OSI Transport and Session Layer Protocols

In this report:

Datapro Summary

With the U.S. government's emphasis on GOSIP, many major vendors now have certified OSI products. The pivotal OSI protocols reside at the Transport and Session Layers, levels 4 and 5, respectively. Used as a set, these two protocols allow information to pass reliably from one computer to another. Their purpose is to bridge applications and communications networks, and they are similar to the TCP/IP protocol suite defined by the DOD. There are two versions of the protocols, connection oriented and connectionless, and they support different functions and features depending on the upper layer applications and lower layer network types with which they are paired. Transport and Session Layer Protocols are mature standards with certified products now available for most major platforms. In selecting them over a vendor's proprietary protocol or TCP/IP, a user or system integrator is assuming little risk.

The OSI Reference Model and its associated protocols and services have been under development for ten years. One barrier to OSI's acceptance in corporate networks has been the lack of products based on the OSI protocols. In the last year, a signifcant number of major vendor products have been tested for OSI certification.

These products can be attributed to the federal government's aggressive application of a Federal Information Processing Standard (FIPS) called the Government OSI Profile (GOSIP). GOSIP compliance became mandatory in August 1991. Compliance to GOSIP Version 2 became mandatory in October 1992.

An important factor affecting OSI acceptance is the availability of products implementing the protocols defined at the Transport and Session Layers of the OSI Reference Model. The protocols at these two layers provide the glue between applications and communication networks.

The OSI Reference Model's Transport and Session Layers, or middle layers, provide the cornerstone of reliability as information passes from one computer to another. In OSI's upper layers, each application requires its own set of protocols. At the lower layers, each network technology requires its own protocols. In the middle layers, a very small set of protocols is defined to bridge the applications and the communication networks. It is only here, at the Transport and Session Layers, that sufficient commonality exists to allow for a small protocol set.

The central theme of the Transport and Session Layers is reliability. The Transport Layer provides a communication network's end-to-end reliability and ensures that communication errors are not the concern of the remaining upper layers. Reliability and synchronizing the data flowing between applications are two of the Session Layer's concerns. It provides the application tools to ensure the reliable exchange of information.

The OSI standards definition involves both ISO and CCITT. Presently, the Transport and Session Layer standards for connection-oriented communication are joint standards. CCITT has recently adopted the connectionless service descriptions and will, in the next year, adopt the protocol descriptions corresponding to the existing ISO standards.

—By James Moulton
President and Principal Consultant, Open Network Solutions, Inc.

Middle Layer OSI Protocol Standards

	Connection Oriented	Connectionless
Session Layer (Level 5)	IS 8327/CCITT X.225	IS 9548
Transport Layer (Level 4)	IS 8073/CCITT X.224	IS 8602

A major new work area has started within ISO and CCITT for developing multicast services and protocols. Work on the Transport Layer extensions has already begun, and significant progress is expected in 1993.

OSI Concepts

Service Definitions

The OSI Reference Model presents the architecture for exchanging information by distributed applications. The model describes a layered approach where each layer has a defined scope and set of functions.

Since the model is purely an architecture, at each layer a service definition describes the abstract interface to a layer protocol. The service describes the protocol user's interaction to the protocol implementation. Since it is abstract, however, it does not define a specific protocol interface used for implementation. Also, many protocols can support a single service definition.

Protocol Specifications

For each service definition, one or more protocols can be defined. Each protocol will follow the service interface while providing different mechanisms and functions. Conformance can only be tested against the protocol specification.

Communication Modes

The OSI Reference Model describes two distinct types of communication: connection-oriented and connectionless data transmission. At each layer of the model, there are both a connection-oriented and a connectionless service. Additionally, protocols supporting each mode of operation are defined.

Connection-oriented communication was the original focus of the OSI Reference Model. In this mode of operation, there are three phases of communication: connection establishment, data transfer, and connection release. The data transfer phase is simplified by maintaining sufficient state information. A typical example of this mode of operation is X.25 virtual circuits.

Connectionless data transmission was added to the model as a result of emerging datagram services at the Network and Data Link Layers. In a connectionless data transmission, there is only a data transfer phase. No state information is maintained, and each transmission is viewed as independent. Examples of connectionless mode transmissions are Logical Link Control (LLC) and the Connectionless Network Protocol (CLNP). For efficiency, the architecture requires that no segmentation or reassembly take place above the Network Layer. For that reason, the protocols at all of the higher layers are very simple. The majority of the functions involve address mapping from one layer to the next.

The Transport Layer

Moving from the Physical Layer to Application Layer, the Transport Layer is the first layer in the OSI Reference Model that maintains end-to-end data significance. That means that the Transport Layer Protocol provides functions between the data's source and destination. The Transport Layer transparently transfers data between the two systems and masks the communication network's specifics from the Session Layer and above. The Transport Layer's central theme is this masking of communication network differences.

Within the Transport Layer are approved standards for both the connection-oriented case (ISO 8073 and CCITT X.224) and for the connectionless case (IS 8602). The connection-oriented protocol was completed in 1984 and has remained very stable over time. Version 2 of ISO 8073/X.224 has just been completed. It adds several performance enhancements while maintaining backward compatibility. The connectionless protocol was completed several years after the connection-oriented protocol; however, since it is a simpler protocol, it has been very stable since its approval. CCITT is in the process of approving the connectionless protocol for inclusion in its series of Recommendations.

Connection-Oriented Transport Layer Protocol

Purpose of Protocol

The purpose of the Transport Layer Protocol is to bridge the gap between the capabilities provided by the Network Layer and the requirements of the Session Layer. The Transport Layer Protocol is responsible for masking differences in the quality of the underlying network from the Session Layer and above. The Transport Layer's main function is to ensure the reliability of the data transfer. Reliability is defined as the error-free, sequenced transfer of data. A secondary function of the Transport Layer is to optimize the use of the communication facilities found at the Network Layer and below. This is accomplished by matching the communication network's characteristics with the application's requirements and adding the functions necessary to achieve those requirements. It also forces the Transport Layer to decide which communication network to use when a choice is available.

The Transport Layer Protocol consists of several classes; however, all protocol classes provide the same service interface to the Session Layer. This allows the selection of the class of protocol on a connection-by-connection basis without affecting the Session Layer Protocol's operation or the mapping of the Session Layer Protocol to Transport Layer Services.

Functions and Features

The Transport Layer Protocol consists of the following functions:

- Error detection—the Transport Layer Protocol provides mechanisms to detect errors not signaled by the Network Layer.
- Error correction—the protocol can correct for detected errors or network signaled errors.
- Flow control—the destination system can regulate the flow of data from the source to the destination system.
- Addressing and address mapping—functions within the Transport Layer Protocol map transport addresses to the appropriate Network Layer address (NSAP).
- Multiplexing onto network connections—functions and mechanisms within the Transport Layer Protocol optimize the use of the network resources by using a single network connection for multiple transport connections.
- Sequence control—mechanisms in the Transport Layer Protocol ensure the delivery of data to the destination user in the same sequence as it is sent.
- Segmentation, blocking, and concatenation—the transport protocol manipulates the data received at the source system to most effectively utilize the network connection's resources.

- Expedited data transfer—the transport protocol provides a separate, flow-controlled data path for "urgent" data. Expedited data is always delivered to the destination user as quickly as possible.

The Transport Layer Protocol definition consists of several classes and options. Each class implements some or all of the functions, and each uses the mechanisms appropriate to the class for implementing the chosen functions. By tailoring the functions provided by the protocol class to the services provided by the Network Layer, the Transport Layer Protocol can provide constant Transport Layer Service.

Class Structure

The Transport Layer Protocol must efficiently operate over a large number and many types of communication networks. These networks can range from very reliable networks that always provide error-free, sequential data to networks providing no assurance of reliability. To achieve the optimal matching of requirements and functions, the Transport Layer Protocol has five classes.

