INTELLIGENCE BRIEFING
Security Command Center
TLP:CLEAR
2026-04-02 13:39 UTC

Residential Proxies Evade IP Reputation Controls in 78% of 4 Billion Malicious Sessions, Structural Defense Gap Quantified

SECURITY ANALYSIS | HIGH | CVSS 7.5

SCC Item ID	SCC-STY-2026-0043
Type	Security Analysis
Severity	HIGH
CVSS Base Score	7.5
Affected Products	IP reputation-based security controls broadly; enterprise VPN and web application login pages; organizations relying on IP blocklists as primary perimeter defense
Published	2026-04-02T11:21:02
Discovery Source	Rss

Executive Summary

GreyNoise analysis of 4 billion malicious edge-targeting sessions found that IP reputation controls, a front-line defense for most enterprises, failed 78% of the time against attacks routed through residential proxies. Attackers are systematically exploiting a structural gap: residential IP addresses rotate faster than reputation feeds can blacklist them, making credential stuffing, account takeover, and scraping campaigns evade IP reputation controls but detectable through behavioral analysis of authentication patterns. Google Threat Intelligence Group's disruption of the IPIDEA residential proxy network removes one major enabler, but the underlying infrastructure model remains commercially available and widely replicated.

Technical Analysis

GreyNoise's analysis of 4 billion malicious sessions targeting internet-facing edge infrastructure quantifies what the security community has long suspected: IP reputation-based blocking is structurally inadequate against residential proxy infrastructure. Approximately 39% of all malicious sessions in the dataset originated from residential IP space - a significant shift from pre-residential-proxy threat baselines - with addresses assigned to home ISP subscribers across 683 ISPs globally. These addresses are functionally indistinguishable from legitimate user traffic at the network layer, and most remain active for under one month before rotating. Reputation feed update latency - typically measured in hours to days for most blocklist services - cannot match residential IP churn cycles measured in weeks, leaving blocklists perpetually behind the threat.

The MITRE ATT&CK techniques observed across this activity cluster tell a coherent story: attackers used external remote services (T1133) as entry vectors, applied password spraying (T1110.003) and credential stuffing (T1110.004) against login pages, and conducted active scanning of both IP ranges (T1595.001) and exposed services (T1595.002). Residential proxy infrastructure (T1090.002) provided the anonymization and IP churn that neutralized perimeter controls. File and directory discovery (T1083) and credential-based modifications (T1556) indicate post-authentication activity in successful sessions.

The IPIDEA residential proxy network, identified by Google Threat Intelligence Group as a primary infrastructure enabler, operated by recruiting residential devices, likely through software bundles or deceptive app installations, into a commercially marketed proxy pool. GTIG's disruption, detailed in the Google Cloud Threat Intelligence blog, removed one major node in this ecosystem. However, the residential proxy-as-a-service model is broadly replicated; IPIDEA's disruption does not close the structural gap it exploited.

The implications for enterprise defenders are direct. Organizations using IP reputation as a primary control on VPN gateways, web application login pages, or API endpoints have a measurable, systemic blind spot. CWE-799 (improper control of interaction frequency) and CWE-307 (improper restriction of excessive authentication attempts) map precisely to the authentication abuse patterns in this dataset. Rate limiting keyed to IP address is insufficient when the attacking IP rotates every few weeks across tens of thousands of residential addresses. Compensating controls must shift the detection anchor from IP identity to session behavior: device fingerprinting, velocity analysis per user account rather than per source IP, behavioral anomaly detection, and authentication hardening through phishing-resistant MFA.

Action Checklist

1. Step 1: Assess exposure, audit whether your VPN gateways, web application login pages, and API authentication endpoints rely on IP reputation feeds or IP-keyed rate limiting as a primary or sole front-line control
2. Step 2: Review controls, verify that rate limiting is applied per user account and per session, not per source IP; confirm MFA is enforced on all external-facing authentication surfaces; evaluate whether behavioral analytics or device fingerprinting is deployed at authentication chokepoints
3. Step 3: Update threat model, incorporate residential proxy-assisted credential stuffing and account takeover (T1110.003, T1110.004, T1090.002) into your threat register; document the quantified evasion rate (78%) as risk evidence against IP reputation as a standalone control
4. Step 4: Communicate findings, brief leadership that perimeter IP blocklisting has a measured 78% failure rate against this attack class; present behavioral analytics and phishing-resistant MFA as necessary controls to close the identified gap, not as generic hardening
5. Step 5: Monitor developments, track GTIG and GreyNoise for follow-up disclosures on IPIDEA successor infrastructure; watch for new residential proxy network advisories; monitor authentication failure telemetry for velocity patterns consistent with distributed credential stuffing across low-rate residential IPs

