TEST REPORT

Tolly.

#214123

July 2014

Commissioned by
Ruijie Networks Co., Ltd.

Ruijie Networks Newton 18000 Series Switch

Performance & Capacity Evaluation, Power Consumption Measurement and Feature Validation

EXECUTIVE SUMMARY

The demands on cloud data centers and enterprise data centers seem ever increasing. Network architects require switch infrastructure that provides high throughput and capacity as well as high-availability and a range of sophisticated features.

Ruijie Networks Co., Ltd. commissioned Tolly to evaluate its Newton 18000 Series Switch in a single-switch deployment as well as in a stacked configuration of 2 and 4 switches.

Tolly engineers found that the Ruijie N18010 switch delivered line-rate, layer 2 throughput at all frame sizes across a configuration of 192 40GbE ports. See Figure 1.

Additional tests, described below, illustrated that the Ruijie N18000 switches provided high system table and user authentication capacity. Furthermore, tests illustrated a number of high-availability, virtualization and multi-switch stacking features of the Ruijie N18000 series.

THE BOTTOM LINE

The Ruijie Networks RG-N18000 switch:

1. Delivered 100% line-rate Layer 2 throughput with 192*40GbE ports in snake topology for 64- to 9216-byte frames (Aggregated throughput of 7.68Tbps)
2. Delivered latency as low as 0.532 µs
3. Provided MAC table capacity up to 512K, ARP table capacity up to 170K and 802.1x authentication concurrent users capacity up to 100K
4. Illustrated virtualization technology of Ruijie's VSU and VSD technology, with up to 4 RG-N18000 virtualized as a single unit and virtualizing a physical switch into 12 virtual switches
5. Supported the high availability features without frame loss including N+1 fabric module redundancy, power redundancy and in-service software upgrade
6. Supported OpenFlow 1.3 features with Ruijie SDN Controllers

Ruijie RG-N18010 Switch Layer 2 40GbE RFC2544 Throughput
192 40GbE ports in a Snake Configuration across 8 Line Cards
(as reported by Spirent TestCenter v4.33)

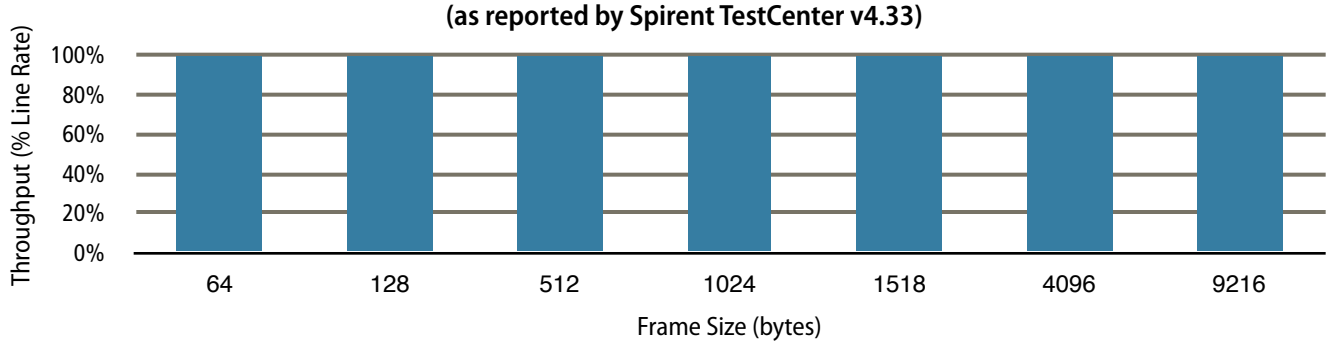

Note: Eight M18000-24QXS-DB line cards with four M18010-FE-DIII fabric modules were used on one Ruijie RG-N18010 switch. 100% line-rate Layer 2 throughput with 192 40GbE ports were verified for all tested frame sizes with zero frame loss. Aggregated throughput is 7.68Tbps. All traffic passed across line cards instead of passing to ports on the same line card.

