128 Technology

128T Session Smart™ Routers

Scale & Performance

Tolly Report #220100

December 2019

Contents

1 Executive Summary

Enterprises recognize the significant economic and performance benefits of migrating away from private WANs to high-speed, "cloud-native" networks that leverage the internet, cloud infrastructure and, ultimately, 5G cellular service. New networks require new approaches. 128 Technology delivers a routing platform that is built from the ground up to improve your network and business by offering a solution that: reduces costs, is built with a focus on session-based routing, and provides zero-trust security with no tunneling.

128 Technology commissioned Tolly to independently evaluate the scalability, resilience and operational features of its 128T Session Smart Router solution. Tests included head-end and branch scalability, resilience in the face of numerous faults including link degradation, link failure, and router failure as well as operational tests that exercised upgrade/rollback and common graphical user interface (GUI) and command line tasks in nearly 20 separate test cases across nine scenarios. Tolly notes that 128 Technology is the first SD-WAN vendor to prove scalability in a Tolly test. 128T demonstrates its commitment to its customers by proving its formidable capabilities in a third-party test.

The Bottom Line

Head-End Router Scaling: Supported over 120,000 user sessions and 4,000 router peer sessions across Gigabit Ethernet on an Intel Xeon platform[1].

Branch Router Scaling: Supported 2,000 user sessions across Gigabit Ethernet & T1 on an Intel Atom platform[1].

Head-End Router Resiliency: Demonstrated rapid recovery and operational continuity in a dual-node configuration across a range of failure conditions that included (for each node): Service restart, node reboot and network interface failures .

Branch Router Resiliency: Demonstrated rapid recovery and operational continuity in a dual-node configuration across a range of failure conditions that included (for each node): Service restart, node reboot, and network interface failures.

Link Resiliency: Migrated traffic to alternate links when packet loss and/or latency exceeded configured levels. Proved quality-of-service enforcement to ensure delivery of mission-critical traffic should a link become oversubscribed.

Operations, Maintenance & Cloud Deployment: Demonstrated effective upgrade/rollback of routers and management system software, configuration backup/restore, intuitive GUI for key status indicators and functions, multi-instance command line implementation and cloud-hosting for router deployment.

Scaling

The scaling scenarios included both head-end and branch router scenarios. It is important to know that the router platform can handle sufficient user and router peer sessions to support required operations. The 128T solution passed both scaling tests.

The head-end scenario involved initiating more than 120,000 user sessions with two test branch routers dividing the traffic across the two. Simultaneously, the head-end router set up 4,000 peer sessions to the test routers simulating the routing

[1] 128 Technology notes that up to 1M session counts can be supported, higher session counts are supported based on the hardware and network topology definitions.

connections present in actual Session Smart network. To demonstrate the ability to sustain this load over time the test was run overnight resulting in a run-time of some 17 hours. Engineers noted trivial packet loss that amounted to less than 0.001% of packets sent.

The remote branch office scenario was run in a similar fashion. This scenario involved initiating more than 2,000 user sessions between the head-end and a branch office router. Here, too, engineers noted trivial packet loss that amounted to less than 0.001% of packets sent.

Resiliency

The resiliency scenarios included both head-end and branch router failure/recovery scenarios. Additionally, tests included link resiliency and quality of service scenarios. It is important to know that the two-node router platform can recover from a variety of outages. Additionally, it is important to know that the router can respond to link degradation conditions and prioritize traffic. The 128T solution passed all resiliency tests.

The head-end and remote branch scenarios were identical and involved deliberately causing software, hardware and system failures in the primary and secondary nodes in succession and confirming that the head-end and branch routers remained operational and traffic flowed with, at most, a brief interruption. The failures introduced were: node routing service shutdown, full node shutdown (reboot) and network interface physical removal.

The first link scenario focused on confirming that the 128T solution could respond to degraded service levels, as specified by the customer, and migrate traffic to an alternate link. One test involved demonstrating this capability first when the link latency (delay) exceeded 300 ms and then when latency exceeded 600 ms. The other degradation test involved packet loss. This test was to show migration to an alternate link when packet loss exceeded 1%, then 5% and finally 10%.

