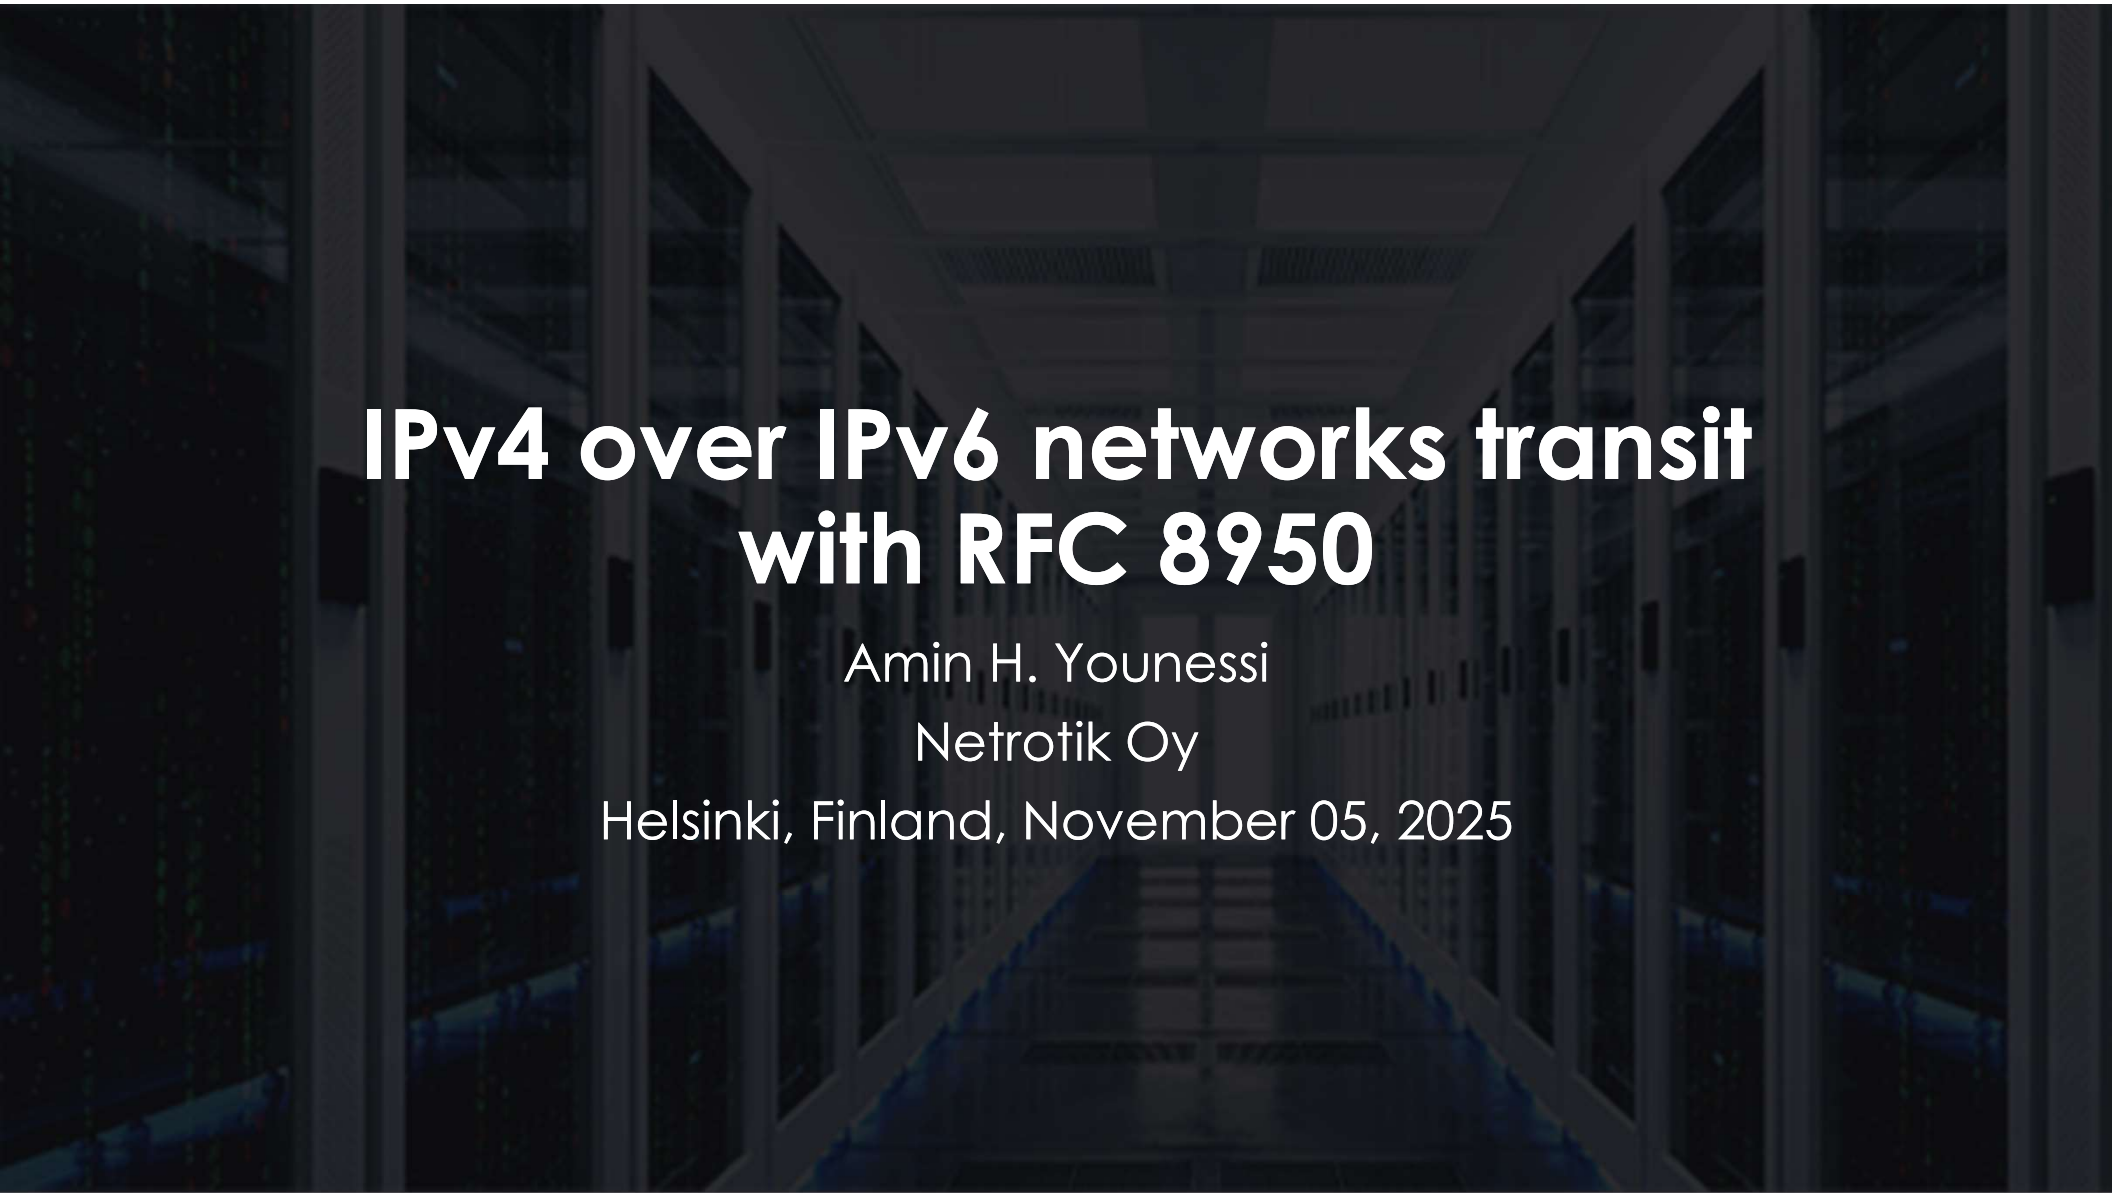

IPv4 over IPv6 networks transit with RFC 8950

Amin H. Younessi
Netrotik Oy
Helsinki, Finland, November 05, 2025

The IPv4 Challenge & a Modern Solution

Enabling IPv4 Services in an IPv6-only World with RFC 8950

1. The Problem: An IPv4-Constrained World
   - IPv4 Address Exhaustion: We've run out of public IPv4 addresses. This is a critical issue for ISPs and enterprises trying to scale and connect new customers.
   - The Transition to IPv6 is Essential: IPv6 offers a massive address space, but the world isn't fully IPv6-native yet. Many legacy services, websites, and clients still require IPv4.
2. The Old Solutions: Complex & Costly
   - Dual-Stack: Running both IPv4 and IPv6 on every router and every link. This is operationally complex, doubles the management overhead, and consumes router resources.
   - Complex Tunneling & NAT: Solutions like NAT444 or DS-Lite create multiple layers of complexity, introduce latency, and are difficult to troubleshoot.

The IPv4 Challenge & a Modern Solution

3. The Modern Solution: RFC 8950
   - A Simple Idea: Use your high-speed, scalable IPv6 core network as the transport for your IPv4 traffic.
   - How it Works: BGP is used to advertise IPv4 routes, but the next-hop address is an IPv6 address. This allows IPv4 packets to traverse your IPv6-only core without any complex tunneling.
   - Key Benefits:
     - Operational Simplicity: Run a single, clean IPv6-only core. No more dual-stack complexity.
     - Reduced Overhead: Less protocol management, fewer configurations.
     - Scalability: Future-proof your network by moving towards a pure IPv6 environment while still serving IPv4 customers.

