

Anuket RA1- OpenStack Releases Highlights

Ussuri, Victoria and Wallaby Releases Highlights- Significant changes

<https://releases.openstack.org/ussuri/highlights.html>

<https://releases.openstack.org/victoria/highlights.html>

<https://releases.openstack.org/wallaby/highlights.html>

O p e n S t a c k S e r v i c e s	Ussuri	Victoria	Wallaby
A o d h			
B a r b i c a n			
B l a z a r			Introduced a framework for enforcing operator-defined limits on reservation usage.
C i n d e r	<ul style="list-style-type: none"> Numerous improvements in current functionality, for example, the ability to set minimum and maximum sizes for volume-types; the ability to filter the volume list using time comparison operators. Support for Glance multistore and image data colocation when uploading a volume to the Image Service. Added some new backend drivers, and many current drivers have added support for more features. 	<ul style="list-style-type: none"> Improved handling around the configured default volume-type and added new Block Storage API calls with microversion 3.62 that enable setting a project-level default volume-type for individual projects. Added some new backend drivers, and many current drivers have added support for more features. For example, the NFS driver now supports volume encryption. Support was added to cinder backup to use the popular Zstandard compression algorithm. 	<ul style="list-style-type: none"> Block Storage API microversions 3.63 and 3.64 add useful information to the volume detail response (volume_type_id and encryption_key_id, respectively). Added new backend drivers: Ceph iSCSI, Dell EMC PowerVault ME, KIOXIA Kumoscale, Open-E JovianDSS, and TOYOU ACS5000. Additionally, many current drivers have added support for features exceeding the required driver functions, with revert to snapshot and backend QoS being particularly popular this cycle. Added a new backup driver for S3-compatible storage. Cinder now stores the format of the backing file (raw or qcow2), for FileSystem backends, in the volume admin metadata and includes the format in the connection_info returned in the Attachments API.

Cyborg	<ul style="list-style-type: none"> Users can now launch instances with accelerators managed by Cyborg, as the Nova-Cyborg integration has been completed. See accelerator operation guide to find which instance operations are supported. New APIs have been implemented to list devices managed by Cyborg and, in general, to view and manage inventory of accelerators. Cyborg has laid the foundations for providing backwards compatibility in future releases by adopting microversions in v2 API. The Cyborg client is now based on OpenStack SDK and supports most Version 2 APIs. Improved quality overall by adding more unit/functional tests and by reducing technical debt. 	<ul style="list-style-type: none"> Users can launch instances with accelerators managed by Cyborg since Ussuri release, this release two more operations * Rebuild and * Evacuate are supported. See accelerator operation guide to find all supported operations. Cyborg supported new accelerator drivers (Intel QAT and Inspur FPGA) and reached an agreement that Vendors who want to implement a new driver should at least provide a full driver report result. (Of course, providing third-party CI is more welcome.) <i>Supported drivers</i> <https://docs.openstack.org/cyborg/latest/reference/support-matrix.html>_ Program API is supported, now users can program FPGA given the pre-uploaded bitstream. <i>program API (PATCH deployable)</i> <https://specs.openstack.org/openstack/cyborg-specs/specs/train/approved/cyborg-api.html>_ And API microversion for existed APIs is improved such as arq APIs. In this release, the policy refresh (RBAC with scoped) for cyborg is partially implemented (Device Profile APIs), we've implemented new default rules in base policy and device_profile policy, and added the basic testing framework for all policies. For the Backward Compatibility, old rules are maintained as deprecated rules with same defaults as today so that existing deployment will keep working as it is. After we implement all the features, we'll give two cycles transition period for operators. See policy default refresh JSON formatted policy file is deprecated; YAML format to be used instead. 	<ul style="list-style-type: none"> Users can launch instances with accelerators managed by Cyborg since Ussuri release, this release more operations such as Shelve/Unshelve are supported. See accelerator operation guide to find all supported operations. Cyborg introduces more new accelerator drivers such as Intel NIC and Inspur NVMe SSD driver which allow user to boot up a VM with such device attached. Cyborg now provides a new configuration for user to configure their devices, for example, user can indicate the vgpu type for their virtualized GPU, user can indicate the specific functions loaded on their NIC, etc. The development of the RBAC policy started in Victoria continued. Support for seven roles, and new defaults for device profiles.
Designate			<ul style="list-style-type: none"> Added the NS1 DNS backend driver for Designate. Designate now supports the Keystone default roles and scoped tokens.
EC2 API			
Freezer			
Glance	<ul style="list-style-type: none"> Enhancement in multiple stores feature, users now can import single image in multiple stores, copy existing image in multiple stores and delete image from single store. New import plugin to decompress the image Introduced S3 driver for glance-store again 	<ul style="list-style-type: none"> Enhancement in multiple stores feature, administrator can now set policy to allow user to copy images owned by other tenants Glance allow to configure cinder multi-stores RBD and Filesystem drivers of glance now support sparse image upload Enhancement in RBD driver chunk upload of image 	<ul style="list-style-type: none"> New API <code>/v2/images/<image-id>/tasks</code> to get tasks associated with image Support for distributed image import Glance's default policies now use the <i>member</i> role on projects to protect writeable and readable image actions. Support was also added for read-only access to image resources when the <i>reader</i> role is granted to users on a project. Administrative operations, like creating public images, is still protected using the <i>admin</i> role on a project. Administrative actions will be updated in the future to consume system-scope. Secure RBAC - Experimental support for project personas Cleanup of stale files in staging area upon service startup
Gnocchi			

