

DECnet

Background

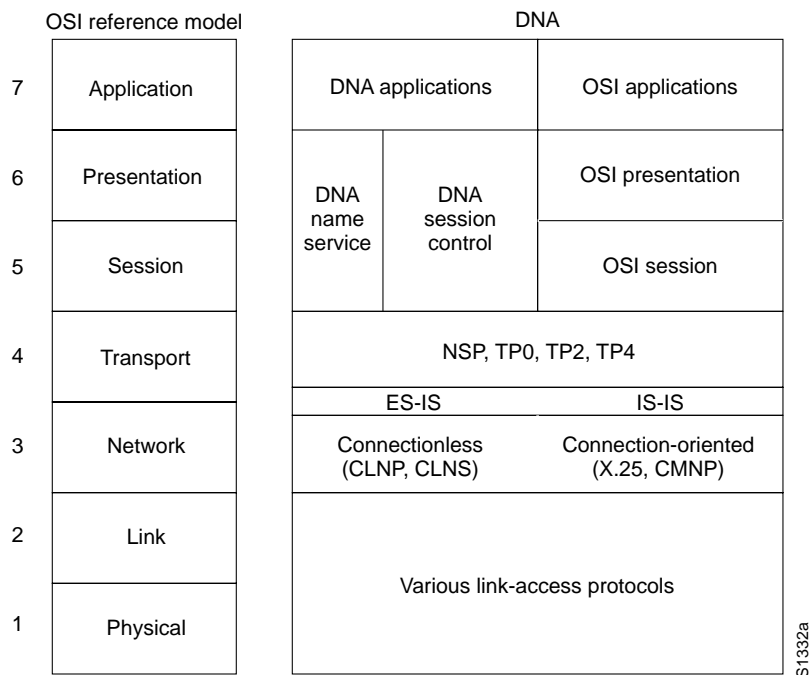
Digital Equipment Corporation (Digital) developed the *DECnet* protocol family to provide a well-thought-out way for its computers to communicate with one another. The first version of DECnet, released in 1975, allowed two directly attached PDP-11 minicomputers to communicate. In more recent years, Digital has included support for nonproprietary protocols, but DECnet remains the most important of Digital's network product offerings.

DECnet is currently in its fifth major product release (sometimes called *Phase V* and referred to as *DECnet/OSI* in Digital literature). DECnet Phase V is a superset of the OSI protocol suite and supports all OSI protocols as well as several other proprietary and standard protocols that were supported in previous versions of DECnet. As with past changes to the protocol, DECnet Phase V is compatible with the previous release (Phase IV, in this case).

Digital Network Architecture (DNA)

Contrary to popular belief, DECnet is not a network architecture at all but is, rather, a series of products conforming to Digital's *Digital Network Architecture* (DNA). Like most comprehensive network architectures from large systems vendors, DNA supports a large set of both proprietary and standard protocols. The list of DNA-supported technologies grows constantly as Digital implements new protocols. Figure 17-1 illustrates an incomplete snapshot of DNA and the relationship of some of its components to the OSI reference model.

Figure 17-1 DNA and the OSI Reference Model



Media Access

As Figure 17-1 shows, DNA supports a variety of media and link implementations. Among these are well-known standards such as *Ethernet*, *Token Ring*, *Fiber Distributed Data Interface (FDDI)*, *IEEE 802.2*, and *X.25*. For more information about these protocols, see Chapter 5, “Ethernet/IEEE 802.3,” Chapter 6, “Token Ring/IEEE 802.5,” Chapter 7, “Fiber Distributed Data Interface,” Chapter 11, “Synchronous Data Link Control and Derivatives,” and Chapter 12, “X.25.” DNA also offers a traditional point-to-point link-layer protocol called *Digital Data Communications Message Protocol (DDCMP)* and a 70-Mbps bus used in the VAXcluster called the *computer-room interconnect bus (CI bus)*.

