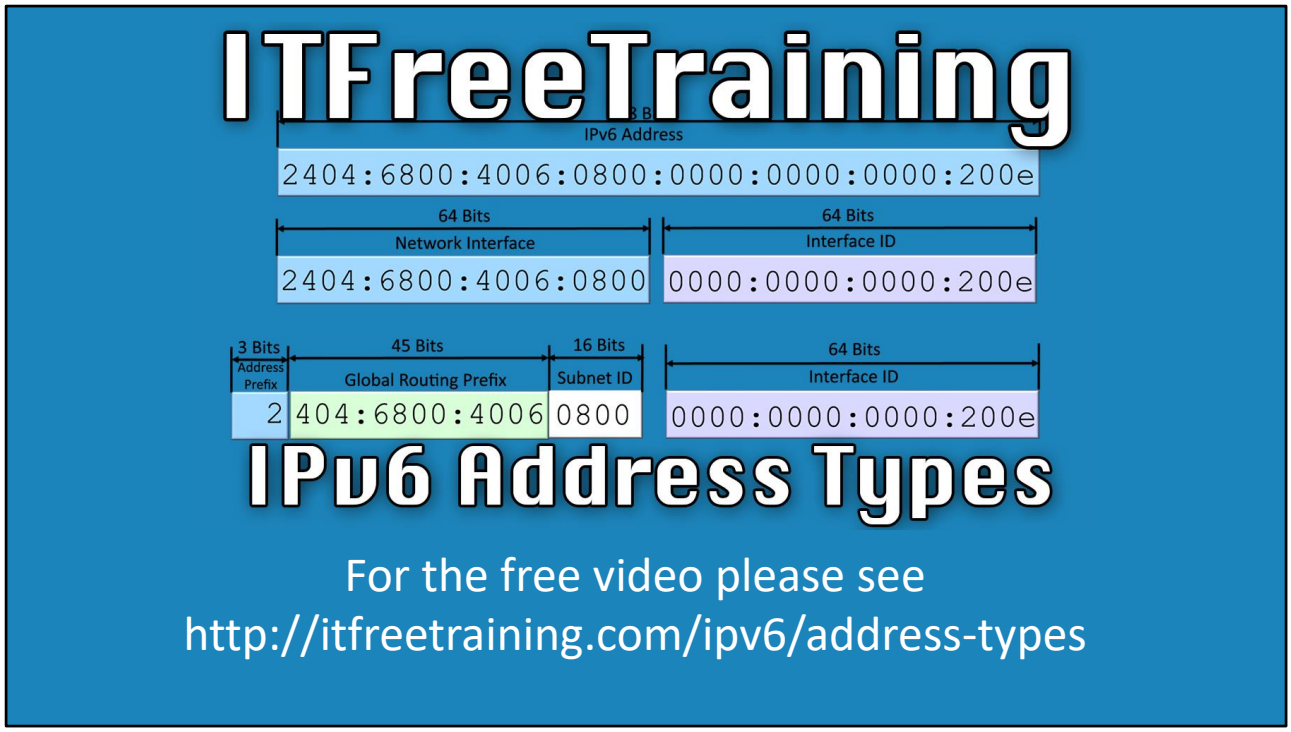

This video looks at the IPv6 address and its different types. The IPv6 address is much larger than the IPv4 address and looks complicated. By end of the video, you will be able to recognize the different types of IPv6 addresses making them less complicated to work with.

