OXFIDA

OXFIDA INSIGHT | AI & CYBER SECURITY

AI Accelerates Discovery. Can Organisations Keep Pace?

PHP2Shell and the New Challenges of Enterprise Resilience

DATE 13 November 2026 **REF** OXF-INS-2026-10 **VERSION** 1.0
INFORMATION AS OF 6 November 2026

1 Executive summary

What happened

Two critical vulnerabilities were identified in PHP, a shared foundation used by many websites, prompting government bodies and companies across several countries to temporarily suspend services while applying fixes and checking their systems. According to the companies involved, AI was used to analyse the vulnerabilities and demonstrate that they could be exploited, and exploitation attempts were observed from four days before the fix was released. Successful exploitation was subsequently observed on decoy servers under specific configurations. The research firm Dream reported that a total of 25 Palo Alto Networks firewalls and management servers had been compromised, three of which were reportedly in an environment that appeared to be affiliated with the US government. Palo Alto Networks also confirmed the exploitation, and says it independently verified the tooling used in the attack and the reported results. As of 6 November, however, the total number of affected organisations, the nature and volume of information stolen and the scale of losses had not been established.

The management challenge this event reveals

At its core is **the gap between the point at which exploitation becomes possible and the point at which organisations complete their countermeasures.** Even as discovery and analysis accelerate, developing and delivering safe fixes, identifying what is affected and coordinating with business units and suppliers all take time. The existence of a fix does not by itself remove the risk.

In that gap, organisations must decide, on uncertain information, whether to continue, restrict or stop services, and then act. A precautionary shutdown is not a failure of response but one way to avoid compromise, though it carries real business impact. And because of dependence on shared foundations, many organisations face the same decision at the same time.

Four response capabilities organisations need

- **Know the scope of impact:** know the systems in operation, their dependencies, critical operations and who is responsible, and identify the impact on the organisation quickly.
- **Shorten the time to complete countermeasures:** prepare, in normal times, to verify and apply updates, migrate end-of-life versions and work with suppliers.
- **Decide and act with incomplete information:** set criteria and authority for continuing, restricting, stopping and restoring services, and reduce time spent waiting for decisions.
- **Detect compromise, limit harm and recover:** be ready to detect compromise, contain damage, and recover and resume operations safely.

AI-accelerated discovery can also widen the opportunity to fix flaws early. Turning that opportunity into better security requires strengthening both technical measures and the organisation's ability to decide and act.

Scope of this paper. For immediate response — affected versions, fixes, interim mitigations and checks for compromise — follow the official guidance from the PHP project, national agencies and vendors listed in Appendix A. This paper focuses on what the event reveals and on the long-term measures it calls for.

2 The event in brief

A flaw in the shared foundation of many websites

PHP is software that works behind websites, processing users' requests and generating pages. PHP2Shell is the name given to two critical vulnerabilities found in the core of PHP. The problem lay not in how individual organisations built their sites, but in a foundation that many sites share. According to W3Techs, as of 5 November, 69.1% of websites whose server-side language is known use PHP. [2][8]

Unauthorised actions and data theft without logging in

The two are separate flaws: CVE-2026-109618 arises when memory runs out partway through handling a request, leaving data in a corrupted state, and CVE-2026-109619 arises when state left over from an earlier request is reused by a later one. [38] Under certain conditions, an attacker can exploit these flaws to run malicious programs on the server without logging in, or to read information from memory while it is being processed. This can include passwords and credentials used to connect to other systems. Both flaws are rated 9.2 under CVSS v4.0, in the highest "Critical" band. [2][4][32] How far the damage reaches depends on the server's privileges and what it connects to. A related vulnerability in the same feature, CVE-2026-115560, has also been disclosed; it requires an uncommon configuration, and in this paper PHP2Shell refers to the two vulnerabilities above. [42]