The five classes fit into a two-tiered hierarchy. The hierarchy assists in ensuring interoperability when systems contain different classes. One tier consists of Classes 0, 1, and 3. This tier contains protocols supporting CCITT applications and uses highly reliable, virtual circuit networks. The other tier consists of Classes 2 and 4, containing protocols required of generalized computer networks. Class 2 is most often used in general networks based upon X.25 technology; Class 4 is analogous to TCP and is used when network reliability is insufficient for the application's requirements.

The use of a particular protocol class is based primarily on network designers' choices. In the U.S., the two protocol classes most often used are Class 4 and Class 0. The Class 4 protocol, with its robust error correction mechanisms, is the most prevalent. It is used in LAN-based environments and when data must traverse an interconnected set of networks. The Class 0 protocol is most often used over public X.25 Value-Added Network (VAN) offerings or in situations where CCITT applications are used. A prime example of a CCITT application is the use of X.400-based mail services on VANs.

Class 0: The Class 0 protocol is the simplest class of operation. It is compatible with the Teletex protocols defined in CCITT. Its correct operation depends upon the services and reliability of the underlying network. It provides no error detection or control, no sequence monitoring, no explicit connection release, and no multiplexing. It achieves its functions by mapping the transport functions directly to the network's procedures.

The Class 0 protocol is defined to be closely tied to X.25. While the Class 0 protocol has an explicit connection establishment procedure, it uses the disconnect of the network connection as the means of breaking the transport connection. During the data transfer phase, the protocol does not acknowledge receiving data and uses the network's error scheme for error detection. If an out-of-sequence packet is detected, the connection is aborted. Flow control is not explicitly provided; the network's flow control is used.

Class 0 is used within the Teletex network and with many CCITT applications that operate over X.25.

Class 1: The Class 1 protocol is a simple extension of Class 0. It operates over networks that do not provide sufficient reliability, but where multiplexing and other complex functions are not required. Class 1 adds procedures for recovering from signaled network errors, such as X.25 RESETS.

The Class 1 protocol utilizes the network's flow control mechanisms. It provides the ability to acknowledge data using positive acknowledgments. During network-signaled error recovery, it provides both positive and negative acknowledgments. There are no error detection mechanisms used in Class 1.

Class 1 is used in many of the same places as Class 0, but especially where the underlying X.25 networks have too many signaled errors for efficient operation or where tariffs drive a solution with very long virtual circuit lifetimes.

Class 2: The Class 2 protocol is the simplest class that provides multiplexing. It provides functions for multiplexing several transport connections onto a single network connection, flow control, and expedited data transfer. It also may detect lost or missequenced data packets by observing a gap in the sequence numbers. No recovery is performed. It is efficient and provides minimal services over those provided by the network.

Class 3: The Class 3 protocol is an extension of the Class 1 protocol to provide multiplexing functions. It builds upon the Class 2 protocol's functionality by adding the capability to recover from network-signaled errors. Therefore, Class 3 provides the multiplexing features of the Class 2 protocol and the error recovery procedures of the Class 1 protocol.

The Class 3 protocol provides for explicit flow control using credit and window techniques. All data is acknowledged using positive acknowledgments. During recovery from errors, negative acknowledgments are used for data resynchronization. The Class 3 protocol does not include any mechanisms for recovering from unsignaled errors, such as lost or missequenced data.

Class 4: The Class 4 protocol is the most sophisticated of the various classes. It operates over any type of network, including networks based on a connectionless service. The Class 4 protocol provides error detection and recovery procedures ensuring that the data is delivered in sequence and without errors. The majority of the protocol functions use timers, so sender and receiver independence is maintained. The Class 4 protocol detects errors in packet sequencing, lost packets, duplicated packets, and errored packets. It uses a positive acknowledgment scheme where each successfully received packet is explicitly acknowledged. Flow control is provided through a complex credit allocation and window scheme. The credit can be both raised and lowered as congestion changes in Transport Layer Protocol buffers. Timers manage retransmitting lost packets and detecting half-open connections.

The Class 4 protocol is the class most often used in the U.S. Functionally, it is very similar to TCP and can be used wherever TCP is appropriate. It is now mandated for government networks through the GOSIP profiles.

GOSIP Considerations

GOSIP requires government agencies to use ISO protocols when implementing a new network. It specifically designates the Class 4 protocol, except in very limited cases. If a network is being designed for the government or to interface with the government's networks, the Class 4 protocol is most appropriate.

The GOSIP profile defines specific Class 4 parameters and options. Users can find these specific design constraints in the OSI Implementors' Workshop agreements.

GOSIP Version 2 makes no changes to the requirements for the Transport Protocol.

Connectionless Transport Layer Protocol

The second mode of operation within OSI is connectionless data transmission (connectionless). The Transport Layer defines a connectionless protocol allowing the mapping of connectionless requests onto the connectionless Network Layer Service.

To achieve a degree of efficiency, the connectionless protocol provides very minimal functionality. The connectionless Transport Layer Protocol is limited to procedures for mapping the transport address to a Network Layer address (NSAP). There are no procedures for segmenting the data into smaller sizes that would fit the packet sizes at the Network Layer. If the data in a request is too large for mapping to a single Network Layer data transfer, the data is discarded and an error message may be returned.

The connectionless protocol includes the facility for a checksum. If the checksum is used, it allows for limited error detection on the data. If the packet is detected to have an invalid checksum, it is discarded.

GOSIP Considerations

GOSIP Version 2 includes the connectionless Transport Protocol as an option. Where appropriate, this protocol can be used in federal networks requiring GOSIP-compliant protocols.

Multicast Transport Protocol Development

In 1992, significant interest was shown in developing extensions to the current Transport Layer Protocol in the area of multicast transmissions. Multicast allows a single data unit to be received by multiple recipients. Current OSI protocols do not provide multicast services; therefore, either new protocols or extensions to the existing protocols will be needed to meet multicast requirements. At the Transport Layer, projects are under way to define the multicast services to be provided. It is expected that for some multicast services, extensions will be made to the existing Transport Layer Protocols, especially the connectionless protocol. For the most difficult multicast operations, a new protocol will most likely be defined. Initial work on the new protocol has already been published, and significant progress is expected in 1993. A draft standard should be available by the first part of 1994.

The Session Layer

The Session Layer's purpose is to provide a means for cooperating applications to manage their dialog and to organize and synchronize the data exchange.

The Session Layer Protocol is stable and has not been changed in the last several years.

Connection-Oriented Session Layer Protocol

The Session Layer deals with ensuring that the dialog between applications remains synchronized and that errors are detected and corrected. The Session Layer Protocol assumes that all errors occurring in the communication network are corrected by the protocols at the Transport Layer. For that reason, errors that may occur are those relating to the system's operation or to the application itself. These would include tape/disk errors, resource allocations, buffer space, or processing errors.

As with all OSI standards, there are two sets of standards: one for the connection-oriented case and one for the connectionless case. Within each set, there are separate standards for services and protocols. The connection-oriented Session Layer Service is defined in IS 8326/CCITT X.215. The protocol is defined in IS 8327/CCITT X.225. For the connectionless case, the service is defined in Addendum 1 to IS 8326, and the protocol is defined in IS 9548.

The Session Layer Service is different from the Transport Layer Service, depending upon the application's requirements. Where the Transport Layer has a common service and different protocols supporting that service, the Session Layer has different services based upon the application's requirements. Each aspect of the Session Layer Service is supported by unique protocol mechanisms.

There are two versions of the Session Layer Protocol: Version 1 and Version 2. The two are differentiated by the amount of user data that is permitted. Version 2 relaxes the limit on the length of user data, permitting it to be unlimited. The negotiation rules for connection establishment ensure that a Version 2 implementation will be capable of communicating with a Version 1 implementation.

Purpose of Protocol

The Session Layer Protocol exists to provide synchronization between applications. The Session Layer controls the dialog between two applications without regard for how the applications communicate. In particular, the Session Layer Protocol deals with the exchange of data between two applications so that each maintains a consistent view of the other application's view.

