

## D-Link DXS-5000-54S Open Network Switch

### 10GbE, Performance, Feature & Network Resilience Evaluation

### EXECUTIVE SUMMARY

Open Network Switches provide network architects with new possibilities with respect to their choices in network hardware and network operating system software.

D-Link commissioned Tolly to evaluate the performance and functionality of its DXS-5000-54S Open Network Switch powered by the Pluribus Netvisor ONE operating system. Testing encompassed throughput performance and latency, resilience in a multi-switch, leaf-spine configuration as well as extensive feature/function testing.

Tolly engineers verified wire-speed L2 throughput and latency across two ports of a 10GbE switch, verified rapid reconvergence times in multi-switch failover scenarios and validated nearly 50 features in areas including L2/L3, automation, VXLAN, QoS, DHCP, mirroring, analytics and management.

### THE BOTTOM LINE


The D-Link switch provides:

- 1 Support for 10/40 GbE LANs
- 2 Support for Pluribus Networks Netvisor ONE open network switch operating system
- 3 Robust L2 support including RSTP/MSTP, VLANs, trunking, LLDP and jumbo MTUs
- 4 Robust L3 support including dual-stack IPv4/IPv6 functionality, VRRP, BGP and OSPF

### Layer 2 10 Gigabit Ethernet RFC 2544 Throughput

Zero-loss Across Two Ports  
(as reported by Xena Valkyrie2544 v2.53)

D-Link DXS-5000-54S Open Network Switch



Source: Tolly, September 2018

Figure 1



## Solution Highlights

Among the areas evaluated, there were several that illustrate important benefits of the Pluribus Networks OS. Simply presenting these as features list without context fails to convey important benefits and capabilities. Thus, this report will instead provide a solution-level discussion for these key areas. All capabilities discussed were demonstrated and tested in the course of the testing.

### Adaptive Cloud Fabric

A key element of the Pluribus solution is what it refers to as its Adaptive Cloud Fabric. In essence, multiple Pluribus Netvisor ONE switches work together forming a distributed switch fabric that functions without a controller.


The distributed nature of the fabric provides many benefits. Among the capabilities illustrated were:

- 1) Single point of management to reduce time to deploy and configure, simplify software upgrades, rapid troubleshooting and reduce human configuration errors.
- 2) Simple and cost effective to deploy: no proprietary controllers and protocols make the Adaptive Cloud Fabric easy to insert into any brownfield network based on standard L2/L3 protocols and any topology.
- 3) Easy to extend across multiple sites and geographies: the fabric control plane can easily be transported across any core/wan network, making it possible to extend the Adaptive Cloud Fabric single pane of glass across multiple locations.
- 4) The ability to run the fabric in-band and/or out-of-band eliminates all single points of failure of centralized controller-based architectures (single switch management interface and single out-of-band management network), making it an ideal fit for mission critical applications.

**D-Link Corporation**


**DXS-5000-54S**

**Features, Performance & Network Resilience**



*Tested September 2018*

- 5) The Adaptive Cloud Fabric distributed architecture is designed for the rapid failover and high-availability even in environments with third party switches and routers.
- 6) Failed nodes can be replaced and fabric will heal automatically recovering old configuration.





**D-Link Open Network Switching Platform/Pluribus Netvisor ONE Operating System  
Tolly Certified Features - Part 1**

<b>Adaptive Cloud Fabric</b>		<b>VXLAN</b>	
✓	Self-healing	✓	VXLAN
✓	ACLs	✓	Auto-Tunnels
✓	Scalability	✓	VXLAN Bridging
<b>Control Plane Traffic Protection (CPTP)</b>		✓	VXLAN Routing
✓	Fabric protection, Hog Protection	✓	Anycast v4/v6/Distributed VRF
<b>Automation/API</b>		<b>Quality-of-Service</b>	
✓	Configuration Automation with Ansible/Unum	✓	Queuing
✓	REST API	✓	Marking
<b>LAN Topologies</b>		✓	Minimum Bandwidth Guarantee
✓	10/40GbE	✓	Policer
<b>L2 Functionality</b>		<b>DHCP</b>	
✓	RSTP/MSTP	✓	DHCP Relay (Dual-stack IPv4/v6)
✓	VLAN and Trunking (Link Aggregation/LAG)	✓	DHCP Snooping (Track/Filter)
✓	Multi-Chassis LAG (vLAG)		
✓	EtherChannel		
✓	BPDU Filter		
✓	SVI (L2/L3)		
✓	LLDP		
✓	Jumbo MTU		
✓	ARP & Broadcast Suppression		
<b>L3 Functionality (Dual-stack IPv4/v6)</b>			
✓	Static Routes		
✓	VRRP Sessions		
✓	BGP Routes		
✓	BGP Neighbors		
✓	BGP BFD Neighbors		
✓	OSPF Routes		
✓	OSPF Neighbors		
✓	vRouters		

