

Dec 15, 2025 | 📅 OVN Community Meeting

Agenda

- (naveen.yerramneni) “Network Function Insertion” Enhancements for OVN 26.03
 - Fallback support (patch out for review)
 - VTAP mode insertion
 - Active/Active Deployment
-

Nov 3, 2025 | 📅 OVN Community Meeting

Agenda

- (dceara) Announcement: [OVS+OVN'25 conference](#) in 2 weeks - hybrid attendance
- (dceara) OVN 26.03 development
 - Soft freeze scheduled for Jan 23
 - Ongoing features:
 - extended EVPN support
 - [distributed load balancers](#)
 - [flow-based tunnel support](#)
 - [OVN Bridge Controller service](#)
 - New features?
- (amusil) Patchwork slowness, is it time for change?
 - is there any desire to switch to some other platform for contributions?
 - options:
 - Github with PRs
 - different patchwork instance
 - gerrit
 - (Mark) different GH or pw instance would be the preference
 - GH PRs reviews and code are in the same interface
 - (Ales) GH PRs might bring more contributors
 - (Felix) Would be interesting to have CI running for each individual commit
 - (Ilya) Reviews in GH PRs are “worse” than in mails because it’s hard to find conversations in GitHub - email notifications are not that useful
 - Jeremy Kerr (maintainer of [patchwork.ozlabs.org](#)) plans to upgrade this instance - may improve performance
 - OVS perspective - not that happy with how the mailing list is managed in general; discussion with LF maintainers to move the mailing list to the same platform as LKM
 - preserves the archives; new archives would be in the lore instance
 - get Linux Kernel patchwork instance instead
 - (Numan) patchwork might make life harder for new contributors
 -
- (dceara) Dedicated ovn-org/ovn(-maintenance) repo for maintainers to run CI?
 - ovn-org has a [GitHub Team subscription](#)

- (Mark) can actions run in parallel if different maintainers apply patches simultaneously?
 - (Ales) yes, the limits we have in ovn-org allow running all the lanes of a CI run simultaneously. ovn-kube CI takes a long time so we would have more capacity to run these lanes on each patch
- (Mark) patchwork checks are done by ovsrobot; should we move patchwork checks to this dedicated repo?
 - (Ales) it's a possibility but we need to ensure we don't hit the limits of the "Team" subscription
- (Felix) Can we use hosted-runners and scale them up when needed?
 - (Ales) bureaucracy might be the problem
 - (Felix) will check if STACKIT can sponsor some runners for OVN

Sep 22, 2025 | 📅 OVN Community Meeting

Agenda:

- (maxlamprecht) We are interested in syncing conntrack between chassis for a smoother failover/maintenance
 - our 2 use cases would be (the first one is more important for us at the moment):
 1. live migration of VMs (Multi-Chassis Port Bindings)
 2. LRP failover on gateway chassis
 - As we plan to move to DPDK in the future, we would like to have a solution that can handle both conntrack implementations(kernel and userspace) - maybe some logic in ovs-vswitchd? We are interested if there are already existing plans to support something like this or if there are any ideas how to build such solution
 - (Ales)
 - one discussed option was using conntrackd
 - we need to sync zone ids between chassis
 - should ovn-controller be the one syncing conntrack data?
 - (Frode)
 - multiple packets of the same session arriving "almost the same time" on multiple chassis
 - scale?
 - (Max) Maybe we avoid syncing conntrack on all members of the HA Chassis group
- (dceara) [26.03 release / dev cycle](#)
 - soft freeze: Jan 23rd 2026 (patches must be posted and have received a reasonable amount of **review on-list** by then)
 - branch created: Feb 20th; release Mar 20th
 - known features under development
 - [Network Function Insertion](#) (slipped from 25.09.0)
 - Extend BGP-EVPN support (L3 support + L2 dynamic FDB learning)
 - [OVN Bridge Controller service](#)

- More ovn-northd incremental processing (routers & switches)
- Distributed load balancers (Alexandra)
- Frode will share potential 26.03 development plans once their planning is done
- anything else the community needs / plans to implement in this dev cycle?

Aug 4, 2025 | 📅 OVN Community Meeting

Agenda:

- (felix.huettner) We are thinking about building a golang wrapper of libopenvswitch.
 - It would link against the existing libopenvswitch for the actual logic
 - We have taken a look at libovsdb by ovn-kubernetes, but we are interested in change tracking as we want to build something similar to northd
 - Is there general interest to have something like this upstream
 - (Felix) will followup with the libovsdb team
- (dceara) [OVN 25.09.0 release/branching](#) - development and review status
- (mmichels) During the past two/three releases we ended up with lots of reviews to do and not enough time to actually review everything
 - proposal1: soft freeze + 4 weeks (reviews and merges) + 4 weeks (testing and bug fixing) before release
 - (Ilya) OVS and OVN release dates will not be “aligned” anymore - new OVS might not yet be accepted in OVS by the time OVN enters soft freeze
 - proposal2: only consider new feature patches that have had at least one round of review before soft freeze
 - (Ales) even if the patch is posted in time people are busy working on their own features and won't have time for the initial review
 - (Ilya) you'll be able to review the patch in the beginning of the next dev cycle + incentivise people to post patches early
- (dceara) Question about the BGP unnumbered e2e test
 - v4 of the patch is on-ist, dceara will review / merge it.

