Network Connectivity

R. Les Cottrell

<cottrell@slac.stanford.edu>

Stanford Linear Accelerator Center (SLAC)

Presented at SLUO Annual Meeting, Jul-15,1998

www.slac.stanford.edu/grp/scs/net/talk/sluo-jul98/

Overview

- SLAC's LAN
- Dial in access
- WAN connectivity & performance

LAN - Topology - Jul98

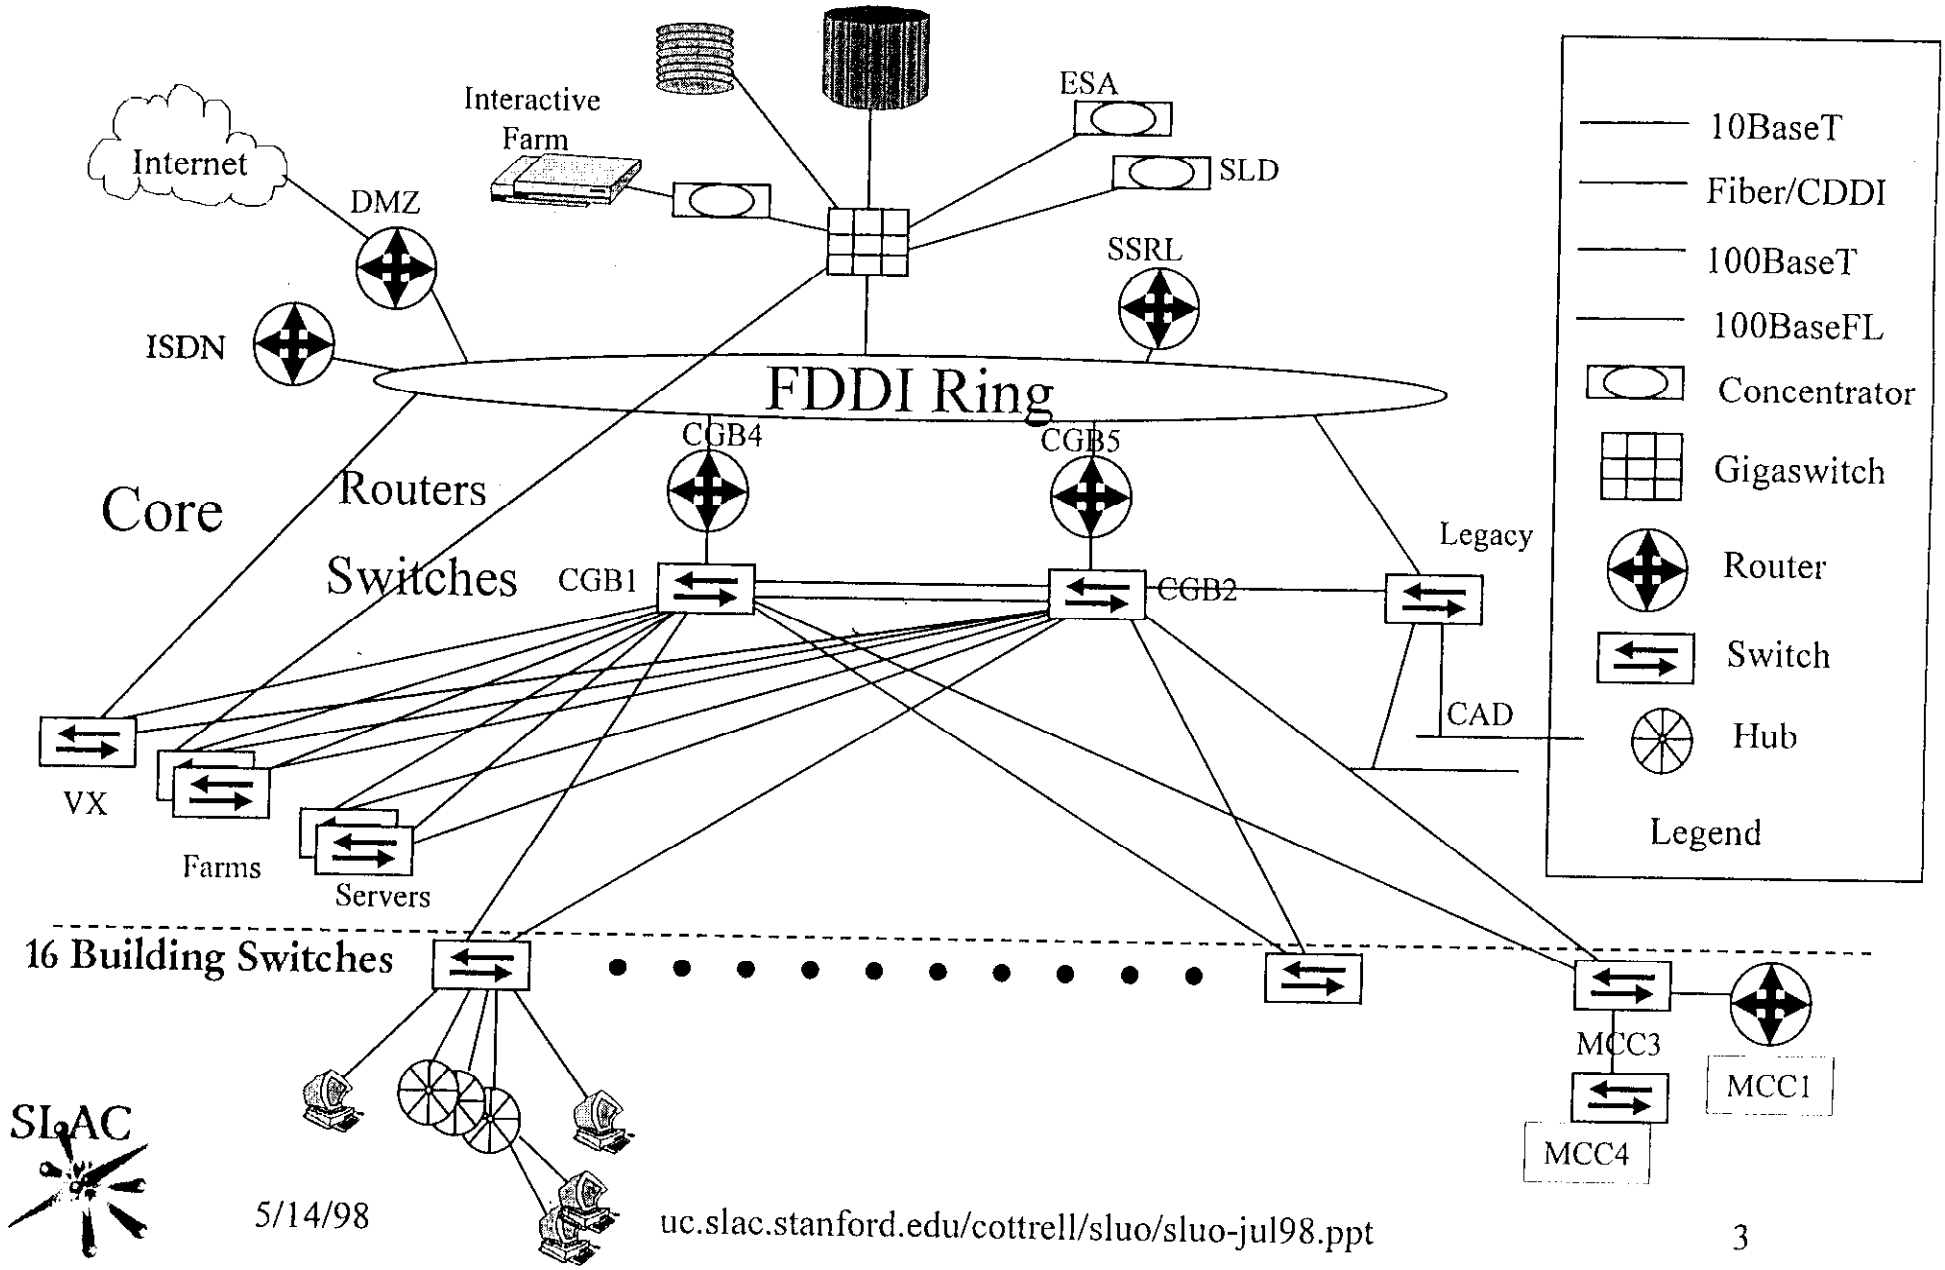

LAN -Status of Structured Wiring

- Imdividual cables with twisted copper wires between desktop & closet. Building closets connect to computer center by fiber
- Move started in 1995
  - Improved management & error isolation
  - Improved installation time
  - Enables switched networking
- About 70% of site completed (i.e. on switches or hubs)
- Plan to complete outside radiation fence in FY99.

LAN - Switched Network

- Based on mass market switched Ethernet
- Standard desktop has 10Mbps shared (via hub)
- Hubs connect to 10Mbps Building switch port
- Building switch connects to core switch at 100Mbps
- Core switches are interconnected at 100Mbps
- Core switches connected to core routers at 100Mbps
- Main servers connect via dedicated 100Mbps
- Use VLANs to provide instances of given subnets across many buildings

LAN - Reliability

- Redundant links with automatic failover to reduce impact of scheduled outages and improve reliability
- UPS for reliability
- Segmentation reduces impact of failure & simplifies id

LAN - Services Highlights

users, 27K msgs/day, notebook volume growing
ó / year

red new mail gateway

P server, evaluating clients

M blocking (1830 blocks, fairly stable)

relieved of entering IP addr/Gwy/DNS/WINS …

ort static & dynamic (e.g. roaming laptops)

