

## Question and Answers N.2

**Question 1:** Technical Response – Appendix, Req. 006 / 101: Would ESM accept the platform’s test-code measure, while detailed line-level coverage remains in ESM’s CI/CD pipelines, as meeting the intent of these requirements?

**Answer 1:** It depends. If the Platform itself can provide line coverage, instruction coverage and branch coverage, this may meet the intent of Requirements 006 and 101; the measures do not necessarily have to be supplied by ESM’s CI/CD pipelines. However, the Candidate should explain how the Platform calculates these metrics if it does not execute ESM’s build and tests, including what source code, compiled artefacts, instrumentation and test-execution data it requires; how that data is obtained; the calculation methodology and denominators used; and any dependencies, supported languages or frameworks, and limitations.

**Question 2:** Technical Response – Appendix, Req. 017: The requirement refers to “the metrics listed #4”. Could ESM confirm whether this means the categories in requirement 004, or another list?

**Answer 2:** The reference to “the metrics listed #4” is a cross-reference error. It should be understood as referring to the categories set out in Requirement 004. The relevant requirement will be interpreted accordingly.

**Question 3:** Technical Response – Appendix, Req. 104: Would a suppression stored in a version-controlled configuration file within the repository satisfy the requirement, or is an inline source-code annotation specifically required?

**Answer 3:** No. Requirement 104 specifically requires suppression by annotation in the source code itself, so that the suppression decision travels with the relevant code under version control. A suppression held only in a separate repository configuration file does not meet this particular requirement, although it may be described as an additional suppression mechanism.

**Question 4:** Technical Response – Appendix, Req. 118: Would an expert-led security review integrated into the same findings and triage workflow be considered an “equivalent” to agent-based deep security analysis, or is an automated agent-based capability specifically required?

**Answer 4:** An expert-led review may complement the Platform’s workflow but is not, on its own, equivalent to the capability requested in Requirement 118. The requirement concerns an automated, agent-based deep security-analysis capability integrated with the findings and triage workflow. Candidates should state clearly whether and how their proposed Platform provides that capability.

**Question 5:** Technical Response – Appendix, Req. 126, 137, 156: These requirements refer to “CQ-GOV-001” and “CQ-IAM-008”. Could ESM provide the mapping of these identifiers to the numbered requirements?

**Answer 5:** “CQ-GOV-001” corresponds to the requirement establishing centrally defined quality-policy conditions, and “CQ-IAM-008” corresponds to the identity and access-management requirement concerning periodic access review. The alphanumeric identifiers were retained inadvertently from an earlier drafting structure. For the purpose of the published response template, Candidates should respond against the numbered requirements; the references in Requirements 126, 137 and 156 should be read in that context.

**Question 6:** Technical Response – Appendix, Req. 128: Would Terraform-based declarative management through the platform’s documented REST APIs be acceptable, or is a dedicated Terraform provider required?

**Answer 6:** Yes. Terraform-based declarative management using the Platform’s documented REST APIs is acceptable, provided it enables the configuration and administrative automation required by Requirement 128, is supported and documented, and is included in the proposed scope. A dedicated vendor-maintained Terraform provider is not required unless expressly stated in the requirement.

**Question 7:** Technical Response – Appendix, Req. 138: Would automated user provisioning and deprovisioning through a scheduled synchronisation between Microsoft Entra ID and the platform’s user-management API be accepted as an alternative to SCIM 2.0?

**Answer 7:** An alternative may be accepted only if it provides an equivalent, reliable and auditable automated provisioning and deprovisioning outcome. The Candidate should describe the synchronisation frequency, treatment of joiners, movers and leavers, failure

handling, monitoring, audit trail and any operational dependency on custom components. ESM will assess the proposed approach against the full intent of Requirement 138; a scheduled API synchronisation is not deemed equivalent solely because it automates account updates.

**Question 8:** Technical Response – Appendix, Req. 139: Must the five roles exist as distinct built-in roles, or is an equivalent separation through roles and repository-level permissions acceptable?

**Answer 8:** The five roles do not need to carry the same names as those stated in the requirement or necessarily exist as five predefined product labels. An equivalent role and permission model is acceptable where it provides the required separation of duties and access at the relevant organisational and repository levels. The Candidate should map its proposed roles and permissions to each required role and identify any limitations.

**Question 9:** Technical Response – Appendix, Req. 142: Would an on-demand CSV export from the platform's user management interface, containing users, roles, group memberships and last activity, satisfy the periodic access-review requirement, or is scheduled report generation and/or a dedicated attestation workflow required?

**Answer 9:** An on-demand CSV export may satisfy the reporting element if it contains the information necessary to conduct and evidence a periodic access review, including users, assigned roles, group memberships and relevant activity information. A dedicated attestation workflow or scheduled report generation is not required unless expressly stated. The Candidate should explain how the export supports a complete and auditable review and how inactive or inappropriate access can be identified and remediated.

**Question 10:** Technical Response – Criteria 2 / Req. 143: Does ESM require a native Microsoft Sentinel integration, or would an ingestion component consuming the platform's documented APIs be acceptable?

**Answer 10:** A native Microsoft Sentinel connector is not mandatory. An integration component that uses the Platform's documented and supported APIs or standard export mechanisms is acceptable, provided it enables reliable ingestion of the required security and audit information into Microsoft Sentinel. The Candidate should describe the architecture, data coverage, latency, authentication, maintenance responsibilities and any additional cost.

**Question 11:** Technical Response – Criteria 3 / Req. 151–158: Is it acceptable for fixes to be generated by an AI coding assistant already approved and governed by ESM, using tools and data supplied by the platform, rather than by a model hosted by the bidder? Which AI coding assistants are approved at ESM?

