CPM-F study supplement

A study supplement for the CPM-F, organised by the seven CPM-F domains, each section built from the CPM-E guide sections mapped to that domain and covering what the CPM-F holder proves.

Domain 1: Distinguishing product responsibility from project and management responsibility

Product management is responsibility for one product: what it promises, what it charges and what it deliberately declines to do. Domain 1 of the CPM-F concerns the holder s ability to distinguish that responsibility from project management and from a management role. The central distinction is the scope of the decision, not the presence of the word manager in a title. The CPM-E guide, section 1.2, establishes the product boundary explicitly. Responsibility for one product is not responsibility for leading an area or managing a portfolio.

The holder must recognise that a product decision concerns what the product does within its boundary. Allocating budget between product A and product B is a portfolio decision. Deciding what product A does with the budget it receives is a product decision. These decisions can affect the same product, but they do not belong to the same scope of responsibility. This boundary prevents responsibility for a product from being mistaken for authority over an area or over competing products. See the CPM-E manual, section 1.2, The unit of responsibility is one product.

The distinction from project management should be made through the product judgement involved, rather than through delivery alone. In the supplied framework, a deadline and available capacity are constraints on a product decision. They do not, by themselves, establish which alternative should be chosen. A product decision also needs a stated objective, real alternatives, an explicit choice and refusal, and a review criterion. Identifying when work can be delivered is therefore not sufficient to explain why that work should be selected for the product.

The distinction from a broader management role is equally concrete. The product responsibility described here does not include deciding the allocation between products or leading an area. It does include choosing within the product s allocated resources. The holder should keep those boundaries visible when analysing a situation. A question about distributing resources across products should not be treated as though it were simply a choice of scope within one product. Conversely, a choice about what one product does with its resources should not be confused with portfolio allocation.

Section 3.1, The anatomy of a product decision, provides the structure for recognising the work. A relevant decision has six parts: the question, the objective, the alternatives, the constraint, the choice and refusal, and the review criterion. The question must admit an answer. The objective supplies the basis for comparison. Alternatives must be real options. Constraints eliminate options that cannot be pursued. The decision records both what is chosen and what is excluded. The review criterion states what would justify reconsideration.

The business objective gives the product choice its direction. Section 4.3 describes a translation from the language of results into the language of scope. The product work identifies the indicator affected by the objective, identifies product levers that act on that indicator, estimates their effects and compares those effects against capacity. This is a further test of role understanding: the holder should distinguish selecting work because it serves the stated objective from selecting work merely because it has been

requested.

Product responsibility also includes the consequences of a choice after delivery. A delivered feature creates maintenance, support, documentation and interface complexity. Those operating costs belong in the decision. The reason for the choice must also be recorded, so that it can later be distinguished from a decision whose basis has expired. This makes the product judgement more than a statement that something will be delivered. It includes why the choice is justified and what continuing cost it creates.

Responsibility does not mean taking every decision personally. Section 3.6, What to do when the decision is not yours, states that preparing a decision for somebody else is part of the work. The product manager supplies the written question, real alternatives, relevant constraints and information that only the product has. The decision remains with its owner. This is the key authority boundary: contributing product knowledge does not transfer the right to decide.

The guide illustrates that boundary with a security requirement. The product manager did not reject the requirement or implement it without raising its consequences. He supplied the estimated effect on usage and alternative mechanisms. The decision remained with the security area. Decision records support this discipline, but they do not authorise someone to decide in another person s place. [S3, S6]

For preparation, use the CPM-F manual alongside the requested CPM-E reading: sections 1.2, 3.1, 3.2, Who decides what, and 3.6. The extracts supplied here support the boundaries explained above through sections 1.2, 3.1 and 3.6. Keep three questions visible: Is this a decision within one product? What objective and alternatives make it a product choice? Does the product manager decide, or prepare the decision for its owner? The CPM-F preparation guide is an additional internal preparation reference. [S1, S5, S6]

Domain 2: Identify users, customers, stakeholders, personas and target audiences

A user is whoever performs the task in the product. A buyer is whoever signs the contract and answers for the budget. An operator is whoever sustains the product through support, implementation or administration. These are the explicit role definitions in CPM-E section 7.2, Core concepts. They provide a starting point for this CPM-F domain: identify the people involved before treating their requests as interchangeable. The supplied sources do not define persona or target audience. They therefore support the distinctions below, but not a complete definition of every term named in the domain. Read the related material in the CPM-E manual.

