



- Goran Slavić
e-mail: gslavic@sox.rs

Blueprint for building IXP

M.Sc.E.E. Goran Slavić
gslavic@sox.rs

What is needed for the IXP ?

- Hardware
- Software
- Staff
- Customers
- Government support
- Network monitoring & Safety
- CDNs
- Services
- Carrier relations
- Help from other IXPs - at the beginning

Hardware

1. L2 Switch

- 48ports x 10Gbps preferred
- \$ 10.000+(unless it is a bare metal / whiteswitch)



2. Route Server hardware

- It does not have to be as powerful as one might think (SOX in the first year = 2 virtual machines)
- \$ 3 000 (Dual Xeon + brand name server)

3. Space for hosting the hardware

- 2-4 RU for start (later, much, much more ...)

Software

Big (ongoing) debate over the Route Server software:

1. Quagga

- Old solution for route servers.
- Old style CLI (similar to configuring of a CISCO router).
- Issues with scalability and stability.

2. Bird

- New solution for route servers.
- CLI only for monitoring. Configuring by editing conf. file
- “Programming” RS rather than “configuring” RS.

Software

3. IXP manager (<https://github.com/inex/IXP-Manager>)

- Allows automated management of customers prefixes, integrated SFLOW statistics, web interface.
- Good software to install when starting the IXP. Migration to it at the later stages of expansion is very hard.

IXP Manager

- portal and RS manager

IXP Manager - Member Dashboard

SIX :: IXP Manager

Home Super User Admin Member Information Peering Documentation Statistics Staff Links Profile Help About [Logout]

SIX Network: Filter

Show 100 entries Search:

Member	Switch	Port	Speed	ASN	Route Server	IPv4 Address	IPv6 Address	Status
Amis	ssix1	Te1/4	1000Mbps	8591	Yes	91.220.194.107	2001:7fb:46:8591	Connected
ARQ	ssix1	Te1/26	1000Mbps	44647	Yes	91.220.194.123	2001:7fb:46:44647	Connected
Ames	ssix1	Te1/7	1000Mbps	2107	Yes	91.220.194.102	2001:7fb:46:2107	Connected
Ames	ssix2	Te1/8	1000Mbps	2107	Yes	91.220.194.2	2001:7fb:46:0:1:2107	Connected
Domenca	ssix2	Te1/7	1000Mbps	43128	Yes	91.220.194.27	2001:7fb:46:0:1:0:4:3128	Connected
Lib	ssix1	Te1/16	1000Mbps	39912	Yes	91.220.194.116	2001:7fb:46:3:9912	Connected
ILQI	ssix2	Te1/18	1000Mbps	198524	Yes	91.220.194.24	2001:7fb:46:0:1:0:19:8524	Connected
IT TEL	ssix1	Te1/17	1000Mbps	39785	Yes	91.220.194.115		Connected
KRS Networks	ssix1	Te1/12	1000Mbps	58056	Yes	91.220.194.120		Connected
Lutjanski Kabel, Telemach	ssix1	Te1/23	1000Mbps	33929	No	91.220.194.113		Connected
Mega M	ssix1	Te1/11	1000Mbps	44549	Yes	91.220.194.121		Connected
Mega M	ssix2	Te1/21	1000Mbps	44549	Yes	91.220.194.21		Connected
Metronet	ssix2	Te1/13	1000Mbps	41427	Yes	91.220.194.16	2001:7fb:46:0:1:0:4:1427	Connected
NETSI.NET	ssix1	Te1/8	1000Mbps	12778	Yes	91.220.194.110	2001:7fb:46:1:2778	Connected

4. Custom software

- Extremely important to develop solutions that are automated and not “human dependent” (some problems need to be addressed quicker than a person can react).
- In-house solutions (some things need to be adapted to actual conditions of the given IXP).
- In-house solutions should be developed only for specific needs and only after the detailed search of the existing solutions.
- IXP manager has (over time) integrated several custom solutions in to it’s program code.

“How many employees does IXP need ?”

1-3 Financially viable (even at the start)

- Inability to maintain 24/7 NOC
- IXP is more equipment dependent than people dependent.

