

Functional and performance comparison of open-source and commercial multicast solutions

Miha Rugelj¹, Janez Sterle², Andrej Kos³

University of Ljubljana, Faculty of Electrical Engineering, Tržaška 25, 1000 Ljubljana, Slovenia

¹ E-pošta: miha.rugelj@tfe.org

² E-pošta: janez.sterle@fe.uni-lj.si

³ E-pošta: andrej.kos@fe.uni-lj.si

Abstract. In this paper we give an overview of open-source and commercial multicast routing solutions. XORP (eXtensible Open Router Platform) is the only system which includes support for the multicast management and multicast routing protocol. We provided support for the multicast routing and multicast routing protocol (PIM sparse mode) in the core of the Linux operating system and developed an advanced open-source compatible interface between the PIM sparse mode for the internet protocol version 6 in kernel and XORP system. We performed tree performance tests to show a comparison between the XORP and the commercial routing Cisco 2811 ISR device. Our test results show that the XORP system achieves better performance results in case of multicast functionality compared to Cisco. Our results also show direct comparison between the internet protocol versions 4 and 6, slightly deviating from each other thus demonstrating maturity of multicast technology transition to the newer protocol IPv6.

Keywords: multicast, open-source and commercial solutions, XORP, Cisco 2811 ISR, multicast performance test

1 INTRODUCTION

In the last years have been developed a lot of new and advanced internet services. These comprise transmission of television, video and audio. The multicast technology [1] [2] is used in order to assure operation of these services by being able to simultaneously send packets to multiple recipients.

IPTV is a typical example of the multicast technology via the IPv4 (internet protocol version 4). The problem of quick internet evolution is the lack of the IPv4 address space. As a solution they have developed and standardized the IPv6 (internet protocol version 6). Transition of IPTV to the IPv6 has not started yet, but the technology is now used in test networks.

Many devices from different manufacturers (e.g. Cisco, Juniper) support multicast and multicast routing protocols [3] for IPv6. Furthermore, there are many open-source routing solutions (e.g. XORP, Vyatta, Quagga) [4] [5].

2 MULTICAST

The multicast technology is used to transmit live contents. Its key benefits are:

- reduced load of servers and network nodes,
- efficient use of the bandwidth and data sources.

2.1 The architecture of the multicast technology

The multicast network includes an access and distribution network (figure 1). The access network consists of multicast clients, multicast servers and edge routers. Multicast clients use the IGMP (internet group management protocol) [6] and MLD (multicast listener discovery) [7] for joining and leaving multicast groups. The distribution network consists of the edge and backbone routers supporting one of the multicast routing protocols (e.g. PIM sparse mode, PIM dense mode, BIDIR-PIM) [3].

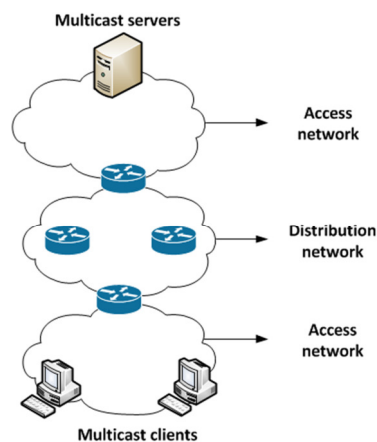


Figure 1. Architecture of the multicast technology

3 OPEN-SOURCE ROUTING SOLUTIONS

Many open-source solutions have been developed to cope with commercial weaknesses, such as dependence on the manufacturer, closed architecture and price. The open-source solutions are:

- XORP (eXtensible Open Router Platform) [4],
- Quagga [5] and
- Vyatta.

These solutions use almost the same open-source unicast routing protocol implementations. The differences among them are in additional functionalities and services as shown in table 1.

3.1 XORP

Though XORP is designed for UNIX, it also fully supports the Linux operating system. Whose characteristics are:

- scalable architecture,
- high performance packet forwarding and
- robustness of the system.

