GDC 2026 — Meeting Report

Vulnerability Handling in the Agentic AI Era:
Sharing, Triage & Multistakeholder Cooperation

Blockchain Governance Initiative Network (BGIN)
Cyber Security / Agentic AI

2 September 2026 | 12:00–12:50 | Geneva

This record is written under the Chatham House Rule. Information may be used; speakers and other participants are not identified.

Agenda

- Opening: BGIN and partners (including Trust Over IP and Linux Foundation Decentralized Trust); two-way exchange with the digital-identity community.
- Scene-setting: AI-speed alert floods versus human-speed handling; wallets, credentials, ledgers, custody; share fast enough without a central authority and without feeding adversaries.
- Working frame: CVE-era pipelines as building blocks, not the whole layer; an AI-agent sharing platform as the working surface; four conditions (schemas/semantics, provenance, validation evidence, review rules); volume and quality both fail inside one country or one stakeholder.
- Demo: knowledge-graph sharing of threat/harm patterns without a required API hop; the shareable object is the observed pathway, compressed, hashed, and comparable across communities.
- Open discussion: graph interoperability; threat intelligence versus system-attached vulnerabilities; weighting anonymous versus credentialed contributions; open-source coordination and application-security communities; European agencies and CRA; next steps.
- Close: BGIN Block 15 (Washington, D.C., October 2026); hybrid.

Session Description

Agentic AI is increasing vulnerability and threat signals—automated findings, chain-visible issues, and third-party reports—faster than many teams can triage, prioritize, and act. For digital wallets, verifiable credentials, and public ledgers, defenders and adversaries often see the same information, and coordination must span regulated institutions, vendors, and jurisdictions without a central authority.

This 50-minute breakout treated next-era vulnerability handling as a move from alert floods to trustworthy, actionable cooperation. It built on BGIN's Cyber Security Information Sharing Framework (now on an ISO/TC 307 pathway), the Security AI Agent workstream (Community Security Agent), and the need for multi-stakeholder pilots. The room was asked to shape triage practice, a minimal high-value data set for sharing, and cooperation structures (ISAC-style and cross-border), with follow-on work at BGIN Block 15.

Speakers

The session was conducted under the Chatham House Rule. Participants may use the information received, but neither the identity nor the affiliation of the speaker, nor that of any other participant, is recorded here.

On the programme: a moderator from the BGIN Cyber Security community; a working-frame presentation on an agent-keyed sharing ecosystem; and a technical demonstration of knowledge-graph pathway sharing.

Floor discussion included open-source vulnerability coordination, application-security credibility ratings of reports, digital identity and verifiable credentials, and graph-index / traversal design.

Named programme listings are omitted in this record.

Slides

Session slides were presented from a BGIN Beamer deck combining the scene-setting frames, the working frame (ecosystem around AI agent systems; four reliability conditions; global collaboration), discussion questions, and a Block 15 call with session, event, and registration links. Session slides (PDF): https://bgin-global.org/documents/meeting-reports/gdc26/GDC26_Vulnerability_Handling_Agentic_AI_Slides.pdf

Notes & Key Points

Human-speed pipelines versus AI-speed findings. Vulnerability handling, incident response, and threat-intelligence sharing were built for human-in-the-loop disclosure. Automated scanning now generates findings continuously across open-source and proprietary stacks, including blockchain infrastructure. Some issues are exploited as soon as they are visible. Classical CVE / NVD-style processes remain necessary, but they slow down and fit poorly when the objects are mixed, the topology is decentralized, and there is no single patch authority.

The core tension. How to share vulnerability intelligence fast enough to matter—without a central authority, across borders, and without feeding adversaries. A recurring theme was a minimal set of information that can grow a shared understanding when another party holds complementary signal, rather than dumping a full vulnerability record.

Why the identity community is in the room. Wallet and credential stacks will inherit decentralized behavior. Co-organizer interest from Trust Over IP, Linux Foundation Decentralized Trust, and other blockchain groups was noted.

BGIN scaffolding. The Information Sharing Framework adapts cybersecurity sharing practice to blockchain and decentralized ecosystems and is being taken up in ISO/TC 307. Work continues with ISAC-style schemes and institutions in Europe, Japan, and the United States. Tools on top of the framework aim to share knowledge while preserving confidentiality, to separate noise from signal, and to support live multi-stakeholder collaboration.

Four conditions for an agent-keyed ecosystem. (1) Shared schemas and semantics, without which LLM / agent systems cannot interoperate. (2) Provenance, because assisted triage will ingest both good and bad stakeholder input. (3) Validation evidence, so assisted output can be checked. (4) Review rules for high-volume data. Connecting existing building blocks (CVE/NVD-class pipelines, national and European standardization bodies, ISACs) through an agent platform was presented as the collaboration problem, not inventing a second clearinghouse.