Source: Tolly, September 2018

Table 1



Tests demonstrated ease-of-use and access via CLI, GUI and/or RESTful API. Automation via Ansible was also demonstrated.

## Control Plane Traffic Protection (CPTP)

CPTP refers to the capability of the switch to protect the flow of traffic on the control plane. This improves switch reliability and availability. Tolly engineers verified three levels of protection: 1) Overall CPU bandwidth, 2) Dedicated queues for critical flows, and 3) DDoS protection from offending flows. This last feature, also known as "hog protection," protects from too much traffic from a given host overloading a critical control queue. Hog protection achieves this by either redirecting traffic to lower bandwidth (100packets/s) hog queues or by dropping them. See Figure below.

## VXLAN

Virtual Extensible LAN (VXLAN) is a standards-based network virtualization technology. Defined in IETF RFC 7348, it allows a Layer 2-like overlay across IP networks. VXLAN lets network architects combine simplicity with scalability. Tunnels are auto-created using the VTEP

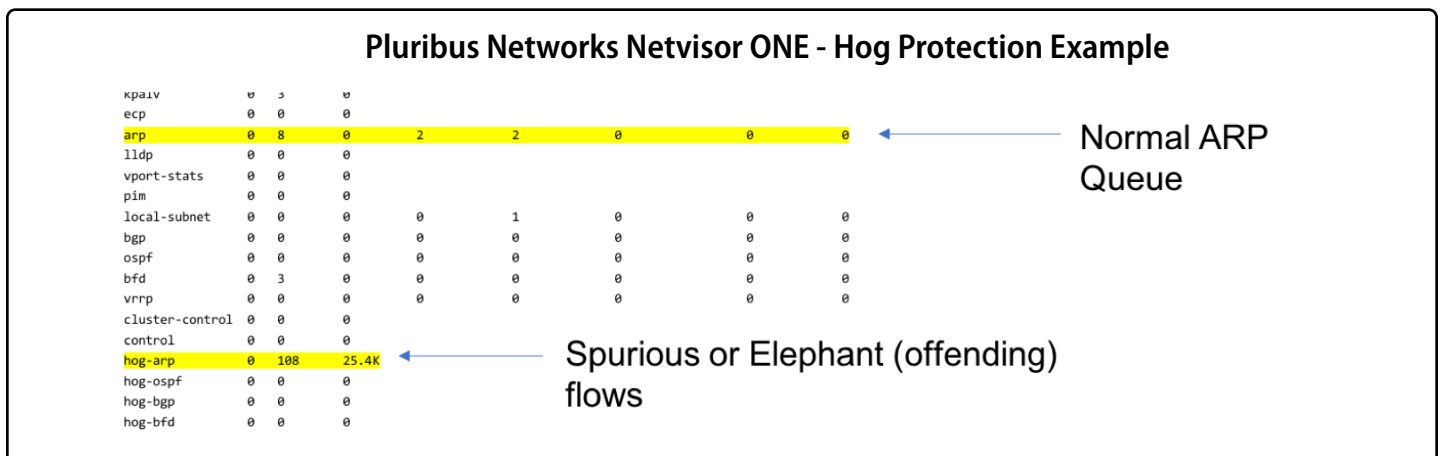
D-Link Open Network Switching Platform/Pluribus Netvisor ONE Operating System	
Tolly Certified Features - Part 2	
<b>Packet Mirroring</b>	
✓	SPAN
✓	RSPAN
✓	SPAN over IP
<b>Analytics</b>	
✓	Unum Analytics UI
✓	Switch Level Information
<b>Management</b>	
✓	SNMP
✓	Logging (Syslog)
✓	AAA (TACACS+)
✓	NTP
✓	DNS
✓	ZTP
✓	Image Management
✓	SSH and Banner

Source: Tolly, September 2018 Table 2

information provided by the Fabric. This eliminates a lot of configuration complexity.

