Atrium Field Trial

Phase 1 - Test Report

NOV 2016

This test report is an evaluation of Open Flow based BGP peering router (Project Atrium) for a field trial requirement.

Atrium is an Open Source SDN project from the Open Networking Foundation. This indepdendent test report from Criterion Network Labs includes assessment of open source routers, controllers and vendor switches participating in the field trial.

CNLabs is a vendor neutral independent network test laboratory for network operators and product vendors in the SDN/NFV domain. The lab also offers interoperability and conformance test programs from ONF and the IPv6 Forums.

The Atrium BGP peering router is a use case requirement from SIFY, an ICT solution provider and an SDN/NFV Alliance member of Criterion Network Labs.

SIFY plans to underake field trial tests for Open Flow based BGP peering routers based on Atrium distribution in Q1 2017. The field trial requirements of the customer are being validated at CNLabs. The solution will additionally undergo formal acceptance testing at SIFY before field trial deployment.

Write to cnlabs@criterionnetworklabs.com for additional questions and clarifications related to the test report.

INTRODUCTION

Project Atrium is an Open Source SDN project from the Open Networking Foundation. The goal of Atrium project is to provide integrated SDN solutions for network operator use-cases while achieving interoperability across multiple switch products. Atrium distributions include concepts like Flow Objectives that allow for interoperability across different implementations of Open Flow switch pipelines.

Atrium distributions are open source, integration tested and prequalified for near production quality. Since 2015, it has been demonstrated at various forums including the SDN world congress and more recently at Open SDN India 2016.

Atrium 16/A is the latest distribution of Atrium for service provider solutions. Along with the routing application, the distribution includes support of flow objectives interoperability module for ONOS and Open Daylight controllers.

The SDN BGP router implementation uses Quagga for control plane learning, a BGP speaker application on the SDN controller and program Open Flow switches using flow objectives.

Additional details on the Atrium project is available at https://github.com/onfsdn/atrium-docs/wiki.

SCOPE AND OBJECTIVE

The objective of the test effort is to ensure that the participating vendors and open source components meet all the core requirements identified for the field trail. The requirements are validated in a phased manner prior to acceptance testing at the customer site.

Gaps identified during each stage of testing are shared with the participants and the Atrium project team atrium_eng@groups.opensourcesdn.org. All open issues must be resolved by field trial participants to be included in the subsequent phase of testing. SIFY, the field trial customer might add additional requirements to be met based on the results of each phase of testing,

Phase 1:

- Validate core functional requirements
- Baseline tests and characterization
- SDN control plane learning/programming

Phase 2:

- IPv6 and MPLS support
- Data plane characterization
- Additional control plane requirements
- System level tests

Phase 3:

- Full test report with member switches

- Phase 3 test report will be disclosed publicly only if the vendors provides explicit consent to share the results

PARTICIPANTS

SIFY is one of the largest integrated ICT Solutions and Services companies in India, offering end-to-end solutions.

The comprehensive range of products from SIFY are delivered over a common telecom data network infrastructure and reach more than 1300 cities and towns in India. SIFY's telecom network connects 38 Data Centers across India including SIFY's 6 Tier III Data Centers across the cities of Chennai, Mumbai, Delhi and Bengaluru.

As part of their next generation network transformation initiatives, SIFY has been evaluating virtual provider edge gateways based on Open SDN solutions. The field trial of the Open Flow based Internet Gateway router is the first step of a phased field trial at SIFY towards achieving this end goal.

The current list of participants in the field trail include switches from Extreme Networks, Noviflow and SDN controllers from ONOS, Open Daylight, Open MUL.

IXIA and Spirent provide traffic generation and network emulation capabilities for the test environment.

DEPLOYMENT OVERVIEW

The test topology is a simple representation of SIFY's provider network. For the first phase of the field trial, SIFY is evaluating the implementation of the Open Flow based Internet Gateway router at their Tier-1 POP. The Atrium internet gateway provides internet connectivity to Enterprise customers on the Internet and VPN Provider Edge routers.

