Remote & Mobile Computing With TCP/IP

Chris Bradley
Product Manager - Remote & Mobile Networking
WRQ
(206) 217-7500

Introduction & Objectives

Remote access of computer resources has been a reality for well over 30 years. It has evolved since its early days of teletype devices interacting with mainframe computers over telephone lines to become a common configuration of all aspects of computing. With advent of the personal computer revolution, the demand for remote connectivity has done nothing but accelerate, while the types of remote and mobile network solutions have expanded to cover all manner of requirements and users. The purpose of this paper is to provide an overview of remote and mobile computing today, with and emphasis on the networking of remote devices and how TCP/IP fits into such configurations

The paper will review the following sub-topics:

- Definitions of remote and mobile computing
- The motivators behind the growth of remote and mobile computing
- An overview of the components that make up a remote computing configuration
- A review of some general considerations and issues

Definitions

For the purpose of this paper, the following definitions apply:

Remote Computing: Access to computing resources from a *fixed* remote location.

This definition applies to a number of user and application types, such as work-at-home programs, virtual office, mobile professionals, road-warriors. It can also utilize different types of remote access configurations. The distinguishing factor in with remote computing is that the mobile computing device is stationary when remotely accessing computing resources, which usually means that the networking medium is provided by the public switched telephone network (*PSTN*), the same network used for the vast majority of phone calls. This network medium, for certain types of remote computing configurations, can also encompass digital services providing by the phone company called Integrated Services Digital Network (*ISDN*). For this reason, this paper will refer to *wireline* infrastructure, which includes both PSTN and ISDN networks, as the networking medium used for remote computing.

Mobile Computing: Access to computing resources from a mobile or moving remote location

This definition usually applies to user or application types such as field services or sales force automation. The user may be in motion while connected to remote computing services, or is

continually moving the location of the mobile computing device. The distinguishing factor in mobile computing is that the transient nature of the configuration requires networking mediums that do not require physical connections, such as cellular or radio networks. This collection of networking mediums will be referred to in this paper as *wireless*.

Motivators For Remote & Mobile Computing

There are both business forces, societal trends, and enabling technologies at work today that are driving the increase in remote computing/mobile computing. In the business environment, greater competition is forcing downward pressure on cost. Companies are downsizing to meet the demand and trying to do more work per worker, while eliminating expensive overhead like real-estate. Results include demand for increasing productivity of smaller work force through automation and technology, longer average work weeks, and alternative work environments such as the 'virtual office' where workers conduct business where ever the are.

Societal trends are partly as a result of business forces, and partly driven by concerns for non-business factors such as environmental concerns. As businesses expect more of the workers, workers are demanding greater flexibility in their work conditions in order to maintain some semblance of balance between work and personal life. This leads to solutions such as work-at-home programs, where workers conduct some or all of their job responsibilities where the live. Also, local and state governments are legislating corporations into reducing the number of cars on congested highways by implementing so-called 'telecommuter' programs.

Enabling technologies that are help meet the challenges of these business forces and societal trends fall into three primary technology areas. The first is the increase in computing power of the mobile platform. Most laptops today can easily match the performance and functionality of desktop systems at comparable pricing. The second evolving technology is the telecommunications services that are facilitating faster communications in more places. These include new digital services such as ISDN for wireline connectivity, as well as the explosion of cellular services in the wireless area. The third area of enabling technology are the hardware and software interfaces between the mobile computing platforms and the telecommuting services that are providing for higher performance and greater flexibility. These include things such as high speed analog modems and wireless data subscriber units.

Technology Components Overview

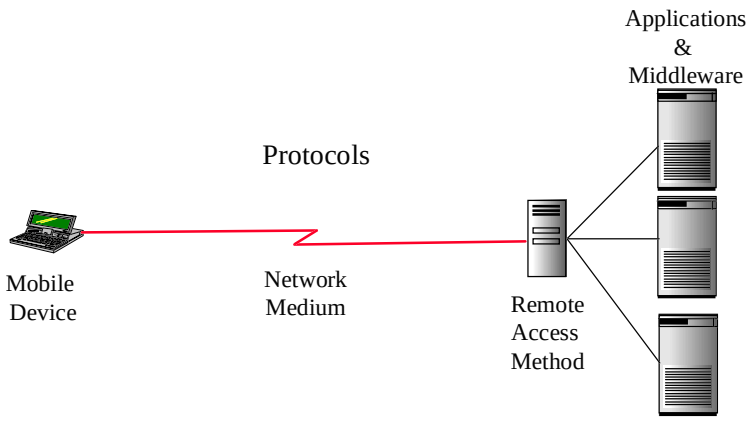

Remote & mobile computing can be accomplished with a number of different types of component technologies, but all usually have 4 common component types, as illustrated above: These are:

- Mobile Computing Devices
- Network Medium
- Remote Access Methods
- Applications/Middleware
- Protocols

The following sections will provide high level overviews of these component types, with descriptions of the sub-types found in these categories and some comparison of the different technologies.

