banix left the chat room. (Quit: banix)
[11:31am] SumitNaiksatam: s3wong: ivar-lazzaro rkukura kevinbenton igordcard: hi
[11:31am] SumitNaiksatam: i was expecting hemanth and subra to be here as well
[11:31am] rkukura: hi
[11:31am] ivar-lazzaro: hi
[11:31am] SumitNaiksatam: let me ping them
[11:31am] kevinbenton: hi
[11:32am] kevinbenton: what was i supposed to review
[11:32am] kevinbenton: i will review it now
[11:32am] SumitNaiksatam: kevinbenton: perfect
[11:32am] hemanthravi joined the chat room.
[11:33am] SumitNaiksatam: kevinbenton: this meeting is mostly to review the following specs and their impl:
[11:33am] SumitNaiksatam: 1. Shared resources: #link https://review.openstack.org/#/c/133603/
[11:33am] SumitNaiksatam: #link https://review.openstack.org/134692
[11:34am] SumitNaiksatam: 2. External connectivity: #link https://review.openstack.org/#/c/136550/
[11:34am] SumitNaiksatam: #link https://review.openstack.org/137267
[11:34am] SumitNaiksatam: if we have time we can review other specs/pending patches as well
[11:35am] SumitNaiksatam: hemanthravi: is here good
[11:35am] SumitNaiksatam: lets get started
[11:35am] hemanthravi: hi
[11:35am] KrishnaK joined the chat room.
[11:35am] SumitNaiksatam: kevinbenton: by the way your homework was the horizon patches!
[11:35am] SumitNaiksatam: KrishnaK: hi, thanks for joining
[11:35am] SumitNaiksatam: just getting started
[11:35am] s3wong: SumitNaiksatam, ivar-lazzaro: as you guys requested, I have the untested patch for update bug: #link https://review.openstack.org/#/c/137439/
[11:35am] SumitNaiksatam: s3wong: sweet
[11:36am] SumitNaiksatam: so s3wong's patch is landing as a bug fix
[11:36am] kevinbenton: SumitNaiksatam: ok
[11:36am] SumitNaiksatam: kevinbenton: just kidding, i know you are swamped!
[11:36am] banix joined the chat room.
[11:36am] s3wong: thanks to ivar-lazzaro. it seems to be much simpler than expected --- as I mostly reused ivar-lazzaro's policy-rule-delete function to perform a delete and add
[11:36am] SumitNaiksatam: kevinbenton: that said, it would be great to get your input on these specs/impl as well
[11:36am] SumitNaiksatam: banix: hi, thanks for joining
[11:37am] banix: SumitNaiksatam: hi
[11:37am] ivar-lazzaro: s3wong: nice!
[11:37am] SumitNaiksatam: s3wong: sweet, yes ivar-lazzaro is definitely the man today!

[11:37am] SumitNaiksatam: ivar-lazzaro: so over to you, can you provide a quick summary to the team on the share attributes proposal?
[11:38am] s3wong: once that works (after unit testing), I am thinking of doing classifier update by getting classifier->policy-rule mapping, then the rest is the same as policy-rule update
[11:38am] ivar-lazzaro: SumitNaiksatam: sure
[11:38am] s3wong: SumitNaiksatam, ivar-lazzaro: but as we talked about yesterday, the action update (which only affects 'redirect'), I would need someone else to pick up, as I am not too familiar with the 'redirect' code in rmd
[11:39am] ivar-lazzaro: So basically the idea is simply to allow some GBP resources to be shared across tenants
[11:39am] ivar-lazzaro: By shared, I mean the Neutron's concept of sharing: all or owner
[11:40am] ivar-lazzaro: Sharing attributes though leads to some restrictions (mostly caused by how Neutron is retrieving resources internally)
[11:40am] ivar-lazzaro: To be specific, whenever you create a Shared resource, it cannot point to non-shared resources
[11:40am] rkukura: ivar-lazzaro: Do non-owners of a shared resource have full read/write access, or just the ability to see and reference the resource?
[11:40am] ivar-lazzaro: For example, a shared EPG must exist on a shared L2P
[11:41am] ivar-lazzaro: rkukura: that depends from the policy.json
[11:41am] ivar-lazzaro: rkukura: by default, non owners cannot write a shared resource
[11:41am] rkukura: ivar-lazzaro: OK, thanks
[11:41am] ivar-lazzaro: rkukura: however, they can always "use" it
[11:41am] KrishnaK: ivar-lazzaro: Can you pls one concrete example ?