AI helped show how the flaws could be exploited

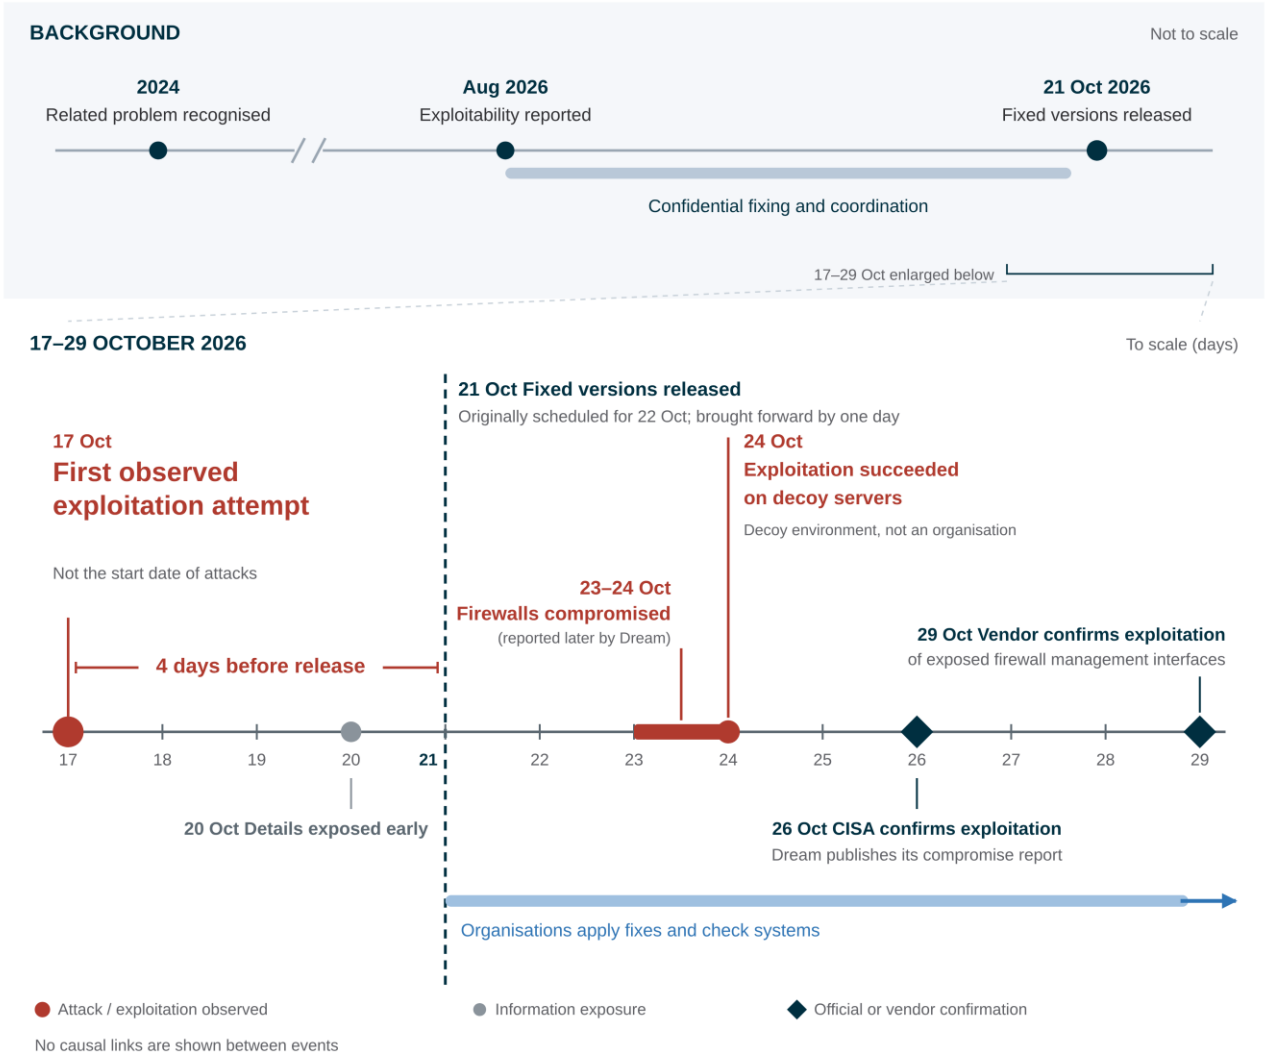
According to Wiz, the vulnerable code had been present for more than a decade, with releases back to PHP 5.4 affected; signs of a related problem were recognised in 2024. Working out how it could be exploited and how to fix it safely, however, proved difficult. In August 2026, researchers at the cybersecurity firm Calif used AI, under an OpenAI support programme, to demonstrate that a PHP flaw could be exploited. OpenAI and Wiz also reported the problem independently, and Wiz says it also used AI. This was not discovery by AI alone: human researchers directed and validated the work. [1][4][5] While AI is reported to have supported discovery and analysis, on the attack side a research firm has reported traces of AI tools in an attacker's infrastructure, but there is no direct evidence that AI was used to write the attack code (see 3.2).

Attempts before the fix, and precautionary shutdowns

Fixed versions were originally scheduled for release on 22 October, but were brought forward by one day to 21 October 2026 after related information leaked and a proof-of-concept video was published by mistake. As a result, some Linux distributors could not provide updated runtime images at the same time as the release; these followed several hours later. Amazon observed an exploitation attempt on 17 October, ahead of the known information leaks; neither the source of the attempt nor how the attacker obtained information about the flaw has been identified. [1][6][18]

Reports indicate that government bodies and companies in many countries temporarily took websites and services offline to apply fixes and check their systems. These shutdowns do not necessarily mean that an intrusion was confirmed; most were precautionary. Successful exploitation has, however, been observed on decoy servers (honeypots) and, subsequently, on internet-exposed firewall management interfaces (see 3.2). [39][40] As of 6 November, the overall scale of compromise and data leakage at ordinary organisations had not been established. [19][20][21][26][33]

Figure 1 Discovery, fixes and attacks over time: exploitation attempts observed before disclosure



Notes on Figure 1. Events shown are based on [1][2][6][33][37][39][40]. The confidential coordination phase lasted 71 days, from 11 August to 21 October; the PHP Foundation and the vendor concerned have given differing accounts of how patched images came to be released early. [1][18] Wiz says it found the first flaw on 19 August and reported the flaws to the PHP security team on 24 August. OpenAI says Calif used GPT-5.6 Sol; Wiz says it used Wiz Atlas, Gemini 4 Argon and GPT-5.6 Sol (company accounts). [4][5]

3 Global impact and the extent of confirmed harm

PHP2Shell led to precautionary service suspensions in several countries, and compromise of devices operated by real organisations has also been confirmed. This section separates the business impact of service suspensions from the state of confirmation of attacks and compromise.

