

DYNAMIC ROUTING OF IP TRAFFIC BASED ON QoS PARAMETERS

Martin Kriška¹, Jozef Janitor² and Peter Fecilak³

¹Computer Networks Laboratory, Technical University of Kosice, Slovakia

²Institute of Computer Technology, Technical University of Kosice, Slovakia

³Department of Computers and Informatics, Technical University of Kosice, Slovakia

ABSTRACT

The article looks into the current state of the art of dynamic routing protocols with respect to their possibilities to react to changes in the Quality of Service when selecting the best route towards a destination network. New options that could leverage information about the ever changing QoS parameters for data communication are analysed and a Cisco Performance Routing solution is described more in detail. The practical part of this work focuses on a design and implementation of a test bed that provides a scalable laboratory architecture to manipulate QoS parameters of different data communications flowing through it. The test bed is used in various use cases that were used to evaluate Cisco Performance Routing optimization capabilities in different scenarios.

KEYWORDS

Performance Routing, PFR, Quality of Service, QoS, Optimized Edge Routing

1. INTRODUCTION

In the field of computer networks, the traditional task of dynamic routing protocols was, and still is, to provide loop free reachability between distant IP [1] networks and communicating parties. In today's world, where the network is not anymore only a file transport service, this traditional approach, while still fulfilling the requirements it was designed for decades ago, is no longer seen as sufficient.

With the invent of new applications of computer networks, converged architectures that integrate Data, Voice, Video and other real-time sensitive services, new requirements are forming which no longer see data communication traffic as only packets passing between source and destination machines. Nowadays, a deeper visibility into data communication is needed, so that different applications inside the data channels passing information between the source and destination machines can be distinguished from each other. Once such a visibility is available, routing protocols can leverage this information and different classes of applications, depending on their usefulness, or priorities, can be treated differently in terms of Quality of Service leading to different routing paths for different classes of communicating applications.

The topic of this work is to give indication and demonstrate on how routing decisions can take advantage of additional information about the current Quality of Service parameters along the whole path. This work also attempts to give multiple views of how such routing can be achieved.

2. GOALS

The goal of this work is to take a deeper look into the possibilities of using currently available technologies and their capabilities, to provide routing decisions based on dynamically changing Quality of Service parameters inside the computer network. The respective work also proposes a design of a laboratory test bed environment, which can manipulate various QoS parameters of data communication traffic flows, flowing through it. Once such an environment exists, the Cisco Performance Routing solution is deployed on top of the test bed to demonstrate its possibilities. Optimization capabilities are illustrated on a simple example.

3. ANALYSIS

A new need to provide IP routing based on the communicating applications, rather than based on the shortest (or cheapest) path, brings the question whether the currently used dynamic routing protocols can address this challenge. Dynamic routing protocols both interior and exterior such as EIGRP, OSPF, IS-IS and BGP were analysed in this work [2].

The OSPF protocol defines its metric as a cost of traversing link and this cost is inversely proportional to bandwidth available on specific interface. This value is configurable by administrator but it neither does change dynamically nor does it have to reflect actual theoretical bandwidth of interface.

Integrated IS-IS also defines its metric as a cost of traversing link, but by default on Cisco devices all interfaces are equal regardless of available bandwidth or any other parameters. It is up to administrator to modify cost values on per-interface basis to achieve suitable routing behaviour. Cisco proprietary (recently with some extent partially open sourced [3]) EIGRP uses a formula for calculating metric, which is very similar to its predecessor protocol IGRP. Parameters such as bandwidth, delay, load and reliability can be taken into an account when calculating composite metric value. Yet, by default only the first two of them – bandwidth and delay – are used. Among the reasons for this default behaviour is that EIGRP, unlike its predecessor IGRP which had been sending updates periodically, generates updates only when there is a significant change in the network. This is also the only time when the load and reliability parameters are read to distil the final composite metric. If these values change over time, which they indeed do, but there is no other change in the network, the metric for passive routes is not recalculated [4].

BGP as the only representative of EGP class of dynamic routing protocols is also known as path-vector protocol. It uses multiple attributes and a rather long decision process which results in selecting the best route according to the routing policy implemented by the network operator [5]. However, none of these attributes represent Quality of Service parameters since the main goal of BGP is to provide a stable and scalable routing between interconnected autonomous systems [6]. If the underlying network is Multiprotocol Label Switching enabled, then MPLS tags are used in routing process instead of destination IP addresses. The ability to stack more than one MPLS tag provides network administrators opportunity to provide additional services such as MPLS Virtual Private Networks, Any Transport over MPLS and MPLS Traffic Engineering and not just loop free reachability. It is the deployment of MPLS Traffic Engineering tunnels that enables network administrators to utilize other available paths for various traffic flows and not just the one with the lowest metric according to routing protocol used. Dynamic routing protocols with extensions for Traffic Engineering such as OSPF or IS-IS carry additional information in routing updates which are necessary so that other routers know about conditions inside the computer network and can compute Traffic Engineering tunnels accordingly. Current implementation of MPLS Traffic Engineering tunnels as described in [7] supports creating of tunnels where the only parameter