Source: Tolly, May 2014

Figure 1

Test Results

Layer 2 40GbE Performance

Throughput & Latency

Tolly engineers tested throughput with a single RG-N18010 switch outfitted with 192 40GbE ports (8x24-port line cards). The traffic flow was in a snake configuration where the test flowed from the test traffic generator and then one port to the next, always crossing the backplane, through all 192 ports and then back to the receive port of the Spirent TestCenter.

Tests were run using 8 different frame sizes ranging from 64-bytes to jumbo frames at 4096 and 9216-bytes. In every test, the Ruijie RG-N18010 delivered the layer 2 traffic at 100% line rate with zero frame loss. See Figure 1.

Next, Tolly engineers benchmarked latency between two 40GbE ports on a single line card. While the default forwarding mode of the switch is store-and-forward, this test was run in both store-and-forward and cut-through mode to illustrate the effectiveness of this mode of operation.

In cut-through mode, the latency ranged from 0.532µs for 64-byte frames to 0.586µs for 9216-byte frames.

Store-and-forward latency was measured in LIFO mode which removes the time required to store a given frame and focuses on the processing time after the frame has been stored.

Here, too, the switch exhibited very low latency with the 0.650µs for 64-byte frames to 0.665µs for 9216-byte frames.

System Capacity

10 GbE Port Buffer

Port buffers provide a way to avoid frame loss in situations where there is momentary congestion (overcommitment) on a given output port.

Using a burst of 500,000 9216-byte frames, engineers created congestion on an outbound port. Tests showed that 36,293 frames were buffered which means 318.98MB buffer size or 268.16ms 9216-byte frame buffer for a 10GbE port on the M18000-40XS-CB line card.

Ruijie Networks Co., Ltd.

Newton 18000 Series Switch

Capacity, Performance, Power Consumption and Features

Tested May 2014

MAC, ARP, FIBv4, FIBv6 and IPv6 Neighbor Cache Table Table Capacity

Data center switches may carry traffic for tens of thousands of stations. Thus, it is important for the switch to be able track layer 2 and layer 3 information for large numbers of devices.

Tests showed that, depending upon the line card installed and the switch mode, the RG-N18000 switch tables could hold information for 512,000 layer 2 MAC addresses or the same number in the

Ruijie RG-N18010 Switch Layer 2 40GbE RFC2544 Average Latency
Two 40GbE Ports in Port-to-Port Traffic Configuration
(as reported by Spirent TestCenter v4.33)

Frame Size	64-byte	128-byte	512-byte	1518-byte	4096-byte	9216-byte
Cut-Through Latency (µs)	0.532	0.545	0.588	0.586	0.586	0.586
Store -and-Forward LIFO Latency (µs)	0.650	0.651	0.666	0.665	0.665	0.665

Note: Port 1 and Port 2 on one M18000-24QXS-DB line card was used for the test. For cut through mode, FIFO latency was captured in the Spirent RFC2544 latency test. For store and forward mode, LIFO latency was captured in the Spirent RFC2544 latency test suite. Thus, store-and-forward results do not include the time required to store the frame.

Source: Tolly, May 2014

Table 1

Ruijie RG-N18000 Series Switch Capacity
(as reported by Spirent TestCenter v4.33)

Port Buffer - 10GbE port[2]	MAC Table Capacity[3]	ARP Table Capacity[3]	FIB Table Capacity[4]	IPv6 Neighbor Cache Table Capacity[3]	802.1x Authentication Capacity[5]	802.1x Authentication Performance	DHCP Server Performance[5]
318.98MB, 268.16ms for 9216-byte frames	512,000	170,000	512,000 IPv4 or 512,000 IPv6	128,000	100,000 concurrent users	1,200 requests per second	2,700 users per second

Note: 1. Capacity varies for different line cards and switch mode. The capacities shown here may not be reached at the same time.
2. Buffer of one 10GbE port on the M18000-40XS-CB line card was evaluated.
3. One M18000-08XS-ED line card was tested.
4. One M18000-44SFP4XS-EF line card was tested.
5. One M18010-CMII control module was tested.

