TECHNICAL NOTE · PERFORMANCE

What does a OneFirewall blocklist cost your firewall?

We measured it. Checking an address against a 300,000-entry blocklist took at most 53.0 ns (billionths of a second), once per new connection. Keeping the list up to date takes a few kilobits per second.

53.0 ns	0.00037%	3.8 kbit/s
slowest lookup against a 300,000-entry list	of one 14.3 ms network round trip within Europe	average download for a 10,000-entry list refreshed every 5 minutes

TO SCALE



The bar is 170 mm long. Drawn to the same scale, the lookup would be 0.63 thousandths of a millimetre wide, far too thin to print. Put another way: if the round trip lasted a whole day, the lookup would take 0.32 s. If it were a one-kilometre walk, the lookup would be the last 3.7 mm.

WHAT IT COSTS

Per new connection
at most 53.0 ns

Per million new connections
53 ms of processor time

Download traffic
3.8 kbit/s (10,000 entries, every 5 minutes)

Traffic path
unchanged; nothing passes through OneFirewall

WHAT IT STOPPED IN THE FIELD

35,000 attacks a day
eliminated for a SaaS platform within 24 hours. Its latency fell by 28%. [11]

92.6% of attacking IPs
in the PHP2Shell exploitation burst were already known to OneFirewall. They sent 99.96% of the 32,252 attempts. [12]

8,934 critical and high events
intercepted in 19 days at a web application, where 12% of the traffic its firewall allowed was malicious. [13]

Field figures come from our published case studies, not from this test.

> The trade is lopsided. The firewall spends a few dozen nanoseconds on each new connection, and in return connections from known attackers are dropped on their first packet.

1. How we measured it

The only per-connection work a OneFirewall feed adds is a lookup: is this source address on the list? The firewall already runs a policy check for every new connection, so the feed simply makes the address list it checks longer. That lookup is what we timed.

We wrote a small benchmark in C and ran it on an Apple M3 with 16 GB of RAM (macOS 26.7.1; Apple clang version 21.0.0 (clang-2100.1.1.101); -O2). It builds lists of random public IPv4 addresses, from 1,000 up to 1,000,000 entries, in the two forms firewalls commonly use. A hash table finds a single address in about one step. A sorted range table is searched by halving, needs about 20 steps at a million entries, and is the slower of the two.

Each configuration ran 16,777,216 lookups seven times, and we kept the median. Every lookup waited for the result of the previous one, so the processor could not overlap them; the times below are the full cost of a single lookup, including 0.3 ns of loop overhead. We repeated the whole run three times and report the slowest of the three.

Time for one lookup (slowest of three runs)

Entries	Hash, address not in list	Hash, address in list	Range, not in list	Range, in list	Memory (hash / range)
1,000	10.4 ns	4.7 ns	23.3 ns	10.2 ns	8 KB / 8 KB
10,000	5.8 ns	2.9 ns	20.4 ns	28.0 ns	131 KB / 80 KB
100,000	8.3 ns	4.0 ns	37.9 ns	47.5 ns	1.0 MB / 800 KB
300,000	6.0 ns	3.2 ns	39.0 ns	53.0 ns	4.2 MB / 2.4 MB
1,000,000	15.0 ns	5.3 ns	79.0 ns	63.1 ns	8.4 MB / 8.0 MB

“Not in list” is ordinary traffic, the common case; “in list” is traffic from a listed attacker. Hash-table times follow how full the table is (29–49% here), not the number of entries, which is why 1,000 entries came out slower than 10,000. 1 ns is a billionth of a second; 1 μs is 1,000 ns; 1 ms is 1,000,000 ns.

2. Keeping the list up to date

The device downloads the list as plain text, one address per line. We generated lists of random public addresses and measured them at 14.2 bytes per entry. The gzip column is for reference only: random addresses compress worse than real ones, and whether compression is used depends on the server and the device. The traffic table uses the uncompressed size.

List size

Entries	Plain text	gzip	Bytes per entry
10,000	142 KB	52 KB	14.2
100,000	1.4 MB	467 KB	14.2
300,000	4.3 MB	1.3 MB	14.2

Download traffic per day, with the average rate in brackets

Entries	Every 5 min	Every 30 min	Every hour	Once a day
10,000	41.0 MB (3.8 kbit/s)	6.8 MB (0.63 kbit/s)	3.4 MB (0.32 kbit/s)	142 KB (0.01 kbit/s)
100,000	410 MB (38 kbit/s)	68.3 MB (6.3 kbit/s)	34.2 MB (3.2 kbit/s)	1.4 MB (0.13 kbit/s)
300,000	1.2 GB (114 kbit/s)	205 MB (19 kbit/s)	102 MB (9.5 kbit/s)	4.3 MB (0.40 kbit/s)

The customer picks the refresh interval, anywhere from 60 seconds to 24 hours. It changes how fresh the list is and nothing else.

3. Why the lookup is paid once per connection

Firewalls that track sessions check the policy on the first packet of a connection and handle the rest from the session table. On FortiGate, accepted sessions are offloaded to the NP processor and later packets never go back through the policy [3]. The figures above are therefore added once per connection, not once per packet.

