OSCRAT

Open-Source Cyber Resilience Act Tools

OSCRAT Deliverable D3.2 – Product release

Version: 29.07.2026

Author: OVES Enterprise

Co-funded by the European Union

Co-funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Cybersecurity Industrial, Technology and Research Competence Centre. Neither the European Union nor the granting authority can be held responsible for them.
The project funded under Grant Agreement No. 101190180 is supported by the European Cybersecurity Competence Centre.

Project acronym	OSCRAT
Project title	Open-Source Cyber Resilience Act Tools
Project Number Name	01190180
Deliverable Number and Name	D.3.2 Product release
Work package	WP3 – Software design and development
Due Date	31/05/2026
Submission Date	08/09/2026
Lead Partner	OVES Enterprise
Author name(s)	OVES Enterprise
Version	1.0
Status	Final version
Type:	☐ R – Document, Report ☐ DEC – Websites, patent filings, videos, etc. ☒ DEM – Demonstrator, pilot, prototype ☒ Other
Dissemination level:	☒ PU – Public ☐ SEN - Sensitive

Co-funded by the European Union

Co-funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Cybersecurity Industrial, Technology and Research Competence Centre. Neither the European Union nor the granting authority can be held responsible for them. The project funded under Grant Agreement No. 101190180 is supported by the European Cybersecurity Competence Centre.

Document History			
Version	**Date**	**Modified by**	**Comments**
0	31/07/2026	OVES Enterprise	First version
0.1	08/2026	PMF	Internal Review
0.2	08/2026	OVES Enterprise	Partner Feedback
1.0	07/09/2026	OVES Enterprise	Final Version

Abstract

This deliverable presents the final release of OSCRAT, an open-source platform supporting SMEs in meeting Cyber Resilience Act (CRA) requirements. It integrates key compliance, vulnerability, risk, incident, and documentation management capabilities into a single solution, ready for public release.

Keywords

Cyber Resilience Act (CRA); OSCRAT; cybersecurity compliance; open-source platform; Software Bill of Materials (SBOM); vulnerability management; risk assessment; incident management; compliance assessment; SMEs.

Co-funded by the European Union

Co-funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Cybersecurity Industrial, Technology and Research Competence Centre. Neither the European Union nor the granting authority can be held responsible for them. The project funded under Grant Agreement No. 101190180 is supported by the European Cybersecurity Competence Centre.

DISCLAIMER

Co-funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Cybersecurity Industrial, Technology and Research Competence Centre. Neither the European Union nor the granting authority can be held responsible for them. The project funded under Grant Agreement No. 101190180 is supported by the European Cybersecurity Competence Centre.

Co-funded by the European Union

Co-funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Cybersecurity Industrial, Technology and Research Competence Centre. Neither the European Union nor the granting authority can be held responsible for them. The project funded under Grant Agreement No. 101190180 is supported by the European Cybersecurity Competence Centre.

Table of contents

Co-funded by the European Union

Co-funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Cybersecurity Industrial, Technology and Research Competence Centre. Neither the European Union nor the granting authority can be held responsible for them. The project funded under Grant Agreement No. 101190180 is supported by the European Cybersecurity Competence Centre.

1. Summary

This deliverable presents the complete and final release of the open-source OSCRAT platform, designed to support SMEs in managing their obligations under the Cyber Resilience Act.

All planned technical activities have been completed, and the platform is ready to be officially launched to the public.[1]

The OSCRAT platform covers the full set of Cyber Resilience Act activities the project set out to support: checking whether the act applies to a product, role-based compliance assessment, SBOM and vulnerability management, incident reporting, risk assessment, configuration management, security awareness training, technical documentation up to the Declaration of Conformity, and a complete audit trail of user actions. The sections below cover each major feature of the finished platform.

