FORTIGATE CONFIGURATION MISCONFIGURATION RISK ANALYSIS

Belsen Group Dataset — Country-Level Profiling

15,474 Devices | 146 Countries | March 2026

Dragon Advance Tech Consulting **(with Claude.ai)**

CONFIDENTIAL — FOR SECURITY RESEARCH & DFIR USE ONLY

Table of Contents

1. Executive Summary

This report presents a country-level misconfiguration risk analysis derived from 15,474 FortiGate firewall configuration files spanning 146 countries. The dataset originates from the Belsen Group disclosure (January 2025), in which configurations were exfiltrated by exploiting CVE-2022-40684 — a pre-authentication administrative bypass vulnerability affecting FortiOS 7.x. The analysis maps 21 misconfiguration indicators per device and aggregates findings at the country level to identify regional security posture patterns relevant to threat intelligence, incident response scoping, and detection engineering.

The headline findings are severe:

- 99.9% of devices (15,452) remain in the CVE-2022-40684 vulnerable FortiOS version window — the same vulnerability exploited to collect this dataset.
- 95.7% of devices expose HTTP management on internet-facing interfaces, enabling plaintext credential capture.
- 91.5% of devices have no configured password policy, and 87.1% retain the default 'admin' account.
- 82.6% of devices have permissive Any/Any/Any firewall rules, indicating flat network architectures with minimal segmentation.
- 78.5% of devices have no syslog or FortiAnalyzer configured, meaning breach activity would generate no off-device telemetry.

2. Dataset and Methodology

2.1 Dataset Provenance

The Belsen Group, a previously unknown threat actor, published 15,474 FortiGate configuration files on January 14, 2025. The configurations were acquired in 2022 via CVE-2022-40684, a critical (CVSS 9.8) authentication bypass in the FortiOS HTTP/HTTPS administrative interface affecting versions 7.0.0–7.0.6 and 7.2.0–7.2.1. Each configuration file contains the complete running configuration of the device at time of extraction, including system settings, interface configurations, VPN parameters, firewall policies, and — in many cases — hashed administrator credentials.

The dataset was structured as a directory tree with ISO 3166-1 alpha-2 country codes as top-level directories, with each device identified by its public IP address and management port. This structure was preserved in the analysis to enable direct country-level correlation.

2.2 Misconfiguration Schema

Each configuration was parsed using a purpose-built Python extractor (**by Claude.ai**) targeting the following indicator categories:

Category	Indicators Extracted	Risk Weight
Admin Interface Exposure	HTTP/HTTPS/SSH/Telnet/FGFM/SNMP on WAN-facing interfaces	2–3 pts
Authentication Controls	Default admin account, password policy, reuse restriction, banners	1–2 pts
SSL-VPN Configuration	Enabled status, split tunneling, cipher strength, source restrictions	1–2 pts
Firewall Policy Quality	Any/Any/Any permissive policies	1 pt
Logging & Visibility	Syslog / FortiAnalyzer configuration	1 pt
CVE Version Exposure	FortiOS version mapped to CVE-2022-40684, CVE-2023-27997, CVE-2024-21762	2 pts/CVE

A composite risk score (0–25+) was computed (**by Claude.ai**) per device by summing weighted indicator scores. Device scores were aggregated to produce per-country averages, maximums, and prevalence percentages. Only countries with five or more devices were included in the ranked country analysis to minimize statistical noise from single-device outliers.

3. Key Findings

3.1 CVE Exposure — Systemic Unpatched Vulnerability Surface

The dataset was collected via CVE-2022-40684 exploitation, yet 99.9% of the 15,474 devices remain in the vulnerable FortiOS version range as of the configuration snapshot. This indicates the vast majority of operators did not apply the emergency patches released by Fortinet in October 2022 (PSIRT FG-IR-22-377), or performed patch cycles significantly later. Critically, the SSL-VPN heap overflow vulnerabilities CVE-2023-27997 and CVE-2024-21762 — disclosed well after this dataset was captured — would also apply to 99.1% of devices given their FortiOS version ranges, representing the extended tail-risk of devices that remain unpatched across multiple generation cycles.

