OPERATING MODERN INTERNET INFRASTRUCTURE

Network operations lessons from a hybrid fiber + wireless ISP

AFNOG Day · Africa Internet Summit 2026 · Nairobi, Kenya

Amin Dayekh
CEO / Network Operator · MegaMore Wireless Broadband Ltd

From infrastructure to operations

The modern operator must design, measure, protect and continuously improve the network.

1. **Operating context**
   What changed from legacy ISP/WISP operations
2. **Modern architecture**
   Access, POP, core, transit, peering, cloud and content
3. **Hybrid access reality**
   FTTH + PTP/PTMP + enterprise links
4. **Operational discipline**
   Monitoring, NOC workflow, power, security, automation
5. **Readiness checklist**
   What any operator must build today

Core Network

A network is no longer a chain of links. It is an operating system for connectivity: routing, power, content, cloud, security, measurement and customer experience must be designed as one system.

Why this matters at AFNOG Day

The summit theme is resilience; network operations is where resilience either exists or fails.

AIS 2026 context

Theme: “A Resilient Internet Ecosystem for an Innovative Digital Africa”

AfNOG Day brings the operator community into the same room: engineers, ISPs, IXPs, researchers, content networks and institutions.

Operator question

Can the network absorb failure, growth, attacks, power loss, fiber cuts, and customer demand without becoming a daily emergency?

Answer

Only when engineering, operations, documentation, monitoring and governance are treated as one discipline.

Source: AIS’26 public information; references at end.

OPERATOR LENS

MegaMore as a practical case study

A hybrid operator view from a secondary-city African market: fiber, wireless, IPv6, transit, customer operations and field realities.

Hybrid access footprint

- Metro fiber and FTTH/enterprise fiber
- PTP backhaul and enterprise links
- PTMP last-mile broadband where fiber economics are not yet mature
- Capacity planning tied to customer use, not marketing claims

Modernization path

- Earlier flexible MikroTik radio-heavy deployments
- Migration toward higher-capacity Ubiquiti/airFiber-class wireless where appropriate
- Better POP discipline, monitoring and escalation
- IPv6 readiness and operational deployment

Operating environment

- Power instability affects topology
- Site security is part of availability
- Middle-mile and IP transit capacity constrain growth
- Local peering/content changes customer experience

Specific operational constraints and the engineering response to each one.

THE SHIFT

Old ISP methods no longer fit the traffic reality

The old model was link-centric. The modern model is service-centric, measured and automated.

Old operating model

- Single upstream; little route policy
- Flat bridge networks and ad-hoc NAT
- Manual configuration without version control
- Radio link uptime treated as service quality
- Reactive support: customer reports before NOC sees
- No IPv6, weak documentation

Modern operating model

- Dual-stack routed network with policy
- Multiple upstreams, IX, caches, content paths
- Telemetry: SNMP, flow, syslog, synthetic probes
- Capacity planning, latency, loss and jitter
- Incident workflow with postmortems and runbooks
- Security and power designed into the topology

SECTION 1

Architecture of a modern operator network

From access to content: every layer must be observable, documented and recoverable.

REFERENCE ARCHITECTURE

The modern ISP is a system of systems

Core, access, content, cloud, transit and operations must be designed together.

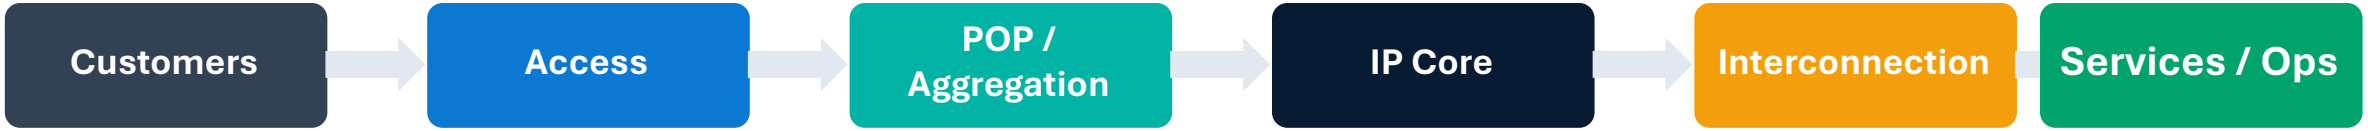

Residential

SME
Enterprise
Campus

FTTH / GPON

PTP / PTMP
CPE / Wi-Fi
Enterprise radio

OLT / switches

ODF / patching
Backhaul
Power + OOB

BGP / OSPF

CGNAT where needed
IPv6 dual-stack
QoS / policy

IP transit

IX peering
CDNs / caches
Cloud on-ramps

NMS / IPAM

RADIUS / DNS
Logs / flow
Ticketing

Design rule: every service path must have a technical owner, a monitoring point, a documentation record and a recovery procedure.

ACCESS LAYER

Hybrid access is not a weakness; unmanaged hybridity is

A modern network may combine FTTH, PTP, PTMP and enterprise links—provided each technology has clear engineering rules.

Fiber access

GPON/XGS-PON design, splitter ratios, ODF discipline, OTDR baselines, GIS/as-built records, restoration spares and customer ONT lifecycle.

Wireless access

PTP and PTMP require link budgets, spectrum scans, alignment records, sector capacity planning, CPE signal thresholds and noise monitoring.

Enterprise links

Dedicated radio/fiber needs SLA thinking: redundancy, latency, jitter, QoS marking, managed CPE and maintenance windows.

Residential broadband

The access layer must be engineered for peak-hour behavior, Wi-Fi support realities, CPE failure rates and clear customer communication.

From flexible radio networks to capacity-engineered wireless

The vendor is less important than the operating discipline: spectrum, throughput, latency, redundancy and visibility.

Earlier stage: flexible deployment

- Low-cost, quick activation
- MikroTik radios useful for early growth and flexible field deployment
- Often driven by immediate customer demand
- Risk: networks become a collection of links instead of an engineered access layer

Scaling stage: engineered wireless

- Higher-capacity PTP backhaul and cleaner PTMP sectors
- Ubiquiti/airFiber-class links where capacity and visibility justify upgrade
- Sector loading, modulation, interference and CPE quality tracked continuously
- Replace “it connects” with “it meets service objectives”

Operating rule

Wireless must be treated as a scarce spectrum asset, not as an infinite access medium. Each link needs baseline throughput, loss, latency, signal, modulation and fallback plan.

Fiber deployment: what it really entails

Fiber is not just trenching. It is civil work, optical engineering, documentation, security and maintenance discipline.

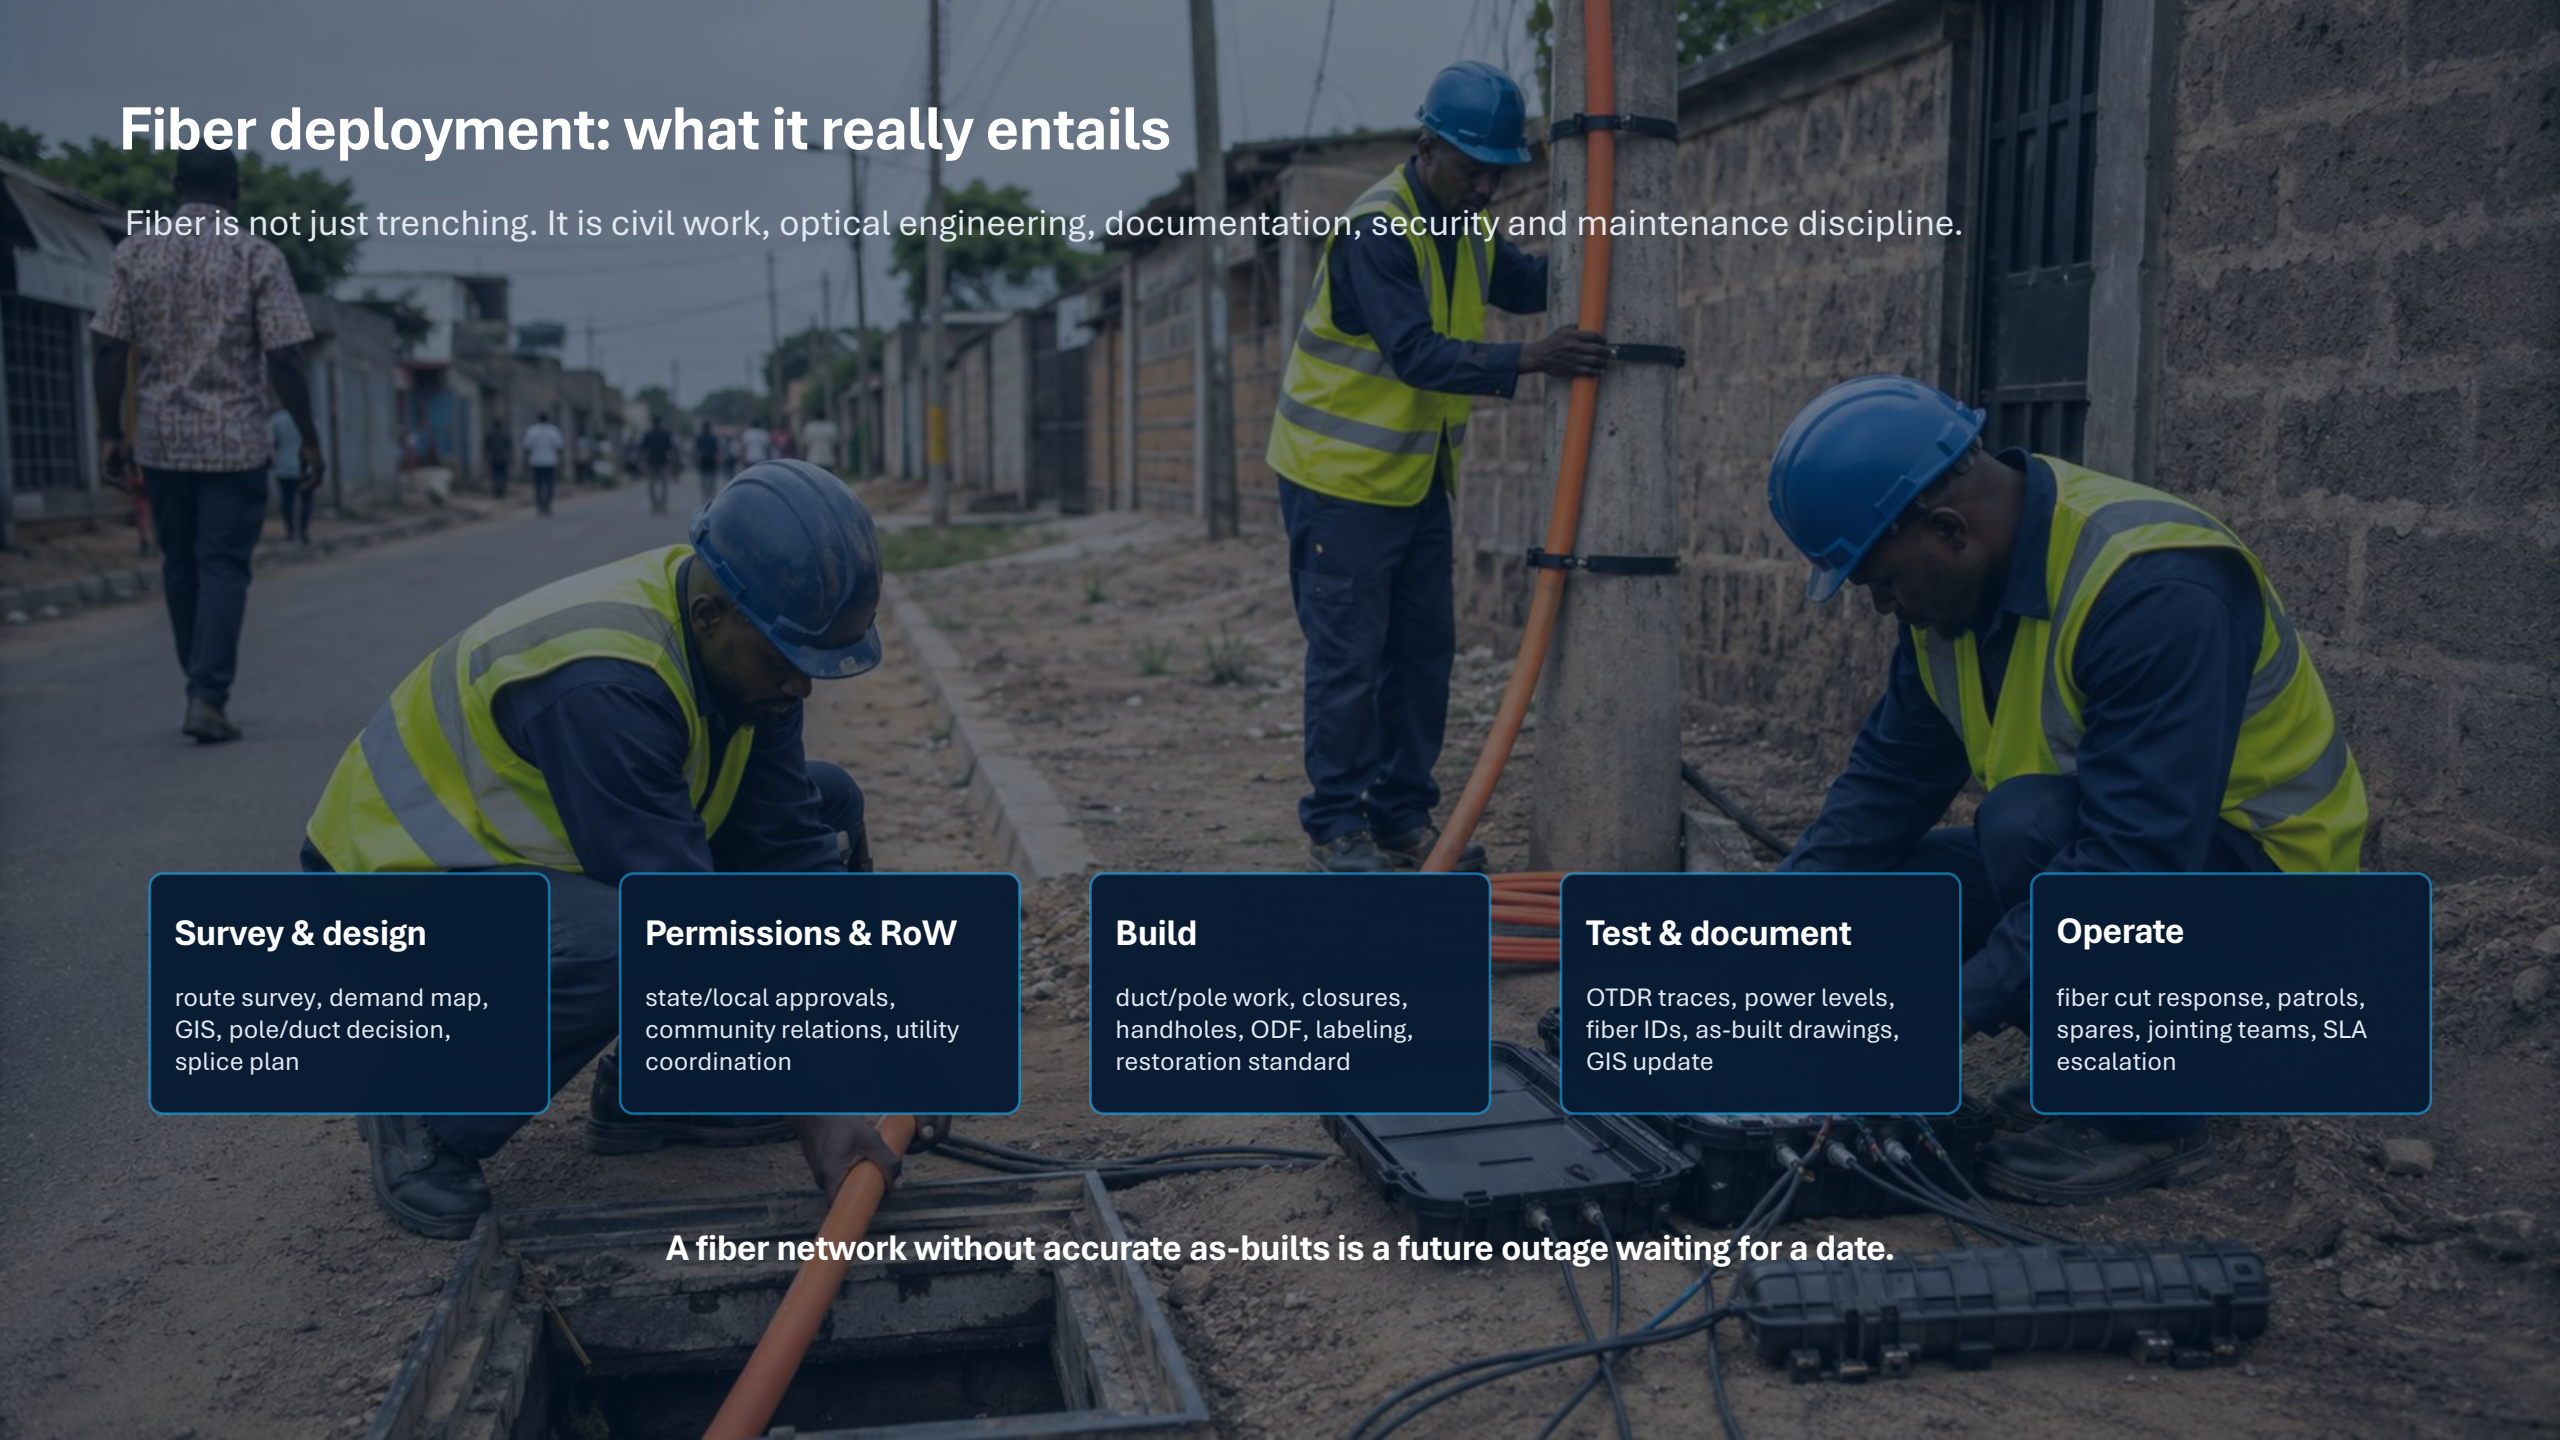

Survey & design

route survey, demand map, GIS, pole/duct decision, splice plan

Permissions & RoW

state/local approvals, community relations, utility coordination

Build

duct/pole work, closures, handholes, ODF, labeling, restoration standard

Test & document

OTDR traces, power levels, fiber IDs, as-built drawings, GIS update

Operate

fiber cut response, patrols, spares, jointing teams, SLA escalation

A fiber network without accurate as-builts is a future outage waiting for a date.

