




# Status and Plans for the Tier 1 Tokyo (Japan)

Hiroshi Sakamoto

International Center for Elementary Particle  
Physics (ICEPP), the University of Tokyo



# Contents

- PC Farm
  - Blade Servers
- Storage System
  - IDE File Server Prototype
- Wide Area Network
  - KEK-ICEPP GbE Connection
  - CERN-ICEPP Connectivity Study
- Short Term Plan





# HP Blade Server

- HP ProLiant BL20p
  - Xeon 2.8GHz 2CPU/node
  - 1GB memory
  - SCSI 36GBx2, hardware RAID1
  - 3 GbE NIC
  - iLO remote administration tool
  - 8 blades/ Enclosure(6U)
- 
- Total 108 blades(216 CPUs) in 3 racks
  - Rocks 2.3.2(RedHat 7.3 base) for Cluster management

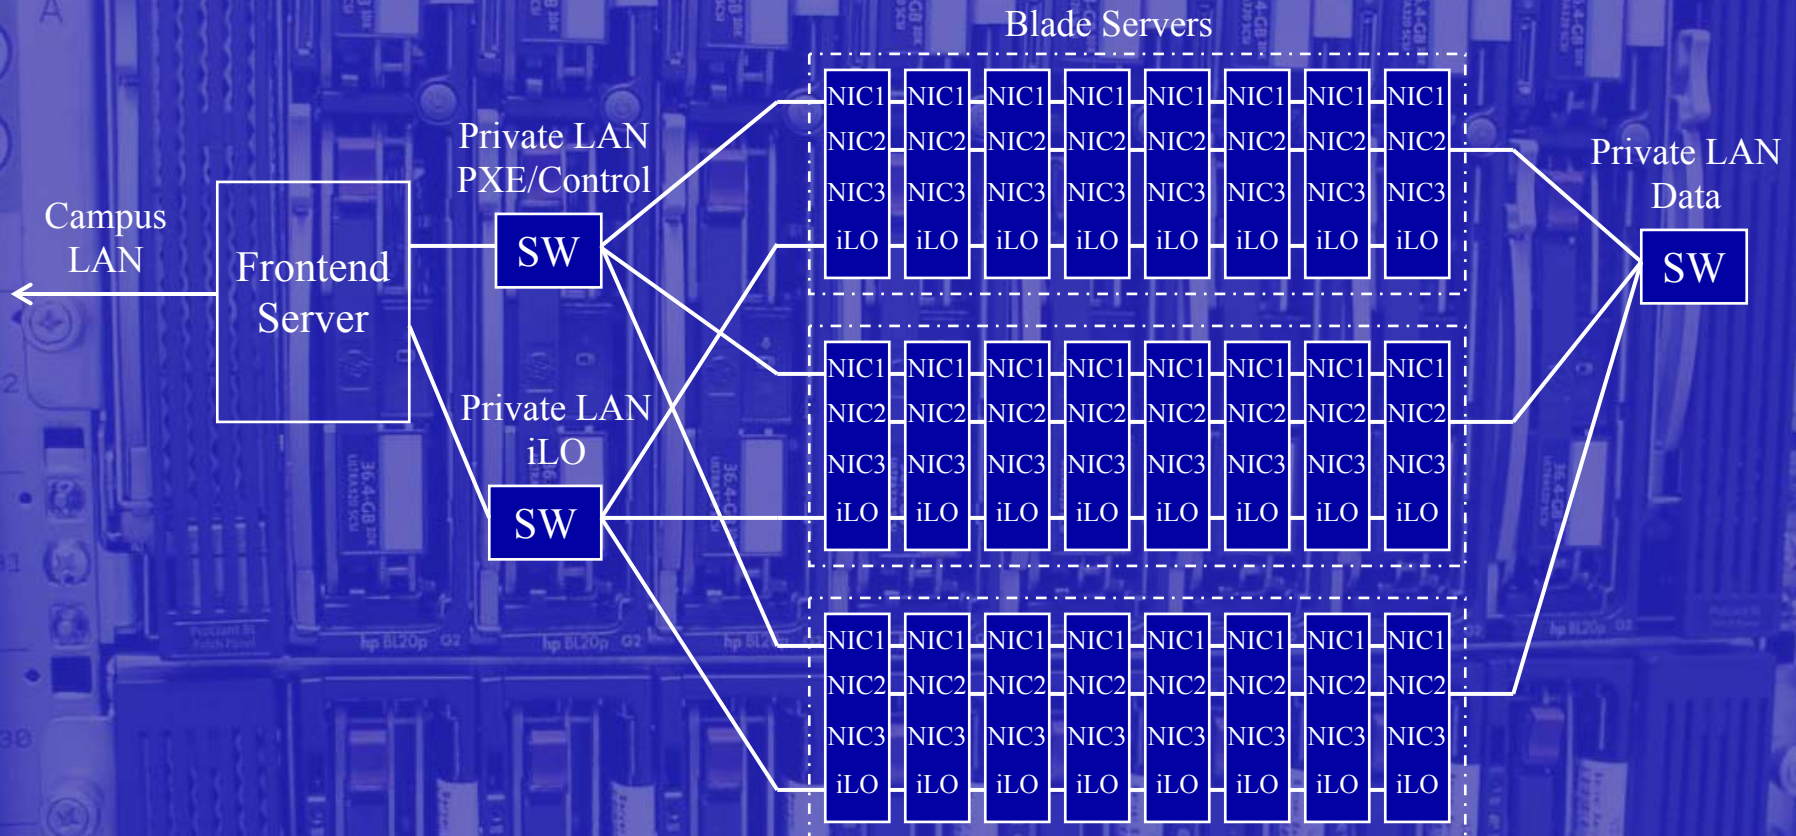









# HP Blade Server Network Configuration





# File Server Prototype

- 2 Xeon 2.8GHz (FSB 533MHz)
- Memory 1GB
- 2 PCI RAID I/F, 3ware 7500-8
- 2 GbE(1000Base-T) LAN
- 4U Rack mount Case, 16 hot-swap IDE HDD bay
- 16 250GB HDD + 40GB HDD
- 460W N+1 Redundant Power Supply
- 7 servers, total 24.5TB in 2003

