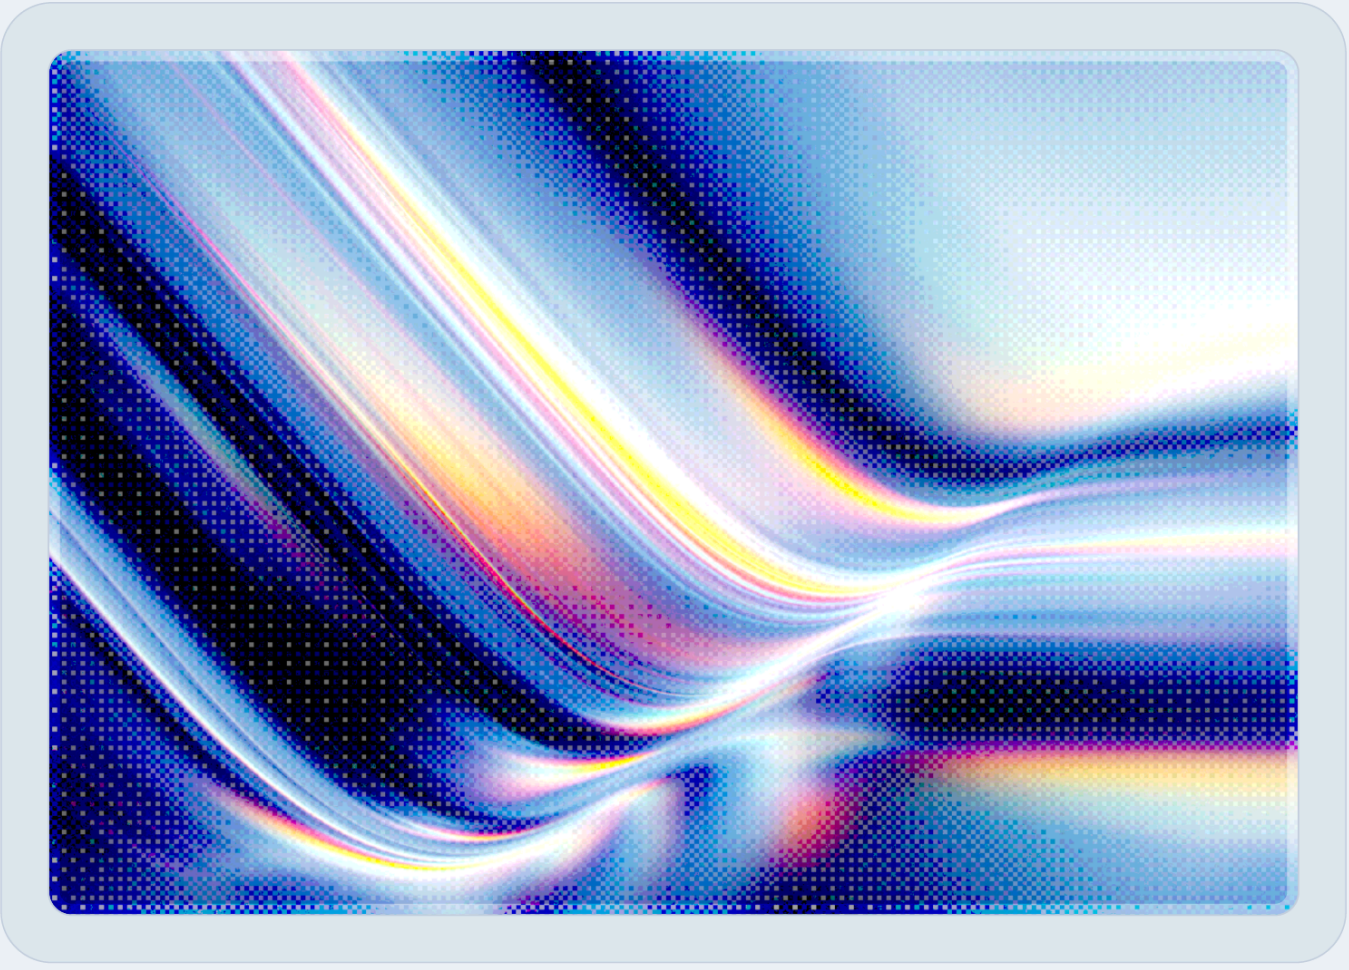
CSAI Foundation | Cloud Security Alliance

Project Zero's Full X Policy: Governing for a 30-Day Disclosure World

Security and governance implications of fixed-deadline vulnerability disclosure in the AI era

2026-11-15

AI-assisted Rapid Research



CSAI

CSA cloud security alliance®

© 2026 Cloud Security Alliance. Some rights reserved.

