Cloud Peak—Virtual Cloud Infrastructure Validation

Problem: production VNFs perform poorly and require costly over-provisioning

Network Functions Virtualization (NFV) is a key building block for 5G mobile networks. In an NFV architecture, various Virtual Network Functions (VNF) and Cloud-Native Network Functions (CNF) run on a common NFV Infrastructure (NFVI) and are controlled by a common management and orchestration (MANO) layer. This has made new use cases such as network slicing and mobile edge computing (MEC) possible.

The flexibility that enables these new 5G use-cases creates major challenges for carriers as they seek to deliver high-quality services to their customers. NFVI must be dimensioned and configured correctly, and thoroughly validated to ensure the performance of the VNFs running on top of it. Complicating matters, the NFVI will be running diverse simultaneous workloads that interact with each other, leading to massive and costly over-provisioning.

How can carriers validate such a complex system—composed of multiple compute, network, and storage resources—and its interaction with the rest of the NFV architecture elements?

Solution: NFVI validation through VNF workload generation

Ixia's Cloud Peak is a web application designed to rigorously benchmark the performance of virtualized network infrastructures. By deploying real Virtual Machine or Docker Container workloads on top of the NFVI system under test (SUT), the application provides key insights into the capability of the NFVI to sustain the required VNF and CNF workloads.

Using an innovative methodology that decomposes system resources into the elementary compute, network, and storage categories, the application can isolate configuration issues and performance bottlenecks. The test results are compared to preconfigured performance levels to generate an easy-to-understand Pass / Fail result, displayed through an intuitive real-time user interface (UI). Whether testing a single small server or validating a whole infrastructure with hundreds of nodes, the product can scale to match the capacity of the SUT.

KEYSIGHT

- Easy-to-use solution for virtual infrastructure validation.
- Benchmark private, hybrid, telco, edge, and public clouds.
- Compare benchmark results across different infrastructures.
- Assess the impact of config changes via repeated testing.
- Find peak system capacity and oversubscription behavior.
- Run workloads in virtual machines or Docker containers.
- Validate the NFVI SUT with VNF / CNF workloads.
- Isolate compute, network, and storage resource categories.
- Start small with testing a single compute node.
- Go large with testing multiple server racks.
- Automate test execution and schedule periodic testing.

