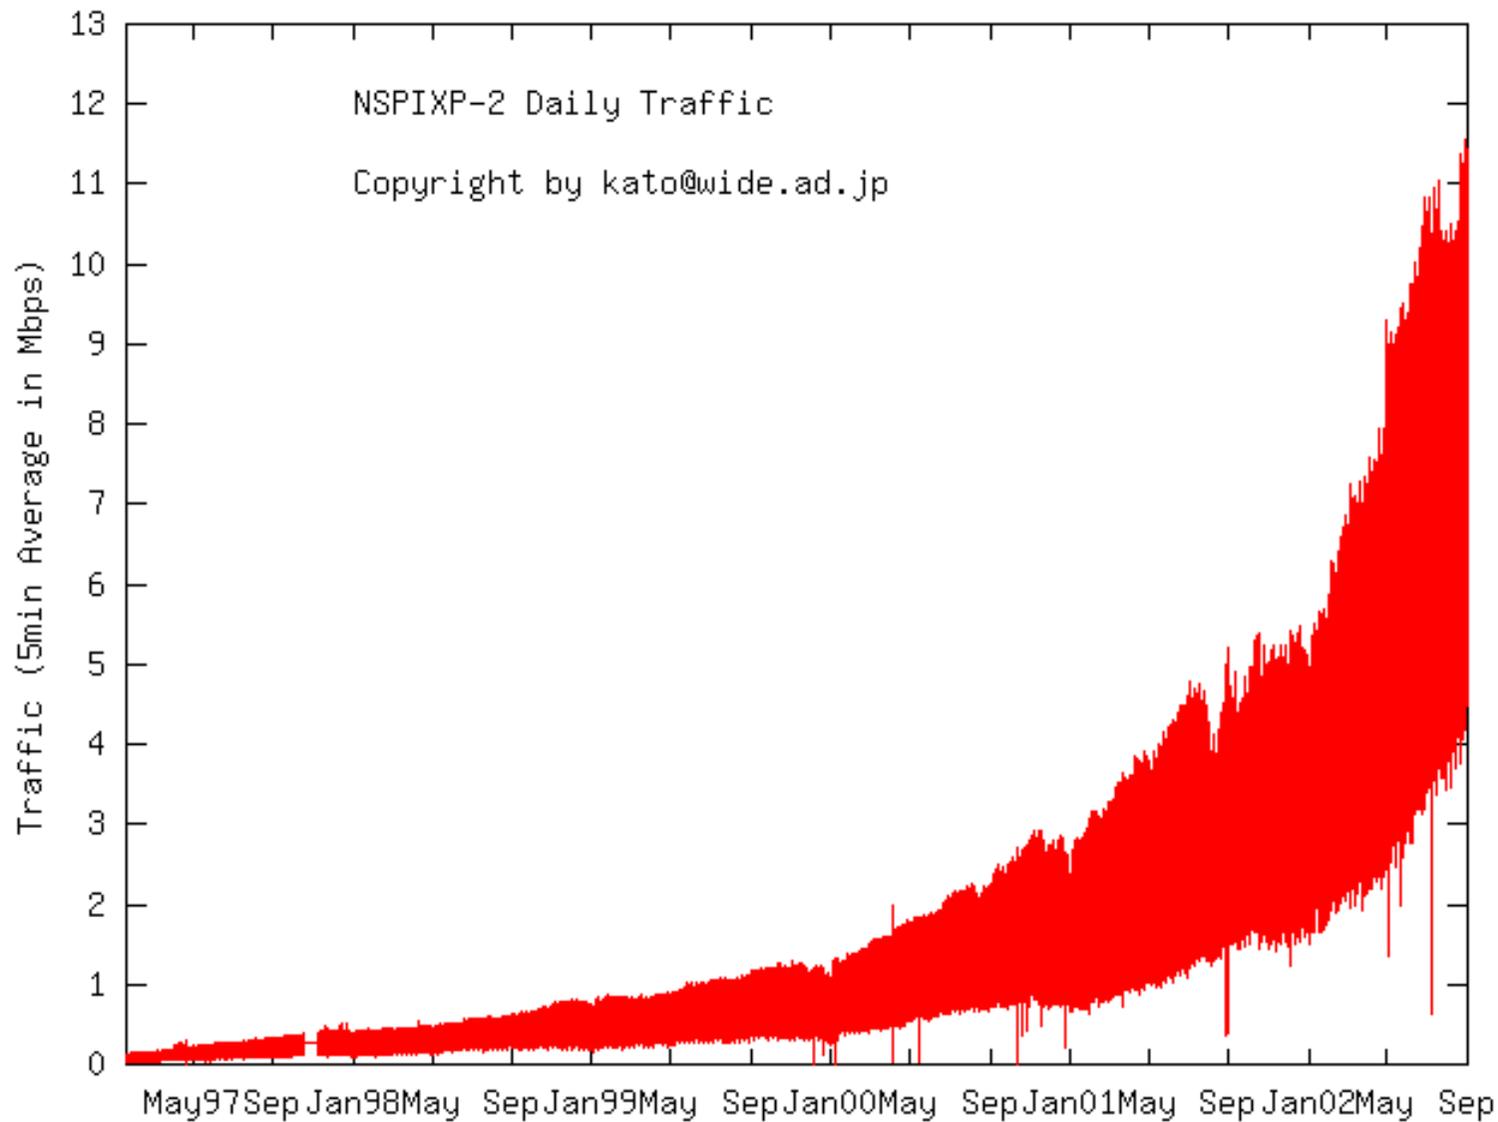
★ **Two Foundry switches**

- BigIron8000 : 2*M4 + 5*8GbE + 24 FE
- BigIron15000 : 2*M4 + 8*8GbE + 2*24FE
- 100 ISP ports, 70 ASes

