TEST REPORT

Tolly.

#216157

November 2016

Commissioned by
Lightfleet Corporation

Lightfleet Data Distribution System (DDS)

Competitive Multicast Performance Evaluation

Lightfleet DDS vs. Arista 7150S-52 10GbE Switch with Solarflare SFN7122F 10GbE Adapters

EXECUTIVE SUMMARY

Lightfleet Corporation commissioned Tolly to benchmark the multicast end-to-end average latency, skew across receiving stations, and jitter for a given station for its Lightfleet Data Distribution System (DDS). This was compared to the performance of a traditional switch as represented by the Arista Networks 7150 in conjunction with Solarflare 7122 network interface cards (NIC) and related software for the end station. Traditional LAN switch architectures add control plane and network protocol stack overhead that negatively impact performance and consistency. The Lightfleet switch-free solution delivered significantly better multicast performance with near zero skew and jitter compared to the Arista/ Solarflare solution.

Applications requiring near instantaneous packet transfer with fairness across all receivers and effectively no jitter could benefit from the Lightfleet end-to-end solution.

THE BOTTOM LINE

Lightfleet's Data Distribution System delivers:

1. End-to-end average multicast latency that is significantly lower for all packet sizes tested
2. Near zero multicast latency skew at all receiving stations
3. Superior and consistent multicast performance with near zero jitter at all receiving stations

Multicast End-to-End Average Latency Comparison
Lightfleet DDS vs Arista 7150 10GbE Switch & Solarflare SFN7122F 10GbE Adapters
(Lower is better)

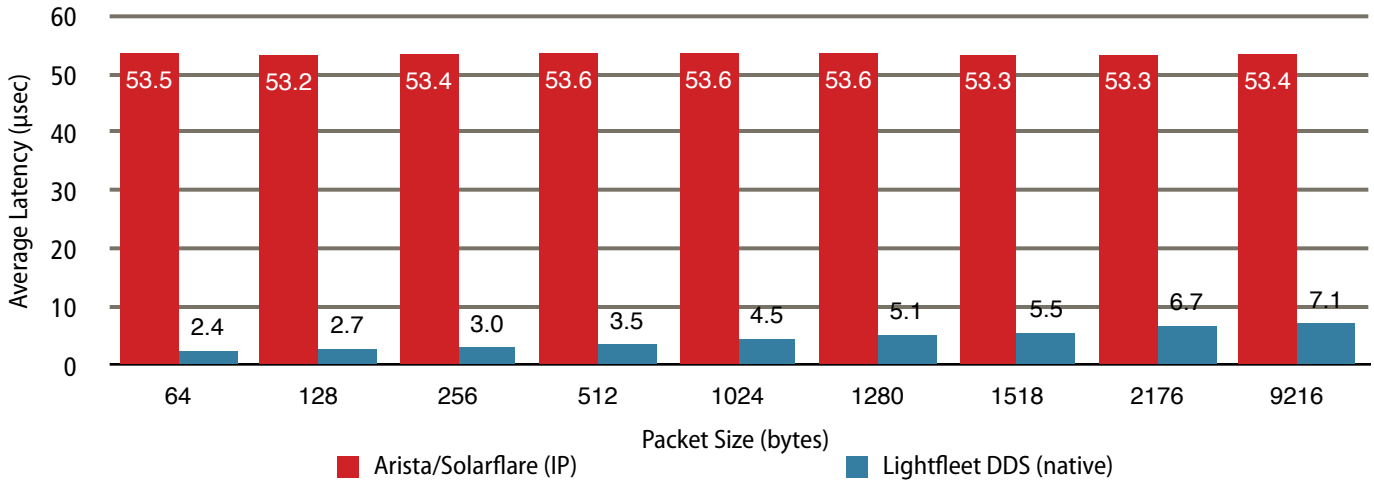

Notes: See Test Methodology section for detailed notes on test.

Source: Tolly, October 2016

Figure 1

Multicast Test Results

Lower latency (delay) is always better. Period. Additionally, packet arrival times should be consistent across stations as well as for a given station.

