


Rapid Remote File System Benchmark on Long Fat Network and User ID Mapping Function

Naoyuki FUJITA, Hirofumi OOKAWA
{fujita,ookawa}@chofu.jaxa.jp

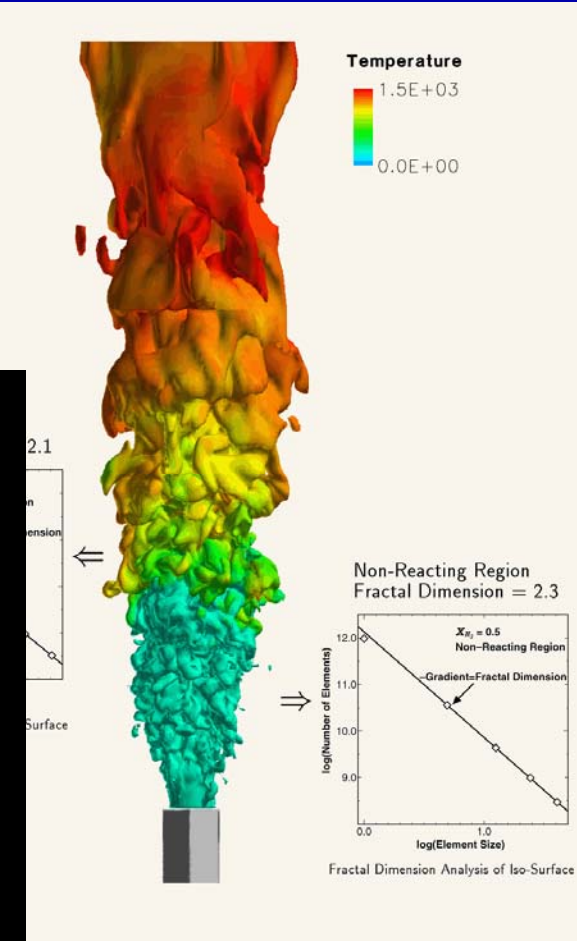

- Application and System Configuration
- Technical Targets
- Former Works
- Benchmark on Long Fat Network
- User ID Mapping
- Summary and Future Works



Application

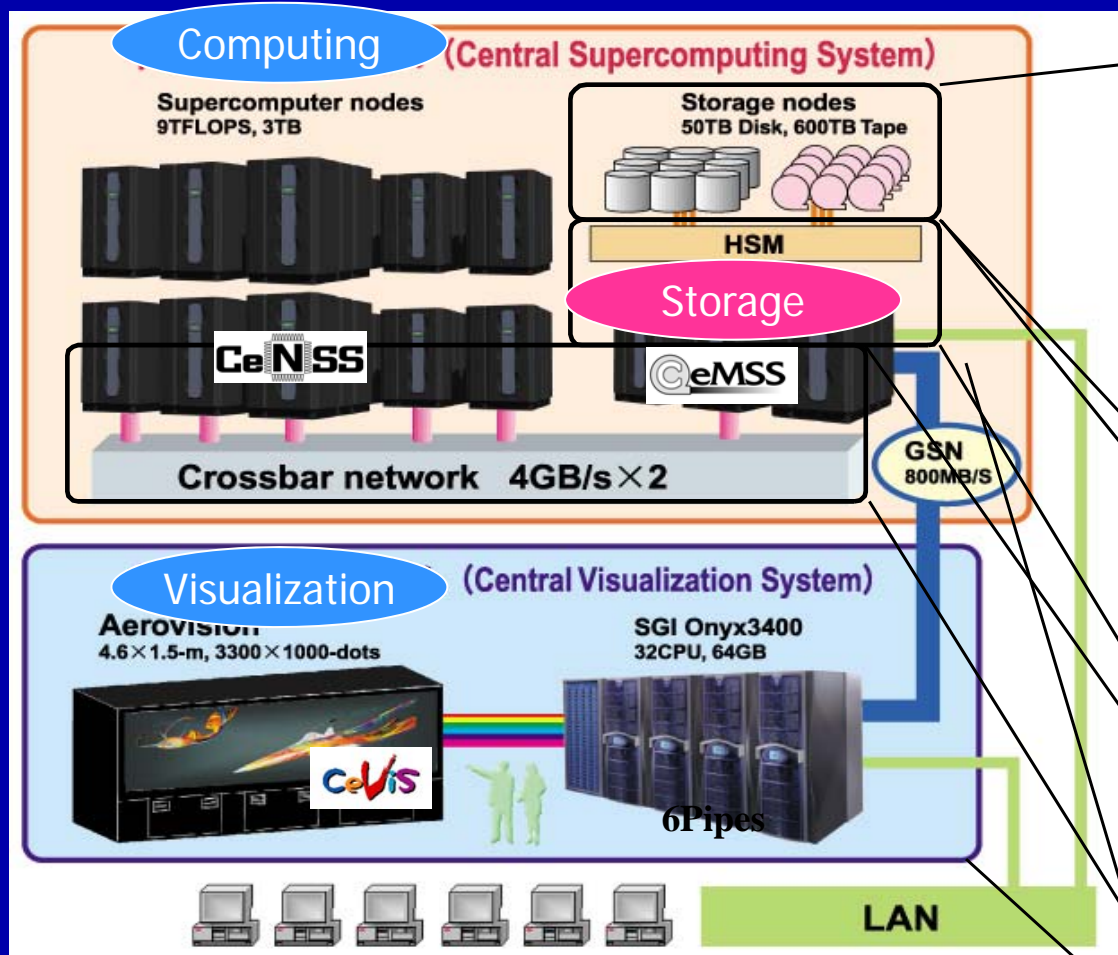


300 [MByte/s]
>1 [TB]



