BNGSOFT

COMPANY PRESENTATION · 2026

Your whole subscriber edge. One commodity server.

BNG, CGNAT, QoS, edge security and low latency in one XDP data plane — run as a fleet from one console, NOC2.

BNG · PPPoE / IPoE | CGNAT | IPv6 | Low latency | Edge security | NOC2

bngsoft.com · sales@bngsoft.com

Every figure in this deck is per server, measured unless labelled otherwise.

A typical subscriber edge is four or five boxes from four or five vendors

BNG / BRAS

CGNAT appliance

DDoS scrubber

QoS shelf

IPv6-transition gear

QoE probe

IPFIX exporter

Every box is its own world

Its own licence, CLI, support contract and failure mode. A fault becomes a meeting between vendors.

Every packet pays several times

Copied, queued and inspected again at each hop. Latency and cost add up; nobody owns the end-to-end path.

Scale means a chassis

The traditional answer is a carrier chassis with per-subscriber licences — or a cheap router that stops scaling at a few thousand PPPoE sessions with QoS on.

Sources: Start Here; One XDP Data Plane; Cisco ASR 9000 and Nokia 7750 SR Alternative.

Classified, policed, translated and queued once — in one program

One XDP data plane on commodity x86

PPPoE + IPoE BNG · RADIUS AAA + CoA · CGNAT NAT44 · Dual-stack IPv6 · DS-Lite · Per-subscriber hQoS · AQM · L4S marking · flow protection · Anti-spoof · per-victim DDoS · Compliance logging

Authorise once, enforce every packet

RADIUS authorises in software; the session's routing, QoS plan, NAT policy and security state go straight into the fast path.

Hardware you choose

1–2U x86 servers with standard Intel, Mellanox or Broadcom NICs, from your usual supplier. A few hundred watts instead of a chassis.

Scale out, not up

Add another inexpensive server instead of another line card. All gateways run as one fleet in NOC2.

Sources: Platform Overview; Carrier-Grade BNG and CGNAT Without the Carrier Price Tag.

PART 1

How it works

Two planes on one node, a hardened OS that boots from RAM, and an in-kernel data path that costs nothing when idle.

Control plane decides. Data plane enforces — at every packet.

Control plane — session service

PPPoE / IPoE termination

RADIUS auth · accounting · CoA · Disconnect

DHCPv4 + Option 82 · DHCPv6-PD

Address pools

FRR routing (BGP, OSPFv2)

SESSION AUTHORISED → WRITTEN INTO PINNED BPF MAPS ↓

Data plane — `bngxdpd`, one eBPF/XDP program in the NIC driver path

Per-subscriber forwarding

Source validation

DDoS + abuse

Firewall

Shaping + AQM / L4S

NAT

IPFIX / syslog export

Runs before the kernel network stack: no round-trip to software after authorisation. Kernel tools — `ip`, ethtool, bonding/LACP — keep working on the same ports.

NOC2 fleet layer

One agent per gateway; every gateway, subscriber and alert in one console. If NOC2 is unreachable, subscribers stay connected.

Sources: RADIUS AAA and CoA; Platform Overview; Full Feature List With Shipping Status; Air-Gapped Procurement.

Every function in one pass, in a fixed order

Parse
Ethernet · S-tag · C-tag · PPPoE · IP

›

Subscriber lookup
plan, pool, policy

›

Source validation
BCP38 / uRPF

›

DDoS + abuse
per victim, per vector

›

Firewall
rules, portal, DNS block

›

Shape + AQM
hQoS · ECN mark · drop

›

NAT
port blocks, logged

›

Uplink
redirect to WAN

Full-CGNAT mode

NAT and firewall run in XDP too. **0 kernel conntrack entries** at 128,000 subscribers, and **24–36% less CPU** than leaving NAT to the kernel. Measured.

Every CPU core works

NICs cannot hash inside PPPoE, so normally one core carries all of it. Full-Core PPPoE opens the envelope and spreads flows across all cores — no reordering.

One program, not a chain

Merging the data path into one XDP program took data-plane CPU from **23% to 2.5%** on the same server. Measured.

Sources: Full Feature List; OrionOS 7.00 and 7.05 Release; Full-Core PPPoE; Carrier-Grade BNG and CGNAT Without the Carrier Price Tag.

Why in-kernel XDP, not DPDK and not an appliance

Versus DPDK / VPP

- "40+ Mpps per core" headlines are 64-byte L2 with no state. VPP's own published CGNAT44 figure is **6.9 Mpps per core** — real subscriber state is what costs.
- A DPDK poll core reads 100% busy at any load. XDP cost follows traffic.
- No hugepages, no separate routing table: XDP reads the kernel FIB directly.

