



Chapter Goals

- Describe Class of Service, what it is, and how to use it.
- Describe how to use subareas to break up the network.
- Describe how peer-to-peer routing works.
- Describe types of PtP routing.

IBM Systems Network Architecture Routing

Background

IBM's networking architecture has evolved considerably as computing in general has evolved away from domination by centralized computing solutions to peer-based computing. Today, IBM Systems Network Architecture (SNA) routing involves two separate kinds of environments, although a number of key concepts are central to all SNA routing situations. This chapter addresses functions and services that make both SNA subarea routing and Advanced Peer-to-Peer Networking (APPN) routing possible. Topics covered include session connections, transmission groups, explicit and virtual routes, and Class of Service (COS). Refer to Chapter 37, "IBM Systems Network Architecture Protocols," for general information about traditional IBM SNA and APPN. Figure 41-1 illustrates the concepts addressed in this chapter in the context of a traditional SNA environment.

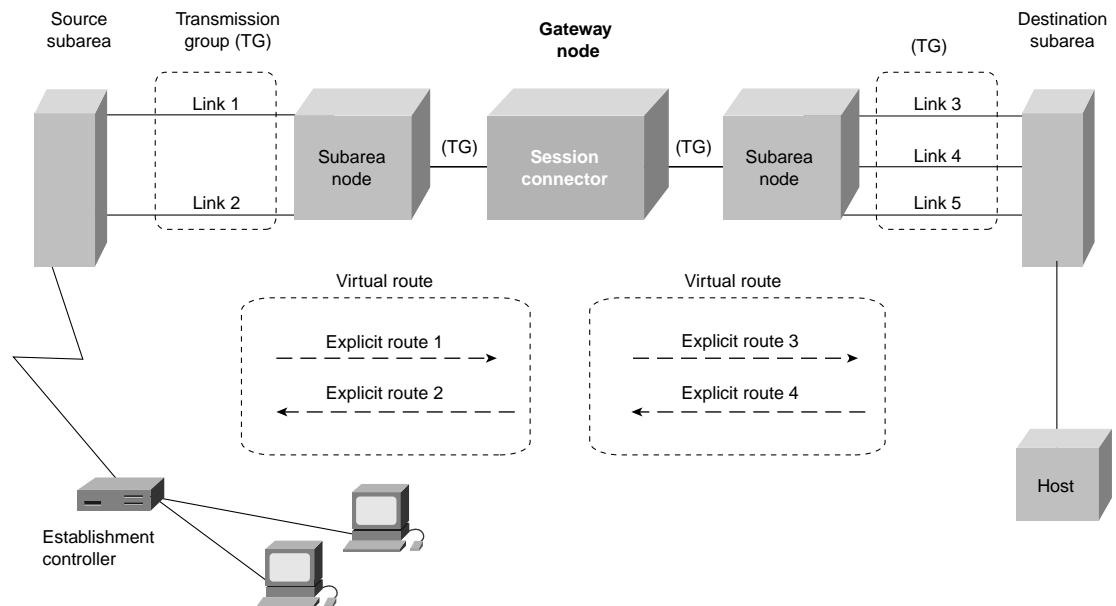
IBM SNA Session Connectors

IBM SNA session connectors are used to bridge address spaces when sessions traverse multiple address spaces. Three types of session connectors exist: boundary functions, SNA network interconnection (SNI) gateways, and APPN intermediate routing functions. Boundary functions reside in subarea nodes and map between subarea and peripheral address spaces. SNI gateways act as bridges between SNA networks, accepting data

from one network and transmitting it to the appropriate destination in another network.

SNI gateways are transparent to endpoint network attachment units (NAUs). APPN intermediate nodes perform intermediate routing within APPN networks. Refer to Figure 41-1 for the relative position of a session connector in a traditional SNA environment.

Figure 41-1 SNA Routing Relies on Transmission Groups to Interconnect Subarea Entities



IBM SNA Transmission Groups

IBM SNA transmission groups (TGs) are logical connections formed between adjacent IBM SNA nodes that are used to pass SNA session traffic. TGs are comprised of one or more SNA links and their assigned transmission priorities. Multilink TGs, which provide added reliability and bandwidth, are used to bundle multiple physical links into a single logical SNA link. Multilink TGs are supported only between T4 nodes. TG sequence numbers are used to resequence out-of-order messages at each hop. Four transmission priorities are supported at each transmission group: low, medium, high, and network-service traffic (the highest priority). Refer to Figure 41-1 for an illustration of the relationship of TGs with respect to other common SNA routing components within the context of a subarea routing environment.