You may download, store, display, view, print, redistribute, and link to this document in its original, unmodified form, provided that attribution to the Cloud Security Alliance is maintained and all trademark and copyright notices remain intact.

This document may not be modified or altered. You may quote portions of the document as permitted by the Fair Use provisions of the United States Copyright Act, provided that attribution is given to the Cloud Security Alliance.

This document may be shared on professional and social media platforms in its original form with attribution.

This document was generated with AI assistance and has not undergone official CSA review and approval processes.

Key Takeaways

On November 13, 2026, Google Project Zero replaced its "90 + 30" disclosure policy, in place since 2021, with a fixed-deadline model called "Full X." Under Full X, Project Zero discloses a vulnerability X days after reporting it to the vendor, whether or not a fix has shipped. The deadline is 90 days now, falls to 60 days on February 2, 2027, and falls to 30 days on January 3, 2028 [1]. The change also removes the 30-day patch-adoption window, the grace period, and the special variant rules. The policy reaffirms that severity does not affect disclosure timing, and whether a bug has been "fixed" is no longer a negotiated decision point [1].

Project Zero's stated rationale is that large language models (LLMs) speed up both discovery and exploitation. They also make it more likely that several parties, some of them malicious, find the same bug. In Project Zero's words, "zero patch-to-exploitation time is a reasonable, though imperfect, model of attacker behavior," and patch soak time "is no longer meaningful" [1]. Vendors face the new deadline directly on every report. This note focuses on an audience that may feel the policy's effects less visibly but just as sharply: the enterprises that consume vendor software.

For those enterprises, the window before public disclosure shrinks to a fixed 30 days by 2028. By early 2028, technical details of vulnerabilities in widely deployed products may be public about 30 days after the vendor is notified. Many of those products may not yet have a patch, and those that do may have one that customers have not yet deployed. The 30-day post-patch adoption buffer is also gone. Once a fix ships, details follow at the fixed deadline, not 30 days later. Organizations that relied on that cushion need faster deployment or compensating controls.

Vendor severity ratings do not buy time. As under the previous policy, Project Zero's deadline applies regardless of how a vendor rates a bug. The shorter, fixed window, however, means a vendor's decision to defer a "low-severity" fix now leaves far less time before publication. Project Zero states that it only reports bugs with substantial severity and demonstrated security impact [1]. A vendor's low rating of such a report is therefore a weaker basis for an enterprise's own prioritization decisions than it may appear.

These pressures make non-patch mitigations a first-class governance requirement. Project Zero explicitly encourages vendors to build feature-flag, filtering, and hotpatch capabilities [1][2], and enterprises should begin asking suppliers for evidence of them. Finally, governance programs should plan for Project Zero to be a leading indicator. Its policy has historically influenced industry norms [3], and regulators are

separately shortening related timelines for exploited vulnerabilities [9][13]. Governance programs should therefore plan for the possibility that 30-day disclosure becomes common, rather than treating it as specific to Google.

Background

From 90 + 30 to Full X

Project Zero's deadline-driven disclosure program has been revised several times. The 2021 edition set the model most of the industry came to know. Vendors had 90 days to fix a bug. If they fixed it within that window, technical details were withheld for a further 30 days so users could install the patch. Vendors could also request a 14-day grace period if they expected a fix by day 104 [3]. The 2021 post explicitly considered a move to 60 + 30 and judged that "jumping straight to a 60+30 policy (or similar) would likely be too abrupt and disruptive." It committed instead to shortening timelines gradually while monitoring patch-time data [3].

In July 2025 Project Zero added a "Reporting Transparency" trial. About a week after filing a report, it began publicly listing the affected vendor, the product, the report date, and the 90-day deadline, without technical details [4]. The trial targeted what Project Zero calls the "upstream patch gap": the period after an upstream vendor has a fix available but before downstream products have integrated it [4]. The current policy keeps that transparency mechanism, with listings published within one week of a report [6], and changes the core timeline.