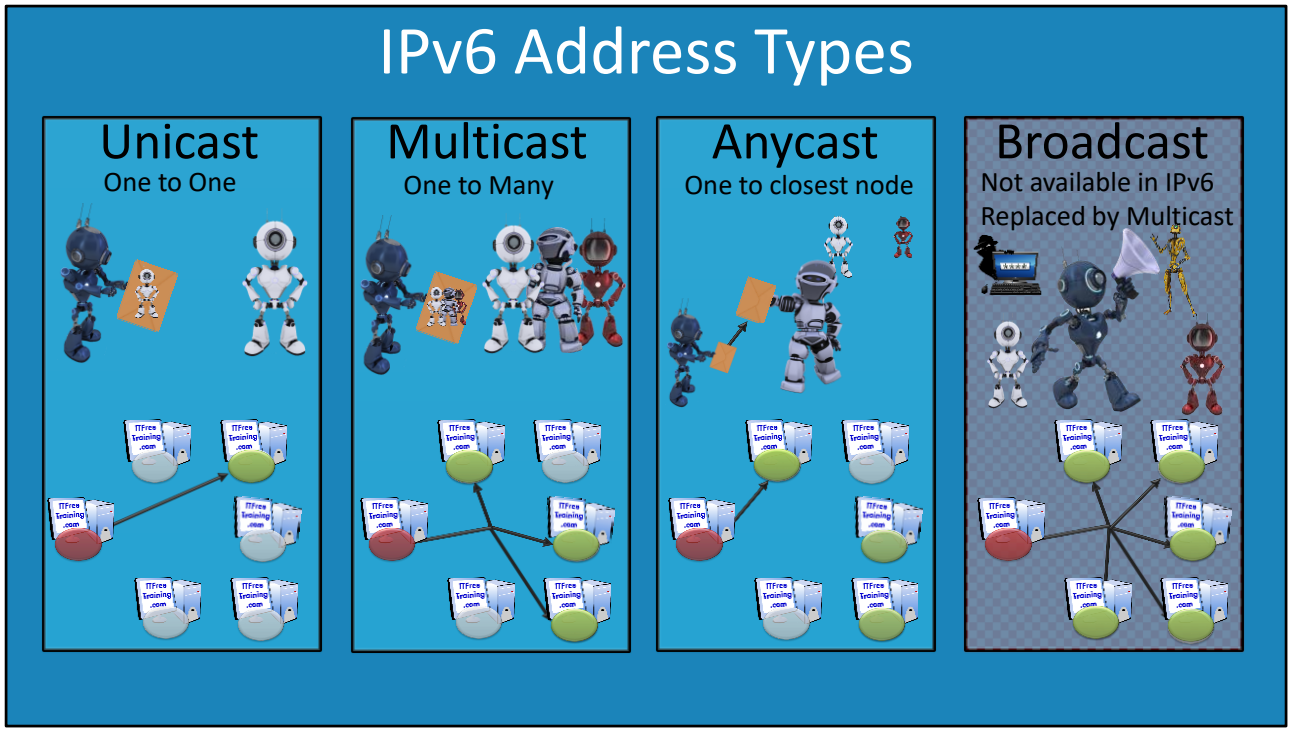

0:18 Looking at basic types of IPv6 address. These are unicast, multicast and anycast. Depending on which type of address is used will determine which node or nodes the data will be routed to.

Unicast: Unicast is the most commonly used communication used on IPv6 and IPv4 networks. Unicast is used when data is sent from one to another node of the network. The other nodes on the network will not receive the data. In order for a data to be intercepted, a person would need to get between the two nodes. The exception to this is older network equipment like hubs which broadcast data to all nodes.

Multicast: Multicast is when the one packet goes to many nodes on the network. The advantage of this is that it reduces the amount of traffic on the network. For example, when deploying operating systems over the network, most deployment software will support multicast. Since operating systems are very large, if you were deploying an operating system to 20 computers at once, the packets would only need to be sent once over the network. In contrast, deploying 20 computers at once using unicast would require the data to be sent 20 times. In order for a node to receive multicast traffic, the node needs to join a multicast group. An intruder on a network could join these multicast groups and thus receive this data. IPv6 does offer additional security so that an administrator can determine who can join particular multicast groups. However, this does need to be set up and the network equipment needs to support it.

Anycast: Anycast is when the data goes from the source node only to the closest node. Anycast works by multiple nodes on the network having the same IP Address. This also means that the router on the network needs to be configured to know where these nodes are. On the internet, it would be difficult to configure anycast as it would require ISPs and other authorities to agree to

change their routing tables. In a company, you would have more control. For example, you could configure multiple DNS servers with the same anycast address. Nodes would then go to the closest DNS server with that anycast address. If this DNS server was to go down, the route would be removed and the nodes would go to the next closet DNS server.

Broadcast: Broadcast is not supported in IPv6. It has been replaced with multicast. To send a message to all nodes on an IPv6 network you would send this message to the multicast group for all nodes. Broadcasts were used a lot in IPv4 and were inefficient. This is because they were often used to communicate with one node, however, all nodes on the network would receive the traffic. This created security problems as it was easy for an intruder to connect to the network and passively listen to all the broadcasts on the network. IPv6 also offers security, if your hardware supports it, to prevent unauthorized nodes from joining multicast groups.

