A Capability Framework and Certification System for Industrial Robot Service Professionals

Six Capability Domains	
Domain 1	*Service Lifecycle Management*
Domain 2	*Fault Diagnosis & Technical Support*
Domain 3	*Remote Service & Digital Operations*
Domain 4	*SLA Design & Service Commitment*
Domain 5	*Cross-Platform Service Procedures*
Domain 6	*AI/ML System Service Methods*

Robotics Service Framework

Whitepaper

Version 1.4

Initiated and Published by

RobotToday.com

About This Document

Document Name	RSF Robotics Service Framework — Whitepaper
Version	v1.4 (Certification System Supplement)
Publication Date	2026
Initiating Organization	RSF Initiative, with RobotToday.com as the initiating platform
Updated Public Release	Chapter 1, 2, 3, 4, 5, 6, 7, 8 with Appendices
Target Audience	Robotics service professionals (including field service engineers and service leads)
Language	English (primary)
Supporting Resources	Certification inquiries · ATC partnership · Framework feedback: framework@robottoday.com
Copyright Notice	Copyright belongs to RSF Initiative. No part of this document may be reproduced in full without written authorization. Citations must acknowledge the source (RSF Robotics Service Framework v1.0, RobotToday.com) and obtain written permission.

RSF Positioning **The Robotics Service Framework (RSF)** is the world's first professional capability framework to systematically cover the full service lifecycle of industrial robots. Built upon the compliance foundation of the ISO/IEC standards landscape, RSF establishes dedicated systems for service-engineer competence certification, service-quality commitment, and cross-platform service procedures — addressing capability-layer needs that sit, by design, outside the scope of existing standards.

RSF takes the six capability domains as its content backbone, the four-tier certification system (Professional · Specialist · Expert · Master) as the capability-development pathway, and the ATC Authorized Training Center network as its implementation vehicle, establishing a common professional vocabulary and capability standard for the robotics service industry.

Standards Reference Notice *International standards cited in this document (ISO, IEC, ANSI/RIA series) are referenced for descriptive purposes only. For any technical limits, parameters or specifications related to cited standards, the official published text of the respective standard organization shall be the authoritative source. This document does not replace any international standard and does not constitute a compliance certification document.*

Table of Contents

Table of Contents

Key Terms Glossary (Selected)

Abbreviation / Term	English Full Name	Definition
RSF	*Robotics Service Framework*	Robotics Service Framework — the professional capability framework system defined in this document
6D	*Six Capability Domains*	The six core capability modules of the RSF framework (Domain 1 – Domain 6)
Professional	*RSF Professional*	Entry tier of the RSF four-level certification; masters execution-layer service capability (Know-How)
Specialist	*RSF Specialist*	Second tier of the RSF four-level certification; masters independent handling of complex faults (Know-Why)
Expert	*RSF Expert*	Third tier of the RSF four-level certification; capable of service-system design and cross-domain integration (Know-Across)
Master	*RSF master*	Highest tier of the RSF four-level certification; capable of industry-standard formulation and technology-roadmap decisions (Know-Beyond)
ATC	*Authorized Training Center*	RSF Authorized Training Center — responsible for course delivery and certification-exam execution
SLA	*Service Level Agreement*	Service Level Agreement — contractual document that quantifies service-quality commitments
OEE	*Overall Equipment Effectiveness*	Overall Equipment Effectiveness = Availability × Performance × Quality
MTTR	*Mean Time to Repair*	Mean Time To Repair — average time from work-order opening to equipment return to normal operation
RUL	*Remaining Useful Life*	Remaining Useful Life — predicted remaining service time of a critical component before maintenance threshold
RDT	*Robot Digital Twin*	Robot Digital Twin — physical-digital virtual mirror system synchronized in real time
UCSP	*Unified Cross-Platform Service Procedures*	Unified Cross-Platform Service Procedure framework
PSC	*Platform Service Card*	Platform Service Card — standardized document format for platform-specific service procedures
MLOps	*Machine Learning Operations*	Machine Learning Operations — methodology for full-lifecycle management of AI models
RaaS	*Robots as a Service*	Robot-as-a-Service — subscription deployment model in which the service provider retains equipment ownership
RAG	*Retrieval-Augmented Generation*	Retrieval-Augmented Generation — AI technique combining knowledge-base retrieval with large language models
CPD	*Continuing Professional Development*	Continuing Professional Development — RSF certification's continuous-learning credit system
OTA	*Over-The-Air*	Over-The-Air update — remote deployment of equipment firmware or software via network
LOTO	*Lockout/Tagout*	Lock-Out Tag-Out — operating procedure for safely isolating energy sources during mechanical maintenance
FTFR	*First-Time Fix Rate*	First-Time Fix Rate — share of work orders whose root cause is correctly identified on the first visit
KFL	*Known Fault Library*	Known-Failure-Mode Library — organizational knowledge asset that stores diagnostic experience in structured form
VDA 5050	*VDA 5050 Interface Specification*	Industry-standard communication protocol for AMR (Autonomous Mobile Robot) fleets

PINN	*Physics-Informed Neural Network*	Physics-Informed Neural Network — AI model fusing physical equations with data-driven learning

Full Glossary in Appendix D · Bilingual (English / Chinese)

Chapter 1

RSF Framework — General Overview

The large-scale industrial deployment of robotics is advancing far faster than expected. The installed base of industrial arms, collaborative robots, autonomous mobile robots, and now humanoid robots is growing rapidly, generating an unprecedented scale of demand for maintenance, upgrade, fault handling, and end-of-life management of these systems. Yet the capability system that should support this demand — who performs the service, how their competence is assessed, and how service quality is quantitatively committed — still lacks systematic framework support.

This chapter builds the general overview of RSF from three progressive perspectives. Section 1.1 analyses the existing international-standards landscape to reveal the five key gaps in robotics service and explains why they fall outside the design scope of those standards. Section 1.2 establishes the precise structural relationship between RSF and the existing landscape through a three-layer nesting model that defines what RSF is, what it co-exists with, and what it fills. Section 1.3 defines the scope and usage of RSF, identifying which scenarios this framework applies to and how readers should use each chapter in practice.

1.1 Five Capability-Layer Needs in Robot Service

International standardisation has been at work in the robotics field for decades, and the results are undeniable. ISO 10218-1/-2 (2025) establishes strict technical boundaries for the safety design and integration of industrial robots; IEC 61508 provides the mathematical framework for quantitative argumentation of safety functions; ISO 9001 and ISO 55001 set management baselines for service organisations from the perspectives of quality management and asset-lifecycle management respectively. Together, these standards form the compliance floor of robotics service operations, and they are the starting point and outer boundary of everything RSF contains.

Understanding a standard's value precisely also requires understanding its boundaries precisely. ISO 10218-2, in its clauses on operational management and information for use, explicitly requires maintenance personnel to be "competent persons," but the standard neither defines competence tiers nor prescribes assessment methods, nor offers a certification pathway — the specification of competence requirements is left entirely to the user organisation. This is by design: a safety standard's job is to set the floor; constructing a competence-certification system is outside its scope. ISO 55001 is similar: it requires the organisation to establish a maintenance strategy and track asset status, but it does not prescribe who executes that strategy or how their competence is assessed.

What is, by design, a reasonable empty space has — under today's rapid evolution of robotics technology — become a systemic industry gap. Robot systems have evolved from purely mechanical devices to mechatronic systems, and now to intelligent systems carrying AI perception and decision layers; the resulting service complexity far exceeds the coverage of any single existing framework. The following five key gaps constitute the fundamental rationale for RSF's existence.

Capability Layer 1 — Building a Service-Engineer Competence Certification System

ISO 10218 calls for "competent persons," but no current robot-specific ISO/IEC technical standard

anywhere defines competence tiers, assessment dimensions, or certification pathways for robot-service engineers. The automotive industry has the ASE certification system; aviation maintenance has the FAA/EASA licensing regime; IT service management has certification pathways such as ITIL® as external reference examples — the robotics-service industry has no equivalent. The direct consequences are several: employers cannot objectively assess engineer competence, customers cannot compare service providers' professional credentials, and the market cannot form effective quality-signalling. The gap leaves the robotics-service market in a state of competence invisibility, which in turn suppresses providers' incentives to invest in professional development.

Capability Layer 2 — Building a Cross-Platform Service Lifecycle Framework

Existing safety standards are highly product-typed: ISO 10218 covers industrial arms, ISO/TS 15066 covers cobots, ISO 3691-4 covers mobile robots, and ISO 13482 covers service robots. This product-type-based structure is technically sound for safety design, but creates fundamental difficulties in service management. In practice, robot application scenarios are almost always heterogeneous — a typical modern manufacturing plant may operate six-axis industrial arms, cobots, and AMRs simultaneously, the three platform classes falling under three separate standards regimes, with no existing standard providing a unified cross-platform service lifecycle framework. There is no normative basis for unifying escalation paths, designing consistent maintenance strategy, or managing engineers' authorisation across platforms.

Capability Layer 3 — Building Service Procedures for AI/ML-Integrated Systems

ISO 10218:2025 made important updates in functional safety, but the logical foundation of its safety framework remains deterministic control systems. Modern robots integrate deep-learning perception layers, reinforcement-learning policy layers, and adaptive control layers — their failure modes are probabilistic: model drift, out-of-distribution inputs, and reward-function misalignment cannot be fully captured by traditional FMEA or fault-tree analysis, nor mapped to deterministic fault codes and troubleshooting trees. ISO/IEC 42001 (2023) establishes a governance framework for AI management systems, but does not reach the operational-layer procedures that field engineers need for diagnosing, recording, and judging AI-layer faults. As commercial deployment of AI-enabled robots accelerates, this gap is widening at a noticeable rate.

Capability Layer 4 — Building an SLA Framework for Robotics Service

Across the entire ISO/IEC standards corpus, no document defines response-time tiers, availability-commitment metrics, escalation triggers, or SLA-performance assessment methods for robotics service. IT service management has SLA-management practices reflected in frameworks such as ITIL® and in the certifiable ISO/IEC 20000-1 implementation standard; aviation MRO has MEL (minimum equipment list) and AOG (aircraft-on-ground) response protocols — the robotics-service field has no equivalent normative language. Service-contract negotiations are conducted on each side's non-standard understanding of the terms, which protects neither the customer nor the provider's ability to build a quantifiable quality-commitment system.

Capability Layer 5 — Building Specifications for Remote Service and Service-Data Management

ISO 27001 provides a robust management-system framework for information security, and IEC 62443 covers industrial control-system cybersecurity — neither addresses the operational specifications of robotic remote diagnostics: who is authorised to initiate a remote connection, how diagnostic data is classified for storage and access, how update authority is allocated for digital-twin models, what the service-verification procedure is for OTA firmware push. ISO 23247 provides a reference architecture for digital-twin manufacturing but again does not descend to the service-procedure layer. Remote service has become one of the principal delivery modalities of robotics service, yet its operational boundaries remain in a standards vacuum.

Table 1.1-1 RSF Five Capability Layers and the Division of Responsibilities with Existing Standards

Gap No.	Gap Area	Capability-Layer Content and RSF's Contribution
Layer 1	**Service-Engineer Competence Certification**	ISO 10218 calls for "competent persons," but no current robot-specific ISO/IEC technical standard defines competence tiers, assessment dimensions, or certification pathways for robot-service engineers. Direct consequence: competence invisibility and market-signal failure.
Layer 2	**Cross-Platform Service Lifecycle Management**	Existing standards are organised by product type (industrial arm / cobot / AMR / service robot); no unified cross-platform service lifecycle framework exists. Direct consequence: fragmented procedures in heterogeneous fleets, and an inability to standardise escalation paths.
Layer 3	**AI/ML System Service Procedures**	ISO 10218:2025 is built on a deterministic-systems framework and cannot cover probabilistic failure diagnosis of AI perception and policy layers. ISO/IEC 42001 stops at the governance level and does not reach field operational procedures.
Layer 4	**Robotics Service SLA Framework**	No ISO/IEC standard defines response-time tiers, availability commitments, or SLA-performance metrics for robotics service. Direct consequence: no common language in service-contract negotiation, and no objective basis for accepting quality commitments.
Layer 5	**Remote Service and Data-Management Specifications**	ISO 27001 / IEC 62443 do not cover remote-diagnostic operational procedures, digital-twin access-authority allocation, or OTA service-verification procedures. Direct consequence: remote-service operational boundaries sit in a standards vacuum.

Structural conclusion These five capability-layer needs share a defining characteristic: each sits, by design, outside the scope of the existing standards landscape. Existing standards precisely answer "what is a safe robot system" — and do so with rigour and authority. The complementary set of questions — "who maintains it, how is competence assessed, how is service quality committed" — lies, from the outset, in the service-capability layer above the compliance baseline. This is a sound and intentional division of responsibilities among standards bodies; it is also precisely where RSF makes its contribution. RSF builds that service-capability layer on top of the compliance baseline, forming an inheritance-and-supplement relationship with the existing landscape — not a replacement, but a purposeful continuation.

1.2 Three-Layer Nesting Model — RSF and Existing Standards

Understanding the relationship between RSF and the existing standards landscape requires a precise structural model. Saying simply that RSF "is based on" or "references" existing standards invites two misreadings: first, that RSF is merely a repackaging of those standards; second, that RSF is in competition with them. Neither is accurate. The relationship is a three-layer nesting structure — an outer layer of compliance baseline, a middle layer of RSF framework content, and an inner layer of operational pattern reference. Each layer has explicit boundaries of responsibility, with no intrusion across them, and together they form the complete robotics-service competence system.

Table 1.2-1 RSF Three-Layer Nesting Model

Tier	Name	Source Authority	Core Responsibilities
Outer Layer	**Compliance Baseline**	ISO/IEC international standards (some legally binding)	Sets the floor. Defines the statutory boundaries for robot-system design and field operations; all RSF content is conditional on conformance with this outer layer.
Middle Layer	**RSF Framework**	RSF Initiative — original framework content	Fills the gaps. Defines service-engineer competence tiers, assessment dimensions, certification pathways, and service procedures — answering "who does it, how it is done, and how the quality is assessed."
Inner Layer	**Operational Pattern Reference**	Automotive 4S service practices; ITIL® and other ITSM references for comparison only	Provides reference points. Operational practice validated at large market scale; RSF's own operational design is robotics-native and cross-checked against these patterns for consistency.

1.2.1 Outer Layer · Compliance Baseline

The outer layer consists of ISO/IEC international standards and constitutes the statutory boundary of the whole system. Its defining feature is that these standards are mandatory: ISO 10218-1/-2 has de facto legal force under the EU Machinery Directive (2006/42/EC); IEC 61508's SIL-level assessment is unavoidable in functional-safety certification; ISO 45001 has been incorporated into occupational-health-and-safety regulations in many countries.

The outer layer's job is to set the floor. Any service engineer performing any on-site robot operation must first operate within the outer compliance baseline. All RSF course content and certification requirements are conditional on conformance with the outer-layer standards; RSF does not modify, bypass, or replace any outer-layer standard. The outer layer also has the property of stability — ISO standards are typically revised on a five-year systematic cycle. RSF takes the currently-valid version of each outer-layer standard as its baseline and synchronously revises related training content when a standard is updated. Citations in this document are based on the current valid version of each referenced standard, with the ISO 10218 series in its 2025 edition.

1.2.2 Middle Layer · RSF Framework

The middle layer is RSF's core contribution and the direct response to the five capability-layer needs described in §1.1. The outer compliance baseline answers "what safety requirements should the robot system satisfy"; the middle RSF framework answers a different set of questions: what competence the service engineer should possess, how that competence is tiered and assessed, what procedures the service process should follow, and how service quality is defined and committed.

The middle-layer framework is built from four core components: a competence-tier system (Professional, Specialist, Expert, Master, organised by deepening cognitive dimension — see §2.3); definitions of the six capability domains (Service Lifecycle Management, Fault Diagnosis & Technical Support, Remote Service & Digital Operations, SLA Design & Service Commitment, Cross-Platform Service Procedures, AI/ML System Service Methods — see Chapters 3–8); the RSF Service Maturity Model (RSF-MM, Levels 0–5, covering an organisation's service-capability evolution path; planned for release in v2.0); and a certification-and-authorisation mechanism that links competence certification with the engineer's authorised scope of work — making certification an authorisation, not just a credential.

The four-tier system organises engineer development along a deepening cognitive dimension — Professional (Know-How), Specialist (Know-Why), Expert (Know-Across), and Master (Know-Beyond) — with parallel Engineer and Service Manager tracks at every tier. Professional covers SOP execution and routine maintenance; Specialist handles complex cross-layer faults and initial team management; Expert leads cross-domain integration, complex-customer operations and ATC-system design; Master sets technology roadmaps, contributes to industry standards, and shapes global service strategy. The complete tier-by-tier definitions — including responsibilities for both tracks — are given in Table 9.2-1 (§9.2).

1.2.3 Inner Layer · Operational Pattern Reference

The inner layer draws comparative lessons from automotive 4S service practices and from IT service-management practices reflected in frameworks such as ITIL® and in standards such as ISO/IEC 20000. These references are used only at the level of general operating logic; RSF does not reproduce ITIL content, claim ITIL compliance, or imply endorsement, certification, or licensing by PeopleCert. Placing them in the inner layer reflects an important RSF design judgement: these systems represent operational thinking validated at large market scale; their value lies in practical maturity rather than regulatory authority — they are consulted as evidence of what works, while RSF's own procedures are designed robotics-first.

The automotive 4S system offers direct reference value to RSF along four dimensions: the structural logic of tiered-technician certification linked to authorised work scope (a convention the ASE system established and RSF's four-tier certification also follows); the closed-loop service process "appointment → reception → diagnosis → repair → delivery → follow-up" (paralleled, for robotics, by RSF's eight-step on-site service procedure); CSI/NPS-based customer-satisfaction tracking (a practice RSF likewise uses for quality assessment of the RSF training system and for ATC reviews); and the ABC parts-inventory model (a standard method RSF Domain 1's spare-parts procedure also uses). Each reference point comes with a clearly drawn boundary — AI-layer fault diagnosis, software/firmware service procedures, and multi-brand fleet management cannot be ported from the 4S system and require original RSF design (see Appendix B for the full adoption logic and transformation boundary).

Table 1.2-2 Three-Layer Nesting — Division of Responsibilities and Inter-Layer Logic

Tier	Core Question Answered	Source of Authority	RSF's Relationship
Outer Layer — Compliance Baseline	What safety requirements should the robot system satisfy?	Mandatory ISO/IEC standards	RSF takes this as a precondition; it conforms but does not modify
Middle Layer — RSF Framework	Who maintains the system? How is competence assessed? How is service quality committed?	RSF Initiative (original framework content)	RSF's core contribution, filling the five gaps that sit above the outer layer
Inner Layer — Operational Pattern Reference	How should service operations be organised and executed?	Automotive 4S service practices / ITIL® and other ITSM references (external reference only)	RSF's operational reference point; RSF's design is robotics-native and cross-checked against it

The overall logic of the three-layer structure is: the outer layer sets the boundary, the middle layer fills the gaps, and the inner layer provides patterns. Together they constitute a complete robotics-service competence system; each layer has its own source of authority and scope of responsibility; the inter-layer relationships are clear, with no overlap and no conflict.

1.3 Scope and Usage of the RSF Framework

RSF is a competence framework and operational reference system aimed at robot-service engineers and service managers. Making the scope and usage of RSF explicit helps readers from different backgrounds use this document in the most effective way and avoids over-interpretation or mis-application.

1.3.1 Robot Platform Types in Scope

The RSF framework applies to service activities for the following types of robot system:

Robot Type	Typical Representatives	RSF Framework Notes
Industrial Arm **Industrial Robot Arm**	Six- or seven-axis industrial robots (Fanuc, ABB, KUKA, Yaskawa)	All six RSF capability domains apply fully; Domains 1–5 form the main framework layer, and Domain 6 applies selectively depending on the degree of AI integration
Cobot **Collaborative Robot (Cobot)**	UR series, FANUC, KUKA Yaskawa、Doosan Robotics、Franka	All RSF framework content applies; force-and-torque verification per ISO/TS 15066 has dedicated procedures in Domains 1 and 5
Autonomous Mobile Robot (AMR)	MiR, Locus Robotics, Amazon/KIVA, Addverb, Geek+, Standard Robots	RSF framework applies; SLAM navigation and path-planning layers receive dedicated treatment in Domains 2 and 6
Service Robot	Commercial service robots, delivery robots, inspection robots	RSF framework applies; ISO 31101:2023 provides the scenario risk-assessment basis for Domain 4 (SLA Design)
Humanoid Robot **Humanoid Robot**	Boston Dynamics Atlas, Unitree H1/G1, UBTECH Walker	RSF framework applies; multi-modal perception and embodied-AI components receive a dedicated service methodology in Domain 6, updated as the technology evolves

1.3.2 Applicable Service Scenarios

RSF covers service activities across the full robot lifecycle from delivery to decommissioning, including:

- **On-Site Service**

  Service activities performed by engineers at the customer site — installation and commissioning, preventive maintenance, fault response and repair, end-of-life equipment disposition. This is the most central application scenario of the RSF framework, with direct coverage across all six capability domains.

- **Remote Service**

  Diagnostic, monitoring, and service-delivery activities conducted via remote connection, digital twins, and predictive-maintenance platforms. Covered specifically by Domain 3 (Remote Service & Digital Operations).

- **Service Management**

  Internal service-strategy formulation, SLA design and performance management, team-competence assessment, training-system construction, and service-quality improvement. Covered specifically by Domain 4 (SLA Design) and the Certification System chapter.

- **Service-Competence Assessment**

  Assessment of an individual service engineer's or an entire service organisation's competence level — used for hiring, training planning, certification, or service-provider qualification verification. The four-tier certification system is designed specifically for this scenario, with the RSF-MM maturity model (planned for release in v2.0) to extend it to organisation-level assessment.

> **Scope boundary (out-of-scope scenarios) RSF does not apply to:** (1) the R&D and manufacturing stages of a robot system; (2) mechanical design or electrical schematic design of a robot; (3) functional-safety certification (SIL/PL assessment must be performed by qualified functional-safety engineers under IEC 61508 / ISO 13849, outside RSF's scope); (4) industrial-automation system integration (the principal business of system integrators, which partially intersects with — but does not overlap — the RSF service-competence framework). For competence frameworks applicable to the above scenarios, RSF recommends consulting the relevant domain-specific standards or certification systems.

1.3.3 Primary Reader Groups and Usage Modes

The principal reader groups for this framework are four, each with a different usage emphasis:

Reader Group	Typical Role / Background	Recommended Usage Emphasis
Robot Service Engineer	Field service engineer, maintenance engineer, service manager	Use Chapters 3–8 (the six domains) as a daily technical reference; use the certification system in Chapter 2 as a personal development pathway; use the appendix's standards-mapping matrix as a compliance-boundary check tool
Service-Organisation Manager	Director of Service, Regional Service Manager, Director of Operations	Use §1.2 as the basis for explaining the service-system positioning to customers and partners; use Domain 4 (Chapter 6) as a framework reference for SLA-system design; use Chapter 9's certification system as a tool for team-competence assessment and talent planning
Customer / End User	Plant operations manager, equipment-procurement lead, maintenance-contract negotiator	Use Chapter 1 to understand the framework as a whole; use the six-capability-domain definitions in Chapter 2 as the reference standard for service-provider qualification; use Domain 4 (SLA) chapter as a reference for contract-negotiation metrics
Training and Certification Institutions	ATC-authorised training centres, university engineering schools, vocational training institutions	Use the full document as the content basis for course development; use Chapter 9 (the certification system) as the alignment reference for course design and assessment methods; use the Appendix glossary as the unified basis for teaching terminology

1.3.4 Document Versioning and Update Mechanism

This document is RSF Framework v1.0, released by the RSF Initiative with RobotToday.com as the initiating platform. The continuing validity of the RSF framework depends on a regular content-update mechanism, driven by three categories of trigger:

- **ISO/IEC standards updates: when a relevant outer-layer standard (especially the ISO 10218 series) publishes a new version, RSF will complete the corresponding chapter synchronisation within a reasonable period after the standard takes effect.**
- **Robotics technology evolution: AI/ML service methodology (Domain 6) and digital-twin / remote-service procedures (Domain 3) are fast-evolving areas; content review and updates are planned at a defined cadence.**
- **Certification-data feedback: based on actual RSF examination data and curriculum feedback from authorised ATCs, the framework content will undergo routine detail revisions; revisions without structural change are released under v1.x numbering.**

> **Chapter Summary** RSF's existence rests on five evidence-backed industry gaps (§1.1) that are not deficiencies of standards quality but competence-layer needs sitting outside the design scope of existing standards. RSF establishes a precise collaborative relationship with the existing landscape through the three-layer nesting structure (§1.2): building the middle-layer framework content on top of the outer compliance baseline, with the inner operational-pattern layer as practical reference. RSF applies to on-site service, remote service, service management, and competence-assessment scenarios for industrial arms, cobots, AMRs, service robots, and humanoid robots; it does not extend to R&D and manufacturing or to functional-safety certification (§1.3). Chapter 2 will establish the structural panorama of the six capability domains, providing the navigational basis for the deep treatments of each domain in Chapters 3–8.

Chapter 2

Six Capability Domains — Overview and Relationship Map

RSF organises the core competencies of a robot-service engineer into six interrelated Capability Domains, which constitute the main body of the RSF framework. These six domains are not a simple knowledge taxonomy — they are a systematic mapping of how robotics service actually operates: they cover the complete lifecycle from delivery to decommissioning, span both the technical-execution and service-management dimensions, and take the robot system's four-layer technical architecture (physical — electrical — software — intelligent) as their internal logical axis.

This chapter, in turn, explains the definitions and boundaries of the six domains (§2.1), the structural relationships among them together with their mapping to typical service scenarios (§2.2), and the capability-growth framework of the RSF four-tier certification system (§2.3). Together these three sections form RSF's structural panorama and provide the navigation basis for the deep treatments of each domain in Chapters 3–8.

2.1 Domain Definitions and Boundaries

The six RSF capability domains are structurally divided into two tiers. The Service-Execution Tier (Domains 1–3) concerns the engineer's technical capabilities acting directly on robot systems in field and remote environments. The Management-Architecture Tier (Domains 4–6) concerns service-quality commitment, multi-platform service procedures, and AI/ML service methodology — the systematic expression and continuous-optimisation mechanism for service competence. The two tiers are not sequential but cooperative: without the execution tier, management-tier design lacks an implementation foundation; without the management tier, execution-tier action lacks systematic direction.

Table 2.1-1 RSF Six Capability Domains — Overview

Capability Domain	English Name	Core Definition	Tier
Domain 1 Service Lifecycle Management	Service Lifecycle Management	The capability to plan, execute, record, and continuously improve service activities at every phase of a robot system's full lifecycle, from delivery and installation through operation and maintenance to decommissioning.	**Service-Execution Tier**
Domain 2 Fault Diagnosis & Technical Support	Fault Diagnosis & Technical Support	The capability to systematically identify and localise the root causes of faults in a robot system's mechanical, electrical, software, and AI layers through a stratified diagnostic method, and to deliver effective technical solutions.	**Service-Execution Tier**
Domain 3 Remote Service & Digital Operations	Remote Service & Digital Operations	The capability to perform diagnostics, monitoring, and service delivery on robot systems off-site through remote connection, digital twins, predictive maintenance, and OTA updates.	**Service-Execution Tier**
Domain 4 SLA Design & Service Commitment	SLA Design & Service Commitment	The capability to design, negotiate, execute, and track robotic service-level agreements, covering response time, availability commitments, escalation processes, performance-indicator definitions, and contractual	**Management-Architecture Tier**

		expression.	
Domain 5 Cross-Platform Service Procedures	Cross-Platform Service Procedures	The capability to formulate and apply unified service procedures across multiple robot platforms (industrial arms, cobots, AMRs, humanoid robots), achieving consistency and portability of service delivery while maintaining safety and regulatory compliance.	**Management-Architecture Tier**
Domain 6 AI/ML System Service Methods	AI/ML System Service Methods	The capability to apply dedicated diagnostic logic, model-performance evaluation methods, and data-driven maintenance procedures to robot systems incorporating AI perception, machine-learning decisioning, or reinforcement-learning policies.	**Management-Architecture Tier**

> **Boundary note: the partition of the six capability domains follows the functional logic of service work rather than disciplinary boundaries of knowledge. There are areas of content overlap among domains — for example, fault diagnosis (Domain 2) involves judging the service-lifecycle phase state (Domain 1), and AI/ML service methods (Domain 6) share data-acquisition infrastructure with remote monitoring (Domain 3). The framework's partition principle maximises practical utility: domains are assigned by the engineer's primary work task, not by a search for logical mutual exclusivity.**

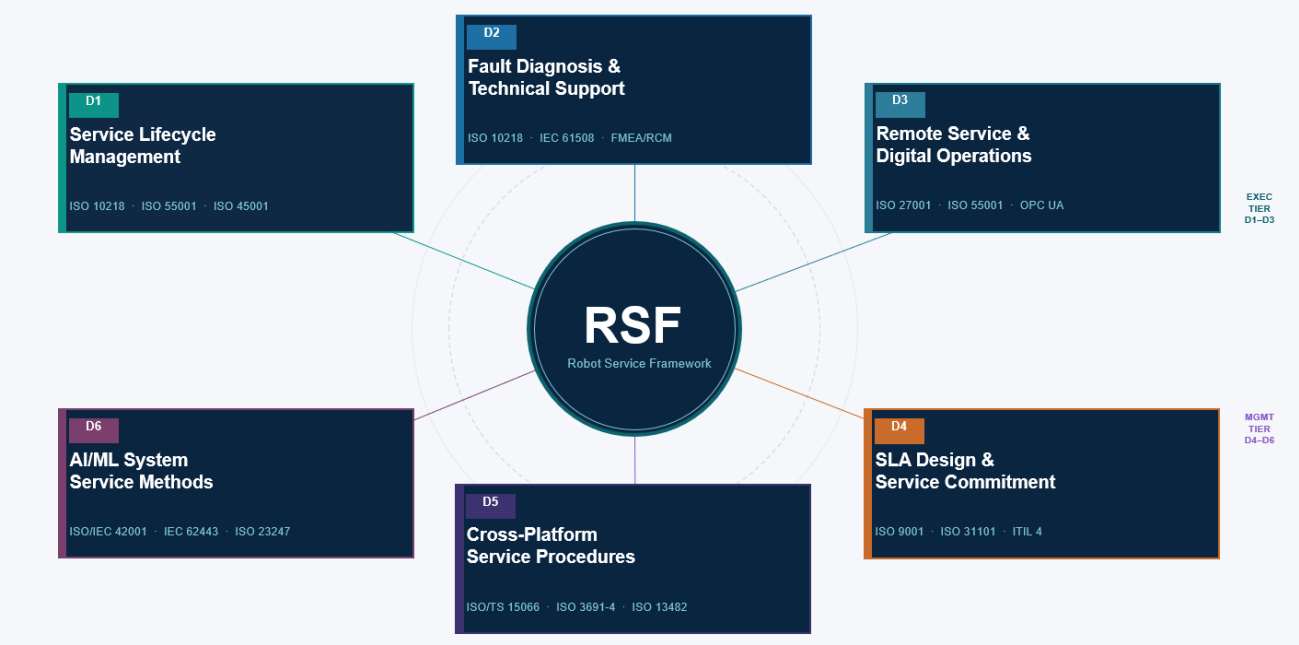

2.1.1 Service Execution Tier: Domains 1–3

The three capability domains in the service-execution tier correspond to the most frequent and most direct work content of the robot-service engineer.

Domain 1 Service Lifecycle Management

Service demand on a robot system is not evenly distributed across its useful life — it shows structural differences across phases: delivery-phase installation and commissioning, operating-phase preventive maintenance and incident response, decommissioning-phase asset disposition. The service objectives, risk priorities, and documentation requirements of each phase differ in nature. Domain 1 requires the

engineer to have a systematic understanding of these phase differences, to take the right service action at the right phase, and to maintain complete records spanning the full cycle.

The reference standards for Domain 1 include ISO 55001 (asset lifecycle management), ISO 10218 (on-site operations safety), and ISO 45001 (occupational health and safety). On top of these standards' compliance framework, RSF establishes dedicated service-phase models, key-document specifications, and milestone-assessment mechanisms for the robotics service context.

Domain 2 Fault Diagnosis & Technical Support

A modern robot system integrates the mechanical, electrical, software, and AI layers, and an anomaly in any one layer can trigger cross-layer fault propagation. Traditional single-layer diagnostic experience — analysing problems only from the angle of mechanical wear or electrical short — fails systemically when faced with AI-perception failure or reinforcement-learning policy drift. Domain 2 establishes a stratified diagnostic methodology: first localise the fault layer (mechanical / electrical / software / AI), then apply the layer's appropriate diagnostic logic, and finally identify the root cause and select the remediation path.

Domain 2 has a significant intersection with Domain 6 (AI/ML System Service Methods): Domain 2 handles the general fault-diagnosis framework for all robot systems, while Domain 6 focuses on the special diagnostic logic for AI components. The division line: when the fault phenomenon stems primarily from AI-model behaviour (such as perception-accuracy degradation or decision-policy drift), enter Domain 6's specialised methods; when the root cause lies in the mechanical, electrical, or classical software layer, Domain 2's framework applies.

Domain 3 Remote Service & Digital Operations

Industrial-robot service delivery is undergoing a structural shift: from reactive response based on on-site dispatch to proactive prediction based on data. The technical foundations of this shift — remote-connection protocols, digital-twin modelling, continuous sensor-data acquisition, and OTA firmware management — form the core of Domain 3. Domain 3 requires engineers not only to use remote-service tools but also to understand the interpretation logic of digital-twin data and the trigger mechanisms behind predictive-maintenance decisions.

On information security, Domain 3 is governed by ISO 27001 — remote-connection authorisation, tiered access permissions for diagnostic data, and third-party-provider access control all operate within the ISO 27001 framework. On top of this baseline, RSF adds a dedicated data-authorisation framework for the robotic-remote-service context.

2.1.2 Management Architecture Tier: Domains 4–6

The three capability domains of the management-architecture tier address service-system-level systematisation — the precondition for execution-tier capabilities to scale, standardise, and continuously improve.

Domain 4 SLA Design & Service Commitment

In the existing international-standards landscape, no ISO/IEC standard defines response time, availability commitment, or escalation processes for robotics service. This gap has left the robotics-service market without a common language for quality commitment: providers cannot make explicit commitments to customers, and customers cannot perform quantitative acceptance of service quality.

Domain 4 fills this gap, establishing structured methods for contractual expression, KPI-system design, and performance tracking in robotics service.

The SLA-metrics framework for robotics service differs fundamentally from that of IT service. IT service centres on RTO (recovery time objective) and RPO (recovery point objective), addressing logical faults that are fully remotely fixable; robotics service addresses physical-system faults, and its metrics should centre on OEE (Overall Equipment Effectiveness), MTTR (mean time to repair), and equipment availability — with response-time definitions distinguishing remote-response and on-site arrival as two separate dimensions.

Domain 5 Cross-Platform Service Procedures

Existing international standards set norms type by type: ISO 10218 for industrial arms, ISO/TS 15066 for cobots, ISO 31101 for service robots — a type-by-type structure that works well in plants with a single platform type, but creates fragmented service procedures in modern manufacturing and logistics scenarios where multiple platform types are deployed together.

Domain 5's core contribution is to establish a unified framework for cross-platform service procedures: it identifies what is shared across platforms (safety-verification workflow, work-order record format, delivery-acceptance standards) and systematically catalogues each platform's differentiated requirements (force-torque sensor verification for cobots, map-consistency verification for AMRs, proprioception calibration for humanoid robots), so that engineers switching between platforms can systematically invoke the right procedure instead of relying on scattered experience.

Domain 6 AI/ML System Service Methods

AI/ML penetration in robot systems has expanded from peripheral perception modules into the core decision layer: motion planning, grasping policy, spatial perception, and task generalisation are increasingly executed by trained neural networks rather than deterministic algorithms. This architectural evolution places entirely new demands on service engineers — the traditional "measure → compare to spec → replace part" diagnostic logic is essentially ineffective against AI-failure categories such as model degradation, distribution shift, and adversarial interference.

