EANTC Tests Prove Hypervisor Speeds Match Bare Metal RAN

Independent Testing: VMware Telco Cloud Platform RAN Performance on Par with Physical Machines

"The architectural approach of VMware ... provided identical performance compared to a bare metal Kubernetes cluster, and much better provisioning and operations support."

EANTC INDEPENDENT TEST REPORT

VMware Surpasses Bare Metal for Automation, Ease, and Simplicity

New third-party testing proves that VMware Telco Cloud Platform RAN performs on par with or better than a bare-metal solution. The test also highlights that VMware simplifies the user experience and automates costly processes. The other platform requires deep technical expertise and lacks built-in checks.

The independent testing, which was performed by EANTC in its virtualized open RAN lab in Germany, solidly debunks the myth that a hypervisor imposes performance overhead on its host machine.

The testing not only proves that the performance of the VMware platform for radio access networks is equivalent to that of a bare metal server but also demonstrates that the VMware solution is easier to deploy, operate, and maintain than its rival.

The facts of the matter are clear from the third-party testing conducted by EANTC, and EANTC's results have been confirmed by VMware's own testing:

- For radio access networks, the VMware hypervisor, ESXi, performs on par with a Kubernetes solution running on a physical server.
- The VMware RAN platform is easier to set up, operate, and maintain.
- The VMware RAN stack simplifies the deployment and management of virtualized and containerized network functions.
- The automation and built-in checks of VMware Telco Cloud Platform RAN set it apart from the other system tested by EANTC.

Summary of EANTC's Key Conclusions

Here are the key conclusions from EANTC's testing; for more information, see the *EANTC Independent Test Report*.

Performance tests: By running the Cyclic and OSLat tests, EANTC confirmed that both the VMware platform and the bare-metal platform passed the O-RAN requirement for a real-time operating system because the maximum latency of both platforms was 9 microseconds (us). (In general, a 9 microsecond latency is well within the 5G RAN latency requirement for low-latency RAN communications.)

Memory test: The VMware platform delivered excellent memory optimization through its Transparent Memory Sharing (TPS) technology. The memory test was unable to be conducted on the rival bare-metal platform. See the report for details.

Noisy neighbor isolation: With both platforms, a noisy neighbor did not have an adverse effect on a latency-sensitive workload.

Powerful simplicity: VMware Telco Cloud Automation, the orchestration framework included with VMware Telco Cloud Platform RAN, shielded the EANTC testers — who played the role of the operator and the second-level supporter — from the complexity

ABOUT EANTC

EANTC is an internationally recognized test center located in Berlin, Germany. It provides objective, vendor-neutral network performance test facilities for manufacturers, enterprise customers, and service providers.

VMware commissioned EANTC to compare the performance and capabilities of VMware Telco Cloud Platform RAN and a rival bare-metal platform.

For more information about the results of the comparison, see the *EANTC Independent Test Report*.

For more information about EANTC, see its website at *https://eantc.de/*

vmware®

"The GUI of VMware Telco Cloud Automation, the orchestration framework included in VMware Telco Cloud Platform RAN, shielded us (playing the operator's and second-level supporter role) from all the complex command-line options. ... It provided a fast-learning curve and did not require much detailed knowledge for simple operations."