The user is identified by the task they perform, not by responsibility for the contract. Section 7.2 states that users usually want fewer steps and less waiting. It also defines context of use as the physical and temporal conditions in which the task happens. These include the environment, device, available time and interruptions. Identifying a user therefore involves more than naming a person or department. The relevant description connects the person to the task and its conditions. The user glossary entry is an internal reference for this term. The source also establishes why the user s request must remain distinct from the buyer s request: performing the task and answering for the budget are different responsibilities.

The buyer is identified by contractual and budget responsibility. Section 7.2 states that buyers usually want control, predictability and proof of compliance. These concerns are not presented as substitutes for the user s concerns. Section 7.1 describes the difficulty directly: the user can want one thing while the contract signer wants another, with a commercial commitment adding a further demand. The distinction is useful when interpreting a request. A requirement expressed by the contract signer should not automatically be attributed to everyone who performs the task. Conversely, a user s preference does not establish what the buyer needs to fulfil contractual or budget responsibilities.

The operator sustains the product. The sources name support, implementation and administration as examples of this role. Identifying the operator makes the consequences of a product choice visible alongside purchase and use. CPM-E describes buyer, user and operator as three parties with legitimate wants. Its arbitration method requires naming what each party wants and what each loses under the alternatives. This is the supported treatment of stakeholders in these excerpts: make the affected parties and their competing interests explicit. The excerpts do not provide a separate, comprehensive definition of stakeholder. They do show why a decision cannot be understood solely through the request of the person who buys or uses the product.

Customer identification must also account for internal and external use. The supplied glossary defines an internal customer as an area of the organisation that uses the product to serve the external customer. Section 14.3 describes a product used both by those who buy it and by the organisation s own operation. Requests from these groups compete for the same capacity. The distinction matters because the gains can differ: speeding up work improves throughput while keeping demand, whereas eliminating work reduces demand permanently. The internal customer glossary entry provides a relevant reference. The sources support identifying this relationship rather than assuming that every customer request comes from outside the organisation.

A customer profile is a description of a group by observed behaviour in relation to a task. Its purpose is decision-making, not illustration. This is the explicit definition in section 7.2 and the glossary. The named study reference, section 7.3, A customer profile used for deciding, should be read in that context. However, its full text was not supplied here. The supported distinction is therefore precise: a customer profile must describe observed behaviour and serve a decision. The supplied excerpts do not establish whether persona is equivalent to customer profile, or how target audience differs from either term. Those definitions should not be inferred from a shared reference to a group. See the customer profile glossary entry.

Arbitration makes role identification useful. The supplied method says to name what buyer, user and operator want, then identify what each loses under the alternatives. It next calls for a solution that localises a requirement where it was requested. The permanent cost of maintaining that localisation must be estimated. The decision must then be recorded and communicated to all three parties. This method supports the study reference to section 7.5, Arbitration between buyer, user and operator, without supplying that section s full text. It also gives a concrete purpose to the distinctions: identifying roles enables an explicit comparison of interests, consequences and continuing costs.

For this domain, the supported foundation is a clear separation of task performance, purchasing responsibility, operational responsibility and internal customer use. The supplied material also establishes the decision-making purpose of a customer profile. It does not fully substantiate the requested coverage of persona and target audience, so those terms require additional source material before their definitions or uses can be stated here. Use the CPM-F manual as the foundation-level study reference, alongside the specified CPM-E sections. The central lesson supported by these excerpts is direct: identify whose task, responsibility, request and consequences are involved before comparing alternatives.