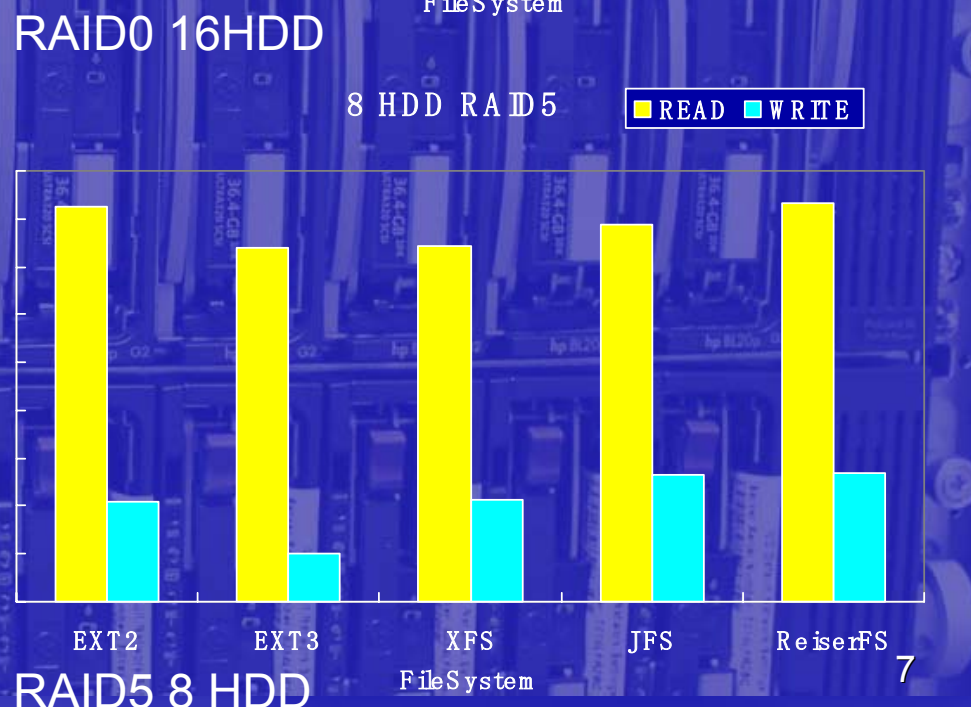
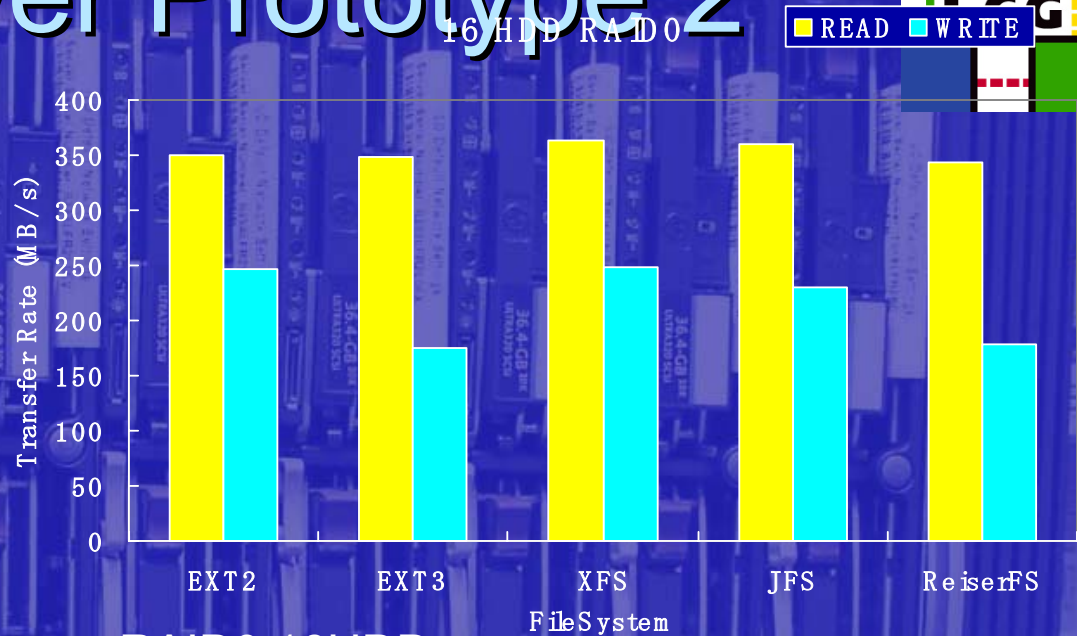




# File Server Prototype 2

## Benchmarks

- Read 350MB/s,  
Write 170-250MB/s  
RAID0 16 HDD
- Read 140-160MB/s,  
Write 20-50MB/s,  
RAID5 8 HDD