admins can quickly register machines via form

The LAN - Growth

/uc.slac.stanford.edu/cottrell/slac/lanutil

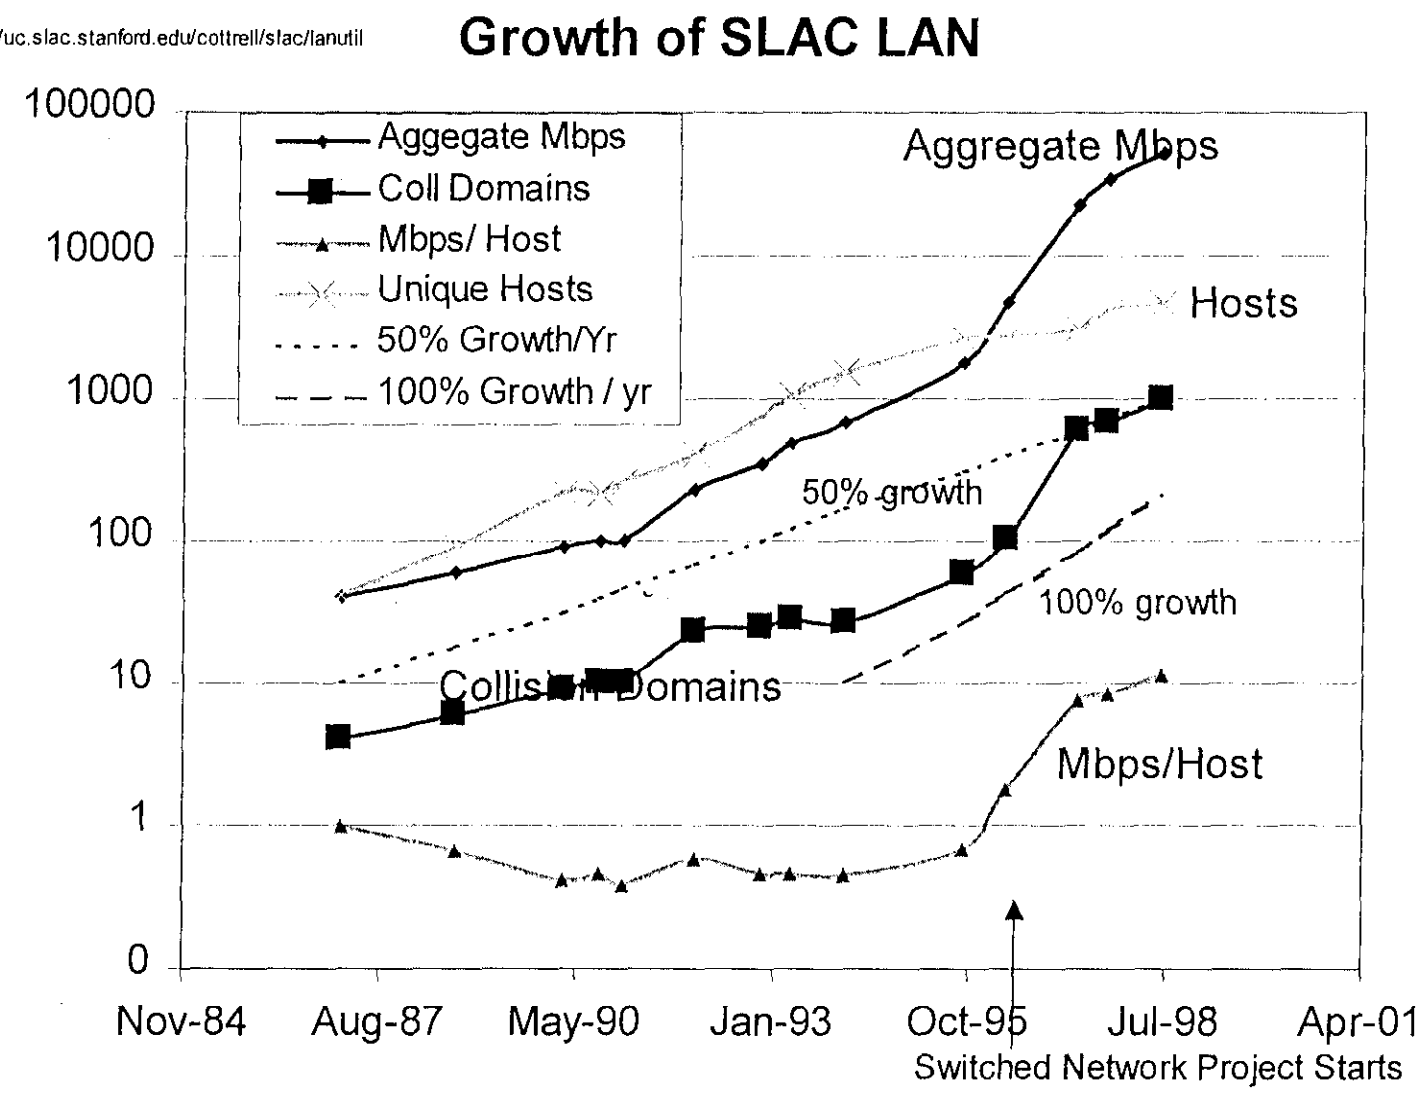

LAN - Next Steps

- Double aggregate bandwidth ~ every 12-18 months
  - more segmentation (hubs => switched ports, 1 host/collision domain)
- Dedicated 100Mbps Ethernet to power user desktops
- Gbps trunks between switches & core routers
- Replace FDDI rings with high speed switched core
- Higher speed routing & more integrated with switching
- Increase UPS & MG backup

Dial-in

- Wireless thru Ricochet
- ISDN
  - > 60 users, production for ~ 9 months
  - typical day 40 different users, 20 simultaneous, 3hrs/user/day
  - high degree of satisfaction
  - startup more expensive than modems
- Voice modem
  - through campus 14.4kbps - getting rusty
  - ISP (e.g. Netcom $20/month) nationwide
  - ARA 33.6kbps ~ 340 accounts - getting rusty

Dial in - Futures

- Disappointed with outsourcing dial-in
- Plan for direct dial-in PPP at < 56kbps
  - ready for pilot users
  - www2.slac.stanford.edu/comp/net/ppp/
- Further out:
  - Have a few users on xDSL thru Stanford
    - higher speeds, leased line, double ISDN cost
  - Couple of users trying cable modems