The OSCRAT platform can be accessed here: https://app.oscrat.eu/auth/login

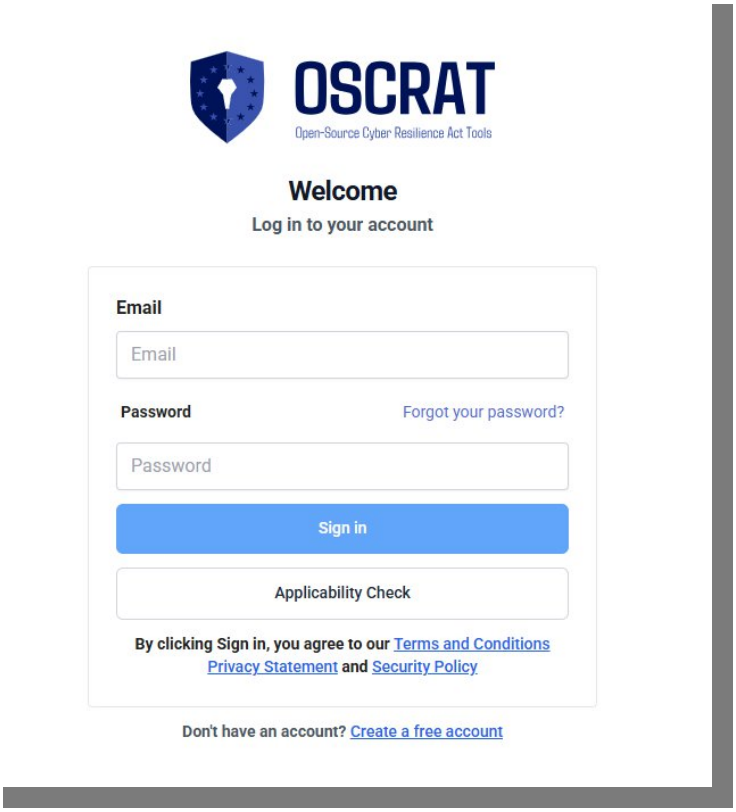

2. Product and Version Management

Users sign up, create or join an organization, and register products. The applicability check, available without an account, walks through the CRA criteria and determines whether a

[1] The OSCRAT platform will be officially launched on 11 September, during the OSCRAT Final event.

Co-funded by the European Union

Co-funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Cybersecurity Industrial, Technology and Research Competence Centre. Neither the European Union nor the granting authority can be held responsible for them. The project funded under Grant Agreement No. 101190180 is supported by the European Cybersecurity Competence Centre.

product falls under the act and in which category: default, important class I or II, or critical. Products carry their classification, conformity procedure, compliance status, and risk level, and split into versions with a draft-to-withdrawn lifecycle plus release and support dates.

The version page is where most day-to-day work happens, with tabs for compliance, configuration, documentation, files, incidents, repository, SBOMs, scans, tasks, version log, and vulnerabilities. Products, versions, vulnerabilities, and incidents each have their own details pages and add/edit flows.

3. Dashboard

The organization dashboard is the landing view: recent activity across the organization, tasks and products with their statuses at a glance, and the state of the applicability check. It gives a quick read on where compliance work stands before drilling into any product; the compliance dashboards cover the assessment results themselves.

