5.8 Public and Private IP Addresses

The stability of the Internet depends directly on the uniqueness of publicly used network addresses. In Figure (5), there is an issue with the network addressing scheme. In looking at the networks, both have a network address of 198.150.11.0. The router in this illustration will not be able to forward the data packets correctly. Duplicate network IP addresses prevent the router from performing its job of best path selection. Unique addresses are required for each device on a network.

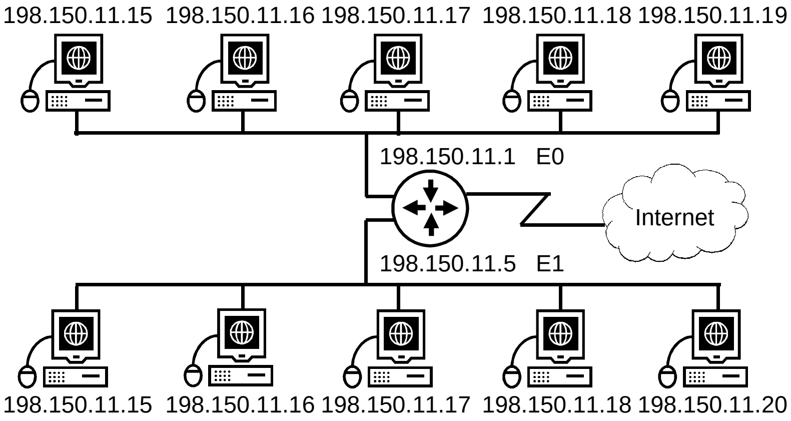

Figure (5) Required Unique Addresses

A procedure was needed to make sure that addresses were in fact unique. Originally, an organization known as the Internet Network Information Center (InterNIC) handled this procedure. InterNIC no longer exists and has been succeeded by the Internet Assigned Numbers Authority (IANA). IANA carefully manages the remaining supply of IP addresses to ensure that duplication of publicly used addresses does not occur. Duplication would cause instability in the Internet and compromise its ability to deliver datagram’s to networks.

Public IP addresses are unique. No two machines that connect to a public network can have the same IP address because public IP addresses are global and standardized. All machines connected to the Internet agree to conform to the system. Public IP addresses must be obtained from an Internet service provider (ISP) or a registry at some expense. With the rapid growth of the Internet, public IP addresses were beginning to run out.

Private IP addresses are another solution to the problem of the impending exhaustion of public IP addresses. As mentioned, public networks require hosts to have unique IP addresses. However, private networks that are not connected to the Internet may use any host addresses, as long as each host within the private network is unique. Many private networks exist alongside public networks. However, a private network using just any address is strongly discouraged because that network might eventually be connected to the Internet. RFC 1918 sets aside three blocks of IP addresses for private, internal use.

Class	RFC 1918 internal address range
A	10.0.0.0 to 10.255.255.255
B	172.16.0.0 to 172.31.255.255
C	192.168.0.0 to 192.168.255.255

Figure (6) Private IP Addresses

These three blocks consist of one Class A, a range of Class B addresses, and a range of Class C addresses. Addresses that fall within these ranges are not routed on the Internet backbone. Internet routers immediately discard private addresses. If addressing a nonpublic intranet, a test lab, or a home network, these private addresses can be used instead of globally unique addresses. Private IP addresses can be intermixed, as shown in the graphic, with public IP addresses. This will conserve the number of addresses used for internal connections.

5.9 IPv4 Versus IPv6

The TCP/IP is sustaining a global network of information, commerce, and entertainment. IP Version 4 (IPv4) offered an addressing strategy that, although scalable for a time, resulted in an inefficient allocation of addresses. Unfortunately, Class C addresses are limited to 254 usable hosts. This does not meet the needs of larger organizations that cannot acquire a Class A or B address. Even if there were more Class A, B, and C addresses, too many network addresses would cause Internet routers to come to a stop under the burden of the enormous size of routing tables required to store the routes to reach each of the network

Over the past two decades, numerous extensions to IPv4 have been developed. These extensions are specifically designed to improve the efficiency with which the 32-bit address space can be used. Two of the more important of these are subnet masks and classless interdomain routing (CIDR).

Meanwhile, an even more extendible and scalable version of IP, IP Version 6 (IPv6), has been defined and developed. IPv6 uses 128 bits rather than the 32 bits currently used in IPv4. IPv6 uses hexadecimal numbers to represent the 128 bits. IPv6 provides 640 six trillion addresses. This version of IP should provide enough addresses for future communication needs.