The November 2026 "Policy and Disclosure: 2026 Edition," authored by Natalie Silvanovich, introduces Full X. The table below summarizes its schedule and the elements it removes [1].

Element	90 + 30 (2021–2026)	Full X (from November 13, 2026)
Base deadline	90 days to fix	90 days now; 60 days from Feb 2, 2027; 30 days from Jan 3, 2028
Disclosure trigger	Day 90 if unfixed; fix date + 30 days if fixed	Fixed X days after report, regardless of fix status

Element	90 + 30 (2021–2026)	Full X (from November 13, 2026)
Patch-adoption window	30 days after fix	Removed
Grace period	14 days on request (3 days for in-the-wild bugs)	Removed
Variant handling	Special rules	Removed
Role of severity	Not a factor in timing	Explicitly not a factor in timing (unchanged)
"Is it fixed?" determination	Decision point for timing	Tracked and communicated, but not a decision point
In-the-wild exploitation	7-day deadline + 30 days	7 days, handled under Google's application security disclosure policy [5]
Other exceptions	Case by case	Already-public bugs, vendor request for earlier release, occasional hardware exceptions (e.g., Spectre/Meltdown-class issues) [1][6]

Why Project Zero Moved Now

The 2026 post gives four connected reasons [1]. First, LLM-assisted tooling has accelerated both discovery and exploit development. That raises the probability that independent parties find the same bug, and it means an embargo no longer guarantees secrecy. Second, attackers can reverse-engineer patches almost immediately, so the 30-day soak period "can create a false sense of secrecy" without protecting users. Third, Project Zero reports finding several vulnerabilities a week in targets where it once found a few a year [1]. That volume makes negotiation-heavy steps, such as agreeing whether a fix is complete, operationally unmanageable and pushes the team toward automatable rules. Fourth, severity measures impact on users, not the likelihood that attackers will find the bug. A vendor may reasonably accept a low-severity risk, but the bug remains available to adversaries.

External data is consistent with this reasoning. Mandiant's analysis of vulnerabilities exploited in 2023 found an average time-to-exploit of five days, down from 32 days across 2021–2022 and 63 days in 2018–2019. Seventy percent of the 138 exploited vulnerabilities it studied were first exploited as zero-days [7]. AI-assisted discovery has also moved from research into operations. Google reported that its Big Sleep agent, built jointly by DeepMind and Project Zero, identified a SQLite memory-corruption flaw (CVE-2025-6965) that it assessed was known to threat actors and at risk of exploitation [8]. More recently, CSA's analysis of Anthropic's Project Glasswing documented that only 97 of 1,596 AI-discovered vulnerabilities disclosed to open-source maintainers had been confirmed patched as of May 22, 2026, roughly six percent [11]. This suggests that, at least for open-source projects receiving high volumes of AI-generated reports, discovery capacity can outrun remediation capacity, though the confirmed-patch figure may also reflect reporting and confirmation lag [11].

Project Zero says the phased schedule came out of vendor consultation. Vendors worried that an immediate change would expose unfixed bugs while they were still improving their processes. Some also warned that strict deadlines could push vendors to prioritize by deadline over severity, or to stop trying to meet deadlines at all [1]. Those concerns did not stop the change, but they shaped its pace, and they point to the failure modes governance teams should watch for.

The policy also raises a broader concern. Publishing technical details before patches are widely deployed, or under a 30-day deadline possibly before a fix exists, widens the pool of attackers able to act on a bug. Project Zero itself acknowledges that delayed disclosure still has short-term benefits [1], and its 2021 post judged a faster move "too abrupt and disruptive" [3]. Whether the policy reduces net user risk will depend on whether vendors respond with faster remediation, as Project Zero intends, or by accepting missed deadlines. This note does not take a position on that question; it focuses on how enterprises should prepare.

Security Analysis

The Exposure Window Is Now Defined by the Reporter's Clock

For enterprises, the most significant change is that the worst-case exposure window is now set by a timer the vendor does not control. Under 90 + 30, a vendor that shipped a fix on day 60 bought its customers 30 days of relative obscurity. Under Full X at 30 days, the same vulnerability could be fully documented in public before many enterprises have finished change-advisory review for the patch. If a vendor cannot ship a fix within 30 days, enterprises may face a public, technically detailed vulnerability with no patch at all. In that case the defensive options are configuration changes, feature disablement, network controls, and detection.