# Network Connectivity



## ● 10Gbps SuperSINET Backbone

- Connect major collaborating institutes

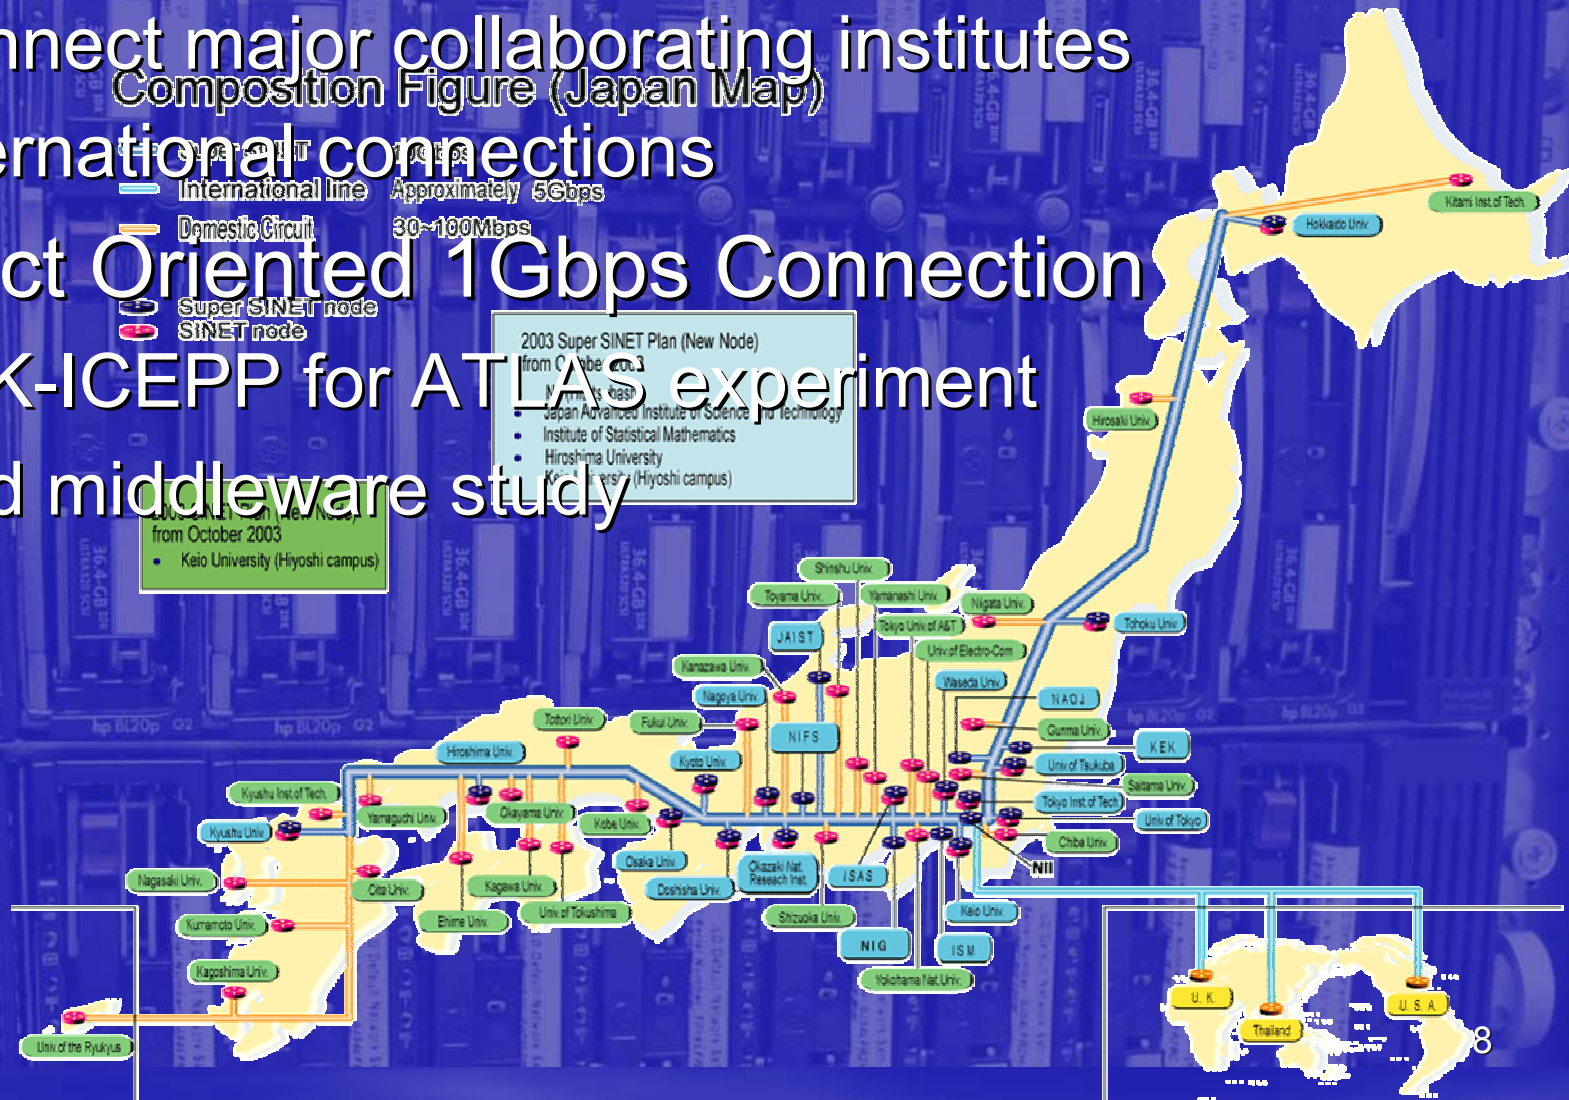
Composition Figure (Japan Map)