Jun 23, 2025 | 📅 OVN Community Meeting

Agenda:

- (dceara) Update on progress of adding support for EVPN dynamic routing in OVN.
 - [initial proposal](#)
 - (Red Hat internal) tracking stories for this feature:
 - [Support for interacting with linux neighbor tables \(fdb/arp/nd\)](#)
 - [Support to learn remote EVPN VTEPs](#)
 - L2 EVPN [FDB learning](#) and [MAC Binding learning](#)
 - L2 EVPN [FDB advertising](#) and [VIF IP advertising](#)
 - L3 EVPN should already work once remote VTEPs are learnt and dynamic-routing-redistribution/learning is configured
- (numans) OVN Bridge controller (or OVN Provider Controller)
 - [Initial proposal](#)

- (Frode) What if the integration bridge handles the BGP routing? (single NIC chassis - managed by OVS)
 - centralized provider database vs distributed
- (numan) Should it be part of ovn-org/ovn or as a separate project?
- (Ilya) What are we missing from ovn(-controller) in order to fully replace the integration bridge with a logical switch (with direct localnet support).
 - (Numan) the CMS can design their own pipeline (flexibility) + custom rule programming
 - (Han) the abstraction is different in the provider bridge, with different requirements from the original ovn requirements; but OTOH how convenient is using logical flows instead of openflows?
 - e.g., it might be useful to support ACLs in the provider bridge but how can we abstract that?
 - (Ales)
 - if the code is part of the ovn project:
 - what do we do with actions that don't make sense for provider bridges?
 - what about actions that depend on specific registers?
 - what if we just reuse the parsing and define new ovn-pr-controller actions?
- (sragdhara) Network function insertion
 - [Design overview](#)
 - (dceara) related [service function chaining feature request](#) from Tim Rozet (for ovn-kubernetes)
 - would it be possible to enhance the “network function insertion” feature so that we accomodate both use cases?
 - (sragdhara) RFC was sent out last year in fall
 - Mark pointed out potential overlap with composable services
 - After discussions the proposal and design was agreed upon
 - The patch was sent out for review in February + rebase since
 - (Numan) Why are we tying the feature to ACLs?
 - (sragdhara) we reuse the ACL matching
- (dceara) Soft freeze for 25.09.0 is approaching (scheduled for Jul 25) - what other new features is the community planning on adding?
 - (Ilya) if ovn features require OVS features then OVS soft freeze is **~next week**
- (amusil) [Performance regression](#) on main (due to garp_rarp incremental processing changes)
 - the problem seems to be due to engine_runs happening with “recompute disallowed (e.g. sb readonly)” which triggers full force recompute in the next run.
 - only on main so we have some time to find a good, generic solution
- (Alexandra) Is anyone working on adding incremental processing in ovn-ic

May 12, 2025 | 📅 OVN Community Meeting

Agenda:

- (dceara) Adding support for EVPN to the dynamic routing support in OVN 25.09
 - [initial proposal](#) TBD
 - (Felix) this might be used significantly more than the current dynamic routing support
 - (Felix) should we consider using the zebra API?
 - (frode) we don't want to depend on FRR - we're doing implementations with other daemons
 - maybe we can revisit the plug-in idea for routing daemon interactions
 - (Frode) do we care about different EVPN control plane implementations (e.g., Juniper vs Cisco)
 - As an example, this 67 page document details differences in implementation between Juniper and Cisco switches:
<https://www.juniper.net/documentation/us/en/software/nce/evpn-vxlan-interop-between-nxos-and-junos-os/evpn-vxlan-interop-between-nxos-and-junos-os.pdf>
 - As highlighted by odivlad@gmail.com in the meeting, OVN's responsibility ends at the Netlink interface, the rest is up to the routing daemon and deployment tooling.
 - (Ilya) can we avoid a new SB table and instead re-use existing IC infrastructure?
 - instead of a new SB.Learned_VTEP table dynamically manage SB.Port_Bindings representing remote VTEPs (most of the infra used by ovn IC can be reused)
 - (Felix) there's also EVPN over SRv6 or MPLS; let's keep the naming open
- (felix) optimization suggestions
 - scaling OVN environments
 - Assumptions
 - There are no FDB or Mac_Binding updates to southbounds
 - There are large network groups on the Northbounds that can be mostly separated to individual Northbounds
 - Idea
 - Individual Northbound shards based on related networks, there should be minimal cross-shard connectivity
 - Southbound per X amount of Chassis
 - Ovn-controller connects to the southbound responsible for its chassis
 - Northd writes all data to all southbounds that have chassis that need that data
 - Alternatives
 - Multiple ovn-controller per chassis that connect to different ovn clusters. This idea here would keep the complexity in northd instead of potentially needing a new ovn-controller when spawning a VM
 - Using ovn-ic with transit switches similar to ovn-k8s

- Drawbacks
 - Northds might be under large load
 - This might only improve efficiency if specific constraints are met and is problematic otherwise (e.g. large amount of FDB, Mac_Binding updates)

Mar 17, 2025 | 📅 OVN Community Meeting

Agenda:

- OVN 25.09 development cycle / what features does the community wish and plan to add to the next OVN release?
 - (dceara) Extend dynamic routing / BGP support for the EVPN case.
 - (dceara) Follow up on TODOs for the dynamic routing / BGP support added in 25.03.0
 - (dceara) Incremental processing in northd for LS/LR changes
 - (numans) local datapath optimization
 - port mirroring - v7 is on list for review
 - (Ilya) remove deprecated features - STT support, older TLS versions (submodule bump to OVS main)
 - (numans) ovn northd logical flow documentation is often out of sync
 - remove the man page entirely?
 - some information might still be useful
 - or document that it's out of sync
 - Numan will start a discussion on-list
 - ovn.org
 - (Alin) we should add a link to the technical meeting on ovn.org
 - (Ilya) remove references to OVS-OVN split
 - (numans) Composable services support (Mark Michelson)
 - (dceara) Update ovn-fake-multinode version (to latest) and update OVN CI to use the latest version.
 -
- OVN (conntrack) HWOL
 - (dceara) NVIDIA [reported a breakage wrt HWOL](#)
 - one of the discussed potential fixes was to [commit all traffic in the SNAT zone](#)
 - this was [implemented in 25.03](#) with an additional LR option
 - but it seems it doesn't fix the NVIDIA use case - what are the next steps?
 - (numans) if we have an ovn-fake-multinode test case that offloads to TC it might be representative enough
 - (Alin) have an internal CI setup that does nightly regression testing wrt HWOL - Alin to follow up internally with Han
 - (Alin) ovn-kubernetes use case when accessing the pod via node port
 - suggestion - revert original fix + commit-all feature
 - reverting these two patches re-introduces two ovn-kubernetes issues the patches were fixing

