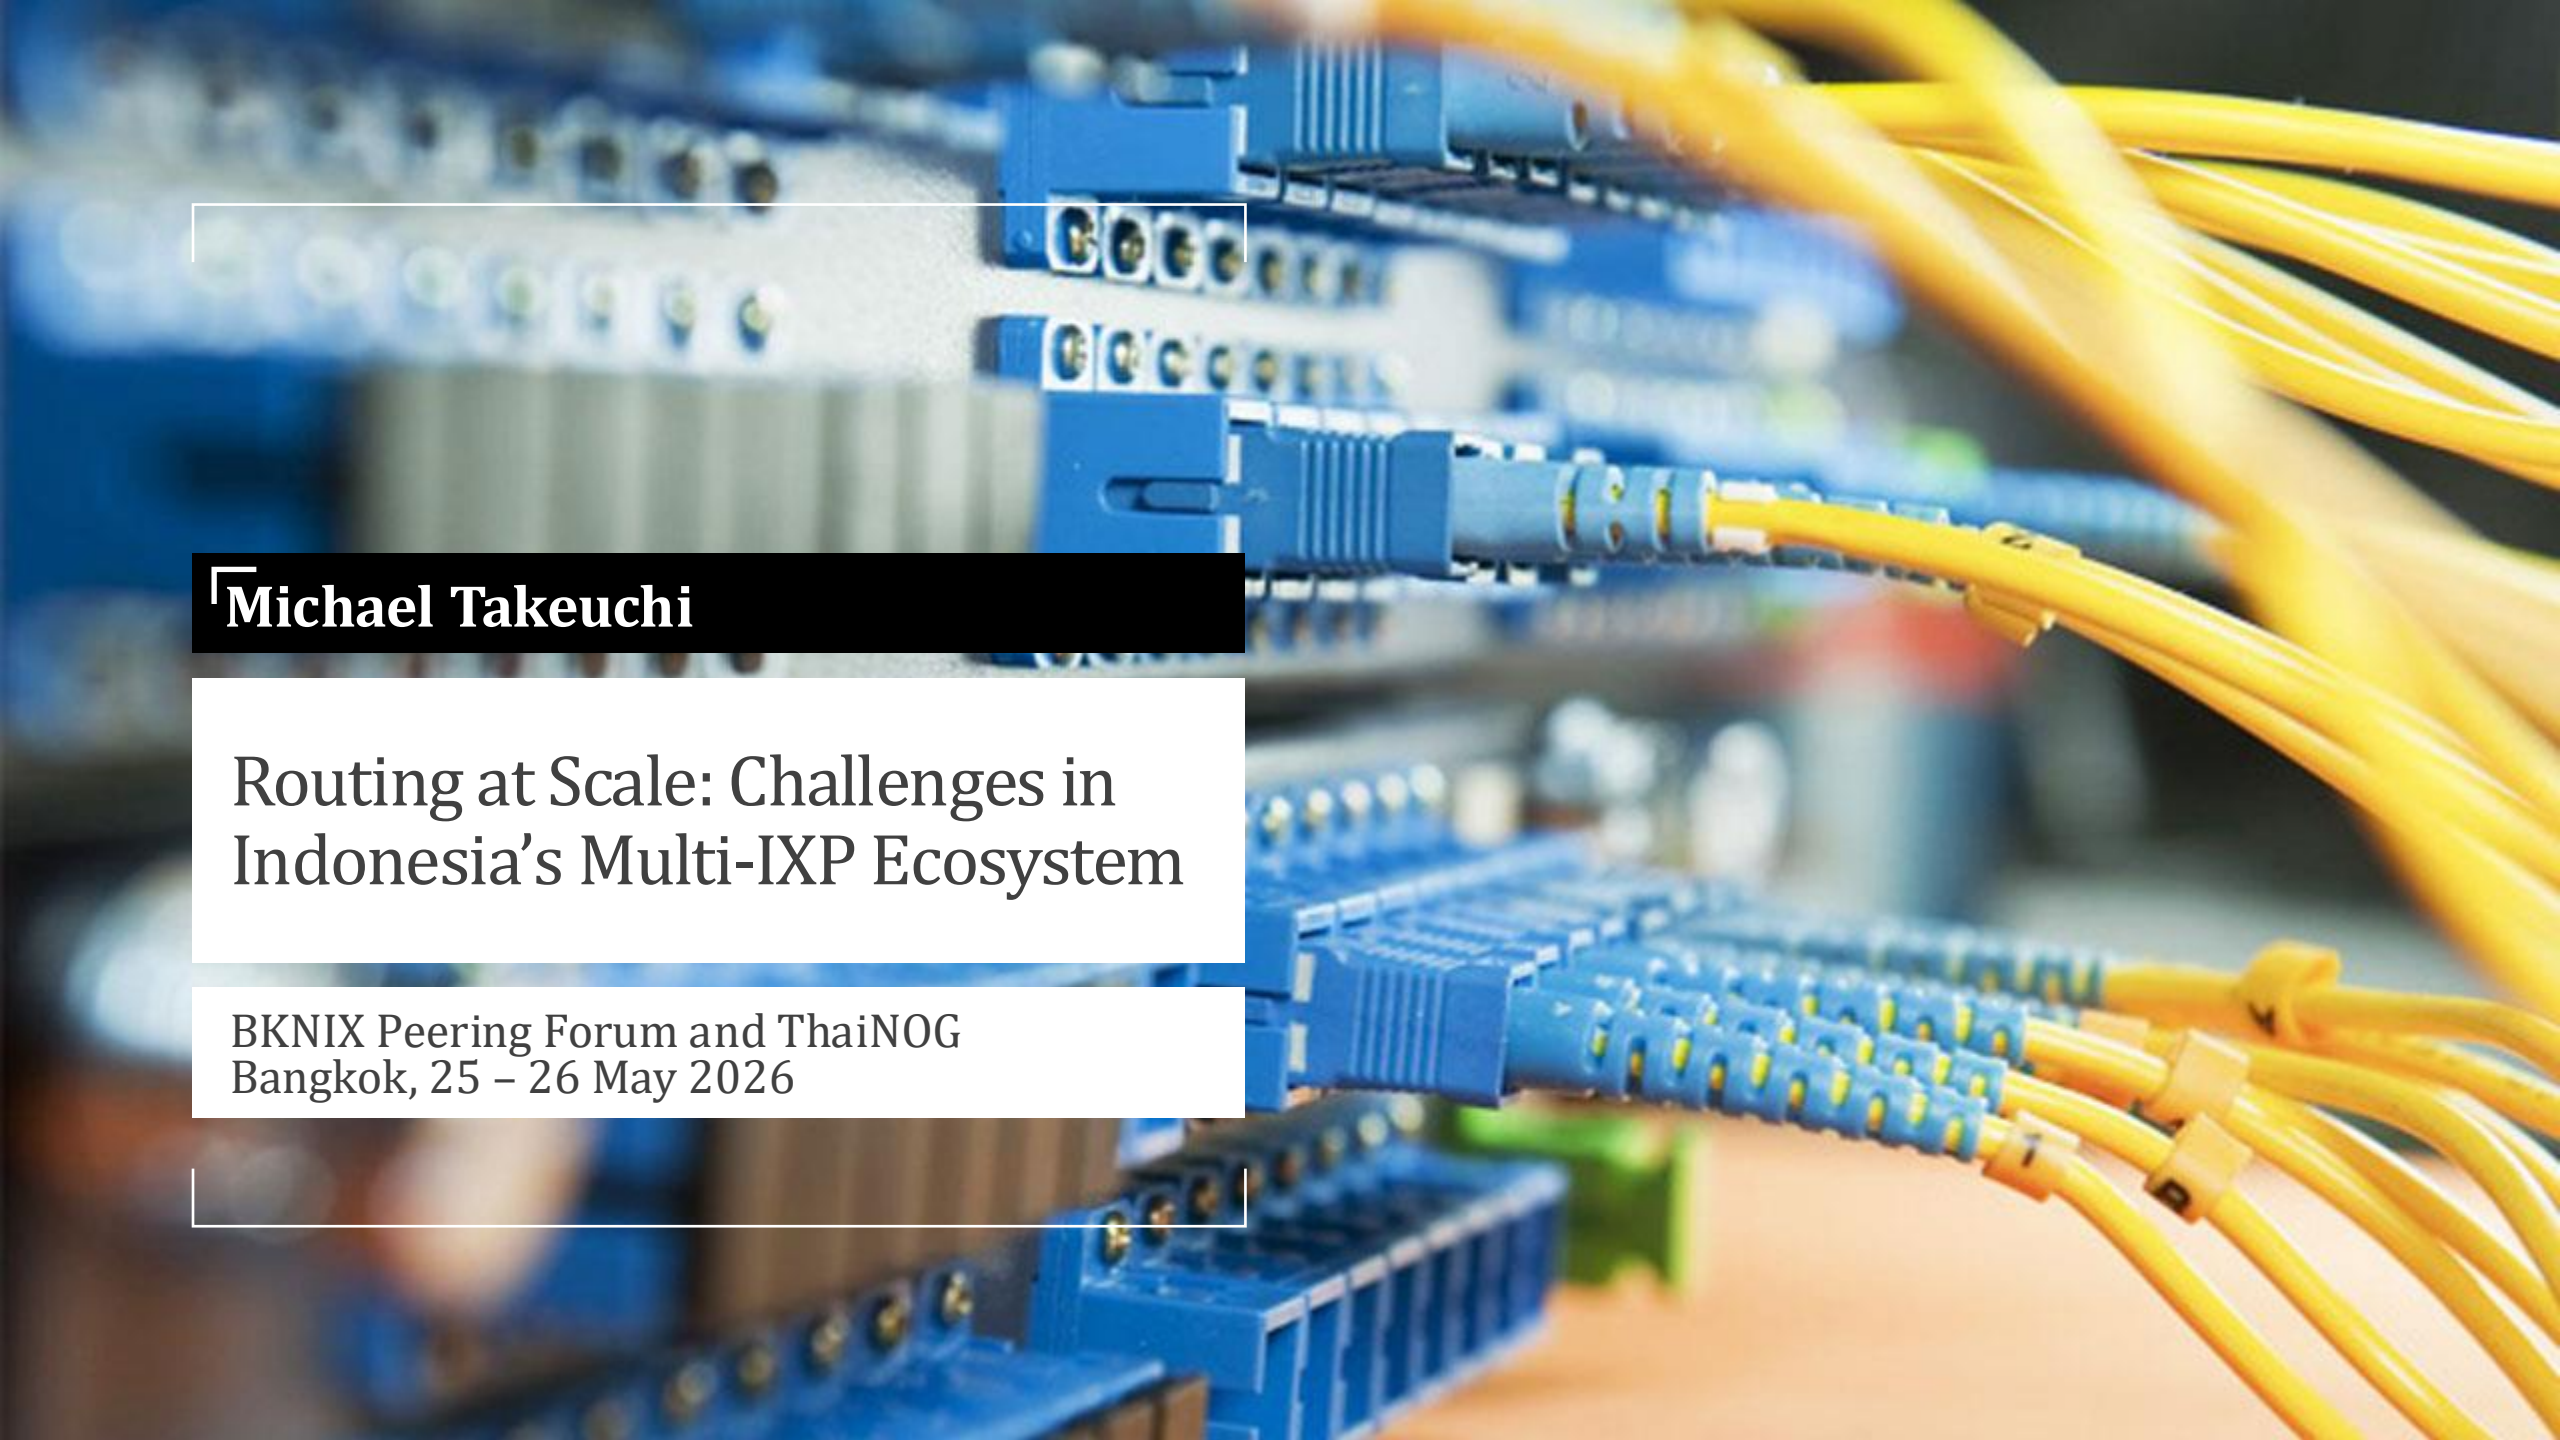

Michael Takeuchi

Routing at Scale: Challenges in Indonesia's Multi-IXP Ecosystem

BKNIX Peering Forum and ThaiNOG
Bangkok, 25 – 26 May 2026

Hi, I'm **Michael Takeuchi**

- Chief Technology & Operating Officer
  of Media Cepat Indonesia, PT (RAPIDNET)
- Deputy Chief of Indonesia Internet Exchange (IIX)