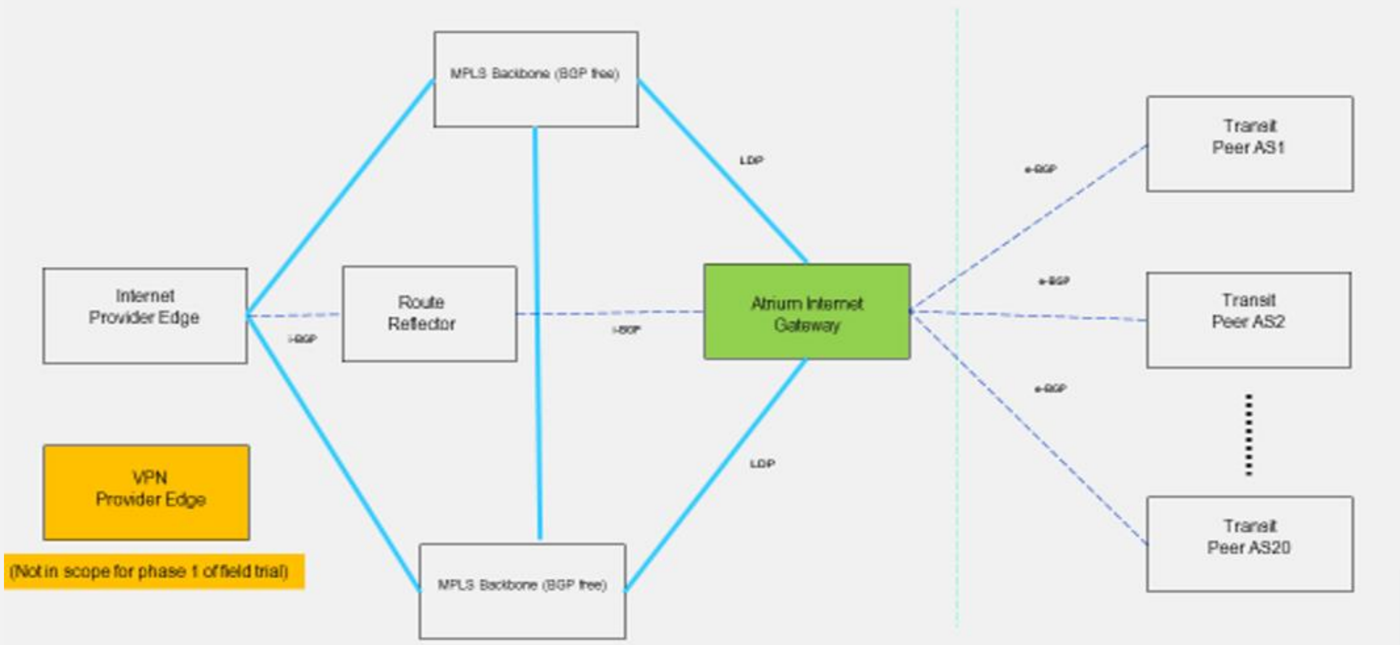

- Each Tier-1 POP has 1 Internet Gateway Router and 1 Route Reflector to provide internet connectivity for several internet provider edge routers connected to the MPLS backbone
- Each Internet Gateway Router is required to peer with
    - Up-to 20 transit peers using eBGP
    - 1 route reflector using iBGP
    - 2 MPLS neighbors using LDP
- Full internet feed and default route (IPv4 + IPv6) is received from different transit partners. Internet routes are reflected to all Internet edge routers through Route Reflectors
- Route reflectors handle control plane traffic only. Best path routing decision is done by Route reflectors using traditional BGP algorithms
- The Internet Gateway and Provider Edge devices are required to handle full RIB and FIB

TEST CONFIGURATION

Phase 1 test efforts is focused on functional testing and baseline characterization of the control plane. Vendor switches are not included in this phase of testing and OVS (Version 2.3.1) is used instead as the Open Flow Switch. Cisco ASR 9000s are used as the provider edge and route reflectors and additionally to emulate backbone network.

The control plane applications (Quagga + SDN controller) are hosted on a general-purpose rack server. All servers in the test-bed have identical hardware configuration and use Intel® Xeon® CPU E5620, 2.4 GHz quad core processors with dual CPUs. The next phase of test efforts will use actual server hardware configuration recommended by the customer.

