



OVN

Open vSwitch

December 8-9, 2020

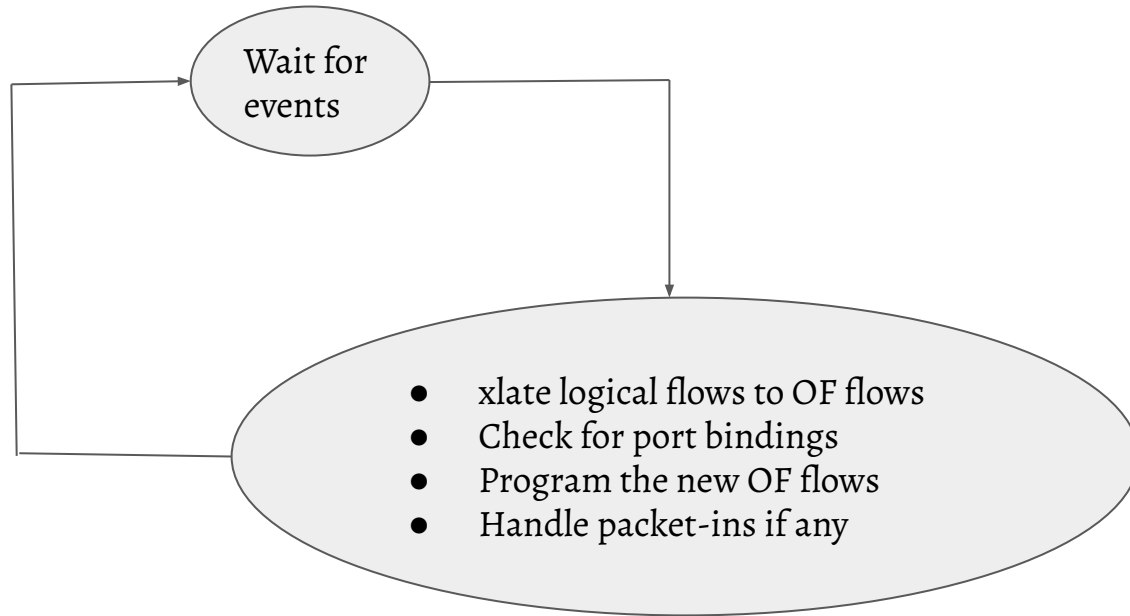
**OVN Controller improvements
in incremental processing**

Numan Siddique, Red Hat

What we'll be discussing today

- Incremental processing (I-P) changes in ovn-controller
 - 20.06 and 20.09
- Improvements

ovn-controller design



ovn-controller design

Before incremental processing (I-P)

- Main while loop which handles events
- In each run
 - Translates logical flows to OpenFlow rules - `lflow_run()`
 - Generate OF flows which connects physical-to-logical and logical-to-physical (physical flows)
 - `lflow_run()` is called even for `pinctrl` (packet-ins) events.

After incremental processing (I-P)

- Main while loop which handles events
- In each run
 - Translates only required logical flows to OpenFlow rules*.
 - `Pinctrl` events doesn't cause flow translation.

I-P Engine

- Has engine nodes.
- Each engine node has inputs. An engine node can be an input to another node.
- Has a handler for each input
- A handler processes the input
- If any handler can't process a change, a full recompute (`engine_run()`) is triggered for that node.

Limitations with initial IP support

- Initial I-P support handles some southbound db changes.
- Falls back to full recomputation of logical flows to OF flows for many changes it cannot handle.

Limitations with initial IP support

- Resulted in full recompute for
 - Port binding changes
 - OVS interface and Port changes
 - Datapath binding changes
 - Chassis addition/deletion.

Further improvements in OVN 20.06

- No more full recomputes for
 - Port binding changes
 - OVS interface and Port changes
 - Datapath binding changes
- Added input handlers to handle these changes.
- Resulted in faster processing and very less recomputations.
- Scales better.