  and Data Center at the Indonesian ISP Association (APJII)
- Certified Juniper JNCIP-DC|ENT|SP, Cisco CCNP-EI,
  EC-Council CEH|CEI|CND|CSA|CCT|CHFI|CTIA|ECIH,
  Fortinet FCP|FCSS, MikroTik Certified Trainer & Engineer
- Based in Jakarta, Indonesia and living around the world

https://www.linkedin.com/in/michael-takeuchi
michael.takeuchi@rapid.net.id / michael.takeuchi@apjii.or.id

Internet eXchange Growth in Indonesia

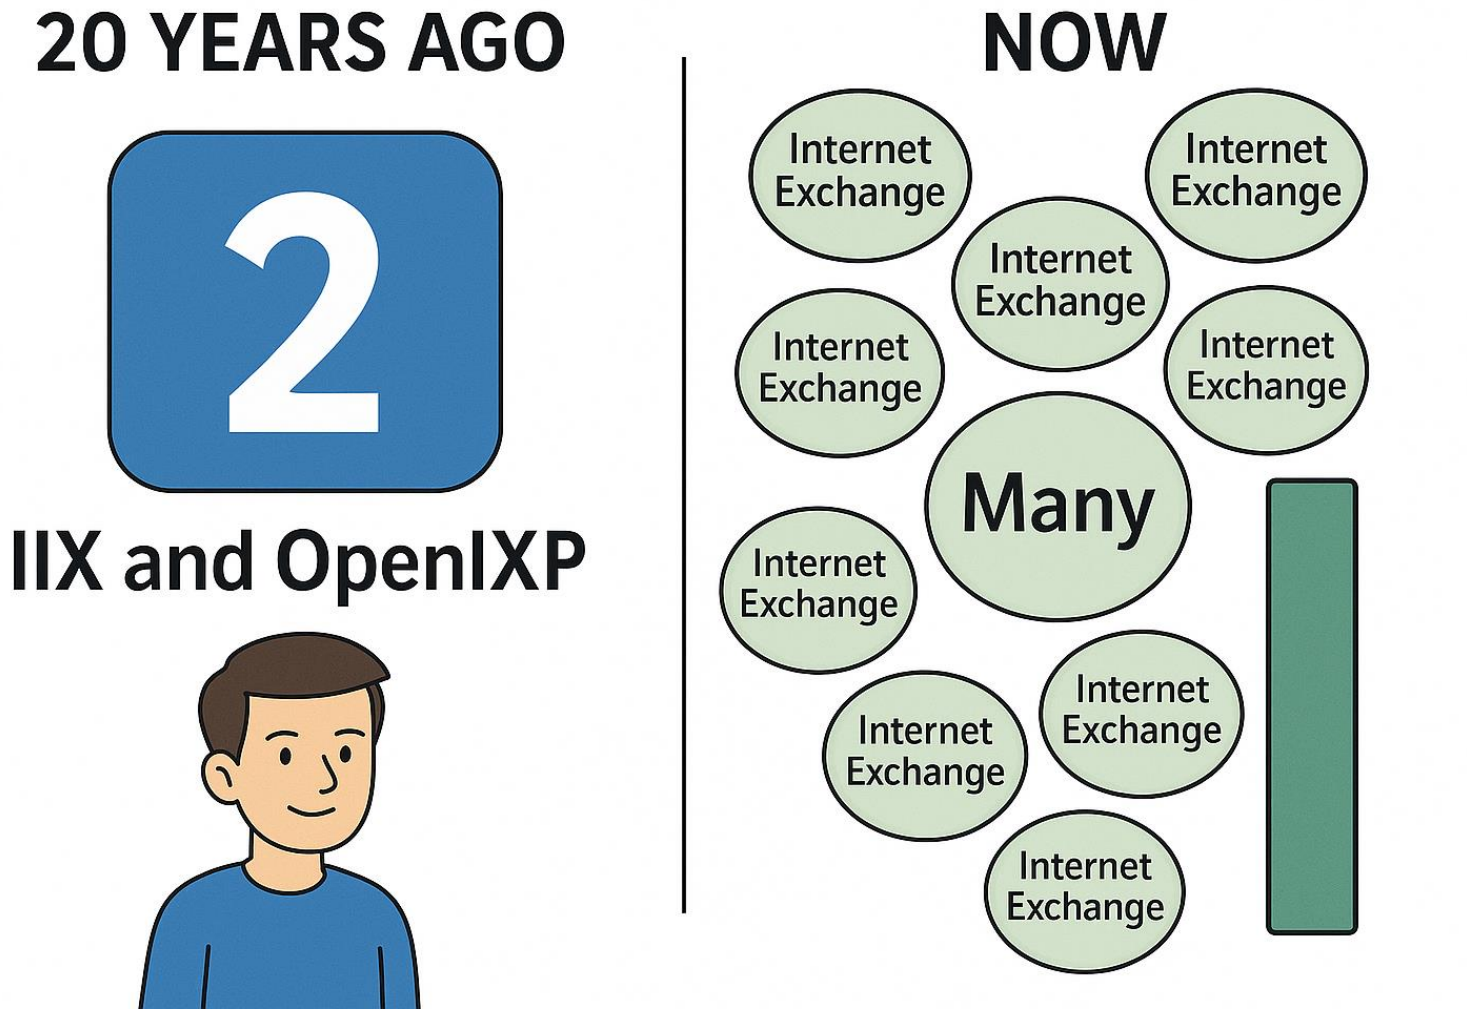

https://www.peeringdb.com/advanced_search?country__in=ID&reftag=ix

1997

IIX

Indonesia Internet Exchange

2000

Peak Traffic of Local Indonesian Internet is…

3Mbps

2005

OpenIXP

National Inter Connection Exchange (NiCE)