The final resiliency test was run to illustrate how the 128T solution could allocate bandwidth according to customer-specified traffic priorities of high, medium, low and best effort. This quality-of-service feature is essential for situations when there is more demand on the WAN link than there is bandwidth available (i.e. when the WAN is oversubscribed). This test was run twice with two different sets of bandwidth allocations across the four traffic classes.

Operations, Maintenance & Cloud Deployment

Essential to any organization is knowing that the router infrastructure supports key operational and maintenance capabilities. Tolly engineers validated a range of upgrade/rollback, backup/restore and graphical user interface (GUI) and command line interface (CLI) operator functions essential to ongoing operations. Finally, engineers verified that the 128T routers can be deployed on a cloud platform, in this case AWS. The 128T solution passed all the tests.

The upgrade/rollback tests demonstrated that the head-end router, branch routers and the Conductor management system could not only be upgraded to a new version but also rolled back to the prior version should that need arise.

The configuration backup/restore tests illustrated that the network configuration could be easily and quickly backed up and restored easily and quickly as well.

The Conductor system management operations tests exercised both the GUI and command line interfaces. The tests demonstrated that common commands and functions could be executed with good response time. (Some functions such as validating and committing an update can take several minutes to complete as the information is checked for validity.) The tests also illustrated support for four simultaneous command line administrative sessions.

The final test verified that 128T routers can be deployed on a cloud platform. In this test, the Amazon Web Services platform was used but 128T notes that other cloud platforms are supported as well.

The following table provides an overview of all the tests conducted in this Tolly test.

Test Results Overview

Focus Area	Scenario	Description	Result
Scaling	Head-end Router Session & Peer Scale	Demonstrate sustained (12+hours) support for 120,000 user sessions and 4,000 router peer sessions with <0.001% packet loss.	**PASS**
	Remote Branch Router Session Scale	Demonstrate sustained (12+hours) support for 2,000 user sessions with <0.001% packet loss.	**PASS**
Resiliency	Head-end Router Failover/Recovery	Demonstrate failover/recovery of each node in a two-node system under the following conditions: 1) Service restart, 2) node reboot, 3) network cable failure.	**PASS**
	Remote Branch Router Failover/Recovery	Demonstrate failover/recovery of each node in a two-node system under the following conditions: 1) Service restart, 2) node reboot, 3) network cable failures.	**PASS**
	Link Degradation: Latency & Packet Loss	Demonstrate that traffic could be automatically redirected to an alternate link when latency and/or packet loss exceeded customer-configured levels.	**PASS**
	Link Quality-of-Service: Prioritized Traffic Classes	Demonstrate that service levels could be applied to traffic to allocate available link bandwidth in accordance with customer-configured requirements.	**PASS**
Operations, Maintenance & Cloud Deployment	Upgrade/Rollback: Head-end and branch routers, Conductor management system	Demonstrate the software upgrade process for each system along with the capability to rollback to the prior system version. Demonstrate simultaneous upgrade of multiple branch systems.	**PASS**
	Configuration Backup & Restore	Demonstrate the capability to backup the router configuration information along with restoring the prior system configuration.	**PASS**
	Conductor Management System Operations	Demonstrate basic GUI functions, single-user command line functions as well as simultaneous, four-user command line functions.	**PASS**
	Router Cloud Deployment	Demonstrate that the 128 Technology routers can be deployed on a cloud platform.	**PASS**

The remainder of this report will summarize additional details of the tests discussed above. For those requiring even more information and/or needing to run the tests themselves, there is a more detailed appendix providing itemized testing procedures that is available from 128 Technology Inc.

2 Scaling

The 128 Technology router solution runs on commercial, off-the shelf (COTS) hardware running Linux. The hardware details of the management and router platforms tested are found in the following table.

System Specifications

System	CPU Type	CPU Specifications	Memory	Notes
Conductor (management system)	Intel Xeon D1548 2.00GHZ	Single socket chip: 8 cores x 2 threads = 16 CPUs	64GB	Hyperthreading enabled
Head-end Router	Intel Xeon D1548 2.00GHZ	Single socket chip: 8 cores x 1 threads = 8 CPUs	64GB	Hyperthreading disabled
Branch Router	Intel Atom C2518 1.74GHz	Single socket chip: 4 cores x 1 threads = 4 CPUs	16GB	
Note: For this test all systems ran on CentOS 7.5 (Linux) plus known critical CVE updates.				