[11:41am] SumitNaiksatam: s3wong: i will check with magesh on the action redirect, thanks for the reminder
[11:42am] ivar-lazzaro: Ok let's say you want to implement a management PTG in your private cloud
[11:42am] ivar-lazzaro: You can create a shared PTG, which will exist on a shared L2P and by reflection on a shared L3P
[11:42am] ivar-lazzaro: All the tenants, when creating their VMs, can then place one of the VM interfaces on the shared PTG
[11:43am] SumitNaiksatam: ivar-lazzaro: by reflection, you mean transitivity?
[11:43am] ivar-lazzaro: Therefore being part of the shared management PTG
[11:43am] ivar-lazzaro: KrishnaK: doe this answer your question?
[11:43am] hemanthravi: ivar-lazzaro, when a PT is created by a tenant on a shared PTG is there any modification to the shared PTG
[11:43am] ivar-lazzaro: SumitNaiksatam: yes, I think I "Italianized" the word
[11:43am] SumitNaiksatam: ivar-lazzaro: np, just so that others are not confused
SumitNaiksatam: ivar-lazzaro: np, just so that others are not confused

[11:44am] SumitNaiksatam: sorry for the distraction!
[11:44am] ivar-lazzaro: hemanthravi: no. Of course you get a back reference to the PT but it doesn't really modify the PTG
[11:45am] ivar-lazzaro: hemanthravi: although today our API doesn't show the PT list starting from a PTG anyway
[11:46am] SumitNaiksatam: okay, so just to level set, this particular scenarion which ivar-lazzaro described is identical to the external-shared network in Neutron
[11:46am] SumitNaiksatam: ivar-lazzaro: accurate to say that?
[11:46am] ivar-lazzaro: SumitNaiksatam: yes
[11:46am] SumitNaiksatam: as in that case, the network id of the port is set
[11:46am] SumitNaiksatam: so in this case the PTG id is set on the port
[11:46am] SumitNaiksatam: sorry on the PT
[11:47am] SumitNaiksatam: the policy mechanism allows this association
[11:47am] KrishnaK: ivar-lazzaro: Iam not clear on placing the tenants VM's vnic interface on the shared management PTG ? Won't they be able to access other's tenants VM's
[11:47am] SumitNaiksatam: although the PTG does not belong to the tenant which is creating the PT
[11:48am] igordcard: SumitNaiksatam, hello, sorry couldn't reach the computer at time
[11:48am] SumitNaiksatam: igordcard: no worries, thanks for joining
[11:48am] ivar-lazzaro: KrishnaK: ofc, but that's a management PTG. Think of an enterprise or private cloud... you just have a common PTG across your teams.
[11:49am] rkukura: ivar-lazzaro: Am I correct that with RMD, a shared L2P’s network is shared, and a shared PTG’s subnets are shared?
[11:49am] KrishnaK: ivar-lazzaro: I see. thanks.
[11:49am] ivar-lazzaro: rkukura: yes that's correct.
[11:49am] SumitNaiksatam: KrishnaK: ivar-lazzaro: in that case the contracts(policy rule sets) are created in such a way that only management traffic is allowed on the management network
[11:50am] ivar-lazzaro: rkukura: if they weren't, no other tenant could access them. _get_by_id in Neutron is checking for the owner of the resource or the "shared" flag
[11:50am] rkukura: ivar-lazzaro: that’s what I thought
[11:50am] ivar-lazzaro: rkukura: thus forcing the workflow in this direction
[11:51am] KrishnaK: SumitNaiksatam: Thx.
[11:51am] rkukura: ivar-lazzaro: but a shared L3P’s router is not shared, right?
[11:51am] ivar-lazzaro: rkukura: No, That's why with the RMD you can only share PTGs
[11:52am] ivar-lazzaro: rkukura: and that's why even though a L2P is shared, no tenant other than the owner can actually use it
[11:52am] ivar-lazzaro: rkukura: we should discuss the possibility to share Routers in Neutron in Kilo
[11:52am] SumitNaiksatam: ivar-lazzaro: so taking a step back, is it

accurate to say that the proposal you have today is purely an artifact of the policy mechanism in Neutron as of Juno, but we expect this to change?
[11:52am] ivar-lazzaro: rkukura: this way we can unlock more interesting use cases
[11:52am] rkukura: Can’t the owner of a shared L3P create an unshared L2P for a specific tenant?
