PPP White Paper

Morning Star Technologies

January 17, 1992

Executive Summary

Morning Star PPP extends the benefits of your LAN across dial-up modem connections. Flexible links can be established quickly and inexpensively, providing immediate productivity increases. Workstation users can access central computing facilities from remote locations, widely separated offices can share resources on demand, and small institutions and companies can have full access to the wider Internet, all without installing costly leased lines and dedicated router hardware. Morning Star PPP minimizes telephone charges by managing the connection in response to user needs. All TCP/IP applications and services operate transparently, just as on the LAN, so no user retraining is necessary to take immediate advantage of the wide-area network link. Morning Star PPP can also replace router hardware in applications requiring high throughput on dedicated lines.

1 The Protocols Compared

1.1 TCP/IP Spans Underlying Technologies

The TCP/IP networking protocol suite is designed to provide the same services to the user across any access media. Applications such as electronic mail, remote login, and file transfer operate identically in every TCP/IP network, regardless of the underlying transport technology. Some popular media for local area networks are Ethernet, FDDI, and Token Ring. Wide-area networks have traditionally required the installation of expensive, dedicated telephone lines and router hardware.

1.2 IP on Dial-Up Connections

Now, it is now possible to connect any computer as a full TCP/IP internetworking peer, using only a voice-grade telephone line and a commodity modem. Two popular methods have emerged in the internetworking community: SLIP and PPP.

SLIP is a self-described "nonstandard", all of whose functions have been subsumed by PPP. PPP is the product of the Internet Engineering Task Force's Point-to-Point Working Group, and has been blessed by the IETF and Internet Activities Board as a standards-track protocol. SLIP was helpful for a time, and PPP's designers learned a lot of lessons from it, but its time has passed. But, since it has attained a degree of popularity, it deserves examination.

1.2.1 The Serial Line Internet Protocol

SLIP, described in RFC-1055, is an extremely simple framing scheme for putting IP packets on a serial line. Figure 1 shows a typical packet.

SLIP's main advantage is simplicity and ease of implementation, but that simplicity leaves out many functions that are useful in a link-level protocol. SLIP contains no provision for error detection, assuming that the upper level services will worry about those sorts of problems. It doesn't provide support for link management tasks, nor for any protocol family besides IP. It contains no mechanism for negotiating values for any of the many things that can be different between hosts, such as network-level addresses or packet size or various sorts of protocol or data compression. It contains no link-level authentication or loop detection mechanisms. These issues must be agreed upon in advance by the end-point systems, and often the decisions are embedded in code that isn't accessible to the user. For example, SLIP implementations that incorporate RFC-1144 "VJ" TCP header compression typically must be told in advance not to do VJ when talking to SLIP implementations that don't understand what a header-compressed TCP/IP packet means.

Figure 1: Typical SLIP Packet Format

SLIP's simplicity of implementation creates difficulties for its users, and for those who must manage and maintain network links. A SLIP link is either open and successfully carrying IP packets, or it's closed. If a SLIP link is malfunctioning, then the reasons are potentially many and ill-defined. Did the administrators of each end fail to agree on frame size? VJ header compression? IP addresses? Considerable expertise and insight is helpful in debugging flaky connections. It helps to have a guru in the house.

1.2.2 The Point-to-Point Protocol

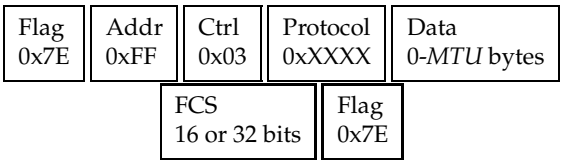

Figure 2: Full PPP Packet Format

By contrast, PPP, defined in RFC-1171 and more recent drafts, is a much more robust, full-featured protocol. Although the full PPP frame specification in Figure 2 contains some per-frame overhead that's unnecessary for many links, PPP peers can agree upon various compression techniques, so the typical packets actually passed on a dial-up link will look more like those in Figure 3.

PPP specifies a link control protocol, used for negotiating things like maximum frame size and HDLC-specific compression techniques. It also allows for control characters to be "escaped" so that they can pass through hazardous datacomm

Protocol 0xXX	Data 0-*MTU* bytes	FCS 16 bits	Flag 0x7E

Figure 3: Typical PPP Packet Format

equipment that might eat XON or XOFF characters, or something else even more unexpected. The link control protocol provides for low-level authentication, link loopback detection, and link quality monitoring that can be run in the midst of other traffic on the link.