Head-end Router Session & Peer Scale

Goal

Demonstrate sustained (12+hours) support for 120,000 user sessions and 4,000 router peer sessions with <0.001% packet loss.

Results

PASS. The test ran for approximately 17 hours. The system was able to maintain all user and router peer sessions. Over the run of the test only a few packets were lost.

Methodology

This test used Topology T2 (See Topologies section for a network diagram.)

Bidirectional traffic simulating 120,000 user sessions was generated using a Spirent traffic generator and multiple systems/routers were used to simulate 1,000 routers having 4,000 router peer paths. Spirent interfaces were connected to 10GbE links which were, in turn, connected to Gigabit Ethernet links in both the head-end router and a pair of branch routers. Unique session pairs were generated by automatically incrementing UDP/IP address information. Approximately 60,000 sessions were directed across each of the two branch routers simulating the a total of 1,000 routers/4,000 peer paths. The frame size used was 1300 bytes with a bidirectional load of 1.02Gbps. Engineers confirmed that traffic was being processed without packet loss and that no system alarms were generated. Traffic was then allowed to run for at least 12 hours to show sustainability. At the end of the test, engineers confirmed the session counts and that traffic loss was minimal (below 0.001% of the total).

The solution was deemed to pass the test if sessions were sustained with minimal packet loss for the duration of the test and no alarms were generated.

Branch Router Session Scale

Goal

Demonstrate sustained (12+hours) support for 2,000 user sessions with <0.001% packet loss.

Results

PASS. The test ran for approximately 17 hours. The system was able to maintain all user sessions. Over the run of the test only a few packets were lost.

Methodology

This test used Topology T2

Bidirectional traffic simulating user sessions was generated by a Spirent traffic generator on branch nodes peered with the head end receiving the traffic. The Spirent generator was connected to 10GbE links which were, in turn, connected to Gigabit Ethernet links in both the head-end router and a pair of branch routers. Unique session pairs were generated by automatically incrementing UDP/IP address information. Approximately 2,000 sessions were directed across the branch router. The frame size used was 250 bytes with a bidirectional load of 400Mbps driving the T1 interfaces on the router. Engineers confirmed that traffic was being processed without packet loss and that no system alarms were generated. Traffic was then allowed to run for at least 12 hours to show sustainability. At the end of the test, engineers confirmed the session counts and that traffic loss was minimal (below 0.001% of the total).

The solution was deemed to pass the test if sessions were sustained with minimal packet loss for the duration of the test and no alarms were generated.

3 Resiliency

Head-end Router Failover/Recovery

Goal

Demonstrate that a two-node, head-end router implementation can maintain operations should the following faults be encountered with either node: 1) Router node service failure (restart), 2) Router node system failure (reboot), 3) network interface failure (cable removed).

Results

PASS. In all scenarios the traffic failed over to the alternate node within several seconds. During the test, engineers noted a few minor alarms but those alarm conditions did not affect traffic flow.

Methodology

This test used Topology T2.

A Spirent traffic generator was used to generate traffic that flowed across the head-end router to/from the branch router. For all tests, statistics and alarms were monitored using the Conductor management system. The Spirent traffic statistics were monitored both to confirm that the failure condition was initiated as well as to confirm that the traffic flow recovered via the secondary node. The tests were run both on the primary (Node 1) and the secondary (Node 2) router nodes.

The router service failure was triggered by using a management command to restart the router service on the node under test.

The router node system failure was triggered by using a management command to reboot the node under test.

The network interface failure was triggered by manually disconnecting the active Ethernet interface from the node under test.

Branch Router Failover/Recovery

Goal

Demonstrate that a two-node, branch router implementation can maintain operations should the following faults be encountered with either node: 1) Router node service failure (restart), 2) Router node system failure (reboot), 3) network interface failure (cable removed).

Results

PASS. In all scenarios the traffic failed over to the alternate node within several seconds. During the test, engineers noted a few minor alarms but those alarm conditions did not affect traffic flow.

Methodology

This test used Topology T2.

A Spirent traffic generator was used to generate traffic that flowed across the head-end router to/from the branch router. For all tests, statistics and alarms were monitored using the Conductor management system. The Spirent traffic statistics were monitored both to confirm that the failure condition was initiated as well as to confirm that the traffic flow recovered via the secondary node. The tests were run both on the primary (Node 1) and the secondary (Node 2) router nodes.