POP DESIGN

Every POP is a miniature data center

A POP is not a cabinet with a switch. It is a controlled service point with power, routing, OOB access, labeling and recovery procedures.

Rack & physical order

ODF, cable management, labels, patch discipline, dust control, access control

Power system

grid, inverter/solar, UPS, DC plant, battery monitoring, generator process

Transport

fiber backhaul, PTP backup, LAG/ECMP where available, capacity headroom

Management

OOB modem or secondary path, device backups, local console plan, remote power

Monitoring

temperature, power, interface errors, optical levels, radio quality, latency probes

Security

locks, cameras/guards where needed, tamper evidence, site access log

The core network must be boring, predictable and policy-driven

Exciting networks fail. A good core is simple to reason about under pressure.

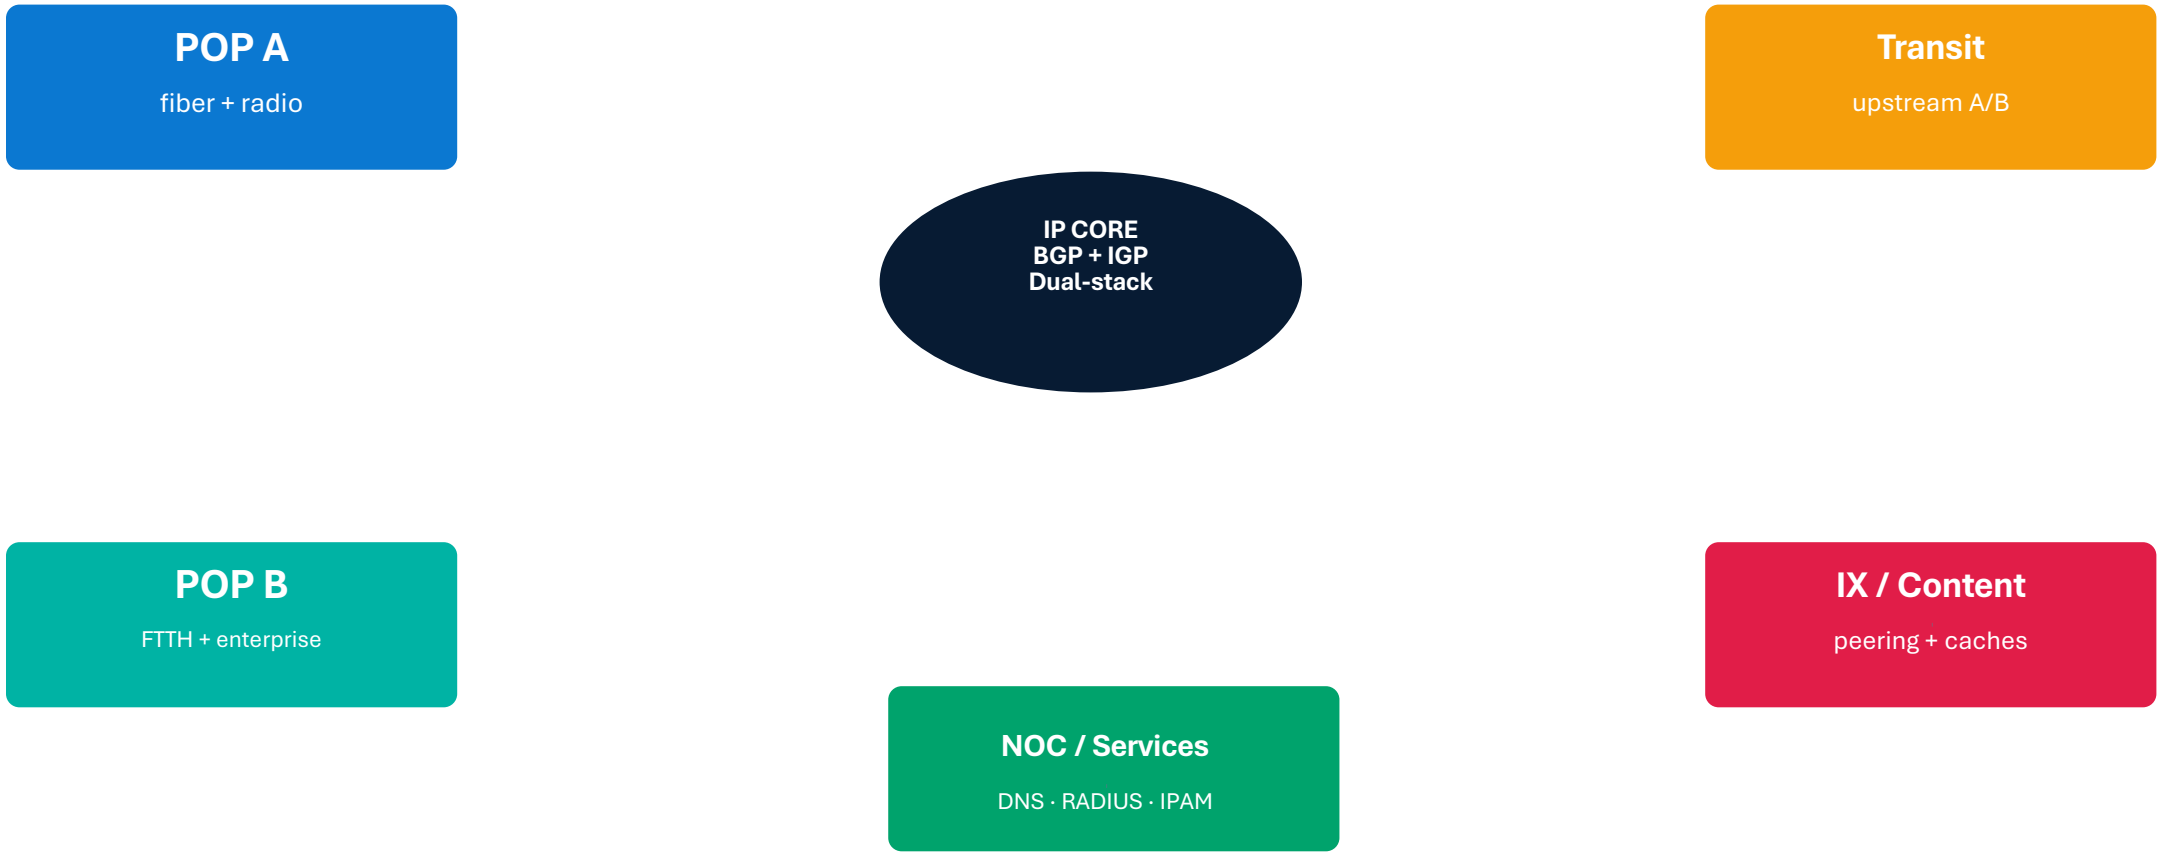

Core baseline

Clear ASN/prefix policy · BGP communities · route filters · IPv6 parity · device backups · config standards · change control · capacity headroom

MIDDLE MILE AND TRANSIT

The hardest bottleneck is often not the last mile

Middle-mile capacity and IP transit availability determine whether local access investments can become real service quality.

Middle-mile problem

- Metro networks can grow faster than national/backhaul capacity
- Secondary cities often depend on distant aggregation points
- A clean access layer cannot compensate for congested middle mile
- Capacity must be planned before demand becomes complaint volume

Transit design

- More than one upstream where economics allow
- BGP local preference and communities
- Provider diversity: physical path and commercial entity
- Blackhole/RTBH capability for large attacks
- Clear IPv6 transit, not IPv4-only scaling

Capacity discipline

- Track p95, peak-hour saturation and packet loss
- Upgrade before customer experience collapses
- Separate “sold speed” from aggregate engineering model
- Measure actual paths to content, not only speed tests

INTERCONNECTION

IXPs and content are part of network operations

Peering is not public relations. It is latency control, cost control, failure-domain control and customer-experience control.

IX value

Local traffic exchange reduces unnecessary tromboning, improves latency and creates a platform for local digital ecosystems.

CDN/cache value

Content closer to users reduces transit pressure and improves streaming, update, education and cloud application experience.

Operator requirement

Maintain PeeringDB records, route-server policy, IRR/RPKI objects, IX monitoring and traffic graphs.

Strategic question

Where should content live: only in coastal hubs, or also in regional exchange points and secondary cities?

Regional IX thinking: localizing traffic is an operational decision before it becomes a policy slogan.

IPV6 READINESS

IPv6 is not a future project; it is an operational requirement

A network is not modern if IPv6 is only present in a policy document or a lab prefix.

Network layer

- RIR allocation and prefix plan
- BGP announcements and route objects
- IPv6 transit and peering
- Dual-stack core, aggregation and access
- Firewall/ACL parity
- Monitoring over IPv6

Customer layer

- CPE support and firmware lifecycle
- DHCPv6-PD / SLAAC / RA design
- Prefix delegation policy
- Helpdesk scripts and customer education
- AAA/accounting compatibility
- IPv6 speed/latency tests

Operational layer

- DNS AAAA readiness
- Abuse handling and logs
- NMS dashboards by protocol family
- PeeringDB/IRR/RPKI consistency
- Security baselines: ICMPv6, RA guard, ND behavior
- Measure real IPv6 traffic share

IPv6 readiness means the customer can use it, the NOC can see it, and the operator can troubleshoot it.

ROUTING SECURITY

A modern network must not pollute the global routing table

Routing hygiene is no longer optional: it is part of operational legitimacy.

RPKI / ROA

Create ROAs for originated prefixes; validate upstream and peer routes where feasible.

IRR hygiene

Maintain route/route6 objects and AS-SETs; align PeeringDB with reality.

Filters

Filter customers, peers and transits; never accept or announce what should not exist.

Communities

Use communities for traffic engineering, blackholing, localpref and operational clarity.

DDoS response

RTBH/scrubbing options, ACL templates, flow visibility, escalation contacts.

MANRS baseline

Prevent propagation of incorrect routing information and maintain contactability.

SECTION 2

Operations: from alarms to service assurance

The NOC must know before the customer knows, and must know what the failure means.