- Alin will try to implement an ovn-only system test
 - Ales - could we at least get OVN/OVS databases?
 - Alin - problem is related to ct zones not being accessed in the correct order in both directions
 - Alin will share the OVN/OVS databases on list
- (Alexandra) Load balancers/stateful ACLs in combination with allow-stateless ACLs
 - Does anyone use stateless ACLs + stateful configs (e.g., LBs)?
 - It doesn't seem to work because there's no CT entry.
 - (numans) we mix stateless + stateful ACLs (different tiers).
Suggestion:
 - use stateless ACLs in tier 1
 - use stateful ACLs in tier 0
 - (Vladislav) stateless ACLs + LB - we don't send traffic to conntrack because it matched stateful ACLs - make an exception if traffic is destined to LB VIP
 - (Vladislav) switches are configured with use_ct_inv_match=false
 - (Ilya) stateless ACLs must be "symmetric" (in original and reply direction)
 - (numans) Could you please share the OVN databases / ACLs that correspond to this problem?
 - Vladislav/Alexandra to share this info on-list

Feb 10, 2025 | 📅 OVN Community Meeting

Agenda:

- (dceara) OVN 25.03.0 release/branching - development and review status.
 - current status of [feature patches](#)
 - items 3, 4, 8 were merged
 - items 5, 6 have new versions incoming
 - the rest have no updates
- (dceara) [\[ovs-dev\] \[RFC ovn v3 4/4\] northd: Advertise Routes for LB and NAT VIPs.](#)
 - small redesign proposal - do we really need to determine the actual VIF that owns private NAT IPs?
 - can we just advertise routes with tracked_port == DGP (distributed gw port)
 - (Martin/Frode): yes, Martin will post v4 as soon as possible
 - (Dumitru): to help out iterating early on this new approach
- (amusil) [OVN port mirroring](#)
 - Mainly to clarify the requirements if there is a possibility to simplify
 - (Ales) would it be enough to have a single filter rule per mirror?
 - (Alexandra) it's probably needed to have multiple filters (for different types of traffic, e.g., tcp or icmp)
 - (Ales) The mirror sink can be on a different chassis. Mirrored traffic just gets sent out (egress) pipeline on the LSP sink. We cannot further process (ACL) the mirrored traffic.

- (Ales) will follow up with detailed suggestions on the mailing list
- (numans) - Local datapath optimization in ovn-controller
 - Not 25.03 candidate
 - https://github.com/numansiddique/ovn/commits/optimize_ldps_original_approach/v1/p3/
 - https://patchwork.ozlabs.org/project/ovn/list/?series=443663&state=*
 - requested to be considered as potential backport for the 25.03 branch (performance optimization)
- (MJ) BGP e2e multinode test ready to be posted to the mailing list

Jan 13, 2025 | 📅 OVN Community Meeting

Agenda:

- (dceara) Soft freeze for 25.03 approaching (Jan 24th). Current list of features or user-visible behavior changing patches posted for review (approx sorted based on current FIFO order in patchwork):
 - [\[ovs-dev,RFC,v2\] northd: Commit all traffic when there is stateful NAT/LB.](#)
 - [\[ovs-dev\] northd: Use lower priority for all src routes.](#)
 - [Introduce the concept of transit router.](#)
 - [\[ovs-dev\] northd: Add ipv6_{src,dst} to selection_fields column in the NB db.](#)
 - [Add "bypass_match_for_established" ACL option.](#)
 - [\[ovs-dev\] Logical Router Policy chains.](#)
 - BGP:
 - [OVN Fabric integration: Northd.](#)
 - [OVN Fabric integration: active-active.](#)
 - [OVN Fabric integration: Ovn-controller](#)
 - [\[ovs-dev,v4\] Allow LR to send RAs through localnet port.](#)
 - LB and NAT announcements (not posted yet formally)
 - [Introduce ECMP_nexthop monitor in ovn-controller](#)
 - [\[ovs-dev\] treewide: Introduce OVN overlay port mirroring support.](#)
 - [\[ovs-dev\] northd: Options for pass connection tracker for stateless lflow with LB.](#)
 - [Incremental processing of IGMP changes in northd.](#)
 - [\[ovs-dev,v0\] ovn-ic: Add support for route prefix filter.](#) (needs v2)
 - [northd I-P for logical switch creation](#)
 - (Ales) are we OK with having the same UUID between NB and SB - because a LS and LR might have the same UUID in NB?
 - (Ilya) "spine-leaf" switch topology
 - Community reviews/tests are very welcome!
- (dceara) BGP update for 25.03
 - OVS patches almost ready (Frode will post a potential final version today)
 - OVN reviews ongoing (dceara)
 - Canoical is adding LB and NAT route announcement - not yet posted for review
- (mmichels) Composable services are unlikely to land in 25.03
 - en-northd/northd.c node needs refactoring because data ownership is not clean enough

- conflicts with other incoming features
- will slip to 25.09 - hopefully not experimental anymore
- (frode) IPv6 prefix delegation
 - Have done some work on this in my spare time over the x-mas break, will make an attempt to get it into a reviewable state, but not sure if time suffices before soft freeze. Main goals for what I want to accomplish:
 - Run one state machine per GW LRP (those with `prefix-delegation` option set) as opposed to one per LRP (those with `prefix` option set) that is the recipient of the prefix.
 - Delegated prefixes are usually shorter than /64 and we usually want exactly /64 per network.
 - Aggregation reduces processing in OVN and the number of leases and routes in the infrastructure.
 - Ideally the implementation should be able to provide a hint for which prefix length is required for a given GW LRP and in the event the prefix length provided does not suffice it should be able to request multiple prefixes, but realistically that will be a future improvement.
 - Support multiple different prefixes per IA_ID:
 - The implementation appears to assume the DHCP server to provide exactly one prefix per DHCPv6 PD request, however implementations exist that will in fact provide multiple prefixes in response to a single request.
 - The practical use case for this is sites that have an RFC 4193 ULA prefix in addition to globally routable prefix and want applications to be able to reach all endpoints using both prefixes.
- (i.maximets) “spine-leaf” switch topology
 - Will attempt to post before soft freeze.

