

FORE
SYSTEMS



The logo for FORE SYSTEMS, featuring the word "FORE" in a large, white, sans-serif font above the word "SYSTEMS" in a smaller, white, sans-serif font. The text is set against a black rectangular background. A thin red horizontal line is positioned below the black rectangle.

FORE
SYSTEMS

The IP Transmission Process

V1.4: Geoff Bennett

Contents

Communication Between Hosts

Through a MAC Bridge

Through a LAN Switch

Through a Router

The tutorial is divided into four sections.

Section 1 looks at direct communication between IP hosts on the same LAN segment.

Sections 2 and 3 describe how this is affected if the hosts are separated by a MAC Layer Bridge or LAN Switch.

The final section looks at communication between hosts that are separated by an IP Router.

Section 1:


Communication Between Hosts

Through a MAC Bridge

Through a LAN Switch

Through a Router


Initially, let me describe the way that IP software makes the decision to transmit an IP datagram *directly* between two hosts located on the *same segment*.




Here we see two computer users connected by a cable in a *local* environment. Let's assume that both computers are running IP communication software.

Sally is using some form of IP application to send datagrams to Harry.

192.32.15.1




192.32.15.2




Here are the IP addresses...

192.32.15.1
255.255.255.0




192.32.15.2
255.255.255.0




Here are the IP addresses.
...and the subnet masks.

192.32.15.1
255.255.255.0



192.32.15.2
255.255.255.0




The IP software in Sally's PC needs to find out if Harry's machine is in the same *Broadcast Domain*. A Broadcast Domain is simply the term that describes the extent of spread of a MAC level broadcast.

In other words, if two hosts are located in the same Broadcast Domain, they will hear broadcast and multicast traffic transmitted by each other.


Transmissions that occur *within* a Broadcast Domain are called *intranetwork*.

Transmissions *between* Broadcast Domains are called *internetwork*.

192.32.15.1
255.255.255.0




192.32.15.2
255.255.255.0




To determine if Harry is in the same Broadcast Domain, Sally's software performs this comparison...

192.32.15.1
255.255.255.0



192.32.15.2
255.255.255.0




Source: (192.32.15.1) *AND* (255.255.255.0) = **192.32.15.0**


To determine if Harry is in the same Broadcast Domain, Sally's software performs this comparison...

She takes her own IP address, and performs a *binary AND* with her subnet mask...

192.32.15.1
255.255.255.0



192.32.15.2
255.255.255.0



Source: (192.32.15.1) *AND* (255.255.255.0) = **192.32.15.0**


Destination: (192.32.15.2) *AND* (255.255.255.0) = **192.32.15.0**

To determine if Harry is in the same Broadcast Domain, Sally's software performs this comparison...


She takes her own IP address, and performs a *binary AND* with her subnet mask...

...then she takes Harry's IP address and performs a *binary AND* with *her own* subnet mask (Sally has no way to find out what Harry's subnet mask is).

192.32.15.1
255.255.255.0



192.32.15.2
255.255.255.0




Source: (192.32.15.1) AND (255.255.255.0) = **192.32.15.0**

Destination: (192.32.15.2) AND (255.255.255.0) = **192.32.15.0**


The resulting two numbers are identical.

This tells Sally that she's in the same Broadcast Domain as Harry.

192.32.15.1
255.255.255.0



192.32.15.2
255.255.255.0



Source: (192.32.15.1) *AND* (255.255.255.0) = **192.32.15.0**

Destination: (192.32.15.2) *AND* (255.255.255.0) = **192.32.15.0**

To transmit to a host in the same Broadcast Domain, Sally needs to determine the MAC address (since she only knows the IP address).

To do this, she uses the Address Resolution Protocol (ARP), which is described in another tutorial.

Communication Between Hosts


Section 2:

Through a MAC Bridge

Through a LAN Switch


Through a Router

Let's see if the basic, direct transmission process changes if we separate the hosts using a MAC Layer Bridge.




In this modified example, a MAC-Layer Transparent Bridge connects two LAN segments. Harry and Sally are attached to different segments. How does this affect the IP Transmission process?

192.32.15.1
255.255.255.0



Bridge


192.32.15.2
255.255.255.0



In fact, Harry and Sally have the same IP addresses and subnet masks as if they were connected to the same LAN segment.


The Bridge is essentially “transparent” to the IP communication.

192.32.15.1
255.255.255.0




Bridge

192.32.15.2
255.255.255.0



The side effect of this transparent operation is that broadcast traffic (such as ARP) must be allowed to pass through the bridge.

Any attempt to filter ARP traffic results in a loss of communication between hosts.



Once again, we can say that these two hosts are within the same Broadcast Domain, even though they are separated by a Transparent Bridge.