Its main disadvantages are its hardware and software. Being developed for a typical computer architecture, the CPU overload and the delay in reading data from the computer memory can limit high routing performance and forwarding packets. Studies [8] [9] prove, that the problem presents 90 % of the total delay.

XORP supports all the typical routing protocols (table 1) for IPv4 and IPv6 and is the only open-source solution to support the multicast routing protocol PIM-SM (PIM sparse mode).

3.1.1 XORP architecture

XORP consist of two subsystems [10]. The user-level subsystem consists of routing protocols and processes providing support for protocols. The kernel-level subsystem manages the forwarding path.

The system consists of a multi-process architecture with one process for each routing protocol and other processes for managing, configuring and coordinating the system. Figure 2 shows the XORP architecture.

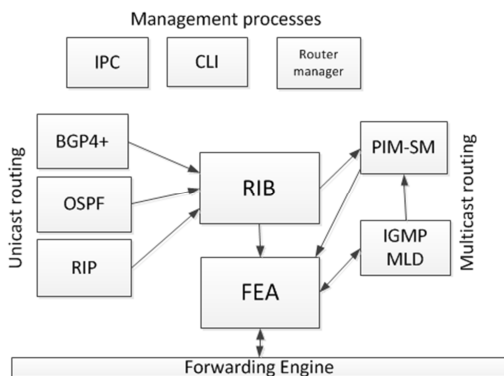


Figure 2. XORP architecture

Table 1. Functional comparison of routing solutions

Functionality	XORP 1.8.3	Vyatta Core 6.2	Quagga 0.99.18	Cisco 2811
Multicast routing protocols	PIM-SM	✗	✗	PIM-SM, PIM-DM, BI-PIM, MBGP
Multicast management protocols	IGMP, MLD	✗	✗	IGMP, MLD
Multicast SSM	✓	✗	✗	✓
Static routing	✓	✓	✓	✓
RIPv2, RIPv6	✓	✓	✓	✓
OSPFv2, OSPFv3 / IS-IS	✓ / ✗	✓ / ✗	✓ / ✗	✓
BGPv4, BGPv6	✓	✓	✓	✓
MPLS	✗	✗	✗	✓
VLAN (802.1Q)	✓	✓	✗ ¹	✓
VPN IPsec, VPN SSL	✗ ¹	✓	✗ ¹	✓
DHCP Server/Client/Relay	✗ ¹	✓	✗ ¹	✓
Route maps	Limited	✓	Limited	✓
Quality of service (VLAN Tag, Priority Queuing)	✗	✓	✗	✓
Virtual Router Redundancy Protocol	✓	✓	✗ ¹	✓

3.1.2 XRLs and Click

XORP provides a mechanism for communication between processes called XRLs (XORP resource locators) [4] and a special modular tool the Click modular router [10]. Click uses a lower subsystem and performs fast packet forwarding. The Click forwarding table is built with elements each having a function of the process.

The XRLs mechanism enables efficient integration of separate components into one coherent system. XRLs resemble the Web's URLs. An example of XRLs for a process Forwarding Engine Abstraction (FEA) is presented in figure 3.




Figure 3. Example of XRLs [4]

4 IMPLEMENTATION OF XORP

We provided supports for the multicast routing (MROUTE) and multicast PIM-SM protocol into the kernel for the IPv6. We used the Ubuntu 10.04 LTS server and the newer Linux 2.6.38.4 kernel (vanilla kernel). We provided an advanced open-source compatible interface between the PIM-SM protocol for IPv6 in the kernel and XORP system. From the kernel

¹ Functionality could be included in the Linux operating system with additional programs.

we excluded the unnecessary modules and support accessories of the Ubuntu system to optimize it.

5 TESTING AND PERFORMANCE ANALYSIS

We performed the below three multicast performance tests [11] between the XORP on Linux and Cisco 2811 ISR device:

- multicast group capacity,
- multicast packet latency and
- multicast group join delay.