Internet Protocol Version 4 (Ipv4) 4 octets

11010001 1001110 1100100 01110001
0 1

209. 156. 201. 113

2^{32}=4,294,967,295 IP addresses

Internet Protocol Version 6 (Ipv4) 16 octets

10100101. 00100100	**01110010. 11010011**	**00101100.10000000**	**11011101.00000010**
A524:	**72D3:**	**2C80:**	**DD02:**
00000000.00101001	**11101100.01111010**	**00000000.00101011**	**11101010.01110011**
0029:	**EC7A:**	**002B:**	**EA73**

2^{128}=3.4x10^{38} IP addresses

Figure (7) Ipv4 and Ipv6 Addresses

Figure (7) shows an IPv4 address and an IPv6 address. IPv4 addresses are 32 bits long, written in decimal form, and separated by periods. IPv6 addresses are 128-bits long and are identifiers for individual interfaces and sets of interfaces. IPv6 addresses are assigned to interfaces, not nodes. Since each interface belongs to a single node, any of the unicast addresses assigned to the interfaces of the node may be used as an identifier for the node. IPv6 addresses are written in hexadecimal, and separated by colons. IPv6 fields are 16 bits long. To make the addresses easier to read, leading zeros can be omitted from each field. The field: 0003: is written: 3:. I Pv6 shorthand representation of the 128 bits use eight 16-bit numbers, shown as four hexadecimal digits.

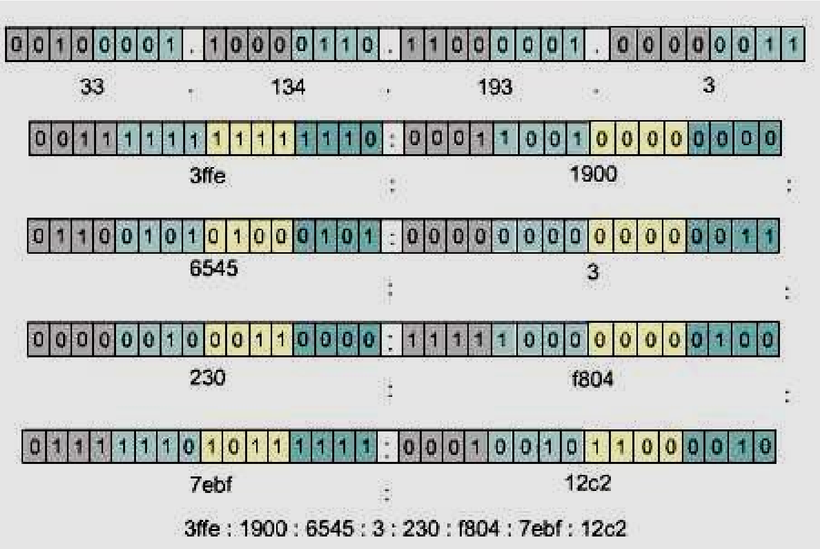

Figure (8) Ipv4 and Ipv6 Addresses

5.3 Obtaining an Internet Address

A network host needs to obtain a globally unique address in order to function on the Internet. The physical or MAC address that a host has is only locally significant, identifying the host within the local area network. Since this is a Layer 2 address, the router does not use it to forward outside the LAN.

IP addresses are the most commonly used addresses for Internet communications. This protocol is a hierarchical addressing scheme that allows individual addresses to be associated together and treated as groups. These groups of addresses allow efficient transfer of data across the Internet.

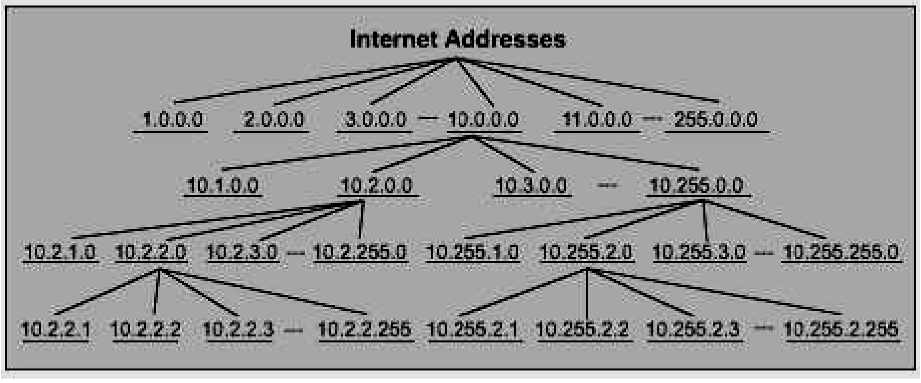

Figure (9) Internet Addresses

Network administrators use two methods to assign IP addresses. These methods are static and dynamic. Later in this lesson, static addressing and three variations of dynamic addressing will be covered. Regardless of which addressing scheme is chosen, no two interfaces can have the same IP address. Two hosts that have the same IP address could create a conflict that might cause both of the hosts involved not to operate properly. As shown in Figure (10), the hosts have a physical address by having a network interface card that allows connection to the physical medium. The figure will focus on static IP address assignments.