3.1 Service suspensions in each country

The table is based on official statements and news reports, as cited in each row.

Country / region	Main targets and suspension	Confirmed impact and caveats
United Kingdom	Central government: 32 services unavailable and four partly unavailable, as the UK technology news site The Register reports a minister telling the House of Lords; the Office for Budget Responsibility and several councils were reported affected [29]; Scotland: according to The National, NHS Scotland's website was offline [30]	Loss of public services, described by the minister as largely precautionary; no compromise of central government systems had been confirmed (report of 22 Oct)
Netherlands	DigiD website closed for most of 21 October [19]; the national cyber security centre's site, the government portal Overheid.nl and the Port of Rotterdam's website also taken offline [20]	Website closures; the DigiD website was accessible again later the same day (21 Oct)
Finland	Government public network services taken offline by the state IT provider [20]	Suspension of public services; detailed scope not published (22 Oct)
Australia	The Australian Signals Directorate reported an outage that may be affecting government websites, including cyber.gov.au; Home Affairs said some websites were temporarily unavailable while risks were mitigated [21]	Temporary unavailability; the Directorate said it was not aware of active exploitation (22 Oct)
Japan	According to Impress Watch, pixiv, DMM.com, niconico, BASE, CoDMON, parts of Sakura Internet and a number of online games suspended services for emergency maintenance [25]; the National Cybersecurity Office urged site operators to prepare [28]	Simultaneous suspension across sectors for emergency maintenance (21 Oct)

Precautionary suspensions are distinct from harm caused by attacks, but interrupting services is a real business impact for the organisations concerned and their users.

3.2 How far attacks and compromise have been confirmed

Exploitation attempts were observed from four days before the fix was released. [6] Among organisations in Check Point's observation base, about one in ten saw attempts, rising to 14% in Europe (published 22 October). These figures show the share of organisations where attempts were detected, not where intrusions succeeded. [7]

Exploitation confirmed by CISA

On 26 October, CISA added both flaws to its Known Exploited Vulnerabilities catalogue, which lists vulnerabilities it has evidence are being actively exploited, under Binding Operational Directive 26-04 for federal civilian agencies. [37] In the information reviewed, CISA has not published the basis for its confirmation or details of any affected organisations, and this paper does not infer affected organisations or losses from the listing. The KEV catalogue sets 29 October as the due date for federal civilian agencies, which are required to apply fixes or take systems offline and to carry out forensic triage. [44]

Commands executed on decoy servers

According to The Register, security researcher Adel Karimi recorded 35 successful shell executions on 24 October on three decoy servers he operated, and published the results on GitHub. The servers ran PHP 8.4.26 inside the Apache web server (the "mod_php" set-up). A shell is the mechanism through which commands are given to a computer, so attackers actually made the servers run their commands. The observations were consistent with exploitation of CVE-2026-109618, with no evidence that CVE-2026-109619 was involved. The decoy servers handled all requests in a single process, a set-up that makes the attack easier (single-worker), so the results do not indicate the success rate in production environments. [33]

Compromise confirmed on network devices

On 29 October, Palo Alto Networks confirmed that PHP2Shell was being actively exploited against PAN-OS management interfaces exposed to the internet. [39] Its threat research unit, Unit 42, says it independently verified the attacker's tooling and the results reported by the research firm Dream. [14]

According to Dream, the attacker combined PHP2Shell with a separate Linux kernel vulnerability to gain root access on a total of 25 firewalls and management servers. Three of these were reportedly in an environment that appeared to be affiliated with the US government. In that environment, Dream also reported the establishment of a channel into the internal network, the extraction of credentials for an LDAP service account and the listing of users and groups; whether subsequent attempts, such as creating accounts, succeeded is not known. Dream also reported that AI tools had been installed in the attacker's infrastructure and that the code suggests AI-assisted development, but states that it has no direct evidence that AI was used to write the exploit. [40]

Limits of confirmation on the scale of harm

The public information reviewed up to 6 November does not establish the total number of affected organisations, the nature and volume of information stolen or the scale of losses. An organisation that finds no traces in its own systems cannot conclude from that alone that it was not attacked. [3]