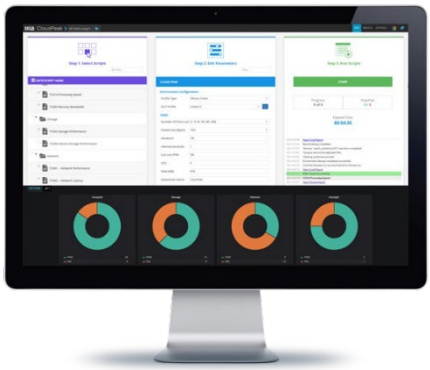

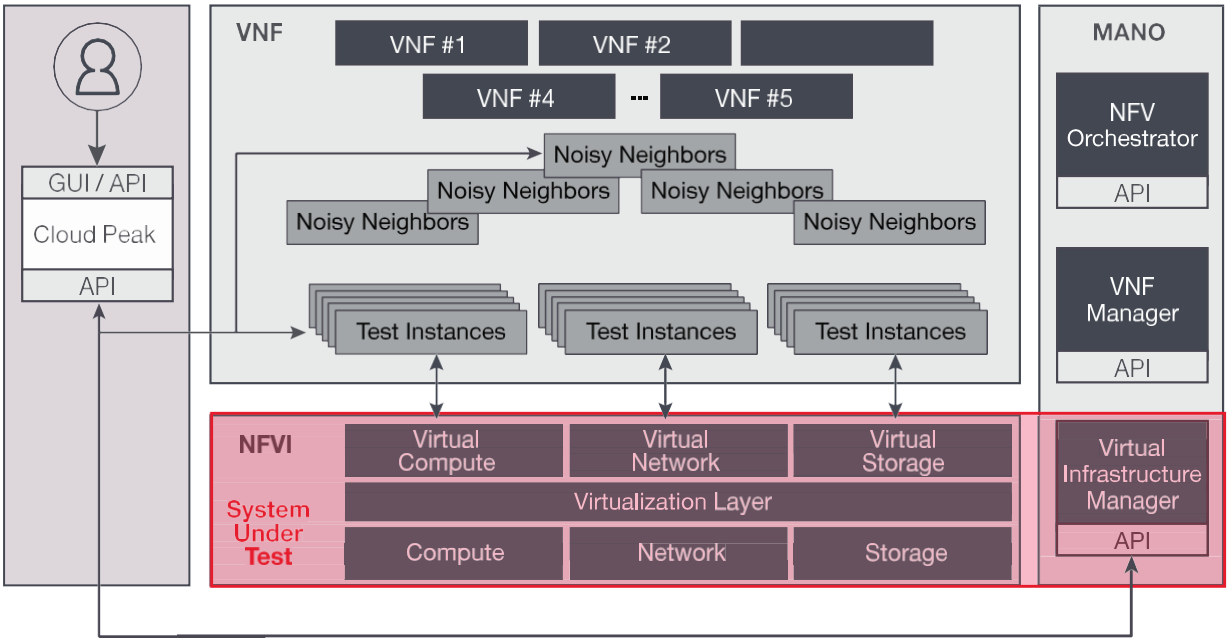

Figure 1. Cloud Peak—architecture diagram

Key features

- Complete virtual infrastructure performance benchmarking for private, hybrid, telco, edge, and public clouds.
- Workloads available as Virtual Machine or Docker Container images with similar benchmark capabilities.
- Automated workload deployment via specific API interfaces on Amazon, OpenStack, VMware, K8S clouds.
- Manual workload deployment via predefined OVA / QCOW2 / AMI images for other types of clouds.
- Predefined test methodologies with unique parameters tuned for virtual infrastructure testing.
- Industry-proven workload emulation based on the open source OPNFV Yardstick portfolio.
- Custom-built workload emulations for complete virtual infrastructure benchmarking.
- Decomposes and individually validates the compute, network, storage, and VIM performance dimensions.
- Measures the VIM performance with custom VM Instantiation and VM Termination test methodology.
- Measures the scheduler capability to isolate the good workloads and the bad noisy neighbors.
- Validates the SUT from application perspective through powerful workload emulation.
- Validates the SUT from network infrastructure perspective through high performance traffic generation.
- Scales from small test beds with a single compute node to large environments with many compute racks.
- Groups the tests into Test Sessions and Test Playlists for simplifying the test configuration management.
- Executes test scenarios sequentially (single-dimensional testing) for initial system characterization.
- Executes test scenarios in parallel (multi-dimensional testing) to assess how they impact each other.
- Controls the test tool via intuitive web-based UI that offers dynamic dashboards and real-time statistics.
- Controls the test tool via REST API to automate the test execution and run repeatable regressions.
- Deploys on standard of-the-shelf virtualization platforms across private / hybrid / telco / edge / public clouds.
- Provides subscription licensing for low startup cost and flexibility of pay-as-you-grow model.