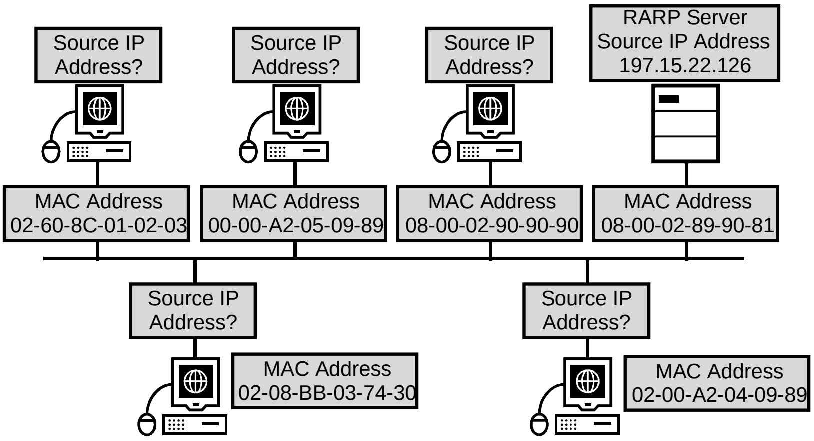
Source IP Address?
Source IP Address?
Source IP Address?
RARP Server Source IP Address 197.15.22.126
MAC Address 02-60-8C-01-02-03
MAC Address 00-00-A2-05-09-89
MAC Address 08-00-02-90-90-90
MAC Address 08-00-02-89-90-81
Source IP Address?
Source IP Address?
MAC Address 02-08-BB-03-74-30
MAC Address 02-00-A2-04-09-89

The hosts have a physical address by having a network interface card that allows connection to the physical medium. IP addresses have to be assigned to the host in some method. The two methods of IP address assignment are static or dynamic.

Figure (10) Assigning IP Addresses

5.4 Static Assignment of an IP Address

Static assignment works best on small, infrequently changing networks. The system administrator manually assigns and tracks IP addresses for each computer, printer, or server on the intranet. Good recordkeeping is critical to prevent problems which occur with duplicate IP addresses. This is possible only when there are a small number of devices to track.

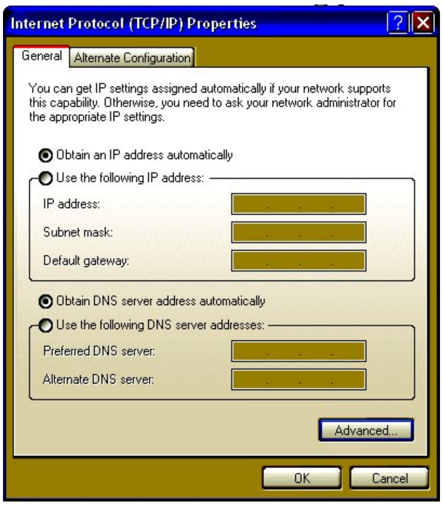

You can get IP settings assigned automatically if your network supports this capability. Otherwise, you need to ask your network administrator for the appropriate IP settings.

● Obtain an IP address automatically

○ Use the following IP address:

IP address: . . .

Subnet mask: . . .

Default gateway: . . .

● Obtain DNS server address automatically

○ Use the following DNS server addresses:

Preferred DNS server: . . .

Alternate DNS server: . . .

Figure (11) TCP/IP Configuration for Windows XP

Servers should be assigned a static IP address so workstations and other devices will always know how to access needed services. Consider how difficult it would be to phone a business that changed its phone number every day. Other devices that should be assigned static IP addresses are network printers, application servers, and routers.

5.5 RARP IP Address Assignment

Reverse Address Resolution Protocol (RARP) associates a known MAC addresses with an IP addresses. This association allows network devices to encapsulate data before sending the data out on the network. A network device, such as a diskless workstation, might know its MAC address but not its IP address. RARP allows the device to make a request to learn its IP address. Devices using RARP require that a RARP server be present on the network to answer RARP requests.

