Zen Dodd
POLICY AND PUBLIC SUBMISSIONS

10 SEP 2026

Submission on the CyberPath Capability Framework Discussion Paper V0.1

CyberPath Program Team, Australian Computer Society

steadytao.com/s/e2f0d150

Executive Summary

I support the proposal to describe professional capability through knowledge, skills, application and task performance. The discussion paper gives appropriate weight to judgement, ethical conduct and the conditions in which work is performed. Its treatment of qualifications as supporting evidence rather than automatic proof of proficiency should remain central to the framework.[1]

My main concern is that several tables and examples could produce different results from the principles stated beside them. The paper distinguishes proficiency from responsibility then includes consulting leadership and strategic influence in general proficiency descriptions. It warns against unnecessary entry barriers then illustrates eligibility through a fixed experience period and a degree or certification. It distinguishes exposure from competence but its highest exposure category assumes that competence has not yet been reached. These ambiguities would affect assessment and recruitment if they become implementation defaults.[2]

CyberPath should retain separate records of demonstrated capability and the scope of responsibility under which it was demonstrated. Technical depth, sound judgement and reliable independent work should be capable of recognition without requiring management authority. Leadership and wider accountability must still be assessed where the work requires them. Knowledge stages, CyberPath proficiency and SFIA responsibility levels should have explicit purposes and should not be converted into one another automatically.

The evidence model also needs revision. Simulation, workplace performance, repetition and validation describe different properties of evidence. A simulation can be independently observed and repeated. A workplace report may reveal little about an individual's contribution. Evidence should be judged against the capability claimed with its origin, authorship, conditions, quality and limitations recorded. This would support credible recognition of workplace practice, open-source contributions, formal learning and other routes without treating any route as sufficient by itself.

Application should describe how knowledge and skills are used during task performance. It should not become a separate score based on the same evidence. Foundational knowledge should likewise be recorded against the knowledge scale rather than treated as practical proficiency before any task performance has been demonstrated.

Role attributes should describe actual working conditions and justified requirements. Candidate interests belong in career exploration rather than capability scoring. Labour-market outlook should remain separate, dated metadata. Personal preferences, career tenure and generic descriptions of the people thought to suit an occupation should not become capability tests.

The pilot should test whether the framework produces consistent judgements across different assessors, backgrounds and working environments. That includes specialist practitioners with limited organisational authority, people working across domains and practitioners whose strongest evidence cannot be published. These cases can test the Capability model without expanding the four-role recognition pilot.[3]

This submission addresses the nine dimensions in Appendix A3. It refers to the Occupations Framework only where shared definitions, role fields or pilot assumptions affect the Capability Framework. Detailed accreditation procedures and learning pathways remain matters for the later Recognition and Pathways consultations.[4]

1. Introduction and scope

I make this submission in my personal capacity. The views are my own and should not be attributed to any organisation, project or other party with which I am associated. I work in cyber threat intelligence and open-source software.

My interest is whether CyberPath can describe work accurately enough to support sound decisions. An employer needs to know what a person can be trusted to do, under what conditions and with what support. A practitioner needs to understand which capability has been recognised, which has not been assessed and what further evidence would resolve a gap.

The paper already contains many of the necessary safeguards. Section 4.1 supports portable skills and transparent mappings. Section 6.3 distinguishes individual capability from organisational conditions. Section 10.1 requires evidence rather than assumptions about suitability. My recommendations seek to apply those safeguards consistently across definitions, tables, examples and eventual implementation.[5]

The relevant scope includes the connection with Occupations V1.0. That document leaves final role-field definitions and data standards to the Capability Framework. References to it here concern that connection. I am not proposing a general redesign of the occupational taxonomy or a new professional licensing regime.[6]

References to the two CyberPath documents use their printed page numbers. Proposed wording and examples in this submission are recommendations for testing. They are not claims that the proposed methods have already been validated.

2. Knowledge

2.1 Preserve the knowledge model without prescribing a single learning sequence

The distinction between technical cyber, adjacent technical and complementary knowledge is useful. So is the distinction between knowing facts, knowing methods and understanding when an approach is appropriate. These categories help explain why familiarity with terminology is insufficient evidence of professional performance.[7]

The definition nevertheless describes knowledge as something a practitioner needs before building skills, applying them or doing the work. Appendix B also recognises tacit and experiential knowledge developed through practice. The final wording should accommodate that relationship without turning the framework into a mandatory sequence of theory, practice and employment.[7]

I suggest:

> *Knowledge is the understanding of facts, principles, methods and context that supports informed practice. It can be developed and refined through study, practical work, feedback and reflection.*

This preserves the requirement to understand the work. A person who reproduces a procedure without understanding its limits has not demonstrated the same knowledge as someone who can explain, adapt and check it. A person should not have to complete a particular course before being allowed to demonstrate knowledge acquired elsewhere.