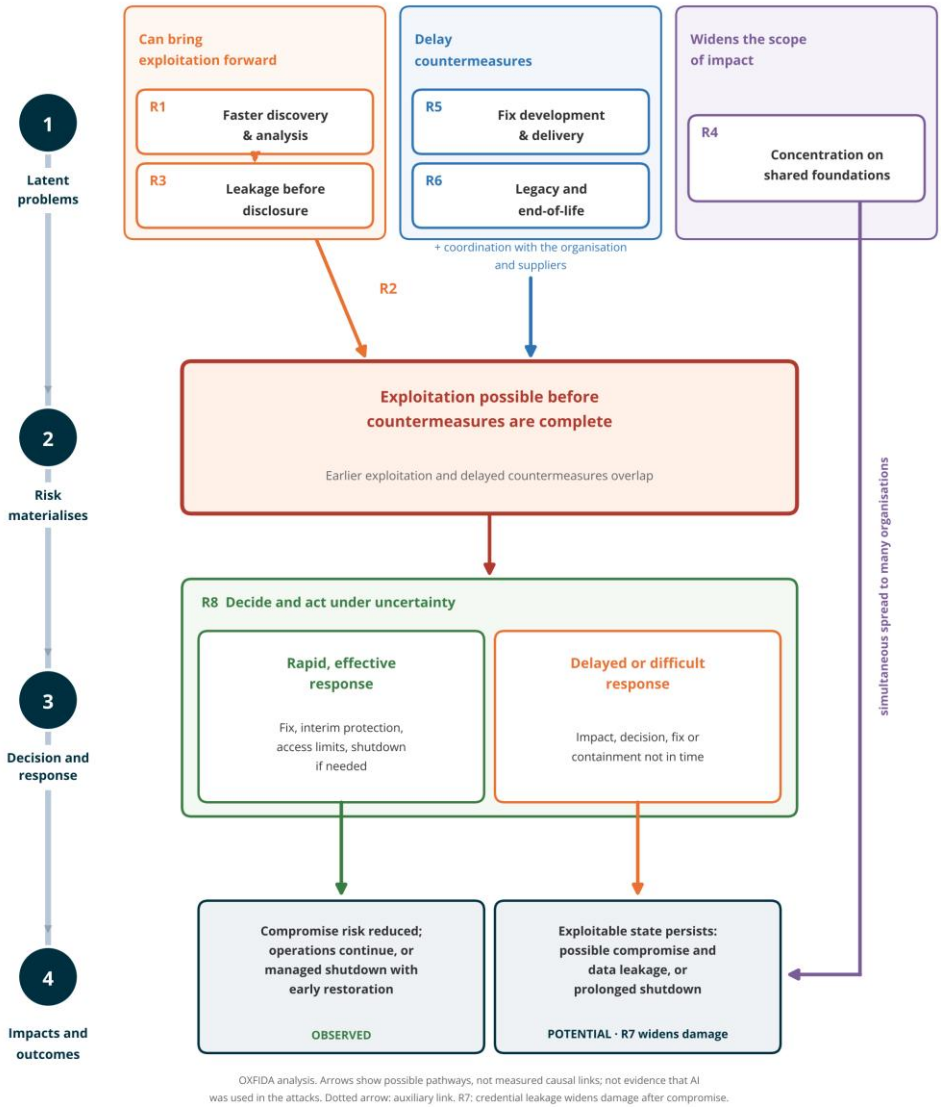
> **Note.** An early assessment by Palo Alto Networks' Unit 42 linking an intrusion at a European government agency to PHP2Shell was retracted by Unit 42 on 25 October. [14][27] Separately, SOCRadar, a threat-intelligence firm that tracks cyber attacks, found exploit tools and a script for selecting target sites on a server used by an attacker (reported 22 October, updated 23 October). The script had checked about 20 sites; none of the 11 whose results were saved met the conditions needed for the exploit to work. This shows that an attack was being prepared, not that a compromise occurred. [36]

4 The structure of risk: what shapes the outcome

What this event exposed is the gap between the point at which a vulnerability can be exploited and the point at which countermeasures are complete. How far harm spreads in that gap depends on how organisations decide, with incomplete information, whether to continue, restrict or stop systems and services, and when to resume them, and how they carry out those decisions. Figure 2 shows how this gap relates to response decisions.

4.1 Overview of the risk

Figure 2 Response decisions and actions, and their relationship to compromise risk and business impact



How risk leads to harm and business impact

The figure makes three points.

- **Countermeasures may not keep pace with faster discovery and analysis.** Faster discovery and analysis of vulnerabilities, and information leaking before disclosure, can bring forward the point at which exploitation becomes possible. Developing and delivering fixes, applying them to older systems and coordinating within the organisation, by contrast, take time. A gap can therefore open between the point at which exploitation becomes possible and the point at which countermeasures are complete.
- **Response decisions and actions change the harm and business impact.** Timely fixes, access restrictions and, where needed, shutdowns can contain compromise risk. If decisions or actions are

delayed, the exploitable state and business disruption may last longer. A precautionary shutdown to prevent harm is not a failure of response, but one of the options for limiting harm.

- **Dependence on shared foundations and leaked credentials can widen the impact.** A problem in a shared foundation used by many organisations can affect several of them at once. If credentials leak, attackers may use them as a foothold to extend harm to the systems that the compromised system connects to.

These are possible pathways to harm and business impact, not a sequence that will necessarily unfold. The table below sets out the facts confirmed in this event and what each indicates from a risk perspective.

4.2 Facts confirmed in PHP2Shell, and the risks they reveal

The table sets out eight factors related to the risk. We do not score them, as there is not enough evidence to assess each one numerically. R1–R7 cover factors that bring exploitation forward, delay countermeasures or widen the damage. R8, by contrast, covers the organisation's response, which shapes the harm and business impact.

