

Gestão de vulnerabilidades na era da IA: como priorizar CVEs que crescem de forma exponencial

Dra. Cristine Hoepers
Gerente, CERT.br/NIC.br
cristine@cert.br


14º Fórum Brasileiro de CSIRTs
São Paulo, SP – 28 de julho de 2026

cert.br nic.br cgi.br

TLP:CLEAR

Year-to-date CVE publications (MITRE CVE List)

Lines showing the daily cumulative count of published CVEs on MITRE's CVE List, <https://cve.mitre.org/cve/>



Fonte: https://www.first.org/epss/data_stats

Source: https://first.org/epss/data_stats, 2026-07-26

Verizon DBIR 2026

Vulnerabilidades críticas/ano

↑ 50%

Exploração de vulnerabilidades é o vetor de **acesso inicial** mais comum em *breaches*.

↑ 31%

Só 26% das Vulnerabilidades no CISA KEV foram completamente **remediadas** em 2025

↓ 38%




Figure 5. Known initial access vectors in non-Error, non-Misuse breaches over time (n for 2026 dataset=19,905)

Fonte: <https://www.verizon.com/business/resources/Tf99/reports/2026-dbir-data-breach-investigations-report.pdf>

Systematic Exploitation of Edge and Core Network Devices

Over the past several years, Mandiant has seen an increase of malicious actors targeting edge and core network devices. Attackers have taken advantage of the fact that many of these devices are not able to run enterprise security tooling that provides increased visibility into threat actor actions and can interdict and prevent malicious activity.

At the same time, Google Threat Intelligence Group (GTIG) has observed that the mean time to exploit (TTE) for vulnerabilities has continually decreased from 63 days in 2018 to -1 day in 2024 and further downward to an estimated -7 days in 2025. A negative number indicates that exploitation of a vulnerability, on average, occurred before a patch was released. This trend is compounded by threat actors' evolving interest in security appliances and networking infrastructure. Because these assets provide critical inspection of data ingress and egress points, they grant significant visibility and permissions to those who gain access to them. Additionally, they often show signs of extended uptime periods, delays in patching, and can often be excluded from vulnerability management altogether. This has allowed threat actors the opportunity to exploit n-day vulnerabilities in these devices repeatedly without needing to identify novel vulnerabilities.

Fonte: M-Trends 2026 – <https://cloud.google.com/blog/topics/threat-intelligence/m-trends-2026>

blog.checkpoint.com

CHECK POINT™

RESEARCHMAY 13, 2026

When the Ransomware Gang Gets Hacked: What the Gentlemen Leak Reveals About Modern Ransomware Risk

By Check Point Team

SHARE

f

in


Key Findings

- The Gentlemen RaaS has 400+ public victims and is the #2 most active ransomware group globally in 2026
- Their internal systems were breached in May 2026, exposing their full operational structure
- The group is run by approximately nine named operators organized around a single administrator (zeta88 / hastalamuerte), who not only manages the platform but personally participates in encryption events

Entry is almost always through an unpatched internet-facing device. The Gentlemen specifically target VPNs and appliances, exploiting CVE-2024-55591 and CVE-2025-32433, buying access from third-party brokers, or using credentials harvested from infostealer log markets. Once inside, they move fast: Active Directory enumeration, NTLM relay attacks (CVE-2025-33073), EDR disablement, lateral movement via legitimate admin tools, browser session harvesting for Microsoft 365 and Okta access, and data exfiltration — all before a domain-wide ransomware deployment via Group Policy that hits every connected endpoint simultaneously.

Source: <https://blog.checkpoint.com/research/when-the-ransomware-gang-gets-hacked-what-the-gentlemen-leak-reveals-about-modern-ransomware-risk/>

Edge and Core Network Devices: Dispositivos com CVEs críticos notificados mensalmente pelo CERT.br



Fonte: <https://stats.cert.br/vulns/>

Fonte: CERT.br — <https://stats.cert.br/> — by Highcharts.com

TLP:CLEAR

cert.br nie.br egi.br

cogent.com

BLOG

RESEARCH

62% of Critical Vulnerabilities Have Exploits Circulating Before Scanners Can Detect Them


See why the race between exploit development and scanner coverage is accelerating, and what the findings reveal about the future of vulnerability detection.

