

NETWORK UPGRADE FOR THE SLC: PEP-II NETWORK

M. Crane, M. Call, S. Clark, F. Coffman, T. Himel, T. Lahey, E. Miller, R. Sass
Stanford Linear Accelerator Center, Stanford University, Stanford CA 94309

Abstract

The PEP-II control system required a new network to support the system functions. This network, called CTLnet, is an FDDI/Ethernet based network using only TCP/IP protocols. An upgrade of the SLC Control System micro communications to use TCP/IP and SLCNET would allow all PEP-II control system nodes to use TCP/IP. CTLnet is private and separate from the SLAC public network. Access to nodes and control system functions is provided by multi-homed application servers with connections to both the private CTLnet and the SLAC public network. Monitoring and diagnostics are provided using a dedicated system. Future plans and current status information is included.

I. INTRODUCTION

The PEP-II control system is a combination of the SLC Control System and EPICS¹. Together these two control systems provide a way to match an application with a control system methodology which best matches the applications needs. The SLC Control System is based on a DEC Alpha OpenVMS central host (MCC) with distributed microprocessor (SLCmicro) front ends and networked displays. The displays use TCP/IP, but the SLCmicros use a SLAC-developed communications system. EPICS for PEP-II is based on Sun UltraSparc Solaris servers with distributed microprocessor front ends (IOC), all based on TCP/IP and Ethernet. A common network is needed to serve these and other PEP-II control system nodes. We expect that this single new network will show a cost savings over time due to the commercial availability of products, new advances to those products, and standard diagnostic tools and experience.

II. COMMON PROTOCOLS

All EPICS IOCs and operator interfaces (OPI)s, as used in the PEP-II control system, use Ethernet as the physical layer. Since the IOCs and local OPIs are distributed around the PEP-II ring, some sort of Ethernet-based technology must be available in these parts of the ring. Parts of EPICS also run on the MCC central host, so MCC must be able to acquire data from IOCs around the ring. The display portions of the SLC control system are TCP/IP-based using the X-windowing system. There are various Ethernet-based X-terminals distributed around the PEP-II ring, so Ethernet must be provided for these terminals. Various other systems such as GPIB controllers and other com-

mmercial instrumentation use Ethernet and TCP/IP to communicate.

An upgrade of the SLCmicros and the SLC Control System to use TCP/IP and Ethernet would provide a common communications path for the PEP-II control system. There are about 70 Intel Multibus I based Intel 386/486 single board computers, or SLCmicros, in the SLC Control System. Communications between the MCC central host and the SLCmicros currently use a SLAC-developed network called SLCNET. The anticipated amount of additional bandwidth required for PEP-II SLCmicros is greater than the current available bandwidth on SLCNET. This, combined with the desire for common physical layer interfaces such as FDDI/Ethernet, and the potential cost savings in maintenance and commercial upgrades, brought us to the conclusion that SLCmicros should be updated to use TCP/IP and Ethernet.

A new Multibus I single board computer from Radi-sys has been chosen as the replacement card for the current SLC micro hardware. This new card, named the Skater card or EPC (Embedded PC), is actually a PC-based computer with an embedded Ethernet interface mounted on a Multibus I card. The EPC is cost effective in that the card is less expensive than the current SLCmicro hardware and allows memory to be upgraded using standard SIMM memory. This hardware/software upgrade has allowed us to modify our SLCmicro code to operate on a PC-like computer, use Ethernet and TCP/IP, but retain our investment in Multibus/CAMAC systems. The software to do this upgrade is the subject of another PAC97 paper.²

III. MACHINE TO MACHINE NETWORK

The network to support these TCP/IP-based nodes must provide enough throughput to run the control system and have available bandwidth for future applications. The fact that this network is compatible with the SLAC public networks and the Internet makes it very important to isolate this network. Bandwidth on this network is reserved for control system usage only. Traffic includes control system computers communicating with control system computers and operators/users communicating with their systems. This excludes someone's everyday work traffic such as e-mail, editing, and WWW access. There is also the need to protect the control system nodes from non-users and would-be hackers, since these nodes provide the basis for a 'factory' with an expected uptime of 99%.