Factor	What happened in PHP2Shell	What it means structurally
R1 Faster discovery and analysis	Several parties reported the flaws independently; AI-assisted analysis and demonstration of exploitability were reported (company accounts)	AI-assisted discovery and analysis may create opportunities for earlier remediation, while also accelerating independent discovery and analysis by others (OXFIDA analysis). The UK NCSC urges organisations to prepare for an AI-driven vulnerability patch wave [10]
R2 Exploitable before or just after disclosure	An attempt was observed four days before the fix [6]; successful exploitation was later observed on decoy servers [33] and on exposed firewall management interfaces [39][40]	The disclosure date cannot be treated as the start of the period of danger, so detection and investigation that assume exploitation before disclosure become more important. AI-driven acceleration on the attacker side is a further concern [9][15]; traces of AI tools in an attacker's infrastructure have been reported, but there is no direct evidence that AI was used to write the attack code [40]
R3 Information leakage before disclosure	Identifiers, a demo video and patched images escaped before the planned release; accounts differ on how the patched images came to be released early	Clues for attackers can become available before defenders are ready
R4 Concentration on shared foundations	A problem in the shared PHP foundation forced organisations in several countries and sectors to respond at once	How far impact spreads depends not only on each organisation's controls but also on shared dependencies
R5 Constraints on fix development and delivery	Maintainer constraints stalled a fix [1]; products embedding PHP took longer to receive fixes [26]. PHP's fix was released on 21 October, while PAN-OS fixes were released by branch between 25 and 29 October; the compromises Dream later reported occurred on 23–24 October, before that period, although mitigations such as restricting management access had been advised [39][40]	Even as discovery accelerates, the capacity to fix, verify and distribute may not grow at the same pace; organisations cannot control on their own when fixes are delivered

Factor	What happened in PHP2Shell	What it means structurally
R6 Legacy and end-of-life constraints	A substantial share of PHP websites run versions that no longer receive official fixes	Environments that cannot obtain or apply fixes prolong the exploitable state; as a past comparable case, exploitation of Log4Shell was observed for a long period after disclosure [16][17]
R7 Credential and secret leakage	Reading passwords, tokens and API keys from memory was demonstrated; the scale of real-world theft is unknown [4]. Separately, after compromise, credentials were reportedly obtained from a device's configuration and used to attack the internal directory, a post-compromise path distinct from memory leakage [40]	Damage can extend beyond the compromised server to the systems its credentials reach
R8 Decide and act under uncertainty	Precautionary suspensions were reported in several countries [19][20][21]; the DigiD website reopened later on 21 October [19]; the UK NCSC advised taking systems offline where mitigations could not be applied [11]; the UK government organised a cross-departmental response before disclosure [29]	Even when a fix exists, delays in deciding and coordinating delay the completion of countermeasures

Notes

R1: The PHP security team published 15 CVEs and handled more than 100 invalid or duplicate reports between 1 July 2026 and the release of the fix. [1] This shows the load on maintainers, but without a comparison with earlier periods or information on AI use, these figures alone do not show that AI accelerated discovery.

R2: A 2024 preprint tested, in an experimental setting, whether an LLM agent (GPT-4) could exploit 15 known vulnerabilities; it succeeded in 87% of cases when given the CVE description and in 7% without. [15] This suggests that public information can help AI-assisted exploitation, but the study has not been peer reviewed and does not show that AI was used in the PHP2Shell attacks.

R4: 69.1% of websites whose server-side language is known use PHP (as of 5 November). [8] Unit 42 reported 22.6 million potentially vulnerable exposed instances. [14] NCSC-NL compiled an inventory of products with a PHP dependency. [13] These show how widely PHP is used and depended on, but 22.6 million is a count of instances, not of unique websites or organisations, nor of actual victims.

R5: Cisco lists fixed releases for Expressway and Duo Access Gateway, and in its update of 6 November it marked fixed release 26.2.1.9 for URWB (Ultra-Reliable Wireless Backhaul) and Unified Industrial Wireless Software as available. [43] Because fixes arrive at different times for different products, the point at which users can complete countermeasures can also differ.

R6: Among PHP websites whose version W3Techs can identify, about 34% run PHP 7 or older (as of 5 November). [8] PHP 8.1 and older are outside the PHP project's regular updates; for this vulnerability, however, minimal source backports were provided for PHP 7.4, 8.0 and 8.1, and some distributions and LTS vendors also provide fixes, so users need to check their supplier's status. PHP 8.2 support ends on 31 December 2026. [2] The share of older versions alone does not show how many sites are affected by this vulnerability or how many were slow to respond.

> **OXFIDA analysis** If AI makes it possible to find vulnerabilities sooner, attackers gain more opportunities to exploit them, but defenders may also be able to fix them earlier. Whether this opportunity can be seized depends on whether organisations have the four response capabilities set out in the next section in place, in normal times.

5 Four response capabilities organisations need

For immediate action, refer to the official guidance listed in Appendix A. Section 4 set out why a vulnerability can be exploited before countermeasures are complete, and the factors that shape harm and business impact. Building on this, this section sets out four response capabilities that organisations should prepare in normal times.

Following the relationships shown in Figure 2, the four are: ① knowing the scope of impact, ② shortening the time to complete countermeasures, ③ deciding and acting with incomplete information, and ④ detecting compromise, limiting harm and recovering. They are not steps to be taken in sequence, but capabilities to strengthen in parallel, in normal times.