WAN - Internet

- ESnet
  - 45Mps to Sprint
  - =>155Mbps
- Stanford 10Mbps

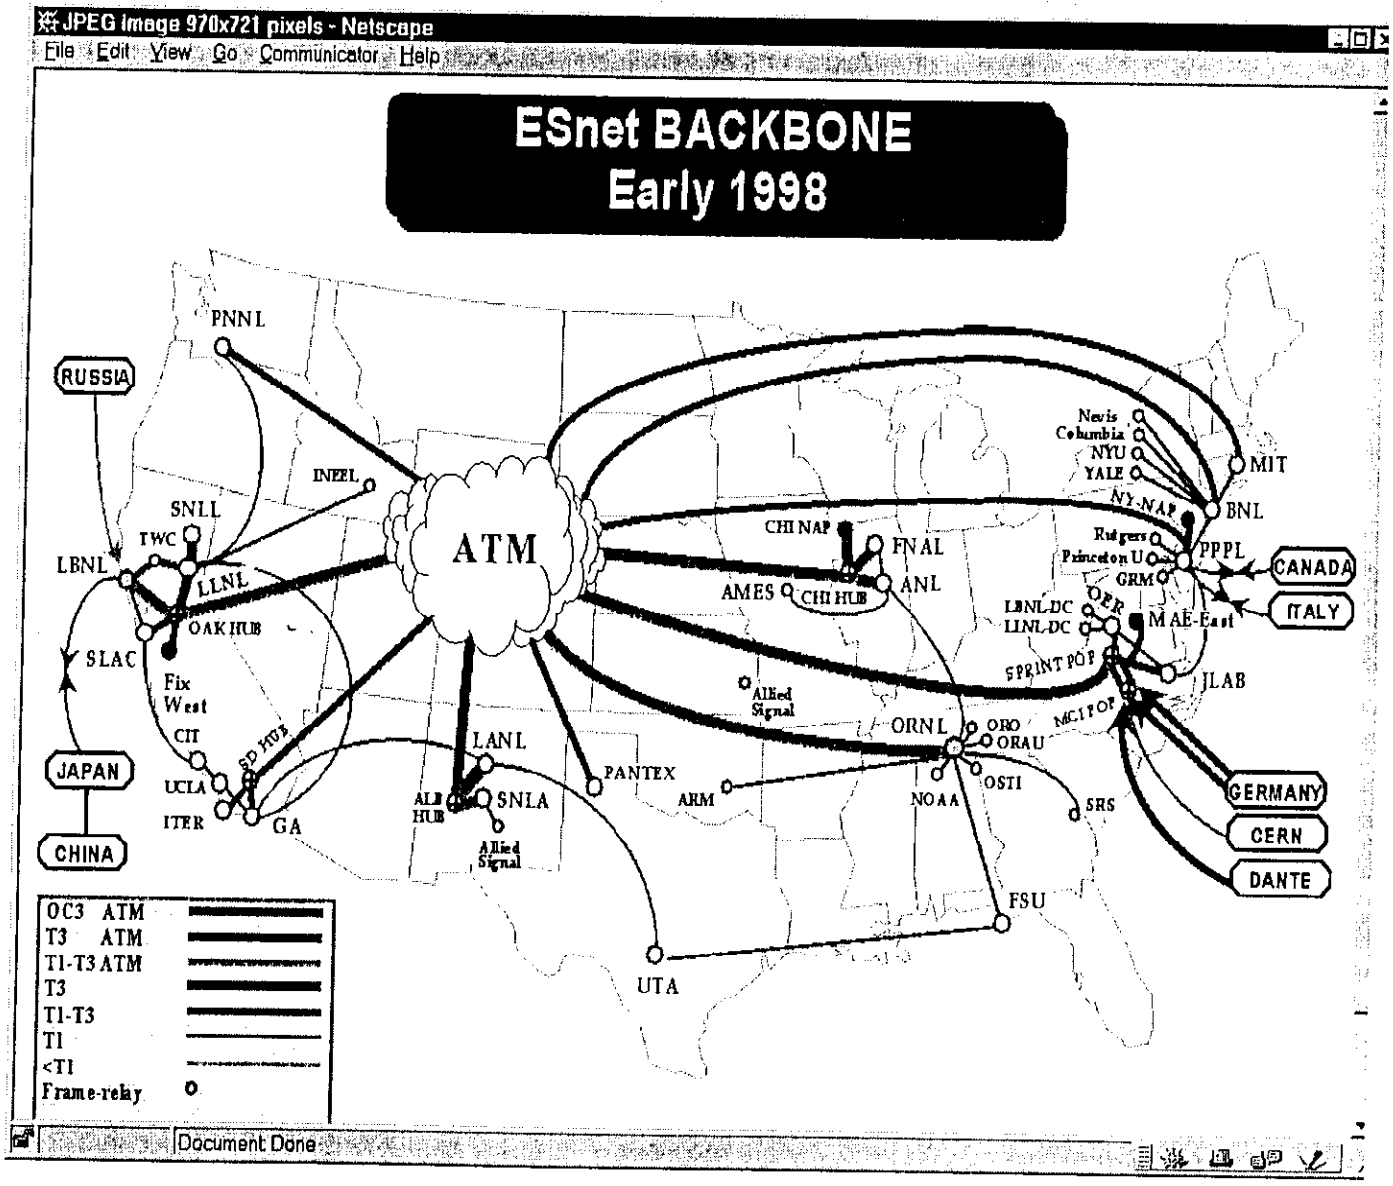

WAN - Performance Environment

- Most European & Japan traffic carried via national A&R nets
- Most US traffic carried by ESnet or Internet MCI
- At least 20 different transatlantic routes with own financial arrangements, packet loss and performance
- Cost of bandwidth varies, US one of cheapest
- HENP only small part of traffic carried
- In most cases no there is priority for HENP

WAN - HENP Use

- ICFA estimates factor 10 traffic growth in 4 years
- HENP (SLAC & CERN) profile different from typical Internet traffic:
  - ~20-60% traffic is data transfer
  - Web 15-40% (c.f. Internet 70%)
  - AFS 6-8%
  - Xwindows ~ 5-10%
  - Remainder mainly Telnet/rlogin, plus email, news, video

WAN - Performance Monitoring

- Internet woefully under-measured, starting to improve. No single path typical of whole
- World-wide HENP participation in measurements
- Set of tools known as PingER, originally developed at SLAC, based on echoing packets
- 15 Esnet/HENP sites in 8 countries monitoring over 900 links in 22 countries
- Data going back years
- Recently defined 50 "beacon sites" that all sites monitor

WAN - Performance Metrics

- Packet loss identified as critical quality indicator
  - below 1% smooth performance
  - > 2.5-5% interactive (telnet, Xwindows, packet video ...) work becomes problematic
  - > 12% interactive unusable
- Fortunately Email & Web not so sensitive

WAN Performance - US <=> US 1/2

- Within ESnet excellent (median loss 0.1%)
- To vBNS/I2 sites very good (~ 2 * loss for ESnet)
- DOE funded Universities not on vBNS/ESnet
  - acceptable to poor, getting better (factor 2 in 6 months)
  - lot of variability (e.g.)
    - Brown^{T}, UMass^{T} = unacceptable(>= 12%)
    - Pitt*, SC*. ColoState*, UNM^{T}, UOregon^{T}, Rochester*, UC*, OleMiss*, Harvard^{1q98}, UWashington^{T}, UNM^{T}= v. poor(> 5%)
    - Syracuse^{T}, Purdue^{T}, Hawaii* = poor (>= 2.5%)
  - * = no vBNS plans, ^{T} = vBNS date TBD, ^{V} = on vBNS