Dec 16, 2024 | 📅 OVN Community Meeting

Agenda:

- (dceara) BGP update for 25.03
 - [Prerequisite patches from Felix](#) merged in main
 - [SB route table schema](#) merged in main
 - Planning to review and potentially apply the following series:
 - Martin's [IPv4 routes over IPv6 next hop addresses](#).
 - MJ's [LRP + RA changes needed for BGP](#).
 - Should we split remaining BGP integration work between multiple parties, e.g., ovn-northd and ovn-controller bits?
 - Frode is working on a new iteration of the OVS patches (Eelco reviewed the latest version)
 - Felix will post the northd side of the BGP support first.
 - Felix will try to update the ovn-controller side of things - Dumitru will try to look at that part too

- (dceara) should we remove the experimental tag on the protocol-redirect feature?
 - (Martin) maybe with the composable services support we don't need protocol-redirect?
 - (Mark) probably 25.03 will likely contain a single experimental composable service - 25.09 is a more reasonable timeline for a "real/actual" support
- (Gurpreet) is there any feature gap between the kernel datapath and the userspace (dpdk) datapath
 - functionality wise there shouldn't be any gap
 - there might be some impact on performance due to BGP control traffic but needs proper testing
- (Vladislav)
 - a few patches sets for 25.03
 - logical router policy chaining
- (mj) Difficulty with Localnet RAs
 - [Context](#)
 - we need periodic RA over localnet
 - ovn-controller drops periodic RAs on localnet ports
 - this breaks BGP unnumbered
 - proposed idea: drop only RAs that are not from gw routers or gw router port but that's hard to determine just from packet contents
 - potential follow up solution: when injecting the periodic RAs from ovn-controller (pinctrl) set a flag in the OF "flags" register to mark whether the packet is generated by a gw router or gw router port.

Dec 2, 2024 | 📅 OVN Community Meeting

Agenda:

- (Frode) BGP update and division of work
 - (Tiago Pires) I'm testing Felix's work (A/A patches and BGP separately). My use case will need OVN-IC to learn the external routes that were learned via BGP. We have a inhouse patch for it that in some moment we can submit.
 - (Frode) multiple streams of work going on:
 - i. BGP unnumbered (Martin/MJ)
 - ii. Patches for OVS
 - export routing table library
 - namespaces support
 - iii. Patches for OVN -
 - northd infrastructure
 - ovn-controller infrastructure
 - (Frode) Should we split up work?
 - i. Felix agreed
 - ii. Frode can pick up the work on OVS patches (routing table)
 - OVS soft freeze is January 2nd - patches must be (at least in some capacity) reviewed
 - iii. Felix can add namespace support on top

- (Ilya) the namespace support might be a bit controversial
 - Eelco is going to review the OVS changes
- (Felix) If we agree on SB schema we can work on northd / ovn-controller changes in parallel
 - (Frode) most of the information to expose/learn is probably per chassis (VIP, NAT, default gw, etc).
 - (Frode) Is there a risk of affecting scale due to DB size explosion?
 - Current DB Structure
 - "Route": {
 - "columns": {
 - "datapath":
 - {"type": {"key": {"type": "uuid",
 - "refTable": "Datapath_Binding"}}},
 - "logical_port": {"type": "string"},
 - "ip_prefix": {"type": "string"},
 - "nexthop": {"type": "string"},
 - "tracked_port": {"type": "string"},
 - "type": {"type": {"key": {"type": "string",
 - "enum": ["set", ["advertise",
 - "receive"]]},
 - "min": 1, "max": 1}}},
 - "indexes": [["datapath", "logical_port", "ip_prefix",
 - "nexthop", "type"]],
 - "isRoot": true}
- (numans) Did removal of the configuration option for centralize_routing break unit/system tests?
 - (Felix) not in my github test runs
- (dceara) Can we/should we change the OVN contribution guidelines to make it easier for new contributors to join the community?
 - should we stick with the “patch on mailing list” approach?
 - should we improve our docs?
 - should we start accepting different ways of contributing code (officially)?
 - (Felix) corporate email addresses/email clients can hardly be made work with this.

Oct 21, 2024 | 📅 OVN Community Meeting

Agenda:

- BGP series discussion
 - Spec doc in progress:
 - 📄 [EXTERNAL] OV013 - OVN layer 3 fabric integration
 - (Frode) mostly an elaboration of the original RFC - everything else on top, in Felix’s series, is a nice-to-have “bonus”
 - (Felix) Status: working implementation was sent; now cleaning up commit history, adding tests etc. Will send a new version this or next week.
 - The direction taken is probably compatible with OVN-K plans. Will also allow to drop the bgp agent from OpenStack.