Versus chassis and CPU routers

- Chassis: purpose-built silicon, per-subscriber licences, several kW. A chassis still wins on single-system capacity — we say so.
- CPU-forwarding routers: fast paths bypass queues and NAT state, so turning on BNG features forces the slow path.
- We are judged on what subscribers feel: latency under load, not price per gigabit.

Native XDP drivers

Intel i40e (X710/XL710), ice (E810), ixgbe (82599/X520/X550) · Mellanox mlx5 (ConnectX-5/6/6 Dx) · Broadcom bnxt_en. Generic mode elsewhere for labs.

Sources: XDP vs DPDK for a BNG; Cisco/Nokia Alternative; MikroTik Alternative; NIC and Driver Compatibility. Competitor figures are vendor-published or community-reported.

What one server carries

128,000
PPPoE subscribers on one server, QoS-only and full-CGNAT modes (OrionOS 7.05 lab)

344 s
from zero to 128,000 online — ~372 new subscribers per second, 0 flaps

~2 of 64
CPU threads busy at 128,000 subscribers with 11 Gbit/s of traffic

50 Mpps
full-feature engine ceiling per box — 33.6 Gbps even at 64-byte frames

~975 ns
per packet on the full-CGNAT path, measured over ~41M packets per arm

4,400+
live subscribers on a 2016 six-core server at evening peak; busiest core 72%, zero loss

252k+
concurrent NAT sessions at ~3% CPU on a 48-core standalone CGNAT node

3 ppm
download loss at the 128,000-subscriber plateau, all streams flowing

Lab server for 128k: AMD EPYC 9375F (32 cores / 64 threads), 128 GB, 2× E810 100G. Per-node design limit 131,072 subscribers. Planning figure for production: ~64,000 subscribers on a 2×100G box at ~3 Mbps busy-hour demand.

Sources: OrionOS 7.05 Release; Packet Rate Is the Number; bngxdpd Datasheet; 4,400 Subscribers on One 2016 Server; Standalone CGNAT.

PART 2

What it does

Subscriber access, address sharing, IPv6, QoS, latency, security, upgrades — then running all of it as a fleet.

Subscriber access, standards-based, on one box

Access

- PPPoE (PAP/CHAP) and IPoE (DHCP) at the same time
- 802.1Q and Q-in-Q (802.1ad) — past the 4,094-VLAN limit; VLANs created on first packet
- DHCPv4 server with Option 82, relay, DHCPv6-PD, IPv6-unnumbered
- Bonding / LACP; static-public and CGNAT subscribers side by side

AAA

- RADIUS authentication and accounting (Start / Interim / Stop, gigawords)
- CoA and Disconnect-Message: change rate, suspend, reactivate — no reconnect
- Your existing RADIUS is the contract: same attributes, same billing
- Local secrets mode when there is no RADIUS

Three models, one image

- **A** — subscriber gateway, no NAT
- **B** — full data plane: BNG + CGNAT + security (default)
- **C** — standalone translation node in front of any BNG or router
- Stateful XDP firewall, captive portal / walled garden, DNS redirect and blocking

Sources: RADIUS AAA and CoA; QinQ With CGNAT; Full Feature List With Shipping Status.

Carrier NAT44 in the fast path — with logs a regulator can use

Engine

- Port-block allocation; full-cone (EIF), symmetric, port-preserving and 1:1 modes
- Hairpinning, PCP / NAT-PMP, exempt destinations (IPTV, on-net), MSS clamp
- A port block is never shared by two subscribers; idle ports reclaimed under pressure without cutting a live connection

Logging and attribution

- One record per port block, not per flow: **~900× fewer log records**
- Deterministic mode (RFC 7422) needs no logs at all
- IPFIX and syslog from the same events; the account name rides in the record, so *public IP:port:time → customer* is one query

Keeping the shared IP clean

- Pool steering: every public IP scored 0–100; subscribers move off pressured or blocklisted addresses
- Abuser isolation: a scanner or proxy-rental line moves to a penalty IP and returns when the hold ends
- Open NAT for gaming; NAT rotate on reconnect

One command

`bngxdpctl check sub <ip|user>` shows a subscriber's public IP, pool, NAT type, port budget, drops — and a Symptom → Cause → Action verdict.

Sources: CGNAT Software; CGNAT Logging for Lawful Intercept; Subscriber Attribution; Pool Steering; CGNAT Abuse; Per-Subscriber CGNAT Troubleshooting; OrionOS 7.05.