System Configuration

Oct. 2002 : Installed



Storage Devices:
[Disk : 50TB]
RAID5:9D+1P(FujitsuPW-D500B1)
80[drives]
[Tape : 600TB]
LTO(IBM3580)
40[drives]

Local File System: SAM-QFS

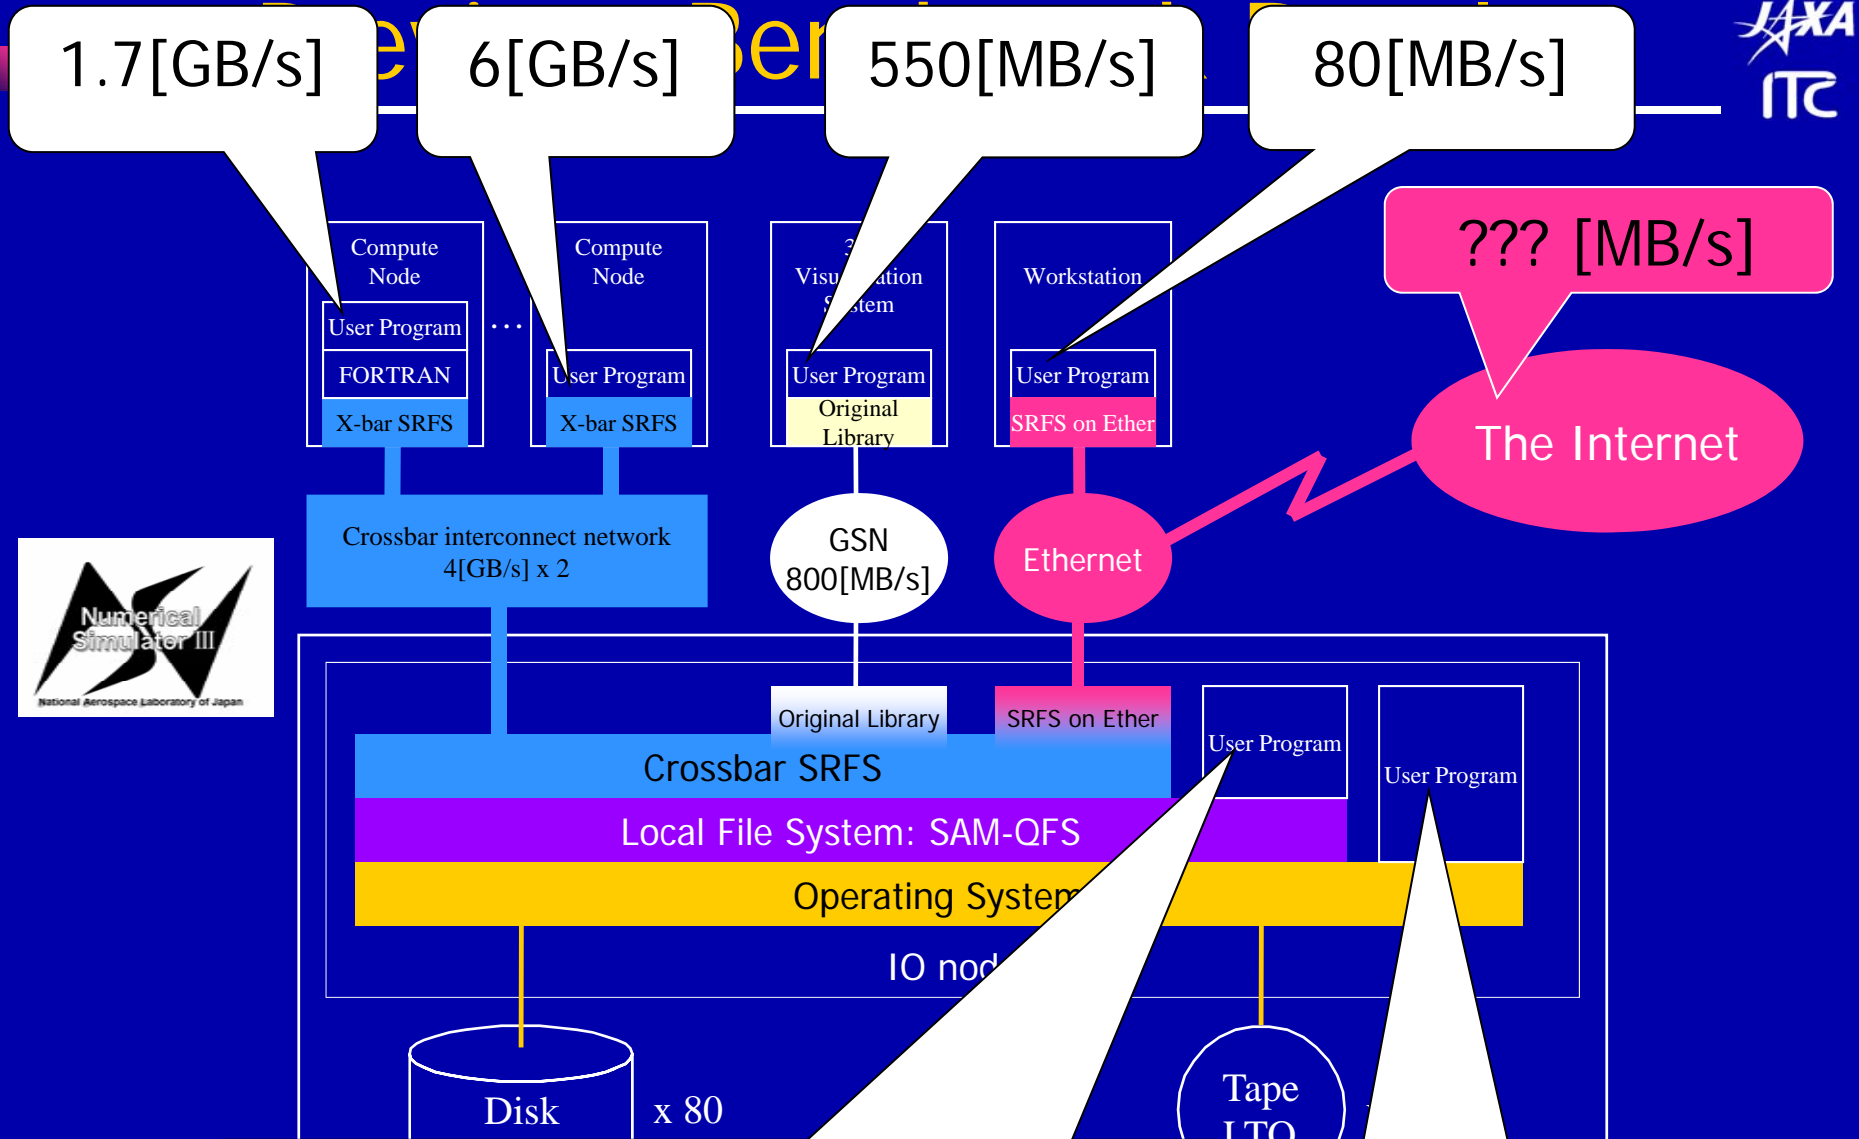
Crossbar Network File System:SRFS

GSN-Lib I/O System

- File Share over IP-network, especially the Internet
- Data coherency under a file-sharing environment
- High-speed without TCP/IP tuning
- Unified user administration
- Security (Working with firewall, Encryption, Intrusion protection)

- File Share over IP-network
 - Added Ethernet communication control module
- Data coherency under a file-sharing
 - Guaranteed by original Crossbar-SRFS
- High-speed without TCP/IP tuning
 - 8 stripes data stream
 - Switch to UDP when I/O length is small
 - 55 [Mbps] on 100Mbps-WAN
 - 610 [Mbps] on GbE-LAN