Distributed virtual routing and forwarding (VRF).

Pluribus demonstrated an impressive range of VXLAN capabilities. These included: 1) Auto tunnels, 2) VXLAN bridging, 3) VXLAN routing, 4) Anycast IPv4/v6, and 5)





# Test Results

This section summarizes tests not covered in the solutions analysis.

## Performance

Industry standard RFC 2544 performance tests were run on the D-Link DXS-5000 using 10GbE ports. All performance testing was performed using one port pair for logistical reasons.

The switch delivered 100%, wire-speed throughput at all frame sizes in the 10GbE test. The tests included the seven frames sizes specified by RFC 2544 as well as the largest jumbo frame of 9216-bytes. See Figure 1.

10GbE average latency ranged from 0.823µs to 8.21µs for frame sizes ranging from 64-bytes to 9216-bytes (jumbo frames). See Figure 2.

## LAN Topologies

(See Tables 1 and 2 for detailed lists of features and functions summarized in this section of the report.)

In addition to the aforementioned 10GbE performance testing Pluribus also demonstrated connectivity at 40GbE.

## Layer 2 Functionality

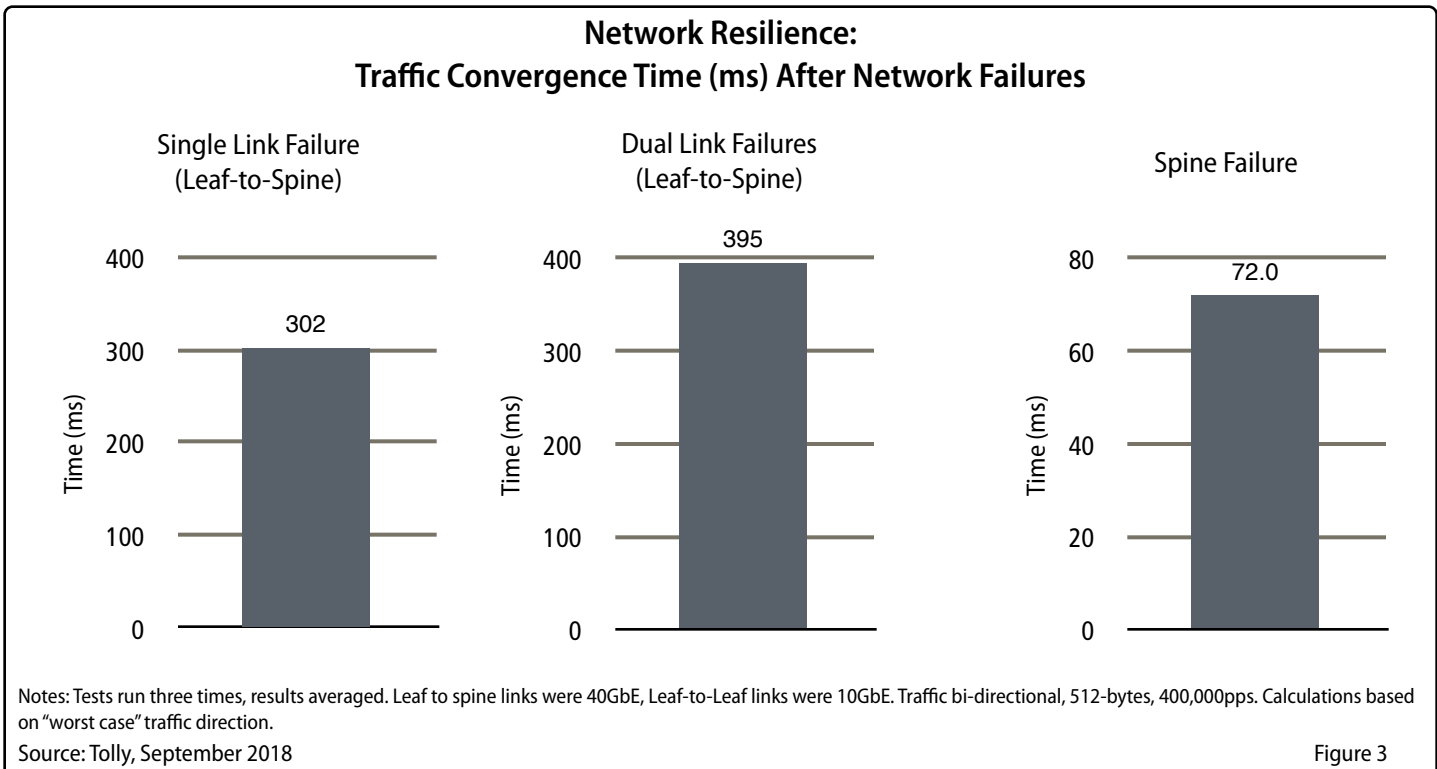
Testing verified a range of features including RSTP/MSTP , VLANs, Trunking/ Link Aggregation, EtherChannel, switch virtual interface (SVI), LLDP, and ARP and broadcast suppression. The Adaptive Cloud Fabric is used to suppress ARP and broadcast to the local nodes.

