


*propagation software* runs on the system and continuously updates the forwarding table to ensure that each datagram follows an optimum route. That is, the software communicates with other systems to learn optimum routes to each destination, and continually checks for network failures that cause routes to change. Ironically, dynamic routing begins exactly like static routing by loading an initial set of routes into a forwarding table when the system boots.

### 26.3 Static Routing In Hosts And A Default Route

Static routing is straightforward, easy to specify, and does not require extra routing software. It does not send extra traffic on networks, and no CPU cycles are required to propagate routing information. However, static routing is relatively inflexible; it cannot accommodate network failures or changes in topology.

Where is static routing used? Most hosts use static routing, especially in cases where the host has one network connection and a single router connects the network to the rest of the Internet. For example, consider the architecture that Figure 26.1 illustrates. A set of hosts that use IPv4 are attached to an Ethernet switch, which connects to the rest of the Internet through router  $R_1$ .



Net	Mask	Next hop
128.10.0.0	255.255.0.0	direct
default	0.0.0.0	128.10.0.100

(b)

**Figure 26.1** (a) A typical connection to the Internet, and (b) the static forwarding table used by IPv4 in each host.

As the figure shows, a static forwarding table with two entries suffices for a typical host. One entry specifies the address of the directly connected network, and the other entry specifies that router  $R_1$  provides a *default route* for all other destinations. When an application generates a datagram for a computer on the local net (e.g., a local printer), the first entry in the forwarding table directs IP to deliver the datagram directly