IR / Forensic Enrichment

Triage Priority	URGENT

Escalation Criteria	Escalate immediately to CISO and legal/compliance if authentication logs show per-account failure patterns consistent with successful account takeover (failed attempts followed by successful login from a new residential IP with no prior authentication history for that account), if any compromised account has access to PII or PHI triggering breach notification obligations under GDPR, CCPA, or HIPAA, or if the security team lacks the capability to implement account-keyed rate limiting and phishing-resistant MFA within 30 days.
Recovery Notes	After implementing account-keyed rate limiting and phishing-resistant MFA, conduct a full review of authentication success logs for the 90 days preceding deployment to identify any accounts that may have been compromised via residential proxy-routed credential stuffing before controls were hardened — specifically look for accounts with successful logins from residential ASNs (Comcast, AT&T, BT, Vodafone) that had no prior authentication history from those ASNs. Force password resets and session invalidation for any accounts meeting that pattern. Monitor post-remediation authentication telemetry for a minimum of 30 days, tracking whether per-account failure rates decline and whether attackers shift tactics to lower-volume patterns or target newly identified gaps at API endpoints not covered by the initial control review.
Forensic Artifacts	Authentication failure logs from VPN gateways keyed to username field with source IP, timestamp, and ASN — specifically Cisco ASA syslog events 113015 (AAA user authentication rejected) and 113005 (AAA user authentication failed) or Palo Alto GlobalProtect authd.log — filtered for accounts with failures from 8+ distinct /24 subnets within a 60-minute window, which is the residential proxy distribution signature | Web application access logs (Nginx access.log, Apache access.log, IIS W3C logs) filtered for HTTP 401 and 403 responses to authentication endpoints (/login, /api/auth, /oauth/token) with full User-Agent strings and source IPs — residential proxy credential stuffing leaves a pattern of high-volume 401s with low per-IP repeat rates and often a narrow set of User-Agent strings across thousands of distinct residential IPs | JA3 TLS client fingerprint logs from WAF or load balancer (Cloudflare, AWS ALB, F5) for failed authentication sessions — commodity credential stuffing tools using residential proxy networks often share TLS fingerprint clusters even when source IPs rotate across thousands of residential addresses, making JA3 hash clustering a proxy-resistant detection signal | IdP or directory service authentication event logs (Azure AD Sign-In Logs event codes 50126/50053/50057, Okta System Log 'user.session.start' failures, or Active Directory Security Event Log Event ID 4625 Type 10 network logon failures) correlated with source IP ASN classification to distinguish residential ISP ASNs from datacenter ASNs — a high ratio of residential ASN failures targeting a small set of high-privilege accounts is the account takeover precursor signature for this attack class | Session token and cookie logs from application-layer authentication (if captured by WAF or application logging) for any authentication sessions that succeeded following a distributed failure pattern — these represent completed account takeovers and should include the session IP, User-Agent, device fingerprint (if available), and the authentication timestamp to establish attacker dwell time before detection

Per-Action IR Details

Step 1: Assess exposure — audit whether your VPN gateways, web application login pages, and API authentication endpoints rely on IP reputation feeds or IP-keyed rate limiting as a primary or sole front-line control

NIST Phase: Preparation

Reference: NIST 800-61r3 §2 — Preparation: Establish IR capability and assess control posture before incidents occur

Controls: NIST IR-4 (Incident Handling) — requires preparation phase that includes assessing existing controls against known attack classes, NIST RA-3 (Risk Assessment) — assess the likelihood and impact of residential proxy-routed credential stuffing given the quantified 78% IP reputation evasion rate, NIST SI-4 (System Monitoring) — verify monitoring controls at VPN gateways and web application login surfaces are not solely IP-keyed, CIS 7.1 (Establish and Maintain a Vulnerability Management Process) — document IP reputation reliance as a structural control gap in the vulnerability register, CIS 4.4 (Implement and Manage a Firewall on Servers) — confirm firewall and perimeter controls at authentication endpoints are not relying on IP blocklists as the sole gate