Source: Tolly, May 2014 Table 2

forwarding information bases (FIB) for IPv4 or IPv6 stations. The IPv6 neighbor cache held 128,000 entries and the address resolution protocol (ARP) table held 170,000 entries. See Table 2.

802.1x Authentication Services Capacity and Performance

802.1x is the commonly implemented, standards-based approach for switch port authentication. Given the large number of users potentially serviced by a data center switch, high-capacity support of 802.1x both in respect to concurrent users and concurrent requests is important.

Tests showed that, depending upon the control module installed, the RG-N18000 switch could support up to 100,000 concurrent users and 1,200 802.1x authentication requests per second. See Table 2.

DHCP Server Performance

As data center switches are often tasked with serving IP addresses to stations, it is important to have sufficient capability to handle users requests especially when there is a surge of requests as frequently is the case at the beginning of the work day.

Tests showed that the RG-N18000, depending upon the control module installed, could serve IP addresses to up to 2,700 users per second via its dynamic host configuration protocol (DHCP) IP address server. See Table 2.

Ruijie RG-N18000 Series Switch VSU Convergence Performance
(as reported by Spirent TestCenter v4.33)

	Uplink Failure	Main and Backup control modules switch over (Main control module failure)	Backup control module failure	Main switch failure	Backup switch failure
Average Convergence Time	35 ms	0	0	11 ms	11ms

Note: Convergence time here considered load balancing. In the uplink failure, main switch failure and backup switch failure tests, due to load balancing, half traffic streams were affected by the failure while half were not. The convergence time for the affected streams is reported here. See the test methodology section for detail.

Source: Tolly, May 2014 Table 3

Virtual Switch Functions

Virtual Switch Unit (VSU) - Resiliency

Multiple physical Ruijie RG-N18000 switches can be virtualized into a single logical switch. Tolly engineers verified that up to 4 switches could be combined in this manner to create a single virtual switch unit (VSU).

To demonstrate the resiliency features of a VSU, Tolly engineers created failures of key components and verified failover/rapid recovery on a 2-unit VSU as follows:

1) Uplink Failure: Single link of two-member LAG failed. Traffic on the failed link reconverged in an average of 35ms. There was no disruption of traffic on the other link of the LAG.

2) Main Control Module Failure: No traffic loss.

3) Backup Control Module Failure: No traffic loss.

4) Main/Backup Switch Failure: When either the main or backup switch was failed, traffic on the remaining was unaffected and the traffic traversing the failed switch reconverged in an average of 11ms.

Ruijie RG-N18000 Series Switch Features

Tolly Certified Features		
✔	4 Member Virtual Switch Unit (VSU) with RG-N18010	
✔	Virtual Switch Device (VSD) - virtualize one RG-N18010 into 12 virtual switch devices	
✔	OpenFlow 1.3 Main Features	
✔	Modular System	In-Service Software Upgrade (ISSU) for Control Modules
✔		Stateless Process Restart
✔		Hot Process Patching
✔	N+1 Fabric Module Redundancy	
✔	Hot Swappable Fabric Module	

Source: Tolly, May 2014 Table 4

Virtual Switch Device (VSD)

Tolly engineers illustrated that a single physical RG-N18000 switch could be configured into multiple virtual switches. In this test, engineers configured and verified 12 virtual switches on a single RG-N18010 physical switch.

Ruijie RG-N18010 Switch Power Consumption with 192 40GbE ports
(as reported by Fluke 317 Clamp Meter and Fluke 15B Digital Multimeter)

	Without Any Line Card	Load All Modules - 0% Traffic	Load All Modules - 30% Traffic	Load All Modules - 100% Traffic	Difference between 100% traffic and without any line card	Average Apparent Power per line card with 100% traffic	Average Apparent Power per 40GbE port with 100% traffic
Apparent Power	1,456 VA	2,573 VA	2,742 VA	3,108 VA	1,652 VA	206.5 VA	8.6 VA