[11:53am] ivar-lazzaro: SumitNaiksatam: yes. There are certainly better way to handle shared resources imho
[11:53am] ivar-lazzaro: rkukura: not in the Neutron mapping driver
[11:53am] ivar-lazzaro: rkukura: oh wait you said the owner for a tenant
[11:54am] ivar-lazzaro: rkukura: I guess that could happen, I didn't look in this case specifically. The only thing I check is the "tenant_id" I get from the context
[11:54am] rkukura: ivar-lazzaro: Yes, I’m wondering if an admin can create an L2P for a specific tenant, and specify a shared L3P that the admin owns?
[11:55am] ivar-lazzaro: rkukura: what does that "tenant_id" represent? The issuer of the call or the destination tenant?\
[11:55am] rkukura: ivar-lazzaro: I’m not sure
[11:56am] ivar-lazzaro: rkukura: however, I don't think it would work at the L2P level. When the tenant then creates his privare PTG, the router interface has to be attached
[11:56am] ivar-lazzaro: rkukura: but there's no way for that tenant to retrieve the Neutron router he doesn't own
[11:56am] ivar-lazzaro: rkukura: so the admin could create a private PTG for another tenant and that should work
[11:56am] rkukura: ivar-lazzaro: Right, the subnet(s) of the PTG get attached to the L2P’s L3P’s router
[11:57am] ivar-lazzaro: rkukura: but I really didn't dig into that case, especially since I'm not sure to who the "tenant_id" I have is referring to
[11:58am] SumitNaiksatam: ivar-lazzaro: “private PTG”?
[11:58am] rkukura: I vaguely recall a neutron use case where the external network is not directly usable by tenants, but tenant’s get their own routers on the external network
[11:58am] ivar-lazzaro: SumitNaiksatam: non shared
[11:58am] rkukura: ivar-lazzaro: I’m OK with following the neutron model right now, and doing only what it allows.
[11:59am] SumitNaiksatam: ivar-lazzaro: so if its not shared how would the tenant be able to create a PT on that PTG?
[11:59am] pgpuany joined the chat room.
[11:59am] pgpuany: any one
[12:00pm] pgpuany: looks like bit late for gbp code review
[12:00pm] ivar-lazzaro: SumitNaiksatam: because it's created by the admin but owned by the tenant! This is possible in Neutron
[12:00pm] ivar-lazzaro: SumitNaiksatam: although I'm not sure what would I get in the context in that case... Therefore it's not supported in the current implementation

[12:01pm] SumitNaiksatam: ivar-lazzaro: i am curious how that is done, but we can take it offline
[12:01pm] pgpuany: good any relation between kesytone RGBP and our neutron gbp?
[12:01pm] SumitNaiksatam: pgpuany: RGBP?
[12:02pm] pgpuany: just a question because I missed a lot could not catchup (pramchan@yahoo.com) in juno
[12:02pm] pgpuany: Role based GBP or something
[12:02pm] SumitNaiksatam: pgpuany: is this prakash? (just asking since we are all on first name basis here :-))
[12:02pm] pgpuany: yes
[12:03pm] SumitNaiksatam: pgpuany: welcome
[12:03pm] SumitNaiksatam: pgpuany: as i mentioned before this proposal is purely in the context of what neutron supports in Juno
[12:03pm] pgpuany: I could not get in due to some id issues here sorry folks but will catchup by next week
[12:03pm] SumitNaiksatam: pgpuany: we will definitely like to leverage the new role based access control that keystone proposes
[12:03pm] SumitNaiksatam: pgpuany: thanks
[12:04pm] SumitNaiksatam: okay, back to the review
[12:04pm] SumitNaiksatam: ivar-lazzaro: magesh-gv has posted some comments
[12:04pm] ivar-lazzaro: SumitNaiksatam: answered just now
[12:04pm] SumitNaiksatam: do we need to discuss those here?
[12:04pm] pgpuany: I wanted to ties these two and Congress at top for policy flow and will work on that during long week end and come up with some suggessions for you Sumit & team
[12:04pm] SumitNaiksatam: ivar-lazzaro: ah good, perhaps they were sitting in your drafts, just saw them!
[12:05pm] SumitNaiksatam: pgpuany: nice, seems like you have your long weekend planned for a good cause!
[12:05pm] ivar-lazzaro: SumitNaiksatam: no I literally pressed "comment" one minute ago
[12:05pm] ivar-lazzaro: pgpuany: that's interesting!
[12:06pm] SumitNaiksatam: hemanthravi: do we need to follow up on magesh-gv's comments (ivar-lazzaro has responded) in case you had a chat with him?