5+ Financially viable only if the IXP has grown to provide more services than a mere IXP

- Ability to maintain quality in customer support and quick reaction time to emergencies (24/7 NOC)

Customers

- Most important and the most problematic asset.
- Large ISP-s don't see the reason to join IXP => need to be persuaded to join IXP.
- Large ISP-s expect to be paid for their service – not to pay for a service.
- Small ISP-s are eager to join IXP but sometimes don't have funds to pay for capacity to connect to an IXP.
- The balanced pricing to customers is the key to success of the IXP

Customer relations

- IXP - ISP relations can be very problematic (who-needs who-more, who is giving who traffic)
- IXP – hosting providers relations can be highly problematic (“we are giving you traffic to sell to other ISP-s”)
- ISP - ISP relations are by definition problematic (more than one horror story of ISPs exchanging traffic over AMS-IX – while located in the same building).

Government support



- Extremely important to IXP when it is the first IXP in that country/region.
- “Keeping local traffic – local” is not just important to the telecom providers but to the government.
- Local IXP solves a lot of issues of important / confidential traffic between firms / government agencies “running around the world” before reaching it’s destination.
- Local government should strongly suggest participation in the local IXP.

CDNs

- GGC, Akamai etc. are important generator of traffic for the IXP.
- Conditions for acquiring one of the major CDN-s (large number of customers - for example) can be hard to satisfy by a new IXP.
- Small CDN node can be underused – if there is a larger one in the network of the large ISP.
- Large CDN node can cause problem for the IXP because it is the preferred choice for the CDN algorithm – now the CDN-s of IXP customers are underused.

Network monitoring

- Proper operation of the IXP is highly dependent on good monitoring of the network status.
- SNMP collectors and graph tools as the valuable asset in prevention and detection of potential problems.




Network monitoring

- Programs like Cacti can send e-mail / SMS alarms to network administrators in case of problems or even warnings before problem manifests itself fully.