For IP, the network control protocol allows the ends to negotiate their IP addresses, and whether they'll use TCP header compression. Dynamic address negotiation and assignment is a major management issue for hubs that provide service to many sporadically-attached remote hosts. PPP doesn't define a single policy, but it provides flexible mechanisms to implement a wide variety of address assignment policies.

As a PPP link comes up, all the variables that may make it unsuccessful are subject to negotiation. A PPP implementation can easily report to its administrator "the link didn't come up because the peers couldn't agree upon their frame sizes," or whatever the cause might have been. The elaborate PPP finite state automaton formalizes many of these link management tasks, distilling the wisdom of a link-level guru into an algorithmic process. This eases the human administrative burden of establishing and running the link, and empowers less-skilled users to solve the most common problems.

1.3 Selecting SLIP or PPP

SLIP has become popular for some dial-up IP applications, and implementations are available for many different types of hosts. Though SLIP would seem to be a de-facto standard, there are sound reasons for migrating existing installations to PPP, and for selecting PPP for new applications.

The SLIP RFC itself describes SLIP as a non-standard, and lists several deficiencies. The PPP RFC is the product of the Internet Engineering Task Force's Point-to-Point Protocol Working Group, and its introduction describes how it addresses each of those deficiencies. PPP is under active development and is experiencing directed evolution at the hands of a group of talented engineers representing a broad cross-section of the networking industry; SLIP has stagnated for several years, having exhausted its potential. PPP's market share is expanding among router and terminal server manufacturers, with new implementations announced frequently. As new vendors enter the PPP marketplace, Morning Star Technologies conducts extensive interoperability tests to ensure that users receive seamless, standards-based connectivity.

New wide-area networking installations should install PPP to carry them into the future. However, to provide backward compatibility with systems that are still unable to run PPP, Morning Star's PPP implementation also provides a SLIP option.

2 Applications of PPP

2.1 Bandwidth Considerations

While by definition all IP-based applications may be transparently used over links traversing any combination of media, some applications are better suited than others to media with the relatively low bandwidth and moderate latency characteristics of asynchronous dial-up connections.

PPP over current-technology dial-up modems[1] provides satisfactory performance for most typical TCP/IP and UDP/IP wide-area applications. Especially with TCP header compression, users will experience useful file transfer throughput[2] and reasonable interactive response in applications such as *telnet*, *rlogin*, and the like. Plenty of bandwidth is available for network infrastructure applications such as SMTP electronic mail, NNTP network news, NTP clock synchronization, domain name service, etc.

The speed limitation of the dial-up medium becomes more apparent when larger amounts of data are transferred, particularly for interactive applications such as GUI-intensive X11 clients. Large bitmap editors and CAD packages will operate correctly, though users generally prefer the fastest modems available to preserve their interactive responsiveness.

Even applications that are extreme users of communications bandwidth, such as Sun's Network File System (NFS), may be used with care over dial-up links. With appropriate configurations, network managers can extend all the services of their LAN to remote locations.

2.2 Typical Configurations

2.2.1 Two Workstations

Figure 4: Two Workstations

Two distant workstations might arrange a private dial-up link with Morning Star PPP, much as they would use UUCP to exchange files, mail, and news. The PPP link would use the same modems as UUCP, and the administrative effort involved in the initial set-up is similar to that for UUCP. They would benefit by interactive access, simultaneous bidirectional activity, and the potential for using the full range of TCP/IP applications. They would even gain in bit efficiency, since although TCP/IP isn't quite as concerned with line bandwidth frugality as Zmodem or Kermit, it's far better than UUCP's "g" protocol.

Some minicomputer and mainframe manufacturers provide a modem or even a small computer for front-end diagnostic access to their systems in the field. That remote front-end can be linked via Morning Star PPP into the diagnostic tools on the manufacturer's home office UNIX network for quicker problem resolution.

2.2.2 Workstation to LAN

Many forward-thinking companies using workstations in the office see the benefits of providing their engineers and support staff with similar workstations for use at their homes, particularly since the home machines can now be easily connected with the office network. It's much less expensive to install a voice telephone line into a home than a dedicated leased data line, and there's less of a wait for installation.

[1] CCITT V.32 (9600) or V.32bis (14400) carrier, with V.42 error correction and V.42bis data compression, and the host serial interface latched at 38400

[2] V.32/V.42bis yields 1.7-2.8 Kbytes/sec, and V.32bis/V.42bis yields 2.4-3.5 Kbytes/sec for a typical large binary data file. Over Telebit T1600 V.32/V.42bis modems latched at 38400, the Morning Star PPP daemon can receive 25224 bits/sec and simultaneously send 26510 bits/sec.

