

# Network Terminal Accelerator

By Eddie Ho, Jim Gallagher, Kent Malave, and David Phipps

Users can be attached to multi-user AIX systems either by terminal adapters or by a LAN. The challenge for any multi-user system is to apply processing power to each user's application efficiently. IBM's Network Terminal Accelerator can improve overall system performance for users on an Ethernet LAN by off-loading the LAN protocol execution from the host.

The growing demand and competition for CPU resources require managers to search for ways to deliver computing services at a minimal cost. One way to reduce costs is to use a Micro Channel multiport adapter to attach terminals to a system. Terminals are often attached across a LAN for convenience and flexibility; however, this method is inefficient because of the software protocol overhead of network connections. Implementing this software in an adapter eliminates this inefficiency, resulting in lower CPU costs.

One question facing many managers is how to increase the number of regular AIX users without overloading the CPU. There are several potential solutions to this problem:

- ◆ **Migrate to a faster system.** Since new systems are often expensive, this solution may not be the most cost-effective.
- ◆ **Add more processors.** Although this method can often meet the needs of end users more efficiently, it can be programming intensive and may increase system administration requirements.
- ◆ **Provide extra capacity through loosely coupled systems on a network.** Such systems are too unstructured for many applications. Some applications can use a client/server model to off-load some processing to distributed client processors. However,

since the client/server model does not fit all applications, this approach may involve extensive programming changes.

- ◆ **Add tightly coupled special-purpose processors to off-load the main processor.** This is often the most desirable approach because it has the least impact on programming environment, current applications, and migration. One example of this approach is the use of communication or storage I/O processors on mainframes. The RISC System/6000 has a variety of intelligent Micro Channel adapters that can perform some level of off-loading. For example, the graphics adapters off-load much of the vector or complex image processing from the graphics subsystem, and the 128-port asynchronous subsystem adapter can off-load all the tty character processing from the AIX tty subsystem. Similarly, the Network Terminal Accelerator off-loads the networks and terminal protocol to the firmware.

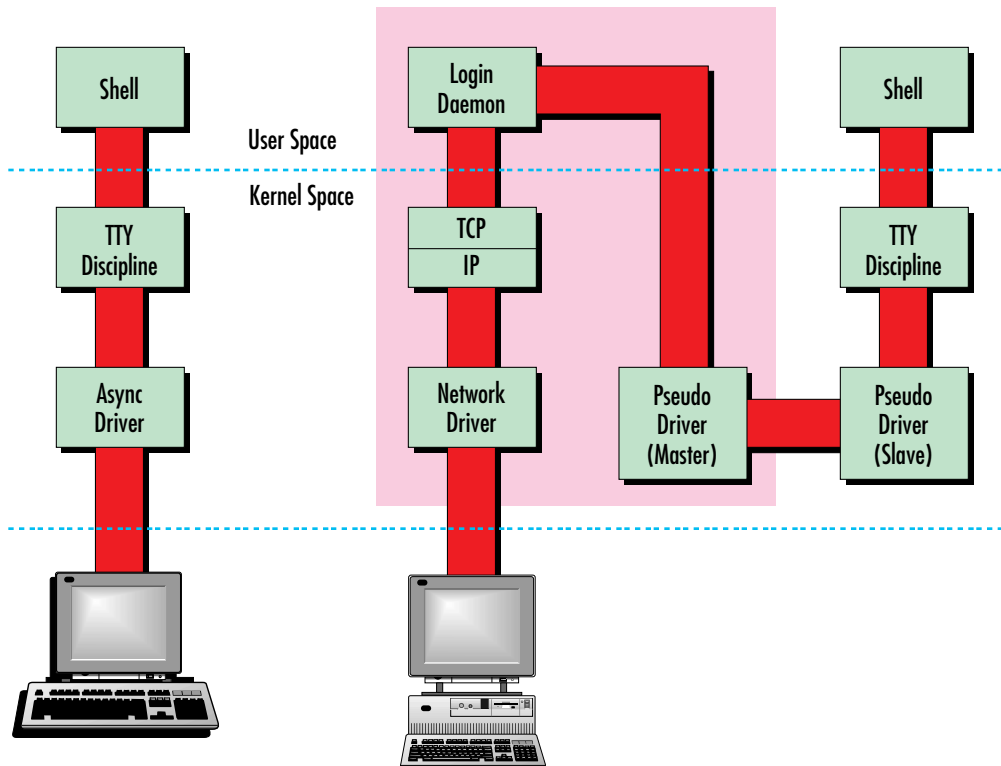
## Network Terminal Acceleration

*Network terminal* is the remote terminal access between two systems in a LAN environment. Traditional network protocols—System Network Architecture (SNA), TCP/IP, and Open Systems Interconnection (OSI)—provide terminal emulation for peer-to-peer host access to data and applications.

