




The OpenBGPD Experience at LINX











Mo Shivji
Senior Network Engineer

May 2024
LINX122



The OpenBGPD experience at LINX

-  **History** 
-  **Re-engagement** 
-  **Deployment** 
-  **What Next?** 





History – Early Days

- 2002 – Beta route server on Foundry/Brocade (now LON1) LAN
 - Old RIPE K-Root server running Red Hat Linux and Vanilla Quagga
 - Started with 8 members peering
- 2003 – 4 Quagga based Route-Servers - 2 on each LAN
 - RS1 and RS3 on LON1
 - RS2 and RS4 on LON2
 - Servers now Gentoo Linux
 - Running Quagga patch for route server and multi-RIB functionality





History – Scaling Pains

- 2009 – All Quagga route servers are suffering from instability
 - Other major IXPs see same issues
 - NIC.CZ release BIRD
 - Euro-IX funding starts for development of OpenBGPD as a route-server
 - Funding and improvements also begin on multi-threaded Quagga fork
 - RSWG formed in Euro-IX to test BIRD, OpenBGPD and Quagga fork release
- 2010 – Deployment of BIRD and Euro-IX Quagga fork into production
 - LINX decides to opt out of using OpenBGPD
 - Quagga using a main thread with additional threads for CLI & 'BGP speaker'





