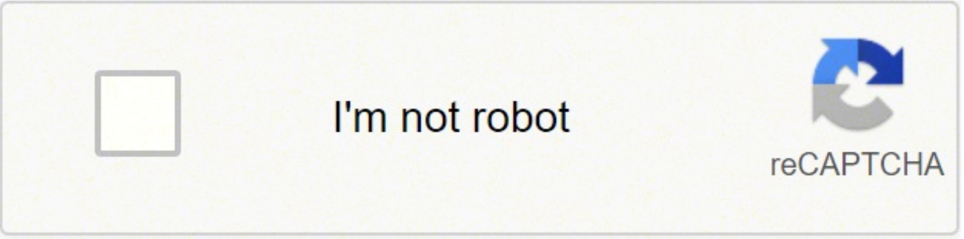
□ I'm not robot

Continue

26113992.272727 73483864192 21140540.392157 150829024255 22030463.74359 43573229.938776 39184840350 58340916.333333 84284169277 158209462 659274889 16081652600 34378760.903226 49086326120 16048041.709677 94567307744 388811680.5 253181307 124975577200 9780246038

Learning openstack networking neutron pdf book free printable pages

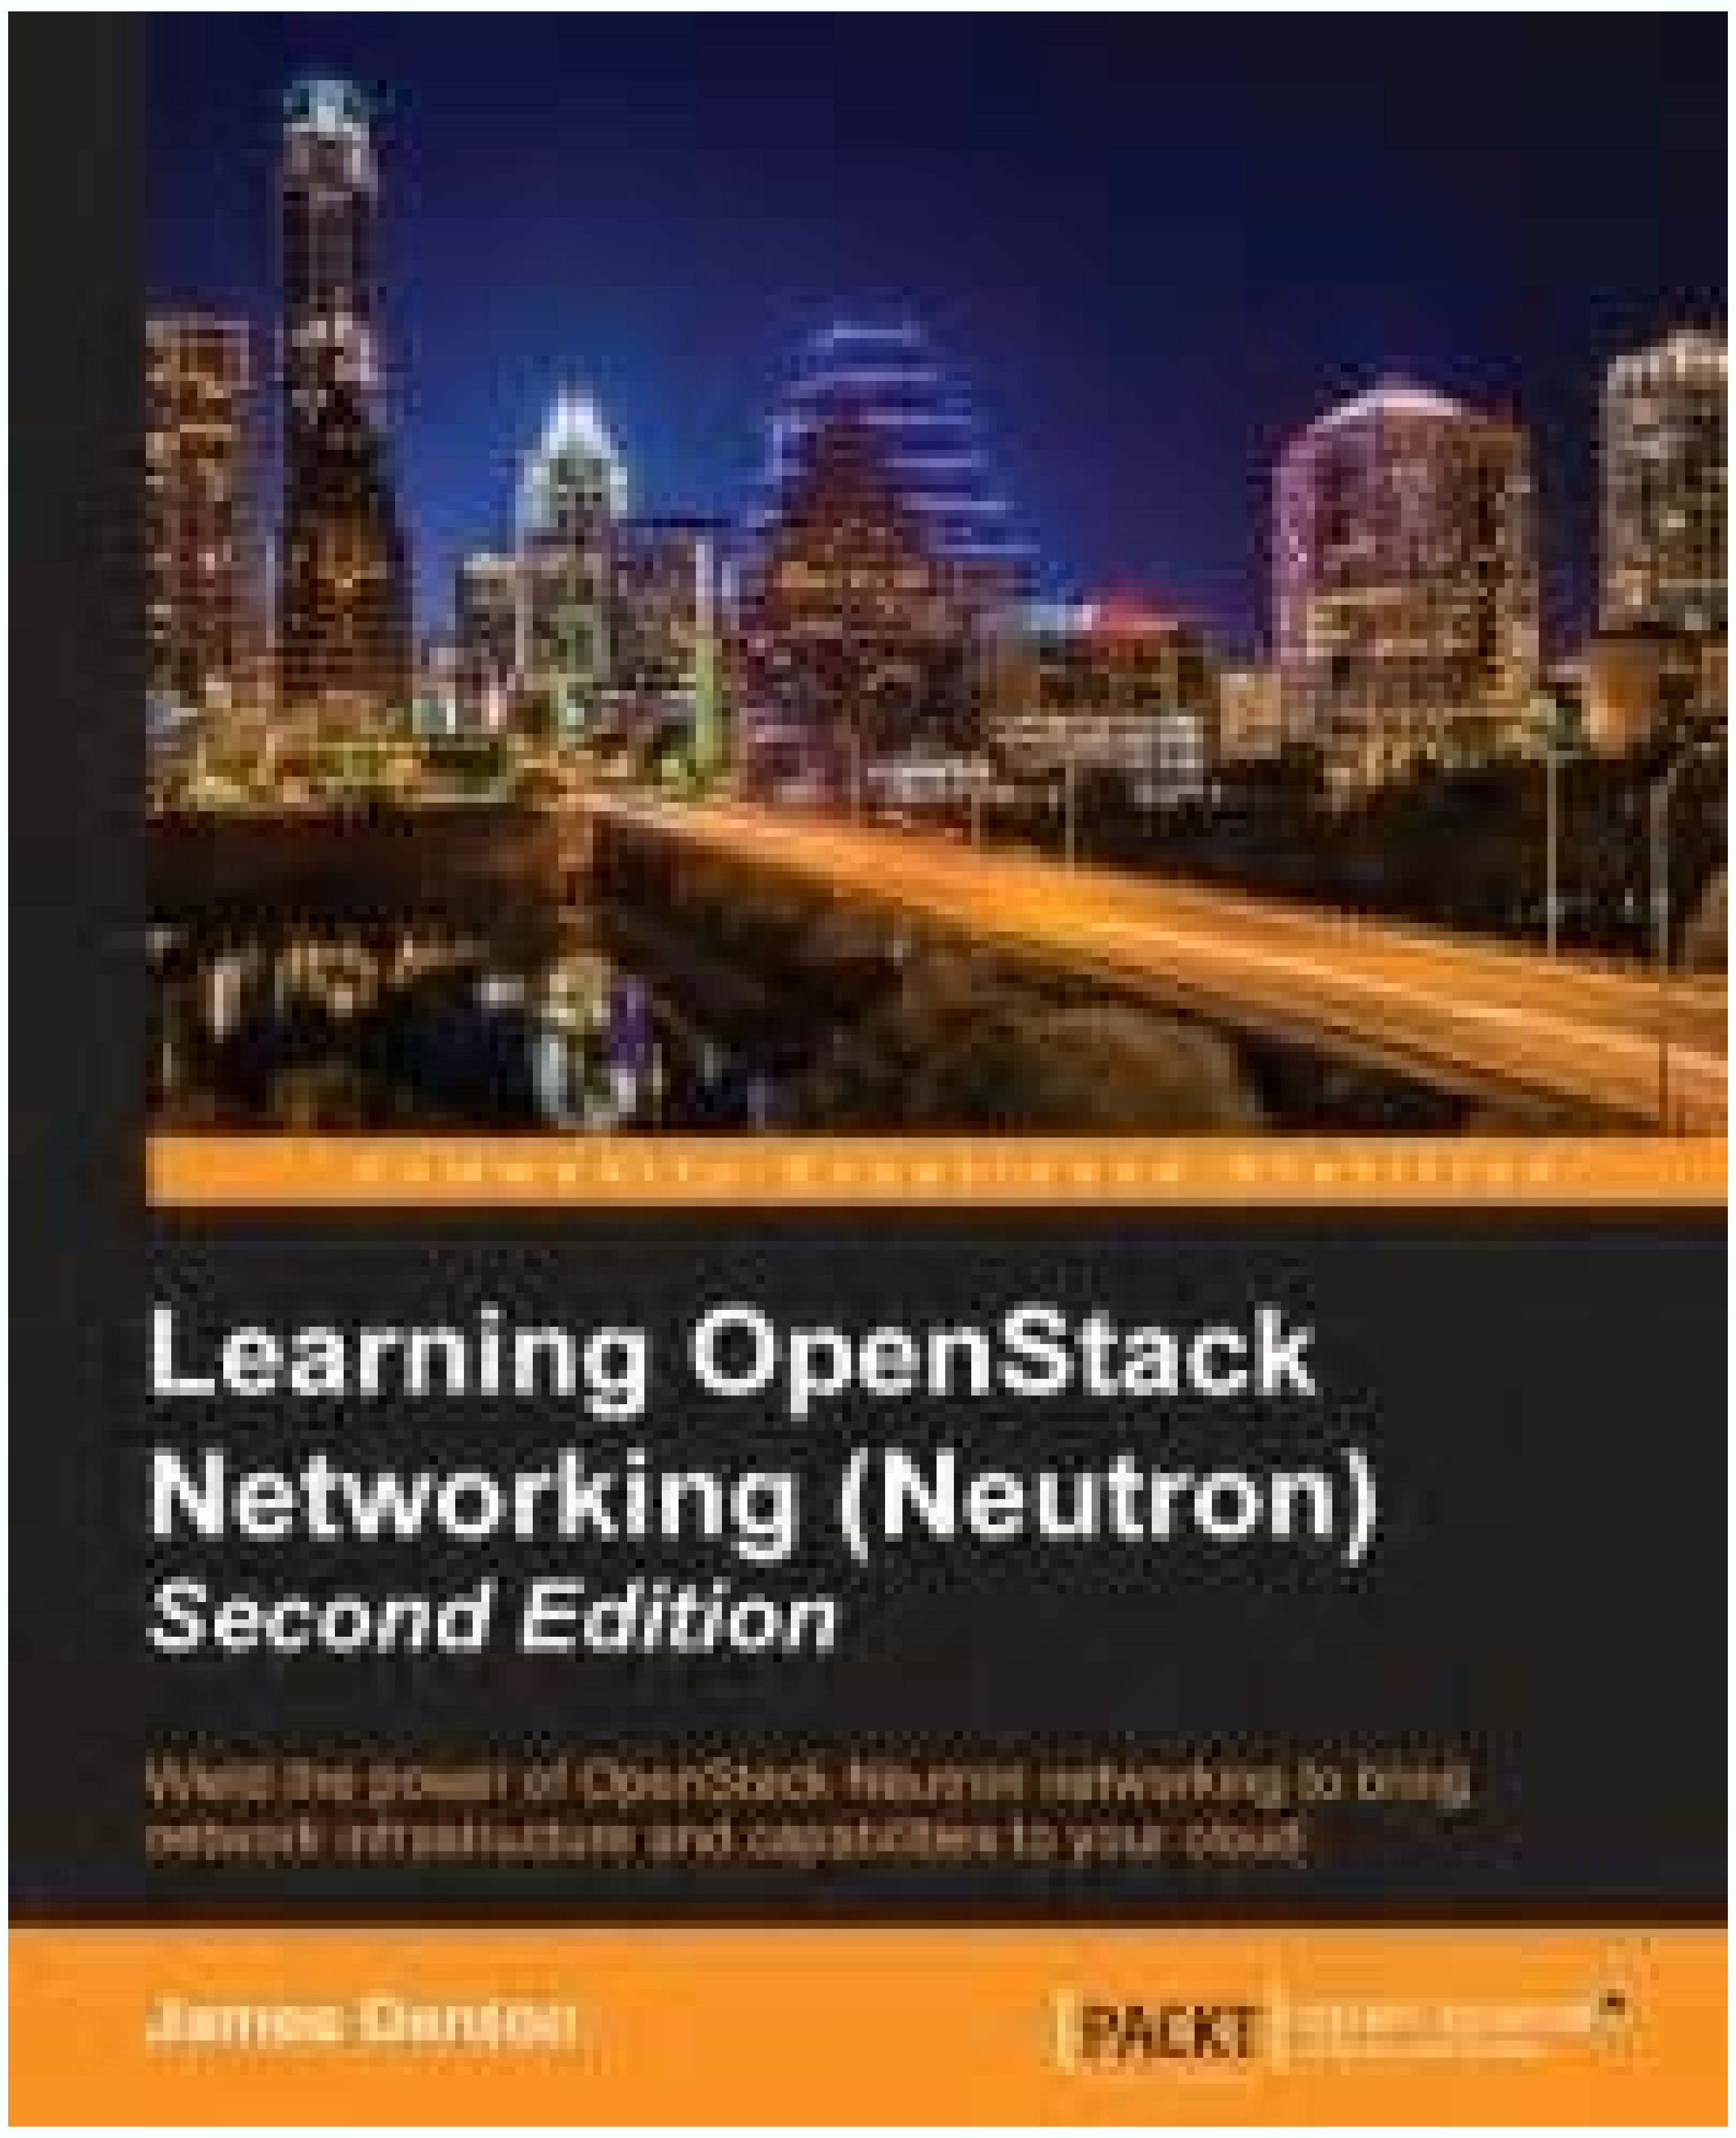
Community Experience Distilled
Learning OpenStack Networking (Neutron)
Second Edition
Wield the power of OpenStack Neutron networking to bring network infrastructure and capabilities to your cloud
James Denton
PACKT open source

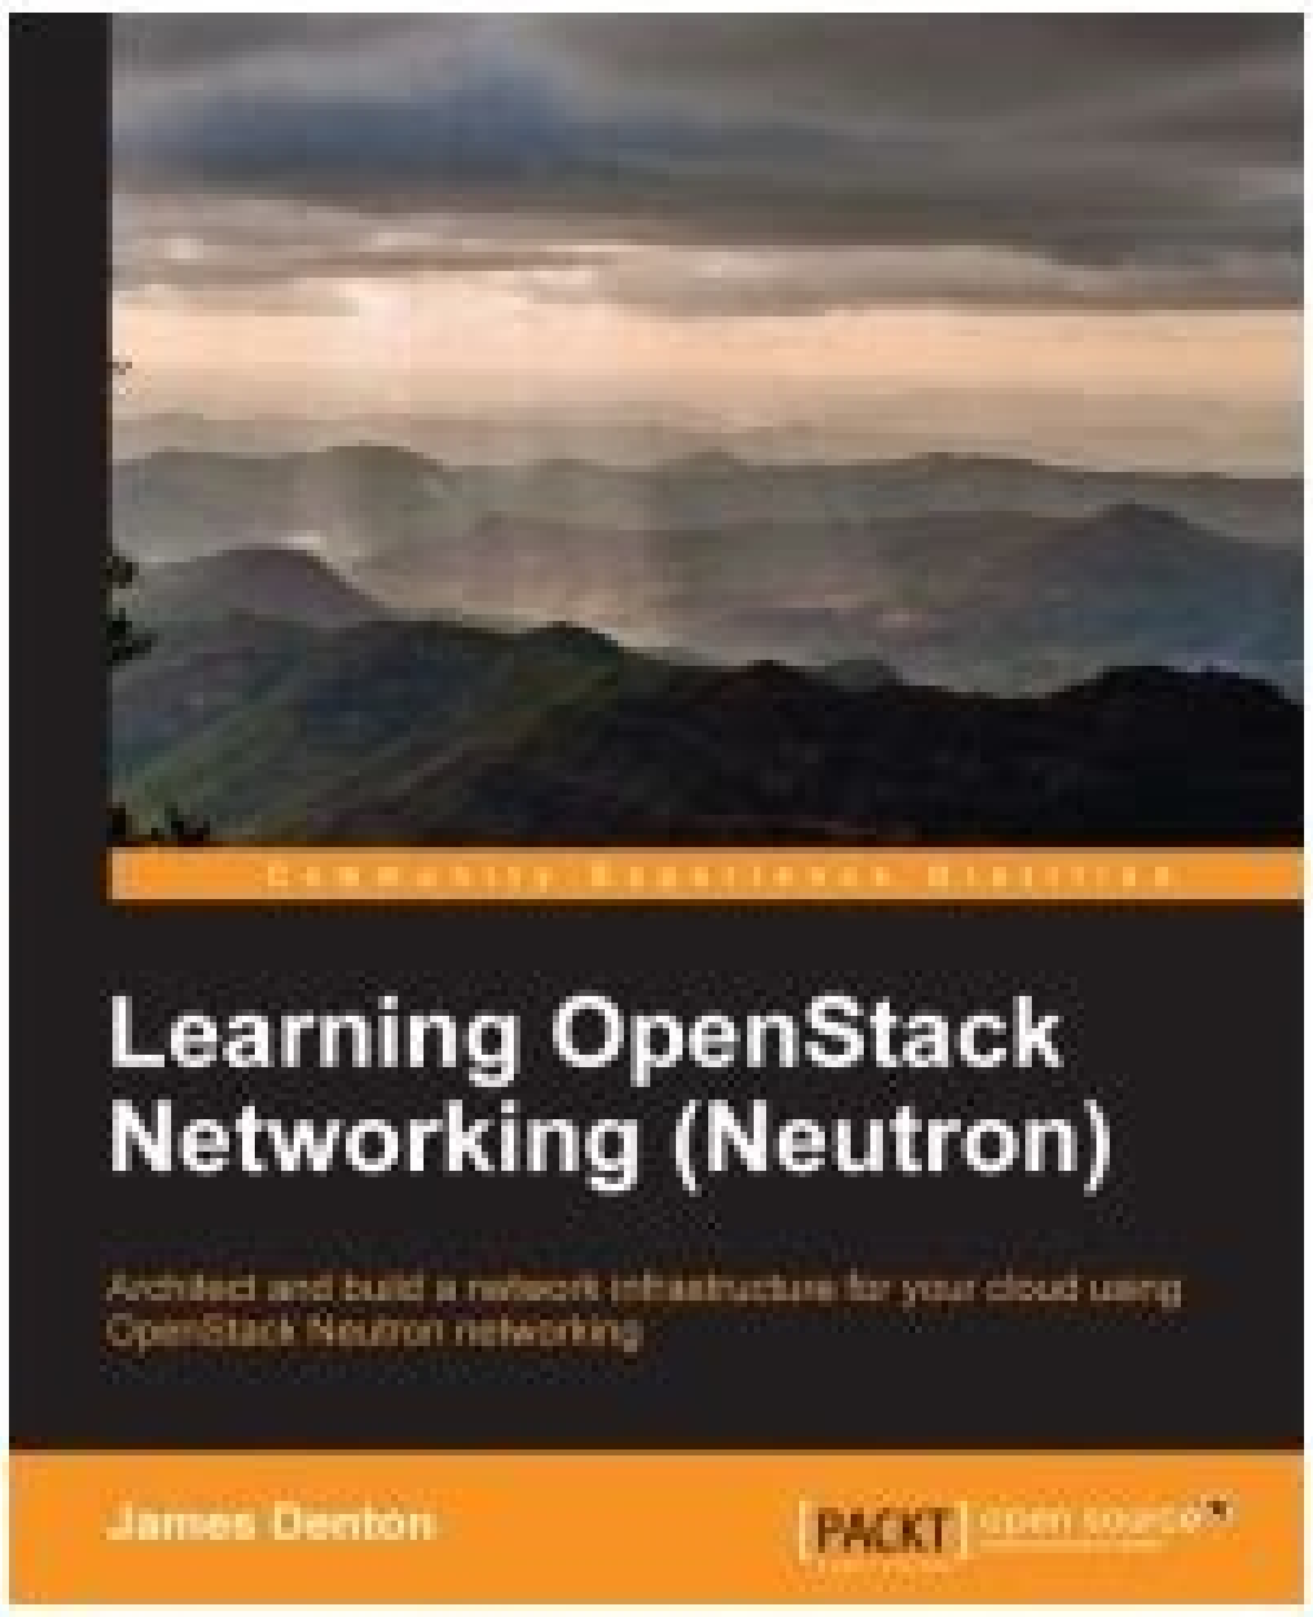
Community Experience Distilled
Learning OpenStack Networking (Neutron)
Architect and build a network infrastructure for your cloud using OpenStack Neutron networking
James Denton
PACKT open source

[52]72 Chapter 3 The qrouter namespace represents a router and routes traffic for instances in subnets that it is connected to. Showing health monitor details To show the details of a health monitor, use the Neutron lb-healthmonitor-show command as shown below: Syntax: lb-healthmonitor-show HEALTH_MONITOR The details returned include delay, expected codes, HTTP method, ID, max retries, pools, tenant ID, timeout, type, and URL path. Traffic that enters physical interface eth1 in the provider bridge is processed by the flow rules on that bridge. To create a local network, use the following syntax: Syntax: net-create --provider:network_type=local [--tenant-id TENANT_ID][--admin-state-down][--shared] NAME When using the LinuxBridge plugin, a bridge is created for the local network, but no physical or virtual VLAN interface is added. LoadBalancerPlugin Firewalling: neutron.services.firewall.fwaas_plugin. The inability to configure the physical infrastructure means that tenants should connect their networks to Neutron routers when external connectivity is required. For GRE packets, the KEY header field is used. The two plugins discussed in this book, LinuxBridge and Open vswitch, implement those features in different ways. Inter-VLAN routing, or routing between VLANs, is only possible through the use of a router. In the next chapter, you will be guided through the installation of Neutron networking services and provided with additional information about the underlying architecture of OpenStack Networking. The --shared flag is optional; it allows the policy to be shared amongst other tenants. For more information on user management in Keystone, please refer to the following URL: Finally, associate the admin role to the admin user when logging in with the admin tenant as follows: # keystone user-role-add --user=admin --tenant=admin --role=admin Define services and API endpoints in Keystone Each OpenStack service that is installed should be registered with Keystone, so its location on the network can be tracked. [15]35 Preparing the Network for OpenStack A single controller with one or more compute nodes In an environment consisting of a single controller and one or more compute nodes, the controller will likely handle all networking services and other OpenStack services, while the compute nodes strictly provide compute resources. Neutron refers to this type of behavior as Source NAT. Chapter 8 The POLICY keyword is used to represent the ID of the policy that should be applied to the firewall. Running iptables-save within a router namespace reveals the iptables rules in place. Traffic not matched by any rule is dropped by the neutron-linuxbri-sg-fallback chain: -A neutron-linuxbri-sg-fallback -j DROP Traffic exiting the tapc2a interface and headed towards an outside network is processed by the neutron-linuxbri-oc2a chain as follows: The first UDP rule allows the instance to send DHCP Discover and DHCP Request broadcast packets. In this installation, the username and password will be keystone: # crudini --set /etc/keystone/keystone.conf sql connection mysql:// Insecure passwords are used throughout the book to simplify the configuration and are not recommended for production use. Neutron provides a set of APIs to allow tenants to create IPSec-based VPN tunnels to remote gateways. Rather than being configured on every compute node, however, firewall rules are implemented using iptables within a Neutron router namespace. Flat networks are assigned a local VLAN ID in the Open vswitch database just like a VLAN network, and instances in the same flat network connected to the same integration bridge are placed in the same local VLAN. For most environments, I recommend the LinuxBridge approach unless integration with OpenFlow controllers or the use of a third-party solution or plugin is required. The following commands are used to manage health monitors in the CLI: lb-healthmonitor-create lb-healthmonitor-delete lb-healthmonitor-associate lb-healthmonitor-disassociate lb-healthmonitor-list lb-healthmonitor-show lb-healthmonitor-update Chapter 7 Creating a health monitor To create a health monitor, use the Neutron lb-healthmonitor-create command as follows: Syntax: lb-healthmonitor-create [--tenant-id TENANT_ID] [--admin-state-down] [--expected-codes EXPECTED_CODES] [--http-method HTTP_METHOD] [--url-path URL_PATH] --delay DELAY --max-retries MAX_RETRIES --timeout TIMEOUT --type {PING,TCP,HTTP,HTTPS} The --tenant-id flag is optional; it allows you to associate the monitor with the specified tenant. The allow_overlapping_ips configuration option specifies whether or not Neutron should allow tenant-created subnets to overlap one another. The --address attribute is required; it is used to specify the IP address of the pool member. As the system boots, these files are used to determine which interfaces to bring up and how they should be configured. Software routers created with Neutron reside on the controller node and handle routing between connected tenant networks [17]37 Preparing the Network for OpenStack A single controller plus network node with one or more compute nodes A network node is one that is dedicated to handling most or all OpenStack networking services, including the L3 agent, DHCP agent, metadata agent, and more. The DHCP driver is specified in the dhcp_agent.ini configuration file found at /etc/neutron/dhcp_ agent.ini. The --port-range-min flag is optional; it allows you to specify the starting port of a range of ports. Managing virtual IPs in the CLI The following commands are used to manage virtual IPs in the CLI: lb-vip-create lb-vip-delete lb-vip-list lb-vip-show lb-vip-update [200]220 Creating a virtual IP To create a virtual IP, use the Neutron lb-vip-create command as follows: Syntax: lb-vip-create [--tenant-id TENANT_ID] [--address ADDRESS] [--admin-state-down] [--connection-limit CONNECTION_LIMIT] [--description DESCRIPTION] --name NAME --protocol-port PROTOCOL_PORT --protocol {TCP,HTTP,HTTPS} --subnet-id SUBNET POOL The --tenant-id flag is optional; it allows you to associate the monitor with the specified tenant. The LinuxBridge and Open vswitch networking plugins and their respective Nova configuration changes will be discussed in further detail in Chapter 4, Building a Virtual Switching Infrastructure. Connections are initially balanced using the round robin algorithm and are then tracked in a table for future lookup with subsequent connections from the same IP address. A pool is a group of pool members that typically serve identical content. [140]160 Within the qrouter namespace, there exists a PREROUTING rule that redirects the HTTP request to a local listener at port 9697: Chapter 5 Using netstat within the namespace, you can see that there is a process that listens on port 9697: The listener in the preceding example is the Neutron metadata proxy service that, in turn, proxies the metadata request to the Nova metadata service: The DHCP namespace When instances are connected to a network that is not connected to a Neutron router, the instance must learn how to reach the metadata service. Additionally, running neutron help from the Linux command line provides a brief description of each command's function. Tunnel bridge The tunnel bridge is a virtual switch, similar to the integration and provider bridge, and is used to connect GRE and VXLAN tunnel endpoints. The use of ML2 solves this issue by creating a common schema for use by all plugins, not just LinuxBridge and Open vswitch: 1. Instead, the physical interface of the host associated with the network is placed directly in the bridge. In this installation, the iptables-based firewall will be used, and Neutron will handle the configuration of the rules on the hosts. The primary commands associated with router management include: router-create router-delete router-gateway-clear router-gateway-set router-interface-add router-interface-delete router-list router-list-on-l3-agent router-port-list router-show router-update Creating routers in the CLI Routers in Neutron are associated with tenants and are available for use only by users within the tenant that created them. Open vswitch relies on flow rules to determine how traffic in and out of the environment should be processed and requires both user-space utilities and kernel modules to perform such actions. Every effort has been made in the preparation of this book to ensure the accuracy of the information presented. The --disabled flag is optional; it allows you to specify whether or not the rule is inserted into the firewall. [16]36 Chapter 1 The following diagram demonstrates a controller node hosting all OpenStack management and networking services, including the Neutron L3 agent. As a result, Neutron is unable to enable VXLAN support using the Open vswitch kernel module. It is not possible to connect a subnet to more than one router at a time. log Using crudini, edit /etc/keystone/keystone.conf, and set the provider value to PKI: # crudini --set /etc/keystone/keystone.conf token provider keystone. [72]92 The following diagram provides a high-level view of a Linux bridge leveraged by Neutron: Chapter 4 eth0 eth0 Single IP address for MGMT & API eth0 MGMT & API Net VM 0 tap0 eth0 VM 1 K V M tap1 br-eht1 (Linux Bridge) eth1 External Networks eth0 tap2 VM 2 Figure 4.1 In the preceding figure, the Linux bridge br-eth1 contains a single physical interface (eth1) and three virtual interfaces: tap0, tap1, and tap2. Navigate to Project Manage Network Networks: From here, notice that there are no actions available next to the networks currently defined. The dhcp_domain configuration option specifies the DNS search domain that is provided to instances via DHCP when they obtain a lease. [174]194 Chapter 6 Attaching internal interfaces in the dashboard In order to attach internal interfaces in the dashboard, perform the following steps: 1. The tenant-id attribute specifies the tenant ID the subnet should be associated with. In the event that more than one VLAN network is needed, another Linux bridge will be created which contains a separate virtual VLAN interface. OpenStack Networking is a standalone service that can be installed independently of other OpenStack services. At the time of writing, however, VXLAN is not supported by the CentOS 6.5 kernel. Users can balance traffic to pools consisting of multiple application servers and can provide high availability of their application through the use of intelligent health monitors. A major limitation to LBaaS can be seen in the inability to create multiple virtual servers using the same IP address and different layer 4 ports. Virtual Ethernet (veth) cables are virtual interfaces that mimic network patch cables. A logical diagram of a load balancer in one-arm mode can be seen in the following diagram: Internet Neutron Router qg: qr: Float: Load Balancer Virtual IP :80 WEB :80 WEB :80 In the preceding diagram, a load balancer is configured in one-arm mode and resides on the same subnet as the servers it is balancing traffic to. Internally, however, Neutron treats flat networks like it does VLAN networks when programming the virtual switches. Flow rules for a particular network will not exist on a bridge if there are no instances or resources scheduled to that node in that network. Options for guest networks include local networks restricted to a particular node, flat or VLAN tagged networks, or the use of virtual overlay networks made possible with GRE or VXLAN encapsulation. Often, network namespaces will exist only on the controller or network nodes (if you have them). Configuring the LinuxBridge plugin Neutron was configured to use the LinuxBridge plugin at the end of the preceding chapter to allow you to access the Neutron command-line interface. This method of NAT allows instances to be reachable from external networks, such as the Internet. The qlbaas namespace represents a load balancer and might contain a load-balancing service, such as HAProxy, which load balances traffic to instances. Use the following commands to change enable_lb from false to true and to restart the Apache web service: # sed -i "/'enable_lb': False,/c\'enable_lb': True," /etc/openstackdashboard/local_settings # service httpd restart Load balancer management in the CLI Neutron offers a number of commands that can be used to create and manage virtual IPs, pools, pool members, and health monitors for load balancing purposes. The --ethertype flag is optional; it allows you to specify whether the rule applies to IPv4 or IPv6 traffic. These features can be configured to leverage open source or commercial software, and provide a cloud operator with all of the tools necessary to build a functional and self-contained cloud. Because of this designed limitation, local networks are recommended for testing purposes only. Both networking plugins are known as monolithic plugins, which means only one of them can be active at any given time. However, some Neutron configuration files must exist on all nodes, and the configuration files can only be installed via packages. Common types include GET and POST. When connecting two Open vswitch bridges, a port on each switch is reserved as a patch port. If using vxlan, set this option to vxlan. Have a look at the following screenshot: The CirrOS image is very limited in functionality and is recommended only for testing connectivity and operational Compute functionality. Chapter 8, Protecting Instances on the Network, will cover the creation and management of security-group rules to secure instance traffic. To clear the gateway of a router, use the router-gateway-clear command as follows: Syntax: router-gateway-clear Neutron includes checks that will prohibit the clearing of a gateway interface in the event that floating IPs or other resources from the network are associated with the router. Port number 5 is named qvo04c49e4a-a6 and corresponds to a Neutron port UUID starting with 04c49e4a-a6. Use crudini to add the bridge mapping to the Open vswitch plugin configuration file on all hosts as follows: # crudini --set /etc/neutron/plugins/openvswitch/ovs_neutron_plugin.ini OVS bridge_mappings physnet1:br-eth1 Configuring the bridges Before the Open vswitch plugin agent can be started, any bridge referenced in the bridge_mappings configuration must exist on the host. The following command will begin the MySQL installation and configuration process: # mysql_secure_installation During the MySQL installation process, you will be prompted to enter a password and change various settings. Then there is the next rule in the FORWARD chain that is processed: -A FORWARD -j neutron-linuxbri-forward This rule causes iptables to jump to the neutron-linuxbri-forward chain as follows: The -m flag followed by physdev is a directive to iptables to use an extended packet-matching module that supports devices enslaved to a bridge device. If the server were to send the response directly to the client, the client would reject the packet. As of this writing, both the Havana and Icehouse releases of OpenStack have a bug that does not allow Neutron to properly determine the version of the installed Open vswitch module in the CentOS and RHEL operating systems. The eth0 interface will serve as the management interface for OpenStack services and API access, and eth1 will serve as the provider bridge and tunnel interface for external and tenant traffic. [63]83 Installing Neutron Out of the box, Neutron utilizes dnsmasq, a free, lightweight DNS forwarder and DHCP server that is used to provide DHCP services to networks. Running instances can be attached to networks using the nova interface-attach command. [75]95 Building a Virtual Switching Infrastructure When GRE or VXLAN networks are created, a unique ID is specified that is used to encapsulate the traffic. A window will appear, allowing you to specify the details of the firewall, including the name, description, and associated policy: [249]269 Protecting Instances on the Network 8. 2. In a LinuxBridge-based network implementation, there are three distinct types of virtual networking devices: Tap devices VLAN interfaces Linux bridges A tap device is how a hypervisor such as KVM implements a virtual network interface card. All other traffic is then processed by the neutron-linuxbri-s c2a chain as follows: The rule above prevents an instance from performing IP and MAC address spoofing. Using the Neutron lb-vip-create command, create a virtual IP with the following attributes: Name: WEB_VIP Protocol Port: 80 Protocol: HTTP Subnet ID: Pool: WEB_POOL Have a look at the following screenshot: [207]227 Load Balancing Traffic in Neutron Once the virtual IP has been created, the state of the VIP and the pool will change to ACTIVE as shown in the following screenshot: The LBaaS network namespace A listing of the network namespaces on the host running the LBaaS agent reveals a network namespace that corresponds to the load balancer just created as shown in the following screenshot: The IP configuration within the namespace reveals a tap interface that corresponds to the subnet of the virtual IP as follows: Neutron creates a haproxy configuration file

specific to every load balancer that is created by users. Packets are processed by sequentially traversing rules in chains within the following tables: Raw: This is a default table that filters packets before any other table. As a result, the pool member would be eligible to receive connections. If issues arise during the installation and test procedures, log messages found in /var/log/nova/, /var/log/glance, /var/log/ httpd, and /var/log/keystone, among others, can be useful in determining and resolving the problem. The most common configuration options will be covered here.165 Creating Routers with Neutron Defining an interface driver Like previously installed agents, the Neutron L3 agent must be configured to use an interface driver that corresponds to the chosen networking plugin. Attaching a gateway interface in the dashboard In order to attach a gateway interface in the dashboard, perform the following steps: 1. An interface file should be created in /etc/sysconfig/ network-scripts/ for both the physical interface and the subinterface you are attempting to create. If this option is specified, a protocol must also be defined. A rule in a chain can cause a jump to another chain, and this behavior can be repeated to whatever level of nesting is required. On the controller node, restart the Neutron server API: # service neutron-server restart On all nodes, restart the LinuxBridge agent: # service neutron-linuxbridge-agent restart Any networks previously created under a monolithic plugin will need to be recreated, as a new database has been built for use with ML2. The relationship between these resources and instances, DHCP, and metadata services can be seen in the following sections. When specifying an endpoint, you must provide URLs for the public API, internal API, and the admin API. [226]246 Creating security group rules in the CLI To create a security group rule, use the Neutron security-group-rule-create command as follows: Chapter 8 The --direction flag is optional; it allows you to specify the direction of traffic that should be affected. The default maximum number of nameservers is five per subnet; this can be modified in the /etc/neutron/neutron.conf file. To update a network, use the Neutron net-update command as follows: Syntax: net-update [--router:external][--shared][--admin-state-up] [116]136 Chapter 5 Provider attributes are among those that cannot be changed once a network has been created. The CirrOS image has a built-in user named cirros with a password, cubswin:): Observe the routing table of each instance. Clicking on the Subnet tab allows you to specify information on the subnet, including the network address, CIDR, and gateway information: 4. Observe the route table within the namespace. Using the brctl show command, the preceding diagram can be realized in the Linux CLI as the following: The bridge id in the output is dynamically generated based on the parent NIC of the virtual VLAN interface. Each network uses its own bridge. [187]207 Load Balancing Traffic in Neutron Because a load balancer in one-arm mode is not the gateway for pool members it is sending traffic to, it must rely on the use of source NAT to ensure that return traffic from the members to the client is sent back through the load balancer. LinuxBridgeInterfaceDriver # crudini --set /etc/nova/nova.conf DEFAULT libvirt_vif_driver nova.virt. As mentioned earlier, the integration bridge is the virtual switch where all virtual machine VIFs, otherwise known as virtual network interfaces, are connected. Flow rules for VLAN networks In the following example, VLAN 30 represents a network in the data center and is trunked down to both the controller and compute nodes. ML2 configuration is outside the scope of this chapter, but I have provided an ML2 plugin file and configuration details relevant to this installation in Appendix B, ML2 Configuration. Both methods are explained in the upcoming sections. Cloud operators can configure network topologies by creating and configuring networks and subnets and then, instructing services such as Nova (Compute) to attach virtual devices to ports on these networks. This book is part of the Packt Open Source brand, home to books published on software built around Open Source licenses, and offering information to anybody from advanced developers to budding web designers. Click on Create Network in the upper right-hand corner of the screen. The following nova boot flag provides a network interface to an instance: --nic net-id= By passing the --nic flag multiple times, it is possible to attach an instance to more than one Neutron network. This verifies that the user account is established in Keystone with the expected credentials. [82]102 Chapter 4 Open vswitch Within OpenStack Networking, Open vswitch operates as a software-based switch that uses virtual network bridges and flow rules to forward packets between hosts. To create the rule, click on the blue Add button. Install the LinuxBridge and Open vswitch plugins with the following commands on all the nodes: # yum -y install openstack-neutron-linuxbridge # yum -y install openstack-neutron-openvswitch These two switching options will be discussed in further detail in Chapter 4, Building a Virtual Switching Infrastructure. When networks are created they are associated with an interface label, such as physnet1. The following diagram demonstrates the use of a second physical interface used for flat networks: Chapter 4 eth0 eth0 Single IP address for MGMT & API eth0 MGMT & API Net VM 0 eth0 VM 1 K V M tap0xxxx tap1xxxx brqyyyy (Linux Bridge) eth1 External Networks eth0 tap2xxxx brqwwww (Linux Bridge) eth2 VM 2 Figure 4.6 On the compute node, the use of two physical interfaces for separate flat networks can be realized as follows: With the two flat networks, the host does not perform any VLAN tagging on the traffic traversing those bridges. First the firewall rules must be created, and then a firewall policy and the firewall itself should be created. [86]106 Chapter 4 Port number 2 is named int-br-eth1 and is one end of a Linux veth cable. In an Open vswitch-based network implementation, there are five distinct types of virtual networking devices: Tap devices Linux bridges Virtual Ethernet cables OVS bridges OVS patch ports Tap devices and Linux bridges were described briefly in the previous