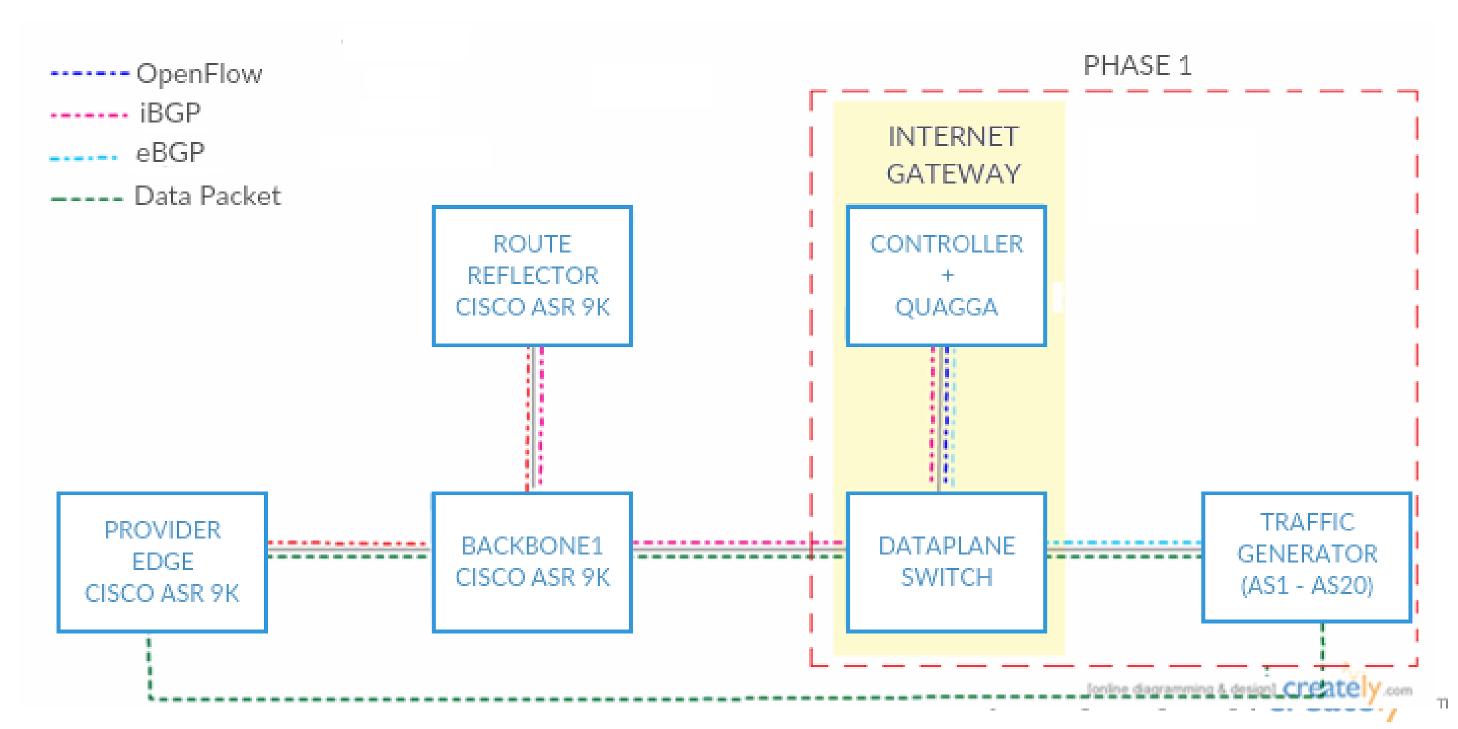

To ensure identical test conditions, all controller distributions are hosted on identically configured virtual machines. VMWare (ESXi 5.5.0) with a base Linux distribution of Ubuntu 14.04 is used in the test configuration. Quagga 0.99.24.1 and 0.99.22.1 with SDN controllers ONOS (1.5.0) and OpenMul (5.10.20) respectively, provide the control plane function. IXIA and Spirent emulate route updates from the internet peers.

TEST SUMMARY AND ANALYSIS

Summarized below are details of tests, results and analysis for three focus areas

- Verify if SDN controllers and routing applications meet core functional requirements; identify gaps in implementation
- Baseline characterization of route learning & flow programming times with Quagga for various traffic profiles
- Compare control plane learning and programming times using SDN controller for key traffic profiles

Given that the focus is on validating control plane functionality, all tests are conducted using Open Virtual Switch 2.3.1 as the Open Flow forwarding plane on identified hardware configurations.

Core requirements

In this test section, we validate if Quagga router and all evaluated SDN controllers meet core functional and scale requirements for the field trial and identify gaps in meeting control plane requirements.

SDN Controller Name	ONOS	OpenMul	ODL
SDN Controller Version	**1.5.0**	**5.10.20**	**—**
Quagga Router Version	**0.99.24.1**	**0.99.22.1**	**—**
Support for IPv4			Not Tested
Support for IPv6	Not supported	Not supported	Not Tested
Support for i-BGP and e-BGP			Not Tested
GTSM and MD5 authentication			Not Tested
BGP route policies			Not Tested
Support for 4 Byte ASN			Not Tested
BGP flow spec (optional)	Not Supported	Not Supported	Not Tested
Route dampening (optional)			Not Tested
Must support LDP processing	Phase 2	Phase 2	Not Tested
Peer with 20 transit peer routers each advertising 30K routes			Not Tested
Learn and program 450K routes from one transit peer			Not Tested
Learn and program 600,000 IPv4 routers and 30,000 IPv6 routes	Only IPv4 routes	Only IPv4 routes	Not Tested

All tested SDN controllers with Quagga as router met most of the core requirements. However, support for IPv6 was not available on both the controllers at the time of testing. The support for missing features is being developed by the member communities and will be included in the scope of Phase 2 testing as the features are made available. Tests were not conducted using the Open Daylight Controller as part of first phase of testing due to resource constraints.