May 27, 2026 5 min read



Geng Sng

Co-Founder and CTO



Fonte: <https://www.cogent.com/blog/2026-q2-detection-gap-report-findings>

Microsoft July 2026 Patch Tuesday fixes massive 570 flaws, 3 zero-days

By **Lawrence Abrams**



VULNERABILITIES

Oracle Patches Over 1,400 Vulnerabilities With Quarterly Security Updates

Many of the vulnerabilities fixed with the July 2026 Critical Patch Update were likely discovered by AI.



By **Eduard Kovacs**

July 22, 2026 (5:33 AM ET)



Fontes:

<https://www.bleepingcomputer.com/news/microsoft/microsoft-july-2026-patch-tuesday-fixes-massive-570-flaws-3-zero-days/>


<https://www.securityweek.com/oracle-patches-over-1400-vulnerabilities-with-quarterly-security-updates/>

TLP:CLEAR

cert.br nic.br cgi.br

Year-to-date CVE publications (MITRE CVE List)

Lines showing the daily cumulative count of published CVEs on MITRE's CVE List, <https://cve.mitre.org/cve/>



Fonte: https://www.first.org/epss/data_stats

Source: https://first.org/epss/data_stats, 2026-07-26

Assimetria em Segurança da Informação

Atacantes precisam achar
uma única fraqueza.

Defensores precisariam proteger
todas as possíveis fraquezas.

Assimetria em Segurança da Informação:

Modelos atuais de IA estão aumentando ainda mais a assimetria

Encadeiam vulnerabilidades com criticidade média e baixa

- Costumava ser um trabalho artesanal, com baixa escala
- Tornou-se trivial permutar sobre milhares de possíveis combinações de vulnerabilidades

Atacam incessantemente até encontrar uma brecha

- OpenAI vs. Hugging Face: 17.000 eventos em um único ataque
- Relatório da Hugging Face fala na dificuldade de “*separate genuine impact from decoy activity*”
 - muitos ataques, mesmo que irrelevantes, até que um tem sucesso
- Os modelos não permitiram seu uso para analisar os ataques, por quebra de compliance


Há a promessa de contribuir para a melhora:

“We believe advanced cyber capable models need to help security teams find weaknesses before attackers do, understand how vulnerabilities can be chained, and remediate them at machine speed.”

Fontes: <https://huggingface.co/blog/security-incident-july-2026> | <https://openai.com/index/hugging-face-model-evaluation-security-incident/>

Assimetria em Segurança da Informação: Mas quem poderá pagar por esses modelos de fronteira?

Trajectories for various AI models on the 32-step “The Last Ones” cyber range—a corporate network attack simulation spanning multiple subnets and hosts, without active defenders.



AISI | AI SECURITY
INSTITUTE

Fonte: <https://openai.com/index/hugging-face-model-evaluation-security-incident/>

TLP:CLEAR

cert.br nie.br egi.br

Gestão de Vulnerabilidades

“A prática de
identificar, priorizar e corrigir
vulnerabilidades de *software*
conhecidas.”

Fonte: <https://arxiv.org/pdf/2302.14172>


Priorização de Vulnerabilidades com Base em Risco

Qual o risco de uma vulnerabilidade
ser ativamente explorada
e
causar um impacto negativo?

Ou seja: risco operacional.

SSVC – Stakeholder-Specific Vulnerability Categorization

- Metodologia para:
 - **priorizar vulnerabilidades**
 - **com impacto operacional e de bem estar**
- Árvore de decisão simples de entender
 - e de explicar para gestores e jurídico
- Usada na nova diretriz do CISA: “BOD 26-04: *Prioritizing Security Updates Based on Risk*”
- A base do **CISA KEV** (*Known Exploited Vulnerabilities*) foi criada para ajudar na decisão sobre “**exploitation: active**”



Fonte: <https://cisa.gov/ssvc/>


Pontos de decisão na árvore do SSVC

CISA KEY
EPSS
Metasploit
Exploit DB
Intelligence Reports
Vendor Analysis
CISA Vulnrichment

CVE + CVSS
Intelligence Reports
Vendor Analysis
CISA Vulnrichment

CVE + CVSS
CISA Vulnrichment

Dados internos
(análise de risco, missão do ativo, etc)



CISA Vulnrichment

- Reduz a árvore para os dois galhos vistos aqui
- Uma base com os CVEs que preenchem estas condições:
 - "Exploitation": "yes"
 - "Exploitation": "poc"
 - "Automatable": "yes"
 - "Technical Impact": "total"
- Para cada CVE, o seu **json** passa a incluir os detalhes de:
 - CWE e CVSS