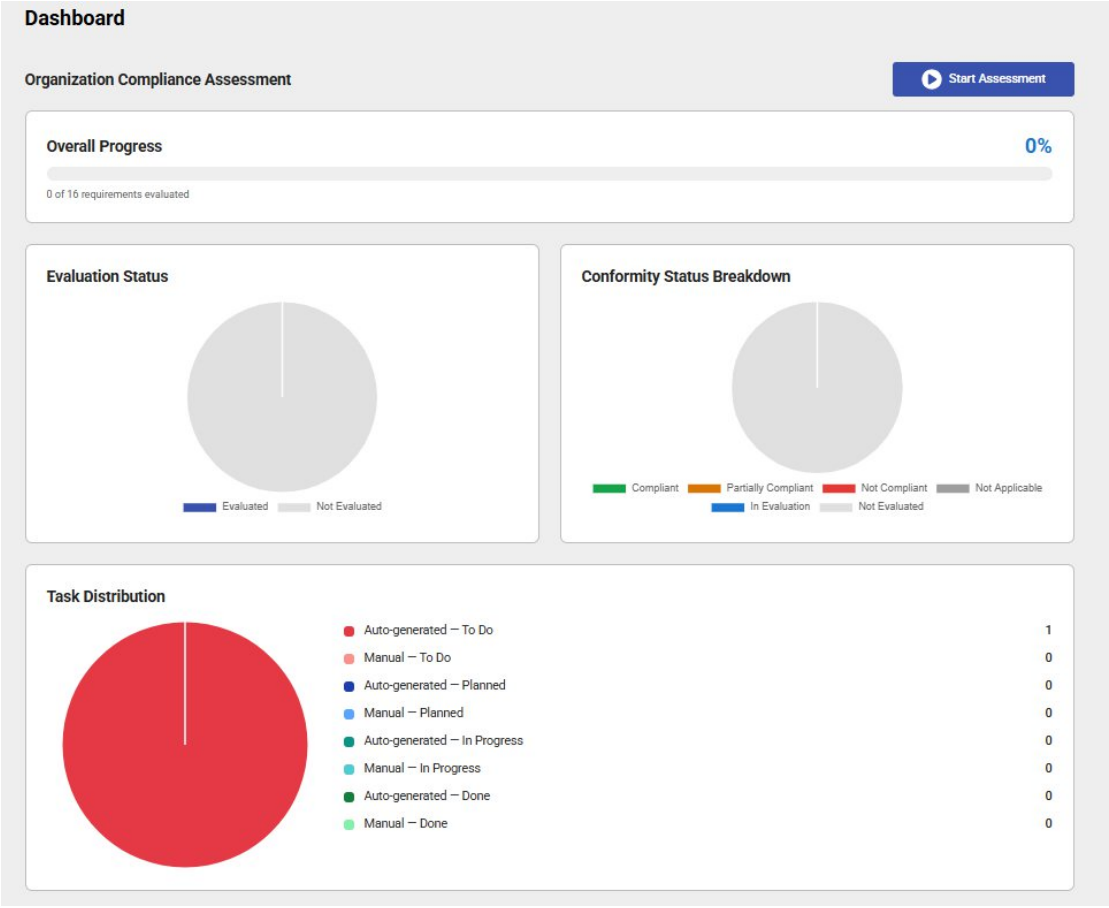

4. Repositories and SBOMs

A version can be linked to a repository on GitHub, GitLab, or Bitbucket, pinned to a branch, tag, or commit. Public repositories work as-is; private ones use a personal access token, stored encrypted. SBOM generation runs as a background job: the job runner clones the repository, generates lockfiles where needed so exact dependency versions are captured, and runs Syft to produce the SBOM. Existing SBOM files can also be imported directly.

Co-funded by the European Union

Co-funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Cybersecurity Industrial, Technology and Research Competence Centre. Neither the European Union nor the granting authority can be held responsible for them. The project funded under Grant Agreement No. 101190180 is supported by the European Cybersecurity Competence Centre.

Reports are kept per version with history, and any report can be sent to vulnerability scanning with one click.

5. Vulnerability Management

Scanning is Grype-based. A scan runs against an SBOM report, whether generated or imported, classifies findings by severity using CVSS scores, and keeps each finding linked to the component that introduced it, so it is always clear where a vulnerability came from. The job runner updates the Grype vulnerability database daily.

Vulnerabilities can also be tracked manually, with CVE reference, severity, advisory information, and affected member states. Each one moves through a handling workflow: pending, preparation, receipt, verification, remediation development, release, and post-release. Remediation plans and other evidence attach directly to the vulnerability record. Advisories and vulnerabilities are kept distinct.

6. Incident Management

Incidents are recorded per product version with the fields CRA reporting expects: classification (confidentiality, integrity, availability, access control, technical failure, theft or loss, and so on), attack type, severity, detection and handling dates, root cause, corrective and preventive actions, suspected unlawful acts, and cross-border impact. Reporting organizations are tracked with their contact emails, and forensic evidence attaches directly to the incident.

Each incident has a title and moves through its own lifecycle from pending through declared, active, and resolved to completed. Incidents surface on the product and version overview pages, so open ones stay visible.

7. Compliance Assessment and Reporting

Assessments exist at two levels: organization-level process assessments and product-version assessments. Both are role-based questionnaires with dedicated content for manufacturers, SME manufacturers, importers, distributors, and authorized representatives, aligned with each role as the CRA defines it. Questionnaire content lives in JSON templates separate from the application logic, with an export mechanism for producing translation templates. Requirements evaluation turns gaps into work: every non-compliant answer proposes a prefilled remediation task that the user accepts with a click.