Horizon	<ul style="list-style-type: none"> This release mainly focuses on bug fixes and improvements from the maintenance perspective including deprecations of old features, cleanup of deprecated features, integration test coverage improvements, migration to mock usage in unit tests and so on. Horizon and all horizon plugins now support Django 2.2 which is the only supported LTS of Django. Django is a framework which horizon depends on. Note that python 2.7 is no longer supported and we have enter into the python3 era. A couple of feature gaps in keystone support are implemented: a feature to allow users to change expired password including first login, password lock option in the user panel, and a support of access rules for application credentials. 	<ul style="list-style-type: none"> Error messages shown in horizon now contains more detail. Previously GUI users cannot know detail reasons of operations, but users can now check detailed information from back-end service so that they can address causes. Added a new tab that shows messages for volumes and volume snapshots. Users can know detail events which happend for corresponding volumes or snapshots. Added support for extending in-use volumes. Users can extend in-use volumes via horizon now. 	<ul style="list-style-type: none"> Horizon supports the registered default policies. Now operators do not need to define all the policies in the policy file instead define only those policy they would like to override. Chinese locales zh-cn and zh-tw have been changed to zh-hans and zh-hant respectively following the change in Django because the new locales decouple what is spoken from specific locations as they are also used outside of China. Added Volume backups support for admin.
Ironic	<ul style="list-style-type: none"> Support for scoped introspection rules which allow to have (and keep) rules per node subsets, such as different hardware deliveries. Support for a hardware retirement workflow to enable automation of hardware decommission in managed clouds. Multitenancy concepts and additional policy options are available for non-administrator usage of Ironic. Addition of authentication of interactions between Ironic and its remote agent enabling deployment over untrusted networks. UEFI and device selection is now available for Software RAID. 	<ul style="list-style-type: none"> The deploy steps work has decomposed the basic deployment operation into multiple steps which can now also include steps from supported RAID and BIOS interfaces at the time of deploy. An agent power interface enables provisioning operations without a Baseboard Management Controller. Ironic can now be configured for HTTP Basic authentication without the need for additional services. Adds initial support for DHCP-less based deployments with Redfish Virtual Media. 	<ul style="list-style-type: none"> Redfish capability enhancements covering Out of Band hardware RAID configuration, and automatic Secure Boot setting enablement. Deployment and Cleaning enhancements including UEFI Partition Image handling, NVMe Secure Erase, per-instance deployment driver interface overrides, deploy time "deploy_steps", and file injection. The System scoped RBAC model is now supported by Ironic along with the admin, member, and reader roles. This work has resulted in over 1500 new unit tests being added to Ironic.
Keystone	<ul style="list-style-type: none"> The user experience for creating application credentials and trusts has been greatly improved when using a federated authentication method. Federated users whose role assignments come from mapped group membership will have those group memberships persisted for a configurable TTL after their token expires, during which time their application credentials will remain valid. Keystone to Keystone assertions now contain the user's group memberships on the keystone Identity Provider which can be mapped to group membership on the keystone Service Provider. Federated users can now be given concrete role assignments without relying on the mapping API by allowing federated users to be created directly in keystone and linked to their Identity Provider. When bootstrapping a new keystone deployment, the admin role now defaults to having the "immutable" option set, which prevents it from being accidentally deleted or modified unless the "immutable" option is deliberately removed. Keystonemiddleware no longer supports the Identity v2.0 API, which was removed from keystone in previous release cycles. 		