taken into consideration is requested amount of bandwidth. Head-end router knowing all available paths to destination and available bandwidth along these paths can select route that fulfils the requirement. If MPLS Traffic Engineering tunnels are used, it is possible to route different traffic classes based on their bandwidth requirements and current link utilization inside the network which is communicated via dynamic routing protocol updates. Despite the fact that no other parameters than bandwidth are currently being used when constructing Traffic Engineering tunnels, drafts of documents that describe support for additional parameters, namely packet loss and delay, exists in RSVP-TE extensions for Loss and Delay Traffic Engineering [8]. If these drafts will ever be incorporated into standard, MPLS Traffic Engineering tunnels will enable network operators to route traffic classes with respect to dynamic and ever-changing QoS parameters that delay, packet loss and required bandwidth without a doubt are.

It is clear that none of the previously named dynamic routing protocols nor MPLS Traffic Engineering tunnels take realtime QoS related parameters into consideration in their decision process. Additional logic is needed to enforce routing changes whenever dynamic QoS parameters change.

Software Defined Networking well known for its idea of decoupling control plane from data plane in networking devices and centralizing control plane logic on a separate controller for whole computer network [9]. This will provide visibility into network traffic flows as well as additional logic required for routing different traffic flows according to their QoS needs. This controller with a complete visibility into traffic flows and knowledge of current conditions inside the computer network can calculate the optimal route for traffic flow based on its specific QoS requirements. This new paradigm shift in networking field is excellent for providing routing based on QoS parameters in greenfield installations, where all networking equipment is new and SDN capable. On the other hand incremental equipment upgrade is far more often the case, which means that new SDN capable hardware has to work alongside legacy equipment which will never support features of Software Defined Networking. This certainly limits benefits gained by deploying SDN in some parts of the network when visibility into traffic flows and routing based on QoS parameters is needed for the whole computer network. Once all network equipment support SDN, network operators will be able to benefit from all new possibilities, to name just a few like Topology Independent Forwarding and Routing for Dollars [10] [11].

The necessity of having all equipment SDN capable as well as lack of publicly available documentation during its early stages of development meant that it was not chosen as a solution used in this work. The additional logic used in this work is the Cisco Performance Routing solution [12] due to support of multiple commonly used router platforms, both brand new as well as older ones. It consists of a centralized master controller which is responsible for evaluating collected performance characteristics and if necessary instructs the so called border routers to apply new changes to the routing processes. Border routers are responsible for collecting performance characteristics, as well as for enforcing routing changes. The Cisco Performance Routing solution can either only passively monitor traffic flows, or actively generate IP SLA probes to collect performance statistics, which are later compared against a predefined optimization policy. If the performance requirements are not met on the current interface, then the master controller can reroute the traffic to another interface which is able to provide performance characteristics in compliance with the predefined optimization policy.

Multiple optimization techniques such as insertion of more specific static route, or BGP route, as well as dynamic policy-based routing for the entire prefix, or just a specific application are available within the analysed solution. It is also possible to manipulate the BGP LOCAL_PREFERENCE attribute for outgoing traffic and AS_PATH attribute for incoming

traffic. Appending communities to BGP updates to manipulate routing inbound requires cooperation with ISP but is also available.

The biggest improvement that the network operator and the network users can get by using the Cisco Performance Routing solution is a per application based different routing behaviour. Scavenger traffic, such as YouTube and other not business critical applications can be routed over a path with a lower bandwidth, while business critical traffic, such as VoIP, CRM, Salesforce can be routed over a high bandwidth link. Should the QoS parameters on the business critical primary link change, the Cisco Performance Routing solution can dynamically detect such a change and apply new routing decisions making sure the business critical applications are always using the best possible path.

4. SOLUTION AND RESULTS

4.1. The design of a test bed to manipulate QoS parameters of data flows

The ambition to use the Cisco Performance Routing solution requires the ability to build a network where we are capable of granularly modifying QoS parameters of data communication flowing through the network, thus simulating a real behaviour of public WAN networks. The test bed was designed to reach this goal in laboratory conditions. This test bed consists of physical networking equipment as well as virtual machines which are interconnected according to physical topology shown in Figure 1.

