Keeping the local traffic local – The NaissIX Internet Exchange Point Success Story

Dusan M. Vuckovic, Aleksandra S. Panajotovic, Vladimir M. Ciric, Marko Z. Zivkovic

***Abstract*—With the increased demand for additional bandwidth, when it comes to the Metropolitan Area Networks, there is also a growing need to keep the local traffic local. Preventing the traffic of local origin from "spilling" over the MAN borders, or maybe even country borders become one of the very important tasks. This is where local Internet Exchange Points can help. This paper illustrates possibilities local IXP implementation has to offer (city of Nis, Serbia in this case) and improvements that can come out of it.**

***Keywords*—Internet exchange point, IXP, BGP, traffic, peering.**

I. Introduction

Internet Exchange Points (IXPs) have emerged as critical components of the global Internet architecture, particularly in managing the Metropolitan Area Network (MAN) traffic. IXPs function as hubs where multiple Internet Service Providers (ISPs) interconnect their networks, enabling data routing between them directly rather than through a third-party network. This simplifies the exchange of traffic between networks, enhances efficiency, reduces latency, and lowers costs for ISPs and their customers. In essence, IXPs can be compared to a busy intersection in a city, where many routes converge, allowing traffic to flow in multiple directions [1].

The volume of data being transmitted across the Internet has been growing exponentially, propelled by the proliferation of Internet-enabled devices, increasing the usage of bandwidth-intensive applications like streaming video, remote work, e-commerce, cloud computing, and online gaming. This trend poses challenges to the efficiency and reliability of network traffic, particularly in metropolitan areas where the demand is exceptionally high. In this regard, IXPs play a pivotal role in managing this data traffic, ensuring the smooth functioning of the Internet in these areas.

To comprehend the role of IXPs in managing MAN traffic, it's important to understand the role of peering. Peering is a process where two or more networks connect and exchange traffic. This arrangement is typically formalized through a peering agreement [2], which outlines the terms and conditions for data exchange. By facilitating peering, IXPs enable networks to directly interconnect rather than routing traffic through a transit provider.

At an IXP, ISPs interconnect at a physical level through network switches, which are devices that channel incoming data from multiple input ports to the specific output port that will take the data toward its intended destination. The interconnection can occur through a public peering arrangement, where all participants connect to a shared switch, or a private peering arrangement, where two participants interconnect directly. The choice between public and private peering is generally based on the volume of traffic exchanged between the participants.

The localization of traffic exchange at IXPs has significant implications for the performance of MAN traffic. It lowers the latency by reducing the physical distance data needs to travel. Moreover, it decreases the load on upstream transit providers, freeing up capacity for other traffic. Furthermore, the direct interconnection of networks at IXPs reduces the dependency on single transit providers, thereby increasing the resilience of the Internet.

In addition to enhancing network efficiency, IXPs contribute to the digital economy in metropolitan areas. They stimulate competition among ISPs by providing a neutral location for traffic exchange, making it easier for smaller providers to interconnect with larger networks. This increased competition can lead to better services and lower prices for end users.

Furthermore, IXPs can foster innovation by facilitating the deployment of new services. For example, content delivery networks (CDNs) often host their servers at IXPs to deliver high-quality streaming services. By enabling direct peering with CDNs, IXPs improve the quality of streaming services and enable the deployment of latency-sensitive applications.

This paper presents the implementation of the IXP within the Nis city limits in Serbia, which can be treated as a Metropolitan Area Network. This implementation helped local providers to exchange content more efficiently and to reduce cost by saving bandwidth on the intercity lines.