Actions	Name	ID	Type	Trigger	Duration	Repeat	Warn H/L	Alert H/L	BL H/L	Current	Triggered**	Enabled
	(64-bit) SRV RTR PALATA - Traffic - team0.19 [traffic_in_hc]	58	High/Low	5 Minutes	N/A	Every Hour	-/128	-/64	N/A	0.6957	yes	Enabled
	Palata UPS 1 - Temperature [snmp_oid]	7	High/Low	5 Minutes	N/A	Every Hour	29/18	30/16	N/A	20	no	Enabled
	Gandjeva SMC - Combrige Global (3G) (OUT)	4	High/Low	5 Minutes	N/A	Every 5 Minutes	125000000/1000	130000000/100	N/A	1,074,763.3333	no	Enabled
	(64-bit) SRV RTR PALATA - Traffic - team0.304 [traffic_in_hc]	44	High/Low	5 Minutes	N/A	Every Hour	-/1024	-/512	N/A	524,557,399.1806	no	Enabled
	Gandjeva SMC - Combrige Global (3G) (IN)	25	High/Low	5 Minutes	N/A	Every 5 Minutes	125000000/10000	130000000/1000	N/A	657,417.0733	no	Enabled
	RTR PALATA - Traffic - team0.681 B-IX-regional IN	38	High/Low	5 Minutes	N/A	Every Hour	-/1024	-/256	N/A	183,665,021.6087	no	Enabled
	Palata UPS 2 - Temperature [snmp_oid]	8	High/Low	5 Minutes	N/A	Every Hour	30/18	31/16	N/A	22	no	Enabled
	(64-bit) Sofija - Telepoint - Ethernet1/0/1 OUT NOVATEL	96	High/Low	5 Minutes	N/A	Every Hour	-/256	-/128	N/A	0	no	Disabled
	RTR PALATA - Traffic - team0.681 B-IX-regional OUT	39	High/Low	5 Minutes	N/A	Every Hour	-/1024	-/256	N/A	149,196,301.2843	no	Enabled
	Palata UPS 2 - Battery Percent Charge [charge]	19	High/Low	5 Minutes	N/A	Every Hour	-/90	-/90	N/A	100	no	Enabled
	SRV SOX-3 - Ukupan broj BGP ruta [snmp_oid]	17	High/Low	5 Minutes	N/A	Every Hour	-/180000	-/150000	N/A	359723	no	Enabled
	SRV SOX-1 - Broj aktivnih BGP sesija [snmp_oid]	14	High/Low	5 Minutes	N/A	Every Hour	-/36	-/35	N/A	-	no	Enabled
	SRV SOX-3 - Broj aktivnih BGP sesija [snmp_oid]	16	High/Low	5 Minutes	N/A	Every Hour	-/34	-/33	N/A	42	no	Enabled
	SRV SOX-1 - Ukupan broj BGP ruta [snmp_oid]	15	High/Low	5 Minutes	N/A	Every Hour	-/45000	-/45000	N/A	-	no	Enabled
	Palata UPS 1 - Battery Percent Charge [charge]	18	High/Low	5 Minutes	N/A	Every Hour	-/90	-/90	N/A	100	no	Enabled
	Palata x650 - Neterra (Google) [IN]	20	High/Low	5 Minutes	N/A	Every Hour	-/100000	-/100000	N/A	678,728,486.6333	no	Enabled
	Palata x650 - Neterra (Google) [OUT]	21	High/Low	5 Minutes	N/A	Every Hour	-/100000	-/100000	N/A	301,047,399.86	no	Enabled
	(64-bit) Sofija - Telepoint - Ethernet1/0/1 IN NOVATEL	95	High/Low	5 Minutes	N/A	Every Hour	-/256	-/128	N/A	0	no	Disabled
	RTR PALATA - Traffic - team0.44 OMN-IX IN	40	High/Low	5 Minutes	N/A	Every Hour	-/1024	-/512	N/A	31,815,198.9498	no	Disabled
	Palata x650 - 1:23 SOFJA (IN)	31	High/Low	5 Minutes	N/A	Every Hour	-/10240	-/1024	N/A	625,273,303.75	no	Enabled
	RTR PALATA - Traffic - team0.4 [traffic_out_hc]	43	High/Low	5 Minutes	N/A	Every Hour	-/1024	-/512	N/A	81,676,686.5418	no	Enabled
	RTR PALATA - Traffic - team0.4 [traffic_in_hc]	42	High/Low	5 Minutes	N/A	Every Hour	-/1024	-/512	N/A	21,048,924.0702	no	Enabled
	Palata x650 - 1:23 SOFJA (OUT)	30	High/Low	5 Minutes	N/A	Every Hour	-/1024	-/512	N/A	375300215	no	Enabled
	Palata DCN - Beotel Korisnik	27	High/Low	5 Minutes	N/A	Every 15 Minutes	-/1024	-/512	N/A	323,542,1037	no	Enabled
	Palata x650 - Google Telnor (IN)	28	High/Low	5 Minutes	N/A	Every Hour	-/1024	-/512	N/A	480,389,585.6533	no	Disabled
	Palata x650 - Google Telnor (OUT)	29	High/Low	5 Minutes	N/A	Every Hour	-/1024	-/512	N/A	51,521,684.1567	no	Disabled
	RTR PALATA - Traffic - team0.44 OMN-IX OUT	41	High/Low	5 Minutes	N/A	Every Hour	-/1024	-/512	N/A	6,381,151.0134	no	Disabled
	Palata Cisco4948 - 5 Minute CPU	32	High/Low	5 Minutes	N/A	Every Hour	30/-	40/-	N/A	9	no	Enabled
	THQ CISCO - 5 Minute CPU	33	High/Low	5 Minutes	N/A	Every Hour	30/-	40/-	N/A	10	no	Enabled
	(64-bit) Sofija - Netix - Ethernet1/0/27 Sofia - Bac (IN)	34	High/Low	5 Minutes	N/A	Every Hour	-/128	-/64	N/A	3,450,654.7119	no	Enabled


Network monitoring

- SFLOW collector as “eyes and ears” of the network monitoring.
- Good SFLOW analyses is the most important information in directing the further IXP expansions and activation of new connections.



Network monitoring

- Quick SFLOW analyses can point not only to the link but also the origin AS of the problematic traffic (DDoS attack for example).



