CONNECTING ENTERPRISE IPV6 ISLANDS

UK IPV6 COUNCIL, NOVEMBER 2023

WHO AM I?

Currently:

- One of the Lead Infrastructure Architects at the BBC

Previously:

- Principal Architect at an altnet
- Senior Lead Security Architect at a global aviation service provider
- Head of IT Operations at a UK university with global presence

Designing & Building IPv6 networks since circa 2008

SCENARIO (BASED ON A PREVIOUS WORKPLACE)

A company with multiple physical locations globally, three of them are core sites

Some sites have their own Internet connectivity, others are satellites connected via a WAN provided by a managed service provider

The WAN provider either won’t support IPv6, or wants to charge an extortionate fee to enable it

Free to enable IPv6 at sites in any order

EXISTING TOPOLOGY

Three core sites acting as regional Internet gateways

25 sites with local Internet breakout + WAN connectivity

Three sites with Internet access only, no WAN

65 sites with WAN connectivity only, no direct Internet access

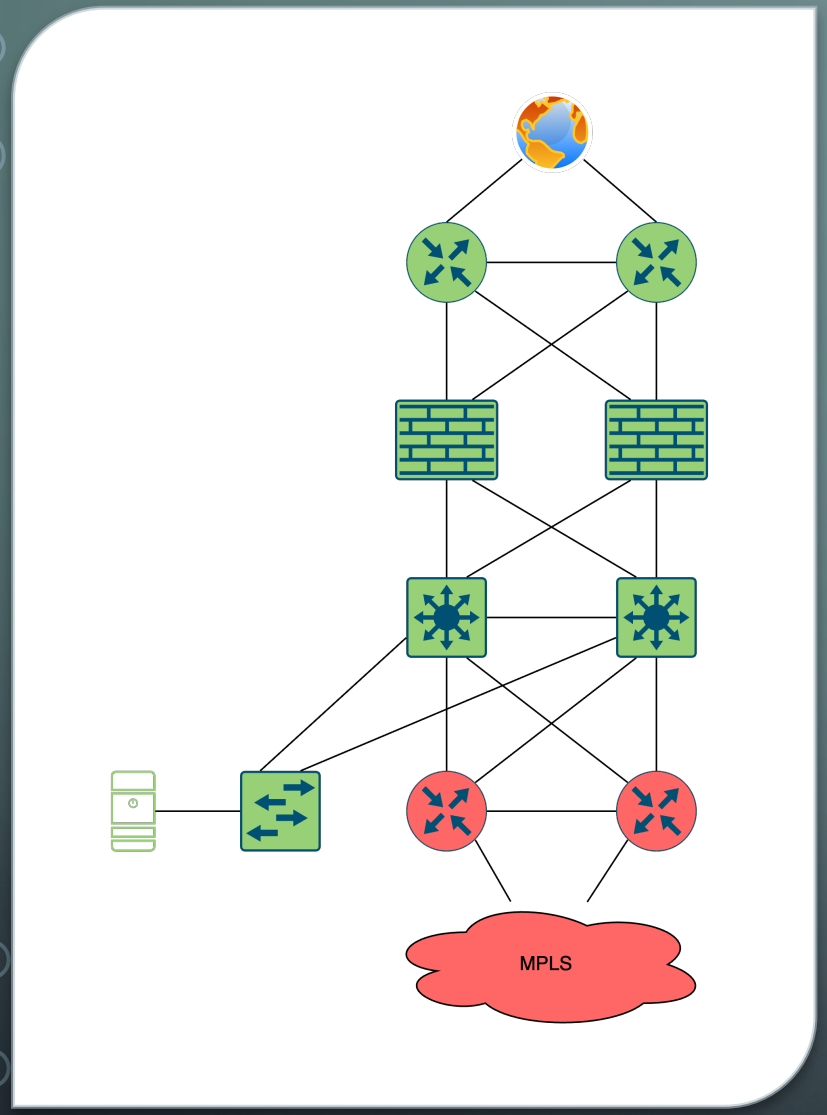

CORE SITES

IPv6 from ISPs

BGP at Internet edge

OSPFv3 through firewalls and on to L3 switches

Server segments configured for IPv6

VPN endpoints on firewalls configured for IPv6

JOINING THE CORE SITES

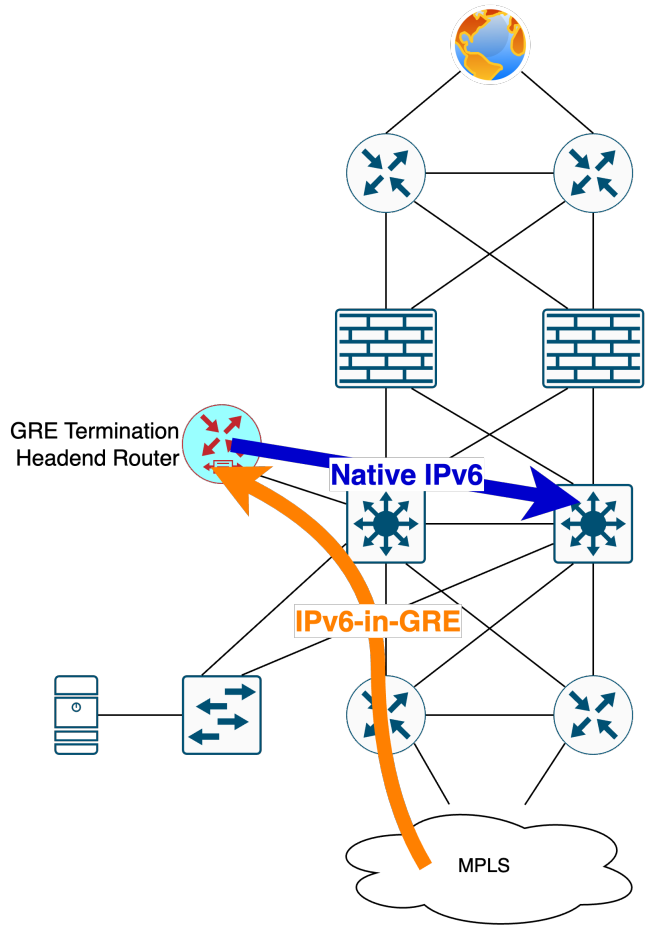

MPLS WAN is not IPv6 capable

Use GRE tunnels to transport v6 over v4

IS-IS for dynamic routing

Export (filtered) routes from OSPFv3 to IS-IS

Export routes from IS-IS to OSPFv3 as E2

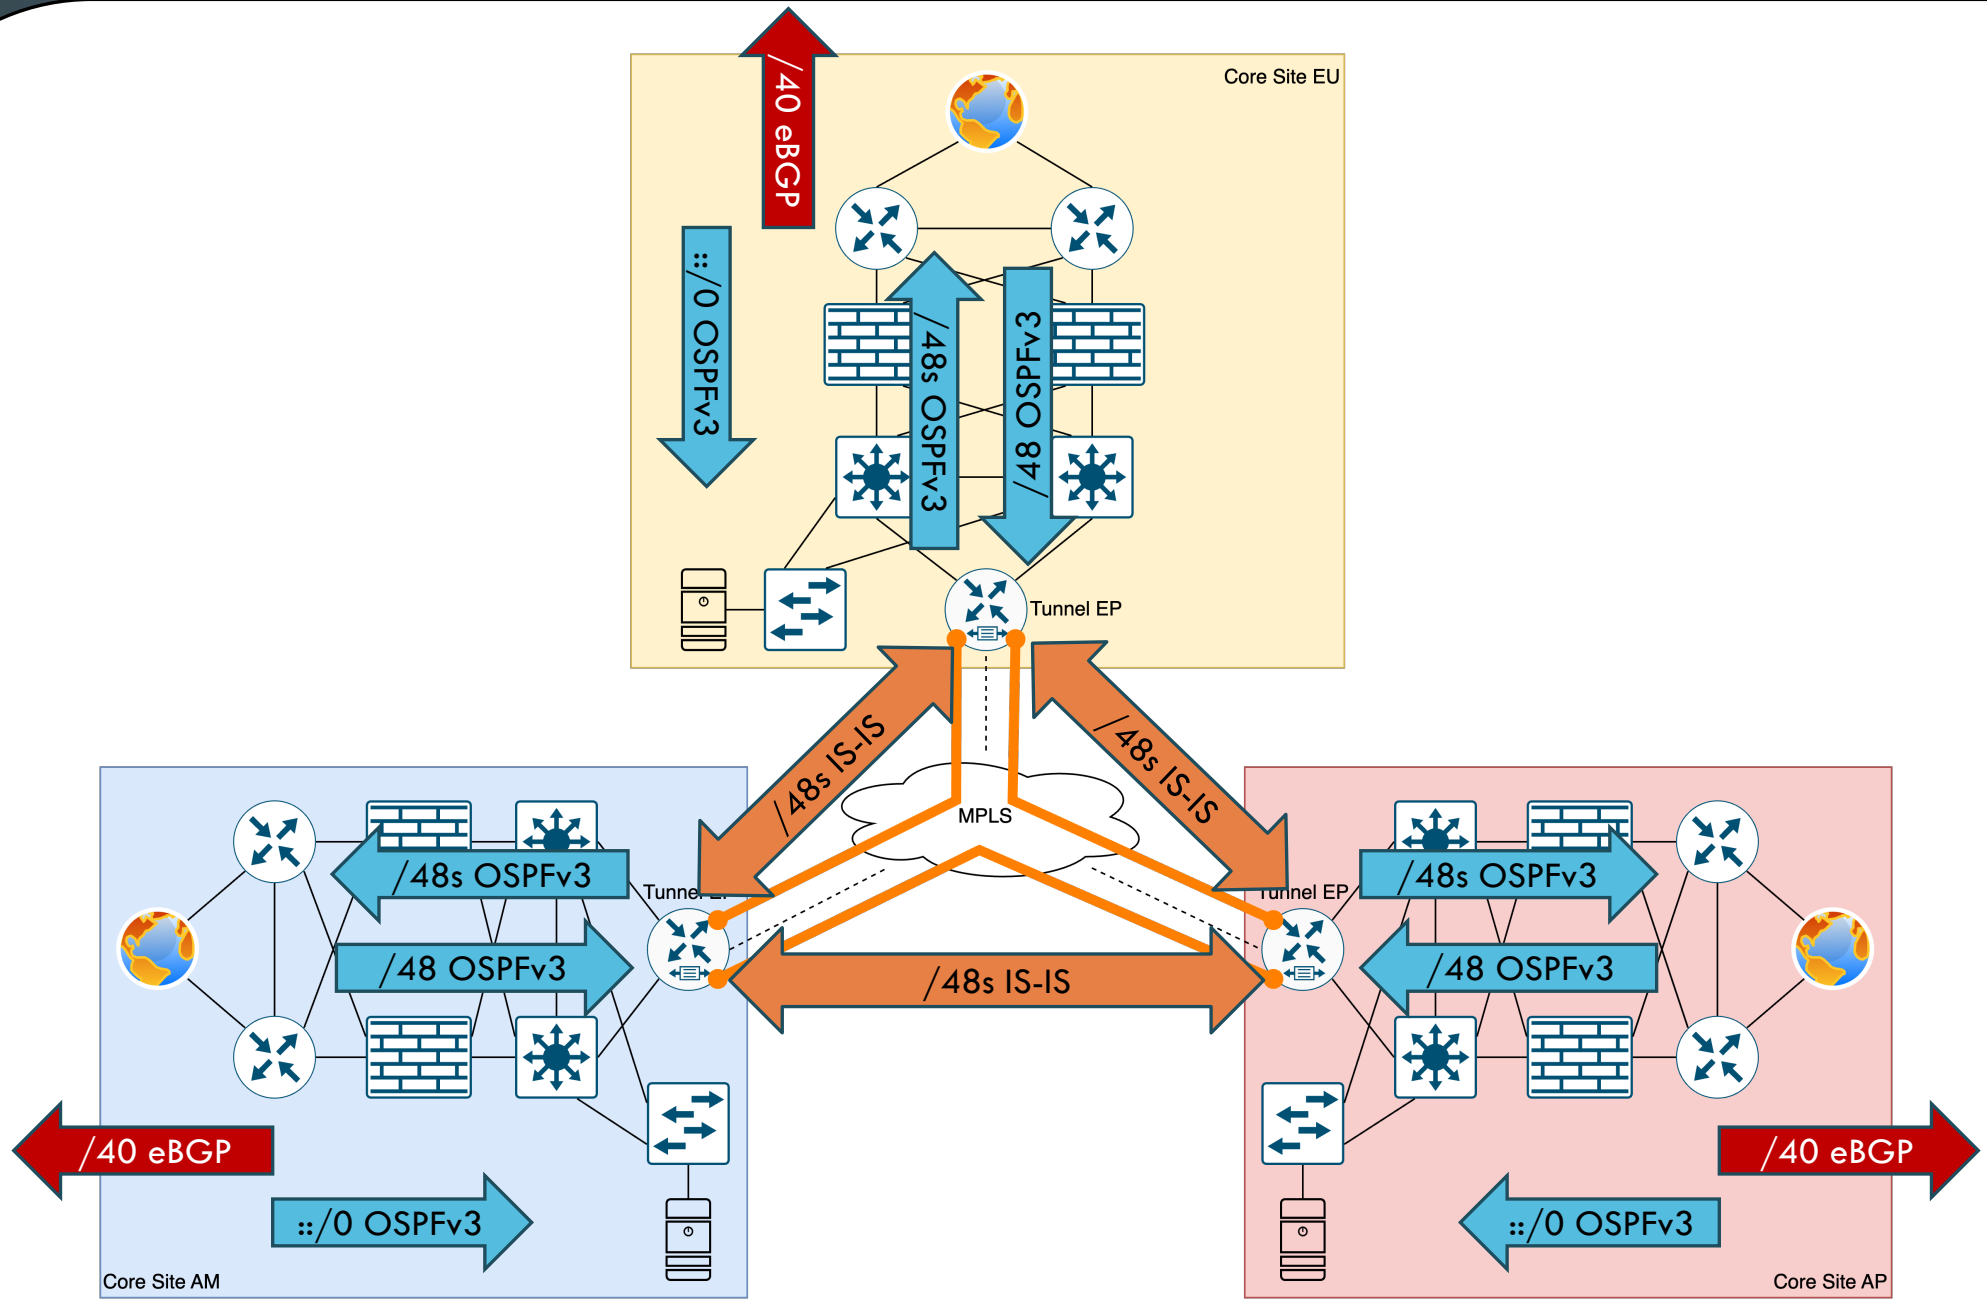
Core Site EU
/40 eBGP
::/0 OSPFv3
/48s OSPFv3
/48 OSPFv3
Tunnel EP
/48s IS-IS
/48s IS-IS
MPLS
/48s OSPFv3
/48 OSPFv3
Tunnel EP
/48s IS-IS
/48s OSPFv3
/48 OSPFv3
/40 eBGP
::/0 OSPFv3
/40 eBGP
::/0 OSPFv3
Core Site AM
Core Site AP