Communication Between Hosts


Through a MAC Bridge

Section 3:


Through a LAN Switch

Through a Router

LAN Switches are an interesting new connectivity tool. How do they effect the IP communication process?




In this diagram we show Harry and Sally are both connected to the same LAN Switch. Switches are multiport MAC Layer Bridges that offer dedicated LAN bandwidth to each port. Typical workgroup switches offer 8, 12, 16, 24 or 64 connections to a single bridging unit.




Hosts that are attached to the same Switch must have the same Network ID and the same subnet mask in order to communicate.

A standard ARP transaction is used to transfer traffic from one host to the other.




Of course, Harry and Sally are not necessarily connected to the *same* switch. However, inter-switch connections are usually just extensions of the Broadcast Domain, with a simple transfer of the Learning Bridge address tables.



An extended switch network is completely transparent to the IP Transmission process, hosts simply ARP directly for the Destination IP address.

To transmit to a host that is in the ***same*** Broadcast Domain, IP software will ARP ***directly*** for the ***destination host*** IP address.

This simple rule summarises the basic operation of the IP Transmission algorithm for hosts on the same LAN segment, and for LAN segments connected by bridges or switches.



One of the side effects of switch designs is that ARP (and other broadcast traffic) *must* be flooded throughout the full extent of the switched network.

Each of these broadcasts will interrupt the host that it reaches, and will waste valuable network bandwidth.

To allow all of these Broadcast Domain networks to scale, we must turn to Router technology.


Communication Between Hosts

Through a MAC Bridge

Through a LAN Switch


Section 4:

Through a Router




In this diagram I show Harry and Sally connected via an IP Router.

192.32.15.1
255.255.255.0




Router

192.32.15.2
255.255.255.0




The IP addresses we have been using so far will not work correctly when Harry and Sally are connected through a router.

192.32.15.1
255.255.255.0



192.32.15.2
255.255.255.0




Router




The first reason for this is that the Router *must not* forward broadcast traffic. This means that the ARP request sent out by Sally will never be received by Harry.

192.32.15.1
255.255.255.0



192.32.15.2
255.255.255.0




Router

192.32.15.15

192.32.15.16

The second problem is that the router ports must be configured with IP addresses. As you can see, in this addressing scheme both ports have the same Network ID (highlighted in red). Internally, the Routing Table would not be able to differentiate between these ports.

192.32.15.1
255.255.255.0



192.32.16.2
255.255.255.0




192.32.15.15


192.32.16.16

This diagram shows a modified IP addressing scheme. As you can see, the Network IDs on the left hand side (192.32.15.0) are different from those on the right hand side (192.32.16.0). All the Network IDs are highlighted in red.

192.32.15.1
255.255.255.0



192.32.16.2
255.255.255.0



Router


Source: (192.32.15.1) *AND* (255.255.255.0) = **192.32.15.0**

Destination: (192.32.16.2) *AND* (255.255.255.0) = **192.32.16.0**

When Sally's IP software makes the same comparison that I described earlier, there is now a difference between the resulting numbers.

Sally's software has determined that Harry is located in a *different* Broadcast Domain.


To communicate between one Broadcast Domain and another, IP hosts must use an IP Router.




Sally's software needs to know the IP address of a router that is located *in the same Broadcast Domain*.

There are several ways the software can find out this address. Let's just assume for the moment that the address is known - 192.32.15.15. The router that an IP host uses to transmit IP datagrams out of the Broadcast Domain is called the *Default Router*.


(NOTE: The older term "Default Gateway" is still used by some IP veterans)



Sally's software performs an ARP *for the Default Router* address.
By definition, this address must be in the same Broadcast Domain.
While her software is waiting for the ARP response, the IP traffic (shown as the blue rectangle) is queued.



Once the ARP request is resolved, Sally sends the traffic destined for Harry to the Default Router.




The router examines the datagrams to see if the destination IP address is known. The procedure used by the router is described in another tutorial.

In this case, the router knows that Harry's LAN segment is directly attached. The router sends out an ARP request to find out Harry's MAC address.

The IP address for Harry is, of course, contained as the Destination Address in the IP datagrams.

The router queues the traffic while awaiting an ARP response.



Once the ARP is resolved, the router sends the traffic to Harry.

To transmit to a host that is in the **same** Broadcast Domain, IP software will ARP **directly** for the **destination host** IP address.

To transmit to a host that is in a **different** Broadcast Domain, IP software will ARP for the **default router** IP address.

So here are the two “Golden Rules” used by host software.



The End

This concludes the tutorial.

If you aren't viewing this on the FORE Systems Web Site, then you can become a member of the ATM Academy free of charge and have access to many more of these tutorials.

You can find details at:

<http://academy.fore.com/>