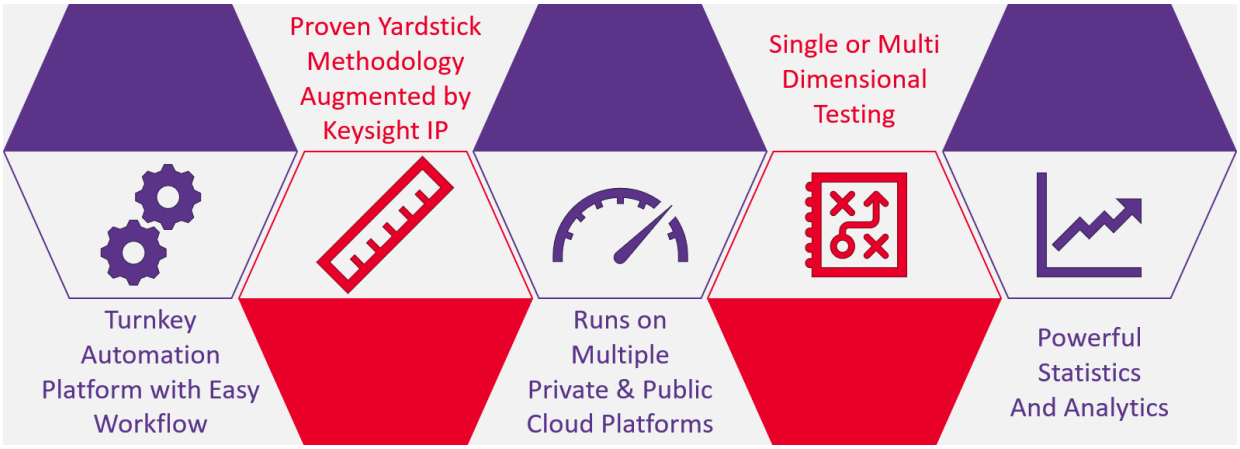

Figure 2. Cloud Peak — key features

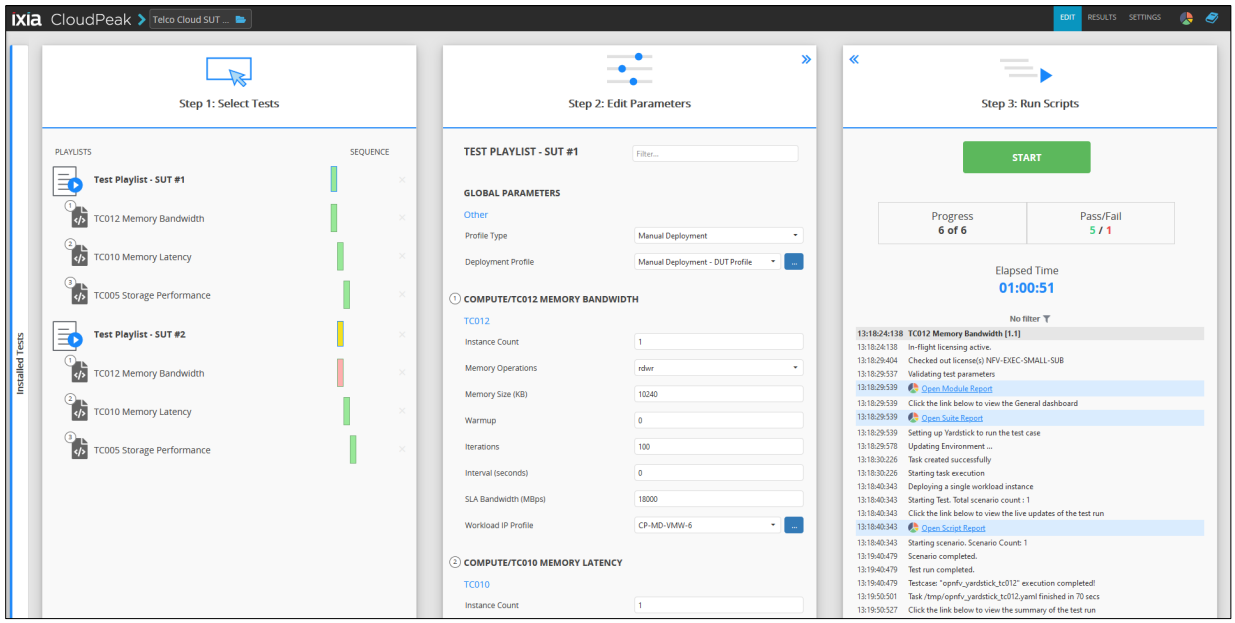

Figure 3. Cloud Peak—configuration interface

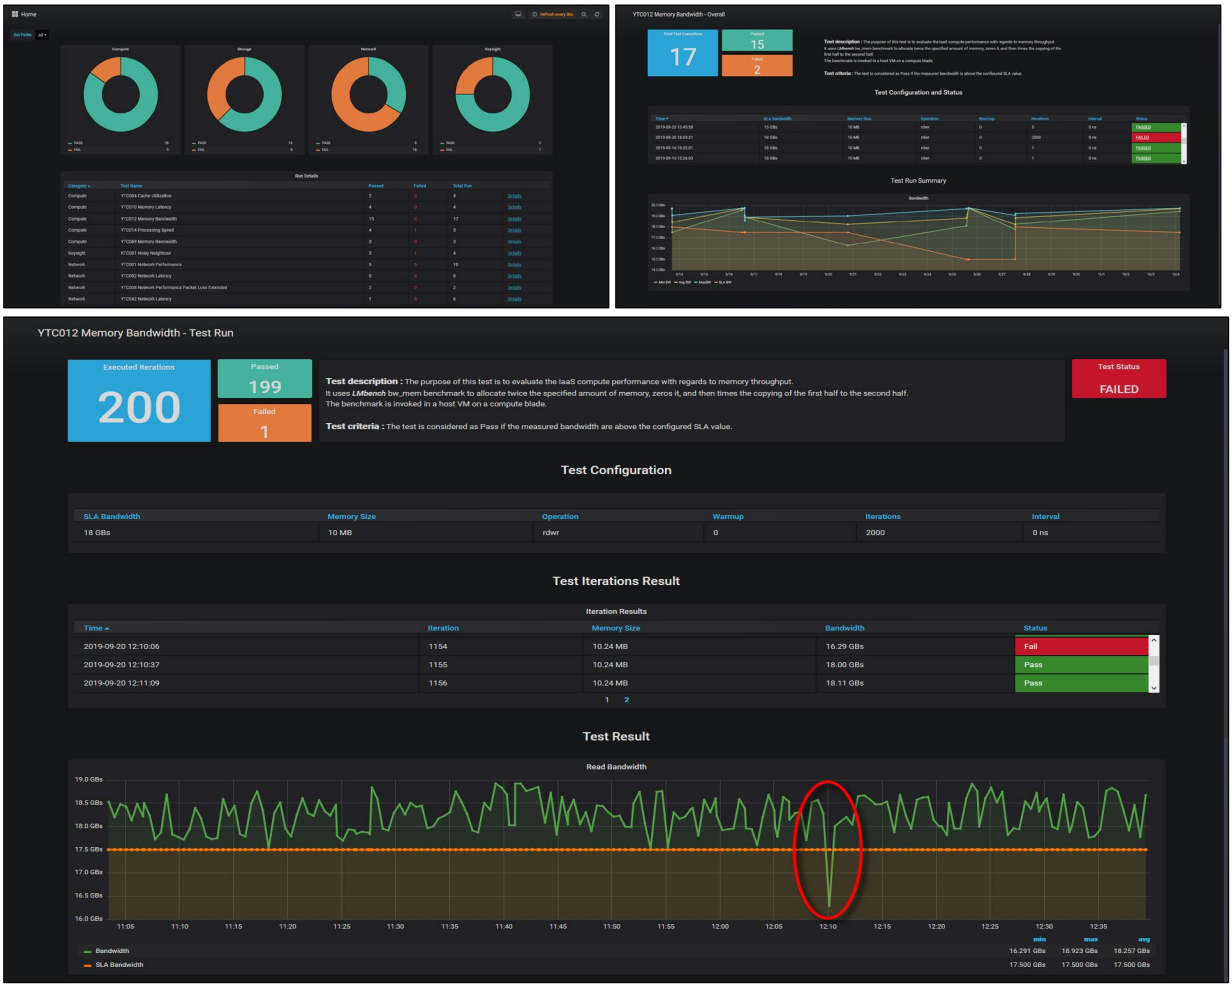

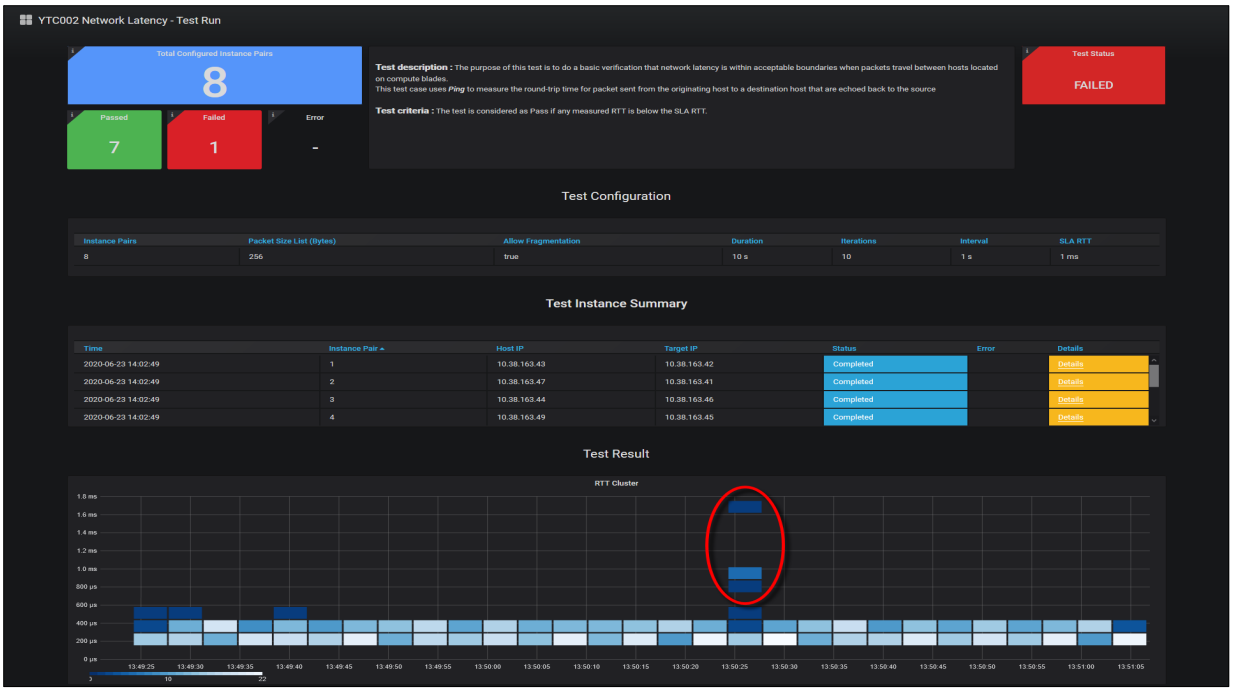

Figure 4. Cloud Peak—reporting interface

Specifications

Cloud Peak contains multiple test methodologies, and each one is specifically designed to validate various performance dimensions of the SUT. The following table describes the category of resources validated, the test methodologies included, the test ID, the key performance indicators (KPI), the underlying tools used for validation, as well as the various environments supported by each test methodology.

Category	Test methodology	ID	Key performance indicators	Tool	Horizontal scale	Automatic deployment				Manual deployment	
						OpenStack	VMware	Amazon	Kubernetes	Amazon	Others *
COMPUTE	CPU cache benchmarking	TC004	CPU cache Hit / Miss / Ratio	CACHESTAT	Y	Y	Y	N	Y	N	Y
	CPU performance benchmarking	TC014	CPU performance score	UNIX BENCH	Y	Y	Y	Y	Y	Y	Y
	Memory latency benchmarking	TC010	Memory latency (ns)	LM BENCH	Y	Y	Y	Y	Y	Y	Y
	Memory bandwidth benchmarking	TC012	Memory bandwidth (GBps)	LM BENCH	Y	Y	Y	Y	Y	Y	Y
	Memory bandwidth benchmarking	TC069	Memory bandwidth (GBps)	RAM SPEED	Y	Y	Y	Y	Y	Y	Y
NETWORK	Network loss benchmarking	TC001	Packet loss (PPM)	PKTGEN	Y	Y	Y	Y	N	Y	Y
	Network latency benchmarking	TC002	Packet latency (RTT)	ICMP	Y	Y	Y	Y	Y	Y	Y
