



Peering with AWS – 2022 Updates




LINX117 - London

Fredrik Korsbäck – Senior Infrastructure BD IP & Interconnect. Aka “BGP guy”

2022-11-23

Amazon at a glance

- Amazon is an American technology company
- Amazon.com; online shopping
- Amazon Web Services; cloud computing
- Prime Video; video streaming **and** content.
- Amazon Music
- Fire Tablets; Fire TV
- Echo and Alexa
- Kindle E-readers
- ...and much more



What does the AS16509 Network Serve?

AWS Cloud regions. The AWS Cloud spans **87** Availability Zones within **27** geographic regions around the world, with announced plans for **21 more** Availability Zones and 7 more AWS Regions in Australia, Canada, India, Israel, New Zealand, Spain and Switzerland

Home to EC2, S3 and DynamoDB but ~~the service catalogue now has over 200 fully featured services for a multitude of IT Workloads~~ **Just Launched!**




IoT



AI/ML



Game Tech



End-User
Compute



Developer Tools



Databases




Containers



Media Services



Quantum
Technologies



Compute



Business
Applications




Blockchain



App integration



Networking &
Content Delivery




Storage



Security, Identity
& Compliance



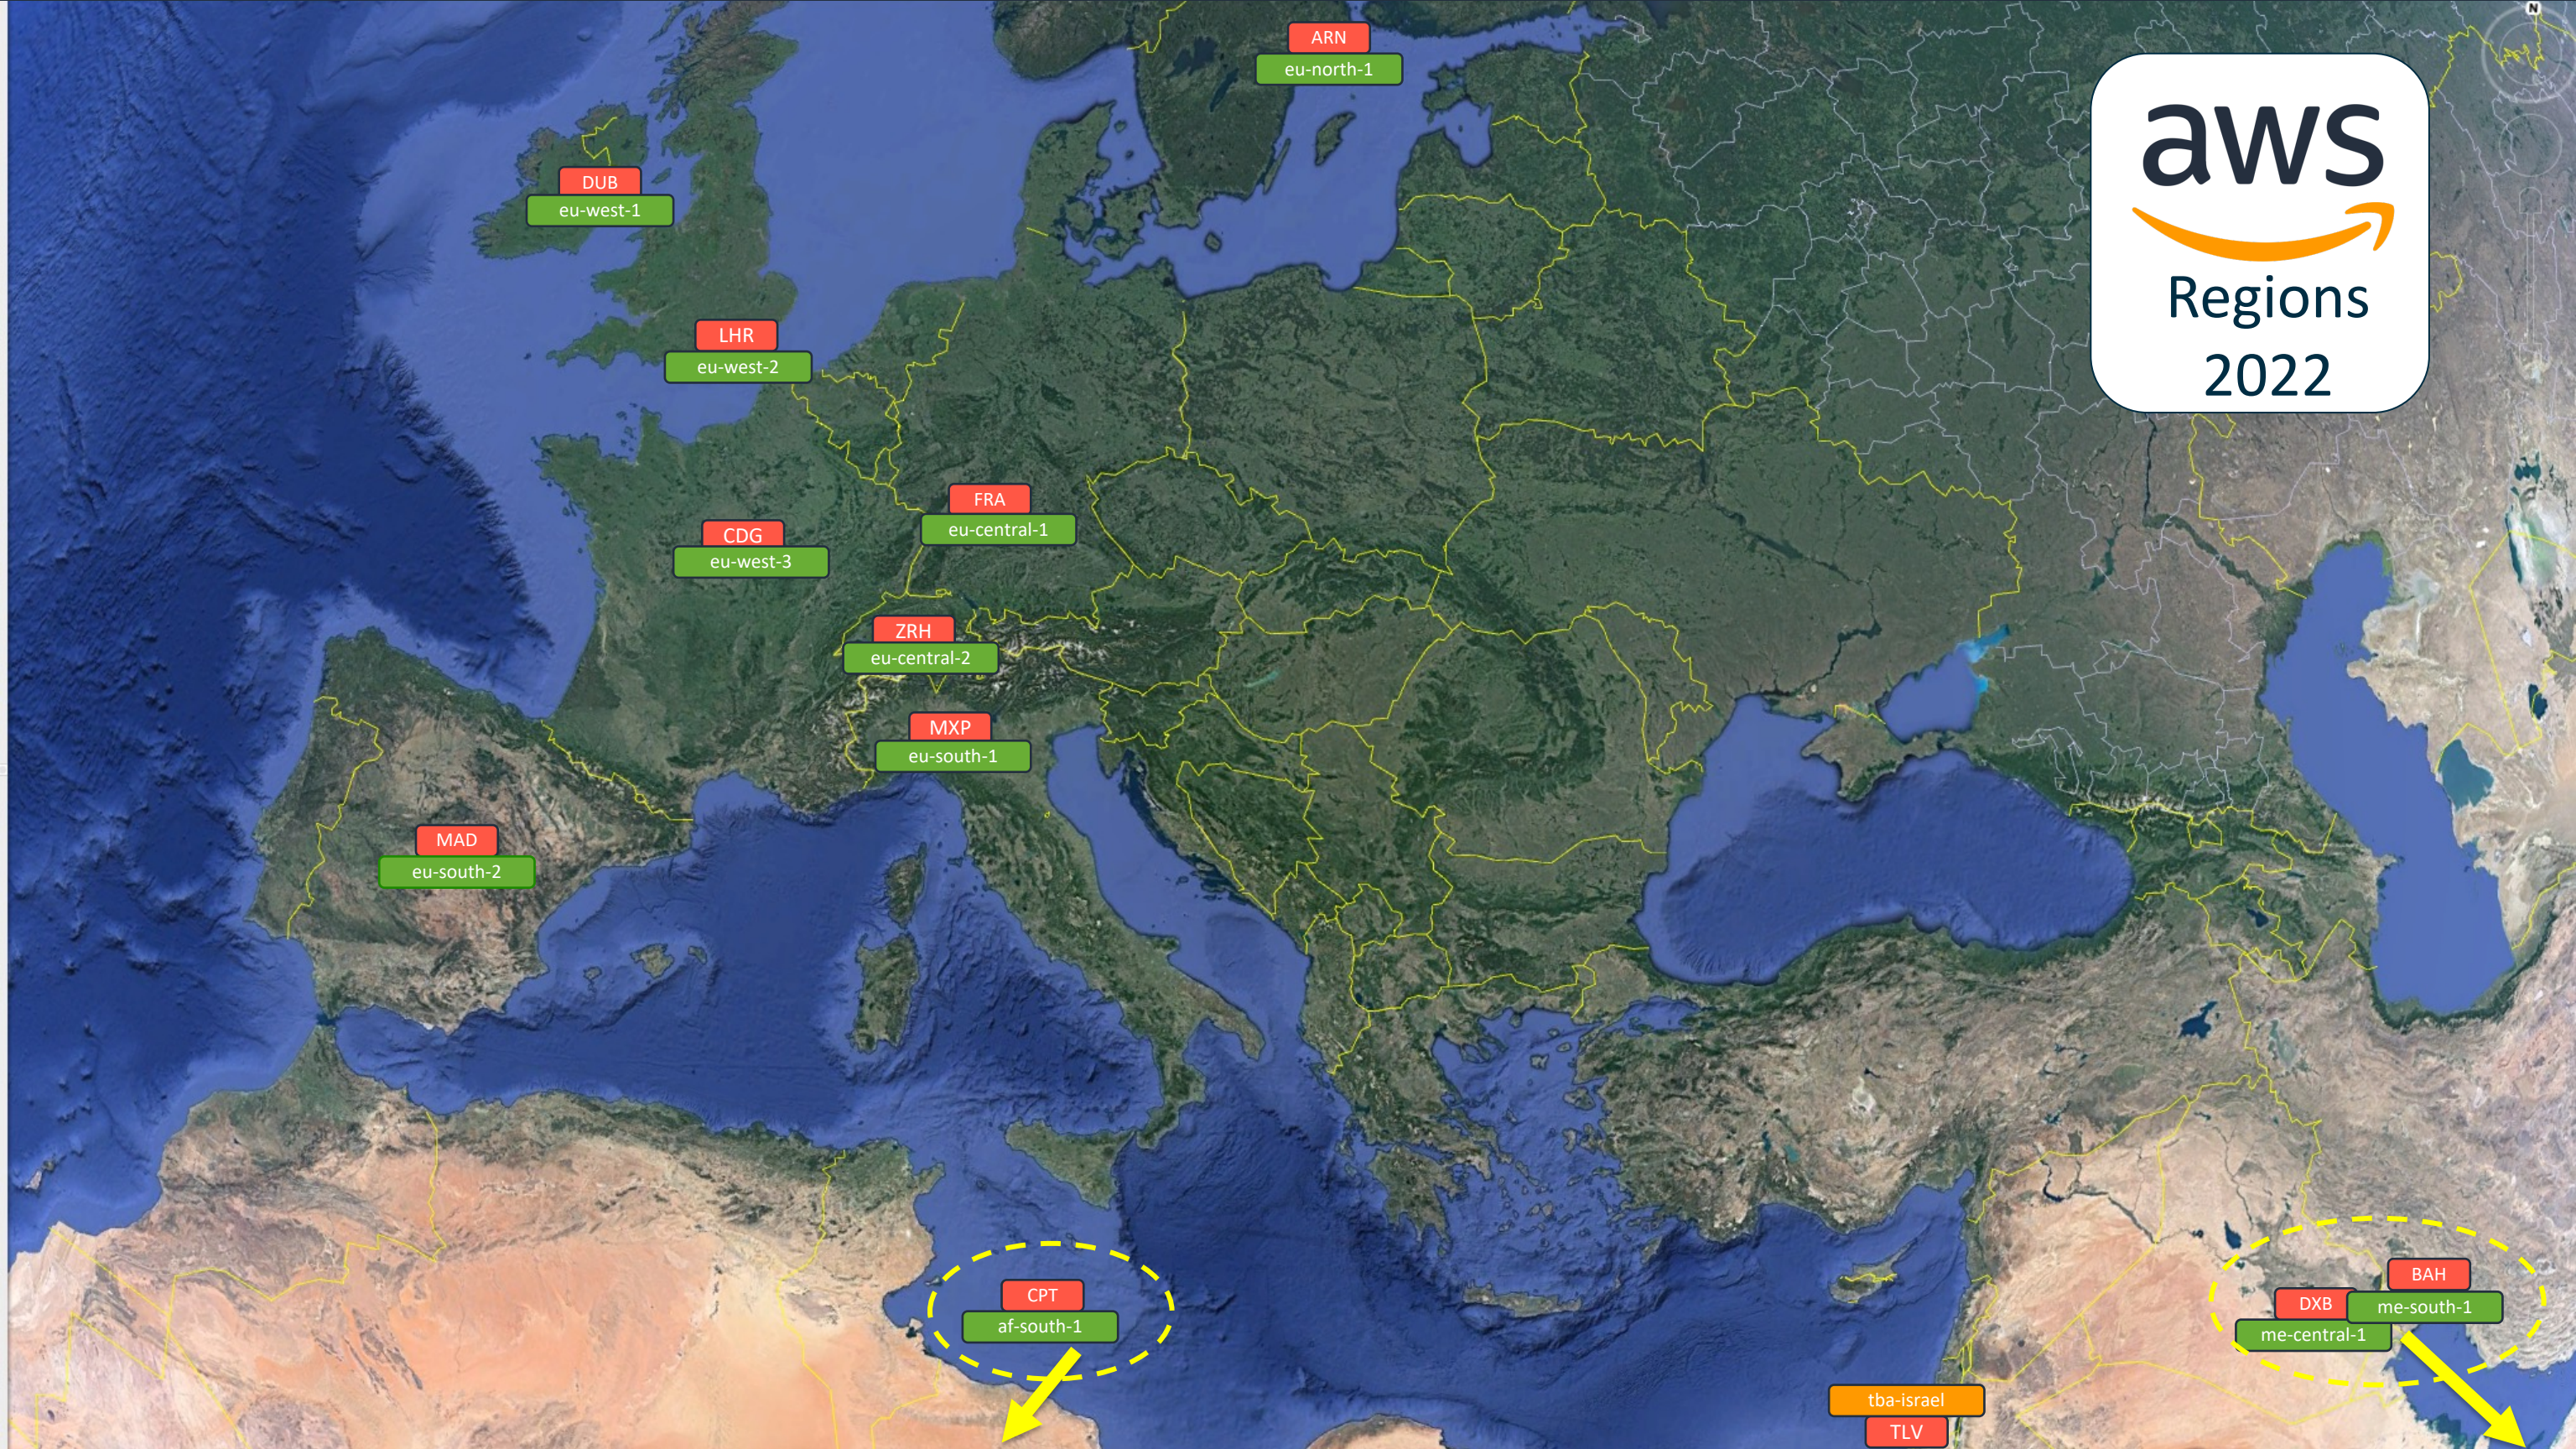
Management
& Governance



Analytics



Regions
2022








What does the AS16509 Network Serve?