Our tests include measurements of certain performance characteristic of the device and do not measure parameters of specific multicast components (e.g. building a multicast distribution tree), which could consequently affect results. In cases where we measured a single device and not a group of devices, the impact of the specific multicast components is negligible.

5.1 Simplification of the used testing devices

When analyzing results, it is necessary to take into account the following assumptions and simplifications [12]:

- the methodology assumes an uniform and stable medium between the test device and the test measurement equipment,
- all the tests were performed at the default setting on the Cisco 2811,
- functionalities, such as flow control, quality of service, protocol CDP (cisco discovery protocol) were excluded from the default setting, due to the possibility of an impact on test results and
- unnecessary modules and other functionalities were removed from the kernel.

5.2 Measurement parameters

We performed tests for one input and one output interface on the test device, for both internet protocol on the Ethernet interface and for different frame sizes:

- for IPv4 (64, 128, 256, 512, 1024, 1280 and 1518 bytes),
- for IPv6 (86, 128, 256, 512, 1024, 1280 and 1518 bytes).

We chose different minimum frame size (86 bytes) for the IPv6. 86 bytes is the minimum frame size with an added UDP (user datagram protocol) header and the signature of the Spirent Test Centre [13].

The multicast management IGMPv2 protocol was used for the IPv4 and the MLDv1 protocol for the IPv6.

5.3 Test network

The test network was placed in a controlled laboratory environment so as to eliminate any external factors that could affect our results. The test network is shown in figure 4.




Figure 4. Test network

The tests were performed by using the dedicated commercial test equipment Spirent Test Center [13].

Properties and characteristics of the used test devices are:

1. Open-source solution XORP on Linux:
 - Operating system: Ubuntu server 10.04 LTS 32-bit,
 - Linux kernel: Linux 2.6.38.4-kernel,
 - CPU: Intel Pentium 4 CPU 3.20 GHz,
 - Memory: DDR2 (2 GB) and
 - Software: XORP 1.8.3.
2. Cisco 2811 ISR:
 - IOS: C2800NM-ADVENTERPRISEK9-M, version 12.4(24)T1, release software (fc3),
 - CPU: Unknown information and
 - Memory: DRAM (512 MB).

5.4 Multicast group capacity

Our test determines the maximum number of the multicast groups a device under test (DUT) can support while maintaining the ability to forward multicast frames to all multicast groups registered to that DUT.

5.4.1 Test results

Our test results are shown in table 2. The following configuration parameters were used:

- test duration: 60 seconds,
- number of the tested egress interfaces on DUT: 1,
- IGMP and MLD version: IGMPv2 (IPv4) and MLDv1 (IPv6) and
- offered load: 100 Mbit/s.

Table 2. Results of the multicast group capacity

Frame size [byte]	size	Multicast group capacity (XORP)		Multicast group capacity (Cisco)	
		IPv4	IPv6	IPv4	IPv6
64	86	122	138	98	115
128		164	154	116	134
256		223	190	165	179
512		260	212	181	210
1024		283	234	201	231
1280		310	281	231	283
1518		360	344	285	319

5.4.2 Analysis of test results

Our test results show that XORP supports more multicast groups for the IPv4 compared to IPv6 for all frame sizes. Such results are expected because of the size of the IPv6 address space which is four times larger than IPv4. The larger address space affects the size and the processing power of maintaining unicast and multicast routing table. Compared to IPv6, XORP supports on average 24 multicast groups more (representing ~11 %) for the IPv4.

The results show that Cisco supports more multicast groups (on average 27 multicast groups, representing ~15 %) for IPv6 compared to IPv4. The reason is in the lower packets loss for the IPv6, which increases the probability of a successful iteration for a certain number of the multicast groups.

Table 2 shows that the XORP system supports more multicast groups compared to Cisco:

- on average 63 multicast groups (representing ~35 %) for the IPv4 and
- on average 11 multicast groups (representing ~6 %) for the IPv6.