Capability and purpose	Main actions	Question to test readiness	Related factors
① Know the scope of impact **Quickly determine which of the organisation's systems and operations are affected**	● Know which systems run where, who manages them and which operations they support ● Use the SBOM and asset inventory as a starting point, including PHP bundled in vendor products and container images; confirm the versions and fix status used by running processes, and set priorities according to external exposure, the conditions for exploitation and the impact on critical operations	When a critical vulnerability emerges, how many hours does it take to identify the affected systems and who is responsible for them?	R4, R6 (and all factors)
② Shorten the time to complete countermeasures **Keep the period in which a vulnerability can be exploited as short as possible**	● Build an organisation-wide capability to verify, distribute and apply updates quickly and continuously, and to confirm the results on live systems [10][34] ● Identify where end-of-life versions are used and retire them on a plan with budget, owners and deadlines ● Identify and manage supply-chain and open-source dependencies, check contract terms for fixes and notification, and support maintainers; apply the same policy to AI-related components [24]	How many critical systems still lack a fix or containment, and who manages their deadlines?	R1, R2, R3, R5, R6

Capability and purpose	Main actions	Question to test readiness	Related factors
③ Decide and act with incomplete information **Limit delays while waiting for decisions, and unnecessary business impact**	● Assume that critical vulnerabilities can emerge without warning; agree in advance who decides to continue, restrict, stop or resume services, on what criteria, and how much authority is delegated to those on the ground; prepare procedures for temporary suspension ● To make decisions possible with incomplete information, set procedures for checking primary sources, recording the basis for decisions and exceptions, conditions for re-evaluation, coordination with stakeholders, and escalation to senior management ● Confirm the reporting and notification duties that apply, with owners and decision procedures (e.g. EU CRA reporting duties for actively exploited vulnerabilities apply to manufacturers from 11 Sep 2026) [22][23]	Can the necessary restriction or shutdown be decided and carried out at night, or when the person responsible is unavailable?	R8 (secondary: R2, R3, R5)
④ Detect compromise, limit harm and recover **Identify compromise early and limit its spread and the burden of recovery**	● Be able to detect anomalies and compromise without relying only on specific CVE information; set policies for collecting and preserving logs so that activity before disclosure can be investigated retrospectively ● Prevent one compromise from spreading to other systems through network segmentation, least privilege, and limiting the exposure and lifetime of credentials ● Define the technical conditions for containment, removing the cause and traces of compromise, restoring a safe state and resuming services [35], and test them in exercises	If one system is compromised, how far could the harm spread? And what must be confirmed to judge that operations can safely resume?	R2, R4, R7, R8

The four capabilities are OXFIDA's regrouping of measures for an executive audience, drawing on the functions of NIST CSF 2.0 [31]; they are not an official CSF classification. Each measure spans several functions and factors; the mapping is not one-to-one.

6 Perspectives from AI experts

The comments below are the personal views of the individuals concerned.

> "Although the recent PHP2Shell vulnerabilities were discovered using AI-assisted techniques, their identification also highlights the growing risk of AI-accelerated exploitation. The same capabilities that enable security researchers to identify complex vulnerabilities can be leveraged by attackers to rapidly analyse flaws, develop working exploits and automate attacks at scale. This dual-use nature of AI risks narrowing the gap between vulnerability disclosure and widespread exploitation, placing ever greater pressure on organisations to detect, prioritise and mitigate security weaknesses before attackers can capitalise on them."
>
> **Stephen Roberts** Professor of Machine Learning, University of Oxford

> "PHP2Shell was a preview of how AI will change the governance of risk. AI-assisted discovery, particularly of software vulnerabilities, will greatly accelerate the exploitation of flaws. Governance built around periodic review and a static register of known risks cannot keep up with the new pace of change. What is needed instead is dynamic governance that is capable of managing uncertain risk landscapes, providing human overseers with honest and reliable risk assessments. PHP2Shell is a warning shot that all organisations should heed as they transform their governance of risk."
>
> **Maike Osborne** Professor of Machine Learning, University of Oxford

Appendix A Official guidance for immediate response

Follow these sources for current versions, interim mitigations and checks for compromise; they are updated as the situation develops.

- The PHP Group: security advisories and release announcements [2]
- JPCERT/CC: JPCERT-AT-2026-0033 (Japanese) [3]
- UK NCSC: "Critical security issue affecting PHP" [11]
- Canadian Centre for Cyber Security: AL26-026 [12]
- NCSC-NL: inventory of software products with a PHP dependency [13]
- Palo Alto Networks: security advisory on the impact on PAN-OS and other products [39]
- Cyber Security Agency of Singapore: AL-2026-157 (network and security products) [41]
- Red Hat: RHSB-2026-013 [42]

Glossary

Term	Meaning
CVE	A unique identifier assigned to a publicly known vulnerability.
CVSS	Common Vulnerability Scoring System: a 0–10 scale of a vulnerability's technical severity; 9.0 and above is rated Critical.
End of life (EOL)	A version that no longer receives security fixes from its maintainer.
Memory disclosure	An attacker reading data from a program's memory, such as passwords or keys.
SBOM	A list of the software components in a product or system.
Remote code execution (RCE)	An attacker running their own commands on another computer over the network.
Honeypot	A decoy server set up to observe attacks.

Zen Hamada | Senior Vice President, OXFIDA

Zen Hamada leads business and product strategy at OXFIDA, including AI governance and model risk management. A former Vice President in Goldman Sachs' Investment Banking Division, he draws on experience advising senior management and building AI businesses to examine how technological change affects corporate risk management and decision-making.

Disclaimer. This paper is for general information and is not professional security, legal or investment advice.