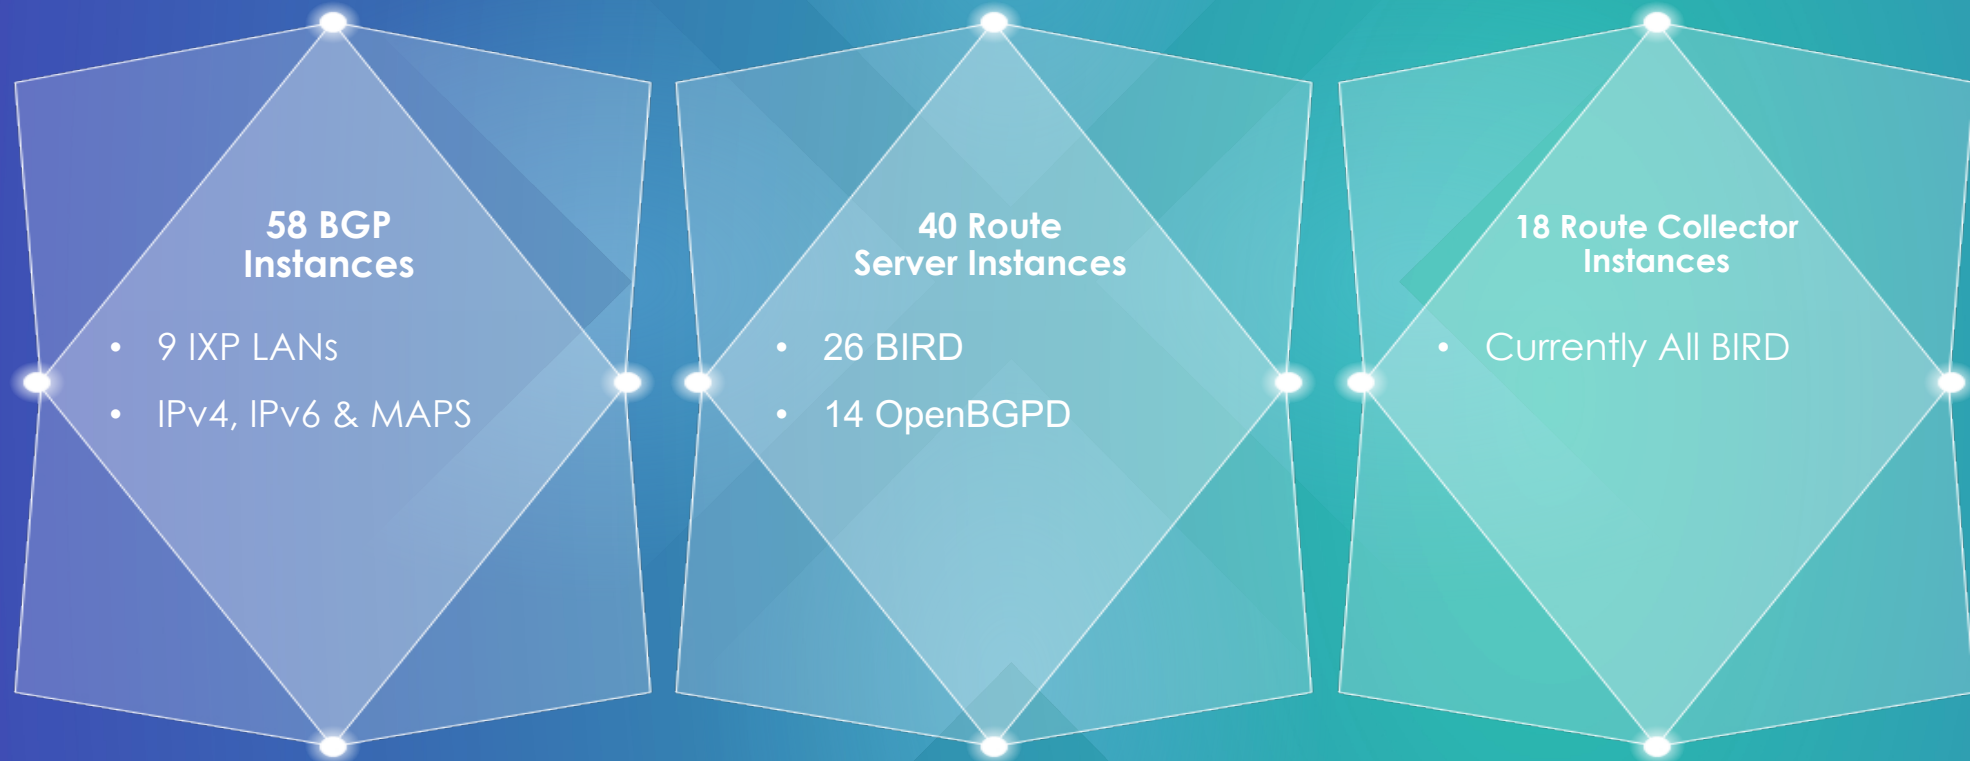
History – Adding Features

- 2015 – Added IRRDB Filtering to all route servers
 - Config built using in-house tooling
- 2020 – Added RPKI Support to all route servers
 - Quagga had no RPKI support
 - Replaced Quagga route servers with BIRD
 - Servers changed to Ubuntu Linux
 - OpenBGPD no longer used by major IXPs
- 2022 – LINX begins testing OpenBGPD from RSSF
- 2023 – LINX deploys OpenBGPD from RSSF
 - Also used in MAPS (Microsoft Azure Peering Service)
- 2024 – Remote Triggered Black Hole feature launched on OpenBGPD route servers






LINX Route Server and Collector Instances – May 2024






Re-engagement - LINX & RSSF



RSSF

Route Server Support
Foundation



Founding Sponsors

AMS-IX, DE-CIX, LINX &
NETNOD




Technical Advisory Board

Mo Shivji & Tim Preston

Meet monthly to discuss bugs,
problems and deployments
with OpenBGPD.

Discuss any new features and
possible feature requests



Advisory Board

Richard, Ruben, Matte and
Thomas/Marcos

Meet quarterly, review
of roadmap, funding,
performance and
direction






Deployment

- Initial Lab testing was successful
 - Route server functionality and features working in version 7.2
 - Could not successfully simulate LINX LON1 route servers ~ 700 peers
 - Simulation of route servers for regional LANs was successful
- Deployment of OpenBGPD into production at LINX Wales in March 2023
 - Not a busy IXP
 - 20 members
 - Allowed us to concentrate on deploying the daemons and supporting services rather than performance scaling





