REMOTE ACCESS MARKET AND PRODUCT EVOLUTION

by Shuang Deng, Ph.D

Aptis™

Aptis Communications, Inc.
White Paper

This white paper provides a brief history on the remote access market and product evolution, and implications for future efforts related to the development of remote access applications. This evolutionary process is presented through a description of the three stages that remote access has experienced since its earliest origins, along with the new generation applications and architecture which will characterize a fourth stage of development. Technological innovations are presented in conjunction with the factors that have influenced the remote access market at various stages of its evolution.

An ever-increasing demand for scalability, manageability and performance has been a recurring theme over the course of this evolutionary process. As one of the most dynamic segments of the high-technology market in recent years, remote access has been characterized by constant, rapid and pervasive change, with dramatic market shifts often occurring over a period of several months. While there are some signs that the marketplace may eventually begin to stabilize as a result of consolidation and other factors, remote access promises to remain a dynamic, challenging, and rapidly evolving market segment for the foreseeable future.

First-Generation: Modem Bank and Terminal Server

The first remote access market evolved in response to a need for host resources such as a mainframe host to access remote resources — such as a shared printer. In this model, remote access was viewed as an extension of local host-to-host or host-to-peripheral connections. Sessions were conducted on an individual basis, with a single remote terminal accessing a single host. These first-generation applications were characterized by extremely limited functionality, resulting in a capability which functioned more as an extension of a discrete task, rather than as a full-featured, mission-critical solution.

First-generation remote access capabilities consisted of a combination of a terminal server and a modem pool. Prior to the development of more sophisticated remote access technology, terminal servers were used in conjunction with serial ports for host-to-host or host-to-peripheral connections. Programmers discovered that inserting a pair of modems in the serial link of this configuration allowed remote access to the host. As the demand for this capability grew, stand-alone modems were packaged together to create a modem bank for direct connection to terminal servers. This first-generation architecture, which is characterized by the separation of the modem function from the host communication device, is illustrated in Figure 1.

First-generation solutions were normally deployed in institutions, corporations and universities with branch and home office users who needed to access the organization's computing resources from remote locations. Typical applications included terminal emulation, telnet and uucp, which utilized data rates ranging from 300 to 2,400 bps.

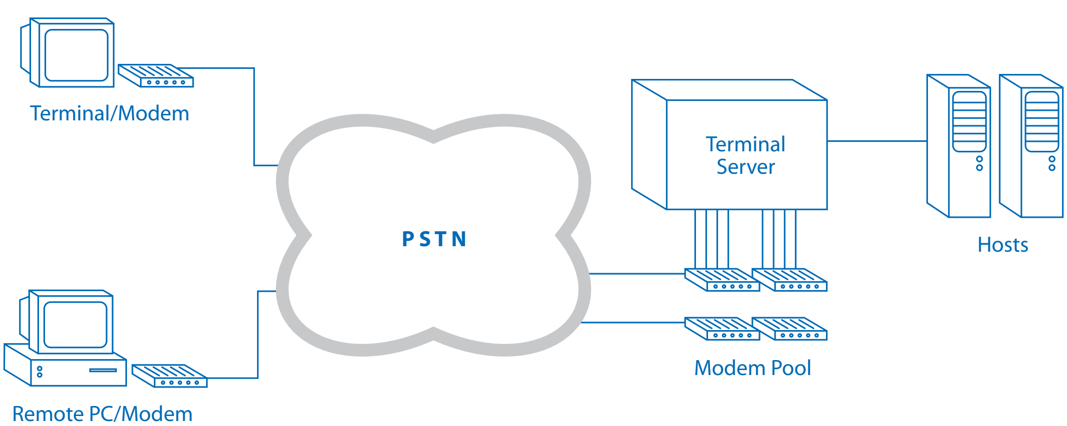

Figure 1: First-generation remote access architecture

Second-Generation: LAN Access With Integrated Modem

Beginning in the late 1980s, the mainframe/terminal paradigm of enterprise computing began to be eclipsed by PC- and LAN-based client/server solutions. As a result, corporate remote access requirements began to shift from the one-to-one connection model of terminal-based remote access to a model based on simultaneous connections conducted in multiple locations. A new paradigm evolved in which remote access requirements focused on providing connectivity between a corporate LAN and a remote PC or branch-office LAN. This produced a proportional increase in the number of nodes which a particular solution had to support, as well as a need for an expansion of functionality to accommodate new protocols, such as LAN bridging and TCP/IP routing.