Domain 6 establishes a dedicated service methodology for AI/ML systems, covering AI-layer fault-symptom recognition (distinguishing the surface differences between AI faults and hardware faults), model-performance evaluation protocols (benchmark-test design, performance-metric interpretation), model-update and rollback procedures, and AI-service compliance foundations (the service-engineer-applicable clauses of the ISO/IEC 42001 AI-management system).

2.2 Inter-Domain Relationships and Service Scenario Mapping

In actual service work the six capability domains do not operate independently — they form a structured cooperative relationship. Understanding these relationships has direct operational meaning for correctly judging which capability domains a specific service task needs to invoke, and in what order.

2.2.1 Principal Inter-Domain Relationships

The relationship structure among the six domains can be understood along three dimensions:

- **Domain 1 ↔ Domain 2 — Phase-and-Diagnosis Relationship**

  Service Lifecycle Management (Domain 1) provides context to Fault Diagnosis (Domain 2): in judging a fault phenomenon, the engineer must first confirm the equipment's lifecycle phase — the same vibration anomaly has completely different diagnostic paths during commissioning (likely a parameter issue) versus during stable operation (likely mechanical wear).

- **Domain 2 ↔ Domain 3 — Field-and-Remote Cooperation**

  Fault Diagnosis (Domain 2) and Remote Service (Domain 3) form a dual-track structure of diagnostic capability: remote diagnostics is the preferred path (faster response, lower cost), while on-site diagnostics is the necessary supplement (when physical inspection, part replacement, or safety-procedure execution cannot be skipped). Domain 3's digital-twin and predictive-maintenance data give Domain 2's on-site diagnosis a historical baseline that significantly shortens diagnostic time after arrival.

- **Domain 4 ↔ Domains 1–3 — Commitment-and-Execution Relationship**

  The SLA (Domain 4) is the quality-framework constraint on the service-execution tier (Domains 1–3): Domain 4 defines what "fast enough" response time and "good enough" availability mean, and the execution quality of Domains 1–3 determines whether those commitments are met. Root-cause tracing of an SLA breach must traverse from Domain 4 back to the specific execution touchpoints in Domains 1–3.

- **Domain 5 ↔ All Domains — Procedure-Standardisation Relationship**

  Cross-Platform Service Procedures (Domain 5) is the multi-platform adaptation of execution methods from Domains 1–3: Domain 5 does not independently define diagnostic logic; instead, it systematically catalogues each platform's commonalities and differences in Domain 1 (maintenance phases), Domain 2 (diagnostic methods), and Domain 3 (remote-access protocols), so that the service procedures become switchable between platforms.

- **Domain 6 ↔ Domains 2 & 3 — AI-Specific Relationship**

  AI/ML Service Methods (Domain 6) is the deepening extension of Domain 2 (Fault Diagnosis) in the direction of AI components, and it shares data-acquisition infrastructure with Domain 3 (Remote Service). When Domain 2's stratified diagnostic locates a fault to the AI layer, Domain 6 supplies the further specialised analysis tools; Domain 3's continuous-monitoring data stream is the main data source for Domain 6's model-performance evaluation.

2.2.2 Typical Service Scenarios and Capability-Domain Invocation Matrix

The table below maps eight of the most common service scenarios for a robot-service engineer to the combinations of capability domains they principally invoke. The matrix is not a complete operations manual but a reference tool to help readers build the "scenario → domain" mapping intuition.

Table 2.2-1 Typical Service Scenarios and Capability-Domain Invocation Matrix

Service Scenario	Primary Domains Invoked	Trigger Mechanism and Key Tasks
New-Equipment Delivery and Commissioning	Domain 1 (primary), Domain 2 (supporting)	Customer procurement and on-site arrival → unpacking acceptance, mechanical mounting, electrical wiring, parameter configuration, safety-function verification, operator training delivery
Preventive Maintenance Execution	Domain 1 (primary), Domain 3 (data support)	PM cycle due or predictive-maintenance trigger → checklist inspection, lubrication, consumable replacement, performance verification, record update
Fault Response and Repair (On-Site)	Domain 2 (primary), Domain 1 (context), Domain 5 (platform procedures)	Customer reports downtime or performance anomaly → symptom record, stratified diagnosis, root-cause localisation, repair operation, verification test, work-order closure
Fault Response and Repair (Remote)	Domain 3 (primary), Domain 2 (diagnostic logic), Domain 4 (SLA tracking)	Monitoring alarm or remote customer report → remote-connection authorisation, data extraction, remote parameter adjustment or firmware push, disposition record
AI-Component Performance Degradation	Domain 6 (primary), Domain 2 (stratified localisation), Domain 3 (data acquisition)	Vision recognition rate drop, abnormal policy behaviour, generalisation failure → AI-layer fault recognition, model evaluation, data-annotation audit, model rollback or update
SLA Performance Monitoring and Customer Communication	Domain 4 (primary), Domains 1/2/3 (data sources)	Monthly / quarterly SLA review cycle → KPI data aggregation, MTTR / availability calculation, reporting to customer, variance analysis, improvement plan
Mixed-Platform Service Environment	Domain 5 (primary), Domain 2 (diagnosis), Domain 1 (lifecycle)	Plant simultaneously operates industrial arms + cobots + AMRs → procedure-switching recognition, platform-specific safety verification, unified work-order format execution
OTA Firmware Update and Version Management	Domain 3 (primary), Domain 1 (change records), Domain 2 (verification testing)	Manufacturer firmware release → update-package validation, current-version backup, staged rollout testing, function verification, rollback contingency

2.2.3 Mapping Between the Four-Layer Technical Architecture and the Capability Domains

The design of RSF's capability domains takes the robot system's four-layer technical architecture as its inner logical axis. These four layers — Mechanical Body, Electrical Control, Software Middleware, AI Decision — provide a systematic analytical framework for fault diagnosis and maintenance decisions. The mapping below explains the correspondence between the four layers and the six capability domains, and is the structural premise for understanding each domain's technical content.

Table 2.2-2 Mapping Between the Robot Four-Layer Technical Architecture and the RSF Capability Domains

Technical Layer	Layer Content	Principal Associated Domains	Typical Service Tasks

L1 Mechanical Body Layer	Joints, transmission systems (harmonic and RV reducers), end-effectors, chassis structure	Domain 1 (install / maintain), Domain 2 (mechanical-layer diagnosis), Domain 5 (platform procedures)	Lubrication, backlash measurement, TCP-accuracy verification, joint replacement, mechanical-wear assessment
L2 Electrical Control Layer	Servo drives, safety controllers (PLC), power systems, sensors (torque / position / vision), communication interfaces	Domain 1 (safety procedures), Domain 2 (electrical-layer diagnosis), Domain 3 (remote I/O monitoring)	Servo-parameter tuning, safety-circuit testing, LOTO operations, sensor calibration, communication-anomaly troubleshooting
L3 Software Middleware Layer	Robot operating system (ROS / ROS2), motion-controller firmware, coordinate-system management, communication protocols (EtherCAT / OPC UA)	Domain 2 (software-layer diagnosis), Domain 3 (OTA / remote debug), Domain 5 (cross-platform procedures)	Firmware updates, parameter backup / restore, log analysis, communication-cycle inspection, coordinate-system calibration verification
L4 AI Decision Layer	Vision-perception models, motion-planning networks, reinforcement-learning policies, natural-language command parsing, multi-modal fusion inference	Domain 6 (AI/ML methods), Domain 2 (AI-layer localisation), Domain 3 (model-monitoring data)	Perception-accuracy benchmark testing, model-performance evaluation, distribution-shift detection, model rollback, data-annotation quality audit

2.3 Four-Tier Certification and Capability Growth Pathway

The RSF certification system uses a four-tier progressive structure, with deepening cognitive dimension as its core logical axis: from Professional (Know-How — knows how to do it) to Specialist (Know-Why — knows why it works), then to Expert (Know-Across — knows how to integrate across domains), and finally to Master (Know-Beyond — knows how to define the future). Advancement between tiers is not merely an accumulation of knowledge but a systematic transformation of cognitive mode and work perspective.

The design logic of this certification architecture references the structure of the automotive ASE certification system (tier linked to authorised scope) and uses ITIL® certification pathways only as an external comparative reference for balancing knowledge and practical capability; the content is reconstructed entirely around the specific needs of robotics service — neither is replaceable by the other.

The four-tier system organises engineer development along a deepening cognitive dimension — Professional (Know-How), Specialist (Know-Why), Expert (Know-Across), Master (Know-Beyond) — with parallel Engineer and Service Manager tracks at every tier. The complete tier-by-tier definitions, including responsibilities for both tracks and coverage across the six capability domains, are given in Table 9.2-1 (§9.2).

2.3.1 Depth of Six-Domain Coverage Across Certification Tiers

The four certification tiers do not correspond to different capability domains; rather, they progress within the same six-capability-domain system, with continuously deepening cognitive depth and expanding scope of application. The table below sets out the core competence requirements of each capability domain at each tier, clarifying the assessment focus and the expected competence boundary at each level.

Table 2.3-2 Coverage Depth of the Six Capability Domains Across the Four Certification Tiers

Capability Domain	Professional	Specialist	Expert	Master
Domain 1 Service Lifecycle	Execute standard SOP, complete service tasks of the designated phase, fill in work-order records	Design phase-based maintenance strategies, assess decommissioning timing, optimise the service-process closed loop	Build a cross-organisation service-lifecycle management system that integrates ISO 55001 requirements	Define the best-practice standard for robotics-industry service lifecycle
Domain 2 Fault Diagnosis	Recognise common fault phenomena, execute stratified diagnostics per the decision tree, invoke standard remediation procedures	Independently handle cross-layer complex faults, lead root-cause analysis, deliver improvement recommendations	Design diagnostic-methodology frameworks, build platform-specific fault databases	Define industry-level diagnostic standards and assessment specifications for AI-assisted diagnostics
Domain 3 Remote Service	Use authorised remote tools to execute designated operations, report diagnostic data	Independently manage digital-twin data interpretation, formulate predictive-maintenance trigger rules	Design remote-service-platform architecture, integrate OTA and digital-twin systems	Define industry specifications for robotics-remote-service data governance and secure access
Domain 4 SLA Design	Understand and execute existing SLA terms, record and report SLA-related performance data	Participate in SLA negotiations, design KPI systems, manage SLA-violation handling processes	Lead SLA-contract-system design, establish service-pricing and cost models	Define industry-benchmark indicators and standard-clause frameworks for robotics-service SLAs
Domain 5 Cross-Platform Procedures	Recognise the procedural differences across robot platforms, execute operations per the designated procedure	Formulate unified service procedures for multi-platform environments, manage procedural version updates	Design cross-platform service-capability assessment matrices and authorisation-procedure systems	Participate in formulating international standards for cross-platform service procedures
Domain 6 AI/ML Service	Recognise AI-layer fault symptoms, distinguish AI faults from hardware faults, escalate per procedure	Execute model-performance evaluation, participate in data-annotation audits, manage model-update workflows	Design AI-service methodology frameworks, establish model-service-capability benchmark-test systems	Define industry-governance frameworks and ethics-compliance standards for AI-robot service

2.3.2 Advancement Requirements and Authorization Mechanism

RSF certification is directly linked to the engineer's authorised scope of work: the certification tier determines which types of service tasks the engineer is authorised to perform independently. This design references the authorisation logic of the ASE certification system — certification is not only proof of knowledge, but also a definition of the work boundary.

Table 2.3-3 RSF Four-Tier Certification — Advancement Requirements and Authorisation Boundaries

Advancement Path	Prerequisites	Assessment Method	Scope of Authorisation
Entry → Professional	No prior certification required; basic electromechanical background or relevant work experience is recommended	Written examination (specific pass threshold per the ATC Examination Guide)	May execute standard preventive-maintenance tasks under supervision; may independently complete basic work-order records; may

			participate in fault response under Specialist supervision
Professional → Specialist	Professional certification held for the prescribed period, plus sufficient independently completed work-order records (specifics per the ATC Examination Guide)	Written examination, practical assessment, and submission of a real service-improvement case (DMAIC format)	May independently handle complex faults without supervision; may lead SLA communication with customers; may audit and mentor Professional-level operations
Specialist → Expert	Specialist certification held for the prescribed period, plus leadership of multiple service-improvement projects with quantifiable outcomes (specifics per the ATC Examination Guide)	Submission of a white-paper-level case, plus oral defence before the RSF peer-review committee	May design service-system solutions; may serve as an ATC trainer; may submit revision recommendations for the RSF framework
Expert → Master	Expert certification held for the prescribed period, plus verifiable industry-level contribution (standards participation, peer-reviewed publication, or systemic innovation)	Nomination by RSF Board of Trustees, review of contribution dossier, plus peer review	Represents RSF in industry-standards formulation; may serve on the RSF Certification Committee; authorised to publish revision recommendations for the RSF framework

Note: this table summarises prerequisites, assessment method and authorised scope at a glance. Quantitative work-experience windows, performance evidence requirements and special-case notes for each advancement step are detailed in Table 9.2-2 (§9.2.2).

Chapter Summary **RSF's six capability domains** form an internally coherent service-competence system: the service-execution tier (Domains 1–3) provides the technical capabilities that act directly on robot systems, and the management-architecture tier (Domains 4–6) provides the systemic framework that allows execution capabilities to scale and standardise. The six domains have explicit cooperative relationships rather than being independent knowledge modules.

The four-tier certification system (Professional → Specialist → Expert → Master), with deepening cognitive dimension as its core logical axis, achieves tiered competence assessment and authorisation management within the same six-domain framework. Chapters 3–8 deepen the treatment of each domain in turn.

Chapter 3

Domain 1 · Service Lifecycle Management

3.0 Chapter Introduction

The value of a robot system is not "delivered" at a single moment in time — it is delivered continuously across its entire useful life. This judgement is the starting point of Domain 1, and the core of what differentiates it from a traditional equipment-maintenance manual: lifecycle management is not a taxonomy of repair work; it is dedicated to the systematic organisation and delivery of service value along the time axis.

From unpacking and power-on to controlled decommissioning, an industrial robot passes through service phases each with their own value objective, risk focus, and documentation requirements. The same vibration anomaly may point to a parameter-configuration issue during commissioning, mechanical wear during stable operation, or be the central evidence for a replacement decision during decommissioning assessment. Correct service action must be preceded by correct phase judgement — this is the fundamental logic of the Domain 1 capability system.

The chapter is organised along the five-phase service lifecycle: §3.1 establishes the phase model and value framework; §3.2–§3.6 unfold the service procedures phase by phase; §3.5 separately treats the incident → RCA → improvement three-tier loop that spans the operating phases; §3.8 defines the configuration-baseline system and documentation standards; §3.9 builds the Domain 1 KPI framework; §3.10 provides the chapter summary and inter-domain interface definitions.

3.1 Robot Service Lifecycle — Five-Phase Model

RSF defines the robot-service lifecycle as five sequential phases, each with explicit service objectives, value outputs, risk focus, and entry/exit gate conditions. The five-phase model takes the ISO 55001 asset-lifecycle-management framework as its principled basis, and at the operating layer is constructed around the technical characteristics of robot systems.

Table 3.1-1 RSF Robot Service Lifecycle — Five-Phase Model

Phase	Name	Core Service Objective	Value-Delivery Definition	Principal Risk Focus
Phase 1	**Delivery and Installation**	Safely and accurately complete the full process from unboxing to power-on readiness	System reaches a safe operable state; initial technical baseline is established	Installation error / electrical wiring error / omitted safety-function verification
Phase 2	**Run-in and Commissioning**	Verify that system performance meets application specifications; establish traceable accuracy and parameter baselines	Confirm the equipment meets the application commitment; establish the comparison baseline for future anomaly judgement	Lifetime erosion from improper parameter configuration / residual integration-interface defects
Phase	**Preventive**	Convert random fault risk into	Transform uncertain	PM strategy

3	**Maintenance Cycle**	planned maintenance activity, sustaining high system availability	downtime risk into a budgetable maintenance cost	mismatched to operating intensity / overlooked wear parts
Phase 4	**Corrective Maintenance and Overhaul**	Restore or enhance system performance through targeted repair or refurbishment, extending effective asset life	Restore or enhance system performance; extend the effective asset life	Over-maintenance / missing safety-function re-verification after overhaul
Phase 5	**Decommissioning Assessment and Disposition**	Close the asset lifecycle safely and compliantly based on combined technical and economic assessment	Safe, compliant closure of the asset; freeing up subsequent capital allocation	Non-compliant disposal of hazardous materials / incomplete data erasure

3.1.1 Stage-Gate Criteria for Phase Transitions

Transitions between phases are not triggered automatically by time but determined by explicit state conditions (stage gates). Meeting a gate condition indicates that the value objective of the preceding phase has been achieved and that the technical and documentary preconditions for entering the next phase are in place. Proceeding without satisfying a stage gate is the systemic root cause of many service problems.

Table 3.1-2 Stage-Transition Gate Conditions

Transition Node	Required Gate Conditions	If Gate Is Not Satisfied
Phase 1 → Phase 2	① Safety-function verification report completed and signed by the customer ② Initial Parameter Record archived (including load-inertia values) ③ All wiring identification and grounding verification completed	Stop. Without verified safety functions, dynamic commissioning is forbidden — violates ISO 10218.
Phase 2 → Phase 3	① Accuracy-acceptance report meets application specifications (TCP error within allowable range) ② Performance-baseline file established (per-axis current baselines, vibration characteristic frequencies) ③ Delivery Acceptance Certificate (DAC) signed by both parties	Stop. Entering operations without baseline files leaves no reference for future anomaly judgement.
Phase 3 → Phase 4	① MTBF declines over three consecutive PM cycles, or	Continue Phase 3 with intensified monitoring; do not trigger overhaul on subjective judgement alone.

	② Annualised maintenance cost exceeds a defined percentage of original equipment value (economic-assessment trigger), or ③ The manufacturer announces EOL on a critical part	
Phase 4 → Phase 5	① Decommissioning-assessment matrix score lands in the "recommend retirement" band, or ② Availability still fails to meet SLA targets after overhaul, or ③ Production-line strategic change triggers a proactive retirement decision	Return to Phase 4 for re-assessment; or request a temporary SLA-term adjustment; do not execute decommissioning directly.

3.1.2 Lifecycle Extension Pathways

The lifecycle of a robot system does not terminate linearly at retirement. After Phase-5 assessment is complete, there are three possible asset pathways depending on system state and business need, each placing different demands on service activity:

Pathway	Applicable Conditions	Service Requirements	Key Risks
Pathway A — Continue Operation	Technical condition supports downgraded duty; maintenance cost remains controllable	Define a downgraded-operation procedure; adjust the PM strategy; update SLA parameters	Over-extension of asset life raises safety risk
Pathway B — Refurbish and Redeploy	Mechanical body in good condition; principal issues concentrated in the control system or a specific joint	Follow the Phase-4 overhaul procedure; complete a full accuracy re-verification; update the equipment passport	Incomplete safety-function re-verification after refurbishment
Pathway C — Controlled Decommissioning	Decommissioning-assessment matrix indicates no economic repair value; or strategic exit	Execute the full decommissioning procedure in §3.7; complete data erasure and compliant hazardous-material disposal	Data-security and hazardous-material disposal violations

3.2 Phase 1 · Delivery and Installation

Delivery and installation is the highest-risk-density phase of the robot lifecycle: the work done here both lays the foundation for years of stable operation and rapidly cements any errors into systemic defects that are very hard to trace. Quality deviations at this stage — from mechanical-mounting torque error to omitted safety-function verification — often do not manifest immediately but surface months later as recurring faults whose root cause is difficult to attribute.

On top of standardised installation procedures, this section introduces a structured project-management perspective: installation is a temporary engagement with an explicit start, end, deliverables, and quality-acceptance criteria, not a simple sum of independent technical operations. This perspective allows the multi-party installation process (service engineer, customer engineer,

system integrator, electrical contractor) to manage risk, control quality, and define responsibility within a unified framework.

3.2.1 Pre-Installation Preparation: Scope Definition and Risk Anticipation

Before any on-site installation work begins, three preparatory activities must be completed: scope and boundary definition, site-condition verification, and risk anticipation with response plan. The quality of these three directly determines whether the installation process is controllable.

Scope and Responsibility-Boundary Confirmation

Installation projects typically involve multiple parties: the robot OEM / service provider (body installation and commissioning), the electrical contractor (power supply and wiring), the system integrator (peripheral-device interfaces), and the customer engineering team (process confirmation and acceptance). Unclear responsibility boundaries are the most common root cause of installation delays and quality disputes.

Table 3.2-1 Installation-Phase Multi-Party Responsibility Matrix (RACI)

Work Package	Robot Service Engineer	Electrical Contractor	System Integrator	Customer Engineer
Unboxing acceptance and damage record	R (Responsible)			A (Approves)
Foundation / base installation	R		C (Consulted)	A
Power supply and main breaker	C	R		A
Internal control-cabinet wiring	R	C		A
Robot-to-peripheral I/O wiring	R	C	R	A
Safety-circuit integration and testing	R	C	R	A (Must be present)
Parameter configuration and software initialisation	R		C	I (Informed)
TCP accuracy acceptance test	R			A
Safety-Function Verification	R	C	C	A (Must be present)
Operator training delivery	R			A

Key: R = Responsible for execution; A = Approves / accepts; C = Consulted; I = Informed

Pre-Installation Site Survey

Before the engineer arrives on site, site conditions must be verified. Unsuitable site conditions are a high-probability root cause of installation interruption, schedule slippage, and later faults; pre-arrival verification is far cheaper than rework on-site.

Table 3.2-2 Pre-Installation Site Survey Checklist

Inspection Category	Specific Items to Verify	Acceptance Criteria	Action if Non-Conforming

Civil Foundation	Floor flatness; anchor-bolt position and specification; load-bearing capacity	Flatness ≤ 2 mm/m; anchor-bolt position deviation ≤ 3 mm; load capacity ≥ robot rated weight × dynamic-load factor	Require customer rectification; suspend installation entry
Power Conditions	Three-phase voltage stability; short-circuit capacity; ground resistance; breaker rating of distribution panel	Voltage deviation ≤ ±5%; ground resistance ≤ 4 Ω; breaker rating matches controller nameplate	Require rectification by the electrical contractor and re-inspect
Workspace	Clearance within the robot's working envelope; safety-fence placement; safety-door reserved location	No obstructions intrude into the working envelope; safety distance per ISO 10218-2 Annex C	Adjust equipment layout or negotiate an enlarged protection zone
Environmental Conditions	Temperature and humidity; dust and chemical corrosives; noise; lighting	Controller ambient temperature 0–45 °C; humidity ≤ 75 %RH non-condensing; no corrosive gases	Evaluate protection upgrades (higher IP rating / air-conditioned cabinet)
Lifting Access	Clear width of the body-moving path; rated capacity of lifting points; floor capacity	Path clearance ≥ robot's maximum dimension + 500 mm; lifting-point ratings verified	Plan the lifting path in advance; prepare auxiliary tools

Installation-Phase Risk Register

Installation-phase risks are highly predictable: the root-cause categories of installation problems are well-documented in the robotics industry. Completing risk identification before on-site work converts reactive response into proactive prevention.

Table 3.2-3 Installation-Phase Risk Register

Risk Description	Likelihood	Impact	Mitigation
Incorrect load-inertia entry overloads joints and shortens reducer life	**High**	**High**	Installation checklist requires entering correct Load Inertia before power-on; on FANUC systems, verify via iHMI
Safety-function verification skipped or performed perfunctorily	**Medium**	**Very High**	Stage-gate condition: do not enter commissioning without dual-signed safety-verification report
Cobot torque-parameter baseline not recorded; subsequent verification of safety-parameter drift impossible	**Medium**	**High**	Mandatorily perform ISO/TS 15066 torque-baseline measurement at end of commissioning and archive it
Disordered electrical-wiring identification leads to wrong-circuit isolation during later maintenance	**High**	**Medium**	After all wiring is complete, perform 100 % identification check and archive photos
Misaligned on-site arrival times of the multiple parties disrupt sequencing	**High**	**Medium**	Hold a pre-installation coordination meeting to confirm each party's timing; maintain a rolling 48-hour plan
Shipping damage not detected before installation completion makes later fault traceability difficult	**Medium**	**Medium**	Unboxing acceptance must be performed under full documentation (photos + notes); installation is suspended on damaged equipment
EtherCAT cycle configuration	**Medium**	**High**	Conduct communication-link integrity pre-

incompatible with the controller			testing before system-integration commissioning; confirm all slaves respond

3.2.2 Delivery Acceptance Procedures

Delivery acceptance is the first quality gate of the installation and the only opportunity to assign responsibility for shipping damage. Once unboxing acceptance is signed off without objection, any later-discovered damage faces serious obstacles in liability attribution. Unboxing must be performed under full documentation; "outer packaging intact" does not substitute for physical inspection.

The four steps of delivery acceptance:

> **Step 1 Outer-Packaging Inspection —** Before unboxing, inspect the shipping packaging: deformation, water staining, state of tilt-indicator labels (if equipped). Photograph all packaging faces as original record. Any anomaly must be raised with the logistics provider on-site before unboxing.
>
> **Step 2 Physical Inventory —** Inventory each item against the packing list: robot body, controller cabinet, teach pendant, connection cables (power and signal), end-effector (if shipped), standard tool kit, product documents (manuals, conformity certificates, calibration certificates). Shortages must be reported immediately to the OEM service contact.
>
> **Step 3 Visual Inspection —** Conduct a comprehensive visual inspection of the robot body: joint interfaces, reducer outputs, cable entry/exit points, controller cabinet panels, and terminal blocks. Record any mechanical deformation, paint scratches (deep scratches may indicate impact damage), or loose parts.
>
> **Step 4 Sign-Off Documentation —** If acceptance has no objection: sign the Delivery Acceptance Confirmation, recording the inspector, time, and equipment serial number. If acceptance has objections: archive photos, fill out the Goods Damage Record, suspend installation, and await OEM instructions. Installation may not proceed while damage remains unresolved.

3.2.3 Mechanical Installation and Safety Protection

The mechanical-installation phase must complete body fixing and safety-protection device installation in parallel — the two cannot be separated. Safety-protection devices are a system-integration requirement of ISO 10218-2, not an optional add-on.

Base-Installation Accuracy Requirements

Base-installation accuracy directly affects TCP accuracy and joint load distribution. After installation, measure and record using a precision level as the mechanical baseline for subsequent accuracy tracking. The principal accuracy requirements are:

Parameter	General Standard	Notes
Base flatness	≤ 0.1 mm/100 mm (high-accuracy applications ≤ 0.05 mm/100 mm)	Deviations must be corrected with shim adjustments; deformed bases may not be used

Mounting-bolt torque	Per OEM manual (e.g., FANUC R-2000 series: M20 bolt ≈ 385 N·m)	Torque values vary significantly by model; estimation is forbidden; verify with a torque wrench
Bolt anti-loosening measures	OEM-specified thread locker (e.g., Loctite 243) or spring washers	Anti-loosening type per OEM manual; substitutions are not permitted
Ground connection	Ground resistance ≤ 4 Ω; ground-wire cross-section ≥ 6 mm²	The controller cabinet and robot body must be grounded independently — not series-connected

LOTO (Lockout / Tagout) Procedure Establishment

Before first power-on, a LOTO (Lockout / Tagout) procedure file must be established for the equipment, and all engineers involved in installation and subsequent maintenance must be briefed on it. LOTO is a foundational on-site safety procedure required by ISO 45001 and ISO 10218 — it may not be omitted under any circumstances.

> **Minimum requirements for the LOTO procedure file: ① identify all energy sources (main power, pneumatic, hydraulic, gravitational potential); ② define the isolation and locking method for each energy source; ③ define verification steps (after tag-out, verify de-energisation with a voltmeter); ④ specify the responsible custodian for each energy lock; ⑤ post the file in a visible location on the controller cabinet and include it in the equipment passport.**

Safety-Protection Device Verification

Safety-protection device placement must conform to the safety-distance calculation method of ISO 10218-2 Annex C, ensuring that when entry into the hazard zone is detected, the device completes its stop response before the robot reaches that zone. Principal verification items include:

- Correct installation position and electrical wiring of safety-door interlock sensors
- Safety-light-curtain protection height and resolution match application requirements
- E-stop button layout covers all operating positions; configured in the controller as Category 0 or Category 1 stop
- Speed and position limit software functions (for cobots: also verify torque-limit setpoints and record the ISO/TS 15066 torque baseline)
- Safety-fence grounding: metal safety fences must be connected to the main grounding system

3.2.4 Electrical Wiring and System Integration

Electrical wiring is the installation step with the densest technical interfaces and the highest multi-party-coordination demand. This section focuses on wiring-quality control and system-integration verification; detailed electrical-design content is governed by project engineering drawings and is not within scope here.

Wiring Quality-Control Highlights

- All cables must be identified per the OEM wiring diagram (wire numbers / terminal numbers) with permanently durable labels; temporary tape labels are forbidden
- After controller-cabinet wiring is complete, perform torque-check sampling on all terminals (random 20 %; if any fails, re-check all)
- Power and signal cables must be routed in separate ducts with ≥ 100 mm spacing (≥ 200 mm between high-voltage power and low-voltage signal) to prevent EMI
- Robot-body cable routing: J1–J6 motor-power cables and encoder cables must follow OEM bend-radius specifications to prevent fatigue failure during motion
- Before first power-on: insulation-resistance test all power cables with a Megger (≥ 1 MΩ @ 500 V DC) and record the values

I/O Interface and Communication-Link Verification

After interface commissioning between the robot and peripherals (PLC, vision system, conveyor, safety controller), perform full signal-mapping verification:

Verification Category	Verification Method	Recording Requirements
Digital I/O signals	Point-by-point forced output (Force Output) with confirmation of correct response at the receiving device	I/O signal-mapping table recording the correspondence between physical terminal numbers and logical addresses
EtherCAT bus communication	Confirm all slaves in the EtherCAT topology are in OP state; measure cycle timing (PDO) with no packet loss	Communication-topology screenshot; communication diagnostic log
Safety circuit (Safety Bus)	Simulate triggering of each safety source (E-stop, safety door, light curtain) and confirm the correct stop category is executed	Safety-function test record (each safety source recorded separately)
Vision-system interface	Run the standard vision-calibration procedure; confirm the transformation-matrix error between vision frame and robot frame is within specification	Vision-calibration report; hand-eye calibration-error record

3.2.5 Software Initialisation and Parameter Configuration

Software initialisation is one of the installation steps with the most far-reaching effect on long-term equipment health. The parameter configuration completed here — especially load-inertia setting — directly determines the force state borne by joint reducers across the entire useful life.

Controller First-Boot Procedure

- Complete initial configuration of language, time zone, unit system, network IP address
- Confirm controller firmware version matches the application software-package version (record both)
- Run the initial system-diagnostic routine; confirm all axis servo drives report normal status

- Set the system clock and synchronise with the customer's network time server (accuracy ±1 s — important for downstream log-timestamp consistency)
- Complete software licence activation and archive a copy of the licence file in the equipment passport

Correctly Setting Load Inertia (Critical)

> **Industry Alert FANUC explicitly notes that entering the correct Load Inertia is "one of the most important settings affecting robot life." An incorrect value causes torque estimation bias in the servo controller, exposing joint reducers to over-spec shock loading on every motion cycle, which accelerates reducer wear within thousands of operating hours — wear that is not directly visible in day-to-day operation. The installation engineer must complete this setting before first program execution, and must base it on measured values of the actual end-effector and payload, not estimates.**

Correct procedure for load-inertia setting:

- Obtain the end-effector weight and centre-of-gravity position (source: tool manufacturer's specification sheet)
- Obtain the weight and centre of gravity of the maximum-payload workpiece
- Enter payload parameters via the controller's specified interface (FANUC: System > Motion > Payload; ABB: ABB > Jogging > Payload)
- Run the automatic load-identification routine (if the controller supports it); compare the measurement with the entered value (deviation ≤ 10 %)
- Record the load-inertia parameter values in the Initial Parameter Record

Initial Parameter Record (IPR)

The Initial Parameter Record is one of the most important technical documents handed to the customer at the end of commissioning, and the reference baseline for all subsequent performance tracking and anomaly judgement. It must include the following core parameters:

Parameter Category	Required Items	Recording-Format Notes
Payload parameters	Payload weight, CG X/Y/Z offsets, and Ix/Iy/Iz inertias for each Tool frame	Numeric value + units (kg, mm, kg·m²); tool-name identifier
Per-axis current baselines	Rated and peak current of J1–J6 under the standard test program (no-load and full-load sets)	Record measurement time and test-program name; values to 0.1 A precision
Accuracy baseline	TCP position accuracy (mm), repeatability (mm), path accuracy (mm/m)	For each item, record measurement-instrument model, number of measurement points, and ambient temperature
Software versions	Controller firmware version, application software-package version, last update date	Record version strings in full — no abbreviation
Safety parameters	Setpoints of each safety function (E-stop category, safe-speed limit, cobot torque limit)	Safety parameters must bear the engineer's signature confirming verification completion

3.2.6 Safety Function Verification

Safety-function verification is the statutory precondition for moving from installation to commissioning — it may not be omitted or reduced to visual inspection. ISO 10218-2 requires the system integrator to ensure overall safety-system function, a responsibility discharged only through concrete verification operations and written records. Safety-function verification records are technical documents of legal significance and must be preserved permanently in the equipment passport.

Table 3.2-5 Safety-Function Verification Checklist

Safety Function	Verification Method	Acceptance Criteria	Recording Requirements
Main E-Stop	Press any E-stop button and observe robot response	Category 0 stop (immediate power-off); cannot restart from controller after stop; manual E-stop reset required	Button ID + test result + tester signature
Safety-door interlock	Open the safety door while the robot is running	Robot enters stop state within 1 s of door opening (Category 1 stop)	Door ID + response-time record
Safety light curtain	Trigger the curtain with a test rod (diameter per curtain resolution)	Robot responds immediately (Category 0 or 1 stop per design)	Curtain ID + test-rod spec + result
Speed limit	Operate the robot in teach mode (T1)	TCP speed does not exceed 250 mm/s (ISO 10218 requirement)	T1-mode speed-measurement record
Cobot torque limit (if applicable)	Apply force in the TCP direction with a calibrated force gauge to trigger torque protection (test-location selection and force criteria per ISO/TS 15066:2016 Annex A — consult the original standard)	Protective stop must be triggered by external force within the limits specified by ISO/TS 15066	Compare measured trigger force (N) to setpoint at each test location; record must note the referenced ISO/TS 15066 version
STO (Safe Torque Off)	Issue STO command via the safety PLC	All axis motor torques drop to zero immediately; no axis movement under gravity	STO trigger method + response time

3.2.7 Delivery Milestone and Project Closure

Installation completion is marked by formal project-closure documentation, not by the engineer's departure. The closure documents fulfil two functions: confirming that Phase 1's value objective has been achieved (system reaches a safe operable state); and providing the evidence required by the Phase 2 stage-gate.

Operator Basic-Training Delivery

At installation closure, basic operator training must be delivered to the customer's operating staff — a critical step in ensuring the equipment runs safely after the service engineer leaves. The training must cover: normal operation (start, stop, program call); emergency handling (E-stop activation and reset); daily visual-inspection points; prohibited operations (in particular, no entry into the safety-fenced area in Auto mode). Training is recorded on a Training Delivery Confirmation form signed by participants.

