



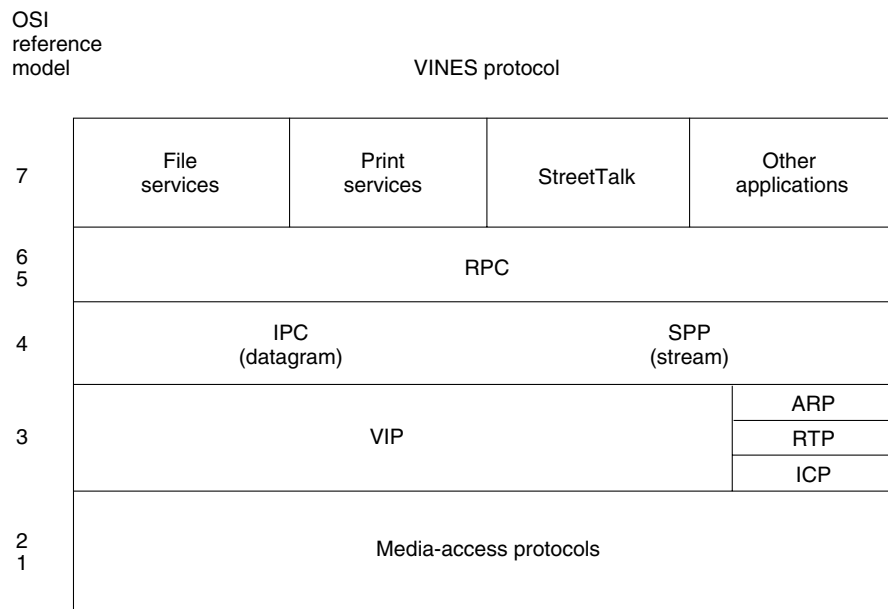
Chapter Goals

- Introduce the VINES protocol, used primarily by Banyan VINES networks.
- Describe how VINES uses each layer of the VINES Protocol Stack.

Banyan VINES

Background

Banyan Virtual Integrated Network Service (VINES) implements a distributed network operating system based on a proprietary protocol family derived from the Xerox Corporation's Xerox Network Systems (XNS) protocols. VINES uses a client/server architecture in which clients request certain services, such as file and printer access, from servers. This chapter provides a summary of VINES communications protocols. The VINES protocol stack is illustrated in Figure 36-1.

Figure 36-1 The VINES Protocol Stack Consists of Five Separate Levels

Media Access


The lower two layers of the VINES stack are implemented with a variety of well-known media-access mechanisms, including High-Level Data Link Control (HDLC), X.25, Ethernet, and Token Ring.

Network Layer

VINES uses the VINES Internetwork Protocol (VIP) to perform Layer 3 activities (including internetwork routing). VINES also supports its own Address Resolution Protocol (ARP), its own version of the Routing Information Protocol (RIP)—called the Routing Table Protocol (RTP)—and the Internet Control Protocol (ICP), which provides exception handling and special routing cost information. ARP, ICP, and RTP packets are encapsulated in a VIP header.