Consider an example where a source device wants to send data to another device. In this example, the source device knows its own MAC address but is unable to locate its own IP address in the ARP table. The source device must include both its MAC address and IP address in order for the destination device to retrieve data, pass it to higher layers of the OSI model, and respond to the originating device. Therefore, the source initiates a process called a RARP request. This request helps the source device detect its own IP address. RARP requests are broadcast onto the LAN and are responded to by the RARP server which is usually a router.

RARP uses the same packet format as ARP. However, in a RARP request, the MAC headers and operation code are different from an ARP request.

0-15 bits		**16-31 bits**
Hardware Type		Protocol Type
HLen (1 byte)	PLen (1 byte)	operation
Sender HA (bytes 1-4)		
Sender HA (bytes 5-6)		Sender PA (bytes 1-2)
Sender PA (bytes 3-4)		Target HA (bytes 1-2)
Target HA (bytes 3-6)		
Target PA (bytes 1-4)		
RARP Header Structure		

Figure (12) ARP/RARP Message Structure

The RARP packet format contains places for MAC addresses of both the destination and source devices. The source IP address field is empty. The broadcast goes to all devices on the network. Figures (15) & (16), and depict the destination MAC address as FF:FF:FF:FF:FF:FF. Workstations running RARP have codes in ROM that direct them to start the RARP process. A step-by-step layout of the RARP process is illustrated in Figures (14) through (21).

Field	**Description**
Hardware Type	Specifies a hardware interface type for which the sender requires a response.
Protocol Type	Specifies the type of high-level protocol address the sender has supplied.
HLen	Hardware address length.
PLen	Protocol address length.
Operation	The values are as follows: 1 ARP request. 2 ARP response. 3 RARP request. 4 RARP response. 5 Dynamic RARP request. 6 Dynamic RARP reply. 7 Dynamic RARP error. 8 InARP request. 9 InARP reply.
Sender (HA) Hardware Address	HLen bytes in length.
Sender (PA) Protocol Address	PLen bytes in length.
Target (HA) Hardware Address	HLen bytes in length.
Target (PA) Protocol Address	PLen bytes in length.

Figure (13) ARP/RARP Message Structure Field Descriptions

Computer FE:ED:F9:23:44:EF needs to get its IP address for internal operation.

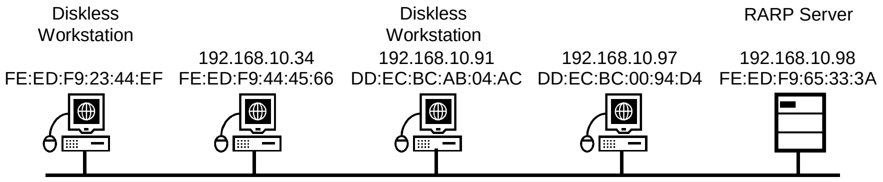

Figure (14) RARP: Network Segment

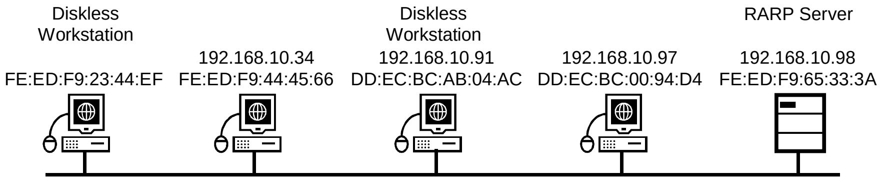

Computer FE:ED:F9:23:44:EF generates a RARP request.

Frame header	1		0800_{16}
Source MAC	48	32	3
FE:ED:F9:23:44:EF	FE:ED:F9:23:		
Destination MAC	44:EF		Undefined
FF: FF: FF: FF: FF: FF	Undefined		FF:FF:
Field Type	FF: FF: FF: FF		
0X8035 (Ethernet)	Undefined		

Figure (15) RARP: Request Generation

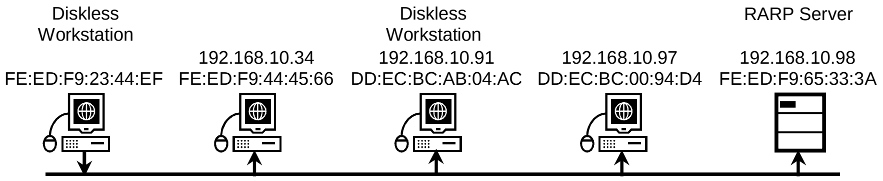

Computer FE:ED:F9:23:44:EF transmits RARP request.