## Layer 3 Functionality

Extensive functionality at Layer 3 was verified. The Pluribus Networks Netvisor ONE OS provides dual-stack, IPV4 and IPV6 support across all areas tested. In addition to static routes, VRRP, and vRouters the test verified support for BGP and OSPF routers and neighbors.

## Quality-of-Service (QoS)

QoS assures that traffic is handled in accordance with priorities set by the network manager. The test verified important QoS features. In addition to core queueing capability, Engineers verified packet marking, Queuing using custom maps by modifying the COS to DSCP Mapping Table, and Minimum bandwidth guarantee.



## DHCP

DHCP functions are essential as they provide the service for dynamic address acquisition by IP endpoints. In addition to core DHCP functions, engineers verified DHCP snooping (both track and filter). Also verified was the important IPv6 router advertisement (RA) and RA guard security features. RA guard prevents unauthorized endpoints from advertising router services.

## Packet Mirroring

For a variety of reasons including troubleshooting, traffic and security analysis, it is important to be able to capture traffic traversing specific ports. Port mirroring is the feature that provides this capability and it is also referred to as switch port analyzer or SPAN.

Engineers verified both local and remote port mirroring as well as SPAN over IP.

## Analytics

As networks become more complex, analytics becomes more and more important. Tolly engineers verified extensive analytics functions provided by the Pluribus Networks Unum analytics environment.

Additionally, engineers verified availability of statistics and related information individually at the switch level.

## Management

Management encompasses many functions. Tolly verified the extensive suite of management features available in the Pluribus Networks environment. The device can be accessed and configured

using CLI, REST and the UNUM GUI. See Figure below.

Tolly engineers verified support for core SNMP, Syslog and authentication (AAA/TACACS+).

Zero-touch provisioning (ZTP) saves significant deployment time and Tolly verified support for this feature along with extensive image management support that eases upgrades and provides fallback in the event that a downgrade is desired. Software upgrades can be performed on one device at a time or all devices in the fabric using the Fabric Software Upgrade.

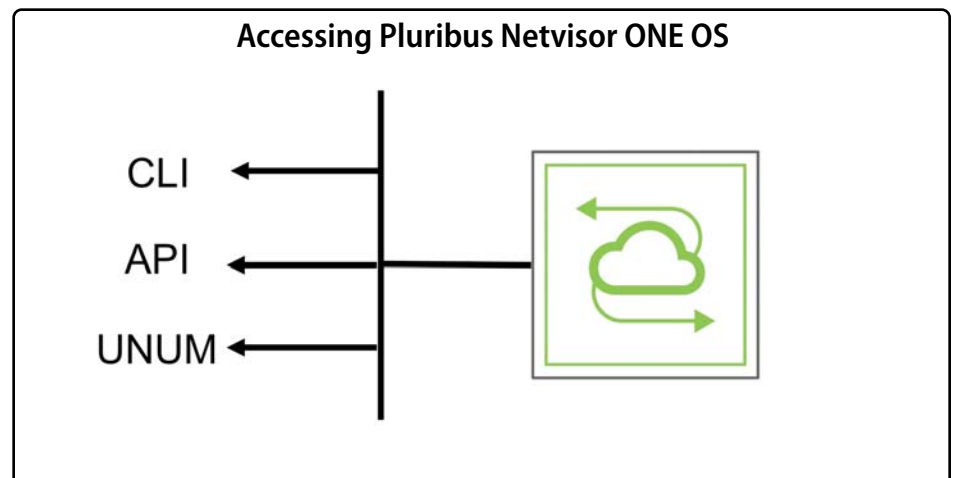
Furthermore, Tolly verified support for SSH capabilities.