Unicast IPv6 Address Types (One to One)

- Global Unicast Addresses
  - Globally routable on the internet (Public IPv4 address)
- Unique Local Addresses
  - Administrator assigns (IPv4 private addresses)
- Link-Local Addresses
  - Communicates with nodes on same link (IPv4 APIPA)
- Special Addresses
  - No address (Unspecified) or loopback
- Compatibility Addresses
  - Used for migration from IPv4 to IPv6

07:08 Unicast traffic is one to one traffic. That is, traffic goes from a source node to a destination node. There are five different types of unicast addresses. Each has a different property and are used in different circumstances. These are:
Global Unicast Addresses: These are IP Addresses that are registerable on the internet. If you register an IP Address this means that anyone on the internet can connect using that IP address. Thus, a global unicast address is routable on the public internet. This is the same as a public IPv4 address.
Unique Local Addresses: These addresses are not routable on the internet. They are used on internal networks. They are equivalate to IPv4 private addresses. For example, the addresses starting with 192.168.* and 10.*
Link-Local Addresses: These addresses are only able to communicate on local networks. That is, they are not able to travel over a router to different networks. These addresses are like the APIPA addresses in IPv4. APIPA addresses start with 169.254.* APIPA addresses are assigned in IPv4 when a DHCP server cannot be contacted. The difference with IPv6 is that link-local addresses are always assigned to a network interface. These are used for basic communication on an IPv6 network and are needed for services like neighbor discovery.
Special Addresses: There are two special addresses. The no or unspecified address is represented by 0:0:0:0:0:0:0:0 or ::. In IPv4 this address is 0.0.0.0. The loopback address is a special address in which any traffic sent to it will be returned back to the interface that sent it and is used for troubleshooting. In IPv6 this is ::1. In IPv4 this is 127.0.0.1
Compatibility Addresses: These are global unicast addresses that are used in the migration from IPv4 to IPv6. For example, IPv6 may have an IPv4 address embedded in the address itself. These addresses are routable on the internet and are designed for the migration of IPv4 devices to IPv6.

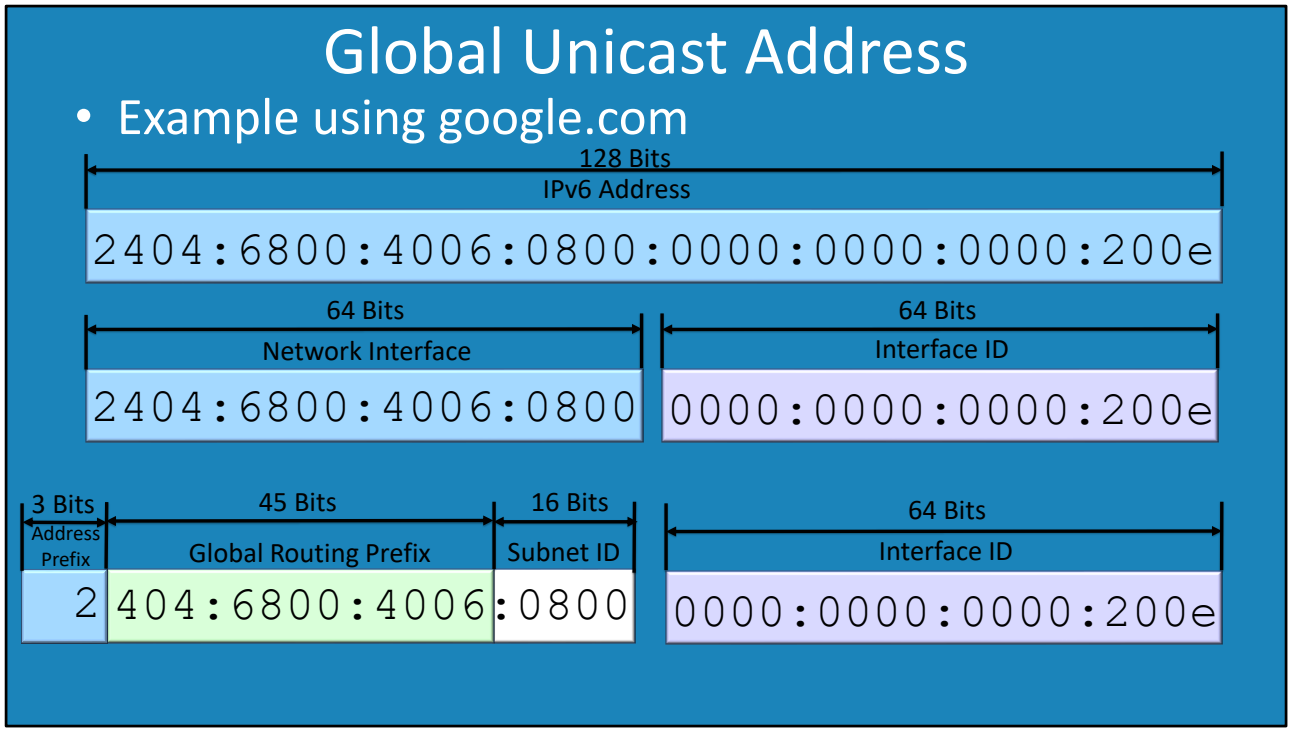

09:45 A global unicast address at first glance looks very complicated. The address can be divided into two parts. The first 64 bits is the network interface. The second 64 bits is the interface ID. In IPv4 a subnet mask was used to divide an IPv4 address into the network and interface portion. In IPv6, the first 64 bits is always the network interface and the last 64 bits is always the interface ID. For this reason, subnet masks are not required in IPv6. Global unicast addresses can be further broken down into the following.