[12:06pm] SumitNaiksatam: or anything interesting from those questions that will be helpful for everyone here
[12:08pm] hemanthravi: i haven't talk to him about these, will talk to him if there is any more clarification reqd
[12:08pm] SumitNaiksatam: okay, are we good on the current spec then?
[12:08pm] ivar-lazzaro: There was a question about overlapping IPs
[12:08pm] ivar-lazzaro: so just for clarification
[12:08pm] SumitNaiksatam: ivar-lazzaro: yes, please
[12:08pm] ivar-lazzaro: Overlapping IPs can exist across L3 policies. Different L3 policies can use the same address pools
[12:09pm] ivar-lazzaro: However, overlapping IPs can't exist within the same L3 policy
[12:09pm] ivar-lazzaro: This is regardless of sharing

[12:10pm] hemanthravi: ivar-lazzaro, does sharing attr for servicechainnode, servicechainspec belong in this spec
[12:10pm] SumitNaiksatam: ivar-lazzaro: got it
[12:10pm] ivar-lazzaro: hemanthravi: those are not considered in this initial spec. We can have a separate meeting to decide what we want to share in that context
[12:11pm] ivar-lazzaro: hemanthravi: and have a different blueprint for that
[12:11pm] SumitNaiksatam: i would suggest that in the interest of making progress, lets roll with what we have in the current spec
[12:11pm] SumitNaiksatam: we can add the services' resources separately
[12:11pm] hemanthravi: ok
[12:11pm] SumitNaiksatam: although it would have been good to get done with everything in one shot
[12:11pm] ivar-lazzaro: hemanthravi: I'm not confident in bulking the services here without the feedback of cores who worked on it
SumitNaiksatam: but i would prefer that we dont block on the current effort
[12:12pm] SumitNaiksatam: so keeping that spec/design discussion in mind - the impl patch: #link https://review.openstack.org/#/c/134692
[12:12pm] hemanthravi: ok, will create a separate spec for that
[12:12pm] SumitNaiksatam: hemanthravi: awesome, thanks!
[12:14pm] SumitNaiksatam: perhaps we can quickly discuss the impl here as well
[12:14pm] SumitNaiksatam: ivar-lazzaro: so for one you have added the share attr and update all the extension and db boilerplate to account for this
[12:14pm] ivar-lazzaro: SumitNaiksatam: yes
[12:15pm] SumitNaiksatam: ivar-lazzaro: i do see a fair number of tests (including negative), thats great!
[12:15pm] ivar-lazzaro: SumitNaiksatam: that's mostly testing probably, this is the kind of feature that needs a lot of those to not fall apart
[12:15pm] SumitNaiksatam: ivar-lazzaro: yes
[12:16pm] SumitNaiksatam: ivar-lazzaro: you have introduce something new in the plugin
[12:16pm] SumitNaiksatam: ivar-lazzaro: usage grapg
[12:16pm] SumitNaiksatam: *graph
[12:16pm] ivar-lazzaro: SumitNaiksatam: yeah that is the most important part
[12:16pm] SumitNaiksatam: ivar-lazzaro: so this extends the policy.json semantics?
[12:16pm] ivar-lazzaro: In the plugin I validate the base constraint
[12:16pm] ivar-lazzaro: SumitNaiksatam: policy.json is about RBAC
[12:16pm] SumitNaiksatam: ivar-lazzaro: or does it provide an implemenation for those?
[12:17pm] SumitNaiksatam: ivar-lazzaro: yes
[12:17pm] ivar-lazzaro: SumitNaiksatam: the usage graph makes sure that the objects pointed by shared resources are shared

[12:17pm] ivar-lazzaro: SumitNaiksatam: and that the update happens without breaking those constraints
[12:17pm] ivar-lazzaro: SumitNaiksatam: eg. can't "unshare" a resource when it is pointed by a shared resource
[12:18pm] ivar-lazzaro: So everything is the plugin is basically for the sake of this validations
[12:18pm] SumitNaiksatam: #link https://review.openstack.org/#/c/134692/10/gbp/neutron/services/grouppolicy/plugin.py (for reference)
[12:18pm] ivar-lazzaro: from unshared to shared (False -> True) the usage graph acts
[12:19pm] ivar-lazzaro: vice versa I have specific methods per resource (was harder to generalize)
[12:19pm] ivar-lazzaro: since A knows who it's pointing but it doesn't know from who's pointed
[12:20pm] ivar-lazzaro: traveling the graph in the opposite direction (looking at incoming edges) didn't work either
[12:20pm] SumitNaiksatam: ivar-lazzaro: it might make it easier to understand if you rename “type” to “resource”
[12:21pm] ivar-lazzaro: SumitNaiksatam: where?