Comparison - without I-P improvements

Rally task results

Task overview

Input file

▶ OvnFakeMultinode

▼ OvnNorthbound

create_routed_lport

▶ OvnNorthboundFakeMultin...

▶ OvnSandbox

OvnNorthbound.create_routed_lport (1,108.146s)

Overview Details Input task


Load duration: **1,106.410 s** Full duration: **1,108.146 s** Iterations: **100** Failures: **0**

Service-level agreement

Criterion	Detail	Success
max_seconds_per_iteration	Maximum seconds per iteration 12.22s <= 30.00s - Passed	True

Total durations

Action	Min (sec)	Median (sec)	90%ile (sec)	95%ile (sec)	Max (sec)	Avg (sec)	Success	Count
ovn.create_or_update_network_policy	1.02	1.266	1.522	1.552	2.053	1.284	100.0%	100
ovn.create_or_update_network_policy_address_sets	0.222	0.276	0.289	0.294	0.318	0.272	100.0%	100
ovn.create_port_group_acls	0.462	0.54	0.57	0.576	0.588	0.531	100.0%	50
ovn.create_or_update_name_space	0.346	0.476	0.554	0.569	0.595	0.47	100.0%	100
ovn_network.bind_port	1.235	1.349	1.404	1.425	1.45	1.351	100.0%	100
ovn.bind_ovs_vm	0.394	0.458	0.494	0.501	0.536	0.458	100.0%	100
ovn.bind_internal_vm	0.795	0.894	0.938	0.95	0.978	0.893	100.0%	100
ovn_network.wait_port_ping	7.388	7.669	7.774	7.83	8.6	7.682	100.0%	100
total	10.546	11.007	11.401	11.526	12.222	11.044	100.0%	100



Test uses existing OVN northbound db which has

- 56 logical switches
- 1256 lports
- 45447 logical flows.

Comparison - with I-P improvements

Rally task results

Task overview

Input file

- ▶ OvnFakeMultinode
- ▼ OvnNorthbound
 - create_routed_lport
 - ▶ OvnNorthboundFakeMultin...
 - ▶ OvnSandbox

OvnNorthbound.create_routed_lport (344.808s)

Overview Details Input task


Load duration: 343.190 s Full duration: 344.808 s Iterations: 100 Failures: 0

Service-level agreement

Criterion	Detail	Success
max_seconds_per_iteration	Maximum seconds per iteration 4.18s <= 30.00s - Passed	True

Total durations

Action	Min (sec)	Median (sec)	90%ile (sec)	95%ile (sec)	Max (sec)	Avg (sec)	Success	Count
ovn.create_or_update_network_policy	0.97	1.229	1.451	1.46	1.944	1.24	100.0%	100
ovn.create_or_update_network_policy_address_sets	0.224	0.267	0.28	0.282	0.291	0.266	100.0%	100
ovn.create_port_group_acls	0.462	0.526	0.539	0.544	0.564	0.52	100.0%	50
ovn.create_or_update_name_space	0.33	0.467	0.549	0.556	0.562	0.463	100.0%	100
ovn_network.bind_port	1.231	1.313	1.386	1.4	1.446	1.317	100.0%	100
ovn.bind_ovs_vm	0.392	0.436	0.464	0.475	0.502	0.438	100.0%	100
ovn.bind_internal_vm	0.816	0.876	0.936	0.967	0.997	0.879	100.0%	100
ovn_network.wait_port_ping	0.097	0.141	0.177	0.18	0.203	0.136	100.0%	100
total	2.977	3.406	3.709	3.748	4.181	3.411	100.0%	100



Test uses existing OVN northbound db which has

- 56 logical switches
- 1256 lports
- 45447 logical flows.

Further improvements ??

- We can try to further improve the I-P
- But this would make the code more complex.
- We did see few regressions when we added the improvements.
- `ovn-controller-ddlog`



Thanks!

Questions.