- International connections


## ● Project Oriented 1Gbps Connection

- KEK-ICEPP for ATLAS experiment

- Grid middleware study







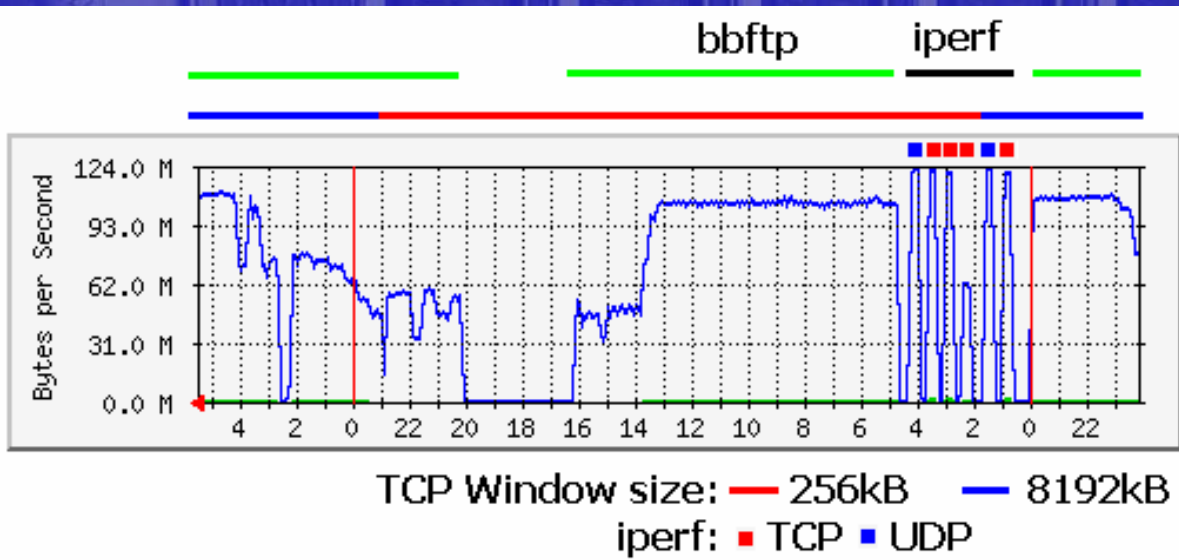

# WAN Performance Measurement

## "A" setting

- TCP 479Mbps -P 1 -t 1200 -w 128KB
- TCP 925Mbps -P 2 -t 1200 -w 128KB
- TCP 931Mbps -P 4 -t 1200 -w 128KB
- UDP 953Mbps -b 1000MB -t 1200 -w 128KB