[12:21pm] SumitNaiksatam: ivar-lazzaro: L44 and L46
[12:22pm] ivar-lazzaro: SumitNaiksatam: I see the resource as an instance of the type
[12:22pm] ivar-lazzaro: SumitNaiksatam: and what is specified in the usage graph is actually the type itself
[12:22pm] SumitNaiksatam: ivar-lazzaro: hmmm, there is no notion of types in the REST parlance
[12:23pm] SumitNaiksatam: ivar-lazzaro: perhaps you can call it resource name if you want to, but that is what it is here the way you are using it
[12:24pm] SumitNaiksatam: ivar-lazzaro: anyway, minor suggestion (but the description is a bit confusing the way it is, the actual code is more readable in comparison)
[12:24pm] ivar-lazzaro: SumitNaiksatam: ok I'll find something that works
[12:25pm] SumitNaiksatam: ivar-lazzaro: okay, what are the other interesting things that reviewers should be looking for in the patch?
[12:25pm] SumitNaiksatam: ivar-lazzaro: btw, this usage_graph is common across drivers?
[12:25pm] ivar-lazzaro: RMD: #link https://review.openstack.org/#/c/134692/10/gbp/neutron/services/grouppolicy/drivers/resource_mapping.py
[12:26pm] SumitNaiksatam: ivar-lazzaro: i mean it is, just trying to think through if it would need to be over-ridden
[12:26pm] ivar-lazzaro: It's important to verify that the validations are those we expect from the spec, and the behavior while mapping to Neutron as well
[12:26pm] ivar-lazzaro: SumitNaiksatam: It is based on our core API
[12:27pm] ivar-lazzaro: SumitNaiksatam: can drivers change it?
[12:27pm] ivar-lazzaro: (actually, it can probably be generated automatically from the GBP core extension, but we can do it later)
[12:27pm] SumitNaiksatam: ivar-lazzaro: i guess my question was

whether this was always definitive or if there were exception use cases which drivers might want to support
[12:28pm] ivar-lazzaro: SumitNaiksatam: I think this is the minimum number of assumptions that should always be valid
[12:28pm] SumitNaiksatam: ivar-lazzaro: okay good
[12:28pm] ivar-lazzaro: SumitNaiksatam: drivers can restrict, but cannot expand
[12:28pm] rkukura: ivar-lazzaro: remind me why _reject_shared() is called for PRS create and update?
[12:29pm] ivar-lazzaro: rkukura: because there are not shared Security Groups in neutron
[12:29pm] SumitNaiksatam: ivar-lazzaro: okay, so in that case the driver would get a reference to the usage_graph from the plugin?
[12:29pm] ivar-lazzaro: rkukura: we can emulate that of course, but it gets troublesome with the little time we have
[12:30pm] ivar-lazzaro: rkukura: maybe I should add a TODO for that?
[12:30pm] rkukura: ivar-lazzaro: makes sense
[12:31pm] rkukura: ivar-lazzaro: At least a REVISIT would make sense - we'd need separate SGs for each tenant using the PRS I guess
[12:31pm] ivar-lazzaro: SumitNaiksatam: no, the "restriction" of the validation is demanded to the driver's pre_commit and post_commit methods
[12:31pm] ivar-lazzaro: SumitNaiksatam: since changing the graph would change the behavior across all the plugins, which is not always what you want
[12:32pm] ivar-lazzaro: rkukura: correct
[12:32pm] SumitNaiksatam: ivar-lazzaro: you mean across all drivers
[12:32pm] ivar-lazzaro: SumitNaiksatam: yes plugins/drivers
[12:32pm] SumitNaiksatam: ivar-lazzaro: so you are saying that the right place to override would be in the driver's pre_commit and post_commit methods where the validation is actually invoked
[12:33pm] ivar-lazzaro: SumitNaiksatam: you can look at # link https://review.openstack.org/#/c/134692/10/gbp/neutron/services/grouppolicy/drivers/resource_mapping.py for reference
[12:33pm] ivar-lazzaro: SumitNaiksatam: that's how drivers restrict the sharing validation
[12:33pm] SumitNaiksatam: ivar-lazzaro: yeah was looking at that
[12:34pm] SumitNaiksatam: ivar-lazzaro: just wanted to bring that out clearly for the information of other folks who will be writing their drivers (or who would need to modify)
[12:35pm] ivar-lazzaro: SumitNaiksatam: great thanks
[12:35pm] ivar-lazzaro: That said, The last important part to check is the testing. They basically define the whole workflow
ivar-lazzaro: That said, The last important part to check is the testing. They basically define the whole workflow
[12:36pm] SumitNaiksatam: ivar-lazzaro: yes
[12:36pm] ivar-lazzaro: So let's make sure what is there makes sense
[12:36pm] SumitNaiksatam: ivar-lazzaro: to summarize - what level of sharing is supported in the "neutron resource mapping driver"?