2.2 Use knowledge stages for a defined purpose

The four knowledge stages should describe understanding of a specified topic. They should not operate as another overall career grade or automatically establish a proficiency level. An individual may explain advanced cryptographic concepts while requiring supervision to change a production authentication system.[7]

Observable indicators should include recognising missing information, identifying the limits of an explanation and knowing when to consult an authoritative source. Appendix B2 already discusses awareness of one's own limitations. This should carry into the main indicators.

The five knowledge attributes should be available where useful rather than requiring a separate test for every attribute of every topic. Assessment should check the understanding relevant to the task. Where immediate recall is necessary for safe performance, the task should state that requirement. Competent use and checking of reference material should be permissible elsewhere.

3. Skills

3.1 Keep the distinction between proficiency and responsibility explicit

Section 4.3 says that SFIA levels describe the scope and responsibility under which a skill is applied while the skill definition describes the capability applied. That distinction should be retained. The unresolved question is how this seven-level structure relates to CyberPath's five proficiency levels and four knowledge stages.[8]

SFIA combines professional skills with generic attributes such as autonomy, influence and complexity. Its professional skills are defined only at the responsibility levels relevant to the work. CyberPath should preserve that meaning when claiming SFIA alignment.[9]

The final framework should specify what each rating measures, what evidence supports it and how it is displayed. A CyberPath proficiency rating should not imply an equivalent SFIA level. A job title, pay grade or reporting position should not establish either rating. Every profile should use only the scales required for its stated decision and should reuse the supporting evidence.

Autonomy and judgement remain relevant to capability. A person who needs help to interpret results cannot be credited with independent performance. A person who independently resolves difficult technical problems should not be denied recognition in that competency merely because another employee holds budget authority or formal sign-off responsibility.

3.2 Make transferability and mapping limits visible

I support the canonical skill catalogue, stable identifiers and distinction between exact, partial and related mappings. The explicit allowance for skills to belong to more than one family is also appropriate.[8]

Each mapping should identify the source version, relevant skill or statement, applicable level and reason for the mapping. A broader or narrower relationship needs a direction. Where there is no equivalent, the catalogue should retain that result rather than manufacture apparent alignment.

A mapping should not confer proficiency, assessment credit or professional recognition by itself. Those conclusions require evidence against the relevant requirements. This distinction would let employers reuse existing framework investments without mistaking a common label for equivalent performance.

Tool-specific experience should remain separately visible. A practitioner may demonstrate sound log-analysis methods while needing induction to an unfamiliar platform. That gap matters operationally but it should not erase the underlying skill or be treated as proof of readiness to administer the unfamiliar product without support.

NICE's current model provides useful implementation precedent. It publishes versioned Task, Knowledge and Skill components in machine-readable formats while keeping relationships explicit.[10]

4. Application

4.1 Define Application as a relationship demonstrated through tasks

Section 5 defines Application as the combined use of knowledge and skills in context. Section 6 then defines a Task as the actual work through which that capability becomes observable. Both concepts are useful but the proposed model does not identify what independent fact would be stored or assessed as Application.[11]

Application should describe how knowledge and skills were used while performing a task. It should be supported by the task evidence and conditions rather than become a second score derived from the same performance. Where the task does not establish an aspect of application, the record should identify the gap rather than infer it from a separate generic rating.

This preserves the distinction between possessing isolated knowledge or skills and using them appropriately. It also prevents the same result being counted twice merely because it appears under two capability dimensions.

4.2 Make evidence properties distinct from evidence strength

Section 5 appropriately supports simulations, scenarios, practical demonstrations, case walkthroughs and oral defence. Detailed assessment methods belong in the Recognition Framework. Section 4.4 already proposes an evidence table that will influence those methods.[11]

Its entries are labelled as evidence levels but knowledge, simulated, workplace, repeated and validated evidence are not mutually exclusive categories. The paper does not expressly prescribe a ranking but the presentation invites that interpretation. A repeated, observed simulation fits several entries at once. A workplace document may be authentic but provide little evidence that the applicant performed the work.

Replace the apparent ladder with a compact evidence record covering:

- the task, knowledge or skill requirement supported and the limits of that support;
- who produced the work, what the practitioner contributed and what assistance or collaboration was involved;
- the environment, complexity, supervision, tools and information available;

- how the work met the relevant performance criteria and any material errors or unresolved issues;
- when the work occurred and whether it shows sufficient consistency and variation for the claim; and
- how the evidence was checked, by whom and what could not be verified.