Address Prefix: This is the first three bits of the address. Currently only address that start with 2 and 3 are being used. There are other ranges that are reserved. These ranges may get used in the future, but currently this is enough to meet the needs of the internet.

Global Routing Prefix: The global routing prefix is 45 bits of the IP Address. This is controlled by the Internet Assigned Numbers Authority (IANA). IANA will allocate ranges from these 45 bits to regional registries. The American Registry for Internet Numbers (ARIN) will be given a range from these 45 bits. ARIN is responsible for allocating IP Addresses in the United States, Canada, several parts of the Caribbean region, and Antarctica. ARIN will then allocate addresses to ISPs who will allocate them to customers. By looking at the global routing prefix, you can determine who the IP Address is registered to and where it is located. This is how routers on the internet know where to route traffic to.

Subnet ID: 16 bits of the address is the subnet ID. In most cases this is controlled by the company and in some cases some of the bits may be controlled by an ISP. These 16 bits are reserved for the subnet ID and this is why a subnet mask is not required. In IPv4, the administrator had to make a choice of how to divide an IP Address up using a subnet mask. In IPv6 since 16 bits are dedicated to the subnet ID, the administrator never needs to worry about using a subnet mask. The subnet ID is used by an organization's internal routers to route traffic. The subnet ID is not registered on the internet.

Interface ID: The interface ID is the last 64 bits. This can be chosen by the administrator or assigned by the operating system. In the case of windows, the interface ID will be random. In the case of operating systems like Linux, the MAC address will be used for the interface ID. In the Google example they have chosen a number, but it could have been anything.