K o l l a	<ul style="list-style-type: none"> • All images, scripts and Ansible playbooks now use Python 3, and support for Python 2 has been dropped. • Added support for CentOS 8 hosts and images. • Added initial support for TLS encryption of backend API services, providing end-to-end encryption of API traffic. Currently Barbican, Cinder, Glance, Heat, Horizon, Keystone, Nova and Placement are supported. • Added support for deployment of Open Virtual Network (OVN) and integration of it with Neutron. • Added support for deployment of Zun CNI (Container Networking Interface) components allowing Docker with containerd to support Zun capsules (pods). • Added support for Elasticsearch Curator to help manage clustered log data. • Added components necessary to use Mellanox networking devices with Neutron. • Streamlined configuration of external Ceph integration, making it easy to go from Ceph-Ansible-deployed Ceph cluster to enabling it in OpenStack. 	<ul style="list-style-type: none"> • Added support for Ubuntu Focal 20.04. • Added support for automatic creation of resources for Octavia. • Added support for container healthchecks for core OpenStack services. • Improved TLS support, covering etcd, RabbitMQ, as well as Ironic, Neutron and Nova backends. Also adds initial support for ACME protocol, as used by Letsencrypt. • Improved performance and scalability of Ansible playbooks. • Added support for integrating Neutron with Mellanox InfiniBand. 	<ul style="list-style-type: none"> • Switched CentOS images to CentOS Stream 8. • Added support for Ubuntu in Kayobe. • Added support for the OpenID Connect authentication protocol in Keystone. • Added Docker healthchecks for several services. • Added support for Prometheus version 2. • Added support for multiple environments in a single Kayobe configuration.
M a g n u m	<ul style="list-style-type: none"> • Support Helm v3 to install all magnum installed charts. Support for Helm v2 client will be removed in X release. • A new config option <code>post_install_manifest_url</code> is added to support installing cloud provider/vendor specific manifest after deploying a kubernetes cluster. • A new <code>--merge-labels</code> boolean flag can be used to merge user labels at cluster/nodegroup scope with cluster template/cluster labels. • Cloud admin users now can do rolling upgrade on behalf of end users to do urgent security patching. • Magnum now cascade deletes all the load balancers before deleting the cluster, not only including load balancers for the cluster services and ingresses, but also those for Kubernetes API/etcd endpoints. • Magnum supports updating the k8s cluster health status via the Magnum cluster update API so that a controller (e.g. magnum-auto-healer) running inside the k8s cluster can call the Magnum update API to update the cluster health status. 	<ul style="list-style-type: none"> • Kubernetes cluster owner can now do CA cert rotate to re-generate CA of the cluster, service account keys and the certs of all nodes. • Label <code>cinder_csi_enabled</code> now defaults to True. • Default storage driver has changed from <code>devicemapper</code> to <code>overlay2</code>. 	<ul style="list-style-type: none"> • Add tags in cluster templates to help operators advertise features in their public cluster templates. • Update versions for kubernetes, containerd and addons.