Phase 1 Delivery-Milestone Document Package

Table 3.2-6 Phase 1 Delivery-Milestone Document Package

Document Name	Core Content	Signature Requirements	Filing Location
Installation & Commissioning Report	Site-survey results, installation parameters, wiring-test results, I/O verification summary	Signed by service engineer; confirmed by customer engineer	Equipment passport (required)
Safety-Function Verification Report	Per-safety-function test result with measured values	Dual signature: service engineer + customer engineer (both required)	Equipment passport (preserved permanently)
Initial Parameter Record (IPR)	Payload parameters, per-axis current baselines, accuracy baseline, software versions, safety parameters	Signed by service engineer	Equipment passport; cloud backup
Delivery Acceptance Certificate (DAC)	System meets contract-specified delivery standards; all Phase 1 stage-gate conditions satisfied	Dual signature: service engineer + customer-authorised representative	Contract file + equipment passport
Training Delivery Record	Training date, participants, content covered	Signed by trainer and all participants	Equipment passport

> **Phase 1 Experience-Capture Note Within five business days after installation closure, the service engineer should complete an Installation Experience Record capturing the non-standard situations encountered during this installation (site-condition deviations, design omissions, customer-specific requirements) and the responses taken. These records are not included in the customer file pack, but are uploaded to the organisation's internal knowledge base for reference by future similar projects — a low-cost, high-return mechanism for reducing recurrence of similar issues.**

3.3 Phase 2 · Run-in and Commissioning

The core task of the run-in and commissioning phase is to establish traceable performance baselines that provide a reference frame for years of subsequent operation and maintenance. This is the lifecycle's densest phase of "investment information production" — every baseline record laid down here yields multiples of diagnostic-efficiency return in future anomaly judgement.

3.3.1 Technical Characteristics of the Run-in Period

Robot-system reliability follows a typical bathtub curve across the lifecycle: the run-in period carries relatively higher early-failure risk, the stable-operating period sustains low failure rate, and the wear-out period rises again.

Principal causes of early run-in failure include: adaptive displacement of manufacturing-assembly tolerances under early dynamic loading (initial contact-stress concentration on harmonic-gear teeth); micro-asperity removal on bearing raceways and rolling elements during early operation; and stabilisation of contact resistance in cables and connectors during early motion cycles. Within normal range these processes are expected — but they must be monitored so that out-of-range run-in failure can be detected early.

Run-in monitoring focus parameters: deviation of each axis's drive current from the no-load baseline; rise trend and steady-state of each axis's operating temperature; abnormal vibration (frequency and amplitude); time-stability of TCP accuracy (whether early drift converges to a stable value).

3.3.2 TCP Accuracy and Repeatability Acceptance

TCP (tool centre point) accuracy verification is the core deliverable of commissioning and the key evidence for whether the system meets application specification. Accuracy verification must be performed under rated load — no-load accuracy does not reflect the real operating condition.

Accuracy Type	Definition	Typical Industry Benchmark	Acceptance Notes
Repeatability	Deviation in reaching the same target point repeatedly from the same direction (per ISO 9283)	Precision assembly: ±0.02–0.05 mm · general material handling: ±0.05–0.1 mm	Measure under rated load; repeat at least 30 cycles; record 3σ value
Absolute Accuracy	Actual error reaching a specified coordinate position (difference from programmed value)	General industrial: ≤ 0.5 mm · precision machining: ≤ 0.1 mm	Requires a laser tracker; costlier than repeatability — perform only if application requires
Path Accuracy	Actual deviation of the TCP from the programmed path during motion	Welding: ≤ 0.5 mm / 500 mm · spray painting: ≤ 1 mm	For applications with high path-accuracy demand (welding / cutting / painting); general handling can omit

3.3.3 Establishing the Performance Baseline

The performance baseline is this phase's most important deliverable. The baseline file's core value is not in proving current equipment state but in providing the reference standard for years of future state comparison. A maintenance system without baseline files cannot execute effective state management.

- **Per-Axis Current Baseline** Under the standard test program (defined trajectory and speed), measure J1–J6 drive currents (rated and peak) under no-load and rated-load conditions. This is the most direct comparison baseline for later judgement of reducer wear — reduced efficiency manifests directly as current rise.
- **Temperature-Distribution Baseline** After one hour of continuous operation, record each reducer's casing temperature and the controller-cabinet internal temperatures (near control board and exhaust). Abnormal temperature rise is the early signal of blocked cooling or grease aging.
- **Vibration-Frequency Baseline** For controllers with built-in vibration analysis (e.g., FANUC iRCalibration), record the vibration-frequency signatures of the principal motion axes. For systems without built-in vibration analysis, take a one-time baseline using accelerometers mounted at key joints.
- **Acoustic Baseline (Optional)** Under a defined motion program, record each axis's operating

sound signature as a subjective reference for future abnormal-noise judgement — low cost, high value, especially useful for early detection of reducer wear.

3.3.4 Phase 2 Closure Documentation

- Accuracy-acceptance report (measurement method, measured values, acceptance conclusion)
- Performance-baseline file (per-axis current baselines, temperature baseline, vibration / acoustic baselines)
- Integration-test report (I/O signal mapping, vision-calibration results, external-axis coordination test conclusion)
- Delivery Acceptance Certificate (DAC) — gate document for Phase 2 → Phase 3, dual-signed by customer and service party

3.4 Phase 3 · Preventive Maintenance Cycle

Preventive maintenance is the longest phase in the lifecycle, typically over 80 % of the robot's useful life. Service quality in this phase ultimately decides the robot's effective life and economic value. The core logic of PM is to replace unplannable, high-cost reactive fault response with plannable, low-cost proactive service action. Industrial-practice data show that a structured PM plan can reduce unplanned downtime by 50–70 %.

3.4.1 Three Models of PM Strategy

PM Strategy	Trigger Mechanism	Applicable Scenarios	Limitations
Calendar-Based	Fixed time interval (e.g., every 12 months)	Stable utilisation, single-shift operation	Structural shortcoming: a three-shift robot uses the same interval as a single-shift robot, but the effective PM frequency differs threefold
Runtime-Based	Cumulative operating hours reach threshold	The OEM-recommended mode; applicable across all operating intensities	Requires a reliable hour-meter (built into most modern controllers)
Condition-Based (CBM)	Sensor data or diagnostic metrics exceed threshold	Critical equipment, high-OEE scenarios (e.g., ABB Connected Services)	Requires remote-monitoring infrastructure; high upfront investment; demands data-analysis capability

> **OEM-Recommended PM Intervals (Standard Practice — Major Brands) FANUC: every 3,840 operating hours or 12 months (whichever comes first) for routine daily maintenance; every 11,520 operating hours (~3 years) for critical-axis reducer deep grease replacement. KUKA / ABB: maintenance systems have a clearly tiered structure — routine system inspection and condition assessment typically at 6,000–10,000 operating hours; full gearbox / reducer lubricant replacement at 20,000–24,000 operating hours (or 5–6 years). Adverse-environment adjustment: in dust, high-temperature, or heavy-shock-load environments, industry best practice is to shorten thresholds to 80 % or below (e.g., FANUC routine PM from 3,840 h to ~3,000 h). Operating-regime adjustment: robots running three-shift (24/7) continuous operation typically accumulate more than three times the annual hours of single-shift operation.**

> **Organisations must use actual operating hours — not calendar months — as the primary PM trigger to prevent assets from failing catastrophically due to overdue service.**

3.4.2 Stratified PM Operating Standards

Table 3.4-1 Stratified PM Operating System

Maintenance Level	Execution Frequency	Core Work Content	Executor Level
Shift Inspection	Every shift change-over (~8 hours)	Visual inspection: cable wear / dust accumulation / oil leakage Listen for abnormal noise / vibration (compare to acoustic baseline) Controller-vent cleanliness E-stop button visual confirmation Status indicators normal	Professional (within 15 minutes)
Monthly Check	Monthly	Lubricant-seepage check (joint-to-reducer interface) Check bend state of all exposed cables Sensor (photoelectric / proximity) lens cleaning I/O response sampling test (5 points) Mechanical zero-mark consistency check	Professional / Specialist
Quarterly Check	Quarterly	Per-axis backlash measurement compared to baseline Brake-holding torque test (J1–J3) Controller cooling-fan speed / noise check Record per-axis drive currents and compare to baseline Cabinet seal condition check	Specialist
Annual / Milestone PM	Annual, or upon reaching OEM milestone hours	Full reducer-grease replacement (OEM-specified type, no substitutes) Servo-motor brake inspection and brake-torque measurement Replace CPU battery and encoder backup batteries Mechanical-unit cable integrity check (focus on high-bending segments) Full software backup of teach pendant and controller (duplicate) Re-measure TCP accuracy and compare to commissioning baseline	Specialist (with dedicated spares)

3.4.3 Adjusting PM Strategy in Special Environments

Table 3.4-2 Special Operating Environments and Their Influence on PM Strategy

Environment Type	Principal Influence Mechanism	PM Strategy Adjustment	Additional Key Inspection Items
High Dust (foundry / welding shop)	Dust ingress accelerates joint-seal wear; cooling-	Shorten PM thresholds to 80 % of OEM standard;	Clean cabinet vents every shift; monthly seal-integrity check

	duct blockage raises temperature	double cooling-check frequency	
High-Temperature Environment (steel / casting / oven)	Accelerated grease oxidation (heat-driven aging); risk of electronic-component over-temperature	Shorten grease-replacement interval to 60 % of standard; add temperature-sensor monitoring	Quarterly measurement of internal controller temperature; check grease colour at each PM (blackened oxidation requires immediate replacement)
Cleanroom / Food-Grade (pharma / food)	IP rating may be compromised by cleaning operations; non-food-grade grease poses contamination risk	Confirm robot in safe position before cleaning; verify IP-rating integrity after maintenance	Use OEM-specified food-safe grease (NSF H1 / H2); inspect all seal points
High-Frequency / High-Speed Cycling (packaging / pick-and-place)	Joint fatigue accumulates far faster under short-cycle high-speed motion than under low-speed applications	Use cycle count (rather than hours) as a secondary trigger; continuously monitor joint temperature	Quarterly torque check on end-effector connection bolts; annual backlash check on J4 / J5 / J6 reducers

3.5 Incident Mitigation & Problem Resolution Loop

Running through the robot-service lifecycle, alongside scheduled maintenance, is a three-tier closed-loop mechanism that converts single fault events into systemic improvement inputs. The closed-loop logic is: Incident Response solves the immediate downtime; Root Cause Analysis (RCA) prevents recurrence of the same fault mode; and Service Improvement converts the root-cause finding into systemic change embedded in PM strategy, SOPs, or spare-parts policy.

The value of the closed loop lies in distinguishing "fixing" from "improving": completing only incident response solves a problem; completing the full closed loop elevates the system. A service organisation's maturity is largely measured by how many fault events trigger the full loop instead of stopping at the first tier.

3.5.1 Level 1 · Incident Response Procedure

Table 3.5-1 Fault Response — Standard Eight-Step Procedure

Step	Node Name	Core Operation	Key Cautions
S1	**Arrival / Remote Connection**	After receiving a fault work order, confirm response mode (on-site / remote); record response start time (SLA-timing starts here)	Response time counts toward SLA performance and must be recorded immediately
S2	**Symptom Recording**	Collect customer description, review alarm codes, observe actual robot state; photograph or record video	Record must objectively describe the phenomenon; subjective judgement is excluded
S3	**Safety Confirmation**	Before any operation, confirm the system is in a safe state (E-stop engaged / energy isolated); apply LOTO if physical contact is required	Do not touch the robot until safe state is confirmed
S4	**Preliminary Stratified Localisation**	Using Domain 2 diagnostic methods, identify the fault layer (mechanical / electrical / software / AI); confirm the remediation path	Stratified localisation determines subsequent direction and may not be skipped
S5	**Repair Operation**	Execute remediation per the localisation; for spare replacements, record both old and new part serial numbers	Spare replacements require part numbers and replacement date logged in the equipment passport
S6	**Function Verification**	After repair, run a function-verification test	Safety-function re-verification

		(executes the test program and confirms symptom elimination); safety-related repairs require safety-function re-verification	after safety-related repair may not be skipped
S7	**Work-Order Documentation and Closure**	Complete the work order in full (hours / disposition / spares used / outstanding items); customer countersigns	Work orders are the basis of SLA settlement and knowledge accumulation
S8	**Recurrence-Risk Assessment**	Decide whether to trigger RCA (criteria in §3.5.2); if not triggered, record the rationale	This step determines whether the second tier of the closed loop is engaged

3.5.2 Level 2 · Root Cause Analysis Triggers and Methods

RCA triggering must follow explicit criteria — neither triggered for every fault (over-analysis wastes resources) nor never triggered (faults recur). Initiate RCA if any of the following is satisfied:

- The same fault mode (same alarm code + same symptom location) recurs at least twice within 12 months
- A single fault causes unplanned downtime exceeding 4 hours
- The fault involves safety-function failure or safety-related component damage (regardless of frequency)
- The fault causes a customer SLA breach
- The symptom does not match any known fault mode (a novel fault type requiring a new knowledge entry)

RSF-recommended RCA-method combinations (select by fault complexity):

Method	Applicable Scenarios	Application Notes in Robot Service
5-Why Analysis	Single-path faults; reasonably clear causal chain	Suited to common mechanical / electrical layer faults (e.g., reducer noise → grease shortage → improper PM interval → incorrect hour-counting method → …)
Fishbone (Ishikawa) Diagram	Multi-factor compound faults requiring systematic search of the cause space	Suited to intermittent faults and AI/ML-layer faults (broad cause space — structured search along mechanical / electrical / software / environmental / operational axes)
Fault-Tree Analysis (FTA)	Safety-related faults; high-consequence events	Suited to root-cause tracing of safety-function failure; requires logic-gate analysis; results must be added to the equipment safety file

3.5.3 Level 3 · Controlled Change and Service Improvement

After RCA, the root-cause conclusion must be converted into systemic service-improvement action through a controlled-change mechanism. The core requirement of "controlled change" is: any modification to PM plans, SOPs, spare-parts policy, or parameter settings must follow an explicit approval workflow and be applied synchronously across all relevant equipment — not only on the unit that failed.

Table 3.5-2 Service-Improvement Actions — Classification and Change-Control Requirements

Improvement Action Type	Typical Trigger Scenarios	Change-Control Requirements	Synchronisation Scope
PM-Interval Adjustment	RCA reveals that an overly long PM interval has accumulated wear	Approval by Specialist-or-above engineer; update the PM-plan document version	All equipment of the same model and environment
SOP Revision	Discovery of omitted steps or incorrect descriptions in an existing SOP	SOP update must note reason and date of change; previous version archived	All engineers using that SOP must receive an update notice
Spare-Parts Policy Change	Emergency procurement of a spare type persistently exceeds 10 %	Update safety-stock setting; synchronise to inventory-management system	All machines that consume this spare
Parameter-Setting Change	RCA identifies a parameter-setting deviation as the root cause	Parameter changes must be logged in the controller operation log (with before / after values)	All equipment of the same model and same application scenario

3.6 Phase 4 · Restorative Maintenance and Overhaul

Once a robot system enters the wear-out period, its performance indicators undergo systemic decline: TCP-accuracy drift exceeds baseline tolerance, MTBF shortens persistently, drive currents climb. Service decisions then shift from "how to fix this fault" to "by what means to restore asset value with the best cost-effectiveness."

3.6.1 Quantitative Recognition of Performance Degradation

Table 3.6-1 Phase 4 Trigger Indicators and Judgement Method

Monitored Indicator	Phase 4 Trigger Threshold	Data Source	Judgement Method
MTBF Trend	MTBF declines over three consecutive PM cycles (not an isolated event)	Work-order historical-records system	Plot MTBF time series; linear-regression slope significantly negative
Annualised Maintenance-Cost Ratio	Annual unplanned maintenance cost / original equipment value > 15 %	Work-order cost records	Aggregate by fiscal year; multi-year trend is more meaningful than absolute value
Accuracy Degradation	TCP repeatability exceeds 150 % of baseline value	Quarterly accuracy-measurement records	Compare to commissioning baseline; rule out the effect of tool wear
Critical-Part EOL	OEM announces end-of-life for a critical part	OEM technical notice	Estimate remaining life supportable by current stock; build an EOL transition plan

3.6.2 Economic Assessment Model for Overhaul

> **Three-Way Economic Judgement "Continue Maintenance" — when the annualised maintenance-cost ratio is below 15 % and accuracy still meets application requirements, continue Phase 3 PM with intensified monitoring. "Overhaul / Refurbishment" — when the ratio is 15–30 % and the mechanical body (joint structure) remains in good condition, targeted overhaul of reducers or control system typically restores over 80 % of performance at roughly one-third of new-equipment cost. "Retirement Decision" — when the ratio exceeds 30 %, accuracy can no longer meet application requirements, or critical spares are unavailable, enter Phase 5 decommissioning assessment.**

3.6.3 Overhaul Procedure Highlights

- Reducer-Replacement Procedure: follow the OEM service manual for removal and refitting; after replacement, redo joint mastering and TCP-accuracy acceptance
- Control-System Refurbishment: after replacing control or drive boards, reload system software and backup parameters; confirm firmware-version compatibility of new parts
- Full Accuracy Re-Verification After Overhaul: perform accuracy measurements equivalent to those at commissioning; update the accuracy records in the equipment passport
- Safety-Function Re-Verification After Overhaul: any overhaul involving physical change to the safety circuit requires re-execution of the full safety-function verification checklist of §3.2.6
- Overhaul Document Archiving: the overhaul report (with replaced-part list and serial numbers) must be added to the equipment passport; update the configuration-baseline snapshot

3.7 Phase 5 · Decommissioning Assessment and Asset Disposal

Decommissioning is the statutory closure of the lifecycle, and its importance is often underestimated. Non-compliant decommissioning brings not only legal and environmental risk but also loss of service history and interruption of knowledge accumulation. Controlled decommissioning is a demonstration of professional service-organisation capability and a data source for the next-generation equipment installation plan.

3.7.1 Decommissioning Assessment Matrix

Table 3.7-1 RSF Decommissioning Assessment Matrix (Four-Dimension Scoring)

Assessment Dimension	Score 1–3 (Recommend Continue Operation)	Score 4–7 (Recommend Overhaul / Refurbish)	Score 8–10 (Recommend Retirement)
Availability Trend	Availability ≥ 95 % with stable trend	Availability 85–95 % with slow decline	Availability < 85 % with continued deterioration
Maintenance-Cost Ratio	Annualised < 10 %, dominated by planned maintenance	Annualised 10–25 %, mixed planned and unplanned	Annualised > 25 %, dominated by unplanned maintenance
Spare-Parts Availability	Full OEM spare-parts supply available	Some critical spares with extended lead time or EOL warning	Critical spares declared EOL with no substitute available
Task Suitability	Fully meets current and planned task requirements	Requires downgraded use or supplementary equipment	Cannot meet task requirements; constrains line capability

Composite score ≤ 12: recommend continue operation; 13–24: combined judgement recommended; 25–40: recommend retirement. The final retirement decision must be confirmed jointly by a Specialist-or-above engineer and customer management.

3.7.2 Controlled Decommissioning Procedure

Decommissioning must be executed strictly in the following order; no omission or reordering is

permitted:

- Data Backup: full backup of controller programs, parameter files, and system image (duplicate, off-site stored); archived to the equipment passport
- Full LOTO Execution: main power disconnected and locked out; confirm full discharge of high-voltage capacitors (verify with voltmeter; do not rely on estimated discharge time); pneumatic / hydraulic sources de-pressurised
- Hazardous-Material Identification and Disposal: grease (waste industrial grease as hazardous waste); controller lithium batteries (separately collected and handed to a licensed hazardous-waste processor); capacitors; lead-soldered electronic boards (per WEEE Directive)
- Secure Data Erasure: factory-reset controller programs and data; confirm full removal of user data; log the operator and time of erasure
- Decommissioning Report: record equipment ID, retirement reason, retirement date, disposition route (resale / parts recovery / scrap), responsible-engineer signature
- Sealed Archive of Equipment Passport: the decommissioned-equipment passport must be fully archived; retention of at least 10 years is recommended (statutory requirement for traceability of safety incidents)

3.8 Service Documentation and Configuration Baseline System

The Configuration Baseline System is Domain 1's information-management backbone across the lifecycle. Its core logic: capture "snapshots" of the robot system's technical state at key state nodes, so that the system's state at any later point can be precisely compared with the most recent baseline snapshot — enabling rapid deviation detection, change-source localisation, and anomaly-evolution tracing.

3.8.1 Trigger Points for Configuration Baseline Snapshots

Table 3.8-1 Configuration-Baseline Snapshot — Trigger Nodes and Contents

Trigger Node	Timing	Snapshot Content	Significance Note
Installation-Completion Baseline	When the Phase 1 → 2 gate is passed	Full Initial Parameter Record + software versions + safety-parameter setpoints	The original reference baseline of the whole lifecycle; must not be overwritten
Commissioning-Completion Baseline	When the Phase 2 → 3 gate is passed	Accuracy-measurement report + per-axis current baselines + temperature / vibration baselines	The core reference for subsequent state comparisons; the starting point for judging accuracy deviation
Annual-PM Completion Baseline	After each annual PM	PM execution record + re-measured accuracy + re-measured currents + software-backup version	Time-series data points for tracking long-term performance drift
Post-Overhaul Baseline	After each Phase 4 overhaul	Replaced-part list (with serial numbers) + post-overhaul accuracy report + safety-function re-verification report	Records the post-rebirth new technical-state baseline of the system
Firmware / Software-Upgrade Baseline	After each major software-version upgrade	Versions before and after + upgrade date + regression-test conclusion	Change-control record for software change; reference basis for rollback
Redeployment	After equipment is	Pre-installation survey of the	New original baseline in the

Baseline	migrated from one workstation to another	new environment + TCP re-calibration result + safety-function re-verification	new environment; the prior baseline is archived but no longer used for comparison

3.8.2 Equipment Passport — Standard Format

The Equipment Passport is the robot system's "lifetime medical record" and must be maintained across the entire lifecycle. Its completeness directly affects diagnostic efficiency, resale value, and audit readiness. A minimum-standard equipment passport must contain the following layers:

- Basic-Information Layer: robot model, serial number, controller model / serial number, installation date and location, customer information, dealer / service-provider information
- Technical-State Layer: configuration-baseline snapshots by phase (in chronological order); all PM execution records (date / hours / replaced parts); major repair records (including spare-part serial-number updates)
- Statutory-Documents Layer: safety-function verification reports (permanently preserved); delivery-acceptance certificates; controlled-change records of major changes; incident reports of any safety event
- Knowledge-Accumulation Layer: non-standard fault records for this equipment (symptom, root cause, remediation); service-improvement records applied to this equipment

3.8.3 Work-Order System and Traceability Requirements

The work order is the smallest traceable unit of service activity and the primary data source for SLA calculation, performance analysis, and knowledge accumulation. Work-order quality is the foundation of Domain 4 (SLA management) and the continuous-improvement loop; low-quality work orders (missing key fields, inaccurate timestamps, vague disposition descriptions) collapse the entire service-data system.

Work-Order Field	Content Requirement	Why It Cannot Be Omitted
Timestamps	Fault-report time / response-start time / on-site arrival time / repair-completion time (minute precision)	Base data for SLA response time and MTTR; missing this field makes SLA calculation impossible
Work Type	PM / fault response / overhaul / installation / training (single-select)	Distinguishes planned vs unplanned; affects PM-compliance-rate calculation
Fault Description	Objective description of the fault (alarm codes + observed anomalous behaviour); no subjective judgement	The original observation is the starting point of RCA; subjective judgement replacing observation makes later traceability fail
Disposition Record	Operations actually performed (with parameter before / after values, spare-part replacement records)	The disposition record is the basis for updating the configuration baseline; missing it breaks equipment-state traceability
Executor	Service-engineer name + certification tier	Basis for audit trace and competence assessment; anonymous work orders lack traceability
Work-Order Closure Standard	Customer representative confirms restored functionality and signs off (engineer-only closure is prohibited)	Unilateral closure is a common source of SLA dispute

3.9 Domain 1 Key Performance Indicator Framework

Domain 1's KPI framework covers four dimensions: availability, maintenance efficiency, service quality, and asset management. The meaning of an indicator lies in trend analysis, not a single-point reading: a one-time 98 % availability does not mean the service system is healthy; a sustained drop from 99 % to 96 % may be an early signal of system degradation.

Table 3.9-1 Domain 1 Key Performance Indicator Framework

Indicator	Formula	Industry Benchmark Target	World-Class Target	Management Significance
Equipment Availability	(scheduled hours − unplanned downtime) / scheduled hours	≥ 95%	≥ 99%	Core SLA indicator; below 95 % typically triggers SLA-compensation clauses
OEE (Overall Equipment Effectiveness)	Availability × Performance × Quality	≥ 75%	≥ 85%	Composite indicator watched by manufacturing leadership; 85 % is the world-class manufacturing benchmark
MTTR (Mean Time to Repair)	Total repair time / number of repairs	< 4 hours	< 2 hours	Directly reflects response and diagnostic efficiency; the design basis for SLA response commitments
MTBF (Mean Time Between Failures)	Operating time / number of failures	Stable trend	Continuous improvement	Absolute value varies by application; a declining trend signals entry into Phase 4
PM Compliance Rate	PM executions as planned / total planned PMs	≥ 90%	≥ 95%	Measures PM discipline; below 80 % portends a rise in unplanned downtime
Planned-to-Unplanned Maintenance Ratio	Planned maintenance hours / total maintenance hours	≥ 0.75	≥ 0.90	A high ratio reflects a proactive maintenance culture; below 0.3 indicates "fire-fighting mode"
Emergency Procurement Ratio (EPR)	Emergency procurements / total procurements	≤ 10%	≤ 5%	A high EPR signals insufficient safety stock or failed PM forecasting

3.10 Chapter Summary and Inter-Domain Interfaces

Chapter Summary Domain 1 (Service Lifecycle Management) organises all robot-service activity into a value-delivery sequence through the five-phase model: Delivery and Installation lays the technical baseline; Run-in and Commissioning establishes the performance reference; Preventive Maintenance converts risk into a budgetable cost; Corrective Maintenance and Overhaul extends asset value; Controlled Decommissioning closes the lifecycle and frees subsequent capital. Phase transitions are driven by stage-gate conditions so that service quality — not calendar progress — is the principal basis for advancement. Spanning the

> lifecycle, the three-tier closed loop of Incident Response → RCA → Service Improvement converts single fault events into systemic improvement inputs, continuously refining PM strategy and service procedures. The Configuration Baseline System and the work-order system together provide the traceable information foundation for the entire framework.

Interfaces with Other Capability Domains

Interface Direction	Interfacing Domain	Content Provided by / to Domain 1
Domain 1 → Domain 2	Fault Diagnosis & Technical Support	Lifecycle-phase context (the same symptom has different diagnostic paths in commissioning vs stable operation); per-phase technical baselines (current / accuracy / vibration baselines used by Domain 2 for anomaly judgement)
Domain 1 → Domain 3	Remote Service & Digital Operations	Phase 3's PM data stream (sensor data, operating hours) is the principal data source for Domain 3's predictive-maintenance models; configuration-baseline snapshots provide the parameter baseline for the digital-twin model
Domain 1 → Domain 4	SLA Design & Service Commitment	Domain 1's KPIs (MTTR / availability) are the technical-design basis for SLA clauses; PM history is the cost-analysis basis for SLA pricing
Domain 1 ← Domain 5	Cross-Platform Service Procedures	Domain 5 supplies platform-specific service procedures that overlay Domain 1's generic five-phase framework (e.g., cobot run-in torque-parameter verification procedures are added by Domain 5)

Chapter 4

Domain 2 · Fault Diagnosis & Technical Support

4.0 Chapter Introduction

Modern industrial robots are tightly integrated systems combining mechanical, electrical, software, and AI layers. This architectural property fundamentally changes the nature of fault diagnosis: the traditional "swap parts and test" logic of equipment repair fails systemically against a four-layer integrated system, because an anomaly in any one layer can produce symptoms in any other, and the layer where the symptom appears is often not the layer where the root cause lies. An input-distribution shift in an AI perception model can disguise itself as a trajectory anomaly; a poor electrical ground can manifest as a software communication timeout; early reducer wear can be "compensated for" by the position loop until one day a large-error alarm suddenly fires.

Domain 2 (Fault Diagnosis & Technical Support) establishes RSF's stratified diagnostic capability system, taking the four-layer architecture as its principal axis and converting diagnostic work from experience-driven "guess and check" into logic-driven "systematic localisation." The chapter's focus is methodology — the stratified framework, the diagnostic process, the tool-capability requirements — not the operations manual of any specific brand. The methods described apply to any brand and any type of industrial robot system.

Chapter Structure: §4.1 establishes the four-layer stratified diagnostic framework; §4.2 defines the diagnostic-tool capability system; §§4.3–4.6 expand the diagnostic methods of each layer; §4.7 addresses intermittent faults; §4.8 defines the complete fault-handling workflow; §4.9 provides a typical case library; §4.10 defines escalation and knowledge-precipitation mechanisms; §4.11 establishes the Domain 2 performance-metrics system.

4.1 RSF Four-Layer Stratified Diagnostic Framework

RSF's fault-diagnosis framework takes the robot system's four-layer technical architecture as its main axis to establish a systematic diagnostic logic. The core significance of the four-layer architecture is "search-space definition" — before any diagnostic operation begins, the primary task is to narrow the possible root-cause space to a specific technical layer and then localise precisely within that layer, rather than searching the whole system indiscriminately.

4.1.1 Four-Layer Architecture and Inter-Layer Fault Propagation

Table 4.1-1 Robot Four-Layer Technical Architecture and Diagnostic Characteristics

Technical Layer	Layer Content	Typical Fault Modes	Diagnostic Characteristics	Required Tool-Capability Type
L1 Mechanical Body Layer	Joints, reducers (RV / harmonic), bearings, end-effector, chassis structure	Backlash increase, accuracy degradation, abnormal noise, lubrication failure, mechanical wear	Physically visible and measurable; symptoms persist; strongly correlated with operating hours	Vibration analysis, precision displacement measurement, thermal imaging, oil analysis
L2 Electrical Control Layer	Servo drives, safety controllers, power systems, sensors,	Servo alarms, communication timeouts, safety-	Usually accompanied by an alarm code; may appear	Servo-waveform acquisition, insulation testing,

	communication interfaces	circuit false triggers, sensor drift	intermittently; strongly heat-correlated	thermal imaging, protocol analysis
L3 Software Middleware Layer	Controller firmware, motion controller, OS middleware, communication protocols, coordinate-frame management	Parameter anomalies, coordinate-frame drift, firmware-compatibility errors, accumulated log errors	Usually logged; may transiently disappear after reboot; version-dependent	System-log analysis, parameter-backup diff, offline-simulation verification
L4 AI Decision Layer	Vision-perception models, motion-planning networks, reinforcement-learning policies, multi-modal fusion inference	Perception-accuracy degradation, decision drift, out-of-distribution failure, generalisation failure	Probabilistic failure; no traditional alarm codes; strongly correlated with input-data distribution	Inference-confidence analysis, benchmark testing, model-performance evaluation (see Domain 6)

Inter-Layer Fault Propagation — Two Typical Misdirection Patterns

Propagation Pattern	Typical Scenario	Misdirection and Detection
Upward Propagation (lower-layer cause → upper-layer symptom)	Reducer wear (L1) → abnormal drive current (L2) → trajectory-error alarm (L3)	Symptom at L3, root cause at L1. Detection: after ruling out software causes for the L3 alarm, trace downward to L2 current state and then L1 mechanical state
Downward Misdirection (upper-layer cause looks like a lower-layer fault)	Vision-model distribution shift (L4) → grasp-position deviation → misjudged as TCP calibration problem (L1)	Symptom at L1, root cause at L4. Detection: problem persists after TCP re-calibration; strongly correlated with specific workpiece or lighting; disappears with a standard reference workpiece

> **Core Diagnostic Principle** Before any physical action (part replacement, parameter change), layer localisation must be completed first. Replacing parts is not diagnosing — replacing a part before the fault layer is confirmed is the most common waste of time and cost. Layer localisation costs time (minutes to hours); a wrong replacement costs money (parts) and credibility (the fault remains).

4.1.2 Diagnostic Funnel Model

The RSF Diagnostic Funnel describes the diagnostic process as a systematic narrowing from broad to precise, in four progressive steps. Each step has explicit inputs and outputs so that the diagnostic process can be recorded, audited, and taught.

Step	Name	Operations and Output	Key Cautions
Step 1	**Symptom Classification and Preliminary Layer Localisation**	Collect fault phenomena (alarm codes + observed behaviour + operator description); against the four-layer symptom table in §4.1.3, complete preliminary layer localisation. Output: "This fault is preliminarily judged as L_ layer."	Classify strictly by symptom; do not skip this step on intuitive experience
Step 2	**In-Layer Precise Localisation**	Within the target layer, systematically investigate every possible root cause in that layer using its diagnostic tool capabilities, narrowing to a specific component or functional unit. Output: "The root	In-layer investigation must systematically cover all possible causes — avoid the cognitive bias of "test the

		cause is most likely ____."	easiest thing first"
Step 3	**Cross-Layer Verification**	Confirm that the located root cause fully explains every observed symptom. If any symptom is unexplained, consider a cross-layer cause or a compound-fault scenario. Output: "Root cause confirmed" or "Extend investigation to an adjacent layer."	Compound faults (problems present in multiple layers simultaneously) make up about 20 % of complex faults; check for them deliberately
Step 4	**Remediation-Path Selection**	Based on the confirmed root cause, choose a path: ① Workaround — restore production while awaiting a permanent fix; ② Permanent Fix — solve the root cause; ③ Escalate — follow the escalation procedure in §4.10	Workaround and permanent fix must be explicitly distinguished; a workaround may not substitute for the permanent fix (see §4.8 for details)

4.1.3 Fault Symptom Classification and Quick-Lookup for Initial Layer Localisation

Symptom Description	Preliminarily Indicated Layer	Confounding Layer to Rule Out	Recommended First Diagnostic Action
TCP repeatability exceeds baseline; trajectory jitter	**L1 Mechanical Layer**	L3 (coordinate-frame parameter error)	Joint backlash measurement; compare drive currents to baseline
Abnormal joint noise (metallic friction / periodic clack)	**L1 Mechanical Layer**	L2 (brake not fully released)	Vibration-frequency analysis; grease-state inspection
Servo-class alarm (position error / current limit)	**L2 Electrical Control Layer**	L1 (overload triggering servo protection)	Review alarm history; capture servo-current waveform
Industrial-Ethernet communication-timeout alarm	**L2 / L3 boundary**	L1 (physical cable damage)	Bus-topology diagnosis; protocol-layer packet capture
Safety-circuit false trigger (no obvious source)	**L2 Electrical Control Layer**	L3 (Safety-logic configuration error)	Segment-by-segment safety-circuit test; thermal-imaging inspection
Problem disappears after reboot; intermittent occurrence	**L3 Software Layer (first choice)**	L2 (heat-related electrical fault)	Time-series analysis of system logs
Specific program fails; other programs run normally	**L3 Software Layer**	L2 (electrical issue on a specific axis)	Program verification; coordinate-frame integrity check
Vision-guided task fails under specific conditions	**L4 AI Layer (first choice)**	L2 (camera hardware / light-source fault)	Three-step hardware exclusion; substitute a standard test workpiece
Grasp-position deviation while TCP verification is normal	**L4 AI Layer**	L1 (small deviation from mechanical backlash)	AI-model benchmark test; check input-image quality
Accuracy drops after hours of operation; cold start normal	**L1 (thermal expansion) + L2 (thermal drift)**	L3 (no temperature compensation)	Continuous-operation temperature monitoring; full-robot thermal-imaging scan

4.1.4 Fault Priority Classification (P1–P4)

Priority	Definition	Typical Scenario	Response-Time Target	Handling Requirements
P1	**Safety fault or complete**	Safety-function failure; uncontrolled	Remote: 15 minutes; on-site: 2 hours	Initiate major-incident response; Specialist-led; escalated to service

	production stoppage	hazardous motion; complete production stop		management within 1 hour; all operations require a second person on site for confirmation
P2	**Functional fault severely affecting production**	Loss of principal function; availability below SLA threshold; no effective workaround	Remote: 1 hour; on-site: 4 hours	Prioritise temporary restoration to minimise downtime; simultaneously initiate RCA; Specialist-led
P3	**Performance-degradation fault**	Accuracy declines but within spec; cycle time lengthens; partial function degradation	Same-day response; planned repair	May be addressed in the next planned-maintenance window; Specialist reviews the remediation plan
P4	**Warning and data anomaly**	Trend warning (slowly rising current); non-critical alarms; accumulated log errors	Response within the week; incorporated into the PM plan	Record and track the trend; address in the next PM

4.2 Diagnostic Tool Capability Framework

RSF defines diagnostic tools by capability rather than by product: the framework specifies what capability the service engineer needs in each diagnostic task, not which brand of tool to use. This principle keeps the framework neutral with respect to all robot brands and tool suppliers, while giving engineers a clear capability-development target.