Remote & Mobile Computing Platforms

Of the common components involved in remote/mobile computing, the computing platform has the highest degree of diversity and rate of change. There are all manner of computing platform types that can be used to remotely access systems and networks, and the number of types continues to accelerate bringing on more and more specialization. Some of the more common platforms are:

Desktop Computers: used in remote computing for telecommuting or work-at-home programs. The device remains fixed and uses PSTN or ISDN connections to temporarily connect to central hosts or networks.

Laptop Computers: Laptops can be found in both remote and mobile configurations. They are the largest and fastest growing of the mobile computing types. Recent acceleration in component technology has blurred the distinction between laptop and desktop systems, and with price/performance ratios continuing to improve, it is no surprise that an ever-increasing number of corporate decisions are being made to provide laptops as the single computing platform for all users.

Specialized Hand Helds: These devices are found in mobile computing configurations, typically in vertical industry implementations such as transportation, health care, and public safety. Recently, they have begun to move away from proprietary components towards industry standards such as Intel processors and DOS/Windows operating systems. The most defining aspect to their design is ruggedization, with the objective of allowing to the device to perform in adverse environments and the ability to withstand impacts sustained from being drop from heights. Recent hand-held designs usually employ stylus input devices and graphical applications. Communications is usually wireless, and these devices functions in both local (wireless LAN) and wide area (wireless WAN) configurations.

Personal Digital Assistants (PDAs): Initially oversold with inflated promises of high function and low cost, PDAs continue to struggle for acceptance, although adoption rates have been steadily improving. The come in a wide variety of designs and capabilities and are intended to provide productivity applications to individuals through horizontally oriented applications such as electronic mail, personal information managers, spread sheets, etc. Ironically, a key contributor to the improvement of PDAs fortunes has been their implementation into vertical application solutions.

One of the biggest promises, and disappointments, has been the PDAs wireless communications abilities. Although some have implemented paging capabilities, full function networking is just now becoming a reality, through the use of cellular modems.

Superphones: A new device that his on the immediate horizon combines cellular phone functionality with computing capabilities. Vendors have implemented limited software clients within a phone that allow interaction with specialized Internet services, including e-mail and web-browsing. The user interface is the LCD panel display and phone keypad, so interaction is somewhat limited. Vendors are designing their solutions to work with special text-based Internet applications. A second wave of these devices will bring greater graphical capabilities to a cell phone/PDA hybrid platform.

Communication Link - Wireline

The communication link is the medium that his used to transmit information across the wide area. It is owned and operated by a service provider, either public or private, and its usage is tariffed. All communications links used for remote and mobile computing are temporary in nature and are described as switched connections (as opposed to dedicated). Wireline communications links have fixed points of termination and usually involve wires. The two most common types of wireline communications links are PSTN and ISDN.

The Public Switched Telephone Network (PSTN) is the switched network used for telephone calls. It is an analog medium that provides relatively reliable but slow speed remote communications. Data communications over the PSTN has benefited in recent years from advancements in modem technology that have allowed throughput rates to grow to 28.8Kbps, with compression standards enabling effective throughput rates for some kinds of communications to exceed 56Kbps across a switched connection. PSTN communications links are point to point, and allow for only a single data stream across one connection.

Integrated Services Digital Network (ISDN) is an international standard for public switched networks based on digital signaling. The digital nature of ISDN provides for a number of advantages over

standard analog PSTN infrastructures, chief among these are enhanced throughput through channel aggregation and simultaneous data streams across a single connection.

ISDN can provide for multiple channels of communications across the same pair of wires. Through inverse multiplexing, it is possible for an ISDN connection to provide aggregate throughput of 128Kbps with very high reliability. ISDN can also allow for simultaneous voice and data communications.