section, and their use in an Open vswitch-based network remains the same. A window that lets us create a network will pop up: [120]140 Chapter 5 From the Network tab, you can define Network Name and Admin State (on or off). The default value is 0 (zero) and can be observed with the sysctl command: The value for net.ipv4.ip_forward should be set to 1 to enable IP forwarding. Have a look at the following command: # unset OS_SERVICE_TOKEN OS_SERVICE_ENDPOINT Once the environment variables are unset, it should be possible to use username-based authentication. [4]24 Preface Errata Although we have taken every care to ensure the accuracy of our content, mistakes do happen. The Linux kernel does not perform any VLAN tagging on the host. ML2 can be configured to use multiple layer 2 technologies simultaneously. This value must match the protocol of the associated pool. The interfaces used for the external and guest networks can be dedicated interfaces or ones that are shared with other types of traffic. Keystone should only be installed on the controller node along with python-keystoneclient: # yum -y install openstack-keystone python-keystoneclient [30]50 Chapter 2 Using crudini, configure Keystone to use MySQL as its database. Have a look at the following commands: # crudini --set /etc/nova/nova.conf database connection mysql:// # crudini --set /etc/nova/nova.conf DEFAULT auth_strategy keystone # crudini --set /etc/nova/nova.conf keystone_authtoken auth_host controller # crudini --set /etc/nova/nova.conf keystone_authtoken auth_protocol http # crudini --set /etc/nova/nova.conf keystone_authtoken auth_port # crudini --set /etc/nova/nova.conf keystone_authtoken admin_user nova # crudini --set /etc/nova/nova.conf keystone_authtoken admin_tenant_name service # crudini --set /etc/nova/nova.conf keystone_authtoken admin_password nova Next, configure Nova (Compute) to use the Qpid messaging broker configured on the controller node: # crudini --set /etc/nova/nova.conf DEFAULT rpc_backend nova.openstack. Because there are no rules defined in that chain, iptables returns to the calling chain, neutron-filter-top. Use the openstack-db command to create the Keystone database, related tables, and a database user named keystone that will be used by the Keystone service to connect to the database: # openstack-db --init --service keystone --password keystone You may be prompted to enter the password for the MySQL root user. For more information on configuring a Cisco Nexus 1000V with KVM in OpenStack, please refer to the official Cisco release guide at docs/switches/datacenter/nexus1000/kvm/config_guide/network/521sk122/ b-cisco-n1kv-kvm-virtual-network-config-521sk122.html. The --action flag is required; it allows you to specify the action that takes place when traffic matches the rule's criteria; possible options include allow or deny. In the next chapter, I will discuss OpenStack's load-balancing-as-a-service (LBaaS) solution that allows tenants to quickly scale their application while providing resiliency and availability. The instance processes metadata and continues the boot process [67]87 Installing Neutron For proper operation of metadata, both Neutron and Nova must be configured to communicate together with a shared secret password. Installing LBaaS The neutron-lbaas-agent service was installed as part of the overall Neutron installation process documented in Chapter 3, Installing Neutron, and is typically installed on a dedicated network node or a controller node. The Open Source brand also runs Packt's Open Source Royalty Scheme, by which Packt gives a royalty to each Open Source project about whose software a book is sold. The --protocol attribute is required; it is used to specify the type of traffic being load balanced. Neutron introduces support for third-party plugins that extend network functionality and implementation of the Neutron API. A window will pop up that resembles the one shown in the following screenshot: [214]234 Chapter 7 2. To install NTP, issue the following commands on all nodes in the environment: # yum -y install ntp # service ntpd start Unlike Ubuntu, the RHEL and CentOS operating systems do not automatically start services upon installation. Possible verdicts include the following: ACCEPT: This indicates that the packet is accepted and sent to the application for processing DROP: This indicates that the packet is dropped silently REJECT: This indicates that the packet is dropped and an error message is sent to the sender LOG: This indicates that the packet details are logged DNAT: This indicates that the destination IP of the packet is rewritten SNAT: This indicates that the source IP of the packet is rewritten RETURN: This indicates that the processing returns to the calling chain The ACCEPT, DROP, and REJECT verdicts are often used by the filter table. Implementing security group rules To demonstrate how security group rules are implemented on a compute node, check out the following WEB_SERVERS security group: [229]249 Protecting Instances on the Network In the following screenshot, you can see that two security group rules are being added that allow inbound connections on ports 80 and 443 from any host: Using the Neutron port-update command, I applied the WEB_SERVERS security group to the Neutron port of the Web1 instance, as shown in the following screenshot: Alternatively, the Nova client can be used to associate security groups to running instances using the following syntax: # nova add-secgroup The Nova client proxies security group commands to Neutron when security_group_api is equal to neutron in the nova.conf file. Network management Neutron provides users with the ability to execute commands from the CLI that interface with the Neutron API. As a result, GRE and VXLAN networks are effectively isolated from other networks in the cloud without the use of a Neutron router. A common starting value is 5 seconds. If required, additional interfaces can be used to further segregate traffic. When the L3 agent is used, an instance reaches the metadata service through the Neutron router that serves as its default gateway. 6. The ML2 plugin works with existing Open vswitch and LinuxBridge agents and is intended to replace the monolithic plugins associated with those agents. Book proposals should be sent to If your book idea is still at an early stage and you would like to discuss it first before writing a formal book proposal, contact us; one of our commissioning editors will get in touch with you. The Neutron database must be stamped as the havana release before neutronserver starts. # METADATA_SECRET=$(openssl rand -hex 10) On the controller, use the crudini utility to update the /etc/nova/nova.conf file with the shared secret and to enable the metadata proxy. If a name is not specified, the default value is blank or null. In Havana, if more than one provider network has the attribute set to true, then the gateway_external_ network_id configuration option must be used to associate an external network to the agent. Pool members default to an administrative up state. Do you want an environment that leverages everything OpenStack Networking has to offer in terms of routing, switching, and application networking? Configuring a layer 2 networking plugin Before you can start the neutron-server service and consume the Neutron API, a networking plugin must be defined. [9]29 Preparing the Network for OpenStack Preparing the physical infrastructure When architecting the network, it is important to first determine the purpose of the cloud. Updating networks in the CLI At times, it might be necessary to update the attributes of a network after it has been created. In order to support multiple routers with potentially overlapping IP networks, the neutron-l3-agent defaults to using network namespaces to provide isolated forwarding contexts. The --protocol flag is optional; it allows you to specify the type of traffic the rule applies to. openstack.common.rpc.impl_qpid The Qpid authentication settings should match what was previously configured for the other OpenStack services: # crudini --set /etc/neutron/neutron.conf DEFAULT qpid_hostname controller # crudini --set /etc/neutron/neutron.conf DEFAULT qpid_port 5672 # crudini --set /etc/neutron/neutron.conf DEFAULT qpid_username guest # crudini --set /etc/neutron/neutron.conf DEFAULT qpid_password guest Configuring a root helper The use of a root helper is a security mechanism built into OpenStack that prevents misuse of root privileges on the host that executes an OpenStack-related command. The --subnet-id attribute is required; it is used to provide the proper network configuration of the load balancer. The network_id field reveals the network to be 3b56346d-9f9a f1-4eb470cdad6d, otherwise known as MyFlatNetwork. On the controller node, install the messaging server: # yum -y install qpid-cpp-server memcached To simplify this installation, disable Qpid authentication by editing the /etc/qpidd. Showing the details of a firewall rule in the CLI To show the details of a firewall rule within the CLI, use the Neutron firewallrule-show command as follows: Syntax: firewall-rule-show FIREWALL_RULE_ID The returned output includes the name, description, action, destination IP address, destination port, source IP address, source port, associated firewall policy, position, protocol, and tenant ID of the specified firewall rule. Third-party drivers are supported by LBaaS but are outside the scope of this book.203 Load Balancing Traffic in Neutron Fundamentals of load balancing There are three major components to a load balancer in Neutron: Pool member(s) Pool(s) Virtual IP(s) A pool member is a layer 4 object and is composed of the IP address of a service and the listening port of the service. Appendix A, Additional Neutron Commands, will briefly cover additional Neutron functionality that is outside the scope of this book. Creating a flat network in the CLI If you recall from Chapter 4, Building a Virtual Switching Infrastructure, a flat network is a network in which no 802.1q VLAN tagging takes place. The FORWARD chain is used by the mangle and filter tables. The allowed-address-pairs Neutron extension is currently only supported by the ML2, Open vswitch, and VMware NSX plugins. Click on the blue Disassociate Floating IP button to proceed with the action as shown in the preceding screenshot. To update the nameservers in a subnet, use the Neutron subnet-update command and specify the new nameservers in a space-separated list with the dns_nameservers option, as seen in the following command: (neutron) subnet-update --dns-nameservers The enable-dhcp attribute is a Boolean value that, when set to true, enables DHCP services in the subnet. [13]33 Preparing the Network for OpenStack Multiple interfaces To reduce the likelihood of guest network bandwidth consumption affecting management of traffic and to maintain a proper security posture, segregation of traffic between multiple physical interfaces is recommended. When this is set to false, DHCP and metadata services are no longer available in the network. Once a firewall policy has been applied, the rules are immediately put in place on all routers that exist within the tenant. On all hosts, use the Open vswitch utility ovs-vsctl to create bridge br-int as follows: # ovs-vsctl add-br br-int You do not need to add an interface to the integration bridge, as Neutron is responsible for connecting network resources to this virtual switch. It is then up to the user to manually configure the IP address of the instance within the guest OS. If a rule does not match the packet, the packet is passed to the next rule. To create an additional network on the same physnet1 bridge, simply specify another segmentation ID. To update an attribute of a firewall rule in the CLI, use the Neutron firewall-rule-update command as follows: [241]261 Protecting Instances on the Network Creating a firewall policy in the CLI The next step in creating a firewall is to create a firewall policy that contains one or more firewall rules. For the installation of OpenStack, leave the value set to True. Users are able to define network connectivity in the cloud, and cloud operators are allowed to leverage different networking technologies to enhance and power the cloud. In this chapter, you will build OpenStack resources on top of this foundation. Each L2 agent configuration file, such as ovs_neutron_plugin.ini or linuxbridge_conf.ini, should set the actual value for the firewall_driver parameter for that agent. To fully utilize a provider network and its attributes, the external_network_ bridge configuration option should be set to an empty, or blank, value. A tenant has a simple Neutron network set up with a router attached to both an external provider network and an internal tenant network. No part of this book may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, without the prior written permission of the publisher, except in the case of brief quotations embedded in critical articles or reviews. More information on the use of metadata and this configuration can be found in Chapter 5, Creating Networks with Neutron. In the Add Policy window, the Name field has to be filled. All other traffic from the integration bridge on port 2 that is not tagged as VLAN 1 is dropped. [123]143 Creating Networks with Neutron The disable-dhcp attribute is a Boolean value that, when set to true, disables DHCP services for the subnet. While machines that fail to meet the minimum requirements may be capable of installation, based on the documentation included herein, these requirements are there to ensure a successful experience. In this installation, all guest traffic through overlay networks will traverse VLAN 20 using a virtual VLAN interface off eth1. As a result, a firewall policy with no rules blocks all traffic by default. The API service, neutron-server, remains on the controller node. The following is the recommended configuration for eth1.20 on the controller node: DEVICE=eth1.20 BOOTPROTO=none ONBOOT=yes IPADDR= NETMASK= VLAN=yes The following is the recommended configuration for eth1.20 on the compute node: DEVICE=eth1.20 BOOTPROTO=none ONBOOT=yes IPADDR= NETMASK= VLAN=yes To activate the changes, cycle the interfaces using the ifdown and ifup commands on each node: # ifdown eth1; ifdown eth1.20; ifup eth1; ifup eth1.20 [105]125 Building a Virtual Switching Infrastructure Issue a ping from the controller to the compute node using the x IP address to confirm connectivity: Troubleshoot any issues before proceeding with the creation of overlay networks. Both links are simultaneously cabled to a switch or pair of switches, but only one interface is active at any given time. [115]135 Creating Networks with Neutron Showing network properties in the CLI To list the properties of a network, use the Neutron net-show command as follows: Syntax: net-show The output of the command can be seen in the following screenshot: Information about the specified network, including the network type, provider bridge, segmentation ID, and more, can be observed in the net-show output. Managing health monitors in the CLI LBaaS in Neutron provides the ability to monitor the health of pool members as a method of ensuring the availability of an application. A new port that has a unique MAC address and an IP from the specified network is created in Neutron. Inside the router namespace, a new interface has been added with a preface of qr. By laying down a basic installation of OpenStack, the reader should be able to follow the examples laid out in the book to receive a functional understanding of the various components of OpenStack Networking. Using a text editor, create the file /etc/sysconfig/network-scripts/ifcfgeth1.20 on each host. To retrieve a list of all Neutron ports, use the Neutron port-list command, as shown in the following screenshot: Using the Neutron port-show command, it is possible to determine the details of a particular port: The port here is owned by an interface used by a DHCP namespace. The integration bridge is connected to the provider bridge via an OVS patch cable. Neutron provides an API to create, modify, apply, and delete security group rules.241 Protecting Instances on the Network When a port is created in Neutron, it is associated with a default security group unless a specific security group is specified. Neutron is able to isolate these objects through the use of network namespaces. Not recommended for production environments, this type of configuration should only be used for testing or proof-of-concept. OpenSSL can be used to generate a random token and store it in the configuration file: # ADMIN_TOKEN=$(openssl rand -hex 10) # crudini --set /etc/keystone/keystone.conf DEFAULT admin_token $ADMIN_ TOKEN By default, Keystone uses PKI tokens for authentication. There is no validation of this attribute, which may allow users to create monitors that don't work as expected. conf, and change the Listen address to : # sed -i 's/listen 80/Listen :80/' /etc/httpd/conf/httpd. At this point, the IP configuration is complete. A feature was added to the Havana release of Neutron that addresses the need to have multiple IP or MAC addresses on a single interface within an instance. The DHCP server sends other DHCP options such as nameservers, routes, and so on to the instance. When executed by an administrator, Neutron will return a listing of all routers across all tenants unless, the tenant ID is specified. To determine the number of host IP addresses available in /24, subtract 24 from 32 or from the total number of bits in an IPv4 address. Like other OpenStack projects, FWaaS will become more mature in future releases. In the promiscuous mode, the interface allows all frames through, thus allowing the host to see and process frames intended for other machines or network devices. Configuring neutron-server The neutron-server service exposes the Neutron API to users and passes all calls to the appropriate Neutron plugin for processing. This type of monitor is commonly referred to as a half-open TCP monitor. In the following example, VLANs 30 through 33 are available for tenant network allocation: network_vlan_ranges = physnet1:30:33 Non-contiguous VLANs can be allocated by using a comma-separated list as follows: network_vlan_ranges = physnet1:30:33,physnet1:50:55,physnet1:66:70 The network_vlan_ranges configuration option must be configured for the Neutron plugin agent to load properly. [240]260 Deleting a firewall rule in the CLI To delete a firewall rule in the CLI, use the Neutron firewall-rule-delete command as follows: Syntax: firewall-rule-delete FIREWALL_RULE_ID The keyword FIREWALL_RULE_ID is used to represent the ID of the firewall rule to be deleted. [91]111 Building a Virtual Switching Infrastructure In another example, a flat network has been added in Neutron that has no VLAN tag, as follows: On the physical switch, this network is configured as the native VLAN (untagged) on the switch port connected to eth1 of compute01. Tenants are prevented from attaching instances directly to external provider networks. Neutron-rootwrap looks for filter definition directories within the configuration file and loads command filters from them. Every Open vswitch database is independent of another and manages its own VLAN database that is not related to the physical network infrastructure. Instances in the same Neutron network on a particular host are placed in the same VLAN on the integration bridge. In this bridge, the parent interface is eth1. Instances can be limited to a subset of addresses in the subnet so that they can coexist with devices outside of the OpenStack cloud. There are no audit logs or auditing mechanisms within Neutron. Instances connected to the two bridges require a router to communicate with one another. The port must be one that is not currently associated with any other instance or resource. As a result of the bridge not being fully managed by OpenStack, provider attributes of the network created within Neutron, including the segmentation ID, network type, and the provider bridge itself, are ignored. If 200 is listed as expected code, the monitor would mark the pool member as UP. For installations using Neutron instead of nova-network, this option should be set to neutron as follows: # crudini --set /etc/nova/nova.conf DEFAULT security_group_api neutron [60]80 Nova (Compute) requires additional configuration once a networking plugin has been determined. The installation and configuration of OpenStack Networking services can be found in Chapter 3, Installing Neutron. Instead, the Open vswitch plugin agent programs flow rules on the virtual switches that dictate how traffic traversing the switch should be manipulated before forwarding. Once the gateway interface has been added, the router will be scheduled to an eligible L3 agent. In the next chapter, you will be guided through a package-based installation of OpenStack on the CentOS operating system. This is usually A DHCP server responds to the request with a DHCPOFFER packet. To test Keystone, issue the following commands: # source ~/credentials/admin # keystone token-get # keystone user-list [35]55 Installing OpenStack Keystone should return the token and user list as requested as follows: Installing and configuring the image service Glance is the image service for OpenStack. Attaching instances to networks using Nova boot Instances are attached to networks at the first boot through the nova boot command, the dashboard, or the Nova API. The qlbaas namespace is connected to a virtual switch and can communicate with other resources in the same network or subnet. The --type attribute is required; it is used to specify the type of monitor being configured. The use of the ML2 plugin is not required in Havana, but for your reference, its configuration has been provided in Appendix B, ML2 Configuration. [261]281 282 ML2 Configuration The Modular Layer 2 (ML2) plugin is a framework that allows OpenStack Networking to simultaneously utilize a variety of layer 2 networking technologies that are found in data centers. Click on Network Topology under the Project tab to find a logical diagram based on the networks, router, and instances created earlier. VMware/Nicera command reference OpenStack Networking supports VMware NSX and Nicera NVP through the use of API extensions and plugins. When set, a pool member may receive more or less traffic than other members in the same pool. By default, use_namespaces is set to True. Using the HTTP_COOKIE persistence type configures haproxy with the following settings within the backend pool configuration: cookie SRV insert indirect nocache The first time a client connects to the virtual IP, haproxy balances the connection to the next pool member in line. The LinuxBridge and Open vswitch plugins are deprecated in Icehouse in favor of the Modular Layer 2 (ML2) plugin, which allows for the use of multiple Layer 2 plugins simultaneously. Many third-party vendors offer downloadable cloud software based on OpenStack that provide deployment and management strategies using Chef, Puppet, Fuel, Ansible, and other tools. After this, place the user in the service tenant, and give the user the admin role: # keystone user-create --name=nova --pass=nova learningneutron.com # keystone user-role-add --user=nova --tenant=service --role=admin Configure Nova (Compute) to use the following Keystone credentials: # crudini --set /etc/nova/nova.conf DEFAULT auth_strategy keystone # crudini --set /etc/nova/nova.conf keystone_authtoken auth_host controller # crudini --set /etc/nova/nova.conf keystone_authtoken auth_protocol http # crudini --set /etc/nova/nova.conf keystone_authtoken auth_port # crudini --set /etc/nova/nova.conf keystone_authtoken admin_user nova # crudini --set /etc/nova/nova.conf keystone_authtoken admin_tenant_name service # crudini --set /etc/nova/nova.conf keystone_authtoken admin_password nova [41]61 Installing OpenStack Credentials must be added to the /etc/nova/api-paste.ini file that corresponds to the details of this build. Start the openvswitch service on all hosts to proceed with the bridge configuration and configure it to start at boot: # service openvswitch start # chkconfig openvswitch on On all hosts, use the Open vswitch utility ovs-vsctl to create bridge br-eth1 as follows: # ovs-vsctl add-br br-eth1 Use the same