INTERNET-ONLY BRANCH

- Quick win!
- Add IPv6 to the existing IPSEC tunnels to the hub sites
- Use BGP across the IPSEC tunnels, AS path prepend for non-preferred tunnel

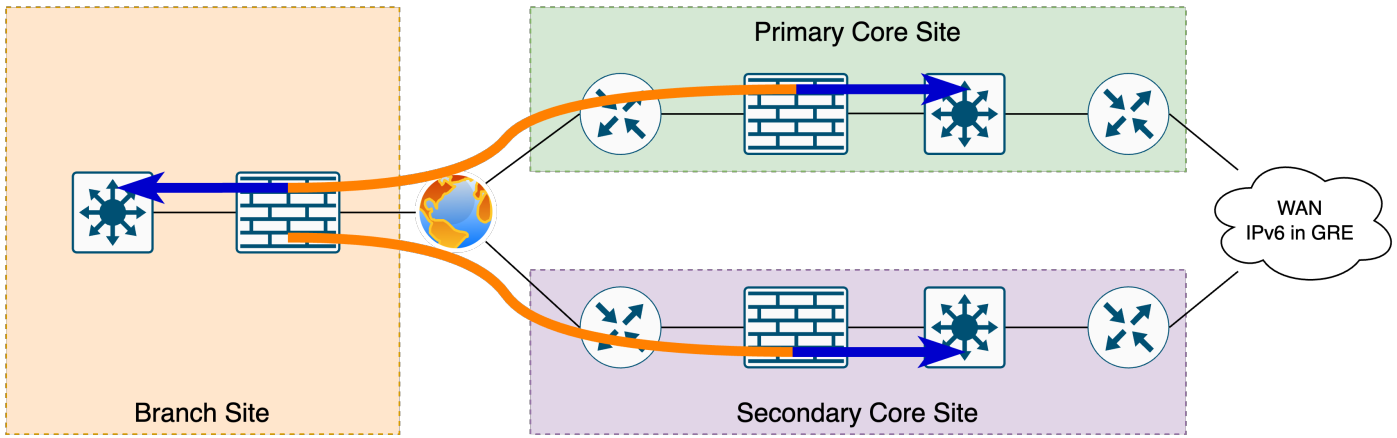

WAN-ONLY BRANCH

Challenges

Cannot change configuration on CE router

Dynamic routing:

- Site's L3 switch doesn't support BGP or IS-IS
- Don't want to run RIPng for security reasons
- Bad practice to extend OSPFv3 area across WAN links (high latency)

WAN-ONLY BRANCH

A solution

Create GRE tunnel from site's L3 switch to two regional core sites

Static outbound default route + BFD for rapid fault detection

Static inbound /48 route + BFD, exported into regional hub OSPFv3 area and IS-IS between regional hubs

