

# Lakelse Release Highlights

Project, Workstream, or committee	Best contact for followup questions  (Use @ notation)	What has changed?	Why is this significant/beneficial to a user/consumer of Anuket?	Notes, links to more information, etc.
Barometer	<a href="#">Emma Foley</a>	<h2>Release Summary</h2> <p>Added unixsock plugin to one-click install.</p> <p>Add ansible playbook for building the containers locally.</p> <p>Since the anuket dockerhub repository was created, and containers are being pushed to there, instructions and build scripts have been updated to reflect this.</p> <p>Testing playbooks were added to compare collectd5 vs collectd6, for the purpose of helping to review new PRs by comparing the generated metrics between versions.</p> <p>Remove dpdkstats and dpdkevents from Barometer.</p> <p>Enable the Logparser plugin by default when using one-click install.</p> <h2>Testing Notes</h2> <ul style="list-style-type: none"><li>Added a playbook to compare collectd 5 and collectd 6. The playbook uses existing ansible roles to build both collectd 5 and collectd 6 container images, creates a common configuration, then runs the containers and shows the outputs to let the user inspect the metrics and whether they match.</li></ul> <h2>Documentation Updates</h2> <ul style="list-style-type: none"><li>Docs have been updated to use anuket/ repository in dockerhub. Container build instructions now use anuket/ prefix to tag images.</li></ul> <h2>Container updates</h2> <ul style="list-style-type: none"><li>Containers are now pulled from anuket/ repository in dockerhub.</li><li>Add a flask app for testing collectd using metrics sent via write_http plugin.</li><li>Grafana container was updated to support both jiffies and percent for cpu metrics.</li></ul> <h2>Ansible playbook updates</h2> <ul style="list-style-type: none"><li>Added <a href="#">unixsock</a> plugin to one-click install, which allows the user to interact with collectd using the <code>collectdctl</code> command in the bar-collectd-* containers. The unixsock plugin is useful for debugging issues in collectd, and can be used to verify that metrics are being collected without having to create CSV files or log into the container.</li><li>Added a playbook and role for building the collectd containers locally. This automates the actions described in the docker install guide. The <code>barometer-collectd</code>, <code>barometer-collectd-latest</code> and the <code>barometer-collectd-experimental</code> containers are now easier to build locally. The <code>barometer-collectd-6</code> and <code>barometer-collectd-experimental</code> containers can also be built with arbitrary PRs applied, to aid in testing locally.</li><li>Containers are now pulled from anuket/ repository in dockerhub.</li><li>The logparser plugin is now rendered for all flavours. The Logparser plugin has been part of collectd since 5.11, however, the ansible playbooks had it marked as experimental, and would not deploy it by default.</li></ul> <h2>Build script updates</h2> <ul style="list-style-type: none"><li>Update <code>collectd_apply_pull_request.sh</code> to rebase only if multiple chanegs are selected. The script will checkout the PR branch if there's only one PR_ID passed.</li></ul> <h2>Normal Bug Fixes</h2> <ul style="list-style-type: none"><li>Update the grafana dashboard to show metrics in both jiffies and percent, depending on what is configured.</li></ul>	<p>The main efforts this cycle were in improved automation, enabling easier testing and configuration of the metrics collection on the host.</p> <p>The containers were updated to more modern OSes, and reduced in size to make them quicker and easier to deploy.</p>	