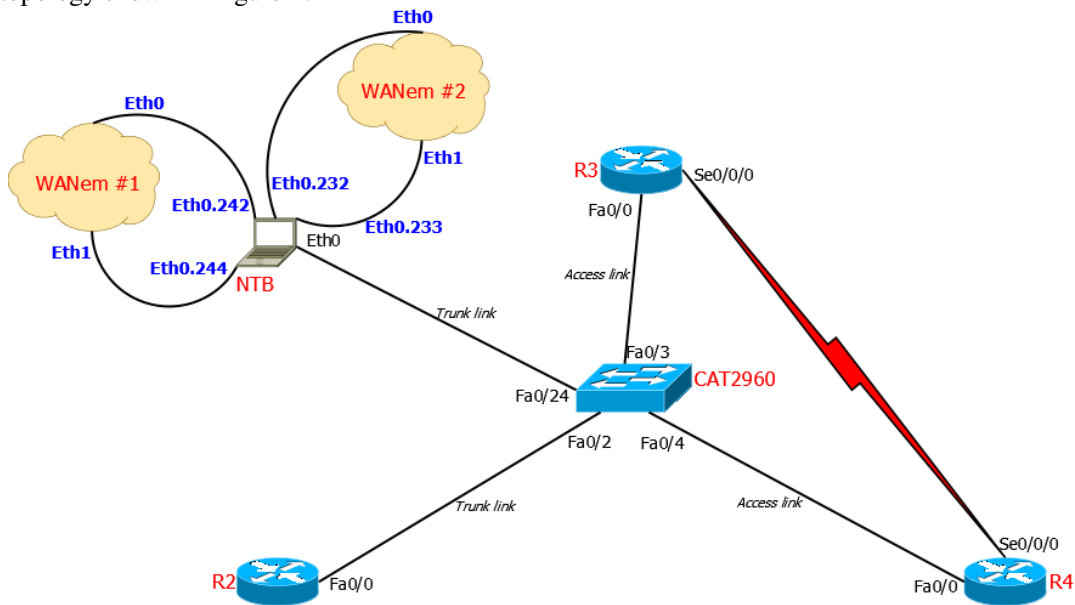


Figure 1. Example shows physical topology of proposed test bed which is used throughout this work

IEEE 802.1Q tagging and VLANs are used to create logical separation of traffic passing through a single physical interface on the physical host named NTB, as well as to provide subinterfaces to which virtual machines Ethernet network interface cards can be bridged. Virtual machines, in Figure 1 named WANem #1 and WANem #2, run a live Knoppix Linux distribution [13] without the need for installation. Since no dynamic routing protocol is by default available in the used Live Linux distribution and a hard disk installation was not desired for this purpose, the problem of routing data traffic passing through these virtual machines had to be solved. Adding static routes for all prefixes in the test bed network is possible but time-consuming and error-prone as more and more networks are added and as the topology grows. When choosing a solution, a

decision was made to build an overlay network on top of the physical topology shown in Figure 1. With such a design the virtual machine only needs to have routing information for its directly connected interfaces, which are always present anyway. Generic Routing Encapsulation was selected as the carrier protocol for data traffic using IPv4 as the transport protocol in this test bed. The physical topology as well as the logical topology with the overlay network built on top of it is shown in Figure 2. This way, routers are unaware of the presence of virtual machines in between the path at the cost of having additional logical tunnel interface and increased overhead associated with GRE encapsulation.