Seven categories of diagnostic-tool capability provide complete coverage of the four-layer architecture's diagnostic needs. Each capability can be realised through the robot maker's bundled tools, generic third-party instruments, or specialised measurement equipment; the specific tool selection is made by the service organisation based on the installed-base brands and business scale.

4.2.1 Definitions of the Seven Diagnostic Tool Capabilities

Capability ① Servo Data Acquisition & Waveform Analysis

Tool Definition	Real-time acquisition of position error, velocity, current, and torque data from the servo system, presented as waveforms, with support for overlay comparison against historical baselines.
Diagnostic Value	Servo waveforms are the core data source for L2 electrical-layer diagnosis, revealing mechanical wear (current anomalies), parameter-configuration issues (waveform-convergence characteristics), and sensor faults (feedback-signal jitter) — diagnostic dimensions invisible to the naked eye and to alarm codes.
Applicable Scenarios	Diagnosis of servo-class alarms; electrical-layer verification of accuracy degradation; early quantitative assessment of reducer wear; gain-parameter tuning.
Tier Requirements	Required for Specialist; Professional must be able to read basic waveform reports.
Corresponding Diagnostic Layer	L1 mechanical layer (current trend); L2 electrical layer (primary)

Capability ② System Event Log & Alarm History Analysis

Tool Definition	Read, filter, and analyse the system event logs and historical alarm records produced by the controller, with support for time-series correlation analysis and cross-type event comparison.

Diagnostic Value	System logs are the primary data source for L3 software-layer diagnosis, and the only reliable information source for diagnosing intermittent faults that cannot be reproduced on demand. Time-series analysis of alarm history reveals the triggering conditions and regularities of a fault — the key bridge from phenomenon to root cause.
Applicable Scenarios	Software-layer fault diagnosis; regularity recognition in intermittent faults; alarm-density trend tracking; pre-fault event-sequence analysis.
Tier Requirements	Professional must be able to access and read basic alarm records; Specialist must be able to perform full time-series correlation analysis.
Corresponding Diagnostic Layer	L3 software layer (primary); L2 electrical layer (auxiliary)

Capability ③ Offline Simulation & Program Verification

Tool Definition	Verify robot programs, detect collisions, and analyse reachability on a virtual controller without occupying the real robot; supports diff-based comparison between the current controller state and a known-good backup.
Diagnostic Value	L3 software-layer parameter and program issues can be safely reproduced and verified offline without the downtime risk of repeated tests on the real robot. Backup diff is the most effective method for identifying "unauthorised parameter modification" — a hidden category of fault.
Applicable Scenarios	Safety verification after parameter changes; offline reproduction and analysis of collision positions; regression testing for program-version updates; checks for accidental parameter modification.
Tier Requirements	Required for Specialist; Expert must be able to design offline test scenarios.
Corresponding Diagnostic Layer	L3 software layer (primary)

Capability ④ Remote Condition Monitoring & Data Acquisition

Tool Definition	Continuously monitor robot-system state over a network connection, collecting time-series data for key parameters (current, temperature, operating hours, alarm frequency), with threshold alarming and trend reporting.
Diagnostic Value	Continuous state monitoring advances diagnostics from "respond after the fact" to "warn before the fact" — the technical foundation of condition-based maintenance (CBM). Remote data acquisition is also a key tool for intermittent-fault diagnosis: background recording while no fault is present ensures that data is automatically captured when the next fault occurs.
Applicable Scenarios	Predictive maintenance of critical equipment; background data capture for intermittent faults; centralised state monitoring across multiple machines; automated calculation of KPIs such as OEE and MTBF.
Tier Requirements	Specialist must be able to deploy and configure monitoring rules; Professional must be able to read monitoring reports and alerts.
Corresponding Diagnostic Layer	L1–L3 (cross-layer data aggregation)

Capability ⑤ Precision Physical Measurement

Tool Definition	Precision measurement of physical quantities on a robot system, covering: mechanical (displacement, clearance, torque); electrical (voltage, current, insulation resistance, waveforms); thermal (temperature distribution, hotspot localisation); vibration (frequency, amplitude, spectrum); optical (illuminance, position accuracy).
Diagnostic Value	Physical measurement is the "confirmation means" of L1 mechanical and L2 electrical diagnosis — after layer localisation, precision measurement converts the qualitative

	judgement into quantitative confirmation (e.g., suspect reducer wear → measure backlash → compare to commissioning baseline → confirm extent of wear). Diagnostic conclusions without quantitative data are not traceable.
Applicable Scenarios	Quantitative confirmation of mechanical-layer accuracy degradation; assessment of electrical-insulation condition; temperature-distribution analysis of heat-related faults; vibration-signature assessment of reducer condition; absolute TCP-accuracy measurement.
Tier Requirements	Professional needs basic electrical measurement (multimeter, clamp meter); Specialist needs thermal-imaging and vibration analysis; Expert needs high-precision optical measurement.
Corresponding Diagnostic Layer	L1 Mechanical Layer (primary); L2 Electrical Layer (secondary)

Capability ⑥ Industrial Communication Protocol Analysis

Tool Definition	Capture and parse communication packets on industrial-Ethernet networks (EtherCAT, PROFINET, EtherNet/IP, etc.) and analyse cycle stability, packet integrity, slave-state transitions, and protocol-level errors.
Diagnostic Value	Industrial-Ethernet faults have a complex mapping between symptoms (alarm codes) and root causes (physical / configuration / timing layers); the alarm code alone cannot guide the diagnostic path. Protocol-layer analysis directly observes the mechanism of the anomaly and is indispensable for precise localisation of L2 / L3 communication faults.
Applicable Scenarios	Industrial-Ethernet timeout / disconnect faults; in-depth analysis of slave-state anomalies; data capture for intermittent communication faults; quantitative verification of EMC interference.
Tier Requirements	Required for Specialist; protocol analysis requires protocol-specific specialised training.
Corresponding Diagnostic Layer	L2 / L3 boundary (communication layer)

Capability ⑦ AI Inference Diagnostics & Model-Performance Assessment

Tool Definition	Monitor and analyse the inference process of AI / ML components (vision perception, motion planning, reinforcement-learning policy), including: real-time inference-confidence tracking; input-data quality assessment; model-performance benchmark testing; anomaly-pattern recognition in inference logs.
Diagnostic Value	AI-layer faults do not produce traditional alarm codes; their failure modes (distribution shift, generalisation failure, decision drift) cannot be identified with mechanical or electrical tools. Specialised AI diagnostic capability is the precondition for accurately distinguishing "AI-layer problem" from "hardware problem" and for avoiding misdiagnosis of AI faults as mechanical or calibration issues.
Applicable Scenarios	Decline in success rate of vision-guided tasks; inconsistent robot behaviour under specific conditions; AI-performance verification for new workpieces or scenes; performance comparison before and after model updates.
Tier Requirements	Expert tier requires basic AI-diagnostic capability; deep AI-layer diagnosis is treated in Chapter 8 (Domain 6).
Corresponding Diagnostic Layer	L4 AI decision layer (dedicated)

4.2.2 Mapping the Seven Tool Capabilities to the Four-Layer Architecture

Table 4.2-1 Diagnostic-Tool Capability × Four-Layer Architecture Matrix

Diagnostic-Tool Capability	L1 Mechanical	L2 Electrical	L3 Software	L4 AI

	Body Layer	Control Layer	Middleware Layer	Decision Layer
① Servo Data Acquisition & Waveform Analysis	◑ Auxiliary	● Primary	○	○
② System Event Log & Alarm History Analysis	○	◑ Auxiliary	● Primary	○
③ Offline Simulation & Program Verification	○	○	● Primary	○
④ Remote Condition Monitoring & Data Acquisition	◑ Auxiliary	◑ Auxiliary	◑ Auxiliary	◑ Auxiliary
⑤ Precision Physical Measurement	● Primary	◑ Auxiliary	○	○
⑥ Industrial Communication Protocol Analysis	○	◑ Auxiliary	● Primary	○
⑦ AI Inference Diagnostics & Model-Performance Assessment	○	○	○	● Primary

● Primary tool ◑ Auxiliary / supplementary tool ○ Not applicable

4.2.3 Tool-Capability Configuration Requirements by Certification Tier

RSF defines tool-capability requirements stratified by certification tier. "Possessing the capability" means the engineer can operate the tool independently and correctly interpret the output, not merely follow the operating steps mechanically.

Diagnostic-Tool Capability	Professional (minimum capability to attain)	Specialist (must fully possess)	Expert (extended capability)
① Servo Data Acquisition & Waveform Analysis	Can read basic waveform reports and identify obvious anomalies	Can independently capture and analyse; can judge by comparison to baseline	Can design custom test programs; can mentor others in analysis
② System Event Log & Alarm History Analysis	Can access and read alarm history; can look up alarm codes in the manual	Can perform time-series correlation analysis; can identify regularities	Can establish organisation-level log-analysis standards
③ Offline Simulation & Program Verification	Understands the purpose and basic operation of offline verification	Can independently perform program verification and backup diff	Can design complex test scenarios; can assess system integration
④ Remote Condition Monitoring & Data Acquisition	Can read monitoring reports and alert notifications	Can configure monitoring rules and alert thresholds	Can design the monitoring architecture; can interpret predictive-maintenance models
⑤ Precision Physical Measurement	Basic electrical measurement (voltage / current / insulation resistance)	Thermal-imaging analysis; vibration-spectrum analysis; precision-displacement measurement	High-precision optical measurement (laser tracking); oil analysis

⑥ Industrial Communication Protocol Analysis	Understands how to read bus-topology status	Can perform protocol-layer packet capture and basic analysis	Can diagnose complex communication-architecture problems; EMC analysis
⑦ AI Inference Diagnostics & Model-Performance Assessment	Can recognise surface symptoms of AI-layer faults	Can execute the three-step hardware exclusion; can compile an escalation package	Can execute benchmark testing; can evaluate model-performance indicators

> **Note on Brand Tools** Robot manufacturers typically bundle diagnostic software with the controller, covering the core functions of capabilities ① ② ③ ④ above; specialised third-party instruments (vibration analysers, thermal cameras, protocol analysers) provide the precision measurement of capabilities ⑤ ⑥. Once engineers internalise the seven capability definitions, they can select specific tools that realise these capabilities based on their installed-base brands and service scenarios. RSF makes no recommendation or exclusion of specific tool brands.

4.3 L1 · Mechanical Body Layer Diagnosis

Mechanical-layer diagnosis is characterised by "physical measurability": symptoms typically persist and can be quantified by physical quantities (clearance, accuracy, current, temperature, vibration) and compared against installation and commissioning baselines. The biggest challenge of mechanical-layer diagnosis is not method but "baseline-data completeness" — without commissioning-period baseline records, quantitative comparison has no reference frame.

4.3.1 Mechanical-Layer Fault Quick-Recognition Guide

Fault Category	Symptom Characteristics	High-Probability Root Cause	Preferred Diagnostic Action
Accuracy Degradation	TCP repeatability exceeds 150 % of commissioning baseline; trajectory straightness declines	Reducer-backlash increase; TCP-calibration drift	Joint-backlash measurement (against commissioning baseline); TCP recalibration test
Abnormal Noise (continuous / periodic)	Metallic friction sound; periodic impact noise; heavy abnormal noise at start-up	Reducer tooth-flank wear; bearing-race damage; grease depletion	Vibration-frequency analysis (Capability ⑤); visual + tactile grease inspection
Leakage	Grease seepage at joint; oil traces at the reducer output end	Aged or damaged reducer seal; previous regrease over-filled	Visually inspect the leakage location; check the last regrease record
Thermal Drift	Cold-start accuracy normal; declines after continuous	Joint-temperature rise causing thermal expansion;	Thermal imaging before and after continuous operation;

	operation	grease viscosity drops at high temperature	correlate temperature rise with accuracy drift
Load Sensitivity	Normal under no load; accuracy drops or alarms appear under load	Incorrect load-inertia parameter setting; insufficient reducer stiffness	Check the load-inertia value in the IPR; re-run load identification

4.3.2 Reducer Diagnostic Methods (RV / Harmonic Reducers)

Reducers are the most expensive parts in the mechanical layer and the strongest determinant of overall robot life — and the primary source of accuracy degradation. Systematic assessment of reducer condition must combine multiple diagnostic means:

- **Joint Backlash Measurement**

  With no drive torque applied, move the joint a small amount forward and backward and measure the difference in the position reached. Tool: dial indicator (0.001 mm precision); the servo-waveform tool can assist by capturing the position-error waveform at the direction reversal. Criteria: an increase over the commissioning baseline by more than 50 % is a warning; more than 100 % suggests replacement.

- **Drive-Current Trend Analysis**

  A drop in reducer efficiency directly raises the current needed to drive the same load. At quarterly PM, with a standard test program (defined trajectory and speed), record each axis's current and compare to the commissioning baseline. Tool: servo data-acquisition tool (Capability ①). Criteria: under equivalent conditions, a current rise above 20 % of baseline is a warning; above 40 % requires combined assessment with backlash measurement.

- **Vibration-Frequency Analysis**

  Reducer wear produces vibration signals at characteristic frequencies early in the wear process, before visible symptoms or accuracy degradation appear. Tool: vibration analyser (Capability ⑤) with accelerometer mounted on the joint housing to acquire spectral data. Note: vibration diagnosis must be referenced to the commissioning vibration baseline — establish the baseline file at first measurement.

- **Grease-State Inspection**

  At annual PM, sample the grease and perform visual inspection and trend analysis. Visual condition assessment: normal grease is uniform beige or light-coloured (per OEM model spec); heavy iron contamination shows grey-black; oxidative aging shows dark brown, oil separation, or caked hardening. Quantitative testing: primarily visual; when feasible, use a portable oil analyser or send a sample to a laboratory for iron spectral analysis (Capability ⑤). Iron-particle concentration (Fe) warning thresholds (based on mainstream reducer industry practice): Normal / break-in accumulation:

< 1,000 ppm — normal operating wear accumulation, continue monitoring. Caution / Warning: ≥ 1,000 ppm (or a doubling trend vs the prior test), indicating possible abnormal wear, overload, or lubrication deterioration — shorten PM interval and monitor closely. Abnormal / Immediate action (Critical): ≥ 2,500 ppm (some OEM red lines are 5,000 ppm), indicating destructive wear of internal gears or bearings — stop operation immediately, arrange grease flush, overhaul, or reducer replacement.

4.3.3 Brake and End-Effector Diagnosis

Brake-holding-torque test: with each joint powered off and locked, use a torque wrench to measure each axis's actual static brake-holding torque and compare to the OEM rated value. Pass: measured torque ≥ OEM specified value. Critical action threshold: any axis whose measured torque falls below the OEM specified value, or any gravity axis that shows any visually perceptible physical drift, requires immediate shutdown (gravity-axis brake failure will cause catastrophic equipment drop or personnel injury).

End-Effector (EOAT) Diagnostic Note: with the end-effector removed, re-measure TCP accuracy under the standard program — if accuracy is restored, the fault originates in the end-effector (mounting-flange accuracy, tool wear, quick-change coupler wear) rather than the robot body, preventing misattribution of EOAT faults to the robot's accuracy.

4.4 L2 · Electrical Control Layer Diagnosis

The core insight of electrical-layer diagnosis: "alarm codes give clues, not answers" — the alarm points to a symptom, and the root cause must be confirmed by systematic measurement. The most common misstep is treating the alarm hint as the root cause. For example, a "servo overload" alarm may come from mechanical wear increasing resistance, from incorrect parameter setting, or from motor degradation — three root causes, one alarm code, three completely different remediations.

4.4.1 Servo System Diagnosis

Servo Waveform Analysis (Core Method)

The servo data waveform is the most powerful tool in electrical-layer diagnosis (Capability ①), simultaneously displaying position error, velocity ripple, current ripple, and torque output, revealing anomalies the mechanical layer cannot directly observe.

Waveform Type	Normal Characteristics	Anomalous Characteristics	Possible Root Causes

Position-error waveform	Converges to zero rapidly after motion ends (typically < 20 ms); no sustained oscillation	Slow convergence; persistent small oscillation after settling	Improper gains; reduced mechanical stiffness; incorrect load-inertia parameter
Current waveform	Smooth at low constant-velocity motion; symmetric peaks during acceleration and deceleration	Periodic current ripple at low speed; asymmetric peaks	Reducer wear (non-uniform friction); motor-winding degradation
Velocity waveform	Commanded and actual velocity curves overlap closely	Actual velocity lags command; step appears in velocity waveform	Poor encoder-signal quality; increased communication latency
Torque waveform	Torque peaks during acceleration / deceleration match the load	No-load torque exceeds baseline by 20 %	Increased mechanical friction (reducer / bearing wear); lubrication failure

Required Capability: ① Servo Data Acquisition & Waveform Analysis — realised through the robot's bundled diagnostic software or third-party data-acquisition instruments.

Motor Electrical-Integrity Tests

- Motor Insulation-Resistance Test: using an insulation-resistance tester (megohmmeter), with the motor powered off and disconnected from the drive, measure each phase winding-to-case insulation resistance (apply 500 V DC test voltage). Quantitative assessment criteria (aligned with robot servo-system standards): Excellent / Normal: ≥ 20 MΩ (new or well-maintained motors typically ≥ 100 MΩ); no electrical-safety or signal-interference risk. Caution / Warning: 2 MΩ ≤ R < 20 MΩ — possible internal moisture, carbon buildup, or mild insulation aging; increase monitoring frequency and inspect motor sealing. Abnormal / Immediate action (Critical): < 2 MΩ (or per OEM absolute red line) — high risk of drive overcurrent alarm or encoder signal interference; stop operation immediately and arrange motor drying, re-impregnation, or motor replacement. Tool: insulation-resistance tester (Capability ⑤)
- Three-Phase Current Balance: measure each phase current separately with a clamp ammeter. Pass: imbalance rate ≤ 3% — torque output stable, no abnormal vibration. Caution / Warning: 3% < imbalance ≤ 5% — possible mild hardware degradation, encoder feedback drift, or inconsistent contact resistance in wiring; increase monitoring and inspect terminal connections. Abnormal / Immediate action (Critical): imbalance > 5% (or triggering a hardware-level overcurrent alarm) — typically indicates inter-turn winding short, partial demagnetisation of permanent magnets, or inverter module hardware failure; take offline immediately and do not run in a degraded state to avoid burning the drive controller. Tool: clamp ammeter (Capability ⑤)
- Encoder-Signal Quality: observe encoder-feedback stability via the servo-diagnostic tool; count jumps (non-monotonic increments) are typical of encoder or cable faults
- Controller Power Quality Test — test method: while the robot executes a high-load, multi-axis simultaneous acceleration programme, use the peak-hold or oscilloscope function of a digital

multimeter (Capability ⑤) to measure dynamic voltage sag on both the incoming three-phase AC supply and the internal 24 V DC logic bus. Quantitative best-practice thresholds: Three-phase AC power supply (drive layer): Excellent: voltage sag ≤ 3% at maximum load. Critical threshold: sag > 5% — readily triggers servo-drive DC bus undervoltage alarm; inspect shop supply cable gauge, transformer capacity, or install a voltage regulator. 24 V DC control supply (logic and safety layer): Excellent: voltage sag ≤ 3% at maximum load (minimum voltage not below 23.3 V), ensuring absolute stability of core boards and safety circuits. Warning: 3% < sag ≤ 5% — within component tolerance but indicates insufficient SMPS headroom or high contact resistance in 24 V distribution; risk of intermittent communication packet loss; monitor closely and inspect wiring. Abnormal / Immediate action (Critical): sag > 5% (minimum voltage below 22.8 V) — can trigger spurious system resets or false safety alarms; take offline immediately for rectification.

4.4.2 Communication Network Diagnosis

Industrial-Ethernet faults (EtherCAT / PROFINET / EtherNet/IP) typically appear as alarm codes, but root causes are highly diverse, including physical-layer issues (cables / connectors), configuration issues (address conflict / duplicate device names), and timing issues (mismatched cycles). Systematic diagnosis must proceed from physical layer upward.

- **Physical-Layer Check (Step 1)**

  Inspect link-status indicators at every node; verify shield continuity at cable-bend segments (measure with a multimeter — shield resistance should be < 1 Ω and unaffected by bending); check connector contact quality. Tool: digital multimeter (Capability ⑤). Cable-shield problems are the most common physical-layer root cause of intermittent communication faults, accounting for roughly 40 % of communication-class intermittent faults.

- **Configuration-Layer Check (Step 2)**

  Verify that bus-topology configuration matches the physical connection order; verify uniqueness of all addresses on the network (IP-address or device-name conflicts cause periodic anomalies); verify slave-firmware version compatibility with the master. Tool: bus-configuration tool (Capability ⑥).

- **Timing-Layer Analysis (Step 3)**

  If the fault is periodic, use a protocol analyser to capture packets during the fault and analyse cycle jitter; cycle jitter > 10 % requires master scan-cycle reconfiguration. Tool: industrial communication-protocol analyser (Capability ⑥).

4.4.3 Safety Circuit Diagnosis

> **Special Handling for Safety-Class Faults Safety-circuit diagnosis is P1. The process must meet: at least two authorised engineers on site (one operating, one supervising); all operations performed under LOTO; diagnostic records signed by both; any safety-function test conclusion must be verified in the presence of the customer engineer. This is an ISO 10218 compliance requirement, not a procedural suggestion.**

- Safety-Circuit Segment-by-Segment Isolation: read safety-function state segment by segment via the safety controller's diagnostic interface and identify which safety signal is in the triggered state
- False-Trigger Judgement: measure the safety-input terminal voltage with an oscilloscope; a genuine trigger is a complete pulse; millisecond-scale narrow pulses typically arise from EMC interference and indicate grounding or shielding issues
- Safety-Controller Logs: modern safety controllers retain detailed safety-event logs with millisecond-precision trigger times and sources — the preferred information source for intermittent safety-fault diagnosis

4.5 L3 · Software Middleware Layer Diagnosis

The defining feature of software-layer faults is that "the symptom may temporarily disappear after reboot" — both a diagnostic challenge (the problem is hard to reproduce) and a diagnostic clue (post-reboot symptom disappearance almost certainly points to a software or memory-class problem). The core tool is the system event log (Capability ②); the core method is time-series correlation analysis.

4.5.1 Controller Log Analysis Methodology

System-log analysis must follow these steps and should not stop at simply searching for an alarm code:

- Time Localisation: pinpoint the exact timestamp of the fault (minute precision); use it as the anchor to trace backward through the 15–30 minutes of preceding events
- First-Occurrence Identification: find the first occurrence of this fault type in the log; distinguish first occurrence from later repeats; root-cause analysis should focus on the first-occurrence context
- Associated-Event Identification: what were the last events before the fault? Was there a parameter change? A software update? Alarms on other equipment?
- Density Trend Analysis: track the frequency trend of similar alarms; sustained increases typically signal that the root cause is worsening
- Cross-System Time Alignment: align the controller-log times with production records and equipment-operation logs to identify external triggers

Required Capability: ② System Event Log & Alarm History Analysis — all controller brands provide log-access interfaces; log files can typically be exported as CSV for external analysis.

4.5.2 Parameter and Coordinate-System Integrity Verification

Software-Layer Hidden Faults: parameter or coordinate-frame modifications produce no alarm; the symptom is accuracy drift or anomalous behaviour. The investigation method is item-by-item comparison against the Initial Parameter Record established at installation.

- TCP-Frame Verification: re-measure the current tool centre point (TCP) values with the standard verification programme, compare item by item to the IPR, and calculate the spatial deviation. Universal red-line principle: actual TCP deviation must not exceed 2–3 times the robot's officially rated repeatability. Application-level quantitative thresholds (for rapid screening): High-precision / vision-guided applications (3C assembly / semiconductor / fine collaborative tasks): deviation > 0.1 mm–0.2 mm requires immediate investigation — these systems are extremely sensitive to spatial geometry, and small drift can cause process failure or vision-collaboration collapse. Standard trajectory applications (arc welding / dispensing / general assembly): deviation > 0.5 mm requires immediate investigation — typically indicates tool collision micro-deformation or assembly tolerance at tool change. Heavy-duty / coarse applications (palletising / large-payload handling / spot welding): deviation > 1.0 mm requires immediate investigation. Corrective action: once the threshold for the application tier is exceeded, stop for investigation of tool-change records, collision history, end-effector fastener loosening, or any unauthorised temporary parameter adjustment.
- User-Frame Verification: re-verify the orientation and origin of the user frame using an alignment tool; frame drift is a frequent cause of accuracy deviation in multi-robot collaboration
- Load-Parameter Check: compare against the load-inertia values in the IPR; confirm no accidental modification
- Backup Diff: file-level comparison between the current controller-parameter backup and the most recent known-good backup; all controller brands provide a backup-export feature, and a generic text-diff tool can perform the comparison. Tool: Capability ③ (Offline Simulation & Program Verification)

4.5.3 Firmware and Software-Version Issue Diagnosis

- Post-Upgrade Anomaly Rollback Decision: function anomaly within 24 hours of upgrade → consider version-compatibility first → consult OEM Release Notes for known issues → if a version issue is confirmed, roll back to the previous version
- Version-Compatibility Matrix: compatibility constraints exist between controller firmware and application-software-package versions; consult the OEM-supplied compatibility matrix before upgrade (usually available on the OEM technical-support portal)
- Patch-Installation Verification: after hotfix installation, perform basic-function regression testing (run the standard test program, confirm the patched function works, and confirm no new anomalies are introduced)

4.5.4 Diagnostic Highlights for Emerging Robotic Middleware

Newer robots (AMR, cobot, humanoid) running ROS / ROS2 middleware have unique software-architecture diagnostic needs and must be diagnosed using the middleware's native toolchain:

- Node Health Check: list all nodes, inspect subscription / publication relationships, visualise the node graph. An isolated node (no subscribers / no publishers) is typically a configuration-error signal
- Topic-Communication Diagnosis: check the publication frequency of key topics (compare to expected); measure end-to-end latency. Control-topic latency exceeding 50 % of the control cycle causes performance degradation
- Data Recording and Playback: record key topic data and replay it in an offline environment — the most effective method to reproduce intermittent faults
- QoS-Configuration Check: mismatched QoS policy is a common but easily overlooked root cause of communication instability in newer middleware

4.6 L4 · AI Decision Layer Initial Diagnosis (Domain 6 Interface)

Domain 2's role in the AI decision layer is preliminary recognition and classification — distinguishing "this is an AI-layer problem" from "this is not." Deep diagnosis and remediation of AI-layer faults are treated by Domain 6 (Chapter 8). Accurate preliminary classification is the key to avoiding misjudgement of AI faults as hardware (or vice versa).

4.6.1 Three Manifestations of AI-Layer Faults

Fault Type	Symptom Description	Key Distinction vs Hardware Faults	Preliminary Disposition Path
Perception Failure	Vision-recognition rate drops suddenly under specific conditions; detection box shifts; objects misclassified	Hardware faults are typically total failures; AI-perception failure is typically "conditional failure" — appears with a workpiece change or lighting change	Hardware exclusion → if hardware is normal, trigger a Domain 6 model evaluation
Decision Drift	Robot behaves inconsistently under identical conditions; completion rate slowly declines	Hardware-accuracy issues correlate with operating time; decision drift correlates with "specific input conditions"	Record failure conditions → trigger a Domain 6 RCA
Generalisation Failure	Complete failure on unseen workpieces or scenes; returns to normal with a known workpiece	Easily misjudged as a mechanical problem; the key recognition feature is that the issue disappears when the known workpiece is restored	Swap to the standard test workpiece to verify → confirm AI layer → trigger Domain 6

4.6.2 Three-Step Hardware Exclusion Protocol (Precondition for AI-Layer Confirmation)

- Step 1 Sensor-Hardware Check: inspect camera-lens contamination / scratching; check light-source intensity uniformity; check sensor power and cabling; observe whether the symptom changes after sensor restart

- Step 2 Standard-Workpiece Test: test using a standard workpiece (known features) under a standard environment; if the standard test passes, AI-layer fault is confirmed (generalisation failure); if the standard test also fails, hardware-layer issues remain suspect
- Step 3 Offline-Inference Verification: run the model on the same input images offline; if the offline result is normal but online fails, the issue is the inference environment or hardware interface; if the offline result is also anomalous, the model itself is confirmed as the issue

> **Conditions for Escalation to Domain 6 After the three-step hardware exclusion confirms an AI-layer fault, submit an "AI Fault Escalation Package" to the Domain 6 queue containing: failure-scenario input-image samples (no fewer than 20), text description of the failure mode, hardware-exclusion record, current model-version number. See Chapter 8 for Domain 6's detailed handling.**

4.7 Intermittent Fault Specialised Diagnostic Methods

Intermittent faults are the hardest to diagnose because they cannot be reproduced on demand, partially disabling every diagnostic means that depends on real-time measurement. RSF's approach to intermittent faults is a "data-driven systematic collection strategy" rather than "wait until next occurrence to measure."

4.7.1 Classification and Priority Diagnostic Hypotheses for Intermittent Faults

Type	Pattern Characteristics	Priority Diagnostic Hypothesis and Entry Point
Heat-Related	Normal at cold start; appears after 1–3 hours of running; frequency increases in summer / hot season	Cooling-system issues; controller internal temperature too high; thermal expansion causing poor contact
Load-Related	Normal under no load; appears under load; appears in specific high-speed / high-torque motions	Reducer near wear limit; drive overload-protection boundary; load-parameter error
Time-Pattern	Appears every N hours; concentrated in specific time windows	Periodic communication timeout; memory leak (software layer); scheduled-task conflict
Environment-Related	Frequent in humid environments or when high-power equipment starts	Grounding / shielding issues causing EMC interference; poor contact (moisture)
Operation-Related	Frequent on specific operator shifts; appears after specific program calls	Unauthorised parameter modifications; collision or singular point on a specific program path

4.7.2 Data Collection Strategy

- Continuous Data Recording: enable the diagnostic tool's background continuous-recording feature (Capabilities ① ④); configure auto-save of the current waveform on baseline-threshold exceedance — when the fault next occurs, data is already captured automatically
- Triggering-Condition Records: build a "trigger-condition card" for each occurrence (time, ambient temperature, continuous-run duration, program name, last parameter-modification time); after five or more occurrences, perform condition-correlation analysis
- Cable Flex Test: actively bend suspect motion-flex cable segments — slowly bend each segment while observing signal changes on the servo-diagnostic tool or oscilloscope; signal jitter during bending localises the damaged segment

- Thermal-Imaging Hot Scan (Capability ⑤): perform a full-robot thermal-camera scan close to the fault occurrence window (after 1.5–2 hours of continuous running); hotspots outside the normal range are direct evidence of a heat-related intermittent fault

4.8 Complete Fault Handling Process: From Report to Root Cause Closure

Fault handling is a complete flow from detection to full closure, in three independent but sequenced stages: Fault Response (restore production) → Root-Cause Disposition (eliminate the cause) → Process Closure (knowledge capture). Explicit separation of these three stages is the core mechanism for preventing the industry's common problem of "workaround substituted for permanent fix."

4.8.1 Eight-Step Fault Response Procedure

Step	Phase	Core Operation	Key Output
S1	**Detection and Recording**	On receiving the fault notification, immediately record: timestamp (minute precision), alarm code, preliminary symptom; confirm priority (P1–P4)	Fault work order opened; SLA response timer begins
S2	**Safety Confirmation**	Confirm safe state on-site or remotely; P1 faults require full LOTO before any further action	Safe-state confirmation record
S3	**Preliminary Layer Localisation**	Use the symptom-classification table of §4.1.3 to complete preliminary layer localisation; confirm the diagnostic path	Layer-localisation conclusion (L1 / L2 / L3 / L4)
S4	**Known-Fault-Mode Match**	Query the organisation's KFL: any matching or similar record? Hit → fast track; miss → continue the full diagnostic funnel	Match result (hit / miss)
S5	**Workaround Evaluation**	Determine whether a workaround can rapidly restore production; if so, execute and explicitly tag as "Temporary — awaiting permanent fix"	Workaround record (if used); production-restoration time
S6	**Stratified Diagnosis**	Execute systematically following the funnel model; use the diagnostic-tool capabilities corresponding to the layer (§4.3–§4.6)	Root-cause localisation conclusion
S7	**Repair Execution**	Permanent fix: eliminate the root cause; record the operation (replaced-part serial numbers / parameter before-and-after values)	Repair-operation record
S8	**Verification Testing**	Run the standard test program; safety-related repairs require full safety-function re-verification	Verification-result record; customer confirms and signs

> **Workaround vs Permanent Fix — Must Be Distinguished Explicitly** When a workaround is used, the work order must clearly state "workaround in effect, permanent fix pending" and a follow-up time must be set (recommended within 48 hours). Open workarounds must appear on the weekly service-review agenda until the permanent fix is complete. "Workaround becoming permanent" is a textbook symptom of low service maturity.

4.8.2 Root Cause Analysis Triggers

Not every fault requires RCA. RCA must be initiated after fault response if any of the following is satisfied:

- The same fault type recurs at least twice within 12 months (same alarm code and same symptom location)
- A single fault causes unplanned downtime exceeding 4 hours
- Involves safety-function failure or safety-related component damage (regardless of frequency)
- This fault triggers an SLA breach
- A workaround was used (all such faults must have a permanent fix and RCA scheduled)
- The symptom does not match any known fault mode (novel fault type — a knowledge entry must be created)

4.8.3 Known Fault Library (KFL)

The Known Fault Library (KFL) is Domain 2's core knowledge asset: it converts each successful diagnosis into a reusable "fast track." Next time the same fault appears, the engineer queries the library and skips the full diagnostic-funnel process.

Table 4.8-1 KFL Entry — Standard Format

Field	Content and Completion Requirements
Entry Number	KFL-XXX (sequential numbering for citation and version traceability)
Fault-Feature Keywords	Platform category / alarm type / affected joint or function / short symptom description (for full-text search)
Symptom Description	Objective description of the phenomenon: alarm code + observed anomalous behaviour; no subjective judgement
Preliminary Layer Localisation	L1 / L2 / L3 / L4 (should match the actual root-cause layer)
Root Cause	Verified root cause (direct cause + underlying cause)
Fast-Track Diagnostic	Three-step short version of "check what first, with which tool, looking for what result" (initial confirmation within 30 minutes)
Permanent Disposition	The complete remediation steps that eliminate the root cause
Workaround Plan	Temporary measures to restore production (with explicit validity window)
Preventive Measures	Systemic improvement actions to prevent recurrence
Entry-Creation Date / Contributing-Engineer Tier	Date format: YYYY-MM-DD; tier provides confidence in entry quality

4.9 Typical Fault Diagnosis Case Library

The following six cases cover typical fault scenarios across the four layers. Each case follows a uniform structure: Background Overview (platform category / symptom / layer localisation) → individual Diagnostic Steps (step name, operation, required tool capability, step conclusion) → Root

Cause and Disposition → Preventive Measures. Cases describe platforms by category; specific brands appear in parenthetical notes only when necessary.

Case 01 Six-Axis Industrial Arm — Excessive Joint Position Deviation

Platform Category	Six-axis industrial arm (heavy-duty, for welding / material handling / assembly)
Preliminary Layer Localisation	L1 mechanical layer (preliminary) → L2 electrical layer (current verification)
Symptom	In auto mode, the J3 axis frequently triggers excessive-position-deviation alarms, pronounced in high-load steps; manual low-speed jogging shows increased resistance; commissioning accuracy was passed and the planned annual-PM interval has been exceeded.

