

Glossary of Technical Practice

61 terms from the TPBOK, in plain language, each tied to its domain.



Glossary of Technical Practice

61 terms from the TPBOK, in plain language, each tied to its domain.

International Insignia of Certified Technical Practitioners (IICTP)

22 October 2026

Glossary of Technical Practice

IICTP Library · Version 0.2. Public draft for comment. 22 October 2026.

© 2026 International Insignia of Certified Technical Practitioners (IICTP).

Licensed under Creative Commons Attribution-NoDerivatives 4.0 International (CC BY-ND 4.0), <https://creativecommons.org/licenses/by-nd/4.0/>. You may copy and share this book unchanged, for any purpose, with credit to IICTP. Modified versions, including translations, need IICTP's written permission.

This book is drawn from *The Standard for Technical Practitioners* (TPBOK™), version 0.2.

TPBOK™, IICTP, ATP® and CTP® are marks of IICTP and are not licensed by the above. IICTP examination items are confidential and are not part of this book.

Status. This is a draft, published to be criticised. Nothing in it is final. IICTP is a new independent credentialing body and is not ISO/IEC 17024 accredited. Products and organisations named are cited as examples or sources; no endorsement is implied.

Contents

ABOUT THIS GLOSSARY	How to use it	5
A to Z		5

ABOUT THIS GLOSSARY

How to use it



61 terms from *The Standard for Technical Practitioners* (TPBOK™), in plain language.

Each entry shows the domain where the term is used most: **D1** Architecture & Systems Design, **D2** Engineering Governance & Agile Delivery, **D3** Applied AI & Emerging Tech Integration, **D4** Infrastructure & Security, **D5** People Leadership in Engineering. The matching Domain Guide explains the term in context.

A to Z

A

Architecture decision record (ADR) [D2](#)

A short, numbered record of one significant decision: its context, the decision, and its consequences. Superseded, not edited.

B

Bounded context [D1](#)

A part of a system inside which every term has one meaning and one model. Boundaries between contexts are explicit.

Bulkhead [D1](#)

Separating resources (threads, connections, instances) so that exhaustion in one area cannot drain another.

Bus factor [D5](#)

The number of people who would have to leave before a part of the system could no longer be maintained.

C

Career ladder [D5](#)

A written description of expectations at each level of an engineering career, used for growth and promotion.

Change fail rate [D2](#)

The share of deployments that cause a failure needing immediate remediation (DORA).

Circuit breaker [D1](#)

A wrapper around calls to a dependency that stops calling it for a while after repeated failures, so callers fail fast instead of waiting.

Cognitive load [D5](#)

The mental effort a task demands; a team's capacity for it is limited.

Concept drift [D3](#)

A change in the relationship between a model's inputs and the outcome it predicts.

Context window [D3](#)

The maximum number of tokens a model can consider at once, including instructions, inputs and its own output.

Cost of delay [D2](#)

The value lost per unit of time while an item waits to be delivered.

Cycle time [D2](#)

Elapsed time from when a work item starts to when it finishes.

D**Data drift** [D3](#)

A change in the distribution of a model's inputs compared with its training data.

Data leakage (ML) [D3](#)

Information in training or evaluation that would not be available when the model is used, making scores falsely high.

Dead-letter queue [D1](#)

A place where messages that repeatedly fail processing are moved, so they stop blocking the queue and can be inspected.

Definition of Done [D2](#)

The shared, checkable quality bar an increment must meet before it counts as done.

Dual track [D5](#)

A career structure in which engineers can progress as individual contributors or as managers, with equal recognition.

E**Embedding** [D3](#)

A vector of numbers representing text so that texts with similar meaning are close together.

Error budget [D4](#)

The amount of failure an SLO allows over its window (e.g. 0.1% of 30 days = 43.2 minutes).

Evaluation set [D3](#)

A curated set of inputs with expected outcomes or grading criteria, used to test an AI feature the way a test suite tests code.

Excessive agency [D3](#)

Giving a model more tools, permissions or autonomy than its task needs.

Expand and contract [D2](#)

Making a breaking change safely: add the new form beside the old, migrate consumers, then remove the old.

F**Faithfulness (groundedness)** [D3](#)

The degree to which every claim in a generated answer is supported by the sources provided.

Feature flag [D2](#)

A runtime switch that turns behaviour on or off without deploying new code.

Fine-tuning [D3](#)

Further training of a pretrained model on examples of desired behaviour.

G**Generative culture** [D5](#)

A performance-oriented culture (Westrum) in which information flows freely, messengers are welcomed and failures lead to inquiry.