The router service failure was triggered by using a management command to restart the router service on the node under test.

The router node system failure was triggered by using a management command to reboot the node under test.

The network interface failure was triggered by manually disconnecting the active Ethernet & T1 interfaces from the node under test.

Link Degradation: Latency & Packet Loss

Goal

Demonstrate service integrity by automatically migrating traffic to an alternate link when latency (delay) and/or packet loss conditions exceed degradation levels specified by the customer via the router system configuration.

Latency and packet loss were tested separately in order to isolate test conditions but 128 Technology notes that both can be configured simultaneously. The latency degradation test was run twice using different trigger settings: 1) 300ms, and 2) 600ms. The packet loss degradation test was run three times using different trigger settings: 1) 1%, 2) 5%, and 3) 10%.

Results

PASS. In all scenarios the traffic failed over to the alternate link within several seconds of the latency/pack loss impairment being introduced on the link.

Methodology

This test used Topology T5.

A Spirent traffic generator was used to generate traffic that flowed across the head-end router to/from the branch router. The NetEm network emulation program was used to generate latency and (separately) packet loss impairments into the link at the levels noted above. The impairments generated were deliberately set slightly above the trigger levels in each case.

Link Quality-of-Service: Prioritized Traffic Classes

Goal

Demonstrate the capability of defining and enforcing multiple traffic queues (also known as QoS or traffic engineering) by allocating a specified percentage of link bandwidth to a particular class. This allows the customer to specify priority for different types of traffic should a situation occur where traffic demands exceed available link bandwidth (i.e. if the link becomes oversubscribed). Traffic was classified as one of the following classes: 1) high, 2) medium, 3) low, or 4) best-effort. The QoS test was run twice with different allocations across the classes to illustrate the flexibility of configuration.

The test was run first with the following bandwidth percentage allocations:

1) High = 50, 2) Medium=25, 3) Low=24, 4) Best-effort=1.

The test was run a second time with different bandwidth percentage allocations:

1) High = 60, 2) Medium=25, 3) Low=14, 4) Best-effort=1.

Results

PASS. In both scenarios the traffic measured across the link matched the percentage definitions exactly or very closely. In the first scenario, the results differed by no more than 0.89% of bandwidth between configured and actual bandwidth. In the second scenario, the results differed by no more than 01.35% of bandwidth between configured and actual bandwidth..

Methodology

This test used Topology T3.

A Spirent traffic generator was used to generate traffic that flowed across the head-end router to/from the branch router. Four traffic flows were set up that corresponded to each of the traffic classes to be tested.

The 128 Technology router configuration traffic profile was updated to reflect the desired allocation for the given test run. Traffic was initiated and run until approximately 1,000 packets of each traffic class had been generated. Engineers then documented the receive counts for each of the traffic types and calculated the test results to match against the configured bandwidth allocations.

4 Operations, Maintenance & Cloud

Upgrade/Rollback: Head-end and Branch Routers, Conductor Management System

Goal

Demonstrate the capability of managed upgrades to new software versions of the router and management software without disrupting operations as well as proving the capability of rolling back to the previous version should that be desired, again without disrupting operations. This test was run individually on: 1) Head-end router, 2) branch router, and 3) the Conductor management system. Additionally, a test case was run upgrading three branch routers simultaneously.

Results

PASS. All scenarios.

Methodology

This test used Topology T1 for the Conductor upgrade and T2 for the router upgrades.

The process was very similar for all scenarios. The current version of the system was verified. The systems were confirmed to be operational as appropriate. Any alarm states existing prior to the upgrade were documented.

The upgrades were performed via the Conductor user interface. Engineers verified the resulting software version level to confirm that it was a newer version. Engineers verified that the system was operational on the new release.

The Conductor rollback used the 128T install utility. The router rollback functions were performed via the Conductor PCLI (command line) "rollback" command.

After completion, engineers verified that the system software version matched the version existing prior to the upgrade. Engineers verified that the system was operational on the rolled-back release.

For the multiple branch router upgrade scenario, engineers used the Conductor GUI to select and execute the upgrade process on three branch routers simultaneously.