ISDN achieved greater acceptance outside of the United States initially, but it has recently enjoyed an increase in adoption within the US. Close to 90% of the US is now ISDN capable, and the public phone companies who offer the service have been working hard to address initial issues of ordering complexity and interoperability. Costs have also become more attractive, both for service and equipment. ISDN makes a good solution for remote corporate access, especially for high frequency usage, or for applications with high bandwidth requirements like CAD.

In addition to PSTN and ISDN, new wireline communications link technologies are arriving from public service providers. Two standout technologies recently receiving a good deal of attention are cable modems and ADSL. Cable modems, which are data devices that hook to coax cable provided by cable TV services, promise to provide near LAN bandwidth (2-10Mbps) directly into the home. Proof of concept pilots have been conducted and initial reports are fairly positive, but many issues remain - including the issue of how quickly cable companies can convert their existing infrastructure, which only provides for one-way data communications, to allow for full-duplex transmissions.

Asymmetric Digital Subscriber Line (*ADSL*) is a switched technology that allows for bandwidths up to 5Mbps over existing phone lines. It is being championed by local telephone companies partly as a response to the cable TV industry's foray into data communications through cable modems. The primary issues with ADSL are its asynchronous nature, allowing for transmissions of 1.5Mbps to 9 Mbps downstream but only 16Kbps to 640Kbps upstream, and interoperability issues that have not yet been resolved between various service providers and equipment manufacturers.

Communication Link - Wireless

Wireless communications links do not have fixed points of termination, and provide the ability to communicate information across distances using radio frequencies. Wireless communications has been in existing for many years, but its high costs and low performance could only be justified by a few particular applications. The growth of cellular phone industry has motivated evolution if technologies that are reducing the cost and improving the performance of wireless data transmissions. Wireless data solutions come in many forms, but most can be broken down into four main categories:

- Circuit Switched Analog
- Circuit Switched Digital
- Packet Switched Analog
- Packet Switched Digital

Circuit-switched connections are similar to conventional wireline modem connections. A phone number is dialed by modem, and a temporary, switched connection is set up between two end points. The difference in wireless circuit switched is that the transmission medium, rather than being a pair

of copper wires as in the PSTN, consists of a pair of radio frequencies (or channel pairs) that are transmitted between some kind of portable transceiver and a base station.

The majority of today's public cellular phone networks still use analog signaling. Standard wireline modems will work in this environment, but because of the unique environment of the medium, the do not work very well. The reasons have to do primarily with the increased latencies involved in radio transmissions, the irregularities of signal strength, and the disruptions that occur in a cellular system when a signal is passed between one cell (or base station) and another. To compensate for these conditions, modem manufactures have developed specialized error correction protocols and implemented them in special cellular modems. These protocols include *ETC*, *MNP10ec*, and *TX-CEL*. Because these protocols work best when both modems on either end of a cellular connection use them, cellular service providers have begun implementing devices known as *modem pools* in their infrastructures. These devices implement cellular protocols, and will redirect incoming calls to standard modems automatically, providing the additional benefit of error correction on each end of the connection.

Digital signaling is rapidly being deployed in public cellular networks. Their are many benefits, but the most significant one is that the digital signaling can multiplex multiple connections over an existing frequency, thereby greatly expanding the call handling capacity of the cellular carriers infrastructure. Circuit switched digital connections tend to be more robust than analog, but they require special digital modems or data cards. Unfortunately, there are a number of different standards for digital signaling and all are vying for acceptance. They differ primarily in the method in which the divide or allocate frequency, which is mostly transparent to end-users. The three most popular standards are:

- Code Division Multiple Access (CDMA)
- Time Division Multiple Access (TDMA)
- Global Messaging System (GSM)

GSM has the widest implementation of digital cellular today, but is mostly found outside of the United States, and has become the standard in Europe. TDMA is the predominate digital standard in the US, but is an older technology than CDMA and lacks features compared to the newer specification. All three standards will be employed by different carriers as the new Personal Communications System (PCS) infrastructure is built out in the US.

In contrast to circuit-switched connections, packet switched systems do not use telephone numbers to place calls, but behave more like virtual LANs, where a device registers itself to a wireless network and then maintains constant connectivity until deregistered. The key attribute of packet-switched architectures is in the protocols that organize the transmission of data across the wireless medium. In circuit-switching systems, a virtual pipeline is opened between two end points in a stream through the exclusive connection. In packet switching, the data is assembled into individual units, or packets, each containing address information for the origin and destination of the packet. This allows the transmission to be shared, with packets from different origin points bound for different destinations all traveling over the same medium.