WAN - Performance - US <=> US 2/2

- A year ago we looked at Universities with large DOE programs
- Identified ones with *poor (>2.5%)* or **worse (>5%)** performance
  - **Harvard**[1q98] = very poor (>= 5%)
  - **JHU**[V], **UOregon**[*], **Duke**[V], *UCSD*[V], *UMD*[V], *UMich*[T], *UColo*[V], *UPenn*[T], *UMN*[V], *UCI*[T], **UWashington**[T], *UWisc*[V] = acceptable (>1%)/good
  - * = no vBNS plans, [T] = vBNS date TBD, [V] = on vBNS

WAN - Performance - Canada

- 23 of 50 major universities connected to CA*net2 (incl. 8 of 10 HENP major sites)
- Seems to depend most on the remote site
  - UToronto bad to everyone
  - Carleton, Laurentian, McGill poor
  - Montreal, UVic acceptable/good
  - TRIUMF good with ESnet, poor to CERN

WAN - Performance - Europe

- Divides up into 2
  - TEN-34 backbone sites (de, uk, nl, ch, fr, it, at)
    - within Europe good performance
    - from ESnet good to acceptable, except nl, fr (Renater) & .uk are bad
  - Others
    - within Europe performance poor
    - from ESnet bad to: be, es, il, hu, pl acceptable for cz