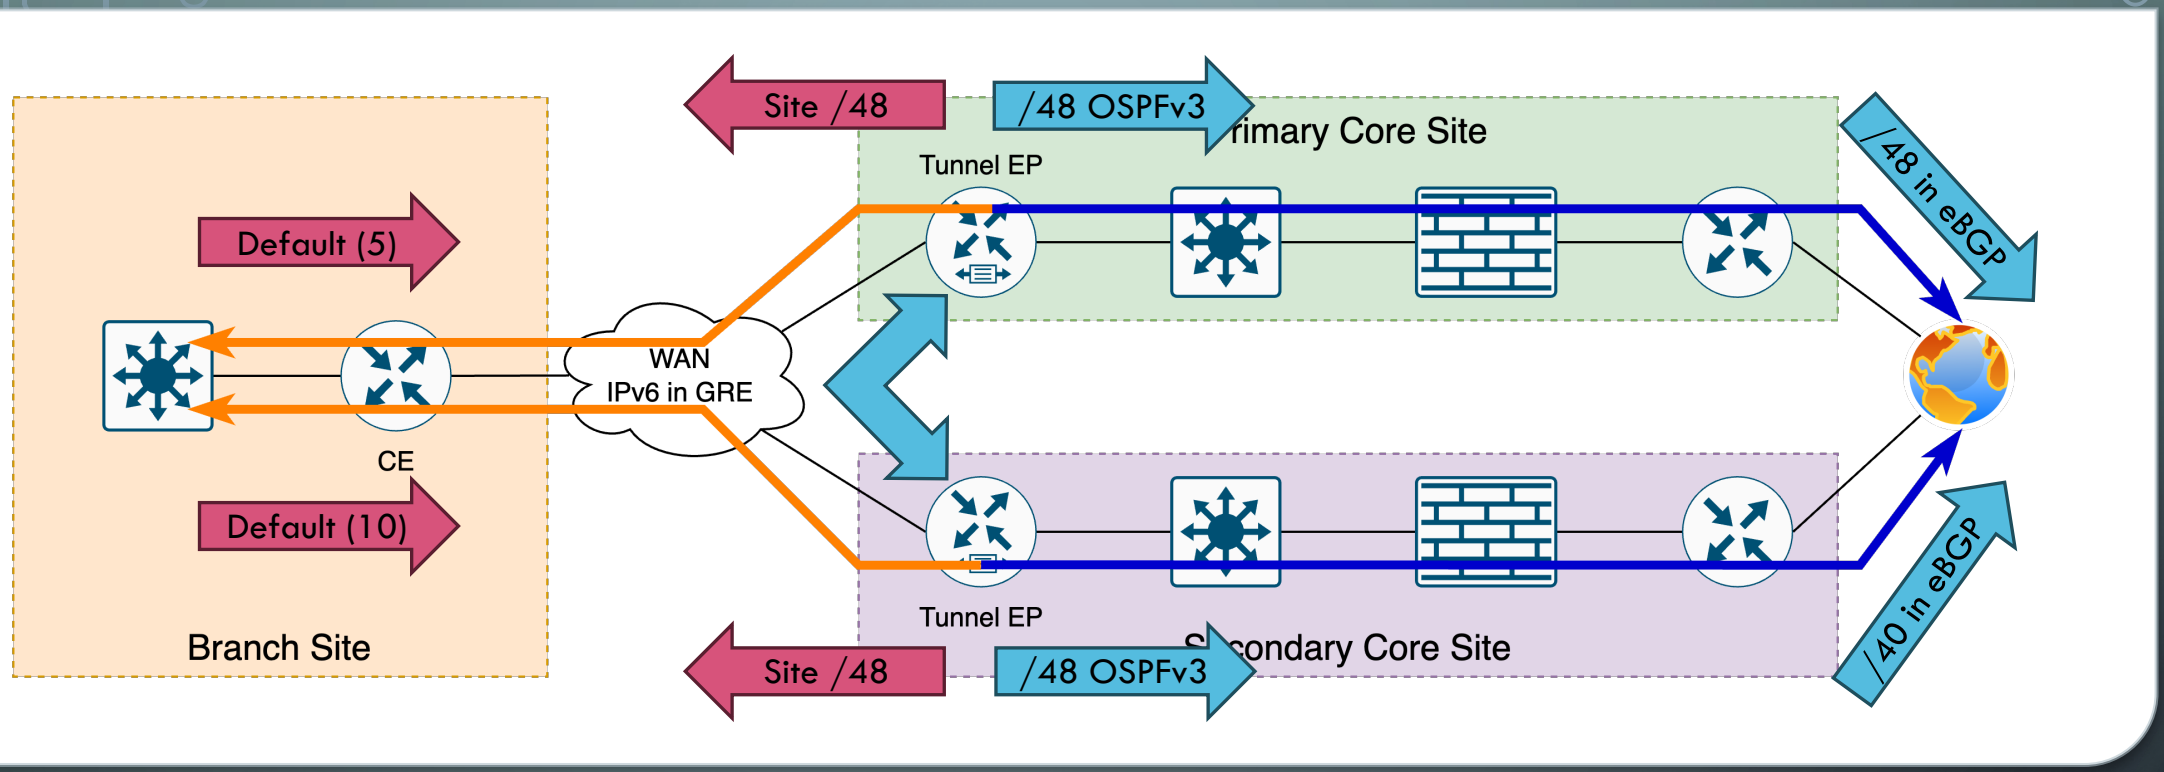

WAN-ONLY BRANCH

WAN & LOCAL INTERNET BREAKOUT

Challenges

Default route should be via local ISP with backup via WAN