	Network loss benchmarking	TC008	Packet loss (PPM)	PKTGEN	Y	Y	Y	Y	N	Y	Y
	Network jitter benchmarking	TC011	Packet jitter (µs)	IPERF	Y	Y	Y	Y	Y	Y	Y
	Network mixed benchmarking	TC038	Packet loss (PPM) Packet latency (RTT) CPU utilization (%)	PKTGEN ICMP MPSTAT	Y	Y	Y	Y	N	Y	Y
	Network latency benchmarking	TC042	Packet latency (RTT)	PKTGEN DPDK	Y	Y	Y	N	N	N	Y
	Network mixed benchmarking	TC070	Packet loss (PPM) Packet latency (RTT) Memory utilization (RAM)	PKTGEN ICMP FREE	Y	Y	Y	Y	N	Y	Y
	Network latency benchmarking	TC083	Packet latency (TCP / UDP)	NETPERF	Y	Y	Y	Y	Y	Y	Y

Category	Test methodology	ID	Key performance indicators	Tool	Horizontal scale	Automatic deployment				Manual deployment	
						OpenStack	VMware	Amazon	Kubernetes	Amazon	Others *
	Network loss benchmarking	KCT03	Packet loss (PPM)	PKTGEN DPDK	Y	Y	N	N	N	N	Y
	Network TPUT benchmarking	KCT04	Packet TX rate (Mbps) Packet RX rate (Mbps)	PKTGEN DPDK	Y	Y	Y	Y	N	Y	Y
	Network loss and throughput	KTC05	Packet loss (PPM) Packet TX rate (Mbps) Packet RX rate (Mbps)	IPERF	Y	Y	Y	Y	Y	Y	Y
STORAGE	**Storage benchmarking**	TC005	BW / IOPS / Latency (read) BW / IOPS / Latency (write)	FIO	Y	Y	Y	Y	Y	Y	Y
STORAGE	**Storage benchmarking**	TC006	BW / IOPS / Latency (read) BW / IOPS / Latency (write)	FIO	Y	Y	N	N	Y	N	N
VIM	**Noisy neighbor benchmarking**	KTC01	Noisy neighbor success rate Noisy neighbor resource usage	STRESS NG	Y	Y	Y	Y	Y	Y	Y
VIM	**VM lifecycle management**	KTC02	VM deployment success rate VM deployment speed	N / A	N / A	Y	N	N	N	N / A	N / A