Note: 1. Apparent power = Current * Voltage. Real power = Apparent power * power factor. Power factor <= 1. So the real power in Watts, which is what the utility company charges the customers, is less than the apparent power reported here.
2. Results in white cells were measured. Results in green cells were calculated from the measured results.
3. One Ruijie RG-N18010 switch was fully loaded with two M18010-CM control modules, four M18010-FE-D III fabric modules, eight power supplies, four fan modules and eight M18000-24QXS-DB line cards. Each line card has 24 40GbE ports.
4. iMIX 4-point traffic in Spirent TestCenter was used as the test traffic.
5. The per 40GbE port's power consumption does not count the base power consumption of the switch chassis (apparent power without any line card). See Test Methodology section for all calculations.

Source: Tolly, May 2014 Table 5

Management, Modularity & High-Availability

OpenFlow 1.3

Engineers verified select OpenFlow features on the RG-N18010 switch. The switch could be managed by one SDN controller cluster which includes two Ruijie IONP SDN controllers (one as active and one as standby) via the OpenFlow channel. Engineers verified that the OpenFlow channel was negotiated with OpenFlow protocol version 0x04, which is defined in OpenFlow 1.3 specification. The OpenFlow channel was composed of a main connection and an auxiliary connection.

Tolly engineers verified that the switch could match traffic with source/destination MAC address, source/destination IP address and DSCP value. Two actions were evaluated - forwarding to a specific port and modifying the DSCP value.

Modular System

The modular nature of the switch operating system provides for continued operation in a number of scenarios that would normally require a switch to reboot. Among the scenarios tested were:

1) In-Service Software Upgrade: The backup and then the main control modules were upgraded were both upgraded while the system was handling traffic and no frame loss was observed.

2) Stateless Process Restart: Engineers demonstrated that a halting the "bridge" process and restarting it did not disrupt existing traffic.

3) Hot Process Patching: Engineers verified that a hotfix patch could be applied to a process without disrupting the traffic controlled by that process.

High-Availability

The RG-N18000 series also provides high-availability with respect to fabric modules.

A switch containing 4 fabric modules, configured as 3+1, delivered line-rate, zero-