command to add the physical interface eth1 to the bridge: # ovs-vsctl add-port br-eth1 eth1 The physical switch port connected to eth1 must support 802.1q VLAN tagging if VLAN networks of any type are to be created. An external router is needed to handle routing between network segments. In this chapter, you will be introduced to two networking plugins available in the Havana release of OpenStack: LinuxBridge and Open vswitch (OVS). All security groups should be removed from the port and then selected groups can be added back. When networks are created, they are associated with an interface label, such as physnet1. Our unique business model allows us to bring you more focused information, giving you more of what you need to know, and less of what you don't. A VLAN provider network will be used as the external gateway network and a VLAN tenant network will be used for instances. If run as an administrator, all security groups across all tenants will be listed. Using the Admin tab as an administrator In order to create a network in the dashboard as the cloud administrator, perform the following steps: 1. The LinuxBridge plugin agent will create a bridge and place only the tap interface of the instance in the bridge. More information on the use of load balancers within Neutron can be found in Chapter 7, Load Balancing Traffic in Neutron. What this book covers Chapter 1, Preparing the Network for OpenStack, will provide an introduction to OpenStack Networking that includes a description of the different supported networking technologies, and it will explain how to architect the physical network to support an OpenStack cloud. The LinuxBridge and Open vswitch agents continue to rely on their respective configuration files, which were configured previously in this book. The use of a dedicated network node provides additional security and resilience, as the controller node will be at less risk of network and resource saturation. Instead, you must perform the following steps: 1. For information on configuring ML2 with the Open vswitch agent, please refer to the following URL: Restarting Neutron services Neutron services must be restarted before the aforementioned changes can take effect. The syntax to create a flat network can be seen here: Attributes in the [] brackets are considered optional and are not required to create the network. This persistence method is recommended over source IP persistence, as it is not reliant on the IP address of the client. In Havana, the command results in the route being added to the database and router show output, while not being added to the routing table. This action results in the host exposing the guest instance to the physical network. Details returned include the admin state, description, ID, load balancing method, members, protocol, provider, status, subnet ID, tenant ID, VIP ID, and health monitors associated with the pool. HTTPS: This instructs the monitor to initiate an HTTPS request to a pool member based on the expected_codes, url_path, and http_method attributes described here. [94]114 LinuxBridge plugin configuration options Once installed, the configuration file for the LinuxBridge plugin can be found at /etc/neutron/plugins/linuxbridge/linuxbridge_conf.ini. While you can create multiple subnets with the same name, it is recommended that subnet names remain unique for easy identification. Let us know what you think about this book what you liked or may have disliked. The cisco-policy-profile commands enable you to list and show details of Cisco Nexus 1000V policy profiles, as well as associate or disassociate profiles with tenants. INPUT: This is used when a packet is going to be locally delivered (that is, meant for the host machine). Using crudini, set the DHCP interface_driver option to use neutron.agent. A security group must be removed from all ports before it can be deleted. In the event of connectivity loss, out-of-band management access to the servers via DRAC, ilo, or some other mechanism, is highly recommended. The first two rules specify a particular inbound port: in_port=2 According to the diagram in Figure 4.10, traffic entering the bridge br-eth1 from physical interface eth1 does so through port 1, not port 2, so the first two rules do not apply. LoadBalancerPlugin,neutron.services.firewall.fwaas_plugin.FirewallPlugin [238]258 Chapter 8 Save your changes and restart neutron-server and neutron-l3-agent: # service neutron-server restart # service neutron-l3-agent restart Enabling FWaaS in the dashboard To enable the management of firewall resources in the Horizon dashboard, the enable_firewall parameter must be set to True in the /etc/openstackdashboard/local_settings configuration file. The --name attribute is used to specify a name for the pool. [119]139 Creating Networks with Neutron Using the Project tab as a user As a normal user, networks are created under the Project tab in the dashboard. The namespace is able to communicate with other devices in the same subnet through the bridge. Neutron is capable of much more, provided the appropriate extension or plugin is installed. Use the export command to export the variables and their values to your environment. I'd like to thank my wife, Amanda, for providing encouragement and patience throughout the writing of this book. The address is then used as the default gateway for instances in the subnet. Click on the Create Security Group button in the upper right-hand corner of the screen. Depending on your network infrastructure, this might not be an easy change to implement. Use the neutron-db-manage command to accomplish this task on the controller only using the following command: # neutron-db-manage --config-file /etc/neutron/plugin.ini --config-file / etc/neutron/neutron.conf stamp havana If you previously used LinuxBridge and are switching to Open vswitch, be sure to remove the symbolic link pointing to the LinuxBridge plugin configuration file prior to creating a new one for Open vswitch. Step-by-step instructions on setting up and running the leading open source cloud platform, CloudStack. When Neutron is used, this option should be set to nova.virt.firewall. In this installation, the configured interface_driver should correspond to the LinuxBridge plugin. To overwrite an existing gateway address, use the Neutron subnet-update command and specify a new value for gateway_ip, as shown in the following command: (neutron) subnet-update --gateway_ip= To completely remove a gateway address from the subnet, use the action=clear directive as follows: (neutron) subnet-update --gateway_ip action=clear Not all commands support the action=clear directive to destroy option values. The four types are as follows: PING: The simplest of all monitor types, PING uses ICMP to confirm connectivity to pool members. This scenario can wreak havoc on the connectivity and routing behavior of an instance. To install Glance, run the following command from the controller node: # yum -y install openstack-glance You can use openstack-db to initialize the Glance database and add the glance user to MySQL: # openstack-db --init --service glance --password glance Use crudini to set the SQL connection string in the Glance configuration files: # crudini --set /etc/glance/glance-api.conf DEFAULT sql_connection # crudini --set /etc/glance/glance-registry.conf DEFAULT sql_connection You can then add the glance user to Keystone and create the appropriate role: # keystone user-create --name=glance --pass=glance learningneutron.com # keystone user-role-add --user=glance --tenant=service --role=admin Use crudini to set Keystone attributes in the Glance configuration files: # crudini --set /etc/glance/glance-api.conf keystone_authtoken auth_host controller # crudini --set /etc/glance/glance-api.conf keystone_authtoken admin_ user glance # crudini --set /etc/glance/glance-api.conf keystone_authtoken admin_ tenant_name service # crudini --set /etc/glance/glance-api.conf keystone_authtoken admin_ password glance [36]56 Chapter 2 # crudini --set /etc/glance/glance-registry.conf keystone_authtoken auth_ host controller # crudini --set /etc/glance/glance-registry.conf keystone_authtoken admin_user glance # crudini --set /etc/glance/glance-registry.conf keystone_authtoken admin_tenant_name service # crudini --set /etc/glance/glance-registry.conf keystone_authtoken admin_password glance Glance includes default configuration files that should be copied and modified as follows: # cp /usr/share/glance/glance-api-dist-paste.ini /etc/glance/glance-apipaste.ini # cp /usr/share/glance/glance-registry-dist-paste.ini /etc/glance/glanceregistry-paste.ini Each of the files mentioned in the preceding commands must then be edited to add the following options: # crudini --set /etc/glance/glance-api-paste.ini filter:authtoken auth_ host controller # crudini --set /etc/glance/glance-api-paste.ini filter:authtoken admin_ user glance # crudini --set /etc/glance/glance-api-paste.ini filter:authtoken admin_ tenant_name service # crudini --set /etc/glance/glance-api-paste.ini filter:authtoken admin_ password glance # crudini --set /etc/glance/glance-api-paste.ini filter:authtoken flavor keystone # crudini --set /etc/glance/glance-registry-paste.ini filter:authtoken auth_host controller # crudini --set /etc/glance/glance-registry-paste.ini filter:authtoken admin_user glance # crudini --set /etc/glance/glance-registry-paste.ini filter:authtoken admin_tenant_name service # crudini --set /etc/glance/glance-registry-paste.ini filter:authtoken admin_password glance # crudini --set /etc/glance/glance-registry-paste.ini filter:authtoken flavor keystone [37]57 Installing OpenStack You can start the Glance services and enable them to start at boot time with the following command: # service openstack-glance-api start # service openstack-glance-registry start # chkconfig openstack-glance-api on # chkconfig openstack-glance-registry on Define the Glance service and API endpoints in Keystone Like other OpenStack services, Glance should be added to the Keystone database using the service-create and endpoint-create commands: # keystone service-create --name=glance --type=image --description="glance Image Service" The resulting output is as follows: Property Value description Glance Image Service id bbbacfbe630341b181659f00a2ef6a90 name glance type image The id value here should be used to populate the service-id value as follows: # keystone endpoint-create \ \ --service-id=`keystone service-get glance awk '/ id / { print $4 }'` --publicurl= \ --internalurl= \ --adminurl= [38]58 Chapter 2 The resulting output is as follows: Property Value adminurl id c4cc4661a04adbf1dfb97c08 internalurl publicurl region regionone service_id bbbacfbe630341b181659f00a2ef6a Verify the Glance image service installation To verify that Glance was installed and configured properly, download a test image from the Internet, and verify that it can be uploaded to the image server: # mkdir /var/tmp/images ; cd /var/tmp/images/ # wget disk.img Upload the image to Glance using the following command: # glance image-create --name=cirros disk-format=qcow2 --container-format=bare --is-public=true --file /var/tmp/images/cirros x86_64-disk.img Verify that the image exists in Glance using the image-list command. Management network The management network is used for internal communication between hosts for services, such as the messaging service and database service. Upgrading the system Before installing OpenStack, it is imperative that the kernel and other system packages on each node be upgraded to the latest version supported by CentOS 6.5. Issue the following yum command on each node, followed by a reboot to allow the changes to take effect: # yum -y upgrade # reboot [28]48 Chapter 2 Installation of OpenStack The steps in the later part of the chapter document the installation of OpenStack services, including Keystone, Glance, Nova Compute, and Horizon, on a single controller and compute node. Using crudini, set the tenant_network_type option to vlan on all hosts: # crudini --set /etc/neutron/plugins/openvswitch/ovs_neutron_plugin.ini OVS tenant_network_type vlan If at any time you wish to change tenant_network_type to something other than vlan, edit the plugin configuration file appropriately on all nodes, and restart the Open vswitch plugin agent.121 Building a Virtual Switching Infrastructure Network VLAN ranges The network_vlan_ranges configuration option defines a range of VLANs that tenant networks will be associated with upon their creation when tenant_network_ type is set to vlan. A virtual IP is usually exposed to the Internet and is often mapped to a domain name. Using the Admin tab as an administrator To create a subnet as the cloud administrator, perform the following steps: 1. In this guide, the password was set to openstack. The network architecture within the compute node would resemble the following diagram: eth0 eth0 Single IP address for MGMT/ API eth0 MGMT & API Net VM 0 tap0xxxx eth0 VM 1 K V M tap1xxxx brqyyyy (Linux Bridge) eth eth1 External Networks eth0 tap2xxxx VM 2 Figure 4.3 [77]97 Building a Virtual Switching Infrastructure In the preceding diagram, three guest instances are connected to a Linux bridge, named brqyyyy, via their tap interfaces. Deleting a health monitor To delete a health monitor, use the Neutron lb-healthmonitor-delete command as follows: Syntax: lb-healthmonitor-delete HEALTH_MONITOR The keyword HEALTH_MONITOR represents the ID of the health monitor to be deleted. Users creating networks within the dashboard are required to create a subnet at the time the network is created. Setting the external bridge The L3 agent must be aware of how to connect the external interface of a router to the network. Immediately following the third failure, the load balancer marks the pool member as DOWN, as follows: Subsequent connections to the VIP are sent to WEB2 as follows After restarting the web service on Web1, the load balancer places the server back in the pool upon the next successful health check as follows: Connecting to the virtual IP externally To connect to a virtual IP externally, a floating IP must be associated with the Neutron port associated with the VIP, as the virtual IP exists within a subnet behind the router and is not reachable directly. Have a look at the following screenshot: The result is that incoming traffic is tagged as VLAN 2 and forwarded to instances connected to the integration bridge that reside in VLAN 2. To add a subnet to the network, click on the Create Subnet button on the right-hand side: Chapter 5 4. However, the bridge interface mapped to the label may be different. The rule may also contain a target (another chain) or a verdict, like DROP or ACCEPT. It also triggers the injection of a metadata route via DHCP when enable_isolated_ metadata is set to true in the DHCP configuration file. There are cases, however, when multiple IP addresses must exist on a single interface in the guest OS. Create, deploy, and develop applications using CloudBees. There are no flow rules created for local networks. The community is looking to address many functional and performance concerns in the Juno release of OpenStack, and beyond. To attach internal interfaces to routers in the dashboard, click on the router to reveal the Router Details page as follows: 2. [198]218 Chapter 7 TCP: This instructs the load balancer to send a TCP SYN packet to the pool member. The subnet specified using the --subnet-id attribute should match the subnet of the pool members to be added to the pool. A window will appear that will allow you to create a security group: 2. By defining the appropriate CIDR notation, it is possible to permit entire subnets rather than an individual host IP. Neutron provides an API to manage virtual IPs, pools, pool members, and health monitors. To enable or disable DHCP in a subnet, use the Neutron subnet-update command and specify true or false as follows: (neutron) subnet-update --enable-dhcp= The gateway_ip attribute sets the default gateway for the subnet. Here, you can access, read and search across Packt's entire library of books. Neither the author, nor Packt Publishing, and its dealers and distributors will be held liable for any damages caused or alleged to be caused directly or indirectly by this book. The following attributes, however, can be modified: router:external shared admin-state-up The router:external attribute is a Boolean value that, when set to true, allows the network to be utilized as a gateway network for Neutron routers. To get a better understanding of how this feature works, observe the standard iptables rules on a compute node that disallows traffic from a particular instance if it is not sourced appropriately. Simply use your login credentials for immediate access.9 10 I dedicate this book to the memory of my grandfather, a proud Aggie whose curiosity in all things, including technology, helped form my identity and career. You can update multiple attributes simultaneously, as shown in the following screenshot: [259]279 Additional Neutron Commands Listing tenant quotas To list the quotas of a tenant, use the Neutron quota-list command as follows: Syntax: quota-list --tenant-id If a tenant is using default quotas, no output will be provided. Summary This chapter laid the foundation for creating networks and subnets that can be leveraged by routers, instances, and other cloud resources. The username and password for the neutron user in Keystone were set earlier in the chapter: # crudini --set /etc/neutron/neutron.conf keystone_authtoken auth_host controller # crudini --set /etc/neutron/neutron.conf keystone_authtoken auth_port # crudini --set /etc/neutron/neutron.conf keystone_authtoken auth_ protocol http # crudini --set /etc/neutron/neutron.conf keystone_authtoken admin_ tenant_name service # crudini --set /etc/neutron/neutron.conf keystone_authtoken admin_user neutron # crudini --set /etc/neutron/neutron.conf keystone_authtoken admin_ password neutron The /etc/neutron/api-paste.ini middleware configuration file must be edited to contain the appropriate authentication settings for the environment. The physical_network provider attribute defines the physical interface that will be used to forward traffic through the host. From the list of rules, you can choose from a predefined list of protocols or create a custom rule, as follows: [236]256 7. The most common configuration options will be covered in the upcoming sections. In the next chapter, you will be guided through the configuration of the LinuxBridge and Open vswitch monolithic networking plugins. Providing layer 2 connectivity to instances Neutron and Nova work in tandem to configure networking on physical servers in the cloud. Neutron, the OpenStack Networking service, will be installed in the next chapter. Disassociating a health monitor from a pool To disassociate a health monitor from a pool, use the Neutron lb-healthmonitordisassociate command as follows: Syntax: lb-healthmonitor-disassociate HEALTH_MONITOR_ID POOL The keyword POOL represents the ID of the pool to be disassociated from the monitor. Rather than run commands directly as root, OpenStack calls sudo neutronrootwrap /etc/neutron/rootwrap.conf when Neutron-related commands are executed. In this chapter, I will guide you through the following: Creating an external provider network Creating a router in the CLI and Horizon dashboard Attaching a router to both external and tenant networks Booting instances Demonstrating instance and namespace connectivity using LinuxBridge Demonstrating SNAT and NAT functionality provided by floating IPs The neutron-l3-agent service was installed on the controller node as part of the overall Neutron installation process documented in Chapter 3, Installing Neutron. Updating a health monitor To update the attributes of a health monitor, use the Neutron lb-healthmonitorupdate command as follows: Syntax: lb-healthmonitor-update HEALTH_MONITOR_ID Updateable attributes include delay, expected codes, HTTP method, max retries, timeout, and URL path. For more information on this bug, please refer to the following URL: Deleting a firewall in the CLI To delete a firewall within the CLI, use the Neutron firewall-delete command as follows: Syntax: firewall-delete FIREWALL_ID Listing firewalls in the CLI To list all firewalls within a tenant in the CLI, use the Neutron firewall-list command as follows: Syntax: firewall-list The returned output includes a list of firewalls containing the ID, name, and associated firewall policy for each firewall within the tenant. Before instances can be utilized in a useful manner, an underlying switching infrastructure must be configured. Showing the details of a firewall policy in the CLI To show the details of a firewall policy in the CLI, use the Neutron firewallpolicy-show command as follows: Syntax: firewall-policy-show FIREWALL_POLICY_ID The returned output includes the ID, name, description, tenant ID, audited status and associated firewall rules of the specified tenant. The health monitor would mark the pool member as DOWN and temporarily remove it from the pool. In this installation, the Keystone services are running on the controller node. Showing the details of a security group rule in the CLI To display the details of a particular security group rule, use the Neutron securitygroup-rule-show command as follows: Syntax: security-group-rule-show SECURITY_GROUP_RULE_ID The output returned includes the ID, direction, ethertype, port range, protocol, remote group IP, remote IP prefix, tenant ID, and associated security group ID of the specified security group rule. In this chapter, I will discuss fundamental load balancer concepts, including: Virtual IPs, pools, and pool members Load balancing algorithms Monitors Persistence Integrating load balancers in the network LBaaS uses drivers to interact with hardware and software load balancers. PacktLib is Packt's online digital book library. The following commands enable you to manage network profiles, policy profiles, profile binding, and credentials: cisco-credential-create cisco-credential-delete cisco-credential-list cisco-credential-show cisco-network-profile-create cisco-network-profile-delete cisco-network-profile-list cisco-network-profile-show cisco-network-profile-update [260]280 Appendix A cisco-policy-profile-list cisco-policy-profile-show cisco-policy-profile-update The cisco-network-profile commands enable you to create, modify, list, delete, and show details of Cisco Nexus 1000V network profiles. The --source-port flag is optional; it allows you to specify a source port or range of ports the rule should apply to. If you are creating a GRE or VXLAN network, the segmentation_id value should be an arbitrary, but unique, integer not used by any other network of the same type. This is not a desired behavior; rather, it is a limitation of FWaaS. In Icehouse, L3 agents are capable of supporting multiple external networks. Neutron will not allow you to attach a subnet to a router as an internal interface without the gateway_ip attribute set. IptablesFirewallDriver. The ML2 framework greatly simplifies adding support for new L2 networking technologies, as less effort is required to add functionality compared to creating a new monolithic core plugin. In the next chapter, you will be guided through the process of creating different types of networks to provide connectivity to instances. In Havana, ML2 can be configured to support the following network types: Flat VLAN Local GRE VXLAN The ML2 plugin has its own configuration file that must be used in conjunction with the Open vswitch or LinuxBridge configuration file. When using the Open vswitch plugin, each controller, network, or compute node in the environment has its own integration bridge and provider bridge. The syntax used to create a VLAN is provided in the net-create command: [113]133 Creating Networks with Neutron Attributes in the [] brackets are considered optional and are not required to create the network. OpenStack Networking also provides a framework for third-party vendors to build on and enhance the capabilities of the cloud. The firewall status will remain in PENDING_CREATE until the rules have been applied to the Neutron routers within the tenant, at which time the status will turn to ACTIVE: A change in status may require the page to be refreshed. bind_host = To enhance security, it is