Average Latency

Tests evaluated a range of packet sizes from 64-bytes up through two "jumbo" sizes as large as 9216-bytes. The average, end-to-end latency for Arista/Solarflare was from 7.5X (9216-byte) to 22X (64-byte) times longer than Lightfleet. Where Arista/Solarflare required 53 µsec across the range of packet sizes, Lightfleet's times ranged from 2.4 to 7.1 µsecs. See Figure 1.

Latency Skew

Reliable, consistent delivery times are critical for certain applications. Latency skew looks at comparing the average results at each receiving station. Such an analysis is best depicted visually. A flat, horizontal, Lightfleet result, illustrates complete consistency across all receivers. A rising line, the Arista/Solarflare result, illustrates that the higher-numbered receiving stations experienced longer delivery times (i.e. higher latency), than the lower-numbered stations. This would indicate that where Lightfleet transmits all data simultaneously to receivers, Arista/Solarflare appear to transmit sequentially. See Figures 2 and 3. For individual data points see Table 1.

Multicast Latency Skew (22 Receivers)
Lightfleet DDS vs Arista 7150 10GbE Switch & Solarflare SFN7122F 10GbE Adapters
(Flatter is Better)

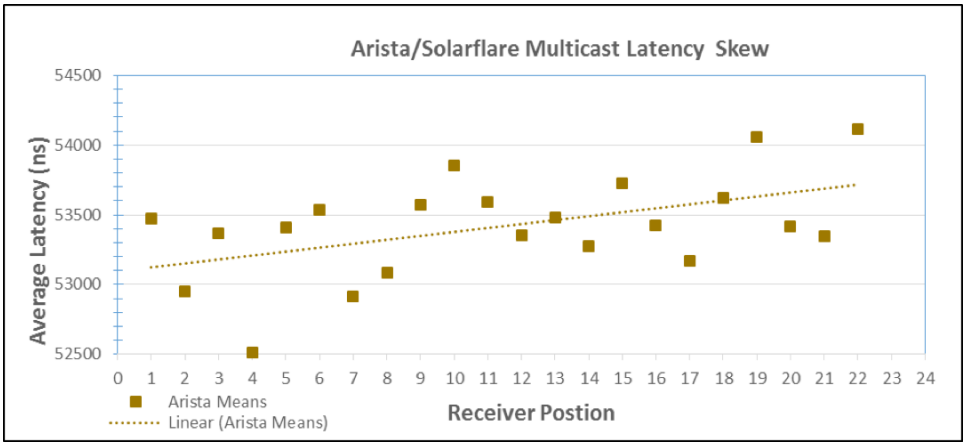

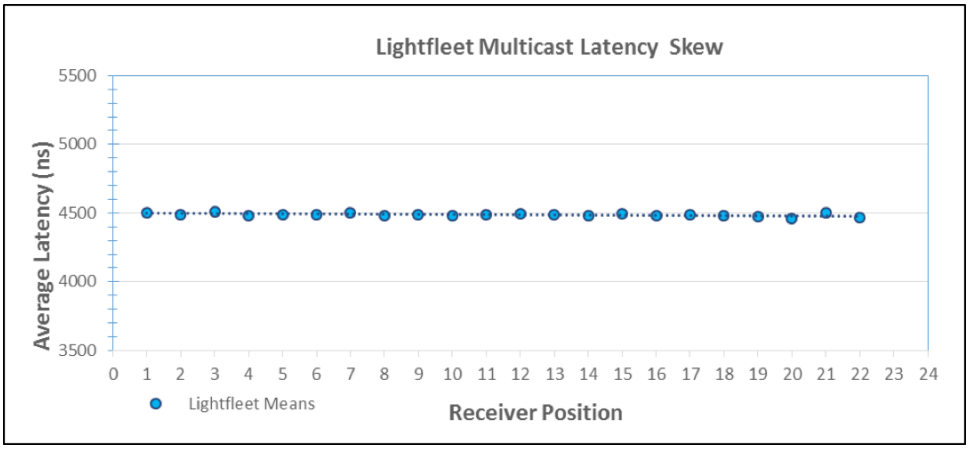

Notes: See Test Methodology section for detailed notes on test. Individual receiver results can be found in Table 1.

Source: Tolly, October 2016

Figures 2 and 3

Latency Jitter

Lastly, test results were analyzed to determine the consistency of delivery times at each receiving station. Jitter is the variation (standard deviation) in those delivery times.