Domain 3: Classifying competitors by customer use

Competitor classification is the assessment of an offering against the use a customer actually makes of a product. In this CPM-F domain, the holder proves that they can classify a competitor as direct, indirect, substitute or potential from a described case. The governing distinction is customer use, not the competitor s origin or the length of its feature catalogue. The CPM-E guide defines direct rivalry as full coverage of actual customer use. It defines indirect rivalry as partial coverage or coverage for another

audience. A substitute meets the same need through another route. These definitions provide the basis for reading the evidence in a case. [S15, S17]

Direct rivalry exists when an offering fully covers the use a customer makes of the product. Full coverage does not mean that both products have identical features. A competitor can lack features and still cover everything the customer concerned uses. Its original audience does not settle the classification either. An offering that began by serving smaller organisations can become a direct rival for another segment when it covers that segment s actual use. The relevant question is therefore specific: does this offering cover the use described for these customers? See the definition of direct rivalry. [S15, S17]

Indirect rivalry exists when an offering covers only part of the customer s use, or serves the same use for another audience. The CPM-E guide states that this offering only threatens if one of those limits falls. A case must therefore be read for both coverage and audience. An offering s earlier position is not enough to preserve an indirect classification after the evidence changes. If it now fully covers the use of the customers under consideration, the basis for calling it indirect no longer holds. The indirect rivalry definition makes that boundary explicit. [S17, S18]

A substitute meets the same need by another route. The alternative does not have to be another product. The guide includes work done by hand and the decision to stay as things are. Its map of undeclared rivals also includes manual processes, spreadsheets, outsourcing and doing nothing. This matters when a case describes how the customer solves the problem rather than naming a competing supplier. The absence of a named product does not remove the alternative from the analysis. The classification must follow the route used to meet the need. See substitute. [S15, S17]

Potential is one of the four classifications in the guide s rivalry grid. The supplied material does not give it a separate glossary definition. It does, however, establish an important limit on its use: a competitor must not be treated as merely potential when the case already demonstrates full coverage of actual customer use. In the guide s clinic case, classifying the offering as potential would have delayed the response until the first customer loss. The accepted classification was direct rivalry. For this domain, the supported lesson is to avoid using a future-threat label to dismiss evidence of a present direct rival. [S15, S18]

The four-faced rivalry grid provides a structured way to examine the case. It records use covered, audience served, route used to solve the problem and relative price. The classification comes from their combination, not from the offering s origin. Read the case against those columns before selecting a category. Identify what the customer uses. Identify whom the competing offering serves. Check whether it meets the need through another route. Retain relative price as part of the comparison, without allowing it to replace the evidence about use. The guide s model demonstrates that use coverage can overturn an initial classification based on origin.

The clinic case shows the distinction in practice. A clinic management system faced an offering that had previously served small clinics and had introduced a plan for mid-sized ones. The competing plan cost much less and lacked three central features. Six mid-sized customers cited it at renewal. All renewed, but none used the three missing features. The guide rejects classification by catalogue and requires checking real usage first. It then classifies the offering as direct rivalry in the mid-sized segment. Neither the missing features nor the absence of customer losses justified calling it indirect or potential.

For further study, use the CPM-E guide, section 5.3, Classifying rivalry by customer use. Read it alongside section 5.4, Cutting segments by behaviour variation. A segment groups people who behave similarly in the decision at hand. A grouping that does not separate behaviour is only a label. This distinction keeps the classification tied to the customers described. In the clinic case, the direct-rivalry

conclusion applies to the mid-sized segment whose actual usage was examined. The guide s verification summary reinforces both principles: behaviour above declaration, and recognition that the aggregate hides differences. [S14, S17, S18]

The same guide connects this reading to section 5.6, Estimating by comparison with analogous offerings, and section 5.7, Sizing total, reachable and target market. These are related concepts, not replacements for rivalry classification. Analogous offerings provide an estimate base when there is no history of one s own. The model aligns comparable offerings by time in market and adjusts for price and commercial reach, producing a range. Market layers distinguish the total market, the portion the current operation can serve and the immediate target, with constraints separating the layers. Together, these sections keep comparisons and market boundaries explicit. For Domain 3, the central proof remains narrower: choose the rivalry category supported by the described customer use, rather than by origin, catalogue or an assumed future loss. [S13, S15, S16]