Configuration Backup & Restore

Goal

Demonstrate the capability of configuration backup (export) and configuration restore (import) of the system configuration.

Results

PASS.

Methodology

This test used Topology T2.

Using the Conductor GUI, engineers exported, downloaded and reviewed the existing configuration. This initial configuration would be used for the restore step of this test. Engineers then made an update to the existing configuration, validated that change and committed that change. After completion, engineers verified that the change was present in the live configuration.

To demonstrate the import capabilities, engineers imported the original configuration, overlying the changed configured. After validating and committing that configuration, engineers verified that the restore had been successful by noting that the prior changes were no longer present.

Conductor Management System Operations

Goal

Demonstrate responsiveness for basic GUI operations and more granular command line operations using a single user as well as four users sessions.

Results

PASS. The response time for the various command scenarios was deemed reasonable for each task being performed. These results are subjective but Tolly found the response times for Conductor GUI, single-user command and line and four simultaneous command lines to be very usable.

Conductor GUI. New devices/services were added in approximately 90-100 seconds. Changes were made in just two seconds. The validation and commit of the the changes are combined in the Conductor GUI and require approx. 11 minutes to complete.

Command Line, Single-user. Commands relating to a single router responded very quickly generally under one second to less than two seconds. Network wide commands could require more time to gather and present information to the user. "Show alarms router all" that displayed all systems with active alarms required 56 seconds to complete. Displaying the running configuration candidate required 74 seconds and showing the config candidate required 65 seconds. The only command that required significantly more than a minute to complete was a network-wide "show sessions" command. This required five minutes and 14 seconds to gather and display the information.

Command Line, Four Simultaneous Users. Adding new devices and interfaces required approximately 90 seconds each. Updating existing devices/interfaces required approx. 20 seconds each. Showing the running configuration required approx. 2 minutes and 30 seconds where showing running configurations and displaying stats required one minute and 21 seconds. As expected committing the configuration required the most time at just over five minutes. Finally, the configuration import was completed in one minute and 15 seconds.

Methodology

This test used Topology T1.

For the Conductor GUI, the scenario included logging and checking the initial dashboard that contains info about routers, alarms, configuration and a health overview. Basic updates were made including adding new devices and services as well as changing existing devices and services. These were validated and committed.

For the command line, single-user test the response time for 15 different commands were validated. These included a range of "show" commands for alarms, assets, fibs, arp, etc. at both the system level and for single routers.

In the four-user tests, four command line sheets were opened simultaneously. Various tasks were performed including adding and changing devices and services along with validating and committing changes and configuration import.

Router Cloud Deployment

Goal

Demonstrate that the 128 Technology router could be deployed on a public cloud platform.

Results

PASS.

Methodology

This test used Topology T4.

For this test, the Amazon Web Services (AWS) platform was used. Engineers deployed two 128T routers in the AWS environment. Once deployed, engineers used the Spirent test tool to generate a traffic stream across the two routers in order to illustrate that both were operational.

5 Test Topologies

This section contains illustrations of the various network topologies used throughout the testing. These topologies are referenced by individual test cases.