The primary advantage of packet systems is that because the resource is shared, service providers can charge for usage based on data sent rather than time connected, a significant advantage for certain types of applications which need to be able to send and receive constantly very small amounts of information. Another advantage of packet-switching is that the time required by circuit-switching systems to set up and establish a connection is greatly reduced or eliminated.

Today, most packet-switching wireless systems run over analog infrastructures. A popular standard in the US is called Cellular Digital Packet Data (CDPD) and uses the analog cellular telephone network for transmission. There are also packet-switching services built on private radio networks, such as the Mobitex standard used world-wide and know by the RAM Mobile Data brand name in the US. Digital packet-switched networks are in development, and expected to be deployed in the next couple of years.

Access Methods

Access methods describe the way in which the remote/mobile computer interacts with the resources it is connecting to. The type of access method used depends on the application requirements, and the type of resources that the remote computer is accessing. Remote access methods have been evolving, beginning in the early days with fixed function devices access mainframe computers, to today's implementations of remote client/server systems. All are still in use today. The four main methods are:

- Remote Dedicated Application Access
- Remote Control Access
- Remote Node Access

Remote Dedicated Application Access

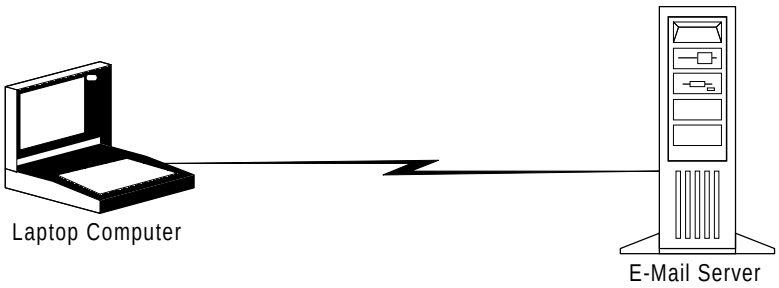

Remote application access involves a remote device accessing a host computer that is dedicated to a specific application. This dedicated computer may be a mainframe, a server running on a LAN, or a stand-along PC. Usually, remote application access configurations involve the use of a proprietary communications protocol. The most common application type used in Remote Application access is electronic mail. While still very prevalent, the trend in remote access is moving away from dedicated application access, where users are restricted to a single service, to other access methods that have greater flexibility.

Remote Control Access

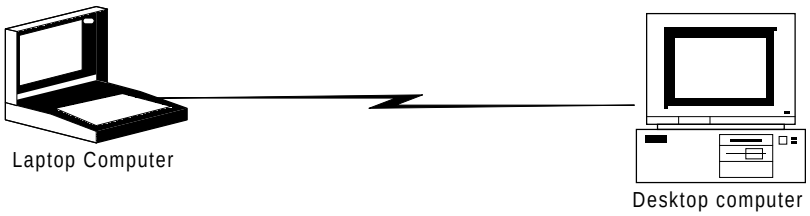

Remote control is a PC based access method. In this configuration, a connection is made between to PCs, and one takes over the input and control of the other. The remote PC is allowed to manipulate the keyboard and mouse of the host PC, and views the output of the host PCs monitor on its monitor, usually in a window. Only screen images, keystrokes, and mouse motions are sent across the remote connection. Remote control can be an efficient way of running applications that generate significant data traffic, such as database reporting. Remote control is also a popular way for help-desk services to provide support to remote users. Help desk staff can take over the remote machine and conduct problem determination processes remotely.

Most remote control configurations are peer-to-peer; that is, on PC takes over another PC. A relatively new variation on this access method, however, is a configuration where a single application runs in multiple instances on a multi-tasking server. Concurrent users can log in to the server and use remote control to execute the application.

Remote Node Access

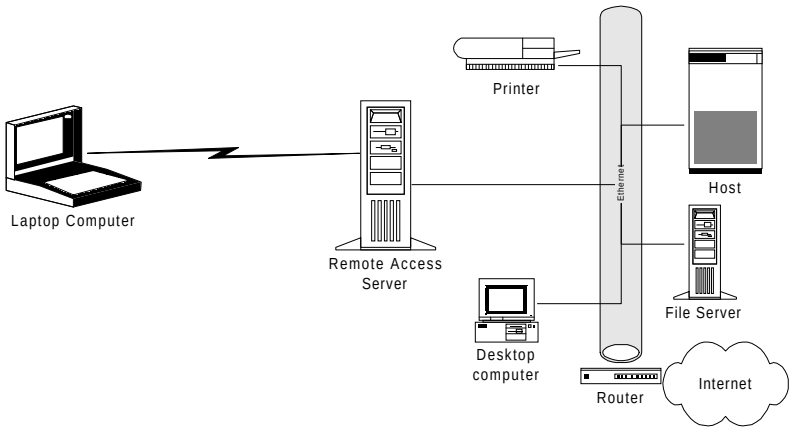

With this access method, the remote computer is accessing a network, not another computer. The remote device , utilizing a remote access server, emulates a node on a local area network. The remote

computer is literally on the network, with access to all resources that a locally attached PC would have, including any interconnected networks.