Diagnostic Steps

Step	Step Name	Operation	Required Tool Capability	Step Conclusion
S1	**Alarm-History Analysis**	Review the controller alarm history and analyse the time distribution of position-deviation alarms. Observe whether alarms concentrate in high-load steps and whether low-load steps are normal.	Capability ② System Event Log Analysis	Confirms alarm concentration in high-load steps; preliminarily rules out program errors; points to L1 mechanical or L2 servo layer
S2	**Servo-Current Waveform Capture**	Under the standard test program (rated load), capture the J3 servo-current waveform and overlay-compare it against the current baseline in the commissioning IPR.	Capability ① Servo Data Acquisition & Waveform Analysis	J3 current is ~40 % above the commissioning baseline; mechanical-friction increase confirmed; L1-mechanical hypothesis strengthened
S3	**Joint-Backlash Measurement**	Place the dial indicator on the J3 output end and slowly alternate small forward and reverse torque; record the indicator swing (backlash). Compare to the commissioning baseline.	Capability ⑤ Precision Physical Measurement (dial indicator)	Backlash is 2.3× the commissioning baseline, exceeding the replacement threshold (> 2×); abnormal reducer backlash confirmed
S4	**Grease-State Inspection**	Open the reducer grease port and take a sample; inspect visually for colour (grey-black?) and texture (gritty?).	Capability ⑤ Precision Physical Measurement (oil analysis)	Grease is grey-black with iron-particle texture; reducer internal wear has entered the accelerated phase; root cause located
S5	**Cross-Layer Verification (rule out software / electrical)**	L3 check: confirm load-inertia parameter matches the IPR; rule out unauthorised modification. L2 check: no servo-drive hardware alarm; control supply voltage normal.	Capabilities ② ③ Log analysis + parameter verification	L3 parameters normal; L2 drive has no hardware fault; root cause confirmed as L1 reducer mechanical wear

Root Cause	The J3 reducer exceeded its grease-replacement cycle (actual operating hours over the limit); the internal grease was depleted and iron-contaminated, reducing transmission efficiency and raising friction torque, so under high load the servo could no longer maintain position accuracy. Underlying root cause: the PM schedule is calendar-month-based and did not account for three-shift operation, which accumulates hours at 3× the standard rate.

Permanent Fix *Permanent Fix*	① Replace the reducer per the OEM service manual (including full grease replacement); ② re-execute joint mastering; ③ perform the full accuracy acceptance; ④ change the PM trigger from calendar year to operating hours (OEM-recommended values); ⑤ update the equipment-passport configuration-baseline snapshot.
Workaround *Workaround*	Temporarily reduce J3 motion speed to 50 %, sustaining ~2–4 hours of limited production. Tag as workaround; schedule permanent fix within 48 hours.
Preventive Measures	① Synchronise all three-shift machines in the same environment to operating-hour PM triggers; ② establish quarterly current-trend tracking — any axis current diverging 20 % from baseline triggers a warning; ③ add this case to the KFL (entry KFL-001).

Case 02 Cobot — Force-Torque Sensor Temperature Drift Causing Frequent Safety Stops

Platform Category	Cobot equipped with an end-effector force-torque sensor; light-load precision application
Preliminary Layer Localisation	L2 electrical layer (sensor layer)
Symptom	During normal light-load operation, protective stops trigger repeatedly — rare on the morning shift (~22 °C) and frequent on the afternoon shift (~28 °C); changing workpiece or reducing speed has no effect.

Diagnostic Steps

Step	Step Name	Operation	Required Tool Capability	Step Conclusion
S1	**Trigger-Condition Correlation Analysis**	Tabulate recent protective-stop occurrences against time, ambient temperature, and continuous-run duration.	Capability ② System Event Log Analysis	Stop frequency strongly correlates with ambient temperature, unrelated to load or program; the temperature hypothesis holds
S2	**Real-Time Sensor Zero-Point Monitoring**	With no external force and the robot stationary, read the force-torque sensor zero-point reading via the teach pendant service interface at 22 °C and 28 °C; record the difference.	Capability ② System diagnostic-interface read	Zero-point at 28 °C is ~10 N offset from 22 °C; sensor temperature drift confirmed beyond normal range
S3	**Temperature-Drift Spec Verification**	Consult the robot's spec sheet for the sensor's temperature-drift coefficient; compute the theoretical drift over the 6 °C difference and compare to measurement.	No measurement tool needed (documentation + calculation)	Measured drift far exceeds spec; the sensor zero-point was calibrated at a non-working temperature, leaving a systemic offset
S4	**External-Heat-Source Exclusion**	Use the thermal camera to scan the area around the sensor mount; confirm no abnormal external heat source (welding equipment, heaters) is causing local heating.	Capability ⑤ Thermal-imaging analysis	No abnormal external heat source; drift is a sensor temperature-characteristic issue requiring recalibration at the working temperature
S5	**Working-Temperature Recalibration Verification**	After the ambient temperature stabilises at the actual working range (25–28 °C), perform zero-point recalibration per ISO/TS 15066; after calibration, repeat step 2 to verify the drift is eliminated.	Capability ⑤ Force-torque calibration	Post-calibration zero-point returns to within spec at working temperature;

			procedure	verification passes and root cause confirmed

Root Cause	The force-torque sensor zero-point was calibrated at a temperature lower than the actual working temperature, producing a systemic offset at working temperature that exceeded the protective-stop threshold. The underlying root cause is that no "working-temperature torque baseline" was established at commissioning.
Permanent Fix *Permanent Fix*	① Re-execute ISO/TS 15066 force-torque zero-point calibration after the actual working temperature stabilises; ② update the IPR with the calibration ambient temperature; ③ re-perform full safety-function verification at working temperature after calibration.
Workaround *Workaround*	Temporarily raise the protective-stop trigger threshold (must verify the adjustment still meets ISO/TS 15066 safety limits and never exceeds the standard's upper bound for human-contact safety torque).
Preventive Measures	① Institutionalise quarterly cobot torque-baseline checks after thermal steady state (30 minutes of continuous running); ② commissioning documents must record the ambient temperature when the torque baseline was established; ③ add this case to the KFL (entry KFL-002).

Case 03 Industrial Ethernet — Periodic Disconnect (- every 10 minutes)

Platform Category	Six-axis industrial arm (any brand) connected via industrial Ethernet to downstream PLC and I/O slaves
Preliminary Layer Localisation	L2 / L3 boundary (communication-network layer)
Symptom	Industrial-Ethernet communication between the robot and the downstream PLC drops about every 10 minutes; the furthest slave shows an anomalous state; restart restores communication; physical-cable visual inspection shows no obvious damage.

Diagnostic Steps

Step	Step Name	Operation	Required Tool Capability	Step Conclusion
S1	**Bus-Topology State Analysis**	Use the bus-configuration tool to observe real-time topology; wait for the next disconnect and record the order in which slaves change state (which goes anomalous first).	Capability ⑥ Industrial Communication Protocol Analysis (configuration tool)	The 3rd slave from the end goes anomalous first; preliminary indication is a physical-layer or configuration issue between slaves 2 and 3
S2	**Protocol-Layer Packet Capture**	Tap a protocol analyser between slaves 2 and 3; capture packets in real time and analyse the communication cycle for timeout spikes; record their time pattern.	Capability ⑥ Industrial Communication Protocol Analysis (protocol analyser)	A timeout spike appears in the communication cycle every ~10 minutes, triggering the slave-state change; timing-timeout problem confirmed
S3	**Time-Pattern Correlation Analysis**	Compare communication-timeout times to the work cycles of other high-power equipment in the shop; identify any strong EMC sources (welder / press / large VFD) running on a ~10-minute cycle.	No measurement tool needed (log-correlation analysis)	The 10-minute period coincides closely with the welding-equipment cycle; EMC-interference hypothesis locked

S4	**Cable Shield-Continuity Test**	Power down and use a multimeter to measure shield continuity segment by segment on the slave 2–3 cable. Bend each segment along the entire length and watch for sudden changes.	Capability ⑤ Precision Physical Measurement (multimeter)	At the conduit-exit bend, the shield is open (resistance jumps from < 1 Ω to > 10 MΩ); the shield was abraded during cable routing, defeating EMC protection
S5	**Root-Cause Confirmation (workaround verification)**	Temporarily repair the shield at the breakpoint and run for 30+ minutes; observe whether disconnects disappear.	Capability ⑤ (temporary materials for verification)	Disconnects disappear; root cause verified as broken shield defeating EMC protection; cable replacement confirmed as permanent fix

Root Cause	The communication cable's shield was broken at the routing bend; high-frequency EMC radiation from the welding equipment coupled through the unshielded segment into the communication signal, causing periodic timing timeouts and slave-state transitions.
Permanent Fix *Permanent Fix*	① Replace the damaged communication cable (industrial-grade shielded twisted pair, shield coverage ≥ 85 %, bend radius per OEM spec); ② re-plan the routing so all bends have radius ≥ 5× cable OD; ③ 100 % retest of shield continuity after replacement; ④ accept after 2+ hours without disconnects.
Workaround *Workaround*	Temporarily repair the shield with aluminium-foil tape and modestly raise the communication retry-count setting. Permanent fix required within one week.
Preventive Measures	① All industrial-Ethernet cables must have 100 % shield-continuity test and record at installation; ② installation SOPs must specify the minimum cable bend radius; ③ add this case to the KFL (entry KFL-003).

Case 04 Welding Robot — Systemic TCP-Accuracy Drift After Continuous Operation

Platform Category	Six-axis industrial arm, precision welding application, TCP-accuracy requirement within ± 0.3 mm
Preliminary Layer Localisation	L1 mechanical layer (thermal expansion)
Symptom	Cold-start TCP error < 0.1 mm (normal); after 2 hours of continuous running, the arc-start position shows -0.8 mm systemic offset that continues to grow; compensating in one direction merely shifts the offset to another.

Diagnostic Steps

Step	Step Name	Operation	Required Tool Capability	Step Conclusion
S1	**Software / Parameter-Layer Exclusion**	Compare current TCP-parameter values to the IPR; confirm no modification. Switch to a backup program and observe whether the issue persists.	Capability ③ Offline simulation and parameter comparison	Software parameters unmodified; the backup program also exhibits the offset; L3 software layer ruled out; problem is in the physical layer
S2	**Cold vs Hot TCP Comparison**	Measure TCP deviation precisely with the standard verification program in cold-start state (within 10 min of power-on) and after 2 hours of continuous running.	Capability ⑤ Precision Physical Measurement (TCP-measurement tool)	Cold-start error < 0.1 mm (normal); hot error 0.82 mm (out of process tolerance); heat-related accuracy issue confirmed

S3	**Full-Robot Thermal-Distribution Scan**	Immediately after 2 hours of continuous running (thermal steady state), thermal-image-scan the entire robot, record each joint's casing temperature, and compute the temperature rise of each part.	Capability ⑤ Precision Physical Measurement (thermal camera)	Major joints (J1–J3) rise ~18 °C; wrist joints ~8 °C; major-joint thermal expansion is the dominant offset source
S4	**Theoretical Thermal-Expansion Verification**	Using the measured temperature rises, the arm material's linear-expansion coefficient, and the arm lengths, compute each joint's thermal expansion and compare to the measured TCP offset.	No measurement tool needed (engineering calculation)	Computed thermal expansion ~0.7 mm closely matches the measured 0.82 mm offset; root cause confirmed as thermal expansion (a physical characteristic, not a fault)
S5	**Warm-Up-Program Effectiveness Verification**	Write a standard warm-up program (10 minutes on a standard path); re-measure TCP accuracy once the robot reaches thermal steady state.	Capability ③ Program writing and verification	After 10-minute warm-up, TCP error stabilises at ± 0.15 mm (meets process requirement); warm-up plan validated

Root Cause	During warm-up, metal thermal expansion produces systemic TCP drift (a physical characteristic, not equipment fault). The underlying cause is that the operating procedure does not require a warm-up and the commissioning accuracy acceptance was done from cold, with no hot-steady-state baseline established.
Permanent Fix *Permanent Fix*	① Create a standard warm-up program (10-min standard-path cycle) and incorporate it into the shift-start SOP; ② enable the controller's temperature-compensation function if available; ③ update the accuracy-acceptance procedure so the final acceptance is performed at thermal steady state.
Workaround *Workaround*	Mandate "10-minute warm-up before production mode" as an immediate operating rule; use the warm-up window for production preparation to minimise idle time.
Preventive Measures	① For all precision welding / assembly applications, commissioning documents must note "accuracy data measured at thermal steady state"; ② quarterly-PM accuracy measurements must be taken at thermal steady state; ③ add this case to the KFL (entry KFL-004).

Case 05 Vision-Guided Picking — Sudden Failure-Rate Rise on a New Workpiece Type

Platform Category	Industrial arm or cobot with vision-guided system, 2D / 3D vision-pick application
Preliminary Layer Localisation	L4 AI decision layer (vision perception) + L2 electrical layer (to be ruled out)
Symptom	After introducing a new workpiece (high-gloss metal surface), pick success rate dropped from 98 % to 67 %; on failure, the TCP-arrival position is off by ~15–30 mm; restoring the original workpiece restores success rate.

Diagnostic Steps

Step	Step Name	Operation	Required Tool Capability	Step Conclusion
S1	**Sensor Hardware Exclusion**	Inspect camera lens (contamination / scratches); measure illuminance with a photometer (compare to commissioning record); check camera power voltage; run the test on the original workpiece and record success rate.	Capability ⑤ Precision Physical Measurement (photometer + multimeter)	Original-workpiece success rate 97 % (normal); no hardware anomaly; the problem is workpiece-related, preliminarily locking L4 AI perception layer

S2	**Input-Image Quality Analysis**	Capture raw images of both workpiece types; compare in the vision-debug interface: check for saturated / overexposed highlights; count the difference in valid feature points.	Capability ⑦ AI Inference Diagnostics (vision-debug interface)	The new workpiece's metallic gloss produces large saturated regions; valid feature points drop by -70 %; vision-model input quality severely degraded
S3	**Model-Match-Confidence Analysis**	Inspect the vision-system match-confidence log; compare confidence ranges for the original workpiece (Score > 0.90) and the new workpiece; identify the confidence distribution of failed picks.	Capability ⑦ AI Inference Diagnostics (confidence log)	New-workpiece confidence concentrates in 0.45–0.65 (below the reliable threshold of 0.75); insufficient AI-perception generalisation confirmed
S4	**Lighting-Adjustment Plan Verification**	Change the light source from frontal direct to 45° side lighting to reduce saturated highlights; retest the new workpiece and record success rate and confidence.	Capability ⑦ (parameter-adjustment verification)	Confidence rises to 0.81–0.92; success rate returns to 93 %; lighting optimisation works as a short-term measure, but generalisation must still be addressed fundamentally
S5	**Domain 6 Escalation Trigger**	Compile the AI-fault escalation package: collect ≥ 30 failure-image samples of the new workpiece (covering different failure poses); document the failure mode; submit to the Domain 6 queue and request model retraining.	Capability ⑦ (Domain 6 escalation flow)	AI-fault escalation package submitted; Domain 6 will complete model evaluation and the retraining plan

Root Cause	The vision-perception model was trained without sufficient samples of high-gloss metallic surfaces; the new workpiece's reflectivity is outside the model's training distribution, producing low feature-matching confidence (generalisation failure from out-of-distribution input). The direct trigger is that the lighting design did not account for surface-gloss differences across workpiece types.
Permanent Fix *Permanent Fix*	① Domain 6 retrains the model using an augmented dataset including the new workpiece; ② update the lighting to a diffuse source (softbox) suitable for multiple gloss types; ③ create a "vision-compatibility check before workpiece change" procedure: new workpieces must pass 100 offline tests with confidence > 0.85 before production.
Workaround *Workaround*	Temporarily switch lighting to 45° side illumination (restores success rate to ~93 %); also institute manual visual sampling (1 in 20) until permanent fix is complete.
Preventive Measures	① Vision-system commissioning documents must record the workpiece type and surface gloss used during verification; ② add a "vision-system compatibility verification" step to the new-workpiece introduction process; ③ add this case to the KFL (entry KFL-005).

Case 06 AMR Laser-SLAM — Localisation Drift in a Loaded-Shelf Zone

Platform Category	Laser-SLAM autonomous mobile robot (AMR), warehouse / logistics application
Preliminary Layer Localisation	L3 / L4 boundary (SLAM algorithm + sensor-environment adaptation)
Symptom	In a particular zone with metal, fully-loaded shelves, the AMR's localised pose diverges from

	its actual pose (up to ~40 cm); re-initialisation temporarily restores; navigation in empty-shelf zones is normal.

Diagnostic Steps

Step	Step Name	Operation	Required Tool Capability	Step Conclusion
S1	**LiDAR Hardware Exclusion**	Inspect LiDAR lens for contamination / scratches; check supply voltage; test AMR localisation accuracy in an open, low-reflectance zone (a clear aisle) and record the error.	Capability ⑤ Precision Physical Measurement (visual + voltage)	Open-zone localisation error < 2 cm (normal); no hardware anomaly; the issue appears only in the high-reflectance shelf zone, pointing to environment / algorithm layer
S2	**SLAM Point-Cloud-Match Visualisation Analysis**	Use the AMR debug tool to visualise real-time point-cloud scan against the pre-built map; observe the match quality (confidence score) and the presence of spurious reflection points in the high-reflectance zone.	Capability ⑦ AI Inference Diagnostics (SLAM visualisation tool)	Point-cloud match confidence drops from 0.92 in a normal aisle to 0.58 in the high-reflectance zone (below the reliable threshold of 0.70); large numbers of spurious reflection points appear
S3	**Map vs Actual-Environment Discrepancy Analysis**	Pull up the original mapping (empty shelves) and compare with the current loaded-shelf point-cloud features; measure the discrepancy between actual point-cloud density and the corresponding region of the map.	Capability ③ (map-file analysis)	The map was built with empty shelves; current loaded-shelf point-cloud features differ from the map by > 60 %; SLAM cannot reliably match current environment to stored map
S4	**Reflectivity-Impact Verification**	Temporarily apply matte (anti-reflective) material to the shelves; retest AMR localisation in the same zone and record error and confidence change.	Capability ⑤ (temporary materials for verification)	After matte treatment, confidence rises to 0.84 and localisation error returns to < 5 cm; metallic-mirror reflection confirmed as the key SLAM-degradation factor
S5	**Reflective-Target Anchor Plan Verification**	Install three laser-reflective targets at the shelf ends in the zone; update the map with target positions; have the AMR navigate through the zone and record localisation accuracy with target assistance.	Capability ⑦ (SLAM-parameter adjustment)	Target-assisted localisation error < 3 cm (meets warehouse-accuracy requirement); plan validated

Root Cause	The SLAM map was built with empty shelves; with shelves loaded, the actual point-cloud features differ fundamentally from the map. Strong reflection from glossy metal shelves causes systemic range-measurement errors that further worsen point-cloud matching. The SLAM algorithm's robustness is insufficient for the significant feature change between empty and loaded shelves.
Permanent Fix *Permanent Fix*	① Rebuild the map with shelves loaded (actual operating state); ② install laser-reflective targets as stable localisation anchors in high-reflectance zones; ③ establish a quarterly map-accuracy review (navigate to known precise coordinates and measure arrival-position error).
Workaround *Workaround*	Designate the high-reflectance zone as a low-speed zone (speed reduced to 0.3 m/s) and provide manual-guidance assistance to maintain basic operation.
Preventive Measures	① The mapping procedure must explicitly require "map under actual operating state (loaded shelves, typical lighting)"; ② first-commissioning documents must include localisation-accuracy tests for all warehouse zones; ③ add this case to the KFL (entry KFL-006).

4.10 Technical Escalation and Support Reporting Mechanism

4.10.1 Engineer Authorised Diagnostic Boundaries

Certification Tier	May Perform Independently	Must Escalate to Specialist	Must Escalate to Expert / OEM
Professional	◦ Standard remediation when a KFL entry is matched ◦ Alarm-code lookup and SOP-defined remediation ◦ Shift inspections and data collection for P3 / P4 faults ◦ Wear-part replacements specified by SOP	Complex diagnosis at L1 / L2; any parameter modification; P1 / P2 priority faults	AI-layer faults; safety-circuit-related work; controller-firmware modification
Specialist	◦ Complex cross-layer diagnosis ◦ Parameter modifications (within known scope of impact) ◦ Lead root-cause analysis ◦ Handle P2 faults independently	Hardware faults requiring OEM-original support; parts beyond local-spares scope	Deep AI-layer diagnosis (→ Domain 6); systemic safety-function modification

4.10.2 OEM Technical-Support Reporting Standards

Before contacting the robot maker's technical support, prepare a complete information pack to reduce support round-trips:

- Robot model and serial number; controller model and serial number
- Controller firmware version (down to minor version)
- Complete alarm-history file (50+ most recent entries)
- Video of the fault occurrence (30+ seconds clearly showing the symptom)
- Summary of diagnostic steps already performed (to avoid duplication)
- List of parts already replaced (if any)
- Date and summary of the most recent PM
- Remote-access authorisation confirmation (which management level may approve remote connection)

4.10.3 Knowledge-Entry Quality Standards

The value of the knowledge base depends on entry quality. The following entry types are not admitted:

incomplete description (missing the fast-track field); unverified root cause (workaround succeeded but root cause unconfirmed); > 80 % overlap with an existing entry. After successfully diagnosing a new fault type, an entry must be created within 24 hours of work-order closure (format per §4.8.3).

4.11 Domain 2 Key Performance Indicator Framework

Indicator	Definition	Industry Benchmark Target	World-Class Target	Management Significance
First-Time Fix Rate (FTFR)	Work orders with correct first-visit root-cause localisation / total work orders	**≥ 70%**	**≥ 85%**	Reflects diagnostic capability and tool proficiency; low FTFR indicates training need
Mean Time to Diagnose Anomaly (MTDA)	Average time from diagnosis start to root-cause confirmation	**< 3 hours**	**< 1.5 hours**	A primary component of MTTR; MTDA analysis localises diagnostic bottlenecks
Layer-Localisation Accuracy	Share of first-attempt layer localisations that match the eventual root-cause layer	**≥ 80%**	**≥ 90%**	Measures funnel-model application quality; low accuracy indicates engineers are skipping the layer-localisation step
Workaround Usage Rate	Share of work orders closed using a workaround	**< 20%**	**< 10%**	A high rate indicates insufficient permanent-fix capability or time — investigate the cause
Workaround-to-Permanent Conversion Rate	Share of workarounds permanently fixed within 48 hours	**≥ 85%**	**≥ 95%**	Measures the discipline of not letting workarounds become permanent
Repeat-Fault Rate (RFR)	Share of faults that recur within 12 months	**< 15%**	**< 5%**	A high RFR indicates poor RCA quality or weak preventive-measures execution
KFL Hit Rate	Share of work orders for which a KFL lookup hits	**≥ 30%**	**≥ 50%**	Measures KFL accumulation effectiveness; a low hit rate means either entry quality or lookup habits need improvement
Work-Order Completeness Rate	Share of work orders complete across four fields (phenomenon / layer / root cause / disposition)	**≥ 90%**	**≥ 98%**	Prerequisite for KFL quality; incomplete work orders cannot produce useful knowledge entries

4.12 Chapter Summary and Inter-Domain Interfaces

Chapter Summary Domain 2 (Fault Diagnosis & Technical Support), built on the four-layer stratified diagnostic framework, transforms robot-fault diagnosis from experience-driven random investigation into logic-driven systematic localisation. Defining seven diagnostic-tool capabilities lets engineers clarify their competence-development targets and tool-configuration direction in a brand-neutral way.

The three-stage fault-handling process explicitly separates workaround from permanent fix, and the Known Fault Library converts diagnostic experience into a reusable organisational asset. The methodology of this chapter applies to any brand and any type of industrial robot system.

Interfaces with Other Capability Domains

Interface Direction	Interfacing Domain	Content
Domain 1 → Domain 2	Service Lifecycle Management	Domain 1's commissioning baselines (current / accuracy / vibration) are the quantitative reference for Domain 2 diagnosis; Domain 1 phase information determines Domain 2's diagnostic context
Domain 2 → Domain 1	Service Lifecycle Management	Domain 2's RCA conclusions trigger Domain 1's PM-strategy adjustments (through controlled change); diagnostic conclusions update the technical-state record in the equipment passport
Domain 2 → Domain 6	AI/ML System Service Methods	After AI-layer preliminary diagnosis (§4.6) completes hardware exclusion, the fault is escalated to Domain 6; Domain 2's failure-sample data is the raw input for Domain 6's model evaluation
Domain 2 → Domain 4	SLA Design & Service Commitment	MTTR / FTFR / MTDA are the technical-design basis for SLA response-commitment clauses; fault history is a key input to SLA pricing models
Domain 2 ← Domain 5	Cross-Platform Service Procedures	Domain 5 supplies platform-specific diagnostic procedures (e.g., cobot force-torque sensor diagnosis differs from industrial-arm reducer diagnosis); Domain 2's four-layer framework is the universal skeleton, with Domain 5 adding platform-specific content

Chapter 5

Domain 3 · Remote Service & Digital Operations

5.0 Chapter Introduction

The delivery model of robotics service is undergoing structural change. In the traditional model, the engineer's value is concentrated in the "moment of arrival" — turning up at the customer's site with a toolbox, performing diagnosis and repair, and leaving. That model is being supplanted by a new service logic: engineers maintain a continuous, proactive digital connection to the robot system, identify degradation trends before faults occur, intervene before downtime occurs, and arrive on site (when required) already armed with precise diagnostic conclusions.

The technical foundations of this shift — robot digital twins, multi-sensor predictive maintenance, secure remote-access architecture, OTA firmware management — form the core of Domain 3. Domain 3 is not the online version of Domain 2 (Fault Diagnosis) but an independent capability system: it advances the engineer's work timeline from "after the fault" to "before the fault," extends the delivery location from "the customer site" to "anywhere with network coverage," and upgrades the basis of maintenance decisions from "experience-based judgement" to "data-driven."

This chapter draws methodology from the most recent academic research (2023–2025) and industrial practice. §5.1 defines the remote-service capability framework; §5.2 establishes the secure remote-access architecture; §5.3 develops the digital-twin methodology; §5.4 builds the data infrastructure for remote operations; §5.5 establishes predictive and prescriptive maintenance; §5.6 governs OTA firmware and software management; §5.7 defines edge-cloud computing; §5.8 sets out the complete remote-service-delivery process; §5.9 establishes the Domain 3 KPI framework; §5.10 provides the chapter summary and inter-domain interfaces.

5.1 Remote Service Capability Framework — From Reactive Response to Proactive Anticipation

Domain 3's capability system takes "service-model evolution" as its core logical axis. Understanding this evolution is the precondition for correctly assessing the value and boundaries of each technical means. RSF organises remote-service capability into four evolution stages, each representing a different level of technical maturity and service value.

Table 5.1-1 Four-Stage Evolution Model of Remote-Service Capability

Phase	Service Mode	Technical Characteristics	Service Value	Typical Trigger
S1	S1 — Reactive	Customer phones in a fault → engineer dispatched → on-site diagnosis and repair	Production restoration (after downtime)	Triggered after customer fault report
S2	S2 — Remote Response	Remote connect → remote diagnose → remote parameter / firmware push; reduces on-site needs	Reduces downtime; cuts travel cost	Triggered by customer report but with remote-first response

S3	S3 — Proactive Monitoring	Continuous state monitoring → threshold alarms → proactive customer contact → early intervention	Action taken before the fault; raises share of planned maintenance	Triggered by monitoring-system alarms; periodic state reports
S4	S4 — Predictive & Prescriptive	ML models → remaining-useful-life (RUL) prediction → auto-generated maintenance prescriptions → parts pre-scheduled	Transforms maintenance from cost centre into value-creating activity; optimal-timing decisions	AI model continuously outputs health index; auto-triggers maintenance work orders

Note on the Evolution Path The four stages are not stepwise leaps but cumulative capability layering: each higher stage builds on the previous. A service organisation should first establish a solid secure-remote-access foundation at S2, then progressively develop the data-collection and analytics capabilities of S3 and S4. Skipping the security foundation and deploying AI predictive models directly leaves no data source and introduces unmanageable cybersecurity risk.

5.1.1 Six Core Capabilities of Domain 3

ID	Core Capability	Definition	Corresponding Stage
C1	**Secure Remote Access**	Establish, manage, and terminate authorised remote connections to robot control systems in compliance with IEC 62443, including authentication, access control, and session audit	S2 and above
C2	**Real-Time State Monitoring**	Continuous data collection to monitor key parameters (current, temperature, vibration, operating hours) in real time, and to generate alarms via threshold rules and statistical models	S3 and above
C3	**Robot Digital Twin (RDT) Construction & Operations**	Build the robot's digital-twin model with bidirectional physical-digital synchronisation; perform state assessment, scenario simulation, and performance prediction in the digital environment	S3 and above
C4	**Predictive & Prescriptive Maintenance**	Use ML and physics-informed hybrid models, with multi-sensor data, to predict the RUL of critical components and output specific maintenance prescriptions (when, where, what)	S4
C5	**OTA Firmware & Software Management**	Remotely update controller firmware, application software, and AI models over secure channels, including version management, staged rollout, rollback, and update verification	S2 and above
C6	**Remote-Service Data Governance**	Establish governance for the collection, storage, access control, ownership, and compliance of robot-operations data, securing data assets while enabling their use	S2 and above

5.2 Secure Remote Access Architecture

Remote access is the physical foundation of every Domain 3 capability — without secure remote connectivity, digital twins cannot synchronise data, predictive models cannot ingest sensor input, and OTA updates cannot reach target devices. Yet "secure" is not the easiest precondition of "reachable" to satisfy. Robot controllers are industrial control systems (IACS), and their cybersecurity needs differ

fundamentally from IT: real-time requirements do not permit complex encryption latency; OT devices update on multi-year cycles and cannot reboot frequently; production continuity makes a "disconnect first, then harden" strategy unacceptable.

5.2.1 IT/OT Convergence Security Framework — Aligning IEC 62443 and ISO 27001

RSF uses a dual-framework approach — IEC 62443 (OT cybersecurity) together with ISO 27001 (information-security management) — rather than a single framework. The two standards are complementary, not substitutes: IEC 62443 addresses device-layer and control-system-layer security; ISO 27001 addresses organisation-level information-security governance.

Dimension	IEC 62443 (OT Security)	ISO 27001 (Information-Security Management)	Application in Robot Remote Service
Focus Area	Cybersecurity of industrial control systems, robot controllers, and field devices	Comprehensive security management of organisational information assets (including IT systems)	IEC 62443 covers OT-data security from controller to cloud; ISO 27001 covers the overall information-security management of the service organisation
Core Concepts	Zone-and-Conduit segmentation; Security Levels (SL 1–4)	Information-Security Management System (ISMS); risk assessment and treatment	Remote-access channel = "Conduit"; controller network = "Zone"; access authorisation = ISO 27001 identity management
Patch Management	IEC 62443-2-3 (2024) expands patch-management requirements with explicit patch-assessment cycles	ISO 27001 control A.12.6: system patch management	OTA-update procedures must satisfy both IEC 62443's patch-authorisation flow and ISO 27001's change-management requirements
Remote Access	Remote access must be isolated through a "DMZ conduit" and must produce access logs	All remote-access sessions (who, when, what) must be logged	Remote sessions must traverse an industrial DMZ; all sessions are fully logged and periodically audited
2024 Update Highlights	IEC 62443-2-1:2024 adds organisation-level cybersecurity-program requirements; clearer SL definitions	ISO 27001:2022 includes cloud-security and remote-work controls	Both standards now cover new threat scenarios of remote operations; service organisations should synchronise internal policy updates

5.2.2 Zone-Conduit Remote Access Architecture

The core architectural concept of IEC 62443 is the Zone-and-Conduit model: the robot-control-system network is partitioned into zones of different security levels, and zones may communicate only through controlled conduits. This model provides a clear security-boundary definition for remote access.

Standard Four-Tier Remote-Access Architecture

Zone 0 — Robot Control Layer Robot controllers, servo drives, safety controllers, field I/O devices

Security Level: SL 2–3 (critical operations system)	Access Rule: only authorised engineer workstations may access through encrypted channels; any direct internet-originated access must be rejected	**Risk Focus: direct access could allow motion-control commands to be tampered with — a personal-safety hazard**

Zone 1 — Manufacturing Network Layer Shop-floor industrial Ethernet, SCADA, MES, historian

database

Security Level: SL 2	Access Rule: authorised engineer workstations may access; an industrial firewall sits between Zone 0; a DMZ sits between Zone 2	**Risk Focus: overly broad permissions could leak production-line data or enable lateral movement into control systems**

Zone 2 — Industrial DMZ (critical isolation layer) Data-acquisition servers, remote-access gateways, protocol converters, data buffers

Security Level: SL 2	Access Rule: the DMZ is the "airlock" of remote access — external connections terminate at the DMZ and never pass directly through to Zone 0/1; data flows unidirectionally (OT → IT)	**Risk Focus: misconfigured DMZ is the most common OT-cybersecurity vulnerability; conduct periodic penetration tests**

Zone 3 — Enterprise IT and Service-Cloud Layer Service-engineer terminals, remote-access portal, digital-twin cloud platform, predictive-maintenance analytics engine, AI-training infrastructure

Security Level: SL 1 (managed under ISO 27001)	Access Rule: engineers connect to the remote-access portal after MFA authentication; all access transits the DMZ; direct connection to Zone 0/1 bypassing the DMZ is forbidden	**Risk Focus: weak authentication is the most common entry vector; enforce MFA and least-privilege**

5.2.3 Zero-Trust Remote Access Principles

The traditional "castle-and-moat" security model (trusted inside the perimeter, untrusted outside) no longer fits the remote-service context — engineers may connect from a café, home office, or third-party network on a customer site, and the castle-and-moat model cannot handle these origins safely. Zero Trust Architecture (ZTA) reduces the granularity of trust from "network location" to "each access request," better suited to remote-service security needs.

Zero-Trust Principle	Implementation in Robot Remote Service	Relationship to IEC 62443
Never Trust, Always Verify	Every remote session requires MFA; authenticated sessions are continuously verified (periodic re-authentication + behavioural-anomaly detection)	IEC 62443 SL2: all remote access requires strong authentication
Least-Privilege Access	Engineers are authorised only for the specific controllers needed for the current task; read and write privileges are separated; parameter-modification privilege requires additional authorisation	IEC 62443-3-3: permissions managed on a need-to-know basis
Assume Breach	Enable full operation logging on every remote session; deploy an IDS in the DMZ; conduct periodic penetration tests	IEC 62443-2-3:2024 requires periodic security assessment; ISO 27001 A.16 requires a security-incident response plan
Micro-Segmentation	Different customers' robot networks must be strictly isolated; one customer's remote-access credentials must never reach another customer's control system	Extension of IEC 62443 zone-isolation: multi-customer service-provider scenarios must achieve logical isolation

5.2.4 Remote-Access Session Management Procedure

Step	Node Name	Operation	Compliance Requirement
R1	**Access Request**	Engineer submits an access request (target robot	ISO 27001: all remote access

	and Authorisation	S/N, purpose, expected duration); the named customer-authorised representative approves; the authorisation record is archived	requires explicit customer authorisation, named and recorded
R2	**Authentication**	Engineer completes MFA (password + hardware token or mobile OTP); the system verifies the engineer's access-permission list and authorisation for the specific controller	IEC 62443-3-3: SL2 requires strong authentication; SL3+ requires hardware tokens
R3	**Encrypted Channel Establishment**	Establish an encrypted tunnel (IPSec VPN or TLS 1.3) through the industrial DMZ; connect to the target controller; log connection time, engineer ID, target-device ID	All transmission end-to-end encrypted; cleartext protocols (Telnet / FTP) prohibited
R4	**Least-Privilege Access Execution**	Only authorised functions are accessed (log read / diagnostic / specific parameter modifications); operations outside scope require suspending the session and requesting additional authorisation	IEC 62443 least-privilege; parameter modifications log before-and-after values in real time
R5	**Real-Time Operation Logging**	All session operations (commands, parameter changes, file transfers) are logged in real time to a tamper-evident system; screen capture or video per contract	ISO 27001 A.12.4: operation logs tamper-resistant; retention ≥ 12 months
R6	**Session Termination and Audit**	Terminate session actively (do not wait for timeout); generate session-summary report (operations, modification log, next-step recommendations); send to the customer-authorised representative	Every remote session has a written summary; the customer may review the full operation log at any time

5.3 Robot Digital Twin (RDT)

Robot Digital Twin (RDT) is a core industrial technology that the 2025 Gartner Hype Cycle predicts will "enter mainstream adoption within 1–3 years." Its essence is to create a real-time-synchronised virtual mirror of the physical robot system, letting engineers observe current state, assess historical trends, and simulate future scenarios in the digital environment — without travelling to the physical site.