In the AIX TCP/IP environment, the network terminal access commands are `telnet` and `rlogin`. The overhead for TCP/IP protocols and `telnet/rlogin` processing can decrease performance of a business application. Figure 1 shows the data path differences between a locally attached ASCII terminal and a network-attached device. The overall path length for a direct attach-

## Locally Attached Terminal

## Network-Attached Terminal



**Figure 1. Local versus network terminal data path**

ment is about one-third that of the network-attached terminal.

Network-attached devices using TCP/IP over Token-Ring, Ethernet, or Fiber-optic Data Distribution Interface (FDDI) networks can provide higher bandwidths than a directly attached async terminal. However, network attachment causes a heavy protocol processing load for the host, including the following:

- ◆ Kernel-level processing
  - Network driver
  - TCP/IP protocol
  - The tty protocol
  - Pseudo-drivers
- ◆ User-level processing
  - Login daemon execution

This heavy load can be reduced by shortening the code path and off-loading the network protocol and daemon processing to firmware (the shaded area in Figure 1). The shell environment in


traditional AIX uses pseudo-tty (pty) as I/O simulation for remote access. In the accelerated environment, a new type of reengineered pseudo-device, high-performance tty (hty), can provide equivalent functions with special firmware assistance.

Because the Network Terminal Accelerator adapter and software can provide this capability and efficiency, they are ideal for a host with many users.

### Highlights of the Terminal Accelerator

The Network Terminal Accelerator adapter can be attached to any standard Ethernet with an appropriate transceiver. Since the adapter supports both IEEE 802.3 and Ethernet Version II packet formats, it can replace the existing Ethernet adapter. Figure 2 shows the network topology with this adapter.

The software support includes a device driver, configuration methods, and diagnostics. The Network Terminal Accelerator requires AIX 3.2.5. The adapter can support two IP addresses: one



**Figure 2. Large multi-user environment topology with terminal acceleration**

for the off-loaded TCP/IP protocol layer with login support, and the other for existing AIX TCP/IP protocol running in AIX kernel space. The system administrator must be aware of both protocol stacks with their functional placements, capabilities, and limitations.

In a typical terminal server host environment, the off-loaded IP address will handle all the terminal users, while the other IP address and resources are intended for other uses, such as NFS, network routing, and so on. Figure 3 shows functional placements and packaging of both TCP/IP protocol layers. The adapter can be configured using System Management Interface Tool (SMIT) to define the IP address for the Ethernet connection, and also to define the number of hty devices available at system boot. Each hty device maps onto a terminal user access shell and becomes available for login.

### The telnet/rlogin Sessions

The `telnet` or `rlogin` sessions can be initiated from any TCP/IP host on the network: a PC, another RISC System/6000, or a terminal server (such as the 7318 S20 described on page 58). The adapter has two models that support 256 or 2048 maximum sessions. Each hty device can be configured as a “call out” or “call in” port. If configured for call out, a `telnet` or `rlogin` session is

automatically established to a destination IP host during the open system call. In a distributed workgroup environment, this support can connect to a modem pool or shared printers that are directly connected to the network.

Ethernet packets for protocols other than `telnet` or `rlogin` will be passed on to the protocol stacks in the host for processing. The address assignment is set up using SMIT. This enables the system to support all existing protocols, yet the adapter can still efficiently handle the LAN terminal protocol processing for the users attached to the network.

The adapter also supports the Simple Network Management Protocol (SNMP) agent function and can be centrally managed by a manager product such as NetView/6000.

### Conclusion

With a dedicated microprocessor, the adapter can efficiently off-load the LAN protocol stacks for TCP/IP. Each CPU interrupt from the adapter to AIX can contain multiplexed data from many sessions. The adapter, which supports the 32-bit busmaster feature, can move large volumes of data efficiently. The protocol off-loading also reduces the number of context switches required to process incoming data. This solution, transparent to both users and application software, can




Figure 3. Structural diagram for both TCP/IP protocol stacks

increase telnet and rlogin efficiency by 300%, freeing the CPU for more application processing.

an advisory engineer in the Async Communication Development area. He has a BS in Electrical Engineering from Auburn University in Alabama.



**Eddie Ho**, IBM Corporation, 11400 Burnet Road, Austin, TX 78758. Internet: eddieho@chang.austin.ibm.com. Mr. Ho is a senior programmer in AIX Communications. He has a BS in Computer Science from the University of Wisconsin and an MS in Computer Science from North Dakota State University.

**Kent Malave**, IBM Corporation, 11400 Burnet Road, Austin, TX 78758. Internet: kent@bach.austin.ibm.com. Mr. Malave is a staff programmer in the TCP/IP Communications area. He has a BS in Electrical Engineering and Computer Science from the University of Oklahoma.

**Jim Gallagher**, IBM Corporation, 11400 Burnet Road, Austin, TX 78758. Internet: jrg@bach.austin.ibm.com. Mr. Gallagher is

**Dave Phipps**, IBM Corporation, 11400 Burnet Road, Austin, TX 78758. Internet: dlphipp@bach.austin.ibm.com. Mr. Phipps is a staff programmer in the AIX Communications area. He has a BS in Applied Mathematics, Engineering, and Physics and an MS in Computer Science from the University of Wisconsin.