This is a minimum description of evidence rather than a proposal to assign numerical scores to every property. One record should be reusable across the requirements it genuinely supports.

4.3 Assess different routes against the same requirements

Workplace evidence can show sustained responsibility, consequences and interaction with other people. A controlled assessment can provide clearer observation of individual decisions and conditions that ordinary work has not exposed. Neither should automatically outrank the other.

The framework should expressly accommodate relevant open-source contributions, independent projects, research and practical competition work. The contribution, review history, conditions and outcomes determine what such evidence supports. Repository popularity, a competition medal or a job title does not establish competence across an entire role. The same discipline should apply to degrees and certifications. Use the assessed outcomes they establish without automatic acceptance or unnecessary repeat assessment.

For consequential duties, credible evidence of sustained judgement and performance may be necessary. Where a simulation cannot test an essential feature, that limitation should be recorded and the additional evidence identified. A general preference for workplace evidence should not become a requirement to have previously held the job being sought.

4.4 Preserve a route for confidential evidence

A public portfolio must not become the default route to recognition. The paper's examples include incident records and forensic case files. Some evidence will be unsuitable for external disclosure.[11]

The Capability model should permit evidence references and verification summaries without requiring central storage of the underlying material. Later recognition procedures can specify redacted samples, authorised attestations, controlled review or alternative demonstrations. An attestation should identify what the verifier actually observed. It should not be treated as proof merely because of the verifier's seniority.

These routes must preserve confidentiality and assess the same requirements. A lack of permission to disclose evidence should be recorded as an evidence-access limitation rather than automatically as a capability failure.

4.5 Establish minimum evidence governance

The Capability Framework should state the minimum information needed for a defensible assessment and reserve everything else as optional. Recognition records should identify the framework and requirement versions, evidence provenance, individual contribution, assessor basis, result, limitations and review status.

The Recognition Framework should then define access, retention, correction, review or appeal, revocation and supersession. Sensitive source material should remain with its authorised custodian wherever a verification record can establish the necessary result. A person should be able to correct factual errors without being able to rewrite an assessor's documented judgement.

These requirements matter because the proposed model may refer to incident material, forensic evidence, employment history, clearances, accessibility needs and observed behaviour. The framework should not create a larger personal-information store merely because more information could be collected.

5. Tasks

5.1 Retain outcome-based tasks and define the minimum standard

Sections 6.1 and 6.2 are among the strongest parts of the paper. They specify intended outcomes, context, required conditions and quality criteria. They also state that completing an activity is insufficient evidence of competence. These provisions should carry into the task catalogue.[12]

The next step is to make the criteria usable. A task should identify the minimum acceptable performance, any safety-critical failures and the circumstances requiring the practitioner to stop, seek approval or escalate. Faster completion or polished presentation should not compensate for a material failure to protect evidence, remain within authority or communicate an important uncertainty.

For analytical work, the standard should assess the reasoning supported by the information available at the time. An inconclusive result can be professionally sound where evidence is insufficient. A confident conclusion reached by luck should not receive the same recognition as a well-supported finding. The framework's existing critical-thinking and accountability behaviours provide a basis for this distinction.

A task should therefore allow a reasoned decision not to act or to defer a recommendation pending further evidence where that meets the stated objective and authority conditions.

5.2 Record individual contribution and integrated performance

Team outputs need clear attribution. A report may combine collection, analysis, technical work, editing and approval by different people. The evidence record should identify the practitioner's contribution without assuming that authorship of the final document means every underlying skill was demonstrated.

Assessment should not require every practitioner to perform an entire organisational workflow alone. A person can competently complete an assigned part, identify dependencies and make an effective handover. Independent judgement within that assignment is compatible with collaboration and required approval.

Some requirements must also be tested together. Separate evidence of log analysis, writing and tool use does not necessarily show that a practitioner can combine them in an investigation. Include integrated tasks where the intended capability depends on sequencing, changing evidence, competing explanations or coordination. This addresses the Occupations Framework's concern about reducing deep tradecraft to interchangeable task lists.[3]

5.3 Separate performance from the conditions supplied

Section 6.3 correctly states that the organisation remains responsible for suitable governance, access, systems and oversight. Assessment should record material constraints alongside the result. A missing data source, unavailable approval or unsuitable assessment setup is different from a practitioner failing to use available evidence correctly.[12]

This should not excuse unsafe work. Recognising a limitation, reporting it and staying within authority are themselves relevant behaviours. The record should distinguish an unmet standard from a task that could not be assessed adequately under the conditions provided.

