

hemanthravi joined the chat room.
[09:04am] SumitNaiksatam: hemanthravi: hi
[09:04am] hemanthravi: SumitNaiksatam, hi
[09:04am] mageshgv: SumitNaiksatam: hi
[09:04am] rkukura: hi
[09:04am] ivar-laz_ joined the chat room.
[09:04am] SumitNaiksatam: rkukura: was earlier asking if we are doing this over hangout as well
[09:04am] SumitNaiksatam: ivar-laz_: hi
[09:05am] SumitNaiksatam: so for now, let roll with the IRC, and we can resort to the hangout if we need a higher bandwidth channel
[09:05am] ivar-laz_: hi
[09:05am] SumitNaiksatam: rkukura: hemanthravi mageshgv ivar-laz_: sound okay?
[09:05am] rkukura: +1
[09:05am] SumitNaiksatam: rkukura: okay
[09:05am] hemanthravi: ok
[09:05am] mageshgv: ok
[09:05am] ivar-laz_: SumitNaiksatam: +1
[09:05am] SumitNaiksatam: ok lets get started
[09:06am] SumitNaiksatam: firstly mageshgv thanks for putting fanstastic effort on these patches
[09:06am] hemanthravi: +1
[09:06am] SumitNaiksatam: it remains to be seen if they stand the test of our reviews
[09:06am] mageshgv:
[09:06am] SumitNaiksatam: but at least from a functionality perspective, i was able to test it out yesterday
[09:07am] ivar-laz_: SumitNaiksatam: nice!
[09:08am] SumitNaiksatam: so i have a setup in which i was able to create a NSP (network service profile), create FW and LB service nodes, create a service chain spec from it
[09:08am] SumitNaiksatam: and then add the chain spec to a redirect action
[09:08am] ivar-lazzaro left the chat room. (Ping timeout: 265 seconds)
[09:08am] mageshgv: ivar-laz_: Thanks for your comments earlier. Was able to incorporate a few of then earlier
[09:08am] SumitNaiksatam: a contract/policy-rule-set was created with this policy-rule
[09:09am] SumitNaiksatam: mageshgv: one sec
[09:09am] SumitNaiksatam: when an EPG/PTG consumed this policy-rule-set, a service chain instance was created
[09:09am] SumitNaiksatam: along with the services instances
[09:09am] SumitNaiksatam: FW and LB
[09:10am] SumitNaiksatam: i did not get a chance to verify the datapath
[09:10am] SumitNaiksatam: but this is mighty cool!
[09:10am] ivar-laz_: SumitNaiksatam: I had the impression that only VPN and FW were supported
[09:10am] SumitNaiksatam: mageshgv: thanks again for the instructions

[09:10am] SumitNaiksatam: ivar-laz_: we are also doing FW, LB
[09:11am] SumitNaiksatam: for this first iteration we will focus on testing that [FW, LB] chain works
[09:11am] ivar-laz_: SumitNaiksatam: I was referring to <https://review.openstack.org/#/c/128551/5/gbp/neutron/extensions/servicechain.py>
[09:11am] ivar-laz_: SumitNaiksatam: L65
[09:11am] ivar-laz_: SumitNaiksatam: I guess it had been added later
[09:12am] ivar-laz_: s/had/has
[09:12am] SumitNaiksatam: i guess mageshgv can respond to that, but ivar-laz_, per mageshgv's earlier comment - thanks for the initial set of reviews!
[09:13am] mageshgv: ivar-laz_: Yes, That is not validated today, because in the absense of service framework, it is only an indication. There is no way to enforce it
[09:13am] SumitNaiksatam: mageshgv: but perhaps we need to add LB to that list as well
[09:14am] SumitNaiksatam: mageshgv: also a comment to the effect of what you just said might be helpful
mageshgv: SumitNaiksatam: Ok. Will see if we can validate that as well
[09:14am] ivar-laz_: mageshgv: ok clear
[09:14am] SumitNaiksatam: so before we dive into the specifics of the reviews
[09:15am] SumitNaiksatam: for the benefit of everyone here, mageshgv do you want to quickly summarize what you are doing in these patches
[09:15am] SumitNaiksatam: and how the patches are structured?
[09:15am] mageshgv: ok
[09:15am] mageshgv: So basically these Service Chain patches extends the policy based networking model introduced by Group Policy.
[09:16am] mageshgv: Any deployment would have Advanced Network services such as Firewalls, LBs, VPN etc in addition to basic endpoints
[09:16am] mageshgv: The Service Chain extension of Group Policy allows us to represent these Network Elements and how to act on them based on user defined set of policies
[09:16am] mageshgv: The servicechain patches add three resources to achieve this:
[09:17am] mageshgv: 1) ServiceChain Node - Represents an advanced service definition (eg FW, LB, VPN template)
[09:17am] mageshgv: 2) ServiceChain Spec - Represents a servicechain definition (List of Nodes)
[09:17am] mageshgv: 3) ServiceChain Instance - Represents the instance of a Service Chain
[09:17am] mageshgv: The initially supported Policy model using ALLOW action is now enhanced to support Redirect Action.
[09:17am] mageshgv: The Redirect action takes a Service Chain Spec ID as action value. As part of Redirect action, GBP Driver creates a Service Chain Instance.
[09:17am] mageshgv: A servicechain provider driver then goes and actually creates the Services defined in the Spec.