Demo: pathway, not only the prompt result. A knowledge graph was shown whose nodes and edges use a schema agents can learn, then communicate by pathing. In the agentic era, what matters is the observed pathway through the knowledge base. A human watching the runtime for a period of time supplies a proof of presence that an agent runtime alone does not. Pathways can be recorded as JSON (commitment and hash); changing the path changes the hash. A compressed visualization (described as a star-tetrahedron / lattice walk, a "floating QR code" of the agent runtime) lets two parties compare keys: has my agent consumed the same patterns as yours?

The graph shown was a real research corpus on agent privacy, not MITRE ATT&CK/D3FEND and not live vulnerability tickets. Nodes can hold compressed documents. The stated next step is to replace that corpus with real threat-intelligence and vulnerability patterns and to accept submissions from participating ecosystems.

How sharing would work across bodies. Each community keeps its own graph. Parties compose artifacts, look for overlaps, fill documentation into the local graph, and return a key for that pattern—offline and browser-based if desired, or via an API or chat. Interoperability lives in shared edge encodings (in the demo: defines, proves, implements, narrates, compresses), plus graph-size statistics, more than in identical node payloads. Detecting overlap was also offered as a way to cross language and vocabulary barriers among communities that do not name the same signal the same way. Selective sharing was described as a "fog of war": each party chooses what to observe and encode before revealing it.

Threat intelligence versus vulnerability. In classical cybersecurity the closer analogue is threat-intelligence feeds (independent signals of what others have seen) rather than vulnerabilities, which are attached to a system, package, or component. The human-observation layer was welcomed as a way to add veracity and expert weight, including for contributors who remain anonymous or pseudonymous.

Weighting and credentials. It was asked how anonymous expert contributions are weighted against certified authorities, and whether verifiable credentials can carry property values. One reply linked this to a trust graph as a knowledge graph plus a promise graph: hashes today, with

an intended path into verifiable credentials and a ceremony when a recipient accepts artifacts. Stakeholder mix, contribution weighting, and correlation of intel across "islands" to raise a maturity / trust score were treated as an open community challenge.

Open-source coordination and application security. From the floor, a Linux Foundation-hosted clearinghouse for open-source vulnerability coordination was described, together with application-security work on FAST-style analysis and community credibility ratings of reports. Collaboration was offered so that that practice and this agent-based backbone can compress the AI-era flood without a single geographic power. Alignment was welcomed: blockchain protocol communities are often skeptical or afraid of classical "threat intelligence" framing, so the work is sometimes presented as sharing knowledge of unusual behavior, while still coordinating with MITRE, ISACs, and other mature processes.

European agencies. It was asked what the work implies from an ENISA, NC, and BSI point of view. The reply was that those bodies have been made aware but are not yet in the room. Crypto-specific ISACs exist in Japan and the United States, with European discussion as well; national agencies are expected to engage once other regions are visibly present. In the meantime European needs, including the CRA, are being taken into account so that designs remain compatible. Specific use cases with stakeholders are being defined in order to implement the solution; a proof or validation path was judged a matter of time rather than a conceptual block.

Determinism of the demo path. Audience questions established that the lattice walk is a bit-flip cycle (a deterministic circuit in the browser), not an API ping of "I am on node X," and that the operator chooses the pathway pattern (for example a defence strategy assigned to a vertex). Documentation for the demonstrated corpus is public; the corpus mixes factual knowledge with a narrative layer used both as a human-readable defence and as a method for managing hallucination pathways.

Session outcomes

- A shared problem statement for the identity and blockchain rooms: AI-speed findings outrun human-speed handling; sharing must work without a central authority and without feeding adversaries.
- A working frame for an agent-keyed ecosystem, with four reliability conditions (shared schemas and semantics; provenance; validation evidence; review rules) and the claim that neither volume nor quality can keep up inside one country or one stakeholder.
- A demonstrated sharing analogue: the shareable object is an observed pathway through a knowledge graph (compressed, hashed, comparable), not a full vulnerability dump or a required API hop. Interoperability was located in shared edge encodings and overlap of keys, including as a way across vocabulary and language barriers.
- Distinctions that the room actually used: threat-intelligence feeds versus system-attached vulnerabilities; proof of presence (human observation of an agent runtime) as a source of

veracity; anonymous or pseudonymous weight versus verifiable credentials / a trust graph.

- Floor alignment that classical open-source vulnerability coordination and application-security credibility ratings can sit beside this backbone, while blockchain protocol communities may need a framing of "knowledge of unusual behavior" rather than classical threat-intelligence language.
- Confirmation that European cyber agencies are aware but not yet in the room, and that CRA-compatible design is a constraint rather than a later add-on.
- A named continuation: BGIN Block 15 (15–16 October 2026, Washington, D.C.; hybrid), including Cyber + IKP: Security AI Agent — Information Sharing & Vulnerability Handling (16 October, 09:50–11:20).

Event: https://bgin-global.org/events/20261015-block15
Registration: https://www.eventbrite.com/e/bgin-block15-tickets-1990274282957
Session: https://bgin-global.org/events/20261015-block15/sessions/2-4