EANTC INDEPENDENT TEST REPORT

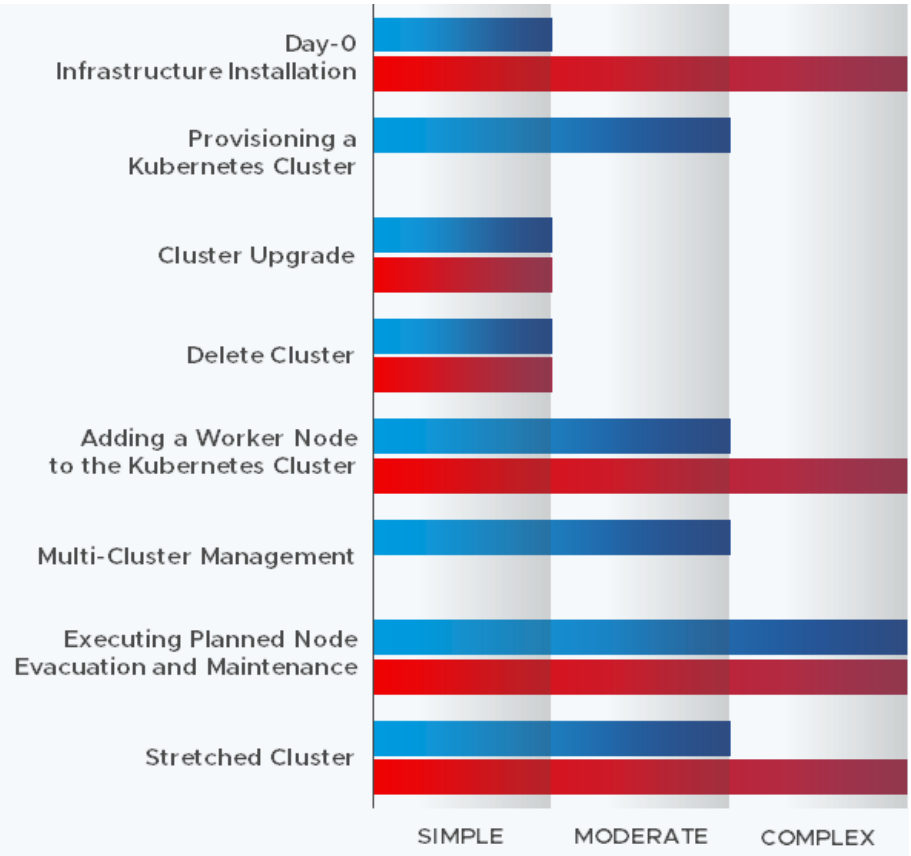

FIGURE 1: The CaaS LCM complexity ratings from the EANTC test results. The blue bars represent VMware Telco Cloud Platform RAN; the red bars represent the other platform..

of command-line options. "It provided a fast-learning curve and did not require much detailed knowledge for simple operations," the EANTC report says. The provisioning and management operations of VMware Telco Cloud Automation were straightforward and well abstracted on the cluster, making it easy and efficient to execute first-time and daily tasks. CaaS migrations can be executed with no downtime. For lifecycle management of the containers as a service (CaaS) infrastructure, the VMware platform is less complex than the bare-metal platform.

Reduction in error risk: The manual configuration that was required on the rival bare-metal platform needs more technical knowledge and increases the complexity of performing an activity. This more complex configuration process leads to a higher risk of probable errors. The VMware RAN platform eliminates the manual overhead and the need for more profound technical knowledge by automating the configuration process and making it more straightforward.

Easy customization through automation: EANTC noted that node customization is partially automated on the VMware platform while the procedure is completely manual on the rival bare-metal platform. On the bare-metal platform, the configuration files had to be manually created and applied.

As for the measured complexity and execution time to customize a node when deploying CaaS infrastructure, the VMware platform performs better because it automatically allocates the right resources — automation that is in effect enabled by the VMware hypervisor.

"[The cluster provisioning operation] on VMware was fully automated and simple to conduct."