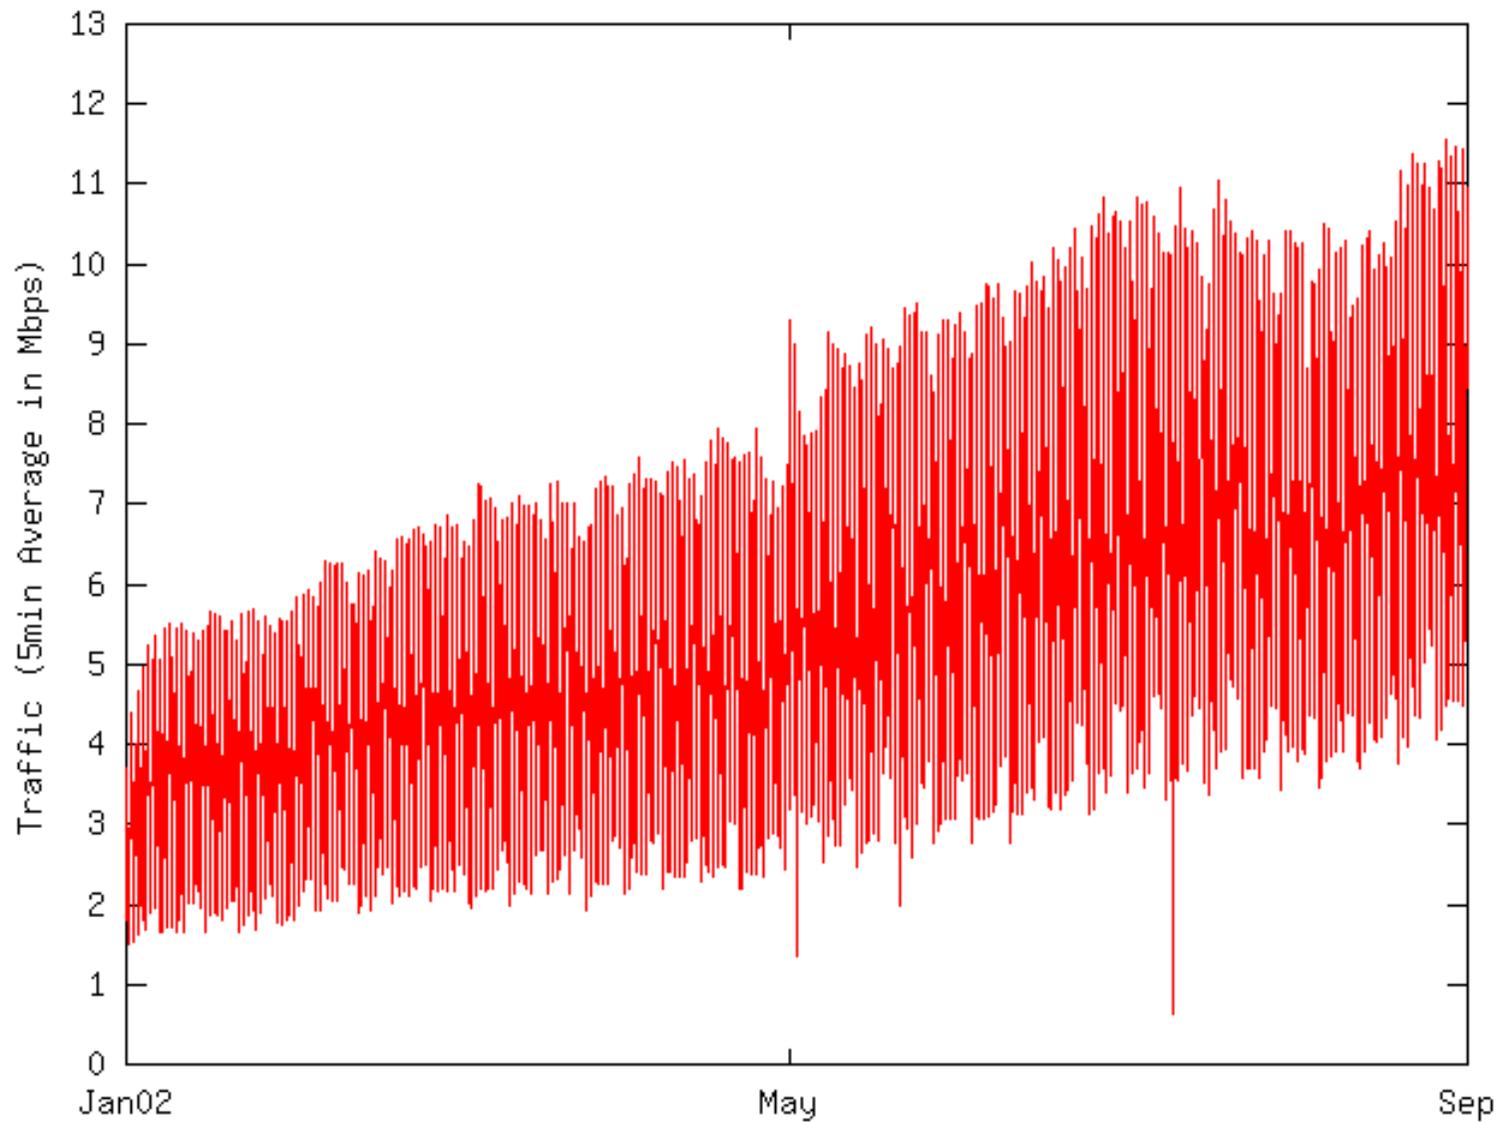
★ **4*GbE Inter-switch link**

- Enhanced to 8*GbE in May 2002
- Not ideally distributed
 - Foundry hashes based on destination MAC
- Ask ISPs to avoid inter-switch peering
 - Big ISPs touch down on both switches