Frame header	1		0800_{16}
Source MAC	48	32	3
FE:ED:F9:23:44:EF	FE:ED:F9:23:		
Destination MAC	44:EF		Undefined
FF: FF: FF: FF: FF: FF	Undefined		FF:FF:
Field Type	FF: FF: FF: FF		
0X8035 (Ethernet)	Undefined		

Figure (16) RARP: Request Transmission

All computers pass the packet up to the network layer. If IP numbers do not match, the packet is discarded except for the RARP server, which detects the RARP request field.

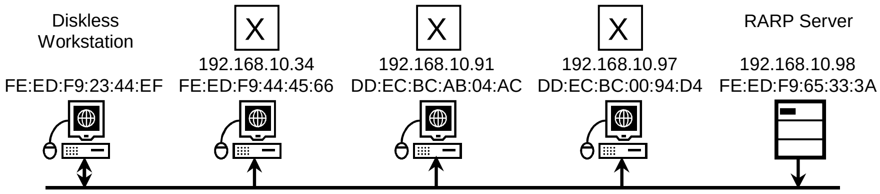

Frame header	1		0800_{16}
Source MAC	48	32	3
FE:ED:F9:23:44:EF	FE:ED:F9:23:		
Destination MAC	44:EF		Undefined
FF: FF: FF: FF: FF: FF	Undefined		FF:FF:
Field Type	FF: FF: FF: FF		
0X8035 (Ethernet)	Undefined		

Figure (17) RARP: Request Verification

The RARP server creates a RARP reply message for the requesting client.

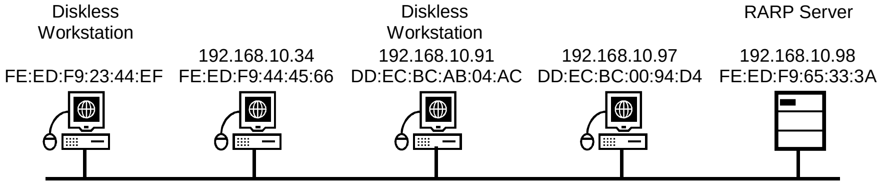

Frame header	2		0800_{16}
Source MAC	48	32	4
FE:ED:F9:65:33:3A	FE:ED:F9:23:		
Destination MAC	44:EF		192.168.
FE: ED: F9: 23: 44: EF	10.36		FE:ED:
Field Type	F9: 65: 33: 3A		
0X8035 (Ethernet)	192.168. 10.98		

Figure (18) RARP: Reply Generation

All computers copy the frame and examine it.

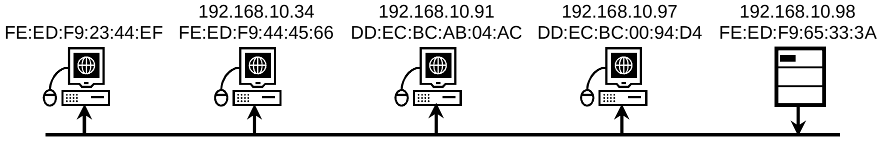

Frame header	2		0800_{16}
Source MAC	48	32	4
FE:ED:F9:65:33:3A	FE:ED:F9:23:		
Destination MAC	44:EF		192.168.
FE: ED: F9: 23: 44: EF	10.36		FE:ED:
Field Type	F9: 65: 33: 3A		
0X8035 (Ethernet)	192.168. 10.98		

Figure (19) RARP: Reply Transmission

If MAC addresses do not match, the packet is discarded.

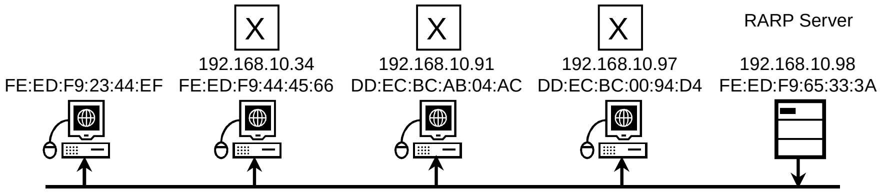

Frame header	1		0800_{16}
Source MAC	48	32	4
FE:ED:F9:65:33:3A	FE:ED:F9:23:		
Destination MAC	44:EF		192.168.
FE: ED: F9: 23: 44: EF	10.36		FE:ED:
Field Type	F9: 65: 33: 3A		
0X8035 (Ethernet)	192.168. 10.98		

Figure (20) RARP: Reply Evaluation

Computer FE:ED:F9:23:44:EF stores the IP address received in the RARP reply for later use.