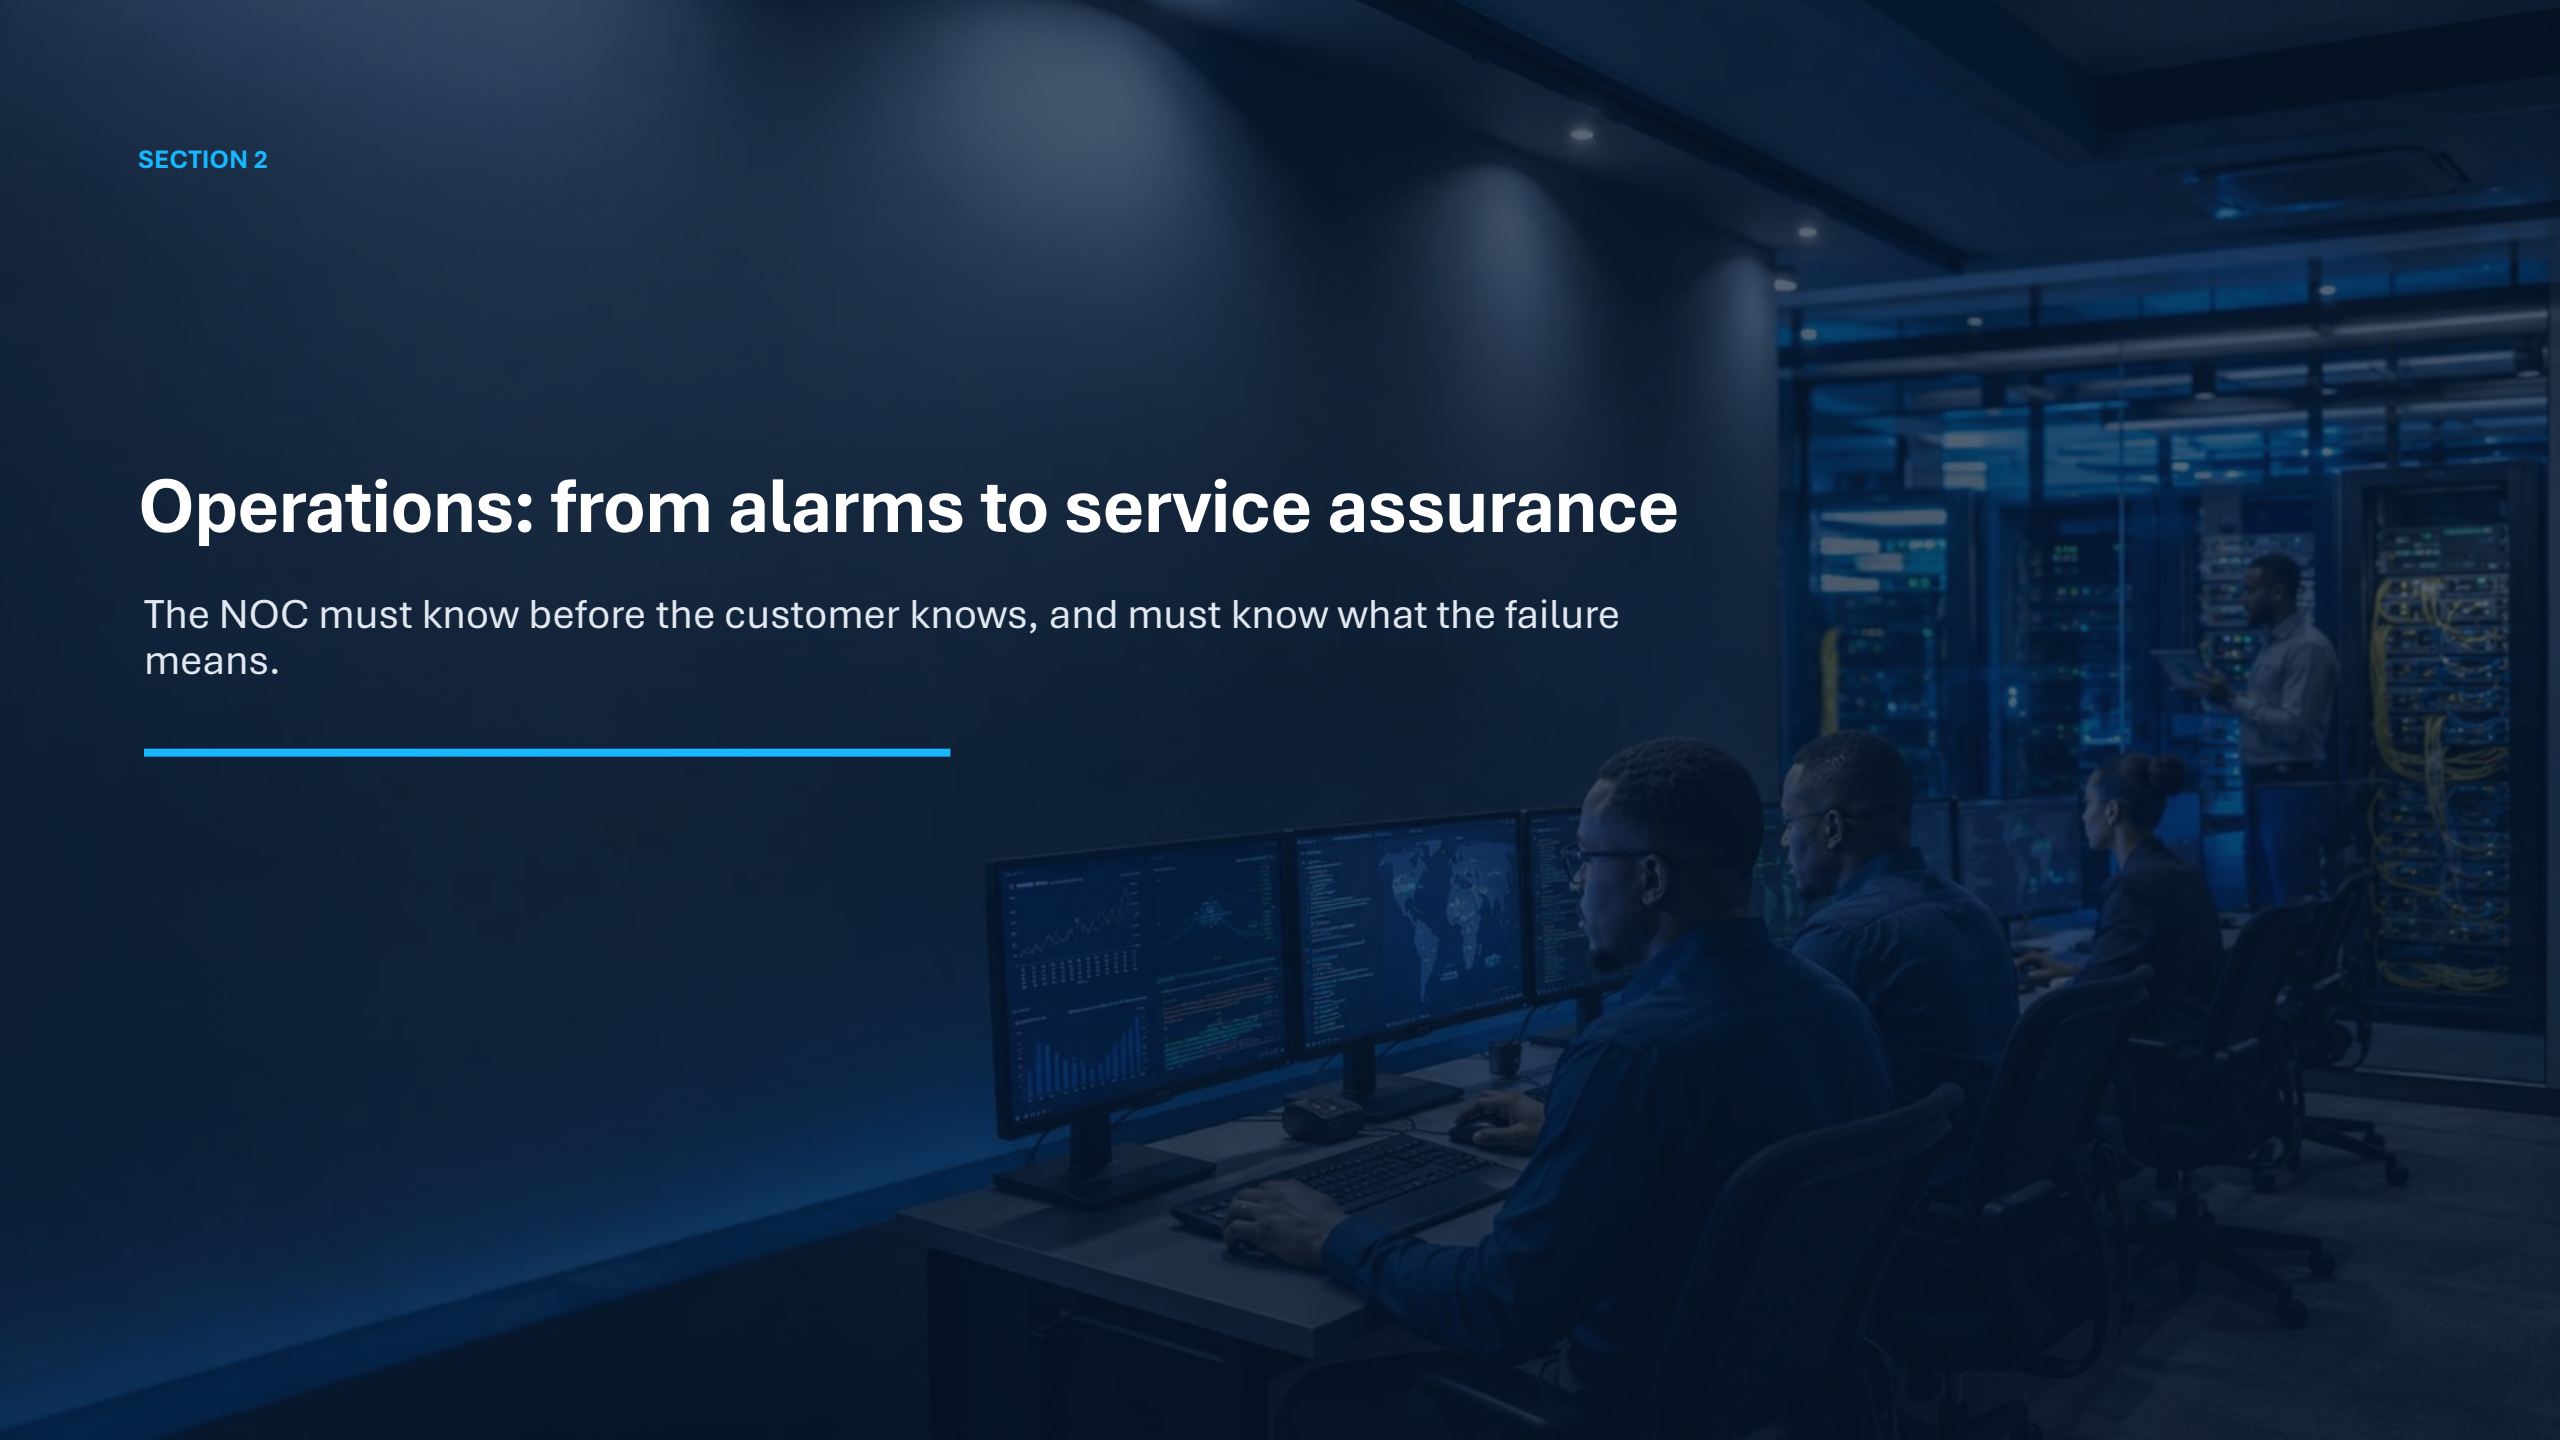

MONITORING

Monitoring must see layers, not just devices

Device uptime is not service uptime. The monitoring system must observe power, links, routing, capacity, traffic and customer experience.

Infrastructure

Power status, temperature, CPU/RAM, storage, interface errors, optical Rx/Tx, radio signal/modulation

Routing

BGP session state, route changes, prefix visibility, RPKI validity, transit/IX reachability

Traffic

SNMP counters, NetFlow/sFlow/IPFIX, p95 utilization, top talkers, CDN/cache-fill behavior

Experience

Latency, jitter, packet loss, DNS timing, HTTP tests, speed tests, RIPE Atlas probes/anchors

Events

Syslog, traps, security alerts, abuse signals, customer ticket correlation

Operations

Alert routing, escalation, maintenance windows, postmortems, SLA/customer communications

The NOC process is as important as the NMS tool

A good incident process compresses confusion into structured action.

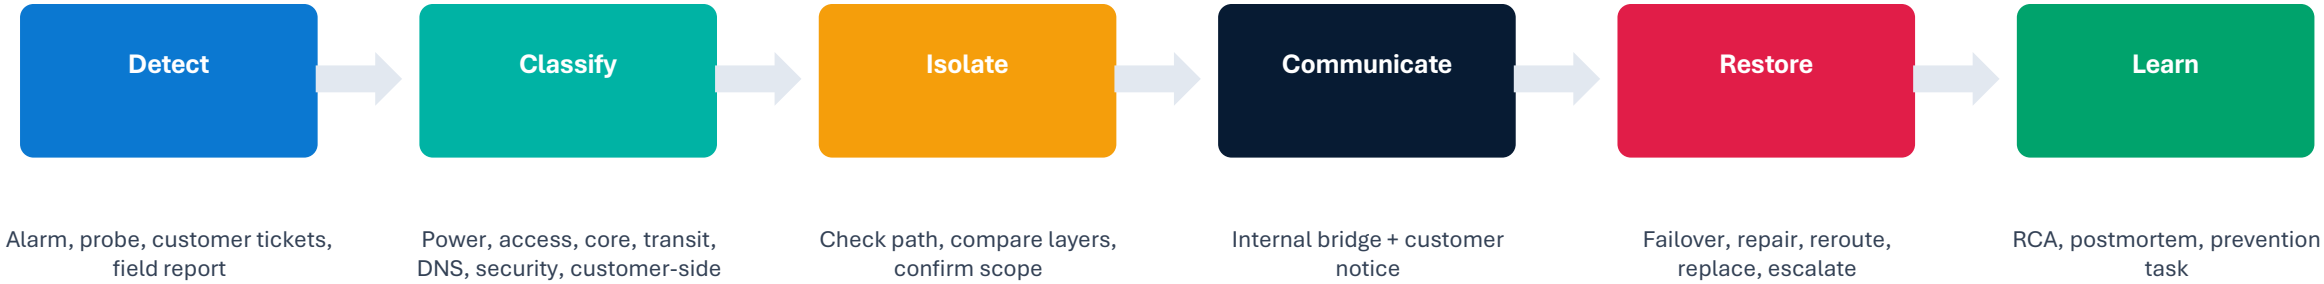

Severity model

Define severity by customer impact, affected geography, service class and duration—not by who shouts loudest.

Runbooks

Have repeatable procedures for power failure, fiber cut, transit outage, BGP leak, DDoS, OLT fault and radio sector failure.

Postmortems

Write what happened, what failed, what detected it, what restored service and what will prevent recurrence.

CAPACITY ENGINEERING

Capacity planning is an operating rhythm, not an upgrade panic

Measure before congestion becomes reputation damage.

p95 traffic trend

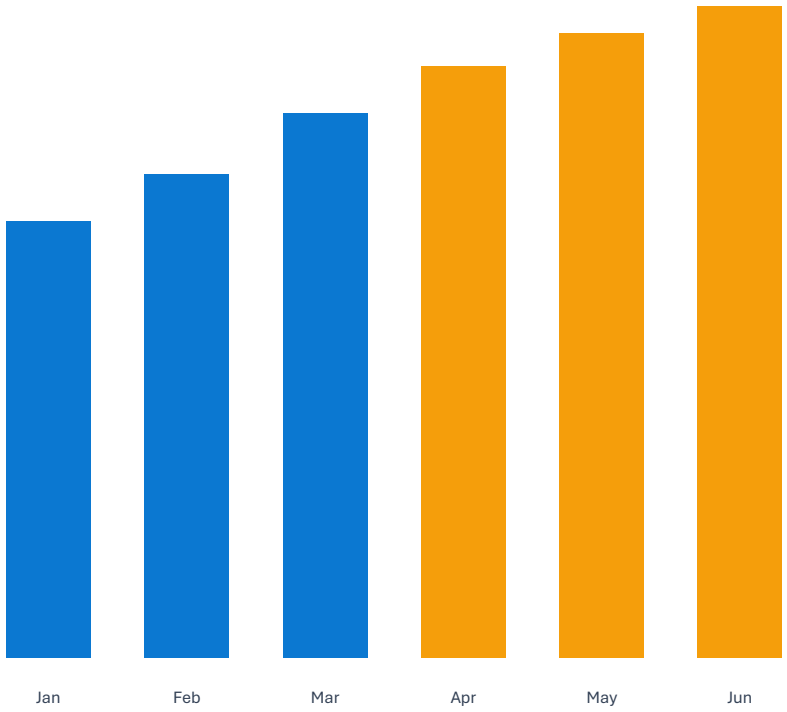

Measure

p95, peak-hour loss, latency to content, sector loading, optical power, queue drops, CPU and customer ticket correlation.

Forecast

Growth by customer class, package mix, video/cloud usage, CDN behavior and enterprise events.

Trigger

Predefined upgrade thresholds, budget cycle alignment, vendor lead time, civil-work lead time.

Validate

After upgrade: verify p95 relief, latency improvement, user complaints, flow distribution and failover behavior.

Example: upgrade trigger when p95 exceeds 70–80% or packet loss appears at peak hour.

POWER REALITY

Power is part of the network architecture

In many environments, power failure is not an exception; it is a design input.

Design inputs

- Grid availability by POP
- Load budget per router, switch, OLT, radio and cooling device
- Battery autonomy target
- Solar contribution where feasible
- Generator/refueling procedure
- Remote power monitoring

Operational controls

- Battery SOC alerting
- UPS/inverter telemetry
- DC vs AC conversion losses
- Load shedding priorities
- Prevent uncontrolled shutdowns
- Maintenance records for batteries and gensets

Network consequence

A POP without power design is a routing failure waiting to happen. Redundancy is fake if both paths depend on the same weak power domain.

Security is not only cyber; it is also physical continuity

A network can be logically sound and operationally fragile if sites, cables and towers are not protected.

Physical risks

- Fiber cuts from construction
- Cable theft and vandalism
- Tower/site access risk
- POP cabinet tampering
- Fuel theft / battery theft
- Weather and water ingress

Cyber/routing risks

- DDoS against customers or infrastructure
- BGP leaks/hijacks
- Customer CPE compromise
- DNS abuse
- Weak management access
- Unpatched edge devices

Controls

- Access logs, locks, cameras/guards where justified
- Fiber route patrols and local community relationships
- OOB access and config backups
- Mgmt VRF/VPN/ACLs/MFA
- Flow visibility and RTBH/scrubbing plan

SECTION 3

What operators need now

A practical checklist for building, scaling and protecting modern Internet infrastructure.