RSF adopts the latest academic four-layer RDT system architecture (based on a 2024 ScienceDirect review) to organise digital-twin functions by technical layer, avoiding the common conceptual conflation of "digital twin" with "equipment-monitoring dashboard."

5.3.1 RDT Four-Layer System Architecture

Table 5.3-1 Robot Digital Twin (RDT) Four-Layer Architecture

Tier	Layer Name	Technical Content	Function in Robot Service
L1	**Physical Data-Acquisition Layer**	Multi-sensor network (vibration / current / temperature / joint position / torque); industrial-protocol interfaces (OPC UA, MQTT, EtherCAT acquisition); edge-gateway preprocessing	Establishes the real-time data channel from the physical robot to the digital twin; data quality directly determines model accuracy. Sensor selection and sampling rate must match the diagnostic objective (vibration diagnosis needs ≥ 1 kHz; temperature monitoring needs only 100 Hz)
L2	**Digital-Twin Model Layer**	Kinematic / dynamic simulation models; physics-informed (PINN) models; ML health-assessment models; finite-element-analysis (FEA) structural models	Translates sensor data into engineering meaning: "current up 3 %" → "early signal of reducer-efficiency decline." Models fall into two classes: physics-equation-based white-box (highly

			interpretable, strong generalisation); data-based black-box (high accuracy, requires large training datasets). Hybrid (grey-box) is current best practice.
L3	**Service-Function Layer**	State assessment and alarming; RUL prediction; virtual commissioning and scenario simulation; maintenance-prescription generation; remote-diagnosis assistance	Converts the model-layer outputs into service instructions an engineer can use directly: "J3 reducer RUL: - 750 h (confidence interval 650–850 h); recommend grease replacement at the next maintenance window (- 23 days)."
L4	**Application-Interface Layer**	Service-engineer dashboard; customer state-report portal; CMMS interface; AR / VR remote-assistance interface; API integration	Presents the service outputs to different audiences (service engineer, service manager, customer); ensures the right information reaches the right person at the right time in the right format; AR remote collaboration is a rapidly maturing 2024–2025 use case

5.3.2 Physical-Digital Synchronisation Mechanism

The "twin-ness" of a digital twin lies in the continuous synchronisation between physical and digital state. The synchronisation design determines model real-timeliness and accuracy:

- **Real-Time Sync**

  Use Cases: remote diagnosis (millisecond response) and motion-trajectory monitoring. Data are synchronised directly through the edge gateway over a low-latency path to the digital-twin server, with target latency < 100 ms. Requirements: low-latency industrial network (5G or deterministic Ethernet / TSN); edge preprocessing to reduce bandwidth.

- **Near-Real-Time Sync**

  Use Cases: state monitoring and health-index updates. Sync frequencies of seconds to minutes are suitable for trend analysis. This covers most predictive-maintenance needs and is the right starting point for digital-twin deployment.

- **Batch Sync**

  Use Cases: historical analysis, AI model training, monthly state reporting. Per-shift or per-day sync is bandwidth-friendly and suits factories with limited network conditions (unstable network or 4G-only coverage).

- **Event-Triggered Sync**

  Use Cases: detailed recording of anomalous events. When sensor data exceed the normal range, high-frequency acquisition is auto-triggered (from 1 Hz to 100 Hz), capturing the full waveform around the anomaly to provide high-quality raw data for downstream RCA.

5.3.3 Service Function Applications of the Digital Twin

Virtual Commissioning

Before the physical robot is installed, the digital-twin simulation environment is used to validate control programs, optimise process parameters, and test safety scenarios. 2024 industrial practice shows that virtual commissioning can reduce on-site commissioning time by 40–60 % and

significantly reduce collision and safety risk during commissioning. In production, the digital twin likewise enables "offline simulation validation" of new processes — testing new robot programs without halting production.

Real-Time Health-Index Tracking

The Health Index (HI) is the digital twin's principal output, fusing multi-dimensional sensor data into a single visualised indicator. HI is computed by weighted fusion of key indicators (percentage deviation of drive current from baseline, vibration energy, temperature-trend slope, ratio of operating hours to PM cycle) into a single value on the 0–100 scale:

> **Health-Index Definition HI = 100 means equipment fully matches commissioning baseline (as-new state); HI > 80: normal operating range, planned PM; HI 60–80: warning range, increase monitoring frequency; HI 40–60: trigger predictive-maintenance assessment, build a concrete maintenance plan; HI < 40: trigger emergency maintenance prescription, set priority by RUL prediction.**

Scenario Simulation and What-If Analysis

One of the most distinctive service values of a digital twin is supporting "what-if analysis": simulating in the digital environment the effect of different operating conditions or maintenance decisions on robot health. For example: if cycle speed increases 10 %, how much does reducer life shorten? If the grease-replacement interval is cut by 20 %, how much improvement in HI? Such analyses give quantitative economic basis to maintenance decisions, upgrading service from "rule-based execution" to "data-driven optimisation."

5.4 Remote Operations Data Infrastructure

The analytical quality of digital-twin and predictive-maintenance systems depends entirely on the quality of the underlying data infrastructure — garbage data in, garbage predictions out. This section defines the three core components of a robot remote-service data infrastructure: the multi-sensor data-acquisition layer, the data-preprocessing pipeline, and the data-governance framework.

5.4.1 Multi-Sensor Data Acquisition Strategy

Table 5.4-1 Sensor-Configuration Matrix for Robot Health Monitoring

Sensor Type	Monitored Target Parameter	Recommended Sampling Rate	Primary Diagnostic Value	Cost-Value Rating
Vibration Accelerometer	Joint / reducer vibration frequency and amplitude	≥ 1 kHz (diagnostic); 100 Hz (trend monitoring)	Early detection of reducer wear (characteristic-frequency analysis); bearing-fault identification; mechanical-resonance diagnosis	★★★★★ High value; most industry-recommended
Current Sensor (drive current)	Per-axis servo current (RMS and transient)	1–10kHz	Reducer-efficiency degradation (rising current); overload detection; motor-winding	★★★★★ Usually obtainable free via

			health assessment	controller API
Temperature Sensor (joint / ambient)	Reducer casing, controller interior, ambient temperature	1–10Hz	Cooling-fault early warning; grease-aging assessment; thermal-drift prediction	★★★★ Low cost, easy deployment
Joint-Position Sensor (encoder data)	Per-axis angle, velocity, acceleration	Controller native rate (typically ≥ 1 kHz)	Accuracy-degradation tracking; backlash estimation; trajectory consistency	★★★★★ Directly from controller, zero incremental cost
Torque Sensor (joint torque)	Per-axis output torque	1–100Hz	Reducer transmission efficiency; load-identification accuracy; cobot safe-torque compliance	★★★★ Mandatory on cobots; optional on industrial arms
Acoustic / Ultrasonic Sensor	Acoustic emissions from mechanical motion	20Hz–20kHz	Early detection of bearing wear; cavitation diagnosis (in liquid-cooled systems)	★★★ Specialised application, mid-cost
Power-Quality Sensor	Supply voltage, current harmonics, power factor	10kHz	Power-quality anomalies (EMC root cause for motor vibration); energy-efficiency tracking	★★★ Recommended in environments with unstable power quality

Best Cost-Start Recommendation For service organisations beginning digital-twin deployment, start with "zero-cost sensors": most modern controllers expose APIs for vibration estimates, joint currents, and encoder positions, yielding high-value diagnostic data with no added hardware. Using these free data to build baselines and simple trend monitoring is the lowest-risk path to validating digital-twin ROI. Add accelerometers and other hardware sensors after ROI is proven.

5.4.2 Data Preprocessing Pipeline

Raw sensor data cannot feed predictive models directly — noise, outliers, missing values, and unit inconsistencies severely degrade prediction accuracy. A standardised four-stage preprocessing pipeline is the core mechanism for ensuring data quality:

Phase	Processing Step	Method Description	Key Quality Indicator
P1	**Denoising**	Butterworth low-pass filter for high-frequency electrical noise; wavelet decomposition on vibration signals to retain fault-frequency components; sliding-median filter on current signals to remove impulse noise	SNR ≥ 20 dB; the base frequency is preserved after filtering
P2	**Normalisation**	Map data from different sensors to a common scale (0–1 or z-score); under varying operating conditions, use load normalisation to remove the effect of conditions on health indicators	Standard-deviation variation of normalised features across operating conditions < 10 %
P3	**Feature Extraction**	Time-domain features: RMS, crest factor, kurtosis, margin factor; frequency-domain features: FFT characteristic-frequency amplitudes, power spectral density; time-frequency features: STFT or wavelet-packet decomposition coefficients	Balance feature dimension against compute: 30–50 features typically cover the principal fault modes

P4	**Health-Indicator Construction**	Fuse the extracted features into single or multi-dimensional HI; use PCA or autoencoders for dimension reduction; construct a composite indicator that monotonically reflects equipment-health decline	Pearson correlation of HI with actual degradation ≥ 0.8; monotonicity index (MI) ≥ 0.7

5.4.3 Data Governance Framework

Robot-operations data touches on customers' production-process secrets, supply-chain information, and equipment capabilities, and is also the core asset for the service provider's predictive-model improvement. The data-governance framework must answer three key questions explicitly: who owns the data, who may access it, and how it is stored and protected.

- Data Ownership: operational data produced by the robot system (production cadence, process parameters, product information) belongs to the customer; the service provider may not use it for purposes beyond service improvement without explicit authorisation. The provider owns its analytic models and baseline databases. This boundary must be made explicit in writing in the service contract
- Data Classification and Access Control: robot data is tiered by sensitivity (L1 public technical parameters; L2 equipment-state data, provider-accessible; L3 production-process data, customer-only; L4 safety-related data, regulated); each tier has different storage requirements and access permissions
- Data-Residency Compliance: cross-border data transfer must comply with the local regulations of the data origin (China: Data Security Law, Personal Information Protection Law; EU: GDPR; US: sector-specific). Data from robots operating in China is in principle to be stored on servers inside China
- Data-Retention Policy: raw sensor data retained 6–12 months (storage cost vs analytic value); preprocessed feature data retained 3–5 years (long-term trends); model training datasets retained permanently; service records retained per the equipment-passport rules (recommended ≥ 10 years)

5.5 Predictive Maintenance and Prescriptive Maintenance

Predictive Maintenance (PdM) and Prescriptive Maintenance (PsM) represent two different levels of intelligent robot service: PdM answers "when will the problem occur," while PsM further answers "what should be done, when, and how to achieve the optimal outcome." 2024–2025 industrial research shows PsM becoming the new benchmark of Industry 4.0 maintenance practice — no longer a distant future vision.

5.5.1 Remaining Useful Life (RUL) Prediction Methodology

Remaining-Useful-Life (RUL) prediction is the core analytical task of PdM: given the equipment's current state and historical degradation trajectory, predict how much longer (in time or cycles) it can operate before reaching a maintenance threshold. The three method classes below represent current mainstream practice in industrial-robot RUL prediction:

Method 1 — Physics-Based RUL Prediction (White-Box)

Build degradation prediction models from the physical equations of the robot's mechanical system (kinematic equations, Hertzian contact-stress model, fatigue-life S-N curves). Physics-based models offer strong interpretability (engineers can understand the physical basis of the prediction), good generalization (reasonable predictions on unseen operating conditions), and low training-data requirements. Typical applications: tooth-flank fatigue-life prediction for harmonic-drive reducers (based on the Hertzian contact-stress model); rolling-element fatigue-life prediction for bearings (based on the L10 life equation). Main limitation: accuracy is bounded by the precision of the physical-model assumptions, and complex multi-factor degradation mechanisms are difficult to model completely.

Method 2 — Data-Driven RUL Prediction (Black-Box)

Use ML models to learn degradation patterns directly from sensor data without explicit physical equations. Current best-performing methods include:

Model Type	Technical Features	Applicable Robot-Service Scenarios
LSTM (Long Short-Term Memory)	Excellent at time-series data; learns long-range dependencies; strong at modelling temporal degradation patterns in sensor data	Reducer-grease aging trend; motor-winding cumulative-temperature effect; joint-current-drift RUL estimation. 2025 published research validates LSTM with < 5 % MAE in industrial-robot PdM
Random Forest	Low feature-engineering dependence; robust on small samples; computationally efficient; native feature-importance output	Suited to early-stage prediction when data is scarce (hundreds of samples can be effective); feature importance helps engineers see which sensors matter most
Physics-Informed Neural Network (PINN)	Embeds physical constraints into the loss function; performs best when physics is known but data is scarce; predictions respect physical constraints and avoid physically impossible outputs	Reducer-RUL prediction (physics gives constraints, data gives individualised calibration); the most actively researched method in 2024–2025 and the optimal fusion of white-box and black-box
Autoencoder	Unsupervised, no labelled fault samples needed; detects anomalies through reconstruction error; especially suited to newer robots lacking historical fault data	Anomaly detection in robotic systems; especially suited to humanoid robots and newer AMR types lacking historical fault data

Method 3 — Physics-Informed Hybrid Model (Grey-Box, Current Best Practice)

Combines the interpretability and generalisation of physics with the accuracy advantage of data-driven models: physics supplies prior knowledge and boundary constraints; data-driven models correct the residual of the physical model's prediction. Multiple 2024 studies confirm that the grey-box approach improves RUL accuracy by 30–50 % over pure physics models in industrial-robot prediction, while exhibiting stronger cross-equipment generalisation than pure black-box methods — the recommended industrial practice.

5.5.2 Predictive Maintenance Implementation Pathway (Four Phases)

Phase	Name	Principal Work Content	Maturity Judgement Indicator
Phase 1	**Data Foundation**	Deploy sensors (start from controller-API currents and positions); build the data-preprocessing pipeline; establish the commissioning baseline dataset; define the data-storage and access architecture	Acquisition completeness ≥ 95 % across all target machines; baseline dataset established

Phase 2	**Rule-Based Monitoring**	Build threshold-alarm rules against commissioning baselines (e.g., "current 20 % above baseline triggers warning"); set alarm-severity tiers and notification flow; engineers review trend reports periodically	False-alarm rate < 15 %; all alarms responded to within the defined time
Phase 3	**Machine-Learning Prediction Models**	Train RUL prediction models on accumulated history (typically 6–12 months of normal-operation data); validate accuracy (MAPE < 15 % for deployment); set up a model-update mechanism	MAPE < 15 %; the RUL error window fits within acceptable maintenance-planning latitude
Phase 4	**Prescriptive Maintenance Closed Loop**	Integrate RUL prediction with the maintenance-scheduling system to auto-generate prescriptions for planning; integrate with the spare-parts supply chain to pre-trigger procurement; establish continuous prediction-accuracy assessment and model iteration	Prescription execution rate ≥ 85 %; planned-to-emergency-maintenance ratio ≥ 0.90 (world-class)

5.5.3 Prescriptive Maintenance

Prescriptive maintenance is the next evolution beyond predictive. If predictive maintenance tells the engineer "when the problem will occur," prescriptive maintenance further outputs the concrete action recommendation "what to do, when to do it, and how to do it for the optimal outcome (lowest cost, highest availability, shortest downtime)."

> **Prescriptive-Maintenance Output Example "System detection: J3 reducer current HI = 62 (warning band); RUL prediction 450 ± 80 hours (85 % confidence). Combined with current production schedule, the optimal maintenance window is day 18 (Saturday night-shift idle period). Recommended action: J3 grease replacement (estimated 2 h labour, cost - USD 180) + vibration-baseline retest. If not executed within the optimal window, HI will drop below 40 between days 23–27, then requiring full overhaul (estimated 8 h labour, cost - USD 2,400). Please confirm the maintenance plan."**

Implementing prescriptive maintenance requires integration of three subsystems: the predictive-analytics engine (RUL and HI outputs); the production-scheduling system (optimal maintenance windows); the spare-parts supply-chain system (inventory and procurement cycles). Integration complexity is high — defer the Phase-4 integration until the Phase-3 prediction model is mature.

5.6 OTA Firmware and Software Management

OTA (Over-The-Air) update is an important delivery means of remote service: version management and update delivery for controller firmware, application packages, and AI inference models over a secure remote channel. Unlike IT-software updates, robot-controller OTA carries three special risks: (1) update failure can leave the controller in an un-bootable state; (2) the wrong firmware version may introduce safety-function defects; (3) robots in production cannot accept the traditional "reboot then verify" update model. A standardised OTA-management process is the systemic response to these three risks.

5.6.1 Technical Requirements for Secure OTA Updates

▸ Code Signing

Every OTA update package must carry a digital signature from a trusted CA; the controller verifies signature validity before installation; unsigned or invalid packages are auto-rejected and raise a security alarm. This is the foundational mechanism against tampering and supply-chain attack, and an explicit IEC 62443-4-1 requirement for the software development lifecycle.

▸ Secure Boot

At controller boot, firmware-image integrity is verified (cryptographic hash comparison); on failure, the controller falls back to the last known-good version rather than refusing to boot (to avoid line stoppage). Secure Boot minimises the firmware-level supply-chain attack surface.

▸ Transport-Layer Encryption

OTA package transmission must use TLS 1.3 or above for end-to-end encryption; transmission over unencrypted HTTP is forbidden; download integrity must be verified by SHA-256 or stronger.

▸ Least-Privilege Update Packages

OTA updates should use delta updates rather than full updates: only changed portions are transmitted, reducing bandwidth (friendly to bandwidth-limited factory networks) and shrinking the attack surface. The trade-off is greater patch-management complexity, requiring full version-dependency tracking.

5.6.2 Tiered Rollout Strategy for OTA Updates

Immature OTA pushes new firmware to all machines simultaneously — this strategy, if the firmware has an unknown defect, takes down the entire fleet at once. Staged rollout is the standardised method for managing update risk:

Rollout Batch	Target Scope	Duration	Pass Criteria	Failure Action
Batch 1 — Canary	1–2 non-critical production machines (dedicated test machines or low-priority lines)	48–72 h observation	Zero severe alarms; all safety-function tests pass; no degradation of key performance indicators	Roll back immediately to the previous version; suspend all further rollout; escalate to the OEM
Batch 2 — Small-Scope Validation	5–10 % of target fleet (covering different platforms and application types)	7-day observation	Same as Batch 1; on-site engineer confirms functional completeness	Depending on fault scope: local issues fixed and re-pushed; systemic issues — full rollback
Batch 3 — Main Deployment	50 % of target fleet (prioritising non-critical production machines)	7–14 days	Batch 2 data shows zero new issues; Batch 1 monitoring data stable	Issues must be traceable to a specific platform or configuration combination; precise fix then continue
Batch 4 — Full Deployment	All remaining target machines	Push during planned maintenance windows	Batch 3 meets all pass criteria	Rare issues at this stage require a dedicated contingency plan

5.6.3 Rollback Mechanism Design

- Dual-Partition Storage (A/B Partitions): controller storage carries two firmware partitions — currently running (A) and backup / new (B). Updates write to B; after verification, the boot partition is switched;

on verification failure, the system continues to boot from A and the B update is discarded. This is the most reliable rollback mechanism, independent of update-process integrity

- Auto-Rollback Triggers: on the first post-update boot, automatic rollback if any of the following is detected: safety-function test failure (E-Stop, STO); critical bus failure (EtherCAT slaves do not enter OP); boot time exceeds 200 % of normal (typically indicating firmware anomaly)
- Manual Rollback: engineer-initiated remote rollback requires double confirmation (to prevent misclicks); rollback produces full logging (engineer ID, time, before-and-after versions); upon completion, notify the customer with a written explanation

5.7 Edge-Cloud Collaborative Computing Architecture

Remote-service data-processing tasks should not all be done in the cloud, nor all at the edge — "edge-cloud collaboration" is the current best industrial architectural pattern. Each paradigm has its strengths; the layer at which each task runs should be decided by its real-time requirements, data volume, and computational complexity.

5.7.1 The Role of Edge Computing

Task Type	Reason to Execute at the Edge	Typical Edge-Computing Tasks
Real-Time Control Response	Cloud round-trip latency (typically 100–500 ms) cannot meet real-time-control needs (typically < 1 ms)	Hard real-time response of safety functions (E-Stop triggering); local closed-loop servo control; real-time collision detection
Local Anomaly Detection	Anomalies require immediate response (cannot wait for cloud confirmation); local detection latency < 10 ms	Immediate stop on threshold exceedance; real-time alarm on vibration spikes; safety-boundary monitoring
Data Preprocessing and Compression	Preprocess raw sensor data at the edge and transmit only valuable features, reducing bandwidth by 80–95 %	Signal denoising and feature extraction; data compression (transmit ~ 5 % feature data to cloud); quality filtering (reject bad samples)
Local Model Inference	Deploy lightweight AI inference (TensorFlow Lite or ONNX) on the edge gateway, with no dependence on cloud connectivity	Local anomaly scoring (real-time computation); real-time local HI update; graceful degradation during network outage

5.7.2 The Role of Cloud Computing

- AI Model Training: large ML models (LSTM, PINN, etc.) must be trained on GPU clusters; the compute load exceeds edge-gateway capability (typical edge gateway: 4–8 CPU cores; training requires GPU clusters)
- Cross-Device Data Fusion: aggregate data from multiple same-type machines to build a generic degradation baseline and enable cross-device transfer learning (transferring fault knowledge from machine A to machine B's predictive model)
- Long-Term Historical Analysis: trend analysis over years of sensor history, seasonal-pattern recognition, and long-cycle degradation modelling — large storage (TB-class) and complex computation — suit cloud batch processing
- Simulation and Scenario Analysis: high-fidelity digital-twin simulation (FEA, collision detection, multi-body dynamics) requires compute well beyond the edge gateway and must run in the cloud

- Global Service Coordination: cross-region, cross-time-zone multi-machine scheduling, global spare-parts inventory optimisation, and aggregated service-performance analysis must be handled centrally in the cloud

5.7.3 5G Applications in Robotic Remote Service

5G's three signature features — enhanced mobile broadband (eMBB), ultra-reliable low-latency communication (uRLLC, < 1 ms), and massive machine-type communication (mMTC) — match the needs of robot remote service closely and are becoming the key network infrastructure for S3 / S4 remote service:

5G Capability	Specification	Robot-Remote-Service Use Cases
uRLLC (Ultra-Low Latency)	< 1 ms air-interface latency; < 10 ms end-to-end	Remote real-time control: engineers operating robots over 5G from hundreds of kilometres away with latency below local-LAN perceptual threshold; AR remote collaboration (engineer gestures synced live to on-site AR glasses)
eMBB (Enhanced Mobile Broadband)	Peak downlink 20 Gbps; peak uplink 10 Gbps	High-definition video streaming (4K / 8K live site video for remote diagnosis); high-rate sensor-data backhaul (parallel monitoring of many machines); high-frequency state sync of digital twins
mMTC (Massive Machine-Type Communication)	≥ 1 million connections per km²	Simultaneous online monitoring of thousands of robots in a large plant; dense deployment of IIoT sensor nodes (temperature, vibration, position); no need to wire each sensor individually
Network Slicing	Virtual private network (VPN-level isolation)	Allocate independent 5G slices for robot control data (real-time critical, must be isolated) and video data (bandwidth-priority); physical-level isolation of critical control from general business data improves security

5.8 Remote Service Delivery — End-to-End Process

Remote-service delivery must be more structured than on-site service — without physical presence, the engineer cannot rely on physical perception (sound, sight, vibration) to fill information gaps, and every judgement must rest on digital data. A standardised process is the core mechanism for compensating for this information limitation.

5.8.1 Active-Monitoring-Triggered Remote Service Workflow (S3/S4 Stages)

Step	Node	Core Operation	Outputs and Deliverables
M1	**Monitoring-Alarm Triggered**	The state-monitoring system detects HI below the warning threshold (HI < 80) or a specific sensor value outside the normal range; the system auto-generates an alarm and notifies the responsible engineer	Alarm notification (device, parameter, current vs threshold, historical trend chart)
M2	**Alarm Classification and Priority Confirmation**	The engineer receives the alarm and confirms priority (P1–P4) based on type and historical context; P3 / P4 alarms enter planned handling; P1 / P2 enter the emergency remote-response flow	Priority confirmation; response-time commitment; for P1, immediately notify the customer's on-site safety lead
M3	**Deep Remote-Data Analysis**	Pull the target device's full historical data via the digital-twin platform; perform trend	RUL prediction report (with confidence interval); HI trend

		analysis (past 30/60/90 days); compare to peer-device baseline; run the RUL model	chart; preliminary maintenance recommendation
M4	**Customer Communication and Plan Confirmation**	Send the analysis and recommendation to the customer in a standardised report format; confirm the optimal maintenance window with the customer (aligned with production schedule); if on-site work is required, launch dispatch planning	Maintenance-prescription confirmation form (customer-signed); planned maintenance date; reserved spares
M5	**Remote Intervention (if applicable)**	For issues resolvable remotely (parameter tweak, firmware update, configuration fix), perform and log the remote operation; for issues requiring on-site work, prepare technically (pre-stage spares, package diagnostic data, write the action plan)	Remote-operation log; on-site dispatch technical pack (diagnostic conclusion and action recommendation)
M6	**Resolution Verification and Report Closure**	After remote intervention, monitor for 48 hours and confirm HI is trending up; after on-site work, remotely verify that key parameters return to normal; produce monthly / quarterly state reports	Resolution-verification report; state report (sent to customer management); maintenance record updated in the equipment passport

5.8.2 AR-Enhanced Remote Collaboration (Emerging Capability)

AR remote collaboration is the fastest-maturing emerging service means in Domain 3: the remote expert engineer sees the field engineer's (Professional's) view live via video call, and uses AR overlay to annotate operating instructions directly in the field engineer's view ("remove this bolt" → arrow overlaid on the actual location). This mode has significant value in the following scenarios:

- First service of a new robot model — the field Professional lacks experience and the Expert tier can guide live without travelling
- Cross-language technical support: the Expert speaks their native language and the AR system live-translates and overlays text in the field engineer's view
- Real-time supervision of complex assembly / disassembly: the Expert confirms each step of the Professional's operation in real time, reducing mis-operation risk
- Network requirements for AR remote collaboration: 5G or low-latency Wi-Fi 6 (latency < 30 ms, bandwidth > 10 Mbps); otherwise video-stream latency mis-registers AR annotations to actual objects, increasing operational difficulty

5.9 Domain 3 Key Performance Indicator Framework

Indicator	Definition	S2 Target	S4 Target	Management Significance
Remote Response Rate (RRR)	Work orders resolved remotely without on-site visit / total work orders	**≥ 40%**	**≥ 70%**	Measures remote-service maturity; higher means lower on-site cost
Mean Remote-Diagnosis Time (MRDT)	Average time from remote-connection establishment to root-cause localisation	**< 2 hours**	**< 45 minutes**	Composite measure of digital-twin data completeness and the engineer's remote-diagnostic capability
Prediction Accuracy	MAPE of RUL prediction	**Not applicable (predictive models**	**< 15%**	Measures predictive-model quality; continuously tracked

		not deployed in S2)		to drive model iteration
Planned-Maintenance Share	Prediction-triggered planned maintenance / total maintenance work orders	**< 40%**	**≥ 85%**	Core indicator of the service-mode transformation; 85 % at S4 represents a very high level of proactive service
Data-Acquisition Completeness	Actual data points acquired / data points expected	**≥ 90%**	**≥ 99%**	Data-infrastructure stability; low completeness directly degrades predictive-model quality
OTA-Update Success Rate	OTA updates succeeding on first push / total OTA pushes	**≥ 95%**	**≥ 99.5%**	OTA-process maturity; a low success rate implies high rollback risk
False-Alarm Rate	Number of false alarms / total alarms	**< 20%**	**< 5%**	A high false-alarm rate causes alarm fatigue and risks engineers ignoring real warning signals
Digital-Twin Model Drift	Persistent deviation between digital-twin predicted state and actual measurement	**Not applicable (model not yet deployed)**	**< 5%**	Monitors long-term twin-model accuracy; drift beyond threshold triggers retraining

5.10 Chapter Summary and Inter-Domain Interfaces

Chapter Summary Domain 3 (Remote Service & Digital Operations), structured around the four-stage evolution model, progressively advances robot-service delivery from reactive (S1) to predictive-and-prescriptive (S4). The secure remote-access architecture (IEC 62443 Zone-Conduit model + zero-trust) is the security foundation for every remote capability.

The Robot Digital Twin (RDT four-layer architecture) maps physical-equipment state in real time into the digital environment, supporting state assessment, scenario simulation, and maintenance prescription. Multi-sensor data infrastructure with a standardised preprocessing pipeline secures data quality. The physics-informed hybrid (grey-box) model represents current best practice in RUL prediction, with prescriptive maintenance as the next evolutionary target. The staged-rollout policy and dual-partition rollback minimise OTA-update risk. The methodology of this chapter is grounded in the latest 2024–2025 academic research and industrial practice and applies to robots of any brand.

Interfaces with Other Capability Domains

Interface Direction	Interfacing Domain	Content
Domain 1 → Domain 3	Service Lifecycle Management	Domain 1's equipment passport (commissioning baselines) is the initial parameter input for Domain 3's digital-twin modelling; Domain 1's PM records are key labelled data for training predictive models ("which data corresponds to which health state")
Domain 3 → Domain 1	Service Lifecycle Management	Domain 3's RUL predictions are the data basis for Domain 1's dynamic PM-strategy adjustment; predictive-maintenance triggering directly replaces Domain 1's fixed-time PM triggers
Domain 3 → Domain 2	Fault Diagnosis & Technical Support	Domain 3's continuous-monitoring data provides historical context for Domain 2's intermittent-fault diagnosis ("sensor data from the 100 hours preceding the fault"); the scenario-replay capability of the digital twin supports Domain 2's root-cause analysis.

Domain 3 → Domain 4	SLA Design & Service Commitment	Domain 3's remote-response rate and prediction-triggered maintenance data form the technical basis for Domain 4's SLA-clause design; digital operations capability directly supports higher availability commitments.
Domain 3 → Domain 6	AI/ML System Service Methods	Domain 3's multi-sensor data infrastructure is the data channel for Domain 6 AI-model-performance monitoring; Domain 3's OTA-management framework is the mechanism for Domain 6's AI-model update deployment

Chapter 6

Domain 4 · SLA Design & Service Commitment

6.0 Chapter Introduction — From Maintenance Contract to Service Commitment

In the perception of most industry participants, the "robot service contract" is still a repair-rate sheet plus a response-time commitment. That perception is becoming a bottleneck for the industry's development.

Customers are evaluating their robot investments by a completely different logic: they care less about how quickly an engineer arrives than whether the production line's actual availability meets the assumption in their investment case; not what response time the contract specifies, but how many of the year's 8,760 hours the robot is actually creating output for them. When a robot's acquisition cost runs between USD 30,000 and USD 300,000 and the hidden loss of an hour's downtime can be USD 10,000 to 50,000, the customer's "quantitative expectation" of service quality has long since crystallised — most service providers simply have not yet responded in the same language.

Domain 4 (SLA Design & Service Commitment) fills exactly this gap of perception and language. This chapter builds the SLA system from three perspectives: §6.1 exposes the essential differences between robotic and IT SLAs and the current industry gap; §6.2 establishes the RSF robotic SLA metrics framework (with OEE at its centre, rather than availability alone); §6.3 designs the four-tier SLA product architecture so that providers can deliver structured commitments differentiated by customer; §6.4 examines specifically how new entrants design credible, sustainable SLAs under limited capability; §6.5 builds the SLA performance-tracking and compensation mechanism; §6.6 looks ahead to how the RaaS (Robots-as-a-Service) model restructures the SLA system; §6.7 defines Domain 4 KPIs; §6.8 closes with chapter summary and inter-domain interfaces.

6.1 Industry Status — How Supplier Commitments and Customer Expectations Are Evolving Together

Before systematically designing the SLA, it is worth grounding the work in an honest reading of where the industry stands today. The robotics-service market is at an important inflection point: the language customers use to evaluate service quality — ROI, OEE, TCO, availability targets — has matured considerably, and service commitments are naturally evolving to meet it. This evolution is a shared development task for the whole industry, not a failure on anyone's part, and it is precisely the kind of structured progress that RSF Domain 4 aims to support.

6.1.1 Service-Commitment Status of Established OEMs

Taking the leading industrial-robot OEMs (FANUC, ABB, KUKA, Yaskawa, etc.) as the reference, the typical structure of a standard service contract is as follows:

Contract Tier	Typical Commitment Content	Real Customer Demand (Unmet Part)
Basic Warranty (1–2 years)	Hardware-defect warranty; 24-hour technical hotline; next-business-day (NBD) response time	No uptime (Availability) metric constraint; no Service Credits mechanism for service delays; accuracy degradation and cumulative thermal drift-induced process failures are typically excluded from warranty

		coverage.
Maintenance Contract (annual fee)	Preventive-maintenance delivery (PM plan); technical hotline; dedicated account manager; priority parts supply	Still a "service-behaviour" commitment, not a "service-outcome" commitment. Contracts typically contain no quantitative availability metric, no MTTR (Mean Time to Repair) target, and no associated KPI Action performance tracking.
Premium Service Package	Four-hour response, remote monitoring, parts pre-staging, extended warranty	OEMs such as KUKA advertise "maximum uptime" on their websites, but contract text is typically not linked to service-fee discounts (lacking financial enforceability); response time is almost universally defined as "Time to Arrive (engineer on-site)" rather than "Time to Restore (fault resolved)."
RaaS / Subscription Model (emerging, some OEMs)	Usage-based or subscription fee; equipment ownership remains with the service provider; provider monitors and maintains proactively	The model closest to a "results commitment." However, it remains rare in the traditional industrial-arm market due to heavy asset ownership; more prevalent in the AMR space. The deployment standards for embodied-AI and software-defined robots in the after-market remain in a standards vacuum.

Industry Context The SLA frameworks that mature OEMs have built centre on service-behaviour commitments — response time, on-site arrival time — which were the right design choice for the industry's earlier stage of development. As customers' procurement language has matured toward result-oriented metrics (availability targets, OEE, TCO), the natural next step is for service commitments to evolve in parallel — from "I will dispatch an engineer" toward "your equipment will achieve availability X%." This is a shared evolution for the whole industry, and RSF Domain 4 is designed to provide the structured methodology that makes that transition practical for service providers of all sizes.

6.1.2 Quantitative Standards of Customer Expectation

Customer expectations of robotic-service quality have, across industries, evolved from "timely response" to specific quantitative KPI targets. The data below come from industry benchmark research and public procurement requirements of major end users:

Industry	Typical Robotic Application	Customer-Expected Availability	Actual Economic Impact of Downtime
Automotive Manufacturing	Welding / assembly / material handling	**≥ 98 % (some lines require 99 %)**	Critical welding-line downtime cost: USD 10,000–50,000 / hour; full final-assembly-line stoppage: USD 300,000+ / hour
Electronics / Semiconductor	SMT placement / packaging / inspection	**≥ 99%**	Cleanroom-robot downtime affects yield and cycle time; downtime cost USD 5,000–20,000 / hour
Logistics / Warehousing	AMR sorting / palletising	**≥ 95 % (higher in peak season)**	Peak-season downtime directly affects order-fulfilment rate; e-commerce peak-season delays are worth tens of thousands of yuan per hour
Food and Beverage	Packaging / filling	**≥ 95%**	Downtime breaks production batches; most plants cannot accept unplanned downtime exceeding 2 hours
Cobots (SMEs)	Machine tending / assembly	**≥ 90%**	SMEs have limited resources; downtime over 4 hours forces temporary manual back-up and

			significantly raises operating cost

6.1.3 The SLA Dilemma Facing New Entrant Robot Companies

The new-entrant robot companies that flooded the market over the past five years (cobot start-ups, AMR companies, humanoid-robot companies) face a severe dilemma in SLA design:

- **Aggressive SLA — Win the Order but Cannot Deliver**

  To win customers (especially against mature OEMs), new entrants tend to commit to aggressive SLA terms (98 % availability, four-hour response). The problem: they lack the OEMs' global service network, local parts inventory, decades of MTBF data, and financial cushion to absorb SLA-breach compensation. Such commitments are, in effect, trading future service credit for current orders; once breaches scale, the company's survival is fundamentally threatened.