Two-Member Stack Convergence Performance Test Bed

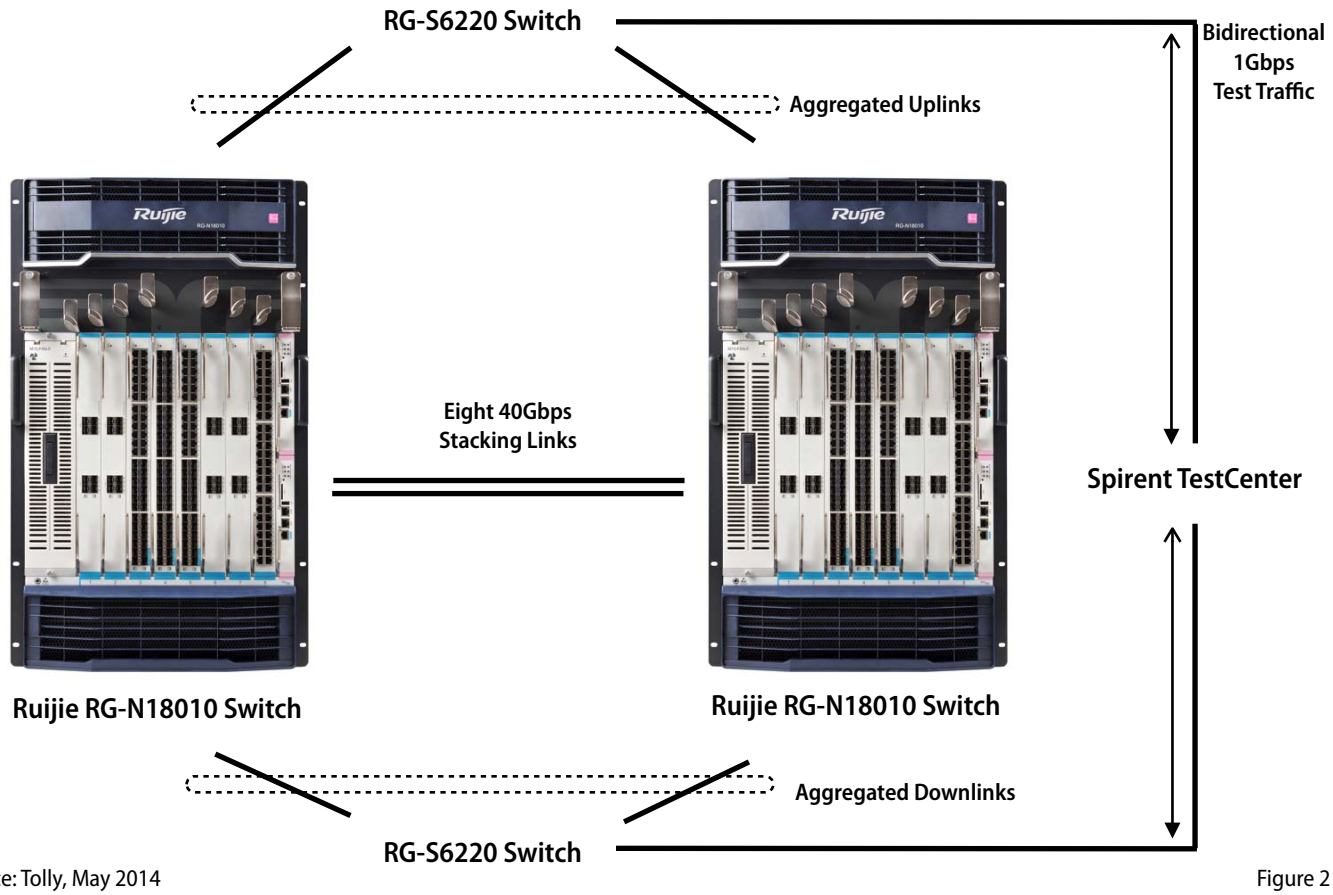

Source: Tolly, May 2014

Figure 2

loss traffic, at all frame sizes tested from 256-byte to 9216 byte frames, across all 192 40GbE ports delivering an aggregate of 7.68Tbps of traffic.

Engineers also verified that fabric modules were hot-swappable by under the maintenance mode and that no traffic was lost during the swapping process.

Power Consumption

Finally, Tolly engineers measure the power consumption of the switch under a variety of traffic conditions. This information helps calculate longer term operational costs of running the switch. Those results are documented in Table 5.

Test Methodology

Layer 2 40GbE Performance

Throughput

One RG-N18010 switch with 8 M18000-24QXS-DB line cards and 4 M18010-FE-DIII fabric modules was under test. 192 40GbE ports on the switch was tested with snake topology. All traffic went across line cards instead of passing to ports on the same line card. RFC2544 Throughput test in Spirent TestCenter was used to run the test. Three iterations were run for each frame size. Each iteration was with 1 minute.

Latency

Port1 and Port2 of the M18000-24QXS-DB line card were under test with 100% line-rate traffic. Cut through as well as store and forward modes were both tested using Spirent TestCenter RFC2544 latency test suite. Cut through latency used FIFO latency measurement. Store and forward latency used LIFO latency measurement. Three iterations were run for reach frame size. Each iteration was with 1 minute. A development command was used to change the latency mode to cut-through.

System Capacity

Port Buffer

Three 10GbE ports (designated A, B and C) on the M18000-40XS-CB line card were used during the test. Engineers used Spirent TestCenter to send line-rate traffic into port A and out of Port B. One burst with 500,000 9216-byte frames were sent into the switch from Port C and out of Port B. 463,707 frames were lost. Thus 36,293 frames were buffered which means a 318.98MB buffer size or a 268.16ms 9216-byte frame buffer for 10GbE port B.

MAC, ARP, FIBv4, FIBv6 and IPv6 Neighbor Cache Table Capacity

Each table's capacity was evaluated individually. Traffic with destination addresses matching the table under test was passed by the switch without loss. For the MAC table test, a listening port was used to make sure there was no broadcast.