Deployment - Challenges

- Training the NOC
 - Different than BIRD, using 'bgpctl' from the Linux command line, not a dedicated CLI

```
moyaze@rs2-la1-lab1:~$ sudo bgpctl -s /var/run/bgpd.socket show summary
Neighbor          AS      MsgRcvd   MsgSent  OutQ Up/Down  State/PrfRcvd
195.66.224.10     22410   0         0        0  01:56:35 Active
moyaze@rs2-la1-lab1:~$
```

- NOC struggled with the correct arguments, particularly specifying the correct socket
 - Created a helper script

```
moyaze@rs2-la1-lab1:~$ obgpd show summary
Neighbor          AS      MsgRcvd   MsgSent  OutQ Up/Down  State/PrfRcvd
195.66.224.10     22410   0         0        0  02:00:43 Active
moyaze@rs2-la1-lab1:~$
```

- Running dual 'bgpd' processes for IPv4 and IPv6
 - Some Ansible work required





Deployment - Tooling

- Looking-Glass
 - Alice Looking-Glass (Alice-LG) used for route servers and route collectors
 - Initial support was for BIRD with birdwatcher providing an API
 - openbgpd-state-server released to provide an API for OpenBGPD
 - Noticeable differences in Alice-LG neighbors list
- Statistics
 - Re-write of old Perl snmp helper scripts to Python - who uses Perl nowadays?
 - Still using RRD

BGP SESSIONS ESTABLISHED

| Neighbour | ASN | State | Uptime | Description |
|---------------|--------|-------|----------|------------------------------|
| 195.66.236.3 | 201838 | up | 2 months | Community Fibre |
| 195.66.236.4 | 42611 | up | 2 months | Zzoomm PLC |
| 195.66.236.6 | 211393 | up | 2 months | Vitrifi Limited |
| 195.66.236.7 | 202596 | up | 2 months | G.Network Communications |
| 195.66.236.8 | 59811 | up | 2 months | Voneus Ltd |
| 195.66.236.9 | 51553 | up | 2 months | Utility Warehouse Ltd |
| 195.66.236.12 | 203372 | up | 2 months | GSA CAPITAL SERVICES LIMITED |


BIRD

BGP SESSIONS ESTABLISHED

| Neighbour | ASN | State | Uptime | Description |
|---------------|--------|-------------|----------|-----------------------------|
| 195.66.236.3 | 201838 | established | 3 months | PEER AS201838 195.66.236.3 |
| 195.66.236.4 | 42611 | established | 4 months | PEER AS42611 195.66.236.4 |
| 195.66.236.6 | 211393 | established | 2 months | PEER AS211393 195.66.236.6 |
| 195.66.236.7 | 202596 | established | 2 months | PEER AS202596 195.66.236.7 |
| 195.66.236.8 | 59811 | established | 4 months | PEER AS59811 195.66.236.8 |
| 195.66.236.9 | 51553 | established | 4 months | PEER AS51553 195.66.236.9 |
| 195.66.236.12 | 203372 | established | 4 months | PEER AS203372 195.66.236.12 |

OpenBGPD





Deployment - Problems

- Found bugs early after deployment where OpenBGPD would send a bad community attribute when non-transitive extended-communities are present
 - Caused Alice-LG to crash
 - Additional bug with extended-community filtering meant we couldn't strip these
 - Had to ask peers to modify what they were advertising
 - Playing whack-a-mole
 - Both issues now resolved
- Addition issue found in not being able to change the BGP next-hop
 - Not able to provide Remote Triggered Black Hole service with OpenBGPD
 - Now resolved
- Some issues with our Ansible playbooks discovered
 - Deploying to a new server different than migrating an existing one from BIRD





What next? The future of OpenBGPD at LINX

Immediate goal is to launch OpenBGPD onto the LINX LON1 LAN

Replace some BIRD Collectors with OpenBGPD

Lab testing to reach 1001+ Peers

Continue with RSSF to ensure roadmap is met and development continues

Sharing of test bed for the wider community to help adoption

Promotion of service contract to make the project sustainable



A white geometric diagram consisting of a pentagon with an internal star shape. The vertices of the pentagon are marked with white circles. The word 'Questions?' is centered within the diagram in a white sans-serif font.

Questions?