- It's a large patch set, the sooner it's posted on ML the better.
- OVN 25.03 development cycle
 - what new features is the community working on?
 - mmichelson: OVN composable services (copy of document: [here](#))
 - No code written right now
 - ~Nov to post a scaffolding patch to register services, create bindings + test/no-op/null service
 - External service to push packets elsewhere for decision / processing before returning back to OVN
 - Building block for function chaining requirements from OVN-K (?)
 - The hope is this is mostly an additive work, should not affect those who don't opt in
 - numans: I-P for logical switch and router in ovn-northd
 - amusil: stateful LR - commit all
 - vladislav:
 - changes in the logical router policy pipeline - Aleksandr sent an RFC patch to support a new action "jump" (resubmit to the same pipeline) to allow chaining policies
 - support for port mirroring inside OVN: mirror any (multiple) port inside the VPC to a monitoring VM
 - Ilya: ideas to improve OVN-IC
 - support transit routers to avoid having to use the transit switch every time packets need to cross AZs
 - for L2 topologies with interconnect ovn-k configures transit switches, i.e., for each container create remote ports on all remote AZ transit switches
 - we're limited on the number of logical switch ports per switch (32K)
 - multicast causes long resubmit chains
 - idea: instead of a single large TS create a "spine-leaf" topology with "leaf switches" connected to transit "spine switches"

Sep 16, 2024 | 📅 OVN Community Meeting

Agenda: (

- BGP series discussion
 - Context:
 - [Original RFC from Frode](#)
 - [Current development thread](#)
 - WIP (?) development branches from Frode/Felix/Roberto:
 - https://github.com/rbartzen/ovn/commits/test_bgp_and_ovn_ic/ (Roberto)
 - https://github.com/felixhuettner/ovn/commits/test_active_active_routing_v4/ (Felix)

- (Frode) summary:
 - about to start a new development cycle
 - bgp support falls into this new development cycle
 - initial rfc added support for redistribution of nat addresses on gw routers but not for distributed routers
- (Felix)
 - switched the approach to distribute info between northd and ovn-controllers
 - new SB table (a place to mainly receive routes)
 - patched the code to support network namespaces
 - patches need “clean up” but are mostly feature complete
 - (Martin) why use namespaces instead of VRFs?
 - (Frode) could we combine using namespaces/vrfs in the same deployment?
- (Frode) what is the reason behind the new “route” table in the SB?
 - handle failovers within OVN
 - it’s also useful for the route learning use case (also for OVN-IC)
 - ovn-controller only has a local view
- (Vladislav) Did graceful shutdown stop ovn-controller (and remove routes)?
 - No, just chassis failover (priority change of the chassis), which means no new traffic should enter the node anymore
- (Vladislav) To learn routes from BGP in a single AZ do we still need a transit switch?
 - yes
- (Frode) How do we address the use case when NAT is not needed?
 - we might need to distribute routing information between logical routers
 - The new SB table has a “tracked port” column that identifies the chassis that should advertise at a higher prio
- (dceara) should we come up with a plan to split the work in smaller chunks?
- (Roberto) How do we describe what routes to be advertised/learnt? (i.e., how do we configure that?)
 - come up with the set of use cases that we want to support and be able to configure
- (Han) Should we start a google design doc for the feature?
 - it could also be used as a starting point for defining the smaller (individual) tasks people can work on (potentially in parallel)

Jul 29, 2024 | 📅 OVN Community Meeting

Attendees:

Agenda:

- BGP series discussion
 - (Ales) are we in a rush for BGP support to land in 24.09? Maybe it’s good to have a “complete” feature in a release instead of getting the whole functionality in multiple releases
 - (Ales) concerns about netlink code in OVN. Options:

- reuse the OVS netlink code (expose the code in OVS and bump submodule version)
 - (Frode) no direct consumer in OVS, is that OK for OVS maintainers?
 - (Ilya/Numan) it might be better if ovn-specific wrapper functions live in OVN
- or OpenFlow extension to allow programming routes from the controller
- (Jakub) can the netlink implementation still be used with the “DPDK” datapath?
 - (Frode) yes, netlink is used just to exchange information in an “universal” way
- (dceara) Could the netlink OVN code live inside of a plugin?
- (dceara) Can the bgp-mirror patch be accepted standalone in 24.03?
 - (Vladislav) it might be good to allow experiments with FRR configured externally
 - (Mark) should we document what “experimental” means?
- (Martin/Vladislav) Supporting forwarding BFD traffic to the routing daemon (similar to BGP traffic).
- (Ales) is there any plan for phase2 (learning routes)?
 - (Frode) yes, it simplifies the configuration in the cluster but it's early
 - (Ales) phase 3 for Red Hat might be EVPN support (creating vxlan tunnels)

Jun 24, 2024 | 📅 OVN Community Meeting

Attendees: Dumitru Ceara Numan Siddique frode.nordahl@canonical.com
 i.maximets@ovn.org Ihar Hrachyshka Daniel Alvarez Sanchez Tim Rozet
 Xavier Simonart Ales Musil hzhou@ovn.org

Recording: <https://www.youtube.com/watch?v=bonLhUj2oGg>

Transcripts:

https://drive.google.com/file/d/1ioVXcHiuD-xr5RQTh5OSGbpmhUtlDiO_/view?usp=sharing

Agenda:

- Discussion: Tighter BGP integration between OVN and BGP
 - (Frode) will show what they are playing with
 - (Frode) Should it be a generic L2 LB mechanism or specific to BGP?
 - (Gurpreet) What's the benefit of native integration compared to other/existing solutions?
 - (Frode) First-class citizen status of the feature; may help with the load; less complexity; hardware offload; may be a good starting point to add support for things like EVPN
 - Dumitru: related: mirror filtering feature; L2 forwarding by match
 - (Ilya) potential use case for BGP (maybe not uniquely for BGP?): use BGP to advertise routes from VIF attached to network, e.g. k8s in OpenStack VMs

- (Ilya) maybe instead of hacking fr into ovn, we should have in-tree implementation. It would be challenging to re-implement the state machine; on the other hand, it's a limited set of features that we'd actually need?
 - Supportability may be a concern
- (trozet) FRR has a "zebra" API that OVN can talk to
 - (numan) if it's not optimal, we could also extend it upstream for OVN use cases
 - Disadvantage: tying OVN to FRR API
- [dceara] Should we start having a public way of tracking the roadmap for upcoming releases?
 - Potential ways of doing that:
 - GitHub milestones
 - A public online document (google doc?)
 - Keep the TODO.rst up to date in the repo or maybe add a ROADMAP.rst and link to it from the readme/website?
 - others..
- OVN 24.09 planning: [dceara]
 - soft freeze is approaching (July 15th), what new features are being worked on for the 24.09.0 release?
- Action Items
 - (Frode) will send email with bgp implementation details to dev@