Fonte: <https://github.com/cisagov/vulnrichment>

Exemplo de uso do SSSC:

CISA BOD 26-04: *Prioritizing Security Updates Based on Risk*

- Metas pragmáticas e simples
- Inclui as decisões do SSSC
- Mas adiciona mais prioridade se está exposto na Internet
- Assume que em alguns casos pode já estar comprometido
 - *patches & forensic triage*

	Publicly Exposed?†	In the KEV?*	Automatable by Adversary?*	Technical Impact*	Agency Timeline (Calendar Days) for Remediation
1	Yes	Yes	Yes	Total control	3 days & forensic triage
2	Yes	Yes	Yes	Partial control	3 days
3	Yes	Yes	No	Total control	3 days & forensic triage
4	Yes	Yes	No	Partial control	14 days
5	Yes	No	Yes	Total control	3 days
6	Yes	No	Yes	Partial control	14 days
7	Yes	No	No	Total control	14 days
8	Yes	No	No	Partial control	60 days
9	No	Yes	Yes	Total control	3 days & forensic triage
10	No	Yes	Yes	Partial control	14 days
11	No	Yes	No	Total control	14 days
12	No	Yes	No	Partial control	14 days
13	No	No	Yes	Total control	60 days
14	No	No	Yes	Partial control	60 days
15	No	No	No	Total control	Fix on system upgrade
16	No	No	No	Partial control	Fix on system upgrade

Fonte: <https://www.cisa.gov/news-events/directives/bod-26-04-prioritizing-security-updates-based-risk>

Refinando para o Seu Ambiente

Tenho o CVE?

É relevante para mim?

Qual o impacto operacional?

O que existe no meu parque?

Versões de *software*, incluindo bibliotecas e dependências leia-se “*supply chain*”

Padrões existentes:

- **SBOM** – *Software Bill of Materials*
 - lista de ingredientes
 - obrigatório na União Europeia a partir de 2027 (norma CRA)
- Identificação de *Software*
 - **CPE** – *Common Platform Enumeration*
 - **PURL** – Package URL
 - **OmniBOR**
 - SHA256 p/ artefatos de *software*

Quais ativos tem vulnerabilidades exploráveis?

Não basta ter um CVE

- precisa ser alcançável

Padrões existentes:

- **CSAF** – *OASIS Common Security Advisory Framework*
 - advisories para serem consumidos de forma automatizada
- **VEX** – Vulnerability Exploitability eXchange
 - um atestado do *vendor* de quais vulnerabilidades estão presentes e são alcançáveis
 - deve ser usado junto com SBOM

O que pode causar impacto negativo?

- Missão do ativo
- Análise de risco aka, “Joias da Coroa”
 - tanto dados, quanto sistemas
- Não esquecer daquilo que é desenvolvido localmente
 - APIs, aplicativos, bancos de dados, personalizações, etc.
 - Incluindo *frameworks*, pacotes e dependências.

Padrão existente:

- **SSVC** – *Stakeholder-Specific Vulnerability Categorization*
 - Uma metodologia para priorizar vulnerabilidades

Exemplo de como colocar em prática

- Os seus sistemas terem um SBOM publicado/gerado
 - incluindo containers, API, nuvem
- O *vendor* possuir um atestado VEX para todos os CVE publicados
- Ferramentas de gestão de vulnerabilidades consumirem
 - os atestados VEX dos vendedores
 - o SBOM do sistema como implantado no seu ambiente
 - um SSVC “base” criado por você ou de serviços como o CISA Vulnrichment
- A ferramenta permitir incluir as **suas métricas** de “Mission & Well-being”

Vendors como RedHat já mantem atestados VEX:

“The VEX file for a single, public CVE ID covers all the component statuses defined in the CSAF standard:

- ***Fixed***: *Contains the same fixed component versions and other details (product_tree objects) that the CSAF advisory reports for that CVE*
- ***Known Affected***: *Confirmation that the specific component and product is affected by a particular CVE*
- ***Known Not Affected***: *Confirmation that the specific component and product is not affected by a particular CVE*
- ***Under Investigation***: *Information that the Red Hat Product Security team is verifying the applicability and impact of a specific CVE to the specific component and product”*

<https://www.redhat.com/en/blog/red-hat-vex-files-cves-are-now-generally-available>