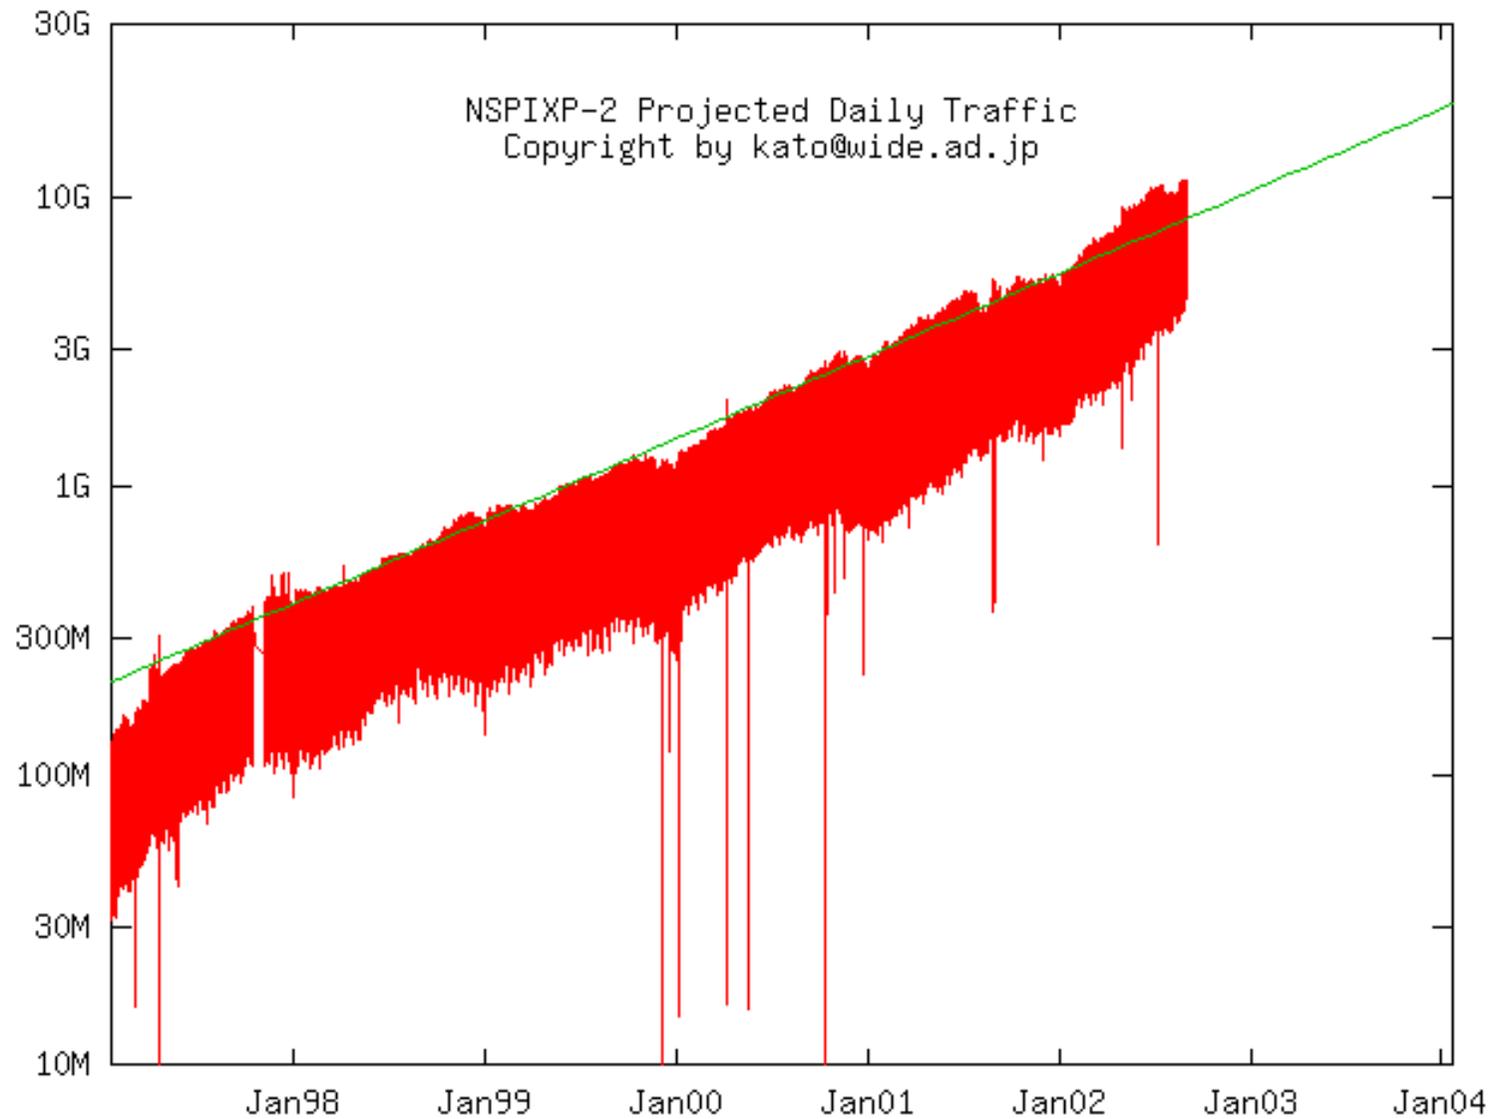
Aggregated traffic growth



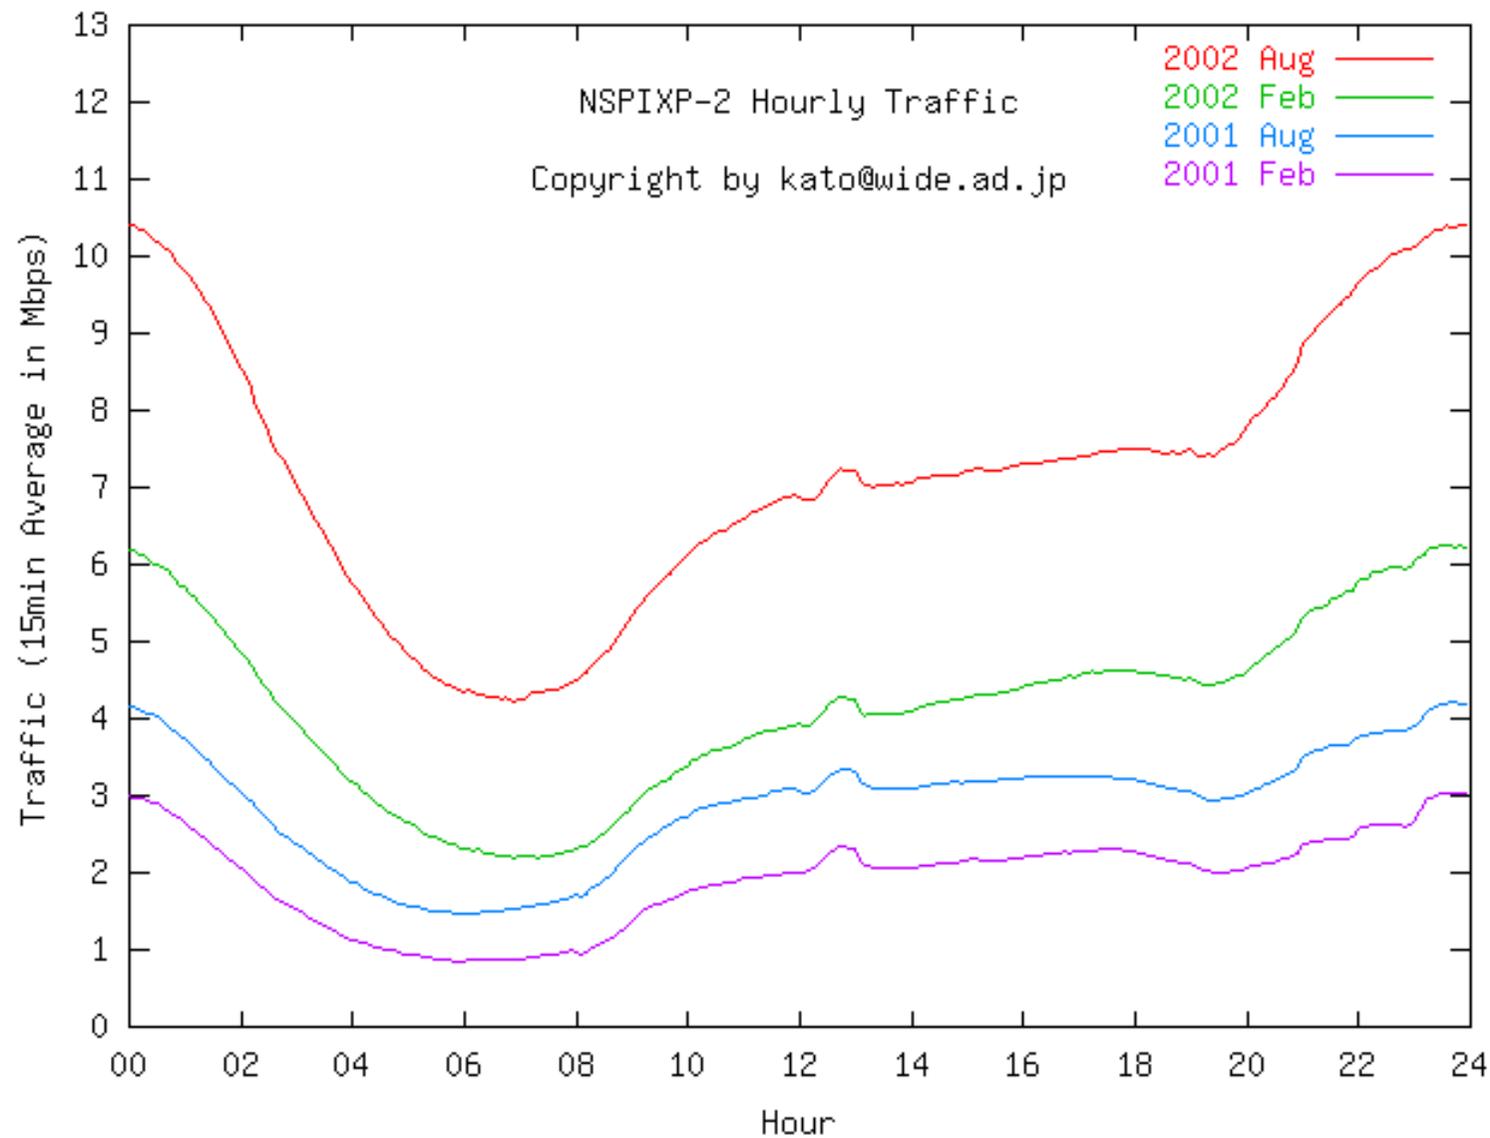
Aggregated traffic growth



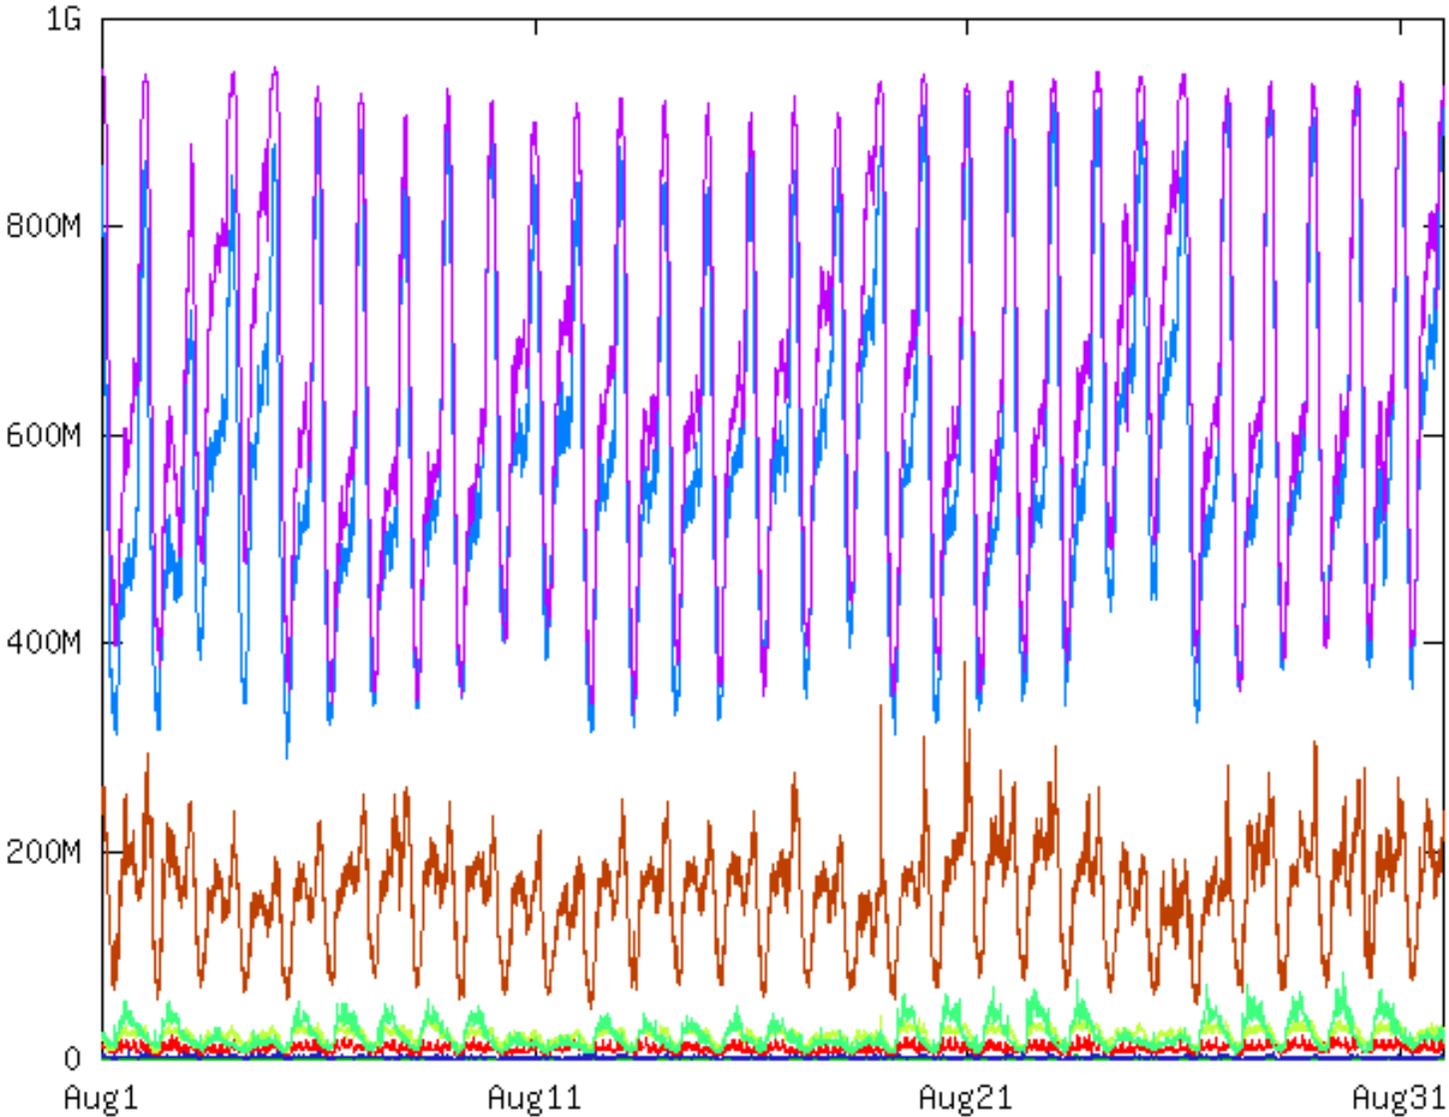