The CTLnet is a separate network from all other networks at SLAC. To allow users to access their systems on this network, there are three gateway machines which the user can log onto as shown in figure 1. CTLnet uses an RFC 1918 Private Internet Address as an added measure

*Work supported by Department of Energy, contract DE-AC03-76SF00515

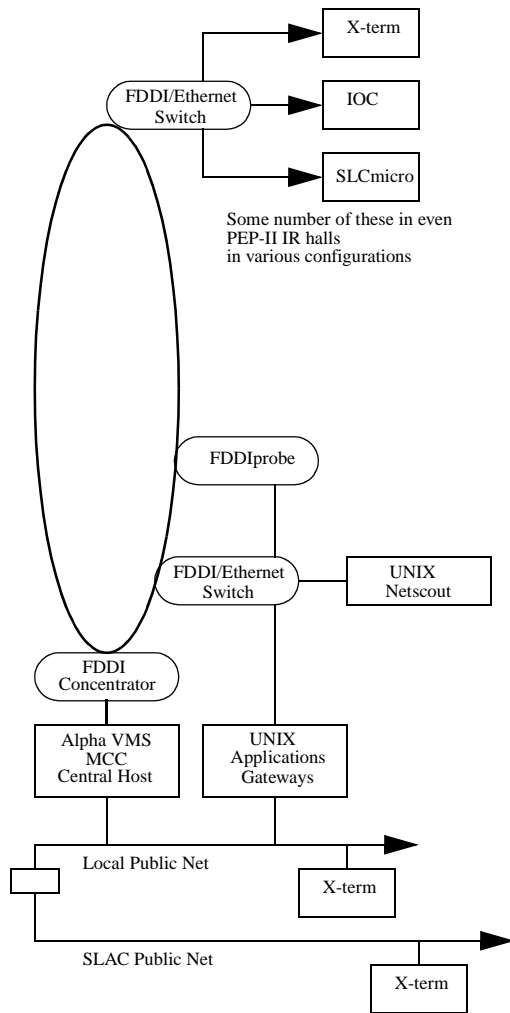


Fig. 1. PEP-II CTLnet.

to promote the idea that this is not a normal public network. The Private IP subnet addresses for CTLnet were chosen so that they would overlay a real SLAC subnet, and that SLAC subnet was reserved for future use. This precaution was taken to allow us to move from the Private IP address space to normal SLAC address space if problems should arise. This capability has already been utilized during commissioning. A router has been installed to allow packets from select internal nodes to be presented to the SLAC public network. This allows users to initially connect their EPICS OPI and IOCs in the control network while maintaining their normal development environment. As the UNIX application gateways are improved, we will implement the unrouted private addresses.

Making this network separate also allows a better understanding of the local network traffic since all traffic is generated and consumed within the local address space.

This also keeps CTLnet traffic from “cluttering up” the public networks.

IV. APPLICATION GATEWAYS

There are three application gateway machines which provide user access to CTLnet and the control systems resident there. These gateways are multi-homed systems with a network connection to the local SLAC public network and a separate connection to CTLnet. These gateways are not configured as routers, so traffic does not pass directly through them from one network to another. The user must first log onto a gateway and then log onto an internal CTLnet node. Note that the normal SLAC network environment is not available on CTLnet, including public printers and file servers, since this traffic is blocked by the gateways.

MCC is the OpenVMS central-host gateway for the SLC Control System. This gateway provides the X-window touch panel interface for access to the SLC Control Program (SCP). The SCP X-window traffic can be directed to a user’s X-terminal on either network, allowing access from a user’s office at SLAC or from somewhere on the CTLnet. When a user is logged onto MCC, internal CTLnet nodes are accessible with such tools as telnet and ftp. Some EPICS functions are also available on MCC, such as the launching of EPICS displays.

There are two UNIX gateways which provide access to the EPICS control system. Like MCC, these EPICS OPIs allow a user to log on and start a display to an X-terminal on either network. The EPICS application then gathers data from IOCs on CTLnet and displays the information to the user. The full complement of tools to run EPICS reside locally on these gateways, so they can operate with no outside references such as the SLAC UNIX environment. This allows the PEP-II control system to be completely removed from the SLAC public network if such a need arises. The tools include support for IOC booting, math tools, display tools, and basic UNIX administration applications.

The gateways are also the common interface for other types of network-based instrumentation such as spectrum analyzers, GPIB instruments, etc. This list will most certainly grow as time passes.

V. FDDI AND ETHERNET

The requirements for the new network made the physical layer choices relatively straightforward. There are clusters of Ethernet-based nodes spread around the PEP-II ring and in the Main Control Center. There was an existing fiber optic cable plant connecting nearly all the cluster locations. The VMS host already had an FDDI interface and software. A high speed backbone was needed which could provide control system growth over the next few years and provide a flexible local interface to Ethernet-based nodes. Fault tolerance is required for single area power failures. The availability of commercial diagnostic and monitoring systems was required to keep failures and

downtime to a minimum. There is also a need to continue the backbone down the linac as a future expansion adding another four miles to the total backbone length.