Here, too, the data analysis showed virtually non-existent delivery differences for Lightfleet with the maximum standard deviation of a miniscule 15.7 billionths of a second (nanoseconds).

By contrast, the standard deviation for the Arista/Solarflare results ranged from 479.4 to 676.6 ns.

The Lightfleet jitter improvements over Arista/Solarflare was always at least 97% better. See Figure 4.

Thus, in every analysis: overall average latency, skew across receiving stations and jitter for a given station, the Lightfleet solution demonstrated dramatically better and consistent multicast performance than the Arista/Solarflare solution.

Applications requiring near instantaneous transfer with fairness across all receivers and effectively zero jitter could benefit from the end-to-end Lightfleet solution.

Lightfleet Corp.

Lightfleet Data Distribution System (DDS)

Multicast Performance Evaluation

Tolly. Certified

Tested October 2016

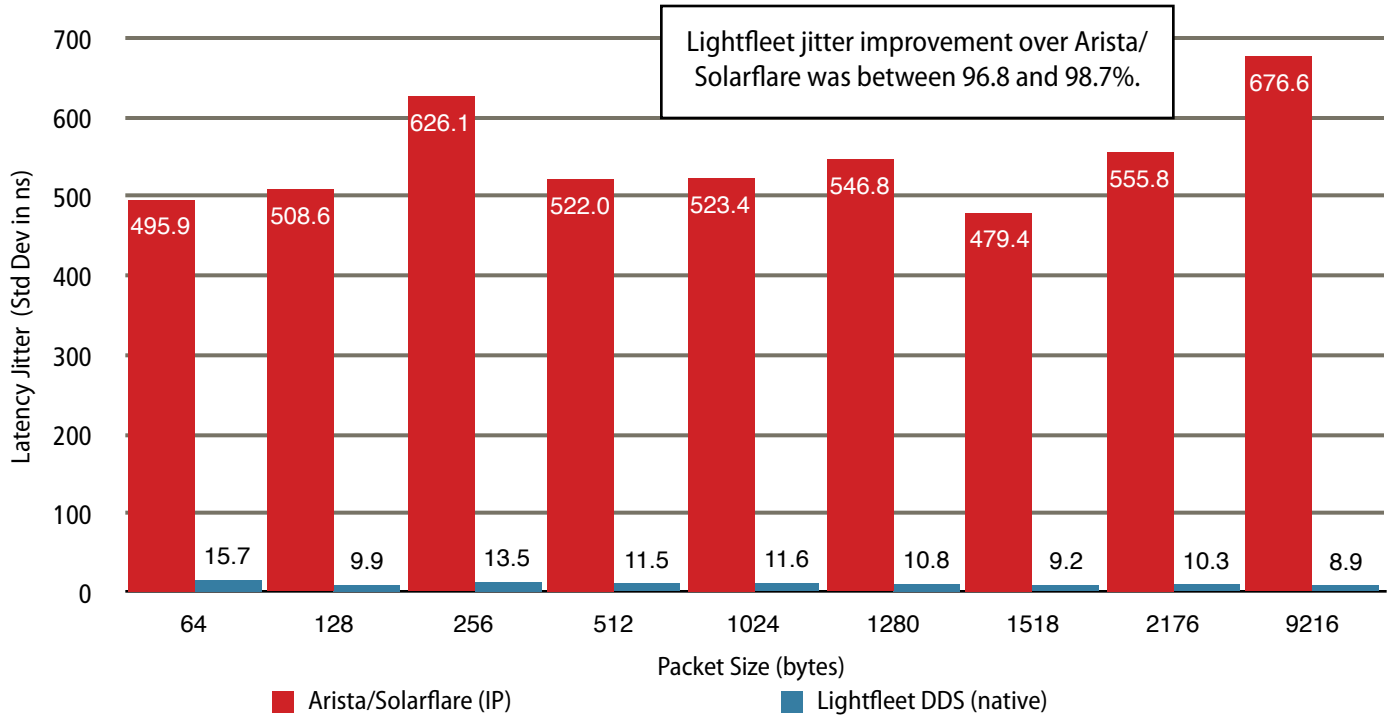

Notes: See Test Methodology section for detailed notes on test.