Compensating: Run a configuration audit using a bash one-liner or PowerShell script against your VPN gateway and WAF configs to grep for IP-keyed rate limiting rules: on Linux, 'grep -rn "limit_req_zone.*\$binary_remote_addr" /etc/nginx/' or equivalent for Apache/HAProxy. For cloud-hosted login pages, review AWS WAF or Cloudflare rule sets for rules keyed solely on IPSets or IP reputation lists. Document each authentication endpoint (VPN portal URL, /login, /api/auth) and map its active controls in a simple spreadsheet — flag any endpoint where IP reputation or IP-keyed rate limiting is the only pre-authentication gate.

Evidence: Before auditing, capture point-in-time snapshots of: (1) current IP blocklist or reputation feed configuration files from your WAF, VPN concentrator, and API gateway (e.g., Nginx limit_req_zone directives, Palo Alto EDL configurations, Cloudflare IP lists); (2) authentication log samples from the past 30 days showing source IP distribution for failed logins — export from VPN auth logs (e.g., Cisco ASA syslog event 113015/113005, Palo Alto GP auth failure logs) to establish a baseline of IP diversity before residential proxy traffic is formally scoped; (3) any existing rate-limiting threshold configuration to document what the current IP-keyed ceiling is.

Step 2: Review controls — verify that rate limiting is applied per user account and per session, not per source IP; confirm MFA is enforced on all external-facing authentication surfaces; evaluate whether behavioral analytics or device fingerprinting is deployed at authentication chokepoints

NIST Phase: Preparation

Reference: NIST 800-61r3 §2 — Preparation: Implement and validate preventive controls to reduce incident probability and impact

Controls: NIST IA-5 (Authenticator Management) — verify MFA enrollment and enforcement status across all externally-exposed authentication surfaces including VPN and web app login pages, NIST AC-7 (Unsuccessful Logon Attempts) — confirm lockout and rate-limiting policies are keyed to user account identifier, not source IP, to resist distributed residential proxy stuffing, NIST SI-10 (Information Input Validation) — evaluate whether authentication endpoints validate session and behavioral signals beyond source IP for anomaly detection, NIST SI-4 (System Monitoring) — assess whether behavioral analytics or device fingerprinting tools are active at authentication chokepoints to detect low-and-slow credential stuffing patterns, CIS 6.3 (Require MFA for Externally-Exposed Applications) — enforce MFA on all externally-exposed enterprise applications, CIS 6.4 (Require MFA for Remote Network Access) — enforce MFA for all VPN and remote access authentication, CIS 6.5 (Require MFA for Administrative Access) — enforce MFA for administrative accounts on all enterprise assets

Compensating: For teams without enterprise behavioral analytics: deploy fail2ban configured with 'findtime' and 'maxretry' keyed to usernames rather than IPs (use 'failregex' patterns that extract the username field from VPN or web app auth logs). For device fingerprinting on a budget, implement browser fingerprinting via the open-source FingerprintJS community edition on login pages, or enforce client certificate authentication for VPN access using an internal CA (e.g., OpenVPN with mutual TLS). To audit MFA enforcement gaps, run: 'Get-MsolUser -All | Where-Object {$_.StrongAuthenticationMethods.Count -eq 0}' in PowerShell against Azure AD, or equivalent LDAP query against your IdP to enumerate accounts missing MFA enrollment on externally-accessible services.

Evidence: Capture before remediating: (1) MFA enrollment reports from your IdP (Azure AD, Okta, Duo) exported to CSV — flag all accounts with access to external-facing VPN or web apps that lack phishing-resistant MFA (FIDO2/WebAuthn); (2) current rate-limiting rule exports from WAF and VPN gateway configs confirming whether 'maxretry' or 'threshold' values are keyed to $remote_addr (IP) vs. username/account fields — this documents the structural gap before remediation; (3) authentication session logs showing User-Agent strings and TLS fingerprints (JA3 hashes if your proxy or WAF captures them) for recent failed login attempts, to establish whether sessions are already showing residential proxy behavioral signatures (high IP diversity, uniform User-Agent, consistent attack timing).