Aggregated traffic growth



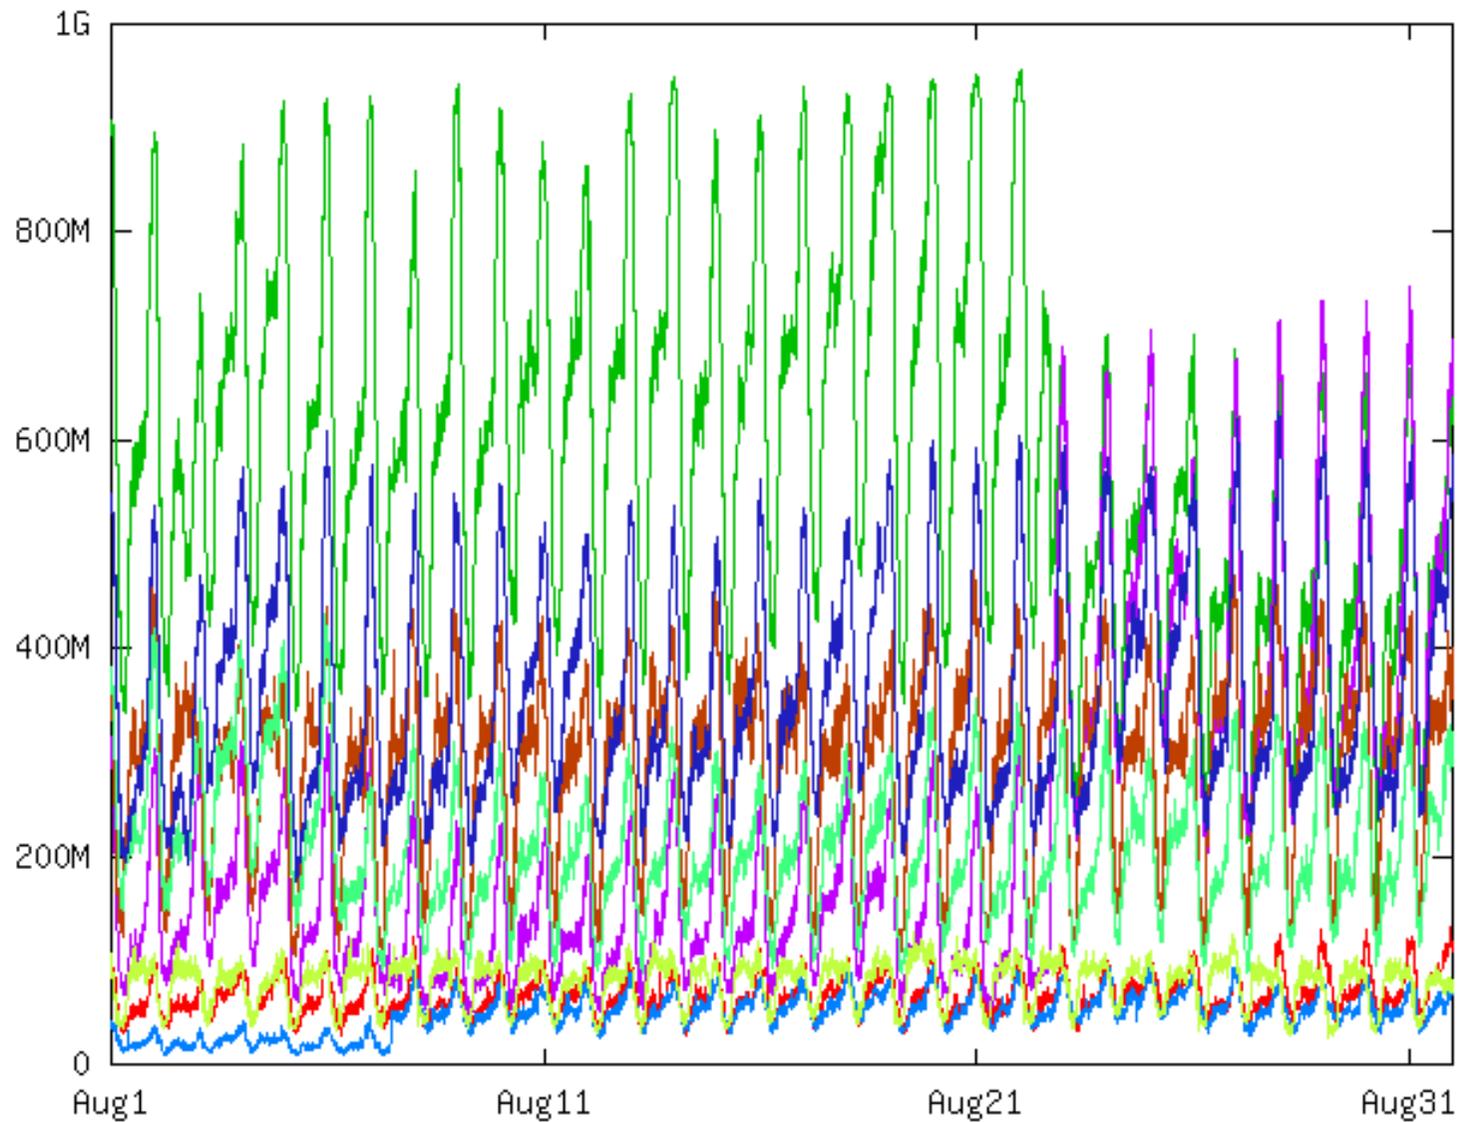
Daily traffic pattern changes



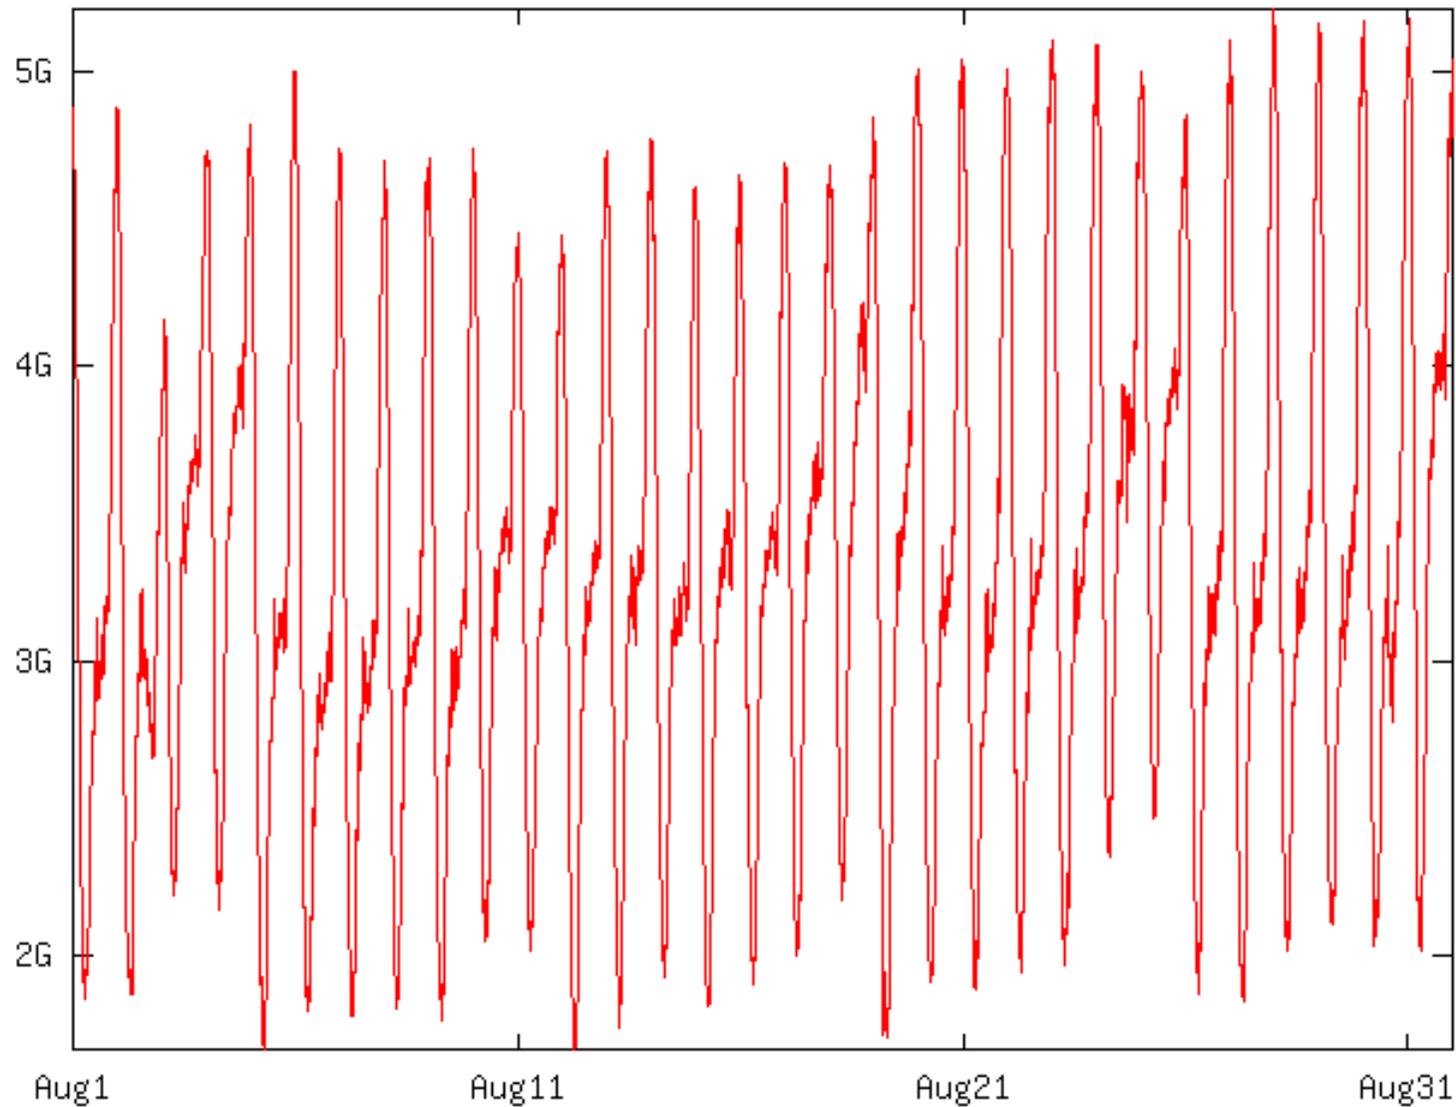
Interswitch Traffic (Aug 2002)