The accessible demonstration principles in Appendix D3 should remain. Alternative interfaces, assistive technology or written responses should be accepted where they preserve the essential outcome. Any required time pressure, physical method or communication mode should be justified by the task.[12]

6. Competency Areas

6.1 Carry cross-domain flexibility through to competency areas

Section 7 defines a competency area as practice within a functional domain. Elsewhere, the paper permits skills in multiple families and describes competency areas supporting roles that span domains. Figure 2 also depicts competency areas spanning functional domains. The final definition should resolve this difference in wording.[13]

I suggest defining a competency area as a structured grouping of knowledge, skills, application and task requirements relevant to one or more functional domains. A capability should have one stable entry that can be linked wherever it is used.

The Occupations Framework's distinction between a national occupation, a CyberPath role description and an organisation's job should be retained consistently in those links. Section 4.3's use of occupation and role interchangeably makes it harder to know what a capability profile is being matched against. This is a terminology correction at the interface rather than a request to reopen the occupational classification.

6.2 Define what must be met without averaging away gaps

The statement that all capability groups must be satisfied needs an explicit rule for applicability. Each competency area should specify its required technical, professional and adaptive capabilities and their relevant level. It should not imply that every capability listed in section 4.2 is required at the same level in every area.[13]

Professional obligations remain essential. Technical ability should not compensate for unsafe information handling or unreliable reporting. Executive influence should not be imported into a narrowly scoped technical competency merely because leadership appears in the general catalogue.

A profile should preserve uneven capability. A strong result in one area should neither establish another area nor conceal an essential gap through an average score. Shared evidence may support several requirements but it should not be counted repeatedly as though it demonstrated independent repetition.

The definition also refers to behaviours and attitudes. Any such requirement should be job-relevant, observable and expressed through conduct. It should not authorise assessors to infer a preferred personality, identity or belief from general impressions.

6.3 Test hybrid profiles without adding a new pilot occupation

Occupations V1.0 excludes defining generalist competency standards from the initial pilot and selects Architect, GRC Analyst, CISO and SOC Analyst roles. That is a reasonable scope constraint.[3]

The Capability model can still be tested against a blended profile such as a SOC analyst who also develops detections and produces threat assessments. Include at least one such case in validation clinics alongside a deep specialist. The test is whether shared capability and its limits can be represented without duplication, forced averaging or inventing a new occupation for each combination. This would not confer recognition outside the pilot's agreed scope.

7. Exposure

7.1 Keep exposure useful after competence has been attained

Exposure is a useful distinction from demonstrated performance. Someone may observe an incident without being able to manage one. A capable practitioner may also have encountered only a limited range of environments.[14]

The proposed Level 4 exposure description includes repeated engagement while stating that the person is not yet fully competent. That limits exposure to a pre-competence stage. It leaves no clear treatment for an experienced practitioner whose exposure remains relevant to the scope of a capability claim.[14]

Exposure should instead describe what has been encountered, in what setting and with what participation. A compact record could identify the activity or environment, observation or hands-on involvement, supervision or responsibility, relevant dates and breadth of encounters. Repetition should be described where it matters. Attendance hours should not become a substitute for this information.

The current four categories could remain as an optional guide to early learning. They should not be an eligibility ladder, require sequential completion or prevent an otherwise capable person from being assessed. Remove the assumption that the highest exposure category excludes competence. Do not publish an exposure level as a general hiring credential where the underlying context is unavailable.

7.2 Distinguish opportunity, learning and demonstrated improvement

Figure 2 links exposure to improved proficiency and page 26 suggests progression to Level 5 through continued exposure. Both should be qualified. Exposure creates opportunities to practise and receive feedback. An improvement in proficiency still needs evidence.[14]

A person repeatedly performing one familiar activity has not necessarily demonstrated capability across materially different conditions. A person should not receive a lower proficiency rating merely because their employer has not experienced a particular incident. Where relevant exposure is missing, the profile should state the gap and what supervised work or controlled exercise could address it.

Record unknown exposure separately from no exposure. Missing records, a career break or confidential work do not establish that an encounter never occurred. Detailed verification of exposure can be developed in the Recognition Framework without treating an exposure claim as demonstrated competence.

