

Broadcasts in Switched LAN Internetworks

To communicate with all or part of the network, protocols use broadcast and multicast datagrams at Layer 2 of the Open Systems Interconnection (OSI) model. When a node needs to communicate with all of the network, it sends a datagram to MAC address 0xFFFFFFFF (a *broadcast*), which is an address to which the network interface card (NIC) of every host must respond. When a host needs to communicate with part of the network, it sends a datagram to address 0xFFFFFFFF with the leading bit of the vendor ID set to 1 (a *multicast*). Most NICs with that vendor ID respond to a multicast by processing the multicast to its group address.

Because switches work like bridges, they must flood all broadcast and multicast traffic. The accumulation of broadcast and multicast traffic from each device in the network is referred to as *broadcast radiation*.

Because the NIC must interrupt the CPU to process each broadcast or multicast, broadcast radiation affects the performance of hosts in the network. Most often, the host does not benefit from processing the broadcast or multicast—that is, the host is not the destination being sought, it doesn't care about the service that is being advertised, or it already knows about the service. High levels of broadcast radiation can noticeably degrade host performance.

The following sections describe how the desktop protocols—IP, Novell, and AppleTalk—use broadcast and multicast packets to locate hosts and advertise services, and how broadcast and multicast traffic affects the CPU performance of hosts on the network.

IP Networks

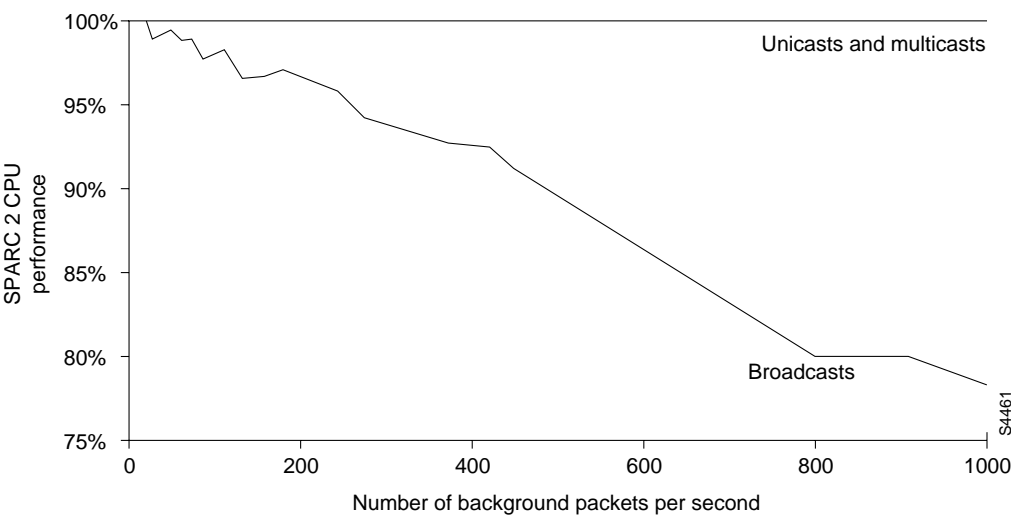
There are three sources of broadcasts and multicasts in IP networks:

- *Workstations*—An IP workstation broadcasts an Address Resolution Protocol (ARP) request every time it needs to locate a new IP address on the network. For example, the command **telnet mumble.com** translates into an IP address through a Domain Name System (DNS) search, then an ARP request is broadcast to find the actual station. Generally, IP workstations cache 10 to 100 addresses for about 2 hours. The ARP rate for a typical workstation might be about 50 addresses every 2 hours or 0.007 ARPs per second. Thus, 2000 IP end stations would produce about 14 ARPs per second.
- *Routers*—An IP router is any router or workstation that runs RIP. Some administrators configure all workstations to run RIP as a redundancy and reachability policy. Every 30 seconds, RIP uses broadcasts to retransmit the entire RIP routing table to other RIP routers. If 2000 workstations were configured to run RIP, and if 50 packets were required to retransmit the routing table, the workstations would generate 3333 broadcasts per second. Most network administrators configure a small number of routers, usually five to ten, to run RIP. For a routing table that requires 50 packets to hold it, 10 RIP routers would generate about 16 broadcasts per second.

- *Multicast applications*—IP multicast applications can adversely affect the performance of large, scaled switched networks. While multicasting is an efficient way to send a stream of multimedia (video data) to many users on a shared-media hub, it affects every user on a flat switched network. A particular packet video application can generate a 7-megabyte (MB) stream of multicast data that, in a switched network, would be sent to every segment, resulting in severe congestion.

Figure E-1 shows the results of tests that Cisco conducted on the effect of broadcast radiation on a Sun SPARCstation 2 with a standard built-in Ethernet card. The SPARCstation was running SunOS version 4.1.3 without IP multicast enabled. If IP multicast had been enabled, for example, by running Solaris 2.x, multicast packets would have affected CPU performance.