This makes enterprise time-to-patch a governance metric with an external benchmark. Project Zero's choice of disclosure date does not stop attackers who already know about a bug. The concern is that publication widens the set of adversaries able to act on a bug, and AI-assisted patch diffing reduces the effort needed to turn a published fix into a working exploit [1][10]. CSA's CISO Community briefing on the "AI Vulnerability Storm" reached a similar conclusion earlier in 2026. It recommended that organizations prepare triage capacity for "patch floods" and pre-authorize containment actions so they do not have to negotiate them during an incident [10].

Severity-Blind Disclosure Breaks Common Prioritization Assumptions

Many enterprise vulnerability management programs implicitly assume that severity predicts urgency and that low-severity issues can wait for routine maintenance cycles. Project Zero's deadline was already severity-independent under 90 + 30. By shortening that deadline and removing the fix-date + 30 path, Full X sharpens the problem with the second half of that assumption for any product in Project Zero's research scope. Project Zero states that it only reports vulnerabilities with substantial severity and demonstrated security impact [1]. Even so, a bug that a vendor rates as low severity, and therefore schedules for a later release, will still be published at the same deadline as one the vendor rates critical [1]. Once it is public, its exploitability depends on the attacker and the environment, not the vendor's label. A low-severity bug can still be an important link in an exploit chain.

This does not mean every disclosed bug deserves emergency treatment. It does mean prioritization should rest on exploitability, exposure, and chaining potential, assessed against the organization's own environment, and not only on vendor-supplied severity. CSA has documented the limits of CVSS-only scoring and recommended context-aware assessment that incorporates exploit prediction and asset context [12]. Full X sharpens that argument with a time dimension: the date on which a vulnerability's technical details become public is now largely predictable in advance, which allows teams to schedule assessment work instead of reacting to it.

Patch Engineering Becomes a Supplier Risk Factor

Project Zero's October 2026 post on emergency patching argues that most vendors' patch speed is limited by testing and delivery, not triage or development. It describes four non-patch remediation mechanisms: feature flags, input filtering, alternate update channels, and hotpatching [2]. The post stresses that "rapid update mechanisms do not need to be heavyweight or be capable of fixing every possible bug" [2]. The 2026 policy refers vendors to these strategies directly [1].

For enterprise governance, this means a supplier's ability to mitigate a vulnerability within 30 days is now a measurable risk factor in third-party risk management. Vendors with tested feature-flag infrastructure or dynamically updatable filtering rules can protect customers before a full patch ships. Vendors that depend solely on monolithic quarterly releases generally have far fewer options. The difference is likely to be most acute for hardware, firmware, embedded systems, and operational technology, where update mechanisms are typically slowest. Project Zero itself acknowledges hardware as a case where exceptions have occasionally been warranted [1], but enterprises should not assume those exceptions will apply to their own estate.

Each mechanism also carries risk. Project Zero itself warns that alternate update channels can introduce critical vulnerabilities if downloads are not verified, and that hotpatching can create opportunities to bypass exploit mitigations [2]. A procurement or assurance process that asks about emergency remediation capability should also ask how those channels are authenticated, tested, and logged.

Downstream and Open-Source Dependencies Bear Concentrated Risk

Full X removes the buffer that allowed downstream integrators to pick up an upstream fix before details were published. Project Zero's 2025 Reporting Transparency trial was designed to give downstream parties early notice that a bug existed in a component they consumed [4]. Under Full X, that notice is still the main signal available. Organizations that build on open-source components, Linux distributions, browser engines, or mobile platforms should treat Project Zero's public reporting-transparency list as an early-warning feed. Under a 30-day deadline, the time between that listing and full publication will shrink to roughly three to four weeks, since listings can appear up to a week after the report [6].

The open-source maintainer base is the most exposed part of this chain. CSA's Project Glasswing analysis documented a structural gap between machine-speed discovery and volunteer-paced remediation [11]. A fixed 30-day disclosure deadline does not create that gap, but it makes it visible on a schedule. Enterprises that depend heavily on a small number of maintainers may need to contribute directly to remediation by funding fixes, providing engineering time, or maintaining internal patches, instead of waiting for upstream releases.