[09:18am] mageshgv: The patches themselves are broken down into API, DB, Plugin, Reference Driver

[09:18am] mageshgv: And finally Redirect from a Policy Action to a Chain

[09:19am] mageshgv: Any questions on this ?

[09:19am] SumitNaiksatam: mageshgv: nice summary!

[09:19am] ivar-laz_: SumitNaiksatam: +1

[09:19am] SumitNaiksatam: so the GBP Driver here is the resource mapping driver?

[09:20am] mageshgv: SumitNaiksatam: yes

[09:20am] SumitNaiksatam: mageshgv: okay

[09:20am] SumitNaiksatam: lets pause for a minute and let everyone absorb this

[09:21am] SumitNaiksatam: mageshgv: also, the service chain resources are managed by a separate service plugin, right?

[09:22am] mageshgv: yes right. The Service Chain Resources are managed by a separate service plugin

[09:22am] ivar-lazzaro joined the chat room.

[09:23am] mageshgv: And a reference driver implementation has also been added for this service chain plugin

[09:23am] SumitNaiksatam: mageshgv: okay

[09:23am] SumitNaiksatam: mageshgv: so looking at the code in the plugin patch: <https://review.openstack.org/#/c/128555>

[09:24am] SumitNaiksatam: you seem to have a driver manager

[09:25am] SumitNaiksatam: mageshgv: so this framework is similar to what we have for the GBP drivers, right?

[09:25am] mageshgv: SumitNaiksatam: Yes, we are following the same architecture as the other service plugins

[09:25am] SumitNaiksatam: which in turn is borrowed from ML2

[09:25am] mageshgv: right

[09:25am] SumitNaiksatam: ok

[09:25am] rkukura: mageshgv: How does this driver's datapath work? Is it part of a neutron router?

[09:25am] hemanthravi: SumitNaiksatam, driver_manager is a cut-paste of GBP driver manager with minor mods

[09:25am] SumitNaiksatam: hemanthravi: okay

[09:25am] ivar-laz_ left the chat room. (Ping timeout: 244 seconds)

[09:26am] mageshgv: rkukura: Yes, in the reference driver, firewall would be instantiated at the neutron router

[09:27am] mageshgv: And so would be VPN when we support it

[09:27am] SumitNaiksatam: rkukura: the way i understand it, the existing services' reference implementation is leveraged as is

[09:27am] hemanthravi: rkukura, the datapath depends on the service-chain-driver and service. The ref driver implements fw/vpn in the router but LB is realized as a port in the provider-epg

[09:28am] hemanthravi: rkukura, this could be different based on the backend realizing gbp/service-chain

[09:28am] rkukura: mageshgv, SumitNaiksatam: OK. I'm not that familiar with these services' implementations, but have a general idea.

[09:28am] SumitNaiksatam: mageshgv: may be you want to provide a high

level overview of the logic in the "simplechain_driver"?

[09:28am] SumitNaiksatam: i mean how you are spinning this services' up

[09:28am] SumitNaiksatam: and what does chaing mean here

[09:28am] mageshgv: SumitNaiksatam: ok

[09:28am] SumitNaiksatam: *chaining

[09:28am] rkukura: With the reference driver, do the services apply between different internal subnets (EPGs) attached to the same router, or just to external traffic?

[09:29am] SumitNaiksatam: mageshgv: may be you can take rkukura's question first

[09:29am] mageshgv: rkukura: In its current form, the LB is applied at the internal EPGs

[09:30am] mageshgv: The Chain creation itself takes EPG as one of the API input.

[09:31am] mageshgv: To give a overview of the implementation, lets go over the APIs first

[09:31am] SumitNaiksatam: mageshgv: sure

mageshgv: First a user creates a ServiceChainNode which takes a Servicetype - which is at present FW, LB (although there is no backend to leverage/validate it)

[09:33am] SumitNaiksatam: mageshgv: so just to clarify

[09:33am] SumitNaiksatam: mageshgv: by saying - there is no backend to leverage/validate it - you mean that the "flavors" framework is missing?

[09:33am] mageshgv: SumitNaiksatam: Right

[09:34am] SumitNaiksatam: mageshgv: okay, so that is outside the scope of this work, so go on

[09:34am] mageshgv: The Service Chain Node is just a definition of what makes up a Service

[09:34am] mageshgv: It takes a parameter Config - Which is a json representation of a service definition

[09:35am] mageshgv: The Reference driver interprets this as a Heat Template

[09:36am] SumitNaiksatam: mageshgv: so we dont validate the "config" in the resource model

[09:36am] SumitNaiksatam: mageshgv: we just treat is as an opaque, and let the driver do the validation?

[09:36am] mageshgv: SumitNaiksatam: Yes, we do not validate it.

[09:37am] SumitNaiksatam: mageshgv: ok

[09:37am] mageshgv: So any other driver is free to interpret this Config as it sees fit

[09:37am] SumitNaiksatam: mageshgv: so if one were to configure the reference driver

[09:38am] mageshgv: Then we create a Service Chain Spec which is basically a list of service chain nodes.

[09:38am] rkukura: mageshgv: Where could I find a realistic example of a service definition?