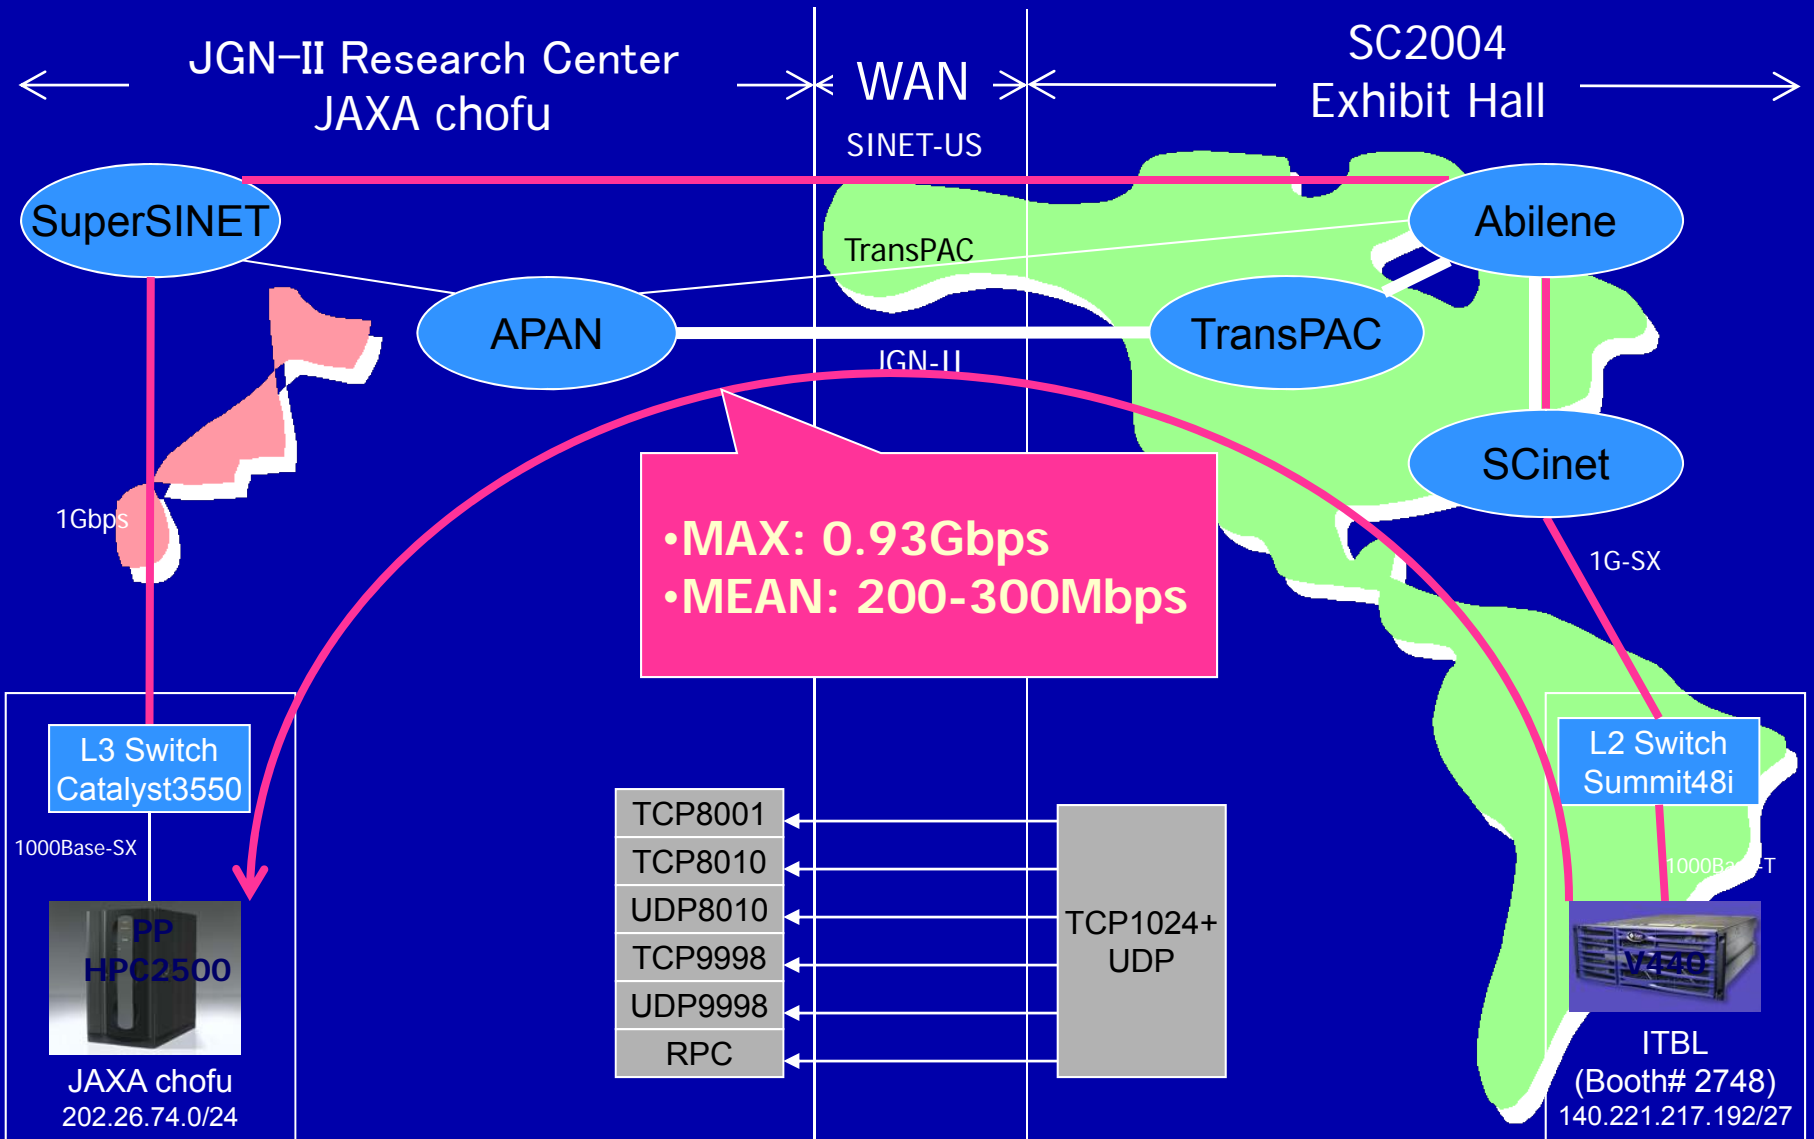
- File Share over IP-network, especially the Internet
- Data coherency under a file-sharing environment
- High-speed without TCP/IP tuning
- Unified user administration
- Security (Working with firewall, Encryption, Intrusion protection)