* Refers to virtualization platforms running Virtual Machines supplied in any of the OVA / QCOW2 formats

Cloud Peak supports benchmarking multiple types of cloud infrastructures across various performance dimensions. These benchmarks can be executed with Cloud Peak Workloads which have one or multiple network interfaces. When multiple network interfaces are used, the first interface is used for management functions (labeled as MNG) while the remaining interfaces are used for test functions (labeled as TST). The MNG network interfaces can only be configured with IPv4 addresses while the TST network interfaces can be configured with either IPv4 or IPv6 addresses.

There is a very large number of possible test combinations involving the deployment type (Automatic or Manual), the cloud infrastructure type (Amazon, Kubernetes, OpenStack, VMware, Others), the test type (Compute, Network, Storage, VIM), the number of interfaces (One or Multiple), and the address type (IPv4 or IPv6). Some of these test combinations are not yet supported by Cloud Peak. The compatibility matrix below contains detailed information regarding the supported capabilities as per the following color-coding rules:

… Test supports complete network capabilities (Multiple Interfaces and IPv6 Addresses where applicable)
… Test supports partial network capabilities (lacking support for Multiple Interfaces or IPv6 Addresses)
… Test not supported (although applicable to the particular cloud infrastructure)
… Test not supported (does not apply)

Category	Test methodology	ID	Automatic deployment				Manual deployment	
			OpenStack	VMware	Amazon	Kubernetes	Amazon	Others
COMPUTE	CPU cache benchmarking	TC004	MAX IF # = 1 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = N/A TST IPv6 = N/A	MAX IF # = 1 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = N/A TST IPv6 = N/A	MAX IF # = 0 MGN IPv4 = NOT MNG IPv6 = NOT TST IPv4 = N/A TST IPv6 = N/A	MAX IF # = 1 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = N/A TST IPv6 = N/A	MAX IF # = 0 MGN IPv4 = NOT MNG IPv6 = NOT TST IPv4 = N/A TST IPv6 = N/A	MAX IF # = 1 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = N/A TST IPv6 = N/A
	CPU performance benchmarking	TC014	MAX IF # = 1 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = N/A TST IPv6 = N/A	MAX IF # = 1 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = N/A TST IPv6 = N/A	MAX IF # = 1 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = N/A TST IPv6 = N/A	MAX IF # = 1 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = N/A TST IPv6 = N/A	MAX IF # = 1 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = N/A TST IPv6 = N/A	MAX IF # = 1 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = N/A TST IPv6 = N/A
	Memory latency benchmarking	TC010	MAX IF # = 1 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = N/A TST IPv6 = N/A	MAX IF # = 1 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = N/A TST IPv6 = N/A	MAX IF # = 1 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = N/A TST IPv6 = N/A	MAX IF # = 1 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = N/A TST IPv6 = N/A	MAX IF # = 1 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = N/A TST IPv6 = N/A	MAX IF # = 1 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = N/A TST IPv6 = N/A
	Memory bandwidth benchmarking	TC012	MAX IF # = 1 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = N/A TST IPv6 = N/A	MAX IF # = 1 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = N/A TST IPv6 = N/A	MAX IF # = 1 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = N/A TST IPv6 = N/A	MAX IF # = 1 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = N/A TST IPv6 = N/A	MAX IF # = 1 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = N/A TST IPv6 = N/A	MAX IF # = 1 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = N/A TST IPv6 = N/A
	Memory bandwidth benchmarking	TC069	MAX IF # = 1 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = N/A TST IPv6 = N/A	MAX IF # = 1 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = N/A TST IPv6 = N/A	MAX IF # = 1 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = N/A TST IPv6 = N/A	MAX IF # = 1 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = N/A TST IPv6 = N/A	MAX IF # = 1 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = N/A TST IPv6 = N/A	MAX IF # = 1 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = N/A TST IPv6 = N/A
NETWORK	Network loss benchmarking	TC001	MAX IF # = 2 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = YES TST IPv6 = YES	MAX IF # = 2 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = YES TST IPv6 = NOT	MAX IF # = 1 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = NOT TST IPv6 = NOT	MAX IF # = 0 MGN IPv4 = NOT MNG IPv6 = NOT TST IPv4 = NOT TST IPv6 = NOT	MAX IF # = 2 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = YES TST IPv6 = YES	MAX IF # = 2 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = YES TST IPv6 = YES
	Network latency benchmarking	TC002	MAX IF # = 2 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = YES TST IPv6 = YES	MAX IF # = 2 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = YES TST IPv6 = NOT	MAX IF # = 1 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = NOT TST IPv6 = NOT	MAX IF # = 1 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = NOT TST IPv6 = NOT	MAX IF # = 2 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = YES TST IPv6 = YES	MAX IF # = 2 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = YES TST IPv6 = YES
	Network loss benchmarking	TC008	MAX IF # = 2 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = YES TST IPv6 = YES	MAX IF # = 2 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = YES TST IPv6 = NOT	MAX IF # = 1 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = NOT TST IPv6 = NOT	MAX IF # = 0 MGN IPv4 = NOT MNG IPv6 = NOT TST IPv4 = NOT TST IPv6 = NOT	MAX IF # = 2 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = YES TST IPv6 = YES	MAX IF # = 2 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = YES TST IPv6 = YES