[09:38am] ivar-lazzaro left the chat room. (Remote host closed the connection)

[09:38am] SumitNaiksatam: mageshgv: and if the heat template provided as config has some errors, would it be caught at the service node creation stage or at the service chain instantiation stage?

[09:38am] ivar-lazzaro joined the chat room.

[09:38am] hemanthravi: mageshgv, in response to rkukura question - taking fw as an example the fw service will apply to all traffic through the router attached to the provider-epg

[09:39am] mageshgv: rkukura: Do you mean a sample definition of a service

[09:39am] SumitNaiksatam: rkukura: what i used yesterday was configs in the form of heat templates that mageshgv provided

[09:40am] rkukura: Are these in the patches or somewhere else?

[09:40am] SumitNaiksatam: rkukura: example of firewall config here: <http://paste.openstack.org/show/126032>

[09:40am] mageshgv: SumitNaiksatam: At present there is no validation at service chain Node create. The errors are caught at runtime.

[09:40am] SumitNaiksatam: mageshgv: okay got it

[09:41am] rkukura: OK

[09:41am] SumitNaiksatam: rkukura: an example of HAProxy-based LBaaS config here: <http://paste.openstack.org/show/126033/>

[09:42am] rkukura: SumitNaiksatam: thanks

[09:42am] SumitNaiksatam: mageshgv: perhaps in the future, we can leverage the configured driver to validate the config?

[09:42am] SumitNaiksatam: mageshgv: so that we dont have to wait until run time?

[09:43am] hemanthravi: mageshgv, will the fw service in the current form would apply to all eggs sharing an L3 policy?

[09:43am] rkukura: Is the plan to use these heat templates as the API long term, or would direct API calls be possible too?

[09:43am] mageshgv: SumitNaiksatam: yes, we can add validations because a provider would know what is supported by it

[09:43am] hemanthravi: mageshgv, since there is a single router for the L3 policy

[09:44am] SumitNaiksatam: rkukura: i believe that the current thinking is that the config needs to be more flexible so that different service chain providers can be supported

[09:44am] mageshgv: hemanthravi: Since we are leveraging the existing firewall implementation, at present it would apply to all the EPGs on the tenant

[09:44am] SumitNaiksatam: rkukura: of course this introduces the issue of validation

[09:44am] SumitNaiksatam: rkukura: hence i was asking if the provider/driver can be leveraged to do the validation

[09:45am] SumitNaiksatam: rkukura: that way, i guess we can support both objectives

[09:45am] ivar-lazzaro: SumitNaiksatam: not sure you can do a full validation at Node level, the resource pool may be drained at instantiation time

[09:45am] SumitNaiksatam: ivar-lazzaro: validity of the integrity of the config

[09:46am] SumitNaiksatam: ivar-lazzaro: beyond just the fact that this is proper json

[09:46am] SumitNaiksatam: ivar-lazzaro: schema check basically

[09:46am] rkukura: Sure, the validation can be improved. I am not yet 100% convinced this style of API makes long term sense yet.

[09:46am] ivar-lazzaro: SumitNaiksatam: does HEAT provide this capability?

[09:46am] rkukura: It seems to me that this approach is not as "intent based" as we might strive for.

[09:47am] mageshgv: ivar-lazzaro: HEAT allows us to validate if a template is valid or not, but not the resource issue

[09:47am] ivar-lazzaro: rkukura: I actually have the same impression

[09:47am] rkukura: Its fine for now.

[09:47am] SumitNaiksatam: rkukura: we want to be able treat the service config as an external thing

[09:47am] rkukura: I'd prefer that the existing ALLOW rules are what driver the FWaaS behind the scenes.

[09:48am] ivar-lazzaro: mageshgv: the resource issue is ok to be validate at runtime imho. you can just go in error state of tell the user to free some resources at instance creation time

[09:48am] ivar-lazzaro: rkukura: the problem with that approach is that describing a LB or a IDS gets tricky

[09:48am] SumitNaiksatam: rkukura: so as far as GBP is concerned, it only knows that a service chain exists, and that it needs to allocate resources (like ip addresses for what is requested)

[09:48am] ivar-lazzaro: rkukura: how do you define the order across multiple contracts?

[09:49am] ivar-lazzaro: rkukura: also, you could have a "BALANCE" rule, but then you hide all the providers' value by normalizing the functionalities

[09:49am] SumitNaiksatam: ivar-lazzaro: my suggestion was only for template validation

[09:49am] rkukura: I think the GBP model would need to capture the notion of an EPG containing a set of replicated application VMs, and take care of configuring load balancing across these.

[09:50am] hemanthravi: ivar-lazzaro, yes that's correct fw (allow/deny) overlaps with ALLOW action but won't apply to any extended firewall func

[09:50am] SumitNaiksatam: rkukura: that goes more into the realm of the "balance" action that ivar-lazzaro is referring to here

[09:50am] ivar-lazzaro: hemanthravi: I don't think we want deny actions

[09:50am] SumitNaiksatam: rkukura: however, service chains are more than just LBs and FWs

[09:50am] rkukura: I'm fine with the current approach for the short term, but would like to see if we could improve the GBP model to capture the intent at a higher level

[09:51am] SumitNaiksatam: rkukura: fair enough

[09:51am] ivar-lazzaro: rkukura: +1

[09:51am] hemanthravi: rkukura, the idea of service intent is to

configure the action_value with a service-chain-spec which describes the services and the services are rendered by the service-chain-driver

[09:51am] SumitNaiksatam: ivar-lazzaro: i dont think hemanthravi is suggesting deny actions

[09:51am] SumitNaiksatam: hemanthravi: yes, well articulate

[09:51am] mageshgv: rkukura: User would declare his intent of applying a group of Network Services (chain). And how it is achieved is abstracted by the Service chain drive

[09:51am] mageshgv: r

[09:51am] SumitNaiksatam: *articulated

[09:51am] hemanthravi: ivar-lazzaro, didn't mean to add deny just referring to that cap of a net fw

[09:51am] s3wong joined the chat room.