EANTC INDEPENDENT TEST REPORT

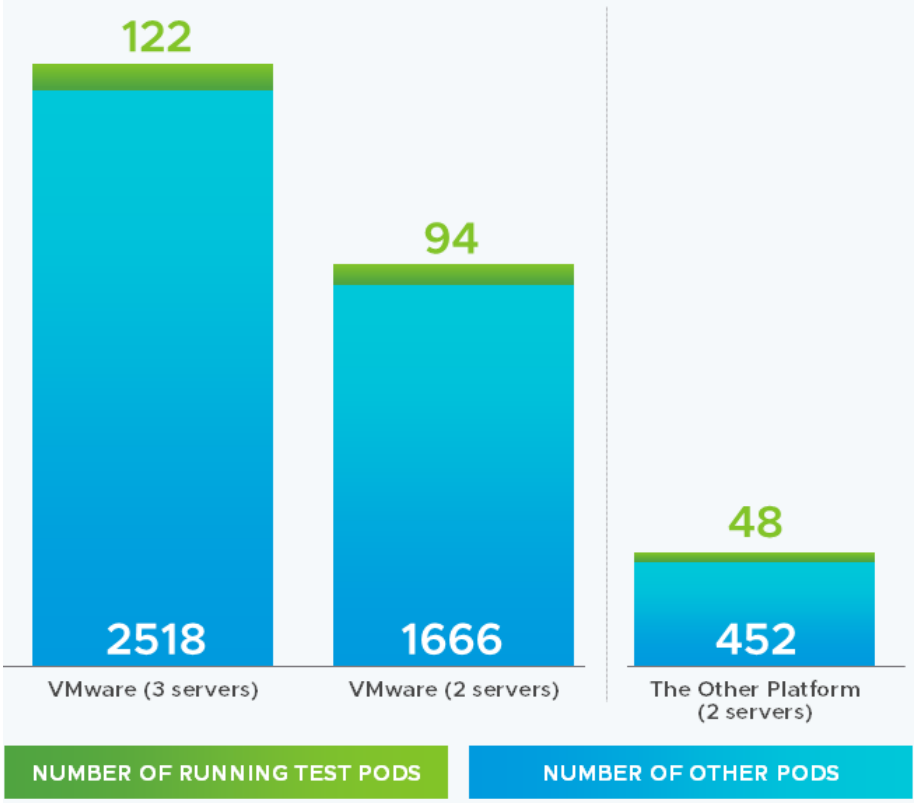

FIGURE 2: Pod density. EANTC executed the test on the VMware platform with three physical worker nodes. Because EANTC had only two physical worker nodes for the other platform's test setup, the VMware results had to be adapted for two physical worker nodes to enable a fair comparison. The first row in the figure shows the original VMware results on three servers; the second row shows the adapted VMware results, as calculated by EANTC; for information about the calculation, see the EANTC test report.

Pod Density Check

Although typical RAN deployments are characterized by the use of a small number of huge pods deployed on single worker nodes, pod density is essential to maximize the use of resources and the return on investment for infrastructure for non-RAN workloads, such as near edge applications. As the figure above illustrates, EANTC's testing shows that a CSP can achieve a much higher pod density with the VMware platform than with the other platform.

According to EANTC's test of pod density, if many simple non-RAN workloads need to be deployed, VMware Telco Cloud Platform RAN is more cost-efficient because it can deploy a higher number of pods on the same hardware compared to the other platform, which implies that the VMware platform would use fewer resources if both platforms were to run the same number of pods.

On the VMware platform, it is also easier to increase the total maximum number of pods because additional worker node virtual machines can be deployed. On the other platform, however, the only option is to increase the default limit in the worker nodes' kubelet configuration file. On the other platform, increasing the default limit and deploying more pods could slow down pod scheduling, increase management overhead and result in higher CPU utilization, or overcommit resources and lead to a reduction in workload performance.

THE SYNERGY OF CONTAINERS AND VIRTUAL MACHINES

VMs solve infrastructure-related problems by better utilizing servers, improving infrastructure management, and streamlining IT operations.

Containers solve application-related problems by streamlining DevOps, fostering a microservices architecture, improving portability, and further improving resource utilization.

Running CNFs on VMware Telco Cloud Platform RAN produces a synergy that helps CSPs transition from 4G to 5G RAN networks with ease.

BENEFITS OF HYPERVISORS AND VIRTUAL MACHINES FOR CNFS

- Onboard, deploy, and manage CNFs at scale through automation
- Establish strong security boundaries for containers
- Isolate workloads and apply built-in security measures like micro-segmentation
- Select the best Linux kernel version for your workload
- Optimize the performance of large Kubernetes clusters and mixed workloads on shared infrastructure
- Automate lifecycle management of Kubernetes clusters and RAN functions
- Optimize the placement and performance of CNFs with automatic resource provisioning
- Scale CNFs without the pain of adding, configuring, and managing physical hardware
- Streamline Day 0 through Day 2 operations and reduce OpEx

Time Check: Operations and Lifecycle Management

The VMware ESXi virtualization layer powers fast, easy, and efficient operations and lifecycle management. In some cases, the automation and built-in checks performed by VMware Telco Cloud Platform RAN result in slower operations than the other platform but easier and simpler operations with a lower risk of manual error.