Intra-company traffic should be via WAN

Cannot change configuration on CE router

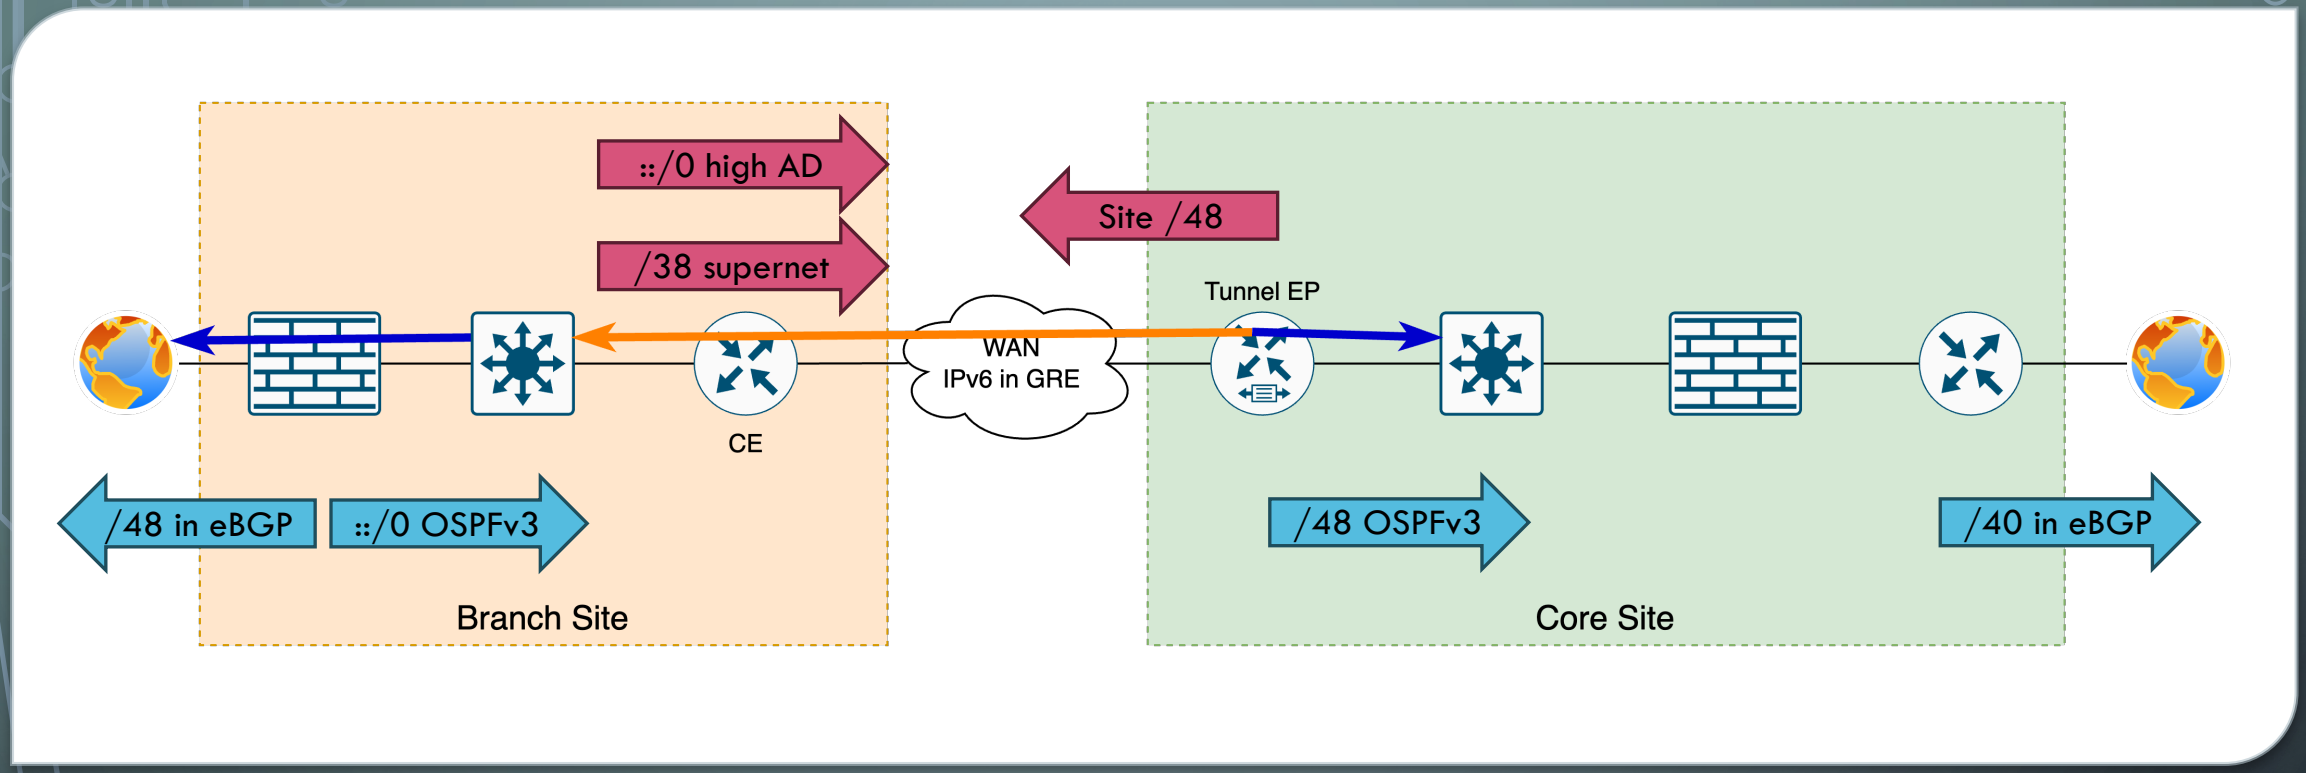

WAN & LOCAL INTERNET BREAKOUT

BRINGING IT ALL TOGETHER: INTERNET ACCESS