Functions and Features

The Session Layer Protocol provides functions that can be tailored to the needs of a specific application. The mechanisms used to tailor the functions in use are *functional units.* A functional unit is a set of protocol mechanisms that, when used together, perform one or more protocol functions.

The Session Layer Protocol supports the following:

- *Normal data transfer*, the transfer of user data between applications.
- *Token management*, where tokens are used to implement turn control. A token's owner can invoke certain services and functions, whereas the other side cannot perform the function until it has the token. An example is the data token used to enforce half-duplex operation.
- *Exception reporting*, where the protocol returns particular status and error information to the user.
- *Typed data transfer*, where the protocol can pass a limited amount of data against the data token. This is analogous to a reverse channel available in certain half-duplex operations.
- *Minor synchronization*, where the protocol supports a type of data synchronization in which sync points are inserted into the datastream. With minor synchronization, data can continue to be sent while waiting for the acknowledgment of the sync point.
- *Major synchronization*, where the user must suspend data transfers until the receiver acknowledges the major sync point. Resynchronization cannot be performed to a point in the datastream before the last acknowledged major sync point.
- *Resynchronization*, where the protocol supports resynchronization, in which the sender and receiver back up to a specified synchronization point. The selected point can be any synchronization point up to and including the last acknowledged major sync point.
- *Expedited data transfer*, where the protocol supports sending expedited data between users.
- *Activity management*, where the protocol supports the concept of activities that delineate a dialog between the sender and receiver of data. An activity comprises several major sync points and is often used in the Teletex service.
- *Capability data*, where the protocol supports the exchange of data concerning the capabilities of terminal equipment.

Functional Units

The Session Layer Protocol is defined as a set of functional units. Each functional unit contains protocol functions and mechanisms supporting specific Session Layer Services. The functional units are selected in a negotiation process during connection establishment.

There are 12 functional units defined:

1. Kernel
2. Negotiated release
3. Half duplex
4. Duplex
5. Expedited data
6. Typed data
7. Capability data exchange
8. Minor synchronize
9. Major synchronize
10. Resynchronize
11. Exceptions
12. Activity management

At connection establishment, the two Session Layer users negotiate which functional units will be present during the life of the connection. The selected functional units define the Session Layer Protocol's functions for that connection.

GOSIP Considerations

The Government Open Systems Interconnection Profile (GOSIP) Version 1 specifies support for both Version 1 and Version 2 of the Session Layer Protocol. In addition, GOSIP specifies that the negotiated release and capability data functional units are not required. All other functional units are allowed. GOSIP states that selecting which functional units to request for a given application will be specified by the application.

GOSIP Version 2 makes no major modifications to the requirements for the Session Layer.

Connectionless Session Layer Protocol

The connectionless Session Layer Protocol is very simple to ensure the efficiency of connectionless operation. In particular, IS 9548 provides limited procedures for address mapping and for mapping the Session Layer Protocol to the connectionless Transport Layer Protocol. The Session Layer Protocol does not provide any functions for segmenting large data requests; rather, the data is discarded, and the user is informed.

Applying Protocols to Problems

The Transport and Session Layer standards are not used independently of other protocols to provide an integrated solution. These protocols form a foundation for developing distributed applications in a network environment. How do these standards relate to the rest of the problem, such as selecting a protocol suite and implementing specific applications?

Relationship to TCP/IP

The protocols defined by the DOD, TCP/IP, have gained significant acceptance as network standards. TCP/IP protocols were designed before the wide acceptance of the OSI Reference Model and are not layered according to the same principles. For that reason, TCP/IP encompass many of the features of the Network, Transport, and Session Layers of the OSI Reference Model (see Figure 1).

Figure 1.
OSI and Internet Protocol Stacks

	OSI	Internet
	Application Layer	Applications Telnet FTP SMTP
	Presentation Layer	
Middle Layers	Session Layer	TCP
	Transport Layer	
	Network Layer	IP
	Data Link Layer	Link
	Physical Layer	Physical

The Transport Layer Protocol Class 4 was designed to be nearly equivalent to TCP. Some differences exist between the two protocols, but in most cases where TCP is used today, Class 4 can be used.

The differences between the OSI protocol suite and the TCP/IP suite are more fundamental. The interface to TCP/IP is different from that of the Class 4 protocol, making it difficult to operate the OSI upper layer protocols over TCP/IP or to run the DOD protocol over the Class 4 protocol. A major issue in deciding between using TCP/IP and OSI is the possibility of creating parallel protocol stacks and subsequent migration to a single stack.

It is possible to implement both protocol stacks, however, and to pass data between the stacks using application gateways. This approach allows users to keep using protocols and applications with which they are comfortable. It is also possible to implement an OSI environment by implementing the OSI upper layer protocols over TCP/IP. This allows users to gain familiarity with the new protocols without requiring the change-out of existing applications. Finally, it is possible to implement the full OSI protocol suite. This requires a complete change-out of protocols and applications. It requires that all network nodes successfully migrate to the new protocols on the day of the change. This approach is the riskiest for downtime and network outages; however, it does create a clean break to the new protocols.

Network and Presentation Layer Issues

Network Layer Issues

The entire focus of the Transport Layer is to hide the differences of network technologies and approaches. The prime reason for the Transport Layer is to allow applications design without regard to the communication characteristics. For that reason, the Transport Layer Protocol can operate over many different networks.

The connection-oriented protocol maps its requirements onto the connection-oriented Network Layer Service. Different types of networks, such as X.25 and ISDN, provide a mapping from the network-specific interface to the Network Layer Service. This provides a very clean interface between the two layers. The actual interface implementation will be specific to each system and may require custom software to ensure that the Transport Layer Protocol interface matches the appropriate Network Layer interface.

Due to the sophistication of the Class 4 protocol, there is a mapping defined to the connectionless Network Layer Service. This mapping provides the capability to operate Class 4 over the OSI version of IP, sometimes called the Connectionless Network Protocol (CLNP). Of the connection-oriented classes, only Class 4 is permitted to operate over the connectionless Network Layer Service.

The restriction on the Transport Layer connectionless protocol's functionality, and the desire for efficiency, limits the protocol's mappings to Network Layer Services. It is restricted to operate only over the connectionless Network Layer Service. The connectionless Transport Layer Protocol provides only that single mapping.

Presentation Layer Issues

The Session Layer is the only OSI layer where the services substantially change based upon the application requirements. This has a major impact upon the Presentation Layer in that the services it can provide to the application depend upon the services available at the Session Layer.

The Presentation Layer standards define the precise mapping of Presentation Layer Services offered to the application based upon the services provided after establishing a session connection. The appropriate mappings are carefully defined in the standards and in documents such as GOSIP.

The range of services offered by the Presentation Layer requires a close coupling of the Presentation Layer Protocol and the Session Layer Service. Therefore, there are few interface issues relating to mapping the Presentation Layer Protocol onto the Session Layer Service.

Specific Applications

Each of the OSI upper layer protocols is specified with mappings onto the Presentation Layer Service. As previously described, the Presentation Layer Protocol is tightly coupled to the Session Layer Service. In practice, mapping the upper layer protocols to the Session Layer Service is the only practical way to determine which Session Layer Services are required.

The services required of each upper layer application are tailored according to the type of interactions the applications require. While there is commonality between many of the services and applications, it is typically necessary to implement the full Session Layer Protocol.

File Transfer, Access, and Management

File Transfer, Access, and Management (FTAM) is the OSI file transfer protocol. It provides the mechanisms for defining, accessing, and managing a virtual file system in the OSI environment. FTAM defines not only the protocol for exchanging files; it also provides extensive facilities for defining the filestore structure. The filestore definition allows for complex file operations to occur on dissimilar systems.

FTAM is often used to retrieve or copy files from one system to another. The capability to move files is characterized by a large amount of data that must be reliably transferred. The transfer is basically one direction; however, controlling the transfer requires the capability to send information in the "reverse" direction. Reliability requires that the transfer be checkpointed and that recovery procedures be available upon a detected transfer error.