- **AWS Edge Services. 400+ Global PoPs**

- **Cloudfront** CDN: Multi-Petabit-scale feature rich CDN used by thousands of customers, everything from Slack to PrimeVideo
- Security: WAF, Shield (DDOS), Route53 DNS
- Acceleration: Global Accelerator, Anycast all the things!
- Computing: Cloudfront Functions, Lambda@Edge (Serverless)

- **AWS Direct Connect**


- Available in select PoPs for direct connectivity to the cloud-regions.
- Can be consumed direct, or through a connectivity partner.
- Comes with **MACSEC**
- Comes with SLA's and QoS.
- Is not the same as “regular” peering



New things for **AS16509**

- **Local zones**

- Select core-features moved into smaller single-AZ cloud-deployment in Edge-locations closer to the end-user to lower latency
- Supports services such as **EC2, EBS, ECS, EKS** and **VPC**
- Ties back to a parent-region for certain services
- Uses existing connectivity



- **Wavelength-Zone** *(The product, not the DWDM-wave...)*

- Similar model to Local zone but for 5G MEC
- Installed into ISP-locations connected into 5G aggregator nodes for ultra-low-latency access to 5G Customers
- Day1 Partners are **Verizon, Vodafone, KDDI, Bell** and **SK-Telecom**





Edge-PoP
2022Q2



General updates

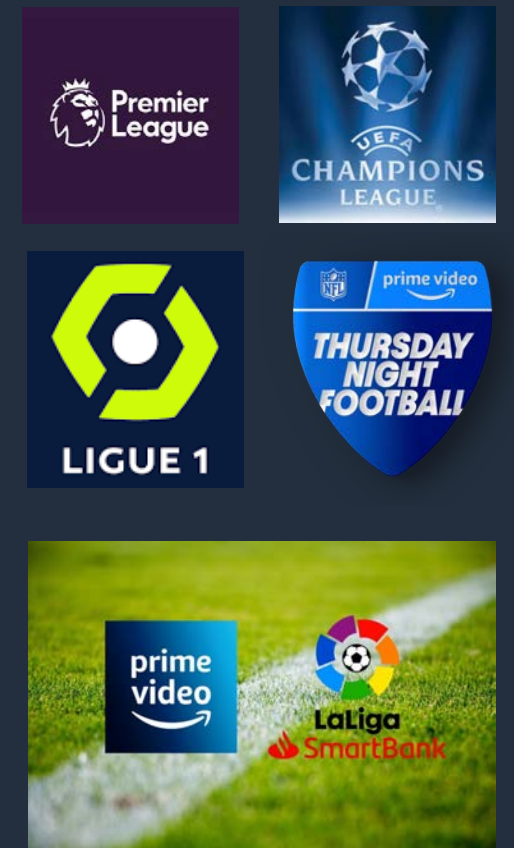
- **Peering concentration:** Especially for IX-based peers. Potential disconnect of IX:es with low local growth and low local coverage. Our smallest capacity-type but one of the largest cost-wise, specifically adjusted for the low growth.

Still not remote-peering, still not using routeservers.

- **Event-focus:** Amazon PrimeVideo is continuing its venture into live-event broadcasting, content acquisition and licensing. **Ligue 1** in France, **Premier League** in UK, **Thursday Night Football** in the US, **UEFA Champion League** in the UK/DE, **LaLiga** in Spain and many more. Ongoing forecasting and planning with all ISPs in respective markets.

Ontop of this we have customers such as DAZN, Hulu (and tons more) using AWS CDN for similar event with live-streaming.

4K UHD HDR is the new norm for both our own and our customers events



UK&IE Updates

- **London**

- **Equinix Slough:** 400G enabled, primary scaleout-site
- **Telehouse North:** 400G enabled, primary scaleout-site. Linx PI
- **Interxion City Campus:** Optical extension, 100G only, not preferred
- **Interxion Sov House:** Small IP PoP, not a preferred site, sunset tba
- **Equinix Harbour Exchange:** Optical extension, 100G only, not preferred

- **LINX LON1** = Nx400G upgrades underway,
- **LINX LON2** = Nx400G upgrades as soon as LINX is ready ;)

- **Manchester**

- **Equinix MA3:** Not a great site...
- **Equinix MA5:** To be launched H1-2023, much more coverage within <10km Darkfibre

- **Dublin**

- We will try to come pick you up anywhere on darkfibre...
- **INEX:** Or peer with us over INEX