Domain 4: recognising the customer development stages and their outputs

The stated requirement for Domain 4 is recognition of the four stages of the customer development model and what each stage produces. The supplied numbered sources do not name those stages or define their outputs. They therefore cannot support a complete stage-by-stage account of this domain. The material below explains the related evidence practices that the sources do support. It should not be read as a substitute for the missing model definition.

A product thesis test is a structured statement with four fields: who, what changes, what you expect to observe and by when. The deadline makes the thesis verifiable. When the test is reapplied at that deadline, it produces one of three readings: confirmed, refuted or insufficient parameter. These are results of the thesis test, not names or outputs of the customer development stages. The distinction matters when preparing an account of the domain. See the product thesis glossary entry for the related term. The supplied support is S22.

For the requested reading on section 4.5, Building and testing the product thesis, the available supporting extract is the product thesis test in S22. Its example concerns a thesis that reducing implementation time would release contracting. After implementation time was reduced, the objection expressed in negotiations moved from time to trust. The model classified the result as refuted, rather than insufficient parameter. The useful distinction is between recording an intended change and checking whether the expected consequence actually appeared. S22 supports that distinction, but does not identify a customer development stage.

A value proposition verification method checks the chain of the reason to buy against evidence from real usage. Section 12.5 instructs the reader to write that chain in four links, seek usage evidence for each link, and mark the weakest one. The next step is to check whether that weakness explains behaviour observed at renewal. The resulting decision is between recovering the link and changing the product's position. These are the method's stated activities and decision, as described in S19. The four links must not be confused with the four stages requested for this domain.

The corresponding verification script makes the evidence record explicit. It contains the four links of the chain, the usage evidence supporting each link, and the date on which each was checked. A link without recent evidence is identified as a natural candidate for explaining a fall in renewal. This instrument, described in S21, complements the method in section 12.5. For the requested section 6.3, Formulating and verifying the value proposition, the supplied extracts support verification through S19 and S21, but do not reproduce that section's full treatment. The CPM-E manual is the relevant internal reading

destination.

Declared and observed evidence are distinct sources: what customers say and what customers do. Section 7.7 states the weighting rule as declaration indicates, behaviour decides . It also warns against treating weighting as exclusivity. A feature used by few accounts can still matter to accounts that represent a relevant share of revenue within a segment. Conversely, low use can result from poor implementation rather than an absent need. The source calls for crossing the accounts that report a concern with those that use the feature before choosing an explanation. This account is supported by S24.

The declared against observed matrix provides a related analysis structure. Its two axes are what customers say and what they do. The quadrants where these diverge call for a specific cross-check before a conclusion. S20 describes this instrument, while S24 explains why the divergence can be information rather than error. Together, they support careful interpretation of customer evidence. They do not establish the names, sequence or outputs of the customer development stages.

Use these readings as supporting material, not as an unsupported reconstruction of the model. The supplied extracts establish a verifiable thesis test, a method for checking the value proposition, a dated evidence script, and a way to compare declarations with behaviour. A source that explicitly defines the four customer development stages and their respective outputs is still required to complete the requested Domain 4 explanation. The CPM-F manual is the appropriate internal destination for the foundation-level reading.

Domain 5: Distinguishing product experiments by purpose

Domain 5 concerns the distinction between a minimum viable product, a low-fidelity prototype, an alpha test, a beta test and an A/B test. The requested competency is to distinguish them by purpose. The supplied CPM-E extracts support preparation on minimum scope, release constraints and success criteria. They do not define these five terms or provide the internal-to-external verification passage. This section therefore explains the supported decision principles without presenting those principles as definitions of the five approaches. The CPM-F manual is the relevant internal reference for checking the foundation-level terminology.