Manila	<ul style="list-style-type: none"> Share groups have graduated from being an experimental feature to being generally available. Starting with API version 2.55, the <i>X-OpenStack-Manila-API-Experimental</i> header is no longer required to create/update/delete share group types, group specifications, group quotas and share groups themselves. Shares can be created from snapshots across storage pools when compatible. This new feature allows better utilization of back end resources by spreading workloads that were previously confined to the back end that hosted the snapshot. New quota control mechanisms have been introduced to constrain projects and their users to the number and size of share replicas they can create. It is now possible to query asynchronous user messages with time intervals. 	<ul style="list-style-type: none"> Tenant driven share replication, a self-service aid to data protection, disaster recovery and high availability is now generally available and fully supported. Starting with API version 2.56, the <i>X-OpenStack-Manila-API-Experimental</i> header is no longer required to create/promote/resync/delete share replicas. Share server migration is now available as an experimental feature. Share servers provide hard multi-tenancy guarantees by isolating shared file systems in the network path. In this release, cloud administrators are able to move share servers to different backends or share networks. 	<ul style="list-style-type: none"> OSProfiler support has been added for tracing and observability. Users may now add and update security services on share networks that are in use. Operators may now set maximum and minimum share sizes as extra specifications on share types. It is also possible to limit the maximum size of shares via project and share type quotas. The number and size of shares can be limited on share servers for load balancing. The service provided default RBAC policies for all API endpoints have been adjusted to accommodate system scoped and project scoped personas with admin, member and reader roles where appropriate. The service now supports a healthcheck middleware that is enabled by default. Several driver improvements have been committed. The Container share driver now supports user defined LDAP security services that can be added to share networks or modified at any time. The NetApp driver supports setting up FPolicy events on shares. It also now allows users to add/update Kerberos, LDAP or Active Directory security services on their share networks at any time. The CephFS driver has been refactored to interact with the ceph manager daemon to create and manage shares. It also supports efficiently cloning snapshots into new shares. A new share driver has been added for Zadara Cloud Storage and supports NFS and CIFS protocols.
Masakari		<p>Adds ability for operators to override, per failure type, the instance metadata key controlling the behaviour of Masakari towards the instance. This makes it possible to differentiate between instance- and host-level failures per instance.</p>	<ul style="list-style-type: none"> Support for disabling and enabling failover segments. This way operators are able to put whole segments into maintenance mode instead of having to do it for each single host. Support for smoothing-out the decision about whether to consider a host down or not. Operators can configure host monitors to consider a chosen number of probes before sending the notification about host being down. Support for running host monitors in environments without systemd, such as app containers. Support for using system-scoped tokens when contacting Nova.
Mistral			
Murano			

<p>Neutron</p> <ul style="list-style-type: none"> The OVN driver is now merged into Neutron repository and is one of the in-tree Neutron ML2 drivers, like linuxbridge or openvswitch. OVN driver benefits over the openvswitch driver include for example DVR with distributed SNAT traffic, distributed DHCP and possibility to run without network nodes. Other ML2 drivers are still in-tree and are fully supported. Currently default agent is still openvswitch but our plan is to make OVN driver to be the default choice in the future. Support for stateless security groups has been added. Users can now create security group set as stateless which means that conntrack will not be used for any rule in that group. One port can only use stateless or stateful security groups. In some use cases stateless security groups will allow operator to choose for optimized datapath performance whereas stateful security groups impose extra processing on the system. Role Based Access Control (RBAC) for address scopes and subnet pools has been added. Address scopes and subnet pools are usually defined by operators and exposed to users. This change allows operators to use more granular access controls on address scopes and subnet pools. Support for tagging resources during creation has been added in Neutron API. User can now set tags for resources like e.g. ports directly in POST requests. This will improve the performance of kubernetes network operations a lot. The number of API calls which e.g. Kuryr has to send to Neutron are greatly reduced. 	<ul style="list-style-type: none"> Metadata service is now available over IPv6. Users can now use metadata service without config drive in IPv6-only networks. Support for flat networks has been added for Distributed Virtual Routers (DVR). Support for Floating IP port forwarding has been added for the OVN backend. Users can now create port forwardings for Floating IPs when the OVN backend is used in Neutron. Added support for router availability zones in OVN. The OVN driver can now read from the router's <code>availability_zone_hints</code> field and schedule router ports accordingly with the given availability zones. 	<ul style="list-style-type: none"> New subnet type <code>network:routed</code> is now available. IPs of such subnet can be advertised with BGP over a provider network. This basically achieves a BGP-to-the-rack feature, where the L2 connectivity can be confined to a rack only, and all external routing is done by the switches, using BGP. In this mode, it is still possible to use VXLAN connectivity between the compute nodes, and only floating IPs and router gateways are using BGP routing. Now it is possible to define a gateway IP when creating a subnet using a subnet pool. If the gateway IP can be allocated in one of the subnet pool available subnets, this subnet is created; otherwise a <code>Conflict</code> exception is raised. A port already bound with a QoS <code>minimum_bandwidth</code> rule can now be updated with a new QoS policy with a <code>minimum_bandwidth</code> rule. It will change the allocations in placement as well. A new <code>vnictype</code> <code>vdpa</code> has been added to allow requesting port that utilize a <code>vHost-vDPA</code> offload. It is supported by ML2/OVS and ML2/OVN mech drivers currently. Deletion of the ML2/OVN agents is now supported. New resource address-groups can be used in the security group rules to add group of the IP addresses to the rule. The OVN Octavia provider driver now supports Stream Control Transmission Protocol (SCTP) load balancing. Better co-existence with floating IP port forwarding load balancers. Introduce the attribute <code>port_device_profile</code> to ports that specifies the device profile needed per port. This parameter is a string. This parameter is passed to Nova and Nova retrieves the requested profile from Cyborg: Device profiles. Operators can turn on this feature via the ML2 extension_drivers configuration option. Fixed a number of bugs so we better reflect load balancer status via the Octavia API.
---	--	---