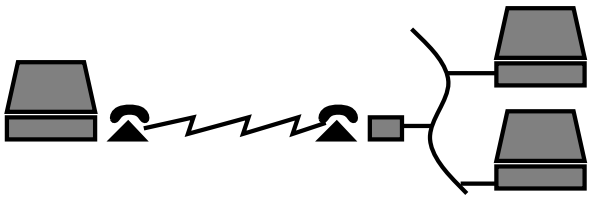

Figure 5: Workstation To LAN

Students who purchase low-cost workstations upon matriculation should be able to connect them with the campus network, thence to the rest of the Internet. This is becoming common in university-provided housing where it is practical to extend fiber or twisted pair to each dormitory room, but off-campus residents can now receive the same services.

With the advent of "laptop" or at least "portable" workstations, migratory professionals like field support and marketing personnel can have full-function access to the home office databases, perhaps even over cellular telephone calls. As commercial IP network points of presence (POPs) become more ubiquitous, the traveling workstation user may not even need to dial directly into the home office network, but will use local Internet access points with cheaper telephone charges.

2.2.3 LAN to LAN

Figure 6: LAN to LAN

Widely separated offices of the same organization can use Morning Star PPP to form an enterprise-wide network at much lower cost than if dedicated leased lines and routers were required. As companies enter into cooperative product development agreements, they can link their private LANs for technology transfer. Technical customer support organizations can provide their services, and new software and database revisions are easily delivered over on-demand network links. While moving equipment to new facilities, Morning Star PPP can provide a quick, cheap, flexible lifeline of connectivity before permanent networking facilities are fully installed.

2.2.4 LAN to Internet

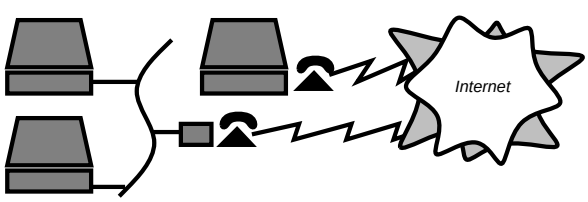

Figure 7: LAN to Internet

Small colleges and universities can "get their feet wet" in networking by connecting to their regional NSFnet tributary via an inexpensive dial-up link, planning to commit the resources later to install a dedicated leased line (the same Morning Star PPP software will serve them then, too). The low cost of entry should be an effective marketing tool for the commercial IP connectivity vendors. Companies can use a workstation with a modem to provide an initial Internet gateway for their entire network, through which they have full access to all the Internet's facilities. A small company will outgrow a dial-up link in about a year, but during that time they'll see the benefits and their employees will become dependent upon it, and they'll never be able to unplug again.

2.3 Case Studies

2.3.1 Medical Information

A company sells a medical information system that uses workstations for data collection at the patient bedside and database servers in the hospital MIS center. When a patient has a problem and an intern calls a doctor at home at 3:30am, the doctor needs access to the most current data about that patient. Using Morning Star PPP, the company is extending their information services, formerly available only on the hospital LAN, into the doctor's home. Each doctor has a low-cost color workstation running Open Windows to display the same picture the intern sees at the bedside.

The database query engine and graphics generator currently run on large Sun systems at the hospital, and the PPP connection carries PostScript images to the doctor's home. Future revisions of the software may support remote queries and generate the images on the workstation, so that the PPP connection would carry the database transactions rather than the resultant graphics.

2.3.2 Résumé Processing

A company provides database software for automated résumé handling. Clerks optically scan incoming paper résumés for entry into a database, with keywords, educational and work history, etc. When a position becomes open, the requisition is pattern-matched against the database to screen for qualified applicants, which are then forwarded to personnel officers for consideration.

Some users of this system are geographically diverse, with multiple employment and hiring locations. The widespread hiring-database servers are linked by dial-up modems. The system was initially implemented with a custom serial protocol designed specifically for queries by this particular database. When the query protocol was converted to run over TCP streams over IP networks, the dial-up links were converted to Morning Star PPP and became useful for customer support, remote printing, electronic mail, and everything else alongside the database queries.

Rather than being restricted to one application, the modem links carrying IP traffic have become part of the corporate infrastructure. Morning Star PPP allows the system's developers to separate the networking issues from their central product focus, and yields increased flexibility for their customers.