Dual-stack at line rate, and a path to switching IPv4 off

SHIPPING

Dual-stack and DS-Lite

- Each subscriber: an IPv4 identity (CGNAT or static) plus an IPv6 /64 and a delegated /56 or /60
- IPv6 forwarding, shaping, AQM / L4S, firewall and anti-spoofing all in XDP
- Native IPv6 uses no CGNAT port blocks — every v6 byte relieves the pool

AVAILABLE FOR EVALUATION

MAP-T, MAP-E and NAT64

- Selected by one key, `cgnat.engine`, in the same data path
- MAP-T / MAP-E border relay: IPv4 and port-set computed from the IPv6 prefix — no session table to exhaust
- NAT64 (RFC 6146) sharing the CGNAT pool; needs your DNS64
- Implemented and in the image; not yet field-proven

The arithmetic

IPv4 costs $20–45 per address (2026). For 10,000 subscribers, one address each is $200k–$450k; at 40–60 subscribers per address it is ~167–250 addresses — about **97% less**.

Sources: IPv6 Dual-Stack and DS-Lite; IPv4 Exhaustion; MAP-T and MAP-E; NAT64 and DNS64; IPv6 Rollout Sequence; Full Feature List.

Hierarchical shaping without a QoS appliance

Global cap
licensed capacity

›

VLAN group
up to 64 groups — reseller caps

›

Subscriber
per class, both directions

Nine classes per subscriber

A default plus eight operator-defined classes by destination prefix and port, each with its own rate and counters from RADIUS.

Plans that sell

Burst (factor + byte ceiling), off-peak boost schedules, up to four peak-period congestion windows, per-VLAN limits.

Changes without redial

Rate changes by RADIUS CoA or a runtime override — applied in the fast path, no reconnect, no per-subscriber QoS licence.

Sources: Hierarchical QoS per Subscriber; Full Feature List With Shipping Status. Classes are destination- and port-defined; we do not do DPI.

"Slow" is usually queueing, not capacity

An idle line answers in ~1 ms; a full one in tens of milliseconds. We fix the queue at the gateway — no new hardware, nothing on the customer's side.

AQM

Per-subscriber queue delay measured continuously; drops start gently before the buffer bloats, steered to a latency target.

L4S-style marking

ECN-capable flows are marked instead of dropped (RFC 9332 coupling) — senders slow down with no loss.

Interactive Flow Protection

DNS, connection setup, ACKs, game and voice packets never wait behind a download. On by default.

Fair dropping

When a line congests, slow the flow causing it — not the video call. Off → observe → enforce.

~0.6ms
average queuing latency held under load on one node

6M+
no-loss ECN marks on one node in a 2-hour peak

1.23M
interactive packets protected in 40 seconds on one node

Sources: Low-Latency Broadband Suite; L4S and AQM on the BNG; Interactive Flow Protection; Bufferbloat. Single-queue L4S signalling; we do not claim a DualQ.

Stop the attack. Keep the customer online.

Inbound: per-victim DDoS

- Traffic classified per victim and per attack vector; only unsolicited excess is dropped
- Never keys on source IP — CDNs and other carriers' CGNAT can't be blocked by accident
- Surgical blocking: only the exact destination IP:port:protocol under attack, armed and released automatically

Outbound: your own subscribers

- Anti-spoofing (BCP38 / MANRS): each line may only send from its own IPv4 and delegated IPv6
- Botnet containment: volumetric, scanner and fan-out detectors rate-limit the infected CPE at its own port
- Proxy-rental lines detected and isolated to a penalty IP

Safe rollout

Every protection runs **off → observe → enforce**, node by node: full measurement, zero packets dropped in observe mode, one command to reverse. Packet capture and a bounded decision log show why anything was dropped.

Sources: Per-Victim DDoS Mitigation; DDoS Protection at the Subscriber Edge; Anti-Spoofing and uRPF; Stop Botnet Traffic; Block the Attack Without Blocking the Customer; Eleven Feature Switches.

Know which line is bad — and where — before the customer calls

Flow Intelligence

- Passive per-flow loss and round-trip time, split into **access** (your network) and **transit** (the internet)
- QoE score 0–100 per subscriber; each line's own normal learned as a baseline
- No probes, no CPE agents

Fault localisation

- Reads the PPPoE keepalive every CPE already sends: **94%** of sessions carry the signal
- Places the fault on a segment — home, access, or us
- Answers the practical question: send a technician, advise the customer, or don't know yet

Lines nobody complained about

Loss and reordering tracked over many windows. One real line ran at **30.96% loss** against a 0.50% box median — with no ticket raised.