## Troubleshooting

Adaptive Cloud Fabric CLI makes troubleshooting extremely simple by providing a single pane of glass of view through CLI, REST and UNUM. For example when using the CLI the user can troubleshoot a single switch, a group of switches or all switches in a fabric in a single SSH session.

## Network Resilience Testing

Tests of recovery times for single and dual link failures as well as spine switch failures confirmed that recovery from every scenario was sub-second. See Figure 3.





# Test Setup & Methodology

## Performance Testing

Throughput and latency tests were running using Xena Networks hardware and software. The benchmarking software was the Xena Valkyrie2544 test suite v2.53 running on a Microsoft Windows 10 Enterprise computer controlling the Xena Networks test chassis. The Valkyrie Layer 2-3 test chassis was outfitted with Xena Networks QSFP28 100G CR4 cards. These cards could be configured for 10GbE LANs.

Tests were run using eight different frame sizes: 64-, 128-, 256-, 512-, 1024-, 1280-1580-, 1518-, and 9216-bytes. The last frame size is a "jumbo" MTU that is larger than the Ethernet specification maximum.

All tests were run three times for 60 seconds each and results averaged. Throughput tests were zero-loss tests. Latency tests were configured in "Last-to-Last" mode.

## Network Resilience Testing

Networks are typically built with multiple paths (links) either directly or indirectly between switches to provide for secondary paths in the event that a primary link fails. The time required for traffic to start

transiting the new, backup path is referred to as convergence or reconvergence time.

This test was designed to illustrate the worst case (i.e. longest amount of time) required for traffic to flow across a new, backup path after the original, active path was deliberately terminated.

The test environment was in a "Leaf-Spine" configuration with multiple leaf switches and multiple spine switches. There were multiple paths between each leaf switch and multiple spine switches. Three failure scenarios were benchmarked: 1) Single-link failure (leaf to spine), 2) Dual-link failures (leaf to spine), and 3) Spine switch failure. Leaf-to-spine links were 40GbE. Leaf-to-Leaf links were 10GbE.

To measure convergence time, bi-directional test traffic of 512-byte frames at a rate of 400,000fps was sent between two ports. Traffic traversed the active path. Engineers introduced the failure and then measured how many frames were lost before traffic resumed using the newly converged path. The known packet rate was used to convert the lost packets into elapsed time. The tests were run three times and the results averaged.

## Feature Testing

Features were evaluated as appropriate for each feature/function. Tolly engineers used terminal commands and logs, GUI, and system analytics among other methods, to document each feature test as appropriate.

### Solutions Under Test

Open Network OS: Pluribus Networks Netvisor ONE OS	v3.1.0
10 GbE Switch Hardware	D-Link DXS-5000-54S Open Network Switch (48x10GbE + 6x40GbE)

Source: Tolly, September 2018

Table 3





### About Tolly


The Tolly Group companies have been delivering world-class IT services for nearly 30 years. Tolly is a leading global provider of third-party validation services for vendors of IT products, components and services.

You can reach the company by E-mail at [sales@tolly.com](mailto:sales@tolly.com), or by telephone at +1 561.391.5610.

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

### Test Equipment Summary

The Tolly Group gratefully acknowledges the providers of test equipment/software used in this project.

Vendor	Product	Web
<b>Xena Networks</b>	<b>Valkyrie Layer 2-3 Test Platform Software: Valkyrie2544 v2.53</b>	 <a href="http://www.xenanetworks.com">http://www.xenanetworks.com</a>

### 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. No part of any document may be reproduced, in whole or in part, without the specific written permission of Tolly. All trademarks used in the document are owned by their respective owners. You agree not to use any trademark in or as the whole or part of your own trademarks in connection with any activities, products or services which are not ours, or in a manner which may be confusing, misleading or deceptive or in a manner that disparages us or our information, projects or developments.