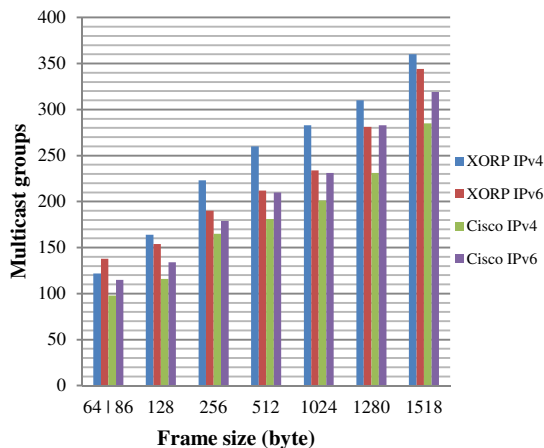


Figure 5. Multicast group capacity (100 Mbit/s)

5.5 Multicast packet latency

Our test determines the multicast latency from a single multicast ingress interface of a DUT through a single egress multicast interfaces of the same DUT.

5.5.1 Test results

Our test results are shown in figures 6 and 7. They were performed for 10 and 100 multicast groups. The following configuration parameters were used:

- test duration: 120 seconds,
- number of the tested egress interfaces on DUT: 1,
- IGMP and MLD version: IGMPv2 (IPv4) and MLDv1 (IPv6) and
- offered load: 10 Mbit/s.




Figure 6. Multicast packet latency (10 multicast groups)

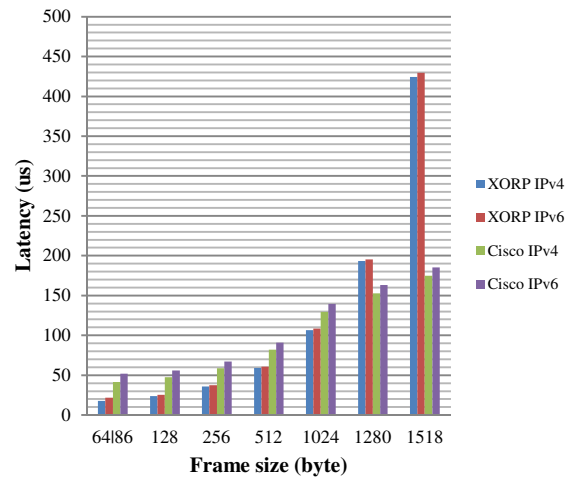


Figure 7. Multicast packet latency (100 multicast groups)

5.5.2 Analysis of test results

XORP achieved minimum packets latency (17.214 us) at the frame size of 64 bytes (10 multicast groups, IPv4) and maximum packet latency (429.355 us) at the frame size of 1518 bytes (100 multicast groups, IPv6).

Our results show, that the packet latency is independent of the multicast groups (the difference in the maximum latency between 10 and 100 multicast groups is 6.423 us, which represents ~1.5 % of the

average latency at 1518 bytes and IPv6) for the XORP system. The impact on the latency of the internet protocol is the same (according to the results obtained for IPv4 the latency is increased by ~2.512 us at IPv6 and 100 multicast groups, representing ~2 % of the average latency). The largest impact on the delay is that of the frame size (the difference in the latency between 86 and 1518 bytes at the 100 multicast groups and IPv4 is 407.454 us).

Cisco achieved a higher latency compared to XORP. A lower latency was achieved only for the frame sizes of 1280 and 1518 bytes. Our results show that the number of multicast groups has a little impact on the latency for Cisco. The major impact has the internet protocol, but much less than the frame size.

The maximum difference in the latency between the devices (~2.46 times) was achieved at the frame size of 1518 bytes, 10 multicast groups and IPv4 protocol.

5.6 Multicast group join delay

With our test we determined the time it takes a DUT to start forwarding multicast frames after a successful IGMP group membership report is issued to it. The multicast group join delay measurement may be affected by the state of the multicast forwarding database² (MFDB).

The join delay is the difference in the time when the IGMP group membership message is sent and the first frame of the multicast group is forwarded to a receiving egress interface.