Detailed test methodologies are available at https://github.com/onfsdn/atrium-docs/wiki/Atrium-Field-Trial

Baseline Characterization Tests

This purpose of this test section is to baseline the characteristics of control plane and to finalize the test methodology to be employed for future phases of testing. All tests in this section are conducted using Quagga as non SDN router with integrated control and forwarding plane.

The measurements in this section will be used to gather key test parameters, finalize test methodologies and to assess the performance impact when an SDN controller with an Open Flow switch is added to the test environment.

The baseline characterization tests will be repeated every time the reference hardware configuration is changed.

Case 1: Route learning time vs processor cores allotted to virtual machine

Test profile - 600K routes, 20 eBGP peers, 30K routes/peer, non-overlapping route updates, 10000 route updates/sec, 2-8 virtual cores

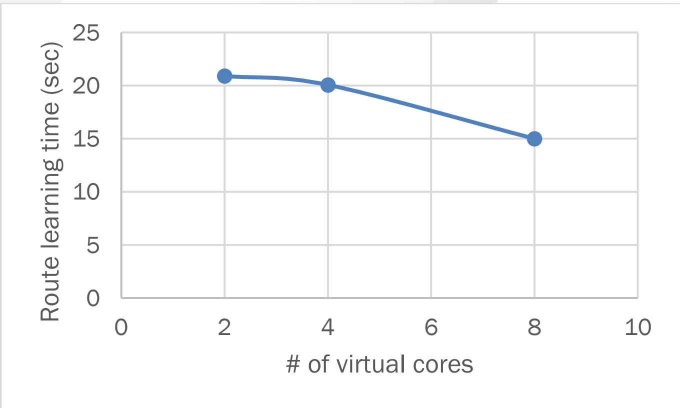

It was observed that that the performance is optimal when 6 or more virtual cores are allocated to the virtual machine that hosts Quagga. Given that in a Quagga + SDN Controller test scenario the same hardware resources will need to be shared by both the router and controller, it was decided to allocate <u>8 virtual cores for control plane VM</u> in the test-bed.

Case 2: Route learning time vs number of route updates for each BGP peer

Test profile - 600K routes, 20 eBGP peers, 30K routes/peer, non-overlapping route updates, 8 virtual cores allocated to the control plane VM, 2000-10000 route updates/sec

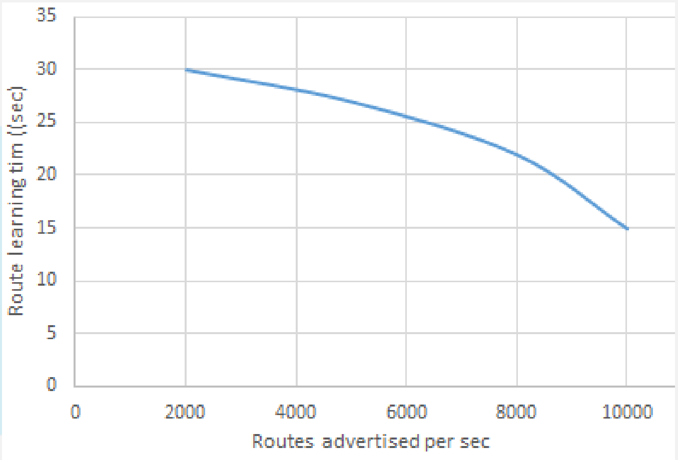

As the number of simultaneous route updates per peer is increased, the route learning times are faster. Based on the observation, the <u>number of route updates per second is set to 10000</u> for all further testing.

