

StarlingX OpenStack Patch Elimination

Pilot Project Supported by the OpenStack Foundation 

A fully featured Cloud for the distributed Edge



Nova Must Haves – Address w/ planned upstream features

- **Live migration patch w/ pinning(`numa_topology`) – this one is table stakes**
 - Live migration has not worked since numa topology was introduced in kilo
 - We have been carrying complicated patches to address this
 - Replace current patches with <https://review.openstack.org/#/c/552722/>
- Shared/Dedicated vcpus on same host
 - Replace with <https://review.openstack.org/#/c/555081/>
- Shared vcpu extension
 - Replace with <https://review.openstack.org/#/c/555081/>
- vTPM support (stx community spec)
 - Replace with spec <https://review.openstack.org/#/c/571111/>
- HPET support (stx community spec)
 - Replace with spec <https://review.openstack.org/#/c/607989/>
 - Merged

Nova Must Haves – Address w/ planned upstream features

- VCPU model (stx community spec)
 - Current StarlingX implementation rejected by Nova
 - Replace with new proposed spec for Stein
- Disk performance fixes (stx community spec)
 - Specless blueprint
 - <https://blueprints.launchpad.net/nova/+spec/io-semaphore-for-concurrent-disk-ops>
- Extra spec validation (stx community spec)
 - <https://blueprints.launchpad.net/nova/+spec/flavor-extra-spec-image-property-validation>

Nova Must Haves – Address w/ upstream features already available

- Replace SR-IOV/PT best effort scheduling policy with Queens feature
 - <https://review.openstack.org/#/c/555000/3..3/specs/queens/implemented/share-pci-between-numa-nodes.rst>
- Replace vswitch affinity with Rocky feature
 - <https://review.openstack.org/#/c/541290/18/specs/rocky/approved/numa-aware-vswitches.rst>
- DB purge
 - Already addressed upstream in Rocky
<https://blueprints.launchpad.net/nova/+spec/purge-db>

Nova Must Haves – Address via Refactoring within STX

- Service create extension
 - Refactor in the VIM
- Low latency compute per-cpu power management
 - Do outside of nova
- RPC timeout handling
 - Leverage upstream enhancements
- Storage aggregate
 - Eliminate auto flavor extra-spec setting
- KVM crash detection
 - Implement with an external monitor
- PCI interrupt affinity
 - Do outside of nova

Nova Must Haves – Address TBD

- Providernet filter
- Live migration windows guest support
 - This did not work in Pike, retest and open bug if still broken

Nova Must Have - Bugs

- High priority bugs that need to be addressed, mainly around live migration
 - 493ae20 rbd disks: convert from source format to raw
 - 248ca72 Migrate code to use request_spec for server groups
 - d19a962 Specify migration type in db filter for anti-affinity migration
 - 1555d82 Do not bypass scheduler when using forced hosts
 - 3143ca8 Exclude current instance when checking server group live migrations
 - 3099f6f Live block migration fails following auth token expiration
 - 05c6b34 Drop support for forced hosts for live migrate
 - These need to be unsquashed:
 - d9f6f1c primary: server-groups best-effort and group size
 - 963dc06 primary: live migration driver fixes
 - 58e885e Fix stale RequestSpec instance numa topology for live-migration

Neutron Must Haves – Address w/ planned upstream features

- Network Segment Range Management (stx community spec)
 - <https://review.openstack.org/#/c/599980>
- DNAT
 - Already available upstream

Neutron Must Haves – Address w/ planned upstream bug fixes/enhancements

- 17 items with the following already in review
 - Populate binding levels when concurrent ops fail
 - <https://review.openstack.org/#/c/606827/>
 - [RFE] Enable other subprojects to extend l2pop fdb information
 - <https://review.openstack.org/#/c/599319/>
 - Enable delete bound trunk for linux bridge agent
 - <https://review.openstack.org/#/c/605589/>
 - Prevent DHCP from processing stale RPC messages
 - <https://review.openstack.org/#/c/609463/>