References

[1] The PHP Foundation, "Today's PHP Security Release, and How We're Strengthening PHP Security", 21 Oct 2026. https://thephp.foundation/blog/2026/10/21/todays-php-security-release-and-how-were-strengthening-php-security/

[2] The PHP Group, security advisory (21 Oct 2026), informational advisory (23 Oct 2026) and PHP 8.2.35 release note. https://www.php.net/archive/2026.php

[3] JPCERT/CC, JPCERT-AT-2026-0033, 22 Oct 2026, updated 29 Oct 2026. https://www.jpcert.or.jp/at/2026/at260033.html

[4] Wiz Research, "PHP2Shell: AI uncovers critical RCE vulnerabilities affecting 70% of the Internet" (title as published), 21 Oct 2026. https://www.wiz.io/blog/php2shell-php-rce

[5] OpenAI, "Patch the Planet: global response coordination as agents uncover critical vulnerabilities", 21 Oct 2026. https://openai.com/index/coordinated-disclosure-of-critical-vulnerabilities/

[6] Amazon Threat Intelligence, "Amazon Threat Intelligence identifies PHP2Shell exploitation prior to public disclosure in MadPot", 21 Oct 2026, updated 22 Oct 2026. https://aws.amazon.com/blogs/security/amazon-threat-intelligence-identifies-php2shell-exploitation-prior-to-public-disclosure-in-madpot/

[7] Check Point Research, "PHP2Shell Update: One in ten organizations impacted as Europe bears the brunt", 22 Oct 2026. https://blog.checkpoint.com/ai-security/php2shell-one-in-ten-organizations-impacted-as-europe-bears-the-brunt/

[8] W3Techs, "Usage statistics of PHP for websites", as of 5 Nov 2026. https://w3techs.com/technologies/details/pl-PHP

[9] UK NCSC, "Impact of AI on cyber threat from now to 2027", 7 May 2025. https://www.ncsc.gov.uk/report/impact-ai-cyber-threat-now-2027

[10] UK NCSC (O. Whitehouse), "Preparing for a 'vulnerability patch wave'", 1 May 2026. https://www.ncsc.gov.uk/blogs/prepare-for-vulnerability-patch-wave

[11] UK NCSC, "Critical security issue affecting PHP", 21 Oct 2026. https://www.ncsc.gov.uk/news/critical-security-issue-affecting-php

[12] Canadian Centre for Cyber Security, AL26-026, 21 Oct 2026. https://cyber.gc.ca/en/alerts-advisories/al26-026-critical-vulnerabilities-impacting-php-cve-2026-109618-cve-2026-109619

[13] NCSC-NL, php2shell operational information and software inventory. https://github.com/NCSC-NL/php2shell

[14] Palo Alto Networks Unit 42, "Threat Brief: Critical Vulnerabilities in PHP", Oct 2026 (updated 25 Oct 2026, retracting the 23 Oct update on a European government agency; updated 29 Oct 2026 on PAN-OS exploitation). https://unit42.paloaltonetworks.com/cve-2026-109618-cve-2026-109619/

[15] Fang, R. et al., "LLM Agents can Autonomously Exploit One-day Vulnerabilities", arXiv:2404.08144 (preprint), 2024. https://arxiv.org/abs/2404.08144

[16] Singh, A. et al., "A Longitudinal Measurement Study of Log4Shell Exploitation from an Active Network Telescope", arXiv:2601.04281 (preprint), Jan 2026. https://arxiv.org/abs/2601.04281

[17] Hiesgen, R., Nawrocki, M., Schmidt, T. C., Wählisch, M., "The Log4j Incident: A Comprehensive Measurement Study of a Critical Vulnerability", IEEE TNSM 21(6), 2024. https://doi.org/10.1109/TNSM.2024.3440188

[18] The Register (J. Lyons), "PHP blames supply chain security biz Chainguard for putting huge chunks of the web at risk", 22 Oct 2026, updated 23 Oct 2026 (includes Chainguard's response). https://www.theregister.com/security/2026/10/22/php-blames-supply-chain-security-biz-chainguard-for-putting-huge-chunks-of-the-web-at-risk/5304611

[19] ANP via Nieuws.nl, "Website DigiD weer toegankelijk", 21 Oct 2026. https://nieuws.nl/algemeen/website-digid-weer-toegankelijk

[20] Al Jazeera, 22 Oct 2026. https://www.aljazeera.com/news/2026/10/22/whats-the-php-coding-flaw-that-could-pull-down-70-percent-of-the-internet

[21] InDaily (F. Farid and L. Garbutt-Young, with AAP), "'Critical' software flaw hits Australian government sites", 22 Oct 2026. https://www.indailysa.com.au/news/just-in/2026/10/22/critical-software-flaw-hits-australian-government-sites

[22] European Commission, "Cyber Resilience Act – Reporting obligations". https://digital-strategy.ec.europa.eu/en/policies/cra-reporting

[23] Regulation (EU) 2024/2847 (Cyber Resilience Act), Articles 14 and 71. https://eur-lex.europa.eu/eli/reg/2024/2847/oj/eng

[24] UK NCSC and partners, "Guidelines for secure AI system development". https://www.ncsc.gov.uk/collection/guidelines-secure-ai-system-development

[25] Impress Watch, [Emergency maintenance and service suspensions at major domestic services in response to PHP vulnerability] (in Japanese), 21 Oct 2026. https://watch.impress.co.jp/docs/news/2150159.html

[26] BankInfoSecurity (M. J. Schwartz), "PHP2Shell Actively Exploited, US Cybersecurity Agency Warns", 27 Oct 2026 (secondary source; reports CISA's KEV listing of 26 Oct and the related federal requirements). https://www.bankinfosecurity.com/php2shell-actively-exploited-us-cybersecurity-agency-warns-a-33219 (The page shows inconsistent date labels; the publication date is taken as 27 Oct 2026, as stated in the article.)