II. The NaissIX Project

RIPE NCC's (fr. Réseaux Internet Protocol Européens – Network Coordination Centre) guidelines on how to "keep the local traffic local" [3] were the main motivator behind the decision to create this local IXP. The idea was to interconnect local internet providers and allow them to exchange traffic using the local exchange node. Up to that point all the traffic was, at least reaching Belgrade's SOX (Serbian Open Exchange- www.sox.rs) Internet Exchange Point. If not that, then the traffic would eventually reach Amsterdam or Frankfurt

Dusan M. Vuckovic, Aleksandra S. Panajotovic, Vladimir M. Ciric are with the Faculty of Electronic Engineering, University of Nis, Aleksandra Medvedeva 14, 18000 Nis, Serbia
(e-mail: dusan.vuckovic@elfak.ni.ac.rs, aleksandra.panajotovic@elfak.ni.ac.rs, vladimir.ciric@elfak.ni.ac.rs)
Marko Z. Zivkovic is with Coming Computer Engineering, Tose Jovanovica 7, Belgrade, Serbia (e-mail:marko.zivkovic@coming.rs)

nodes, at least it would from the Serbian Academic Network. This was congesting intercity and international lines with the traffic of local character. Optimisation had to be done.

Project of this magnitude required the help of the local Internet providers , like, in our case, NiNet (www.ninet.rs) or NetNet (www.netnet.rs) to connect their data centres with the Faculty of Electronic Engineering in Nis, where the NaissIX Internet Exchange Point was hosted.

The Internet Exchange itself was envisioned like the "open type" one, where both internet operators and companies can join. The whole concept was formalized back in 2016 and five local operators decided to join, initially.

The 10Gb backbone was created between the Faculty of Electronic Engineering (where the NaissIX Internet Exchange resides) on the North-West end, JUNIS – the main University datacentre on the West end, and industry partner NiNet, (the local internet provider) with its datacentre in the South-East part of the city.

The project of this magnitude required certain level of synchronisation between partners. The work was split to the "passive" part where Industry partners took the lead by laying down fibre optic cables between IXP nodes and the "active network" part, which was in the domain of Academic partners. Allied Telesys x510 L3 switches were selected as main L3 switching points mainly due to their 10Gb SFP+ support. Their four 10Gb ports allowed 10Gb backbone across the city to be created.

The scalability of the network was increased with the ability to connect potential new clients to the nearest datacentre (either JUNIS, Faculty of Electronic Engineering or NiNet datacentre) inside the Nis city centre. The traffic would then be routed using the BGP protocol. Figure 1 shows the physical connections of the backbone network between these 3 mentioned datacentres across the city of Nis.

III. THE ROUTING SOLUTION

In terms of the routing, the most common "routing reflector" solution was used. There is only one server (DELL R320) acting as a router with CentOS 7 operating system and Quagga [4] as a BGP routing software. Intention was to avoid memory constraints certain "conventional" routers could impose. BGP routing tables can become rather big [5] when internet providers start exchanging them.

The "route reflector" server was then connected with the 10Gb network card to the main L3 switch in order to do the routes announcement in the most efficient way. This is considered to be the efficient way to route traffic between nodes in Internet Exchage Points with low number of members, and it is the rather common implementation [6].

Two Virtual LANs (VLANs) were allocated for Internet exchange purpose. VLAN 464 for the IPv4 traffic and VLAN 646 for the IPv6 traffic.

Every IXP member was then allocated with needed public IPv4 and IPv6 addresses allocated to the IXP by RIPE and would then send their traffic to the same virtual LANs. Figure 2 shows the whole network topology and how the system was envisioned. Important was to form the 10Gb redundancy triangle between three main IXP points: NaissIX, JUNIS and NiNet. Once this was established, the process of adding new clients was just the matter of the extension of the BGP table, once the physical connections were in place.

IV. GOING LIVE

Prior to going live, the internet exchange point had to be registered with the RIPE NCC and assigned with the unique Autonomous System (AS) number. After the application process was over the IXP was registered as a subsystem of the Academic Network of Serbia and assigned with the AS202720 number (along with one $IPv4_{/24}$ and one $IPv6_{/48}$ subnets for this specific purpose).

Fig. 1. The map shows the topology of the network backbone across the city of Nis

These addresses serve as every client's identifier on the combined network. The actual routes are defined using BGP protocol.

The Quagga routing was setup in a way to advertise all of the adjunct networks in the BGP announcements and to establish the "neighboring" structure. Client's routers, on the other hand, are advertising routes for their respective Autonomous Systems. This way the traffic exchange is happening.

It is important to mention that this is a "peering" Internet Exchange type intended to exchange only the traffic directed towards the other party. So called "transit peering" that would exchange all of the traffic towards unknown destinations is not allowed at the moment, due to limited resources. In the foreseeable future it will be possible for local operators to use these lines as "transit" ones and forward all of the customer's traffic to it.