IBM SNA Explicit and Virtual Routes

Routes between subareas assume either an explicit or virtual route form. Explicit routes are the physical connections between two subarea nodes, and they serve as the ordered sequences of subareas and connecting transmission groups. Explicit routes are unidirectional, and two explicit routes are required to create a full-duplex path.

Virtual routes are two-way logical connections formed between two subarea nodes. A virtual route flows over both an explicit route and a reverse explicit route that follows the same physical path. Virtual routes do not cross network boundaries; instead, they use an SNA network interconnect session connector to bridge two virtual routes. Virtual routes include values defining transmission priority and global flow control, which is provided by pacing, in which a receiver with sufficient buffer space grants pacing windows to the sender. Each pacing window enables the sender to transmit a certain amount of information before the sender must request the next pacing window. Refer to Figure 41-1 for an illustration of the relationship between explicit routes and virtual routes, and their relative position in the context of an SNA subarea routing environment.

IBM SNA Class of Service

The *IBM SNA Class of Service (COS)* function designates the transport network characteristics of a given session. Depending on user requirements, different COSs can be specified in an SNA network. COS provides the mechanism to determine all SNA routes and describes acceptable service levels for a session. COS also specifies the collection of session characteristics, including response time, security, and availability. In addition, COS can be established automatically when logging in, or manually (by the user) when the session is initiated. Each COS name is associated with a list of virtual routes that meet the desired service-level requirement. Relevant information for a given session is captured in COS subarea and APPN tables. The differences between COS implementation in subarea and APPN routing are summarized in the following sections.

COS in Subarea Routing

In subarea routing, the user defines COS support required for a particular session. Specific virtual routes are mapped to identified services, while COS characteristics are associated with the underlying explicit routes. The System Services Control Point (SSCP) uses the COS table to provide information on virtual routes and transmission priority to the path control function. Path control, in turn, selects a virtual route and transmission priority for use in a session. Figure 41-2 illustrates the subarea routing COS table-entry format.

Figure 41-2 A Subarea Routing COS Table Holds Data on Virtual Routes and Transmission Priorities

Row 1	COS name	VRN	TPRI
	Row 2	VRN	TPRI
	Row 3	VRN	TPRI

Subarea routing COS table entries include COS name, virtual route number (VRN), and subarea transmission priority (TRPI).

COS name is a standard name, such as SEC3, that is agreed upon by conventions.

The VRN identifies a specific route between subareas. Up to eight virtual route numbers can be assigned between two subarea nodes. Each virtual route can be assigned with up to three different transmission priorities, and up to 24 virtual routes are possible between two subareas.

TPRI identifies the priority of logical unit-to-logical unit (LU-to-LU) session data flowing over an explicit route. Users can select one of three priorities for each virtual route: 0 (lowest), 1, or 2 (highest).

COS in APPN Routing

COS in APPN is defined explicitly with COS table parameters. COS is more granular in APPN than subarea SNA. In particular, COS for APPN allows a route to be defined based on capacity, cost, security, propagation delay, and user-defined characteristics. It extends service to end nodes (ENs) and is not limited to communications controllers, as in subarea SNA. APPN COS permits the topology database to maintain a tree for every COS that tracks all routes and costs. APPN COS also provides a configuration option to control memory dedicated to COS trees. Figure 41-3 illustrates the APPN routing COS table-entry format.

Figure 41-3 An APPN Routing COS Table Can Include Special Character Returns and Route-Weighting Information

		← Characteristics →					
COS Name	Index	TPRI	C ₁	C ₁	→	C _n	WF
	VRN	TPRI	C ₁	C ₂	→	C _n	WF
	VRN	TPRI	C ₁	C ₂	→	C _n	WF

APPN routing COS table entries include COS name, index, APPN transmission priority (TPRI) characteristics, and APPN COS Weighted Field (WF).

The COS name is a standard name, such as SEC3, that is agreed upon by conventions.

The Index field entry enables computed weight values for route components to be stored and retrieved. This entry points to the entry in the COS weight array where the weights for the COS are stored.

APPN TPRI identifies the priority of LU-to-LU session data flowing over an explicit route. It specifies only one TPRI field for each COS table entry. APPN TPRI requires that traffic over a given session with the same COS in a particular APPN network flow with the same transmission priority.