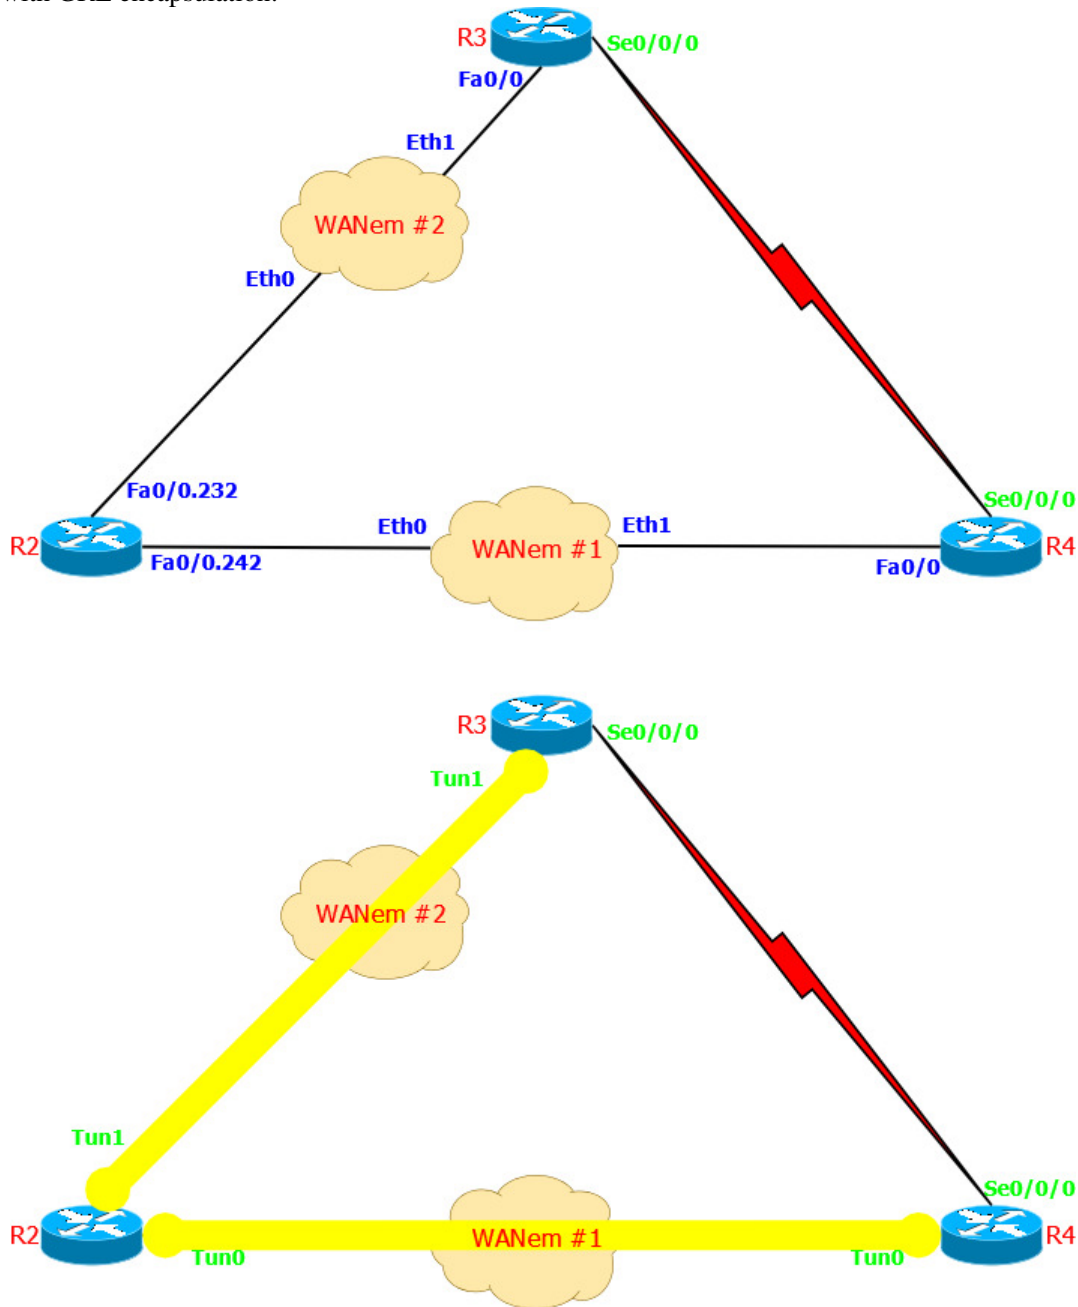


Figure 2. Example shows logical topology of proposed test bed as well as the overlay network using GRE tunnels

Communicating computer workstations are connected to the LAN interfaces of routers R2 and R4. As shown in Figure 2 to communicate with each other there are two possible paths. The first path is using the Tunnel0 interface between routers R2 and R4. The second path is using the Tunnel1 interface between R2 and R3 and then the Serial0/0/0 interface interconnecting routers R3 and R4. In both cases, traffic is flowing through a separate WANem virtual machine which enables to modify and tune QoS parameters of traffic passing through each virtual machine separately and therefore modify and tune QoS parameters of the two existing paths interconnecting computer workstations. If a computer connected to router R2 communicates with a workstation connected to router R4 and this traffic traverses the WANem #1 virtual machine, then the traffic flow through the physical topology of proposed test bed is shown in Figure 3 with green arrows:

1. Traffic flows from router R2 towards the switch CAT2960
2. Traffic flows from the switch CAT2960 towards the physical host NTB
3. Traffic flows from within physical the host NTB towards the virtual machine WANem #1
4. Traffic flows from the virtual machine WANem #1 towards physical the host NTB
5. Traffic flows from physical the host NTB towards the switch CAT2960
6. Traffic flows from the switch CAT2960 towards the router R4