VINES Internetwork Protocol

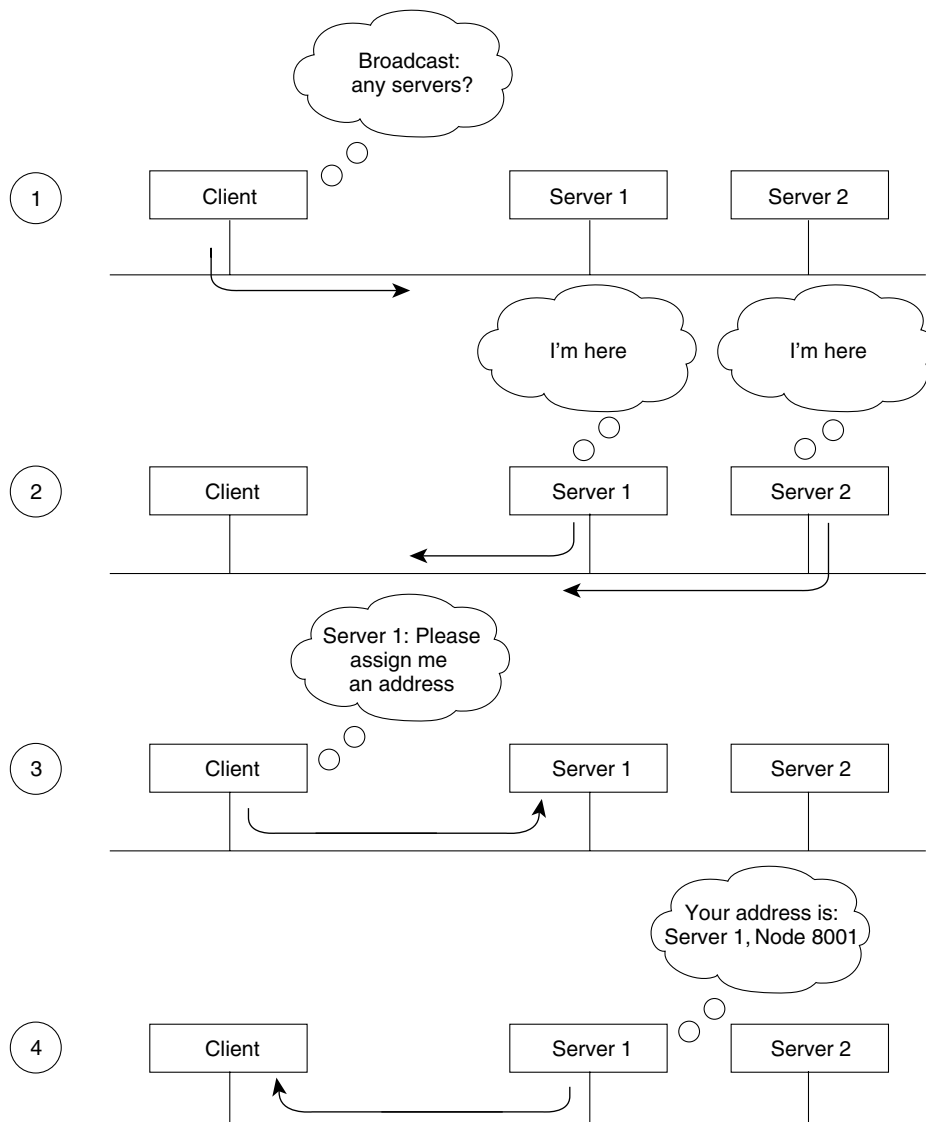
VINES network layer addresses are 48-bit entities subdivided into network (32 bits) and subnetwork (16 bits) portions. The network number is better described as a server number because it is derived directly from the server's key (a hardware module that identifies a unique number and the software options for that server). The subnetwork portion of a VINES address is better described as a host number because it is used to identify hosts on VINES networks. Figure 36-2 illustrates the VINES address format.

Figure 36-2 A VINES Address Consists of a Network Number and a Subnet Number

The network number identifies a VINES logical network, which is represented as a two-level tree with the root at a service node. Service nodes, which are usually servers, provide address resolution and routing services to clients, which represent the leaves of the tree. The service node assigns Vines Internetwork Protocol (VIP) addresses to clients.

When a client is powered on, it broadcasts a request for servers, and all servers that hear the request respond. The client chooses the first response and requests a subnetwork (host) address from that server. The server responds with an address consisting of its own network address (derived from its key) concatenated with a subnetwork (host) address of its own choosing. Client subnetwork addresses typically are assigned sequentially, starting with 8001H. Server subnetwork addresses are always 1. Figure 36-3 illustrates the VINES address-selection process.

Figure 36-3 VINES Moves Through Four Steps in Selecting an Address



Dynamic address assignment is not unique in the industry (AppleTalk also uses this process), but it certainly is not as common as static address assignment. Because addresses are chosen exclusively by a particular server (whose address is unique as a result of the hardware key), very little chance exists for duplicating an address. This is fortunate because duplicate addresses could cause potentially devastating problems for Internet Protocol (IP) and other networks.

In the VINES network scheme, all servers with multiple interfaces are essentially routers. Clients always choose their own server as a first-hop router, even if another server on the same cable provides a better route to the ultimate destination. Clients can learn about other routers by receiving redirect messages from their own server. Because clients rely on their servers for first-hop routing, VINES servers maintain routing tables to help them find remote nodes.