Nova	<ul style="list-style-type: none"> Support for cold migrating and resizing servers between Nova cells. Support for precaching Glance images to Nova compute hosts. Support for creating servers with accelerator devices via Cyborg. Further enhanced support for moving servers with minimum bandwidth guarantees. Support for nova-manage placement audit CLI to find and clean up orphaned resource allocations. Nova API policies are introducing new default roles with <code>scope_type</code> capabilities. These new changes improve the security level and manageability. New policies are richer in terms of handling access at system and project level token with 'Read' & 'Write' roles. This feature is disabled by default and can be enabled by config options. See the Policy Concepts documentation for more details. Improved robustness for cases where high levels of concurrent allocation writes are common, such as a busy clustered hypervisor, by making allocation retry count configurable. 	<ul style="list-style-type: none"> Nova supports mixing pinned and floating CPUs within the same nova server. Nova supports customizing the placement resource inventory of the compute node via a provider configuration file. Nova supports fast cloning of Glance images from the Ceph RBD cluster even if Glance multistore configuration is used. Nova supports creating servers with virtual TPM devices. 	<ul style="list-style-type: none"> Support for accelerators in Nova servers has been improved. Now shelving and unshelving such server is supported. Now Nova supports attaching neutron ports with QoS minimum bandwidth rules for running servers. The Nova scheduler can now ensure that servers with the requested networks or ports related to Neutron routed networks are scheduled to compute hosts where network segments are available. The Hyper-V virt driver can now attach Cinder RBD volumes. The libvirt driver now support changing the default machine type on a compute node safely The libvirt driver now supports UEFI Secure Boot. The libvirt driver now supports vDPA (vHost data path acceleration), a vendor neutral way to accelerate standard virtio device using software or hardware accelerator implementations. A new image metadata property, <code>hw_input_bus</code>, has been added. This allows you to specify the bus used for input devices - a pointer and keyboard - which are attached to the instance when graphics are enabled on compute nodes using the libvirt virt driver. Two values are currently accepted: <code>usb</code> and <code>virtio</code>. This image metadata property effectively replaced the <code>hw_pointer_model</code> image metadata property, which is nonetheless retained for backwards compatibility purposes. Added IP addresses to the metadata in libvirt XML. If an instance has more than one IP address, enumerate those IP addresses. The port attach or detach is performed dynamically after the creation of the instance. Every time there is a change, it is reflected in the contents of the XML.
Octavia	<ul style="list-style-type: none"> Octavia now supports deploying load balancers in specific availability zones. This allows the deployment of load balancing capabilities to edge environments. The Octavia amphora driver has added a technology preview feature that improves control plane resiliency. Should a control plane host go down during a load balancer provisioning operation, an alternate controller can resume the in-process provisioning and complete the request. Users can now specify the TLS ciphers acceptable for listeners and pools. This allows load balancers to enforce security compliance requirements. 	<ul style="list-style-type: none"> Users can now specify the TLS versions accepted for listeners and pools. Operators also now have the ability to set a minimum TLS version acceptable for their deployment. Octavia now supports HTTP/2 over TLS using the new Application Layer Protocol Negotiation (ALPN) configuration option for listeners. Load balancer statistics can now be reported to multiple statistics drivers simultaneously and supports delta metrics. This allows easier integration into external metrics system, such as a time series database. Octavia flavors for the amphora driver now support specifying the glance image tag as part of the flavor. This allows the operator to define Octavia flavors that boot alternate amphora images. Load balancer pools now support version two of the PROXY protocol. This allows passing client information to member servers when using TCP protocols. PROXYV2 improves the performance of establishing new connections using the PROXY protocol to member servers, especially when the listener is using IPv6. 	<ul style="list-style-type: none"> With the addition of ALPN and HTTP/2 support for backend pool members, Octavia now supports the gRPC protocol. gRPC enables bidirectional streaming of Protocol Buffer messages through the load balancer. Octavia now supports Stream Control Transmission Protocol (SCTP) load balancing. The addition of SCTP enables new mobile, telephony, and multimedia use cases for Octavia. Load balancers using the amphora provider will benefit from increased performance and scalability when using amphora images built with version 2.x of the HAProxy load balancing engine. Amphora instances are now supported on AArch64 /ARM64 based instances. Octavia now supports the Keystone default roles and scoped tokens.
Openstack Ansible	<ul style="list-style-type: none"> Ceph Octopus support MariaDB upgraded to 10.4 release Added Centos 8 support Added Ubuntu Focal support 	<ul style="list-style-type: none"> MariaDB upgraded to 10.5 release Ansible bumped to 2.10 release and switched to collections usage Added <code>os_senlin</code> role Added <code>os_adjutant</code> role 	<ul style="list-style-type: none"> Significantly improved Zun role and moved from experimental to stable status Exerimental support for CentOS Stream Experimental support for Debian Bullseye Self-signed SSL will be generated and signed with local Certificate Authority