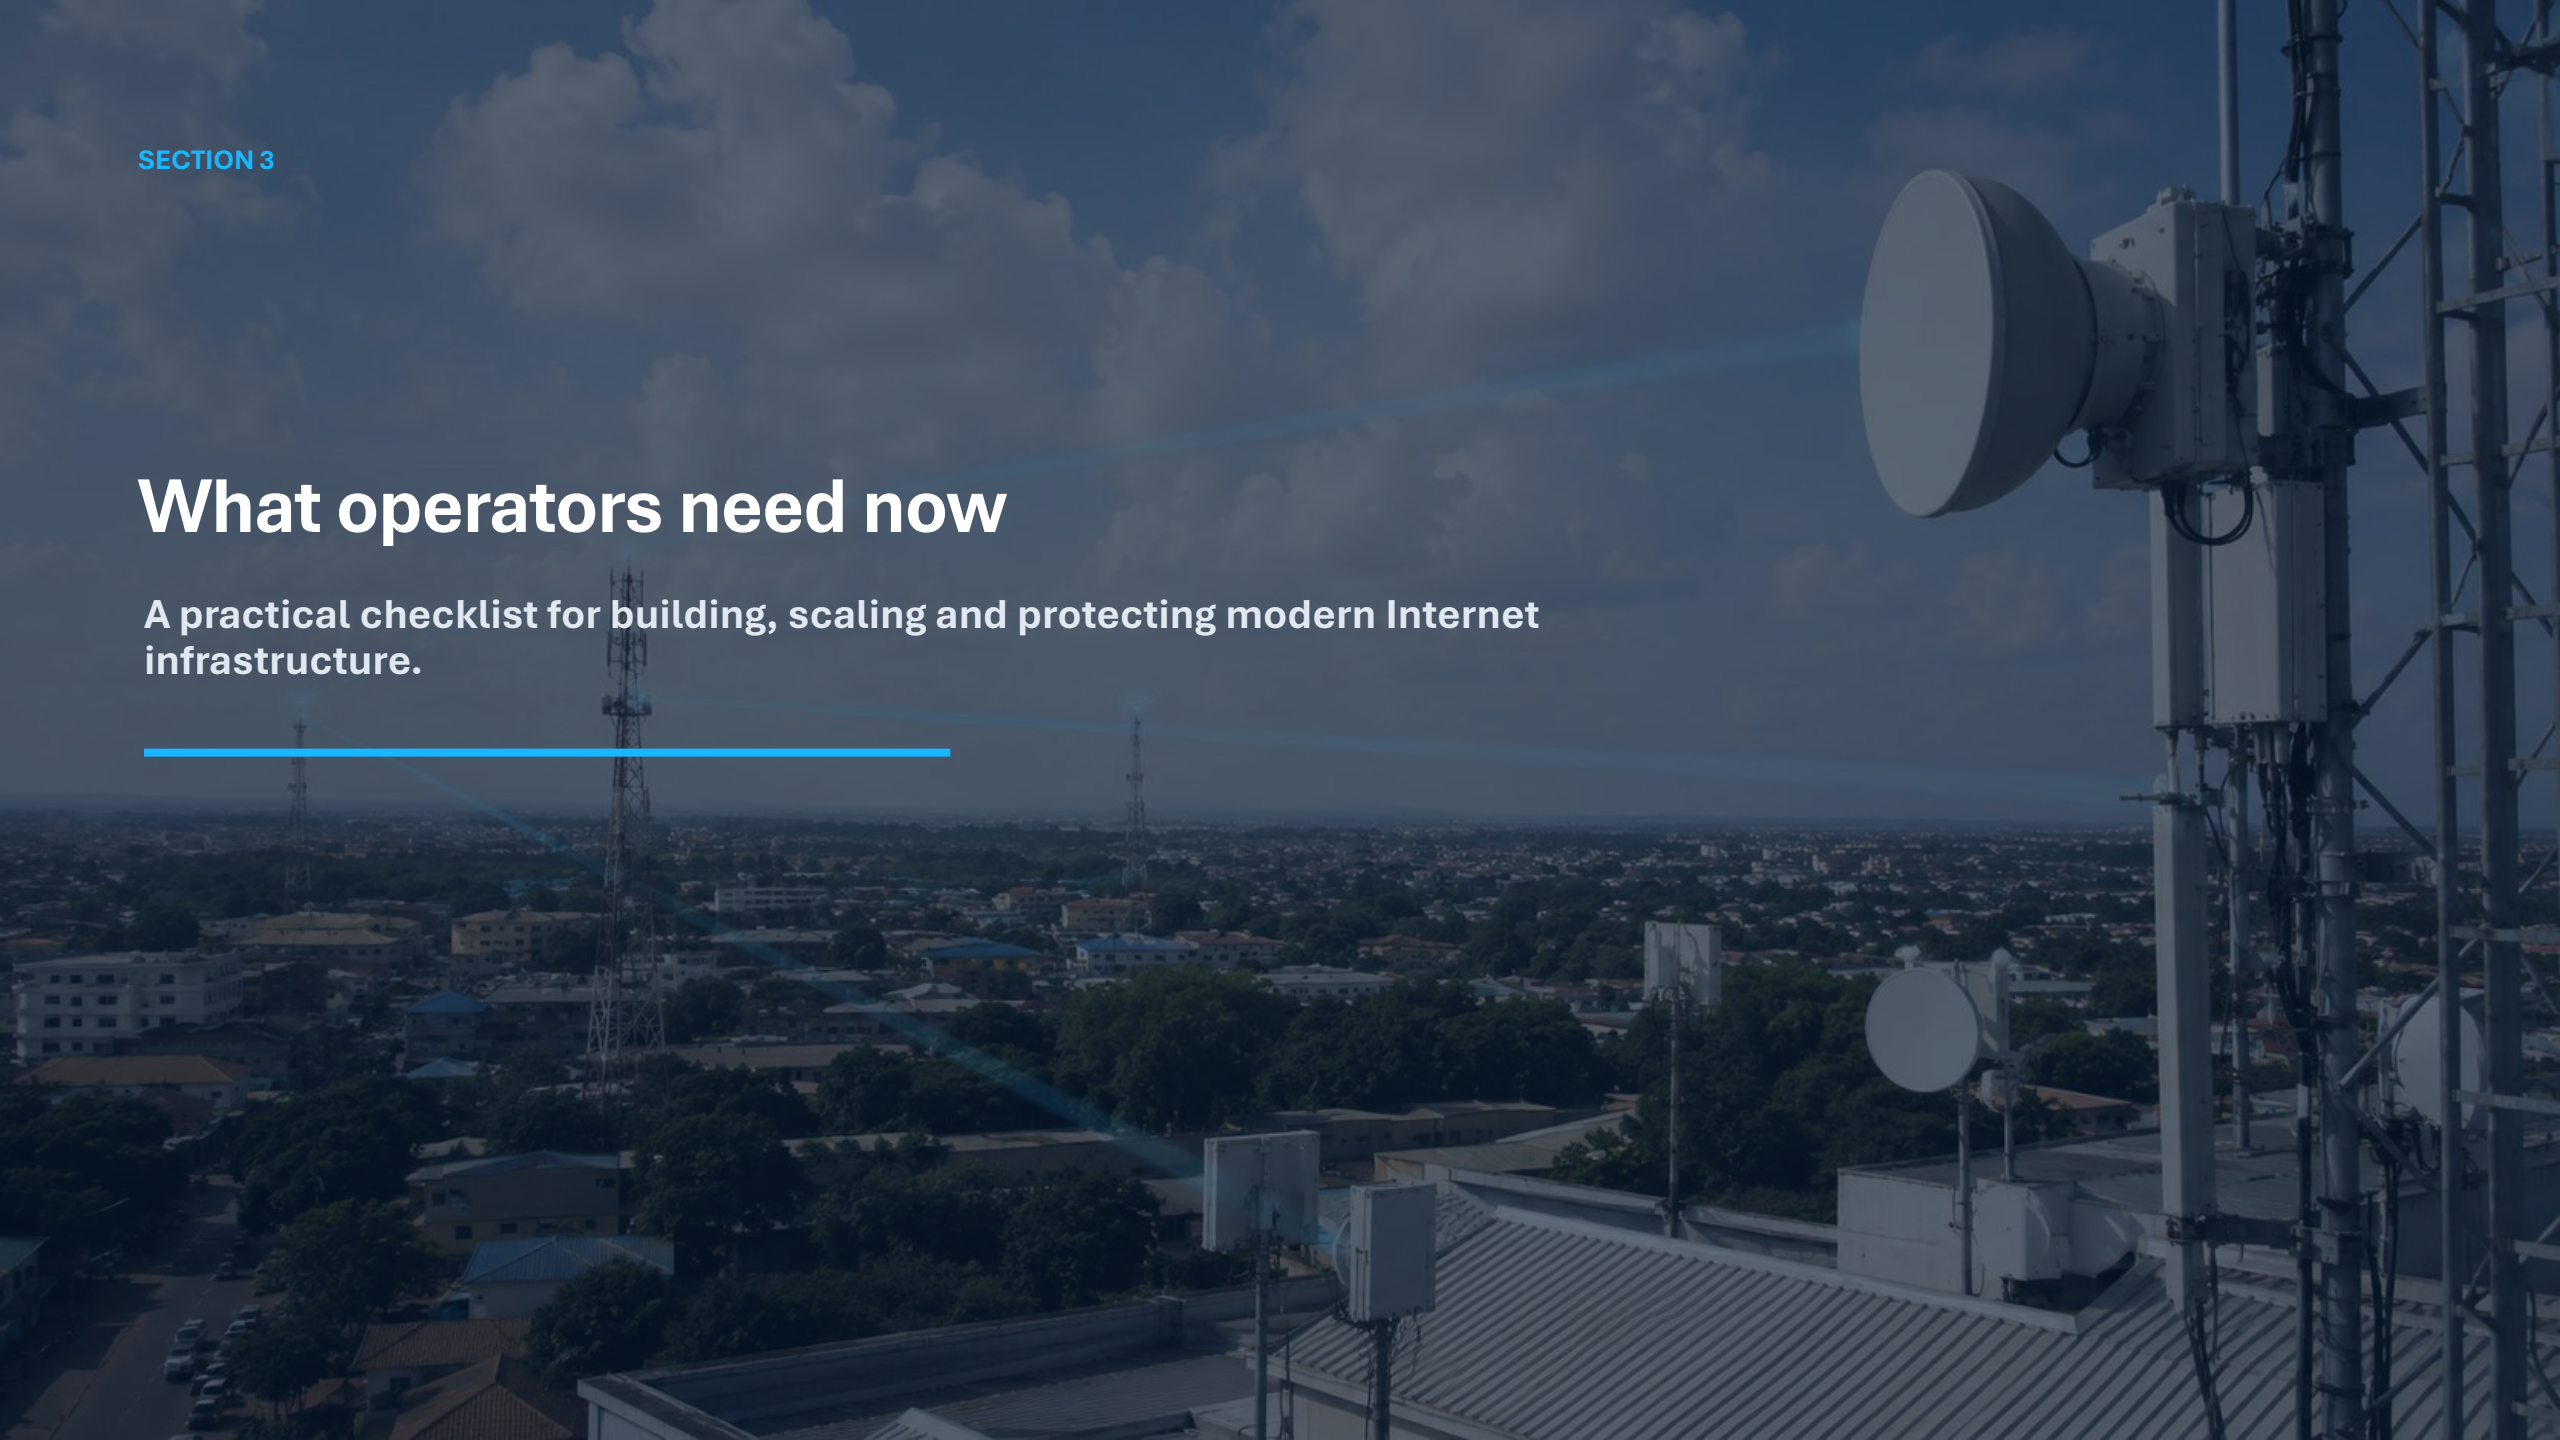

SCALING ROADMAP

Scaling is not only bigger links; it is better architecture

The upgrade path must move access, aggregation, core, content, monitoring and operations together.

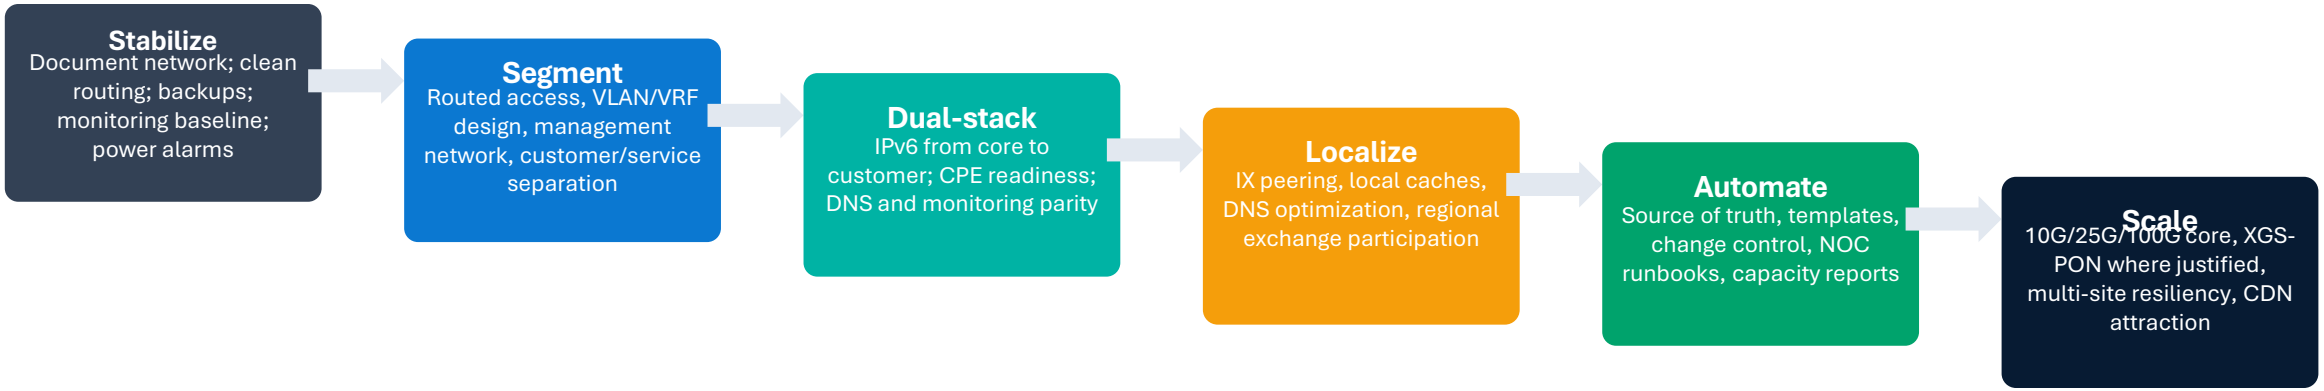

Scaling principle: the customer should feel stability improving as capacity grows. If growth creates more incidents, the architecture is not scaling.

OPERATOR CHECKLIST

If you are running a network today, you need this baseline

This is the minimum serious operating stack for a modern ISP or network operator.

Routing

ASN, BGP policy, IPv4/IPv6 prefix plan, ROAs, IRR, filters, communities

Infrastructure

Routed core, documented POPs, resilient backhaul, power monitoring, OOB access

Access

FTTH/PON design, wireless link budgets, CPE standards, service classes

Interconnection

Multiple upstreams where possible, IX membership, peering records, CDN strategy

Operations

NMS, flow, syslog, ticketing, runbooks, escalation, postmortems

Data

IPAM/source of truth, GIS/as-builts, config backups, capacity reports

Security

Management isolation, DDoS plan, customer filtering, abuse response, site security

People

Field team, NOC team, routing knowledge, fiber skills, customer comms discipline

LESSONS FROM THE FIELD

MegaMore BROADBAND

Five lessons from operating hybrid infrastructure

The modern operator must combine engineering ambition with field realism.

1. **Hybrid is viable only when each access technology has clear limits and measurement.**
2. **Power and security are not support issues; they are availability architecture.**
3. **Middle-mile and transit capacity can neutralize last-mile investment if not planned early.**
4. **IXPs, caches and IPv6 are operational tools, not conference talking points.**
5. **The strongest network is the one the team can understand and recover under pressure.**