VINES routing tables consist of host/cost pairs, where the host corresponds to a network node that can be reached, and cost corresponds to a delay (expressed in milliseconds) to get to that node. RTP helps VINES servers find neighboring clients, servers, and routers.

Periodically, all clients advertise both their network layer and MAC-layer addresses with the equivalent of a hello packet, which indicates that the client is still operating and network-ready. The servers themselves send routing updates to other servers periodically to alert other routers to changes in node addresses and network topology.

When a VINES server receives a packet, it checks to see whether the packet is destined for another server or whether it is a broadcast. If the current server is the destination, the server handles the request appropriately. If another server is the destination, the current server either forwards the packet directly (if the server is a neighbor) or routes it to the next server in line. If the packet is a broadcast, the current server checks to see whether the packet came from the least-cost path. If not, the packet is discarded. If so, the packet is forwarded on all interfaces except the one on which it was received. This approach helps diminish the number of broadcast storms, a common problem in other network environments. Figure 36-4 illustrates the VINES routing algorithm.

Figure 36-5 illustrates the VIP packet format.

The fields of a VIP packet include information on the checksum, packet length, transport control, protocol type, destination network number, destination subnetwork number, source network number, and source subnetwork number.

The Checksum field is used to detect packet corruption. The Packet Length field indicates the length of the entire VIP packet.

Figure 36-4 The VINES Routing Algorithm Determines the Appropriate Path to a Destination