- **Conservative SLA — Compliant but Loses the Market**

  If a new entrant designs a conservative SLA grounded in real capability (e.g., "we provide commercially reasonable efforts rather than quantitative commitments"), it is severely disadvantaged in competition with mature OEMs — customers have no reason to choose a "we'll try but not guarantee" new brand over an established supplier with a global service network.

The way through this dilemma is not to choose between "aggressive" and "conservative" but to design a tiered SLA system aligned with real capability and able to upgrade as capability grows. This is exactly Domain 4's core methodological contribution (see §6.3).

6.2 Robot SLA Metrics Framework — OEE as the Core

Robotic-SLA metrics design must start from the customer's perspective — "is the robot service effective" — not the supplier's perspective — "did the engineer arrive on time." The fundamental difference: supplier-view indicators (response time, arrival time) can be controlled unilaterally by the supplier; customer-view indicators (availability, OEE) are the joint outcome of supplier capability and equipment state. A high-quality SLA should cover both — supplier-view indicators ensure the regularity of service behaviour; customer-view indicators ensure the value of service outcomes.

6.2.1 Why Robot SLA Cannot Be Imported Wholesale from IT SLA

Dimension	IT-Service SLA Logic	Robotic-Service SLA Logic
Core Availability Indicator	System uptime (%) — share of time the server or software is online	Equipment availability (%) is the base, but more important is OEE (Overall Equipment Effectiveness) = availability × performance × quality
Fault-Recovery Method	Typically fully recoverable remotely (restart, rollback, failover); RTO and RPO are core indicators	Physical repair cannot be skipped; reducer replacement takes 2–8 hours; parts wait can be 1–5 days; RTO corresponds to MTTR
Downtime-Loss Quantifiability	Mainly data-loss risk (RPO) and business interruption (RTO)	Production-downtime loss is directly quantifiable: output loss × product margin or absorbed fixed cost, typically USD 10,000–50,000 per hour
Effect of Maintenance on	Planned-maintenance windows typically run during off-peak	Robot PM must be executed within production downtime windows; whether "downtime" during PM

Availability	periods, with minimal user impact	counts in the SLA is an important contract-design point
SLA-Breach Compensation	Typically a small service-fee credit (e.g., 10 % of the monthly fee); rarely covers actual loss	Real downtime losses far exceed the service fee; the compensation mechanism must be carefully designed (or the supplier faces compensation risk far exceeding service revenue)

6.2.2 RSF Core Robot SLA Metrics Framework

The RSF robotic-SLA metrics framework has three classes: outcome indicators (measuring service contribution to customer business objectives); process indicators (measuring regularity of service delivery); and data indicators (measuring the information quality of the SLA system itself). The three together form the complete measurement framework.

Table 6.2-1 RSF Robotic-SLA Core Metrics Definitions

Class	Indicator	Formula	Industry Benchmark	Notes
Outcome	**Equipment Availability**	(planned hours − unplanned downtime) ÷ planned hours × 100 %	**≥ 95 % (general industry); ≥ 98 % (automotive line)**	The definition of "unplanned downtime" must be explicit in the contract (includes / excludes: software issues, operator errors, raw-material issues)
	OEE (Overall Equipment Effectiveness)	Availability × Performance × Quality (all percentages)	**≥ 75 % (industry average); ≥ 85 % (world-class)**	Including OEE in the SLA requires joint management of non-equipment factors (materials, operators) with the customer; usually only applies in high-maturity service relationships
	MTTR (Mean Time to Repair)	Total repair time ÷ number of repairs	**< 4 hours (critical assets); < 8 hours (general equipment)**	The start and end definitions of repair time are critical: typically from "fault work order opened" to "equipment restored to normal operation and confirmed by customer"
	Annual Maximum Downtime	Annual total unplanned downtime (hours)	**≤ 438 hours / year (corresponding to 95 % availability)**	More intuitive than an availability percentage; directly linked to business-loss estimation
Process	**Remote Response Time (RTR)**	Time from fault-work-order opening to the engineer establishing a remote connection	**P1: ≤ 15 minutes; P2: ≤ 1 hour; P3: same day**	Distinguish "remote response" from "arrival"; remote response should always precede arrival
	On-Site Arrival Time	Time from fault-work-order opening (or customer confirmation that an on-site visit is required) to the engineer's arrival on site	**P1: ≤ 2 hours (same city); P2: ≤ 4 hours (same city); cross-city: as contracted**	Specify whether the clock starts at "work order opened" or "customer authorises on-site visit"; distinguish business hours from 24/7
	PM-Compliance Rate	PMs executed as planned ÷ total planned PMs	**≥ 90%**	PM delays (not customer-caused PM postponements) count against the rate; postponements require advance notice to the

				customer and recorded reason
	First-Time Fix Rate (FTFR)	Work orders with correct first-visit root-cause localisation ÷ total work orders	**≥ 75%**	Reflects engineer capability; low FTFR causes repeat visits and raises the customer's hidden downtime cost
Data	**SLA Data Availability**	Share of SLA indicators in the month with supporting data	**100 % (compliance requirement)**	An SLA indicator without data has no legal effect; the data-collection system must be built before contract signing
	Report-Delivery Timeliness	Share of monthly SLA reports delivered before the agreed date	**100 % (compliance requirement)**	Report content is pre-agreed (minimum fields: each indicator's actual value, target value, variance explanation, improvement plan)

6.2.3 Service Controllability Analysis of OEE (Critical Design Consideration)

OEE is the most accepted composite indicator of equipment performance in manufacturing, but in SLA design "controllability" must be handled very carefully: the supplier should only be SLA-responsible for OEE losses it can directly control.

OEE Loss Category	Loss Source	Supplier Controllability	SLA Handling Recommendation
Unplanned Downtime (Availability Loss)	Robot-equipment fault, controller fault, safety-function failure	**High — directly corresponds to MTTR and response time**	Include in SLA, with explicit MTTR targets and availability commitments
Planned-Maintenance Downtime (Partial Availability Loss)	Production time consumed by the supplier's PM execution	**Medium — PM time window and efficiency can be optimised**	PM time should count as "planned downtime," not as an SLA breach, but PM duration must be controlled
Performance-Efficiency Loss (Speed Loss)	Robot runs below rated speed (caused by wear)	**Medium — strongly correlated with PM quality**	May be partly included in SLA (speed-baseline maintenance), provided a commissioning-period speed baseline is recorded
Quality Loss (Defect Rate)	Robot-accuracy degradation produces non-conforming product	**Low — beyond accuracy maintenance, strongly correlated with process parameters and raw-material quality**	Handle with caution; usually included in SLA only when robot accuracy can be measured independently, otherwise responsibility cannot be attributed
Operator-Caused Downtime	Operator misoperation, unauthorised parameter modification	**None — customer responsibility**	Explicitly excluded from SLA scope, but must be defined in the contract with a work-order recording mechanism in place
External Factors (Force Majeure)	Power outage, supply-chain disruption, natural disaster	**None**	An explicit Force Majeure clause in the contract

6.3 RSF Four-Tier SLA Product Architecture

RSF establishes a tiered SLA product architecture that lets suppliers offer differentiated commitments based on customer needs, equipment type, geographic coverage, and the supplier's own capability.

The four tiers progress from basic coverage to comprehensive assurance, each with explicit capability requirements and commercial pricing logic. This tiering simultaneously solves two problems: for customers, it provides comparable, selectable transparent options; for suppliers, it avoids one-size-fits-all high-risk commitments.

Table 6.3-1 RSF Four-Tier SLA Product Architecture

Tier	Product Name	Core Commitment	Service Coverage	Applicable Scenarios and Pricing Logic
L1	**L1 Basic Response**	Response-time commitment (behaviour commitment, no outcome guarantee)	Business hours (8×5); standard response time; no parts pre-staging	Use: SME customers, non-critical machines, capability-limited new entrants; pricing benchmark: 2–4 % of original equipment value per year; this is every supplier's "minimum entry point" and should not fall below it
L2	**L2 Standard Assurance**	Availability target ≥ 92 %; MTTR ≤ 8 hours; P1 on-site response ≤ 2 hours	Business hours + extended hours; basic parts pre-staging; PM-plan coverage	Use: most industrial applications, lines of moderate annual output; supplier must have Professional + Specialist engineer teams; pricing benchmark: 4–7 % of original equipment value per year
L3	**L3 Priority Assurance**	Availability target ≥ 95 %; MTTR ≤ 4 hours; P1 on-site response ≤ 2 hours; 24/7 coverage	24/7 response; critical-parts pre-staging; periodic state reports; basic remote monitoring	Use: critical production lines, multi-shift operation, customers introducing availability KPIs; supplier needs 24-hour on-call engineers and local parts depot; pricing benchmark: 7–12 % of original equipment value per year
L4	**L4 Comprehensive Assurance**	Availability target ≥ 98 %; MTTR ≤ 2 hours; quantitative OEE target; breach-compensation clauses	24/7 dedicated support; complete parts pre-staging; predictive maintenance; monthly SLA reports; annual joint review	Use: automotive lines, semiconductor fabs, high-OEE scenarios; full digital-twin + predictive-maintenance infrastructure required; pricing benchmark: 12–20 % of original equipment value per year; supplier must have Expert-tier engineers and scaled service-operations capability

6.3.1 Core Parameter Design Across the SLA Tiers

Precise Definition of Response Time (Three Key Points to Avoid Ambiguity)

- The "start" of response time: count from "customer submits the fault work order (or monitoring system auto-opens one)" — not from "customer phone connected" or "supplier internally dispatches engineer." The gap can be 30–60 minutes — significant for P1 faults
- Distinguishing response types: "remote response" (engineer establishes the remote connection) should precede "arrival" (engineer reaches the physical site); higher-tier SLAs should define targets for both separately; "arrival time" cannot substitute for "remote response"
- Explicit definition of service windows: "business-hours response (8×5)" must specify the hours (e.g., Mon–Fri, 08:00–18:00, excluding statutory holidays); "24/7 response" must specify whether statutory holidays are included; fault response outside the service window must be separately agreed

Boundary Definitions for the Availability Target Calculation

The availability-target number itself is less important than the calculation-boundary definition — under the same "95 % availability commitment," different boundary definitions can yield permitted downtimes differing by several-fold:

Calculation Parameter	Loose Definition (Supplier-Favourable)	Strict Definition (Customer-Favourable)	RSF-Recommended Balanced Definition
Reference Time Base	Whole contract year (8,760 hours)	Customer's actual planned production time	Customer's actual planned production time (record annual production schedule in the contract as the baseline)
Planned-Downtime Treatment	All PM downtime excluded from the denominator (and not deducted from the numerator)	All downtime included	Supplier-led PM: not counted as breach (window pre-agreed); customer-led shutdown: not counted in the denominator
Downtime-Recognition Criterion	Robot completely non-operational (zero output)	Performance below 80 % of rated value is counted as downtime	Counted as downtime when the robot cannot run the production program normally; speed reduction is recorded but does not equal downtime
Reporting Period	Annual (full-year calculation, no monthly commitment)	Monthly (each month calculated independently)	Monthly calculation with annual accumulation; monthly breach triggers service credits, serious annual breach triggers compensation clauses

6.4 SLA Design and Assurance Strategies for New Entrant Robot Companies

Dozens of new-entrant robot companies emerge each year (AMR start-ups, new cobot makers, humanoid-robot companies). Facing the decades of accumulated service infrastructure at the mature OEMs, how should they design credible, sustainable SLAs? The answer is not to compete with the OEMs on the same dimensions, but to follow a two-track strategy of "differentiated positioning" and "progressive build-up."

6.4.1 Core SLA Challenges for New Entrant Companies

Challenge Dimension	Concrete Manifestation	Implication for SLA Design
Reliability-Data Scarcity	New products lack historical MTBF data; cannot support availability commitments with real failure data	Cannot base availability commitments on data; must use conservative design plus a continuous-learning mechanism
Thin Service Network	Limited number of service engineers; geographic coverage restricted (typically concentrated in a few cities)	Arrival-time commitments must be realistic (some regions may be 4–8 hours, not 2); a partner-service-provider network must be built
Unstable Parts Supply	Small-batch production raises parts-inventory cost; supply-chain risk is high (critical parts single-sourced)	Parts-delivery SLA commitments must be conservative; a safety-stock strategy must be established
Limited Financial Cushion	Start-up funds go mainly to R&D and expansion; large SLA-breach payouts can threaten company survival	Compensation clauses require liability caps (typically a specific fraction of the affected-period service fee, not full production-loss compensation)
Rapid Technology-Maturity Evolution	Software / firmware updates are frequent; each may introduce new stability issues	OTA updates must be incorporated into SLA management (post-update stability commitments); rapid rollback is required

6.4.2 SLA Strategy for New Entrants — Eight Actionable Recommendations

Strategy ① Start at L1, build a real credit record, and progressively upgrade tiers

A new entrant's first commercial customers are the most critical asset in building a service-credit record. Win the first orders with an honest L1 (basic-response) commitment, build a quantifiable service history through high-quality delivery, then use real data to support upgrades to L2 / L3. This builds long-term commercial credit far better than committing to L3 and breaching repeatedly.

Strategy ② Compensate for a thin physical network with strong remote-service capability

New entrants typically cannot deploy a large national field workforce, but can build high-quality remote-service capability (remote diagnosis, OTA updates, digital-twin monitoring) at lower cost. Designing the response commitment as "15-minute remote response + X-hour arrival (tiered by city)" is more honest and more achievable than a single "2-hour arrival" commitment. The remote-first model also lowers unit-service cost and improves service economics.

Strategy ③ Build an Authorized Service Partner network

Engage authorised third-party service providers (local electromechanical service firms, automation integrators) for field service to extend geographic coverage without growing fixed headcount. Key design points: partners must accept the OEM's (new entrant's) technical training and certification; field-service procedures must be standardised (referencing Domain 5's UCSP framework); the customer must see a single service interface (one unified work-order system; the customer need not distinguish OEM staff from partner staff).

Strategy ④ Design "growth-tier SLAs" that upgrade commitments as deployment scales

Insert "growth-tier" clauses into the contract: Year 1 runs at L1 / L2; if the Year-1 record meets target (actual availability above target, MTTR consistently below threshold), Year 2 auto-upgrades to L2 / L3 with a corresponding price adjustment. This binds customer expectations to supplier-capability build-out in the same incentive structure — honest and sustainable for both sides.

Strategy ⑤ Transfer delivery risk through a RaaS (Robots-as-a-Service) model

For new entrants with sufficient capital, the RaaS model (equipment ownership remains with the supplier; customer pays by subscription or usage) is the most natural "outcome commitment": because the supplier owns the equipment, it has the strongest economic incentive to ensure uptime — delivery pressure shifts from "contract obligation" to "commercial self-interest." Locus Robotics' RaaS deployment with DHL (500+ million autonomous picks) is a successful example, but strong capital backing and scaled operations are prerequisites.

Strategy ⑥ Be transparent about technology maturity and earn trust through honesty

The worst sin for a new entrant is to hide technical limitations. A supplier that honestly tells the customer "our product is in rapid iteration, software updates are frequent, and we secure update stability through the following mechanisms: staged rollout + 24-hour rollback capability + dedicated technical support" wins long-term relationships far more effectively than one that promises "99 % availability" and breaches repeatedly. Technical honesty is one of the few differentiators a new entrant can sustain against mature OEMs over the long run.

Strategy ⑦ Build a real-time, customer-visible SLA transparency dashboard

Give the customer a real-time SLA performance dashboard (web-accessible, showing current availability, MTD MTTR, PM-compliance rate, etc.) — this turns the SLA from a year-end report into a continuously visible commitment, pressuring the supplier's operating quality upward while building customer trust. Agility Robotics' Arc platform (showing customers Digit's real-time KPIs, uptime, and MTBF) is a pioneering practice in this direction.

Strategy ⑧ Turn the first cohort of customers into SLA-design collaborators

Early adopters typically have higher tolerance for technical limitations and stronger expectations for improvement. Bringing them into the SLA-design co-creation process ("which indicator is your business most sensitive to? What SLA flexibilities would you trade for a lower price?") yields a better-fitting SLA and converts early customers into a data source for product improvement and a market reference for new sales.

The Fundamental Relationship Between Product Reliability and SLA Commitments The ceiling of an SLA commitment is set by the product's actual reliability. No SLA design can transcend the physical limits of the product — if a robot's real MTBF is 500 hours, no contract design can truly deliver a 98 % annual-availability commitment. The fundamental path for a new entrant to lift its SLA capability is to lift product reliability itself. The service system (rapid response, parts pre-staging, predictive maintenance) can narrow the gap between actual availability and the theoretical MTBF-bound ceiling, but cannot exceed that ceiling. This is the physical constraint of SLA design and the boundary that service engineers and managers must honestly communicate to customers.

6.5 SLA Performance Tracking and Compensation Mechanism

The legal force of an SLA arises from the combination of two mechanisms: quantifiable performance tracking (so that breach can be determined objectively) and a clear compensation mechanism (so that breach produces predictable consequences). An SLA without tracking is an empty promise; an SLA without compensation is an unenforceable moral constraint.

6.5.1 Monthly SLA Performance Tracking Process

Step	Node	Operation	Deliverable
T1	**Data Lock-in (day after month-end)**	Collect all relevant month's data from the work-order system, equipment-monitoring platform, and customer-confirmation records; data must be stored in a system accessible to both parties — not held unilaterally by the supplier	Raw data report (each fault's start time, recovery time, disposition record)
T2	**Metrics Calculation (3 business days after T1)**	Calculate the month's SLA-indicator actual values per contract definitions; automated calculation preferred (work-order-system built-in); manual calculations require a reviewer's signature	Monthly SLA calculation table (each indicator: target / actual / variance)
T3	**Variance Analysis (2 business days after T2)**	Root-cause-analyse every below-target indicator: supplier responsibility (triggers compensation) or customer-operation / external factors (recorded in the report with supporting evidence)	Variance-analysis report (with responsibility attribution)
T4	**Customer-Report Submission (8th business day after month-end)**	Submit the monthly SLA report to the customer's authorised contact; the report must include each indicator's actual value, improvement actions, and next-month expectations	Monthly SLA report (signed or e-confirmed by both parties)
T5	**Compensation Settlement (if applicable)**	If a supplier-responsibility breach exists, calculate and execute compensation per contracted clauses; the form (service credit / cash refund) is pre-agreed in the contract	Compensation notice (customer-confirmed)

6.5.2 SLA Violation Compensation Mechanism Design

Compensation-mechanism design must satisfy two goals at once: provide sufficient incentive to the customer (so that SLA breaches produce real commercial consequences and motivate the supplier to maintain quality); preserve an acceptable risk boundary for the supplier (so the maximum compensation is predictable and absorbable, and a single major incident does not destroy supplier financial health).

Tiered Compensation Design (Recommended)

Breach Severity	Trigger Condition	Compensation Magnitude	Notes
Minor Breach	Availability 1–3 percentage points below target, or MTTR < 50 % over target	5–10 % of the month's service fee as service credit	Service credit (no cash refund); applied to the next month's service fee; noted in the monthly SLA report
Moderate Breach	Availability 3–5 percentage points below target, or more than 3 months of minor breach within a year	10–20 % of the month's service fee as service credit, or equivalent cash refund	Customer may choose credit or cash; an improvement plan must be submitted
Severe Breach	Availability more than 5 percentage points below target, or two consecutive P1 responses at 200 %+ of target time	20–30 % of the month's service fee as cash refund, plus a written RCA and a 30-day improvement plan	Written RCA required; improvement plan submitted within 30 days and customer-confirmed; consecutive severe breaches trigger contract-review rights
Accumulated Breach (Annual Trigger)	Annual availability below the SLA target, or 3+ moderate / severe breaches in the year	Annual service-fee rebate (per contracted ratio), or customer's right to terminate without penalty	An annual breach usually indicates systemic service issues — triggers an annual joint review

> **Liability-Cap Clause (Contract Protection the Supplier Must Retain) SLA-breach compensation must have an explicit liability cap in the contract: typically "X % of the annual service-fee total (recommended 30–50 %) or a fixed amount, whichever is lower." The cap protects the supplier from a single major incident (e.g., 72-hour production stoppage) producing compensation exposure many multiples of service revenue. At the same time, the contract must explicitly exclude "indirect loss" and "loss of profit" — otherwise the customer could claim the production-profit loss during downtime, with amounts far exceeding the service-contract value. This is not an unfair exclusion clause but a necessary design for sustainable service-market operations.**

6.6 RaaS Model — Structural Reconstruction of the SLA System

Robots-as-a-Service (RaaS) fundamentally changes the economic logic of the SLA: when the supplier owns the robot, the supplier's commercial interest is directly bound to high availability — when the equipment is not working, the supplier has no revenue. This transforms SLA delivery from legal obligation to commercial self-interest — the strongest SLA-assurance mechanism to date.

6.6.1 RaaS vs Traditional Sales + Service Contract — SLA Logic Comparison

Dimension	Traditional Sale + Service Contract	RaaS Model
Equipment Ownership	Customer-owned: equipment is the customer's fixed asset after purchase	Supplier-owned: customer subscribes to usage rights, does not hold the asset

SLA-Delivery Incentive	Contractual obligation: breach triggers compensation penalty	Commercial self-interest: equipment down = supplier no revenue — far stronger incentive than breach penalties
Availability Measurement	Defined by negotiation; dispute space exists	Typically billed by "actual operating hours" or "tasks completed"; availability directly affects supplier revenue
Maintenance-Investment Incentive	Supplier PM cost is a profit deduction; an incentive to reduce investment exists	Supplier-owned: good maintenance extends life and preserves asset value — strong incentive to do PM well
Technology-Upgrade Responsibility	Customer pays for upgrades; upgrade willingness is constrained by TCO	Supplier-owned: the supplier decides upgrades; subscribed customers automatically benefit from the latest version
Risk Allocation	Customer bears most of the risk of equipment aging and unexpected downtime	Supplier bears equipment-aging, downtime (revenue loss), and tech-refresh risks; customer mainly bears operational risk
Applicable Stage	Mature products with known reliability data	Especially suited to emerging products (still rapidly iterating; customer unwilling to bear technology risk)

6.6.2 Key SLA Design Considerations for RaaS

The SLA design of a RaaS contract differs from a traditional service contract in several key ways:

- Billing-Base Directly Tied to Service Outcome: the simplest RaaS SLA bills by "available operating hours" — customer pays for the hours the robot works; downtime hours are automatically not billed (which is both SLA compensation and revenue loss — two birds with one stone)
- Precise Definition of Performance Indicators: "available operating hours" must define the minimum bar for "available" (e.g., running at ≥ 90 % of rated speed with accuracy within spec); operation below that bar does not count toward available hours
- KPI Transparency Reporting: because supplier revenue is tied to KPIs, the KPIs must appear in real time in a system both parties can independently verify; a customer-side independent data-capture node is recommended to prevent tampering (rare in real disputes, but the design must guard against it)
- Usage-Based Elastic Pricing: a RaaS contract may use tiered usage pricing (like cloud pay-as-you-go); fleet can flex up in peak season and contract in off-season — one of RaaS's most important customer value propositions (Locus Robotics / DHL case: flexible deployment to peak-season demand)
- Contract Term Matched to Equipment Life: RaaS terms are typically 3–5 years (matching the principal lifecycle of robotic equipment); the supplier amortises depreciation and maintenance over the term; the post-term equipment-disposition path must be explicit in the contract

6.7 SLA Joint Review and Continuous Improvement Mechanism

An SLA is not a static contract executed once and left alone — it is a dynamic system that should be continuously optimised as equipment state changes, business needs evolve, and service capabilities grow. Establishing a periodic SLA joint-review mechanism is key to upgrading the SLA from a "legal constraint" to a "service-improvement engine."

Review Frequency	Review Content	Participants	Typical Output
Monthly Review (Routine)	Monthly SLA-data review; root-cause confirmation for anomalous indicators; next-month focus items	Service manager + customer operations lead	Monthly SLA-report confirmation; short-term improvement-action list
Quarterly Review (Strategic)	Quarterly KPI-trend analysis; service-maturity assessment (S1 → S4 evolution state); spare-parts consumption and supply-chain health	Service manager + customer equipment / production lead	Quarterly improvement plan; spare-parts strategy adjustment; predictive-maintenance roadmap update
Annual Review (Contract Level)	Full annual SLA-performance evaluation; contract-clause applicability review; tier upgrade / downgrade decisions; pricing-renegotiation baseline	Service director + customer leadership (Director of Equipment / VP Operations)	Contract renewal / amendment intentions; SLA-tier adjustment; next-year service plan

6.8 Domain 4 Key Performance Indicator Framework

Indicator	Definition	Industry Benchmark Target	Excellent Target	Management Significance
Overall SLA-Compliance Rate	Share of SLA indicators meeting target (monthly)	**≥ 85%**	**≥ 95%**	Core composite indicator; weighted across all sub-indicators
Customer-Availability Attainment Rate	Months where actual availability ≥ SLA-committed availability / total contract months	**≥ 90 % (monthly)**	**≥ 98 % (monthly)**	The most central outcome indicator; consecutive breaches trigger contract review
SLA-Breach Occurrence Rate	Compensation-triggering SLA breaches in the month / total SLA clauses	**< 10%**	**< 3%**	Measures SLA-design reasonableness; an excessive breach rate indicates over-aggressive targets or insufficient service capability
Compensation-to-Service-Fee Ratio	Month's SLA compensation / month's service-fee revenue	**< 5%**	**< 1%**	Financial-health indicator; above 10 % triggers an urgent service-capability review
SLA-Report Timeliness Rate	Share of monthly SLA reports submitted by the agreed date	**100%**	**100%**	Compliance requirement; deviations must record a reason
Customer Satisfaction (CSAT)	Monthly customer-satisfaction survey related to SLA (1–10 scale)	**≥ 7.5**	**≥ 8.5**	Subjective but important; the gap between CSAT and objective SLA indicators reflects the effectiveness of expectation management
SLA-Clause Dispute Rate	SLA-indicator clauses disputed by the customer in the month / total clauses	**< 5%**	**< 1%**	A high dispute rate typically indicates contract-definition ambiguity — triggers a contract-clause refinement

6.9 Chapter Summary and Inter-Domain Interfaces

Chapter Summary Domain 4 (SLA Design & Service Commitment)'s core contribution is to fill the long-missing "common language" of the robotic-service market — enabling suppliers and customers to discuss service quality, design service commitments, and resolve service disputes using the same quantitative indicators.

The RSF metrics framework, with OEE at its centre, distinguishes supplier-controllable from non-controllable OEE losses, ensuring fair attribution of SLA responsibility. The four-tier SLA product architecture (L1–L4) enables suppliers to honestly match their capability to their market commitments. The eight strategies for new entrants provide a viable path to credible SLA commitments during the capability-build-up phase. The RaaS model, by restructuring ownership, transforms SLA delivery from legal obligation into commercial self-interest — an important direction in the SLA system's evolution.

Interfaces with Other Capability Domains

Interface Direction	Interfacing Domain	Content
Domain 1 → Domain 4	Service Lifecycle Management	Domain 1's KPIs (MTTR, PM compliance, planned-to-unplanned maintenance ratio, etc.) are the core technical-design basis for Domain 4 SLA clauses; the quality of Domain 1 PM execution directly determines whether availability commitments can be achieved
Domain 2 → Domain 4	Fault Diagnosis & Technical Support	Domain 2's FTFR and MTDA directly affect MTTR (diagnosis time is the largest component of repair time); Domain 2's repeat-fault rate affects the stability of annual-availability achievement
Domain 3 → Domain 4	Remote Service & Digital Operations	Domain 3's remote-response rate and predictive-maintenance capability are the technical underpinnings of L3 / L4 SLA commitments; without Domain 3's digital infrastructure, a P1 15-minute remote-response commitment cannot be met; Domain 3's real-time data is the automated-calculation source for the monthly SLA report
Domain 4 → Domain 5	Cross-Platform Service Procedures	Domain 4 SLA indicators must be defined separately by platform type (industrial-arm availability target vs AMR positioning-accuracy target); a heterogeneous-fleet SLA must specify response-priority rules when multiple platforms fault simultaneously
Domain 4 ← Domain 6	AI/ML System Service Methods	Performance degradation of AI / ML components (perception-accuracy drop, policy drift) can degrade OEE quality rate; the attribution of AI-layer performance must be explicit in SLA design; stability fluctuations introduced by AI-model updates are handled under the SLA's "planned-maintenance downtime" clause

Chapter 7

Domain 5 · Cross-Platform Service Procedures

7.0 Chapter Introduction — The Industry's Fragmentation Dilemma

If a robot-service engineer has mastered every skill for an industrial arm, can they directly service a collaborative robot? An AMR? A humanoid robot? In most cases the answer is no — not because the engineer is under-skilled, but because the differences across these four robot families in safety norms, physical mechanisms, software architecture, fault modes, and service-operating procedures far exceed what most people expect.

This is a severely under-appreciated structural difficulty in the robotic-service industry. As manufacturing, logistics, and the service sector deploy multiple robot types at speed, customer fleets are becoming heterogeneous at an unprecedented pace: a single production line may simultaneously run six-axis industrial arms, cobots, and AMRs; a logistics warehouse may deploy three different AMR brands within three years; in 2025, plants such as BMW and Amazon began piloting humanoid robots operating alongside traditional industrial robots.

Procedural fragmentation is the direct consequence of this heterogeneity: each robot family has its own safety standards, operating procedures, maintenance manuals, and fault-code system; service organisations are forced to maintain independent service capabilities for each family; engineers face a "re-learning" cognitive load when switching service objects; and service-quality consistency cannot be guaranteed across families. The dilemma will not ease as platform counts decline — on the contrary, with new types such as humanoid robots accelerating into commercial deployment, heterogeneity will only deepen.

Domain 5's core task is to find unity while acknowledging difference: identify which service procedures are reusable across platforms (the unifiable base), which must be designed separately per platform (the necessarily-differentiated layer), and how to build an engineer-capability system that supports flexible cross-platform scheduling. This chapter also examines the special challenges of humanoid-robot service and the industry's forward-looking response strategy.

7.1 Multi-Platform Robot Fleets — Industry Status and Evolution Trends

Understanding the current industry state is the precondition for designing an effective cross-platform service framework. Fleet heterogeneity is not an accident — it is a structural trend driven by multiple converging industrial forces.

7.1.1 Overview of Major Robot Types and Service Characteristics

Table 7.1-1 Service-Characteristics Overview of the Four Major Robot Types (2025)

Robot Type	Typical Application Scenarios	Principal Reference Safety Standards	Service Complexity	Current Service-Capability Challenges
Industrial Arm	Welding, material handling, assembly, painting, machine tending	ISO 10218-1/-2:2025 ISO 45001 IEC 61508	High maturity ★★★★	Reducer-wear diagnosis is technically demanding; different brands' controllers are completely incompatible; precision requirements in high-end applications keep rising

Collaborative Robot (Cobot)	Human-collaborative assembly, inspection, small-part handling, medical assistance	ISO 10218-1/-2:2025 (Note: core safety requirements have fully incorporated the collaborative technical specifications of the original ISO/TS 15066:2016)	Mid-maturity ★★★	Force-torque sensor temperature drift; compliance verification of safety parameters; electrical compatibility of multi-brand end-effectors; ambiguous application-scope boundaries
Autonomous Mobile Robot (AMR)	Warehouse logistics, factory material flow, hospital delivery, commercial cleaning	ISO 3691-4:2023; ANSI/RIA R15.08-1/-2; VDA 5050 (logistics scheduling standard)	Mid-maturity ★★★	SLAM-map management complexity; multi-brand AMR fleet management (VDA 5050 compliance differences); battery-health management; dynamic-environment adaptation
Humanoid Robot	Factory assembly (BMW / Tesla pilots), logistics sorting (Amazon's Digit), disaster response, care assistance	No current dedicated ISO standard (currently references selected ISO 13482 clauses; ISO/TC 299 dedicated standard in progress)	Extremely high ★★★★★	No service-safety procedure has yet been established for biped dynamic balance; very short battery runtime (90 minutes to 4 hours); unpredictability of AI behaviour; shortage of maintenance technicians; parts supply chain not yet established

Note: ★ count indicates the current maturity of service technology and procedures for that platform type, not of the robot technology itself. For Humanoid Robots, the ★★★★★ rating reflects extremely high service complexity rather than maturity — service procedures are largely unestablished, making this the most challenging platform for field engineers.

7.1.2 Drivers of Heterogeneity and Causes of Service Fragmentation

Fleet heterogeneity is driven by three mutually reinforcing forces:

- **Task diversification** No single robot type can cover every task scenario: industrial arms suit high-precision fixed stations; cobots suit flexible scenarios working alongside humans; AMRs suit large-area material transfer; humanoid robots are exploring general-purpose work in unstructured environments. Complex manufacturing and logistics scenarios inherently require multi-type collaboration.
- **Supplier diversification** For supply-chain risk diversification, large manufacturers typically do not concentrate all robot procurement with a single brand; even within a single type (e.g., industrial arms), the fleet may include 3–5 different suppliers, each with its own service ecosystem.
- **Differing rates of technology iteration** Iteration speed varies sharply across robot types: industrial-arm technology is relatively mature (annual iteration), while AMRs and humanoids remain in rapid iteration (quarterly or faster). Equipment procured early may have generational

gaps in service procedure relative to newly procured units.

> **RSF's Core Response Logic** Facing heterogeneity, RSF does not try to establish a single procedure applicable to every robot type — that is technically infeasible and, on safety grounds, dangerous. RSF's logic is: identify and systematise the unifiable service layer (safety-verification flow, work-order format, documentation standards, technology-upgrade path), and at the same time build necessarily-differentiated procedures per family (platform-specific operating steps, safety parameters, fault modes), so engineers can switch between robot types at the lowest possible additional learning cost.

7.2 Systemic Challenges of Cross-Platform Service

Systematic analysis of cross-platform service challenges is the prerequisite for constructing an effective response. RSF organises these challenges along five dimensions; the nature of the challenge and the direction of response differ for each dimension:

7.2.1 Five-Dimension Challenge Framework

Dimension	Challenge Category	Core Problem	Typical Manifestation
①	**Safety-Standard Fragmentation**	Different robot types follow entirely different safety standards; safety-operating procedures cannot be ported directly across families	Industrial-arm engineers are accustomed to entering the working envelope after LOTO; cobot engineers may believe "LOTO is unnecessary when torque limits are active" — each conception, transplanted into the other's context, can cause an accident
②	**Technical-Capability Gap**	Structural tension between depth in a single robot family and breadth across families	An engineer who is expert on a FANUC controller may drop from Specialist to Professional when facing a ROS 2-based AMR; most service organisations have no engineer who can independently service all platform types
③	**Tool and Process Incompatibility**	Different robot families have entirely different diagnostic tools, work-order systems, and parts-coding schemes; the same service organisation maintains multiple parallel processes	Industrial-arm spares are managed by OEM part numbers; AMR spares by modular components; cobot end-effectors are managed separately by the tool supplier — three coding systems, three procurement processes, three inventory systems
④	**Non-Comparable SLA Indicators**	Availability definitions, downtime-recognition criteria, and response-time expectations differ entirely across robot families; cross-platform service quality cannot be quantified uniformly	Industrial arm: downtime = line stoppage (directly visible); AMR: downtime = task failure (partial failure may be routed around); cobot: downtime = protective stop (may auto-recover in seconds) — the same "response-time target" carries entirely different practical meaning across the three
⑤	**Standards Vacuum and Regulatory Uncertainty**	New robot types (such as humanoid robots) lack dedicated safety standards; service engineers face a technical entity with a complete service-procedure blank	2025 reality: humanoid robots have no dedicated safety standard equivalent to ISO 10218; service engineers face a triple blank — "what to reference, what to do, how to demonstrate compliance."