Node and transmission group (TG) characteristics consist of a user-specified list of characteristics acceptable for an identified COS. Each row defines either a set of node characteristics or a set of TG characteristics. Entries can include security, cost per connect time, and available capacity. The field representing a characteristic contains a range of acceptable values.

The APPN COS WF enables routes-selection services (RSS) to assign a weight to a given possible route component (node or TG). WF is used by RSS to determine relative desirability of a particular route component. The WF can contain a constant or the name of a function that RSS uses in weight calculation.

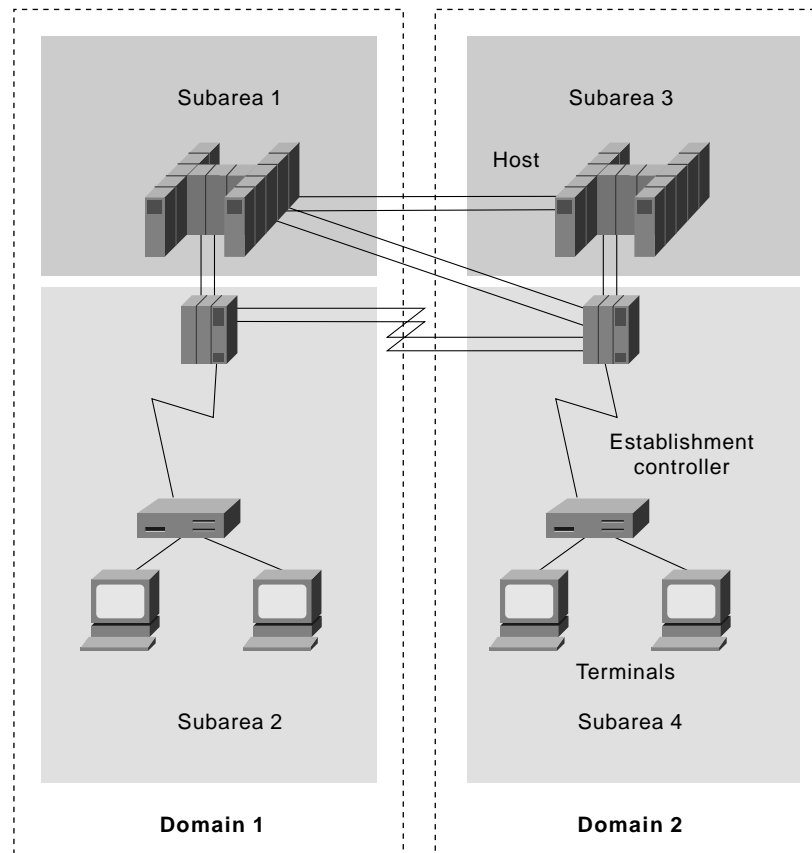
IBM SNA Subarea Routing

SNA logical areas and node addressing are two central components of traditional routing in SNA environments. This section addresses these topics in the context of traditional SNA networking.

SNA networks are divided into logical areas: subareas and domains. Subareas consist of a subarea node and its attached peripherals. Domains consist of a system services control point (SSCP) and the network resources that it can control. SSCPs in different domains can cooperate with each other to compensate for host processor failures. Figure 41-4 illustrates the relationship between subareas and domains in the context of SNA subarea routing.

Node addresses are categorized as subarea-node and peripheral-node addresses. Subarea-node addresses are global and must be unique within the entire network. These addresses are assigned to NAUs when activated. Subarea-node addresses generally consist of a subarea portion and an element portion. All NAUs within a given subarea share the same subarea address but have different element addresses.

Figure 41-4 Subareas Exist Within Domains in SNA Subarea Routing




Peripheral-node addresses, which are considered local addresses, differ depending on whether the node is T2 or T2.1. T2 addresses refer to NAUs and are statically assigned, while T2.1 addresses are dynamically assigned for the duration of a session and identify the session rather than the NAU. Peripheral-node addresses are referred to as local form-session identifiers.

IBM Advanced Peer-to-Peer Networking Routing

IBM Advanced Peer-to-Peer Networking (APPN) routing is dynamic and is based on a least-weight path calculated from input received from all APPN network nodes. Each APPN network node is responsible for reporting changes in its local topology (that is, the node itself and the attached links). Topology information is passed until all APPN nodes receive it. When a node receives data that it already has, it

stops forwarding the data to other nodes. Duplicate information is recognized via a check of update sequence numbers. Figure 41-5 illustrates where APPN network nodes fit into the general scheme of an APPN environment with ENs and low-entry network (LEN) nodes.