The Magic behind RFC8950

1. The Core Concept: MP-BGP in Action
   - RFC 8950 isn't a new protocol; it's a new way to use an existing one: **Multiprotocol BGP (MP-BGP)**.
   - MP-BGP allows BGP to carry routing information for protocols other than IPv4, such as IPv6 or even VPN routes.
   - RFC 8950 specifically defines how to advertise **IPv4** routing information using an **IPv6** next-hop.

The Magic behind RFC8950

2. How it Changes BGP
   - Standard BGP: When an IPv4 prefix (10.0.0.0/24) is advertised, the BGP next-hop must be an IPv4 address. This requires a full dual-stack network.
   - With RFC 8950: The router advertises the same IPv4 prefix, but the next-hop is a **global IPv6 address**. The IPv4 traffic is then encapsulated and transported over the IPv6-only core network.

The Magic behind RFC8950

3. The Formal Name
   - **RFC 8950:** Advertising IPv4 Network Layer Reachability Information (NLRI) with an IPv6 Next Hop.

The IXP and ISP Lab Topology

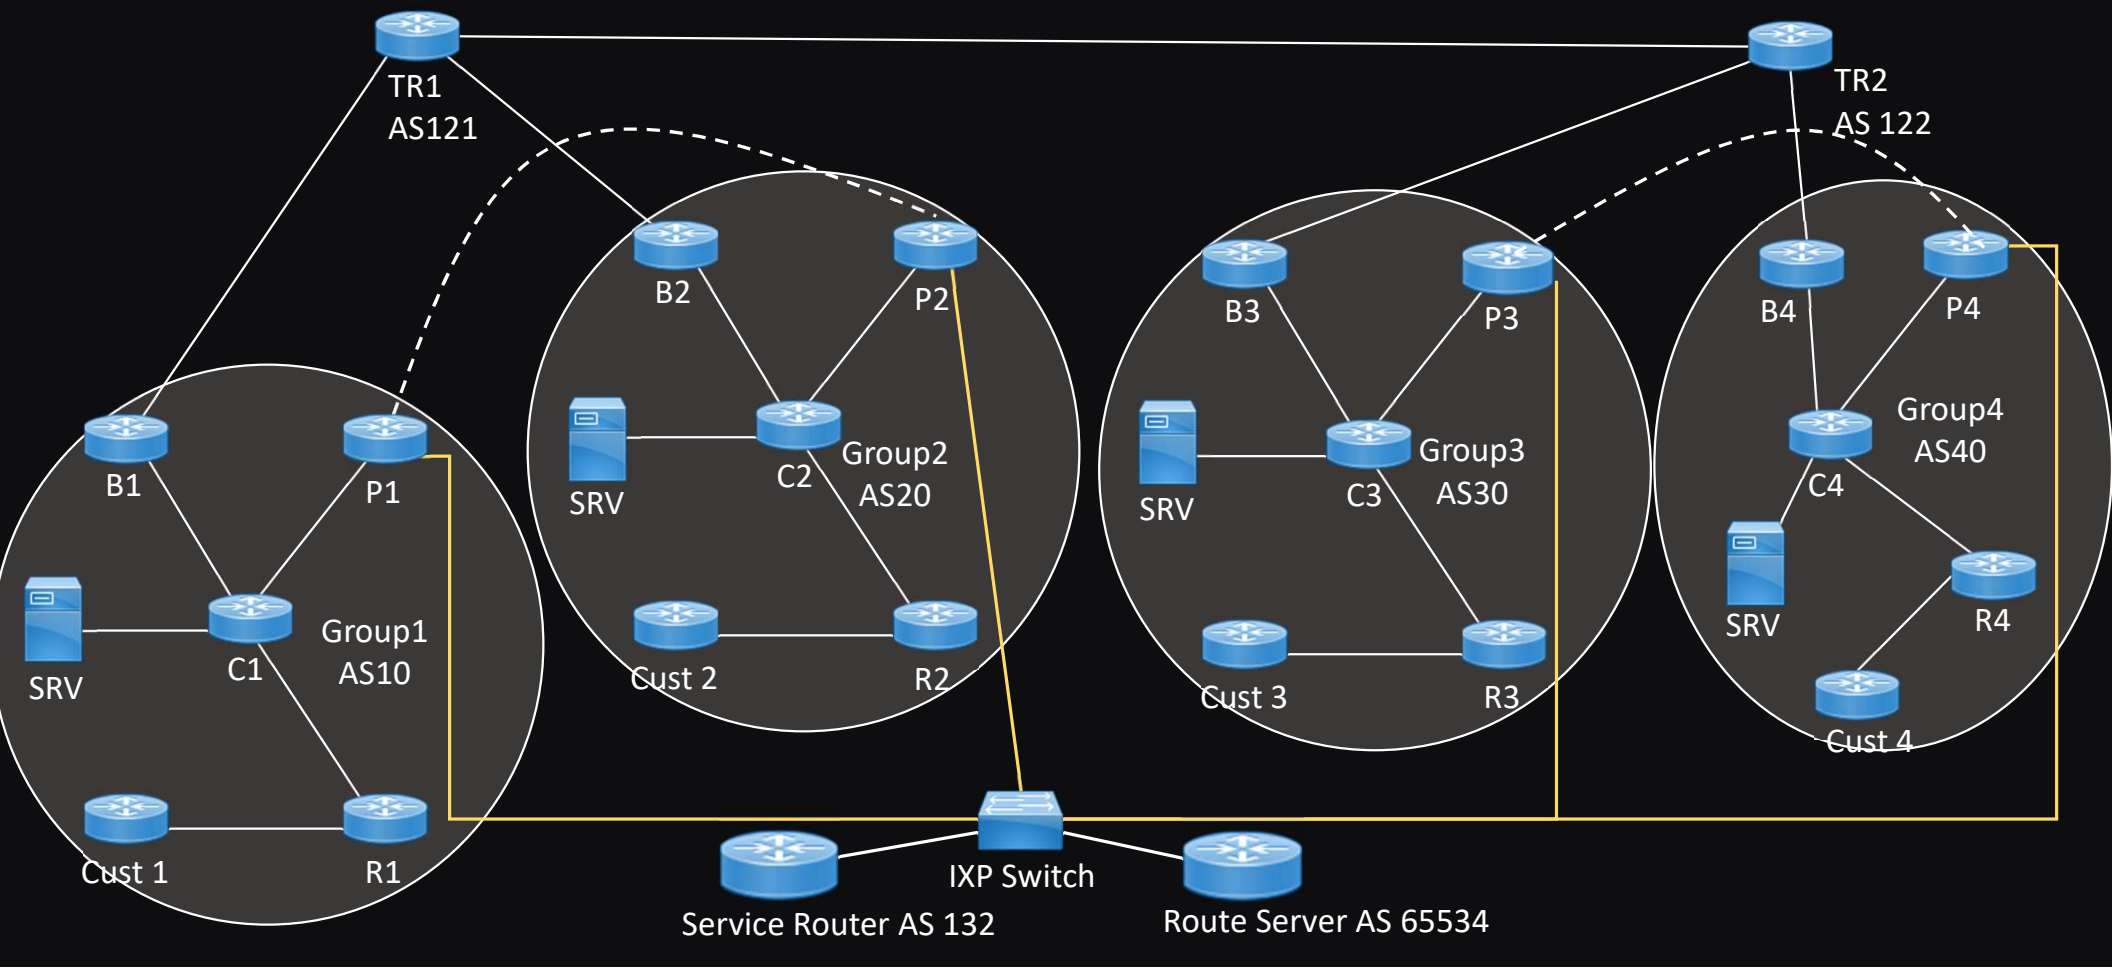

Practical Demonstration in EVE-NG

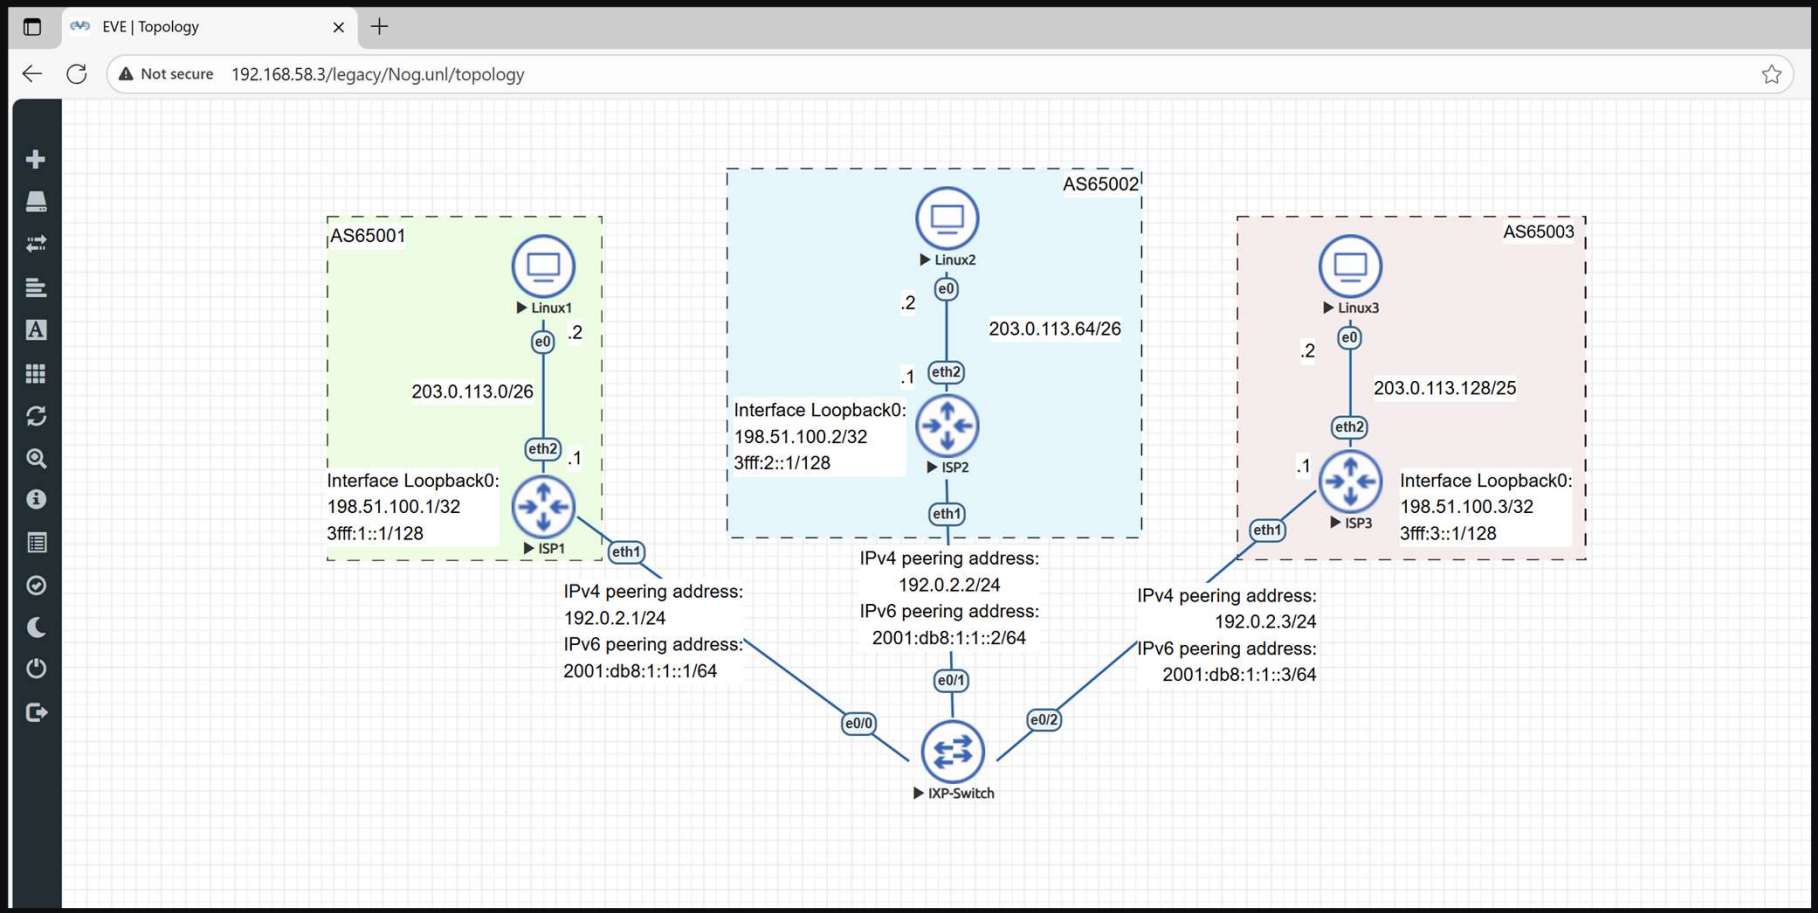

BGP configuration in MikroTik RouterOS 7.20.rc3