Interswitch Traffic (Aug 2002)



Aggregated Interswitch Traffic



Future plan

★ Introduction of 10GE

- Performance 7-8Gbps
- Expected to be better than 8*GbE
 - Load does not distribute ideally

★ Two 10GE cards will be delivered soon

- Reconfig is to be scheduled

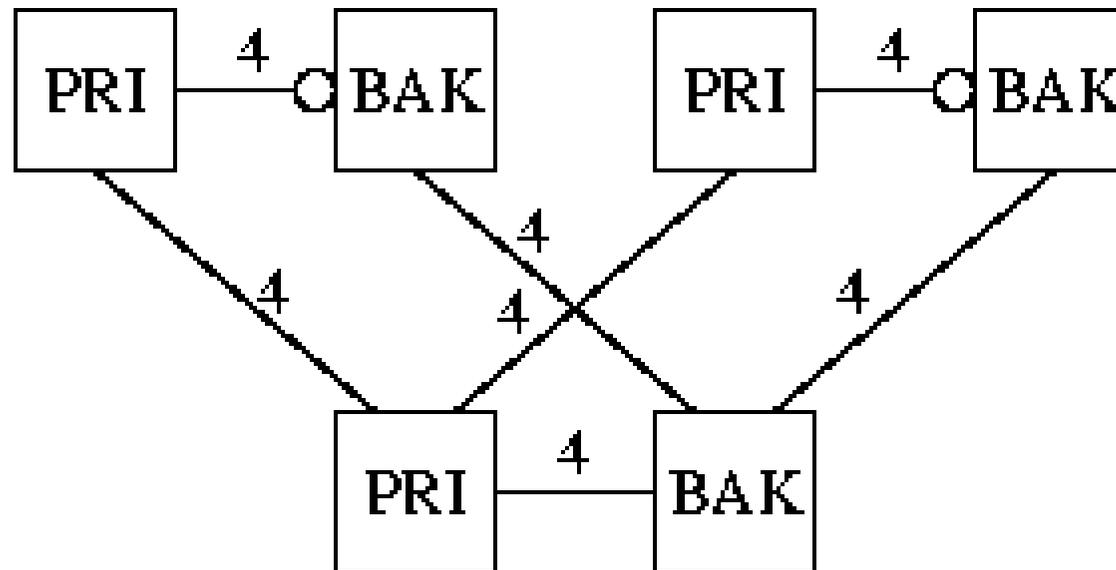
★ IPv6 promotion is available toward 2005