Step 3: Update threat model — incorporate residential proxy-assisted credential stuffing and account takeover (T1110.003, T1110.004, T1090.002) into your threat register; document the quantified evasion rate (78%) as risk evidence against IP reputation as a standalone control

NIST Phase: Preparation

Reference: NIST 800-61r3 §2 — Preparation: Update threat intelligence and risk documentation to reflect current adversary TTPs

Controls: NIST RA-3 (Risk Assessment) — formally document residential proxy-routed credential stuffing as a threat scenario with quantified likelihood evidence (GreyNoise 78% evasion rate across 4B sessions) in the organizational risk register, NIST IR-8 (Incident Response Plan) — update the IR plan to include residential proxy-assisted account takeover as a named threat scenario with detection indicators and response actions specific to T1110.003, T1110.004, and T1090.002, NIST SI-5 (Security Alerts, Advisories, and Directives) — formalize intake of GreyNoise and GTIG advisories as authoritative external threat intelligence sources for this threat class, NIST PM-16 (Threat Awareness Program) — incorporate residential proxy evasion as a documented threat pattern in organizational threat awareness materials, CIS 7.1 (Establish and Maintain a Vulnerability Management Process) — register IP reputation reliance as a documented control gap with the 78% failure rate as quantified risk evidence

Compensating: Use the MITRE ATT&CK Navigator (free, browser-based at attack.mitre.org/resources/attack-navigator) to create a layer file highlighting T1110.003 (Password Spraying), T1110.004 (Credential Stuffing), and T1090.002 (External Proxy) with annotation linking to the GreyNoise 4B-session finding. Export as JSON and attach to your risk register entry. For the threat register itself, a structured Markdown or CSV file with fields: Threat Name, ATT&CK TTP IDs, Evidence Source, Quantified Risk (78% evasion), Affected Controls, Compensating Controls, Owner, Review Date — is sufficient for a 2-person team without a GRC platform.

Evidence: Before finalizing the threat model update, collect: (1) a pull from GreyNoise Community API (free tier) for your externally-facing IP ranges to check whether residential proxy traffic has already been observed targeting your assets — query: 'https://api.greynoise.io/v3/community/{ip}' for each VPN gateway and web app IP; (2) authentication failure logs from the past 90 days filtered for accounts with failures from more than 5 distinct /16 subnets within a 24-hour window, which is a behavioral signature of distributed residential proxy stuffing rather than traditional botnet attacks (which show ASN clustering); (3) any existing threat register or risk assessment documentation showing IP reputation as a listed control, to establish the pre-update baseline for audit trail purposes.

Step 4: Communicate findings — brief leadership that perimeter IP blocklisting has a measured 78% failure rate against this attack class; frame compensating control investment (behavioral analytics, phishing-resistant MFA) as risk reduction against a quantified gap, not a generic hardening request

NIST Phase: Post Incident

Reference: NIST 800-61r3 §4 — Post-Incident Activity: Communicate lessons learned and evidence-based control improvement recommendations to leadership

Controls: NIST IR-6 (Incident Reporting) — report the identified structural control gap (IP reputation failure against residential proxy evasion) to organizational leadership with quantified evidence, NIST IR-8 (Incident Response Plan) — update the IR plan based on identified gap and present recommended changes to leadership for approval and resource allocation, NIST CA-7 (Continuous Monitoring) — propose behavioral analytics and phishing-resistant MFA as continuous monitoring improvements to replace or augment IP reputation controls, NIST PM-9 (Risk Management Strategy) — frame the 78% evasion rate as quantified organizational risk requiring documented acceptance or mitigation decision from leadership, CIS 7.2 (Establish and Maintain a Remediation Process) — present the compensating control roadmap (behavioral analytics, FIDO2 MFA) as a risk-based remediation plan with prioritization rationale

Compensating: Build the leadership brief as a one-page risk memo with three sections: (1) Quantified Gap — cite GreyNoise analysis of 4B sessions, 78% IP reputation evasion rate, specific ATT&CK TTPs (T1110.003, T1110.004, T1090.002); (2) Current Exposure — list your specific externally-facing authentication surfaces audited in Steps 1-2 and which rely on IP-keyed controls; (3) Recommended Investment — FIDO2/WebAuthn phishing-resistant MFA (free with Duo Free tier for up to 10 users, or Cloudflare Zero Trust free tier) and account-keyed rate limiting (free in Nginx, HAProxy, or fail2ban). Attach the MITRE ATT&CK Navigator layer as a visual annex. This avoids the need for a GRC platform or reporting tool.