Network Layer

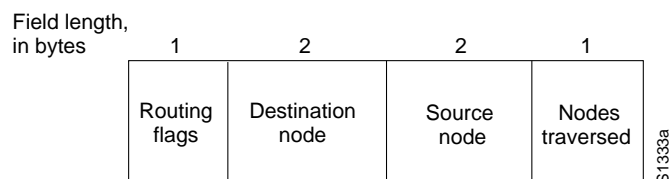
DECnet supports both connectionless and connection-oriented network layers. Both network layers are implemented by OSI protocols. The connectionless implementation uses the *Connectionless Network Protocol (CLNP)* and the *Connectionless Network Service (CLNS)*. The connection-oriented network layer uses the *X.25 Packet-Level Protocol (PLP)*, which is also known as *X.25 Level 3*, and the *Connection-Mode Network Protocol (CMNP)*. These OSI protocols are described more completely in Chapter 20, “OSI Protocols.”

Although most of DNA was brought into OSI conformance with DECnet Phase V, DECnet Phase IV routing was already very similar to OSI routing. Phase V DNA routing consists of OSI routing (*ES-IS* and *IS-IS*), plus continued support for the *DECnet Phase IV routing protocol*. *ES-IS* and *IS-IS* are described in Chapter 28, “OSI Routing.”

DECnet Phase IV Routing Frame Format

The DECnet Phase IV routing protocol differs from IS-IS in several ways. One difference is in the protocol header. The DNA Phase IV routing layer header is shown in Figure 17-2; IS-IS packet formats are shown in Chapter 28, “OSI Routing.”

Figure 17-2 DNA Phase IV Routing Layer Header



The first field in a DNA Phase IV routing header is the *routing flags* field, which includes:

- A *return-to-sender* bit that, if set, indicates that the packet is returning to the source.
- A *return-to-sender-request* bit that, if set, indicates that request packets should be returned to the source if they cannot be delivered to the destination.
- An *intraLAN* bit, which is on by default. If the router detects that the two communicating end systems are not on the same subnetwork, it turns the bit off.
- Other bits that indicate header format, whether padding is being used, and other functions.

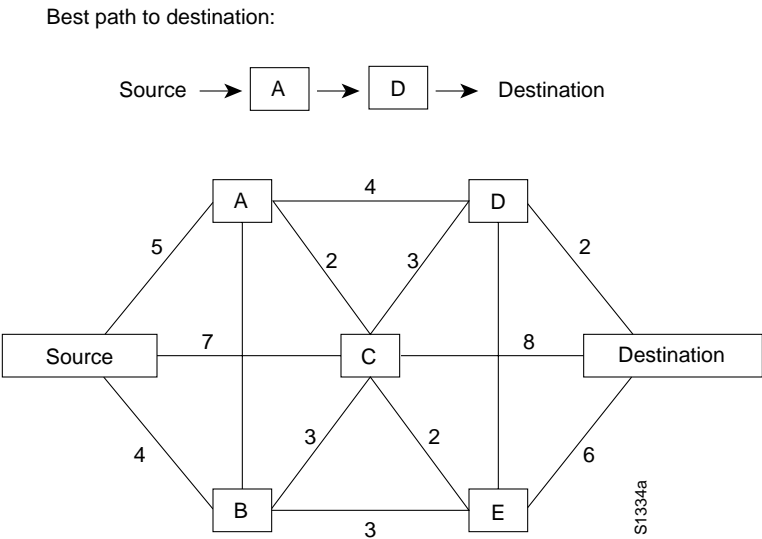
The *destination node* and *source node* fields identify the network addresses of the destination nodes and the source node.

The *nodes traversed* field shows the number of nodes the packet has traversed on its way to the destination. This field allows implementation of a maximum hop count so that obsolete packets can be removed from the network.

DECnet identifies two types of nodes, *end nodes* and *routing nodes*. Both end nodes and routing nodes can send and receive network information, but only routing nodes can provide routing services for other DECnet nodes.

DECnet routing decisions are based on *cost*, an arbitrary measure assigned by network administrators to be used in comparing various paths through an internetwork environment. Costs are typically based on hop count, media bandwidth, or other measures. The lower the cost, the better the path. When network faults occur, the DECnet Phase IV routing protocol uses cost values to recalculate the best paths to each destination. Figure 17-3 illustrates the calculation of costs in a DECnet Phase IV routing environment.