Figure 41-5 APPN Network Nodes Link to ENs, LENs, and Other Network Nodes



Several underlying functions and capabilities enable APPN routing. These include node type 2.1 routing, Dependent Logical-Unit Requester/Server (DLUR/S) routing, connections networks, and border nodes.

IBM APPN Node Type 2.1 Routing

Node type 2.1 routing involves routing traffic between one or more APPN network nodes. Two node type 2.1 routing processes are supported: intermediate session routing (ISR) and high-performance routing (HPR).

Intermediate Session Routing

The *ISR* process involves bind session BIND requests and responses flowing from network node to network node. In this environment, session connectors are built and used in place of routing tables in APPN. With ISR, a map is generated of the session identifier and port from one side of a node to the other. A unique session identifier in the session connector header is swapped for an outgoing identifier and then is sent out from the appropriate port.

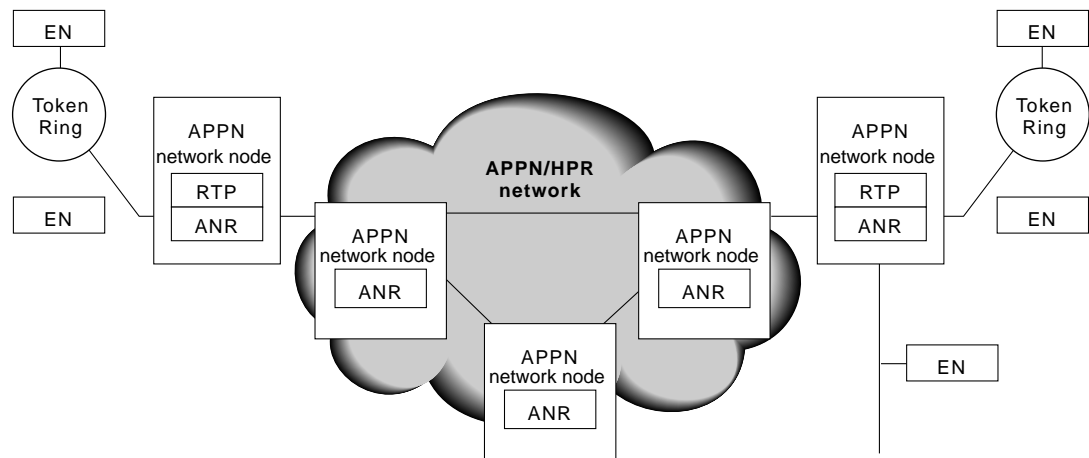
ISR-supported subarea SNA features include node-to-node error and flow-control processing, as well as session switching around network failures. Node-to-node error and flow-control processing are considered redundant and unnecessary because these processes reduce end-to-end throughput.

High-Performance Routing

The HPR protocol, an alternative to ISR, is based on two key components: Rapid-Transport Protocol (RTP) and Automatic Network Routing (ANR). RTP is a reliable, connection-oriented protocol that ensures delivery and manages end-to-end network error and flow control. RTP creates new routes following a network failure. ANR is a connectionless service that is responsible for node-to-node source-routed service.

The RTP layer is invoked only at the edges of an APPN network. In intermediate nodes, only the ANR layer is invoked. RTP nodes establish RTP connections to carry session data. All traffic for a single session flows over the same RTP-to-RTP connection and is multiplexed with traffic from other sessions using the same connection. Figure 41-6 illustrates the overall architecture of an HPR-based routing environment.

Figure 41-6 RTP Is Supported Only in APPN Edge Network Nodes



A typical HPR routing process involves several stages. First, a route is selected while it is using ISR. To establish a connection between the edge RTP nodes, either an existing RTP-to-RTP connection is used or a route-services request (RSR) is sent. The returned route-services reply (RSP) carries information showing forward and reverse paths through the network.

Paths represent the forward and reverse port lists and include the port identifier used in each ANR node. These lists are carried on every message, eliminating the need for routing tables or session connectors in the ANR nodes.

HPR provides for link-failure recovery. If a link fails and an alternate path exists between the RTP endpoints for a particular COS, a new RTP-to-RTP connection can be selected, and a session can be moved without disruption. If a connection does not exist along the new path, RSR and RSP messages are sent to obtain the new port lists. Sending a new BIND is not required because the session is not disrupted.

Flow control in an HPR environment uses a technique called adaptive rate-based (ARB) flow control. ARB flow control monitors and controls the amount of traffic introduced into the network. Under ARB flow control, the sending and receiving RTP nodes exchange messages at regular intervals. Traffic introduced into the network is modified to adapt to network conditions.