8. Proficiency Level

8.1 Apply levels to a stated capability and context

I support section 9's rule that proficiency applies to competency areas rather than an entire practitioner or occupation. That rule should govern summaries, examples and later recognition products.[15]

The general descriptions need to follow it more closely. Level 4 includes leading consulting engagements or managing security service delivery without supervision. Level 5 is introduced through strategic cyber security leadership and its quality standard includes strategic influence. Those may be appropriate requirements for particular competency areas. They should not be universal requirements for high proficiency in every technical area.[15]

A specialist who handles difficult malware analysis, checks conclusions and recognises the limits of available evidence may demonstrate high proficiency within that scope without managing a consultancy. A senior leader may demonstrate high proficiency in governance while holding a lower or unassessed rating in that technical area. Neither assessment diminishes the importance of the other work.

This does not justify removing judgement, communication or accountability from technical assessment. It requires expressing them at the scope of the competency being claimed. Wider leadership should be demonstrated when wider leadership is the requirement.

8.2 Do not treat knowledge-only attainment as proficiency

The current Level 1 says that competency is demonstrated through comprehension rather than practical execution. That conflicts with the framework's definition of competency through task performance. It also duplicates the knowledge-stage model.[15]

Foundational knowledge should be recorded against the relevant knowledge stage. A CyberPath proficiency rating should begin only when relevant task performance has been demonstrated, even if the task is tightly bounded and performed under direct guidance.

If five proficiency levels are retained, an illustrative sequence for pilot testing is:

1. Foundational practice: performs tightly bounded tasks with direct instruction and explains the relevant concepts and limits.
2. Supported practice: performs defined tasks safely with specified guidance, handles routine variation and escalates appropriately.
3. Independent practice: reliably completes defined tasks in familiar conditions and recognises the limits of the assessed scope.
4. Advanced practice: handles complex or unfamiliar conditions, justifies trade-offs and checks the quality of the result with limited guidance.
5. Expert practice: demonstrates sustained depth and sound judgement in highly complex or novel work within a clearly defined area.

These labels require task-specific indicators and calibration before adoption. Expertise should not require novelty for its own sake. Selecting a well-established method can be the expert decision. Required review or sign-off should not automatically mean a person lacks technical independence.

8.3 Remove assumed graduate ratings and a universal entry threshold

The statement that most formal graduates will be at Level 3 should be removed unless supported by assessment results for defined pathways and competencies. It does not follow from a qualification title. Qualifications may establish particular outcomes but the paper's own evidence model rejects automatic full proficiency.[15]

Workplace readiness should be tied to specified duties and available support rather than one minimum level across the workforce. A person at Supported practice may be ready for genuinely supervised work. An experienced practitioner may be ready for some duties immediately and need induction or supervision for others.

I suggest:

> *Readiness for work depends on the requirements of the specified duties and the support available. A practitioner may be ready to undertake some duties independently and others with supervision. Proficiency must be supported by relevant evidence. It is not assigned automatically from a qualification, job title or period of employment.*

8.4 Keep evidence status separate from proficiency

Not assessed, insufficient evidence and currency requiring review should be separate statuses. They should not default to Level 1, which makes a positive claim. A demonstrated failure against a criterion should also be distinguishable from a criterion for which no adequate evidence was available.

Section 4.6 describes current proficiency by skill while section 9 applies proficiency to competency areas. The final framework should state how skill-level observations support a competency-area judgement. It should publish rules for essential requirements, unresolved gaps and shared evidence rather than leave implementers to invent numerical averages.[8]

9. Role Attributes

9.1 Make the eligibility safeguards govern the examples

Section 10.1 and Appendix H7 distinguish capability from genuine eligibility conditions and warn against gatekeeping. The example schema nevertheless includes a fixed three-year incident-response requirement and a bachelor's degree or CISSP without identifying a necessity specific to the role.[16]

The Occupations Framework separately includes general experience, expected experience and qualifications fields. Its job-description guidance presents experience and qualifications as core advertising fields. These fields need the same safeguards as the Capability paper. Otherwise an implementer could reproduce the example requirements while overlooking the caveats.[6]

For each proposed mandatory condition, the record should state its source, why the work requires it, when it must be met and whether an equivalent route or supervised arrangement is available. This should distinguish external legal or contractual conditions from employer preferences and requirements justified by the work itself. Recording an external condition does not establish that it is a capability measure or a CyberPath-endorsed default.

For experience, describe the required performance and circumstances wherever possible. Repeated responsibility for coordinating an incident involving several teams communicates more than an unexplained duration. Where elapsed experience is retained as mandatory, explain why capability evidence alone is insufficient.

For qualifications, identify the specific requirement and acceptable alternatives rather than populating a generic degree-or-certification field. Distinguish requirements needed at appointment from access conditions that can be completed before restricted duties start and development needs that can be addressed during employment.