Case 3: Route learning times vs number of peer routers vs number of routes

Test profile - 30K routes/peer, 10000 route updates/sec, non-overlapping route updates, 8 virtual cores allocated to the control plane VM, 1K to 1M routes, 1 to 50 eBGP peers

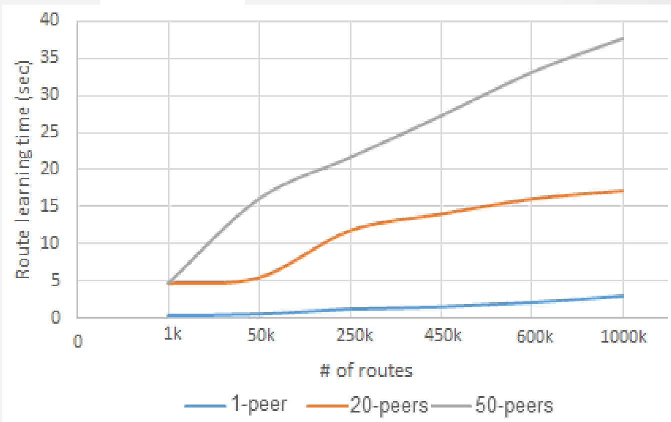

For future tests, <u>20 eBGP peers with 600K routes</u> will be used as a base profile for core measurements.

Case 4: Route learning vs route programming time using Quagga/Zebra and Linux kernel for forwarding

Test profile - 20 eBGP peers, 30K routes/peer, 10000 route updates/sec, non-overlapping route updates, 8 virtual cores allocated to the control plane VM, 1K to 1M routes

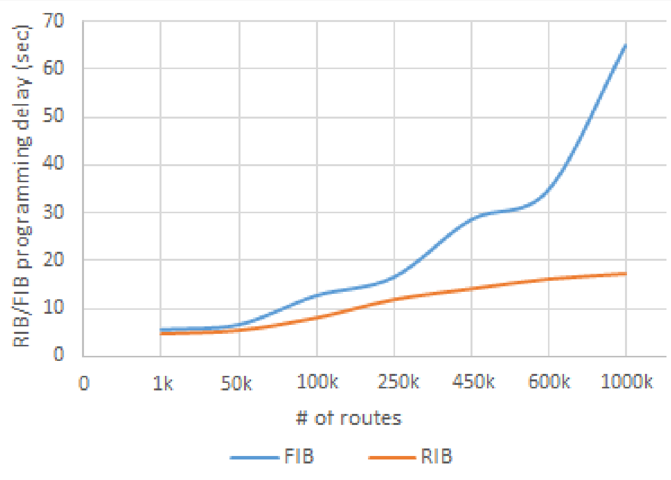

The reference data above for route learning vs route programming times for a non-SDN router will be used to compare test results when identical configurations with SDN controllers are deployed.

Flow programming time with SDN routers

This section characterizes difference in performance observed when SDN test controllers with an external Open Flow Switch is added to the base test configuration.

Test scenario: Route learning times vs route programming time using SDN controller and data plane

Test profile - 20 eBGP peers, 30K routes/peer, 10000 route updates/sec, non-overlapping routes, 8 virtual cores allocated to the control plane VM hosting both Quagga and SDN controller, 1K to 1M routes.

Flow programming times are measured on the OVS after the flows are installed and active on the Open Flow switch.

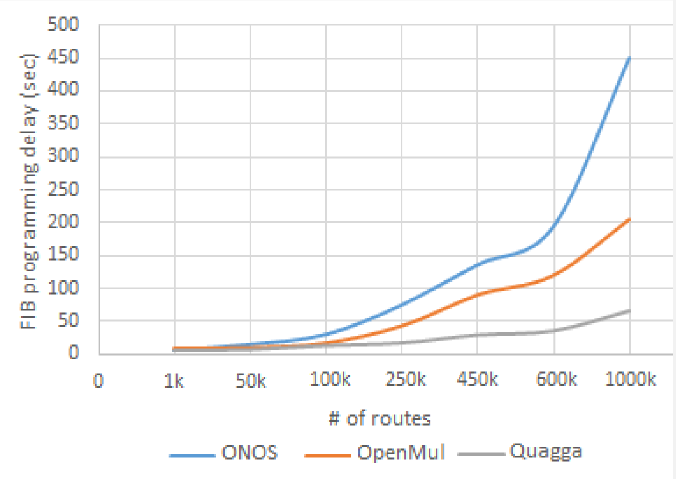

The table below captures the average time taken to program the switch for boundary conditions identified in the core requirements.

# of advertised routes	Route Learning time for ONOS (in sec)	Route Learning time for Open MUL (in sec)
600k Non-Overlapping Route Updates	194	120
600K with 30% Overlapping Route Updates	175	90

CONCLUSION

This concludes the first phase of the pre-testing efforts for the Atrium Field Trial.

The next phase of testing will validate support for IPv6 and MPLS. These two key requirements were not tested due to non-availability of these features on the SDN controllers at the time of Phase-1 testing. Phase 2 tests would also involve system level testing, characterization of the data planes and additional test scenarios based on feedback from SIFY.

Subscribe to atrium_eng@groups.opensourcesdn.org to participate or to get additional information on the progress of the field trial.