[09:52am] SumitNaiksatam: okay folks, lets converge back to the reviews

[09:52am] rkukura: SumitNaiksatam: +1

[09:52am] SumitNaiksatam: mageshgv: thanks for the overview of the "current" driver

[09:52am] ivar-lazzaro: mageshgv: what rkukura is saying is that we are nicely hiding L2/L3 but deploying network services is not that different from the past... You still need to specify all the configuration and such

[09:52am] rkukura: This discuss was really helpful for me, though. I would not have made much sense of the code without it.

[09:52am] mageshgv: ivar-lazzaro: got it

[09:53am] ivar-lazzaro: mageshgv: for instance, we don't automatically understand that a member pool for LBaaS instance is the providing EPG

[09:53am] rkukura: s/discuss/discussion/

[09:53am] SumitNaiksatam: rkukura: good

[09:53am] SumitNaiksatam: mageshgv: okay so moving on

[09:53am] SumitNaiksatam: so your next patch is redirect implementation

[09:53am] s3wong: Oops... I miss the beginning of the review session...

[09:53am] s3wong: *missed

[09:53am] SumitNaiksatam: s3wong: no worries

[09:54am] SumitNaiksatam: i am capturing the logs

[09:54am] hemanthravi: ivar-lazzaro, for LB the providing-epg is the member pool

[09:54am] mageshgv: ivar-lazzaro: Actually with the current implementation it does understand it

[09:54am] ivar-lazzaro: hemanthravi: aren't the members specified in the heat template?

[09:55am] mageshgv: ivar-lazzaro: Yes, but the member information itself is fetched from the Providing EPG

[09:55am] SumitNaiksatam: in terms of logistics - some of us having a team meeting starting in about 5 mins

[09:55am] ivar-lazzaro: mageshgv: I probably missed that

[09:55am] SumitNaiksatam: however i will continue to be here

[09:55am] mageshgv: So only the member definition goes into the template

[09:55am] hemanthravi: ivar-lazzaro, the idea is for the heat-template to have a parameter that can be configured for members...and the driver passes this info based on the epg

[09:56am] SumitNaiksatam: rkukura: ivar-lazzaro: i think you can also stay here if you feel comfortable

[09:56am] ivar-lazzaro: mageshgv: what do you mean? that on HEAT you create 10 members and then the driver fills the members?

[09:56am] ivar-lazzaro: mageshgv: with their information (like IP address)

[09:56am] rkukura: SumitNaiksatam: I'll probably listen in on the team meeting but continue to follow this as well.

[09:56am] SumitNaiksatam: rkukura: same here

[09:57am] ivar-lazzaro: SumitNaiksatam: I have another hard stop unfortunately

[09:57am] SumitNaiksatam: ivar-lazzaro: no worries

[09:57am] hemanthravi: ivar-lazzaro, trying to use a resource-group in heat to start variable number of members but running into issue with this

[09:57am] mageshgv: ivar-lazzaro: We were exploring the possibility of leveraging some inbuilt functionality in HEAT to specify the number of members and List of IPs dynamically

[09:58am] mageshgv: But that didnt work due to some limitations unfortunately. So the idea is that the service chain provider would be free to interpret and modify the json/template on the go

[09:58am] SumitNaiksatam: ivar-lazzaro: so the point is that the underlying service-chain provider has to support it

[09:58am] ivar-lazzaro: mageshgv: I see, that'd be helpful to see an example somewhere on how this works. But I get the point

[09:58am] SumitNaiksatam: ivar-lazzaro: this is not a limitation of the model

[09:59am] SumitNaiksatam: ivar-lazzaro: because the mode incorporatest the notion of a "network service policy"

[09:59am] SumitNaiksatam: ivar-lazzaro: and that drives ip address resource allocation for these dynamic resources

[09:59am] ivar-lazzaro: SumitNaiksatam: got it

[10:00am] SumitNaiksatam: mageshgv: so in the redirect patch you are extending the DB schema

[10:01am] mageshgv: SumitNaiksatam: Yes, so in Redirect implementation, RM driver would create a Service Chain Instance from the Policy Action Value(Service Chain SPec ID)

[10:01am] ivar-laz_ joined the chat room.

[10:01am] mageshgv: To store this, I have extended the DB scema

[10:02am] SumitNaiksatam: mageshgv: okay

[10:02am] s3wong: mageshgv: I have a question with that --- how do you perform 'redirect' only on the traffic matching a classifier in RM?