According to GOSIP Version 1, File Transfer, Access, and Management (FTAM) requires Version 2 of the Session Layer Protocol. It requires using the kernel and duplex functional units. FTAM may optionally request the use of the resynchronize and minor synchronize functional units. GOSIP Version 2 extends FTAM into new areas. It does not require additional Session Layer Services.

At a minimum, GOSIP requires that FTAM be operated with the kernel and duplex functional units. With these selected, the transfer would not have error control and, upon detecting an error, the entire transfer would be aborted. This simple requirement is based upon the idea of requiring the minimum Session Layer Services possible in order to have a useful file transfer capability. Adding the optional functional units for resynchronize and minor synchronize allows FTAM to include error recovery procedures. With that addition, detected errors require only a resynchronization to known sync points. Utilizing minor synchronization points allows the transfer to proceed while awaiting the acknowledgment, speeding the overall transfer.

X.400—Message Handling System

The X.400 protocol suite comprises the CCITT Message Handling System. X.400 provides a standard electronic mail environment and the protocols required to exchange mail messaged between distributed mail systems. The X.400 electronic mail system consists of several protocols, including protocols that only deal with exchanging mail messages between systems. Other protocols address issues such as the user interface and storing and retrieving messages.

The X.400 protocol suite is defined by CCITT. As part of the specification, CCITT specifies the complete CCITT protocol stack used for communication. That protocol stack includes the Class 0 Transport Layer Protocol and most of the Session Layer Protocol's functional units. This protocol suite is applicable to public data networks, however, or for systems exchanging information with public data networks. X.400 systems that are not used over public networks need not follow these recommendations.

According to GOSIP Version 1 and GOSIP Version 2, X.400 requires Version 1 of the Session Layer Protocol when the X.400-1984 version is used. GOSIP Version 2 also allows for the use of X.400-1988; in this case, Version 2 of the Session Layer Protocols may be used. It requires using kernel, half-duplex, exceptions, activity management, and minor synchronize functional units. That set of functional units is equivalent to those specified by CCITT.

GOSIP recognizes the need to maintain capability with the mappings of application services to Session Layer Services. This is maintained in the consistent definition of Session Layer Services. (In the case of the Transport Layer Protocol, GOSIP specifies the Class 0 protocol, where interconnection with public systems is required. Otherwise, it allows the use of the Class 4 protocol.)

Virtual Terminal Protocol

The Virtual Terminal Protocol (VTP) provides a consistent terminal interface to applications across a distributed environment. The VTP provides a "virtual" terminal session that can be mapped to a variety of real terminals. VTP's scope is a simple two-dimensional terminal, such as a Digital VT100. It does not provide editing capabilities found on some complex forms-type terminals. VTP can control both the placement and look of characters (and graphics) within the viewing area.

VTP is an interactive protocol, where one user sends data and then waits for the response. A critical VTP component is to ensure that the view of the two users remains in sync. That is, both users have a consistent view of the screen information.

The Basic Class VTP is not included in GOSIP Version 1 and is now optional in GOSIP Version 2. The Basic Class VTP requires using the following functional units:

- kernel
- half duplex
- resynchronize
- major synchronize
- typed data
- expedited data

The functional units required of VTP mirror its transfer mode. To ensure that the data is written by only one user at a time, the VTP operates in a half-duplex mode. This is enforced through the half-duplex functional unit. The typed data and expedited data functional units provide a means of exchanging data outside the normal data path and permit the users to exchange status information at any time. The synchronization mechanisms are in place using the resynchronize and major synchronize functional units. The use of the major synchronize functional unit ensures that the data is acknowledged before more data is sent.

Emerging Application Protocols

As OSI continues to evolve, new protocols are being defined to take advantage of its distributed processing capabilities. Of particular interest is the work on Transaction Processing (TP). TP implements a distributed transaction model where single, atomic transactions can be processed in a distributed manner. TP implements significant mechanisms to ensure that no data modifications are made until all processing is complete. To ensure that the underlying protocols supply the needed services, TP requires that the Session Layer Protocol provide both major and minor synchronization as well as resynchronization. TP utilizes the duplex functional unit, since data flows are bidirectional. Finally, TP requires that the expedited data functional unit be implemented to ensure that synchronization is not lost due to flow control.

Practical Implementation Considerations

Since the emergence of the OSI Reference Model and the U.S. GOSIP program, most major vendors now have certified OSI-compliant products. A major issue in applying the OSI protocol suite is interoperability among vendors' products and how OSI fits into proprietary architectures. With a certified product, interoperability is not an issue. Vendors with major network architectures, such as IBM and Digital Equipment Corp., have announced products and plans to migrate portions of their networks to OSI. IBM has a complete set of middle layer products allowing the interconnection of systems operating the OSI protocols. These products also allow users to exchange data with SNA users. Digital has announced support for OSI protocols under DECnet Phase V. These products will implement the OSI model's middle layers.

Vendors with a commitment to Open Systems, such as Sun, Hewlett-Packard, and NCR, have certified complete OSI protocol stacks that operate across their hardware platforms. These products are integrated into their standard product offerings.

There are also vendors with "portable" products, typically UNIX based, which can be ported to a variety of systems. These allow users to select from appropriate vendors based on features, platforms, and prices.

System integration considerations concern which classes of the Transport Layer Protocol are required for the application and which of the functional units are supplied in a Session Layer Protocol. In most domestic cases, vendors are supplying GOSIP-certified products. This means that the products contain at least Class 4 of the Transport Layer Protocol.

For the Session Layer Protocol, it has become clear that no set of functional units is sufficient for all applications. For that reason, all products include all of the functional units defined in Version 2. Many applications now come bundled with the appropriate Presentation and Session Layer Protocols, thus ensuring that the application receives the necessary approach. It can also lead to duplicate protocol implementations based upon common features found in the Session Layer Protocol.

Using OSI in PC-based LANs has often been dismissed as infeasible, due to overhead and cost. Today, solutions are available that incorporate OSI into the server technology and allow the PCs and LAN operating systems to operate normally. As PCs have grown more sophisticated, OSI on PCs is a viable solution. (One can argue that if TCP/IP-based solutions are acceptable, then OSI-based solutions will also be acceptable.) The strategy for implementing OSI on PCs or in PC-based LANs must be tailored to the environments, many of which intentionally restrict users to the server platform. Others prefer to allow PCs to operate as peers in the LAN. An additional consideration is the protocols selected for use on the LAN. If OSI is destined to be used only for remote communication, a server solution may be more appropriate. If OSI will supply the protocols used on the LAN, then, by necessity, all PCs must use OSI.

Conclusion

The OSI Transport and Session Layer Protocols have reached a level of maturity and stability demonstrated in the marketplace with numerous commercial offerings. In selecting the OSI protocols over a vendor's proprietary offering or over the DOD TCP/IP suite, a user or system integrator is assuming little risk. The system integrator's role will be to incorporate vendor products into a cohesive network structure allowing for full protocol functionality but tailored to provide the required flexibility.

OSI's acceptance within both the commercial and government markets is leading to fierce competition from vendors of other solutions, which has tended to polarize the marketplace and drive vendors into providing full-functioned OSI solutions at very attractive prices. ■

This report was prepared exclusively for Datapro by James Moulton, president and principal consultant for Open Network Solutions, Inc. (ONS). Mr. Moulton has over 18 years' experience in data communication and telecommunication research and development. He has been involved in designing and standardizing the OSI protocols and has designed networks based on DOD protocols, TCP/IP, and GOSIP protocols. He is an active participant in international and national standards committees defining the network standards of the future.

ONS, based in Sterling, VA, is a diversified company providing solutions to data communication and networking problems. It has designed major network systems, including WANs and LANs for both industry and government. ONS is heavily involved in researching the next generation of network protocols. It is experienced in both GOSIP and TCP/IP networks as well as the range of LAN operating systems.