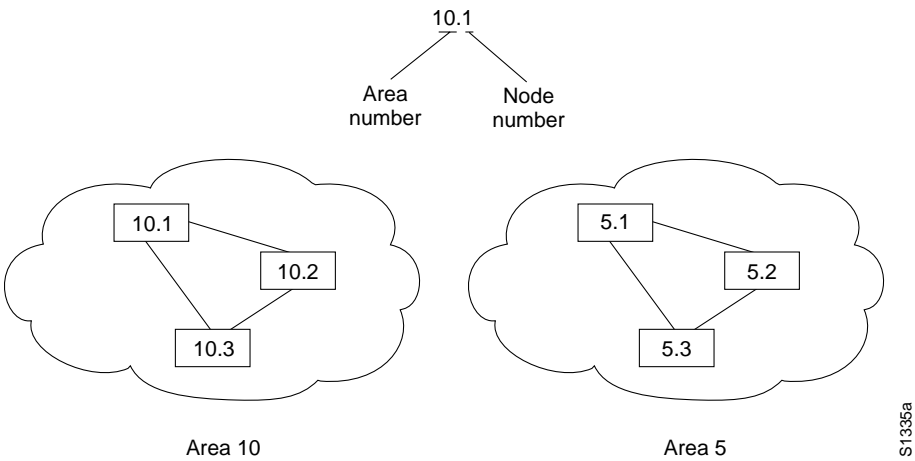
Figure 17-3 DECnet Phase IV Routing Protocol Cost Calculation



Addressing

DECnet addresses are not associated with the physical networks to which the nodes are connected. Instead, DECnet locates hosts using *area/node address* pairs. An area's value ranges from 1 to 63, inclusive. A node address can be between 1 and 1,023, inclusive. Therefore, each area can have 1,023 nodes, and approximately 65,000 nodes can be addressed in a DECnet network. Areas can span many routers, and a single cable can support many areas. Therefore, if a node has several network interfaces, it uses the same area/node address for each interface. Figure 17-4 shows a sample DECnet network with several addressable entities.

Figure 17-4 DECnet Addresses



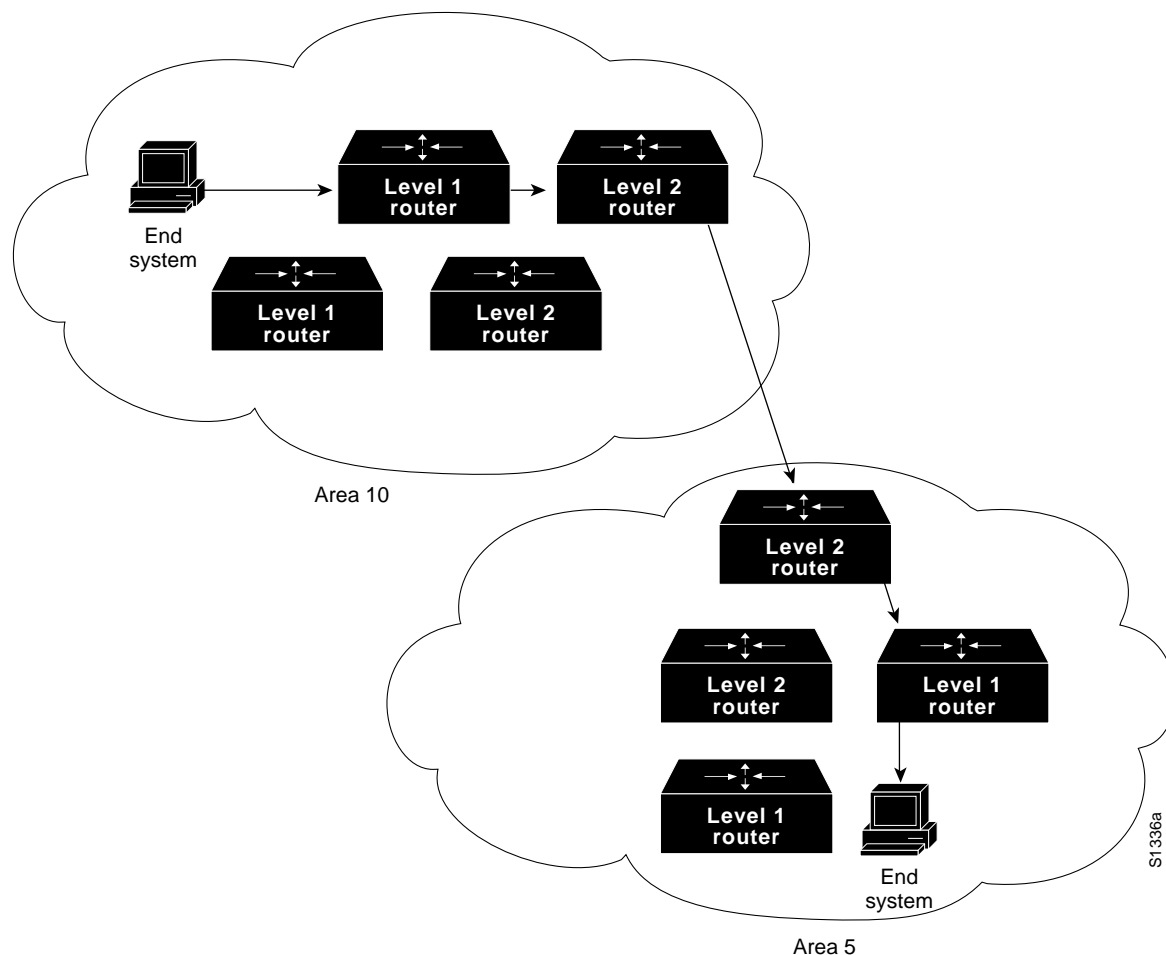
DECnet hosts do not use manufacturer-assigned *Media Access Control* (MAC)–layer addresses. Instead, network-level addresses are embedded in the MAC-layer address according to an algorithm that multiplies the area number by 1,024 and adds the node number to the product. The resulting 16-bit decimal address is converted to a hexadecimal number and appended to the address