P l a c e m e n t			<ul style="list-style-type: none"> Added IP addresses to the metadata in libvirt XML. If an instance has more than one IP address, enumerate those IP addresses. The port attach or detach is performed dynamically after the creation of the instance. Every time there is a change, it is reflected in the contents of the XML. If you do not override policy with custom rules you will have nothing to do. If you do override the placement default policy then you will need to update your configuration to use the standard <code>[oslo_policy] /policy_file</code> config option.
S a h a r a			
S e n l i n			
S o l u m			
S w i f t	<ul style="list-style-type: none"> Added a new system-namespace for Swift containers and objects. Added a new Swift object-versioning API using the new namespace. Added support for S3 versioning using the new API. Added the ability to use SIGUSR1 to perform "seamless" reloads, where the WSGI server socket never stops accepting connections. 	<ul style="list-style-type: none"> Improved time-to-first-byte latencies when reading erasure-coded data. Increased isolation between background daemons and proxy-servers when running with a separate replication network. We're beginning to see non-trivial production clusters transition from running Swift under Python 2 to Python 3. 	<ul style="list-style-type: none"> Static large object segments can now be deleted asynchronously; multipart uploads deleted through the S3 API will always be deleted asynchronously. Numerous sharding improvements, including the ability to cache shard ranges for listings and support for operator-driven shrinking. Several part-power-increase improvements, which ensure small clusters are capable of growing to be large clusters.
T r o v e			<ul style="list-style-type: none"> Support image tags for the datastore version. When using image tags, Trove is able to get the image dynamically from Glance for creating instances. Added custom container registry configuration for trove guest agent, it's now possible to use images in private registry rather than docker hub. Added a new field <code>operating_status</code> for the instance to show the actual operational status of user's database. In multi-region deployment with geo-replicated Swift, the user can restore a backup in one region by manually specifying the original backup data location created in another region.
Z a q u a r	<ul style="list-style-type: none"> Support querying queues with a 'with_count' to return the amount of the queues. Help users to quickly get the exact total number of queues which they own. Introduce new resource called Topic which is a concept from SNS. User can send message to a Topic and then subscribers will get the message according to different protocols like http, email, sms, etc. 		
Z u n	Starting from this release, Zun adds support for CRI-compatible runtime. Zun uses CRI runtime to realize the concept of capsule (pod). As a result, users can use Zun API to create pods in Kata container via a CRI runtime.		<ul style="list-style-type: none"> Introduce the python-binding for interacting with CRI runtime via GRPC Introduce CNI plugin for container network