Middle Layer OSI Protocols

In this report:

Datapro Summary

Since the emergence of the OSI Reference Model and the U.S. government's GOSIP program, most major vendors have announced plans to supply OSI-compliant products. The pivotal OSI internetworking protocols consist of the Transport and Session Layers, levels 4 and 5, respectively. Used as a set, these two protocols allow information to pass reliably from one computer to another. Their purpose is to bridge applications and communications networks, and they are similar to the TCP/IP protocol suite defined by the DOD. There are two versions of the protocols, connection oriented and connectionless, and they support different functions and features depending on the upper layer applications and lower layer network types with which they are paired. Transport and Session Layer protocols have matured and stabilized in today's marketplace. In selecting them over a vendor's proprietary protocol or TCP/IP, a user or system integrator is assuming considerably less risk than in previous years.

Development of the OSI Reference Model and its associated protocols and services has been under way for ten years. A major barrier to OSI's acceptance in corporate networks has been the lack of mature protocols. It is only in the last few years that a number of OSI protocols have reached a sufficient level of maturity. Recognizing this, the federal government has mandated the use of OSI protocols through a Federal Information Processing Standard (FIPS) called the Government OSI Profile (GOSIP). GOSIP compliance became mandatory in August 1991.

An important part of accepting OSI is the availability of products implementing the protocols defined at the middle layers of the OSI Reference Model. The protocols at these two layers provide the glue between applications and communication networks.

The OSI Reference Model's middle layers consist of the Transport and Session Layers. Taken as a set, these two layers provide the cornerstone of reliability as information passes from one computer to another. In OSI's upper layers, each application requires its own set of protocols. At the lower layers, each network technology requires its own protocols. In the middle layers, a very small set of protocols is defined to bridge the applications and the communication networks. It is only here, at the Transport and Session Layers, that sufficient commonality exists to allow for a small protocol set.

The central theme of the Transport and Session Layers is reliability. The Transport Layer provides a communication network's end-to-end reliability and ensures that communication errors are not the concern of the remaining upper layers. Reliability and synchronizing the data flowing between applications are two of the Session Layer's concerns. It provides the application tools to ensure the reliable exchange of information.

The OSI standards definition involves both ISO and CCITT. Presently, the Transport and Session Layer standards for connection-oriented communication are joint

—By James Moulton
President and Principal Consultant
Open Network Solutions, Inc.

standards. CCITT is developing connectionless service and protocol descriptions corresponding to the existing ISO standards.

Middle Layer OSI Protocol Standards

	Connection Oriented	Connectionless
Session Layer (Level 5)	IS 8326/CCITT X.225	IS 9548
Transport Layer (Level 4)	IS 8073/CCITT X.224	IS 8602

OSI Concepts

Service Definitions

The OSI Reference Model presents the architecture for exchanging information by distributed applications. The model describes a layered approach where each layer has a defined scope and set of functions.

Since the model is purely an architecture, at each layer a service definition describes the abstract interface to a layer protocol. The service describes the protocol user's interaction to the protocol implementation. Since it is abstract, however, it does not define a specific protocol interface used for implementation. Also, many protocols can support a single service definition.

Protocol Specifications

For each service definition, one or more protocols can be defined. Each protocol will follow the service interface while providing different mechanisms and functions. Conformance can only be tested against the protocol specification.

Communication Modes

The OSI Reference Model describes two distinct types of communication: connection-oriented and connectionless data transmission (called connectionless). At each layer of the model, there is both a connection-oriented and a connectionless service. Additionally, protocols supporting each mode of operation are defined.

Connection-oriented communication was the original focus of the OSI Reference Model. In this mode of operation, there are three phases of communication: connection establishment, data transfer, and connection release. The data transfer phase is simplified by maintaining sufficient state information. A typical example of this mode of operation is X.25 virtual circuits.

Connectionless Data Transmission was added to the model as a result of emerging datagram services at the Network and Data Link Layers. In a connectionless data transmission, there is only a data transfer phase. No state information is maintained, and each transmission is viewed as independent. Examples of connectionless mode transmissions are Logical Link Control (LLC) and the Connectionless Network Protocol (CLNP). For efficiency, the architecture requires that no segmentation or reassembly take place above the Network Layer. For that reason, the protocols at all of the higher layers are very simple. The majority of the functions involve address mapping from one layer to the next.

The Transport Layer

Moving from bottom to top, the Transport Layer is the first layer in the OSI Reference Model that maintains end-to-end data significance. That means that the Transport Layer Protocol provides functions between the data's source and destination. The Transport Layer transparently transfers data between the two systems and masks the communication network's specifics from the Session Layer and above. The Transport Layer's central theme is this masking of communication network differences.

Within the Transport Layer are approved standards for both the connection-oriented case (ISO 8073 and CCITT X.224) and for the connectionless case (IS 8602). The connection-oriented protocol was completed in 1984 and has remained very stable over time. Because 8073/X.224 is a joint ISO/CCITT standard, modifications to the protocol have followed a path of backward compatibility. The connectionless protocol was completed several years later; however, since it is a simpler protocol, it has been very stable since its approval. CCITT is evaluating the connectionless protocol for inclusion in its series of Recommendations.

Connection-Oriented Transport Layer Protocol

Purpose of Protocol

The purpose of the Transport Layer Protocol is to bridge the gap between the capabilities provided by the Network Layer and the requirements of the Session Layer. The Transport Layer Protocol is responsible for masking differences in the quality of the underlying network from the Session Layer and above. The Transport Layer's main function is to ensure the reliability of the data transfer. Reliability is defined as the error-free, sequenced transfer of data. A secondary function of the Transport Layer is to optimize the use of the communication facilities found at the Network Layer and below. This is accomplished by matching the communication network's characteristics with the application's requirements and adding the required functions to achieve those requirements. It also forces the Transport Layer to decide which communication network to use when a choice is available.

The Transport Layer Protocol consists of several classes; however, all protocol classes provide the same service interface to the Session Layer. This allows the selection of the class of protocol on a connection-by-connection basis without affecting the Session Layer Protocol's operation or the mapping of the Session Layer protocol to Transport Layer Services.

Functions and Features

The Transport Layer Protocol consists of the following functions:

- Error detection, where the Transport Layer Protocol provides mechanisms to detect errors not signaled by the Network Layer.
- Error correction, mechanisms by which the protocol can correct for detected errors or network signaled errors.
- Flow control, mechanisms by which the destination system can regulate the flow of data from the source to the destination system.
- Addressing and address mapping, functions within the Transport Layer Protocol that map transport addresses to the appropriate Network Layer address (NSAP).

- Multiplexing onto network connections, functions and mechanisms within the Transport Layer Protocol to optimize the use of the network resources by using a single network connection for multiple transport connections.
- Sequence control, mechanisms in the Transport Layer Protocol to ensure the delivery of data to the destination user in the same sequence as it is sent.
- Segmenting, blocking, and concatenation, where the transport protocol manipulates the data received at the source system to most effectively utilize the network connection's resources.
- Expedited data transfer, where the transport protocol provides a separate, flow-controlled data path for "urgent" data. Expedited data is always delivered to the destination user as quickly as possible.

The Transport Layer Protocol definition consists of several classes and options. Each class implements some or all of the functions, and each uses the mechanisms appropriate to the class for implementing the chosen functions. By tailoring the functions provided by the protocol class according to the services provided by the Network Layer, the Transport Layer Protocol can provide constant Transport Layer Service.

Class Structure

The Transport Layer Protocol must efficiently operate over a large number and many types of communication networks. These networks can range from very reliable networks that always provide error-free, sequential data to networks providing no assurance of reliability. To achieve the optimal matching of requirements and functions, the Transport Layer Protocol has five classes.

The five classes fit into a two-tiered hierarchy. The hierarchy assists in ensuring interoperability when systems contain different classes. One tier consists of Classes 0, 1, and 3. This tier contains protocols supporting CCITT applications and uses highly reliable, virtual-circuit networks. The other tier consists of Classes 2 and 4, containing protocols required of generalized computer networks. Class 2 is most often used in general networks based upon X.25 technology; Class 4 is analogous to TCP and is used when network reliability is insufficient for the application's requirements.