Source: Tolly, October 2016

Figure 4

Multicast Latency Average (Mean) Results By Receiver

Lightfleet DDS vs Arista 7150 10GbE Switch & Solarflare SFN7122F 10GbE Adapters

(Lower is better)

Average Latency (ns)	Receiver Position										
	1	2	3	4	5	6	7	8	9	10	11
Arista/ Solarflare	53472	52949	53369	52512	53407	53538	52914	53081	53574	53853	53593
Lightfleet	4499	4487	4508	4483	4488	4490	4499	4484	4488	4482	4492

Average Latency (ns)	Receiver Position										
	12	13	14	15	16	17	18	19	20	21	22
Arista/ Solarflare	53354	53481	53276	53730	53425	53171	53617	54056	53415	53348	54111
Lightfleet	4495	4488	4484	4493	4481	4491	4481	4477	4461	4502	4469

Note: This table contains individual data points plotted in Figure 2. The average (mean) is across all packet sizes tested. For Arista/Solarflare, individual packet size results ranged from ~51,000 to 55,000ns. For Lightfleet, individual packet size results ranged from ~2,400 to 7,100ns.

Source: Tolly, October 2016

Table 1

Test Stack Configurations

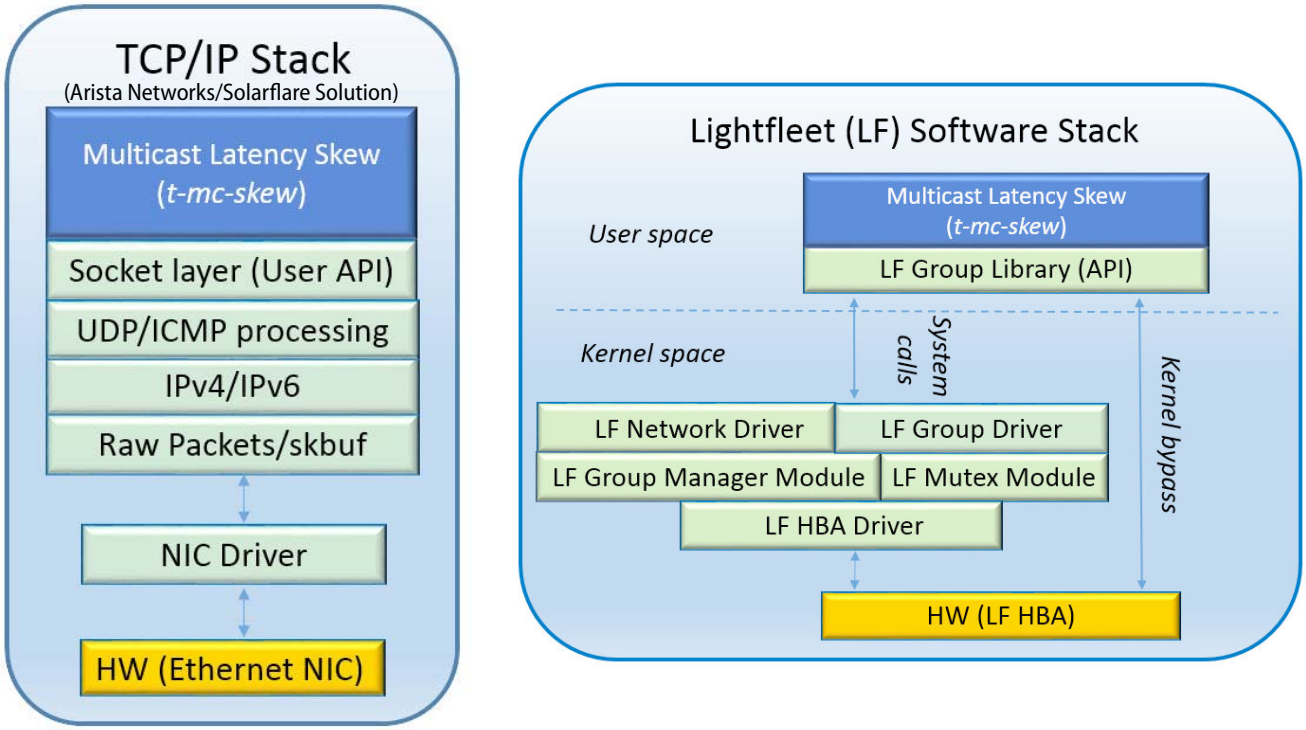

Source: Lightfleet, October 2016

Figure 5

Test Setup & Methodology

Test Environment

Testing was "end-to-end" and thus benchmarked not only the switch/data fabric but also the end-station network adapter and drivers.