I recommend a common rule:

> *Capability decisions must be based on sufficient, relevant and verified evidence against the stated requirements. Qualifications, job titles, organisational grade and elapsed experience may contribute evidence but must not substitute for that assessment. Mandatory access or appointment conditions must be recorded separately and justified for the work concerned.*

This concerns the integrity of a CyberPath capability judgement. It does not promise employment, override a genuine restriction or require recognition where evidence is insufficient.

9.2 Describe the work without prescribing a type of person

The paper's commitment to individual needs and accessible assessment should remain. Appendix H2 also provides useful guidance on describing communication, sensory conditions and working routines.[16]

The example tables should be more carefully qualified. Descriptions such as low social demand for threat intelligence or predictable routines for GRC are possible configurations rather than defining occupational properties. State the actual meeting load, interruptions, working hours, communication modes, on-call duties and available adjustments for the job concerned. Do not infer these from its title.

Integrity and ethical conduct remain professional requirements. They should be judged through relevant behaviour rather than an assumed personality profile. Assessment should not require disclosure of a diagnosis or private health history to establish technical capability. Adjustments should preserve the essential performance standard.

There is also a direct field-definition inconsistency. Worker characteristics describes operating conditions in Occupations V1.0 but personal and professional characteristics in Capability Appendix H6. Give these separate names and purposes.[16][6]

9.3 Keep pathway and market information outside capability decisions

Candidate interests can be useful for voluntary career exploration but they are not evidence of capability. Maintain them in the Pathways Framework or a separate user-controlled profile. Exclude them from recognition credentials, capability scores and automated candidate filtering.

Workforce outlook is also not a property of a person's capability. It should be dated and sourced as external labour-market metadata. Changes in employer demand should not change the meaning of an earlier assessment or what counts as competent performance.

This separation reduces the risk that inferred personality, private preferences or changing market conditions become durable labels attached to a practitioner.

9.4 Record operational authority and context directly

Appendix H3 presents a military-civilian comparison with general claims about legal constraints, collateral impact and reactive intelligence. The main Role Attributes section is broader and already recognises defence, public-sector, critical-infrastructure and other settings. The appendix should follow that approach.[16]

Describe the mission, applicable authority, information restrictions, service criticality, consequences of error and escalation arrangements directly. A sector label should not establish permission to act or an acceptable level of harm. Any specialist legal or doctrinal example retained should identify its source and applicable setting.

The framework can recognise mission-aware, proactive defensive capability across settings without suggesting that authority transfers with the skill. It should preserve the distinction between understanding an obligation, making a decision within delegated authority and holding the authority to approve that decision.

10. Future Proofs

10.1 Use a stable core with technology-specific detail where needed

The paper is right to account for AI-assisted work and changing tasks. The common model should describe the assistance used, the human contribution, required checks and accountability. Technology-specific capabilities and failure modes can then be added where relevant.[17]

Occupations V1.0 includes an AI Confidence Threshold field defined as the reliability required before an output can be accepted without human review. It does not explain how that reliability would be established. Replace or supplement this with task-specific acceptance criteria: what is being checked, what evidence supports reliance, what failures matter and when review is required.[6]

These properties should attach to tasks or workflows with role descriptions summarising them. The same practitioner may use different levels of assistance for collection, analysis and final advice. A role-wide automation label cannot describe those differences adequately.

For capability evidence, distinguish producing an acceptable tool-assisted result from demonstrating every underlying skill unaided. The record should identify what the practitioner did, how they checked the output and which capability is being claimed. Additional explanation or a changed scenario may be appropriate to test understanding. A blanket ban on ordinary

tools is unnecessary unless unaided performance is itself a justified requirement. Fallback capability should be assessed where the role requires it.

10.2 Distinguish changing requirements from fading capability

Section 11.2 proposes different refresh expectations for enduring and perishable capabilities. The distinction is useful but changes to the work, possible loss of individual capability through disuse and lack of recent evidence should be recorded separately. They require different responses.[17]

An old assessment should not automatically establish current competence. The passage of time or loss of access to an employer's system does not by itself demonstrate that a skill has disappeared. Use a review trigger followed by proportionate evidence requirements. Recent verified work may meet a refresh need without a new course or examination.

Keep enduring principles and change-sensitive details within technical capabilities visible rather than treating technical knowledge as generally short-lived and professional judgement as permanent. Refresh expectations should identify the affected capability, the relevant change or risk and an accessible way to demonstrate currency. Training hours or payment for renewal should not stand alone as evidence.