The use of a particular protocol class is based primarily on network designers' choices. In the U.S., the two protocol classes most often used are Class 4 and Class 0. The Class 4 protocol, with its robust error correction mechanisms, is the most prevalent. It is used in LAN-based environments and when data must traverse an interconnected set of networks. The Class 0 protocol is most often used over public X.25 Value-Added Network (VAN) offerings or in situations where CCITT applications are used. A prime example of a CCITT application is the use of X.400-based mail services on VANs.

Class 0: The Class 0 protocol is the simplest class of operation. It is compatible with the Teletex protocols defined in CCITT. Its correct operation depends upon the services and reliability of the underlying network. It provides no error detection or control, no sequence monitoring, no explicit connection release, and no multiplexing. It achieves its functions by mapping the transport functions directly to the network's procedures.

The Class 0 protocol is defined to be closely tied to X.25. While the Class 0 protocol has an explicit connection establishment procedure, it uses the disconnect of the network connection as the means of breaking the transport connection. During the data transfer phase, the protocol does not acknowledge receiving data and uses the network's error scheme for error detection. If an out-of-sequence packet is detected, the connection is aborted. Flow control is not explicitly provided; the network's flow control is used.

Class 0 is used within the Teletex network and with many CCITT applications that operate over X.25.

Class 1: The Class 1 protocol is a simple extension of Class 0. It operates over networks that do not provide sufficient reliability, but where multiplexing and other complex functions are not required. Class 1 adds procedures for recovering from signaled network errors, such as X.25 RESETS.

The Class 1 protocol utilizes the network's flow control mechanisms. It provides the ability to acknowledge data using positive acknowledgments. During network-signaled error recovery, it provides both positive and negative acknowledgments. There are no error detection mechanisms used in Class 1.

Class 1 is used in many of the same places as Class 0, but especially where the underlying X.25 networks have too many signaled errors for efficient operation or where tariffs drive a solution with very long virtual circuit lifetimes.

Class 2: The Class 2 protocol is the simplest class that provides multiplexing. It provides functions for multiplexing several transport connections onto a single network connection, flow control, and expedited data transfer. It also may detect lost or missequenced data packets by observing a gap in the sequence numbers. No recovery is performed. It is efficient and provides minimal services over those provided by the network.

Class 3: The Class 3 protocol is an extension of the Class 1 protocol to provide multiplexing functions. It builds upon the Class 2 protocol's functionality by adding the capability to recover from network-signaled errors. Therefore, Class 3 provides the multiplexing features of the Class 2 protocol and the error recovery procedures of the Class 1 protocol.

The Class 3 protocol provides for explicit flow control using credit and window techniques. All data is acknowledged using positive acknowledgments. During recovery from errors, negative acknowledgments are used for data resynchronization. The Class 3 protocol does not include any mechanisms for recovering from unsignaled errors, such as lost or missequenced data.

Class 4: The Class 4 protocol is the most sophisticated of the various classes. It operates over any type of network, including networks based on a connectionless service. The Class 4 protocol provides error detection and recovery procedures ensuring that the data is delivered in sequence and without errors. The majority of the protocol functions use timers, so sender and receiver independence is maintained. The Class 4 protocol detects errors in packet sequencing, lost packets, duplicated packets, and errored packets. It uses a positive acknowledgment scheme where each successfully received packet is explicitly acknowledged. Flow control is provided through a complex credit allocation and window scheme. The credit can be both raised and lowered as congestion changes in Transport Layer Protocol buffers. Timers manage retransmitting lost packets and detecting half-open connections.

The Class 4 protocol is the class most often used in the U.S. Functionally, it is very similar to TCP and can be used wherever TCP is appropriate. It is now mandated for government networks through the GOSIP profiles.

GOSIP Considerations

GOSIP requires government agencies to use ISO protocols when implementing a new network. It specifically designates the Class 4 protocol, except in very limited cases. If a network is being designed for the government or to interface with the government's networks, the Class 4 protocol is most appropriate.

The GOSIP profile defines specific Class 4 parameters and options. Users can find these specific design constraints in the OSI Implementors' Workshop agreements.

Connectionless Transport Layer Protocol

The second mode of operation within OSI is connectionless data transmission (connectionless). The Transport Layer defines a connectionless protocol allowing the mapping of connectionless requests onto the connectionless Network Layer Service.

To achieve a degree of efficiency, the connectionless protocol provides very minimal functionality. The connectionless Transport Layer Protocol is limited to procedures for mapping the transport address to a Network Layer address (NSAP). There are no procedures for segmenting the data into smaller sizes that would fit the packet sizes at the Network Layer. If the data in a request is too large for mapping to a single Network Layer data transfer, the data is discarded and an error message may be returned.

The connectionless protocol includes the facility for a checksum. If the checksum is used, it allows for limited error detection on the data. If the packet is detected to have an invalid checksum, it is discarded.

The Session Layer

The Session Layer's purpose is to provide a means for cooperating applications to manage their dialog and to organize and synchronize the data exchange.

Connection-Oriented Session Layer Protocol

The Session Layer deals with ensuring that the dialog between applications remains synchronized and that errors are detected and corrected. The Session Layer Protocol assumes that all errors occurring in the communication network are corrected by the protocols at the Transport Layer. For that reason, errors that may occur are those relating to the system's operation or to the application itself. These would include tape/disk errors, resource allocations, buffer space, or processing errors.

As with all OSI standards, there are two sets of standards: one for the connection-oriented case and one for the connectionless case. Within each set, there are separate standards for services and protocols. The connection-oriented Session Layer Service is defined in IS 8326/CCITT X.215. The protocol is defined in IS 8327/CCITT X.225. For the connectionless case, the service is defined in Addendum 1 to IS 8326 and the protocol is defined in IS 9548.

The Session Layer Service is different from the Transport Layer Service, depending upon the application's requirements. Where the Transport Layer has a common service and different protocols supporting that service, the Session Layer has different services based upon the application's requirements. Each aspect of the Session Layer Service is supported by unique protocol mechanisms.

There are two versions of the Session Layer Protocol: Version 1 and Version 2. The two are differentiated by the amount of user data that is permitted. Version 2 relaxes the limit on the length of user data, permitting it to be unlimited. The negotiation rules for connection establishment ensure that a Version 2 implementation will be capable of communicating with a Version 1 implementation.

Purpose of Protocol

The Session Layer Protocol exists to provide synchronization between applications. The Session Layer controls the dialog between two applications without regard for how the applications communicate. In particular, the Session Layer Protocol deals with the exchange of data between two applications so that each maintains a consistent view of the other application's view.

Functions and Features

The Session Layer Protocol provides functions that can be tailored to the needs of a specific application. The mechanisms used to tailor the functions in use are *functional units.* A functional unit is a set of protocol mechanisms that, when used together, perform one or more protocol functions.

The Session Layer Protocol supports the following:

- *Normal data transfer,* the transfer of user data between applications.
- *Token management,* where tokens are used to implement turn control. A token's owner can invoke certain services and functions, whereas the other side cannot perform the function until it has the token. An example is the data token used to enforce half-duplex operation.
- *Exception reporting,* where the protocol returns particular status and error information to the user.
- *Typed data transfer,* where the protocol can pass a limited amount of data against the data token. This is analogous to a reverse channel available in certain half-duplex operations.
- *Minor synchronization,* where the protocol supports a type of data synchronization where sync points are inserted into the datastream. With minor synchronization, data can continue to be sent while waiting for the acknowledgment of the sync point.
- *Major synchronization,* where the user must suspend data transfers until the receiver acknowledges the major sync point. Resynchronization cannot be performed to a point in the datastream before the last acknowledged major sync point.
- *Resynchronization,* where the protocol supports resynchronization, where the sender and receiver back up to a specified synchronization point. The selected point can be any synchronization point up to and including the last acknowledged major sync point.
- *Expedited data transfer,* where the protocol supports sending expedited data between users.