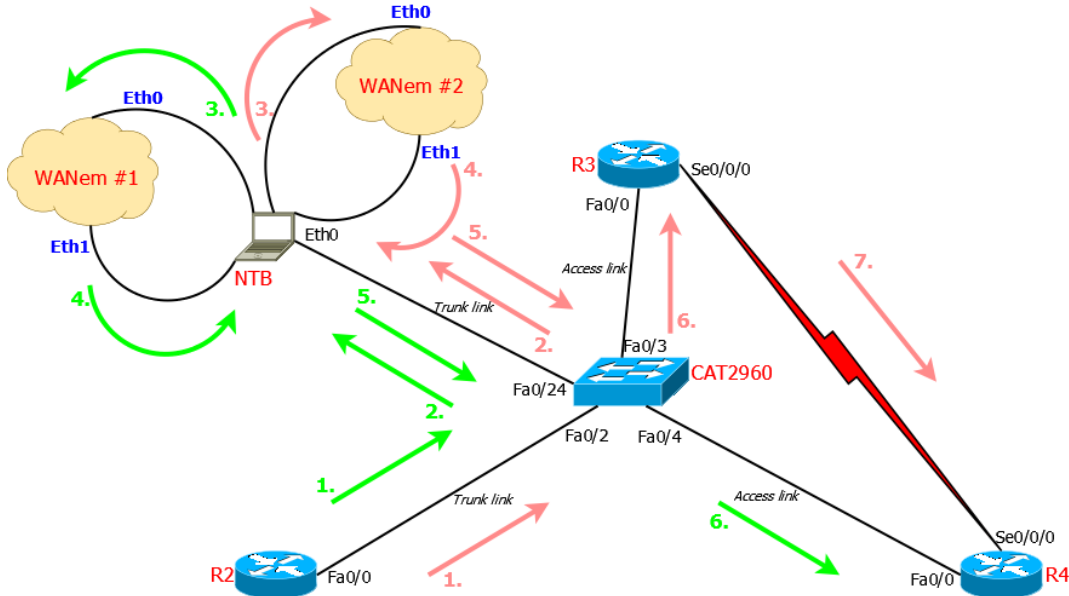


Figure 3. Example shows traffic flow through physical topology of proposed test bed from workstations connected to router R2 towards workstations connected to router R4

If the workstation connected to router R2 communicates with the computer connected to router R4 and this time this traffic flow traverses the WANem #2 virtual machine, then the traffic flow through physical topology of proposed test bed is shown in Figure 3 with pink arrows. The difference when compared to previous detailed description of traffic flow via WANem #1 is obviously flowing through different virtual machine, WANem #2 and then in the last two steps where traffic from switch CAT2960 flows towards router R3 and then it traverses the serial interface towards router R4.

Using WANem virtual machines enables us to modify and tune QoS parameters on ingress interfaces separately within virtual machine. This means that the proposed test bed allows the modification of different QoS parameters for requests and replies if these are flowing through the

same virtual machine. If this is not the case and asymmetrical routing exists then different QoS parameters for requests and replies can be enforced by setting these on different virtual machines.

4.2. Optimization of a destination prefix with usage of dynamic routing protocol

This section provides simple example how the proposed test bed can be utilized to demonstrate the Cisco Performance Routing optimization capabilities with topology shown in Figure 4. The branch office location is connected to company HQ via two separate links from two different ISPs for redundancy reason and is using the Cisco Performance Routing solution to provide routing based on QoS parameters. From the branch office point of view, the link via ISP_1 is preferred to reach company HQ and the other link is used as a backup link. Optimization policy dictates to monitor the latency of a communication between the branch office and company HQ and the threshold for latency is set to 250 milliseconds.