T1 Conductor Use Cases

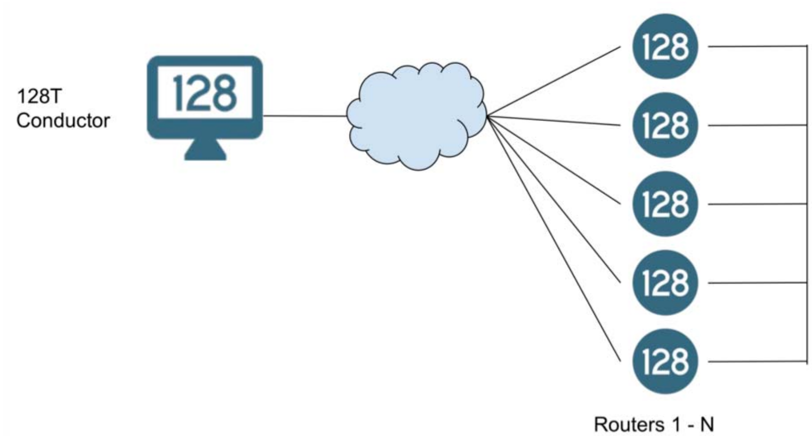

T2 Head-end/Branch Router Use Cases

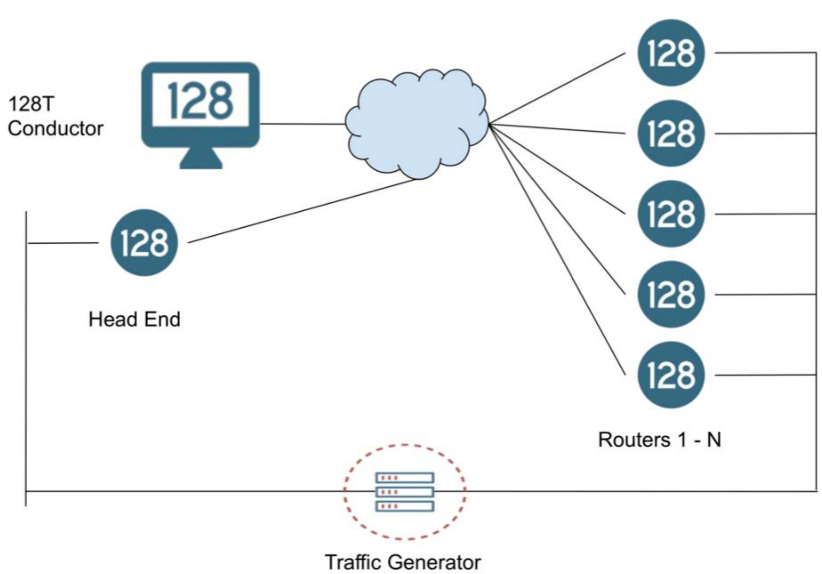

T3 Traffic Engineering Use Cases

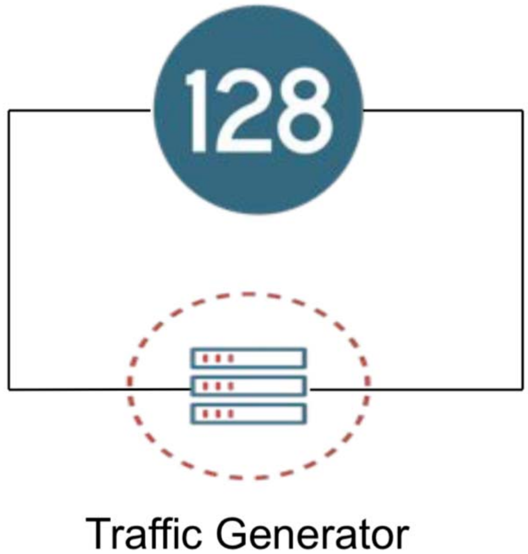

T4 Cloud & Router Cases

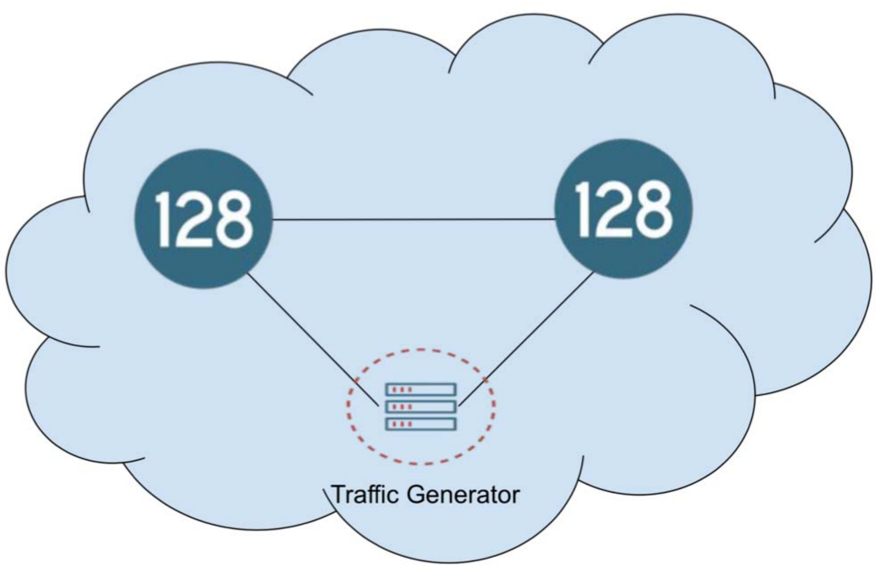

T5 Link Loss & Latency Test Cases

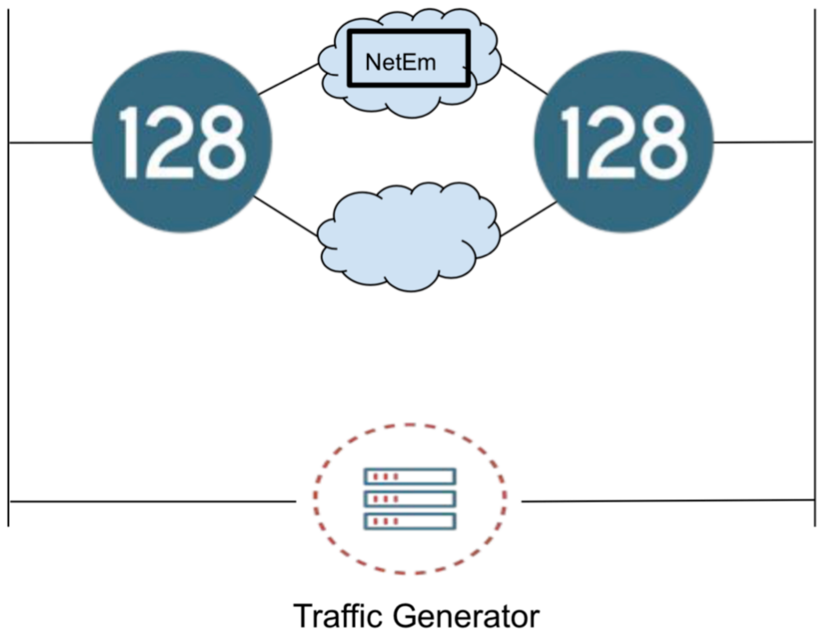

About Tolly…

The Tolly Group companies have been delivering world-class IT services for 30 years.
Tolly is a leading global provider of third-party validation services for vendors of IT products, components and services.
Tolly also assists medium-sized businesses and large enterprises evaluate, benchmark and select IT products for deployment.

You can reach the company by email at sales@tolly.com, or by telephone at
+1 561.391.5610.

Visit Tolly on the Internet at:
http://www.tolly.com

Terms of Usage

This document is provided, free-of-charge, to help you understand whether a given product, technology or service merits additional investigation for your particular needs. Any decision to purchase a product must be based on your own assessment of suitability based on your needs. The document should never be used as a substitute for advice from a qualified IT or business professional. This evaluation was focused on illustrating specific features and/or performance of the product(s) and was conducted under controlled, laboratory conditions. Certain tests may have been tailored to reflect performance under ideal conditions; performance may vary under real-world conditions. Users should run tests based on their own real-world scenarios to validate performance for their own networks.