SCENARIO	RESULT
Upgrading a Kubernetes cluster	The VMware platform took 24 minutes; the other platform, 77 minutes. Because cluster upgrades are repeated multiple times a year, the time it takes to upgrade a cluster affects maintenance windows.
Adding a worker node	The VMware platform took 8 minutes; the other platform, 17 minutes. Fully automated execution makes VMware less complicated than the other platform.
Instantiating a CNF	The VMware platform took 48 seconds; the other platform, 10 seconds. Although the command-line execution time of the other platform is faster than execution through the VMware GUI, the VMware platform automatically validates that the infrastructure can support the CNF. The other platform performs no infrastructure checks, increasing error risks.
Updating or upgrading a CNF	The VMware platform took 27 seconds and 25 seconds, respectively; the other platform, 7 and 14 seconds. Although the rival's command-line execution is faster, it is error prone, and it can extend upgrade times for inexperienced users.
CNF lifecycle management	Although the complexity of operations was seen by EANTC as equal, the VMware platform takes seconds more because of its automated prechecks and data synchronization. With the other platform, the checks must be performed manually.
Scaling a CNF	VMware took 23 seconds; the other platform, 21 seconds. The VMware platform uses automation to perform prechecks before scaling the CNF; the other platform does not.
Rolling back a CNF	VMware took 25 seconds; the other platform, 2 seconds. The extra time taken by the VMware platform stems from automated checks that simplify operations and reduce manual errors; the other platform performs no checks or automation.

TELCO-GRADE KUBERNETES

The CaaS functionality of VMware Telco Cloud Platform RAN simplifies the operation of Kubernetes, centralizing management and governance for clusters. The platform furnishes telco-grade CaaS enhancements, such as the following:

- Multus to attach multiple container networking interfaces to Kubernetes pods through its plug-ins
- Topology Manager to optimally allocate CPU memory and device resources on the same NUMA node to support performance-sensitive applications
- Kubernetes cluster automation to simplify deployments and management of Kubernetes master and worker nodes.
- With these enhancements, CSPs can take advantage of a telco-grade Kubernetes platform to address emerging 5G use cases at the RAN.

Day 0 through Day 2 Operations

According to the EANTC report, the minor differences in execution times between the other platform's command-line interface and the VMware platform are overshadowed by several clear and distinct advantages of the VMware platform:

- Simple setup and configuration
- Easy yet powerful user experience
- Automation and built-in prechecks and safeguards

Looking at VMware Telco Cloud Platform RAN through the lens of Day 0, Day 1, and Day 2 operations sets it apart from the other platform and reveals its simplicity, ease of use, and powerful automation.

The results of the EANTC report show that, by using a hypervisor, VMware Telco Cloud Platform RAN streamlines key Day 0, Day 1, and Day 2 operations, including the following:

- Installation and setup
- Kubernetes cluster deployment
- Node customization
- Cluster upgrade

Day 0: Set up the Platform and Deploy Kubernetes Clusters

Although the installation time took longer with the VMware platform, the VMware installation was less complicated because of the fully automated installation process.

With VMware, a zero-touch installation automatically installed all the required components, including ESXi, VMware vCenter Server Appliance, VMware Telco Cloud Automation, and Harbor, a container image registry. During setup, VMware Telco Cloud Automation automatically applied the configuration for the groups, data stores, and SR-IOV interfaces.

As for deploying a Kubernetes cluster, one key difference between the VMware platform and the other platform is that for the VMware platform, provisioning a new cluster is an operation carried out with complete independence from Day 0 installation. This separation fosters the use of templates to deploy Kubernetes clusters on any available host and at any time.

In contrast, on the other platform, provisioning a cluster is part of the Day 0 installation procedure, and thus tightly coupled to the cluster's characteristics and to the host being activated. This dependency means that if a new cluster needs to be provisioned or if a catastrophic cluster problem occurs, the cluster cannot be recovered — recovery will require undergoing the whole Day 0 installation procedure again, which takes about 77 minutes.

VMware separates Day 0 installation from new cluster creation by using the web interface of VMware Telco Cloud Automation to provision a Kubernetes cluster. VMware Telco Cloud Automation executed all the steps automatically —- including cluster creation and configuration, configurations for add-ons, VIM registration, and inventory update. The provisioning of a cluster was completed in 19 minutes.