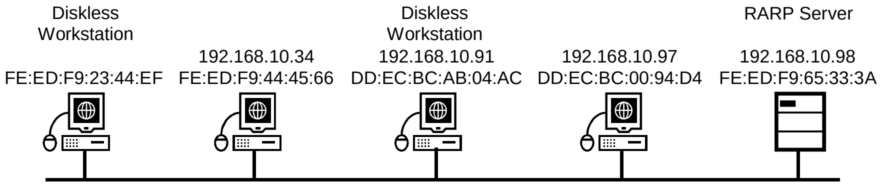

Figure (21) RARP: Data Storage

5.6 BOOTP IP Address Assignment

The bootstrap protocol (BOOTP) operates in a client-server environment and only requires a single packet exchange to obtain IP information. However, unlike RARP, BOOTP packets can include the IP address, as well as the address of a router, the address of a server, and vendor-specific information.

0-7 bits	8-15 bits	16-23 bits	24-31 bits
Op (1)	Htype (1)	HLen (1)	Hops (1)
Xid (4 bytes)			
Seconds (2 bytes)		Unused	
Ciaddr (4 bytes)			
Yiaddr (4 bytes)			
Siaddr (4 bytes)			
Giaddr (4 bytes)			
Chaddr (16 bytes)			
Server Host Name (64 bytes)			
Boot File Name (128 bytes)			
Vendor Specific Area (64 bytes)			
BOOTP Message Structure			

Figure (22) BOOTP Message Structure

One problem with BOOTP, however, is that it was not designed to provide dynamic address assignment. With BOOTP, a network administrator creates a configuration file that specifies the parameters for each device. The administrator must add hosts and maintain the BOOTP database. Even though the addresses are dynamically assigned, there is still a one to one relationship between the number of IP addresses and the number of hosts. This means that for every host on the network there must be a BOOTP profile with an IP address assignment in it. No two profiles can have the same IP address. Those profiles might be used at the same time and that would mean that two hosts have the same IP address.

Field	**Description**
Op	Message operation code. Messages can be either BOOTREQUEST or BOOTREPLY.
Htype	Hardware address type.
HLen	Hardware address length.
Hops	Client places zero, this field is used by BOOTP server to send request to another network.
Xid	Transaction ID.
Secs	Seconds elapsed since the client began the address acquisition or renewal process.
Ciaddr	Client IP address.
Yiaddr	“Your” (client) IP address.
Siaddr	IP address of the next server to use in bootstrap.
Giaddr	Relay agent IP address used in booting via a relay agent.
Chaddr	Client hardware address.
Server Host Name	Specifies particular server to get BOOTP information from.
Boot File Name	Allows for multiple boot files to be used allowing hosts to run different operating systems.
Vendor Specific Area	Contains optional vendor specific information that can be passed to the host.

Figure (23) BOOTP Message Structure Field Descriptions

A device uses BOOTP to obtain an IP address when starting up. BOOTP uses UDP to carry messages. The UDP message is encapsulated in an IP packet. A computer uses BOOTP to send a broadcast IP packet using a destination IP address of all 1s, 255.255.255.255 in dotted decimal notation. A BOOTP server receives the broadcast and then sends back a broadcast. The client receives a frame and checks the MAC address. If the client finds its own MAC address in the destination address field and a broadcast in the IP destination field, it takes and stores the IP address and other information supplied in the BOOTP reply message. A step-by-step description of the process is shown in Figures (24) through (31).

Computer FE:ED:F9:23:44:EF needs to obtain its IP address for Internet and Internet operation.

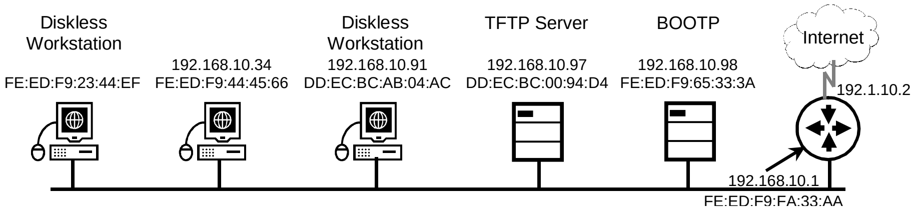

Figure (24) BOOTP: Network Segment

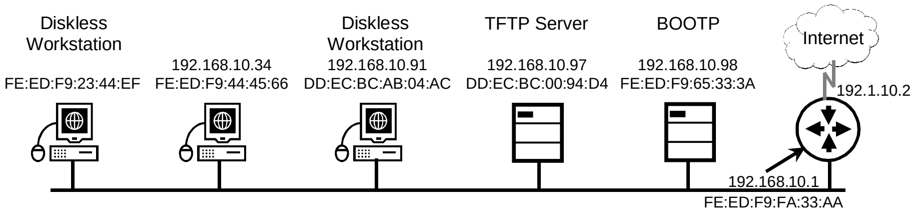

Workstation FE:ED:F9:23:44:EF generates a BOOTP request.