Remote node access is the most intuitive form of remote access for end-users, and provides the highest degree of flexibility and functionality. These advantages come at the price of performance. As a node on a LAN, the remote device must participate in all LAN traffic, including overhead associated with LAN protocols that were designed to run over mediums 100 times faster than the remote link. Certain features have been developed by remote access server vendors to help minimize this overhead, but remote node may not meet the performance requirements for all application types or users.

Protocols

The networking protocol used for remote access is perhaps the most defining of the configuration components, and has the most far reaching impact on performance, flexibility, and cost. There are a number of different protocols that can be considered, and to make a proper selection, it is very important to consider all the services and applications that are required by use of the remote user. Protocols function at all levels of a communications configuration, from the electronic signaling that occurs at a hardware level, to the format and presentation of characters and graphics on a display. For this discussion, to areas of protocols will be review; *remote access protocols* and *networking protocols*.

Remote Access Protocols

These protocols manage the communications as it is formatted and transmitted across the switched communications link. There are two basic types of remote access protocols: proprietary and open. Proprietary remote access protocols have been developed by different vendors to provide solutions for their specific products. Most are associated with remote dedicated application access, but some have been provided by vendors of LAN network operating systems for remote node access. Proprietary remote access protocols include:

- AppleTalk Remote Access Protocol (ARAP)
- Asynchronous NetBEUI (AsyBEUI)

Open remote access protocols have been developed by standards committees and are available from all vendors on many different types of platforms. Specifically, the Internet Engineering Task Force, which governs the development of technologies for the Internet, has codified standards for remote access protocols that are generally associated with the network protocol, TCP/IP (see below), but some of which have evolved to support network protocols other than TCP/IP. The three most common open remote access protocols are:

- Serial Line Interface Protocol (SLIP)
- Compressed SLIP (CSLIP)
- Point To Point Protocol (PPP)

Remote Access Protocols Compared

Remote Access Protocol	
SLIP	Added in the early 1980s to Berkley UNIX, SLIP uses a very simple framing scheme for TCP/IP packets over serial connections
CSLIP	An enhancement to SLIP added to Berkley UNIX in 1988, it compresses the 20 byte IP packet header to as little as 3 bytes
PPP	A major improvement to SLIP, first proposed in 1989, it allows multiple protocols to share a single serial line, including TCP/IP, IPX, and Appletalk. PPP also provides for automated negotiation of performance parameters, compression, and authentication
ARAP	A framing scheme for Appletalk packets over serial links, it also provides for encapsulation of IP over Appletalk
AsyBEUI	Microsoft developed protocol originally used for remote connections to its Network Operating System LAN Manager

Network Protocols

Network protocols are used to set up and manage connections between applications, and manage the process of addressing and routing traffic across multiple networks. Again the distinction of open and proprietary can be applied to network protocols, with the predominant open network protocol today being TCP/IP. The two most common proprietary network level protocols found in remote access configurations are Novell's IPX and Microsoft's NetBEUI..

Networking protocols, while all providing similar functional capabilities, differ greatly in there performance attributes. A key distinction between TCP/IP and both IPX and NetBEUI is the environments that they were originally designed to operate in. Both IPX and NetBEUI were designed to connect PCs across local area network environments. In order to achieve performance advantages in a local area network, both IPX and NetBEUI made certain design decisions that are not optimal in switched connection, WAN environments with their slower bandwidths and higher latencies and delays. These design attributes include the use of broadcast packets for address discovery, where at a high frequency, packets are sent across the entire network to interrogate status of network stations. LAN network protocols also tend to use simple acknowledgment routines, waiting for an acknowledgment to each and every packet sent before transmitting the next packet. These attributes combine to make very inefficient use of remote communications links.