Minimum scope is scope selected to answer a question. The supplied material makes this relationship explicit through the scope, question and deadline triangle. The question sits at the centre. Available scope, deadline and capacity form the vertices. The purpose is to cut scope without losing what answers the question. This is the supported connection to CPM-E section 6.5, Minimum scope in service of a question, although the section itself is not included in the extracts. The domain summary also identifies defining minimum scope in service of a question as a value-domain capability. Preparation should therefore focus on the question that must remain answerable, rather than treating a smaller delivery as sufficient evidence of sound scoping. [S29, S30]

The example in the value models shows how this principle changes a scope decision. A reduction in the available time preserved the path for which the product charged and removed the administration panel. The important distinction is the reason for retaining or removing work. The retained path served the question at the centre of the decision. The example supports question-led scope reduction. It does not establish a definition of minimum viable product, nor does it establish when a prototype or a particular test should be used. The minimum scope glossary entry provides an internal route for reviewing that concept alongside the guide.

Release scoping under constraint is a separate decision. CPM-E section 8.5 states that when deadline and capacity do not move, scope is the variable. It orders cuts by the cost of breaking different

commitments. Faults affecting an active base have an immediate and visible cost. Public promises carry a credibility cost. Internal requests are the most negotiable of these categories. This principle adds a commitment constraint to scope decisions. The question a release must answer matters, but so do the commitments affected by removing work. A reduced release should not be described solely through what remains. The withdrawal of previously expected work also requires attention.

Advance notice distinguishes an announced cut from a delivery failure in the supplied release-scoping passage. An item withdrawn with notice is a cut. An expected item that simply does not appear is a delivery failure. For preparation, keep this distinction separate from the names of product experiments. The passage explains how to manage a scope reduction under fixed deadline and capacity. It does not classify that release as a minimum viable product, an alpha test or a beta test. Read section 8.5 in the CPM-E manual for the requested connection between scope, constraints and commitments.

A success criterion specifies the threshold, deadline and consequence before the result is assessed. The supplied practice summary names these elements under Defining the success criterion beforehand. Another extract describes a new module assessed using the same indicators as established product lines. Without an agreed success criterion, discussion of the result became a negotiation of interpretation. These passages support the preparation link to CPM-E section 9.7. They establish why the basis for judging a result must precede that result. They do not provide the definition or operating requirements of an A/B test. The success criterion glossary entry is an additional internal reading route. [S26, S27, S30]

The remaining reading requirement is CPM-E section 10.6, The transition between internal and external verification. That passage is not present in the supplied sources. A source-grounded explanation of the alpha and beta distinction therefore remains incomplete. The same limitation applies to the purpose-specific definitions of low-fidelity prototype and A/B test, and to the precise distinction between minimum scope and minimum viable product. The supported preparation outcome here is narrower: preserve the question when reducing scope, account for commitments when cutting a release, and establish the success criterion before interpreting results. Use the named manual sections to complete the five-way distinction rather than inferring definitions from these adjacent principles. [S25, S26, S29, S30]

Domain 6: User-story clarity and the delivery concepts behind iterative and incremental work

Domain 6 concerns identifying the elements of a user story and the marks of an iterative and incremental approach. The supplied CPM-E references support the delivery concepts used to examine that work: an ordered product queue, a declared ordering criterion, observed capacity, explicit completion conditions and control of work in progress. They do not provide a user-story template or explicit definitions of iterative and incremental development. Those details should therefore not be inferred from these extracts. Use the CPM-F manual as the reference point for the foundation qualification. The discussion here stays within the concepts documented in the supplied CPM-E sources.

A product queue is the ordered list of items accepted as candidates for execution. It has one order and one person responsible for that order. Section 8.2, Core concepts, establishes this definition. The glossary also distinguishes a backlog from any general list of things to do: in this work, it is a single ordered queue with one owner and a published criterion. For preparation, keep the distinction precise. Inclusion in the queue identifies an item as a candidate for execution. Its position must be explained by the ordering rule. Read these concepts alongside the CPM-E manual, particularly section 8.2.