FDDI matches these needs very well. FDDI is a high speed backbone to connect stations nearly 1000 meters apart using fiber optics, connected as a dual FDDI ring to provide single area fault tolerance. There is an abundance of commercially available support tools and experience for both the FDDI hardware and software. The VMS host supports DEC hardware and TCP/IP using Multinet for the FDDI interface.

The network backbone initially includes the Main Control Center (MCC) and the periphery of PEP-II with the total ring distance of approximately one and one half miles. Nearly all of the Ethernet-based nodes are located in the even-numbered interaction regions. A Cisco Catalyst 1200 switch, which has one dual-access FDDI port and eight 10baseT (twisted-pair) Ethernet ports, is installed in each even-numbered interaction region to support the Ethernet requirements. Each Catalyst switch is located so that most or all of its Ethernet customers can be reached within the standard 100-meter limit of a 10baseT twisted-pair cable run. We assume that 10baseT Ethernet connections will generally be used, unless there is some overriding need for a specific node to use some other type of connection. We will use a dedicated port from the Catalyst for each SLCmicro and EPICS IOC. For other devices, we may choose to use either a dedicated port from the Catalyst or shared port from a multiport repeater (hub) connected to the Catalyst. This will depend mainly on the number of Ethernet devices and their geographical distribution relative to the Catalyst. The hubs provide 12 10baseT Ethernet connections from a single Catalyst 10 baseT port and can be monitored remotely.

The MCC central host has an FDDI interface connected to the PEP-II network via an FDDI concentrator. The UNIX gateway machines will be connected via Ethernet ports on a Catalyst switch in the Main Control Center.

VI. MONITOR AND DIAGNOSTICS

The monitoring and diagnostic system for this network is modeled after the SLAC network monitoring facilities. A separate, dedicated system must be implemented for this private network rather than using the public SLAC services.

The system uses the SNMP protocol to access information from Ethernet hosts and network devices such as the Cisco Catalyst switch, hubs, and others. The FDDI information comes from a NetScout FDDIprobe using the

RMON protocol. NetScout management software running on a SUN station completes the system.

Other tools are also available to look at specific traffic on the CTLnet. These tools include portable analyzers and the normal UNIX-style network utilities.

VII. FUTURE PLANS

The first expected expansion will be the extension of the FDDI backbone down the SLC linac. Fiber is already installed and reserved for this usage. The expected plan would be to install FDDI/Ethernet switches and add these new sections to the existing FDDI ring. Then as time, schedule (and budget) permit, new EPC SLCmicros would be installed and enabled to use Ethernet. There are nearly 70 existing micros, so an incremental installation plan with accompanying SLCNET decommissioning must be devised for each micro.

VIII. SYSTEM STATUS

The network has been installed for six months now and is being used for PEP-II commissioning. As mentioned earlier, the network is currently routed as a public network to allow developers easy access to their normal development environment while commissioning. The router was installed since the EPICS application gateways are not yet on-line to support the developers. This has fallen behind partially because the new Sun hardware uses Solaris 2.5, whereas all of the EPICS work so far at SLAC has been with SunOS.

The monitoring systems are installed but not completely configured, so accurate information is not yet available.

Delays in developing software and hardware for the EPC SLCmicro has caused PEP-II to use the standard SLCNET SLCmicro for commissioning. The Ethernet-based software will be released when more testing has been completed. It is expected to install the new EPC SLCmicros as time permits.

IX. REFERENCES

- [1] L. Dalesio et.al., The Experimental Physics and Industrial Control System Architecture: Past, Present, and Future, Proc. ICALEPCS, Berlin, Germany, 1993.
- [2] M. Crane et.al., Network Upgrade for the SLC: Control System Modifications, PAC, Vancouver, BC, Canada, 1997.

X. TRADEMARKS

- *OpenVMS is a trademark of Digital Equipment Corporation.
- *Alpha ia a trademark of Digital Equipment Corporation.
- *UltraSPARC is a trademark of SPARC International, INC.
- *Solaris and SunOS are trademarks of Sun Microsystems, INC.
- *Catalyst and MultiNet are trademarks of Cisco Systems, INC.
- *NetScout is a trademark of NetScout Systems, INC.