WAN Performance - Asia

rael bad

EK & Osaka good from US, very poor from anada

okyo poor from US

pan-CERN/Italy acceptable, Japan-DESY bad

SU bad to Moscow, acceptable to Novosibirsk

iina is bad with everywhere

WAN Performance - Intercontinental

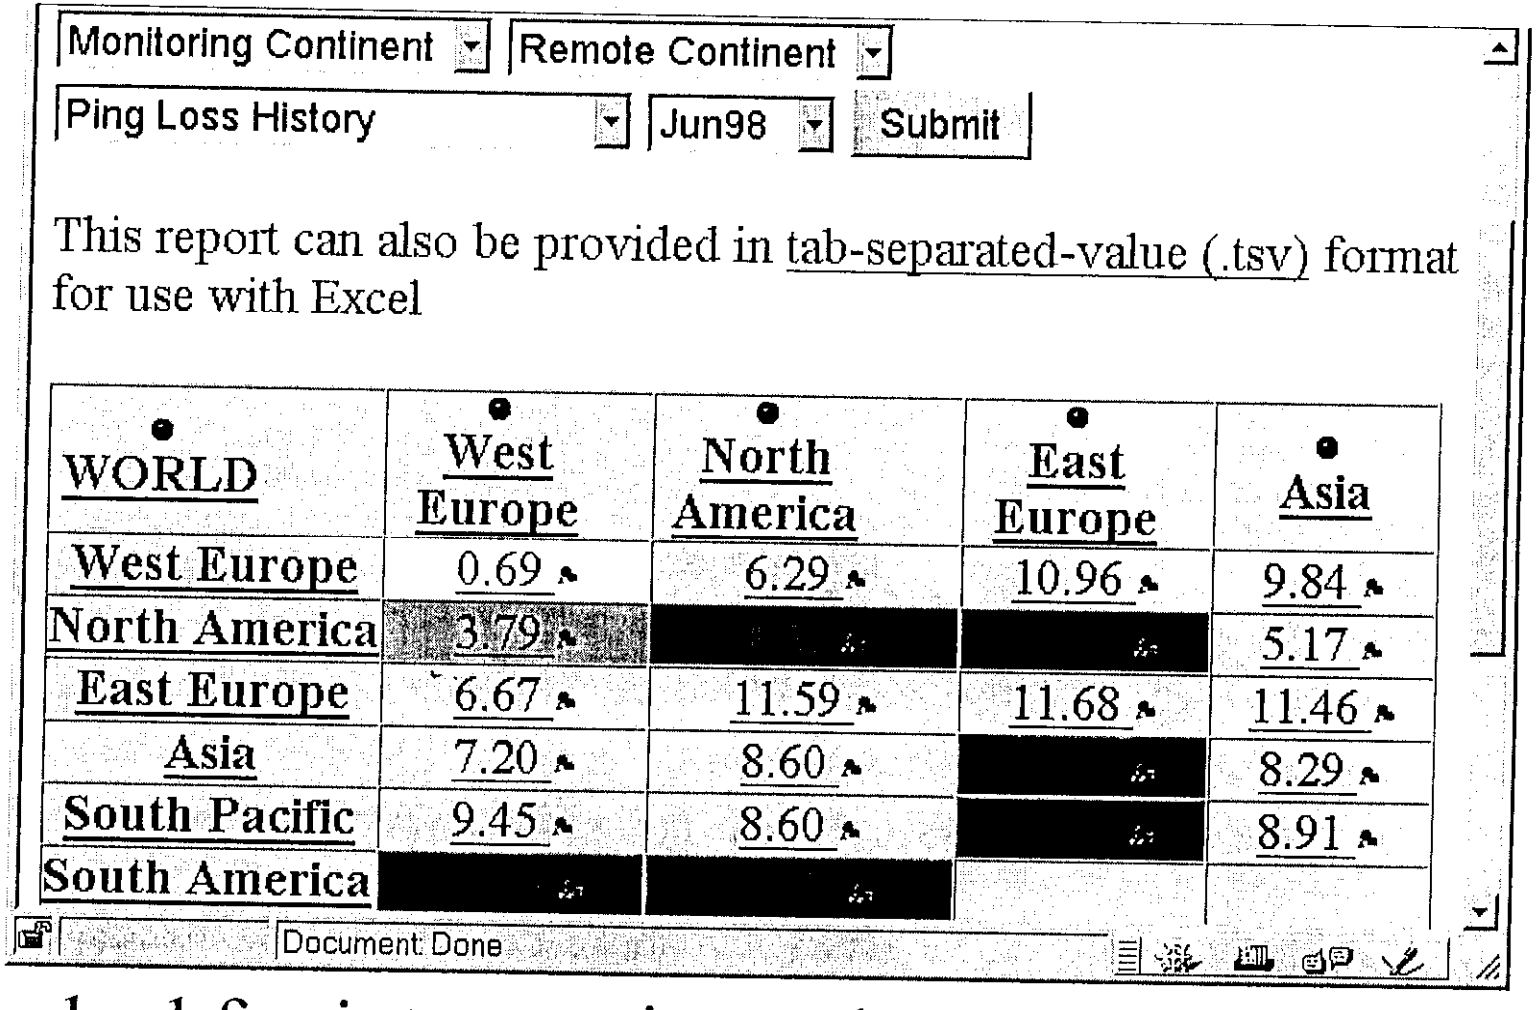

WORLD	West Europe	North America	East Europe	Asia
West Europe	0.69	6.29	10.96	9.84
North America	3.79			5.17
East Europe	6.67	11.59	11.68	11.46
Asia	7.20	8.60		8.29
South Pacific	9.45	8.60		8.91
South America				

Looks pretty bad for intercontinental use

Improving (about factor of 2 in last 6 months)

WAN - Performance - Summary

- Performance worse when source & destination on different ISPs, nets need to interconnect
- Some interconnects are very bad
  - e.g. MAE-West, MAE-East, but changes with time
  - Private peering to avoid congestion points
- Transatlantic important & bad

WAN - Performance Futures

- Increased bandwidth
  - WDM (factor 4-16 today, going to 100)
  - Competition to traditional carriers (e.g. Qwest)
  - Intra continent
    - US: More sites on I2, second I2 backbone (Abilene)
    - Europe TEN-34 => TEN-155
  - Inter continent more problematic
- Differentiated services: policy tag packets and prioritize through Internet (premium class service)
- Improved understanding: increased measurement of end-to-end performance & identifying bottlenecks

Further Information

- DHCP at SLAC
  - www2.slac.stanford.edu/comp/net/dhcp/dhcp.htm
- Direct dial-up PPP pilot at SLAC
  - www2/comp/net/ppp/
- Email: www/comp/net/email/
- ICFA Monitoring WG home page (links to status report, meeting notes, how to access data, and code)
  - www/xorg/icfa/ntf/home.html
- WAN Monitoring at SLAC has lots of links
  - www/comp/net/wan-mon.html