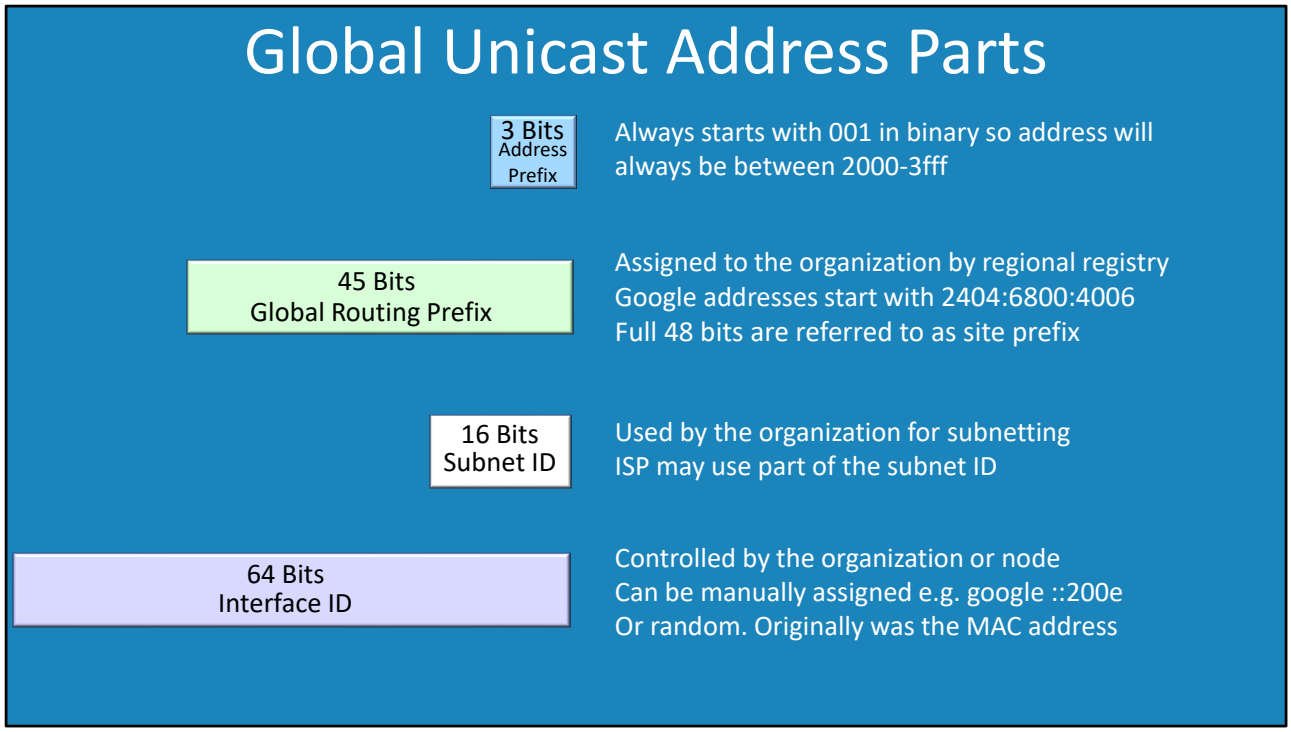

13:18 As seen, a global unicast IPv6 address can be broken up into four parts. I will go into more detail about each part.
Address Prefix: The first three bits of the address determine that it is a global unicast address. This means the address will always start with a 2 or a 3. To understand this, consider that the address will always start with the binary value 001. To convert binary to hexadecimal you need four bits. So any address starting with 0010 or 0011 will be converted to 2 or 3 in hexadecimal. Thus, a global unicast address will always start with a 2 or a 3. In order words, the address will always start with a value between 2000-3ffff.
Global Routing Prefix: The global routing prefix is assigned to a customer. It is unique on the internet. In Google's case this is 2404:6800:4006. The first 48bits of a global unicast address is referred to as the site prefix.
Subnet ID: This is allocated by the organization and in some cases, some of it may be allocated by the ISP. For example, the ISP may allocate four bits of the subnet ID and leave the organization 12 bits. The administrator of the organization is free to divide the subnet ID up however they like. For example, they may use a hierarchic approach, where the first eight bits is for countries and the next eight bits for the network in that country.
Interface ID: This can be manually assigned, assigned by the operating system or allocated by a DHCP sever. In this example, Google has assigned it the value of ::200e. In the original IPv6 implementation, the interface ID was the MAC address. The MAC address is unique to the network adapter. Thus, if a person was to travel between different networks, for example different wireless hot spots, the user could be tracked by using the MAC address. For this reason, Windows randomizes the interface ID rather than using the MAC address. Linux still uses the MAC address. The user is free to manually assign any interface ID they want, as long as it is unique on that local network.

Unique Local Address

- Non routable privately assigned addresses
  - Equivalent to IPv4 private addresses
    - 10.0.0.0/8, 172.16.0.0/12, 192.168.0.0/16
- Starts with fd

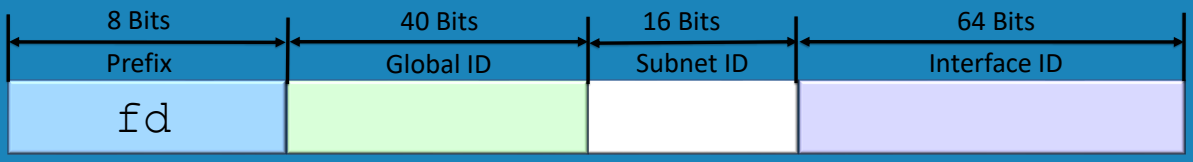

15:45 These IPv6 addresses always start with fd and assigned by the administrator. They are functionality equivalent to IPv4 private IP Addresses, such as IPv4 addresses starting with 10. The administrator is free to configure unique local address however they see fit in their organization. The fd addresses contain a global ID and a subnet ID. The administrator is free to configure their internal network using the global ID and subnet ID and thus make a fully routable internet network. Since unique local addresses are not registered on the internet, they are not routable on the internet. Previously in IPv4, private IP Addresses had to be used as there was a shortage of IP Addresses. In IPv6, there are plenty of IP Addresses and thus the administrator has the choice of using unique local addresses on their internal network or publicly routable addresses on their internal networks. Even with the larger address space, an administrator may elect to use unique local addresses. This is because having all addresses on an internal network publicly routable can cause potential security problems. For example, a computer on an internal network can now directly be connected to any computer on the internet. It is matter of the administrator having security such as firewalls in place to prevent this. However, with a unique local address, the device would first have to make a connection to the internet in order for an attacker to access the device. Thus, unique local address offer additional security by design.