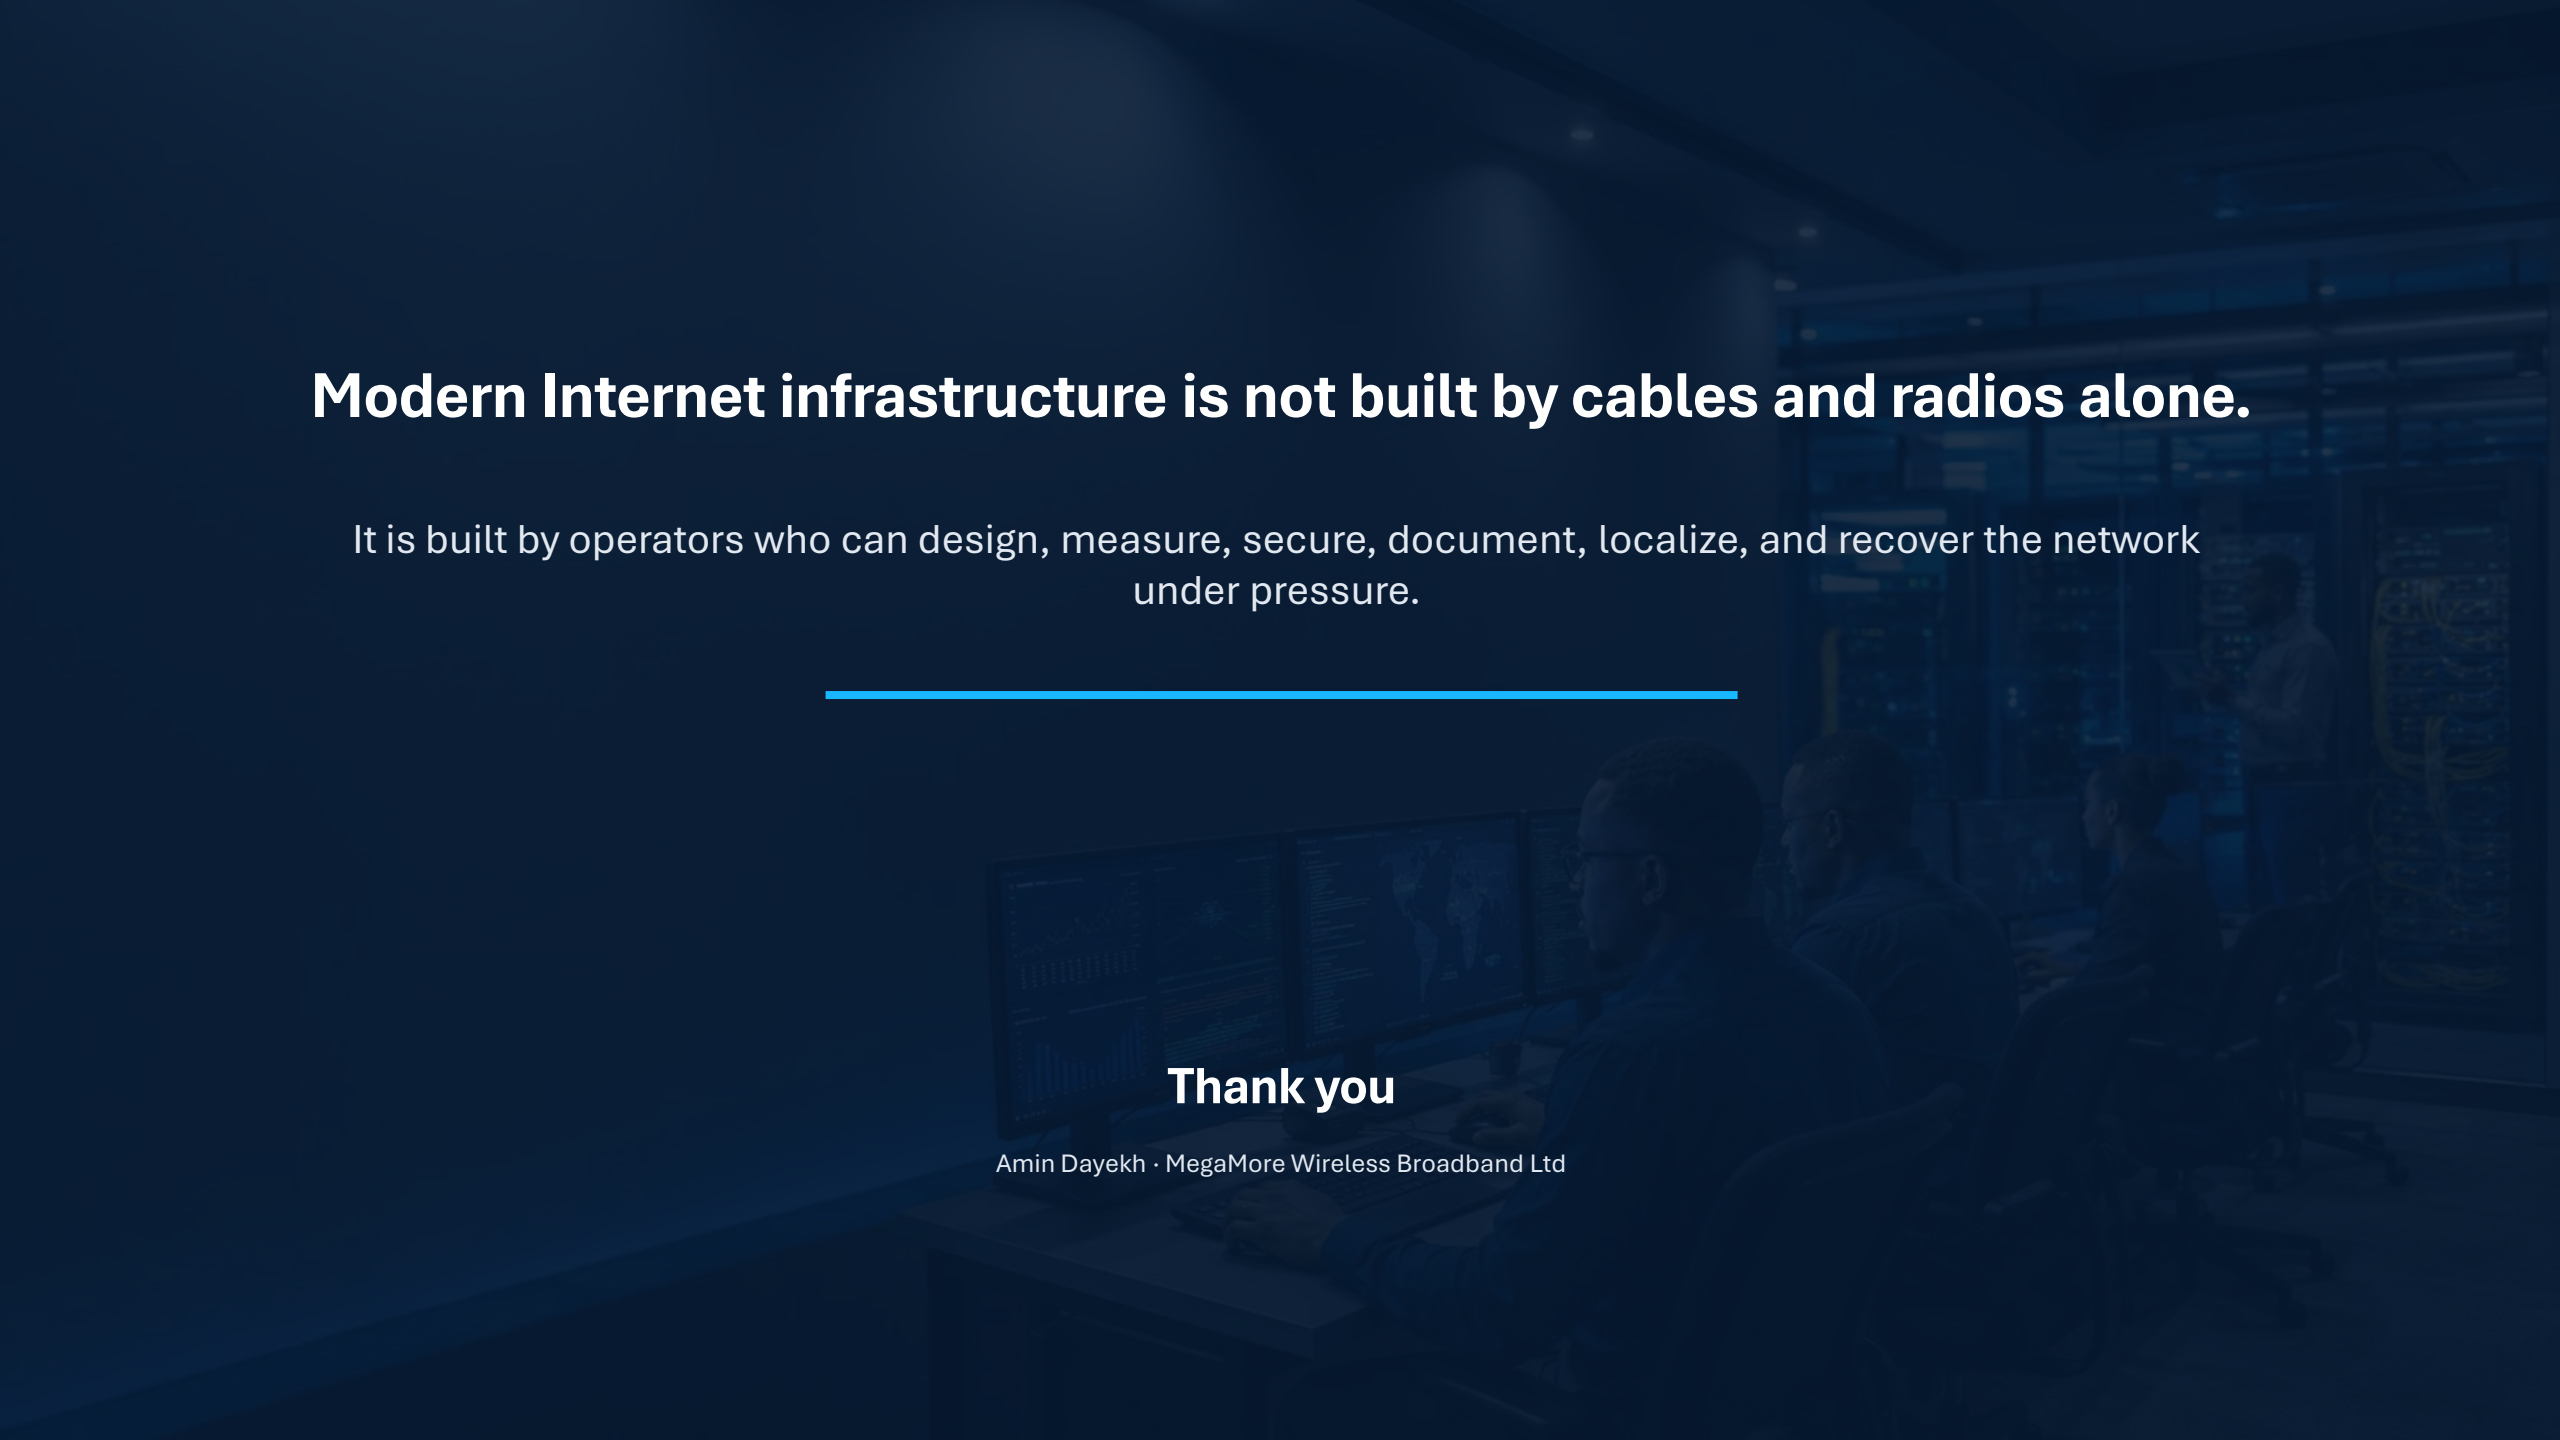

Modern Internet infrastructure is not built by cables and radios alone.

It is built by operators who can design, measure, secure, document, localize, and recover the network under pressure.

Thank you

Amin Dayekh · MegaMore Wireless Broadband Ltd

SECTION FIELD

MegaMore operational evidence

Real screenshots turn the talk from theory into field practice: fiber routes, BGP, monitoring, DNS security, virtualization and external measurement.

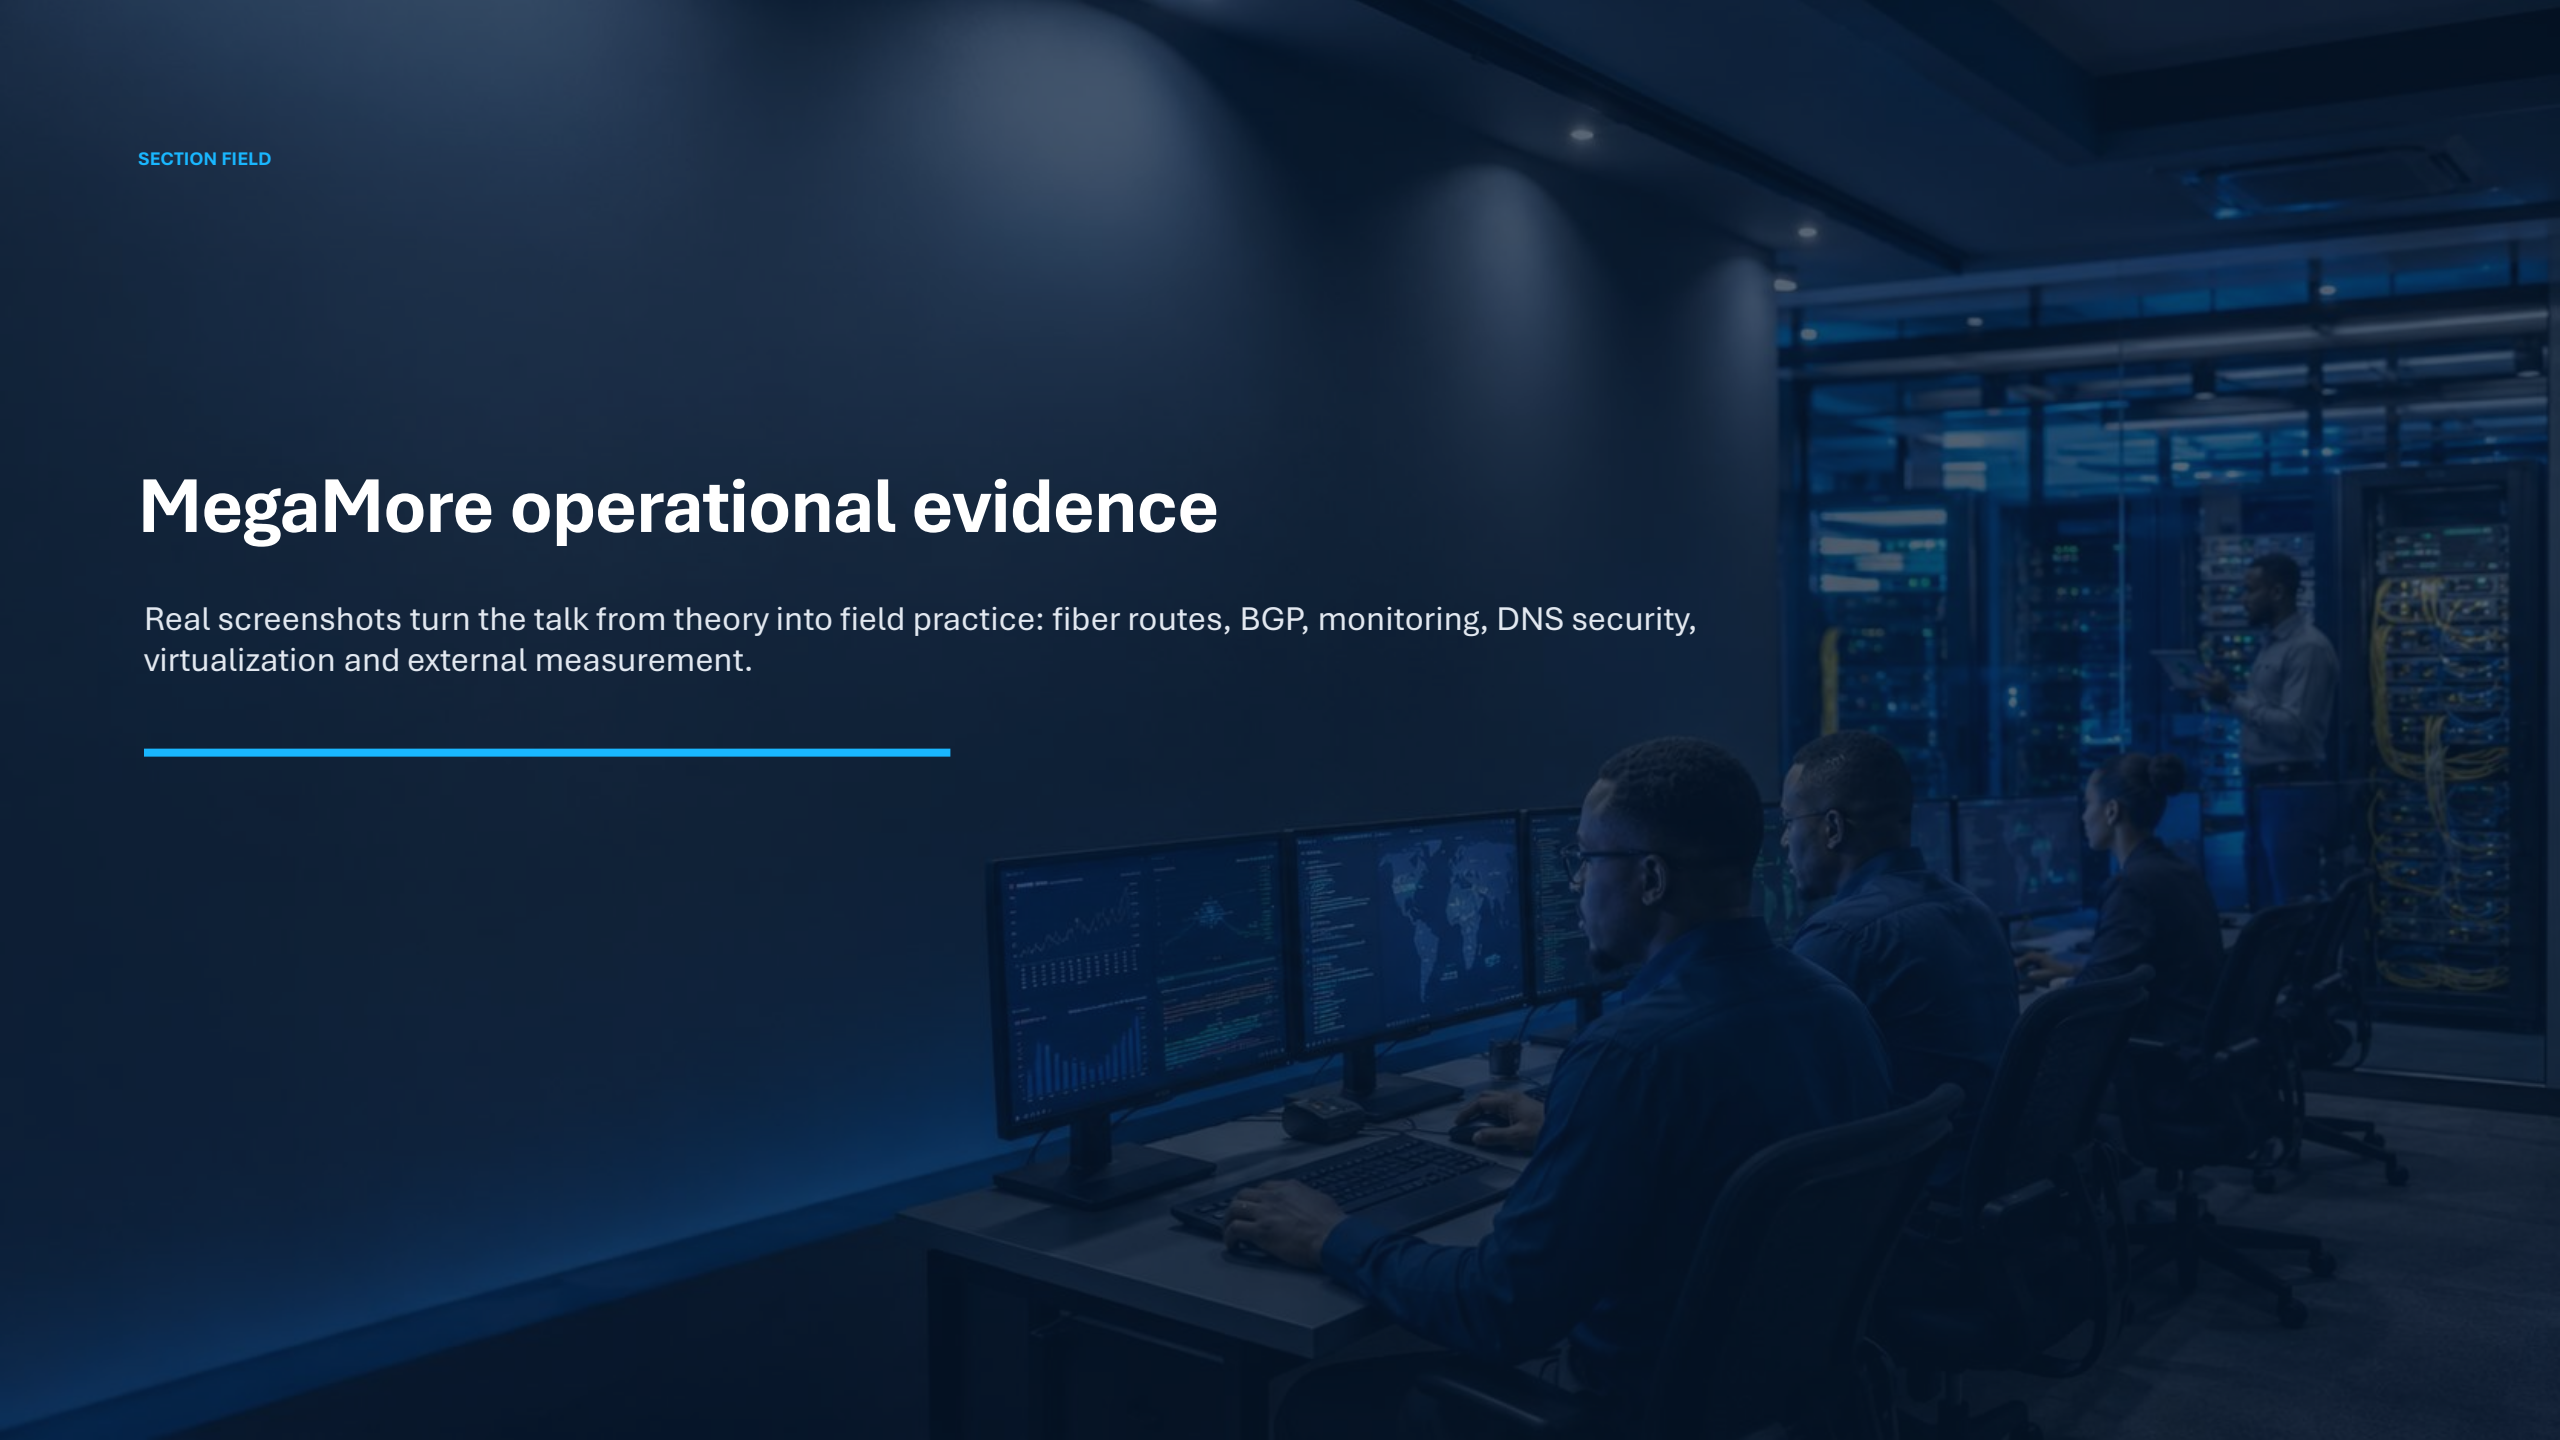

FIELD EVIDENCE

MegaMore BROADBAND

Fiber route as-built: the physical network is the first database

A fiber map is not a decoration. It is the operational source of truth for capacity, outage isolation, route risk and expansion.