		<h2>Deprecations</h2> <ul style="list-style-type: none"> <li>The dpdkstats and dpdkevents plugins were removed from Barometer. These plugins are still available in collectd, however, will not be deployed by Barometer. It is recommended that the DPDK telemetry plugin be used instead.</li> </ul> <h2>Other Notes</h2> <ul style="list-style-type: none"> <li>Add <a href="#">reno</a> and corresponding tox jobs (compile notes and add new notes) to make compiling release notes easier</li> </ul>		
VinePerf	<a href="#">Sridhar Rao</a>	<ul style="list-style-type: none"> <li>Traffic generator: <ul style="list-style-type: none"> <li>Support for DPPD-Prox is added. This is the sixth traffic generator supported by ViNePerf.</li> <li>T-Rex statistics are improved to provide better latency results.</li> </ul> </li> <li>Kubernetes: <ul style="list-style-type: none"> <li>ViNePerf container and corresponding pod-deployment file is included.</li> <li>A tool to deploy pods via helm-charts is added. This tool also extracts all the deployed pods and corresponding service information.</li> <li>Deployment of different CNIs (userspace, sriov, calico, cilium, multus, and danm) are supported.</li> <li>Custom DPPD-Prox and T-Rex containers are added - source files.</li> <li>L2l3fwd and VPP containers are included to support service-chain testing usecases.</li> </ul> </li> <li>Additional Features <ul style="list-style-type: none"> <li>Run ViNePerf as pod in Kubernetes. Results-only on-screen output mode is added for this use case.</li> <li>Dedicated 'k8s' mode is included for Kubernetes data plane testing.</li> </ul> </li> <li>Miscellaneous <ul style="list-style-type: none"> <li>More changes made to be consistent with the name change from VSwitchPerf to ViNePerf.</li> <li>K8S Dataplane Benchmarking work is accepted to be published in IEEE Globecom 2021.</li> </ul> </li> </ul>	<ul style="list-style-type: none"> <li>User can use DPPD-Prox with ViNePerf. DPPD-Prox users can take advantage of ViNePerf's ability to setup DUTs, manage Trafficgenerator, manage tests, manage load-generators, and collect and analyze results.</li> <li>ViNePerf has significantly enhanced its framework to support Kubernetes usecases. Users interested in Kubernetes Network performance testing can use ViNePerf.</li> </ul>	
CIRV	<a href="#">Sridhar Rao</a>	<ul style="list-style-type: none"> <li>Reorganization and Enhancement of Platform Description File - The single large PDF is broken down into multiple logical entities to include both Openstack and Kubernetes support.</li> <li>Kubernetes Post-Deployment Validation. Includes 12 custom post-deployment validations.</li> </ul>	<p>Kubernetes Post-Deployment validation can be used at different stages of Kubernetes cluster lifecycle.</p> <p>The enhanced PDF can be used for multiple automation purposes.</p>	
Airship				
RM	<a href="#">Walter Kozlowski</a>	<ul style="list-style-type: none"> <li>Hybrid Multi-cloud actor and interaction model has been significantly evolved</li> <li>Storage section has been further developed improving consumption and production models</li> <li>Security alignment with industry standards, in collaboration with GSMA FASG and LFN Security Forum</li> <li>Alignment of Acceleration Abstraction model with ORAN</li> <li>Reference Model Kali release contents was published as GSMA NG Permanent Reference Document NG.126 ver 2.0</li> </ul>	<p>The telco cloud is evolving towards a hybrid multi-cloud deployment model. The RM Lakelse defines a generic abstract powerful model that is of value to the Reference Architectures and other standard bodies engaged in this endeavor. The current model and related reference architectures provide and will continue to provide advanced guidelines for multi-cloud deployments in the industry. The growing interest in security, especially related to 5G and automation, was reflected in the Lakelse release. The infrastructure security guidelines in RM have been referenced by GSMA 5G documents making RM work widely accessible for the telecommunication industry participants.</p>	
RA1	<a href="#">Karine Sevilla</a>	<ul style="list-style-type: none"> <li>The latest release of RA1 specifications conforms to the OpenStack Wallaby release.</li> <li>This has allowed the introduction of Cyborg, the general-purpose acceleration resources (i.e., various types of accelerators such as GPU, FPGA, ASIC, NP, SoCs, etc.) management framework, as an optional service.</li> <li>Submitted RA-1 Kali release specifications were submitted for publication as a GSMA Permanent Reference Document.</li> </ul>	<p>RA1 now supports a more recent OpenStack release in addition to the widely used Train release version. The increasing use of accelerators (software and hardware-based), required for 5G, IoT, and ML among others, requires support for their management. Cyborg provides an abstraction for the specific devices (a la the Acceleration Abstraction Layer in O-RAN) and, thus, allows for their generic management.</p> <p>Publishing of the RA-1 by GSMA makes it more widely available as an industry standard.</p>	
RA2	<a href="#">Riccardo Gasparetto Stori</a>	<ul style="list-style-type: none"> <li>The specifications were updated to the latest release of Kubernetes, 1.22</li> <li>Many improvements on critical telco areas like Multitenancy, CNF and cluster Lifecycle management</li> <li>Alignment with RM on security and other topics</li> <li>Linked specifications with RC2/ Kuberef tests in order to track what requirements drive conformance testing of Kubernetes clusters.</li> </ul>	<p>As Kubernetes releases have a fast pace, users can refer to RA2 specifications to increase their conformance to an ever evolving platform.</p> <p>The new chapters and specifications on security and multitenancy are completing core parts of the Cloud Native Telco Cloud architecture.</p> <p>Linking to Conformance testing helps final users benchmark implementations of the platform and ensure consistency with the Architecture.</p>	
RC1				
RC2				
RI2	<a href="#">Rihab Banday</a>	<ul style="list-style-type: none"> <li>Update of Kubernetes components and features</li> <li>Update of RI2/Kuberef Cookbook to reflect new changes</li> <li>Support for RI2 deployment on VMs</li> </ul>	<p>Automated e2e deployment of a reference platform based on RA2 Kali on baremetal, infrastructure provider (Equinix Metal) and virtualized infrastructure using opensource tooling.</p>	