All end stations used in the testing were Dell PowerEdge R220 Rack Servers outfitted with a single network adapter. The end-stations ran CentOS/Red Hat Enterprise Linux Version 7.2.

Timing Source: PTP

Because of the importance of exact time measurements in this test, a precision time protocol (PTP) server was used. The grandmaster clock was a Trimble, Inc. GM100. This clock received its time via GPS signal. The GM100 clock was connected to the end-stations by a dedicated switch and the clocks on the end stations were synced to within 50ns of each other. See Figure 6.

Systems Under Test

The Arista Networks 10GbE switch was paired with 10GbE network adapters from Solarflare. Solarflare was chosen for the network adapter as the company is recognized as producing high-performance, low-latency adapters that are in broad use throughout the industry.

Lightfleet's Data Distribution System (DDS) was tested. This consisted of several components, including the company's Data Distribution Module (DDM) and multiple Host Bus Adapters (HBAs), mounted in computer servers and connected to the DDM by means of optical fibers.

Lightfleet implements its switch-replacement data fabric in its Data Distribution Module - referred to as the Lightfleet DDM.

The Lightfleet HBA provides the network adapter function for the end station along with various standards-based and

Test Environment - Precision Time Protocol (PTP)

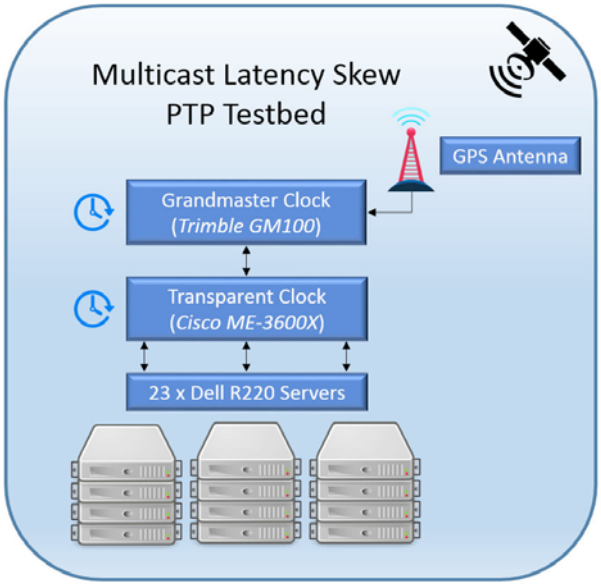

Source: Tolly, October 2016 Figure 6

Solutions Under Test

	Switch/Data Fabric	Networking Adapter
Arista Networks	Arista DCS-7150S-52-CL-F (Hardware: 2.01, Software: 4.17.0F) Tested using 24 10GbE ports (SFP-10G-SR)	Solarflare SFN7122F 10GbE Adapter (Drivers: 4.8.2.1004)
Lightfleet	Lightfleet Data Distribution Module (Firmware v2016/8/12)	Lightfleet HBA (Drivers lf_hbadrv.ko 0.0.1-test)

Source: Tolly, October 2016 Table 2

Lightfleet protocol stack modules. See Table 2 for details and version levels for the systems under test.

The Arista switch was outfitted with 24 10GbE ports (out of a possible 52) using SFP+ interfaces. The switch had two power supplies. The switch was configured to run cut-through switching to reduce packet latency.

The Lightfleet DDM was outfitted with 24 ports and two power supplies.

Test Methodology

Multicast Latency

Tests focused on measuring the multicast end-to-end latency of the systems under test. A custom program, t-mc-skew, was used on both systems to generate multicast traffic from one sender to 22 receivers.

The sending program inserted a timestamp in the outgoing packet immediately before transmission. Each receiving station timestamped the arrival of the packet and recorded the value. The delta from send to receive time was used as the result. Packets were sent at a low rate of approximately four packets-per-second as throughput was not a focus of the test.