[10:03am] mageshgv: s3wong: The idea is that Redirect implicitly means Allow + Redirect

[10:04am] mageshgv: So traffic itself is filtered by security groups

and we apply Redirect on top of it

[10:04am] hemanthravi: s3wong, in the reference impl any traffic between the eggs (consumer, provider) will go through the chain

[10:05am] ivar-lazzaro left the chat room. (Ping timeout: 255 seconds)

[10:05am] rkukura: mageshgv: I think s3wong is asking if certain traffic between two EPGs can be redirected and other traffic allowed but not redirected

[10:05am] s3wong: mageshgv: yes --- but if I have two SGs (for different contracts) on the same EP (port), and both have 'redirect' to different chain, the SG is per port --- so therefore traffic would be allowed and go through the chains?

[10:05am] s3wong: rkukura: putting the question better than I could, thanks!

[10:06am] mageshgv: rkukura: the current reference driver doesnt have this capability

[10:06am] SumitNaiksatam: s3wong rkukura: this reference implementation does not leverage any traffic steering

[10:07am] SumitNaiksatam: because those constructs dont exist in neutron today

[10:07am] rkukura: mageshgv, SumitNaiksatam: that's what I thought

[10:07am] SumitNaiksatam: that said, for the combination of services' we are using

[10:07am] s3wong: mageshgv, hemanthravi: in essence, on the same egress port, we can only have ONE 'redirect' -- that could be documented as a limitation on the resource mapping driver, I am fine with that... just make sure it is documented

[10:08am] SumitNaiksatam: and their specific reference implementation, we dont need the traffic steering

[10:08am] hemanthravi: s3wong, ok

[10:08am] s3wong: SumitNaiksatam: agreed --- the vendor / 3rd party drivers can --- I am sure -- do the right thing

[10:09am] SumitNaiksatam: s3wong: true, so the claim is that the model allows that, but if it does not we need to identify that

[10:10am] s3wong: SumitNaiksatam: sure, as long as the APIs and model are sound. I am good with that

[10:10am] SumitNaiksatam: s3wong: can you add your comment to the review?

[10:10am] s3wong: SumitNaiksatam: the 'redirect' reference implementation? sure

[10:10am] SumitNaiksatam: mageshgv: regarding the reference driver, you have a DB schema change as well

[10:10am] SumitNaiksatam: s3wong: yeah

[10:11am] SumitNaiksatam: mageshgv: that is in the context of the heat "stacks"?

[10:12am] mageshgv: SumitNaiksatam: Right. The reference driver leverges heat to instantiate the services. So the db schema to manage the stack instances is added

[10:12am] SumitNaiksatam: mageshgv: so the lifecycle of the heat deployment is handled with this stack id reference?

[10:14am] mageshgv: SumitNaiksatam: yes, When a service chain is

instantiated (Service Chain Instance create), we end up creating one stack per service chain node. These stacks lifecycle is managed with this reference

[10:15am] SumitNaiksatham: mageshgv: ah okay, so how are the stack ids correlated with the service chain?

[10:17am] SumitNaiksatham: mageshgv: so may be that is happening in the redirect implementation patch

[10:17am] mageshgv: SumitNaiksatham: The stack IDs are stored against Service Chain Instance IDs. Each service chain instance may have more than one stack

[10:18am] SumitNaiksatham: mageshgv: ok, i was trying to understand which patch has what

[10:18am] mageshgv: SumitNaiksatham: In the Redirect implementation, the Resource mapping driver creates a service chain instance from the Service Chain Spec ID in action value

[10:18am] SumitNaiksatham: mageshgv: ah thanks, just saw that

[10:19am] mageshgv: So it only holds a reference to the Chain Instance in GBP DB

[10:19am] SumitNaiksatham: mageshgv: the next patch is: "Group Policy Service Chain Implementation change for updated Spec": <https://review.openstack.org/#/c/130004>

[10:19am] SumitNaiksatham: mageshgv: so what is the "updated" part here?

[10:20am] mageshgv: The major change is that initial service chain spec did not know about other GBP Resources

[10:21am] mageshgv: The new spec allows us to pass the provider and consumer EPGs and classifier to the Service Chain Instantiation

[10:22am] SumitNaiksatham: mageshgv: okay

[10:22am] mageshgv: This way the service chain driver has the complete understanding of the insertion context

[10:22am] SumitNaiksatham: mageshgv: i think ivar-laz_ has a question on this, perhaps you can provide a summary response

[10:22am] mageshgv: SumitNaiksatham: You mean in the patch?

[10:23am] SumitNaiksatham: mageshgv: yes, in patchset 3

[10:23am] mageshgv: SumitNaiksatham: ok

[10:24am] SumitNaiksatham: mageshgv: hemanthravi: can you also quickly summarize here, why we pass the provider/consumer EPGs to service chain?

[10:24am] SumitNaiksatham: mageshgv: and also the classifier

[10:25am] SumitNaiksatham: i believe we need this so that the service chain can be anchored correctly on the relevant subnets

[10:26am] mageshgv: SumitNaiksatham: That would be usecase for the reference driver.