H**Hallucination** [D3](#)

Fluent model output that is false or not supported by the provided sources.

Hybrid retrieval [D3](#)

Combining keyword (sparse) and vector (dense) search, then merging the ranked results.

I

Idempotent [D1](#)

An operation that has the same effect whether it is performed once or many times.

Information hiding [D1](#)

Putting each design decision likely to change inside one module, behind an interface that does not reveal it.

Infrastructure as code (IaC) [D4](#)

Defining infrastructure in versioned, reviewed code that a tool applies, instead of manual changes.

L

Least privilege [D4](#)

Giving every person and workload only the access its task needs, for only as long as it needs it.

Liveness / readiness probe [D4](#)

Orchestrator health checks: liveness decides whether to restart a container; readiness decides whether to send it traffic.

M

Model card [D3](#)

A short record of a model's intended use, performance (including by subgroup), limitations and ethical considerations.

O

Outbox (transactional) [D1](#)

Writing an outgoing event to a table in the same database transaction as the state change, then publishing it separately, so the state and the event cannot disagree.

Overfitting [D3](#)

A model fitting its training data so closely that it performs poorly on new data.

P

PACELC [D1](#)

If there is a Partition, choose Availability or Consistency; Else, choose Latency or Consistency.

Precision / recall [D3](#)

Precision: the share of flagged cases that are correct. Recall: the share of real cases that are found.

Prompt injection [D3](#)

Text that reaches a model and tries to override its instructions; *indirect* injection arrives through content the system processes, not from the user.

Psychological safety [D5](#)

A shared belief that a team is safe for interpersonal risk-taking: asking, admitting mistakes, disagreeing.

R

Recovery point objective (RPO) [D4](#)

The maximum acceptable data loss, measured in time, after a failure.

Recovery time objective (RTO) [D4](#)

The maximum acceptable time to restore service after a failure.

Release train [D2](#)

A fixed release cadence where the date is fixed and the scope varies.

Reranking [D3](#)

Re-scoring retrieved candidates with a more precise model before passing the best ones on.

Retrieval-augmented generation (RAG) [D3](#)

Retrieving relevant passages at question time and instructing the model to answer from them.

S

Service level expectation (SLE) [D2](#)

A forecast of how long a work item should take, with a probability (e.g. 85% within 8 days).

Service level indicator (SLI) [D4](#)

A measure of user experience, such as the share of requests that succeed within a time limit.

Service level objective (SLO) [D4](#)

A target for an SLI over a time window, such as 99.9% over 30 days.

Software bill of materials (SBOM) [D4](#)

A machine-readable list of the components in a piece of software.

Sponsorship [D5](#)

Using one's influence to give someone opportunities and visibility, as distinct from mentoring (advice) or coaching (questions).

Strangler fig [D1](#)

Replacing a legacy system by routing one capability at a time to new code until the old system can be removed.

Stream-aligned team [D5](#)

A team that owns a flow of change for a product, service or user journey, end to end (Team Topologies).

STRIDE [D4](#)

Threat categories: Spoofing, Tampering, Repudiation, Information disclosure, Denial of service, Elevation of privilege.

T

Tail latency [D1](#)

The latency of the slowest requests (e.g. the 99th percentile), which users of large fan-out systems feel most.

Technical debt [D2](#)

The extra future effort caused by code or design that no longer fits what is now known.

Token [D3](#)

The unit of text a model reads and writes; limits, latency and price are counted in tokens.

Training-serving skew [D3](#)

A difference between how features are computed in training and in production.

Trunk-based development [D2](#)

All developers integrate into one main branch at least daily, using short-lived branches if any.

Trust boundary [D4](#)

A point where data or control passes between parts of a system with different levels of trust.

W

Work item age [D2](#)

Elapsed time since an unfinished work item started.

Z

Zero trust [D4](#)

A security model in which no request is trusted because of its network location; every access is authenticated and authorised.

Glossary of Technical Practice

61 terms from the TPBOK, from bounded contexts to zero trust, defined in plain language and tied to their domain.

Part of the IICTP Library. Every book is free to download at iictp.org.

- Domain Guide 1: Architecture & Systems Design
- Domain Guide 2: Engineering Governance & Agile Delivery
- Domain Guide 3: Applied AI & Emerging Tech Integration
- Domain Guide 4: Infrastructure & Security
- Domain Guide 5: People Leadership in Engineering
- Illustrative Questions
- Casebook
- Code of Ethics and Professional Conduct