May 13, 2024 | 📅 OVN Community Meeting

Attendees:

Recording: <https://youtu.be/k448ada9aFQ>

Transcripts: https://drive.google.com/file/d/1VuF5l9wSR9z4rIH_LVF3PJBScj4Mb8JT

Agenda:

- Discussion: [Frode/Vladislav] explore tighter integration between OVN and BGP
 - Goals
 - Re-use existing BGP implementations as much as possible
 - BiRD, FRRouting, Holo Routing, others?
 - The running of a host level BGP daemon is up to the end user and outside the scope of this work.
 - It must be possible to construct a hardware offloaded data path.
 - Build infrastructure for route consumption / leaking from / to OVN at the chassis level
 - Allow the fabric to route traffic directly to resources such as:
 - Floating IPs
 - Load Balancer VIPs
 - Fixed IPs
 - Use cases for OVN consumption of fabric provided routes?
 - Different external routers
 - How to leak routes to BGP daemon?

- The netlink route approach that ovn-bgp-agent uses does appear appealing
 - Standard API
 - Mechanisms exist to isolate the routes (VRF)
 - A bit abuse of the kernel as a database
 - Are there any alternatives that do not involve implementing some route exchange protocol in OVN?
 - How to consume routes from BGP daemon?
 - Does this fit within core OVN or is a plugin layer ala OVN-VIF required?
 - Start with inside ovn-controller with the netlink interface to leak routes
 - If we get multiple approaches we can branch out to plugin framework later
 - How to make BGP session from OVN Router IP?
 - Use cases?
 -
 - Rewriting any part of the packet does not work well due to how BGP MD5 authentication works.
 - Would it be possible to do some sort of L2 level LB and transmit messages to/from TCP port 179 on the OVN router IP to an internal port?
 - To investigate
 - pinctrl forwarding of TCP port 179 over a unix socket?
 - Hardware offloaded data path
 - The ovn-bgp-agent currently suggests a topology with additional OVS bridges and a second OVN cluster for managing the external networking in a hardware offloaded way.
 - Does this serve any other purpose than to make OpenStack work this way without changing the core topology of the Neutron OVN driver?
 - Could we do with a single OVN cluster if the CMS set up the external network differently?
- OVN 24.09 planning: [dceara]
 - improved OVN observability: revive and expand Adrian's [RFC series ACL Sampling using per-flow IPFIX](#)

Notes:

- OVN+BGP
 - We should reuse existing projects
 - OVN agent is OpenStack specific
 - OVN controller already knows needed info, no need for SB connection
 - Not sure if there's use case to consume routes from fabric
 - Can have different default gw routes advertised from fabric
 - With ECMP no need for BGP - use multiple routes in OVN
 - Useful when fabric routing is dynamic, unrelated to ICMP
 - (Vladislav) is it just default routes or more?
 - They have a use case to install BGP routes in multiple routers
 - chained routers

- Right now doing with CMS but would be nice if OVN does
 - (Han) what's the major disadvantage of CMS approach? Running BGP between routers is itself complex.
 - (V) no answer
 - Maybe contribute CMS integration upstream instead
 - Run BGP sessions between routers as peers to propagate?
- How to push routes into an agent
 - Openstack agent pushes through netlink, using dummy interface
 - (Dumitru) it's a feature not a bug; using interfaces that can be used by other workloads on the host
 - OVN will reuse the approach (VRF and set routes there, just not external daemons)
 - Plugins may be helpful because then we can have different implementations (one through netlink, another through e.g. grpc to FRR or wherever else)
 - This can be done later, factor out when the 2nd implementation is really needed
- (Frode) prefer having leaking routes from OVN as V1; leave consuming from fabric for V2
- (Dumitru) L2 policies could be used to pass BGP frames (match on port etc.) downstream to next routers
 - Would need to somehow automatically configure the policies

Mar 19, 2024 | 📅 OVN Community Meeting

Attendees: frode.nordahl@canonical.com numans@ovn.org Ales Musil

Lorenzo Bianconi hzhou@ovn.org Ilya Maximets Dumitru Ceara

Recording: <https://youtu.be/mE1ByXeCQ7A>

Transcripts:

Agenda:

- Recurring item: anyone opposing this instance of the meeting being recorded?
- [felix.huettner] The openstack large-scale sig is looking for interested parties for live sessions to share their openstack deployments. If someone here is running openstack and wants to share it, let me know
- [dceara] OVN 24.09 planning (soft freeze July 15th 2024):
 - what are the features the community needs to get into OVN 24.09.0?
 - IPv6 prefix delegation?
 - DHCP relay agent
 - Geneve provider networks
 - Any more northd I-P? E.g. special cases of port group, etc.
- [dceara] ovn-heater OpenStack scenarios:
 - we had a community effort to figure out scale requirements for various OSP deployments:

- [Questionnaire](#)
- [Survey results](#)
- [ovn-heater re-architecture](#) for OSP/neutron-like CMS [DONE]
- should we define “official” ovn-heater test scenarios that use these?