[10:26am] SumitNaiksatham: the above explanantion in the context of the provider/consumer EPGs

[10:26am] mageshgv: Without the EPG/Classifier context, the model does not allow for traffic steering to be performed by any other driver implementation

[10:27am] hemanthravi: yes, this will allow any providers with traffic steering to address the issue raised by rkukura, s3wong

[10:27am] hemanthravi: to only forward traffic between the eggs matching the classifier to the service-chain

[10:30am] mageshgv: SumitNaiksatam: Hope this clarifies the usecase for these parameters

[10:30am] SumitNaiksatam: mageshgv: hemanthravi: yeah, thanks

[10:31am] SumitNaiksatam: next patch is: Add Network Service Policy support for service chains <https://review.openstack.org/#/c/129545/>

[10:31am] SumitNaiksatam: that is my patch

[10:31am] SumitNaiksatam: i posted another rev to the spec to explain this better

[10:33am] SumitNaiksatam: but i am not sure if it explains well

[10:33am] SumitNaiksatam: i am happy to clarify here

[10:35am] mageshgv: SumitNaiksatam: Can you briefly explain about the PolicyLabels/tags

[10:35am] SumitNaiksatam: mageshgv: we are not implementing that

[10:36am] SumitNaiksatam: mageshgv: but other than that, the PolicyTags are just literals

[10:36am] mageshgv: SumitNaiksatam: ok. The spec the patch refers to was having that. thought we are implementing it

[10:36am] SumitNaiksatam: so clarification

[10:36am] ivar-laz_: mageshgv SumitNaiksatam: just dropping in a quick question about the previous patch... Why do we need providing and consuming EPGs for the ServiceChainInstance? Can't this be retrieved from the contracts implementing the rules?

[10:36am] SumitNaiksatam: i meant not implemented before Paris

[10:37am] SumitNaiksatam: but yes during and after Paris

[10:37am] SumitNaiksatam: ivar-laz_: the service chain does not have reference to the contract AFAIK

[10:38am] ivar-laz_: The contract knows the spec though

[10:38am] ivar-laz_: and the instance knows the spec as well

[10:38am] ivar-laz_: shouldn't a join be enough?

[10:38am] SumitNaiksatam: ivar-laz_: spec - you mean service chain spec?

[10:39am] ivar-laz_: yes

[10:40am] SumitNaiksatam: ivar-laz_: the contract can be associated with multiple provider and consumer EPG/PTGs

[10:40am] SumitNaiksatam: ivar-laz_: how would you know which one?

[10:41am] ivar-laz_: SumitNaiksatam: let's say that a contract has 1 redirect rule only

[10:42am] ivar-laz_: SumitNaiksatam: that redirect rule refers to a ServiceInstance ultimately

[10:42am] ivar-laz_: SumitNaiksatam: and the same ServiceChainInstance can't be referenced by multiple contracts

[10:43am] ivar-laz_: SumitNaiksatam: so all the provider/consumer EPGs of that contract are the provider/consumer EPGs of the ServiceChainInstance

[10:43am] ivar-laz_: SumitNaiksatam: unless I'm missing the actual use case

[10:43am] hemanthravi: ivar-laz_, the same spec can be used in different contracts

[10:43am] hemanthravi: mageshgv, isn't that correct?
[10:44am] mageshgv: hemanthravi: yes. you are right.
[10:44am] ivar-laz_: hemanthravi: yeah but the "stack" table on the driver gives the contract -> instance one to one relationship
[10:44am] SumitNaiksatam: ivar-laz_: also, one service chain instance is created per EPG-pair, at least thats what i thought it was
[10:45am] mageshgv: ivar-laz_: The mapping is for service chain instance -> Heat stack
[10:45am] SumitNaiksatam: i dont think we are sharing service chain instances between EPG-pairs
[10:45am] ivar-laz_: SumitNaiksatam: I see, but why doing this instead of simply using the contract? I'm probably just missing the use case
[10:46am] ivar-laz_: If a contract pointed to a ServiceChainInstance, would that be different?
[10:46am] hemanthravi: SumitNaiksatam, you are right chain instances are per epg-pair + rule
[10:46am] ivar-laz_: I'm afraid of putting GBP knowledge in the ServiceChain model
[10:46am] SumitNaiksatam: ivar-laz_: same reason, we need to know exactly which EPG-pair we need to instantiate this service chain instance for
[10:47am] SumitNaiksatam: ivar-laz_: either GBP has to know about the service chain or the service chain needs to know about some part of the GBP model
[10:47am] SumitNaiksatam: ivar-laz_: we have tried as much separation as possible
[10:47am] ivar-laz_: SumitNaiksatam: GBP already know about service chain... The driver at least
[10:48am] SumitNaiksatam: ivar-laz_: not really, it only treats is as UUID
[10:48am] ivar-laz_: SumitNaiksatam: the API is agnostic (eg. action_value doesn't know about it)
[10:48am] SumitNaiksatam: one thing we did realize while we were implementing this though, is that to realize the redirect to its full potential and in a flexilbe way ...
[10:48am] ivar-laz_: SumitNaiksatam: however I see the point, basically we are missing away to actually have 1 chain instance for all the providing/consuming EPGs of a given contract
[10:49am] SumitNaiksatam: ...the GBP driver and the service chain drivers probably need to be more tightly integrated
[10:49am] ivar-laz_: So let me see if I understand
[10:49am] SumitNaiksatam: the current separation of the driver implementation is pretty artificial
[10:50am] ivar-laz_: Given NxM EPGs related by a single contract
[10:50am] ivar-laz_: with redirect on a SC
[10:50am] SumitNaiksatam: ivar-laz_: logically its always a new service chain instance per EPG-pair
[10:50am] ivar-laz_: we create a SCI for each pair of EPGs
[10:50am] ivar-laz_: is that correct?
[10:50am] SumitNaiksatam: ivar-laz_: your implementation may or may

not want to map it to the same set of services

[10:51am] SumitNaiksatam: ivar-laz_: correct

[10:51am] ivar-laz_: SumitNaiksatam: ok, but isn't this a driver implementation detail more than an API need?