Category	Test methodology	ID	Automatic deployment				Manual deployment	
			OpenStack	VMware	Amazon	Kubernetes	Amazon	Others
	Network jitter benchmarking	TC011	MAX IF # = 2 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = YES TST IPv6 = YES	MAX IF # = 2 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = YES TST IPv6 = NOT	MAX IF # = 1 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = NOT TST IPv6 = NOT	MAX IF # = 1 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = NOT TST IPv6 = NOT	MAX IF # = 2 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = YES TST IPv6 = YES	MAX IF # = 2 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = YES TST IPv6 = YES
	Network mixed benchmarking	TC038	MAX IF # = 2 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = YES TST IPv6 = YES	MAX IF # = 2 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = YES TST IPv6 = NOT	MAX IF # = 1 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = NOT TST IPv6 = NOT	MAX IF # = 0 MGN IPv4 = NOT MNG IPv6 = NOT TST IPv4 = NOT TST IPv6 = NOT	MAX IF # = 2 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = YES TST IPv6 = YES	MAX IF # = 2 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = YES TST IPv6 = YES
	Network latency benchmarking	TC042	MAX IF # = 3 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = YES TST IPv6 = NOT	MAX IF # = 3 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = YES TST IPv6 = NOT	MAX IF # = 0 MGN IPv4 = NOT MNG IPv6 = NOT TST IPv4 = NOT TST IPv6 = NOT	MAX IF # = 0 MGN IPv4 = NOT MNG IPv6 = NOT TST IPv4 = NOT TST IPv6 = NOT	MAX IF # = 0 MGN IPv4 = NOT MNG IPv6 = NOT TST IPv4 = NOT TST IPv6 = NOT	MAX IF # = 3 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = YES TST IPv6 = NOT
	Network mixed benchmarking	TC070	MAX IF # = 2 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = YES TST IPv6 = YES	MAX IF # = 2 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = YES TST IPv6 = NOT	MAX IF # = 1 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = NOT TST IPv6 = NOT	MAX IF # = 0 MGN IPv4 = NOT MNG IPv6 = NOT TST IPv4 = NOT TST IPv6 = NOT	MAX IF # = 2 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = YES TST IPv6 = YES	MAX IF # = 2 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = YES TST IPv6 = YES
	Network latency benchmarking	TC083	MAX IF # = 2 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = YES TST IPv6 = YES	MAX IF # = 2 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = YES TST IPv6 = NOT	MAX IF # = 1 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = NOT TST IPv6 = NOT	MAX IF # = 1 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = NOT TST IPv6 = NOT	MAX IF # = 2 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = YES TST IPv6 = YES	MAX IF # = 2 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = YES TST IPv6 = YES
	Network loss benchmarking	KCT03	MAX IF # = 2 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = YES TST IPv6 = NOT	MAX IF # = 0 MGN IPv4 = NOT MNG IPv6 = NOT TST IPv4 = NOT TST IPv6 = NOT	MAX IF # = 0 MGN IPv4 = NOT MNG IPv6 = NOT TST IPv4 = NOT TST IPv6 = NOT	MAX IF # = 0 MGN IPv4 = NOT MNG IPv6 = NOT TST IPv4 = NOT TST IPv6 = NOT	MAX IF # = 0 MGN IPv4 = NOT MNG IPv6 = NOT TST IPv4 = NOT TST IPv6 = NOT	MAX IF # = 2 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = YES TST IPv6 = NOT
	Network TPUT benchmarking	KCT04	MAX IF # = 2 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = YES TST IPv6 = NOT	MAX IF # = 2 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = YES TST IPv6 = NOT	MAX IF # = 1 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = NOT TST IPv6 = NOT	MAX IF # = 0 MGN IPv4 = NOT MNG IPv6 = NOT TST IPv4 = NOT TST IPv6 = NOT	MAX IF # = 2 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = YES TST IPv6 = NOT	MAX IF # = 2 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = YES TST IPv6 = NOT
	Network loss and throughput	KTC05	MAX IF # = 2 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = YES TST IPv6 = YES	MAX IF # = 2 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = YES TST IPv6 = NOT	MAX IF # = 1 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = NOT TST IPv6 = NOT	MAX IF # = 1 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = NOT TST IPv6 = NOT	MAX IF # = 2 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = YES TST IPv6 = YES	MAX IF # = 2 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = YES TST IPv6 = YES