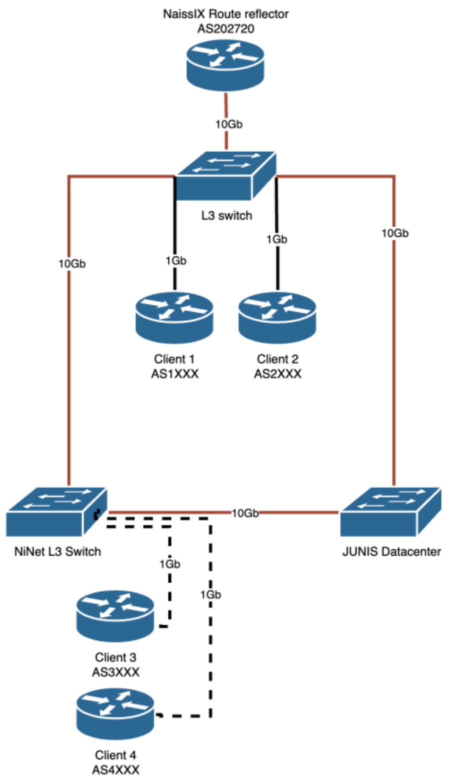

Fig. 2. Network topology of the Internet Exchange Point.

V. The Effects

In order to stay aligned with the primary goal of "keeping the local traffic local" the implementation of the NaissIX internet exchange point offered shorter pathways between the local Internet Providers. There is also obvious improvement in speed and bandwidth illustrated with the Fig. 3. Here we can see the traceroute from one of the academic network nodes to the webserver of NiNet internet provider.

Initially this route involved nodes in Belgrade, SOX Internet exchange and AMRES (Academic Network of Serbia). The total time needed for package to reach the destination between two places physically less than kilometer away was around 87ms. With the IXP in place this was reduced to only 6 hops and 7ms round time.

The improved network capabilities and bandwidth have proven to be very useful during the COVID-19 pandemic and lockdown that followed, when Internet as a whole experienced overload [7]. Our students, located around the City of Nis, connected to any of the member internet providers, were able to follow on-line classes, and multiple simultaneous live streams without any problems, exactly because of low latency and 10Gb network backbone.

VI. Conclusion

The importance of IXPs in managing MAN traffic cannot be overstated. As the volume and complexity of Internet traffic continue to grow, their role will only become more vital.

```
#Traceroute before the IXP implementation
>tracert www.ninet.rs

Tracing route to ninet.rs [212.200.45.19]
over a maximum of 30 hops:

  1    <1 ms    <1 ms    <1 ms  192.168.98.1
  2    <1 ms     1 ms    43 ms  160.99.35.11
  3     1 ms     1 ms    <1 ms  10.10.192.11
  4     1 ms     1 ms     1 ms  160.99.34.101
  5     2 ms     1 ms     1 ms  160.99.8.1
  6    <1 ms    <1 ms    <1 ms  160.99.1.12
  7     2 ms     1 ms    <1 ms  147.91.6.177
  8     4 ms     3 ms     4 ms  172.18.12.1
  9     3 ms     4 ms     3 ms  172.18.13.1
 10     3 ms     3 ms     3 ms  147.91.6.58
 11     3 ms     3 ms     3 ms  147.91.6.57
 12     4 ms     3 ms     3 ms  147.91.6.190
 13     3 ms     3 ms     3 ms  193.105.163.38
 14    15 ms     8 ms    10 ms  212.200.29.42
 15    14 ms     *        9 ms  212.200.45.19


#Traceroute after the IXP went online
Tracing route to ninet.rs [212.200.45.19]
over a maximum of 30 hops:

  1    <1 ms     2 ms    <1 ms  160.99.35.11
  2    <1 ms    <1 ms    <1 ms  10.10.192.11
  3     2 ms     2 ms     2 ms  160.99.34.101
  4     1 ms     1 ms    <1 ms  185.96.208.1
  5     1 ms     1 ms     1 ms  10.18.0.14
  6     1 ms    <1 ms     1 ms  212.200.45.19
```

Fig. 3. Traceroute showing differences in number of hops before the IXP implementation and after

However, realizing their full potential requires a supportive policy environment that encourages the development and operation of IXPs. This includes, among other things, regulations that foster competition, facilitate the sharing of infrastructure, and promote transparency in the operation of IXPs. The type of traffic is also changing over time and implementations such this one can significantly reduce cost [8].