Link-local Addresses

- Used to communicate with neighboring nodes
  - Not routable
- Equivalent to APIPA in IPv4
- All nodes are assigned one of these addresses
- Start with fe80

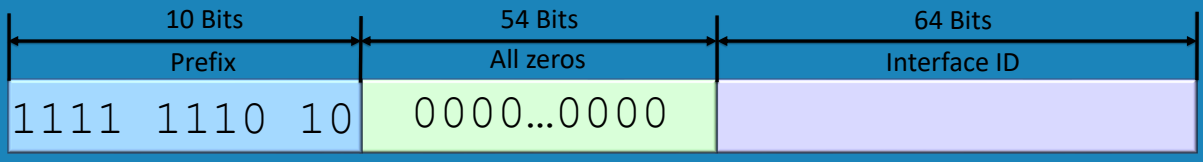

18:17 One of the goals of IPv6 was to remove the need for broadcasts. In order to find other nodes and services on the network, IPv4 used broadcasts. In IPv6, every node on the network is assigned a link-local address. This address is used on the local network only and thus is not routable. This address is used for services like neighbor discovery. Since every node on the network has a link-local address, the link-local address is used between different nodes on the same network to communicate rather than using broadcasts. Thus IPv6 does not require broadcasts. In IPv4, a link-local address would be assigned to a node if an IP Address was not statically assigned and the node could not reach a DHCP server. In IPv4 the idea was that nodes could communicate with each other on the same network when a DHCP server was not present. If the node later got an IP Address, IPv4 use this IP Address and the APIPA address would be removed. In IPv6, every node is allocated a link-local address when networking is started up. It is the first task IPv6 will perform. The network part of the address, the first 64 bits, will be fe80 followed by zeros. The interface part of the address, the last 64 bits, will be random on Windows and for Linux will be the MAC address. Once a link-local address has been allocated to the node, the node will keep using the link-local address for services like neighbor discovery regardless of whether the node gets allocated an IP Address later on. By using the link-local address, this eliminates the need for the node to use broadcasts on the network.

Special Addresses

- Unspecified address
  - 0:0:0:0:0:0:0:0 or ::
  - Equivalent of IPv4 0.0.0.0
- Loopback address
  - 0:0:0:0:0:0:0:1 or ::1
  - Equivalent of IPv4 127.0.0.1

20:34 There are two specials address in IPv6. These addresses have a special function when used. Unspecified Address: This address is 0:0:0:0:0:0:0:0 or ::. Some interfaces in IPv6 do not support the use of :: so you may need to use 0:0:0:0:0:0:0:0 instead. It is the same as the address 0.0.0.0 in IPv4. It is used in routing to indicate the default gateway.
Loopback address: Anything sent to the loopback address is returned back to the node that sent it. It does not go out on the network, so essentially is a way of testing if the software on the network connection is working. It is also a way or accessing the local host, as the traffic will always be diverted back to the node that sent it. In IPv6, the address is 0:0:0:0:0:0:0:1 or ::1. It is a single address in IPv6. In IPv4 the loopback address was any address starting with 127, however, 127.0.0.1 was the most commonly used address.

Compatibility Addresses

- Used for the migration of IPv4 to IPv6
- IPv4 address is added to IPv6 address
- Often IPv4 address represented by w.x.y.z in IPv6 address
  - E.g. 0:0:0:0:0:0:*w.x.y.z*

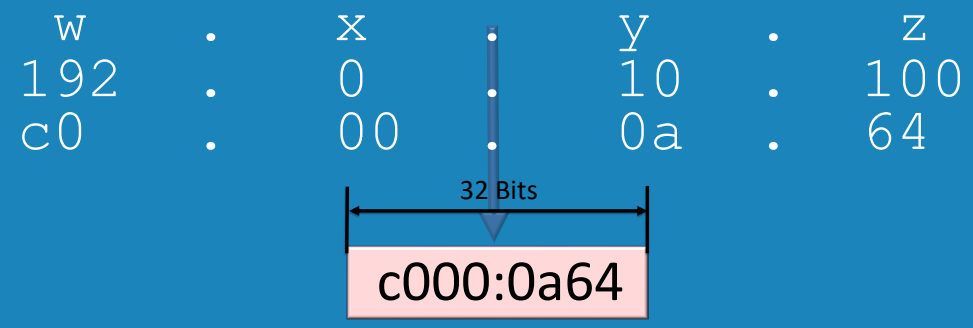

21:50 With the migration from IPv4 to IPv6, in some cases the IPv4 address will need to be embedded in the IPv6 address. This is because the data may need to travel over a mixture of IPv4 and IPv6 networks in order to get to the destination. In compatibility addresses that embed an IPv4 address in the IPv6 address, you will see this written as w.x.y.z. This translates to the four values in the IPv4 address. In some cases, you may see the values converted into hexadecimal, and in other cases you will see them converted into binary. Either way is correct. Generally, it is easier to leave the values as IPv4 octets as it is easy to read.