[10:51am] ivar-laz_: SumitNaiksatam: let's say we have traffic steering in place at some point

[10:52am] ivar-laz_: SumitNaiksatam: in that case 1 instance per contract should be enough... right?

[10:53am] SumitNaiksatam: ivar-laz_: i am not sure how one logical SCI would be enough

[10:53am] SumitNaiksatam: ivar-laz_: you need to be able to provide traffic isolation, right?

[10:54am] SumitNaiksatam: ivar-laz_: that is being captured the logical notion of a separate SCI

[10:54am] SumitNaiksatam: *captured by

[10:55am] SumitNaiksatam: mageshgv: hemanthravi: feel free to correct me if i am misrepresenting the rationale

[10:55am] ivar-laz_: SumitNaiksatam: well I would say that 1 chain per providing EPG would be enough in this case

[10:55am] ivar-laz_: SumitNaiksatam: as opposed to n^2 instances

[10:56am] mageshgv: SumitNaiksatam: you are correct

[10:56am] ivar-laz_: SumitNaiksatam: consumers will all go through the same load balancer in order to reach the given EPG

[10:56am] mageshgv: ivar-laz_: To answer your question, it is sufficient for reference implementation of what we are doing today

[10:56am] SumitNaiksatam: ivar-laz_: but the LB will have different contexts

[10:57am] hemanthravi: ivar-laz_, i think the mapping to the same instance of LB is upto the service-provider, multiple SCI can map to the same service instance

[10:57am] SumitNaiksatam: ivar-laz_: also, one is free to map to a single instance

[10:57am] hemanthravi: mageshgv, is that correct

[10:57am] ivar-laz_: SumitNaiksatam: I feel also that it depends on the driver implementation. Maybe 1 LB slice can host up to 10 E providing/consuming EPGs and you want to schedule instances somehow

[10:57am] SumitNaiksatam: ivar-laz_: yeah hemanthravi's point

[10:57am] SumitNaiksatam: ivar-laz_: we should not mandate that it always has to be one SCI, some providers may not be able to support it

[10:58am] SumitNaiksatam: btw, just to be clear, this is "policy driven service redirect/chain"

[10:58am] ivar-laz_: SumitNaiksatam: ok so to be completely generic we may need a one to many relationship between SCI and EPGs

[10:59am] mageshgv: ivar-laz_: But limiting to one Contract per Instance is limiting for other use cases, especially traffic steering. And providing EPGs explicitly to instantiate a chain handles this well for all use cases

[10:59am] SumitNaiksatam: we are not trying to solve the generic service chain problem that is applicable everywhere

[10:59am] ivar-laz_: But I know get the use case much better

[10:59am] SumitNaiksatam: we are doing this very much in the context of GBP and service redirect

[10:59am] ivar-laz_: s/know/now

[11:00am] SumitNaiksatam: so the service chain API knowing about GBP constructs is natural here

[11:00am] SumitNaiksatam: i am not trying to say that we make bad choices

[11:00am] SumitNaiksatam: i am just saying that there is a separation of concerns between GBP and services

[11:01am] SumitNaiksatam: however that does not necessarily manifest at the GBP-Service-Chain API level

[11:01am] ivar-laz_: SumitNaiksatam: I see, as long as it is an implicit process I'm fine with it (eg. provider/consumer EPG added by the driver at SCI creation time based on the redirect rule)

[11:01am] ivar-laz_: SumitNaiksatam: thanks for the explanation much clearer now

[11:02am] SumitNaiksatam: the separation of concerns manifests at the level where the individual service is defined (outside of GBP) and where the service chain provider implementation interfaces with each of the services (southbound API from the provider)

[11:02am] SumitNaiksatam: ivar-laz_: sure, but good discussion, and perhaps a topic for the summit discussion as well

[11:03am] SumitNaiksatam: so before that detour we were discussing: Add Network Service Policy support for service chains <https://review.openstack.org/#/c/129545/>

[11:03am] ivar-laz_ left the chat room. (Remote host closed the connection)

[11:04am] hemanthravi: ivar-laz_, the eg, you are referring to what epg get passed to SCI creation?

[11:04am] ivar-lazzaro joined the chat room.

[11:07am] SumitNaiksatam: hemanthravi: ivar-lazzaro might have missed that question - but perhaps we can pursue it offline

[11:07am] SumitNaiksatam: i will move on to the final patch

[11:07am] hemanthravi: will do, i need to break now

[11:07am] ivar-lazzaro: SumitNaiksatam hemanthravi: yeah I d/cd sorry

[11:07am] SumitNaiksatam: yeah we will wrap up shortly

[11:07am] SumitNaiksatam: *shortly

[11:08am] SumitNaiksatam: its pretty late for mageshgv as well

[11:08am] SumitNaiksatam: so Network Service Policy Driver Implementation <https://review.openstack.org/#/c/130920/>

[11:08am] SumitNaiksatam: mageshgv: is this still WIP?