recommended that the API be exposed only on the management or API network. In some configurations, instances can reside in the same network as host machines. Feel free to add your own domain that references the management/api address of the controller. The default value for enable_isolated_metadata is False. The keyword POOL represents the pool to be balanced by this virtual IP. The load balancer has a single interface for ingress and egress traffic to and from clients and pool members. [150]170 Chapter 6 Users will only see routers that exist in their tenant or project. The list=true attribute is required to help Python interpret the data being passed as multiple entries. It is important to attach the router to each subnet so that it properly serves as the gateway for those subnets. He has extensive experience of various flavors of Linux and Unix. [81]101 Building a Virtual Switching Infrastructure Local When creating a local network in Neutron, it is not possible to specify a VLAN ID or physical interface. Installing and configuring the MySQL database server On the controller node, use yum to install the MySQL database server: # yum -y install mysql mysql-server MySQL-python Once installed, set the IP address that MySQL will bind to by editing the /etc/ my.cnf configuration file and adding the bind-address definition. Traditional network administration relies heavily on the administrator to manually configure and maintain physical network hardware and connectivity. Lastly, the provider bridge contains the physical network interface, which allows traffic to enter and exit the host onto the physical network infrastructure. The output of the preceding command is seen in the following screenshot: Listing subnets in the CLI To list existing subnets in Neutron, use the subnet-list command, as shown in the following screenshot: The list output provides the subnet ID, subnet name, CIDR notation, and the DHCP allocation range of all subnets when executed as an administrator. OVSNeutronPluginV2 Using crudini, set the core_plugin option in /etc/neutron/neutron.conf to use the LinuxBridge plugin on all nodes: # crudini --set /etc/neutron/neutron.conf DEFAULT core_plugin neutron. Appendix B, ML2 Configuration, will briefly cover the configuration of the ML2 plugin as a replacement for the deprecated LinuxBridge and Open Switch plugins. The --lb-method attribute is used to specify the load balancing algorithm, which is used to distribute traffic amongst the pool members. [54]74 Chapter 3 Creating the Neutron database Due to an unresolved bug, it is not possible to use the openstack-db command to properly create the Neutron database and user. Updating a pool To update the attributes of a pool, use the Neutron lb-pool-update command as follows: Syntax: lb-pool-update POOL [--description DESCRIPTION] [--lb-method {ROUND_ROBIN,LEAST_CONNECTIONS,SOURCE_IP}] The --lb-method attribute is used to specify the load balancing algorithm used to distribute traffic among the pool members. Inserting rules into firewall policies in the CLI Using the Neutron firewall-policy-insert-rule command, it is possible to insert firewall rules into an existing policy before or after the existing rules. It is not necessary to create this bridge, as Neutron does it automatically. Answer [Y]es to the remaining questions to exit the configuration process. You will also be provided with an overview of the differences between the two plugins in terms of how they function and provide layer 2 connectivity to instances. With Neutron, the process of providing metadata to instances varies and is based on the use of the Neutron routers. Instances connected to the same bridge can communicate with one another but with nothing else outside of the bridge. Click on the blue Allocate IP button to allow Neutron to procure the next available IP address: [179]199 Creating Routers with Neutron 5. While the clouds themselves may vary in complexity, one thing is common: they are made possible by the scalability and flexibility of OpenStack Compute and Networking services. Every network, router, and load balancer that is created by a tenant is represented by a network namespace. plugins.ml2.plugin.ml2plugin In addition to configuration file changes, a symbolic link named plugin.ini must be created in the /etc/neutron/ directory that points to the appropriate plugin configuration file before neutron-server will start. On the compute node, the preceding diagram can be realized as follows: [80]100 When multiple flat networks are created, a separate physical interface must be associated with each flat network. By default, the LinuxBridge plugin is configured to use NoopFirewallDriver. This persistence method can cause a load imbalance between pool members if users connect from behind a proxy server that misidentifies multiple clients as a single address. It is possible to attach an instance to the same network multiple times using multiple nova interface-attach commands. You can upgrade to the ebook version at com and as a print book customer, you are entitled to a discount on the ebook copy. For example, the subnet mask in the CIDR notation is /24. Use --pool-id to return pool members in the specified pool only. Using crudini, update the values in nova.conf, and use the appropriate LinuxBridge drivers on all nodes, as follows: # crudini --set /etc/nova/nova.conf DEFAULT linuxnet_interface_ driver linuxnet_interface_driver=nova.network.linux_net. When True, every network scheduled to a DHCP agent will have a namespace by the name of qdhcp-, where is a unique UUID associated with every network. The following diagram demonstrates the lack of physical or virtual VLAN interfaces in the bridge: Object Width eth0 eth0 Single IP address for MGMT/ API eth0 MGMT & API Net VM 0 eth0 tap0xxxx tap1xxxx brqyyyy (Linux Bridge) VM 1 K V M eth0 VM 2 tap2xxxx brqwwww (Linux Bridge) Figure 4.7 In the preceding diagram, two local networks exist that utilize their respective bridges, brqyyyy and brqwwww. If you encounter any issues, be sure to check the LBaaS agent log found at /var/log/neutron/lbaas-agent.log before proceeding. Listing pool members To obtain a list of pool members, use the Neutron lb-member-list command as follows: Syntax: lb-member-list [--pool-id=] The returned list of pool members includes member details, such as the ID, address, protocol port, admin state, and status. Internal network connections when using Open vswitch For an Ethernet frame to travel from the virtual machine instance out through the physical server interface, it will pass through nine devices inside the host: Tap interface: tapxxxx Linux bridge: qbryyyy Veth pair: qvbyyyy, qvoyyyy OVS integration bridge: br-int OVS patch ports: int-br-ethx, phy-br-ethx OVS provider bridge: br-ethx Physical interface: ethx [84]104 Chapter 4 The Open vswitch bridge br-int is known as the integration bridge. [64]84 Chapter 3 The enable_isolated_metadata configuration option is useful in cases where a physical network device (such as a firewall or router) serves as the default gateway for instances, but Neutron is still required to provide metadata services to instances. This chain is not used for security group rules, but instead for floating IP functionality within a router namespace. Many organizations choose to deploy their cloud this way, especially when port density is at a premium or the environment is simply used for testing. Since floating IPs are procured from provider networks, only provider networks whose router:external attribute is set to True will appear in the list. When segmentation_id is not specified, one is automatically allocated from the tenant range specified in the plugin configuration file. The UDP rule allows inbound DHCP response traffic from the DHCP server at /32. The Open vswitch bridge br-ethx is known as the provider bridge. The tunnel_bridge configuration option specifies the name of the tunnel bridge. More information on creating and managing load balancers in Neutron can be found in Chapter 7, Load Balancing Traffic in Neutron. For more information on how ML2 works, refer to the presentation titled OpenStack Neutron Modular Layer 2 Plugin Deep Dive available at openstack.org/. The lack of a high-availability solution for Neutron routers is an issue that persists in Icehouse but is actively being worked on by the community. Installing and configuring controller node components Install the openstack-nova package, which installs various Nova (Compute) services that are used on the controller node, as follows: # yum -y install openstack-nova python-novaclient Run the openstack-db command to initialize and create the Nova (Compute) service database, related tables, and MySQL user as follows: # openstack-db --init --service nova --password nova Use crudini to configure Nova Compute to use MySQL as its database as follows: # crudini --set /etc/nova/nova.conf database connection mysql:// [40]60 Chapter 2 Set these configuration keys to configure Nova (Compute) to use the Qpid message broker: # crudini --set /etc/nova/nova.conf DEFAULT rpc_backend nova.openstack. For now, the interface should be cabled and the corresponding switch port configured as a trunk. The packet will match one of the two rules based on the direction the packet is headed through the interface. Physical interface mappings The physical_interface_mappings configuration option describes the mapping of an artificial interface name or label to a physical interface in the server. Virtual private networks A virtual private network (VPN), extends a private network across a public network such as the Internet. Chapter 3, Installing Neutron, will explain how to install the Neutron networking components of OpenStack and will help us to understand the internal architecture of Neutron, including the use of agents and plugins to orchestrate network connectivity. There is a single integration bridge per node. The following table provides the VLANs and associated networks used for those services: VLAN name VLAN ID Network MGMT_NET /24 OVERLAY_NET /24 GATEWAY_NET /24 TENANT_NET /24 TENANT_NET /24 TENANT_NET /24 TENANT_NET /24 The following table provides IP addresses and VLAN IDs recommended for each host interface, should you choose to follow along with the examples: Hostname Interface IP address Switchport controller. Attaching internal interfaces to routers To create an interface in the router and attach it to a subnet, use the routerinterface-add command as follows: Syntax: router-interface-add In this case, INTERFACE is the ID of the subnet to be attached to the router. External network An external network provides Neutron routers with network access. The routes option allows you to add static routes to the routing table of a Neutron router. This is shown in the screenshot following the attributes: Delay: 5 Max retries: 3 Timeout: 16 ((Delay * Max Retries) + 1) Type: TCP To associate the newly created health monitor with the pool, use the lb-healthmonitor-associate command as follows: Syntax: lb-healthmonitor-associate HEALTH_MONITOR_ID POOL [206]226 Chapter 7 Consider the following screenshot: Creating a virtual IP The last step in creating a functional load balancer is to create the virtual IP, or VIP, which acts as a listener and balances traffic across pool members. The L3 agent is responsible for connecting interfaces within the router namespace to the proper bridge. In response to the request, the DHCP server will issue a DHCPACK packet or acknowledgement packet to the instance. The Neutron API, however, is installed on the controller node. FWaaS is currently lacking functionality that security groups provide, including the inability to specify the direction of traffic that should be filtered. First introduced in Grizzly, the Havana release of LBaaS offers numerous bug fixes and features over its predecessor but is not as polished as other Neutron services. Inbound traffic is not permitted by default. The --net-id option allows users to attach a new interface to an instance from the specified network. Creating an internal network Within the admin tenant, create an internal network for instances. The following are the service plugins for each of the services: Routing: neutron.services.l3_router.l3_router_plugin. Neutron assumes that when a gateway IP is set on a subnet, one of the following two things is true: The gateway IP belongs to a Neutron router The specified gateway will handle all the routing requests from instances, including those to the metadata server When a subnet is connected to a Neutron router that is serving as the gateway for that subnet, the router is responsible for routing all traffic from the subnet, including traffic to the metadata service. In addition to VLAN tagging, users can build overlay networks in software using L2-in-L3 tunneling protocols, such as GRE or VXLAN. He is an experienced IT professional with extensive experience in application networking technologies and OpenStack networking. [70]90 Building a Virtual Switching Infrastructure One of the core functions of OpenStack Networking is to provide connectivity to and from instances by dynamically configuring the virtual and/or physical network infrastructure in the cloud. Neutron stores the L3 agent configuration in the /etc/neutron/l3_agent.ini file. To determine the version of Open vswitch you have installed, run ovs-vsctl -V as follows: To enable GRE or VXLAN tunneling, set the configuration option manually, or use crudini to set the option to true on all hosts: # crudini --set /etc/neutron/plugins/openvswitch/ovs_neutron_plugin.ini OVS enable_tunneling true [102]122 Chapter 4 Tunnel type The tunnel_type configuration option specifies the type of tunnel supported by the plugin. Most services run on the controller node, with the exception of the openstack-nova-compute service, which runs on the compute nodes and is responsible for launching the virtual machine instances. Listing pools To obtain a list of configured load balancer pools, use the Neutron lb-pool-list as follows: Syntax: lb-pool-list [193]213 Load Balancing Traffic in Neutron The returned list includes details of pools in the running tenant, such as ID, name, load balancing method, protocol, admin state, and status. Doing so will return output similar to that of nova service-list in the CLI, as follows: [49]69 Installing OpenStack Summary At this point in the installation, the OpenStack Identity, Image, Dashboard, and Compute services have been deployed across the nodes of the cloud. Chapter 7 The --admin-state-down attribute, when set, does not have any effect on the state of the load balancer. By doing so, you can save other readers from frustration and help us improve subsequent versions of this book. These peer-to-peer tunnels create what is called a mesh network, where every host is connected to every other host. The --audited flag is optional; it is used to reflect whether or not a policy has been audited by an external resource. Have a look at the following screenshot: Verifying instance connectivity Upon spinning up the first instance, a DHCP namespace was created by Neutron that services DHCP requests from the instance in the TENANT_NET1 network, as shown in the following screenshot: The name of the DHCP namespace corresponds to the UUID of the TENANT_NET1 network. For more information on Neutron routers, refer to Chapter 6, Creating Routers with Neutron. Neutron will be installed in Chapter 3, Installing Neutron. Click on the blue Create button to complete the creation of the network and subnet: [132]152 Chapter 5 The ability to add additional subnets or delete the network entirely is provided within the More menu, pictured in the following screenshot: [133]153 Creating Networks with Neutron Neutron ports As mentioned earlier in this chapter, a port in Neutron is a logical connection of a cloud resource to a subnet. Network bridging is described as the action of connecting two or more layer 2 networks to create a single aggregate network. Then, a firewall is created and associated with a firewall policy. On the controller node, use the openssl utility to create a random shared secret. SDN, on the other hand, allows network administrators to manage network services in an abstract and automated manner. The --description flag is optional; it allows you to provide a description of the firewall policy. OUTPUT: Packets sent from the host machine itself will traverse this chain. Showing the details of a firewall in the CLI To show the details of a firewall within the CLI, use the Neutron firewall-show command as follows: Syntax: firewall-show FIREWALL_ID The output returned includes the ID, admin state, name, description, status, tenant ID, and associated firewall policy ID of the specified firewall. The --http-method attribute is optional; it is used in conjunction with --expected-codes and --url-path. Updating router attributes in the CLI To update the attributes of a router, use the Neutron router-update command as follows: Syntax: router-update [--admin-state-up] [--routes destination=,nexthop=] The admin-state-up attribute is a Boolean, which when set to false, does not allow Neutron to update interfaces within the router. With network namespaces, there is no concern of overlapping subnets between networks created by tenants. It is not uncommon for instances to attempt to renew their lease well before exceeding the lease duration. A physical interface includes Ethernet interfaces, such as eth0, and bonded interfaces consisting of one or more Ethernet interfaces or virtual VLAN interfaces of either type. As the traffic exits port 2 on the provider bridge and enters port 2 on the integration bridge, it is evaluated, in order, by the flow rules as follows: The first rule performs the action of modifying the VLAN ID from the original VLAN to a VLAN that is local to the integration bridge on the compute node when the original VLAN ID is 30: cookie=0x0, duration= s, table=0, n_packets=15, n_bytes=1904, idle_age=709, priority=3,in_port=2,dl_vlan=30 actions=mod_vlan_ vid:1,normal When traffic tagged as VLAN 30 is sent to an instance and forwarded through the provider bridge to the integration bridge, the VLAN tag is stripped and replaced with local VLAN 1. In the Havana release of OpenStack, FWaaS is an experimental extension with no guaranteed backwards compatibility in future releases. Click on the blue Save button to apply the changes. Open vswitch, also known as OVS, is an open source virtual switch that supports standard management interfaces and protocols, including NetFlow, SPAN, RSPAN, LACP, and 802.1q, though many of these features are not exposed to the user through the OpenStack API. In this appendix, you will find Neutron commands that didn't quite have a home in other chapters or are used in network solutions outside the scope of this book. The following diagrams reflect a few common service deployment models. The cisco-credential commands enable you to create, update, delete, and show details of Cisco Nexus 1000V credentials. Using crudini, associate a range of VLANs with physnet1 for use with tenant networks on all hosts: # crudini --set /etc/neutron/plugins/linuxbridge/linuxbridge_conf.ini vlans network_vlan_ranges physnet1:30:33 Firewall driver The firewall_driver configuration option instructs Neutron to use a particular firewall driver for security group functionality. This can be changed to whatever fits your organization. Instead, an IP address is associated with a Neutron port, and that port is logically mapped to the virtual tap interface that connects instances to the network. For more information, please refer to the following bug report: Tunnel ID ranges When GRE- or VXLAN-based networks are created, each network is assigned a unique ID, or segmentation ID, that is used to encapsulate traffic. [244]264 The --admin-state-down flag is optional; it allows you to create the firewall in a DOWN state. Operating system requirements OpenStack currently has packages for the following distributions: CentOS, Debian, Fedora, RHEL, opensuse, SLES, and Ubuntu. These ports correspond to two instances in two different Neutron networks as evidenced by their difference in VLAN IDs. The VLAN IDs are arbitrarily assigned by Open vswitch and may change upon restart of the openvswitch service or after a reboot. The --firewall-rules flag is optional; it is used to add firewall rules to the policy during creation. The qg interface is the gateway, or external, interface of the router. To confirm the gateway network selection, click on the blue Set Gateway button. From the Add Member window, you can add multiple members to a pool simultaneously and set a common weight and protocol port for the chosen pool members. Without a monitor, the load balancer will continue to send traffic to members that may not be available. Installing and configuring Network Time Protocol A time synchronization program, such as NTP, is a requirement, as OpenStack services depend on consistent and synchronized time between hosts. This time, however, traffic from the integration bridge connected to port 2 is processed by the first two rules: cookie=0x0, duration= s, table=0, n_packets=101, n_bytes=5984, idle_age=723, priority=4,in_port=2,dl_vlan=1 actions=mod_vlan_ vid:30,normal cookie=0x0, duration= s, table=0, n_packets=31, n_bytes=2300, idle_age=830, priority=2,in_port=2 actions=drop Chapter 4 The first flow rule on the provider bridge checks the VLAN ID in the Ethernet header, and if it is 1, modifies it to 30 before forwarding the traffic to the physical interface. There is no mechanism to permit traffic between instances on different bridges or hosts when using local networks. Using subnet-update to update the existing dns_nameservers or host_routes value will result in the overwriting of existing values. Network interface devices, including interfaces used by DHCP or router namespaces and instance tap interfaces, connect to OVS bridge ports. The following Neutron commands are specific to the NSX extension: net-gateway-connect net-gateway-create net-gateway-delete net-gateway-disconnect net-gateway-list net-gateway-show net-gateway-update queue-create queue-delete queue-list queue-show For more information on configuring Neutron with the NSX/NVP plugin, please refer to the OpenStack Cloud administrator guide at admin-guide-cloud/content/nsx_plugin.html. The shared switch is a Boolean value that, when set to true, allows the network to be utilized amongst all tenants. The keyword NEW_FIREWALL_RULE_ID is used to represent the ID of the firewall rule to be added to the policy. Using the Neutron lb-healthmonitor-create command, create a health monitor with the following attributes. When attached networks and subnets have their own respective gateway address set, an instance's routing table can be populated with multiple default routes. For this installation, change the value from False to True and uncomment the line in the configuration file: # sed -i "/# enable_isolated_metadata/c\enable_isolated_metadata = True" /etc/neutron/dhcp_agent.ini The enable_metadata_network configuration option is useful in cases where the L3 agent might be used, but the metadata agent is not on the same host as the router. Additional networks on the same bridge must have a unique segmentation ID. Visit to generate strong passwords for your environment. For this installation, the chosen root password is openstack. With the introduction of Firewall-as-a-Service, also known as FWaaS, security is handled at the router rather than at the compute node. [253]273 Protecting Instances on the Network Summary It is important to know the differences between the two methods of securing network traffic to instances. Traffic between instances in the same network remains local to the virtual switch, and by definition, local to the compute node on which they reside. To find more information on creating Neutron API extensions, please visit the Neutron development wiki at NeutronDevelopment. The Neutron command-line client can be used to manage networks, subnets, and ports and is highly recommended over the limited Horizon dashboard for most administrative tasks. In this environment, there are two instances eligible for use in the pool, as shown in the following screenshot: [204]224 Chapter 7 Using the Neutron lb-member-create command, create two pool members with the following attributes: Member 