Project like this, on a MAN scale showed its true potential in times of hardship, like COVID-19 pandemic was. Suddenly, everyone had to move to a new on-line system, and regular infrastructure was simply not prepared for that. Interconnection, like the one NaissIX provided, allowed students to follow classes without interruption, exchange big files, stay in touch with their peers and teaching staff. Consequently, academic infrastructure received certain level of redundancy, which was not present before. Instead of backup

ADSL lines, there was now an option of just rerouting the traffic towards any of the NaissIX member service providers, and creating "transit peering" connection, in order to keep the systems running. Faculty of Electronic Engineering managed to run its operations for 52 days of lockdown without single access to the physical equipment. This illustrates the best level of redundancy IXPs have to offer.

In conclusion, IXPs play an indispensable role in managing MAN traffic. They improve the efficiency, reliability, and resilience of the Internet while stimulating competition and innovation in the digital economy. As the Internet continues to evolve and grow, the role of IXPs in shaping its future will be increasingly central.

Future work will be focused mostly in enabling the transit peering for all IXP members and introducing measurement systems like RIPE Atlas software probes [9]. This will help quantify IXPs efficiency and further improve traffic exchange.

Acknowledgment

The authors would like to acknowledge the help of NiNet Internet Service provider, University Network staff and Allied Telesyn Serbia for their contributions and volunteer work on this project.

References

[1] G. Nomikos, P. Sermpezis, and X. Dimitropoulos, "Re-mapping the internet: Bring the IXPs into play: www.inspire.edu.gr/ixp-map," in *2017 IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, May 2017, pp. 910–915. doi: 10.1109/INFCOMW.2017.8116497.

[2] M. Z. Ahmad and R. Guha, "Internet exchange points and Internet routing," in *2011 19th IEEE International Conference on Network Protocols*, Oct. 2011, pp. 292–294. doi: 10.1109/ICNP.2011.6089065.

[3] Emile Aben, "IXP Country Jedi," presented at the RIPE NCC::Educa, on-line, Oct. 2017. Accessed: Sep. 06, 2023. [Online]. Available: https://www.ripe.net/support/training/ripe-ncc-educa/presentations/ixp-country-jedi.pdf

[4] P. Jakma and D. Lamparter, "Introduction to the quagga routing suite," *IEEE Netw.*, vol. 28, no. 2, p. 42, 2014.

[5] Geoff Huston, *ISP Survival Guide: Strategies for Running a Competetive ISP*, illustrated ed., vol. 4. Wiley, 1999. [Online]. Available: https://books.google.rs/books/about/ISP_Survival_Guide.html?id=aQypjlwxtI0C&redir_esc=y

[6] D. Ó Briain, D. Denieffe, Y. Kavanagh, and D. Okello, "Rebuilding the Internet Exchange Point in Uganda," in *2017 28th Irish Signals and Systems Conference (ISSC)*, Jun. 2017, pp. 1–6. doi: 10.1109/ISSC.2017.7983601.

[7] A. Feldmann *et al.*, "Implications of the COVID-19 Pandemic on the Internet Traffic," in *Broadband Coverage in Germany; 15th ITG-Symposium*, Mar. 2021, pp. 1–5.

[8] M. Soobron, C. Soobron, S. Soobron, A. Sukhoo, and R. H. Hawabhay, "Connectivity within Indian Ocean Islands (Mauritius, Seychelles, Comoros, Reunion and Madagascar) — A case for a Regional Internet Exchange," in *2014 IST-Africa Conference Proceedings*, May 2014, pp. 1–18. doi: 10.1109/ISTAFRICA.2014.6880667.

[9] A. Davies and P. Homburg, "RIPE Atlas Software Probes," Sep. 10, 2023. https://labs.ripe.net/author/alun_davies/ripe-atlas-software-probes/