IPv6 addressing and routing in Openstack

Enumeration of possibilities

We can configure IPv6 addresses on VMs in Openstack using the following mechanisms:

1. SLAAC from a neutron router, using RAs (on networks with routers, providing we also configure the route)
2. DHCPv6 from a neutron server (also works on unrouted networks)
3. static config on a config drive (on any VM)
4. link local
5. SLAAC from an external router on a provider network
6. DHCPv6 from an external server on a provider network

We can configure routes on VMs in Openstack using the following mechanisms:

1. RAs from a neutron router
2. RAs from an external router on a provider network
3. static config on a config drive

External information may be propagated via DHCP server. For instance, we can configure MTUs on VMs in Openstack using the following mechanisms:

1. DHCPv6 from a neutron server
2. DHCPv6 from an external server on a provider network
3. static config on a config drive

Note that, although static config always works, we cannot use it in bare metal systems, so we must not rely on it. Also note that, at present, metadata is not available over IPv6. This blueprint does not address that issue, which is one of a lack of well known address.

SLAAC based addressing

SLAAC addresses come from a router (neutron or otherwise) and are associated with a single port and a single subnet. The router may be neutron controlled or external. Port firewalls should block RAs from unapproved sources. Port firewalls should respond to the address the router assigns to the VM, recognising it as a source address for that port.

1. for Neutron routers, they may advertise a route with SLAAC, one without SLAAC, or they may have their advertisement turned off.
2. any subnet may have a maximum of one gateway address. This may be from a neutron router - and should be automatically assigned when a neutron router is attached to that network - or it may be external - so the value may be overridden by a network owner. The value may be unset. RAs will be accepted from this source for this subnet.

3. RAs are independent of DHCPv6. None of these options shall affect DHCPv6 in any way.

DHCPv6 based addressing

DHCPv6 addresses and further information may be sent over a network. Neutron may provide a DHCPv6 server or one may be enabled externally.
1. for Neutron servers, they may advertise an address and additional data, additional data only, or they may be disabled for a network.
2. any subnet may have a list of DHCPv6 server addresses. This may be from a neutron server - and should be automatically replaced when a neutron network is created or if the server is turned on or off - or it may be external - so the value may be overridden by a network owner. The value may be unset. DHCPv6s will be accepted from this source for all VM ports on the subnet.
3. RAs are independent of DHCPv6. None of these options shall affect RAs in any way.

Properties on objects supporting ipv6, and their uses

Routers

Routers have an IPv6 external address but also an external IPv6 prefix that is delegated to them. This was not necessary in IPv4 due to the NAT component of routers.

Attribute	Default	Notes
external_ipv6_address		Nexthop on which ipv6 packets received
external_ipv6_prefix		Prefix delegated to this router

Networks

Unlike IPv4, DHCP is not enabled on the entire network but on each subnet individually.

Subnets

In IPv4, a single address is taken from one subnet on a network and assigned to a port. Subnets are intended to extend the range of available addresses.
In IPv6, an address is taken from *every* subnet on a network.
Neutron-provided DHCPv6 may be enabled or disabled for a subnet using dhcpv6_enabled. In either case, the dhcpv6_address is set to indicate which address is used to source DHCPv6 packets, and may also be set to “”. This is used by the firewall; “” implies no firewall opening for incoming DHCPv6 packets, otherwise a specific hole is punched to allow for incoming and outgoing packets to (and only to) the known server.

Attribute	Default	Notes

dhcpv6_enabled	False	When set to true, sets dhcp6_address; when set false, clears dhcp6_address
dhcpv6_address	“”	set by dhcpv6_enabled; affects automatic port firewalling
ra_enabled	ra-slaac; also available: ra-only, off	Affect router; apply only when router port attached to subnet
ra_address	“”	Set to the router’s port address when router is attached; cleared when removed

Ports

The IPv6 source address is a map of (subnet ID -> IPv6 addresses) assigned to this port. There will be one address per subnet plus one link local address (subnet ID = “”). Depending on the nature of the subnet addressing, the address may be user modifiable, or it may be derived from the MAC and therefore a read-only value. Privacy extensions are not supported at this time as port firewalling requires a predictable address; machines with privacy extensions turned on will find their outgoing traffic is blocked.

Depending on the address assignment mechanism it may be that an address from a subnet has been allocated, but the VM never learns of this address. As with IPv4, the API does not prevent such misconfigurations from occurring.

Attribute	Default	Notes
fixed_ips	{}	List including both v4 and v6 address allocations. For v6, one per subnet and the link local address. Read only.

Address selection

Addresses are selected by creating subnets on Neutron networks. Traditionally, any address has been selectable as a subnet, but this causes issues with routing, where an ipv6 network address that is expected to be routeable must be within the range of routeable addresses.

IPv6 address choices for subnets are therefore constrained.

Subnets may be created with any routeable (non-link-local, non-reserved) address, either globally routeable or ULA.

In ipv6, routers may be designated the external routers. This router will have an external port with an address and a routeable subnet range, which the external networking will ensure is used as the nexthop for all traffic. Unlike ipv4, the external network is meaningless.

Unlike ipv4, an ipv6 router interface may only be connected to a network if the router gateway has already been connected and a routeable address range assigned. If the router gateway is connected to an external network, the interface attach to the subnet will only succeed if the subnet is within the routeable address range of the external network port, *and* the range has not previously been attached. To facilitate this, subnets may be created with a new flag that ensures that their are immediately attached to a router with an unused subnet.