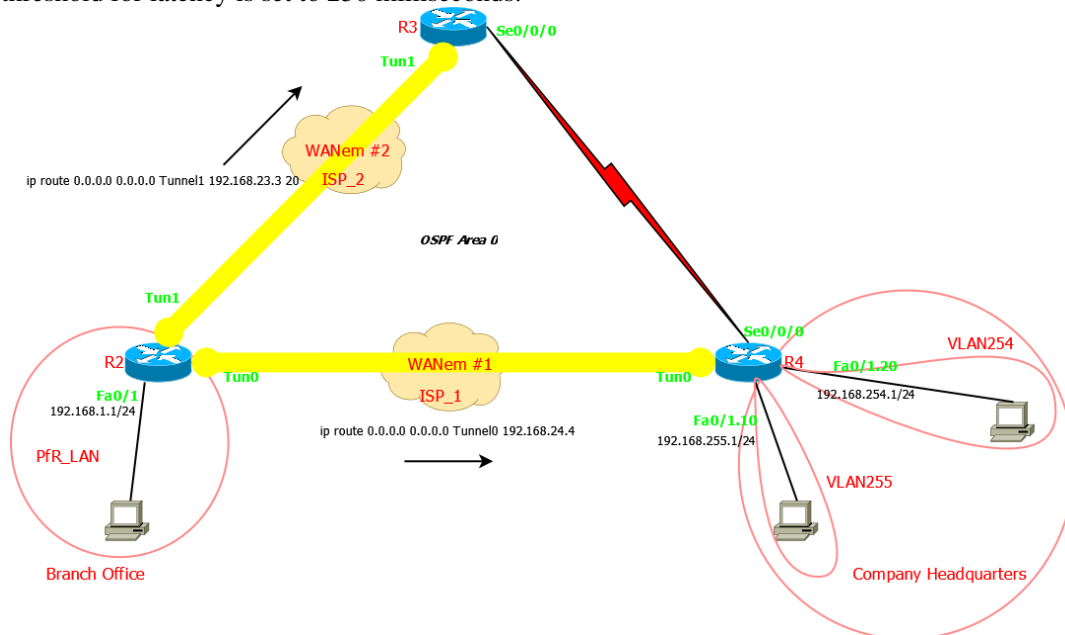


Figure 4. Example shows simple topology used to demonstrate traffic flow optimization

At the beginning, a data communication stream towards host company HQ is initiated from the branch office location. During normal operation of network, both links on R2 are in compliance with the optimization policy. Both provide latency of roughly the same value - 100 milliseconds as shown in Figure 5. Since the link via ISP_1 is preferred and both links provide acceptable Quality of Service, this link is used to route communication towards company HQ. Cisco Performance Routing is monitoring delay on both links, but as they both comply with the policy, it does not enforce any routing changes and OSPF is responsible for routing now.

```

Most recent data per exit
Border      Interface  PasSDly  PasLDly  ActSDly  ActLDly
*1.1.1.1    Tu0        0         0        104      104
1.1.1.1     Tu1        0         0        108      108
    
```

Figure 5. Output shows latency via both exit interfaces during normal operation

To simulate an anomaly inside the ISP_1 network, which causes customer traffic from the branch office location towards the company HQ to suffer from increased latency, the proposed test bed is used to manipulate this specific QoS parameter. Delay on an ingress interface Eth0 of the WANem #1 virtual machine is set to the value of 300 milliseconds as shown in Figure 6.

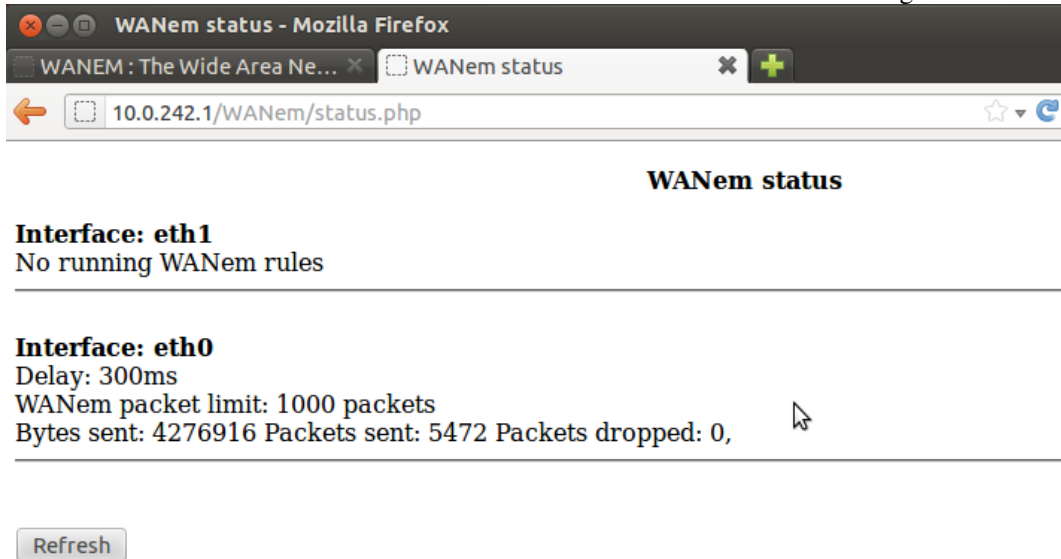


Figure 6. Output shows latency manipulation on WANem #1 virtual machine

Cisco Performance Routing keeps actively monitoring performance characteristics of traffic flow from the branch office towards company HQ, while at the same time the test bed manipulates the latency with the WANem #1 virtual machine. After detecting that latency on the currently chosen primary link by OSPF increased beyond the value acceptable by the optimization policy, it tries to evaluate if the backup link can provide a better path based on the specified QoS parameters - in this case delay only. The backup link is indeed capable of providing better latency compared to the link via ISP_1 that is experiencing problems, or so called brownout, as shown in Figure 7.

```

Most recent data per exit

Border      Interface    PasSDly    PasLDly    ActSDly    ActLDly
1.1.1.1     Tu0          0          0          404        404
*1.1.1.1    Tu1          0          0          108        108
    
```