Evidence: Assemble as supporting evidence for the leadership brief: (1) the point-in-time control configuration snapshots captured in Steps 1-2 (WAF rules, VPN rate-limit configs) demonstrating IP-keyed reliance; (2) authentication failure log analysis showing IP diversity patterns consistent with residential proxy distribution (high unique IP count, low per-IP failure count, targeting specific high-value accounts); (3) the GreyNoise advisory citation with the 4B session dataset reference as the external evidence anchor — do not present the 78% figure without this citation to maintain credibility with technically-literate leadership.

Step 5: Monitor developments — track GTIG and GreyNoise for follow-up disclosures on IPIDEA successor infrastructure; watch for new residential proxy network advisories; monitor authentication failure telemetry for velocity patterns consistent with distributed credential stuffing across low-rate residential IPs

NIST Phase: Detection Analysis

Reference: NIST 800-61r3 §3.2 — Detection and Analysis: Monitor for indicators of adverse events and correlate authentication telemetry for residential proxy-routed attack patterns

Controls: NIST SI-4 (System Monitoring) — implement continuous monitoring of authentication failure telemetry specifically tuned to detect distributed low-rate credential stuffing patterns that evade per-IP thresholds, NIST AU-6 (Audit Record Review, Analysis, and Reporting) — establish recurring review cadence for authentication logs analyzing per-account failure velocity, IP diversity ratios, and ASN distribution anomalies, NIST SI-5 (Security Alerts, Advisories, and Directives) — formalize GTIG and GreyNoise as tracked external intelligence sources with an intake process for new residential proxy network advisories, NIST IR-5 (Incident Monitoring) — track and document any authentication anomalies meeting residential proxy-pattern thresholds as potential incidents, CIS 8.2 (Collect Audit Logs) — ensure authentication logs from VPN gateways, web application login pages, and API authentication endpoints are collected with sufficient fidelity (source IP, username, User-Agent, timestamp, success/failure, session ID) to support residential proxy pattern detection

Compensating: Deploy the following free detection stack for a 2-person team: (1) Sigma rule targeting authentication logs — create a rule matching on: per-account failure count > 10 within 60 minutes AND unique source IP count > 8 for the same account within the same window AND no single IP accounting for more than 3 failures (this is the residential proxy signature: high IP diversity, low per-IP rate, account-focused). Run via sigma-cli converting to Splunk, Elastic, or grep-compatible format against exported auth logs on a daily cron job. (2) For real-time alerting without SIEM: configure fail2ban with a custom jail that tracks per-username failure counts across all source IPs using the 'ip6tables-multiport' action combined with username extraction regex from your VPN or Nginx auth logs. (3) Subscribe to GreyNoise Community RSS or their free API to poll weekly for new tags matching 'residential-proxy', 'ipidea', or 'credential-stuffing' against your externally-facing IPs.

Evidence: Capture and retain as ongoing monitoring evidence: (1) Authentication logs from VPN concentrator (Cisco ASA: syslog events 113015, 113005, 716001; Palo Alto GlobalProtect: authd.log entries with 'Authentication failed' for the username field), web application login pages (Nginx access.log HTTP 401/403 responses to /login or /api/auth paths with full User-Agent and source IP), and API gateways (AWS API Gateway access logs with identity.caller and sourceIp fields); (2) JA3/JA3S TLS fingerprint logs if your load balancer or WAF captures them — residential proxy tooling often shares TLS fingerprint clusters even when source IPs rotate; (3) ASN distribution reports from authentication failure logs — residential proxy traffic will show high ASN diversity (ISP ASNs: Comcast, AT&T, Verizon, Vodafone) rather than the datacenter ASN clustering typical of traditional credential stuffing botnets; (4) User-Agent string frequency analysis from failed authentication sessions — commodity credential stuffing tools often use a narrow set of User-Agent strings even when rotating IPs.