2026

>15 IXP in Jakarta and >60 across Indonesia

with Peak Local Traffic in **Tbps**

2026

More than 1400++ Internet Service Providers

And more than 4100++ non-ISP networks (5500+ in total)

One ISP connects to 12 IXP only in Jakarta

Organization	PT Media Cepat Indonesia
Also Known As	RAPID Networks
Long Name	Media Cepat Indonesia (RAPIDNET)
Company Website	https://rapid.net.id
ASN	141137
IRR as-set/route-set	AS141137:AS-RAPIDNET
Route Server URL	
Looking Glass URL	
Network Types	NSP
IPv4 Prefixes	1000
IPv6 Prefixes	1000
Traffic Levels	5-10Tbps
Traffic Ratios	Mostly Inbound
Geographic Scope	Asia Pacific
Protocols Supported	✓ Unicast IPv4 ○ Multicast ✓ IPv6 ○ Never via route servers
Last Updated	2025-05-28T08:47:39Z
Public Peering Info Updated	2025-05-28T08:47:33Z
Peering Facility Info Updated	2025-03-17T09:16:39Z
Contact Info Updated	2024-05-10T11:44:58Z
Notes	
RIR Status	ok

Public Peering Exchange Points

Filter

Exchange IPv4	ASN IPv6	Speed Port Loc…	RS Peer	BFD Su…
DE-CIX Jakarta 103.159.71.67	141137 2001:df5:7880:0:2:2751:0:1	100G	✓	✓
Digital Edge EPIX Jakarta 202.159.50.35	141137 2402:ac00:de00:de00:0:14:1137:1	100G	✓	✓
EdgeNXT 103.61.232.69	141137 2001:df1:d440:1::69	100G	✓	✓
IIX-Jakarta 123.108.8.195	141137 2001:7fa:2:5::c3	100G	✓	✓
JKT-IX 119.11.184.192	141137 2404:c8:0:a:0:14:1137:1	100G	✓	✓
OpenIXP / NiCE 218.100.36.65	141137 2001:7fa:f::1b1	100G	✓	✓
C2IX 218.100.72.55	141137 2001:df0:b1:100::72:55	20G	✓	✓
CDIX 103.30.172.251	141137	20G	○	✓
CXC Jakarta 103.225.173.119	141137 2400:9c80:0:173::1411:37	20G	✓	✓
DCI Indonesia DCI-IX 103.142.207.86	141137	10G	✓	✓
Equinix Jakarta 202.3.136.8	141137 2001:dee:c000::14:1137:1	10G	✓	✓

Many Challenges!!!

1. Routing Complexity

- Too many BGP sessions = harder to maintain stability, especially for small router resources.

```
> show bgp summary | match peers
Groups: 21 Peers: 387 Down peers: 41
```

- Network engineers need to manage more granular policies to filter and maintain the routing, which increases the risk of fat-finger errors. Even by automation?