[12:37pm] ivar-lazzaro: Sharing PRs, PAs, PCs and PTGs

[12:38pm] SumitNaiksatam: ivar-lazzaro: sharing of PAs requires sharing of service chain specs?
[12:38pm] ivar-lazzaro: When sharing a PTG, also L2P and L3P are shared by constraint defiunition. However, they cannot be used outside the owner tenant (due to Neutron's sharing limit)
[12:38pm] ivar-lazzaro: SumitNaiksatam: that's a grey area. It's probably true but out of scope of this blueprint
[12:39pm] SumitNaiksatam: ivar-lazzaro: i know sharing of services' is out of scope here
[12:39pm] SumitNaiksatam: ivar-lazzaro: however if its a gray area we should not support the sharing of PAs until that gray area is resolved, right?
[12:40pm] ivar-lazzaro: SumitNaiksatam: or we could validate that a PA can't be shared if pointing to a SC ?
[12:40pm] SumitNaiksatam: ivar-lazzaro: exactly, i was just typing that
[12:40pm] pgpuany left the chat room. (Quit: http://www.kiwiirc.com/ - A hand crafted IRC client)
[12:40pm] ivar-lazzaro: SumitNaiksatam: fair enough, please add a comment in the code
[12:41pm] SumitNaiksatam: ivar-lazzaro: okay, i think this shuold probably be called out in the spec as well (and later the spec can updated if it changes)
[12:41pm] ivar-lazzaro: SumitNaiksatam: +1
[12:41pm] SumitNaiksatam: ivar-lazzaro: looking at the spec it should be clear what is shared in the "neutron resource mapping" (even if it means nothing is shared, thought that is not the case)
[12:42pm] SumitNaiksatam: ivar-lazzaro: thanks
[12:42pm] SumitNaiksatam: okay so we have 20 mins left
[12:42pm] SumitNaiksatam: and a mountain to climb in terms of the external access spec!
[12:42pm] SumitNaiksatam: so lets start
[12:42pm] SumitNaiksatam: #link https://review.openstack.org/#/c/136550/
[12:43pm] SumitNaiksatam: External connectivity in GBP ^^^
[12:43pm] SumitNaiksatam: ivar-lazzaro: over to you again to summarize
[12:43pm] ivar-lazzaro: SumitNaiksatam: yeeey
[12:43pm] ivar-lazzaro: Ok the idea is of course to model external connectivity in GBP
[12:44pm] ivar-lazzaro: This is a tricky area... As a matter of fact, we want to make easy the base case
[12:44pm] rkukura: ivar-lazzaro: happened to the idea of simply defining pseudo-PTGs representing external subnets (or the entire Internet)?
[12:44pm] ivar-lazzaro: But since the external world has a lot of "manual" configuration, we need to provide APIs to represent all the useful cases
[12:45pm] SumitNaiksatam: rkukura: i believe this is along similar lines
[12:45pm] ivar-lazzaro: rkukura: the EAP is exactly that

[12:45pm] rkukura: ivar-lazzaro: Haven’t looked yet, but this seems to introduce several new resources
[12:45pm] rkukura: s/looked/reviewed/
[12:46pm] ivar-lazzaro: rkukura: but you need more information! Because the boundary between the could and the outside world could be already existing and manually configured
[12:46pm] SumitNaiksatam: ivar-lazzaro: so i believe the disclaimer is that this aspect of the model bleeds a little more on the imperative side (but that is unavoidable on account of the nature of the problem)
[12:46pm] ivar-lazzaro: SumitNaiksatam: yes! thanks for the great rephrasing
[12:47pm] rkukura: I’m wondering if we need to congigure that through GBP APIs, or could just use neutron APIs for admin-only connectivity detail
[12:47pm] ivar-lazzaro: rkukura: we do because we still need policies (eg. contracts) to be applied. Also not all our drivers will use Neutron in the backend
[12:48pm] SumitNaiksatam: rkukura: my personal preference would be supplementing with GBP model, since its still closer to the intent than neutron APIs
[12:48pm] SumitNaiksatam: rkukura: actually what ivar-lazzaro said
[12:49pm] ivar-lazzaro: Anyway, here is a brief overview of all the new objects:
[12:49pm] SumitNaiksatam: rkukura: i meant the first part of what ivar-lazzaro said (we still use the notion of provide/consume contracts here and the claim is that its easier for the use to understand these)
[12:49pm] SumitNaiksatam: ivar-lazzaro: yes please shoot
[12:50pm] ivar-lazzaro: EAP (External Access Policy): provides and consumes PRSs, and refers to a specific portion of the external world (described by the External Access Segment)
[12:50pm] ivar-lazzaro: an EAP can group multiple EASs, and the EAS can have multiple EAPs
[12:51pm] ivar-lazzaro: The EAS (External Access Segment) is a "segment" through which a L3 Policy can reach the external world.