Notes:

- OVN 24.06
 - Frode uses IPv6 prefix delegation in personal lab deployments, not necessarily a business need for it yet
 - Canonical to complete roadmap planning
 - host-based networking / BGP
 - maybe we need tighter integration with OVN for BGP
 - Red Hat might also need tighter integration with BGP
 - Frode: “anyone doing C dev should consider learning rust”
 - should we consider adding some new features/modules in Rust?
 - make it as little-effort as possible
 - Numan asks: what’s the ultimate goal? Are we trying to eventually remove the C part?
 - Frode: it made into the kernel, safer language
 - Frode will try to come up with a proposal, e.g.:
 - rewrite ovn-vif (provider side from OVN) in rust
 - Numan: Geneve provider networks
 - currently not working because of certain assumptions
 - Han: more northd I-P?
 - first port in port group triggers recompute
 - more frequent in OpenStack (first time a VM is added to a Tenant logical switch)
 - it’s actually an lflow recompute for each port group modification
 - ACL I-P:
 - dceara plans to look into this in the next release cycle
- ovn-heater for OpenStack:
 - there’s still a need for scale testing ovn in OSP-like scenarios
 - we need to continue trying to figure out the right configurations to test
 - Ilya: we should have a large scale realistic scenario upstream to allow us to detect regressions (downstream) and report back upstream
 - dceara: will open a PR to add large scale OSP-like scenarios to ovn-heater and share it on the ovs-dev mailing list too
 - Han: it’s more practical to have several (2-3) most critical profiles with different focus

Feb 19, 2024 | 📅 OVN Community Meeting

Attendees: Martin Kalcok Ales Musil frode.nordahl@canonical.com Xavier Simonart
Dumitru Ceara

Recording: <https://youtu.be/QjKV64EWxZ0>

Agenda:

- Recurring item: anyone opposing this instance of the meeting being recorded?
- [felix.huettner] The openstack large-scale sig is looking for interested parties for live sessions to share their openstack deployments. If someone here is running openstack and wants to share it, let me know
- ...
- Next meeting?

Notes:

- Recording
- Canonical put a snapshot of the ovs/ovn release in their development release for testing
- Move Felix's topic to the next instance of the meeting, in case more OpenStack folks are available.

Jan 8, 2024 | 📅 OVN Community Meeting

Attendees: Dumitru Ceara felix.huettner@mail.schwarz Mark Michelson Ales Musil @Vladislav Odintsov frode.nordahl@canonical.com Martin Kalcok Ihar Hrachyshka Ilya Maximets Xavier Simonart

Recording: <https://drive.google.com/file/d/1uB5mQapsqxMkG5BwtrGHmeUgN9K28XBI/>

Transcripts: <https://drive.google.com/file/d/1mNb84ffQ2cGhUcaRUUUzfF9DSxgcURIL>

Agenda

- [mmichelson] OVN 24.03 soft freeze is scheduled for 19 January.
- ..
- Next meeting?

Notes

- [Recording](#) and [chat](#)
- 24.03 soft freeze:
 - Any problem with the date? What's the state of pending planned patches?
 - What's the state of Numan's patches for I-P?
 - (dceara) seems in advanced state, should be possible to get in time
 - (Ales) CT limit for ports
 - Not sure if it will be on time though
 - Doesn't matter if it's in this release
- (Martin) Issues with OpenStack SNAT networks
 - External to internal traffic gets snatted mid-stream (after handshakes)
 - Numan: should use snat_and_dnat, not snat
 - https://bugzilla.redhat.com/show_bug.cgi?id=2149021
 - Martin to report on ML with a reproducer
- (Vladislav) Discuss BGP integration with OVN
 - Use case: public cloud
 - announce FIPs to leaf switches (ovn-bgp-agent in openstack)
 - They are not using OpenStack, having their own custom solution - wonder if this would fit into a native OVN solution
 - AWS DirectConnect
 - BGP peer a hw router with your virt infrastructure

- Run BGP speaker on hw router
- Running FRR via DNAT
 - Problem: packets are modified, so can't use authenticated BGP (?)
- They then populate NB from BGP info
- Anyone else using BGP without OpenStack? Thoughts?
 - (Felix) also using it for FIPs, but they use OpenStack + bgp-agent
 - (Felix) were also thinking about DirectConnect scenario but nothing specific
 - (Dumitru) should OVN natively talk BGP?
 - (Vladislav) FRR has plugins system; lots of risks in re-implementing BGP since it's very complex
 - (Dumitru) could we add a proper plugin to FRR to update NB?
 - (Vladislav) how to send matched traffic from ovn-controller to FRR? Not sure how it would look like networking environment wise.
 - (Dumitru) Add router; add policy to send it elsewhere
 - We could probably add policies to LSs and avoid LRs
 - (Numan) should it be a separate project under OVN org?
 - (Vladislav) Next steps?
- Ilya & Frode - discussion about cooperative multi-threading support in OVS (for OVSDDB):
 - Frode to post RFC patches/exception request.
- Next meeting in ~6 weeks, Dumitru will schedule

Nov 27, 2023 | 📅 OVN Community Meeting

Attendees: Ihar Hrachyshka Ilya Maximets Mark Michelson Lorenzo Bianconi
 Martin Kalcok Ales Musil Han Zhou Dumitru Ceara Numan Siddique Frode Nordahl

Recording: <https://drive.google.com/file/d/1sK3-KVT726FXc0kLrKcS83jwh5gPYAuc/view>

Agenda

- Recurring item: anyone opposing this instance of the meeting being recorded?
- 24.03.0 release:
 - Progress update?
 - Any new features that might be requested/are in the process of being requested for 24.03?
- ..
- Next meeting?
- [mmichelson] OVN Service Extraction proposal
- Register for OVS+OVN Conf'23:

<https://mail.openvswitch.org/pipermail/ovs-announce/2023-November/000332.html> !!!