Day 1: Instantiate and Deploy Network Functions

The normal way to manage package templates in Kubernetes is by using Helm. Templates that potentially consist of multiple YAML files and dependencies are called Helm charts. Both VMware and the other platform use such helm charts. The other platform uses them directly, whereas VMware Telco Cloud Automation provides a graphical front-end and a consistency validator for the configuration.

The GUI-based approach by VMware has more steps, but each of them is

straightforward or even trivial. The other platform's experience requires operators to fully understand the creation and maintenance of such Helm charts.

In some cases, software manufacturers might provide the Helm charts with their software, but that cannot be taken for granted in all cases. In such cases where a Helm chart does not accompany a manufacturer's software, VMware Telco Cloud Automation makes it easy and requires less expertise to create a Helm chart.

VMware Telco Cloud Automation took 48 seconds of execution time to complete the instantiation; the other platform took 10 seconds. The EANTC report notes that there are no hard, well-defined goals for instantiation; in many cases, an instantiation that takes a minute or less is reasonable. If an instantiation is highly time-critical (because, for example, a large number of CNFs needs to be instantiated), the CLI-based approach is faster. However, the likelihood of failures with the CLI-based approach is higher because the checks happen only at instantiation time. The other platform does not differentiate between onboarding and instantiation as the VMware platform does.

Node Customization

The EANTC report showcases the advantages of automatically allocating resources during node customization. VMware Telco Cloud Platform RAN dynamically configures the underlying infrastructure resources to meet the requirements of the CNFs being deployed.

This approach prevents overprovisioning of hardware resources because VMware Telco Cloud Platform RAN allocates only pass-through resources when a CNF is deployed. The VMware ESXi hypervisor virtualizes the underlying server resources to enable VMware Telco Cloud Platform RAN to dynamically allocate resources during CNF instantiation and to ensure the right node customization at instantiation time, regardless of the type of CNF or the CNF's vendor.

This automated approach significantly shortens pre-deployment configuration, customization, and validation times. An automated approach also greatly simplifies the deployment process, especially in multi-vendor RAN environments in which different CNFs require different infrastructure properties.

Day 2: Manage, Scale, and Optimize Network Functions with Automation

This section looks at a sampling of the Day 2 operations highlighted in the test report.

CaaS Lifecycle Management

The duration of CaaS lifecycle management operations differ between the two platforms, and the level of expertise of the administrator can affect the results.

When it comes to the complexity of executing these operations, VMware automates most of the procedures by default. Following a few steps using the GUI is easier than executing different CLI commands. It reduces the probability of human errors, simplifies the steps for non-experienced users, and eases troubleshooting.

CNF Update and Upgrade

It's not uncommon to change the configuration of a CNF during its operation. One option is to terminate the CNF and instantiate another one with the new configuration. That approach, however, can interrupt service. In cases in which CNFs are long-lived, complex, or both, a common practice is to update the configuration.

Kubernetes manages the update with a "helm upgrade" command that supplies a new Helm chart. Helm then calculates the differences between the old and new configuration and runs a minimal modification for a smooth upgrade. While the helm upgrade covers all possible configuration and software image changes, "CNF Update" refers to configuration changes, and "CNF Upgrade" refers to upgrading to new software versions. EANTC tested both scenarios, which differ on the VMware platform but are identical on the other platform.

THE BOTTOM LINE FROM EANTC'S INDEPENDENT TESTING

EANTC's extensive tests of VMware Telco Cloud Platform RAN and a rival platform compared the installation, functionality, and performance of the two solutions. The validation specifically focused on the requirements for disaggregated RAN solutions, which are being intensively evaluated by mobile operators.

- VMware Telco Cloud Platform RAN excelled in reliable low-latency performance and user-friendly GUI-based management procedures.
- The command-line operations of the other platform increase the risk of manual errors; the built-in checks and automation of the VMware platform safeguard operations.
- EANTC did not witness any performance overhead of the VMware hypervisor that is included with VMware Telco Cloud Platform RAN.