CVE	Description	CVSS	Affected Devices	% of Dataset
CVE-2022-40684	Admin auth bypass (HTTP/HTTPS)	9.8	15,452	99.9%
CVE-2023-27997	SSL-VPN heap overflow pre-auth RCE	9.8	15,335	99.1%
CVE-2024-21762	SSL-VPN OOB write pre-auth RCE	9.6	15,335	99.1%

3.2 Administrative Interface Exposure

95.7% of devices (14,806) allow HTTP access on internet-facing interfaces. While FortiOS may redirect HTTP to HTTPS, this configuration exposes devices to credential interception via SSL stripping, HSTS bypass[1], or downgrade attacks. In addition, 78.2% of devices (12,099) expose the FortiManager fabric management protocol (FGFM, TCP/541) ***on public interfaces*** — a high-value vector for lateral movement if FortiManager instances are also reachable. ***SSH exposure on WAN*** interfaces affects 34% of devices, and ***Telnet — a cleartext protocol*** — remains enabled on 2.5% (386 devices).

3.3 Authentication and Credential Hygiene

87.1% of devices retain the ***default 'admin' superuser account***, providing a predictable username target for credential attacks. ***Password policy enforcement*** is absent on 91.5% of devices — meaning minimum length, complexity, and lockout thresholds are unenforced. 281 devices (1.8%) contain ***login banners*** explicitly advertising instance-ID-based default passwords, creating a trivially exploitable credential pattern for cloud-deployed FortiGate-VM instances where the instance ID is often discoverable via cloud metadata APIs.

3.4 SSL-VPN Attack Surface

99.1% of devices (15,335) have SSL-VPN enabled, consistent with the enterprise remote access use case that dominates this device segment. Given that CVE-2023-27997 and CVE-2024-21762 are both pre-authentication SSL-VPN RCEs affecting the same FortiOS version range, this represents a substantial unauthenticated remote code execution surface. ***SSL-VPN configuration*** quality is additionally poor: source-address restrictions are rarely implemented, and split tunneling enables attacker ***pivoting through the VPN tunnel*** rather than all traffic being forced through the corporate perimeter inspection stack.

3.5 Firewall Policy and Network Segmentation

82.6% of devices (12,780) have at least one accept firewall policy matching source 'all', destination 'all'. This suggests the use of FortiGate primarily as an SSL-VPN terminator or basic NAT gateway rather than as a stateful inspection device with meaningful policy enforcement. Flat network architectures indicated by these policies would offer minimal lateral movement resistance to an attacker post-exploitation.

3.6 Logging and Detection Visibility

78.5% of devices (12,148) have ***neither Syslog nor FortiAnalyzer configured***. For these devices, all authentication attempts, VPN connections, firewall policy hits, and administrative actions exist only in the device's local circular log buffer — typically overwritten within days. This represents a systemic blind spot: even if exploitation occurred, no off-device forensic evidence would exist unless memory forensics was performed on the device itself. This finding directly impacts DFIR triage prioritization for any investigation involving FortiGate devices.

[1] https://www.1kosmos.com/resources/security-glossary/ssl-stripping

4. Country-Level Analysis

4.1 Highest Risk Countries (≥5 devices)

The following countries exhibit the highest average composite risk scores across their device populations:

Rank	Country	Devices	Avg Risk	HTTP WAN %	Telnet %	SSH WAN %	CVE-40684 %
1	ML	8	20.38	100%	37.5%	62.5%	100%
2	HN	5	20.20	100%	0%	80%	100%
3	FJ	6	20.17	100%	0%	66.7%	100%
4	MZ	6	20.17	100%	0%	50%	100%
5	UG	21	20.05	100%	0%	95.2%	100%
6	CL	8	20.00	100%	25%	62.5%	100%
7	LU	6	19.83	100%	0%	66.7%	100%
8	GT	113	19.71	98.2%	33.6%	53.1%	99.1%
9	MQ	7	19.71	100%	0%	42.9%	100%
10	GD	20	19.70	100%	0%	35%	100%

4.2 Highest Volume Countries

Countries with the largest device populations in the dataset, indicating high FortiGate market penetration and correspondingly large attack surfaces:

Country	Device Count	Avg Risk	SSL-VPN %	Any/Any Rule %
MX	1,603	18.43	99.1%	84%
AE	1,081	19.05	99.3%	86.6%
TH	1,043	19.24	98.8%	89.6%
MY	798	18.18	98.7%	80.1%
BR	679	18.85	99.4%	82.5%
US	679	18.62	98.7%	83.5%
AU	553	18.84	99.3%	88.2%
CO	529	18.62	98.9%	81.9%
DO	513	19.60	100%	92.4%
SA	460	18.70	99.6%	79.8%

Mexico (MX) holds the largest device count at 1,603 devices, followed by UAE (AE) at 1,081 and Thailand (TH) at 1,043. Despite the large sample sizes, risk scores remain consistently high across all top-volume countries, indicating these patterns represent systemic industry-wide misconfiguration norms rather than outlier behaviour.