[27] The Hacker News, "PHP2Shell Exploitation Attempts Continue Amid Suspected N. Korea Global Scanning Activity", 25 Oct 2026 (updated with Unit 42 retraction). https://thehackernews.com/2026/10/php2shell-exploitation-hits-government.html

[28] GIGAZINE, [Niconico, pixiv and other websites and apps go into maintenance to apply emergency PHP patch] (in Japanese), 21 Oct 2026. https://gigazine.net/news/20261021-php-php2shell-maintenance/

[29] The Register (C. Page), "PHP2Shell puts dozens of UK government services out of service", 22 Oct 2026 (reporting House of Lords proceedings). https://www.theregister.com/security/2026/10/22/php2shell-puts-dozens-of-uk-government-services-out-of-service/5304840

[30] The National via Yahoo News UK, "NHS Scotland website remains offline as critical internet flaw sparks security fears", 22 Oct 2026. https://uk.news.yahoo.com/nhs-scotland-website-remains-offline-070029113.html

[31] NIST, "The NIST Cybersecurity Framework (CSF) 2.0", NIST CSWP 29, Feb 2024. https://nvlpubs.nist.gov/nistpubs/CSWP/NIST.CSWP.29.pdf

[32] Fastly, blog post on critical RCE vulnerabilities in core PHP (CVE-2026-109618, CVE-2026-109619), Japanese edition, 21 Oct 2026. https://www.fastly.com/jp/blog/critical-rce-vulnerabilities-discovered-core-php-cve-2026-109618-cve-2026-109619

[33] The Register, "CISA confirms PHP2Shell under attack as researchers clock successful exploits", 27 Oct 2026 (links to the researcher's report at github.com/0x4D31/honeypot-data; only the repository and its description of 35 case folders were confirmed, not the detailed report). https://www.theregister.com/cyber-crime/2026/10/27/cisa-confirms-php2shell-under-attack-as-researchers-clock-successful-exploits/5305681

[34] Souppaya, M. and Scarfone, K., "Guide to Enterprise Patch Management Planning: Preventive Maintenance for Technology", NIST SP 800-40 Rev. 4, Apr 2022. https://doi.org/10.6028/NIST.SP.800-40r4

[35] Nelson, A. et al., "Incident Response Recommendations and Considerations for Cybersecurity Risk Management: A CSF 2.0 Community Profile", NIST SP 800-61 Rev. 3, Apr 2025. https://doi.org/10.6028/NIST.SP.800-61r3

[36] SOCRadar, "PHP2Shell Exploit Tooling Found on Exposed Attacker Server", 22 Oct 2026, updated 23 Oct 2026. https://socradar.io/blog/php2shell-exploit-tooling-attacker-server/

[37] CISA, "CISA Adds Two Known Exploited Vulnerabilities to Catalog", 26 Oct 2026. https://www.cisa.gov/news-events/alerts/2026/10/26/cisa-adds-two-known-exploited-vulnerabilities-catalog

[38] OffSec, "PHP2Shell – Critical Vulnerabilities in PHP Core (CVE-2026-109618/CVE-2026-109619)", 22 Oct 2026. https://offsec.com/blog/php2shell-critical-vulnerabilities-in-php-core/

[39] Palo Alto Networks, "CVE-2026-109618 Impact of Critical Vulnerabilities in PHP (CVE-2026-109618, CVE-2026-109619) on Palo Alto Networks Products", security advisory, 21 Oct 2026, updated 29 Oct 2026. https://security.paloaltonetworks.com/CVE-2026-109618

[40] Dream, "Chained PHP2Shell Exploit Roots 25 Servers", 26 Oct 2026, updated 29 Oct 2026. https://dreamgroup.com/blog/chained-php2shell-exploit-roots-25-servers

[41] Cyber Security Agency of Singapore, "Active Exploitation of PHP2Shell Vulnerabilities Affecting Network and Security Products" (AL-2026-157), 3 Nov 2026. https://www.csa.gov.sg/alerts-and-advisories/alerts/al-2026-157/

[42] Red Hat, "RHSB-2026-013 Request handling memory corruption - php (CVE-2026-109618, CVE-2026-109619, CVE-2026-115560) - "php2shell"", 21 Oct 2026, updated 23 Oct 2026. https://access.redhat.com/security/vulnerabilities/RHSB-2026-013

[43] Cisco, security advisory cisco-sa-php-october26-Jj384Daq on PHP vulnerabilities in Cisco products, version 1.11, updated 6 Nov 2026. https://sec.cloudapps.cisco.com/security/center/content/CiscoSecurityAdvisory/cisco-sa-php-october26-Jj384Daq

[44] CISA, Known Exploited Vulnerabilities Catalog, entries for CVE-2026-109618 and CVE-2026-109619 (added 26 Oct 2026). https://www.cisa.gov/known-exploited-vulnerabilities-catalog