2.3.3 Air Traffic Control

A group at NASA is developing an Air Traffic Control automation system for the FAA, consisting of a truly distributed set of tools running on high speed color graphics UNIX workstations, and communicating via TCP sockets. Morning Star PPP enabled developers spread across the continent to collaborate and develop the application. Many of those developers are contractors without Internet access, so Morning Star PPP was

the only practical solution for providing connectivity to their offices.

Morning Star PPP was deployed in three gradual phases. First, it simply replaced UUCP for electronic mail and file transfer. Second, the various parts of the distributed application were linked together so that widely-separated engineers could participate in air traffic simulations and cooperative debugging sessions. This also provided a realistic approximation of the operational environment the system will encounter when deployed. Third, certain ATC monitoring services, currently provided only over a dedicated leased line from one location to another, will soon be made available to any authorized machine via Morning Star PPP.

3 Morning Star PPP

3.1 Other Implementations

Other implementations of PPP for UNIX systems fall into two general categories: those that give the kernel's networking device drivers too much specialized knowledge about the PPP protocol, and those that put all the protocol handling into STREAMS modules. Each approach presents its own problems. Anything compiled into the kernel is inflexible, nonportable, and difficult to debug. Not all UNIX variants provide both BSD-style networking *and* STREAMS for I/O, and those that do often differ in their API.

3.2 The Tunnel Driver

The "IP tunnel" device driver as implemented by Morning Star Technologies serves as the key to our PPP implementation's flexibility. It separates transport protocol handling from the kernel's knowledge of IP.

A tunnel driver makes network traffic available on a file, such as `/dev/tun0`. There, packets may be read by a user-space daemon process, which has the responsibility for delivering them via the desired encapsulation scheme on whatever transport mechanism is available. One daemon might deliver IP traffic via X.25, another via PPP, another via SLIP, another via SNA, etc. The tunnel device driver knows nothing of the network transport issues; it is only responsible for presenting an interface to the UNIX kernel's IP.

3.3 Advantages of This Strategy

When compared with other methods, Morning Star's daemon implementation is superior both in the functions it provides and the ease with which it may be ported to new platforms. The PPP protocol is entirely handled in a process in user space, and the tunnel driver provides a convenient and narrow interface to the kernel's IP layer. The daemon is not dependent upon new line disciplines or kernel architectures or STREAMS modules, which can hinder portability to different UNIX systems.

The daemon handles all aspects of the PPP protocol, and it manages the link state as well. It contains timers to support hanging up the telephone when the line goes idle, and it can dial the modem and reestablish connectivity when a packet needs to cross the link. Its state can be observed with `ps` because it manipulates its argument vector (`argv`). It can log its progress in the filesystem. It can respond sensibly to software signals, just like any other process. It can provide capabilities that would be difficult or dangerous to put into a kernel, or that would cause kernel processing to block inconveniently.

3.4 Features of Morning Star PPP

- Standards Support
  - The Internet standard Point-to-Point Protocol (PPP), as defined in RFCs 1171 and 1172 and updated in more recent drafts, providing
    - Link-level error detection
    - IP address negotiation and assignment
    - PPP Address/Control field and Protocol field compression
    - Link-level authentication by PAP (RFC 1172) and CHAP (October 1991 Draft)
    - Link status monitoring by LQM (August 1991 LCP Draft)
  - All PPP features except 32-bit FCS negotiation
  - The popular Serial Line Internet Protocol (SLIP), as described in RFC 1055
- Optimal Performance
  - 'VJ' TCP packet header compression as described in RFC 1144
  - TCP interactive 'fast queue' management
  - Asynchronous speeds as high as the underlying UNIX provides, usually 38400
  - Synchronous speeds to T1 (1.544 Mb/sec) or CEPT (2.048 Mb/sec) using our SnapLink SCSI-attached serial interface
- Powerful security features including
  - Packet filtering by protocol, source, destination, traffic direction
  - Multiple connection-time authentication methods
  - Highly configurable packet tracing
- Ease of Management
  - Easy and extremely flexible configuration
  - On-demand link establishment and idle line disconnection
  - Flexible "chat script" for connection establishment
  - Flexible daemon implementation
  - Operates as either the remote (calling) or hub (answering) site
  - Uses most asynchronous modems or dedicated lines, using the workstation's native serial ports
  - Shares modems with other activities, such as UUCP or interactive users
  - Monitors link status, reliability, and performance
  - Multiple line failover for redundancy and high availability

4 For More Information

Morning Star Technologies
1760 Zollinger Rd
Columbus OH USA 43221
+1 614 451 1883 (voice)
+1 800 558 7827 (voice)
+1 614 459 5054 (FAX)
marketing@morningstar.com (e-mail)