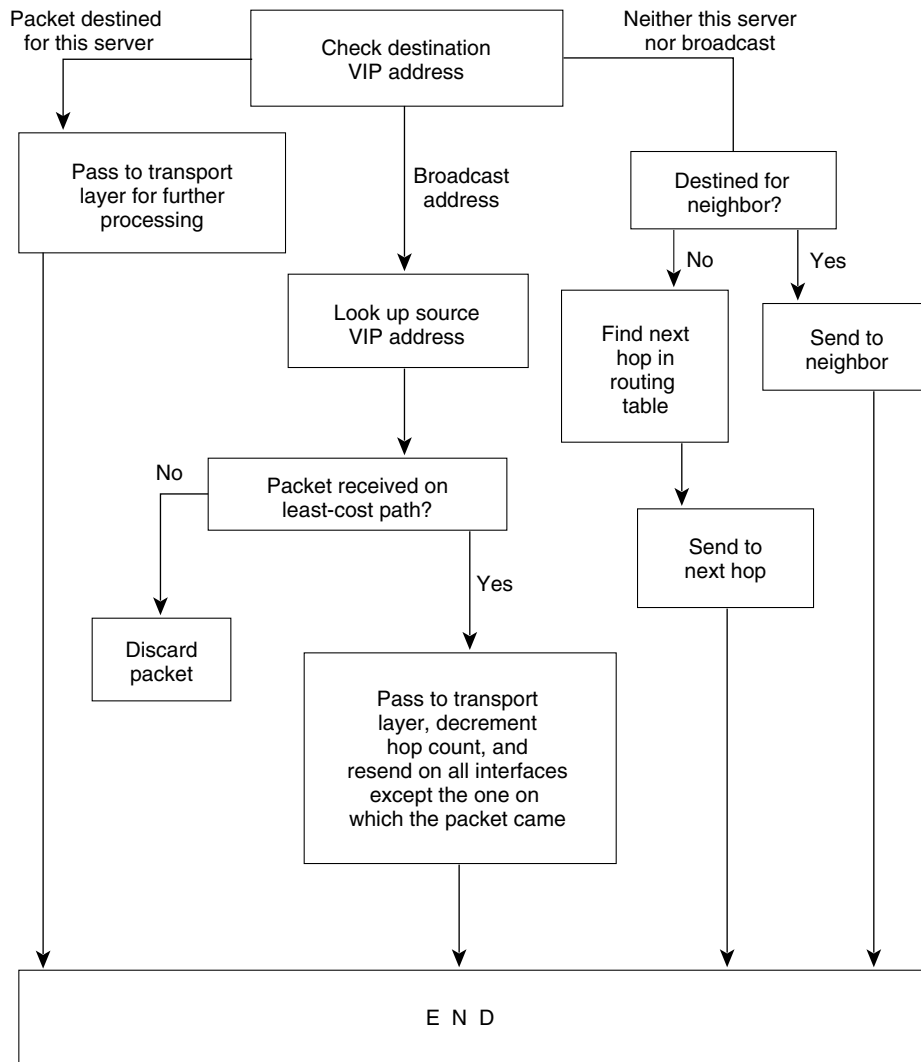


Figure 36-5 A VIP Packet Consists of Nine Individual Fields

Field length, in bytes	2	2	1	1	4	2	4	2	Variable
Check-sum	Packet length	Transport control	Protocol type	Destination network number	Destination subnetwork number	Source network number	Source subnetwork number	Data	

The Transport Control field consists of several subfields. If the packet is a broadcast packet, two subfields are provided: Class (bits 1 through 3) and Hop Count (bits 4 through 7). If the packet is not a broadcast packet, four subfields are provided: Error, Metric, Redirect, and Hop Count. The Class subfield specifies the type of node that should receive the broadcast. For this purpose, nodes are broken into various categories according to the type of node and the type of link on which the node is found. By

specifying the type of nodes to receive broadcasts, the Class subfield reduces the disruption caused by broadcasts. The Hop Count subfield represents the number of hops (router traversals) the packet has been through. The Error subfield specifies whether the ICP protocol should send an exception-notification packet to the packet's source if a packet turns out to be unroutable. The Metric subfield is set to 1 by a transport entity when it must learn the routing cost of moving packets between a service node and a neighbor. The Redirect subfield specifies whether the router should generate a redirect, when appropriate.

The Protocol Type field indicates the network layer or transport layer protocol for which the metric or exception-notification packet is destined.

Finally, the Destination Network Number, Destination Subnetwork Number, Source Network Number, and Source Subnetwork Number fields all provide VIP address information.

Routing Table Protocol

Routing Table Protocol (RTP) distributes network topology information. Routing update packets are broadcast periodically by both client and service nodes. These packets inform neighbors of a node's existence and also indicate whether the node is a client or a service node. In each routing update packet, service nodes include a list of all known networks and the cost factors associated with reaching those networks.

Two routing tables are maintained: a table of all known networks and a table of neighbors. For service nodes, the table of all known networks contains an entry for each known network except the service node's own network. Each entry contains a network number, a routing metric, and a pointer to the entry for the next hop to the network in the table of neighbors. The table of neighbors contains an entry for each neighbor service node and client node. Entries include a network number, a subnetwork number, the media-access protocol (for example, Ethernet) used to reach that node, a local-area network (LAN) address (if the medium connecting the neighbor is a LAN), and a neighbor metric.

RTP specifies four packet types: routing update, routing request, routing response, and routing redirect. A routing update is issued periodically to notify neighbors of an entity's existence. Routing requests are exchanged by entities when they must learn the network's topology quickly. Routing responses contain topological information and are used by service nodes to respond to routing-request packets. A routing-redirect packet provides better path information to nodes using inefficient paths.

RTP packets have a 4-byte header that consists of the following 1-byte fields: Operation Type, which indicates the packet type; Node Type, which indicates whether the packet came from a service node or a nonservice node; Controller Type, which indicates whether the controller in the node transmitting the RTP packet has a multibuffer controller; and Machine Type, which indicates whether the processor in the RTP sender is fast or slow.

Both the Controller Type and the Machine Type fields are used for pacing.

Address Resolution Protocol

Address Resolution Protocol (ARP) entities are classified as either address-resolution clients or address-resolution services. Address-resolution clients usually are implemented in client nodes, whereas address-resolution services typically are provided by service nodes.

ARP packets have an 8-byte header that consists of a 2-byte packet type, a 4-byte network number, and a 2-byte subnetwork number. Four packet types exist: a query request, which is a request for an ARP service; a service response, which is a response to a query request; an assignment request, which is sent

to an ARP service to request a VINES internetwork address; and an assignment response, which is sent by the ARP service as a response to the assignment request. The Network Number and Subnet Number fields have meaning only in an assignment-response packet.