Reasonable efforts were made to ensure the accuracy of the data contained herein but errors and/or oversights can occur. The test/audit documented herein may also rely on various test tools the accuracy of which is beyond our control. Furthermore, the document relies on certain representations by the sponsor that are beyond our control to verify. Among these is that the software/hardware tested is production or production track and is, or will be, available in equivalent or better form to commercial customers. Accordingly, this document is provided "as is", and Tolly Enterprises, LLC (Tolly) gives no warranty, representation or undertaking, whether express or implied, and accepts no legal responsibility, whether direct or indirect, for the accuracy, completeness, usefulness or suitability of any information contained herein. By reviewing this document, you agree that your use of any information contained herein is at your own risk, and you accept all risks and responsibility for losses, damages, costs and other consequences resulting directly or indirectly from any information or material available on it. Tolly is not responsible for, and you agree to hold Tolly and its related affiliates harmless from any loss, harm, injury or damage resulting from or arising out of your use of or reliance on any of the information provided herein.

Tolly makes no claim as to whether any product or company described herein is suitable for investment. You should obtain your own independent professional advice, whether legal, accounting or otherwise, before proceeding with any investment or project related to any information, products or companies described herein. When foreign translations exist, the English document is considered authoritative. To assure accuracy, only use documents downloaded directly from Tolly.com.

220100—wt-2019-12-09-VerGbis