1. Routing Complexity

- 1 prefix received from 11 BGP sessions from 5 different IXP

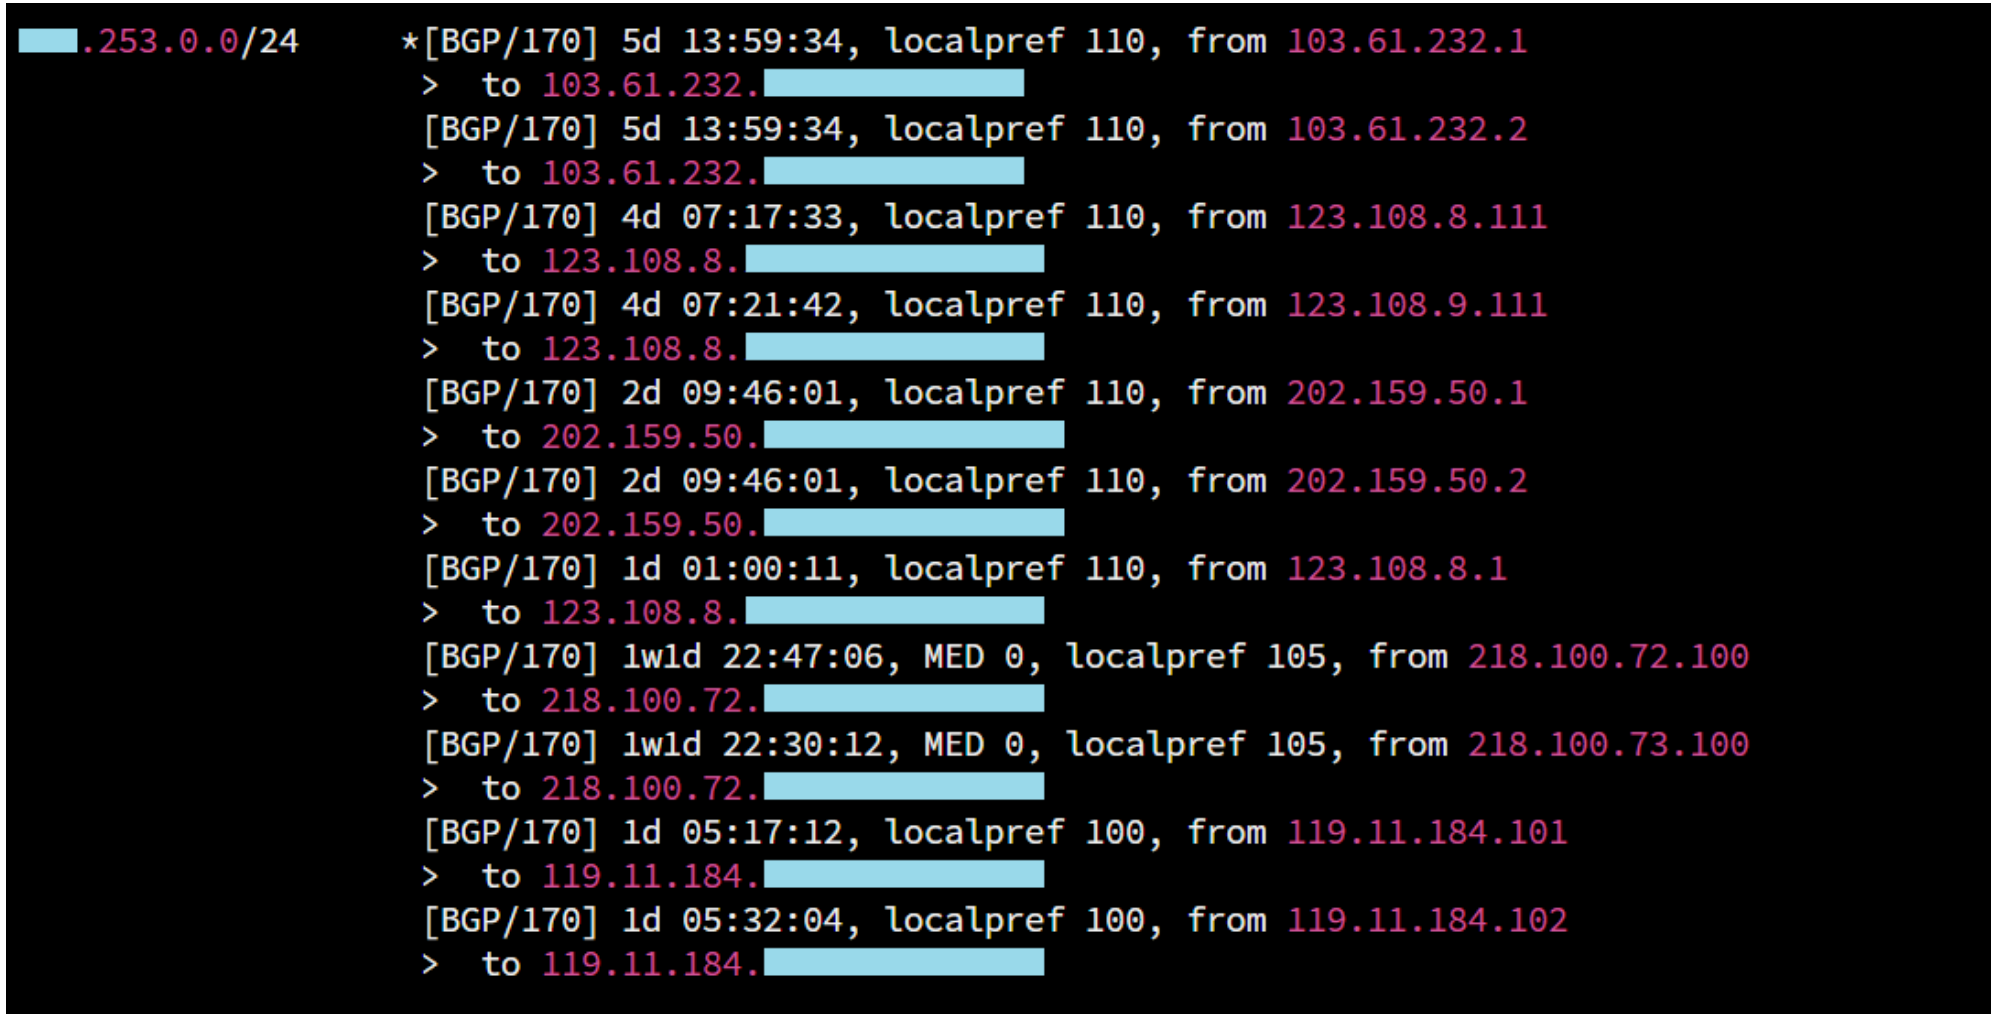

2. Suboptimal Pathing – Case Study 1

- Without topology awareness, BGP may misroute traffic by prioritizing the shortest AS-PATH **over actual latency**. As a result, even local Indonesian IPs might be routed through Singapore first, causing much higher latency.

```
traceroute to 103.167.11.xxx (103.167.11.xxx) from 103.xx.xx.x, 30 hops max, 52 byte packets
 1  * * *
 2  218.100.36.xxx [AS  7717]  1.564 ms  1.445 ms  1.504 ms
 3  184.104.189.xxx [AS  6939]  12.922 ms  16.596 ms  13.477 ms << latency start to increase
 4  * * *
 5  184.104.199.xxx [AS  6939]  21.789 ms  14.187 ms  13.083 ms
 6  184.104.210.xxx [AS  6939]  13.537 ms  13.650 ms  13.803 ms
 7  103.87.186.xxx [AS  136106]  14.273 ms  14.636 ms  14.028 ms
 8  103.87.185.xxx [AS  136106]  13.225 ms  20.647 ms  13.162 ms
 9  103.167.11.xxx [AS  141137]  15.375 ms  14.099 ms  13.768 ms
```