For Arista, tests were run using standard TCP drivers in the end stations. For Lightfleet, tests were run using Lightfleet's application programming interface (API) driver in the end stations. For details of the network stacks tested, please see Figure 5.

Test results detailed latency times (i.e. packet arrival times) across individual receiving stations across a range of packet sizes. Those results were analyzed to illustrate not only raw latency but the variation (skew) in arrival times at different receivers and the variation in arrival time (jitter) of packets at the same receiver.

Understanding Lightfleet's Switch-Free Approach

All existing network architectures are optimized for point-to-point transfers and managed by a network control hierarchy. Lightfleet took a radically different approach. The company conceived and developed a network fabric where data transfers are self-directed, requiring no separate control plane to intervene in setting up transfer paths. The company's innovative Data Distribution System (DDS) eliminates the switching overhead caused by control planes that extends transfer latencies in large fabrics. Multicast transfers are inherent to the DDS architecture, producing no packet skew, with unicast transfers being handled as single-channel multicast transfers. Hence, multicast transfers are as efficient as unicast ones.

Most HPC and Big Data system managers are unaware that it is possible to have truly skew-free multicast transfers. The network switches on the market today generally implement multicast operations as a sequence of unicast transfers, causing an unavoidable skew between times when packets are received by client systems. With the default, skew-free multicast capability of the Lightfleet DDS, a packet-sending application can now be assured that its information is received simultaneously by a multitude of receiving applications. This capability may serve as an enabling technology for new applications.

Source: Lightfleet

Lightfleet Definitions

End-to-end Multicast Latency: is the time it takes for a multicast packet to travel across the network from user application to user application

Multicast Latency Skew: is the variation in arrival times of a packet at multiple stations when a single multicast packet is sent

Jitter: is the variation in arrival times for a packet to move from point A to point B when repeated packets of the same size are transmitted to the same receiver

Source: Lightfleet

The following summarizes key attributes of the test and the analysis:

1) Packet forwarding for the Arista 7150S-52 switch was run in cut-through mode
2) The same latency test application (t-mc-skew) was used for both Arista and Lightfleet testing
3) Multicast packets were sent from 1 producer to 22 receivers for all packet sizes
4) All latency measurements are end-to-end (application-to-application) averages
5) Latency of optical transceivers and fiber cables are included in latency results
6) All receivers have the same cable length to the DUT
7) Latency results are based on the average of three, 260 second test runs
8) A total of 594,000 latency measurements were made (22 receivers x 9 packet sizes x 1000 iterations x 3 test runs) for each DUT
9) The latency skew trend lines (Figures 2 & 3) are the linear best fit of all latencies measured across the 22 receiving stations for all packets sizes
10) Slope is the slope of a least-squares fit to the means
11) Skew is the latency skew across the set of 22 receiving stations
12) Jitter is the latency skew across packets arriving at a single receiving station

Relative Performance Calculation

To calculate how much better one solution is than another, the formula used is 1- (T1/T2) where T1 is the better result and T2 is the slower (worse) result. This is multiplied by 100 to give the percentage benefit. For "X" (times) better calculations, the better result is simply divided by the worse result.

Lightfleet's Data Distribution Module (DDM)

Importance of End-to-End Performance

For many demanding applications, such as HPC and Big Data, it is not just the data fabric performance characteristics that matter but also the performance "up" through the network adapter and even the software stack to that point where the network meets the application.

For this reason, these tests measure not only the data fabric device (e.g. switch), but also the network adapter and the communications protocol stack. Results represented what was measured by the benchmarking applications on top of the relevant protocol stack. Lightfleet provides a data fabric that is a purpose-built alternative to traditional Ethernet-based LAN switches. End-stations use a Lightfleet host bus adapter (HBA) in place of an Ethernet adapter. Applications can be run either over a standard TCP sockets interface or using the high-performance Lightfleet API.

Source: Tolly

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

Interaction with Competitors

In accordance with Tolly's Fair Testing Charter, Tolly personnel invited representatives from Arista Networks to participate in the testing. Arista Networks did not respond to this invitation.

For more information on the Tolly Fair Testing Charter, visit:

http://www.tolly.com/FTC.aspx

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.

216157 In-1 wt-2016-11-08-VerE