4.3 Lowest Risk Countries (≥5 devices)

For comparative baseline, the countries with the lowest average risk scores:

Country	Devices	Avg Risk	HTTP WAN %	No PW Policy %
IS	9	17.00	66.7%	88.9%
NZ	110	17.50	77.3%	89.1%
PY	23	17.61	91.3%	78.3%
LB	6	17.67	100%	100%
BS	26	17.69	100%	100%
EC	60	17.87	78.3%	86.7%
GP	5	18.00	100%	100%
SN	18	18.11	83.3%	94.4%
ZA	57	18.11	96.5%	66.7%
NP	23	18.13	82.6%	100%

Note: Even the 'lowest risk' countries show significant exposure on key indicators. Iceland (IS) at avg risk 17.00 still has 66.7% of devices exposing HTTPS on WAN. The global floor is high — there is no country cluster that demonstrates consistently well-hardened FortiGate deployments in this dataset.

5. Implications for DFIR and Detection Engineering

5.1 Incident Response Triage

When triaging FortiGate-related incidents, the following prioritization framework is informed by the dataset analysis:

- ***Assume no off-device logging***: 78.5% of devices lack syslog/FortiAnalyzer. Default to memory forensics (RAM dump via SSH) and local log extraction as primary evidence sources.
- ***Correlate admin port from config***: 95.7% expose management on non-standard ports (4443, 9443, 4433). Standard port scanning may miss management surfaces — parse config for admin-sport and admin-ssh-port.
- ***Default credential validation***: With 87.1% retaining the 'admin' account and 91.5% lacking password policy, validate credential hygiene as an initial hypothesis. Check for cloud-deployed instances (FGVMA/FGVMC platforms) where instance-ID default passwords are common.

- ***CVE-2022-40684 specific artefacts:*** Check /var/log/httpd/ for requests to /api/v2/cmdb/system/admin with X-Forwarded-For spoofing or Node/Python user-agent strings — hallmarks of the exploitation toolkit.

5.2 Cobalt Strike / Malleable C2 Correlation

The high prevalence of Any/Any/Any firewall rules (82.6%) combined with FortiGate SSL-VPN as the ***remote access gateway creates favorable conditions for Cobalt Strike beacon staging***. The beaconing pattern over HTTPS on port 443 or 4443 — matching the FortiGate management port itself — represents a ***traffic blending opportunity*** that security teams should account for in their malleable C2 detection rules. In Microsoft Sentinel, ***custom analytic rules should filter on FortiGate*** source IPs specifically for outbound HTTPS to non-Fortinet ASNs on management port ranges.

5.3 Microsoft Sentinel Detection Rules

Key KQL detection opportunities derived from this dataset:

- FortiGate admin interface access from non-RFC1918 sources on management ports (4443, 9443, 4433, 443)
- FortiOS authentication events with username 'admin' — 87.1% prevalence makes this a high-signal indicator when combined with geolocation anomaly
- SSL-VPN connections from countries with >95% CVE-2023-27997 exposure where the connecting client is not a known user device
- FGFM (TCP/541) traffic originating from internet IPs — 78.2% of devices expose this; any external FGFM connections warrant immediate investigation
- Absence of FortiGate syslog events for >24h from a known device — 78.5% have no logging, so sudden log silence may indicate active tampering

6. Conclusions

The Belsen Group dataset provides an unprecedented longitudinal view of enterprise firewall misconfiguration at global scale. The findings indicate that the security posture of internet-exposed FortiGate deployments — at least within this cohort — ***is uniformly poor across all geographies***. The average risk score of 18.82/25 across 15,474 devices, with virtually no statistical separation between countries at the macro level, suggests the root cause is not regional security awareness but rather a systemic configuration-complexity and patch-management failure across the FortiGate product deployment lifecycle.

From a threat intelligence perspective, the dataset serves as a ground-truth snapshot of the attack surface available to adversaries operating in 2022 and — given the evidence of persistent non-patching — into subsequent years. Organizations that use FortiGate devices should treat CVE-2022-40684, CVE-2023-27997, and CVE-2024-21762 as high-priority patching obligations regardless of their apparent internal network position, given the widespread WAN exposure of management interfaces documented here.

As a final word for my home residing city of Hong Kong, currently, put the slide as my warm advice even though it may offend some IT folks.