Table capacity varies with different line cards or switch mode. For chassis based switches, the switch capacity is defined by the lowest capacity line card installed. Also, the switch mode may determine the capacity of different tables. See the notes for Table 2 for the specific line cards used in this test. Users should consult the vendor for the capacity of different switch configurations.

802.1x Authentication Capacity

Ten ports of the Spirent TestCenter were connected to the RG-N18010 switch under test. Each port sent 10,000 users to authenticate with the Ruijie RG-SAM Operator 3.82 Build 20130424 Radius

About Ruijie Networks

Ruijie Networks (stock code: SZ: 002396) is a leading network solution supplier of China. We focus on customer benefits and strive to improve the network application experience of our customers through continuous technological innovation. Ruijie Networks provides end-to-end network solutions for telecom carriers, financial services, government agencies, education and enterprises to create values for customer networks.

Ruijie Networks has over 38 branch offices with sales and service covering Asia, Europe, North America, and South America. Ruijie Networks is the only data communication company certified as an innovative enterprise in China. With continuous improvement on the innovative road of independent R&D, Ruijie Networks leads and promotes the development of cutting-edge network technologies worldwide.

For more information on Ruijie Networks enterprise switches solution, please visit our company website, or scan the QR code below:

Learn more about Ruijie Switches

Source: Ruijie Networks, July 2014

server though the switch. User authentication was sent out port by port. All 100,000 concurrent users were authenticated to access the network.

802.1x Authentication Performance

Ten ports of the Spirent TestCenter were connected to the RG-N18010 switch under test. Each port sent users to authenticate with the Ruijie RG-SAM Operator 3.82 Build 20130424 Radius server though the switch. Ten ports sent out user authentication together. After 1 minute, 74,172 users were authenticated.

DHCP Server Performance

The M18010-CMII control module was used in the test. One PC with Cyclone 3.8.0 Beta was used to simulate users to request DHCP addresses from the RG-N18010 switch under test. 10,000 users got their IPv4 DHCP addresses in 3.7 seconds. Thus 2,700 DHCP addresses was assigned by the switch every second.

Virtual Switch Functions

Virtual Switch Unit (VSU) - Two Members

Two RG-N18010 switches virtualized into 1 logical switch using eight 40Gbps links to stack. Two Ruijie RG-S6220 switches were used (one as uplink and one as downlink switch) in the test. There was one GbE link between each RG-S6220 switch to each RG-N18010 switch. RG-N18010 was in default local forwarding mode. 1Gbps bidirectional traffic with 1,000 source IP addresses and 1,000 destination IP addresses was used as the test traffic. So there was 1,000,000 source/destination combinations. See Figure 2 for the test bed.

Uplink Failure

To test the uplink failure, one uplink (tried either one) was unplugged. As there was load balancing in the Link Aggregation Group (LAG), each uplink only has 500Mbps traffic. The traffic was bidirectional. Worst-case convergence time was reported.

According to the number of lost frames, half test streams (500,000 source/destination combinations) were not affected. The other half of test streams required an average 35ms to converge.

Main Control Module Failure

When the main control module fails, the backup control module on the other RG-N18010 became the main control module. Engineers evaluated the failover by pulling the main control module out of the switch chassis. There was no frame loss for existing traffic.

Backup Control Module Failure

When the backup control module failed, there was no frame loss for the existing traffic. Engineers tested this by pulling the backup control module out of the switch chassis.

Main/Backup Switch Failure

When the main switch (the RG-N18010 with the main control module) failed, half of the test streams (500,000 source/destination combinations) were not affected while the other half of the test streams took an average of 11ms to converge.

The same results were captured for the backup switch failure.

Engineers created the switch failure by unplugging all power cables of the switch.

Virtual Switch Unit (VSU) - Four Members

Four RG-N18010 switches interconnected via 10GbE stacking link equipped with the M18010-CM control modules were virtualized into one virtual switch.