2. Suboptimal Pathing – Case Study 2

- Not only Indonesia! I also experiencing the same issue with my POP in Singapore, the routing are going from Singapore to Amsterdam -> Marseille and back to Singapore again
- Here is the traceroute from another network to my network

```
traceroute to 112.7x.xx.x(112.7x.xx.x), max hops: 30 ,packet length: 40,press CTRL_C to break
 1 xxx.x.xxx.13 4 ms  4 ms  4 ms
 2 xxx.x.xx.x96 5 ms  11 ms  5 ms
 3 80.249.209.150 243 ms  239 ms  245 ms << latency start increasing more than 200ms
 4 184.105.81.109 245 ms  245 ms  245 ms
 5 184.105.222.22 167 ms  171 ms  172 ms
 6 184.105.65.13 163 ms  167 ms  169 ms
 7 184.105.80.14 179 ms  192 ms  183 ms
 … … … traceroute are truncated … … …
```

3. Lack of Visibility

… Some networks give you a looking glass.
… Some do not even show basic telemetry.
… No NMS, no traffic stats, no visibility.
… And we are still expected to troubleshoot in the dark.

Network Engineer = shaman? magician? witch?

We had to escalate the routing issue
to an external black magic consultant.

3. Lack of Visibility – Case Study

Due to the lack of visibility, no NMS dashboard, and limited ports on other peering members, I experienced congestion even though my own port wasn't congested. How can we monitor the others?

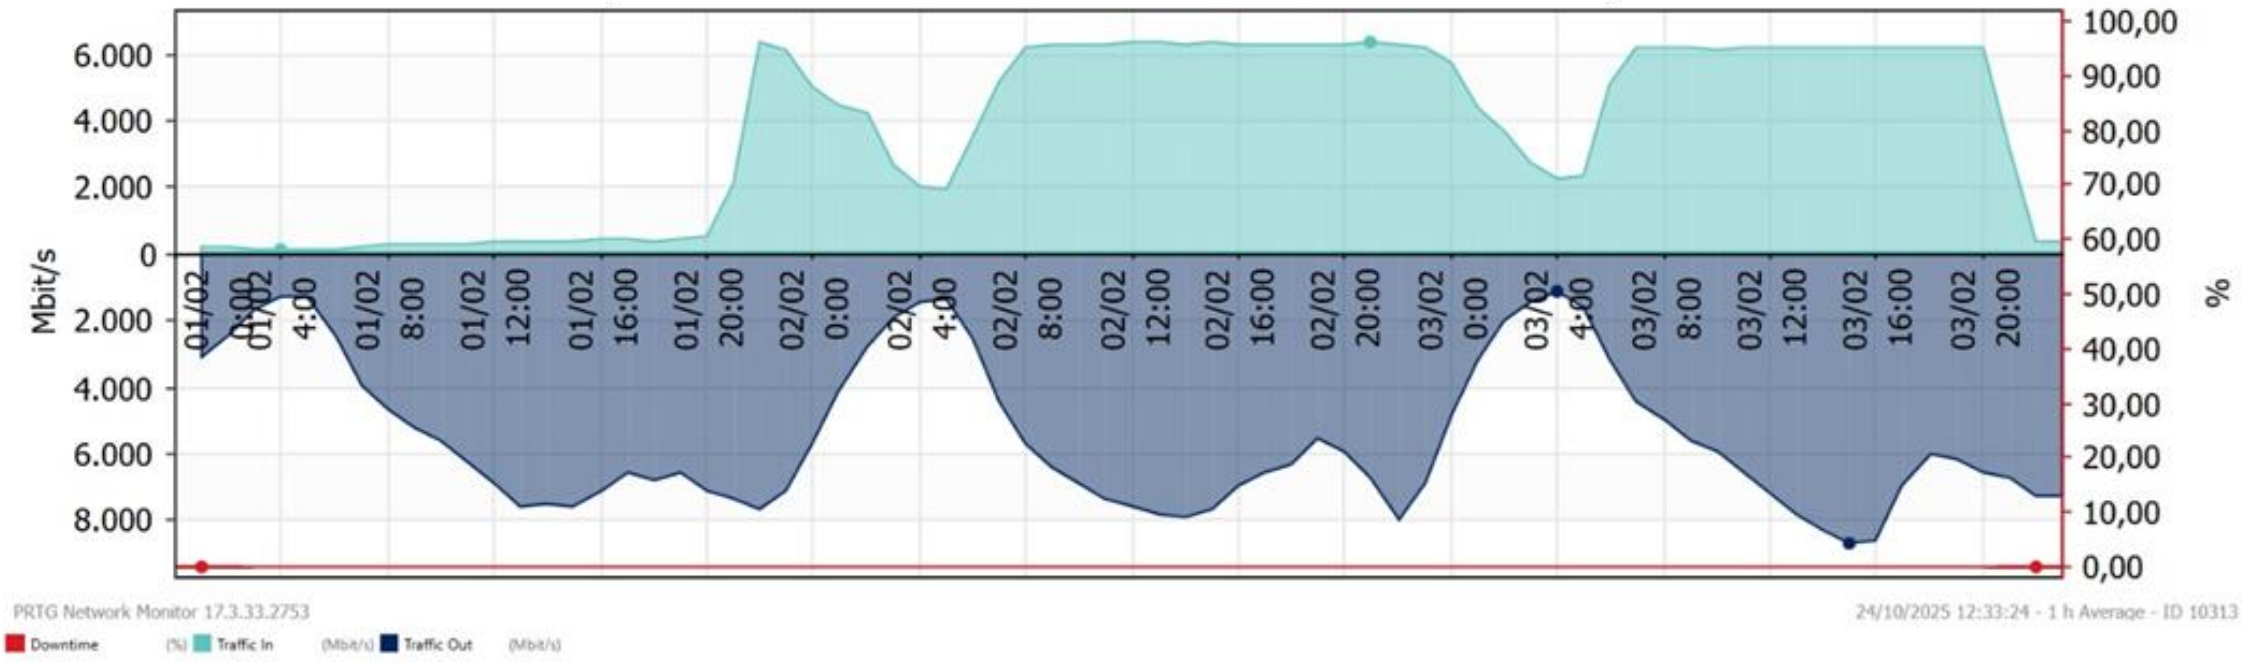

4. No Common Standard

⚙️ **Different Platforms:**
Each IXP runs unique Route Server software and configs