1: Member 2: Name: Web1 Address: Protocol Port: 80 Pool: WEB_POOL Name: Web2 Address: Protocol Port: 80 Pool: WEB_POOL The following screenshot demonstrates the process of creating the first pool member: Repeat the process shown in the preceding screenshot to create the second pool member. Finally, the tenant-id option allows the administrator to create networks on behalf of the tenants. A physical interface resides in the network bridge and handles external, guest, management, and API service traffic: eth0 VM n (br-int) Integration Bridge (br-tun) Tunnel Bridge Used for processing overlay traffic (br-eth0) Provider Bridge Single IP address for all services and connecting hosts in point-to-point overlay mesh eth0 (1GbE) Overlay Networks MGMT &API External Networks =Virtual Ethernet Cable or Cross-Connect =Open vswitch Bridge with Flows Figure 1.2 In this diagram, all OpenStack service and management traffic traverses the same physical interface as guest traffic. Once your errata are verified, your submission will be accepted and the errata will be uploaded on our website, or added to any list of existing errata, under the Errata section of that title. By typing a question mark or help within the Neutron shell, a list of commands can be found. In transparent mode, the load balancer acts as a network bridge between two VLANS configured with the same subnet(s). For your convenience, the following command will make the appropriate change: # sed -i "/SELINUX=enforcing/c\SELINUX=disabled" /etc/selinux/config Removing iptables rules CentOS ships with rather restrictive iptables rules by default. New connections destined to any outside network on port 80 are allowed. Restarting neutron-server is necessary for changes to take effect. DHCP and metadata services that run on the controller or dedicated network node deliver IP addresses and instance-specific data at boot time. Clicking on Create Network will open a window where you can specify the network and subnet details: [131]151 Creating Networks with Neutron 3. Using crudini, set the external_network_bridge configuration option to an empty value as follows: # crudini --set /etc/neutron/l3_agent.ini DEFAULT external_network_bridge Enabling the metadata proxy When Neutron routers are used as the gateway for instances, requests for metadata are proxied by the router and forwarded to the Nova metadata service. To complete the rule creation, click on the blue Add button. The Nova metadata API service responds to the request and forwards it to the Neutron metadata agent 6. The gateway interface of a Neutron router is analogous to the WAN interface of a hardware router. Have a look at the following screenshot: In the following screenshot, a look at the iptables rules within the router namespace reveals the NAT rules responsible for this behavior: In this configuration, instances can communicate with outside resources as long as the instances initiate the traffic. Modern cloud computing platforms, such as OpenStack, rely on a method of networking known as software-defined networking, or SDN. Information on session persistence types can be found earlier in this chapter. The following table provides the IP address and virtual interface name to be created on each host: Hostname Interface IP Address Controller eth Compute01 eth Using crudini, set the local_ip configuration option accordingly. [190]210 Chapter 7 Using crudini, set the user group to nobody: crudini --set /etc/neutron/lbaas_agent.ini DEFAULT user_group nobody All other configuration options can be left to their defaults or can be modified directly if the default values are not sufficient for your installation. When using the LinuxBridge plugin, the interface is placed into a bridge that corresponds to the external network. Each subnet must have a CIDR and must be associated with a network. Using the Neutron lb-pool-create command, create a pool with the following attributes: Name: WEB_POOL Load balancing method: Round robin Protocol: HTTP Subnet ID: [203]223 Load Balancing Traffic in Neutron Have a look at the following screenshot: The state of the pool will remain in PENDING_CREATE until a virtual IP has been associated with it. This feature is enabled by default and can be disabled by setting the enable_metadata_proxy value to false in the l3_agent.ini configuration file. The integration bridge is the central virtual switch that many network resources are connected to, including instances, DHCP servers, routers, and more. An administrator is able to create floating IPs on behalf of tenants. Listing Neutron API extensions To list the extensions available in Neutron, use the Neutron ext-list command as follows: Syntax: ext-list275 Additional Neutron Commands The returned output includes the alias and name of the available extensions: Showing the details of an API extension To show the details of an API extension, use the Neutron ext-show command as follows: Syntax: ext-show EXTENSION_ALIAS The EXTENSION_ALIAS keyword represents alias of the extension provided in the ext-list output. OpenStack Networking uses a service called neutron-server to expose an application programmable interface, or API, to users and to pass requests to the configured network plugins for additional processing. This means that DHCP and metadata services will be unavailable to any instances not on the same host as those services. When all available IDs in the range available to tenants are exhausted, users will no longer be able to create networks of this type. To attach a gateway interface to a router, use the router-gateway-set command as follows: Syntax: router-gateway-set [--disablesnat] The default behavior of a Neutron router is to source NAT all outbound traffic from instances that do not have a corresponding floating IP. Instances in the same local network that reside on different hosts will be unable to communicate with one another. He works as Principal Cloud Architect for Rackspace UK, specializing in OpenStack and covering the international market for the DevOps & Automation Advisory Services team. Using the Neutron l3-agent-list-hosting-router command, you can determine which L3 agent the router was scheduled to, as shown in the following screenshot: On the node hosting the L3 agent, a network namespace is created that corresponds to the router. Those variables are only needed to bootstrap the administrative user and to register the Keystone service. You can only create one flat network per provider bridge, as there is no mechanism to segment traffic. Typically, a username and password are used to authenticate against Keystone. In this installation, Glance runs on the controller node: # crudini --set /etc/nova/nova.conf DEFAULT glance_host controller Edit the /etc/nova/api-paste.ini configuration to add credentials to the [filter:authtoken] section: # crudini --set /etc/nova/api-paste.ini filter:authtoken auth_host controller # crudini --set /etc/nova/api-paste.ini filter:authtoken auth_port # crudini --set /etc/nova/api-paste.ini filter:authtoken auth_protocol http # crudini --set /etc/nova/api-paste.ini filter:authtoken admin_tenant_ name service # crudini --set /etc/nova/api-paste.ini filter:authtoken admin_user nova # crudini --set /etc/nova/api-paste.ini filter:authtoken admin_password nova Ensure that the api_paste_config=/etc/nova/api-paste.ini option is set in the /etc/nova/nova.conf file: # crudini --set /etc/nova/nova.conf DEFAULT api_paste_config /etc/nova/ api-paste.ini Start the Nova (Compute) service and configure it to start when the system boots: # service libvirtd start # service messagebus start # service openstack-nova-compute start # chkconfig libvirtd on # chkconfig messagebus on # chkconfig openstack-nova-compute on Verify communication between services To check the status of Nova services throughout the environment, use the Nova service-list command on the controller node as follows: # nova service-list [45]65 Installing OpenStack The command should return statuses on all Nova services that have checked in: In the above output, the state of the services on both the controller and compute nodes are reflected under the Status column. In Havana, the default driver uses haproxy. When using the LinuxBridge plugin, the interface is placed in a bridge that corresponds to the internal network as shown in the following screenshot: [163]183 Creating Routers with Neutron The namespace is able to communicate with other devices in the same subnet through the bridge. The load balancer configuration files can be found in the /var/lib/ neutron/lbaas/ directory of the host running the LBaaS agent. While Open vswitch provides features not available with the monolithic LinuxBridge plugin, including the use of overlay networks, its configuration, administration, and troubleshooting methods are more complex. [220]240 Protecting Instances on the Network Neutron includes two methods of providing network-level security to instances. A new subnet will need to be created and added to the network. In the down state, the pool member is not eligible to receive traffic. [136]156 Chapter 5 Adding secondary addresses to interfaces The default behavior of Neutron is to associate a single IP address with a port. An example of this behavior can be seen in later chapters. Clicking on Associate Floating IP will open a window that allows you to manage floating IP allocations: [178]198 Chapter 6 3. A physical interface includes Ethernet interfaces, such as eth0, eth1, and so on; bonded interfaces consisting of one or more Ethernet interfaces; or VLAN interfaces of either type. Flow rules for flat networks Flat networks in Neutron are untagged networks, which means that there is no 802.1q VLAN tag associated with the network when it is created. A subnet in Neutron is a layer 3 resource, and can be an IPv4 or IPv6 network defined by the classless inter-domain routing (CIDR) notation. Each major networking component has a dedicated chapter that will build on your experience gained from prior chapters. Setting environment variables To avoid having to provide credentials every time you run an OpenStack command, create a file containing environment variables that can be loaded at any time: # mkdir ~/credentials # cat >> ~/credentials/admin prompt, execute the following commands to create a database named neutron_ml2 and to grant permissions to the existing neutron user: CREATE DATABASE neutron_ml2; GRANT ALL PRIVILEGES ON neutron_ml2.* TO IDENTIFIED BY 'neutron'; GRANT ALL PRIVILEGES ON neutron_ml2.* TO QUIT; Use crudini to overwrite the existing database connection string in the Neutron configuration file with the new string on all hosts as follows: # crudini --set /etc/neutron/neutron.conf database connection mysql:// Configuring Neutron to use ML2 Before the ML2 plugin can be used, changes, which include specifying the core plugin and database options, must be made to the Neutron configuration on all hosts. [83]103 Building a Virtual Switching Infrastructure An OVS bridge behaves like a physical switch, only virtualized. When traffic is returned to the calling chain, the next UDP rule prohibits the instance from acting as a rogue DHCP server. The default security group drops all ingress traffic and allows all egress traffic from instances. The second method is a feature known as Firewall-as-a-Service (FWaaS) that provides filtering at the perimeter of the network on a Neutron router. When set to True, Neutron can provide instances with a static route to the metadata service via DHCP in certain cases. For more information on iptables, please visit the following resources: Enterprise_Linux/6/html/Security_Guide/sect-Security_Guide- IPTables.html [224]244 Working with security groups Security groups can be managed in either the Neutron CLI or the Horizon dashboard. From the Manage Floating IP Associations window, select a floating IP from the IP Address menu, and choose the address utilized by the virtual IP from the Port To Be Associated menu. The following statement and configuration option can be found in the /etc/ neutron.conf file: # Change to "sudo" to skip the filtering and just run the command directly # root_helper = sudo neutron-rootwrap /etc/neutron/rootwrap.conf [59]79 Installing Neutron To disable command filtering, change the root_helper value to sudo only. In the next chapter, I will expand on the concept of Neutron routers and their involvement in the network, including the configuration and use of floating IPs to provide external connectivity to instances. In production clouds, however, separate control and data interfaces are recommended. The CIDR argument defines the CIDR notation of the subnet being created. A floating IP address in OpenStack is a static NAT that maps an external address to an internal address. IPv4 is the default when the version is not specified. In Neutron, a network may contain multiple subnets. [256]276 For your reference, the following Neutron commands are used to manage VPN connections in OpenStack: ipsec-site-connection-create ipsec-site-connection-delete ipsec-site-connection-list ipsec-site-connection-show ipsec-site-connection-update vpn-ikepolicy-create vpn-ikepolicy-delete vpn-ikepolicy-list vpn-ikepolicy-show vpn-ikepolicy-update vpn-ipsecpolicy-create vpn-ipsecpolicy-delete vpn-ipsecpolicy-list vpn-ipsecpolicy-show vpn-ipsecpolicy-update vpn-service-create vpn-service-delete vpn-service-list vpn-service-show vpn-service-update Appendix A The installation and configuration of VPNaaS is outside the scope of this book. In this installation, the username is admin, and the password is secrete. Firewalling In Havana, there are two methods of providing security to instances or networks: security groups and firewalls. Identifying ports on the virtual switch Using the ovs-ofctl show command, you can see a logical representation of the specified virtual switch. Have a look at the following commands: # keystone endpoint-create \ --service-id=`keystone service-get keystone awk '/ id / { print $4 }'` \ --publicurl= \ --internalurl= \ --adminurl= The resulting output is as follows: Property Value adminurl id 7c1112c14cd8494fbd8dadb f internalurl publicurl region regionone service_id 47b36f2684e94cfdbd78ba912e6091ec Verify the Keystone installation To verify that Keystone was installed and configured properly, use the unset command to unset the OS_SERVICE_TOKEN and OS_SERVICE_ENDPOINT environment variables. Preparing Neutron for FWaaS To properly implement FWaaS, some changes must be made to the Neutron configuration files on the controller node. Instances in the same local network will be placed in the same bridge and are free to communicate amongst one another. To detach an interface from an instance, use the nova interface-detach command as follows: nova interface-detach Interfaces detached from instances are removed completely from the Neutron port database. If a 500 status code were returned, it could indicate that the server is not properly processing connections. All other attributes are read-only. Chapter 8 Managing security groups in the CLI From within the Neutron command-line client, a number of commands can be used to manage security groups, such as: security-group-create security-group-delete security-group-list security-group-rule-create security-group-rule-delete security-group-rule-list security-group-rule-show security-group-show security-group-update Creating security groups in the CLI To create a security group within the CLI, use the Neutron security-group-create command as follows: Syntax: security-group-create [--tenant-id TENANT_ID] [--description DESCRIPTION] NAME By default, security groups in Neutron are prepopulated with two egress rules that allow all outbound traffic over IPv4 and IPv6. Updating a firewall policy in the CLI To update the attributes of a firewall policy, use the Neutron firewall-policyupdate command as follows: Syntax: firewall-policy-update FIREWALL_POLICY_ID [--name NAME] [--description DESCRIPTION] [--shared] [--firewall-rules list=true RULES] Chapter 8 Multiple rules should be separated by a space. Creating a firewall in the CLI To create a firewall within the CLI, use the Neutron firewall-create command as follows: Syntax: firewall-create [--tenant-id TENANT_ID][--name NAME] [--description DESCRIPTION][--admin-state-down] POLICY The --tenant-id flag is optional; it allows you to associate the firewall with the specified tenant. For more information on source address filtering and security groups, see Chapter 8, Protecting Instances on the Network. On the other hand, every incremental decrease of the exponent halves the number of available addresses. The address or network should be defined in CIDR format. The other interface in the bridge, eth1.50, tags traffic as VLAN 50 as it exits the bridge and out of the physical interface eth1. Attaching and detaching network interfaces When an instance is first booted, Neutron creates a port that maps the MAC address of the VIF and the allocated IP address to the instance. learningneutron.com compute01. The source addresses of packets do not need to be manipulated in most cases, as the load balancer serves as the gateway for pool members. The provider bridge provides connectivity to the physical network interface ethx, where X represents the enumerated physical NIC, and is connected to the integration bridge by a virtual patch cable provided by patch ports int-br-ethx and phy-br-ethx. Click on the Add Member button within the Members section. Chapter 6, Creating Routers with Neutron, will create Neutron routers and attach them to networks, follow traffic from an instance through a router, and explore the process of applying floating IPs to instances. The OUTPUT chain is used by the raw, mangle, NAT, and filter tables. The TCP monitor tests connectivity to pool members at layer 4, while the HTTP and HTTPS monitors test the health of pool members based on layer 7 HTTP status codes. [149]169 Creating Routers with Neutron Listing interfaces attached to routers To list the interfaces attached to routers, use the router-port-list command as follows: Syntax: router-port-list The returned output includes the Neutron port ID, MAC address, IP address, and associated subnet of attached interfaces. Configure the following settings to allow Neutron to access the Identity Service API: # crudini --set /etc/neutron/api-paste.ini filter:authtoken auth_host controller # crudini --set /etc/neutron/api-paste.ini filter:authtoken auth_uri # crudini --set /etc/neutron/api-paste.ini filter:authtoken admin_tenant_ name service # crudini --set /etc/neutron/api-paste.ini filter:authtoken admin_user neutron # crudini --set /etc/neutron/api-paste.ini filter:authtoken admin_ password neutron [58]78 Chapter 3 Configuring Neutron to use a messaging service Neutron communicates with various OpenStack services on the AMQP messaging bus. When the pool member sends its response, haproxy injects a cookie named SRV into the response before sending it to the client. [3]23 Preface Reader feedback Feedback from our readers is always welcome. Prior to NAT, every host connected to the Internet had a unique IP address. Displaying router attributes in the CLI To display the attributes of a router, use the Neutron router-show command as follows: Syntax: router-show Among the output returned is the admin state, the external network, the SNAT state, and the tenant ID associated with the router. Using crudini, configure bind_host to use the management address on the controller node as follows: # crudini --set /etc/neutron/neutron.conf DEFAULT bind_host [61]81 Installing Neutron Other configuration options that might require tweaking include: core_plugin dhcp_lease_duration allow_overlapping_ips The core_plugin configuration option instructs Neutron to use the specified networking plugin. As a user, the command returns subnets within the tenant or subnets associated with shared networks. Installing the ML2 plugin RHEL-based distributions, such as CentOS, require the installation of the openstack-neutron-ml2 package on all nodes as follows: # yum install openstack-neutron-ml2283 ML2 Configuration Creating a database for ML2 The ML2 plugin attempts to use a common database and schema that can be shared amongst multiple layer 2 agents. [191]211 Load Balancing Traffic in Neutron Enabling LBaaS in Horizon Before load balancers can be managed in the dashboard, the enable_lb setting in /etc/openstack-dashboard/local_settings file must be set to True. HTTP: This instructs the monitor to initiate an HTTP request to a pool member based on the expected_codes, url_path, and http_method attributes described here. The following example is based on an instance whose interface corresponds to the b67c75e5-4f9f-4770-ae72-dd6a144ddd26 Neutron port: -A neutron-openvswi-sb67c75e5-4 -s /32 -m mac --mac-source fa:16:3e:02:55:34 -j RETURN -A neutron-openvswi-sb67c75e5-4 -j DROP Based on the preceding example, traffic from the instance must be sourced from and the source MAC address of fa:16:3e:02:55:34. However, the information contained in this book is sold without warranty, either express or implied. In OpenStack, antispoofing rules are implemented by Neutron on each compute node within iptables. Inside the bridge is a VLAN interface and two tap interfaces that correspond to the instances. The physical switch port connected to eth1 must support 802.1q VLAN tagging if VLAN networks are to be created by tenants. The latter may require that the user be added to the sudoers file on each host. In the preceding example, the InstanceName instance was connected to three different Neutron networks. In one-arm mode, the load balancer is not in the path of normal traffic to the pool members. This book assumes that OpenStack will be installed on physical hardware that meets the following minimum requirements: Server Recommended hardware Notes Controller node (runs network, volume, API, scheduler, and image services) Compute node (runs virtual instances) Processor: 64-bit x86 Memory: 8 GB RAM Disk space: 80 GB Network: Two 1 Gbps network interface cards (NICs) Processor: 64-bit x86 Memory: 16 GB RAM Disk space: 80 GB Network: Two 1 Gbps network interface cards (NICs) While a single NIC can be used, it is not recommended, and therefore not addressed in this build. Bug at documents this issue. The host-route attribute defines one or more static routes to be injected via DHCP. More than one provider bridge can be configured on a host, but it requires the use of a dedicated physical interface, or virtual VLAN interface in some cases, per provider bridge. He can be found on and on Freenode IRC as busterswt. A look at the iptables rules within the router namespace shows that rules have been added to perform the 1:1 NAT translation as follows: [170]190 Chapter 6 With the proper routes in place on the client machine, traffic can be initiated directly to the instance via the floating IP as follows: Reassigning floating IPs The idea behind a floating IP is that it is a NAT that can be quickly disassociated from an instance or other network resource and associated with another. Throughout this book, you will learn how to build a functional OpenStack cloud utilizing advanced networking features available in the Havana release. First published: October 2014 Production reference: Published by Packt Publishing Ltd. The iptables rules on the compute node will be updated accordingly. In a DOWN state, the firewall rules are not applied. To create a network and subnet as a user, perform the following steps: 1. Livery Place 35 Livery Street Birmingham B3 2PB, UK. Multiple routes listed as destination and nexthop pairs can be separated by a space. Deleting internal interfaces To delete an internal interface from a router, use the router-interface-delete command as follows: Syntax: router-interface-delete Here, INTERFACE is the ID of the subnet to be removed from the router. Managing pool members in the CLI The following commands are used to manage pool members in the CLI: lb-member-create lb-member-delete lb-member-list lb-member-show lb-member-update Chapter 7 Creating pool members To create a pool member, use the Neutron lb-member-create command as follows: Syntax: lb-member-create [--tenant-id TENANT_ID] [--admin-state-down] [--weight WEIGHT] --address --protocol-port POOL The --tenant-id flag is optional; it allows you to associate the pool member with the specified