[12:52pm] KrishnaK left the chat room. (Ping timeout: 246 seconds)
[12:52pm] ivar-lazzaro: It is composed by a cidr (subnet) encapsulation info and a PAT flag
[12:52pm] ivar-lazzaro: which specifies whether we want port address translation to happen in that segment or not
[12:53pm] ivar-lazzaro: The EAS is associated to the L3P (n:m) through an IP address
[12:53pm] rkukura: I’d like to make sure we are separating the “what” from the “how”. Why are things like encap_* needed in the tenant API?
[12:53pm] ivar-lazzaro: which describes the allocated address of the L3P in the segment
[12:54pm] ivar-lazzaro: rkukura: because the external router interface may already exist in a specific vlan
[12:54pm] SumitNaiksatam: rkukura: good question, did not catch that
[12:54pm] rkukura: Why would a tenant know or care about VLANs?

[12:54pm] ivar-lazzaro: rkukura: in the RMD you could just associate a Neutron's network to the segment
[12:55pm] ivar-lazzaro: rkukura: because it's external world, you can't make the choice. Even in Neutron you specify the provider extension info when creating the external network
[12:55pm] SumitNaiksatam: ivar-lazzaro: but that is not part of neutron tenant API
[12:55pm] ivar-lazzaro: rkukura: of course that can be made automatic (a configured pool?)
[12:55pm] rkukura: ivar-lazzaro: Only the admin sees any provider info.
[12:56pm] ivar-lazzaro: rkukura: yes. and the same goes for us
[12:56pm] ivar-lazzaro: rkukura: that's the "admin" part of the configuration. Tenants only deal with PRSs
[12:56pm] ivar-lazzaro: rkukura: or EAPs at most
[12:57pm] ivar-lazzaro: everything else on the model is for the admin to configure (and share)
[12:57pm] rkukura: ivar-lazzaro: I don’t mean to be critical. This may be exactly the API that’s needed. I’m just surprised at the amount of detail being exposed (subject to policy.json I guess)
[12:58pm] SumitNaiksatam: ivar-lazzaro: is the enacp type and value meant to be dynamic?
[12:58pm] ivar-lazzaro: rkukura: I know and I agree with you... That's why I disclaimed that this is a very tricky feature
[12:58pm] rkukura: We should at least consider whether admin config stuff is common, or should be vendor-specific extensions (or neutron-based)
[12:58pm] SumitNaiksatam: ivar-lazzaro: perhaps its more appropriate in a conf file (as static conf)?
[12:58pm] ivar-lazzaro: rkukura: the extension framework will help with that
[12:58pm] ivar-lazzaro: rkukura: right now I think encap_type and value is good enough
[12:59pm] ivar-lazzaro: SumitNaiksatam: the config can be good for some things... Like default vlan pool for external connectivity
[12:59pm] rkukura: ivar-lazzaro: I’m wondering if we should split this into two BPs - one that deals with the model seen by normal tenants, and a 2nd BP for the admin side
[12:59pm] ivar-lazzaro: SumitNaiksatam: but at some point you may really need to configure that manually, without restarting everything
[1:00pm] ivar-lazzaro: SumitNaiksatam: that's because we don't fully control the external world like we do under the cloud
[1:00pm] ivar-lazzaro: rkukura: the Tenant model is probably just the EAP, I don't really feel we should have a separate BP for that
[1:00pm] ivar-lazzaro: rkukura: I can make it clearer in the spec though
[1:01pm] SumitNaiksatam: ivar-lazzaro: i am not disputing the manual configuraion, i am just teeing from rkukura’s point, in terms of what we want to expose in the API
[1:01pm] ivar-lazzaro: SumitNaiksatam: what's the actual difference? I

think that for an admin it would be worse to restart everything to add an external vlan
[1:01pm] ivar-lazzaro: SumitNaiksatam: instead of running an API call
[1:02pm] rkukura: I'd like to see a nice clean compelling store of how normal tenants will use GBP to control incoming and outgoing external access.