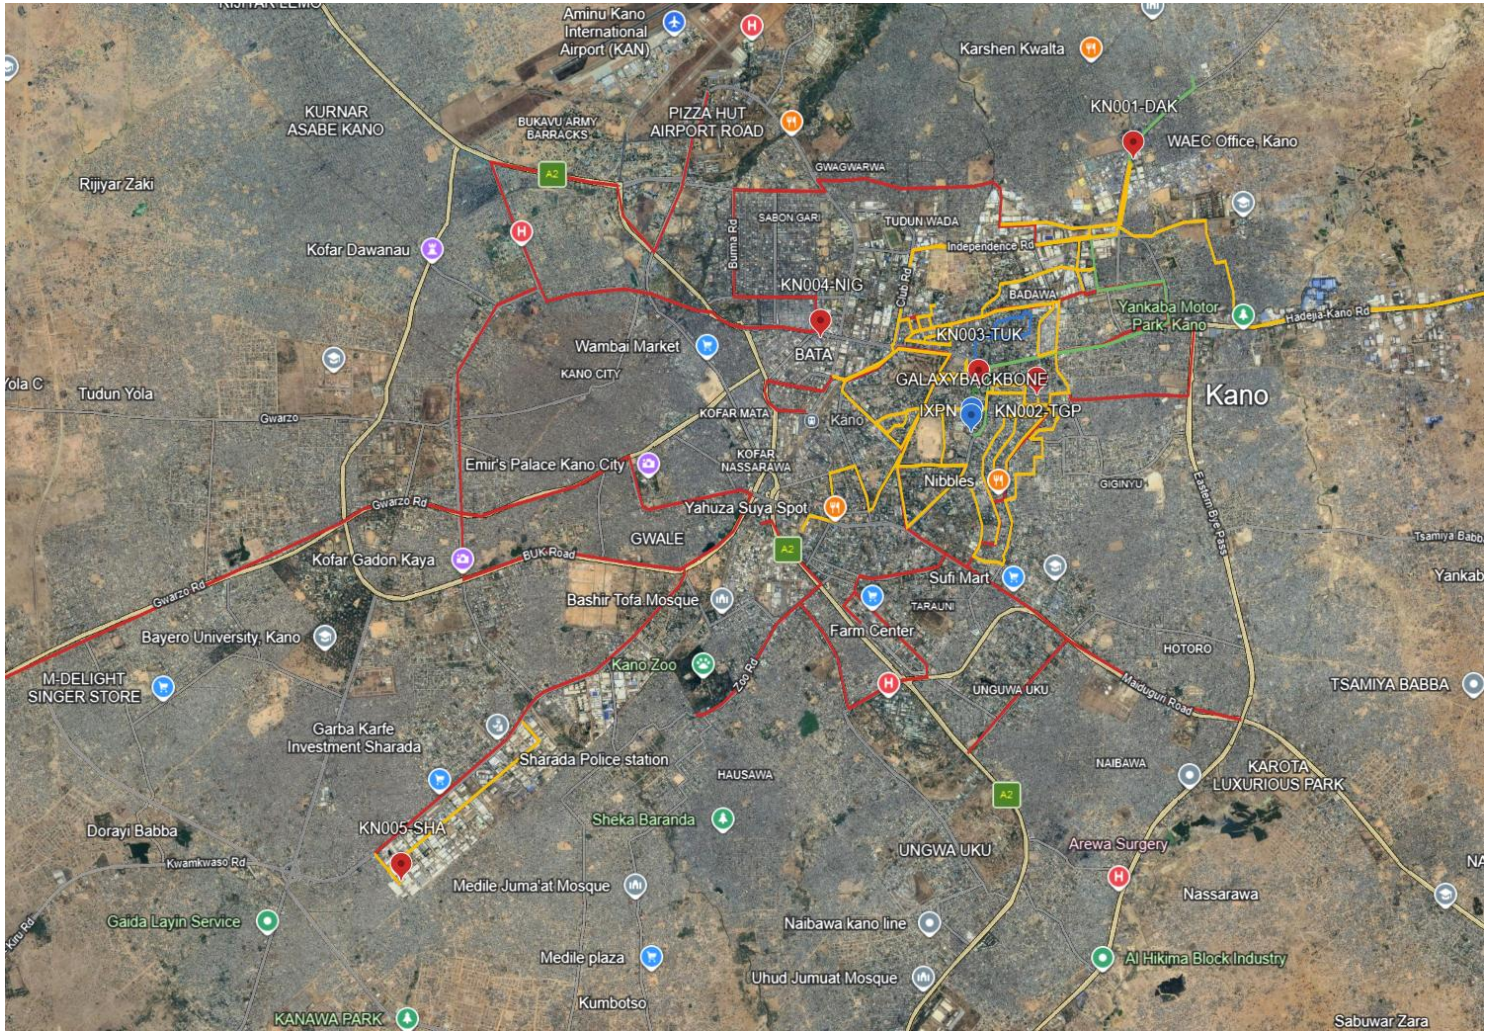

What the map proves

Named POPs and routes turn civil work into an operating system: topology, reachability, expansion and fault localization.

What it enables

Capacity sales, protection planning, maintenance dispatch, splice/closure tracing, and credible restoration estimates.

What must be added

GIS IDs, ODF/closure IDs, splice records, OTDR baselines, power notes, security exposure and restoration access.

“Before traffic reaches BGP, the network already exists as ducts, poles, closures, splices, power, security and right-of-way.”

From coverage drawing to operational topology

MegaMore BROADBAND

The clean topology slide is an abstraction. The real operator must reconcile maps, POPs, wireless, fiber and interconnection points.

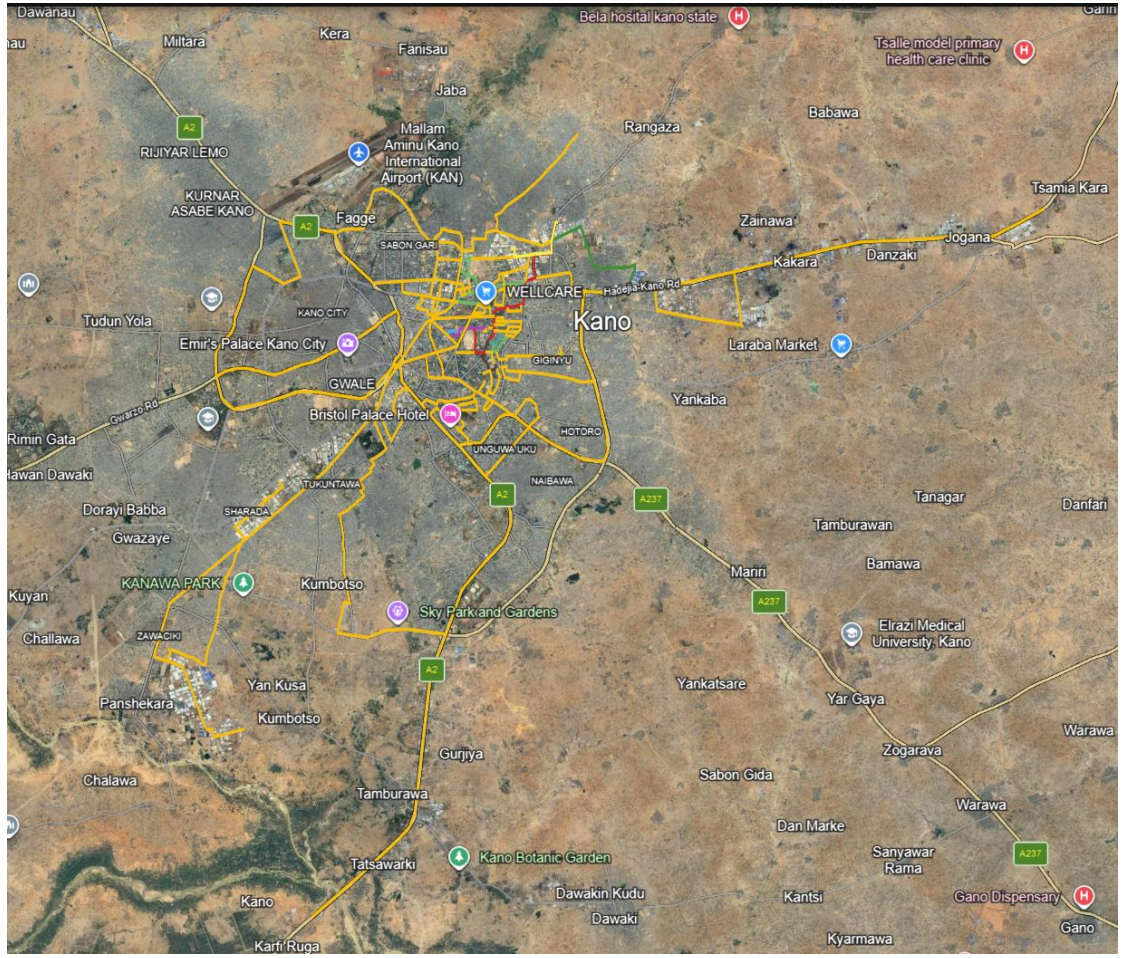

Coverage and route density

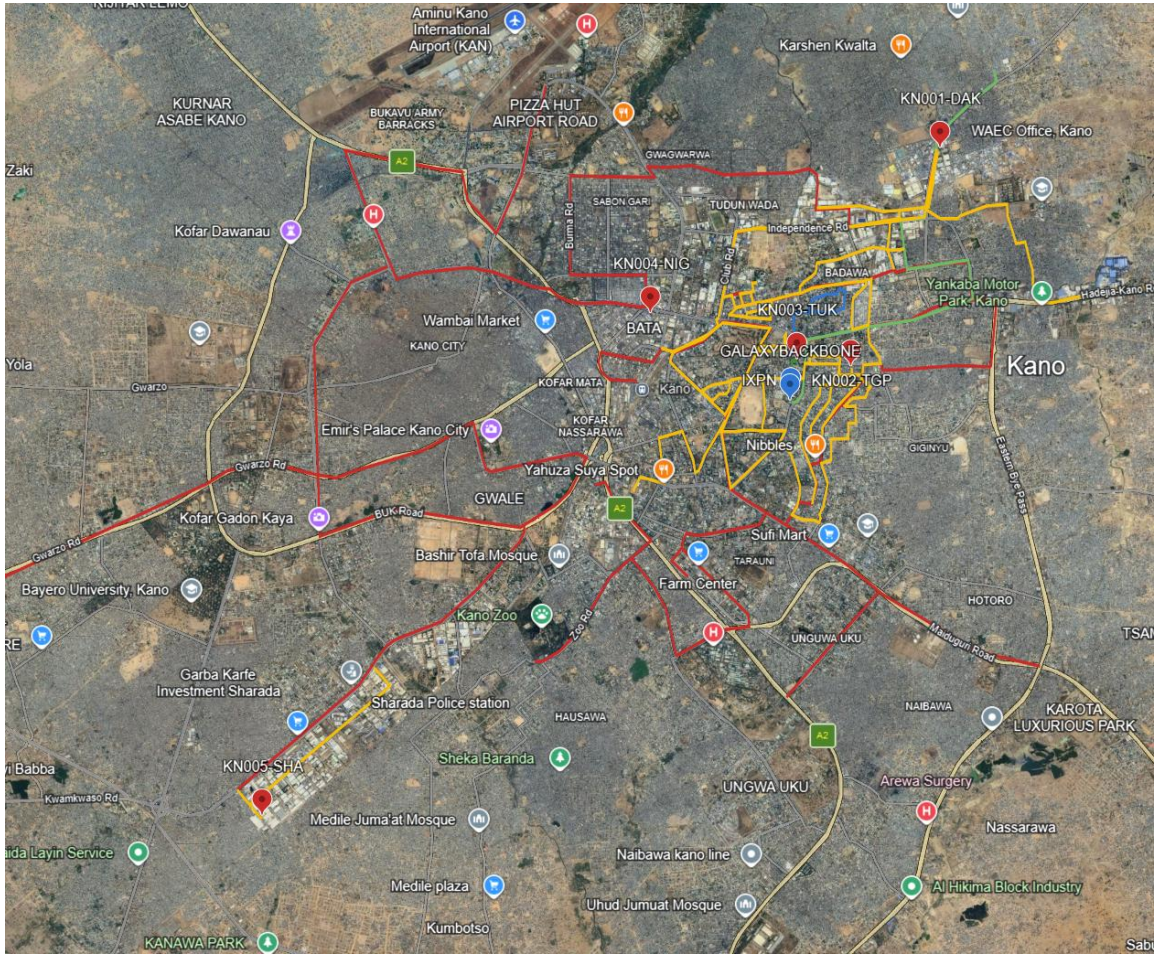

POP/interconnection-aware as-built

FIELD EVIDENCE

IPv6 proof point: external route propagation

MegaMore BROADBAND

IPv6 readiness is not a statement in a policy document. It must be visible in global routing and measurable from outside the network.

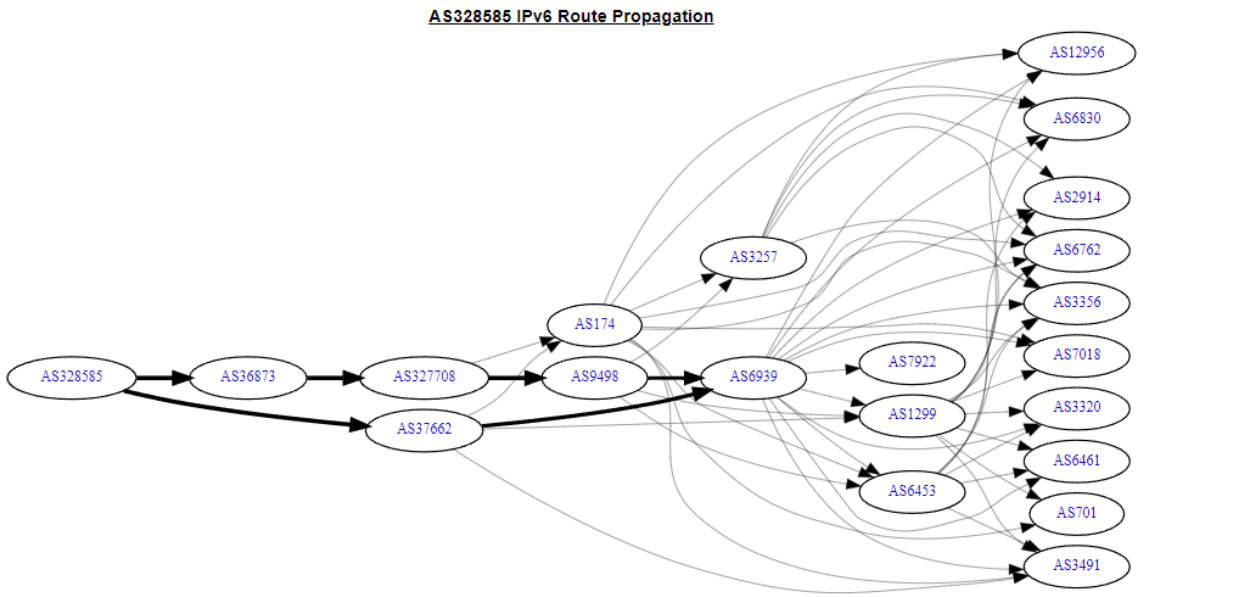

External validation

BGP.he.net route propagation shows AS328585 IPv6 reachability through upstream and onward autonomous systems.

Operational meaning

The prefix is not only configured locally; it is being carried across the interdomain routing system.

Next maturity step

Measure IPv6 traffic separately: counters, graphs, customer uptake, path quality, DNS behavior and trouble tickets.

Slide claim: IPv6 becomes real when routing, customer service, monitoring and support all treat it as production traffic.

IPv6 inside the router: BGP sessions and route table

The NOC must see IPv6 the same way it sees IPv4: sessions, routes, filters, counters, reachability and customer impact.