## "B" setting

- TCP 922Mbps -P 1 -t 1200 -w 4096KB
- UDP 954Mbps -b 1000MB -t 1200 -w 4096KB

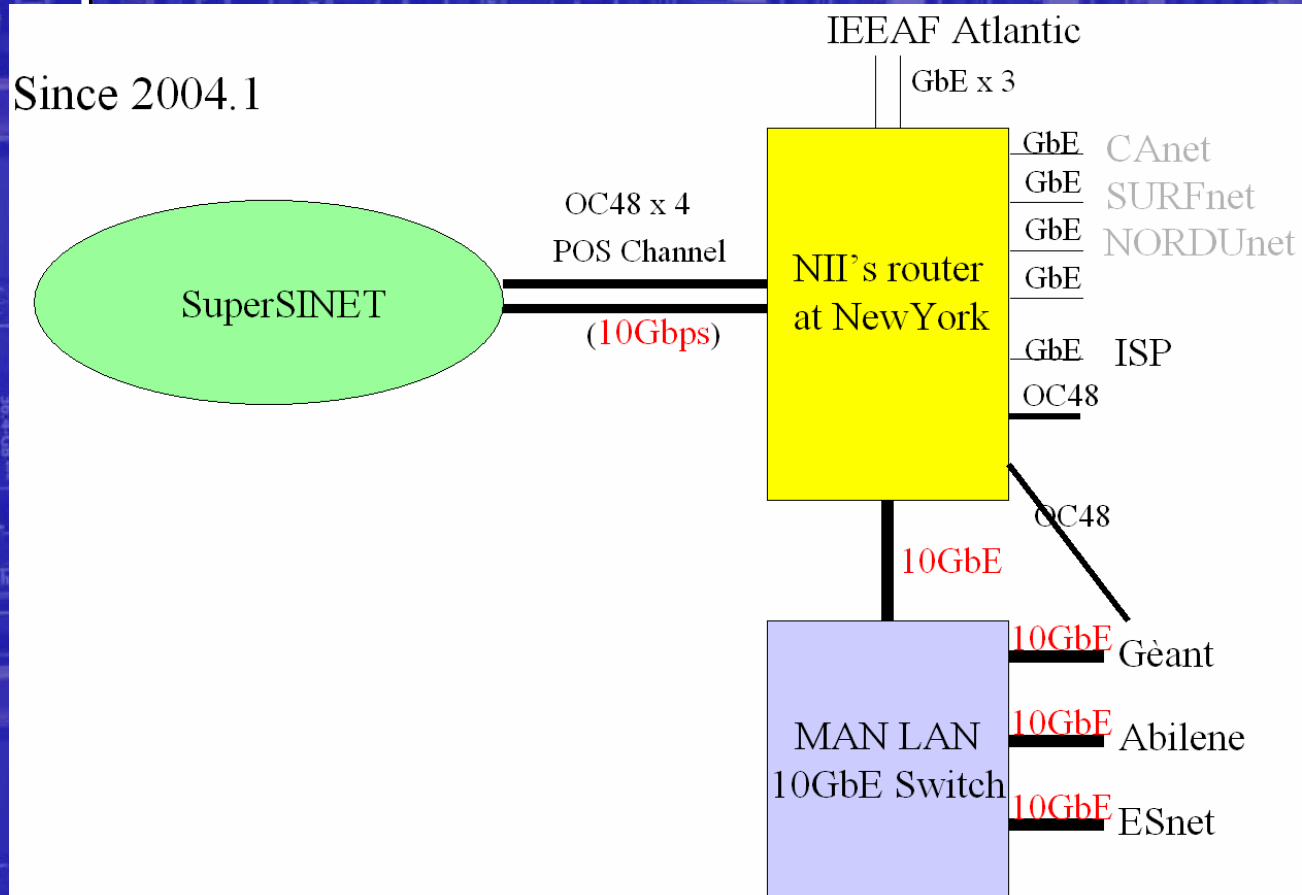
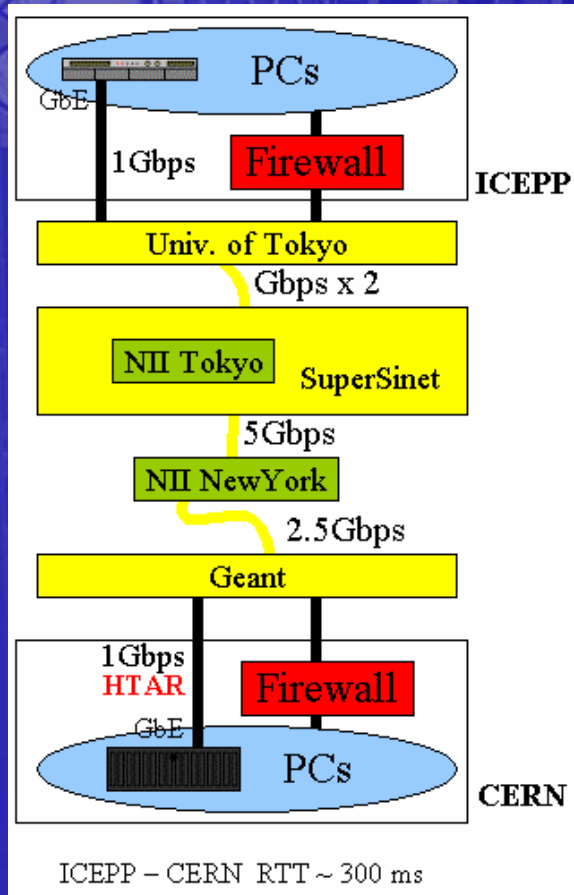