We performed tests for the method A² at 10 and 100 multicast groups and offered 10 Mbit/s load.

5.6.1 Test results

Our test results are shown in figures 8 and 9. The following configuration parameters were used:

- number of the tested egress interfaces on DUT: 1,
- IGMP and MLD version: IGMPv2 (IPv4) and MLDv1 (IPv6) and
- offered load: 10 Mbit/s.

5.6.2 Analysis of test results

Results given in figures 8 and 9 show that the internet protocol and frame size have a relatively small impact on the multicast group join delay for the XORP device. The maximum impact is that of the multicast groups with the difference in the delay between 10 and 100 multicast groups (the frame size of 1518 bytes, IPv4) increased by ~4.22 times.

An additional factor impacting the delay (Cisco) is internet protocol:

- the difference between the internet protocols (64186 bytes, 10 multicast groups) is approximately ~2.15 times and

- the difference between the 10 and 100 multicast groups (64 bytes, IPv4) is approximately ~6.54 times.

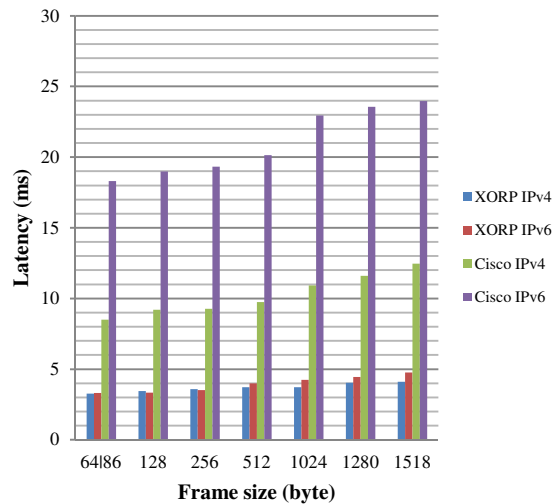


Figure 8. Multicast group join delay (10 multicast groups)

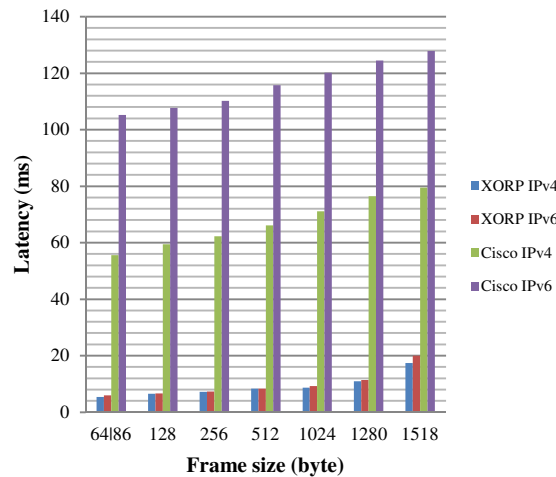


Figure 9. Multicast group join delay (100 multicast groups)

The difference between the two used testing devices is considerable. It is shown that XORP achieves less join delay on all segments. The maximum difference between the devices is at the IPv6 (86 bytes, 100 multicast groups), where XORP achieved a ~17.75 times lower join delay.

6 CONCLUSIONS

By using Linux and XORP we achieve a very useful and performance multicast routing device. Our results show that XORP exhibits better performance in case of multicast functionality compared to Cisco 2811. With the XORP system, implementation of the necessary functionalities in Linux and performance tests we prove that the open-source XORP solution is more appropriate

² Method A: MFDB at startup does not contain any multicast group address.

and a better alternative in terms of multicast for being used in small and medium-sized enterprises.

It should be noted, that the used hardware affects the test results. Irrespective of the fact that the hardware specification of the XORP system is better, at the test maximum load CPU does not exceed 10 % and the used memory space does not exceed 16 % of the available resources. It is therefore more likely that XORP will achieve similar results on a lower hardware specification. Besides its primary functionalities the device is also useful for many other advance services.