6to4

- Nodes have IPv6 and IPv4 address
- IPv4 address is encapsulated in IPv6 address
- IPv6 packets can be routed over IPv4 networks
- Address starts with 2002

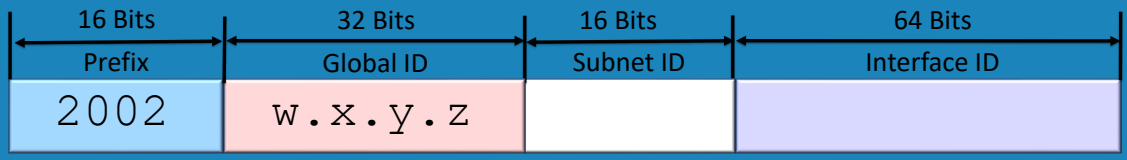

23:30 When 6to4 is used, the node is given an IPv4 address and an IPv6 address. The IPv6 address will have the IPv4 address embedded in it. The embedded IPv4 address needs to be a globally registered IPv4 address. 6to4 addresses start with 2002. They have 16 bits for a subnet ID and the last 64 bits are for the interface ID. Like before, the interface ID is random in Windows and the MAC address is used in Linux. 6to4 is used when you want to connect two IPv6 networks together and there is an IPv4 network between.

ISATAP

- Routes IPv6 packets over IPv4 networks
- IPv4 address is encapsulated in the Interface ID
  - Interface ID starts with 0:5efe
  - Followed by w.x.y.z

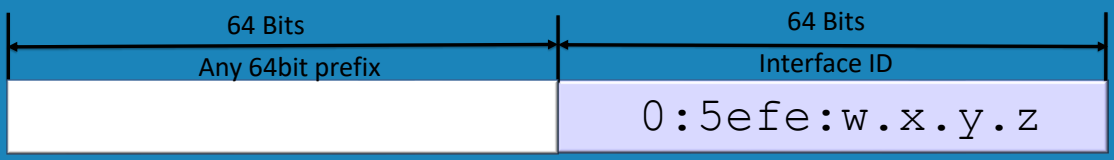

24:48 Intra-Site Automatic Tunnel Addressing Protocol is a protocol that is used on internal networks only. It is not routable on the internet. In the case of ISATAP, the node is assigned an IPv6 address. The node interface ID will start with 0:5efe followed by the IPv4 address. If the node is on an IPv4 only network, the IPv6 packet will be placed in an IPv4 packet. This allows it to be transmitted over the IPv4 network. When it reaches and ISATAP router, the IPv4 packet is removed and the IPv6 packet is sent over the IPv6 network. Since the IPv4 address is embedded in the IPv6 address, this allows the ISATAP router to work out the destination on the IPv4 network and sent it correctly over the IPv4 network.

Teredo

- Allows IPv6 clients to work with IPv4 NAT
- Addresses start with 2001:0000
  - IPv6 address contains
    - Teredo IPv4 server, flags, UDP port, Client IPv4

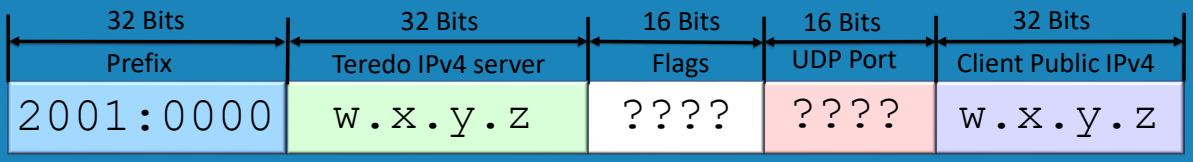

26:24 Teredo allows IPv4 nodes to communicate with IPv6 nodes on the internet. The big difference between it and the other migration protocols is that it supports Network Address Translation (NAT). A Teredo address starts with 2001:0000. It is then followed by the Teredo server that it is using. It is up to the administrator to configure a Teredo server for the node to use. Windows had a default Microsoft Teredo server configured, however this Teredo server has been shut down. The next part is 16 bits of flags which lets Teredo know how to send the traffic. Following this is another 16 bits which is the port to be used. This is how it works around NAT. A specific port will be used that will allow the traffic to be transmitted over NAT. The last 32 bits are the IPv4 address of the node that sent the traffic. So using this and the port, allows traffic to be sent back and to go over NAT.