- *Activity management*, where the protocol supports the concept of activities that delineate a dialog between the sender and receiver of data. An activity comprises several major sync points and is often used in the Teletex service.
- *Capability data*, where the protocol supports the exchange of data concerning the capabilities of terminal equipment.

Functional Units

The Session Layer Protocol is defined as a set of functional units. Each functional unit contains protocol functions and mechanisms supporting specific Session Layer Services. The functional units are selected in a negotiation process during connection establishment.

There are 12 functional units defined:

1. Kernel
2. Negotiated release
3. Half duplex
4. Duplex
5. Expedited data
6. Typed data
7. Capability data exchange
8. Minor synchronize
9. Major synchronize
10. Resynchronize
11. Exceptions
12. Activity management

At connection establishment, the two Session users negotiate which functional units will be present during the life of the connection. The selected functional units define the Session Layer Protocol's functions for that connection.

GOSIP Considerations

The Government Open Systems Interconnection Profile (GOSIP) Version 1 specifies support for both Version 1 and Version 2 of the Session Layer Protocol. In addition, GOSIP specifies that the negotiated release and capability data functional units are not required. All other functional units are allowed. GOSIP specifies that selecting which functional units to request for a given application will be specified by the application.

Connectionless Session Layer Protocol

The connectionless Session Layer Protocol is very simple to ensure the efficiency of connectionless operation. In particular, IS 9548 provides limited procedures for address mapping and for mapping the Session Layer Protocol to the connectionless Transport Layer Protocol. The Session Layer Protocol does not provide any functions for segmenting large data requests; rather, the data is discarded and the user is informed.

Applying Protocols to Problems

The Transport and Session Layer standards are not used independently of other protocols to provide an integrated solution. These protocols form a foundation for developing distributed applications in a network environment. How do these standards relate to the rest of the problem, such as selecting a protocol suite and implementing specific applications?

Relationship to TCP/IP

The protocols defined by the DOD, TCP/IP, have gained significant acceptance as network standards. TCP/IP were designed before the wide acceptance of the OSI Reference Model and are not layered according to the same principles. For that reason, TCP/IP encompass many of the features of the Network, Transport, and Session Layers of the OSI Reference Model (see Figure 1).

The Transport Layer Protocol Class 4 was designed to be nearly equivalent to TCP. Some differences exist between the two protocols, but in most cases where TCP is used today, Class 4 can be used.

The differences between the OSI protocol suite and TCP/IP suite are more fundamental. The interface to TCP/IP is different from that of the Class 4 protocol, making it difficult to operate the OSI upper layer protocols over TCP/IP or to run the DOD protocol over the Class 4 protocol. A major issue in deciding between using TCP/IP or OSI is the possibility of creating parallel protocol stacks, and subsequent migration to a single stack.

It is possible to implement both protocol stacks, however, and to pass data between the stacks using application gateways. This approach allows users to keep using protocols and applications with which they are comfortable. It is also possible to implement an OSI environment by implementing the OSI upper layer protocols over TCP/IP. This allows users to gain familiarity with the new protocols without requiring the change-out of existing applications. Finally, it is possible to implement the full OSI protocol suite. This requires a complete change-out of protocols and applications. It requires that all network nodes successfully migrate to the new protocols on the day of the change. This approach is the most risky for downtime and network outages; however, it does create a clean break to the new protocols.

Network and Presentation Layer Issues

Network Layer Issues

The entire focus of the Transport Layer is to hide the differences of network technologies and approaches. The prime reason for the Transport Layer is to allow applications design without regard to the communication characteristics. For that reason, the Transport Layer Protocol can operate over many different networks.

The connection-oriented protocol maps its requirements onto the connection-oriented Network Layer Service. Different types of networks, such as X.25 and ISDN, provide a mapping from the network-specific interface to the Network Layer Service. This provides a very clean interface between the two layers. The actual interface implementation will be specific to each system and may require custom software to ensure that the Transport Layer Protocol interface matches the appropriate Network Layer interface.

Due to the sophistication of the Class 4 protocol, there is a mapping defined to the connectionless Network Layer Service. This mapping provides the capability to operate Class 4 over the OSI version of IP, sometimes called the Connectionless Network Protocol (CLNP). Of the connection-oriented classes, only Class 4 is permitted to operate over the connectionless Network Layer Service.

Figure 1.
OSI and Internet Protocol Stacks

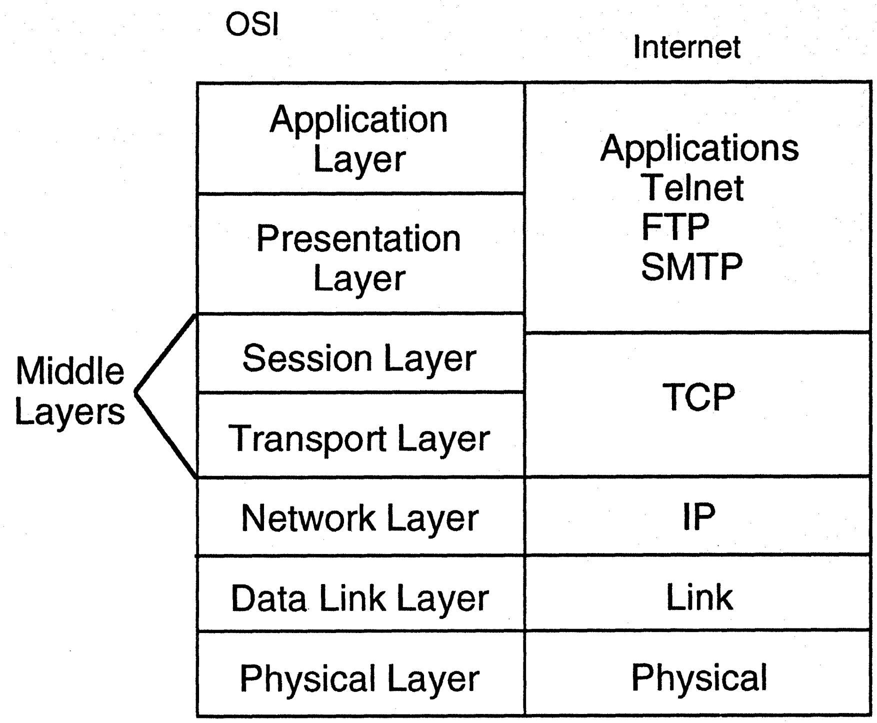

The restriction on Transport Layer connectionless protocol's functionality, and the desire for efficiency, limits the protocol's mappings to Network Layer services. It is restricted to operate only over the connectionless Network Layer Service. The connectionless Transport Layer Protocol provides only that single mapping.

Presentation Layer Issues

The Session Layer is the only OSI layer where the services substantially change based upon the application requirements. This has a major impact upon the Presentation Layer in that the services it can provide to the application depend upon the services available at the Session Layer.

The Presentation Layer standards define the precise mapping of Presentation Services offered to the application based upon the services provided after establishing a session connection. The appropriate mappings are carefully defined in the standards and in documents such as GOSIP.

The range of services offered by the Presentation Layer requires a close coupling of the Presentation Layer Protocol and the Session Layer Service. Therefore, there are few interface issues relating to mapping the Presentation Layer Protocol onto the Session Layer Service.

Specific Applications

Each of the OSI upper layer protocols is specified with mappings onto the Presentation Layer Service. As previously described, the Presentation Layer Service is tightly coupled to the Session Layer Service. In practice, mapping the upper layer protocols to the Session Layer Service is the only practical way to determine which Session Layer Services are required.

The services required of each upper layer application are tailored according to the type of interactions the applications require. While there is commonality between many of the services and applications, it is typically necessary to implement the full Session Layer Protocol.