Safety

- Large capacity network of the IXP is the magnet for multitude attempts for the DDoS and similar type of attacks.
- Prevention and mitigation of the DDoS should be initiated both on the client side and by the features of the route server of the IXP (BGP community activated black hole routing of the DDoS).
- Non malicious mistakes in configuration and poor understanding of BGP can be just as dangerous as purposeful and malicious attacks (“whole BGP table announcement”).

Services

Primary:

- IXP is primary the BGP peering service.
- Exchange of traffic between customers.

“Extended services”

- Private VLAN / Secure VLAN between customers.
- DNS, NTP, Distribution of video streaming / TV channels.

Multiservice exchange

- Long term goal: Migration of IXP to *multiservice exchange* with ability to exchange of TV/Audio/Video streams, telephone traffic and private communications between customers.

Carrier relations

- Important thing when IXP starts to expand/grow.
- Good international links and connection to other IXP-s can generate large amount of traffic and high quality – low latency links to major telecom players.
- Bad choice of connection points or poorly chosen / expensively paid links with small traffic can financially break the IXP.
- Back-to-back / bilateral agreements can be a good way to interconnect with regional carriers.
- Most important thing – strong international link is needed for the IXP (mainly for cash server servicing).

Expansion of the IXP

- Growth of the IXP is primarily customer driven.
- There are different types of IXP-s:
 - 1. There are IXPs with limited scope of services / number of customers.**
 - Good example would be the Toulouse-IXP that is connecting technological firms of Toulouse (FRA).
 - They are only connecting fixed number of geographically located ASs
 - They don't grow – because they don't need to grow.

Expansion of the IXP

2. Some IXP-s remain tied to a single (or couple) POP.

- They expect customers to “come to them”.
- Low costs of maintenance
- Lot of financial pressure on ISPs.
- Lingering problem of dependence on the single DC services.
- Limited number of capacity providers to a single DC.
- Growth is limited by the ability of its customers to connect to the DC.

Expansion of the IXP

3. Some IXP-s remain local places of Internet exchange.

- Good example is VIX
- They are located in major DCs of the “targeted” geographical location (Vienna).
- They stay away from opening POP-s in other countries/regions/cities.
- They do expect customers to “come to them” but are at the same time present in all of the relevant DCs of the region
- Growth is limited by the number of ISP-s of the region.²²

Expansion of the IXP

4. Some IXP-s “go to the customers”

- Good example is SOX (at its beginning)
- SOX had the POP in every major DC in Belgrade !
- Easier for the customers.
- Problem with scaling of the backbone network of the IXP.
- Problem with maintenance and leasing of the capacities that are connecting the POP-s (redundancy routes, MSTP, scaling of links etc.)

Expansion of the IXP

5. Some IXP-s “go abroad”

- They have presence in all of the DCs of the targeted region (for SOX - Belgrade).
- They have international presence / POPs in major DC of the wider geographical region (SEE for example)
- They have links to major IXPs (in case of SOX – NetIX, VIX, OMNIX, B-IX, AMS-IX ...)
- Business model that is very risky and needs to be planed and implemented very carefully