Sized on measured demand

Busy-hour demand measured per subscriber — median **3.21 Mbps** — replaces contention-ratio guesswork for uplink sizing.

Sources: Per-Subscriber Flow Intelligence; Autonomous QoE; Fault Localisation; Find the Broken Line; Should You Send a Technician?; Uplink Sizing Without Contention Ratios.

A hardened carrier Linux that boots from RAM

Immutable

Boots a signed image into a RAM root. A reboot always returns to a known-good state. Kernel, XDP and the BNG — and little else.

Reproducible

One artefact, identical on every node. No ad-hoc package installs; the data-plane binary only ever comes from the boot image.

Patched on a dated record

We own the OS and the NIC drivers. Security fixes ship in days, with a public dated release history.

We fix the driver

When a NIC driver crashed a live router, we root-caused it to one line and hardened **12 defects** across the 10/40/100G family — no downgrade, fast path untouched.

Zero-touch install

Create an install code in NOC2; the live installer sets up the server over remote console or network boot and enrols it in the fleet.

Sources: OrionOS: a Hardened Carrier Linux; What Is in the BNG Image; When the NIC Driver Crashes the Router; OrionOS 7.0x Release.

Upgrade in the middle of the afternoon

1,601
subscribers before and after a forwarding-plane swap — traffic continuous

0.83 s
session-layer gap on upgrade, without re-authentication (was 13.3 s)

24–29 s
for 128,000 sessions to come back after a live upgrade — 0 flaps, same Acct-Session-Ids

kill -9
of the session service: every session stays up, no Accounting-Stop

How

- New program loaded and verified against the live tables, then swapped atomically — replace, never detach first
- Most upgrades are invisible; the rare deeper one is ~30 s with sessions and public addresses kept
- Not dual-supervisor chassis ISSU — and we don't call it that

After the power cut

- Reconnect storm is the real test: 10,000 redialling subscribers back in under a minute
- A full 128,000-subscriber box online in 344 s with 0 flaps
- Soak: 8,000 subscribers × 4 h with constant churn — 0 crashes, 0 leaks

Sources: ISSU for a Software BNG; Zero-Downtime Upgrade; BNG Resilience; OrionOS 7.00 and 7.05 Release.

One control plane for every gateway

An agent on each gateway, no probes. One console on desktop, phone and wall-board.

Triage by impact

Problems ranked by subscribers affected. One search — username, IPv4, IPv6 or MAC — returns a verdict: **the subscriber, the line, or us**, with reconnect history and a live read.

Live, from the data plane

A 5-second live feed across gateways. Session Info shows any subscriber in seven views — CGNAT, flows, QoE, drops, protection — read live every 1–5 s.

Check-Engine

One command scans the whole box — licence, RADIUS, upstream, data path, CGNAT, QoS, security, NIC — and isolates the fault with Symptom → Cause → Action.

Internet Observatory

Link by link from subscriber line to destination: is it the internet, a service, or your access network? Popular services watched against their own baseline.

Config you can trust

Drift detected and diffed, one-click restore, unsaved changes counted. Written, running and enforcing shown as three separate states.

Safe actions

Audited, reversible, time-limited rate changes; bulk jobs; capacity forecasts with a saturation date per resource.

Sources: NOC2: One Control Plane; Live BNG Console; Session Info; One-Search Triage; One Command That Tells You What Is Wrong; Internet Observatory 3; Config Drift; Running State vs Config.

Fewer calls, cleaner hand-offs, room to grow

Self-check and alerts

- Your branded page where customers check their own line — no password at home
- One alert when an incident starts, one when it ends (Telegram, email)
- Public status page; monthly report

Operator Insights

- A to-do board with a button per fix
- Finds subscribers capped by a 100M router port
- Capacity forecast per VLAN; who to call first

Wholesale, multi-tenant

- Platform → reseller → subscriber portal
- Each operator: own gateways, logins, portal and status page
- No cross-tenant visibility in lists, searches, alerts or figures

Scale-out cluster — `bngclusterd`

Gateways share load, capacity and health; each delays its PPPoE answer by its rank, so new sessions land on the lightest box. Peer loss noticed in 6 s. It places new sessions — it does not move live ones, and it is not active/standby HA. Starts in observe mode.

Sources: NOC2 Self-Check & Alerts; NOC2 Operator Insights; Wholesale Multi-Tenant BNG; BNG Cluster.