An ordering criterion is a declared rule explaining why one item sits above another. This is the central concept to follow into section 8.3, Ordering the queue with a declared criterion. The source states that, without the rule, order becomes opinion and is contested each cycle. It also defines prioritisation as ordering by a declared criterion, with the displaced item named. A useful reading question is therefore not just which item comes first. Ask what rule explains its position and which item has moved down as a result. The ordering criterion glossary entry is an additional internal reference for this term.

An acceptance criterion is a verifiable condition specific to one item. Its purpose is to resolve ambiguity in that item's wording. This gives a source-supported way to examine whether an item is sufficiently clear before execution, without treating the supplied extracts as a complete account of user-story structure. Section 8.6, Definition of done and acceptance criterion, is the relevant guide section. The verification questions ask whether ambiguous items receive an acceptance criterion before starting. In preparation, distinguish a statement that leaves interpretation open from a condition that can be checked. Consult the acceptance criterion glossary entry for the linked term.

A definition of done is the set of conditions separating delivered from not delivered. These conditions are agreed before starting and apply across items. An acceptance criterion, by contrast, is specific to one item. The distinction matters when deciding whether the wording of an individual item is clear and whether the work meets the agreed conditions for delivery. Neither concept replaces the other. The domain verification summary identifies an ambiguous item entering execution without a criterion as a sign that the practice is missing. Review section 8.6 with that distinction in mind, supported by the definition of done glossary entry.

Capacity is the amount of work a team completes per period, measured by observation rather than intention. Work in progress is the number of items started and not completed at the same time. Completion time is the interval between an item entering execution and becoming available to whoever requested it. These definitions belong together when reading section 8.8, Controlling work in progress. The source locates the effect of excess work in progress in completion time. Its verification summary identifies rising completion time that nobody notices as a warning sign. Starting work and completing work must therefore remain distinct in any reading of the supplied material.

Use the four guide sections as a connected preparation route. Begin with section 8.2 to establish the vocabulary. Continue to section 8.3 to examine how the queue's order is justified. Read section 8.6 to separate item-specific acceptance conditions from the conditions that apply to delivered work generally. Finish with section 8.8 to examine started but unfinished work and its relationship to completion time. The supplied verification questions reinforce this route: look for one queue, one responsible person, a published ordering criterion, an agreed definition of done and acceptance criteria for ambiguous items before execution. These are the documented delivery concepts relevant to this domain. They should not be presented as a substitute for a sourced definition of user-story elements or of iterative and incremental development.

Domain 7: Recognising product life cycle stages through market signals

Domain 7 of the CPM-F concerns recognising the stages of the product life cycle through the market signals that characterise them. The holder proves that a stage judgement rests on observable evidence, not simply on the product s age. This distinction is central to the CPM-E guide: its verification summary identifies justification by age as a sign that stage recognition is missing. The foundation-level focus here is recognition. The wider CPM-E lifecycle domain also addresses decision authority, traceability and withdrawal. Those responsibilities provide context, but they should not replace the specific capability described for this domain: identifying the stage from observable signs. [S39, S40, S41]

The starting point is section 10.3, Identifying the stage by observable sign, in the CPM-E guide. A stage

diagnosis should make the evidence visible. The guide s analysis instrument records observable signs, their current values, their values in the previous period, the resulting stage and the decision criterion that follows. Repeating that diagnosis each period creates a series that can anticipate a change of stage. For preparation, the important distinction is between naming a stage and supporting that name. A label alone does not show which market observations justify it or whether those observations have changed. [S40, S42]

A mature market competing on price provides a concrete example. The CPM-E guide describes several offerings with equivalent propositions, a falling average price, slow growth in the customer base and a rising acquisition cost. These signs form the evidence for the diagnosis. The guide then explains the implication: each additional unit of volume becomes more expensive, and the useful lever moves from acquisition towards retention, the composition of the offering and price discipline. The relevant capability is to recognise that pattern and connect it to the stage, rather than treat each observation as an unrelated problem.