- Dec 6-7 <http://www.openvswitch.org/support/ovscon2023/>

Notes

- [Recording](#)
- 24.03 release topics
 - Numan's [patches](#) for I-P for Logical Flows
 - Dumitru is looking into I-P for ACLs, pending if to be included in the release
 - Mark will pitch Service Extraction later, if interested, this may be considered for the release too
 - [DHCP Relay Agent support](#) from Nutanix?
 - Han mentioned multi-NIC support for tunnel interfaces
 - Running multiple controllers is not ideal
 - The idea is to define multiple encaps (already supported) but there are issues (?)
 - Patches not posted yet
 - (Ales) Why not localnets with different VLANs?
 - (Han) requirement is to have different **tunnels**
 - (Frode) reiterating, no major requests but
 - Looking into some issues with NAT FIPs
 - Predictable schema versions (?)
- Mark's proposal (OVN Service Extraction)
 - Slides
 - TL;dr split out special logic into "Services" (NAT, ACL, LB)
 - Comments?
 - (Frode) sounds super exciting, need more details how this will be modelled in database
 - Mark had and has some ideas
 - Maybe a "Link" type to link any two things together
 - (Ales) How would tables look like?
 - Mark assumes a similar "jump-to-prev table" approach
 - (Ales) Means we'll still be limited as to the number of services to chain (or use a different metadata?)
 - Mark: not table per service, we should be able to jump back as much as we want (in theory)
 - (Dumitru) I see these as "new logical appliances". We can link them through logical patches. There are limits like resubmits, but it's not a new problem. It's a nice thing. Big gain - simplifying I-P a lot.
 - (Dumitru) Service Extraction is a bad name. We are not deprecating anything, it's new features.
 - Mark: maybe they won't **interact** well with "old" stuff. Don't use both.
 - (Han) like in general. Maybe going further with API to define pipelines? Also, how big of a change / impact would we have to the current implementation? Performance concerns around limits in openflow stack size / resubmits etc. Also, splitting conntrack calls between "appliances" may also be a problem since we usually try to avoid it as much, and go at length to squash stage operations.
 - Mark: not planning to rewrite much, it's new constructs.

- (Ales in gcomments) though dropping old way eventually would be nice
 - Mark: idea to simplify implementation for some new features, as an example - conditional SNAT (?)
- (Ales) was thinking about splitting pipelines for some of these, though not in terms of completely separate “Services”
 - Mark: ACL or LB appliances shouldn’t care where they are attached - to router or switch pipelines
- (Mark) it’s a lot of work to define and implement to make sure in the end it’s better than what we have now
 - 23.04 is LTS - introducing a Completely New Way is a concern
 - But we may start already, pick one, e.g. LB and see how it goes.
 - Maybe in 26.03 we may even consider dropping the “old” ways?
 - btw OVN has no deprecation method, Mark will check it
- (Mark) first impression seems to go well, tech proposal will be posted in ML
- OVS/OVN Conference
 - Next week
 - Registration link above, please register
 - Topics already posted, see links in the Agenda section
- Next meeting - some time first weeks of January

Action items

- ☐ Mark will check deprecation method in OVN in context of Service Extraction
- ☐ Mark to attach link to slides (above)
- ☒ ~~Dumitru to attach video recording (above)~~
- ☒ ~~Dumitru will propose next meeting dates on ML~~

Oct 23, 2023 | OVN Community Meeting

Attendees: [Ales Musil](#) [Dumitru Ceara](#) [Frode Nordahl](#) [Han Zhou](#) [Ilya Maximets](#)
[Ihar Hrachyshka](#) [Lorenzo Bianconi](#) [Mark Michelson](#) [Numan Siddique](#)
martin.kalcok@canonical.com

Recording: https://drive.google.com/file/d/1UyW9czlq2kh4iMWlvEfQt1yDz8_XR6h1/view

Agenda

- Recurring item: anyone opposing this instance of the meeting being recorded?
- (dceara) What new features or relevant changes does the community think are important to have in the next release (24.03.0)?

- (dceara) Decide on a cadence for the technical meetings.

Notes

- 24.03.0 Planning
 - It's LTS - may want to do some hardening
 - Features we'd like to include?
 - (dceara) ACL I-P
 - Waiting for Numan's LogicalFlow I-P
 - (han) ACL is probably not very complex, but in general I'd suggest to first evaluate cost / benefit for new I-P patches
 - (dceara) OVN-K attempt to define a scale test for network policies to clarify how often ACLs are touched there
 - (ilya) a scale test for policies obviously will touch ACLs a lot - not a great representation of actual workloads?
 - (numans) it depends
 - (numan) LF I-P tests are almost ready
 - Scale testing with ovn-heater
 - (fnordahl) Canonical mainly focusing on quality / stability (cue to LTS), but...
 - A large OpenStack user hitting some issue with NAT FIPs, in routed topology
 - Will work on figuring out how to fix/alleviate in core OVN.
 - Another CMS - LXD
 - Interested in SmartNIC DPU support (OVN-VIF)
 - OVN as a standalone product?
 - DPUs as standalone firewall for bare metal
 - Database schema upgrade LTS to LTS may reveal issues too
 - (han) confirming schema updates are still a challenge since all nodes have to re-download the whole database
 - Relays to offload may be an option but looking for simpler deployment options
 - (han) How to make only the leader update the schema?
 - Script currently doesn't support it
 - (ilya) do you use conditional monitoring?
 - (han) yes
 - Unconditional may work better - Sends even a large database quite fast
 - There may be place for improvements for conditional
 - Lately it's been better?
 - (dceara) There [are](#) CI jobs for basic upgrade scenarios
 - Would be nice to expand to at least make sure proper LTS to LTS scenarios are covered

- (mark) Red Hat is also pushed into stability / quality direction these days
 - Looking for a better scheme to test OVN-K against OVN upstream CI - combined effort ongoing
 - Figuring out the list of unstable tests
 - Later - test against a stable hash of OVN-K
 - Goal - more **green** in OVN-K CI jobs
 - Also focus on “OVN proper” CI
 - No plan right now, but higher ups suggested the direction, so...
- (numan) it's LTS. Do we need a longer freeze?
 - (ilya) in OVS, a release becomes LTS **after** the next release (?)
 - -> extended period may not make sense due to that
 - (mark) if we make the change, it's possible
- Cadence
 - (mark) monthly or 6 weeks? Slot WFM
 - Keeping the same slot for now, cadence will be adjusted as needed
 - Dumitru will schedule the next in a month

Action items

- ☒ ~~Dumitru will schedule the next in a month~~