Figure 7. Output shows latency via both exit interfaces during ISP_1 problems

Cisco Performance Routing used IP SLA probes towards the destination of traffic flow residing in company HQ location. Once it has detected noncompliance of primary link through which the traffic is flowing with optimization policy and after identifying another link that can satisfy QoS requirements, it took action and modified routing and therefore now traffic flows from the branch office location towards company HQ location as shown in Figure 8.


```
%PFR_MC-6-ROUTE_EVENT_INFO:      Route      changed      Prefix
192.168.255.0/24, BR 1.1.1.1, i/f Tu1, Reason Delay, OOP Reason
Delay
```

Figure 8. Output shows modification of routing by Cisco Performance Routing. Instead of generating a new routing protocol update information for any dynamic routing protocol deployed in all possible scenarios, whenever Cisco Performance Routing needs to enforce a new routing policy, a Protocol Independent Route Optimization feature is used. This means that instead of influencing a dynamic routing protocol itself and propagating such a change deep into routing domain causing recomputation and when misconfigured, even routing instability, dynamic policy based routing solves the problem without directly involving whatever dynamic routing protocol is used. Dynamic policy based routing can direct traffic flow from the border router selected by a dynamic routing protocol to a different border router that is currently the optimal exit router for such a traffic flow. Once the traffic flow is on correct a border router, then the dynamic policy based routing can divert this traffic flow to the exit interface that is currently optimal for this specific traffic flow. This happens on behalf of instructions from the master controller router and requires no intervention from the network operator's point of view, except the initial configuration of the Cisco Performance Routing solution.

The fact that this dynamic policy based routing, as well as dynamic access list, were created dynamically and not configured by the network operator is confirmed by running various commands as shown in Figure 9, Figure 10 and Figure 11.

```
R2#show running-config | section route-map
R2#
R2#show route-map
R2#
R2#show route-map dynamic
route-map OER_INTERNAL_RMAP, permit, sequence 0, identifier 335544323
  Match clauses:
    ip address (access-lists): oer#3
  Set clauses:
    ip next-hop 192.168.23.3
    interface Tunnell
  Policy routing matches: 299 packets, 34086 bytes
Current active dynamic routemaps = 1
R2#
```

Figure 9. Output shows dynamically created route map by Cisco Performance Routing

```
R2#sh running-config | section access-list
R2#
R2#sh ip access-lists
R2#
R2#sh ip access-lists dynamic
Extended IP access list oer#3
    1073741823 permit ip any 192.168.255.0 0.0.0.255 (299 matches)
    2147483647 deny ip any any
R2#
```

Figure 10. Output shows dynamically created access list by Cisco Performance Routing

```
R2# show running-config interface fastEthernet0/1
Building configuration...

Current configuration : 162 bytes
!
interface FastEthernet0/1
  description INTERFACE TOWARDS LAN ON ROUTER R2
  ip address 192.168.1.1 255.255.255.0
  ip ospf 1 area 0
  duplex auto
  speed auto
end

R2#show ip interface fastEthernet 0/1 | include route map
  Policy routing is enabled, using route map OER_INTERNAL_RMAP
R2#
R2#show interfaces fastEthernet0/1 | include Description
  Description: INTERFACE TOWARDS LAN ON ROUTER R2
```

Figure 11. Output shows dynamically configured policy based routing by Cisco Performance Routing

5. CONCLUSIONS

In this paper we analysed various possible approaches, capable of providing routing of data flows based on their QoS requirements. The proposed design of a laboratory test bed, which allows to dynamically manipulate QoS parameters of data traffic flowing through it provides the opportunity to simulate real WAN network characteristics in a simulated laboratory environment. The proposed test bed provides a scalable solution since multiple instances of virtual machines can run on the same physical machine. If available bandwidth on the interface becomes the bottleneck and it is not possible to install another NIC, then the virtual machines can be distributed among multiple physical hosts, which also lowers the strain on CPU and RAM resources.

The disadvantage of the proposed test bed is that logical tunnel interfaces are needed to be created for every link traversing the virtual machine and also reachability to the far-end tunnel endpoint requires for example a static route entry. It means that some configuration is needed on these routers and it is not totally transparent for the initial use. This could be solved at a cost of installing additional pair of devices per tunnel and using for example L2TPv3 pseudowires.

The proposed test bed for manipulating QoS parameters of traffic flowing through it can be useful in various fields not limited to university sector only. Developers of client-server based applications are given an opportunity to test their applications in various QoS conditions, typical for their application usage scenario over the Internet, instead of testing the application on localhost only. Experimenting with VoIP technologies can also benefit from the test bed, as it allows demonstration of impact of QoS parameters on voice traffic and how different codecs can handle degraded network conditions even in a laboratory environment.