Frame header	Packet Header	1	1	6	0	CRC
Source MAC	Source IP	221				Check
FE:ED:F9:23:44:EF	Unknown	2		Unused		:
Destination MAC	Destination IP	0				
FF: FF: FF: FF: FF: FF	225. 225. 225. 225	0				
Field Type		0				
0X8035 (Ethernet)		0				
		FE:ED:F9:23:44:EF				

Figure (25) BOOTP: Request Creation

Workstation FE:ED:F9:23:44:EF encapsulates the request in a packet header. The header contains an unknown source IP address and a broadcast destination IP address. For the frame header the workstation uses its MAC address as the source and a broadcast for the destination as it does not know the address of the BOOTP server. The workstation then transmits a BOOTP request frame.

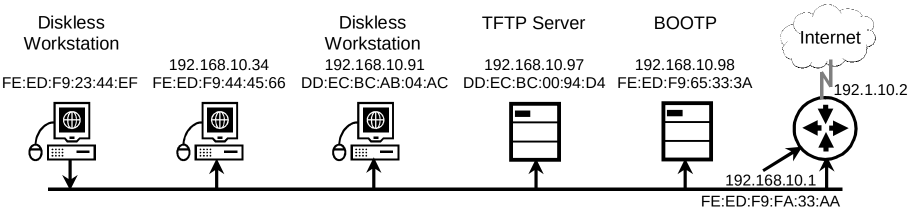

Frame header	Packet Header	1	1	6	0	CRC
Source MAC	Source IP	221				Check
FE:ED:F9:23:44:EF	Unknown	2		Unused		:
Destination MAC	Destination IP	0				
FF: FF: FF: FF: FF: FF	225. 225. 225. 225	0				
Field Type		0				
0X8035 (Ethernet)		0				
		FE:ED:F9:23:44:EF				

Figure (26) BOOTP: Request Transmission

All devices pick up a copy of the frame, detect a broadcast MAC destination, strip off the frame header, and pass the packet up to the network layer. The devices detect that the IP destination is a broadcast IP address, strip off the packet header, and pass the reply data to the transport layer. All of the devices detect the BOOTP request field as being a BOOTP request. All devices except for the BOOTP server discard it.

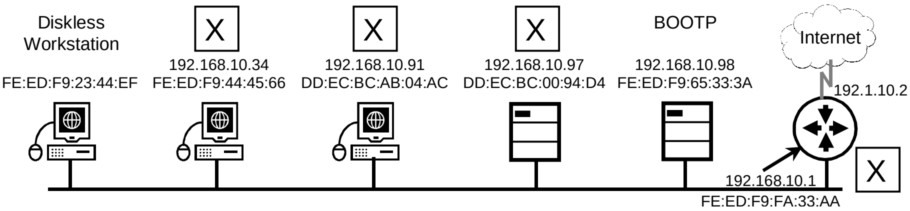

Frame header	Packet Header	1	1	6	0	CRC
Source MAC	Source IP	221				Check
FE:ED:F9:23:44:EF	Unknown	2		Unused		:
Destination MAC	Destination IP	0				
FF: FF: FF: FF: FF: FF	225. 225. 225. 225	0				
Field Type		0				
0X8035 (Ethernet)		0				
		FE:ED:F9:23:44:EF				

Figure (27) BOOTP: Request Verification

The server prepares a BOOTP response from its database to send back to the requesting device. This includes client IP address. TFTP server address, and default Gateway address (other fields are omitted for this example). In the frame header, source and destination addresses are reversed. In the packet header, the BOOTP server places its IP address in the source field and a broadcast address in the destination field. This is done to get the BOOTP response packet back up to the transport layer to be processed. Only a broadcast will be passed since the client still does not know its IP address.

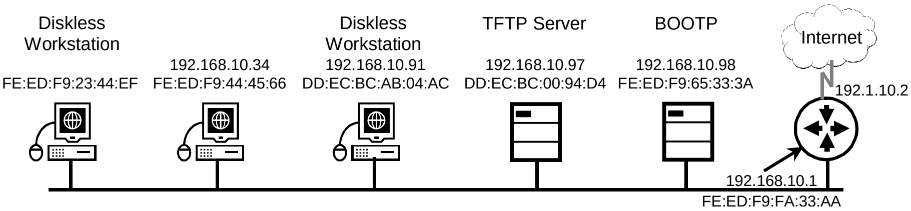

Frame header	Packet Header	2	1	6	0	CRC
Source MAC	Source IP	221				Check
FE:ED:F9:65:33:3A	192.168.10.98	2		Unused		:
Destination MAC	Destination IP	0				
FE: ED: F9: 23: 44: EF	225. 225. 225. 225	192.168.10.36				
Field Type		192.168.10.97				
0X8035 (Ethernet)		192.168.10.97				
		FE:ED:F9:23:44:EF				

Figure (28) BOOTP: Reply Creation

The BOOTP server then sends the BOOTP reply frame back to the requesting device. All devices pick up the packet and examine it.