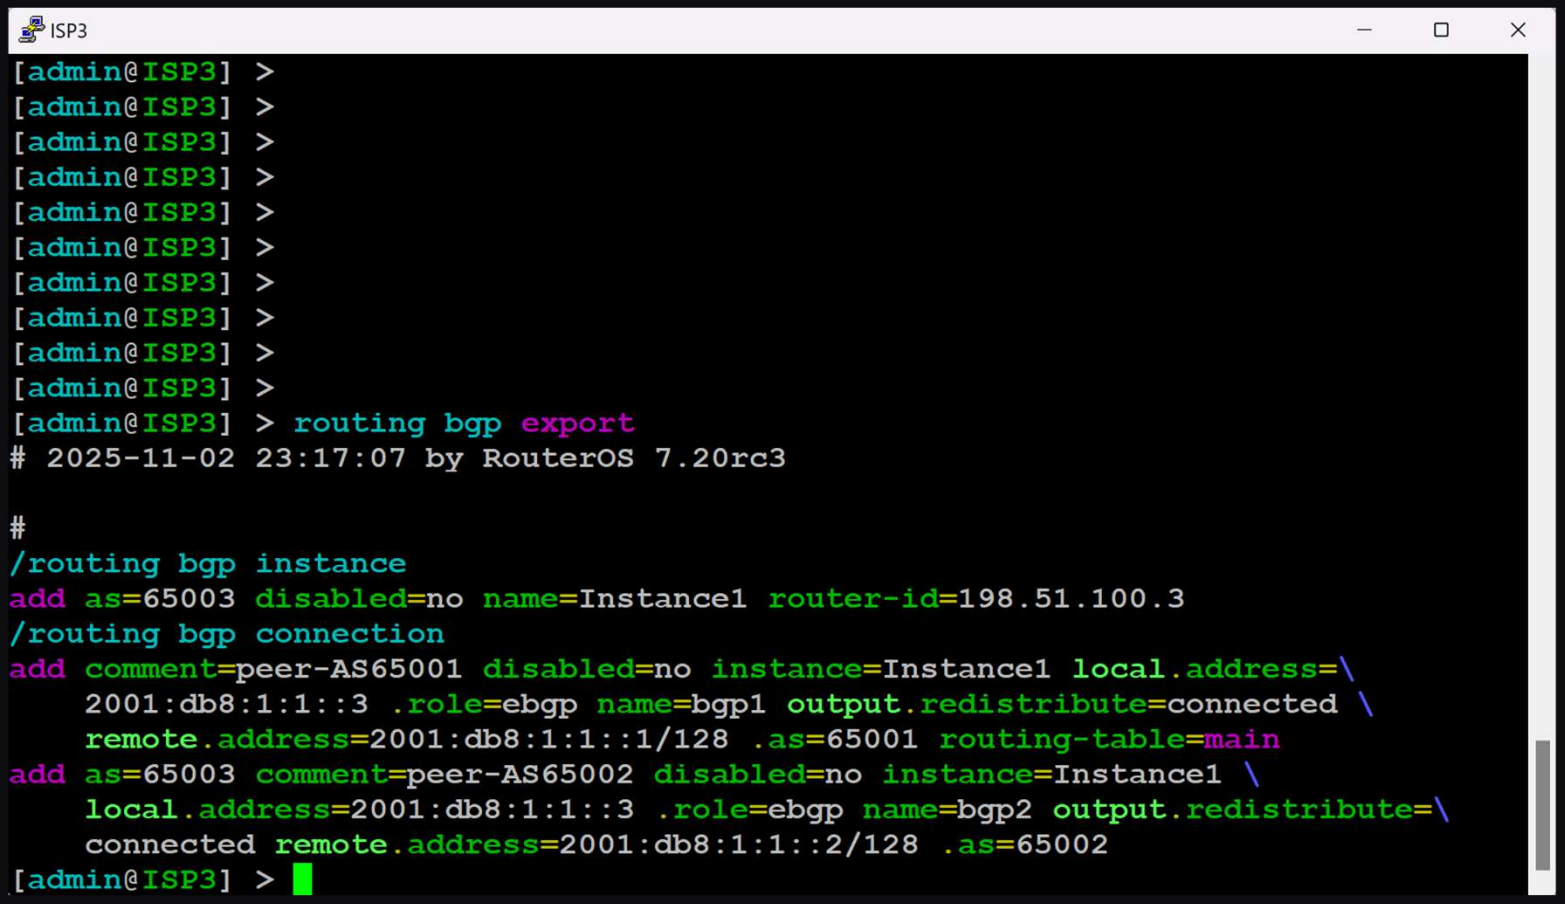

BGP session confirmation

ISP3

```
[admin@ISP3] >
[admin@ISP3] > routing bgp session print
Flags: E - established
 0 E name="bgp1-1" instance=Instance1
     remote.address=2001:db8:1:1::1 .as=65001 .id=198.51.100.1
     .capabilities=mp,rr,enhe,gr,as4 .afi=ip,ipv6 .messages=6 .bytes=398
     .eor=""
     local.address=2001:db8:1:1::3 .as=65003 .id=198.51.100.3
     .cluster-id=198.51.100.3 .capabilities=mp,rr,gr,as4 .afi=ip,ipv6
     .messages=6 .bytes=398 .eor=""
     output.procid=21
     input.procid=21 ebgp
     hold-time=3m keepalive-time=1m uptime=25s980ms
     last-started=2025-11-02 23:25:09 last-stopped=2025-11-02 23:25:03
     prefix-count=3

 1 E name="bgp2-1" instance=Instance1
     remote.address=2001:db8:1:1::2 .as=65002 .id=198.51.100.2
     .capabilities=mp,rr,enhe,gr,as4 .afi=ip,ipv6 .messages=5 .bytes=318
     .eor=""
     local.address=2001:db8:1:1::3 .as=65003 .id=198.51.100.3
     .cluster-id=198.51.100.3 .capabilities=mp,rr,gr,as4 .afi=ip,ipv6
     .messages=6 .bytes=398 .eor=""
     output.procid=20
     input.procid=20 ebgp
```