- One file disk I/O: 80 [MB/s]
- 1 Terabyte file I/O: 3.5 [GB/s]

- DISK: 80 [MB/s]
- TAPE: 14 [MB/s]

Benchmark on Long Fat Network



Benchmark on Long Fat Network (High Speed Function)

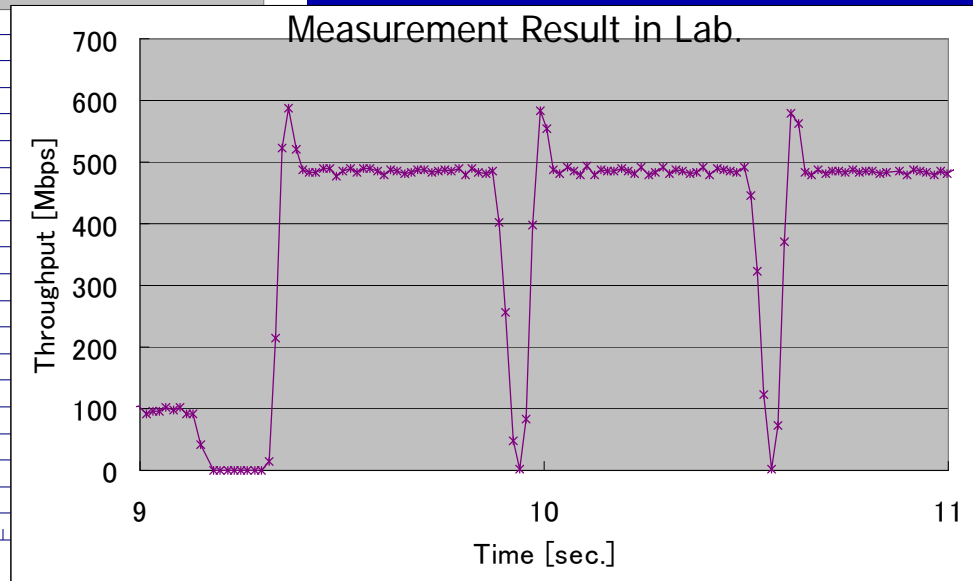
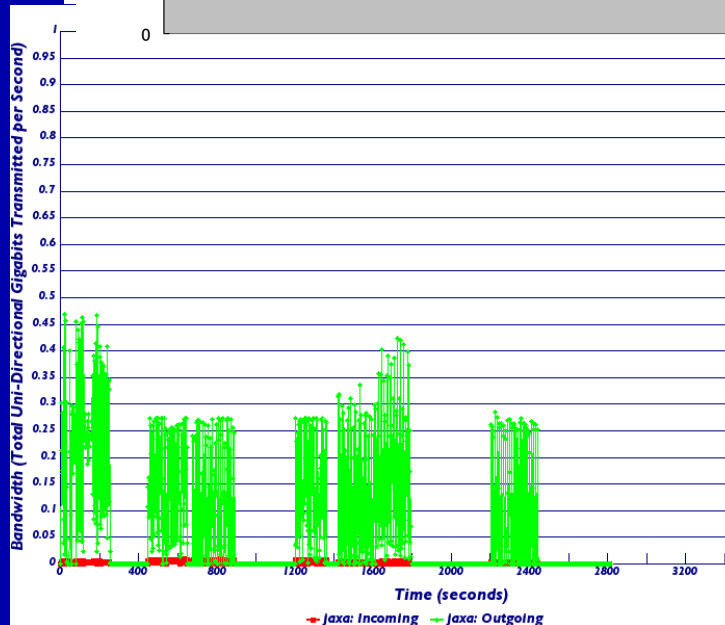
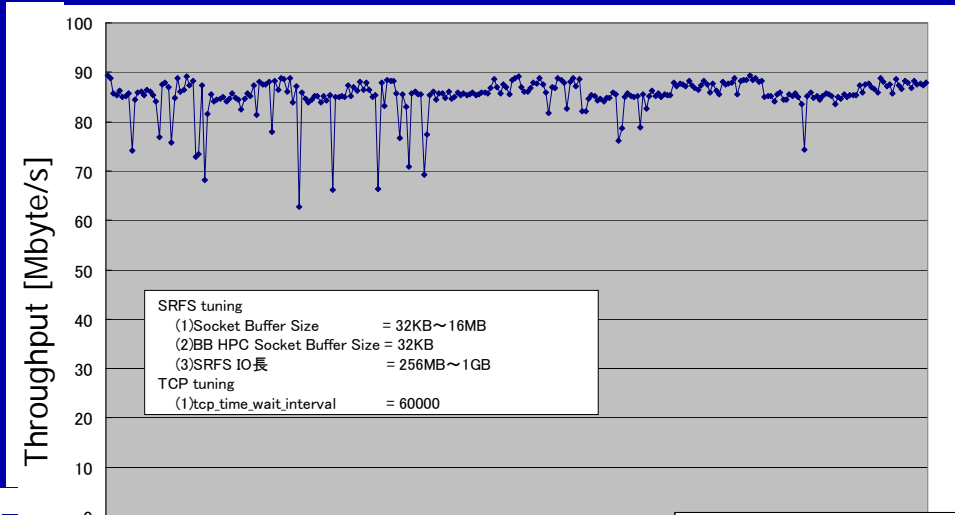
a: Single stream bandwidth
w: TCP window size=16[KB]
r: Round Trip Time=200[ms]