Figure E-1 Effect of Broadcast Radiation on Hosts in IP Networks



As indicated by the results shown in Figure E-1, an IP workstation can be effectively shut down by broadcasts flooding the network. While extreme, broadcast peaks of thousands of broadcasts per second have been observed during “broadcast storms.” Testing in a controlled environment with a range of broadcasts and multicasts on the network shows measurable system degradation with as few as 100 broadcasts or multicasts per second.

Table E-1 shows the average and peak number of broadcasts and multicasts for IP networks ranging from 100 to 10,000 hosts per network.

Table E-1 Average Number Broadcasts and Multicasts for IP Networks

Number of Hosts	Average Percentage of CPU Loss per Host
100	.14
1000	.96
10,000	9.15

Although the numbers in Table E-1 might appear low, they represent an average, well-designed IP network that is not running RIP. When broadcast and multicast traffic peak due to “storm” behavior, peak CPU loss can be orders of magnitude greater than average. Broadcast “storms” can be caused

by a device requesting information from a network that has grown too large. So many responses are sent to the original request that the device cannot process them, or the first request triggers similar requests from other devices that effectively block normal traffic flow on the network.

Novell Networks

Many PC-based LANs use Novell's Network Operating System (NOS) and NetWare servers. Novell technology poses the following unique scaling problems:

- NetWare servers use broadcast packets to identify themselves and to advertise their services and routes to other networks.
- NetWare clients use broadcasts to find NetWare servers.
- Version 4.0 of Novell's SNMP-based network management applications, such as NetExplorer, periodically broadcast packets to discover changes in the network.

An idle network with a single server with one shared volume and no print services generates one broadcast packet every 4 seconds. A large LAN with high-end servers might have up to 150 users per PC server. If the LAN has 900 users with a reasonably even distribution, it would have six or seven servers. In an idle state with multiple shared volumes and printers, this might average out to 4 broadcasts per second uniformly distributed. In a busy network with route and service requests made frequently, the rate would peak at 15 to 20 broadcasts per second.

Figure E-2 shows the results of tests that Cisco conducted on the effect of broadcast radiation on the performance of an 80386 CPU running at 25 MHz. Performance was measured with the Norton Utilities System Information utility. Background traffic was generated with a Network General Sniffer and consisted of a broadcast destination packet and a multicast destination packet, with data of all zeroes. CPU performance was measurably affected by as few as 30 broadcast or multicast packets per second. Multicast packets had a slightly worse effect than broadcast packets.

Figure E-2 Effect of Broadcast Radiation on Hosts in Novell Networks

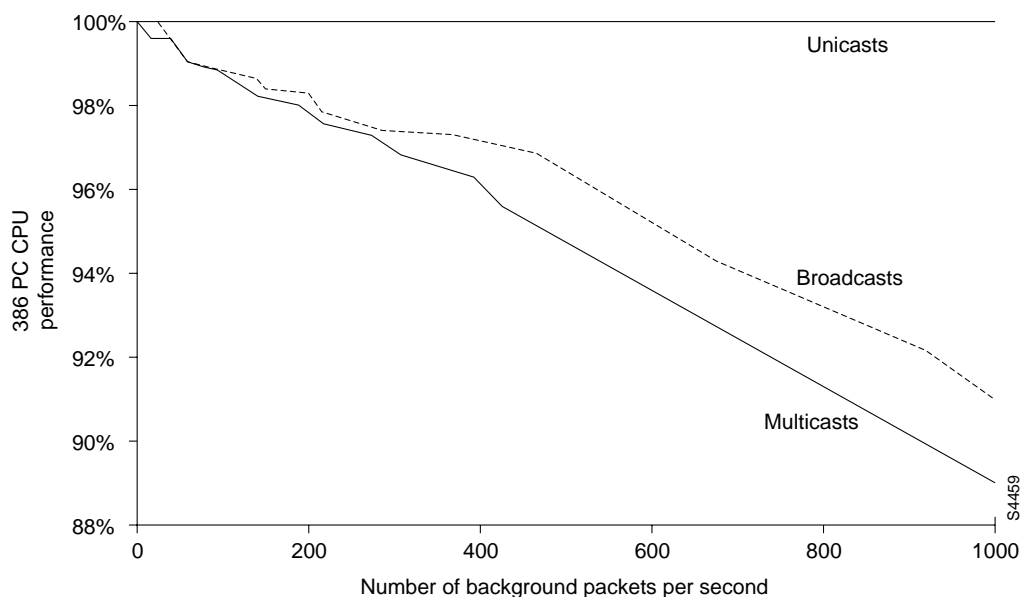


Table E-2 shows the average and peak number of broadcasts and multicasts for Novell networks ranging from 100 to 10,000 hosts per network.