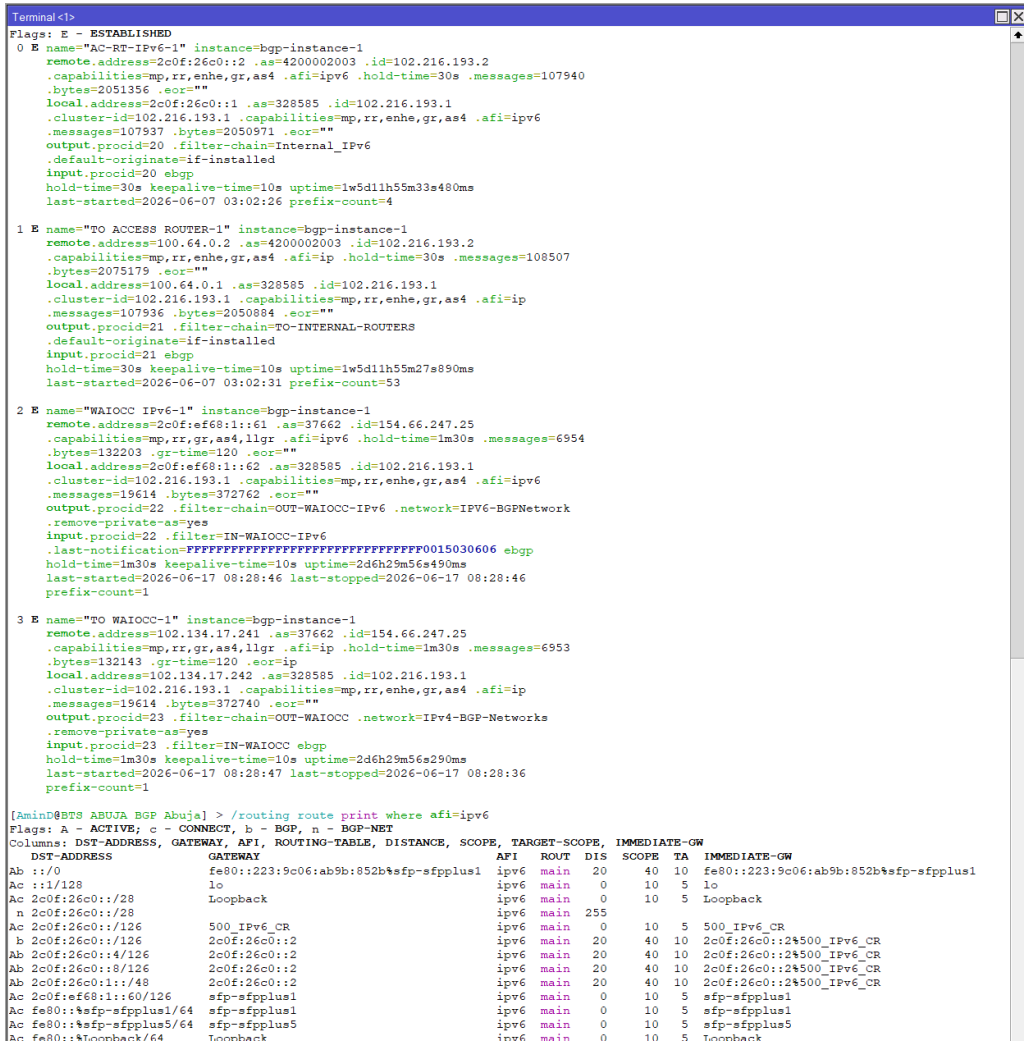

```
Terminal <1>
Flags: E - ESTABLISHED
 0 E name="AC-RT-IPv6-1" instance=bgp-instance-1
     remote.address=2c0f:26c0::2 .as=4200002003 .id=102.216.193.2
     .capabilities=mp,rr,enhe,gr,as4 .afi=ipv6 .hold-time=30s .messages=107940
     .bytes=2051356 .eor=""
     local.address=2c0f:26c0::1 .as=328585 .id=102.216.193.1
     .cluster-id=102.216.193.1 .capabilities=mp,rr,enhe,gr,as4 .afi=ipv6
     .messages=107937 .bytes=2050971 .eor=""
     output.procid=20 .filter-chain=Internal_IPv6
     .default-originate=if-installed
     input.procid=20 ebgp
     hold-time=30s keepalive-time=10s uptime=1w5d11h55m33s480ms
     last-started=2026-06-07 03:02:26 prefix-count=4

 1 E name="TO ACCESS ROUTER-1" instance=bgp-instance-1
     remote.address=100.64.0.2 .as=4200002003 .id=102.216.193.2
     .capabilities=mp,rr,enhe,gr,as4 .afi=ip .hold-time=30s .messages=108507
     .bytes=2075179 .eor=""
     local.address=100.64.0.1 .as=328585 .id=102.216.193.1
     .cluster-id=102.216.193.1 .capabilities=mp,rr,enhe,gr,as4 .afi=ip
     .messages=107936 .bytes=2050884 .eor=""
     output.procid=21 .filter-chain=TO-INTERNAL-ROUTERS
     .default-originate=if-installed
     input.procid=21 ebgp
     hold-time=30s keepalive-time=10s uptime=1w5d11h55m27s890ms
     last-started=2026-06-07 03:02:31 prefix-count=53

 2 E name="WAIOCC IPv6-1" instance=bgp-instance-1
     remote.address=2c0f:ef68:1::61 .as=37662 .id=154.66.247.25
     .capabilities=mp,rr,gr,as4,llgr .afi=ipv6 .hold-time=1m30s .messages=6954
     .bytes=132203 .gr-time=120 .eor=""
     local.address=2c0f:ef68:1::62 .as=328585 .id=102.216.193.1
     .cluster-id=102.216.193.1 .capabilities=mp,rr,enhe,gr,as4 .afi=ipv6
     .messages=19614 .bytes=372762 .eor=""
     output.procid=22 .filter-chain=OUT-WAIOCC-IPv6 .network=IPV6-BGPNetwork
     .remove-private-as=yes
     input.procid=22 .filter=IN-WAIOCC-IPv6
     .last-notification=FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF0015030606 ebgp
     hold-time=1m30s keepalive-time=10s uptime=2d6h29m56s490ms
     last-started=2026-06-17 08:28:46 last-stopped=2026-06-17 08:28:46
     prefix-count=1

 3 E name="TO WAIOCC-1" instance=bgp-instance-1
     remote.address=102.134.17.241 .as=37662 .id=154.66.247.25
     .capabilities=mp,rr,gr,as4,llgr .afi=ip .hold-time=1m30s .messages=6953
     .bytes=132143 .gr-time=120 .eor=ip
     local.address=102.134.17.242 .as=328585 .id=102.216.193.1
     .cluster-id=102.216.193.1 .capabilities=mp,rr,enhe,gr,as4 .afi=ip
     .messages=19614 .bytes=372740 .eor=""
     output.procid=23 .filter-chain=OUT-WAIOCC .network=IPv4-BGP-Networks
     .remove-private-as=yes
     input.procid=23 .filter=IN-WAIOCC ebgp
     hold-time=1m30s keepalive-time=10s uptime=2d6h29m56s290ms
     last-started=2026-06-17 08:28:47 last-stopped=2026-06-17 08:28:36
     prefix-count=1

[AminD@BTS ABUJA BGP Abuja] > /routing route print where afi=ipv6
Flags: A - ACTIVE; c - CONNECT, b - BGP, n - BGP-NET
Columns: DST-ADDRESS, GATEWAY, AFI, ROUTING-TABLE, DISTANCE, SCOPE, TARGET-SCOPE, IMMEDIATE-GW
   DST-ADDRESS             GATEWAY                          AFI   ROUT  DIS  SCOPE  TA  IMMEDIATE-GW
Ab ::/0                    fe80::223:9c06:ab9b:852b%sfp-sfpplus1  ipv6  main   20     40  10  fe80::223:9c06:ab9b:852b%sfp-sfpplus1
Ac ::1/128                 lo                               ipv6  main    0     10   5  lo
Ac 2c0f:26c0::/28          Loopback                         ipv6  main    0     10   5  Loopback
 n 2c0f:26c0::/28                                          ipv6  main  255
Ac 2c0f:26c0::/126         500_IPv6_CR                      ipv6  main    0     10   5  500_IPv6_CR
 b 2c0f:26c0::/126         2c0f:26c0::2                     ipv6  main   20     40  10  2c0f:26c0::2%500_IPv6_CR
Ab 2c0f:26c0::4/126        2c0f:26c0::2                     ipv6  main   20     40  10  2c0f:26c0::2%500_IPv6_CR
Ab 2c0f:26c0::8/126        2c0f:26c0::2                     ipv6  main   20     40  10  2c0f:26c0::2%500_IPv6_CR
Ab 2c0f:26c0:1::/48        2c0f:26c0::2                     ipv6  main   20     40  10  2c0f:26c0::2%500_IPv6_CR
Ac 2c0f:ef68:1::60/126     sfp-sfpplus1                     ipv6  main    0     10   5  sfp-sfpplus1
Ac fe80::%sfp-sfpplus1/64  sfp-sfpplus1                     ipv6  main    0     10   5  sfp-sfpplus1
Ac fe80::%sfp-sfpplus5/64  sfp-sfpplus5                     ipv6  main    0     10   5  sfp-sfpplus5
Ac fe80::%Loopback/64      Loopback                         ipv6  main    0     10   5  Loopback
```

What is visible

Established BGP sessions, AFI IPv6, WIAOCC IPv6 peer, route entries, loopback/core links and prefix count.

What should be monitored

Session state, prefix count, route changes, default route, latency, packet loss and interface counters.

What should be documented

Prefix plan, route filters, ROAs, customer delegation pools, firewall policy and escalation runbook.

Topology monitoring: useful, but not enough by itself

Dude-style maps help visualize reachability, but modern operations need topology plus metrics, tickets, logs and source-of-truth documentation.

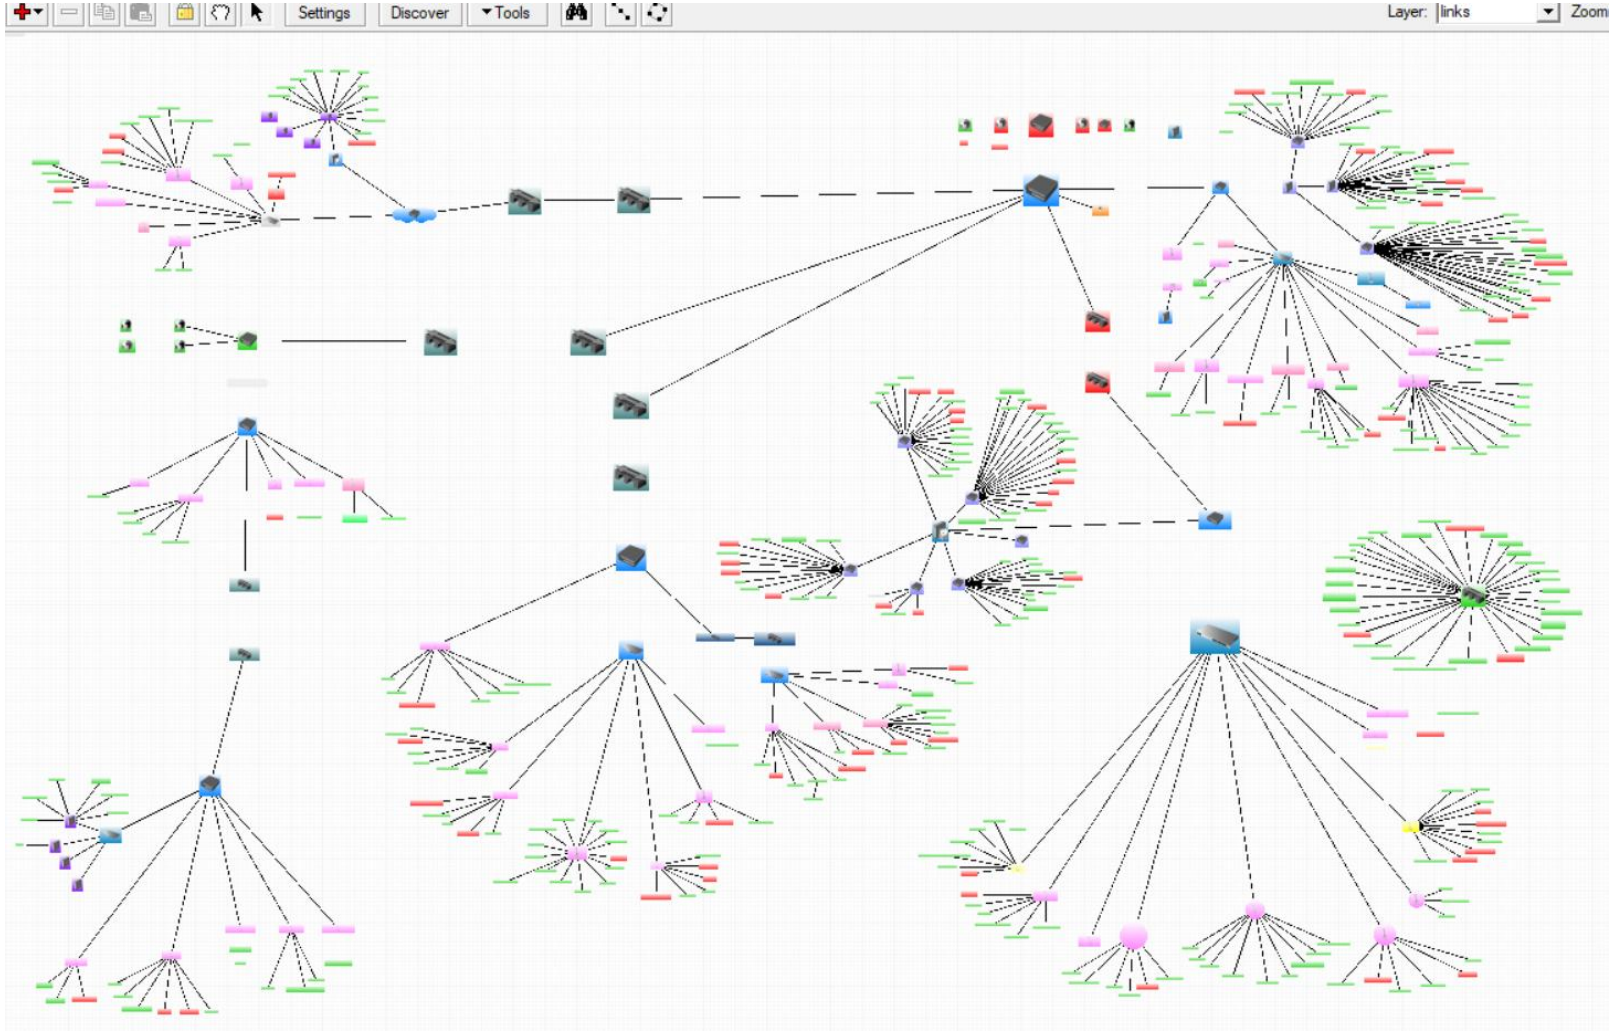

Good use

Reachability map, quick visual fault detection, site/device dependency view.

Risk

If the map is not maintained, it becomes a wall poster, not an operational system.

Upgrade path

Integrate with SNMP metrics, syslog, IPAM, ticketing, capacity graphs and incident workflow.

FIELD EVIDENCE

MegaMore BROADBAND

DNS security telemetry: access networks now see abuse and risk

Modern ISP operation includes customer protection and abuse visibility, not only bandwidth delivery.

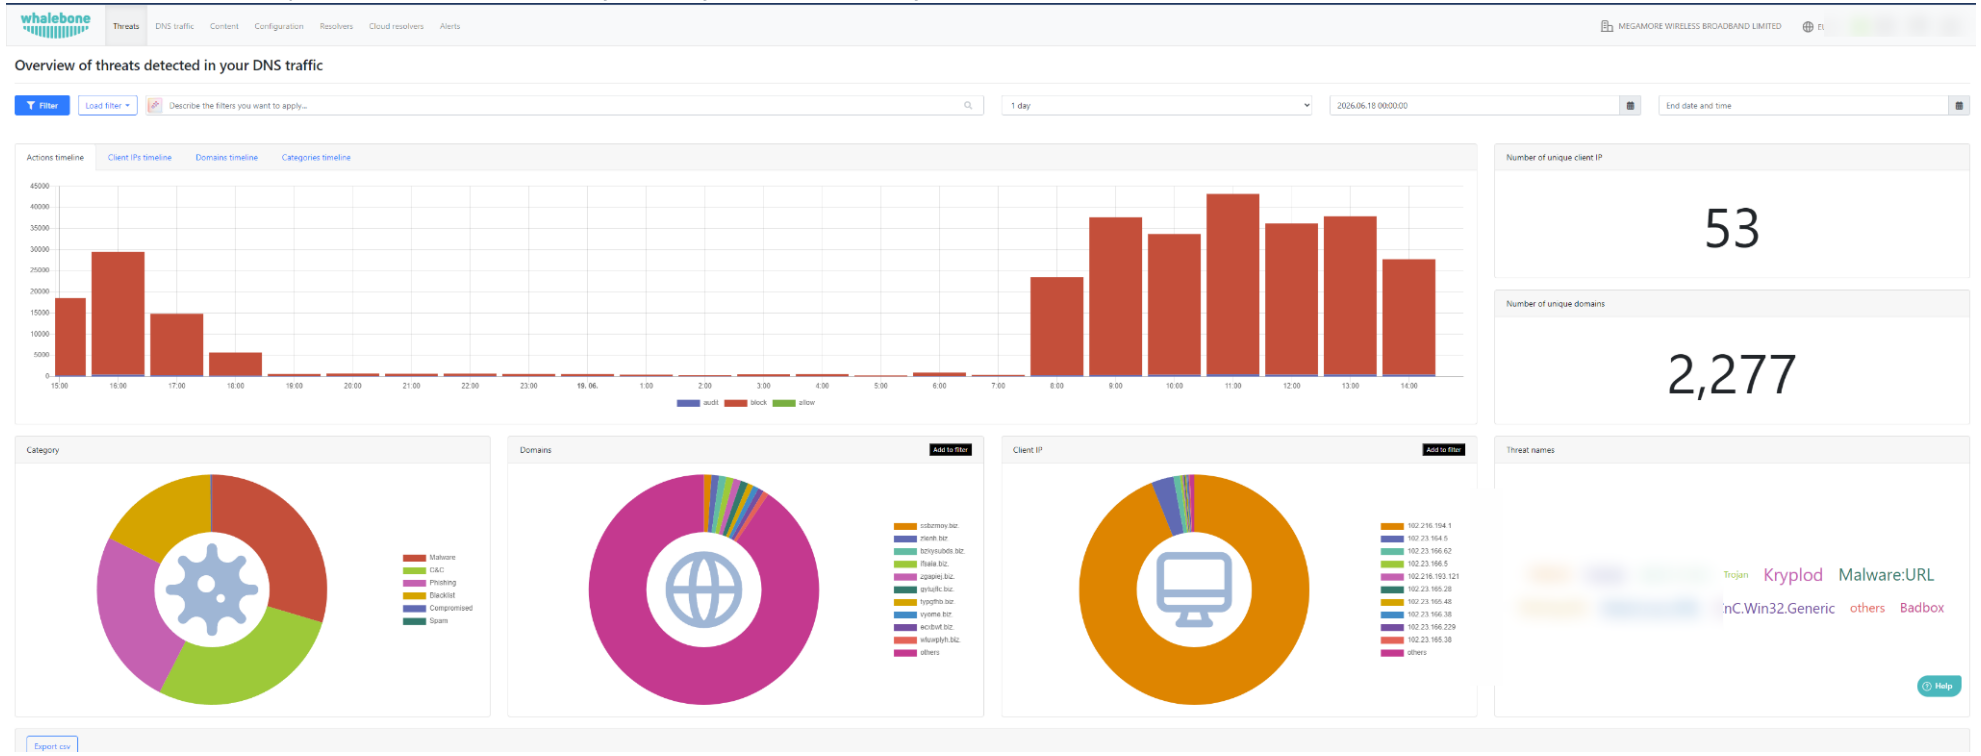

Security layer

Whalebone provides DNS-layer threat visibility and blocking/audit/allow views.

Operational signal

Threat categories such as malware, C&C, phishing and malicious URLs become measurable.

Customer care

Security data can guide customer education, support escalation and abuse handling.

Boundary

This is not a full SOC; it is one layer in the ISP security stack.

Local operations platform: virtualized tools near the network

The operator needs local services for measurement, monitoring, DNS/security, vendor controllers and operational tooling.