- GbE ISP will be offered a free FE for IPv6
- 12 ISPs are connected

Distributed IX

☆ Accomodate several iDC facilities

- Central topology from KDDI building
- NTT-C, @Tokyo, MIND, L3(Reach), WCOM, Above
- Redundant configuration
 - 4*GbE for primary and backup



Distributed IX

★ **Current status: Not operational yet**

- Un-interoperable 802.3ad / 802.1d

★ **Foundry specific problems**

- Adjacent ports form a trunk
- 802.3ad forms a trunk dynamically
- Possible persistent loop
 - when one of the link is down
 - then one of the switch reloads

★ **Fixed software is being released**

NSPIXP-3

★ **The first IX in Osaka area**

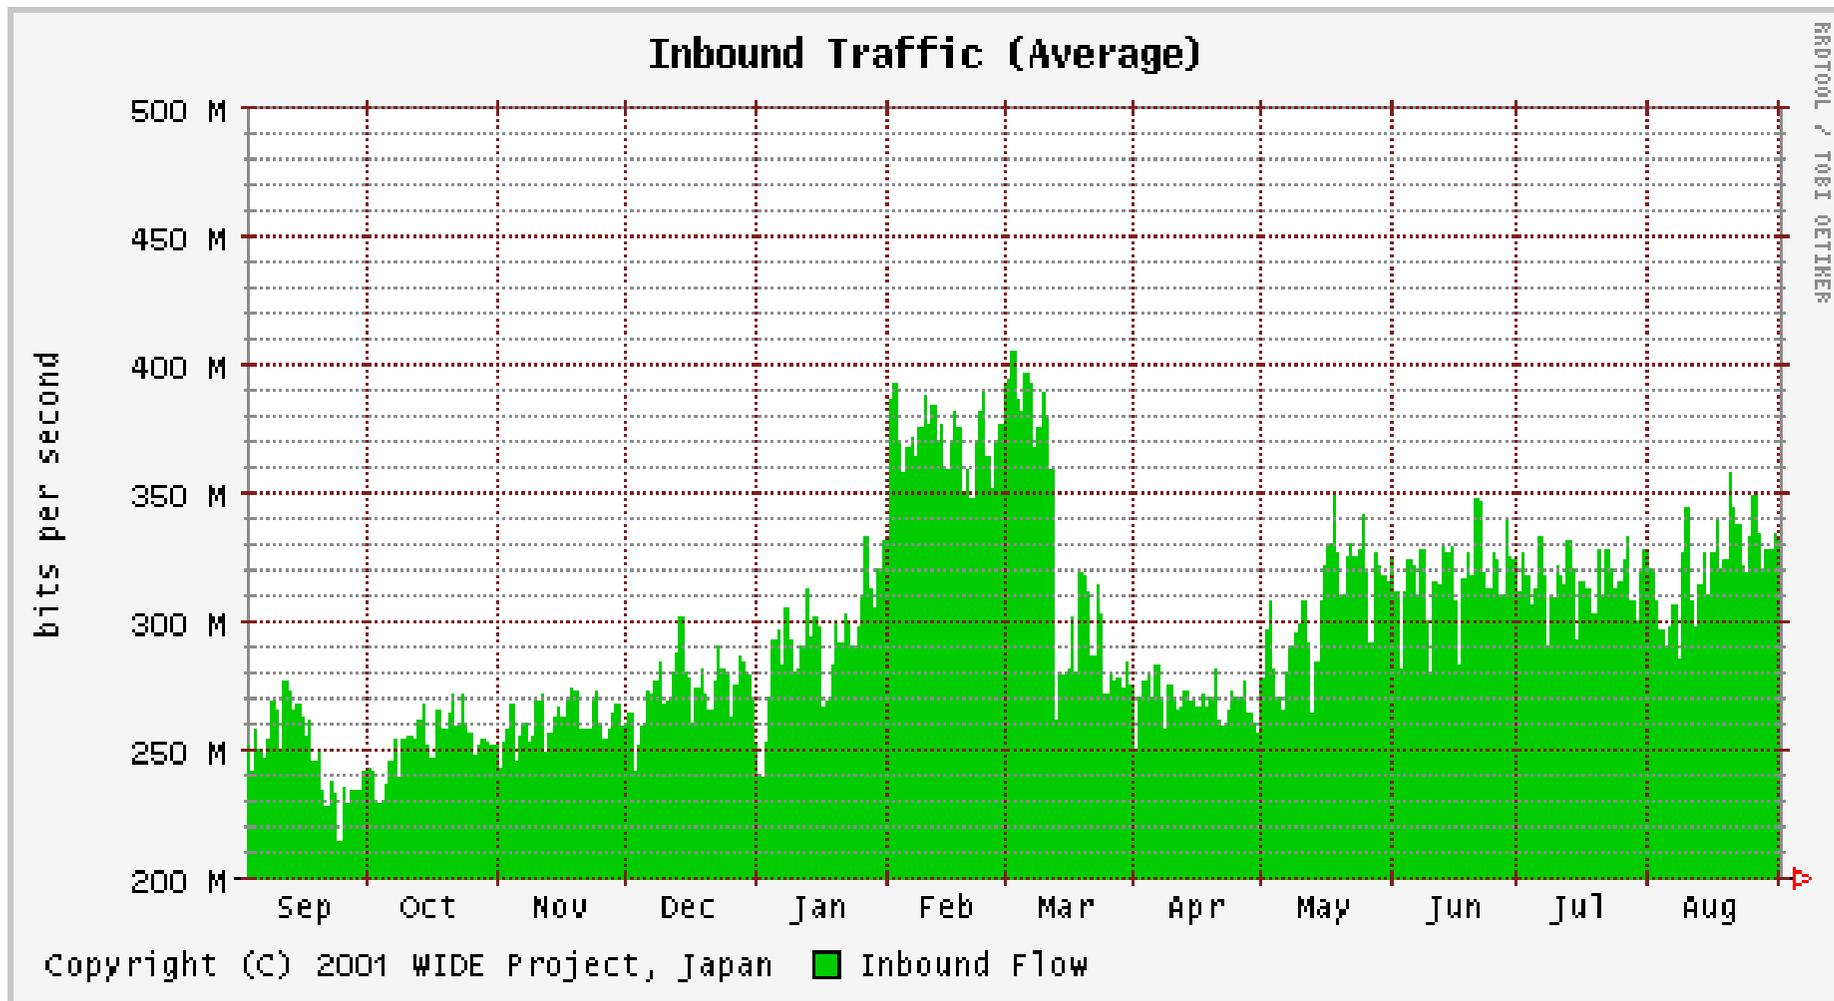
- Booted in Sep 1997
- Since bootstrap, two locations
 - OMP DC and IDC DC
 - Cat5000, two FEs w/ STP
- Later NTT-W DC added in Oct 2000
 - Cat6500, triangle GbE w/ STP
- # of ISP: $8 + 6 + 5 = 19$