Regulatory Timelines Are Converging on the Same Window

Full X arrives while regulatory timelines are also shortening. In the United States, CISA's Binding Operational Directive (BOD) 26-04, issued June 10, 2026, revoked and superseded BOD 22-01 and its default two-week remediation deadline for Known Exploited Vulnerabilities (KEV) [9]. BOD 26-04 requires federal civilian agencies to remediate on risk-based timelines that depend on whether an asset is publicly exposed, KEV status, automatability, and technical impact. Its most urgent tiers require

remediation or mitigation within three days, plus forensic triage of the affected asset [9]. In the European Union, the Cyber Resilience Act's Article 14 reporting obligations have applied since September 11, 2026. They require manufacturers to give authorities an early warning of actively exploited vulnerabilities in their products within 24 hours of becoming aware of them [13][15].

These regimes cover different triggers from Project Zero's policy (known exploitation versus research disclosure), but the effect is cumulative. For a manufacturer whose products are placed on the EU market, a single Project Zero report could be followed by public disclosure, exploitation after publication, and, once exploitation is known, a 24-hour early-warning duty [13]. The table below compares the timing pressures an enterprise or vendor may face at once.

Driver	Trigger	Timeline	Primary obligated party
Project Zero Full X [1]	Report to vendor	90 → 60 → 30 days to public disclosure	Vendor (de facto), enterprise (exposure)
Google in-the-wild policy [5]	Evidence of active exploitation	7 days	Vendor
CISA BOD 26-04 [9]	KEV listing or agency identification, weighted by exposure, automatability, and impact	Risk-based tiers; most urgent tier three days plus forensic triage	U.S. federal civilian agencies
EU CRA Article 14 [13][15]	Manufacturer awareness of active exploitation	24-hour early warning	Manufacturers of products with digital elements

Recommendations

Immediate Actions

Security leaders should brief executive leadership and audit committees on the Full X schedule this quarter, before the 60-day transition on February 2, 2027. The briefing should name the product families in the organization's estate that fall within Project Zero's historical research scope, typically major operating systems, browsers, mobile platforms, hypervisors, and widely used libraries. It should also state whether current patch SLAs for those products fit within a 30-day total window from vendor notification. Many organizations are likely to find that the honest answer differs by business unit and asset class.

Vulnerability management teams should subscribe to Project Zero's reporting-transparency list and its public bug tracker. They should map newly listed vendors and products against their software inventory and SBOM data, and create a tracking record with the expected public-disclosure date. That date is now predictable by default, subject to the policy's few exceptions [1], so it can be scheduled in the same way as a vendor's monthly patch release.

Short-Term Mitigations

Over the next six months, organizations should reset internal patch SLAs so that, for in-scope products, a patch can be tested and deployed to internet-exposed and high-value systems within days of release. When no patch is available at disclosure, pre-approved compensating controls should be ready to apply. Change-management processes should include a pre-authorized emergency path for vendor-issued mitigations such as configuration changes and feature disablement, with rollback plans tested in advance. This follows the CSA CISO Community's recommendation to pre-authorize containment actions [10].

Prioritization logic should be revised so that vendor severity is one input among several. Exploitability, external exposure, presence in known exploit chains, and the scheduled public-disclosure date should carry explicit weight [12]. Third-party risk questionnaires should ask critical suppliers how they plan to operate under a 30-day disclosure regime. Questions should cover feature-flag coverage, filtering capability, hotpatch support, how emergency update channels are authenticated, and historical median time-to-fix for externally reported vulnerabilities [2].

Strategic Considerations

Over a 12-to-24-month horizon, organizations should treat patch velocity as an enterprise resilience metric reported to the board, not as an operational KPI within IT. They should also plan for the possibility that other research organizations, AI-assisted discovery platforms, and regulators adopt timelines similar to Project Zero's. Investment in software minimization, segmentation, and egress control reduces the number of vulnerabilities that become urgent at disclosure. Investment in detection engineering shortens the gap between publication and recognizing exploitation attempts.

Organizations that consume large volumes of open-source software should consider how they contribute to upstream remediation capacity. Under Full X, the public disclosure date no longer depends on whether a maintainer has had time to respond. Finally, vendors and enterprises alike should watch for the behavioral failure modes vendors raised during consultation. These include prioritization by deadline instead of risk, and resignation to missed deadlines [1]. Such behavior could reduce the policy's intended security benefit and would be an early signal of supplier stress.