Expansion of the IXP

6. Some IXP-s rival the large capacity / DC / Hosting providers

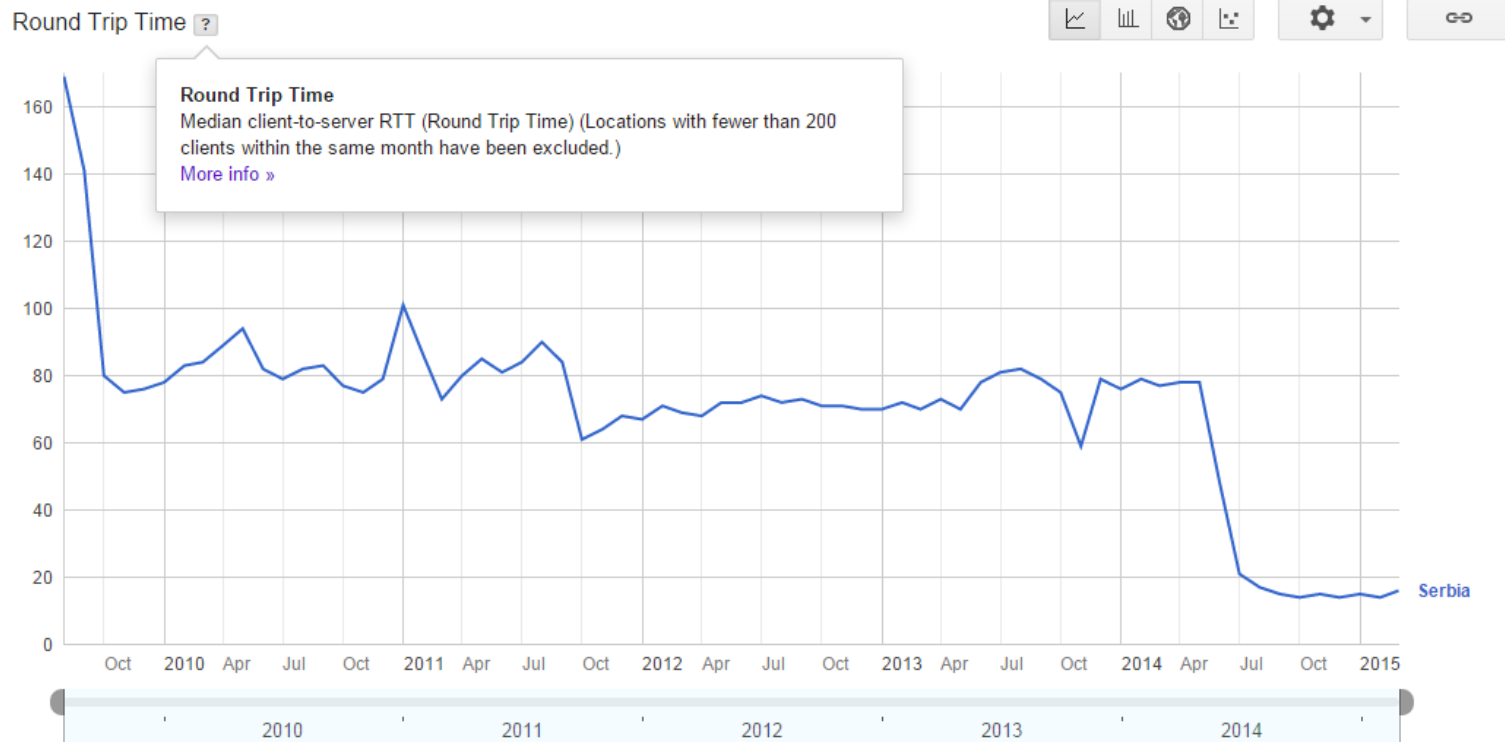
- Good example is NetIX / Nettera.
- They have international presence / POPs / in major DC of the wider geographical region (SEE for example)
- They own high capacity links to major IXPs (VIX, OMNIX, B-IX, AMS-IX ...)
- Business model that is very risky and needs to be planed and implemented very carefully

Expansion of the IXP

- Most important thing for the growth is the amount of traffic that IXP can offer to it's customers.
 - The CDNs can give IXP a major “boost in traffic”.
 - Gaming companies (EA, Riot Games, Capcom ...) are also the large traffic generators.
 - Connection to other IXP-s can bring a very positive results both IXP and it's members (more on that – on the round table).

Expansion of the IXP

- Having DNS root servers hosted in the IXP network can significantly accelerates Internet communication.



Expansion of the IXP

- **Communication...** with partners and customers in order to understand their needs and anticipate future steps (SDN, CDN, NTP, Security)
- **Collaboration...** with Internet giants, Internet associations, technology innovators, government entities
- **Cooperation...** with similar companies, learning from their experience

- No man or a company is an island! 3C is for success

Expansion of the IXP

- Important for the growth is the cooperation and help of international organizations (RIPE, RIPE NCC, ISOC)
- EuroIX the organization of IXP-s of Europe
- Most of the things that are needed for building, maintaining and expanding IXP - is specific for the individual IXP. Consequence ?

There is no universal solution !

Communication with SOX

noc@sox.rs

Zoran Perović
CEO

iConsult@perke.NET

Goran Slavić
Chief Engineer
gslavic@sox.rs

Nenad Krajnović
CTO
krajko@sox.rs