Pipelining internet security · *Ben Smyth · 30 April 2026*

Three thousand mandatory clauses demanded by security-relevant internet protocols, chewing on wire messages with vocabularies of roughly 2^64 possibilities (excluding crypto), compounding into handshake sentences spanning 2^256 shapes. Critical infrastructure rides on this complexity. ‘Good’ isn’t enough. AI is our rude awakening. Billions spent, decades of effort, specifying, auditing, deploying these protocols. Nobody has ever checked they’re right. They’re not. Four hundred and fifty-three causes for concern raised. You break it, you fix it:

L1 **RFC: English-prose governing**
Kernel level : Ethernet, VLAN, ARP, IP, ICMP, NDP, TCP, UDP, fragment reassembly, eBPF
Transport layer : TLS, QUIC, SSH, DTLS, IKEv2
Application layer : HTTP, SMTP, DNSSEC, LDAP, OAuth, Kerberos, MQTT, …

L2 **EXTRACT: Normative clauses**
Protocol specifications mark mandatory requirements with specific keywords — MUST, MUST NOT, etc. A rule-based extractor pulls each clause into a structured row. Thousands of pages of English prose become twenty-nine hundred named obligations.

L3 **ENRICH: RFC as type system**
Each clause becomes a type rule. “Implementations SHOULD bound memory when decompressing,” length refinement, NSS miss: Sixty-four kilobyte wire pins sixteen megabytes, one page load, sixteen-hundred connections BURST through thirty-eight gigabytes — remote kill, Firefox dead.

L4 **INVERT: Probe generation**
Each type rule gets one probe: The smallest wire input that violates it. “Server MUST NOT send a second HelloRetryRequest.” Compliant servers abort; Java loops FOREVER. Surprisingly, one probe is enough — we catch any implementation silently admitting direct-literal rulebook violation.

L5 **COVERAGE: Parametrise probes**
Parametrise probes across six bad-value classes: Sentinel, adjacent-illegal, overflow, wrap-around, sign, upper-boundary. Bertie key_share length 0x8000 caught on sign: size_t cast reads eighteen exabytes, server crashes.

L6 **COMPLETENESS: Close the gap**
Probes catch violations of individual rules; augmentation captures combinations. Non-literals catch: Subtle-valid-parse, cross-constraint, stateful, ordering, multi-message. “ProtocolNameList MUST be non-empty,” Botan misses, subtle-valid-parse catches.

Phase one, done. Exponential handshake depth to linearity; every mandatory clause covered. Implementations reintroduce complexity: Twenty-two stacks, one hundred and eleven production codebases, fourteen languages, one-hundred million lines of code, all from one English-prose RFC. Abstraction handles:

L7 **OBSERVE: Read the wire**
Each implementation runs inside an observation harness — bidirectional relay for transport-and-above, raw sockets plus in-kernel tripwires for kernel. Wire bytes are ground truth.

L8 **VERIFY: Layered type checks**
Bytes parse as wire messages, fields satisfy RFC type rules, values pass runtime type schema.

L9 **PROBE: Sweep the fleet**
Run probes against every implementation under every configuration — 2755 parametrised/augmented probes, a few hours on a laptop — four hundred and fifty-three findings.

First nine layers mechanical; intelligence drives verdict: Probe hands agent wire message forcing non-compliance; trace to use, evaluate attack surface, assess CVSS impact; skip low-severity vulnerabilities, weaponise the rest. Verdict remains: Of every flagged finding, is there security impact?

L10 **CLASSIFY: Trace to use**
When probe breaches implementation, chase through source to consumption site, classify by CVSS.

L11 **WEAPONISE: Attack team**
Build runnable PoC exploit, framed as security game from published cryptographic literature.

L12 **ADJUDICATE: Judge fires**
Judge predicate from literature deterministically evaluates attack validity.

L13 **REVIEW: Disclosure or post-mortem**
Human verdict: System failure, or critical-infrastructure vulnerability?

I’ve considered TLS and QUIC, twenty arena findings, fourteen responsibly disclosed, three confirmed CVE-2026-6772,32883,35582, six under my review, four hundred and thirty-three unopened.