10.3 Test assessment quality rather than framework maintenance alone

The proposed guild, change process, version control and validation clinics are useful. The listed health measures focus on requests, turnaround, unresolved gaps and satisfaction. The pilot should also test whether the descriptions support consistent and defensible judgements. The separate CyberPath Evaluation Framework may address some of these matters. The next release should include them or provide an explicit cross-reference.[17]

Use a manageable set of cases to compare independent assessors' decisions, investigate the reasons for disagreement and check performance on additional tasks not used to support the original rating. Include comparable evidence from different learning routes and examine whether the decision changes because of the route rather than the performance.

Test confidential evidence, reasonable adjustments, supported entry-level work, specialist depth and hybrid profiles. Measure the time, cost and documentation needed from candidates and assessors. A model that is accurate only after an impractical amount of paperwork will be difficult to use consistently.

Report methods, limitations and unresolved disagreements with protections against identifying participants through small groups or distinctive case histories. These tests would assess the framework's usefulness. They would not establish that professional recognition alone reduces national incident rates.

10.4 Keep the framework maintainable and proportionate

Publish a minimum required record, optional context fields and completed examples. Define each term once and cross-reference it. Extend the existing stable skill identifiers to linked records and preserve source versions, change history and the meaning of earlier assessments.

The guild proposal allows material changes to the ontology in its charter description but elsewhere says the ontology remains static. Provide an explicit route for structural corrections where evidence shows that the relationships themselves cause a problem while retaining a higher threshold for disruptive changes.[17]

Publish the catalogue in a documented, reusable format with clear terms for CyberPath-authored material and appropriate treatment of third-party content. Practitioners should be able to understand and use the framework without having to purchase a particular course or software product. That does not prevent charging for an assessment service.

Guild participation should include practitioners from different settings and people outside established training and recognition providers. Publish decision criteria and manage relevant conflicts of interest. Detailed accreditation, appeals, evidence retention and development pathways should be addressed in the announced later consultations with this framework supplying consistent definitions and requirements.

10.5 Publish a completed machine-readable example

Occupations V1.0 assigns final field definitions and data standards to the Capability Framework. The discussion paper describes many proposed fields but does not show one complete record connecting an occupation, role, competency area, task, knowledge, skill, application evidence, proficiency result, role context and unresolved gap.[6]

Publish at least one completed example in a documented machine-readable format. It should show stable identifiers, source and framework versions, required and optional fields, relationship direction, evidence references, applicability, assessment status and supersession. The example should round-trip without changing the meaning of missing, unknown, unassessed and inapplicable values.

This is not a request to finalise the Recognition Framework early. It is a test that the Capability ontology can represent the distinctions the paper already relies upon.

11. Recommendations

I recommend that the CyberPath Program Team:

1. recognise learning through practice as well as study and keep knowledge stages distinct from proficiency;
2. define Application as the contextual use of knowledge and skills demonstrated through task performance rather than a separate score;
3. specify the relationship between CyberPath proficiency and SFIA responsibility while retaining transferable skills, source-versioned mappings and explicit mapping limits;
4. replace overlapping evidence levels with evidence properties, recognise different routes against the same requirements and preserve a confidential-evidence route;
5. establish minimum evidence-governance requirements for provenance, contribution, versioning, access, retention, correction, review and supersession;
6. retain task quality criteria, add clear minimum and safety-critical requirements, record individual contribution and assess integrated performance where needed;
7. permit cross-domain competency links, specify applicable capability requirements and preserve partial or unassessed results without averaging away essential gaps;
8. describe exposure throughout a career without implying that it establishes progression or must precede eligibility for assessment;
9. begin proficiency with demonstrated task performance, remove assumed graduate ratings and define readiness for particular duties with stated support;
10. justify mandatory eligibility conditions, keep candidate interests and workforce outlook outside capability decisions and use actual operational conditions rather than occupational or personality assumptions;
11. define tool-assisted performance and acceptance criteria at task level, use proportionate currency reviews and test assessment consistency, practical usefulness and candidate burden; and
12. publish one completed machine-readable example that connects the Capability dimensions, corrects shared definitions with Occupations V1.0 and identifies decisions reserved for Recognition and Pathways.

The most consequential changes concern proficiency, evidence, eligibility and the boundary between Application and Tasks because these determine how a practitioner may be assessed or screened. The remaining recommendations support their consistent implementation. This is an assessment of implementation risk rather than a claim that the pilot has already produced those failures.

12. Closing

CyberPath already contains useful foundations: observable skills, tasks with quality standards, recognition of different learning routes and a clear statement that organisations remain responsible for the conditions they provide.

Those principles should govern the examples and records as well as the explanatory text. The final framework should make clear what someone has demonstrated, what remains uncertain, what support is needed and which restrictions are separate from capability. This would support high standards without making prior access to a particular employer, qualification or job title the only route to meeting them.