Category	Test methodology	ID	Automatic deployment: OpenStack	Automatic deployment: VMware	Automatic deployment: Amazon	Automatic deployment: Kubernetes	Manual deployment: Amazon	Manual deployment: Others
STORAGE	Storage benchmarking	TC005	MAX IF # = 1 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = N/A TST IPv6 = N/A	MAX IF # = 1 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = N/A TST IPv6 = N/A	MAX IF # = 1 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = N/A TST IPv6 = N/A	MAX IF # = 1 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = N/A TST IPv6 = N/A	MAX IF # = 1 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = N/A TST IPv6 = N/A	MAX IF # = 1 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = N/A TST IPv6 = N/A
STORAGE	Storage benchmarking	TC006	MAX IF # = 1 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = N/A TST IPv6 = N/A	MAX IF # = 0 MGN IPv4 = NOT MNG IPv6 = NOT TST IPv4 = N/A TST IPv6 = N/A	MAX IF # = 0 MGN IPv4 = NOT MNG IPv6 = NOT TST IPv4 = N/A TST IPv6 = N/A	MAX IF # = 1 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = N/A TST IPv6 = N/A	MAX IF # = 0 MGN IPv4 = NOT MNG IPv6 = NOT TST IPv4 = N/A TST IPv6 = N/A	MAX IF # = 0 MGN IPv4 = NOT MNG IPv6 = NOT TST IPv4 = N/A TST IPv6 = N/A
VIM	Noisy neighbor benchmarking	KTC01	MAX IF # = 1 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = N/A TST IPv6 = N/A	MAX IF # = 1 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = N/A TST IPv6 = N/A	MAX IF # = 1 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = N/A TST IPv6 = N/A	MAX IF # = 1 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = N/A TST IPv6 = N/A	MAX IF # = 1 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = N/A TST IPv6 = N/A	MAX IF # = 1 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = N/A TST IPv6 = N/A
VIM	VM lifecycle management	KTC02	MAX IF # = 1 MGN IPv4 = YES MNG IPv6 = NOT TST IPv4 = N/A TST IPv6 = N/A	MAX IF # = 0 MGN IPv4 = NOT MNG IPv6 = NOT TST IPv4 = N/A TST IPv6 = N/A	MAX IF # = 0 MGN IPv4 = NOT MNG IPv6 = NOT TST IPv4 = N/A TST IPv6 = N/A	MAX IF # = 0 MGN IPv4 = NOT MNG IPv6 = NOT TST IPv4 = N/A TST IPv6 = N/A	MAX IF # = 0 MGN IPv4 = N/A MNG IPv6 = N/A TST IPv4 = N/A TST IPv6 = N/A	MAX IF # = 0 MGN IPv4 = N/A MNG IPv6 = N/A TST IPv4 = N/A TST IPv6 = N/A

Cloud Peak is a software product composed of two elements that communicate to rigorously validate the SUT. The Cloud Peak Appliance is used for management functions and for hosting test methodologies and test results. The Cloud Peak Workload generates the test stimuli used to characterize the SUT performance. One single Cloud Peak Appliance can manage one or more Cloud Peak Workloads. We recommend deploying the Cloud Peak Appliance on a separate server to prevent interference with the SUT performance and the test results.

	Cloud Peak appliance		Cloud Peak workload	
File format **File size**	OVA	= 2.81 GB	OVA	= 1.09 GB
	QCOW2	= 2.34 GB	QCOW2	= 1.07 GB
	AMI	= 2.57 GB	AMI	= 1.08 GB
	Docker	= N / A	Docker	= 0.12 GB
vCPU	4 vCPUs		Min. 1 vCPU	
Memory	8 GB RAM		Min. 1 GB RAM	
Disk	40 GB		Min. 8 GB	
vNIC	1		Min. 1	
Software packaging	Image formats for the Cloud Peak appliance – OVA / QCOW2 / AMI Image formats for the Cloud Peak workload (simulated VNFs) – OVA / QCOW2 / AMI Image formats for the Cloud Peak workload (simulated CNFs) – Docker containers			

Cloud Peak validates virtual cloud infrastructures with the characteristics described in the following table:

SUT component	Supported value				
SUT type	OpenStack	VMware vCenter	Kubernetes	Amazon AWS	Others
SUT version	Liberty, Mitaka, Newton, Ocata, Pike, Queens, Rocky, Stein, Train, Ussuri, Victoria, Wallaby, Xena, Yoga	VMware vCenter 6.0 VMware vCenter 6.5 VMware vCenter 6.7 VMware vCenter 7.0	v1.16 v1.17 v1.18 v1.19 v1.20 v1.21 v1.22 v1.23 v1.24	Amazon Region Amazon Outposts Amazon Wavelength	N / A
SUT services	Glance, Heat, Keystone, Neutron, Nova	HTTPS	Kubectl, Kube-API-Server, Kube-Controller, Kube-Proxy, Kube-Scheduler, Calico CNI, Core DNS	Amazon CFT Amazon EC2 Amazon IAM Amazon S3 Amazon VPC	N / A
SUT hypervisor	KVM ESXi	VMware ESXi	Docker Engine CRI-O ContainerD	N / A	N / A

Cloud Peak is available in the Amazon AWS Marketplace under the Keysight Technologies product portfolio and can be used by subscribing to the Keysight Cloud Peak and Keysight Cloud Peak Workload components. The product can be operated under the Bring-Your-Own-License (BYOL) model and can be used to benchmark the performance of various AWS Instance Type across different AWS Regions, AWS Outposts, AWS Wavelength, or AWS EKS platforms.

Technology solutions

Visit www.keysight.com for more information on our virtualization solutions

Cloud Peak	Virtualized cloud infrastructure benchmarking
IxNetwork Virtual Edition (VE)	Virtualized network performance testing
IxLoad Virtual Edition (VE)	Virtualized multiplay services testing
BreakingPoint Virtual Edition (VE)	Virtualized application and security testing

Ordering information

Cloud Peak has an all-inclusive licensing model and it is available as a subscription or a perpetual license. There are three-licensing tiers (Small / Medium / Large), which control the maximum size of the simulated workload (defined based on the total number of CPU cores and amount of memory required). In addition, depending on tier type, each license allows up to 2, 4, or 8 parallel executions. Optional execution licenses can be ordered to increase the number of parallel executions allowed by the system.