The Technical Documentation Checklist covers the documentation elements the act requires. Unlike the rest of the questionnaire it is a checklist rather than questions: each entry pairs a requirement with its regulatory reference and a compliant/non-compliant status. It counts toward the compliance result for manufacturers, importers, and distributors, and is optional and excluded from the result for authorized representatives. Dashboards at both the organization and product level chart the results and export to PDF. Assessments can be updated or reset as requirements change, and the content and export pipeline is built for translation, with English shipping at release. The end of the chain is the documentation workflow: a Conformity Assessment Report as the prerequisite, then the Declaration of Conformity for the version.

Co-funded by the European Union

Co-funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Cybersecurity Industrial, Technology and Research Competence Centre. Neither the European Union nor the granting authority can be held responsible for them. The project funded under Grant Agreement No. 101190180 is supported by the European Cybersecurity Competence Centre.

8. Task Management

Tasks are the platform's unit of work. Each one gets a per-organization task number and moves through a to-do, planned, in-progress, done lifecycle, with comments and attachments along the way. Tasks are created manually or automatically: training generation creates them on its own, while requirements evaluation and configuration scans propose prefilled ones. They link to product versions and documentation. ISO-mapped cybersecurity controls can be attached to a task, tying the work back to the control it satisfies. Risk, training, configuration remediation, and requirements follow-up all run through this system, so open work across the platform is visible in one place.

9. Risk Assessment

Risk assessments are a task type. Stage one records the asset, the threat, the affected security categories, and likelihood and impact; the platform calculates exposure from a defined risk matrix. Stage two, locked until stage one is complete, records the treatment decision (accept, reduce, avoid, or transfer), mitigation measures, residual exposure, and the responsible party. Each assessment is linked to a product and follows the normal task lifecycle, with exposure scores driving prioritization.

10. Configuration Management

The platform processes OpenSCAP results rather than scanning itself. Users upload scan artifacts in ARF, XCCDF, or OVAL format; a background job renders them into a read-only report with per-rule results, severities, and pass/fail counts. Every failed rule offers a one-click, prefilled remediation task carrying the rule reference and severity. After fixing, the user re-scans and uploads fresh results; the report updates and the tasks can be closed. Past results are kept for tracking configuration over time, and artifacts are compressed before storage.

11. Security Awareness Training

Training is task-driven as well. When someone joins an organization they get a training task with a due date and links to the material; completion is recorded on their membership, and a daily background job issues the next task when the annual cycle comes around. The program keeps running regardless of how a user was added and without manual administration.

12. Documentation

The documentation module holds versioned markdown documents, scoped to the organization, a product, or a version. Documents move through draft, published, and archived states and are either private or public; public documents are served on a public route. Documents and tasks link to each other, so a document can point at the work it supports and vice versa.

Co-funded by the European Union

Co-funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Cybersecurity Industrial, Technology and Research Competence Centre. Neither the European Union nor the granting authority can be held responsible for them. The project funded under Grant Agreement No. 101190180 is supported by the European Cybersecurity Competence Centre.

13. Files and Attachments

A unified attachment system handles documents across the whole platform: documentation on products, remediation plans on vulnerabilities, forensic evidence on incidents, supporting files on assessments, and the CAR and DoC on versions. Files are validated for type and size, SBOM files get special handling so the scanning pipeline stays intact, access is confined to the owning organization, and large scan artifacts are compressed before storage.

14. Audit Logging

Audit entries are written in the same database transaction as the change they record, so the trail always matches actual state. Coverage spans product-level activity (files, SBOM operations, vulnerabilities, incidents, repositories, configuration scans), task activity, and organization-level activity (membership and product changes), plus report and file downloads. Each entry stores who did what and when, with the previous value next to the new one so edits read as a before-and-after. The log is filterable by user, date range, and action type, and user identifiers resolve to names even after account changes.

15. Organizations and Access

Organizations carry the CRA-relevant data: roles under the act, size classification (micro, small, medium), tax ID, contact and postal details, plus product acronyms and support end dates for reporting. Sign-up captures the organizational details compliance workflows need, and members join by invitation through a dedicated mail service, which replaced the original email subsystem and simplified member management; magic-link login was switched off by default in favor of the standard sign-in flow. Members act as owner, admin, member, or auditor. For enterprise use the platform ships SAML SSO, SCIM directory sync for automatic user provisioning, webhooks for outbound events, and API keys for programmatic access, and login attempts are rate limited.