CSA Resource Alignment

CSA has already published guidance that addresses the conditions motivating Full X. The most directly relevant is **"The 'AI Vulnerability Storm': Building a 'Mythos-ready' Security Program"** from the CSA CISO Community (May 2026) [10]. The briefing argues that AI-accelerated discovery and patch diffing have compressed the window between fix and exploit. It recommends standing up a Vulnerability Operations (VulnOps) function, preparing triage capacity for patch floods, and pre-authorizing containment actions. Full X turns those recommendations from preparation for a possible future into requirements on a dated schedule. Organizations that have already implemented the briefing's action plan will have several of the necessary capabilities in place before the 60-day and 30-day transitions.

"Project Glasswing: AI Discovery Outpaces Open Source Patching Capacity" (June 2026) [11] documents the asymmetry that makes Full X consequential for open-source dependencies, where machine-speed discovery meets maintainer-paced remediation. Its recommendations to treat available patches as emergency changes, set up direct communication channels with maintainers, and apply segmentation and least privilege during the patch window apply directly to the downstream exposure analyzed above.

CSA's **"Top Concerns With Vulnerability Data"** [12] provides the analytical basis for moving prioritization beyond CVSS. It recommends context-aware assessment and exploit prediction, which become more important as a shorter, severity-independent deadline leaves less time to correct for vendor ratings.

At the control level, the **AI Controls Matrix (AICM) v1.1** [14] provides auditable anchors for these practices. Its Threat and Vulnerability Management (TVM) domain covers remediation timelines and prioritization. Its Change Control and Configuration Management (CCC) domain addresses the emergency change authorization and rollback practices recommended above, while its Application and Interface Security (AIS) domain is relevant to the design of secure update mechanisms. Its Supply Chain Management, Transparency, and Accountability (STA) domain supports the supplier-assurance questions recommended above. Organizations can map their revised patch SLAs and supplier requirements to these control objectives to show they have responded in a way auditors can verify.

References

[1] Silvanovich, Natalie / Google Project Zero. "Policy and Disclosure: 2026 Edition." Project Zero Blog, November 13, 2026.

[2] Silvanovich, Natalie / Google Project Zero. "How to fix a bug in a fix." Project Zero Blog, October 6, 2026.

[3] Willis, Tim / Google Project Zero. "Policy and Disclosure: 2021 Edition." Project Zero Blog, April 15, 2021.

[4] Willis, Tim / Google Project Zero. "Policy and Disclosure: 2025 Edition — Reporting Transparency." Project Zero Blog, July 29, 2025.

[5] Google. "Application Security Vulnerability Disclosure Policy." Google, accessed November 2026.

[6] Google Project Zero. "Vulnerability Disclosure FAQ." Project Zero, accessed November 2026.

[7] Mandiant / Google Cloud. "How Low Can You Go? An Analysis of 2023 Time-to-Exploit Trends." Google Cloud Threat Intelligence Blog, October 2024.

[8] The Hacker News. "Google AI 'Big Sleep' Stops Exploitation of Critical SQLite Vulnerability Before Hackers Act." The Hacker News, July 2025.

[9] Cybersecurity and Infrastructure Security Agency. "BOD 26-04: Prioritizing Security Updates Based on Risk." CISA, June 10, 2026. Supersedes "BOD 22-01: Reducing the Significant Risk of Known Exploited Vulnerabilities," November 2021.

[10] Cloud Security Alliance CISO Community. "The 'AI Vulnerability Storm': Building a 'Mythos-ready' Security Program." Cloud Security Alliance, May 2026.

[11] Cloud Security Alliance. "Project Glasswing: AI Discovery Outpaces Open Source Patching Capacity." Cloud Security Alliance, June 2026.

[12] Cloud Security Alliance. "Top Concerns With Vulnerability Data." Cloud Security Alliance, November 2024.

[13] European Parliament and Council. "Regulation (EU) 2024/2847 (Cyber Resilience Act)." Official Journal of the European Union, November 2024.

[14] Cloud Security Alliance. "AI Controls Matrix (AICM) v1.1." Cloud Security Alliance, June 2026.

[15] Hogan Lovells. "EU Cyber Resilience Act: Preparing for Vulnerability and Incident Reporting." Hogan Lovells, June 10, 2026.