[1:02pm] ivar-lazzaro: SumitNaiksatam: as long as only the specific role can access those resources imho we are in good shape
[1:02pm] rkukura: s/store/story/
[1:02pm] SumitNaiksatam: ivar-lazzaro: hence my earlier question, if this "vlan" conf is dynamic or static?
[1:02pm] ivar-lazzaro: SumitNaiksatam: I think most of the cases is static for a single external router
[1:03pm] ivar-lazzaro: SumitNaiksatam: but it may be different when you add more next hops
[1:03pm] SumitNaiksatam: okay so we are a couple of mins over our planned time
[1:03pm] ivar-lazzaro: rkukura: I'll add some CLI example for that. how does it sound?
[1:03pm] rkukura: Can't we just require admins to use neutron to create "external" networks (or any provider network) and pass those to GBP? I really don't like the idea of GBP knowing about network_types and so forth.
[1:03pm] SumitNaiksatam: ivar-lazzaro: thanks so much for providing the context for reviewing this
[1:03pm] ivar-lazzaro: rkukura: with the RMD that's the workflow
[1:03pm] rkukura: I apologize for not having looked at this before.
[1:04pm] ivar-lazzaro: rkukura: but it can't be forced on all the drivers
[1:04pm] hemanthravi: ivar-lazzaro, is l3p addr allocation addresses of hosts/endpoints on the external network?
[1:04pm] ivar-lazzaro: rkukura: I'm ok about removing extra info like "vlans" as long as we have the extension framework to allow all the drivers to implement their workflows
[1:05pm] ivar-lazzaro: hemanthravi: no
[1:05pm] rkukura: Lets not jump to conclusions too quickly - I would like to read and understand what is currently proposed
[1:06pm] hemanthravi: ivar-lazzaro, what does l3p addr alloc represent
[1:06pm] SumitNaiksatam: rkukura: okay, so can we can circle back on this over emails by EoD?
[1:06pm] rkukura: I just did not expect anything this ambitious in the juno timeframe
[1:06pm] ivar-lazzaro: hemanthravi: the L3P address on the external network is specified when you attach the segment to the L3P
[1:06pm] ivar-lazzaro: hemanthravi: oh I thought you meant the ip_pool on L3P
[1:06pm] ivar-lazzaro: hemanthravi: yeah sorry that's what you said
[1:06pm] rkukura: I need to head out for some errands - its snowing and getting dark here
[1:07pm] ivar-lazzaro: Let's keep in mind two things, one is that we

have to give the admin the ability to do manual configuration for external connectivity (giving the nature of the problem)
[1:07pm] ivar-lazzaro: and two that we can't tie the model to Neutron's. any driver has to be able to implement external connectivity
[1:08pm] ivar-lazzaro: that said, I'm all for simplifying the model!
[1:08pm] SumitNaiksatam: ivar-lazzaro: thanks again
[1:08pm] ivar-lazzaro: so let's work on the spec
[1:09pm] SumitNaiksatam: hemanthravi: any other immediate questions?
[1:09pm] rkukura: I'll argue that external connectivity details may really need to be driver-specific. Some might need an API, others might use config. I'd like to see us start with the minimum API needed from the tenant's view.
[1:09pm] hemanthravi: none now...
[1:10pm] ivar-lazzaro: rkukura: I completely agree with that
[1:10pm] ivar-lazzaro: rkukura: still we need something to exist until we get the extension framework
[1:10pm] ivar-lazzaro: rkukura: I'm not a great fun of configuration files
[1:11pm] rkukura: ivar-lazzaro: I hope to have the extension framework code ready to review by Monday
[1:11pm] ivar-lazzaro: rkukura: nice! In that case we can tweak the implementation as needed... But I have to go on with what we have until then
rkukura: ivar-lazzaro: Thanks.
[1:13pm] rkukura: bye
[1:13pm] SumitNaiksatam: okay lets call in wrap
[1:13pm] ivar-lazzaro: bye! have a great thanksgiving
[1:13pm] SumitNaiksatam: thanks all for joining
[1:13pm] ivar-lazzaro: thanks everyone
[1:13pm] SumitNaiksatam: i will post this session in the logs
[1:13pm] SumitNaiksatam: happy thanksgiving, bye!
[1:13pm] hemanthravi: bye