BGP next-hop validation check

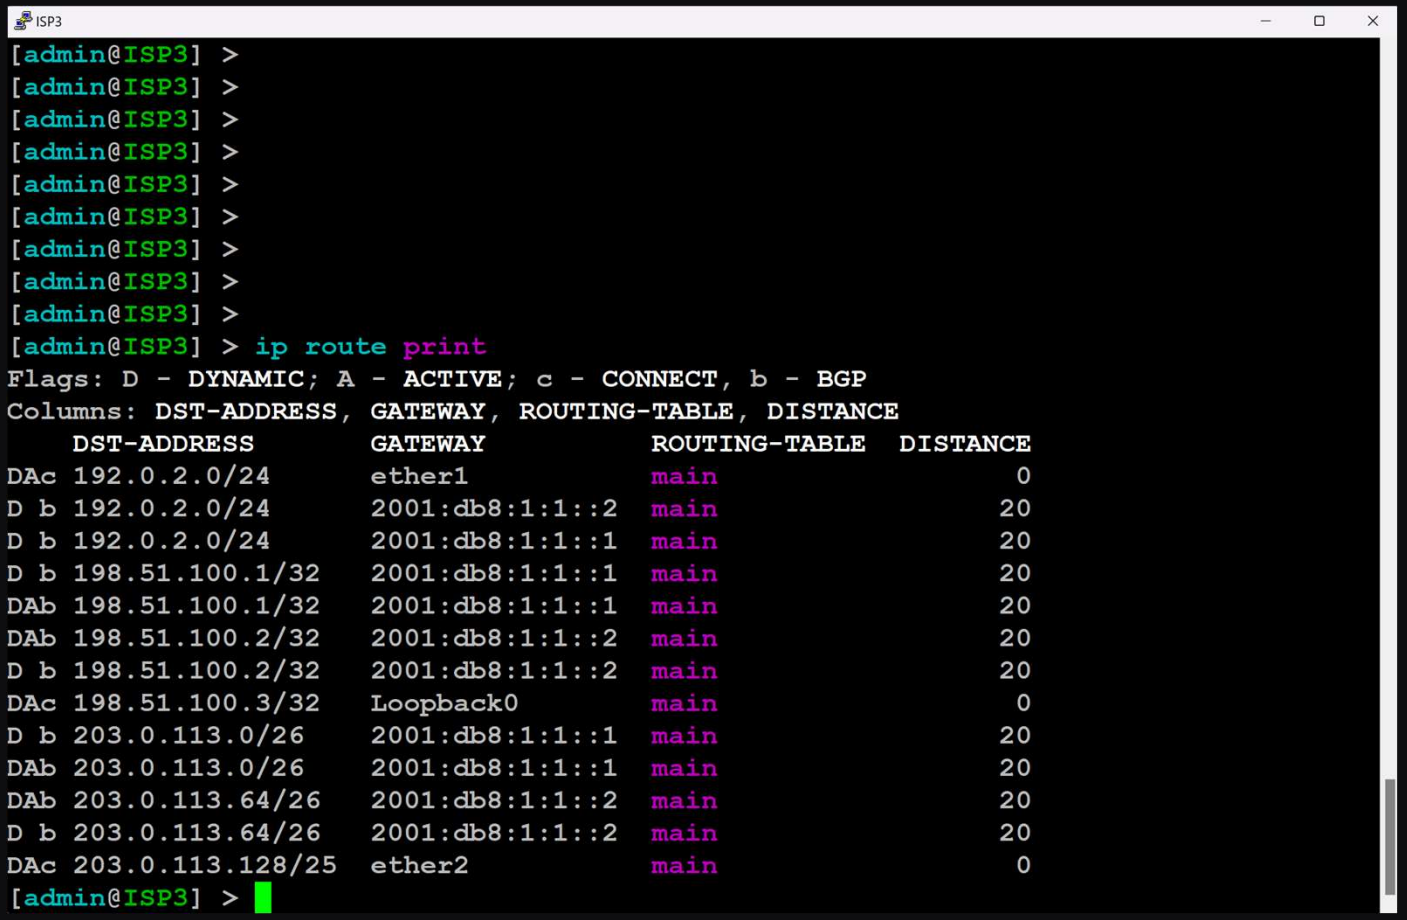

Customer's connection result

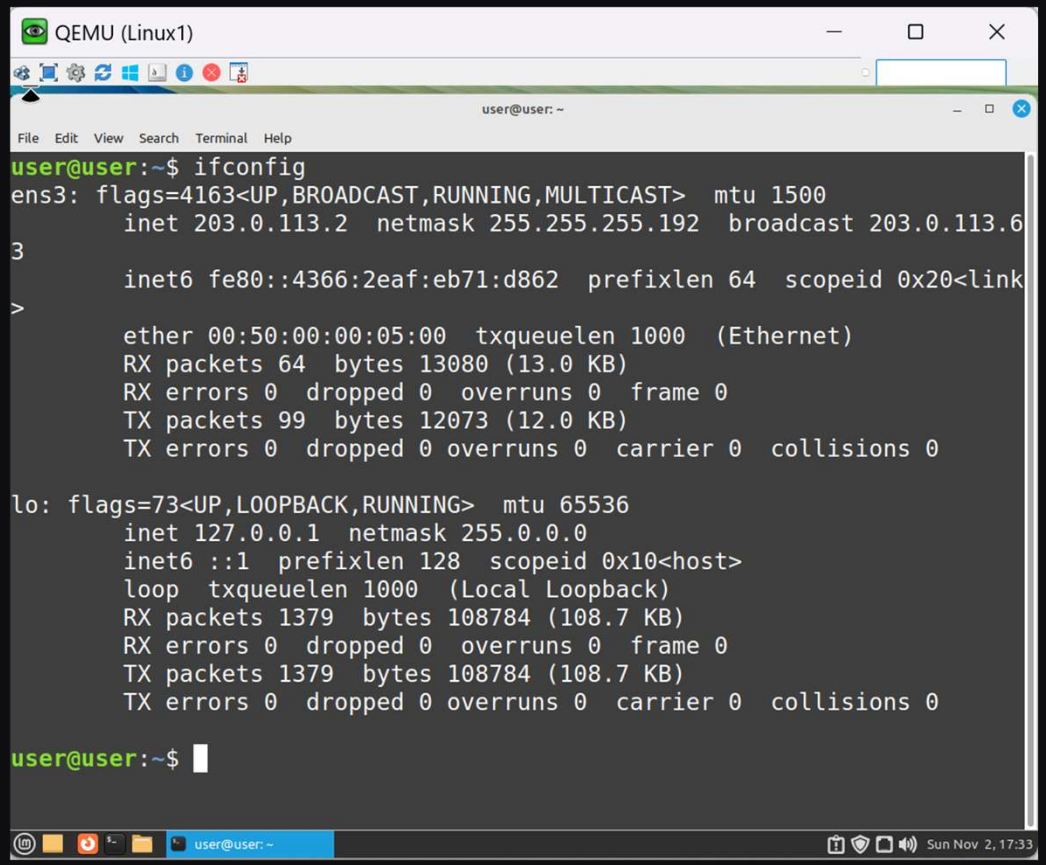

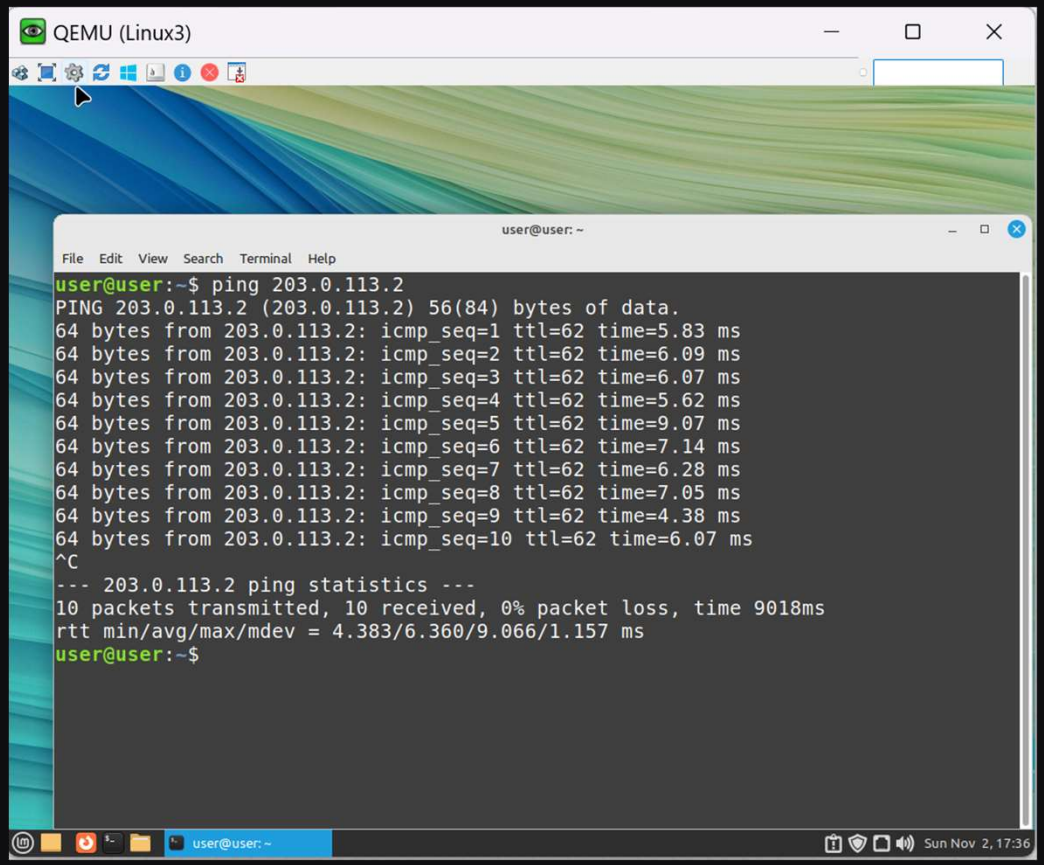

Conclusion: Why this Matters

- **Scalability:** Reduces dependence on IPv4 address space for the core network.
- **Operational Simplicity:** A single IPv6-only underlay network is easier to manage than a dual-stack one.
- **Future-Proofing:** Prepares the network for an IPv6-dominant future without sacrificing existing IPv4 services.