w 16[KB]

a = $\frac{w}{r} = \frac{16 \text{ [KB]}}{200 \text{ [ms]}} = 80 \text{ [KB/s]}$

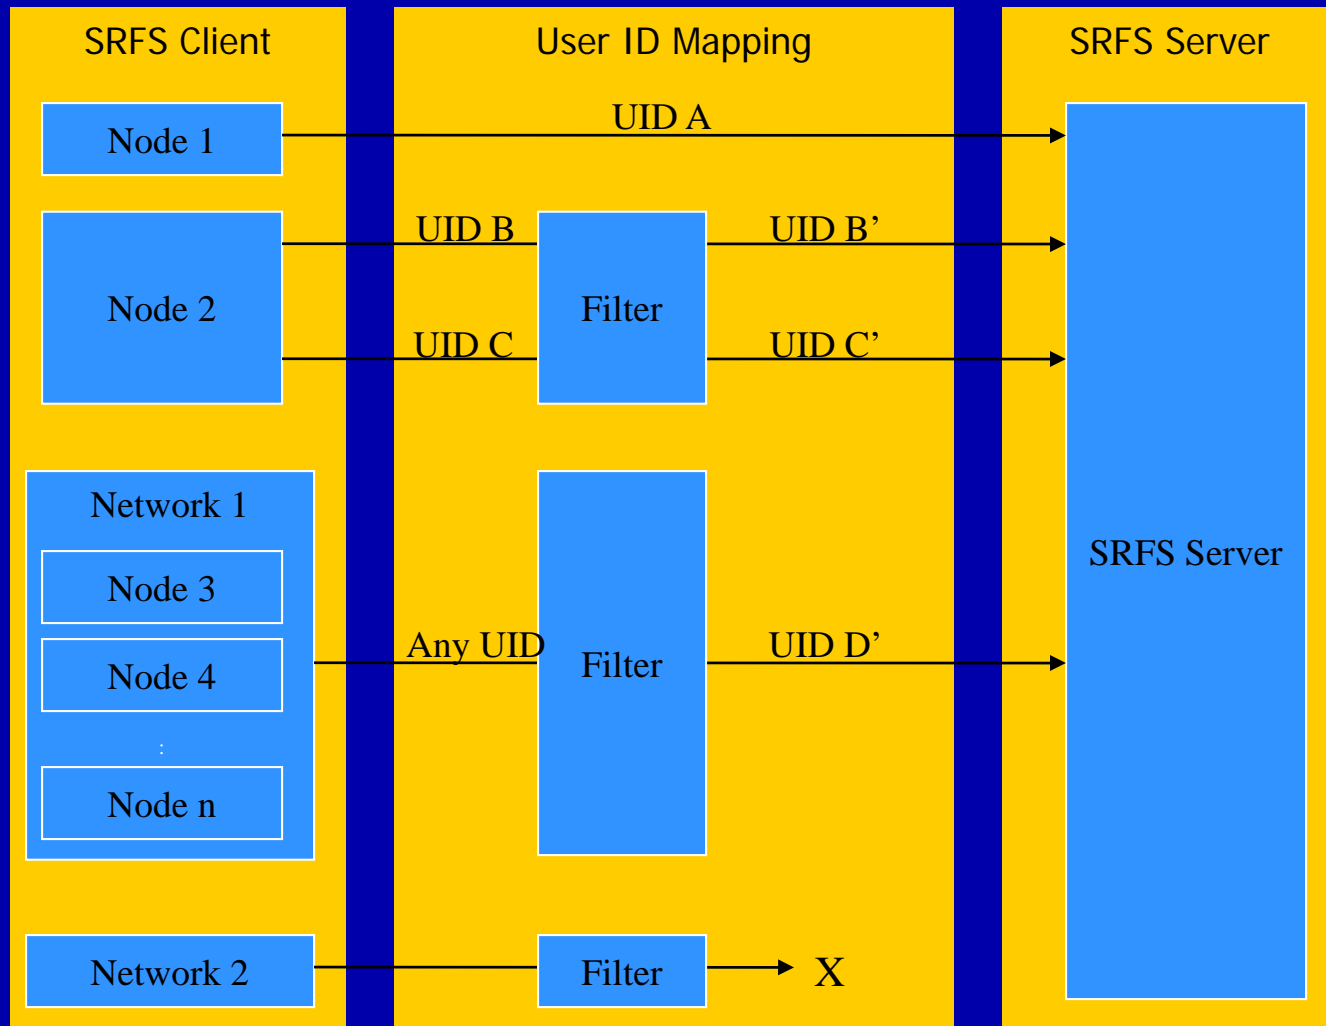
r 200[ms]