[11:08am] SumitNaiksatam: if not can you update the commite title?

[11:09am] mageshgv: yes, Unit tests are not added yet

[11:09am] SumitNaiksatam: mageshgv: ah ok

[11:09am] SumitNaiksatam: mageshgv: and what exactly is happening in this patch, as in which resources are being handled?

[11:09am] mageshgv: Otherwise it is fine. Wanted to test resource reclaim on Policy delete properly by unit tests

[11:10am] SumitNaiksatam: mageshgv: sure

[11:10am] mageshgv: As per the model Network Service Policy can be

attached to multiple EPGs

[11:12am] mageshgv: And at present we support allocating/reserving a single IP Address from the EPG subnet

[11:12am] SumitNaiksatam: so you are essentially reserving certain IPs

[11:12am] SumitNaiksatam: mageshgv: okay

[11:12am] SumitNaiksatam: mageshgv: you dont have to handle the "vip" in this patch?

[11:13am] mageshgv: SumitNaiksatam: Actually we do, but it is handled in a generic way

[11:13am] SumitNaiksatam: mageshgv: okay, where?

[11:13am] mageshgv: The NSP parameter type would be ip-single for VIP.

[11:13am] SumitNaiksatam: mageshgv: correct

[11:14am] mageshgv: So any parameter that has a type of ip_single would need a single address to be reserved.

[11:14am] mageshgv: So the address reservation is generic.

[11:15am] mageshgv: The parameter itself would be differentiated later on while instantiating the service chain

[11:15am] hemanthravi: mageshgv, the addr resv happens when nsp is attached to an epg

[11:15am] mageshgv: vip would be represented by the NSP parameter name which can be any string

[11:15am] hemanthravi: for any param with type ip-single?

[11:16am] mageshgv: hemanthravi: yes, you are correct

[11:16am] SumitNaiksatam: mageshgv: so what if multiple IPs are reserved, how do you know which one to use for the VIP?

[11:17am] hemanthravi: that'll be by the name request from service-chain-spec using parms, correct?

[11:17am] mageshgv: SumitNaiksatam: The reserved ip address is stored against the EPG and NSP id in the DB

[11:18am] SumitNaiksatam: mageshgv: ah ok

[11:18am] SumitNaiksatam: mageshgv: what if the NSP has multiple "ip_single" types?

[11:18am] mageshgv: hemanthravi: yes, it should be provided by fetching the param names from the service chain spec. Although today it is blindly passed

[11:20am] mageshgv: SumitNaiksatam: Such a scenario is not handled today. The reference driver only expects one ip_single. It ignores the repetition

[11:20am] SumitNaiksatam: mageshgv: ah okay

[11:20am] SumitNaiksatam: mageshgv: perhaps good to document this in the driver (wherever relevant)

[11:21am] mageshgv: SumitNaiksatam: ok

[11:22am] SumitNaiksatam: mageshgv: not so much because i see that as a limitation, simply to just understand the logic

[11:22am] mageshgv: SumitNaiksatam: got it

[11:22am] SumitNaiksatam: ok, rkukura hemanthravi s3wong ivar-lazzaro: any more questions on these patches?

[11:22am] SumitNaiksatam: i think we are pretty saturated now!

[11:23am] ivar-lazzaro: +1

[11:23am] rkukura: SumitNaiksatam: A lot to absorb!

[11:23am] SumitNaiksatam: yes true
[11:23am] s3wong: SumitNaiksatam: will re-read the meeting log... missed a lot of stuff as I went in and out of my chair
[11:23am] SumitNaiksatam: s3wong: sure no worries
[11:23am] SumitNaiksatam: i am copy pasting and will post the log once we wrap up
[11:23am] s3wong: but I will take a look at the patches... sounds interesting
[11:23am] SumitNaiksatam: since we dont have a meeting bot
[11:23am] rkukura: SumitNaiksatam: Are you able to put this log on the wiki?
[11:23am] SumitNaiksatam: rkukura: yes absolutely
[11:24am] rkukura: thanks
[11:24am] SumitNaiksatam: i will post this right away
[11:24am] SumitNaiksatam: mageshgv: thanks very much!
[11:24am] mageshgv: SumitNaiksatam: welcome
[11:24am] SumitNaiksatam: mageshgv: can i also request you to be on #openstack-gbp whenever you are working?
[11:25am] SumitNaiksatam: so that you can quickly answer questions from the team here
[11:25am] mageshgv: SumitNaiksatam: Sure, although we do have a timezone issue
[11:25am] SumitNaiksatam: mageshgv: i realize, just whenever you are available/working
[11:25am] SumitNaiksatam: more so for this week
[11:25am] mageshgv: ok
[11:26am] SumitNaiksatam: ok thanks all, if nothing else, lets wrap this for now
[11:26am] SumitNaiksatam: and we can continue to lurk here
[11:26am] SumitNaiksatam: magesh: feel free to take off, good night!
[11:27am] hemanthravi: bye
[11:27am] SumitNaiksatam: rkukura: ivar-lazzaro hemanthravi s3wong: thanks for joining as well and providing feedback!
[11:27am] s3wong: SumitNaiksatam, hemanthravi, mageshgv: thanks for the explanations
[11:27am] rkukura: bye
[11:27am] mageshgv: bye
[11:27am] ivar-lazzaro: bye