📜 **No Common Standards:**
Lacking shared BGP Looking Glass and IRR compliance

🏷️ **Inconsistent Community Tags:**
Different or no even BGP Community support for traffic engineering

🤝 **Different Peering Policies:**
Each IXP applies its own peering and filtering rules

Shall we submit new RFC/standardization on it? Or does it exist?

4. No Common Standard – Case Study

No BGP Community support in this IXP. How can we avoid path congestion? By manually doing selective peering? Yes, if we have the same IP subnet with the others to do bilateral peering

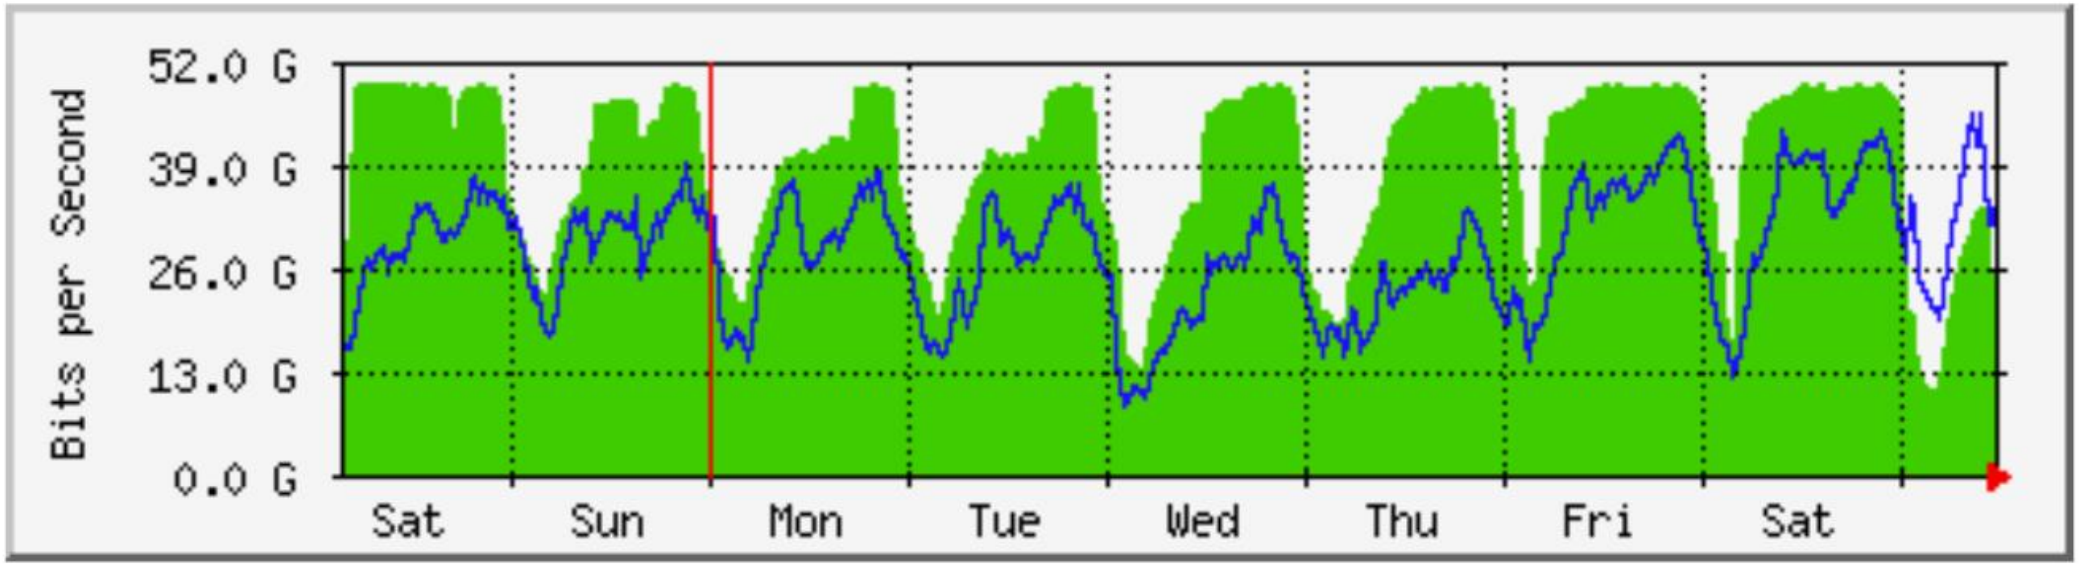

	Max	Average	Current
In	49.2 Gb/s (1.2%)	39.1 Gb/s (1.0%)	24.8 Gb/s (0.6%)
Out	45.2 Gb/s (1.1%)	28.0 Gb/s (0.7%)	29.6 Gb/s (0.7%)

4. No Common Standard – Existing Effort

We are not completely lawless. **MANRS** is already here, and it functions as the closest thing we have to a shared rulebook.

What MANRS Brings to IXPs

- Security baseline through IRR and RPKI validation
- Route Server filtering and routing hygiene requirements
- Shared minimum level of routing safety for all participants
- Requires contact information and escalation paths for members
- Encourages documentation and transparent operational policies

It is help? **Yes, but** MANRS is not the final answer, but it gives us a common foundation to build real improvements on top of.