★ **1/5 -- 1/10 of NSPIXP-2**

★ **Future plan**

- upgrade inter-switch link to 10GE

NSPIXP-3



NSPIXP-6

★ **Experimental IX for IPv6**

- Operational since Aug 1999
- Started with a single BayStack350-24T
- No redundant configuration
- <http://www.wide.ad.jp/nspixp6>

★ **Possible Remote Access**

- W/ ATM/FE bridges through JGN
- Dark Fiber
- WIDE Extended Layer-2 infrastructure

★ **3 Locations**

- KDDI(28), NTT-C(12), NTT-W(Osaka, 4)

NSPIXP-6

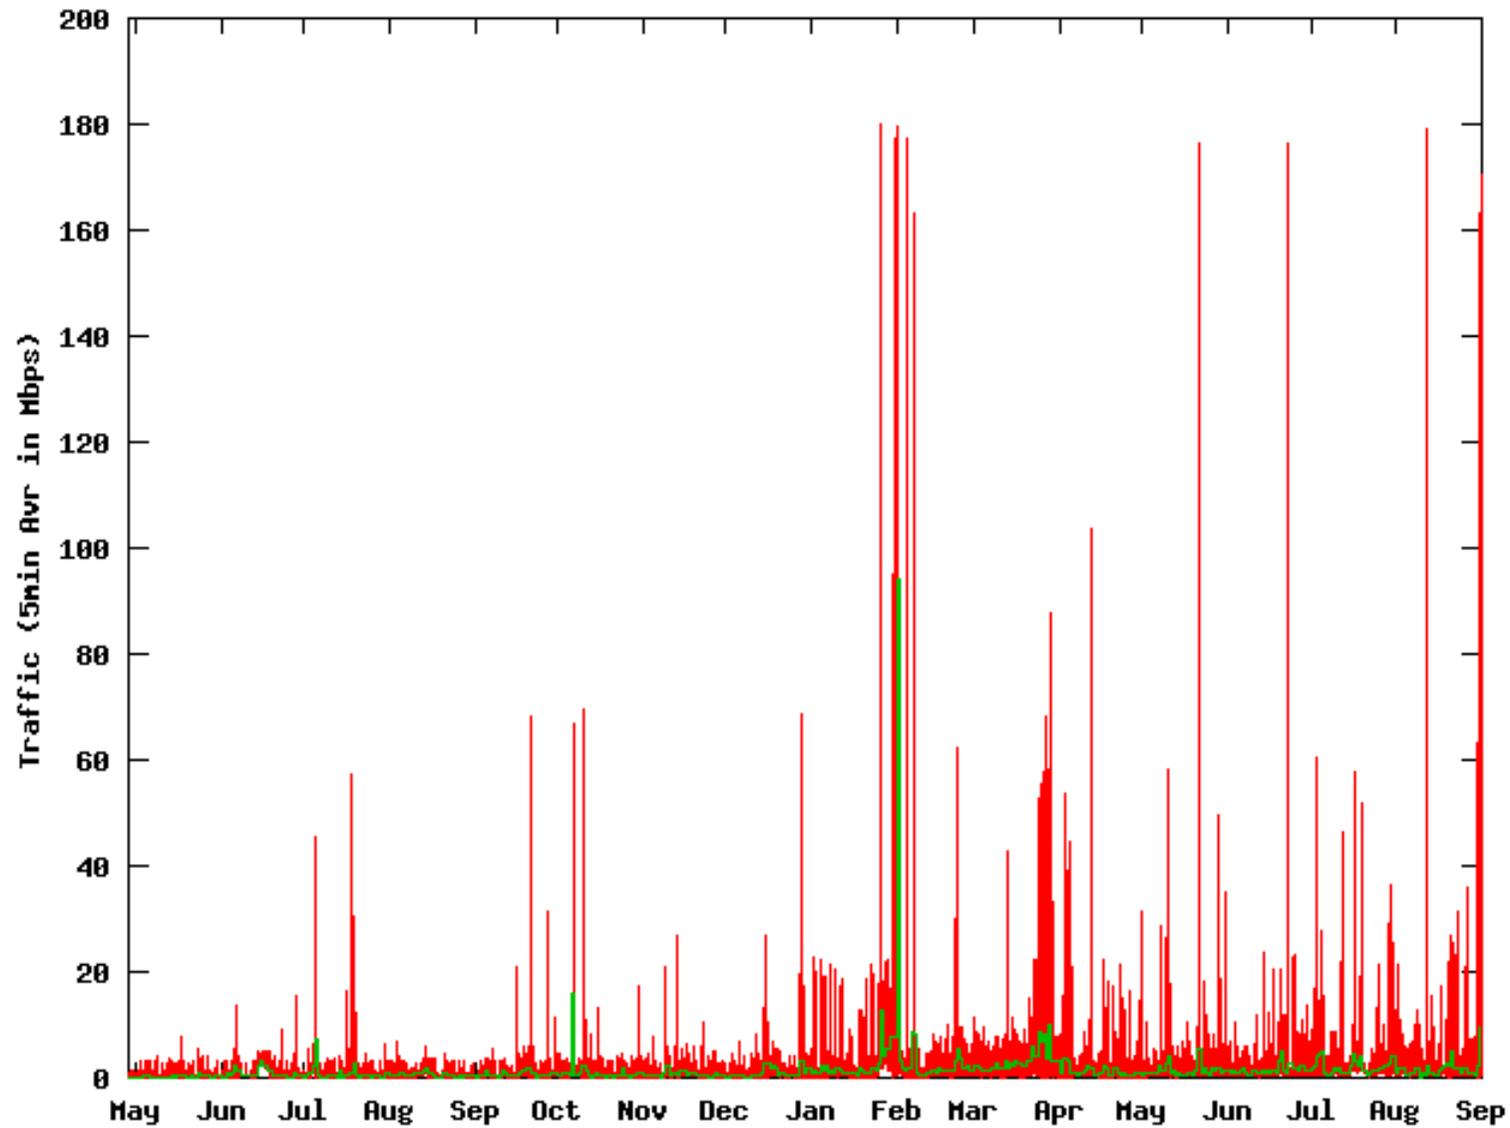
★ **Almost all IPv6 ISPs in JP connected**

- 44 ISPs and non-ISP
- 29 ISPs are sTLA holders

★ **Separate L2 environment to NSPIXP-2**

- WIDE offers free transit
- A GR2000 attaches both at GbE

NSPIXP-6



NSPIXP-6

★ The main switch (Catalyst2948G)