Neutron Must Haves – Address w/ planned upstream bug fixes/enhancements (cont'd)

Neutron

- 0005-pnet-extension-to-expose-providernet-management-at-a.patch 021ae1a
- 0018-patching-scripts-for-neutron-processes.patch 87a8c625
- 0037-add-process-synchronization-around-segment.patch e490b5a
- 0039-dhcp-handle-concurrent-port-creation-error.patch a38f89
- 0040-rpc-removing-timeout-backoff-multiplier.patch 2e4ca00
- 0047-wsgi-prevent-accepting-socket-without-a-gr.patch 8e72491
- 0062-dvr-force-admin-state-update-before-distri.patch edf9731
- 0065-dvr-do-not-create-agent-gateway-ports-unle.patch 2857911
- 0077-api-reject-routes-with-invalid-network-val.patch 6ed8fe6
- 0081-l2pop-added-rpc-query-api-for-resync.patch 88b7bc7
- 0092-l3-add-l2pop-support-for-floatingip-resources.patch 9f926a5

Client

- 0009-allow-admin-password-changes-on-controller.patch c3fa9d9b3303433f09ca16c089b2db52e116dfba
- 0014-add-support-for-querying-quotas-with-usage.patch 71c07d7b22bb5aaea8c39d4d4b97883a0bae5c1b

Neutron Must Haves – Address via refactoring within STX

- Host State Management
- DHCP agent rescheduling / rebalancing
- L3 agent rescheduling / rebalancing
- Patching script rework
- Modeling the provider networks in sysinv

Horizon Must Haves - Address w/ new upstream features

- These are specs that would need to be created and worked through the horizon community
 - Auth plugin extensions
 - Configurable auto page refresh
 - Loading bars
 - Custom Confirmation Messages
 - CSRF AngularJS fixes

Horizon Must Haves - Address w/ refactoring within STX

- Move the following to stx-gui after moving to master
 - Port Forwarding tab
 - Controller services tab
- The following can be moved to stx-gui now
 - Branding extensions
- Refactor FM panel to use angularjs
 - Retires Custom 'fixedwithquery' filter/search

Other Projects Must Haves

- Keystone/Ironic/Magnum/Ceilometer
 - Align w/ master
 - No/minimal loss of functionality vs. master
- Murano
 - Deprecate support in StarlingX
- Heat
 - Align w/ master
 - Most of the remaining patches are in support of other project extensions which are being deprecated
- Glance/Cinder
 - Replace cache-raw extension with upstream cinder caching feature
 - Outside of that no significant loss of functionality aligning w/ master

Strategy

High Level Implementation Strategy

- Rebasing to master would be on containerized openstack only
 - Do not want to start the migration on bare metal openstack as it would require rebasing the openstack puppet modules and be thrown away shortly thereafter
 - **Note we will have to do bare metal keystone/horizon at some point before the end of the release**
- Create parallel development tracks
 - Track 1: Continue the container integration with STX-Pike branches
 - Track 2: Start integrating containers based on master builds
 - STX feature branches required
 - Abandon track 1 when track 2 reaches equivalent functionality/stability from an infrastructure perspective
- Cutover STX master to container branches once sufficient stability has been achieved
 - At cutover time openstack master will be missing some required STX release content
 - Incrementally integrated during the remainder of the release

High Level Implementation Strategy (cont'd)

- Track 2
 - Need to build openstack images off upstream master branches
 - Will need to build any of the dependency services (i.e. rabbit, memcached) if upstream requires a newer version
 - Need to be launched immediately
 - Integrate HELM charts with mater images
 - Implement any required HELM chart changes
 - The following critical refactoring needs to be accelerated as it is the bare minimum to move from track 2
 - Neutron host state management moved to the STX-NFV
 - Nova service create moved to the STX-NFV
 - Modeling the provider networks in sysinv
 - Need Neutron Spec implemented to cutover to track 2

High Level Implementation Strategy (cont'd)