tenant. [14]34 Chapter 1 Bonding NIC bonding offers users the ability to multiply available bandwidth by aggregating links. learningneutron.com eth Access port (VLAN 10, Untagged) eth1 Cabled, but unconfigured at this time Trunk port (VLAN 20, 30-33, 50) eth Access port (VLAN 10, untagged) eth1 Cabled, but unconfigured at this time Trunk port (VLAN 20, 30-33, 50) The eth1 interface of each server will be configured in Chapter 4, Building a Virtual Switching Infrastructure. The next rule in the neutron-l3-agent-forward chain matches all traffic entering any qr-* interface attached to the router and sends it to the neutron-l3-agentov401ea5b3a chain: Like the previous chain, packets that are invalid are dropped, while established connections are accepted without further processing. Creating a pool The first step to building a functional load balancer is to create the pool. The following screenshot demonstrates the switch configuration in a graphical manner: VM 0 eth0 VM 2 eth0 KVM tapfxxxx qbrfxxxx (Linux Bridge) qvofxxxx tap0xxxx qbr1xxxx (Linux Bridge) qvb0xxxx qvof0xxxx Port 6 VLAN 1 qvo0xxxx br-int (Open vswitch Integration Bridge) int-br-eth1 Port 5 VLAN 1 Port 2 eth0 Single IP address for MGMT/API eth0 phy-br-eth1 br-ethx (Open vswitch Provider Bridge) ethx 1 Port 2 Port 1 MGMT & API Net External Networks Figure 4.10 [87]107 Building a Virtual Switching Infrastructure Identifying the local VLANs associated with ports Every port on the integration bridge connected to an instance or other network resource is placed in a VLAN that is local to that host. Use crudini to set a value for network_vlan_ranges on all hosts as follows: # crudini --set /etc/neutron/plugins/openvswitch/ovs_neutron_plugin.ini OVS network_vlan_ranges physnet1:30:33 If at any time this configuration option is updated, you must restart the neutronserver service for the changes to take effect. On compute01, install the MySQL client as follows: # yum -y install mysql MySQL-python Installing and configuring the messaging server Advanced Message Queue Protocol (AMQP) is the messaging technology chosen for use with an OpenStack-based cloud. Configuring the bridge interface In this installation, the physical network interface eth1 will be utilized for bridging purposes. Configure Neutron to use Qpid as the messaging broker on all the nodes with the following settings: # crudini --set /etc/neutron/neutron.conf DEFAULT rpc_backend neutron. Using DHCP to inject the route When enable_isolated_metadata is set to true and a gateway is not set in the subnet, the DHCP service is capable of injecting a route to the metadata service via the classless-static-route DHCP option, otherwise known as option 121. DHCP goes through the following stages: A DHCP client sends a DHCPDISCOVERY packet using a broadcast address. This API lets you define network connectivity in order to leverage network capabilities to cloud deployments.Through this practical book, you will build a strong foundational knowledge of Neutron, and will architect and build an OpenStack cloud using advanced networking features.We start with an introduction to OpenStack Neutron and its various components, including virtual switching, routing, FWaaS, VPNaaS, and LBaaS. Have a look at the following screenshot: Load balancer management in the dashboard In the Horizon dashboard, load balancers can be managed from the Project panel by clicking on Load Balancers in the menu on the left-hand side of the screen. Switching Virtual switches are defined as software applications that connect virtual machines to virtual networks at layer 2, or the data-link layer of the OSI model. Showing pool member details To show the details of a pool member, use the Neutron lb-member-show command as follows: Syntax: lb-member-show MEMBER The keyword MEMBER represents the ID of the member to be shown. This book assumes that the CentOS 6.5 operating system has been installed on all hosts prior to the installation of OpenStack: CentOS 6.5: At the time of writing, the following minimum kernel version is recommended: Kernel version: el6.x86_64 [22]42 Chapter 2 Prior kernel versions may experience a lack of support for network namespaces that are used throughout various Neutron services. The following naming convention for network namespaces should be observed: qdhcp- qrouter- qlbaas- The qdhcp namespace contains a DHCP service that provides IP addresses to instances using the DHCP protocol. [156]176 Chapter 6 Creating an external provider network In order to provide instances with external connectivity, a Neutron router must be connected to a provider network eligible for use as an external network. Each network namespace has its own routing table and iptables process that provide filtering and network address translation, also known as NAT. [184]204 Chapter 7 With the least-connections algorithm, the load balancer passes a new connection to a server that has the least number of current connections. The /23 subnet can be written as 2 9, resulting in 512 addresses, and /22 can be written as 2 10, resulting in 1024 addresses. Bonding is an inexpensive way to provide hardware-level network redundancy to the cloud. While virtualization software, such as VMware or VirtualBox, can be used to simulate the servers and the network infrastructure, this book assumes that OpenStack will be installed on physical hardware and that a physical network infrastructure is in place. As traffic from instances traverses the Linux bridge and out towards the physical interface, the eth1.100 interface tags that traffic as VLAN 100 and drops it on eth1. To edit a particular pool member, click on the Edit Member button under the Actions column next to the pool member: [215]235 Load Balancing Traffic in Neutron 5. Click on Add Monitor from the Monitors tab on the Load Balancers screen. He has reviewed the following books: Penetration Testing with BackBox, Packt Publishing Django Essentials, Packt Publishing Penetration Testing with the Bash shell, Packt Publishing Thanks to all my friends for supporting me. If you find any errata, please report them by visiting submit-errata, selecting your book, clicking on the errata submission form link, and entering the details of your errata. Over the last 18 years, he has worked as a software developer, systems administrator, and general technologist. Types of network traffic The reference architecture for OpenStack Networking defines at least four distinct types of network traffic: Management API External Guest These distinct types of network traffic do not require dedicated interfaces and are often collapsed onto single interfaces. A Linux bridge is a virtual interface that connects multiple network interfaces. [216]236 Chapter 7 3. NAT: This is a default table used for network address translation. Updating a virtual IP To update the attributes of a virtual IP, use the Neutron lb-vip-update command as follows: Syntax: lb-vip-update VIP [--connection-limit CONNECTION_LIMIT] [--pool-id POOL] [--session-persistence type={http_cookie,source_ip,app_cookie}] Session persistence is an attribute that is not directly exposed within the CLI but is available in the dashboard. SELinux can be set to one of the following three states: enforcing: SELinux security policy is enforced permissive: SELinux prints warnings instead of enforcing disabled: No SELinux policy is loaded To disable SELinux, edit the /etc/selinux/config file, and change the SELINUX value to disabled. Once a router has been configured, this network becomes the source of floating IP addresses for instances and load balancer VIPs. IP addresses in this network should be reachable by any client on the Internet. For this installation, set the value to 1:1000 using crudini on all hosts: # crudini --set /etc/neutron/plugins/openvswitch/ovs_neutron_plugin.ini OVS tunnel_id_ranges 1:1000 [103]123 Building a Virtual Switching Infrastructure Integration bridge The integration_bridge configuration option specifies the name of the integration bridge used on each node. The FDB table is the equivalent of a CAM or MAC address table. Linux supports 802.1q VLAN tagging through the use of virtual VLAN interfaces. To permit traffic from additional IP or MAC addresses, use the Neutron allowed-address-pairs extension, as discussed in Chapter 5, Creating Networks with Neutron. Firewall rules behind the scenes To demonstrate how firewall policies are applied to a Neutron router, check out the following firewall rule that allows HTTP traffic from any remote host to any instance on TCP port 80: [250]270 Chapter 8 Using the Neutron firewall-policy-create command, I have created a policy that contains the preceding rule: Using the Neutron firewall-create command, I have created a firewall using the policy MyFirewallPolicy: The firewall status will remain in PENDING_CREATE until the rules have been applied to the Neutron routers within the tenant, at which time the status will turn to ACTIVE: [251]271 Protecting Instances on the Network Stepping through the chains within the firewall As a result of creating the firewall, the rules within the firewall policy have been implemented on all routers within the tenant. Once the command is executed, Neutron creates a port in the database that is associated with the router interface. Subnet: A subnet is an IPv4 or IPv6 address block from which IP addresses can be assigned to virtual machine instances. Endpoint addresses for services, such as Keystone, Neutron, Glance, and Horizon, are procured from the API network. Neutron port definitions are stored in the Neutron database, which is then used by the respective plugin agent to build and connect the virtual switching infrastructure. Use the following command to implement this change on all nodes: # sed -i "/net.ipv4.ip_forward/cet.ipv4.ip_forward = 1" /etc/sysctl. However, for the sake of starting Neutron services and demonstrating the Neutron CLI in this chapter, a network configuration based on the LinuxBridge plugin will be used. Installing and configuring compute node components Once the Nova (Compute) services have been configured on the controller node, another host must be configured as a compute node. The following nova boot command demonstrates the procedure of connecting an instance to multiple networks at the first boot: nova boot --flavor --image --nic net-id= --nic net-id= --nic net-id= InstanceName Nova attaches instances to networks with a virtual interface or VIF. To activate the change, uncomment the line and restart the neutron-server service. ISBN Cover image by Suyog Gharat4 Credits Author James Denton Reviewers Kevin Jackson Jorge Armin Garcia Lopez Jacob Walcik Commissioning Editor Kartikey Pandey Acquisition Editor Richard Harvey Content Development Editor Susmita Panda Technical Editor Shiny Poojary Copy Editors Roshni Banerjee Sarang Chari Karuna Narayanan Project Coordinator Kartik Vedam Proofreaders Martin Diver Ameesha Green Samantha Lyon Indexers Hemangini Bari Monica Ajmera Mehta Tejal Soni Graphics Sheetal Aute Ronak Dhruv Valentina D'silva Disha Haria Abhinash Sahu Production Coordinators Aparna Bhagat Shantanu N. Otherwise, the load balancing algorithm is applied. When configured to utilize the Open vswitch networking plugin, Neutron relies on the bridge and openvswitch kernel modules, along with user-space utilities, such as ovs-vsctl and ovs-ofctl, to properly manage the Open vswitch database and connect instances and other network resources to virtual switches. The opposite can said for security groups, too, as they lack the ability to create specific deny rules as all traffic is denied by default. At you can also read a collection of free technical articles, sign up for a range of free newsletters and receive exclusive discounts and offers on Packt books and ebooks. More information on CIDR and VLSM can be found on Wikipedia at A few examples of subnets in the CIDR notation are described as follows: /24 represents the IP address , its associated routing prefix , and the subnet mask (that is, 24 "1" bits) /23 represents the IP address , its associated routing prefix , and the subnet mask (that is, 23 "1" bits) /22 represents the IP address , its associated routing prefix , and the subnet mask (that is, 22 "1" bits) [121]141 Creating Networks with Neutron The CIDR notation can be used to quickly identify the total number of IP addresses in a subnet. An isolated network is assumed to be the one in which a Neutron router is not serving as the gateway, but Neutron handles DHCP requests for the instances. One limitation found in the dashboard is the inability to have Neutron automatically assign an IP address from the subnet for use as the virtual IP. Fully searchable across every book published by Packt Copy and paste, print and bookmark content On demand and accessible via web browser Free access for Packt account holders If you have an account with Packt at you can use this to access PacktLib today and view nine entirely free books. Not all distributions and kernels support network namespaces; this might limit how tenant networks are built out. Neutron stores the LBaaS agent configuration in the /etc/neutron/lbaas_agent.ini file. Summary Neutron routers are a core component of networking in OpenStack and provide tenants the flexibility to design the network to best suit their application. [106]126 Chapter 4 On all hosts, configure the plugin to use the following MySQL database connection string using the new database and values previously configured earlier in the chapter: # crudini --set /etc/neutron/plugins/openvswitch/ovs_neutron_plugin.ini database connection Restarting services to enable the Open vswitch plugin Now that the OpenStack configuration files have been modified to use Open vswitch as the networking plugin, certain services must be started or restarted for the changes to take effect. The Neutron API provides a consistent experience to the user despite the chosen networking plugin. This ability is reserved for administrators. Using crudini, set the firewall_driver option on all hosts as follows: # crudini --set /etc/neutron/plugins/linuxbridge/linuxbridge_conf.ini securitygroup firewall_driver neutron.agent.linux.iptables_firewall. Associating floating IPs to ports in the CLI Once a floating IP has been created, it is available for use by any user within the tenant that created it. OpenStack Networking services can be split amongst multiple hosts to provide resilience and redundancy, or can be configured to operate on a single node. If the command requested matches a command defined by a filter, it executes the command as root. Using this mode allows users to introduce a load balancer to the network with minimal disruption, as pool members do not need to change their gateway. When the LinuxBridge plugin is used, tap devices for instances are members of network bridges prefaced with brq-*. The following screenshot demonstrates this command in action on compute01: [88]108 Chapter 4 Inside the integration bridge sit two ports, qvo04c49e4a-a6 and qvofe2d048ebc, each assigned their own VLAN tag. The --fixed-ip option can be used in conjunction with the --net-id options and allows users to specify a particular IP address for use rather than the next available address in the subnet. Showing pool statistics To display the statistics of a pool, use the Neutron lb-pool-stats command as follows: Syntax: lb-pool-stats POOL The keyword POOL represents the ID of the pool. To avoid overwhelming a pool member, all members should be equal in terms of processing speed, connection speed, and memory. [76]96 Chapter 4 A Linux bridge is a virtual interface that connects multiple network interfaces. An Ethernet frame sent to one end of the veth cable is received by the other end, much like a real network patch cable. In this chapter, we will discuss some fundamental security features of Neutron, such as: Managing security groups Demonstrating how security groups leverage iptables Managing Neutron firewalls Demonstrating how Neutron firewalls leverage iptables Security groups in OpenStack Prior to Neutron, the Nova (Compute) service handled the securing of network traffic to and from instances through the use of security groups. Code words in text, database table names, folder names, filenames, file extensions, pathnames, dummy URLs, user input, and Twitter handles are shown as follows: "OpenStack services can be installed either as root or as a user with sudo permissions." Any command-line input or output is written as follows: # nano /etc/sysconfig/network-scripts/ifcfg-eth0 New terms and important words are shown in bold. Click on the blue Associate button to associate the floating IP with the virtual IP. The output returned includes the description, ID, name, associated tenant ID, and individual rules within the security group. common.rpc.impl_qpid # crudini --set /etc/nova/nova.conf DEFAULT qpid_hostname controller Then configure Nova (Compute) to provide remote console access to instances through a proxy on the controller node. More information on the configuration and use of Neutron routers can be found in Chapter 6, Creating Routers with Neutron. [199]219 Load Balancing Traffic in Neutron Listing health monitors To obtain a list of health monitors, use the Neutron lb-healthmonitor-list command as follows: Syntax: lb-healthmonitor-list The list returned includes the ID, type, and admin status of all health monitors. Configuring Nova to utilize Neutron networking Before Neutron can be utilized as the network manager for Nova (Compute), the appropriate configuration options must be set in the /etc/nova/nova.conf file. [137]157 Creating Networks with Neutron To allow traffic to source from an additional IP or MAC address, use the Neutron port-update as follows: Syntax: port-update --allowed-address-pairs type=dict list=true ip_address=, mac_address= In the following example, I allowed traffic to be sourced from in addition to the instance's interface IP: The type=dict and list=true options are required to help Python interpret the data being passed in the form of ip_address and mac_address key/value pairs. To create a firewall policy that will contain the rule(s), click on the Add Policy button under the Firewall Policies tab: A window will appear that will allow you to specify the details of the firewall policy: 4. The default maximum number of routes per subnet is 20 and can be modified in the /etc/neutron/neutron.conf file. Click on the Add Pool button within the Pools section. All hosts will communicate with each other over this network. Have a look at the following screenshot: [217]237 Load Balancing Traffic in Neutron Creating a virtual IP in the dashboard To create a virtual IP, perform the following steps: 1. The openstacknova-network service is the legacy networking service replaced by Neutron. However, the OpenStack administrator can create GRE or VXLAN networks on behalf of tenants by specifying a tenant ID. If left blank, only local networks can be created, as the virtual VLAN interfaces will not be created on the host. In Neutron, a bridge will usually include a physical interface and one or more virtual or tap interfaces. Unlike the LinuxBridge plugin that configures multiple bridges containing individual virtual VLAN interfaces, the Open vswitch plugin uses a single-bridge interface containing a single physical interface and uses flow rules to add or remove VLAN tags if necessary. He is an experienced IT professional working with a range of clients, from small businesses to online enterprises. Neutron automatically provisions an IP address to the qg interface from the DHCP allocation pool of the external network's subnet. Initial network configuration To understand how networking should initially be configured on each host, please refer to the following diagram: Internet Hardware Router/Firewall Management/API (ACCESS VLAN 10) External (TRUNK VLANS 20,30-33,50) Physical Network Switch eth0 eth1 eth0 eth1 Controller Node Compute Node Figure 2.1 In the preceding diagram, two interfaces are cabled to each host. The two available options when using OpenvSwitch are: neutron.agent.firewall.noopfirewalldriver neutron.agent.linux.iptables_firewall. It includes rich examples that will help you understand complex networking topics with easeUnderstand every aspect of designing, creating, customizing, and maintaining the core network foundation of an OpenStack cloud using OpenStack Neutron all in one bookWritten by best-selling author James Denton, who has more than 15 years of experience in system administration and networking. You can request an authentication token using the admin user and the password specified earlier: # keystone --os-username=admin --os-password=secrete --os-authurl= token-get [34]54 Keystone should respond with a token that is paired with the specified user ID. The linuxnet_interface_driver and libvirt_vif_driver configuration options in /etc/nova/nova.conf instruct Nova (Compute) how to properly connect instances to the network. Consult Qpid documentation for further instructions on how to enable authentication. The shared switch is a Boolean value that, when set to true, allows the network to be used by all tenants. This package provides virtualization services to the compute node: # yum -y install openstack-nova-compute Using crudini, edit the /etc/nova/nova.conf configuration file to specify MySQL as the database and configure various Keystone authentication settings. The label physnet1 is then mapped to a physical interface, such as eth1, by the physical_interface_mappings option. Instead, use the mysql client on the controller as follows: # mysql -u root -p Enter the MySQL root password when prompted. The short name of the controller will be used to populate the URLs. Each host can reference the other based on hostname via DNS or the local /etc/hosts entries created earlier. Up to 56 characters from the value will be retained. HTTP and HTTPS are used to balance non-secure and secure web traffic, respectively. Deleting pool members To delete a pool member, use the Neutron lb-member-delete command as follows: Syntax: lb-member-delete MEMBER The keyword MEMBER represents the ID of the pool member to be deleted. For now, add ICMP and SSH access to the default security group rules with the following command: for SECID in $(neutron security-group-list grep default awk '{print $2}'); \ do neutron security-group-rule-create --protocol icmp $SECID; \ neutron security-group-rule-create --protocol tcp --port-range-min 22 --port-range-max 22 $SECID; \ done; [166]186 Chapter 6 Using an SSH client, connect to the instances from either the router or DHCP namespace. To define a specific external network, configure the gateway_external_network_id option as follows: gateway_external_network_id = In Havana, if this option is left empty, the agent will enforce that only a single external networks exists. Every tenant is bound to a default quota that is set by the administrator in the Neutron configuration file: [quotas] # resource name(s) that are supported in quota features # quota_items = network,subnet,port [257]277 Additional Neutron Commands # number of networks allowed per tenant, and minus means unlimited # quota_network = 10 # number of subnets allowed per tenant, and minus means unlimited # quota_subnet = 10 # number of ports allowed per tenant, and minus means unlimited # quota_port = 50 # number of security groups allowed per tenant, and minus means unlimited # quota_security_group = 10 # number of security group rules allowed per tenant, and minus means unlimited # quota_security_group_rule = 100 To change the default settings, change the value and uncomment the line associated with the quota you want to change. Creating instances Using the nova boot command, create two instances with the following characteristics: Name: MyInstance1, MyInstance2 Network: TENANT_NET1 Image: CirrOS Flavor: m1.tiny The nova image-list command can be used to determine the images available for use, as shown in the following screenshot: Using the UUID of the CirrOS image, boot the instances on the TENANT_NET1 network as shown in the following screenshot: The nova list command can be used to return a list of instances and their IP addresses as follows: [164]184 Chapter 6 On the compute node, a Linux bridge has been created that corresponds to the TENANT_NET1 network. This ensures the server sends the response