Multicast

- Sends data to multiple nodes
  - All nodes (IPv4 broadcast) ff02::1
  - All routers ff02::2
- Always starts with ff

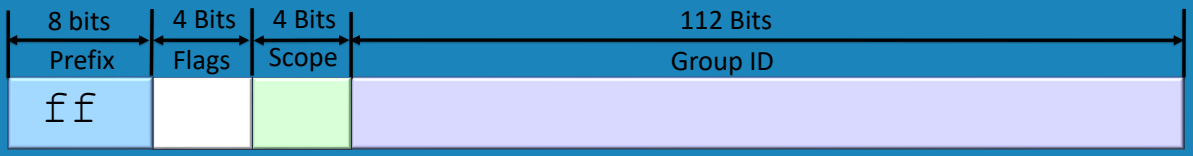

29:06 Data that is sent to a multicast address is designed to be sent to multiple nodes. The idea behind multicast is that it replaces the need for broadcast. Any IPv6 address that starts with ff is multicast. Following this is four bits for flags and four bits for scope. The last 112 bits are for the group ID. A node on the network is able to create its own multicast group, but there are also some well-known multicast addresses. For example, ff02::1 will be sent to all nodes on the network. This replaces the existing broadcast on IPv4 networks. In this case the flags are set to zero. The scope is to 2 which is local network only.

Anycast

- Unicast address that has been made into Anycast
  - Indistinguishable from other Unicast addresses
- Data will go to closest node (Reduces DDOS attacks)
- Nodes advertised through routing protocols

31:00 Any valid unicast address can be made into an anycast address. So looking at an anycast address you will not be able to determine that it is anycast. The difference between an anycast address and a unicast address is the way they are routed. A unicast address that has been made into an anycast address will be present on multiple parts of the network. In other words, multiple nodes on the network will have the same anycast address. When a device attempts to contact a node with an anycast address, the traffic is routed to the device on the network that is closest to the device. This is achieved by using routing protocols. The routing protocols route traffic to the node that is closest. Inside an organization, the administrator will have control over the routers and thus can set up anycast. On the internet, the administrator would need to have changes made on routers controlled by ISP's. For this reason, anycast is difficult to set up on the internet. The advantage of anycast is that since traffic is routed to different nodes on the network, this reduces the effectiveness of distributed denial of service (DDOS) attacks. Since a DDOS works by sending data from all over the internet to a single point, having the nodes spread out across the internet reduces the effectiveness of the attack.

Summary

Global Unicast Addresses 2000-3fff (Routable on the internet) 6to4 starts with 2002 Teredo starts with 2001:0000 Teaching example 2001:db8
Link-Local Address fe80 (Local network only) Example: fe80::a29d:dc6e:7761:8a39
Unique Local Address fd (Private network only) Company can allocate how they see fit and not routable on the internet
Multicast Address ff (Go to multiple nodes) Used for neighbor discovery and other functions previous done with broadcast
Anycast (Go to nearest node) Are a unicast address that has been made into an anycast address

33:38 IPv6 addresses are quite long and look complicated at first. If you look at the start of the address this will tell you what the address is used for.

Global Unicast Address
Always start with 2000-2fff.
These addresses are routable on the public internet. Because they are registered on the internet, you can perform a lookup of the IP Address to find out who it is registered to. There are some common ones to look out for.
6to4 addresses start with 2002.
Teredo addresses start with 2001:000
You may also see 2001:db8 addresses used in documentation. These are not valid on the internet. They are just like phone number in movies that start with 555, they look real but are not and are used for documentation only.

Link-Local Address
These addresses start with fe80. Every node on the network will have one of these addresses regards of whether another IP Address or addresses have been assigned to it. It is used for local network traffic and for the basic functions of IPv6 like neighbor discovery. Every node on the network will have a different link-local address. In Windows this is randomized and in Linux it is based on the MAC Address.

Unique Local Address
These addresses start with fd. They are private addresses that are not routable on the internet. An organization can use these on their internal network in any way they see fit.

Multicast

These addresses start with ff. A node on the network will join a multicast group which means that it wants to receive all traffic sent to that group. In IPv6, there is no broadcast traffic. So, a node will join the multicast group of all nodes when it starts up. If you want to send data to all IPv6 nodes on the network, you would sent it to the multicast address for all nodes. Multicasts are used for basic services like neighbor discovery and can also be used for other purposes. For example, if the administrator was using deployment software, they could have all nodes that they want to receive traffic from the deployment software to join a particular multicast group.

Anycast

Anycast addresses can be any valid unicast address. The difference is that the router on the network will route it to the closest nodes that has been allocated that anycast address. So essentially multiple nodes on the network will have the same address. It is the same principal as emergency services. When you call the emergency service number the call will be diverted to the closest emergency service call center. An anycast address will be diverted to the closed node with the address.