Benefits Of Using TCP/IP

Unlike LAN network protocols TCP/IP was designed from the beginning to efficiently use wide area networks and remote communications links. TCP/IP minimizes the use of broadcast mechanisms for address and state discovery. It also employs more efficient acknowledgment algorithms, allow for continuing transmission of user data packets while waiting for acknowledgments.

In addition to performance efficiencies, there are a number of other reasons why TCP/IP has become the network protocol of choice for remote and mobile communications. These include:

- Standard application API, which helps protect corporate investment in applications by ensuring that that the interface for networking can be used across all kinds of platforms and in different networking environments. TCP/IP uses SOCKETS and WINSOCK as its API, and these are becoming the defacto standard for application development today
- Consistent User Interface, which means that because the same networking protocol can be used for both LAN and remote access, applications and services that users interact with remain the same, no matter where they are.
- Access To Host Systems: Most legacy application platforms today support terminal emulation access via TCP/IP
- Support Of NetBIOS: NetBIOS is the API used by the NetBEUI network protocol. A standard has been developed to run these type of applications directly over TCP/IP
- Enterprise Standardization: The vast majority of all new networking implementations today in corporate environments are being based on TCP/IP. This means that support staff responsibilities can be consolidated when it comes to supporting networking protocols
- Seamless Internet Access: As more business services are being based on the Internet, it is important that remote access strategies provide access to this environment. TCP/IP is the network protocol which the Internet is based on

Considerations For Using TCP/IP In Remote/Mobile Environments

Performance Tuning

TCP/IP implements performance parameters that can be manipulated to achieve different performance profiles depending on your environment. There may be some tuning that will benefit remote/mobile implementations of TCP/IP to achieve optimal throughput, but this needs to be balanced with application behavior. Tuning for optimal performance of an interactive application like telnet will not achieve optimal throughput for batch applications like FTP. Parameters to consider are:

- SLIP/PPP frame size
- Maximum Transfer Unit (MTU) size
- Window size

Security

Remote Access, by definition, represents an increased security exposure to corporate computing assets. Security concerns in remote access fall into two basic areas, with corresponding solutions:

- Unauthorized Access - Authentication
- Privacy - Encryption

Unauthorized access is probably the predominate concern of most corporate network managers. The challenge is to ensure that while enabling remote workers to gain access to resources necessary for them to conduct their business, a company doesn't allow unauthorized individuals to gain access to the same resource for mischievous or destructive purposes. Authentication is the process of ensuring

that the remote user is who he/she says they are, and that there access to corporate computing resources follows business policies.

Authentication for remote access can be implemented on multiple levels with increasing degrees of security. These different levels may also coexist. At a basic level password security. These may be provided by the application or network operating system, or even at an individual file level. The problem with basic password protection is that it can be easily comprised, either through guessing or observation. At a higher degree of security, authentication can be provided by Protocol Enforced Security (PES). The most common example of this is the authentication contained within PPP, which offers two types of security. The first, called Password Authentication Protocol (PAP), requires the remote client to send in a user name and password which is then validated by the PPP server. The second form of PPP security is called Challenged Handshake Authorization Protocol (CHAP), and uses secret passwords, and encrypting their transmission with random number generation

More involved authentication mechanisms can be employed to provide a higher degree of security. Some use smart card technology, wherein a remote users carriers a credit card sized device that has a secret and continually changing password which has been previously synchronized with a security server. Upon login, the user is requested to enter the password currently being displayed on the smartcard, which is then validated against the server. Even higher levels of security can be achieved using key encryption technology, such as Kerberos.

Encryption technologies are used to maintain the privacy of some or all of the data that is being transmitted across the remote communications links. Encryption technologies use some form of public or private key mechanism to allow both ends of the communications system to interpret the data. Encryption can be implemented in either hardware or software, and tends to be fairly costly in terms of computing resources.

Middleware

An increasing number of vendors are providing solutions that help integrate different remote and mobile communications link types into a single solution, and at the same time, try to improve the performance of applications in remote/mobile environments. This solutions are usually provided as software development kits and often come with proprietary application interfaces. They also require client/server configurations, meaning that functionality is provided by specific software features in both the remote client and a local server.

In the area of integration, middleware solutions try and hide the underlying communication link from the application, so that a application can take advantage of both wireline and wireless communications links of different types, usually in order to compensate for the lack of availability of one particular wireless communications infrastructure in all geographic regions.

Middleware also tries to improve performance of applications in remote/mobile environments, typically through a combination of compression and caching mechanisms, with the objective of cutting down on the transmission of redundant or irrelevant data.