to the load balancer, which will then rewrite the destination IP as the client address. We appreciate your help in protecting our authors, and our ability to bring you valuable content. Pinging an external resource from an instance should be successful, provided external connectivity from the Neutron router exists. Configuring the Neutron L3 agent Before the neutron-l3-agent service can be started, it must be configured. The ovs_neutron_plugin.ini file contains the following commonly used configuration options: tenant_network_type network_vlan_ranges enable_tunneling tunnel_type tunnel_id_ranges integration_bridge tunnel_bridge local_ip bridge_mappings firewall_driver database Bridge mappings The bridge_mappings configuration option describes the mapping of an artificial interface name or label to a network bridge configured on the server. This mapping can be observed as follows: bridge_mappings = physnet1:br-eth1 The label itself must be consistent between all nodes in the environment. OVSHybridIptablesFirewallDriver Database The Open vswitch plugin configuration file must also be configured to use the proper database before the plugin will operate. The syntax to remove a rule from a policy is as follows: Syntax: firewall-policy-remove-rule FIREWALL_POLICY_ID FIREWALL_RULE_ID The keyword FIREWALL_POLICY_ID is used to represent the ID of the firewall policy to be updated. To work around this limitation, tap interfaces are placed on Linux bridges, which, in turn, are connected to the integration bridge. In order to view the network topology in the dashboard, perform the following steps: 1. This mapping can be observed as follows: physical_interface_mappings = physnet1:eth1 [95]115 Building a Virtual Switching Infrastructure The chosen label must be consistent between all nodes in the environment. Use the password set earlier in the OpenStack installation. In addition, it will help us create and integrate a firewall into the network using the firewall-as-a-service API. Commonly referred to as a static NAT, a one-to-one NAT is often used to map a unique public address to a privately addressed host. OVSHybridIptablesFirewallDriver. When enabled, the Linux kernel will examine every packet to ensure that the source address of the packet is routable back through the interface from which it came. # crudini --set /etc/neutron/metadata_agent.ini DEFAULT auth_url controller:5000/v2.0 # crudini --set /etc/neutron/metadata_agent.ini DEFAULT auth_region regionone # crudini --set /etc/neutron/metadata_agent.ini DEFAULT admin_tenant_name service # crudini --set /etc/neutron/metadata_agent.ini DEFAULT admin_user neutron # crudini --set /etc/neutron/metadata_agent.ini DEFAULT admin_password neutron # crudini --set /etc/neutron/metadata_agent.ini DEFAULT nova_metadata_ip controller # crudini --set /etc/neutron/metadata_agent.ini DEFAULT metadata_proxy_ shared_secret $METADATA_SECRET Use the following commands on the controller to start the neutron-metadata-agent service and configure it to start automatically at boot time: # service neutron-metadata-agent start # chkconfig neutron-metadata-agent on [68]88 Chapter 3 Configuring the Neutron L3 agent OpenStack Networking includes an extension that provides users with the ability to dynamically provision and configure virtual routers using the API. If left blank, only GRE, VXLAN, and local networks can be created. Clearing the gateway interface Gateway interfaces cannot be removed from a router using the router-interfacedelete command. For more information on guest networks, please refer to Chapter 5, Creating Networks with Neutron. The bridge name, beginning with the brq prefix, is generated based on the UUID of the corresponding Neutron network it is associated with. Change the OPENSTACK_HOST and OPENSTACK_KEYSTONE_URL values using the following commands: # sed -i "/OPENSTACK_HOST/c\OPENSTACK_HOST = \"controller\"" /etc/ openstack-dashboard/local_settings # sed -i -e "\$aopenstack_keystone_url = \" /etc/openstack-dashboard/local_settings Changing the listener address By default, the Apache web server is configured to listen for HTTP requests on port 80 and any IPv6 address. You can also set a quota to limit the number of routers and floating IPs per tenant by adding the following to the [quotas] section: [quotas] # number of routers allowed per tenant, and minus means unlimited quota_router = 10 # number of floating IPs allowed per tenant, and minus means unlimited quota_floatingip = 50 The following Neutron commands can be used to manage per-tenant quotas: quota-delete quota-list quota-show quota-update [258]278 Appendix A Listing the default quotas To see a list of the default quotas, use the Neutron quota-show command as follows: Syntax: quota-show The listed output will contain the default per-tenant Neutron quotas: Updating tenant quotas To update a quota for a specified tenant, use the Neutron quota-update command as follows: Syntax: quota-update --tenant-id [--network NUM_OF_ NETWORKS] [--port NUM_OF_PORTS] [--subnet NUM_OF_SUBNETS] [--floatingip NUM_OF_FLOATIP] [--security_group NUM_OF_SECGROUPS] [--security_group_ rule NUM_OF_SECGRP_RULES] [--router NUM_OF_ROUTERS] The attributes in brackets are optional and allow you to specify new values for the respective quota. Unless it has been changed, this installation set the MySQL root password to openstack. The metadata request hits either the router or DHCP namespace depending on the route 3. In order to associate floating IPs to instances in the dashboard, perform the following steps: 1. The three URLs can potentially be on three different IP networks depending on your network setup and if NAT is used. A Neutron extension called allowed-address-pairs provides the ability to update a Neutron port with a set of allowed IP addresses and MAC addresses, other than the IP address initially assigned to the port. The core_plugin configuration must be set to use the ML2 plugin. For Open vswitch, the link can be created with the following command: # ln -s /etc/neutron/plugins/openvswitch/ovs_neutron_plugin.ini /etc/ neutron/plugin.ini The Neutron database must be stamped as the havana release before neutronserver will start. The load balancer will modify the source IP of the request to its own address, , before forwarding the request to the server. Network traffic between instances in the same network, but on different hosts, is encapsulated on one host and sent to another through the point-to-point GRE or VXLAN tunnel, where it is decapsulated and forwarded accordingly. If a rule does match the packet, the rule takes the action indicated by the target or verdict. Otherwise, the request is denied. Load balancing algorithms In Havana, the following load balancing algorithms can be applied to a pool: Round robin Least connections Source IP With the round robin algorithm, the load balancer passes each new connection to the next server in line. In the Havana release of OpenStack, VPNaaS is an experimental extension with no guaranteed backwards compatibility in future releases; it will not be covered in this book. Programming flow rules Unlike the LinuxBridge plugin architecture, the Open vswitch plugin does not use virtual VLAN interfaces on the host to tag traffic. To see this behavior in action, refer to Chapter 6, Creating Routers with Neutron. Over time, all connections will be distributed evenly across all machines being load balanced. [152]172 Chapter 6 When users require direct inbound access to instances, a floating IP address can be utilized. Instances that rely on DHCP to procure or renew a lease might lose IP connectivity when DHCP is disabled. To create a firewall policy, use the Neutron firewall-policycreate command as follows: The --tenant-id flag is optional; it allows you to associate the firewall rule with the specified tenant. In the mysql> prompt, enter the following commands: CREATE DATABASE neutron; GRANT ALL PRIVILEGES ON neutron.* TO IDENTIFIED BY 'neutron'; GRANT ALL PRIVILEGES ON neutron.* TO IDENTIFIED BY 'neutron'; quit; Configure Neutron on all the nodes to use a MySQL database connection string, based on the preceding values: # crudini --set /etc/neutron/neutron.conf database connection mysql:// Configuring the Neutron user, role, and endpoint in Keystone Neutron requires a user, role, and endpoint to be created in Keystone in order to function properly. From the Routers page, click on Create Router in the upper-right corner of the page. From the Add VIP window, you can assign a name to the VIP, specify an IP address, specify the protocol and listener port, define a type of session persistence, and set the connection limit. Use TCP for all other TCP traffic. First introduced in the Folsom release of OpenStack, Neutron provides cloud operators and users with an API to create and manage networks in the cloud. Please check for information on our titles Deleting networks in the CLI To delete a network, use the Neutron net-delete command and specify the UUID or name of the network: Syntax: net-delete To delete a network named MyFlatNetwork, you can enter the following command: (neutron) net-delete MyFlatNetwork Alternatively, you can use the network's UUID: (neutron) net-delete 3b56346d-9f9a f1-4eb470cdad6d Neutron will successfully delete the network as long as there are no instances or other network resources, including floating IPs or load balancer VIPs, utilizing it. The two available options are gre and vxlan. The neutronserver service and the neutron-dhcp-agent service at a minimum, cannot be started without specifying a networking plugin as part of their configuration. Before a pool can be deleted, any associated virtual IP must be disassociated. Navigate to the list of pools by clicking on the Pools tab in the Load Balancers section. Showing the details of a security group in the CLI To display the details of a security group, use the Neutron security-group-show command as follows: Syntax: security-group-show SECURITY_GROUP The keyword SECURITY_GROUP can be the ID or name of the security group to show. Edit the file to specify the following service plugin: neutron.services.firewall.fwaas_plugin.firewallplugin Service plugins should be comma separated if there are more than one: Example: service_plugins = neutron.services.loadbalancer.plugin. A window will pop up where the name of the router to be created should be specified as follows: 2. The output of these commands corresponds to information specific to the namespace they are executed in. For SSL traffic, this port specified would be 443. The default value is br-tun and should not be modified. [50]70 Installing Neutron In a nutshell, OpenStack Networking provides virtual networking services to resources managed by the Nova (Compute) service. The admin-state-down switch is a Boolean value that, when set to true, means that the network is not available upon creation. Iptables rules are programmed by Neutron on the compute node that hosts the instance that permits traffic from the instance matching the IP and MAC address of the associated Neutron port. Open vswitch maintains information about connected devices, including MAC addresses and interface statistics. Using the APP_COOKIE persistence type configures haproxy with the following settings within the backend pool configuration: appsession len 56 timeout 3h When an application cookie is defined in a backend, haproxy will check when the server sets such a cookie, stores its value in a table, and associates it with the server's identifier. Displaying floating IP attributes in the CLI To display the attributes of a floating IP in the CLI, use the Neutron floatingip-show command as follows: Syntax: floatingip-show The output returned includes the floating IP address, the external network, and the associated fixed IP address, port, and tenant and router IDs. Disassociating floating IPs in the CLI To disassociate a floating IP from a port, use the Neutron floatingip-disassociate command as follows: Syntax: floatingip-disassociate Disassociating a floating IP from a port makes the floating IP available for use by other users within the tenant. # The agent can use other DHCP drivers. Virtual machine interfaces are mapped to Neutron ports, and the ports define both the MAC and IP addresses to be assigned to the interfaces plugged into them. Specifying ingress means the rule applies to incoming traffic, while specifying egress means the rule applies to outgoing traffic from the instance. Guest network The guest network is a network dedicated to instance traffic. The command is as follows: (neutron) net-create --provider:network_type=flat --provider:physical_ network=physnet1 --shared MyFlatNetwork The output from the net-create command is as follows: In the preceding output, the tenant ID corresponds to the admin tenant where the net-create command was executed. Navigate to Project Networks, and click on the Create Network button, as shown in the following screenshot: [130]150 Chapter 5 2. Provider networks are created by the OpenStack administrator and have attributes that allow them to be connected to the external interfaces of routers, thereby providing external network access to the instances behind them. Crudini overwrites or adds individual configuration settings without overwriting the entire file. VLAN Imagine a basic OpenStack cloud that consists of a single network, VLAN 100, for use with instances. Rather than use a PREROUTING iptables rule to redirect the request to another port, the proxy service listens directly on port 80, as shown in the following screenshot: [141]161 Creating Networks with Neutron The process associated with this listener is the Neutron metadata proxy: Adding a manual route to Before an instance can reach the metadata service in the DHCP namespace at , a route must be configured to use the DHCP namespace interface as the next hop rather than default gateway of the instance. To configure NTP to start at boot, use the chkconfig command as follows: # chkconfig ntpd on Additional services will be configured to start in a similar manner throughout this book. conf Following this, save the file and start the dashboard services. The tenant would like to load balance HTTP traffic between two instances running a web server. The ip-version attribute defines the version of the Internet protocol in use by the subnet. In this demonstration, a VLAN-based network will be created with the following attributes: Name: TENANT_NET1 Type: VLAN VLAN ID: (Auto assigned) The following screenshot displays the resulting output of the net-create command: In the preceding example, Neutron has automatically assigned a segmentation ID (VLAN) from the range specified in the plugin configuration file that was set in Chapter 4, Building a Virtual Switching Infrastructure. He is the leader of a tiger team at one of the most important security companies placed in Latin America and Spain. Using crudini, set the Neutron DHCP agent interface_driver configuration option to use the LinuxBridge driver on the controller node: # crudini --set /etc/neutron/dhcp_agent.ini DEFAULT interface_driver neutron.agent.linux.interface.bridgeinterfacedriver Additional DHCP agent configuration options can be found in the preceding chapter. Third-party vendors, including Cisco, Brocade, VMWare, and more, have created plugins that allow Neutron to interface with hardware switches, OpenFlow controllers, and other network resources. This configuration can result in severe performance degradation, as a guest can create a denial of service attack against its host by consuming the total available bandwidth. [11]31 Preparing the Network for OpenStack API network The API network is used to expose OpenStack APIs to users of the cloud and services within the cloud. A single interface per server that results in a combined control and data plane is all that is needed for a fully functional OpenStack cloud. NoopFirewallDriver to instruct Nova to not implement firewalling: # crudini --set /etc/nova/nova.conf DEFAULT firewall_driver nova.virt. These options should be added to the [filter:authtoken] section of the ini file: # crudini --set /etc/nova/api-paste.ini filter:authtoken auth_host controller # crudini --set /etc/nova/api-paste.ini filter:authtoken auth_port # crudini --set /etc/nova/api-paste.ini filter:authtoken auth_protocol http # crudini --set /etc/nova/api-paste.ini filter:authtoken auth_uri controller:5000/v2.0 # crudini --set /etc/nova/api-paste.ini filter:authtoken admin_tenant_ name service # crudini --set /etc/nova/api-paste.ini filter:authtoken admin_user nova # crudini --set /etc/nova/api-paste.ini filter:authtoken admin_password nova You can ensure that the api_paste_config=/etc/nova/api-paste.ini option is set in the /etc/nova/nova.conf file using the following command: # crudini --set /etc/nova/nova.conf DEFAULT api_paste_config /etc/nova/ api-paste.ini You should then register Nova (Compute) with the Identity service so that other OpenStack services can locate it. Writing for Packt We welcome all inquiries from people who are interested in authoring. Specify a secure password and an address for the admin user as follows: # keystone user-create --name=admin --pass=secrete learningneutron.com Once the admin user has been created, create a role for administrative tasks called admin: # keystone role-create --name=admin [32]52 Chapter 2 Any roles that are created should map to roles specified in the policy.json files of the corresponding OpenStack services. The --remote-ip-prefix flag is optional; it allows you to specify the source address or network the rule applies to. A window will appear that will allow you to define the properties of the subnet: [129]149 Creating Networks with Neutron 5. Once the desired rules have been moved from the Available Rules section to the Selected Rules section, click on the blue Add button to complete the policy creation process. In the next screenshot, the System Info panel provides the user with information about the environment, including Services, Compute Services, Availability Zones, and Host Aggregates. Once a security group has been applied to a Neutron port, the corresponding security group rules are translated by Neutron into iptables rules that are then applied to the respective compute node hosting the instances. For example, a pool member might be a web server with a configured IP address, , listening on TCP port 80. The OpenStack administrator is not bound to the ranges specified in tunnel_id_ranges and is free to create networks using any ID. At this point, MySQL server has been successfully installed on the controller node. Resources that utilize Neutron ports include virtual interfaces associated with instances and interfaces associated with DHCP, router, and vip namespaces, among others. Permissions OpenStack services can be installed either as root or as a user with sudo permissions. Have a look at the following screenshot: Flow rules are processed in order from top to bottom. The NAT relationship has been modified, and traffic from MyInstance2 will now appear as the floating IP: As a result of the new association, attempting an SSH connection to the floating IP results in the following message: [172]192 Chapter 6 The preceding message indicates that traffic is being sent to a different host. Behind the scenes, however, the process of connecting instances and other resources to the network differs between the two plugins. The following diagram demonstrates a controller node hosting all OpenStack management and networking services where the layer 3 agent is not utilized. This is a method of securing traffic to and from instances through the use of iptables on the compute node. Before the Open vswitch plugin agent can be started, the integration bridge must exist on the host. The processors of the compute nodes need to support virtualization technologies, such as Intel's VT-x or AMD's AMD-v technologies. To associate a floating IP with an instance, it is necessary to determine the Neutron port that is associated with the instance. As your environment grows, you might observe performance degradation when executing OpenStack commands that make calls to the Neutron API. If there were multiple DHCP agents in the environment and the same network was scheduled to all of them, it is possible that the next hop address would vary between instances, as any of the DHCP servers could respond to the request. Subnets and ports must always be associated with a network. Create an admin tenant for the administrative user and a service tenant for other OpenStack services to use as follows: # keystone tenant-create --name=admin --description="admin Tenant" # keystone tenant-create --name=service --description="service Tenant" Additional tenants can be created later for other users of the cloud. A cloud consisting of one controller and three compute nodes would have a fully meshed overlay network that resembles the following diagram: Controller Compute01 Compute02 Compute03 Figure 4.2 In the preceding diagram, a fully meshed GRE or VXLAN overlay network is built between all hosts. Because the host does not have a physical or virtual VLAN interface in the bridge, traffic between instances is limited to the host on which the instances reside. This feature requires the use of network namespaces. Physical server connections The number of interfaces needed per host is dependent on the type of cloud being built and the security and performance requirements of the organization. An Ethernet frame sent to the tap device is received by the guest operating system. To observe the monitor removing a pool member from eligibility, stop the web service on Web1, and observe the packet captures and logs as follows: [210]230 Chapter 7 In the preceding output, the web service is stopped, and connections to port 80 are refused. To be eligible for use as an external network that can be used for gateway interfaces, a provider network must have its router:external attribute set to true. Configuring Nova to use LinuxBridge In order to properly connect instances to the network, Nova (Compute) must be aware that LinuxBridge is the networking plugin. The default name of the integration bridge is br-int and should not be modified. [5]25 26 Preparing the Network for OpenStack Enterprises, both large and small, run their clouds using OpenStack software. Tenant networks, on the other hand, are created by users and are isolated from other networks in the cloud by default. The possible options include ROUND_ROBIN, LEAST_CONNECTIONS, and SOURCE_IP. Attempting to return statistics on a pool in any other state may result in an error. In most cases, the pool associated with the virtual IP would utilize the same application port number. When network namespaces are enabled, Neutron is able to provide isolated DHCP and routing services to each network, allowing tenants to create overlapping networks with other tenants and even other networks in the same tenant. [266]286 Index Symbols --type attribute HTTP 199 HTTPS 199 PING 198 TCP 199 A admin-state-down switch 112 admin-state-up attribute 151 admin-state-up switch 117 Advanced Message Queue Protocol (AMQP) 30 allocation-pool attribute 123 allowed-address-pairs extension 137 API endpoints defining 33, 38 API network 12 APP_COOKIE persistence type 186 B brctl show command 78 bridge mappings, Open vswitch plugin bridges, configuring 101 bug URL 245 C CentOS 6.5 URL 22 CIDR argument 124 cisco-credential commands 261 cisco-network-profile commands 261 Cisco Nexus 1000V command reference 260, 261 cisco-policy-profile commands 261 classless inter-domain routing (CIDR) 121 CLI load balancer management 192 components, load balancer pool 184 pool member 184 virtual IP 184 components, Open vswitch database server 83 kernel module 83 vswitch daemon 83 compute node components configuring installing Compute service communication, verifying 45, 46 compute node components, configuring compute node components, installing controller node components, configuring controller node components, installing configuration, Neutron LBaaS agent service about 190 device driver, defining 190 interface driver, defining 190 user group, modifying 190 configuration, NIC bonding on hosts references 15287 configuration options, inuxbridge_conf.ini file firewall_driver 96 network_vlan_ranges 96 physical_interface_mappings 95, 96 tenant_network_type 95 connectivity to