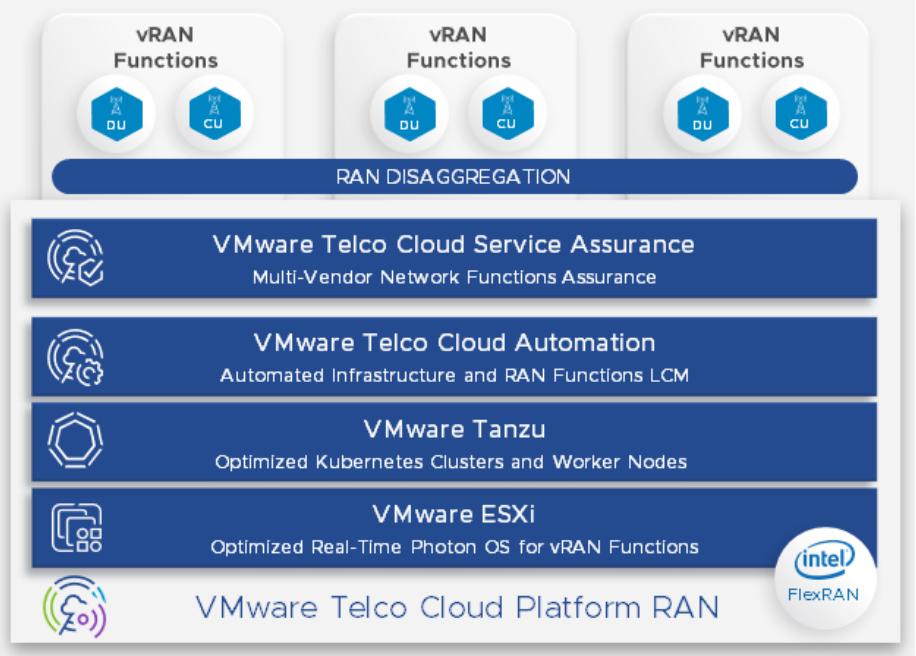

FIGURE 3: VMware Telco Cloud Platform RAN modernizes the radio access network.

CNF Update and CNF Upgrade functions are supported by both platforms. VMware Telco Cloud Automation provides a GUI-based wrapper for the basic helm upgrade function. In the GUI, an operator can change the individual instantiation properties. On the other platform, a new Helm chart needs to be edited by the operator and supplied to the CLI-based upgrade function.

Roll Back CNF

No operator wants to have to roll back a CNF. Rolling back a CNF to a previous state usually means that something went wrong during a previous Upgrade command. The Rollback command returns a CNF to a previous revision smoothly. Both platforms support rolling back CNFs, and the functionality was provided by each of them correctly. The rollback took two seconds on the other platform, executing the corresponding helm rollback command.

The VMware platform conducts additional checks and verifications, and takes a total of 25 seconds to roll back to a previous revision under similar conditions. The times of both platforms were seen as acceptable because they were within less than a minute, which is more than acceptable for such a rare activity.

Summary CNF Lifecycle Management

According to EANTC, on the rival platform, the difficulty of executing a LCM operation increases with the complexity of the CNF; on the VMware platform, however, the difficulty of executing the operation did not change with the complexity of the CNF. The node customization test case illustrates the difference. Although the rival platform performed the operation slightly faster, the VMware platform removed potential manual configuration steps. The advantages of automated hardware configurations are that they are less prone to error and reduce the potential for human errors that lead to long troubleshooting cycles and, in some cases, outages, and the user does not need to know details about the infrastructure to get the desired configuration.

RAN PERFORMANCE IS EQUIVALENT

Industry-standard real-time micro-benchmarks run by VMware, namely cyclictest and oslat, show that RAN performance is equivalent on VMware vSphere and bare metal — results that were confirmed by EANTC's testing.

Cyclictest, which uses a hardware-based timer to measure platform latency and jitter, demonstrated that the latency on both vSphere 7.0U3 and on bare metal was less than 10 microseconds. A 10 microsecond latency is well within the latency requirements of RAN workloads.

The oslat performance test is an open-source micro-benchmark that measures jitter in a busy loop. Instead of using hardware-based timers, this benchmark uses a CPU bound loop as its measurement—which emulates a virtualized RAN workload in a real-world scenario, such as a polling thread using the Data Plane Development Kit (DPDK).