"A" setting: 104.9MB/s  
"B" setting: 110.2MB/s



# Connection to CERN

- 2.5Gbps x 2 (=5Gbps) to NY (2003)
- 2.5Gbps x 4 Now







# CERN-Tokyo (Very preliminary)

- Large RTT ~ 300ms
- Possibly certain amount of packet loss

[UDP] 17:00-17:30


option: -b 1000MB -w 256KB

[ ID] Interval Transfer Bandwidth Jitter Lost/Total Datagrams

[ 3] 0.0-1095.5 sec 116 GBytes 906 Mbits/sec 0.015 ms

4723300/89157875 (5.3%)

[ 3] 0.0-1095.5 sec 134933 datagrams received out-of-order








# Connection with neighboring countries

- SINET and APAN (as is on 02/02/2004)

	AU	CN	HK	JP	KR	TW	US
Australia (AU)		--	--	--	--	--	310
China (CN)	--		622	155	--	--	155
Hong Kong (HK)	--	622		--	--	155	45
Japan (JP)	--	155	--		2000	622	10000
Korea (KR)	--	--	--	2000		--	155
Taiwan (TW)	--	--	155	622	---		5000
US	310	155	45	10000	155	5000	



# Connection with neighboring countries (cont.)

- SINET and APAN (expected in 2004)

	AU	CN	HK	JP	KR	TW	US
Australia (AU)		--	--	--	--	--	20000
China (CN)	--		2500	155	--	--	155
Hong Kong (HK)	--	2500		--	--	155	10000
Japan (JP)	--	155	--		2000	622	30000
Korea (KR)	--	--	--	2000		--	155
Taiwan (TW)	--	--	155	622	---		5000
US	20000	155	10000	30000	155	5000	





# LCG-1

- Very limited resources allocated

Type	Host	System	Middleware	Install	Status Install	Status Operation	Comment	Local Comment
CE	dgce0	RH7.3	LCG-1-1_1_4	LCFGng	configured	running		by diskette
SE	dgse0	RH7.3	LCG-1-1_1_4	LCFGng	configured	running		by diskette
UI	dgui0	RH7.3	LCG-1-1_1_4	LCFGng	configured	running		by diskette
WN	dgwn0	RH7.3	LCG-1-1_1_4	LCFGng	configured	running		by diskette
WN	dgwn1	RH7.3	LCG-1-1_1_4	LCFGng	configured	running		by diskette
LCFG	dglcfg0	RH7.3		manual	configured	running		With TCP wrapper, 2NICs
CE	dgce1	RH7.3	LCG-1-1_1_4	LCFGng	configuring	preparing		PXE, 2NICs
SE	dgse1	RH7.3	LCG-1-1_1_4	LCFGng	configuring	preparing		PXE, 2NICs
WN	ldgwn0000	RH7.3	LCG-1-1_1_4	LCFGng	configuring	preparing		private network
WN	ldgwn0001	RH7.3	LCG-1-1_1_4	LCFGng	configuring	preparing		private network

# Short term plan

- Network Connectivity
  - CERN-Tokyo to 10Gbps now
  - Tokyo-Taipei connectivity study soon
- PC Farm / Mass Storage
  - ~100TB fiber-channel disks installed
  - Hierarchical storage study soon (IBM LTO2)
  - PC farm upgrade
- LCG-2
  - 60 Nodes (Dual 2.8GHz Xeon)
  - ~30TB Disk Space



# News

- Reorganization of our institute
  - 1994.4~2004.3 for LEP OPAL
  - Extension approved
- 2004.4~2014.3 for LHC ATLAS
  - Regional Center (Tier1) for ATLAS
  - Budget for R&D in FY2004
  - One engineer position
- New building
  - Completed in March 2005
  - 9th floor for RC