As discussed in Section 4.2, Cisco Performance Routing was periodically measuring if performance characteristics of data flow are in compliance with defined policy. After manipulating QoS parameters in the network using the proposed test bed, Cisco Performance Routing detected this event and acted accordingly, hence providing routing of traffic based on dynamically changing QoS parameters.

ACKNOWLEDGEMENTS

This article is a result of the Project implementation: University Science Park TECHNICOM for Innovation Applications Supported by Knowledge Technology, ITMS: 26220220182, supported by the Research & Development Operational Programme funded by the ERDF. We support research activities in Slovakia/This project is being co-financed by the European Union.

REFERENCES

- [1] Information Sciences Institute, University of Southern California. RFC 791 INTERNET PROTOCOL - DARPA INTERNET PROGRAM, PROTOCOL SPECIFICATION. s.l. : Internet Engineering Task Force, 1981.
- [2] Cisco Systems, Inc. Route Selection in Cisco Routers. Cisco. [Online] 2008. [Date: 25th of October 2013.] <http://www.cisco.com/image/gif/paws/8651/21.pdf>.
- [3] D. Savage, et. al.: Enhanced Interior Gateway Routing Protocol. IETF. [Online] 2013 [Date: 25th of October 2013.] <http://tools.ietf.org/html/draft-savage-eigrp-00>.
- [4] Teare Diane: Implementing Cisco IP Routing (ROUTE) Foundation Learning Guide. Indianapolis: Cisco Press, 2010. ISBN 1587058820.
- [5] Cisco Systems, Inc. BGP Best Path Selection Algorithm. Cisco. [Online] 2012. [Date: 25th of October 2013.] <http://www.cisco.com/image/gif/paws/13753/25.pdf>.

- [6] Doyle Jeff, Carroll Jennifer: CCIE Professional Development Routing TCP/IP Volume I. Indianapolis: Cisco Press, 2006. ISBN 1587052024.
- [7] D. Awduche, et. al.: RSVP-TE: Extensions to RSVP for LSP Tunnels. IETF. [Online] 2013 [Date: 11th of November 2013.] <http://tools.ietf.org/html/rfc3209>.
- [8] X. Fu, et. al.: RSVP-TE extensions for Loss and Delay Traffic Engineering. IETF. [Online] 2013 [Date: 11th of November 2013.] <http://tools.ietf.org/html/draft-fuxh-mpls-delay-loss-rsvp-te-ext-02>.
- [9] Z. Seils. Defining SDN Overview of SDN Terminology & Concepts. Cisco. [Online] 2013. [Date: 4th of October 2013.] <https://learningnetwork.cisco.com/docs/DOC-21946>.
- [10] Cisco Systems, Inc. onePK Chat and Demo at Cisco Live. SlideShare. [Online] 2012. [Date: 4th of October 2013.] <http://www.slideshare.net/getyourbuildon/onepk-chat-and-demo-at-cisco-live>.
- [11] S. Cadora. Hitchhiker's Guide to onePK. Cisco. [Online] 2013. [Date: 12th of September 2013.] <https://learningnetwork.cisco.com/docs/DOC-22910>.
- [12] R. Trunk. Understanding Performance Routing (PFR). Chesapeake Netcraftsmen. [Online] 2009. [Date: 15th of November 2013.] <http://netcraftsmen.net/archived-documents/c-mug-article-archive/7-20090922-cmug-understanding-performance-routing/file.html?limit=10>.
- [13] Kalita Hemanta Kumar, Nambiar Manoj K.: Designing WANem: A Wide Area Network Emulator tool. Bangalore, 2011. ISBN 9780769546186.
- [14] R. Pandi Selvam, V.Palanisamy: An efficient cluster based approach for multi-source multicast routing protocol in mobile ad hoc networks, International Journal of Computer Networks & Communications (IJCNC) Vol.3, No.1, January 2011, PP. 154-166. ISSN 0974-9322

Authors

Martin Kriška received his Ing. (equiv of M.Sc.) Degree from Technical University of Kosice in 2014. Since 2011 he was an active member of the Computer Networks Laboratory where he was working on several R&D activities connected with computer networks in general, QoS and security. He has spent one semester studying at the Polytechnic in Jyväskylä, Finland. He has been selected to join the Cisco incubator program under the umbrella of the TUKE ASC, where he gone through training in the field of advanced network technologies.



Jozef Janitor received his PhD. Degree from Technical University of Kosice in 2013. Since then, he has been working at the Institute of Computer Technology as a systems engineer, where he is working with the local team on R&D activities in the field of computer networks, IP telephony and cloud technologies.



Peter Fecil'ak received his PhD. Degree from Technical University of Kosice in 2011. Since then, he has been working at the Department of Computers and Informatics as a teacher and researcher in the field of computer networks, security and smartgrid technologies. He is leading the TUKE ASC as well as other educational activities in the region.