4. No Common Standard – Gap Analysis

🔍 Gaps That Still Need Work

- No standard BGP community format for TE or blackholing
- Route Server behavior differs across platforms and regions
- Tools and visibility are not aligned (Looking Glass, APIs, telemetry)
- Peering policies vary widely between IXPs
- Automation is still fragile without unified tagging or policy formats
- We need to have an operational approach, not only the security

4. No Common Standard – Gap Analysis

From my personal perspective…

- **MANRS gives us security.**
  - ✓ Strong baseline
  - ✓ Absolutely useful
  - ✓ A step in the right direction
- **But, in daily operations…**
  - Operational inconsistencies still happen
  - Non-technical factors still slow us down
  - Without shared standards, every IXP speaks its own “BGP accent”

Happy to discuss!

5. Prefix Deaggregation & Route Table Bloat

🔷 Some networks announce /24 separately at each IXP instead of advertising a single /22 to do traffic engineering or having some purpose on it
→ This makes routing tables getting bigger and puts extra CPU and memory load on routers, especially if you are using smaller L3 switches as your border.

🔷 One /22 = four /24 + two /23
→ If those four /24 and two /23 are received across 10 BGP sessions, that becomes 60 routes from just one /22 prefix (excluding the /22 itself).

🔷 Without proper filters at the IXP, this can create routing pollution.
→ Worst case, even /32 leaks can slip through.

5. Prefix Deaggregation & Route Table Bloat

This is my router for domestic routes only. We have **260749** routes in total (RIB), with **38016** currently active (FIB).

```
[redacted]> show route summary table inet.0 | except "(Direct|Local|Static|IS-IS|LDP)"
Autonomous system number: 141137
Router ID: [redacted]

Highwater Mark (All time / Time averaged watermark)
    RIB unique destination routes: 78602 at 2025-12-04 11:11:36 / 65573
    RIB routes                   : 286854 at 2025-12-04 11:27:49 / 275638
    FIB routes                   : 53066 at 2025-12-04 11:27:50 / 42007
    VRF type routing instances   : 0 at 2025-11-26 00:48:01

inet.0: 59353 destinations, 260956 routes (38223 active, 0 holddown, 107542 hidden)
                BGP: 260749 routes,  38016 active

[redacted]>
```

Then… what can we do?

What we can do…

- **Ensure zero internal congestion**
  - Track P95 (95th percentile traffic usage) and upgrade before saturation.
  - Build enough headroom and rely on capacity alerts, not customer complaints.
- **Build real redundancy everywhere**
  - Use dual upstreams, IXPs, and route reflectors with MPLS/EVPN or LAG.
  - Ensure diverse fiber paths, POP-level failover, and full power/route backups.
- **Implement robust routing policies**
  - Use selective advertisements, community-based control, and consistent inbound/outbound traffic engineering.
  - Enforce RPKI/IRR filtering and max-prefix limits for clean routing table.

What we can do…

- **Deploy real observability & monitoring**
  - Use telemetry (sFlow/NetFlow) with Zabbix/Prometheus/Grafana for real alerting.
  - Monitor CDN server routing paths and some low-latency services we used, plus provide public looking glass + status transparency.
- **Plan ahead, not react**
  - Forecast capacity and track lifecycle/refresh cycles proactively.
  - Maintain tested BCP/DR runbooks and plan hardware before demand forces you.
- **Stay aligned with global standards**
  - Follow MANRS, RFC, and actively engaging in IXP & NOG communities.

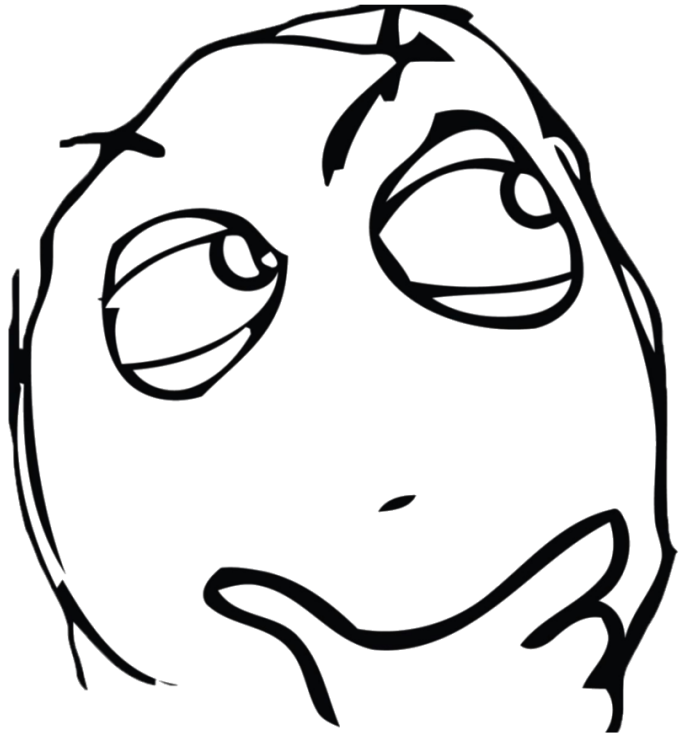

Enough?

Not Really…
We need to…

COLLABORATION AND COMMUNICATION

Why It Matters?

- Multi-IXP routing is not just technical things to do, it's a **communication and coordination problem.**
- Poor communication causes **suboptimal paths routing, route leaks, path congestion or any other problems**.
- Solving it requires **shared responsibility** and **active collaboration** across ISPs, IXPs, and the network community.
- Don’t forget to **update** your PeeringDB and WHOIS **contacts**!

Well… my conclusion is…

Whatever remains unresolved, we resolve it at the dining table!

Thank you, let's eat!

michael.takeuchi@rapid.net.id