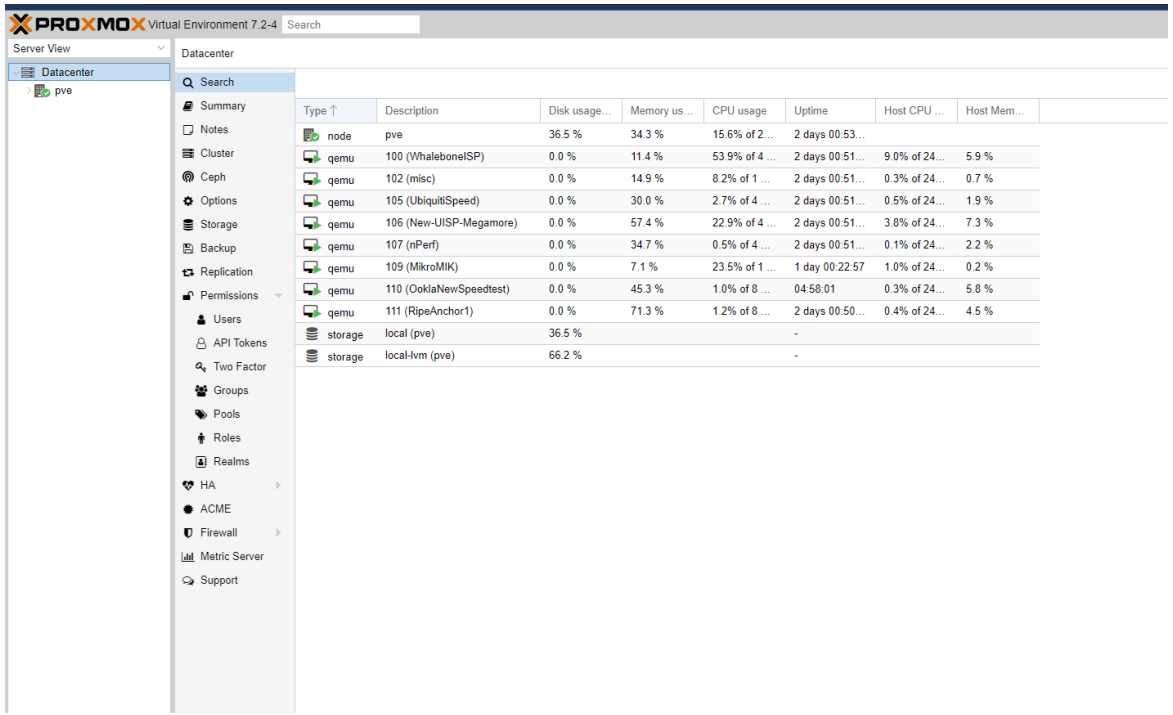

Proxmox: local virtualized service stack

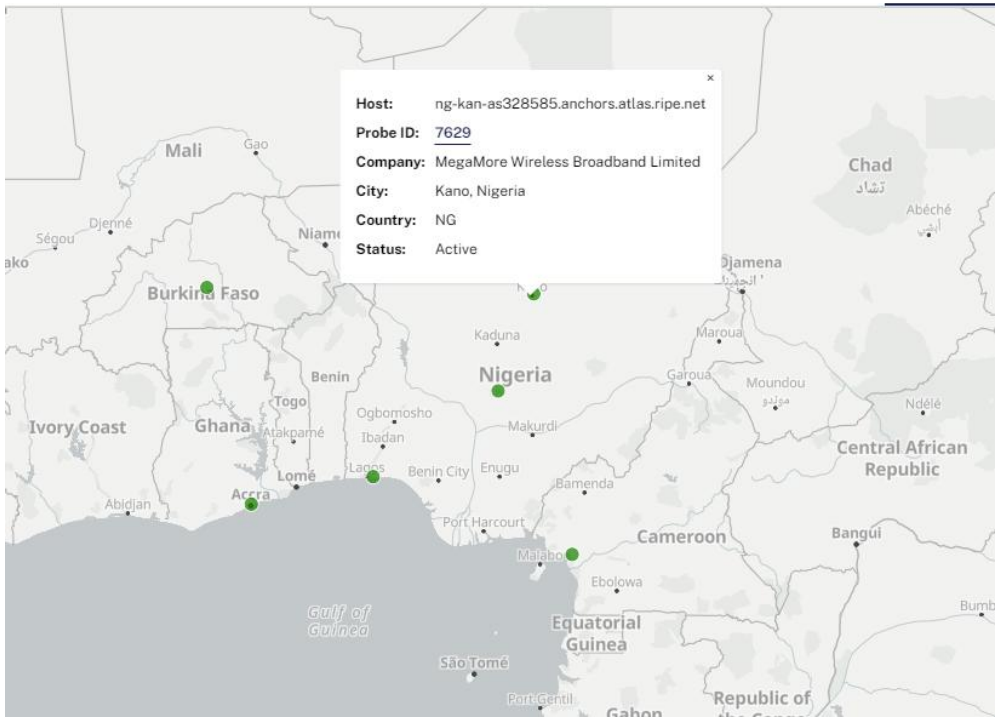

RIPE Atlas Anchor: external measurement from Kano

FIELD EVIDENCE

What the screenshots prove operationally

MegaMore BROADBAND

The evidence is not the individual tool. The evidence is that each operational layer has an instrument.

Physical plant

Fiber route/as-built maps
Routes, POPs, IX proximity, expansion and fault restoration

Routing

BGP.he.net + MikroTik BGP
IPv6 propagation, sessions, route table and prefix visibility

Monitoring

Dude topology
Reachability and dependency visualization

Security

Whalebone dashboard
DNS-layer threat visibility and customer protection

Service platform

Proxmox stack
Local tooling for speed tests, controllers, DNS/security and measurement

External measurement

RIPE Atlas Anchor
Independent active measurement point from Kano

Operator maturity: for every layer, ask what detects failure, what proves performance, who owns response, and where the source of truth lives.

SECTION A

Appendix: operational reference slides

Use these if discussion becomes more technical or for distribution after the talk.

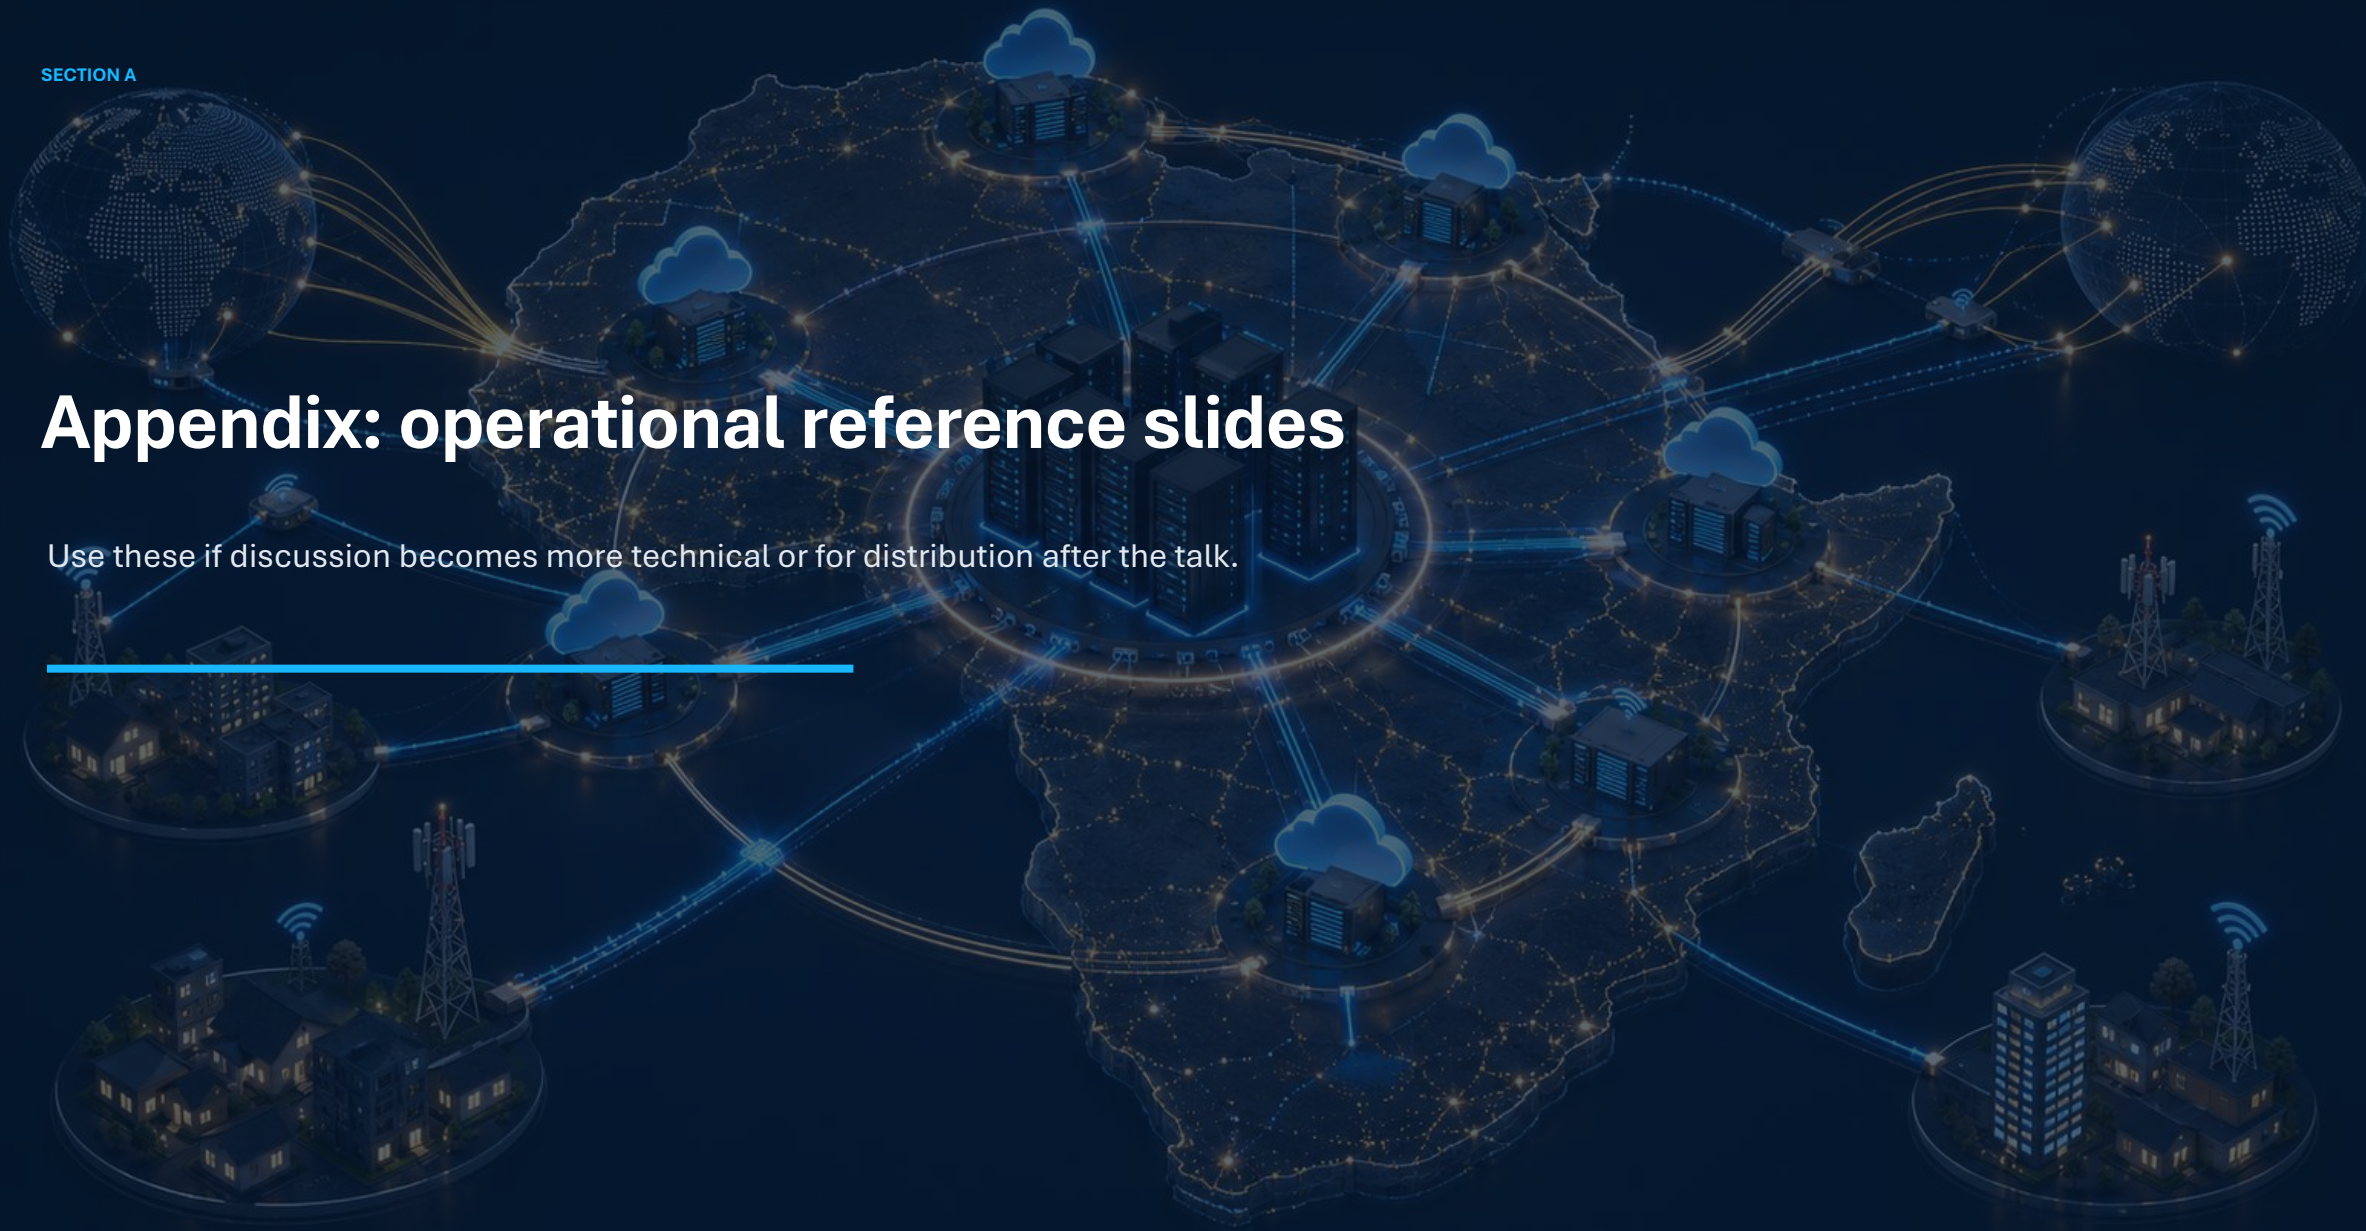

APPENDIX A

Operational metrics to track weekly

MegaMore BROADBAND

A modern network review should use objective indicators, not only complaint volume.

Transit / IX

p95, peak utilization, loss, latency to top destinations, BGP changes

Core

CPU, memory, interface errors, queue drops, route table, failover tests

Fiber

OTDR baseline, optical Rx/Tx, cuts, MTTR, closure/ODF fault history

Wireless

RSSI/SNR, modulation, retries, sector load, interference, uptime by sector

Access / CPE

ONT status, CPE firmware, Wi-Fi complaints, session drops, DHCP/AAA failures

Customer experience

Latency, jitter, DNS time, app tests, ticket volumes, repeat faults, SLA compliance

MegaMore BROADBAND

Fiber deployment checklist

The difference between fiber deployment and fiber infrastructure is documentation plus maintainability.

- Route survey and demand map completed
- RoW/permissions and community engagement completed
- Civil design: duct/pole, handholes, protection and route markers
- Splice plan, splitter design and ODF allocation documented
- OTDR traces and optical power budget recorded before acceptance
- GIS/as-built updated with closure, pole, duct and fiber IDs
- Spares: cable, closures, patch cords, ONTs, splitters, ODF modules
- Maintenance process: fault dispatch, escalation, restoration SLA and post-repair documentation

IPv6 deployment checklist

Deployment means production traffic, customer support and NOC visibility—not only address allocation.

- Prefix plan: aggregation, POPs, loopbacks, infrastructure, customers, DHCPv6-PD pools
- BGP: announce IPv6 prefixes, create route6 objects/ROAs, test transit and peering
- Core/access: dual-stack on routers, OLTs/BNG, management systems and monitoring
- CPE: firmware, RA/SLAAC/DHCPv6-PD behavior, firewall defaults and customer guides
- DNS: recursive resolver IPv6 reachability, authoritative AAAA readiness where needed
- Security: ICMPv6 policy, RA guard where applicable, ND considerations, logs and abuse handling
- Customer validation: IPv6 tests, speed, latency, traceroute, app compatibility and support workflow

MegaMore BROADBAND

Selected references and standards

Use official and operational sources where possible.

- Africa Internet Summit 2026 / AfNOG: AIS'26 public information and AfNOG hosting information
- RFC 8200 Internet Protocol, Version 6 (IPv6) Specification, RFC Editor / IETF
- MANRS Mutually Agreed Norms for Routing Security: network operator actions
- PeeringDB interconnection database for networks, IXPs and facilities
- RIPE Atlas global active measurement platform using probes and anchors
- Operational practice references: BGP route filtering, RPKI/ROA, IPAM/source-of-truth, NOC runbooks, telemetry and incident postmortems

Speaker note: convert references into local operator practice. Standards and tools only matter when they change day-to-day network behavior.