ARP clients and services implement the following algorithm when a client starts up. First, the client broadcasts query-request packets. Then, each service that is a neighbor of the client responds with a service-response packet. The client then issues an assignment-request packet to the first service that responded to its query-request packet. The service responds with an assignment-response packet that contains the assigned internetwork address.

Internet Control Protocol

The *Internet Control Protocol (ICP)* defines exception-notification and metric-notification packets. Exception-notification packets provide information about network layer exceptions; metric-notification packets contain information about the final transmission used to reach a client node.

Exception notifications are sent when a VIP packet cannot be routed properly, and the Error subfield in the VIP header's Transport Control field is enabled. These packets also contain a field identifying the particular exception by its error code.

ICP entities in service nodes generate metric-notification messages when the Metric subfield in the VIP header's Transport Control field is enabled, and the destination address in the service node's packet specifies one of the service node's neighbors.

Transport Layer

VINES provides three transport layer services: unreliable datagram service, reliable message service, and data-stream service.

Unreliable datagram service sends packets that are routed on a best-effort basis but not acknowledged at the destination.

Reliable message service is a virtual circuit service that provides reliable sequenced and acknowledged delivery of messages between network nodes. A reliable message can be transmitted in a maximum of four VIP packets.

Data-stream service supports the controlled flow of data between two processes. The data-stream service is an acknowledged virtual circuit service that supports the transmission of messages of unlimited size.

Upper-Layer Protocols

As a distributed network, VINES uses the remote procedure call (RPC) model for communication between clients and servers. RPC is the foundation of distributed-service environments. The NetRPC protocol (Layers 5 and 6) provides a high-level programming language that allows access to remote services in a manner transparent to both the user and the application.

At Layer 7, VINES offers file-service and print-service applications, as well as StreetTalk, which provides a globally consistent name service for an entire internetwork.

VINES also provides an integrated applications-development environment under several operating systems, including DOS and UNIX. This development environment enables third parties to develop both clients and services that run in the VINES environment.

Summary

This may be the last book that talks about VINES as a protocol. The user community has almost completely disappeared, the server OS and VINES software are not sold anymore, and migration to TCP/IP is no longer provided.

Review Questions

Q—What does a VINES network number identify?

*A—*The network number identifies a VINES logical network, which is represented as a two-level tree with the root at a service node. Service nodes, which are usually servers, provide address resolution and routing services to clients, which represent the leaves of the tree. The service node assigns VIP addresses to clients.

Q—How does a Banyan Vines station discover servers on the network?

*A—*When a client is powered on, it broadcasts a request for servers, and all servers that hear the request respond. The client chooses the first response and requests a subnetwork (host) address from that server.

Q—Does a VINES client always choose the most efficient network path to remote networks?

*A—*No. In the VINES network scheme, all servers with multiple interfaces are essentially routers. Clients always choose their own server as a first-hop router, even if another server on the same cable provides a better route to the ultimate destination.

Q—What protocol within the VINES suite distributes routing information?

*A—*RTP distributes network topology information. Routing-update packets are broadcast periodically by both client and service nodes. These packets inform neighbors of a node's existence and also indicate whether the node is a client or a service node.

Q—What are the four packet types used by RTP?

*A—*RTP specifies four packet types: routing update, routing request, routing response, and routing redirect.

Q—What is VINES an acronym for?

*A—*Virtual Integrated Network Service.

Q—What are the media access layers of VINES?

*A—*The lower two layers of the VINES stack are implemented with a variety of well-known media access mechanisms, including High-Level Data Link Control (HDLC), X.25, Ethernet, and Token Ring.

Q—What are the two routing tables maintained by RTP?

*A—*Two routing tables are maintained: a table of all known networks and a table of neighbors. For service nodes, the table of all known networks contains an entry for each known network except the service node's own network.

Q—What are the transport layer services provided by VINES?

*A—*VINES provides three transport layer services: unreliable datagram service, reliable message service, and data-stream service.

Q—What are the Layer 7 services provided by VINES?

*A—*At Layer 7, VINES offers file-service and print-service applications, as well as StreetTalk, which provides a globally consistent name service for an entire internetwork.

For More Information

<http://products.banyan.com/>