Detection Guidance

Authentication logs are the primary detection surface. Prioritize these signals: (1) failed authentication events grouped by target account rather than source IP, looking for multiple failures from distinct residential ISPs within a time window; (2) successful logins preceded by authentication failures from geographically dispersed IPs within a short window; both patterns are consistent with low-and-slow credential stuffing via residential proxies. Standard brute-force detection tuned to single-source IP thresholds will miss this pattern entirely.

Additional detection vectors: authentication events from residential ISP ASNs (as opposed to datacenter or commercial ASNs) at abnormal hours or volumes; session anomalies where device fingerprint or user-agent does not match the account's established baseline; authentication attempts where the source IP has no prior login history for that account and the ASN is a residential broadband provider (cross-reference with threat intelligence feeds that tag residential proxy exit nodes, noting that coverage will be incomplete given IP churn rates).

Behavioral analytics should track account-level velocity: number of authentication attempts per account per hour regardless of source IP diversity; time-to-first-success ratios; and geographic impossibility flags where the same account authenticates from geographically distant residential IPs within implausibly short intervals.

For organizations with web application firewalls, audit whether challenge-response mechanisms (CAPTCHA, JavaScript challenges) are triggered by session behavior signals rather than IP reputation scores alone. IP-gated challenges will fail to fire against clean residential addresses.

Framework Mappings

MITRE-ATTACK

- **T1083** — File and Directory Discovery
- **T1090.002** — External Proxy
- **T1133** — External Remote Services
- **T1110.001** — Password Guessing
- **T1556** — Modify Authentication Process
- **T1110.003** — Password Spraying
- **T1595.002** — Vulnerability Scanning
- **T1110.004** — Credential Stuffing
- **T1595.001** — Scanning IP Blocks

NIST-800-53R5

- **AC-17** — Remote Access
- **AC-20** — Use of External Systems
- **IA-2** — Identification and Authentication (Organizational Users)
- **IA-5** — Authenticator Management
- **SC-7** — Boundary Protection
- **AC-6** — Least Privilege
- **SI-4** — System Monitoring
- **SI-7** — Software, Firmware, and Information Integrity
- **AC-7** — Unsuccessful Logon Attempts

OWASP-TOP10-2021

- **A07:2021** — Identification and Authentication Failures

CIS-V8

- **6.3** — Require MFA for Externally-Exposed Applications

HIPAA-SECURITY

- **164.312(d)** — Person or Entity Authentication

SOC2-TSC

- **CC6.1** — Logical access security software, infrastructure, and architectures

MITRE ATT&CK Mapping

Technique ID	Technique Name	Tactic
T1083	File and Directory Discovery	Discovery
T1090.002	External Proxy	Command-And-Control
T1133	External Remote Services	Persistence
T1110.001	Password Guessing	Credential-Access
T1556	Modify Authentication Process	Credential-Access
T1110.003	Password Spraying	Credential-Access
T1595.002	Vulnerability Scanning	Reconnaissance
T1110.004	Credential Stuffing	Credential-Access
T1595.001	Scanning IP Blocks	Reconnaissance

Sources

Source	URL	Tier
Security News	https://www.bleepingcomputer.com/news/security/residential-proxies-...	T3
Disrupting the World's Largest Residential Proxy Network	https://cloud.google.com/blog/topics/threat-intelligence/disrupting...	T3
Google Threat Intelligence Group shuts down IPIDEA proxy network	https://blog.google/innovation-and-ai/infrastructure-and-cloud/goog...	T1
Google Disrupts IPIDEA Proxy Network - SecurityWeek	https://www.securityweek.com/google-disrupts-ipidea-proxy-network/	T3
Google Disrupts Major Residential Proxy Network IPIDEA	https://www.esecurityplanet.com/threats/google-disrupts-major-resid...	T3

DISCLAIMER

This intelligence report is produced by Tech Jacks Solutions Security Command Center (SCC) for informational purposes only. It does not constitute professional security advice, legal counsel, or an incident response engagement. The information herein is derived from publicly available sources and AI-assisted analysis; while every effort is made to ensure accuracy, Tech Jacks Solutions makes no warranties regarding completeness or timeliness. Organizations should conduct their own validation and consult qualified security professionals before taking action based on this report. Tech Jacks Solutions is not liable for any damages resulting from the use of this information.

Generated 2026-04-02 13:39 UTC by TJS Security Command Center