- Remaining refactoring can go in parallel with other activities

Key Stein Dates

- Milestone2: Jan11
- Milestone3: Mar8
- RC1: Mar22
- Release: Apr12
- ** MS2 is spec freeze for nova
- ** MS3: Feature Freeze

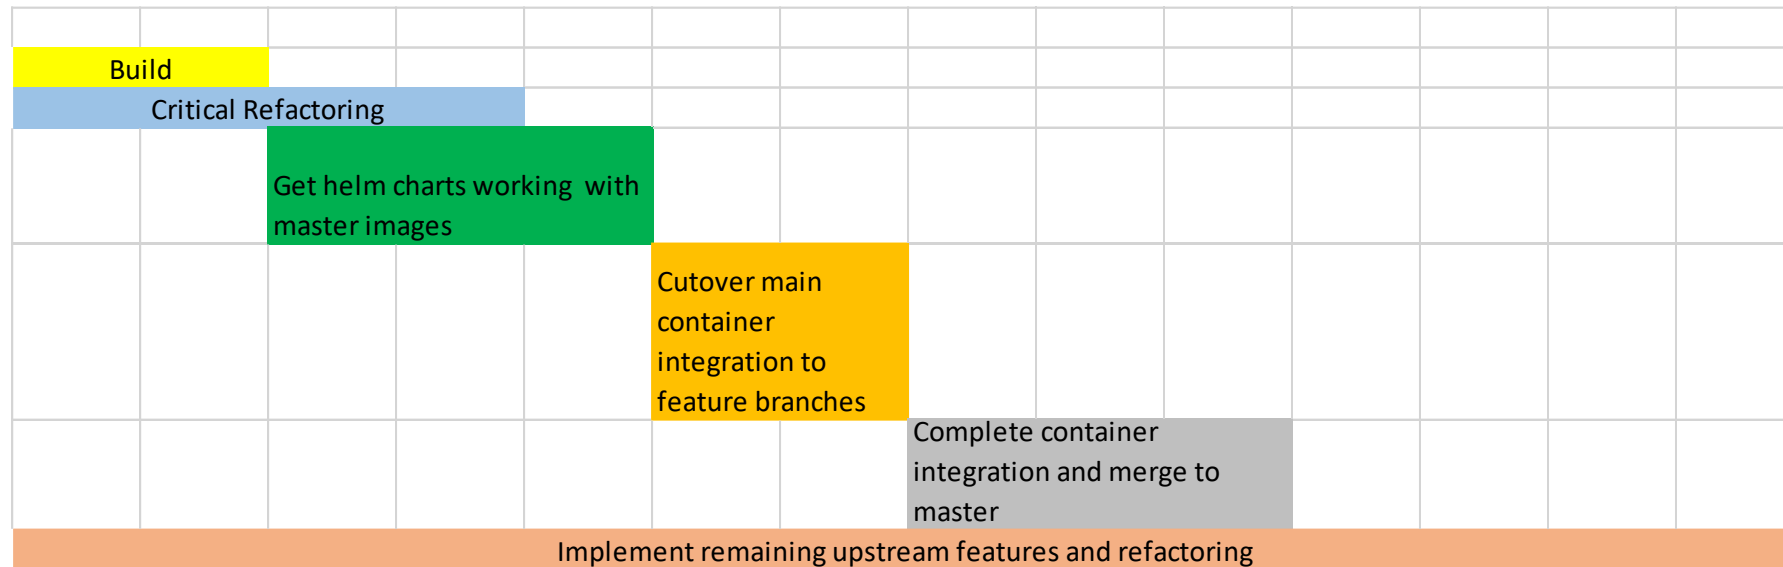
Key Next Steps

- Required feature branches are created
- Create master container builds
 - Includes associated build system changes
- Start helm chart integration w/ master container builds

- Commitment and delivery date from stakeholders for required cut-over content
 - Refactoring and openstack feature deliverables
 - Once this is provided a date to cutover to track 2 can be determined

- Commitment and delivery date from stakeholders for release gating content
 - All content identified is considered release gating

High Level Sequencing



- Sequencing view only, not a time scale