Appendix A. Worked example for pilot validation

This example illustrates the proposed relationships. It is a fictional assessment case rather than an assessment of the author or an account of an employer incident. It does not propose an additional recognised occupation.

A1. Work and context

A SOC analyst is asked to assess whether an unfamiliar security alert warrants escalation. The available material includes selected endpoint and authentication records, an asset description and the organisation's escalation rules. One relevant data source is deliberately unavailable. The exercise is authorised and isolated. No production change is required.

The task is to produce a justified assessment and an appropriate next action. Identifying a named attacker is not required. The practitioner has no authority to isolate a production asset or publish a finding externally.

A2. Capability and task requirements

The task draws on event analysis, source evaluation, communication, judgement under uncertainty and responsible information handling. Relevant knowledge includes the limitations of the supplied records and the distinction between an observation and an inference.

Acceptable performance requires the practitioner to establish what the records support, consider plausible alternatives, explain the effect of missing evidence and follow the escalation rules. A justified decision that the evidence is inconclusive can meet the standard. Unsupported attribution, concealment of a material uncertainty or action outside the stated authority would be assessed against the applicable failure criteria.

A3. Evidence and assistance

The evidence comprises the working notes, cited records, final assessment and an explanation of important decisions. Permitted analysis tools and any AI assistance are recorded. The practitioner must be able to explain the material conclusions and the checks applied to generated output.

An assessor introduces a new record that changes the balance of evidence. The question is whether the practitioner revises the assessment appropriately rather than merely defending the first answer. A separate, varied case tests whether the original performance is repeatable.

A4. Result and limits

The record states which task requirements were met and the conditions under which they were demonstrated. It separately records analytical independence, evidence of consistency and the limited authority supplied by the scenario. It does not infer incident-command capability, unrestricted operational readiness or competence across the whole SOC role.

A public project, authorised workplace sample or controlled exercise could contribute evidence for the same requirements provided its contribution and limitations can be verified. Where a requirement remains untested, the record identifies the gap and further evidence needed.

The validation clinic should ask independent assessors to apply the proposed descriptors to this case. Disagreement about the meaning of a level should lead to refinement of the descriptor or evidence requirement. It should not be resolved by referring to the candidate's age, years employed or preferred career route.

References

[1] CyberPath, *Capability Framework Discussion Paper V0.1*, released for public consultation September 2026, sections 3-6, pp. 6-20.
[2] CyberPath, *Capability Framework Discussion Paper V0.1*, section 4.3, pp. 13-14; sections 8-9, pp. 22-26; Appendix H7, pp. 90-92.
[3] CyberPath, *Occupations Framework V1.0*, final release August 2026, sections 3.2 and 3.4, pp. 11 and 13; Appendix E, pp. 58-65.
[4] CyberPath, *Capability Framework Discussion Paper V0.1*, Appendix A3-A4, pp. 38-39.
[5] CyberPath, *Capability Framework Discussion Paper V0.1*, pp. 10-11, 20 and 28-29.
[6] CyberPath, *Occupations Framework V1.0*, sections 4.2-4.5, pp. 25-32.
[7] CyberPath, *Capability Framework Discussion Paper V0.1*, section 3, pp. 6-8; Appendix B, pp. 40-43.
[8] CyberPath, *Capability Framework Discussion Paper V0.1*, sections 4.1-4.6, pp. 10-15; section 9, pp. 24-26.
[9] SFIA Foundation, *How SFIA works: Levels of responsibility and skills*, accessed 10 September 2026.
[10] National Institute of Standards and Technology, *NICE Framework: Current Versions*, accessed 10 September 2026.
[11] CyberPath, *Capability Framework Discussion Paper V0.1*, pp. 14 and 17-20; Appendix A4, p. 39.
[12] CyberPath, *Capability Framework Discussion Paper V0.1*, pp. 12-13 and 18-20; Appendix D, pp. 60-64.
[13] CyberPath, *Capability Framework Discussion Paper V0.1*, pp. 5, 11 and 21; Appendix E, pp. 65-66.
[14] CyberPath, *Capability Framework Discussion Paper V0.1*, pp. 5 and 22-26; Appendix F, pp. 73-75.
[15] CyberPath, *Capability Framework Discussion Paper V0.1*, section 9, pp. 24-26.
[16] CyberPath, *Capability Framework Discussion Paper V0.1*, section 10, pp. 27-29; Appendix H, pp. 79-92.
[17] CyberPath, *Capability Framework Discussion Paper V0.1*, section 11, pp. 30-34; Appendix I, pp. 97-101.