Part number	Description
Subscription licenses	
939-9851	**IXIA Cloud Peak, NFVI Tier-SMALL Floating Subscription License** Includes all NFVI test suites (Compute / Network / Storage / Custom), all configuration parameters, all reporting capabilities, and quantity 2 execution licenses. Enables parallel execution of up to 2 workload simulations, each with capacity up to the license limit (32 vCPU / 128 GB RAM / SMALL) per test. Must purchase higher tiers to increase workload execution limits. Multiple execution licenses do not increase the size of the simulated workload. Requires license term to be specified (must be purchased in multiples of years, list price is per unit per year).
939-9852	**IXIA Cloud Peak, NFVI Tier-MEDIUM Floating Subscription License** Includes all NFVI test suites (Compute / Network / Storage / Custom), all configuration parameters, all reporting capabilities, and quantity 4 execution licenses. Enables parallel execution of up to 4 workload simulations, each with capacity up to the license limit (512 vCPU / 2048 GB RAM / MEDIUM) per test. Must purchase higher tiers to increase workload execution limits. Multiple execution licenses do not increase the size of the simulated workload. Requires license term to be specified (must be purchased in multiples of years, list price is per unit per year).
939-9853	**IXIA Cloud Peak, NFVI Tier-LARGE Floating Subscription License** Includes all NFVI test suites (Compute / Network / Storage / Custom), all configuration parameters, all reporting capabilities, and quantity 8 execution licenses. Enables parallel execution of up to 8 workload simulations, each with capacity up to the license limit (8192 vCPU / 32768 GB RAM / Tier- LARGE) per test. Multiple execution licenses do not increase the size of the simulated workload. Requires license term to be specified (must be purchased in multiples of years, the list price is per unit per year).
Optional, subscription execution licenses	
939-9861	**IXIA Cloud Peak, optional NFVI EXEC-SMALL Floating Subscription License** Includes quantity 2 execution licenses. Enables parallel execution of up to 2 workload emulations, each with capacity up to the license limit (32 vCPU / 128 GB RAM / SMALL). REQUIRES license term to be specified (must be purchased in multiples of years, list price is per unit per year). Requires prior purchase of Cloud Peak, NFVI Tier-SMALL Floating Subscription License (939-9851).
939-9862	**IXIA Cloud Peak, optional NFVI EXEC-MEDIUM Floating Subscription License** Includes quantity 2 execution licenses. Enables parallel execution of up to 2 workload emulations, each with capacity up to the license limit (512 vCPU / 2048 GB RAM / MEDIUM). REQUIRES license term to be specified (must be purchased in multiples of years, list price is per unit per year). Requires prior purchase of Cloud Peak, NFVI Tier-MEDIUM Floating Subscription License (939-9852).
939-9863	**IXIA Cloud Peak, optional NFVI EXEC-LARGE Floating Subscription License** Includes quantity 2 execution licenses. Enables parallel execution of up to 2 workload emulations, each with capacity up to the license limit (8192 vCPU / 32768 GB RAM / LARGE). REQUIRES license term to be specified (must be purchased in multiples of years, list price is per unit per year). Requires prior purchase of Cloud Peak, NFVI Tier-LARGE Floating Subscription License (939-9853).
Perpetual licenses	
939-9856	**IXIA Cloud Peak, NFVI Tier-SMALL Floating Perpetual License** Includes all NFVI test suites (Compute / Network / Storage / Custom), all configuration parameters, all reporting capabilities, and quantity 2 execution licenses. Enables parallel execution of up to 2 workload simulations, each with capacity up to the license limit (32 vCPU / 128 GB RAM / SMALL) per test. Must purchase higher tiers to increase workload execution limits. Multiple execution licenses do not increase the size of the simulated workload.
939-9857	**IXIA Cloud Peak, NFVI Tier-MEDIUM Floating Perpetual License** Includes all NFVI test suites (Compute / Network / Storage / Custom), all configuration parameters, all reporting capabilities, and quantity 4 execution licenses. Enables parallel execution of up to 4 workload simulations, each with capacity up to the license limit (512 vCPU / 2048 GB RAM / MEDIUM) per test. Must purchase higher tiers to increase workload execution limits. Multiple execution licenses do not increase the size of the simulated workload.
939-9858	**IXIA Cloud Peak, NFVI Tier-LARGE Floating Perpetual License** Includes all NFVI test suites (Compute / Network / Storage / Custom), all configuration parameters, all reporting capabilities, and quantity 8 execution licenses. Enables parallel execution of up to 8 workload simulations, each with capacity up to the license limit (8192 vCPU / 32768 GB RAM / LARGE) per test. Multiple execution licenses do not increase the size of the simulated workload.
Optional, perpetual execution licenses	
939-9866	**IXIA Cloud Peak, optional NFVI EXEC-SMALL Floating Perpetual License** Includes quantity 2 execution licenses. Enables parallel execution of up to 2 workload emulations, each with capacity up to the license limit (32 vCPU / 128 GB RAM / SMALL). Requires prior purchase of Cloud Peak, NFVI Tier-SMALL Floating Perpetual License (939-9856).
939-9867	**IXIA Cloud Peak, optional NFVI EXEC-MEDIUM Floating Perpetual License** Includes quantity 2 execution licenses. Enables parallel execution of up to 2 workload emulations, each with capacity up to the license limit (512 vCPU / 2048 GB RAM / MEDIUM). Requires prior purchase of Cloud Peak, NFVI Tier-MEDIUM Floating Perpetual License (939-9857).
939-9868	**IXIA Cloud Peak, optional NFVI EXEC-LARGE Floating Perpetual License** Includes quantity 2 execution licenses. Enables parallel execution of up to 2 workload emulations, each with capacity up to the license limit (8192 vCPU / 32768 GB RAM / LARGE). Requires prior purchase of Cloud Peak, NFVI Tier-LARGE Floating Perpetual License (939-9858).