Table E-2 Average Number Broadcasts and Multicasts for Novell Networks

Number of Hosts	Average Percentage of CPU Loss per Host
100	.12
1000	.22
10,000	3.15

The results listed in Table E-2 represent multihour, average operation. Peak traffic load and CPU loss per workstation can be orders of magnitude greater than with average traffic loads. A common scenario is that at 9 a.m. on Monday, everyone starts their computers. Normally, in circumstances with an average level of utilization or demand, the network can handle a reasonable number of stations. However, in circumstances in which everyone requires service at once (a demand peak), the available network capacity can support a much lower number of stations. In determining network capacity requirements, peak demand levels and duration can be more important than average serviceability requirements.

AppleTalk Networks

AppleTalk uses multicasting extensively to advertise services, request services, and resolve addresses. On startup, an AppleTalk host transmits a series of at least 20 packets aimed at resolving its network address (a Layer 3 AppleTalk node number) and obtaining local “zone” information. Except for the first packet, which is addressed to itself, these functions are resolved through AppleTalk multicasts.

In terms of overall network traffic, the AppleTalk Chooser is particularly broadcast intensive. The Chooser is the software interface that allows the user to select shared network services. It uses AppleTalk multicasts to find file servers, printers, and other services. When the user opens the Chooser and selects a type of service (for example, a printer), the Chooser transmits 45 multicasts at a rate of one packet per second. If left open, the Chooser sends a five-packet burst with a progressively longer delay. If left open for several minutes, the Chooser reaches its maximum delay and transmits a five-packet burst every 270 seconds. By itself, this does not pose a problem, but in a large network, these packets add to the total amount of broadcast radiation that each host must interpret and then discard.

Other AppleTalk protocols, such as the Name Binding Protocol, which is used to bind a client to a server, and the Router Discovery Protocol, a RIP implementation that is transmitted by all routers and listened to by each station, are broadcast intensive. The system init called AutoRemounter (part of the Macintosh operating system) is also broadcast intensive.

Note The AppleTalk stack is more efficient than the Novell stack because the AppleTalk stack discards non-AppleTalk broadcasts earlier than the Novell stack discards non-Novell broadcasts.

Figure E-3 shows the results of tests that Cisco conducted on the effect of broadcast radiation on the performance of a Power Macintosh 8100 and a Macintosh IIfx. Both CPUs were measurably affected by as few as 15 broadcast or multicast frames per second.

Figure E-3 Effect of Broadcast Radiation on Hosts in AppleTalk Networks

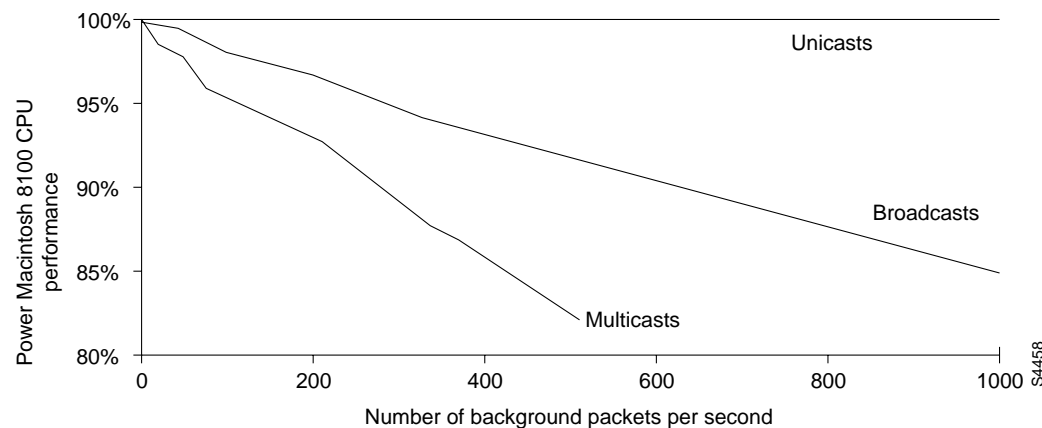


Table E-3 shows the average and peak number of broadcasts and multicasts for AppleTalk networks ranging from 100 to 10,000 hosts per network.

Table E-3 Average Number Broadcasts and Multicasts for AppleTalk Networks

Number of Hosts	Average Percentage of CPU Loss per Host	Peak Percentage of CPU Loss per Host
100	.28	6.00
1,000	2.10	58.00
10,000	16.94	100.00

Slow LocalTalk-to-Ethernet connection devices are a major problem in large-scale AppleTalk networks. These devices fail in large AppleTalk networks because they have limited ARP caches and can only process a few broadcasts per second. Major broadcast storms arise when these devices lose their ability to receive Routing Table Maintenance Protocol (RTMP) updates. Once this occurs, these devices send ARP requests for all known devices, thereby accelerating the network degradation because they cause their neighbor devices to fail and send their own ARP requests.

Multiprotocol Networks

The following can be said about the interaction of AppleTalk, IPX, and IP:

- AppleTalk stacks ignore any other Layer 3 protocol.
- AppleTalk and IP broadcast and multicast packets affect the operation of IP and IPX stacks. AppleTalk and IP packets enter the stack and then are discarded, which consumes CPU resources.

These findings show that AppleTalk has a cumulative effect on IPX and IP networks.