→ # of stripes : 1563[streams]



User ID Mapping

- Unified user administration over multi-organization



- Benchmark on Long Fat Network

- Pacific Ocean 1Gbps Link

- MAX: 0.93Gbps

- MEAN: 200-300Mbps


- Striping data stream is not good for High-speed I/O on Long Fat Network

- User ID Mapping

- Simple User ID Mapping between client and server



■ Security (Working with firewall)



■ Linux Client Deployment

- Server ··· Solaris 8 & 9
- Client ··· Solaris 8 & 9, Linux (Red Hat Enterprise)


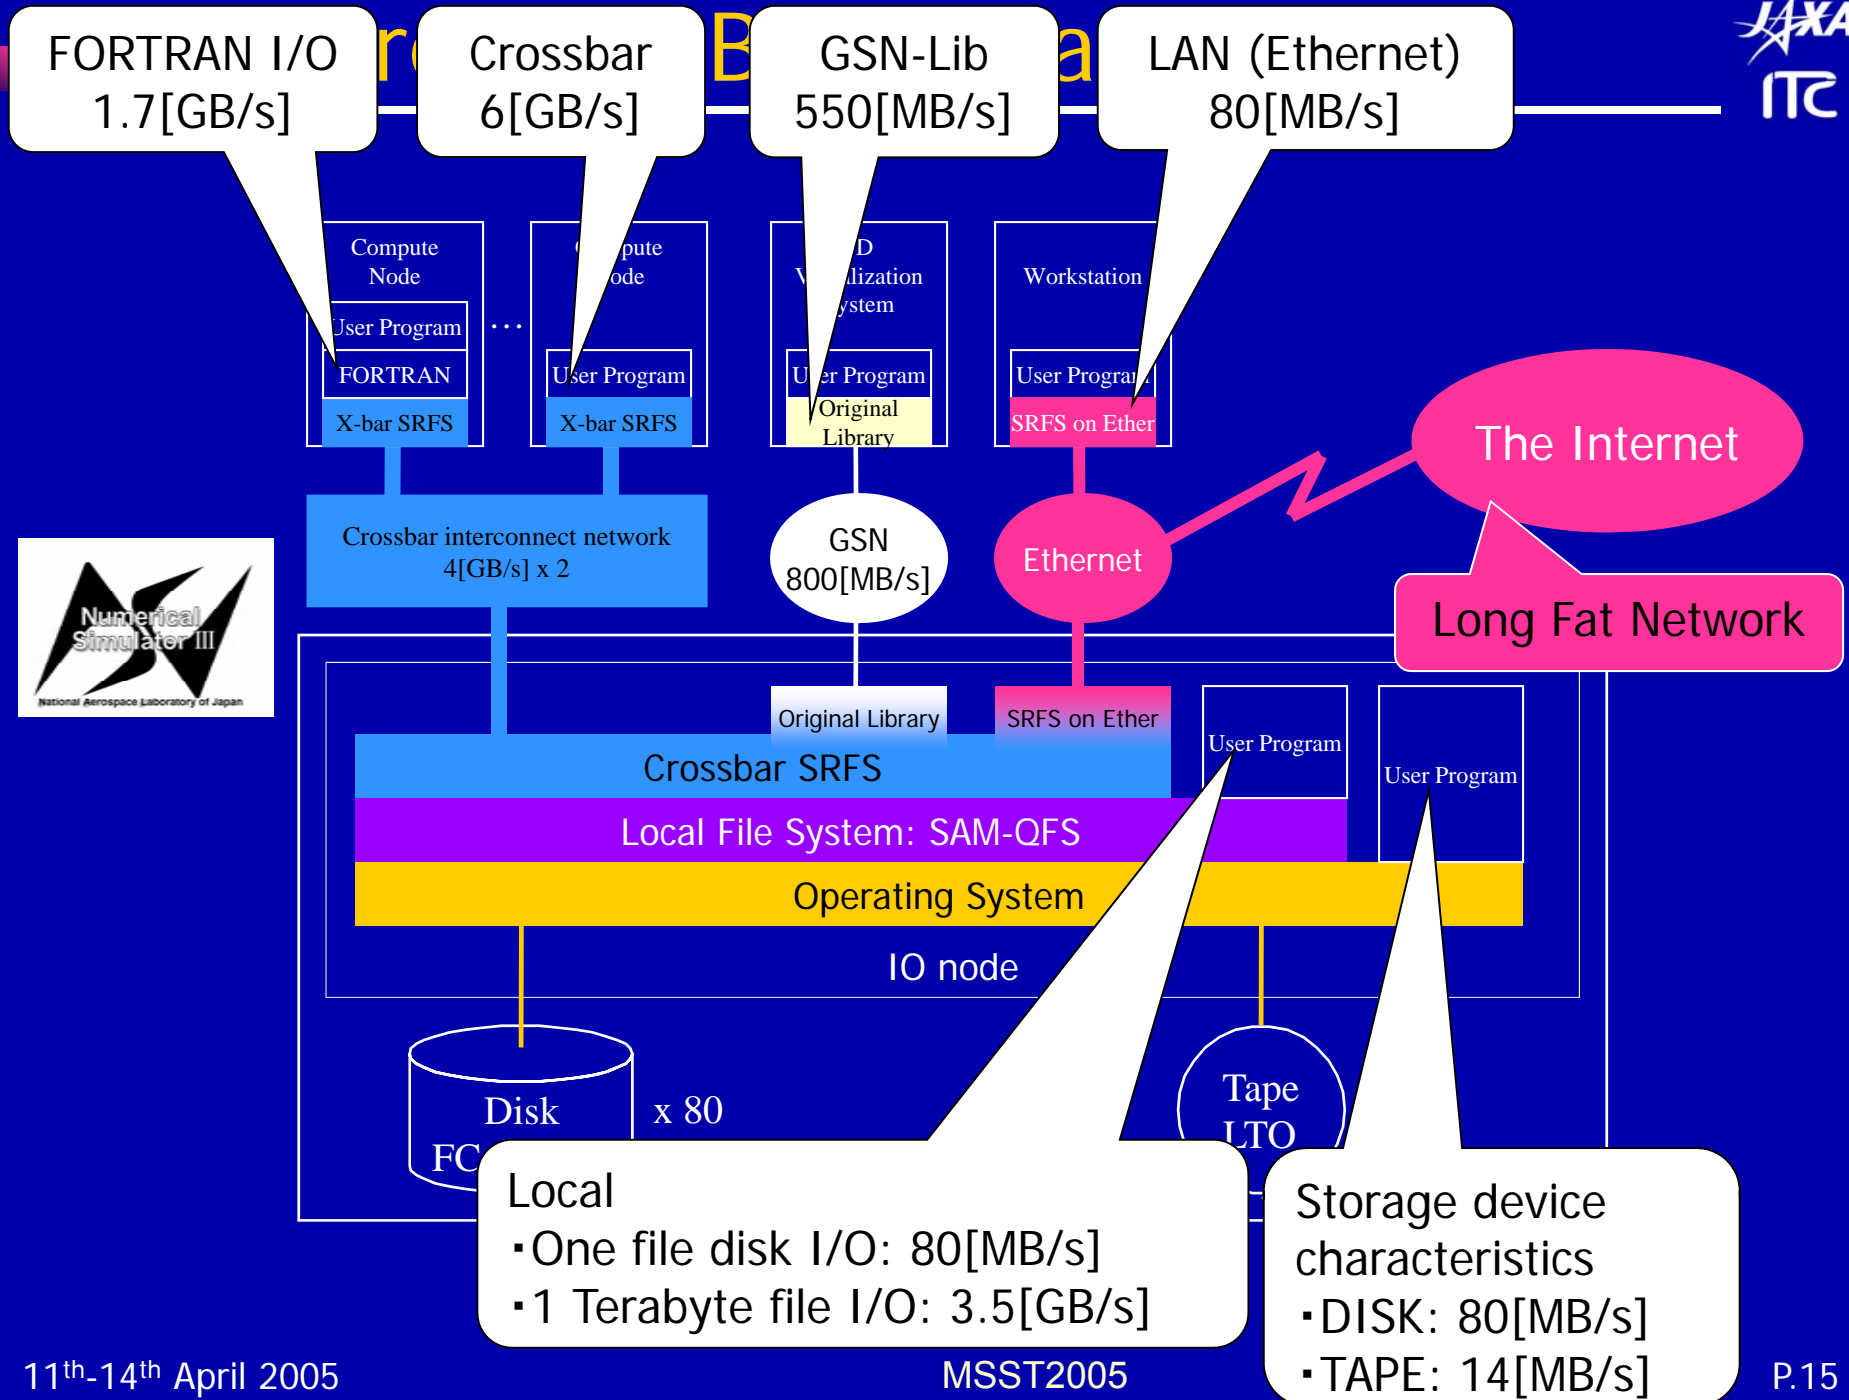
Coming Soon! 2005 May
Contact fujita@chofu.jaxa.jp

Thank you



fujita@chofu.jaxa.jp





Local

- One file disk I/O: 80 [MB/s]
- 1 Terabyte file I/O: 3.5 [GB/s]

Storage device characteristics

- DISK: 80 [MB/s]
- TAPE: 14 [MB/s]