UK&IE Updates

- **London**

- **Equinix Slough:** 400G enabled, primary scaleout-site
- **Telehouse North:** 400G enabled, primary scaleout-site. Linx PI
- **Interxion City Campus:** Optical extension, 100G only, not preferred
- **Interxion Sov House:** Small IP PoP, not a preferred site, sunset tba
- **Equinix Harbour Exchange:** Optical extension, 100G only, not preferred

- **LINX LON1** = Nx400G upgrades underway,
- **LINX LON2** = Nx400G upgrades as soon as LINX is ready ;)

- **Manchester**

- **Equinix MA3:** Not a great site...
- **Equinix MA5:** To be launched H1-2023, much more coverage within <10km Darkfibre

- **Dublin**

- We will try to come pick you up anywhere on darkfibre...
- **INEX:** Or peer with us over INEX

- **Rest**

- Lets talk about embedded caches today or tomorrow.

Network Edge Technologies



Cloud Stuff




Peers



Cloud stuff

1RU 32x400G | 1RU 32x400G | 1RU 32x400G | 1RU 32x400G



1RU 32x400G | 1RU 32x400G | 1RU 32x400G | 1RU 32x400G

Peers

400G

- AWS has been a big user and supporter for 400G for a long time. We even had 400G Instances since 2020! (Based on Nvidia A100 ML/HPC)
- Migrations to 400G goes Datacenter -> Backbone -> Edge. Since im here talking with you, we are already in the Edge-readiness stage.
- In 2022 and onwards more and more sites will have 400G support at the AWS Edge-locations available for peering. We are happy to put in orders for 400G augments anywhere in the world from **today** and prioritize our supplychain to make it happen.
- We will use **400G-LR4** in the Edge for external interconnect. Longer distance-optics is being evaluated. 8-lane options is **not** going to happen, 2km options is **not** going to happen either.
- Speak with your fellow AWS-representative about **YOUR** plans for 400G enablement in the edge and our timeline for enablement. (most likely me)

10G/100G

- 10G peering will be of less interest going forward and will not be offered anymore. In our 400G edge-platform, producing a 10G port means sacrificing 390G to 360G of potential capacity on the port
- 100G continues to be the de-facto standard interconnect-method for us going forward for the foreseeable future. Happy to hear and take note if anyone would be interested in **100G-LR1** (Single Lambda) instead of **100G-LR4** to optimize for cost and simplification in 400G native networks.

We are not fully convinced its a "thing" yet due to lacking HW-support in 100G Native platforms.

AWS and **RPKI**, where we are today

AWS and RPKI, where we are today.

Blogpost for full context: <https://aws.amazon.com/blogs/networking-and-content-delivery/how-aws-is-helping-to-secure-internet-routing/>

- We are dropping RPKI **invalids** in **100%** of our Internet Edge Border, in over 400+ global PoPs on all eBGP-peering sessions of all kinds (Transit, IX, PNI)
- We have signed more than 99% of our announced IP-space.
- We have fully automatic ROA-renewal, creation and maintenance in our “IP-vending machine”.
- Bring-Your-Own-IP (**BYOIP**) Relies on RPKI for Correctness
- RPKI-OV and RPKI-ROA-Creation is a ‘Severity 1’ service with oncall-teams on rotation.

AWS and **RPKI**, where we are going

AWS and RPKI, where we are going

1. Investing and looking more into Delegated RPKI solutions, with our own publication points. (Already live under APNIC and RIPE!)
2. Improve the BYOIP-process for customers. Specifically lookint at upcoming RSC IETF Standards.
3. Work with and reach out to networks that has RPKI invalids to have them fixed.
4. Continue the work on community-projects such as MANRS to launch new initiatives and frameworks to foster the use of RPKI.
5. Bring RPKI into RFPs and RFQs as if it would be a standard feature in all connectivity-business going forward.

Questions to fkback@amazon.com

Get in touch!

- Find me, or any of my colleagues in the hallways or in the AWS meeting room if you want to discuss any of these topics
- Use regional aliases when requesting peering! Regular peering@ will not answer
peering-emea@amazon.com
peering-amer@amazon.com
peering-apac@amazon.com
peering-india@amazon.com
- Use our excellent peering tech-ops for all technical inquiries around peering, maintenance or smaller changes.
peering-to@amazon.com.
- Self-service some troubleshooting using our NLNOG-ring trace-nodes and our EC2 reachability matrix linked at our PeeringDB-page.
- External peering-portal has been launched in a limited trial to select peers. Global launch to follow **shortly!**