File Transfer Access and Management

File Transfer Access and Management (FTAM) is the OSI file transfer protocol. It provides the mechanisms for defining, accessing, and managing a virtual file system in the OSI environment. FTAM defines not only the protocol for exchanging files, it also provides extensive facilities for defining the filestore structure. The filestore definition allows for complex file operations to occur on dissimilar systems.

FTAM is often used to retrieve or copy files from one system to another. The cability to move files is characterized by a large amount of data that must be reliably transferred. The transfer is basically one direction; however, controlling the transfer requires the cability to send information in the "reverse" direction. Reliability requires that the transfer be checkpointed and that recovery procedures be available upon a detected transfer error.

According to GOSIP Version 1, File Transfer Access and Management (FTAM) requires Version 2 of the Session Layer Protocol. It requires using the kernel and duplex functional units. FTAM may optionally request the use of the resynchronize and minor synchronize functional units.

At a minimum, GOSIP requires that FTAM be operated with the kernel and duplex functional units. With these selected, the transfer would not have error control and, upon detecting an error, the entire transfer would be aborted. This simple requirement is based upon the idea of requiring the minimum Session Layer services possible in order to have a useful file transfer capability. Adding the optional functional units for resynchronize and minor synchronization allows FTAM to include error recovery procedures. With that addition, detected errors require only a resynchronization to known sync points. Utilizing minor synchronization points allows the transfer to proceed while awaiting the acknowledgment, speeding the overall transfer.

X.400—Message Handling System

The X.400 protocol suite comprises the CCITT Message Handling System. X.400 provides a standard electronic mail environment and the protocols required to exchange mail messaged between distributed mail systems. The X.400 electronic mail system consists of several protocols, including protocols that only deal with exchanging mail messages between systems. Other protocols address issues such as the user interface and storing and retrieving messages.

The X.400 protocol suite is defined by CCITT. As part of the specification, CCITT specifies the complete CCITT protocol stack used for communication. That protocol stack includes the Class 0 Transport Layer Protocol and most of the Session Layer Protocol's functional units. This protocol suite is applicable to public data networks, however, or for systems exchanging information with public data networks. X.400 systems that are not used over public networks need not follow those recommendations.

According to GOSIP Version 1, X.400 requires Version 1 of the Session Layer Protocol. It requires using kernel, half-duplex, exceptions, activity management, and minor synchronize functional units. That set of functional units is equivalent to those specified by CCITT. GOSIP recognizes the need to maintain capability with the mappings of application services to Session Layer Services. This is maintained in the consistent definition of Session Layer Services. (In the case of the Transport Layer Protocol, GOSIP specifies the Class 0 protocol where interconnection with public systems is required. Otherwise, it allows the use of the Class 4 protocol.)

Virtual Terminal Protocol

The Virtual Terminal Protocol (VTP) provides a consistent terminal interface to applications across a distributed environment. The VTP provides a "virtual" terminal session that can be mapped to a variety of real terminals. VTP's scope is a simple two-dimensional terminal, such as a Digital VT100. It does not provide editing capabilities found on some complex forms-type terminals. VTP can control both the placement and look of characters (and graphics) within the viewing area.

VTP is an interactive protocol, where one user sends data and then waits for the response. A critical VTP component is to ensure that the view of the two users remains in sync. That is, both users have a consistent view of the screen information.

The Basic Class VTP requires using the following functional units:

- kernel
- half duplex
- resynchronize
- major synchronization
- typed data
- expedited data

VTP is not included in Version 1 of the GOSIP specification.

The functional units required of VTP mirror its transfer mode. To ensure that the data is written by only one user at a time, the VTP operates in a half-duplex mode. This is enforced through the half-duplex functional unit. The typed data and expedited data functional units provide a means of exchanging data outside the normal data path and permit the users to exchange status information at any time. The synchronization mechanisms are in place using the resynchronize and major synchronization functional units. The use of the major synchronization functional unit ensures that the data is acknowledged before more data is sent.

Emerging Application Protocols

As OSI continues to evolve, new protocols are being defined to take advantage of its distributed processing capabilities. Of particular interest is the work on Transaction Processing (TP). TP implements a distributed transaction model where single, atomic transactions can be processed in a distributed manner. TP implements significant mechanisms to ensure that no data modifications are made until all processing is complete. To ensure that the underlying protocols supply the needed services, TP requires that the Session Layer Protocol provide both major and minor synchronization as well as resynchronization. TP utilizes the duplex functional unit, since data flows are bidirectional. Finally, TP requires that the expedited data functional unit be implemented to ensure that synchronization is not lost due to flow control.

Practical Implementation Considerations

Since the emergence of the OSI Reference Model and the U.S. GOSIP program, most major vendors have announced plans to supply OSI-compliant products. A major issue in applying the OSI protocol suite is interoperability among vendors' products and how OSI fits into proprietary architectures.

Vendors with major network architectures, such as IBM and Digital Equipment Corp., have announced products and plans to migrate portions of their networks to OSI. IBM has a complete set of middle layer products allowing the interconnection of systems operating the OSI protocols. These products also allow users to exchange data with SNA users. Digital has announced support for OSI protocols on DECnet Phase V. These products will implement the OSI model's middle layers.

There are also vendors with "portable" products, typically UNIX based, which can be ported to a variety of systems. Products of this type allow users to select from appropriate vendors based on features, platforms, and prices.

System integration considerations concern which classes of the Transport Layer Protocol are required for the application and which of the functional units are supplied in a Session Layer Protocol. In most domestic cases, products supplied by vendors contain the Class 4 protocol. Other classes may be supplied, such as the Class 2 protocol. It seems clear that if the OSI protocols are to be used, however, then the Class 4 protocol is most appropriate.

For the Session Layer Protocol, it has become clear that no set of functional units is sufficient for all applications. For that reason, all products include all of the functional units defined in Version 2. Many applications now come bundled with the appropriate Presentation and Session Layer Protocols, thus ensuring that the application receives the necessary approach. It can also lead to duplicate protocol implementations based upon common features found in the Session Layer Protocol.

Using OSI in PC-based LANs has often been dismissed as infeasible, due to overhead and cost. Today, solutions are available that incorporate OSI into the server technology and allow the PCs and LAN operating systems to operate normally. As PCs have grown more sophisticated, OSI on PCs is a viable solution. (One can argue that if

TCP/IP-based solutions are acceptable, then OSI-based solutions will also be acceptable.) The strategy for implementing OSI on PCs or in a PC-based LAN must be tailored to the environment, many of which intentionally restrict users to the server platform. Others prefer to allow PCs to operate as peers in the LAN. An additional consideration is the protocols selected for use on the LAN. If OSI is destined to be used only for remote communication, a server solution may be more appropriate. If OSI will supply the protocols used on the LAN, then, by necessity, all PCs must use OSI.

Conclusion

The OSI middle layer protocols have reached a level of maturity and stability demonstrated in the marketplace with numerous commercial offerings. In selecting the OSI protocols over a vendor's proprietary offering or over the DOD TCP/IP suite, a user or system integrator is assuming considerably less risk than in previous years. The system integrator's role will be to incorporate vendor products into a cohesive network structure allowing for full protocol functionality but tailored to provide the required flexibility.

OSI's acceptance within both the commercial and government markets is leading to fierce competition from vendors of other solutions, which has tended to polarize the marketplace and drive vendors into providing full-functioned OSI solutions at very attractive prices. ■

This report was prepared exclusively for Datapro by James Moulton, president and principal consultant for Open Network Solutions, Inc. (ONS). Mr. Moulton has over 18 years in data communication and telecommunication research and development. He has been involved in designing and standardizing the OSI protocols and has designed networks based on DOD protocols, TCP/IP, and GOSIP protocols. He is an active participant in international and national standards committees defining the network standards of the future.

ONS, based in Sterling, VA, is a diversified company providing solutions to data communication and networking problems. It has designed major network systems, including WANs and LANs for both industry and government. ONS is heavily involved in researching the next generation of network protocols. It is experienced in both GOSIP and TCP/IP networks as well as the range of LAN operating systems.