AA00.0400 in byte-swapped order, with the least significant byte first. For example, DECnet address 12.75 becomes 12363 (base 10), which equals 304B (base 16). After this byte-swapped address is appended to the standard DECnet MAC address prefix, the resulting address is AA00.0400.4B30.

Routing Levels

DECnet routing nodes are referred to as either *Level 1* or *Level 2* routers. A Level 1 router communicates with end nodes and with other Level 1 routers in a particular area. Level 2 routers communicate with Level 1 routers in the same area and with Level 2 routers in different areas. Together, Level 1 and Level 2 routers form a hierarchical routing scheme. This relationship is illustrated in Figure 17-5.

Figure 17-5 DECnet Level 1 and Level 2 Routers



End systems send routing requests to a designated Level 1 router. The Level 1 router with the highest priority is elected to be the designated router. If two routers have the same priority, the one with the larger node number becomes the designated router. A router's priority can be manually configured to force it to become the designated router.

As shown in Figure 17-5, multiple Level 2 routers can exist in any area. When a Level 1 router wishes to send a packet outside its area, it forwards the packet to a Level 2 router in the same area. In some cases, the Level 2 router may not have the optimal path to the destination, but the mesh network configuration offers a degree of fault tolerance not provided by the simple assignment of one Level 2 router per area.

Transport Layer

The DNA transport layer is implemented by a variety of transports, both proprietary and standard. OSI transports TP0, TP2, and TP4 are supported. These are described in greater detail in Chapter 20, “OSI Protocols.”

Digital’s own *Network Services Protocol* (NSP) is functionally similar to TP4 in that it offers connection-oriented, flow-controlled service with message fragmentation and reassembly. Two subchannels are supported—one for normal data and one for expedited data and flow control information. Two flow control types are supported—a simple start/stop mechanism where the receiver tells the sender when to terminate and resume data transmission and a more complex flow control technique where the receiver tells the sender how many messages it can accept. NSP can also respond to congestion notifications from the network layer by reducing the number of outstanding messages it will tolerate.

Upper-Layer Protocols

Above the transport layer, DECnet supports its own proprietary upper-layer protocols as well as standard OSI upper-layer protocols. DECnet application protocols use the DNA session control protocol and the DNA name service. OSI application protocols are supported by OSI presentation- and session-layer implementations. For more information about these OSI protocols, see Chapter 20, “OSI Protocols.”