A Primer on IPv6 Addressing

1. You already know that IPv6 addresses are 128 bits and are written using "**:**"
2. IPv6 eliminates broadcast, instead using multicast extensively.
   - But note multicast FF02**::1** address for "all nodes", equivalent to broadcast!
   - Anycast addresses ("to closest device") not extensively used yet (but DNS)
     - special handling achieved by suitable configuration on routers
     - indistinguishable from unicast addresses (RFC4291)
     - but note the "subnet router" anycast address: all 128 host bits set to 0
3. No need for private IP addresses; ideally no more NAT, all addresses could be globally routable
   - But have "unique [site] local unicast addresses" - RFC4193
4. Generally, IPv6 addresses are:

Global routing or Site prefix 48 bits	Subnet 16 bits	Interface ID ("host") 64 bits

5. IPv6 interfaces usually have at least **2 different addresses** and often more (yes, simultaneously):
   1. link-local address (typically autogenerated, FE80::/10 + 64 bit Interface ID)
   2. globally unique unicast address (eg. 2001::/16 + additional 112 bits)
6. Host-portion of address can be (and often is) automatically derived from L2 interface ID eg. Ethernet MAC address using EUI-64:
   - insert 0xFFFE in middle of 48-bit MAC address to make 64 bit host portion: MAC 00:0C:29:0C:47:D5 becomes 00:0C:29:FF:FE:0C:47:D5
   - See also Switch v7 Lab 5.2, p. 4-5 for a mixed fixed/random set of addrs
7. These (and other protocol changes, such ICMPv6) allow many tasks & features to happen automatically.

Good News: Both NetLab and PT support IPv6. Great for learning & testing!

Good references:

– Ranges & reserved addresses: https://en.wikipedia.org/wiki/IPv6_address
– http://www.rmv6tf.org/wp-content/uploads/2012/11/Implementing-IPv6-the-Nuts-and-Bolts-About-It1.pdf

IPv6 Address ranges (Partial! But most of the important ranges)

- All addresses from the block outlined in black are "global unicast" addresses and often have their lower 64 bits formed from the "interface ID".
- Addresses from the multicast range (FF00::/8) may be reserved or have the remainder of their bits freely assigned (some addresses of each type).

Prefix	10 high-order bits (0 = 0 bit, 1 = 1 bit)	Description
::/128	0000000000	"unspecified", not generally valid (! link-local !)
::1/128	0000000000	loopback
::/96	0000000000	IPv4-compatible (deprecated by RFC4291)
::FFFF:0:0/**96**	0000000000	IPv4 compatible addresses (RFC4291)
2000::/3	001	/3 block currently being assigned by IANA
2002::/16	0010000000	IPv6-to-IPv4 tunnelling (automatic; RFC3056)
4000::/3	010	(reserved for future use)
6000::/3	011	(reserved for future use)
8000::/3	100	(reserved for future use)
A000::/3	101	(reserved for future use)
C000::/3	110	(reserved for future use)
E000::/4	1110	(reserved for future use)
F000::/5	11110	(reserved; range extends to FC00::/7, below)
FD00::/8	1111111	unique [site] local unicast addresses (RFC4193)
FE80::/10	1111111010	Link-local address + lower 64-bit interface ID
FEC0::/10	1111111011	Site-local address (deprecated by RFC3879)
FF00::/8	11111111	Multicast addresses
FF02::1	1111111100	"All nodes" multicast (link-local) = "broadcast"
FF02::2	1111111100	"All routers" multicast (link-local)
(others...)	11111111	Many other reserved multicast addr (see IPv4!)
FF02::1:FFXX :XXXX/128	1111111100	Solicited node multicast (link-local) = "arp" XX:XXXX = 24 low order bits of IPv6 address

IPv6 scheme for mapping mcast Ethernet <—> IP addresses is similar to IPv4:

IPv4: 0x01-00-5E + a 0 bit + low-order 23 bits of IPv4 address

IPv6: 0x33:33 + low-order 32 bits of IPv6 address

IPv6 Features (RFC 4861)

Neighbor **D**iscovery Protocol (**NDP**, or **ND** in Cisco output) is embedded inside IPv6 ICMP and enables plug-and-play features, using the following functions:

1. Router Discovery – Discover the local routers without DHCP.
2. Prefix Discovery – Discover the prefix or prefixes assigned to that link.
3. Parameter Discovery – Discover other parameters such as the link MTU and hop limits for its connected link.
4. Address Autoconfiguration – Determine its full address, without DHCP.
5. Address Resolution – Discover the link-layer addresses of other nodes on the link without the use of Address Resolution Protocol (ARP).
6. Next-Hop Determination – Determine the link-layer next hop for a destination, either as a local destination or a router to the destination.
7. Neighbor Unreachability Detection – Determine when a neighbor on a link, either another host or a router, is no longer reachable.
8. Duplicate Address Detection – Determine if an address it wants to use is already being used by another node on the link. (Similar to gratuitous ARP)
9. Redirect – A router can notify a host of a better next-hop than itself to an off-link destination. The redirect function is a part of basic ICMP functionality in IPv4, but is redefined as part of NDP in IPv6.
10. (Entire) Subnet renumbering (reassignment) – Swap subnet prefix for all devices on the subnet by reconfiguring only the router(s)

Key messages:
RS: Router Solicitation (request) [type=133, code=0]
RA: Router Advertisment (periodic, or in response to RS) [type=134, code=0]
NS: Neighbour Solicitation (request) [type=135, code=0]
NA: Neighbour Advertisment (response) [type=136, code=0]
Redirect: Tell host of better gateway [type=137, code=0]

Good references:
– IPv6 ICMP: defined by RFC 4443; discussed in https://en.wikipedia.org/wiki/ICMPv6
– IPv6 SLAAC: http://www.ciscopress.com/articles/article.asp?p=2154680
– NDP: http://resources.intenseschool.com/ipv6-neighbor-discovery/

Your job is to determine which of these five messages provides the functionality for each of the 10 items in the list above. You might be surprised how useful your IPv4 knowledge is for (guessing?) these items.

IPv6 Tunneling References:

"6in4" https://en.wikipedia.org/wiki/6in4
"6to4" http://packetlife.net/blog/2010/mar/15/6to4-ipv6-tunneling/