Sumarizando

- CVEs vão continuar crescendo exponencialmente
- O “status quo” dos fornecedores de soluções não está acompanhando
- Já existem padrões melhores do que usar CVSS e *scanners* de vulnerabilidades
- Mesmo sem ferramentas ideais, os times podem melhorar a priorização de patches e aumentar a resiliência operacional

Recomendações

É URGENTE conseguir identificar os patches imprescindíveis

– “É pra ontem”

- Se o CERT.br notificar (<https://stats.cert.br/vulns/>)
- Se estiver no CISA KEV (<https://cisa.gov/kev/>)

– Aplicar assim que possível

- Se for em um ativo crítico que pode impedir sua organização de funcionar
- Se atingir os critérios do CISA Vulnrichment (<https://github.com/cisagov/vulnrichment>)
- Se o percentil do EPSS for muito alto (<https://www.first.org/epss/>)

SBOM/AIBOM, CSAF, VEX, Vulnrichment e afins:

Palestras da CVE/FIRST Vulncon são o melhor ponto de partida

- Palestra sobre efetividade de *scanners* de vulnerabilidade para DevOps
“Mind the Match: Why Vulnerability Matching Is Harder Than You Think”, Alexandra Selldorff (Manifest Cyber, US)
https://youtu.be/WYMpvXImHI?si=i8p_FuyursBiugjz
- Vulnrichment: Year One, Art Manion (Tharros Labs, US); Lindsey Cerkovnik (CISA, US)
<https://youtu.be/g5pSVMnWD7k?si=9jFRPT6fG0qu1hye>
- Agendas e vídeos do CVE/FIRST VulnCon & Annual CNA Summit de 2024 a 2026
2024:
<https://www.first.org/conference/vulncon2024/program>
https://youtube.com/playlist?list=PLBAUUhONOrO_aB01IOv6XNRTHD4ueFVTp&si=UU3qHdDOZK1IBQJf
2025:
<https://www.first.org/conference/vulncon2025/program>
https://youtube.com/playlist?list=PLBAUUhONOrO8iOYvs3pAbuzb-A07ZdT9C&si=NWPEVHXUWmA4_Nil
2026:
<https://www.first.org/conference/vulncon26/program>
https://youtube.com/playlist?list=PLBAUUhONOrO_yESOH6JnwWBoRdDRVXDr0&si=2liX3-1N1QVD6_4U

Referências Adicionais

- Software Bill of Materials (SBOM)
<https://www.cisa.gov/sbom>
- CERT.br – Estatísticas de notificações de hosts com serviços potencialmente vulneráveis expostos na Internet
<https://stats.cert.br/vulns/>
- Obsolescência não-Programada: Análise do Uso de Software Desatualizado em Ambiente de Produção, Luan Marko Kujavski, Ulisses Penteado, Paulo Lisboa de Almeida, André Grégio, SBSeg 2024
<https://sol.sbc.org.br/index.php/sbseg/article/view/30046/29853>
- CIRCL *Vulnerability Lookup Tool* – *facilitates quick correlation of vulnerabilities from various sources, independent of vulnerability IDs, and streamlines the management of Coordinated Vulnerability Disclosure (CVD).*
<https://github.com/cve-search/vulnerability-lookup>

Padrões que tratam de CVD, SBOM e afins

- The European Union The Cyber Resilience Act
<https://www.cyberresilienceact.eu>
- ISO/IEC 29147 and 30111
<https://webstore.ansi.org/standards/iso/isoiec3011129147security>
- EthicsfIRST
<https://ethicsfirst.org/>
- FIRST PSIRT and CSIRT Services Frameworks
<https://www.first.org/standards/frameworks/>
- Políticas da OECD e da União Europeia
<https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0482>
<https://www.enisa.europa.eu/news/enisa-news/coordinated-vulnerability-disclosure-policies-in-the-eu>
- IETF security.txt – *RFC 9116: A File Format to Aid in Security Vulnerability Disclosure*
<https://securitytxt.org>
- Ato nº 77 da Anatel, Requisitos de Segurança Cibernética para Equipamentos para Telecomunicações, seção 6.1.6
<https://informacoes.anatel.gov.br/legislacao/atos-de-certificacao-de-produtos/2021/1505-ato-77>

Obrigada

@ cristine@cert.br

@ Notificações para: cert@cert.br

<https://cert.br>

nic.br **cgi.br**

www.nic.br | www.cgi.br

TLP:CLEAR