Virtual Switch Device (VSD)

One RG-N18010 switch was virtualized into 12 virtual switch devices. Engineers accessed each virtual switch via Telnet and

Systems Under Test

Model	Version
RG-N18010 Switch	RGOS11.0(1)B2
M18010-CM Control Module	
M18010-CMII Control Module	
M18010-FE-DIII Fabric Module	
M18000-24QXS-DB Line Card	
M18000-40XS-CB Line Card	
M18000-08XS-ED Line Card	
M18000-44SFP4XS-EF Line Card	

Source: Tolly, May 2014 Table 6

verified the broadcast domain was isolated in each virtual switch.

Management, Modularity & High-Availability

OpenFlow 1.3

Engineers verified select OpenFlow features on the RG-N18010 switch with two Ruijie IONP SDN controllers (one as active and one as standby) via the OpenFlow channel. See the Test Results - OpenFlow 1.3 section for detail.

Modular System

The operating system RGOS11.0(1)B2 on the RG-N18000 switch is a modular system. The fault of one process could be isolated from other processes.

In-Service Software Upgrade (ISSU)

When one RG-N18000 switch is equipped with two control modules, the switch supports ISSU to update the firmware of the control modules. The backup control module was first upgraded and became the master control module. Then the original master control module (now backup control module) was upgraded. There was no frame loss during the update for existing traffic.

Stateless Process Restart

Engineers killed the "bridge" process, the existing traffic could be forwarded without frame loss before and after the process restarted.

Hot Process Patching

One hotfix patch to patch one process was applied to the switch. The existing traffic could be forwarded without frame loss.

High Availability

N+1 Fabric Module Redundancy

The RG-N18010 switch supports 4 fabric modules. With 3 fabric modules, RG-N18010 switch could still support 100% line-rate across board traffic with 192 * 40GbE ports in snake topology for 256-, 512-, 1024-, 1518-, 4096-, and 9216-byte frames with zero frame loss. Across backplane traffic aggregated throughput: 7.68Tbps. Three iterations were run. Each iteration ran for 1 minute.

Hot Swappable Fabric Module

Two fabric modules were used under test with 10Gbps across board traffic. Tolly engineers verified that both fabric modules were forwarding the traffic. Tolly placed the component into maintenance mode and performed the swap without frame loss.

Power Consumption

One Fluke 317 Clamp Meter was used to measure the current and one Fluke 15B Digital Multimeter was used to measure the voltage.

The RG-N18010 switch was fully loaded to measure the power consumption. See the notes of Table 5 for details.

The fully loaded switch has 8 line cards. Each line card has 24 40GbE ports. The calculations used in Table 5 are as follows : (Note: some results are rounded.)

1,652 = 3,108 - 1,456;

206.5 = 1,652 / 8;

8.6 = 206.5 / 24 = (3,108 - 1,456) / (8 * 24)

Ruijie Networks RG-N18000 Overview

Information provided by Ruijie Networks, not necessarily tested by Tolly

RG-N18000 Series-Basic Information

Developed by Ruijie based on cloud network architecture, the RG-N18000 series core switches provide high forwarding performance, high-density 10GE/100GE ports, and FCoE/IP integrated networking to meet application requirements of integrated data center networks and campus networks.

	Campus Network	Data Center
N18000 Series Ultra-Simplified Solution	• L2/L3, QinQ, IPv6 • MPLS, Multicast • Tunnel, ACL • QoS, BFD • IPFIX, sFlow • Hot Backup, ISSU • Unified authentication • Large table size • Unified management mode (One Switch One Network) • SDN and more…	• FCoE • DCB • TRILL • Security policy migration • VSU (4 to 1) • VSD (1 to 12) • High density/High bandwidth • Low latency • Ultra-large buffering • L2-GRE and more…
Platform Features	**Hardware**: CLOS architecture, 1.28Tbps per slot (unidirectional) **Software**: Multiprocessing and modularized OS, ISSU and process backup	

Source: Ruijie Networks, July 2014

About Tolly

The Tolly Group companies have been delivering world-class IT services for more than 25 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, 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
Spirent	**TestCenter v4.33**	http://www.spirent.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. 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.

214123 iv3 yx-wt-2014-08-18-VerK