# NAT, STUN, TURN Deployment Checklist

Companion to: NAT, Firewalls, STUN, TURN, ICE — How WebRTC Reaches a Phone

## 1 · Ports & transports

*Open every transport the spec defines, then verify in production.*

- ☐ STUN: UDP/3478 open inbound and outbound at every TURN host.
- ☐ TURN: UDP/3478, TCP/3478, TLS over TCP/5349 all reachable.
- ☐ turns://host:443 over TLS — the universal corporate-firewall fallback.
- ☐ Verify a real client picks UDP first, falls back to TLS-on-443 cleanly.
- ☐ Block direct egress from the TURN host except to the public internet.

## 2 · Credentials (RFC 8656 §9.2)

*Short-lived HMAC credentials minted server-side per call.*

- ☐ Issue ephemeral usernames (e.g. unixtime:userid) on the application server.
- ☐ Compute password = base64(HMAC-SHA1(username, shared_secret)).
- ☐ Set credential lifetime ≤ 24 h, ideally minutes for sensitive verticals.
- ☐ Rotate the TURN shared secret on a regular cadence; keep an N-1 grace window.
- ☐ Never expose long-lived TURN credentials in client-side code or HTML attributes.

## 3 · Geographic placement

*Latency added by TURN equals client-to-TURN RTT, doubled.*

- ☐ Deploy at least one TURN cluster per continent your users live on.
- ☐ Use anycast IPs or geo-DNS to steer clients to the nearest cluster.
- ☐ Target a P75 client-to-TURN RTT of ≤ 40 ms in each region.
- ☐ Watch the long tail: rural mobile, Starlink, and CGNAT users may add 80–120 ms.
- ☐ Cap allocation lifetime; reclaim idle slots after ≤ 10 minutes of silence.

## 4 · Observability KPIs

*Instrument every call and every TURN allocation; ship the dashboard.*

- ☐ ICE-connected rate per network type (Wi-Fi / 4G / 5G / corporate).
- ☐ TURN-relay fraction — expect 15–20 % for consumer, ~100 % for AI / IoT.
- ☐ Median time to selected candidate — target ≤ 500 ms on Trickle ICE.
- ☐ TURN bandwidth per call — relay traffic = 2 × call bitrate.
- ☐ Consent freshness failures (RFC 7675) — alert at 0.5 % session drop rate.

## Worked TURN bandwidth model — 1,000 concurrent calls, 20 % relayed

200 relayed calls × 3.2 Mbps per call (1.6 Mbps × 2 directions) = 1.28 Gbps sustained at TURN. Monthly egress ≈ 410 TB at a 12.5 % busy-hour ratio.

Cost on AWS EC2 egress at blended ~$0.054/GB above 150 TB: ≈ $22,000 / month. Same throughput on a colocated bare-metal cluster with $0.50–$1.00 per Mbps/month transit: ≈ $700–$1,400 / month. The 20× gap is why every WebRTC product at meaningful scale uses bare metal or a TURN provider whose pricing reflects bare-metal economics.

Controlling specs: RFC 8489 (STUN, Feb 2020) · RFC 8656 (TURN, Feb 2020) · RFC 8445 (ICE, Jul 2018) · RFC 8838 (Trickle ICE, Jan 2021) · RFC 4787 (NAT behavioural requirements) · RFC 7675 (consent freshness).