The expansion of remote access capabilities also created the demand for more robust solutions. The modem/terminal server configuration typical of first-generation remote access capabilities resulted in problems associated with adding ports, troubleshooting, and managing and operating disjointed, incompatible systems. In addition, the external serial links between modems and terminal servers restricted data transfer performance. These were the first manifestations of the difficulties incurred in attempting to formulate "end-to-end" remote access solutions. In an effort to address the problems associated with modem/terminal server functionality, a second generation of remote access products evolved which integrated the modem, terminal server and LAN access function in a single piece of hardware. Over the course of this second-generation phase, the focus of remote access capabilities shifted from analog modems to digital modems and multi-card systems.

Analog Modems

Early second-generation products typically featured one analog modem port and one Ethernet port which supported NetWare and AppleTalk protocols (the predominant networking protocols of that time). Incorporating multiple modem ports which supported the TCP/IP protocol suite improved the functionality of these products considerably, and helped position them as market leaders. However, the scalability, manageability and performance issues associated with analog-based products created a demand for digitally-based modem capabilities.

Digital Modems

As a greater portion of telephony infrastructure began to migrate from an analog to a digital architecture, corporations began to receive their telephone services on digital circuits that could carry multiple phone calls simultaneously. Initially deployed for PBX connections, these applications attracted the interest of remote access users in the early 90s, due to the significant cost savings associated with digital interfaces.

This shift was spurred by the intrinsic efficiency of digital solutions, such as T1 lines, which were more cost-effective than individual analog lines of the same capacity. (A single T1 port can replace 24 analog ports, resulting in greater operational simplicity and reduced costs.) As a result, many companies were able to realize significant decreases in both recurring telephone line charges and operating costs through the deployment of digital technology. In addition, as port density began to increase, analog modem ports were soon viewed as being too difficult to install, manage and troubleshoot.

Beginning late in 1992, remote access servers with digital interfaces began to appear on the market. These new solutions combined the functionality of modems, terminal servers and (to a limited degree) routers within a single piece of hardware. Initially, the modems were analog, and the concentrators performed a digital-to-analog conversion to interface the digital circuitry with the analog modems. This conversion, which was transparent to the end-user, enabled customers to realize the cost benefits associated with this higher-density configuration. Modem suppliers quickly caught up with this trend toward digitally-based solutions and were soon delivering more cost-effective, higher-performance digital modems, which were quickly embraced by the marketplace.

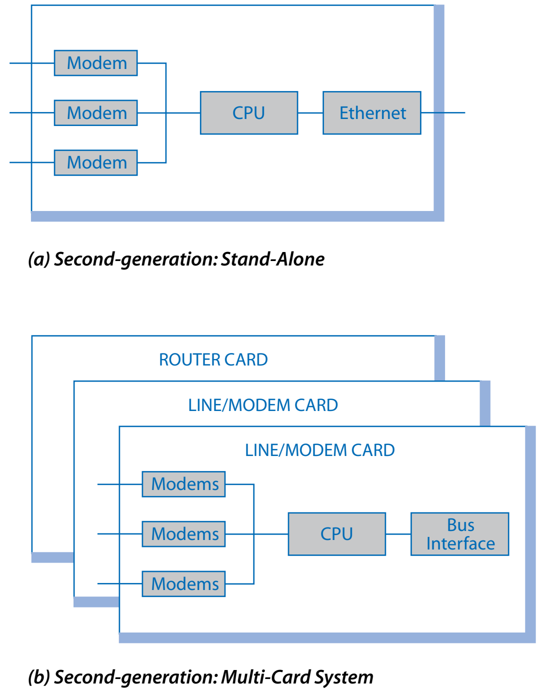

Figure 2: Second-generation remote access architecture

Multi-Card Systems

Remote access vendors soon began to develop multi-card systems for these stand-alone servers, which replaced the LAN interface with a more powerful bus interface. While these systems allowed for the deployment of a greater number of modems, modem density was still limited by the range of interfaces that could be handled by a single card.