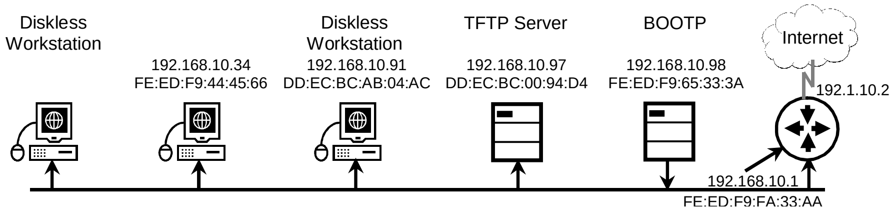

Frame header	Packet Header	2	1	6	0	CRC
Source MAC	Source IP	221				Check
FE:ED:F9:65:33:3A	192.168.10.98	2		Unused		:
Destination MAC	Destination IP	0				
FE: ED: F9: 23: 44: EF	225. 225. 225. 225	192.168.10.36				
Field Type		192.168.10.97				
0X8035 (Ethernet)		192.168.10.97				
		FE:ED:F9:23:44:EF				

Figure (29) BOOTP: Reply Transmission

The destination MAC address is not theirs and not a broadcast, so they discard the packet. The MAC address is matched on the requesting client device, so the source IP and MAC address of the BOOTP server are stored in the ARP table of the diskless workstation. The frame header is stripped off and discarded.

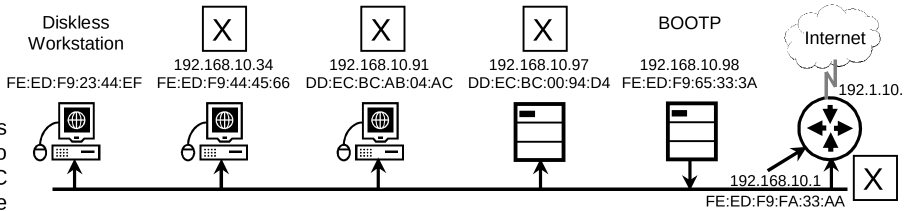

Frame header	Packet Header	2	1	6	0	CRC
Source MAC	Source IP	221				Check
FE:ED:F9:65:33:3A	192.168.10.98	2		Unused		:
Destination MAC	Destination IP	0				
FE: ED: F9: 23: 44: EF	225. 225. 225. 225	192.168.10.36				
Field Type		192.168.10.97				
0X8035 (Ethernet)		192.168.10.97				
		FE:ED:F9:23:44:EF				

Figure (30) BOOTP: Reply Verified

The packet destination IP is a broadcast, so the packet header is stripped off and the BOOTP reply data is passed up to the transport layer, where the OP field data says that this is a BOOTP reply. The reply data is stored in the appropriate memory locations in the workstation. The workstation now has access to the TFTP server for further operating system downloads and to the default Gateway as well as having its own IP address. It can now fully function on the network and the Internet.

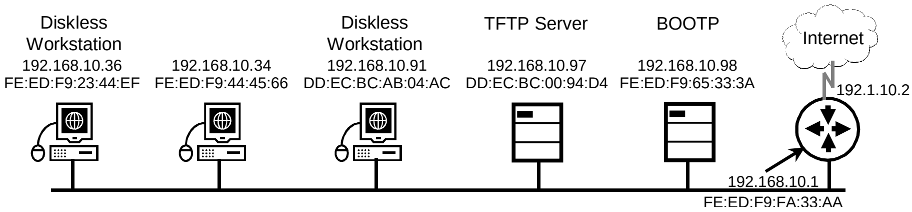

Frame header	Packet Header	2	1	6	0	CRC
Source MAC	Source IP	221				Check
FE:ED:F9:65:33:3A	192.168.10.98	2		Unused		:
Destination MAC	Destination IP	0				
FE: ED: F9: 23: 44: EF	225. 225. 225. 225	192.168.10.36				
Field Type		192.168.10.97				
0X8035 (Ethernet)		192.168.10.97				
		FE:ED:F9:23:44:EF				

Figure (31) BOOTP: Data Storage