OpenStack Operations tooling	Ussuri	Victoria	Walla by
Adjutant			
Ceilometer			

CloudKitty		<ul style="list-style-type: none"> After a period of inactivity, development has been resumed by a group of new contributors. Introduced a Monasca fetcher, to gather scopes to be rated from Monasca. 	
Monasca			
Panko			
Patrole			
Rally			
Tempest			
Vitrage	Added more concise and friendly Template Version 3 syntax .	<ul style="list-style-type: none"> Add new datasource for TMF API 639 Datasource. Complete verification to verify the Vitrage API 	
Watcher	<ul style="list-style-type: none"> Added a new webhook API and a new audit type EVENT. Now Watcher user can create audit with EVENT type and the audit will be triggered by webhook API. The building of the compute (Nova) data model will be done using the decision engine threadpool, thereby, significantly reducing the total time required to build it. 		

OpenStack Integration enablers	Ussuri	Victoria	Wallaby
Kuryr	<ul style="list-style-type: none"> Support for IPv6. DPDK support for nested setups and various other DPDK and SR-IOV improvements. Multiple fixes related to NetworkPolicy support. 	<ul style="list-style-type: none"> Kuryr will no longer use annotations to store data about OpenStack objects in K8s API. Instead a corresponding CRDs are created, i.e. KuryrPort, KuryrLoadBalancer and KuryrNetworkPolicy. Logs on INFO level should be much cleaner now. Added support for autodetection of VM bridging interface in nested setups. 	<ul style="list-style-type: none"> Nested mode with nodes VMs running in multiple subnets is now available. To use that functionality a new option <code>[pod_vif_nested] worker_nodes_subnets</code> is introduced accepting multiple Subnet IDs. Kuryr now handles Services that do not define the <code>.spec.selector</code>, allowing the user to manually manage the Endpoints object. Kuryr can handle egress Network Policy that allows traffic to Pods being pointed by a Service without Selector. Added support for SCTP. Networks can now be created by relying on the default MTU defined in Neutron, regardless of the SDN used and without changing the default configuration value in Kuryr.
Tacker		<ul style="list-style-type: none"> Implement ETSI NFV-SOL standard features (Life-cycle management, Scaling, VNF operation, etc.). Add Fenix plugin for Rolling update for VNFs with Fenix and Heat. Expand Kubernetes support. 	<ul style="list-style-type: none"> Add APIs for scale, update, and rollback operations for VNF defined in ETSI NFV. Add fundamental VNF lifecycle management support for subscriptions and notifications defined in ETSI NFV. Implement VNF package management interface to obtain VNF package, grant interface to allow the VNFM to request a grant for authorization of a VNF lifecycle operation defined in ETSI NFV SOL003 specification compliant operations to cooperate with 3rd-Party NFVOs as VNFM. Add container based VNF support with ETSI NFV-SOL003 v2.6.1 VNF Lifecycle Management. User is able to create, instantiate, terminate, and delete VNF on Kubernetes VIM. Kubernetes resource files are available as VNFD and it's uploaded as a part of VNF Package. Enable VNF vendors to customize configuration methods for applications via MgmtDriver. These customizations are specified by "interface" definition in ETSI NFV-SOL001 v2.6.1.

OpenStack Deployment tools	Ussuri	Victoria	Wallaby

Tripleo		<ul style="list-style-type: none"> · Moving network and network port creation out of the Heat stack and into the baremetal provisioning workflow. · Ceph version upgraded to Pacific. cephadm may be used to deploy/maintain a Ceph RBD cluster but not all Ceph services (e.g. RGW). ceph-ansible may still be used to deploy/maintain all Ceph services but will be replaced with cephadm in next release. This work is described in the TripleO Ceph spec and the Tripleo Ceph Client spec. · Removed Swift from the Undercloud services and removed the deployment 'plan' as described in the Excise swift spec. · Early (beta) support for deploying FRRouter in the Overcloud to support BGP routing as described in the FR Router spec. <p>Moving away from using a dedicated Heat service on the Undercloud for the Overcloud deployment and instead using Ephemeral Heat.</p>
OpenStack-Helm		
Kolla-Ansible	Kolla Ansible backend TLS ensures OpenStack API traffic is encrypted end-to-end.	
OpenStack-Ansible		
OpenStack-Charms		
Bifrost		
OpenStack-Chef		
LOCI		
Puppet-OpenStack	Puppet OpenStack can now bootstrap Keystone using an admin password instead of using the legacy admin token.	
RPM-Packaging		