The various second-generation products described above are characterized by a centralized processing architecture with integrated modems and networking, as illustrated in Figure 2. While these networked modem products enabled remote access systems to deliver a much higher level of scalability to multiple users, customers continued to require more ports and density than was provided by these second-generation capabilities. This ongoing need helped propel remote access into its third stage.

Third-Generation: Distributed Processing

The Internet explosion in the mid-1990s created an accompanying demand for access on a massive scale, which overwhelmed the capacity of second-generation remote access products. Because the centralized processor featured in second-generation architecture necessitated low port density, these products were unable to accommodate large-scale networks with hundreds of ports. As a result, "carrier-class" switches began to appear during the mid-1990s. These switches offered a higher density capacity (approximately 672 calls, equivalent to a DS-3 of bandwidth), which represented a considerable increase over previous solutions offering 288 to 400 modems per shelf.

New Connectivity Requirements

ISPs, whose end-users required dial-up Internet access capabilities, constituted the primary market for these third-generation solutions. A new breed of remote access provider, ISPs had connectivity requirements which were markedly different from those of the traditional enterprise market. Whereas a corporate network might require only a few dozen modems, a typical ISP's POP has anywhere from several hundred to several thousand modem ports and uses a number of channelized T1 or PRI access lines.

Conversely, while ISPs had greater density requirements, their functionality requirements were less complex than those of corporate customers. Instead of having to support an entire suite of LAN protocols to accommodate various enterprise needs, ISPs required support for TCP/IP only. Because the Internet constitutes the sole application driver for an ISP, PPP and IP were the only protocols required for ISP-based remote access servers. As a result, most POPs were able to utilize a relatively straightforward configuration, in which a group of servers were connected to a separate router.

The access concentrators which emerged as third-generation solutions were used to terminate PPP sessions and route IP traffic to local Ethernet ports. Multiple concentrators were aggregated, through shared or switched Ethernet, to a router, which would then send traffic to the Internet NAP.

Third-Generation System Design

A flexible chassis design and modular card system were the primary architectural improvement offered by third-generation remote access products. This modular design allowed for a portion of the processing to be distributed to each individual card, thereby increasing overall system capacity. Figure 3 illustrates a typical third-generation architecture, in which a dedicated T1/PRI card terminates access lines and sends individual calls at DS-0 streams on the TDM bus to modem cards. The modem card then terminates the modem calls and PPP sessions and sends the logon information to the central processor card, which transmits the query to an external RADIUS server for authentication and accounting.

Once authenticated, the modem card queries the central processor for route updates and then transmits PPP data from the modem card to the central processor card. Since a cell bus is used, each PPP packet is segmented into cells for bus transfer and then reassembled at the central processor card. The central processor then re-encapsulates the payload into Ethernet or WAN packets for uplink transmission.

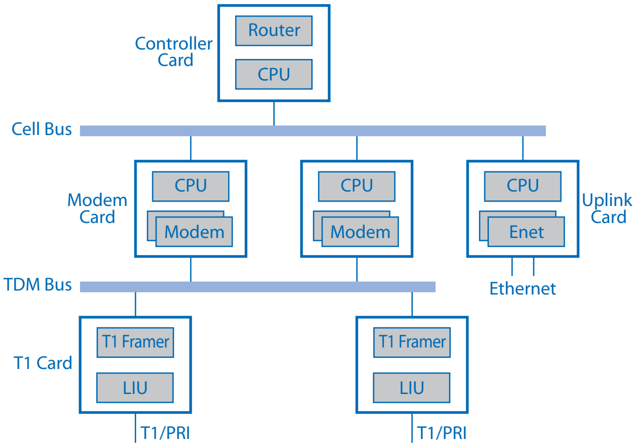

Figure 3: Third-generation remote access architecture

Third-Generation Limitations

Third-generation access concentrators represent a marked improvement over previous generations in both density and operating cost. However, they also incorporate a number of fundamental design flaws, including:

- Foreign bus overhead
- Scalability difficulties with port configurations numbering several hundred or more
- Unreliable design
- Inability to support multiple services
- Lack of suitability for emerging VPN outsourcing and wholesale services

The use of cell bus cards also imposes severe performance limitations on third-generation solutions. For example, while these concentrators are incapable of handling more than several hundred calls per shelf, the burgeoning demand for Internet access is much greater, since these configurations must often support thousands of calls per shelf. As the number of ports increase on an access concentrator, both the hardware reliability and availability of services become increasingly important. While outages on a limited number of connections may not have severe consequences, a loss of 672 ports or more at a time is a matter of serious concern for service providers. In addition, as business customers begin to migrate their mission-critical applications to the PSTN, the need for reliable, high-performance, high-density access capabilities increase proportionally.

The Next Generation: Scalable, High-Density, High-Performance Solutions

As with previous generations of remote access products, the limitations of third-generation solutions have created demand for a new generation of access switches. As network service providers, such as telcos and large ISPs position themselves to deliver a greater breadth of services, their capability requirements have increased proportionally. These needs include:

- Higher Density: Functionality capable of handling more than 2,000 calls per shelf, or 5,000-10,000 calls per rack
- Scalability: Flexible provisioning of multiple ports and
- services on a single platform
- Performance and QoS: Sustainable throughput, regardless
- of traffic load, with multiple QoS levels
  Reliability: Hot-swappable cards, redundancy protection, and separate wiring to avoid re-cabling
- Multiple Services: Capacity to accommodate both VPN outsourcing and ISP wholesaling, by supporting multiple services (including both narrow-band dial-up and broadband xDSL) on the same platform
  Operations System Support: Integration with existing OSS platforms to facilitate customer network management

In order to satisfy industry requirements for the next generation of remote access technology, new techniques must be adopted which are based on the design of telecommunications carrier equipment. In particular, developers must find a way to overcome heat dissipation by designing a vertical card orientation (rather than the horizontal configuration utilized in existing products) for optimal air movement. In addition, providers must also adopt distributed power supplies which not only improve reliability but also increase power efficiency and reduce heat generation.

To address the high-performance requirement, developers need to utilize integrated native buses which provide a packet bus for packet data, and a cell bus for cell data. This eliminates the performance limitations associated with the expensive, time-consuming, and unnecessary process of packet segmentation and re-assembly which results from exclusive cell bus usage. These innovations are needed to satisfy the expectations of end-users who purchase services from retail ISPs, or work for companies that have outsourced their remote access functions.

Outsourcing and Wholesaling

The outsourcing of corporate remote access and the wholesaling of ISP access constitute two of the key market drivers behind the evolution of next-generation access solutions. These services require a level of capability which exceeds the capacity of third-generation remote access concentrators, which are designed to handle DS-0 traffic, and cannot adequately support multiple types of access services. In addition, because third-generation products are designed for access by a single ISP, they are poorly suited to an outsourcing/wholesaling business model, in which multiple ISPs and enterprises are accessing a portion of the server's overall capacity. As a result, delivering multiple services on a third-generation concentrator would require a total redesign of the system's software.

Looking ahead, there are four new market conditions that have resulted in emerging requirements that cannot be satisfied by existing product architecture.

- Businesses are placing an increased demand on network solutions for greater reliability, performance and scalability
- The industry trend toward specialization and consolidation is resulting in the emergence of network service providers offering wholesale ISP services
- Corporations are looking to outsource their remote access functions and other components of their VPN infrastructure to these network service providers
- Telephone companies are seeking to deploy CO-based remote access to relieve congestion on the PSTN

Providers who position themselves to address these next-generation requirements for architecture and applications will be well positioned for success during 1998 and beyond.

Shuang Deng *is director of Product Management at Aptis Communications. He has been credited as an integral member of the GTE team that developed the original concept of ADSL for data services; and at Hewlett-Packard as a co-designer of the industry's first ATM broadband tester. Mr. Deng holds a Ph.D in Computer Science from the University of Alberta in Canada.*

11 Elizabeth Drive
Chelmsford, MA 01824-4111
(888) 278-5699
Fax (978) 250-8022
www.aptis.com

Aptis

Aptis Communications, Inc.