16. Security Hardening & Pentest

Ahead of release the platform went through an external penetration test of the web app, API, authentication layer, and jobrunner worker. It turned up twelve findings, one critical and three high, with a common thread of tenant isolation and too much trust in user-supplied input, which for a multi-tenant compliance product is the guarantee that matters most. The serious ones ranged from a file-write flaw on the shared worker host to ways a regular user could reach another organization's files, take over an account through SSO email matching, or move data into a team they had no access to.

All ten code-level findings were fixed and verified, including the critical and all three highs; the two left open are deployment configuration rather than code defects. Uploads no longer trust user-supplied filenames, file access is scoped to the owning organization at the data layer, SSO no longer links accounts by email alone, writes go through server-side allow-lists so fields like team and owner can't be overwritten, sessions are revoked on logout, sensitive account changes require re-authentication, and standard browser-level security headers are in place, with the content-security policy rolled out in report-only mode first.

Co-funded by the European Union

Co-funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Cybersecurity Industrial, Technology and Research Competence Centre. Neither the European Union nor the granting authority can be held responsible for them. The project funded under Grant Agreement No. 101190180 is supported by the European Cybersecurity Competence Centre.

17. Quality and UX

Several QA passes and feedback rounds after the beta cleaned up the platform along three lines: data is validated up front and survives editing intact, listings and counts update immediately after a change and follow the same ordering and paging rules everywhere, and user-facing feedback is clear and translated instead of internal codes and stuck loading states. A frontend facelift brought a consistent look across dialogs, cards, and emails; dashboard chart colors are managed centrally in the design system, pages keep their layout regardless of content length, and navigation gained breadcrumbs, side panels, semantic status badges, and useful empty states.

18. Technical Approach

The codebase is a monorepo with three parts: a Next.js frontend that also serves the REST API, a separate job runner, and a shared model package (Prisma schema, database operations, validation schemas, audit logger) used by both. Data lives in PostgreSQL. The job runner polls a job queue table with bounded concurrency and shells out to Syft, Grype, and the OpenSCAP tooling; SBOM generation, vulnerability scans, configuration scan processing, and training renewal all run there, keeping the web application responsive, with retry and error handling for resilience.

Two design decisions did a lot of the work. First, risk, training, configuration, and documentation tracking are all built on a single task model rather than four separate systems, so they share the same lifecycle, audit trail, and interface. Second, correctness is enforced at system boundaries and in shared definitions: request validation is strict and uniform across the API, and chart colors, page sizes, sorting, and cache refresh behavior each live in one shared definition used everywhere, so consistency does not depend on manual upkeep. Beyond that, repository tokens are encrypted at rest, audit writes are transactional, scan artifacts are compressed before storage, and the frontend and job runner deploy as separate services with health checks and automated migrations.

19. Conclusion

The project set out to give organizations a single place to handle the Cyber Resilience Act, and that is what shipped. At the center of it is the SBOM pipeline: link a repository, and the platform generates the bill of materials, scans it on demand, and turns the findings into linked, severity-ranked vulnerabilities — an up-to-date view of what is actually in a product, which the rest of the act depends on. Around that core, an organization can determine whether and how the act applies to its products, assess itself and each product against the requirements of its role, and turn the gaps into tracked work; incidents, risks, configuration, and training live next to the products they belong to, and the documentation ends in a Declaration of Conformity with an audit trail that can stand up to a regulator. What used to require stitching together scanners, spreadsheets, and questionnaires is now one platform, with room to grow as the regulation evolves.

Co-funded by the European Union

Co-funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Cybersecurity Industrial, Technology and Research Competence Centre. Neither the European Union nor the granting authority can be held responsible for them. The project funded under Grant Agreement No. 101190180 is supported by the European Cybersecurity Competence Centre.

OSCRAT

Open-Source Cyber Resilience Act Tools

OSCRAT - Deliverable D3.2 – Product release

www.oscrat.eu https://www.linkedin.com/com

Co-funded by the European Union

Co-funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Cybersecurity Industrial, Technology and Research Competence Centre. Neither the European Union nor the granting authority can be held responsible for them. The project funded under Grant Agreement No. 101190180 is supported by the European Cybersecurity Competence Centre.