**Answer 11:** No. Requirements 151–158 concern capabilities to be provided by the Candidate’s Platform. If fixes are generated using ESM’s own AI coding assistants, rather than by the offered Platform, the Platform would not fulfil these requirements. The Candidate must therefore demonstrate that its proposed solution itself provides the required AI-assisted remediation capabilities. ESM currently uses Claude Code and OpenAI Codex; however, their availability does not replace or reduce the Candidate’s obligations under these requirements.

**Question 12:** Technical Response – Appendix, Req. 167–170: Could ESM clarify:

- a. what is meant by “without degradation”;
- b. whether “concurrent analyses” means full analyses or pull-request analyses;
- c. how the 10-minute pull-request feedback time is measured and which analysis types should be included.

**Answer 12:** (a) “Without degradation” means that the Platform must continue to meet its committed functionality and response-time expectations at the stated scale and concurrency, without material deterioration attributable to the concurrent workload. (b) “Concurrent analyses” covers analyses initiated by ESM’s normal development activity, including pull-request analyses and relevant branch or pipeline analyses that may be in flight at the same time. (c) The 10-minute period is measured from the point at which the relevant pull-request analysis is triggered and the Platform has access to the required inputs until the analysis result and applicable pull-request feedback are made available. It includes the analysis needed to provide the quality and security feedback applicable to the pull request; external queueing or processing within ESM-controlled systems should be identified separately where it affects measurement.

**Question 13:** Commercial Response (updated document): Item 1: What do licensed lines of code mean?

**Answer 13:** “Licensed lines of code” means the volume of source-code lines counted for licensing purposes under the Candidate’s proposed commercial model. The Candidate should state its counting methodology, including which languages, branches, repositories, generated code, test code and duplicate files are included or excluded, when the count is measured, and how temporary peaks and overages are treated. The submitted price must cover the ESM scope and assumptions stated in the Commercial Response document.

**Question 14:** Commercial Response (updated document): Item 1: How many teams are grouped/formed by the 75 contributing developers? Are there more (other types of) developers?

**Answer 14:** The figure of 75 contributing developers is the estimated total population of developers expected to contribute code within the scope of the Platform; it is not a commitment to a fixed number of formally constituted teams. Team structures may change during the Contract. Candidates should base their price on the stated population and clearly identify any assumptions or limits linked to the number of teams, organisations, projects or repositories. Other users may require read-only, reporting, administrative or governance access as described in the procurement documents.

**Question 15:** Commercial Response (updated document): Item 1: Can you explain the difference between "named Platform Users" and "read-only/reporting users"?

**Answer 15:** “Named Platform Users” are individually identified users who require interactive access to the Platform and may perform functions permitted by their assigned role, such as administration, configuration, analysis, triage or remediation. “Read-only/reporting users” are users whose access is limited to viewing dashboards, findings and reports, or exporting information, without changing Platform configuration, policies, findings or workflow status. Candidates should state how each category is licensed and identify any functional or numerical restrictions.

**Question 16:** Commercial Response: Item 3 (updated document): What exactly does migration entail within the context of this project?

**Answer 16:** Migration concerns the transition from ESM’s incumbent code-quality platform to the proposed Platform. It includes the implementation, configuration and onboarding activities necessary to establish the ESM environment; connect the relevant Azure DevOps repositories and integrations; configure quality policies, rule sets, user groups, roles and

permissions; and, where technically available and agreed, transfer relevant configuration, findings, triage decisions and historical metric data. Migration of the source-code repositories themselves is out of scope, as the repositories will remain in Azure DevOps. Candidates should describe what is included, any dependencies or limitations, and which migration items cannot be transferred.

**Question 17:** Commercial Response: Item 5 (updated document): What is expected in the KPI reporting?

**Answer 17:** KPI reporting is expected to demonstrate performance against the contractual service commitments at agreed intervals. It should cover, as applicable, service availability; incidents and service interruptions; support requests by severity; response and resolution or workaround times against the committed targets; planned maintenance and relevant notifications; recurring issues and corrective actions; and consumption or entitlement information where relevant. The Candidate should describe its standard report, reporting frequency, delivery format, underlying data and any proposed service-review arrangements.

**Question 18:** Would ESM accept a solution in which the Code Quality Tool's AI functionality relies on models selected by ESM and either operated by ESM or accessed through a model service contracted and paid for directly by ESM? Under this arrangement, the Code Quality Tool vendor would not charge ESM for model usage.

**Answer 18:** We expect this to be a full service, without model selection/operation by ESM.

**Question 19:** The glossary refers to support incident severities S1–S4 being defined in Table 3 of section 4, but this table appears to be missing. Please provide the definitions or confirm that Candidates should propose their own severity classification under requirement 219.

**Answer 19:** Candidates should use the following classification:

| Severity               | Typical meaning                                                                                              | Example                                                           |
|------------------------|--------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------|
| <b>S1<br/>Critical</b> | — Major outage or severe business/security impact; no practical workaround. Immediate, coordinated response. | Production service fully down; data loss; active security breach. |

|                    |                                                                                                                        |                                                              |
|--------------------|------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------|
| <b>S2 — High</b>   | Significant degradation or major feature unavailable; many users affected, perhaps with a workaround. Urgent response. | Core workflow failing for a large customer segment.          |
| <b>S3 — Medium</b> | Limited impact; non-core functionality impaired; workaround available. Normal-priority fix.                            | A report is incorrect or a secondary integration is failing. |
| <b>S4 — Low</b>    | Minor defect, cosmetic issue, or improvement request; little to no operational impact.                                 | UI alignment issue; unclear validation message.              |