Results also show a direct performance comparison between the IPv4 and IPv6 protocols. As seen from the results, there is notice a small deviation due to the internet protocol, thus proving maturity of the multicast technology transition to newer protocol IPv6.

It should be noted that there is still room for further improvements (e.g. optimization of the cache device) which would improve performance and consequently testing of open-source and commercial routing devices.

REFERENCES

- [1] J. Sterle, I. Humar, S. Tomažič, S. Kos, Varnost v IP-multicast, Zbornik ERK, 2006, zv. A, str. 167-170.
- [2] IETF RFC 1112, Host Extensions for IP Multicasting, August 1989, <http://www.ietf.org/rfc/rfc1112.txt>.
- [3] IETF RFC 2362, Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification, June 1998, <http://www.ietf.org/rfc/rfc2362.txt>.
- [4] M. Handley, O. Hodson, XORP: An Open Platform for Network Research, <http://www.xorp.org>, September 2011.
- [5] K. Ishiguro, Quagga 0.99.4, July 2006, <http://www.quagga.net>.
- [6] RFC 2236, Internet Group Management Protocol, Version 2, November 1997, <http://www.ietf.org/rfc/rfc2236.txt>.
- [7] IETF RFC 2710, Multicast Listener Discovery (MLD) for IPv6, October 1999, <http://www.ietf.org/rfc/rfc2710.txt>.
- [8] M. Pustišek, A. Kos, J. Bešter, Architecture-dependent packet switch performance under imbalanced traffic. *Elektrotehniški vestnik*, 2005, letn. 72, št. 1, str. 36-44.
- [9] V. Eramo, M. Listanti, OSPF Performance and Optimization of Open Source Routing Software, *International Journal of Computer Science & Applications*, Vol. 4, No. 1, str. 53 – 68, January 2007.
- [10] M. Handley, A. Greenhalgh, XORP: Breaking the Mould in Router Software, <http://www.xorp.org>, September 2011.
- [11] RFC 3918, Methodology for IP Multicast Benchmarking, October 2004, <http://www.ietf.org/rfc/rfc3918.txt>.
- [12] I. Humar, M. Golja, J. Bešter, Sistem za merjenje prometa v omrežjih z internetnim protokolom. *Elektrotehniški vestnik*, 2005, letn. 72, št. 2/3, str. 71-78.
- [13] Spirent Test Center, <http://www.spirent.com/>, september 2011.
- [14] A. Kos, J. Bešter, Delay characteristic of the synchronous bulk packet switch loaded with real traffic, *Elektrotehniški vestnik*, 2004, vol. 71, no. 1-2, str. 70-75.
- [15] M. Handley, E. Kohler, Designing Extensible IP Router Software, <http://www.xorp.org>, September 2011.
- [16] A. Locker, H. Payer, Linux IPv6 Multicast Routing, June 2005, http://student.cosy.sbg.ac.at/~hpayer/praxis/ipv6mf_pimsm.pd.

Miha Rugelj graduated in 2011 from the Faculty of Electrical Engineering, University of Ljubljana. He currently works towards his PhD at the Laboratory for Telecommunication at the same faculty. His research interests include network technologies and telecommunication services.

Janez Sterle graduated in 2003 from the Faculty for Electrical Engineering, University of Ljubljana, where he is currently pursuing his postgraduate degree. Since 2002 he has been working at the Laboratory for Telecommunications at the same faculty. His educational, research and development work is oriented towards design and development of the next generation networks and services. His current research areas include the next generation internet protocol, network security, traffic analyses, QoE modeling, QoS measuring and development and deployment of new integrated services into fixed and mobile networks.

Andrej Kos, Ph.D., is an assistant professor at the Faculty of Electrical Engineering, University of Ljubljana. He has acquired an extensive research and industrial experience in analysis, modeling and design of advanced telecommunication elements, systems, and services. His current work focuses on managed broadband packet switching and next generation intelligent converged services.