- For the results of the cyclictest and oslat tests run by VMware, see *vSphere Performance Is Equivalent to Bare Metal for RAN Workloads*.
- For the cyclictest and oslat results from the independent testing performed by EANTC, see the *EANTC Independent Test Report*.

LEARN MORE

For more information about VMware Telco Cloud Platform RAN, call 1-877-VMWARE (outside North America, dial +1-650-427-5000) or visit *https://telco.vmware.com/*

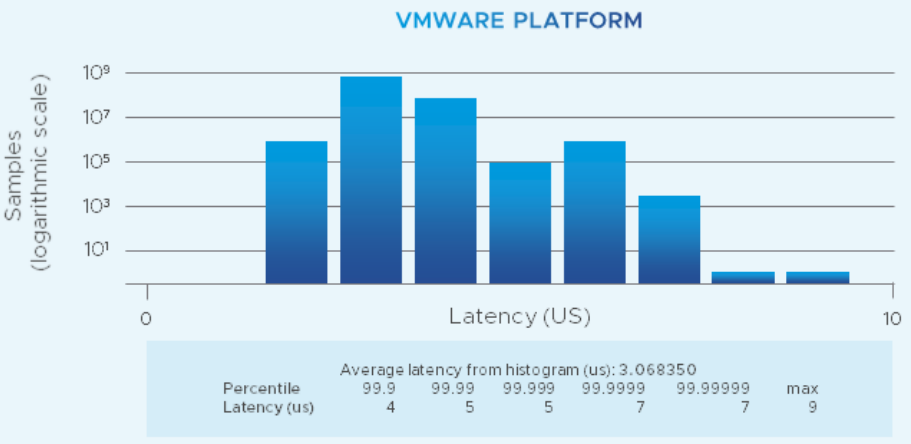

FIGURE 4: Cyclic tests performed by EANTC prove that latency on the VMware platform is less than 10 microseconds (us), which is ideal for RAN workloads. For the test results illustrated in this diagram, hyper-threading was enabled.

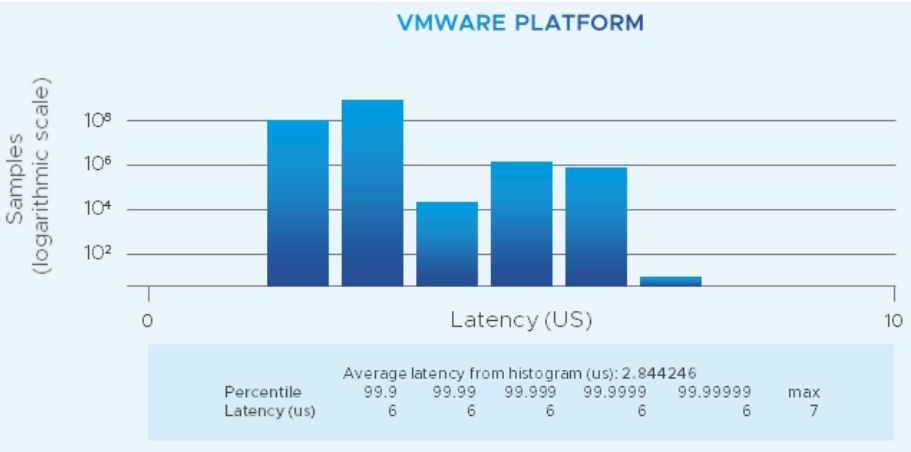

FIGURE 5: Cyclic tests performed by EANTC prove that latency on the VMware platform is less than 10 microseconds (us), which is ideal for RAN workloads. For the test results illustrated in this diagram, hyper-threading was disabled.

RAN Performance Is Equivalent on Bare Metal and VMware

Some vendors claim that bare-metal servers provide better performance for RAN workloads than the VMware platform. Is there a performance tax for real-time RAN workloads on VMware? The answer is no.

EANTC ran the industry-standard real-time micro-benchmarks, namely cyclictest and oslat, to compare the performance of RAN workloads on the VMware platform and a bare-metal platform. EANTC's testing found that **performance is equivalent**.

The tests show that there is no performance penalty or latency tax with the VMware platform. The performance of RAN workloads on VMware vs. bare metal, as measured by the micro-benchmarks, is equivalent. For more information, see the *EANTC Independent Test Report*.