Two ceilings: packets through the NIC, subscribers through CPU and RAM

REFERENCE BUILD	NIC	TRAFFIC	SUBSCRIBERS	POWER	HARDWARE / SUB
Edge · 1U	Intel X710	**~20G**	**~6.7k**	**~250 W**	**~$0.37**
Standard · 1U	Intel E810	**~100G**	**~33k**	**~380 W**	**~$0.15**
High density · 2U	2× Intel E810	**~200G**	**~64k**	**~650 W**	**~$0.14**
SmartNIC · 2U	ConnectX-6 Dx	**~200G**	**~64k**	**~600 W**	**—**

Sizing rules we use

- Packet rate first; Gbps follows from frame size
- Watch PCIe: an E810-CQDA2 has two 100G ports but one controller — 100G total
- Subscribers × measured busy-hour Mbps, not contention ratios; core counts come from the PoC

Illustration: 50,000 subscribers, 36 months

Hardware + power, before licence. Router-appliance fleet: ~19 boxes, ~18U, ~$65,000 (~$1.30 per subscriber). BNGSOFT: 3 boxes, 3U, ~$20,000 (~$0.41 per subscriber). The ~$45,000 difference is the headroom the licence comes out of.

Sources: BNG Server Sizing; Packet Rate Is the Number; OrionOS 7.05 §12; Start Here. Reference builds and costs are planning figures, not measurements.

Built for the networks that need it most

Fixed broadband ISPs

PPPoE / IPoE fibre and cable operators replacing an ageing BRAS or a stack of CPU routers — and adding CGNAT, QoS and protection without more boxes.

Satellite and fixed wireless

On LEO, idle RTT of 25–60 ms jumps 2–4× under load. That is the gateway queue, not the orbit — and the gateway is where we fix it. Every satellite operator is a CGNAT operator.

Wholesale and resellers

One platform, many operators, structural isolation; each reseller with its own caps, logins, portal and status page.

Sovereign and air-gapped

Ed25519-signed updates with the signing key on your server, one switch to freeze updates, role-scoped access. Subscribers stay connected even if our console disappears.

Sources: Satellite Ground-Segment BNG; Satellite Latency; Wholesale Multi-Tenant BNG; Air-Gapped BNG Procurement.

What ships, what is in validation, what we don't claim

ACTIVE ON DELIVERY

PPPoE / IPoE, RADIUS + CoA, DHCPv4/v6-PD, Q-in-Q, LACP · per-subscriber hQoS, 9 classes · Interactive Flow Protection · dual-stack IPv6 · health engine, baselines, self-test · zero-downtime upgrade

CONFIGURED AT COMMISSIONING

CGNAT and compliance logging, RFC 7422 · AQM, L4S-style marking, Flow Intelligence, QoE · anti-spoofing, per-victim DDoS, outbound abuse containment, firewall, captive portal, DNS blocking

IN VALIDATION

MAP-T, MAP-E and NAT64 (implemented, not field-proven) · arena conntrack engine · Autopilot master switch

ROADMAP

BGP border-router role (no MPLS, SRv6, IS-IS or BFD)

NOT CLAIMED

DPI / application identification · SIP ALG · chassis active/standby HA · REST API · Prometheus endpoint · DualQ L4S · install on an existing Linux

Every protective feature is switched **off → observe → enforce**, node by node, and can be reversed with one command. The full list, with conditions, is in *Full BNG and CGNAT Feature List, With Shipping Status*.

Source: Full BNG and CGNAT Feature List, With Shipping Status.

Prove it on your hardware, with your traffic, beside what you run

1 · Size
Tell us subscribers, peak traffic and functions

›

2 · Proof of concept
On your x86, beside the incumbent, carrying no subscribers

›

3 · One VLAN
A few hundred to a few thousand subscribers

›

4 · VLAN by VLAN
POP by POP, at the pace the evidence supports

›

5 · Retire
The old box goes last

BNG migration

- Your RADIUS stays: same attributes, CoA, accounting
- Vendor-specific surprises surface in step 3, not step 5
- Rollback = point the VLAN back

CGNAT migration — no flag day

- Canary one small public subnet, then shift pools subnet by subnet
- The switch is a BGP withdraw / re-advertise; rollback is the reverse
- No maintenance window

Sources: BNG Migration Plan; CGNAT Migration Without a Flag Day; Who Owns the Bug Below Your Software.

BNGSOFT

Let's size your network.

Send your subscriber count, peak throughput and the functions you need. We size it, quote it, and hand you a build to test — a 30-minute demo on request.

sales@bngsoft.com | bngsoft.com | bngsoft.com/brochures

All figures per server; lab measurements are labelled with their release. Competitor figures are vendor-published or community-reported.

© 2026 BNGSOFT