IBM APPN DLUR/S Routing

The *Dependent Logical-Unit Requester/Server (DLUR/S)* is an APPN feature that allows legacy SNA traffic to flow on an APPN network.

Under DLUR/S, a client/server relationship is established between a Dependent Logical-Unit Server (DLUS) and a Dependent Logical-Unit Requester (DLUR). A DLUS is typically an ACF/UTAM4.2 entity, and a DLUR is typically a router. A pair of LU 6.2 sessions is established between the DLUR and DLUS. These LU 6.2 sessions transport legacy SNA control messages. Those messages not recognized in an APPN environment are encapsulated on the LU 6.2 session. Messages then are de-encapsulated by DLUR and are passed to the legacy SNA LU. The DLU session initiation is then passed to the DLUS and is processed by the DLUS as legacy traffic. The DLUS sends a message to the application host, and the application host sends the BIND. Finally, legacy SNA data flows natively with APPN traffic.

IBM APPN Connection Network

An *IBM APPN connection network* is a logical construct used to provide direct connectivity between APPN ENs without the configuration overhead of defining direct connections between every pair of ENs. In general, the process of creating a connection network starts when a LOCATE request is received from an EN.

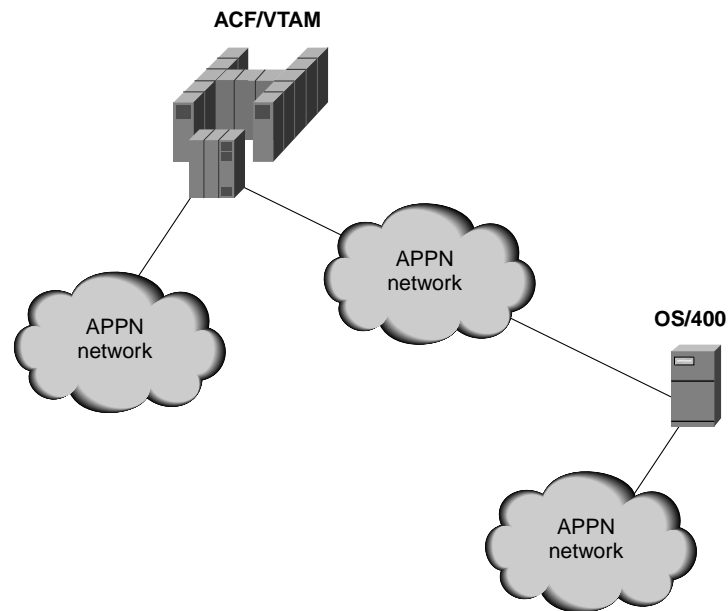
A network node (NN) then is used to locate the destination specified in the LOCATE request. If the NN sees that the two ENs (source and destination) are attached to the same transport medium (such as Token Ring), a virtual node (VN) is used to connect the two endpoints and form a connection network. The NN defines the session path as a direct connection from EN1 to VN to EN2, and then traffic is permitted to flow.

IBM APPN Border Node

A border node is an APPN entity that enables multiple APPN networks to interconnect. Presently, border nodes are implemented only in ACF/VTAM and OS/400. Border nodes are responsible for tying directories and topology databases together for connected networks, and rebuilding BIND requests to show separate routes in each network.

With border nodes, topology and directory databases in NNs are reduced to sizes required for individual subnetworks rather than the composite network. In addition, cross-network sessions are routed through the border node. Figure 41-7 illustrates the position of border nodes (ACF/VTAM and OS/400 devices) in a multinetwork APPN environment:

Figure 41-7 Border Nodes (ACF/VTAM and OS/400 Devices) Can Connect APPN Networks



Review Questions

Q—What are SNA session connectors used for?

A—IBM SNA session connectors are used to bridge address spaces when sessions traverse multiple address spaces. Three types of session connectors exist: boundary functions, SNA network interconnection (SNI) gateways, and APPN intermediate routing functions.

Q—What is created when a network node determines via the LOCATE request that the two end nodes are attached to the same medium?

A—A network node (NN) then is used to locate the destination specified in the LOCATE request. If the NN sees that the two ENs (source and destination) are attached to the same transport medium (such as Token Ring), a virtual node (VN) is used to connect the two endpoints and form a connection network.

Q—True or false: All NAUs within a subarea have the same element address.

A—False. All NAUs within a given subarea share the same subarea address but have different element addresses.