The same example shows why stage recognition matters. The induced error is to read the pressure as temporary noise and respond with more investment in the lever that has become more expensive. A supported interpretation therefore needs more than a statement that growth is difficult. It should identify the combination of equivalent offerings, price pressure, slow base growth and rising acquisition cost. It should also distinguish that observed combination from an unsupported claim that conditions will soon return to normal. The market evidence, rather than the preferred explanation, must support the stage judgement.

Section 10.4, Reading leading indicators, adds a second requirement. The lifecycle chapter calls for noticing that the stage has changed before revenue announces it. Its verification summary identifies reliance on revenue as the only evidence of health as a sign that this practice is missing. Revenue should therefore not be the sole basis for the judgement. Preparation should connect the leading-indicator topic with the guide s explicit warning: the aggregate can hide what matters. The holder should be able to explain why the evidence presented is insufficient when revenue is the only support offered for a claim of product health. [S40, S41]

Section 9.3, Choosing the indicator by the decision it informs, connects measurement to that interpretation. The guide s decision and indicator chain starts with the decision, then identifies the indicator that informs it and the value at which the decision changes. An indicator without that final field leaves the panel. Applied to stage recognition, the preparation question is direct: what decision would change if this signal changed? The guide s indicator sheet makes this relationship explicit by recording the calculation, data origin, informed decision, change value and accountable person. These fields keep the observation connected to its intended use. [S38, S42]

Section 9.4, Distinguishing a vanity indicator from a result indicator, helps assess the evidence selected for that judgement. The guide s vanity filter asks three questions: how much the indicator costs whoever produces it, whether it can vary, and whether it corresponds to a business result. Two negative answers discard it. Its example is the number of people reached by a communication, which failed all three questions despite being presented weekly as the main indicator. The lesson for this domain is not to accept a prominent figure as useful evidence merely because it appears regularly. Its connection to a result and a decision must be examined.

Reading below the aggregate also supports the diagnosis. The guide s segmentation grid places the aggregate above the same measure split by cohort, plan and profile. It includes each segment s weight in the base and in revenue. The decision comes from the asymmetry between those weights. This gives a practical connection between performance measurement and lifecycle recognition: an overall figure is

not the only available view of product conditions. The CPM-E preparation guide provides a route to the broader study context, while sections 9.3, 9.4, 10.3 and 10.4 remain the specific reading points for this domain. [S38, S40]

A disciplined preparation response should therefore state the stage judgement, identify the observable signs supporting it and distinguish those signs from product age alone. It should read changes across periods, examine evidence beyond revenue and explain which decision the selected indicator informs. Where the evidence resembles the mature-market example, it should recognise the combined pressure rather than assume temporary noise. This keeps Domain 7 focused on what the holder proves: recognition supported by market evidence, with the measurement practices needed to make that recognition defensible. [S37, S38, S40, S42]

CPM-E guide, page 5
CPM-E guide, page 20
CPM-E guide, page 84
CPM-E guide, page 11
CPM-E guide, page 14
CPM-E guide, page 16
CPM-E guide, page 41
CPM-E guide, page 80
CPM-E guide, page 89
CPM-E guide, page 109
CPM-E guide, page 41
CPM-E guide, page 41
CPM-E guide, page 27
CPM-E guide, page 33
CPM-E guide, page 72
CPM-E guide, page 108
CPM-E guide, page 27
CPM-E guide, page 95
CPM-E guide, page 80
CPM-E guide, page 73
CPM-E guide, page 85
CPM-E guide, page 72
CPM-E guide, page 44
CPM-E guide, page 51
CPM-E guide, page 79
CPM-E guide, page 93
CPM-E guide, page 73
CPM-E guide, page 113
CPM-E guide, page 49
CPM-E guide, page 109
CPM-E guide, page 49
CPM-E guide, page 55
CPM-E guide, page 111
CPM-E guide, page 55
CPM-E guide, page 90
CPM-E guide, page 74
CPM-E guide, page 113
CPM-E guide, page 69
CPM-E guide, page 63