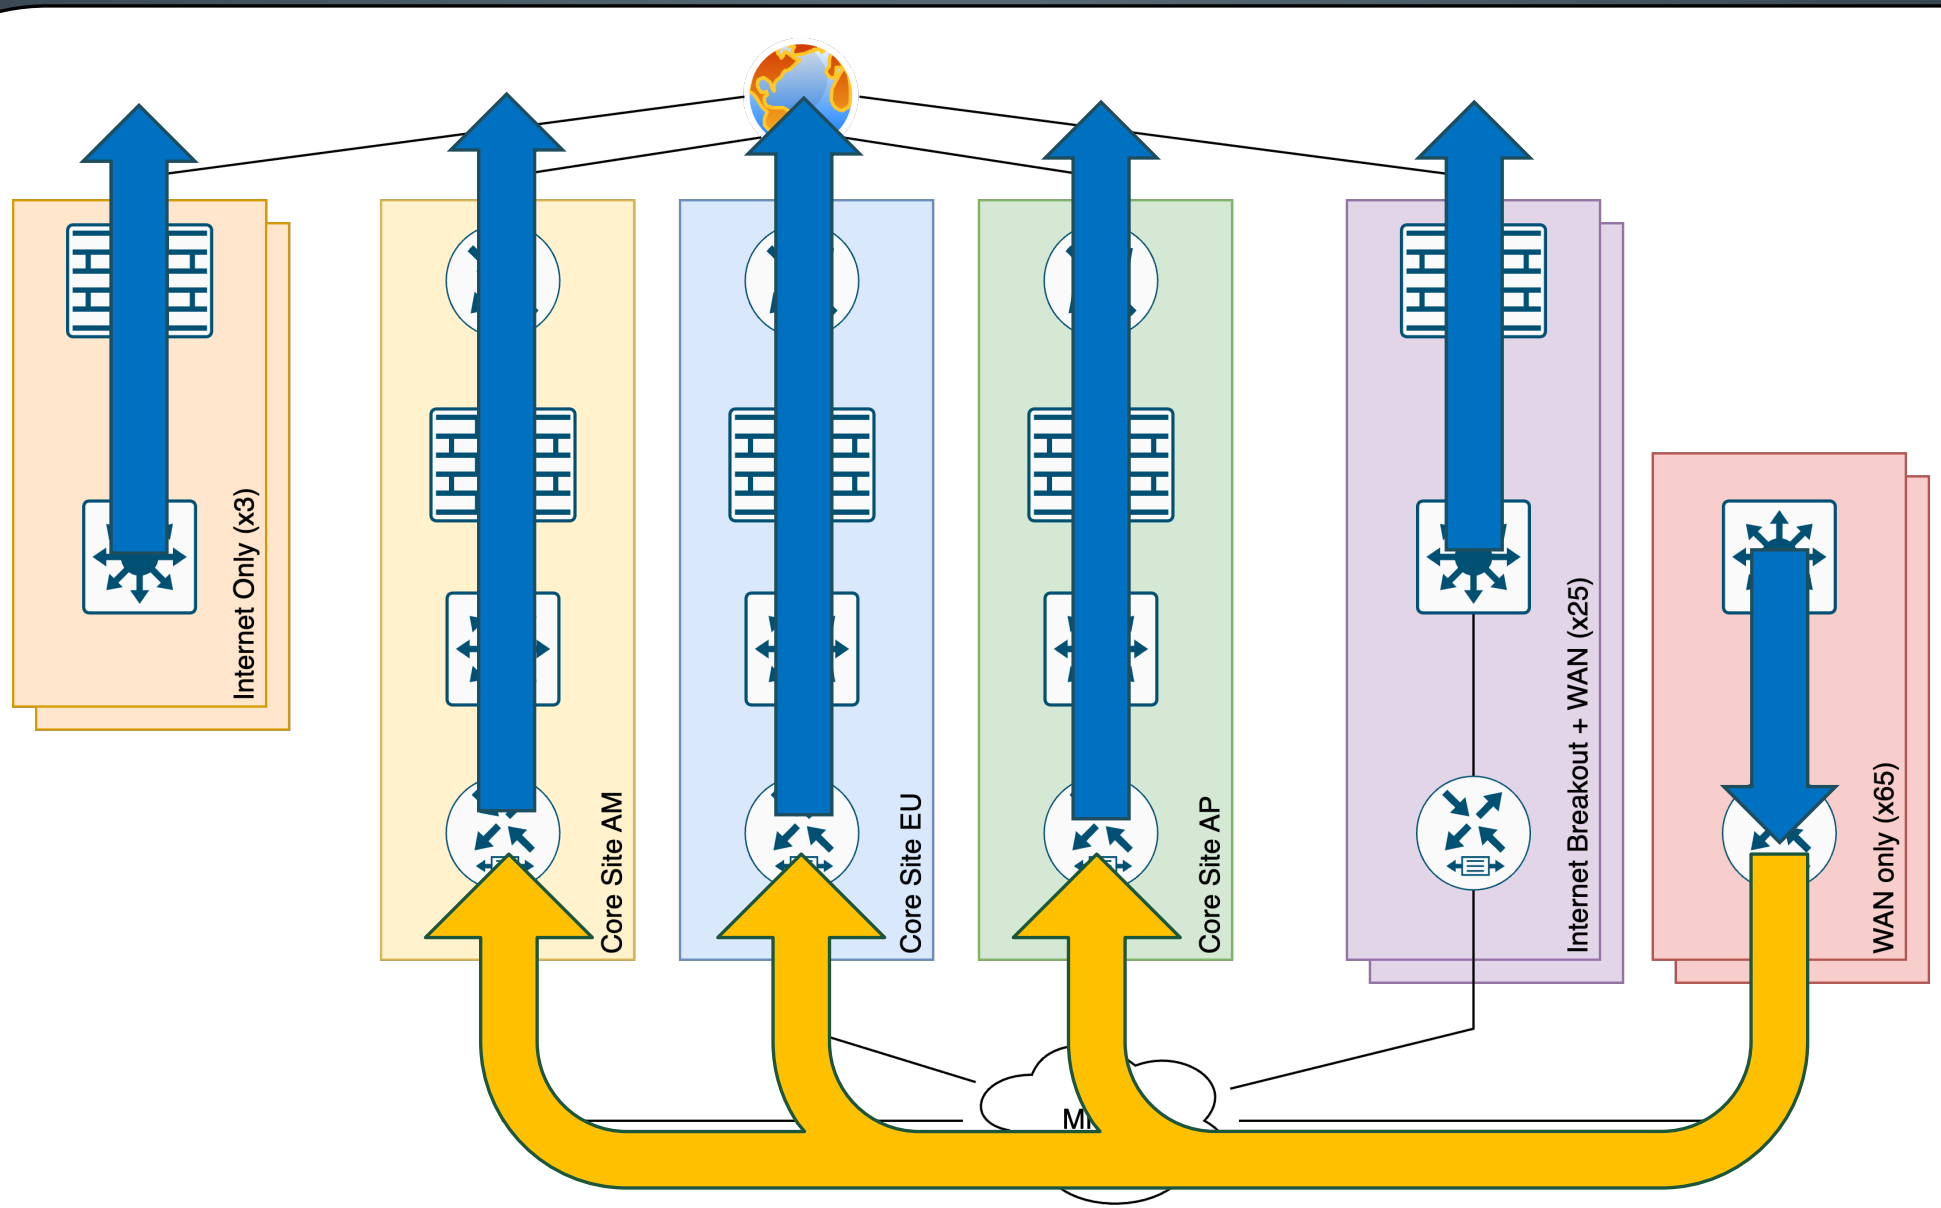

BRINGING IT ALL TOGETHER: CENTRAL RESOURCE

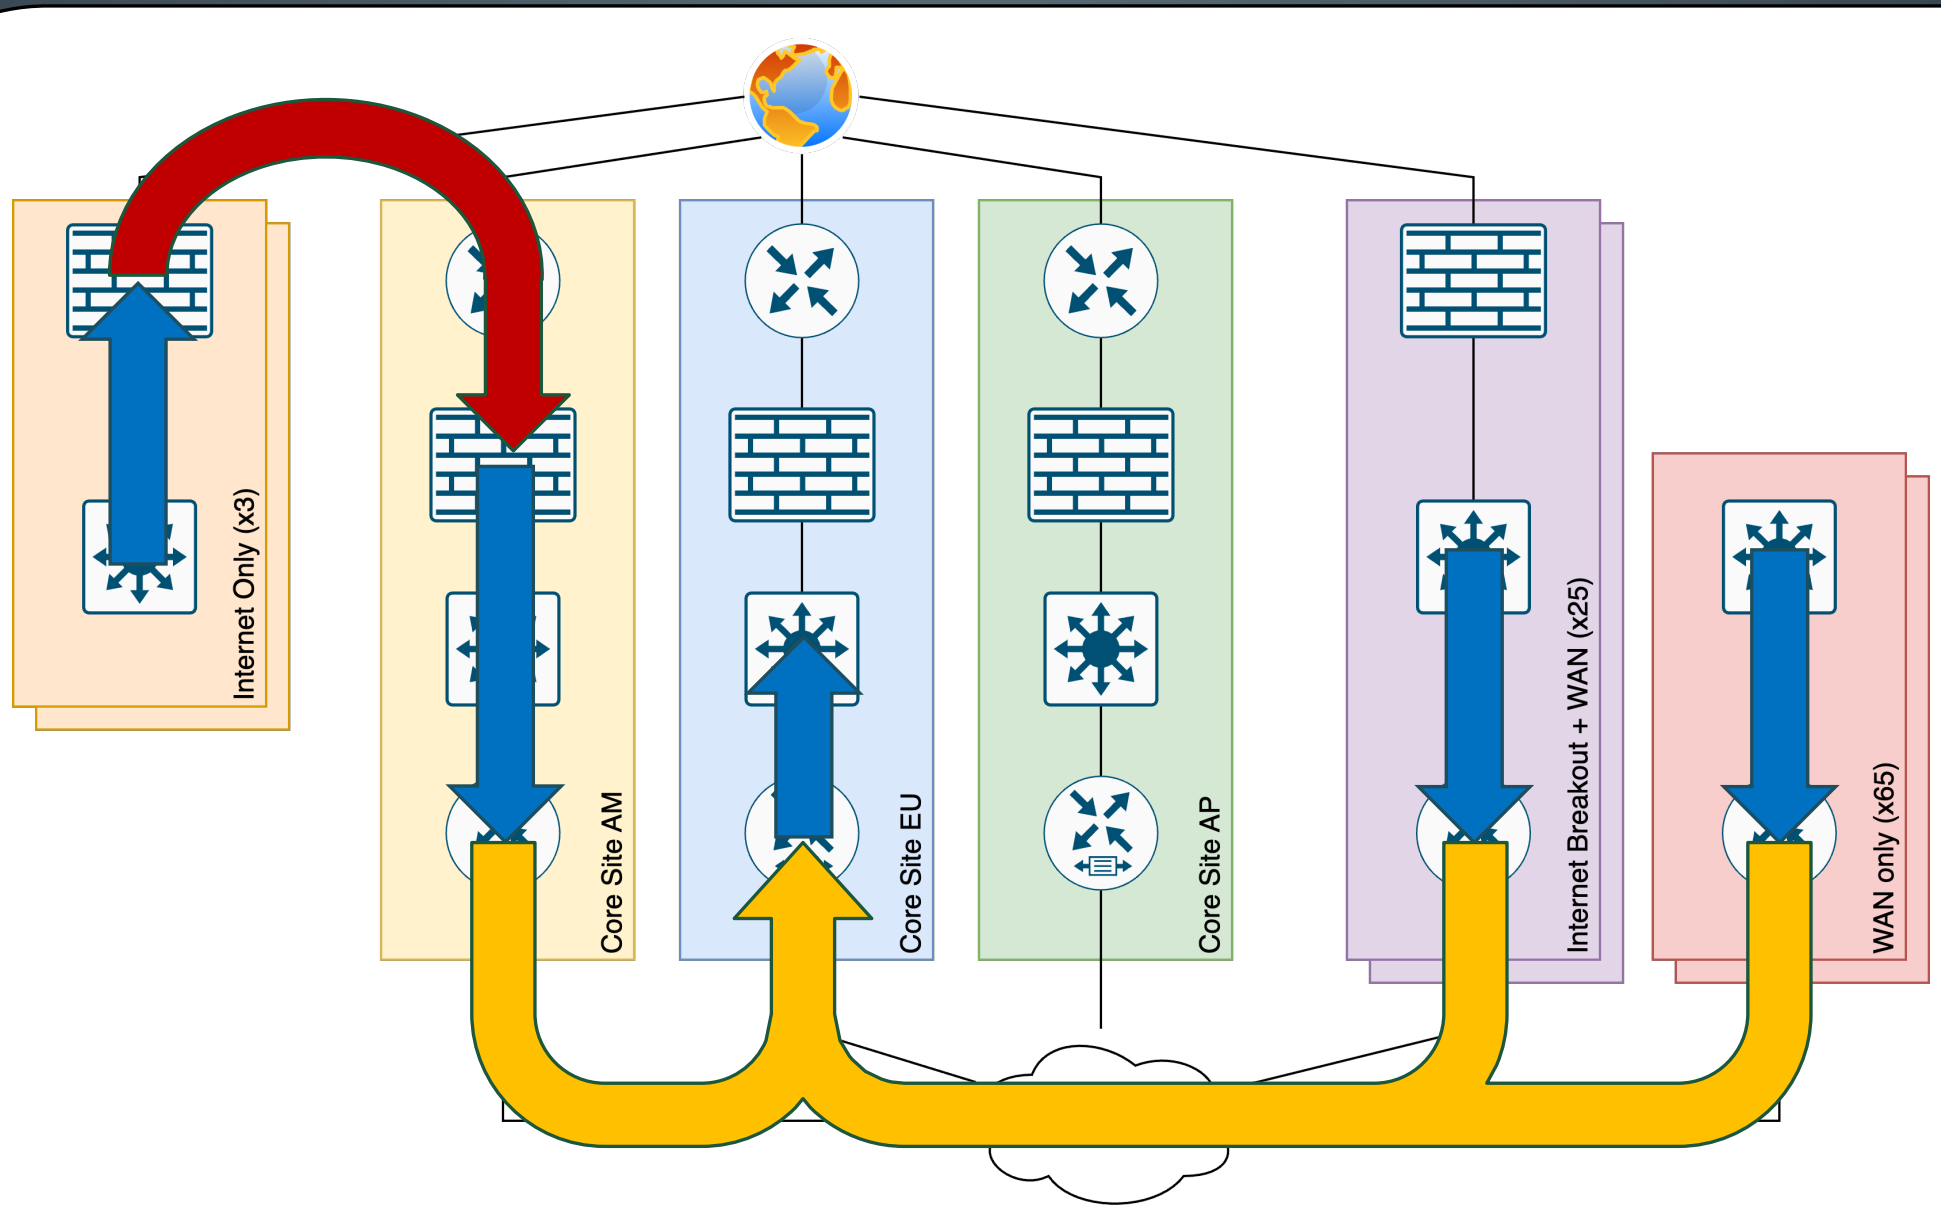

SUBOPTIMAL ROUTING

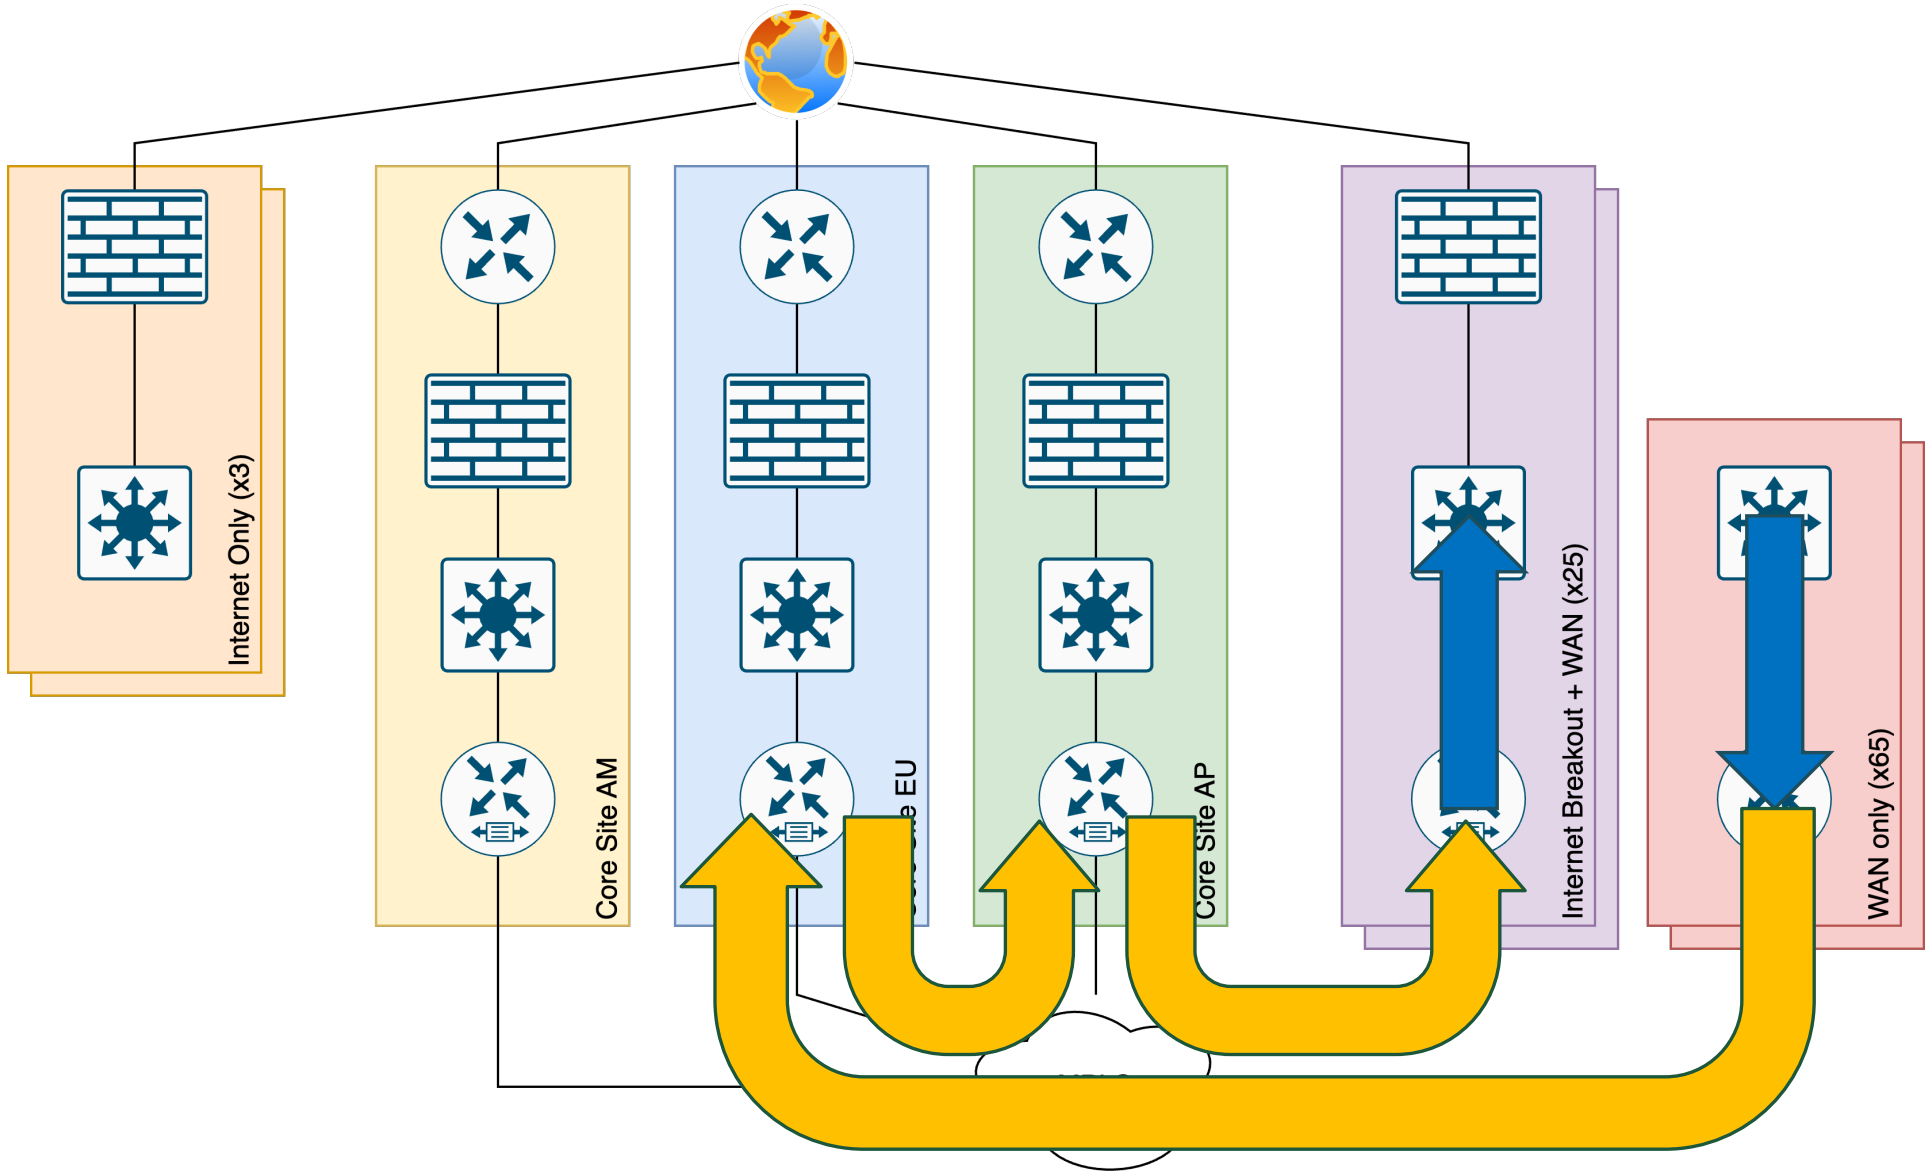

LESSONS LEARNED

MTU size

GRE has 24-byte overhead, so IPv6 MTU becomes 1476 bytes

Can rely on PMTU discovery as well as advertising optimal MTU in RAs

Suboptimal routing

Analysis showed that >90% of traffic was to / from the Internet, so no problem

Malicious attacks against WAN bandwidth

Mitigate with security controls at the perimeter

THANK YOU!