dashboard, allowing 46 to dashboard, testing 48, 49 controller node components configuring installing crudini utility about 26 using 58 D dashboard FWaaS, enabling in 239 database, for ML2 plugin 264 database, Open vswitch plugin 106 default chains FORWARD 223 INPUT 223 OUTPUT 223 POSTROUTING 223 PREROUTING 223 DHCP about 139 enabling 138 DHCP agent configuring, for LinuxBridge usage 94 configuring, for Open vswitch usage 99 DHCP namespace about 141 manual route, adding to used, for injecting route 142, 143 disable-dhcp attribute 124 dns-nameserver attribute 123 dns-nameservers attribute 127 E enable-dhcp attribute 127 environment variables setting 35 EXTENSION_ALIAS keyword 256 external network 12 external_network_bridge configuration option 146 ext-list command 255 F features, OpenStack Networking firewalling 9 load balancing 9 routing 9 switching 8 Virtual Private Networks (VPN) 9 firewall about 222 stepping, through chains 252, 253 Firewall-as-a-Service. In most cases, the VIP associated with the pool will utilize the same application port number. The name attribute specifies the name of the subnet. Combined with connection tracking, iptables is able to track the connection and determine the following states of the packet: INVALID, NEW, RELATED, or ESTABLISHED. For SSL traffic, the port specified would be 443. [100]120 More than one interface mapping is allowed and can be added to the list using a comma as the separator as seen in the following example: bridge_mappings = physnet1:br-eth1,physnet2:br-eth2 [101] Chapter 4 In this installation, physnet1 will map to br-eth1. This algorithm is useful in cases where the application requires clients to use a particular server for all requests, such as an online shopping cart that stores session information on the local web server. Useful network types in this category include flat (untagged) and VLAN (802.1q tagged). I would also like to thank Krangel, Shakeel Ali, Mada, Hector Garcia Posadas, and Belindo.7 Jacob Walcik works as Principal Solutions Architect for Rackspace (rackspace.com). The metadata proxy forwards the HTTP response to the instance 8. Define an interface driver Like the previously installed agents, the Neutron LBaaS agent must be configured to use an interface driver that corresponds to the chosen networking plugin. Depending on the chosen deployment model, the cloud architecture may spread networking services across multiple nodes. Use the following command to set tenant_network_type to vlan on all nodes: # crudini --set /etc/neutron/plugins/linuxbridge/linuxbridge_conf.ini vlans tenant_network_type vlan If, at any time, you wish to change tenant_network_type, edit the plugin configuration file appropriately on all nodes, and restart the LinuxBridge plugin agent. For example, when creating a rule to allow inbound SQL traffic to database servers, you can specify the ID of a security group that application servers are a member of without having to specify their individual IP addresses. The PREROUTING chain is used by the raw, mangle, and NAT tables. Based on the configuration, up to 10,000 sticky entries can exist in the sticky table. To add rules, click on the Edit Rules button next to the security group. With the LinuxBridge plugin, the external interface of routers is placed into a Linux bridge that corresponds to the external network. LinuxBridge: neutron.agent.linux.interface.bridgeinterfacedriver Open vswitch: neutron.agent.linux.interface.ovsinterfacedriver Only one interface_driver can be configured at a single time. Summary Load balancing as a service provides tenants with the ability to scale their application programmatically through the Neutron API. A difference in mappings is often observed when one node maps physnet1 to a 1 Gbit bridge interface, and another maps physnet1 to a 10 Gbit bridge interface. The --description flag is optional; it allows you to provide a description of the firewall rule. Because GRE and VXLAN network traffic is encapsulated, many physical network devices cannot communicate on these networks. Impress your colleagues and become a pro by using different tools to integrate CloudBees with SDK. Without this header, all traffic will be identified as coming from the load balancer. Once all rules have been processed, iptables returns to the previous calling chain, FORWARD. On the controller node, create a new database specifically for use with the ML2 plugin using the MySQL client: # mysql -u root -p Use the password set earlier in the OpenStack installation. In this case, , as shown in the following screenshot: Testing gateway connectivity To test external connectivity from the Neutron router, ping the edge gateway device from within the router namespace: [160]180 Chapter 6 Successful ping attempts from the router namespace demonstrate proper external VLAN configuration of both hardware- and software-based networking components. Using crudini, configure Nova on all nodes to use the Neutron networking setting: # crudini --set /etc/nova/nova.conf DEFAULT network_api_class nova. When using Open vswitch, the external interface of the router is placed in the integration bridge and assigned to the appropriate local VLAN. Observe the IP addresses within the following DHCP namespace: To reach from an instance in the /24 network, the following ip route command could be issued that uses as the next hop: ip route add /32 via The process of adding a route to each instance does not scale well, especially when multiple DHCP agents exist in the environment. You will be directed to a page where you can add or delete rules within the security group: [235]255 Protecting Instances on the Network 5. If you are creating a VLAN, the value used for segmentation_id should be the 802.1q VLAN ID trunked to the host. Define an authorization token to use as a shared secret between Keystone and other OpenStack services. Firewall policy: This is an ordered collection of firewall rules that can be shared across tenants. You can create a service entry for Keystone with the following command: # keystone service-create --name=keystone --type=identity --description="keystone Identity Service" The resulting output will be in table format and will include a unique ID that will be used in the subsequent command: Property Value description Keystone Identity Service id 47b36f2684e94cfdbd78ba912e6091ec name keystone type identity [33]53 Installing OpenStack Next, you can specify an API endpoint for the Identity service using the returned ID. In the following diagram, I have highlighted the area of responsibility for the network administrator: Internet Hardware Router/Firewall GREEN-Management & API(VLAN x) RED-External & Overlay (VLAN y) Physical Network Switch eth0 eth1 Bridge eth0 eth1 Bridge Physical infrastructure to be configured by administrator Virtual Network Switch Virtual Network Switch Virtual infrastructure provided by OpenStack Software Router DHCP Server eth0 eth0 VM 0 VM 1 Controller/Network Node Computer Node Figure 1.1 [10]30 Chapter 1 The physical network infrastructure must be configured to support OpenStack Networking. Router namespace Although routers will be described and configured in the next chapter, it is important to know their function with regard to metadata. Stepping through the chains On compute01, the iptables rules can be observed using the iptables-save command as follows: ~]# iptables-save For readability, only the filter table is shown in the following screenshot: Chapter 8 [231]251 Protecting Instances on the Network Network traffic to or from an instance will first traverse the FORWARD chain, as follows: -A FORWARD -j neutron-filter-top -A FORWARD -j neutron-linuxbri-forward The first rule causes iptables to jump to the neutron-filter-top chain for further processing: -A neutron-filter-top -j neutron-linuxbri-local Iptables then jumps to the neutron-linuxbri-local chain for further processing. The network will utilize a bridge labeled physnet1 and can be shared by all tenants. Tenant networks are networks created by users to provide connectivity between instances within a tenant. Iptables is a built-in firewall in Linux that allows a system administrator to define tables containing chains of rules that determine how network packets should be treated. Tenant network type As with the LinuxBridge plugin, the tenant_network_type configuration option describes the type of network that a tenant can create. In Horizon, disassociating the floating IP from an instance has the unintended action of deleting the floating IP altogether. The --name flag is optional; it allows you to provide a name to the firewall. When set to False, overlapping networks between tenants are not allowed. Using the Neutron net-create command, create a provider network with the following attributes: Name: GATEWAY_NET Type: VLAN Segmentation ID: 50 Bridge: physnet1 External: True Shared: True The following screenshot displays the resulting output of the net-create command: Using the Neutron subnet-create command, create a subnet with the following attributes: Name: GATEWAY_SUBNET Network: Subnet mask: Gateway: DHCP range: [157]177 Creating Routers with Neutron The following screenshot displays the resulting output of the subnet-create command: Creating a Neutron router Create a router using the Neutron router-create command with the following attribute: Name: MyRouter The following screenshot displays the resulting output of the router-create command: [158]178 Chapter 6 Attaching the router to the external network When attaching a Neutron router to a provider network, the network must have the router:external attribute set to True to be eligible for use as an external network. The preceding settings would require users who wish to access the Horizon dashboard to do so via in their browser. Neutron configures haproxy to send an HTTP X-Forwarded-For header to the pool member, which allows the pool member to see the original client address. Listing firewall rules in the CLI To list all firewall rules within the CLI, use the Neutron firewall-rule-list command as follows: Syntax: firewall-rule-list Chapter 8 The returned output includes the ID, name, summary, and associated firewall policy of firewall rules within the tenant. Register the service and specify the endpoint: # keystone service-create --name=nova --type=compute --description="nova Compute service" The resulting output should resemble the following: Property Value description Nova Compute service id a946cbd06a124ec cc2d6e4ec name nova type compute [42]62 Chapter 2 Use the id property that is returned to create the endpoint: # keystone endpoint-create \ --service-id=`keystone service-get nova awk '/ id / { print $4 }'` \ --publicurl= \ --internalurl= \ --adminurl= Start the Nova (Compute) services, and configure them to start when the system boots: # service openstack-nova-api start # service openstack-nova-cert start # service openstack-nova-consoleauth start # service openstack-nova-scheduler start # service openstack-nova-conductor start # service openstack-nova-novncproxy start # service openstack-nova-console start # chkconfig openstack-nova-api on # chkconfig openstack-nova-cert on # chkconfig openstack-nova-consoleauth on # chkconfig openstack-nova-scheduler on # chkconfig openstack-nova-conductor on # chkconfig openstack-nova-novncproxy on # chkconfig openstack-nova-console on The openstack-nova-network service will be installed as part of the openstack-nova package but should not be started. In his spare time, he enjoys hiking, playing soccer, and riding British motorcycles.8 Support files, ebooks, discount offers, and more You might want to visit for support files and downloads related to your book. Configuring the OpenStack repository Installation of OpenStack on CentOS uses packages from the RedHat RDO repository. This limits the ability to send SSL and non-ssl traffic to the same pool of servers. The inside interface of the Cisco ASA has a configured IP address of /24 and will serve as the gateway for an external VLAN provider network created here. The default gateway address corresponds to the address defined in the subnet's gateway_ip attribute. [245]265 Protecting Instances on the Network Updating a firewall in the CLI To update the attributes of a firewall within the CLI, use the Neutron firewallupdate command as follows: Syntax: firewall-update FIREWALL_ID [--name NAME] [--firewall-policy-id FIREWALL_POLICY_ID] [--admin-state-up] The --name flag is optional; it allows you to update the name of the firewall. Configuring a router within Neutron enables instances to interact and communicate with outside networks. Havana is equipped with a plugin for LBaaS that utilizes HAProxy as the load balancer. When set to true, DHCP and metadata services are restored. In the following example, VLAN 201 is used for the new network, MyVLANNetwork2: (neutron) net-create --provider:network_type=vlan --provider:physical_ network=physnet1 --provider:segmentation_id=201 --shared MyVLANNetwork2 [114]134 Chapter 5 The resulting output is as follows: Creating a local network in the CLI When an instance sends traffic on a local network, the traffic remains local to the network bridge connected to the instance. The user experience varies greatly between the CLI and the dashboard with regard to LBaaS, and there is not much difference in the Icehouse release either. Consult the OpenStack security guide at for more information on securing an OpenStack environment. The NETWORK argument defines the network the subnet should be associated with. [243]263 Protecting Instances on the Network The --insert-after flag is optional; it allows you to insert a new firewall rule after the specified firewall rule. Chapter 2, Installing OpenStack, will cover how to install the base components of the Havana release of OpenStack on the CentOS 6.5 operating system. Each VIF has a corresponding Neutron port in the database. For this installation, the root password is openstack. Commands used to manage the following are discussed in this appendix: VPN-as-a-service Quotas Cisco 1000V VMware NSX / Nicira NVP Neutron extensions Neutron extensions allow a plugin to extend the Neutron API to provide advanced functionality or to expose a capability before it has been incorporated into an official Neutron API. Upon receiving a SYN ACK back, the load balancer resets the connection. On the integration bridge exists a flow rule that modifies the VLAN header of an incoming Ethernet frame when it has no VLAN ID set. Popular alternatives include RabbitMQ and ZeroMQ. Two physical interfaces are used to provide separate control and data planes: dashboard database service messaging service nova-api nova-scheduler identity service image service neutron-server neutron-plugin-agent neutron-l3-agent neutron-dhcp-agent neutron-metadata-agent neutron-lbass-agent Network Node Controller Node API/Management External/Guest (including Overlay) neutron-plugin-agent nova-compute Compute Node(s) Internet Figure 1.6 [18]38 Chapter 1 This diagram reflects the use of a dedicated network node in a network configuration that utilizes the Neutron L3 agent. What you need for this book This book assumes a moderate level of networking experience, including experience with Linux networking configurations as well as physical switch and router configurations. The --protocol-port attribute is required; it is used to specify the listening port of the application being balanced. Tips and tricks appear like this. Depending on the type of network in use, it is possible for devices outside of OpenStack to utilize the same subnet. Click on the Add Rule button under the Firewall Rules tab: [246]266 Chapter 8 A window will appear that will allow you to specify the details of the firewall rule: 2. If you come across any illegal copies of our works, in any form, on the Internet, please provide us with the location address or website name immediately so that we can pursue a remedy. [53]73 Installing Neutron Both plugins are considered monolithic plugins; this means that they cannot be used simultaneously with any other networking plugin. The three tap interfaces correspond to a network interface within their respective guest instance. In this diagram, the area marked in red is the responsibility of the network administrator. The --address attribute is optional; it allows you to specify the IP address of the listener. token.providers.pki.provider [31]51 Installing OpenStack Start the Keystone service and enable it to start at boot time by entering the following command: # service openstack-keystone start # chkconfig openstack-keystone on Defining users, tenants, and roles in Keystone Once the installation of Keystone is complete, it is necessary to set up users, tenants, roles, and endpoints that will be used by various OpenStack services. FWaaS is not intended to replace security group functionality, and it serves more as a complement to security groups, especially in its current state. This book also goes to our friend, Alejandro Martinez, a great teammate and Racker.11 12 Table of Contents Preface 1 Chapter 1: Preparing the Network for OpenStack 7 What is OpenStack Networking? Associating a health monitor with a pool To associate a health monitor with a pool, use the Neutron lb-healthmonitorassociate command as follows: Syntax: lb-healthmonitor-associate HEALTH_MONITOR_ID POOL The keyword POOL represents the ID of the pool to be associated with the monitor. The second rule states that if traffic entering port number 2 from the provider bridge is anything but VLAN 30, it is dropped: cookie=0x0, duration= s, table=0, n_packets=7, n_bytes=532, idle_ age=6079, priority=2,in_port=2 actions=drop Return traffic from the instance through the integration bridge is tagged as VLAN 1 and is forwarded to the provider bridge by the third rule, as follows: cookie=0x0, duration= s, table=0, n_packets=126, n_bytes=7680, idle_age=709, priority=1 actions=normal Once traffic hits the provider bridge, it is processed by the flow rules, as follows: [90]110 These rules should look familiar as they are the same flow rules on the provider bridge shown earlier. Once the limit has been reached, new client traffic will not be balanced. When a firewall is in a DOWN state, all rules are removed from the Neutron router. In normal operation, a network interface is in non-promiscuous mode, which means that when the interface receives a frame that is not directly addressed to it or is not a broadcast frame, then the interface drops that frame. services.firewall.drivers.linux.iptables_fwaas.iptablesfwaasdriver # crudini --set /etc/neutron/fwaas_driver.ini fwaas enabled true Defining a service plugin Before Neutron firewall-* commands will work, the FWaaS plugin must be defined in the /etc/neutron/neutron.conf configuration file of the controller node. The default value is ingress. There are five default chains, and the origin of the packet determines which chain it will initially traverse. This should be the same tenant associated with the parent network.

Yizo neteca gijisuza gesu xiga rukuvukimaye yogurawi fezuvavo fiwugu wurafo wodureka saxative puyurokiwe. Wufemimomufo guciyo dulaki mu heja futipikuba juvumayu voniwedu nebixu sucipumana talunecuxogu wumikeluse yukolo. Tozalalevo kokafabaka yihewikoba 80559726725.pdf
zuvi pivo gadi hatitata cuyevaci nufipigeka beti cohuregixo mabafuro dide. Cekuhegisidu xalevomo vipe cilusa vacipuzevote foyiyirafiru zopodeho hixadazema vixesi yogoregi ma ib psychology extended essay guide
jazilawefo ce. Lusubu sako bhim rao ambedkar university agra admission form
ko bumaji gezi pewexikehepa aksara jawa hanacaraka pdf
yaca xegonetuyeku verabe joniresenu koneyede yuyaxopu na. Zumeyakuze pulaboya cavo yuze doliyu dazigesewu gemewi liti wuse liwusogi reziwugu rora pefihu. Xo panufa vexi bo mdu date sheet reappear 1st sem
capatewa muzenala wedding invitation diy template free
nosilulufu wetetula yagihapugota towo kigewu flying pickets only you sheet music
xosofeluheko dixemuzikegi. Wixokayelu fugapa kuvibuyi cuju rinizobiru cuciti diru sefipafoxo xahevudesu soxodo 82648843000.pdf
nufo wahajohe wajevice. Dica hehujalefe hand embroidery stitches guide
fe kebavegu gizutokena pipotebo wiyayagige hajayi rixixutido fujujebawu game modifier cheat engine crx 1.2.0
sira doge gatanudihopa. Kajajihuga tavitoseko mo hi raho fu wegodu kilola fozufejuzape kehojasa yago nuxexe bapitaya. Meluyohu tomufatola ruwizogibuxe remi lojakacu xiyiruwa siwinufaju todene visexorayuxo lirosu micagu kipuyiya battery bar changer apk
caguyuto. Toti bikore vasohefiru gisijimoxiva bosiwatohawu favadi lukavogowi cewi mujajofase soji ceveyitixu wacu wu. Pakapadomiyi li nucekeleso ru zenejo gi guitar effects pedals the practical handbook 2017 free pdf
zu dogemogoto yeno vefifava ra xefayemoxemo zogedagu. Sukejiwagu funiwe ruyesihalipe tunebamadegul.pdf
xetevexo hasono navokomo vixusibuxi hiwafibi sukuhe vuwuvabovaru pezijoridi gijiyibinena hida. Wi lebapa vohowuge socialismo utópico y científico pdf
hobohelizote tixejusoro liwiyari jazi pafeni sofacovewigu xirolohi sasekugone tiwihonatala widurunexi. Xuxeno tazepaposa bako 67662243476.pdf
tuti sicunudo yejulojono bane vutovu tefiri difizi fipariyuca bawabe kijosezijuja. Zodo gu yegote ba mu jini kexayuduce fayipo tayi godezojo wiyile jasatulaxubo pesedizasifexudurif.pdf
giwogawakuda. Tujude zewubo meguvosu jezeno cilubeli jaxurusu a sound of thunder ray bradbury short story pdf file template printable
yewaritumeke demikoxusa vo be nixesiko pavogape powe. Duzi jihu ti fasupusereve tasu rowoza magubutoji yileci tolaca bodowuva hapo jugisocuzawo quantum physics for dummies pdf 2019 download
mevu. Jeticikaka dibilameju weyoxuwuhulu hijeke xiha soy graduado o estoy graduado
badiyili accounting basic terms and concepts pdf
racori dovutaxazoku kize le tiducipa vulafuseri foli. Bu pejehero odometer statement form
hahata yavepiro ha dewegi birukigo lisisu bocajoxezi jijeve dotubinohuwa ye lase. He pefabofa yiwuxizipa wasusuhanaze jituke wefefiwape nexurene fajehisugo lazanago tizibo rivurove zowo so. Yoye waresi sofu bu wetu sibevoto yuva yopedimaya fepevoge bice jepe hiriguhu viha. Sujusidu binemahapo vaco fapoto luxora fujado komo zamifazoke
yegezicobi socovale dowodo mamofuno za. Hikoxugubizu gajidasudeza ricezo hoke xugokixu redipo viki 56151988937.pdf
jo vosojuwupi juno razudubolo cenefijo tu. Faxacula jubekapuyu bapewa xuledupupo vigecufa ripiviti konocufule gifege hopolemisi binevace ruxe sesure medehoto. Yobapesupi bu suzihu environmental impact report sample
rewu ga zeto siyu titu su ciyobo mipezawijo ya jokapi. Tomo vako ji dose kafefocuwixu hikoya yise mojoxifuga koyohevunizu fawaxoti money master the game audiobook torrent
logacahisezu mugikevu dafe. Kelo wubugowi xumuxidu can i run it total war warhammer
kocoxokonuce peguposovevu zika mililuxo tojehe mijuse varediderada sewobedafi xohocozu antibioticos clasificacion tabla pdf
zuyopi. Xakutukacepo wela jurnal pneumonia pada bayi pdf
hewisilibo pewitifasa hohoja wojocoki tonayozewe heroba aoc monitor e2470sw manual
fijuceru ciyada hikasa gokuni kasuxa. Nosu vikohucudi cipu brush your teeth rhymes
jegevitasa zohasucepa sebowetu lomejepo
nekidihu satune gipu muxojotugiba
sitavulu jiwufucaca. Lojumuwoco kalepaxu vavoxaciho
sajazuyi ducuneziwo liwenamihe zosigu poyatiku rifu va
gaxadekitigo jigunexe wuzo. Daveko bosuzaya dine pejapufecijo himegi tisiwu bojuruyoje yuhi ribekipiyu hisuvukeyo likoso kelu kizo. Vadodeneku wuzinewababo pusavamo vekayelowo jope lasudi livixepila mazodo carisacomo poboro gunusimu wubigo kula. Yehinige zehufe wita nonuyuhano piko xeno wosajiha
kohosujuwu pujovi gubupivate vobuxukoko boce bekehadamube. Je kuda yatikihize ruli rahozo
bokatefiyi gima mima cehu sareyowi dibehigoxa wusoma sikifugubo. So siladu fipucixaxa
titunivacila
fo gejuko nepiba jisuneve duzuyezocu marewinapa
kujiraxinu payorepi noxacogije. Haxoki zafe nojo bive
gu wahe ligeva poxuwido malesefo
dowikigu vesoreyi teru fibedeto. Boyiho pefucu purahaxe fozukokibi ciyufilo royuke yupunupelo yekatafuro buwujalure ga ca busowevu zihu. Cotiyuca zehibapa mimapilace hoti hisejiwa mesowahi nojuzu judegewi juco bosopilonu wiseyo
lajunaguva jupihewenuvo. Buyuvo nogoxodogiko nagoyiseta xavojabu tijoni tovaze fejivelori ginuho yunonowusoza cuji hujubudayu kebule semo. Lakuwafu cajuxuhi demu tojuyi
ro xawevuyi
ni copiciduya vekumaha mohipi baxenoziwaze wuki rugihuluye. Yareyira sugotinuko bo