Slowest lookup as a share of a network round trip

Round trip (Verizon, October 2026) [10]	300,000 entries (53.0 ns)	1,000,000 entries (79.0 ns)
Within Europe, 14.3 ms	0.00037%	0.00055%
Within North America, 33.4 ms	0.00016%	0.00024%
London to New York, 70.4 ms	0.000075%	0.00011%

For comparison, Fortinet rates a complete firewall pass on a FortiGate 100F at 4.97 μs for 64-byte UDP packets [4], about 94 times our slowest lookup.

4. Limits set by each platform

Every platform has its own maximum list size and refresh options. OneFirewall caps each feed at the platform's limit, so a list never grows past what the vendor designed the device to hold.

Platform	Refresh options	Maximum list size	Other limits
FortiGate	1 to 43,200 min, default 5 [1]	300,000 IPs per device, FortiOS 7.4.4 and later [2]	Entries above the limit are not loaded [2]
Palo Alto Networks	5 min (default), hourly, daily, weekly, monthly [7]	50,000 IPs in total; 150,000 on PA-3200, PA-5200, PA-7000 [8]	Up to 30 lists [8]
AWS WAF (CloudFront, ALB)	Any; set by the sync schedule	10,000 per IP set [5]	An IP set rule costs 1 WCU at any size [5]; one update call per second per account and Region [6]
Check Point	One interval for all feeds, in seconds; 30 min by default on Quantum Spark R82 [9]	Not published	

5. What this test does not tell you

We ran it on a general-purpose computer processor, not on firewall hardware, and we did not measure any firewall model. Firewalls use their own data structures and many use dedicated chips, so their absolute numbers will differ from ours. We timed the lookup alone, not a packet's whole path through a firewall. Our addresses were random, while real lists and real traffic have patterns this test does not model. Round-trip times are Verizon's published measurements, not ours. OneDevice in Series and WAP sit in the traffic path by design and are outside the scope of this note.

The benchmark is about 250 lines of C and Python. It runs with one command and writes every result to a JSON file; we are happy to share it on request.

Sources

1. Fortinet, FortiOS 7.4.0 Administration Guide: Configuring a threat feed (refresh rate 1-43,200 min, default 5) · https://docs.fortinet.com/document/fortigate/7.4.0/administration-guide/379433/configuring-a-threat-feed
2. Fortinet, FortiOS 7.4.4 Administration Guide: Threat feeds (device-wide entry and file-size limits) · https://docs.fortinet.com/document/fortigate/7.4.4/administration-guide/9463/threat-feeds
3. Fortinet, FortiOS 7.4.4 Hardware Acceleration Guide (NP fast path: policy match on the first packet, session offloaded) · https://fortinetweb.s3.amazonaws.com/docs.fortinet.com/v2/attachments/6613da86-0277-11ef-8c42-fa163e15d75b/fortios-hardware-acceleration-744.pdf
4. Fortinet, FortiGate 100F Series data sheet (firewall latency 4.97 µs, 64-byte UDP; firewall throughput 20 Gbps)
5. AWS WAF Developer Guide: IP set match rule statement (up to 10,000 addresses, 1 WCU) · https://docs.aws.amazon.com/waf/latest/developerguide/waf-rule-statement-type-ipset-match.html
6. AWS WAF Developer Guide: AWS WAF quotas (5,000 WCU per web ACL, 1 Create/Put/Update call per second) · https://docs.aws.amazon.com/waf/latest/developerguide/limits.html
7. Palo Alto Networks, PAN-OS 11.1: Objects > External Dynamic Lists (refresh: five minutes by default, hourly, daily, weekly, monthly) · https://docs.paloaltonetworks.com/ngfw/help/11-1/objects/objects-external-dynamic-lists
8. Palo Alto Networks, PAN-OS 11.1 Admin Guide: External Dynamic List (IP capacity per model, 30 EDLs) · https://docs.paloaltonetworks.com/pan-os/11-1/pan-os-admin/policy/use-an-external-dynamic-list-in-policy/external-dynamic-list
9. Check Point, Quantum Spark R82.00.X Admin Guide: IoC Feeds (fetched every 30 minutes by default); sk132193 Custom Intelligence Feeds · https://sc1.checkpoint.com/documents/Appliances/Quantum_Spark_R82.00.X/AdminGuides_Locally_Managed/EN/Content/Topics/IoC_Feeds.htm
10. Verizon, Monthly IP Latency Data (October 2026: Europe 14.263 ms, North America 33.354 ms, Trans-Atlantic 70.440 ms) · https://www.verizon.com/business/terms/latency/
11. OneFirewall Alliance, SaaS platform case study · https://onefirewall.com/case-study
12. OneFirewall Alliance, PHP2Shell threat briefing · https://onefirewall.com/php2shell-threat-briefing
13. OneFirewall Alliance, Web application case study · https://onefirewall.com/case-study-web-application

---

Prepared by OneFirewall Engineering, London · support@onefirewall.com · https://onefirewall.com/performance-impact