7.2.2 Structural Root Cause — Historical Limitations of Standards Division of Labour

The five challenges share a common structural root cause: the existing international standards system divides responsibility by "robot product type" rather than by "service-engineer work scenario." ISO 10218 serves industrial arms; ISO/TS 15066 serves cobots; ISO 3691-4 serves AMRs — each standard is high-quality within its specific product scope. But when the service engineer stands on a shop floor running all three platform types simultaneously, this product-type-based standards system provides three independent reference books, not one unified work guide.

This is not a standards-body oversight — drafting safety standards by product type is technically the most rigorous approach, because different families have fundamentally different physical characteristics, risk sources, and safety mechanisms. But this reasonable technical choice creates systemic procedural fragmentation at the practical-service layer. Domain 5's design logic is precisely to acknowledge this structural constraint and to build, above the "safety floor specified by standards," a "service-logic layer that operates across robot types."

7.3 RSF Unified Cross-Platform Service Framework — Common Layer and Differential Layer

The RSF cross-platform service framework rests on a core distinction: which service activities share identical logic across all robot types (the unifiable set), and which must be specified per platform (the necessarily-differentiated set). This distinction lets the engineer build cross-platform service capability through a "unified framework + platform-specific plug-ins" model, instead of learning a fresh, independent procedure system from scratch for each robot type.

7.3.1 Common Service Layer (Universal Cross-Platform Procedures)

The following seven categories of service activity, although their operating details vary by platform, share a consistent logical framework across all robot types and can be cast as a unified procedural template:

Unified Procedure 1 Safety-Verification Flow Structure

Service work on every robot type must complete safety-state verification before any physical operation begins (energy-isolation confirmation, clearing of the hazard zone, safety-device function test). The logical structure of the flow is the same across platforms, while the concrete steps differ (industrial arm: LOTO; cobot: torque-limit verification; AMR: stop and secure; humanoid: full-joint lock + physical support).	**Cross-Platform unification: the five-step logic of safety verification (identify → isolate → verify → operate → restore)** **Platform-specific difference: the concrete way each step is executed (platform-specific)**

Unified Procedure 2 Work-Order Format

Regardless of robot type, the work order must contain the same core information elements: timestamps, robot ID, task type, executor, disposition description, and outstanding items. The meanings of work-order fields are identical across robot types; the field content (e.g., the format of "alarm codes" on different platforms) differs.	**Cross-Platform unification: mandatory fields and format standards of the work order** **Platform-specific difference: the data source and completion conventions of each field on different platforms**

Unified Procedure 3 Fault-Escalation Path

The three-tier escalation logic — "Professional handles → escalate to Specialist beyond authorisation → Specialist handles → escalate to Expert / OEM beyond capability" — is identical across robot types. The criteria that mandate escalation (P1 safety class, recurring fault, novel fault type) are also identical across platforms.	**Cross-Platform unification:** definition of escalation triggers and the escalation-notification flow **Platform-specific difference:** the typical P1 scenario list for each platform

Unified Procedure 4 Maintenance Records and Equipment-Passport Standards

The four-layer structure of the equipment passport (basic-information layer, technical-state layer, statutory-document layer, knowledge-accumulation layer) and the six trigger nodes for configuration-baseline snapshots remain the same across all robot types.	**Cross-Platform unification:** structural standards and mandatory content of the equipment passport **Platform-specific difference:** the concrete content items of each layer on different platforms

Unified Procedure 5 Performance-Acceptance Standard Format

The logic "run an acceptance procedure after work and record the result" is identical across robot types; acceptance content must be defined per platform (industrial arm: accuracy test; cobot: torque-baseline test; AMR: navigation-path test; humanoid: motion-stability test).	**Cross-Platform unification:** structure of the acceptance flow (when to perform, who signs, where the document is archived) **Platform-specific difference:** concrete test content and acceptance criteria

Unified Procedure 6 Customer-Deliverable Format

The format and content logic of written deliverables provided to the customer after each service (service report, parameter-change note, next-maintenance recommendation) are identical across robot types, so the customer can manage the service histories of different robot families in a single unified format.	**Cross-Platform unification:** template structure of the service report **Platform-specific difference:** technical content in the report (platform-specific parameters, standards referenced)

Unified Procedure 7 Knowledge-Library Entry Format

The knowledge-library entry format "symptom → layer localisation → root cause → disposition → prevention" is generic across robot types, enabling cross-platform service-knowledge sharing (e.g., an experience disposition for a bus-communication problem discovered on a cobot may be a reference for diagnosis of a similar issue on an industrial arm).	**Cross-Platform unification:** structure and mandatory fields of knowledge-library entries **Platform-specific difference:** entry content (fault characteristics and diagnostic methods are platform-specific)

7.3.2 Differential Service Layer (Platform-Specific Procedures)

Because of fundamental differences in physical characteristics or safety mechanisms across robot families, the following service content cannot be unified and must have a dedicated procedure built for each platform. When switching service objects, the engineer must confirm mastery of the target platform's procedures below:

Table 7.3-1 Necessarily-Differentiated Service Content (Comparison Across Four Robot Types)

Service Content	Industrial Arm	Cobot	AMR	Humanoid Robot
Energy-Isolation Procedure	Disconnect main power + LOTO; entry into the working envelope requires complete lockout of all energy sources	Main-power disconnect + LOTO, or torque-limit verification (distinguished by task type); per ISO/TS 15066	Power disconnect + wheel-system lockout (to prevent free roll on slopes); SLAM-map freeze	All-joint power-off + mechanical lockout; prevent unintended motion induced by gravity; physical support required to prevent collapse
Safety-Function Verification	E-stop / safety door / safety light curtain / STO; full verification per ISO 10218-2	Torque-limit baseline test; speed-limit verification; ISO/TS 15066 torque-contact test	E-stop response; map-boundary limits; collision-detection sensitivity test; personnel-detection sensors	Full-body E-stop response; fall-protection mechanism; balance-recovery test (standard in development)
Accuracy-Acceptance Methods	TCP repeatability (ISO 9283 method); TCP-calibration procedure	TCP accuracy + force-control accuracy; torque baseline established at working temperature; path-deviation measurement	Localisation accuracy (navigation measurement); docking accuracy (charging dock / workstation); SLAM-map-consistency check	Gait stability; joint-angle accuracy; hand-grasping success rate (methodology still developing in the industry)
Key Wear-Part Replacement Procedures	Reducer grease (RV / harmonic, OEM-specified type); brakes; encoder batteries	Force-torque sensor calibration; joint elastic actuators; end-effector flange	LiDAR-lens cleaning; drive-wheel-wear inspection; charging-contact maintenance; battery-health assessment	Battery-pack replacement (short runtime, high replacement frequency); electric actuators; perception-sensor suite
Software and Firmware-Specific Requirements	OEM controller firmware (brand-proprietary); parameter-backup format	Force-control software parameters; safety-PLC program; end-effector configuration	SLAM-navigation software; map-update procedure; VDA 5050 interface compliance	Proprioception calibration; AI-inference model updates; reinforcement-learning policy verification (frontier area; procedures still developing)

7.4 Platform-Specific Service Procedures

7.4.1 Industrial Arm — Key Procedure Highlights

Industrial arms have the most mature service-procedure system, but that does not mean service is easy — the challenges of high precision, high speed, and high payload become more prominent as years of use accumulate. The following are the most easily overlooked key points of industrial-arm procedures:

- Brand-specificity of reducer service: harmonic reducers (common in wrist joints) and RV reducers (common in main-axis joints) have entirely different grease specifications, replacement intervals, and replacement procedures; each OEM selects different reducer models, and grease cannot be substituted across brands or models — OEM-specified grease is non-substitutable, and the wrong grease can cause seal failure within hundreds of hours

- Lifelong importance of the load-inertia parameter: an incorrect Load Inertia is the most common preventable cause of shortened reducer life. After every end-effector change, workstation redeployment, or joint replacement, load identification must be re-executed and the parameter record updated
- Thermal-steady-state requirement for high-accuracy applications: applications such as precision welding and laser cutting with TCP accuracy requirements within ±0.1 mm must perform accuracy acceptance only after thermal steady state (≥ 30 minutes of continuous running) — cold-start accuracy compliance does not imply thermal-steady-state compliance
- Legal status of the safety-function document: ISO 10218-2 requires the integrator to provide a full safety-function verification report; this document has legal force under the EU Machinery Directive, must be archived permanently, and must be re-signed after any service-induced change to the safety circuit

7.4.2 Cobot — Key Procedure Highlights

The core service challenge of cobots is "maintaining the compliance of safety parameters." Cobot safety mechanisms do not rely on physical protection barriers — they rely on software-defined torque and speed limits. This makes drift of safety parameters (induced by temperature, wear, calibration error) a high-frequency service issue, and its danger is not as visually obvious as physical-barrier damage.

Cobot-Specific Service Items	Operating Requirements	Compliance Reference and Key Cautions
Force-Torque Sensor Baseline (mandatory quarterly)	After thermal steady state (≥ 30 minutes continuous running), perform the force-torque-sensor zero-point calibration per the ISO/TS 15066 procedure; record per-axis baseline values and ambient temperature at the time	ISO/TS 15066:2016 specifies maximum allowable torque limits for different body-contact regions (consult the standard for exact values); calibration records must include ambient temperature (re-calibration required if temperature drifts more than ±5 °C)
Safe-Contact Torque Verification (mandatory after overhaul)	Use a calibrated force gauge to apply force to the cobot end-effector, simulating contact scenarios with different body parts, and verify that the protective stop triggers within the specified torque	Maximum allowable torque at each contact region must conform to the limits specified in ISO/TS 15066:2016 Annex A (consult the standard for exact values); any post-maintenance operation that involves safety parameters must be re-verified
End-Effector Mass Update	On every end-effector change, re-enter the tool mass and centre-of-gravity, and execute the load-identification procedure	Incorrect tool-mass parameters cause torque-calculation deviation, allowing actual contact force to exceed the set limit while the controller fails to respond correctly
Collaborative-Space Re-Definition	After workstation relayout, re-run the collaborative-space risk assessment and update the definitions of safe-speed and torque-limit zones	ISO/TS 15066 requirement: safety parameters of the collaborative space must be based on the risk assessment; a workstation change must trigger re-assessment rather than reuse of prior parameters
Verification of Parameter Switching for Multiple End-	When using an automatic tool changer, verify that the controller's tool parameters auto-update with each	Tool-parameter desynchronisation is one of the most common precursors to incidents in cobot applications

Effectors	physical tool change	

7.4.3 AMR — Key Procedure Highlights

AMR service challenges differ fundamentally from those of fixed industrial robots: an AMR moves autonomously in a dynamic environment, and its "operating state" encompasses not only mechanical and electrical state but also map state, navigation state, and charging state. These three are critical to AMR availability and safety, yet rarely covered in traditional robotic-service training.

Map-Management Procedures (AMR-Specific)

- Map-Validity Check: after every significant change in the production environment (shelf movement > 30 cm, addition or removal of fixed obstacles, changes in aisle width), verify AMR-map accuracy by driving the AMR to a known precise coordinate and measuring the localisation error; an error > 5 cm triggers a local map rebuild
- Periodic Map-Consistency Audit: a full-path consistency audit is recommended monthly — have the AMR traverse all working paths and record per-segment localisation confidence; zones with continuously declining confidence should be prioritised for investigation of environmental changes
- Re-Mapping Procedure: after a large-scale environmental change (line relayout) or when cumulative localisation error exceeds threshold, re-map under actual operating state (shelves loaded, normal lighting); other AMRs must be temporarily cleared from the area during mapping

Battery-Health Management (Critical AMR Service Item)

Battery management is the single most critical factor affecting AMR-fleet availability and is often underestimated. The battery is not a simple charge-discharge cycle — its health directly affects range, charging time, and emergency-stop capability:

- Capacity-Degradation Tracking: each quarter, record actual runtime under a standard task cycle and compare to the new-battery baseline; degradation to 80 % (runtime reduced by 20 %) is the planned-replacement threshold; degradation to 70 % is the immediate-replacement threshold
- Charging-Strategy Optimisation: avoid frequent deep discharge (< 20 %) and excessive overcharge (> 95 %) — both significantly accelerate lithium-ion degradation; an opportunity-charging window of 20–90 % is recommended rather than waiting for full depletion followed by full recharge
- Charging-Contact Maintenance: monthly cleaning of AMR and dock charging contacts (copper-oxide layers reduce charging efficiency and raise thermal loss); contact resistance > 10 mΩ requires cleaning or replacement
- Temperature Effects on Charging: low-temperature environments (< 10 °C) reduce lithium-battery charge acceptance; AMRs in cold-storage environments must adjust the charging strategy (lower charge current, accept longer charging time)

Multi-Brand AMR Mixed-Fleet Management (VDA 5050 Protocol)

When the customer runs a multi-brand AMR fleet, VDA 5050 is the key interface standard enabling unified cross-brand dispatch. Service engineers must understand VDA 5050's basic logic and the compliance-check method:

- VDA 5050 is the common AMR-interface protocol jointly developed by the German Association of the Automotive Industry (VDA) and the German Mechanical Engineering Industry Association (VDMA); it

defines a standardised communication protocol between the dispatch system and the AMR (based on MQTT messaging)

- VDA 5050 Compliance Check: during mixed-fleet installation and commissioning, verify the VDA 5050 version compatibility of each AMR (current mainstream: VDA 5050 v2.0); different versions may have functional differences; check the version range supported by the Fleet Management System (FMS)
- Common Integration Problems: coordinate-system inconsistency (different brands use different origins); pose-report frequency mismatch; task-state-machine definition differences — these must be investigated case-by-case during system integration; one cannot assume "VDA 5050 compliance = direct mixed operation works"

7.5 Humanoid Robot Service Procedures — Systemic Frontiers and Response Strategy

Note on the Positioning of This Section Commercial deployment of humanoid robots is accelerating (BMW's partnership with Figure AI has completed 20-hour continuous-shift tests; Amazon's Agility Digit charges 9 minutes after each 30-minute work cycle in warehouses; Tesla Optimus is planned for in-factory deployment within 2025), yet dedicated service procedures for humanoid robots are nearly blank globally. The goal of this section is not to provide a mature operating manual, but to systematise the challenges, propose currently available best-practice directions, and define RSF's frontier framework-construction direction in this field.

7.5.1 Ten Special Service Challenges of Humanoid Robots

ID	Challenge Item	Concrete Manifestation	Essential Difference vs Traditional Robots
H1	**Service-Safety Risk of Dynamic Balance**	Humanoid robots collapse after power-off; any unintended joint actuation during maintenance can unbalance the robot	Traditional industrial arms remain still after power-off; AMRs stop moving after power-off. After power-off, a humanoid's centre of gravity is uncontrolled, and the engineer must provide physical support — a brand-new service-safety need not directly covered by existing ISO standards
H2	**Very Short Battery Runtime**	Current mainstream humanoid runtimes: Agility Digit - 90 minutes; Figure 02 - 2–3 hours; Tesla Optimus Gen 2 projected 4–5 hours	Traditional robots can run for months on continuous external power; humanoid service windows must be embedded in charging-cycle plans, otherwise PM may lead to uncontrolled collapse from depletion mid-task
H3	**Safety-Standards Vacuum**	As of 2025, no ISO standard dedicated to the service-operation safety of dynamic-balance humanoid robots exists	Engineers lack a reference for compliance proof; service organisations bear legal risk in the absence of applicable standards
H4	**Unpredictability of AI Behaviour**	A humanoid's motion is generated by neural networks and reinforcement-learning policies; under certain boundary conditions it may produce untrained "emergent behaviour"	Traditional robots have fully deterministic motion paths; the humanoid-service engineer faces an "AI-controlled physical entity" — traditional behaviour-prediction logic no longer applies
H5	**Scarcity of Maintenance Technicians**	Engineers globally with humanoid-service capability are very few in 2025 and are concentrated inside the robot makers	Traditional robots have a mature third-party service ecosystem; third-party humanoid service providers are nearly non-existent, leaving customers highly dependent on OEM support

H6	**Parts Supply Chain Not Yet Established**	Humanoid robots are not yet in mass production; the parts supply chain (especially actuators and sensor suites) is incomplete; lead times are long and prices high	Traditional robots have a multi-tier parts chain (OEM → authorised distributor → refurbished-parts market); humanoid robots currently have only the OEM channel, with most parts not in stock
H7	**Multi-Modal Perception-System Maintenance**	Humanoid robots integrate cameras, LiDAR, IMUs, tactile sensors, microphones — multi-modal perception suites with time-synchronisation requirements among them	Traditional sensor systems are relatively simple and well-defined; humanoid perception systems are highly integrated, and a single-sensor failure can trigger a cascade of AI-perception failures
H8	**Very High Software-Update Frequency**	Humanoid AI models and control policies typically iterate on a weekly cadence — far above the annual firmware-update cycle of traditional robots	Traditional-robot OTA updates are relatively rare events; humanoid service procedures must treat "continuous software-update management" as core daily-service content
H9	**Work-Environment Adaptability**	Humanoid reliability in unstructured environments is significantly lower than in structured environments; environmental changes (floor state, lighting, obstacles) can induce anomalous behaviour	Traditional robots operate at fixed stations with limited environmental impact; humanoid robots must continuously assess the effect of the working environment on performance
H10	**Opaque Costs**	Lifecycle service costs for humanoid robots (maintenance, battery replacement, software subscription, OEM support) lack reliable industry data as of 2025	Traditional robots have decades of accumulated service-cost data; humanoid SLA pricing and service-cost budgeting lacks a data foundation

7.5.2 Best-Available Practice Framework (2025 Perspective)

In the absence of dedicated standards, RSF recommends that service organisations adopt the following interim best-practice framework, built from currently available norms:

Best Practice 1 Dynamic-Balance Maintenance-Safety Procedure (Interim Framework)

Before power-off: move the humanoid robot to a spacious, level service area; guide it into a service-seated or kneeling posture (lowered centre of gravity); immediately after power-off, install a dedicated stand for physical support (analogous to an aircraft-maintenance support jig). Reference: by analogy to the "stability confirmation" principle of ISO 45001's work-at-height procedures and the airframe-support norms in aviation MRO (maintenance / overhaul / refurbishment), adapted for the humanoid context.

Best Practice 2 Battery Management and Service-Window Planning

Establish a charging-service coordination schedule for the humanoid: planned maintenance starts only at full charge (> 80 %); maintenance expected to exceed 30 minutes either uses external power (if supported) or runs in segments (maintain → charge → maintain); maintain a battery-health log that records actual runtime per charge-discharge cycle and builds a capacity-degradation curve.

Best Practice 3 AI Behaviour-Boundary Verification Procedure

Before maintenance: run a standard behaviour-verification programme in an isolated environment (no other personnel) to confirm the AI control logic is in the expected state; after maintenance: rerun the behaviour-verification programme and compare pre- and post-behaviour consistency; after any software update: run "sandbox verification" (verify new behaviour in a constrained environment,

confirming no emergent anomalies). Reference: drawn from the "Safety Driver Protocol" concept of autonomous-vehicle testing.

Best Practice 4 Multi-Modal Perception-System Maintenance Procedure

Establish a full perception-system calibration procedure: vision sensors → LiDAR → IMU → tactile sensors → time-synchronisation verification, executed in order; full calibration is mandatory after every actuator maintenance or sensor replacement; build a perception-system baseline dataset (sensor-output signatures under "new / normal" state) for subsequent state comparison.

Best Practice 5 Close OEM-Collaboration Service Model

Until a mature third-party humanoid-service ecosystem forms, service organisations should establish a "joint-service agreement" with the OEM: every non-standard maintenance operation has OEM-engineer remote supervision (video call or AR collaboration); a dedicated escalation channel guarantees the organisation can reach OEM technical support within 24 hours; participate in the OEM's "Early Customer Service Partner" (ECSP) programme to jointly accumulate service knowledge.

7.5.3 RSF Roadmap for Humanoid Robot Service Procedures

Table 7.5-1 RSF Humanoid-Robot Service-Capability Roadmap

Time Window	External-Condition Assumption	RSF Response Action	Capability Milestone
2025–2026 (Current Stage)	Humanoid mass production < 10,000 units; dedicated ISO standards still in development; OEM-led service	Build a humanoid-service-challenge knowledge library; submit RSF service-procedure contributions to ISO TC 299; establish an interim best-practice framework based on challenges H1–H10	Release the RSF humanoid "Interim Framework v0.1"; five pioneer service organisations validate it
2027–2028	ISO humanoid safety-standard draft released; mass production reaches tens of thousands; third-party maintenance ecosystem begins forming	Align the RSF interim framework with the ISO draft; develop a Professional dedicated certification course for humanoid robots; establish service-partner agreements with 2–3 OEMs	Release of RSF Humanoid Service Procedures v1.0; first cohort of humanoid-dedicated certified engineers
2029–2030	ISO humanoid safety standard formally released; mass production reaches hundreds of thousands; complete service ecosystem formed	Fully integrate humanoid-dedicated procedures into the main RSF framework; develop Specialist- and Expert-tier humanoid-dedicated certifications; establish an industry service database	Humanoid-service capability sits alongside traditional-robot service capability as a regular item across the six RSF domains

7.6 Cross-Platform Engineer Capability Model

The core of Domain 5 capability is not "master every detail of every robot type" but "build a structured capability framework for cross-platform switching": when switching service objects, the engineer must rapidly identify "the parts identical to my familiar platform" and "the parts requiring dedicated learning," rather than starting from scratch every switch.

7.6.1 T-Shaped Competency Structure

RSF recommends engineers build cross-platform capability in a T-shaped structure: the horizontal breadth represents understanding of the service logic common to all robot types (wide and thin); the

vertical depth represents dedicated capability on one or two primary platforms (deep and precise). The T-shape carries more service value than "a little of everything" (the shallow-dish profile) and gives greater scheduling flexibility than "only one platform" (the I-shape).

Capability Type	Professional	Specialist	Expert	Master
Horizontal Breadth (Common to All Platforms)	Understand the basic safety-operating differences across the four families; able to recognise "this is outside my platform authorisation"	Can independently switch service between industrial arms and cobots; understands the basic service logic of AMRs and humanoid robots	Can independently deliver Specialist-level service across all four robot families; masters a rapid cross-platform capability-gap assessment method	Can design organisation-level cross-platform service-capability systems; can represent RSF in cross-platform procedure standards work
Vertical Depth (Primary Platform Specialisation)	Full Professional-level capability on one primary platform (typically industrial arm or cobot)	Independent Specialist-level capability on 1–2 primary platforms; can mentor Professionals on basic operations across other platforms	Expert-level capability on 2–3 platforms; Specialist-level capability on all platforms	Strategic-level understanding across all platforms; recognised industry authority on at least one specific platform type
Cross-Platform-Switch Authorisation Requirements	On switching to a new platform, completion of that platform's "Switchover Assessment" (2–4 hours) is required before performing basic service tasks on that platform	Switching to a new brand within a familiar platform category: 4 hours of brand-specific training; switching to a new platform category: complete that category's Professional-level dedicated course	After switching to a new platform category, Specialist-level capability can be built rapidly (estimated training time: 1–3 months of intensive practice)	No switching restrictions; can provide strategic guidance on any platform

7.6.2 Platform Switchover Assessment

When an engineer needs to switch to a not-yet-serviced platform, RSF recommends completing the Switchover Assessment before formally beginning service, to confirm the engineer has mastered the platform's key difference points:

Assessment Module	Core Assessment Content	Pass Criteria
Safety-Procedure Differences	Differences between this platform's energy-isolation procedure and that of the engineer's familiar platform; the platform's unique safety functions and their verification methods	Can correctly articulate at least three key safety differences; can independently execute the platform's LOTO or equivalent energy-isolation procedure
Alarm-Code Basics	The platform's alarm-code structure (category prefixes, priority definitions); how to access alarm history and system logs	Can correctly distinguish the platform's principal alarm categories without consulting the manual; can access and export alarm history
Basic Maintenance Operations	Key points of the platform's daily-inspection checklist; 1–2 of the platform's most common maintenance operations (such as grease check or battery-state check)	Correctly perform the daily inspection under supervision; able to identify anomalous states requiring escalation
Documentation Requirements	Platform-specific filling conventions under the RSF unified work-order format (which field content differs from other platforms)	Can correctly complete a full service work order on the platform without omitting key fields
Confirmation of	Operation types that a Professional engineer	Can correctly judge the operating-

Authorisation Boundaries	may not execute independently on this platform; the list of scenarios requiring immediate escalation to Specialist	authorisation tier in three scenario cases

7.7 Heterogeneous Robot Fleet Service Management Methodology

When the customer operates multiple robot types and brands simultaneously, service-management complexity does not stack linearly — it grows exponentially. Domain 5 establishes the framework for "per-robot cross-platform service procedures" at the capability-domain level and must also provide a management methodology for "the heterogeneous fleet as a whole."

7.7.1 Four Core Principles of Heterogeneous-Fleet Service Management

Principle 1 Unified Information Entry, Tiered Technical Implementation The customer and service manager should see the service state of every platform through one unified service-management dashboard (a CMMS platform), without having to log into multiple systems. Technical-layer heterogeneity (different data formats, interface protocols) is handled by the service-management system's integration / adaptation layer — not by asking the customer to accept multiple interfaces. This is the "unified information entry" principle: unified at the interface layer, diverse at the implementation layer.

Principle 2 Risk-Tiered, Not Platform-Tiered Fleet-maintenance priority must be based on the equipment's business importance (the production impact of its downtime) rather than platform preference. A cobot performing a critical production task should outrank an industrial arm performing an auxiliary task, even if the engineer is more familiar with the industrial arm. Build an equipment-priority matrix based on Business Impact Score (BIS) to rank service priority uniformly across platforms.

Principle 3 Shared Parts Management, Platform-Partitioned Storage Establish a unified parts-management system that handles business logic uniformly (uniform stock alerts, procurement triggers, cost tracking) while keeping physical warehousing and system coding partitioned by platform (to prevent mixed use of parts across platforms creating safety problems). Procurement strategy can be optimised across platforms; issuance and use must be strictly platform-bound.

Principle 4 Knowledge Flows Across Platforms, Platform-Specific Knowledge Locked In the Known Fault Library, entries involving generic service logic (e.g., "communication-shield breakage is a high-frequency root cause of intermittent communication faults") are shared across platforms; entries involving platform-specific techniques (e.g., "J3 RV reducer grease-type selection") are locked under the corresponding platform tag to prevent incorrect cross-platform application.

7.7.2 Engineer Dispatch Model in a Heterogeneous Fleet Context

With a limited total of service engineers, how to optimally schedule cross-platform capability is one of the core decisions of service-organisation operating efficiency:

Dispatch Model	Organisational Structure	Applicable Scenarios	Strengths	Limitations
Platform Specialist (Specialist Model)	Each engineer specialises deeply in one platform; dispatch matches by platform	Plants with a single platform or very high industrial-arm share	Strongest technical depth; high diagnostic	High platform-switching cost; poor dispatch elasticity; minority platforms

			efficiency	hard to cover
Generalist Model	Each engineer holds basic service capability across all platforms	SMB service organisations with diverse platforms and small counts of each	Highest dispatch elasticity; high workforce utilisation	Limited technical depth; complex faults require external support
Hybrid T-Model (RSF-Recommended)	Core team is T-shaped (1–2 platform specialties + cross-platform common base), supplemented by specialist advisor support	Mid-to-large service organisations with diverse platforms and meaningful counts of each	Balances depth and flexibility; scalable	Requires deliberate capability planning and training investment; the T-shaped engineer takes time to develop
Fleet Coordinator + Specialist Support	Establish a "service-coordinator" role (non-technical, dispatch-only); outsource non-primary-platform specialist service	Large plants with 1–2 primary platforms and very few of the rest	High service quality on core platforms; clear cost structure	Non-primary-platform response time is bounded by the outsourcing partner

7.8 Looking Ahead — Service Response Strategy for Robot Diversification

Over the next 5–10 years the robotic-service industry will face three near-certain trends: continually rising platform count (especially new AMR brands and humanoid scaling); continually rising AI-component penetration (further increasing diagnostic difficulty); accelerating standards-update cadence (ISO 10218:2025 just released, ISO/TS 15066 revised edition expected 2026, dedicated humanoid standards still being developed). Service organisations must respond strategically, not passively wait for standards and procedures to mature.

7.8.1 Three Strategic Response Directions

Strategic Direction	Core Logic	Concrete Implementation Path
Build "Meta-Procedure" Capability	Rather than chasing the dedicated procedure for each new platform, build the organisational capability to "rapidly formulate dedicated procedures" — able to establish a complete service-procedure system within 6–12 months when a new platform arrives	Establish a "new-platform service-procedure construction template" (based on the RSF framework of seven unifiable layers + five necessarily-differentiated layers); cultivate Expert engineers who can lead new-platform procedure development; establish early relationships with new-platform OEMs (participate in their early-customer-service-partner programmes)
Use ROS 2 as the Middleware Layer to Advance Cross-Platform Service Standardisation	ROS 2, as the de facto open middleware standard, is being adopted by an increasing number of new robots (AMRs, cobots, humanoid robots). Service tools and procedures at the ROS 2 layer carry inherent cross-platform portability	Add ROS 2 middleware-diagnostic capability into service training (already defined in §4.5); develop ROS 2-interface-based cross-platform remote-monitoring tools; encourage customers to weight ROS 2 compatibility as a long-term serviceability indicator in procurement decisions
Active Standards Engagement	As a framework standard in the robotic-service field, RSF is uniquely positioned to feed the "service perspective" directly into international standards bodies such as ISO TC 299. Many current ISO standards have explicit service-procedure gaps in their safety-operating requirements; RSF practice can systematically fill them	Track the revision cycles of ISO TC 299 (Robots) and IEC TC 62443; submit RSF service-procedure recommendations to relevant ISO working groups through national standards bodies; pay particular attention to the development process of humanoid safety standards (the 2025–2028 critical window)

7.8.2 Anticipated Impact of Emerging Technologies on Cross-Platform Service

▸ Cross-Platform Unification Potential of AI-Assisted Diagnosis

Large language models (LLMs) and vision-language models (VLMs) are exploring the feasibility of a "general-purpose robot diagnostic assistant": the engineer describes the symptom and the AI gives cross-platform diagnostic-guidance recommendations. USC's 2024 RobotFleet study demonstrated an LLM-based unified-dispatch framework for heterogeneous fleets. Expected impact: once LLM-assist tools mature, the cognitive load on engineers when switching platforms drops significantly, and parts of Domain 5's dedicated procedures (e.g., "rapid platform-switch decision support") will be replaced by AI tools.

▸ Digital-Twin Enablement of Cross-Platform Service

As each platform builds its own digital twin, integrating multi-type robot twins into a unified fleet-level digital-twin platform becomes possible, enabling cross-platform health-state visualisation and unified predictive-maintenance management. This will resolve heterogeneous-fleet service-information fragmentation at the technical layer.

▸ Self-Maintenance Capability of Humanoid Robots

Long-term vision (post-2030): humanoid robots gain the capability to perform basic maintenance tasks on their own kind (one humanoid maintaining a damaged humanoid). The scenario remains in research today but preliminary exploration exists (Robot-to-Robot Maintenance, R2RM). RSF must think ahead about how "authorisation" and "responsibility" frameworks are defined when robots participate in robot maintenance.

7.9 Domain 5 Key Performance Indicator Framework

Indicator	Definition	Industry Benchmark Target	World-Class Target	Management Significance
Cross-Platform Service-Coverage Rate	Platforms on which the service organisation can independently provide Specialist-level service / platforms the customer actually operates	**≥ 70%**	**= 100%**	Core capability indicator; below 70 % indicates many platforms depend on external experts
Platform-Switch Preparation Time	Average time from when the engineer is assigned a cross-platform service to when they are qualified to execute	**< 4 hours (completing the Switchover Assessment)**	**< 1 hour**	Measures the response elasticity of cross-platform dispatch
Cross-Platform SLA Parity	Spread (max − min) in SLA-attainment rates across platforms	**< 15 % spread**	**< 5 % spread**	Detects platform bias in service capability; "industrial arms well served, AMRs poorly served" is a typical problem
Procedure-Compliance Rate	Share of service operations with full platform-specific procedure-execution records	**≥ 85%**	**≥ 98%**	Service operations without procedure records are a high-risk source of quality incidents
New-Platform Procedure Establishment Cycle	Time from customer introduction of a new robot model to completion of the RSF dedicated procedure	**< 6 months**	**< 3 months**	Measures the maturity of the organisation's meta-procedure capability

Humanoid Readiness Index	Composite score (0–100): procedure coverage × engineer-training completion × OEM-relationship strength × tool configuration	**Target: establish the assessment baseline (2025)**	**≥ 60/100**	A new indicator reflecting the strategic readiness for humanoid service

7.10 Chapter Summary and Inter-Domain Interfaces

Chapter Summary Domain 5 (Cross-Platform Service Procedures), with "find unity while acknowledging difference" as its core logic, builds a two-layer architecture: the unifiable service layer (seven categories of cross-platform common procedures) and the necessarily-differentiated layer (five categories of platform-specific service content). The reducer specialty of industrial arms, the torque-compliance verification of cobots, the map and battery management of AMRs, and the dynamic-balance maintenance-safety procedure of humanoid robots each carry technical specificity that cannot be ported across platforms. Humanoid robots represent the most frontier service challenge; in the absence of dedicated ISO standards, RSF provides an interim best-practice framework and a 2025–2030 procedure-build roadmap.

The T-shape is the recommended path for engineers to build cross-platform capability. The three forward-looking strategies — meta-procedure capability, ROS 2 middleware standardisation, and active standards engagement — are the core means by which service organisations stay ahead under the robot-diversification trend.

Interfaces with Other Capability Domains

Interface Direction	Interfacing Domain	Content
Domain 5 → Domain 1	Service Lifecycle Management	Domain 5 supplies platform-specific operating procedures to Domain 1's five-phase lifecycle (e.g., cobot torque-baseline establishment during commissioning, AMR SLAM-mapping procedure); platform-specific procedures are "plug-in supplements" to Domain 1's generic framework
Domain 5 → Domain 2	Fault Diagnosis & Technical Support	Domain 5 supplies platform-specific differences within Domain 2's four-layer diagnostic framework (the L3 / L4-boundary diagnostic characteristics of AMRs, the safety-circuit specialty of cobots); the platform-specific fault library is the per-platform subset of Domain 2's Known Fault Library
Domain 5 → Domain 4	SLA Design & Service Commitment	Domain 5 exposes the differences in SLA-indicator definitions across platforms (different downtime definitions, different response-time expectations), providing technical basis for Domain 4 to design platform-differentiated SLA clauses
Domain 5 ← Domain 3	Remote Service & Digital Operations	Domain 3's digital-twin and remote-monitoring capabilities must adapt to each platform's data interface (industrial-arm OPC UA, AMR VDA 5050 / MQTT, humanoid ROS 2 topics); Domain 5 defines the remote-service interface specifications per platform
Domain 5 → Domain 6	AI/ML System Service Methods	Humanoid robots (Domain 5) are the most complex application scenario of Domain 6 (AI/ML service methods) — maintenance procedures for embodied AI and reinforcement-learning policies sit at the intersection of Domain 5 (platform-specific) and Domain 6 (AI-system service)

Chapter 8

Domain 6 · AI/ML System Service Methods

8.0 Chapter Introduction — AI's Dual Role in Service

The title "AI/ML System Service Methods" carries an easily overlooked dual meaning. The usual reading is: how to service a robot that carries AI components — how to diagnose vision-perception model failures, how to assess reinforcement-learning policy drift, how to manage AI-model updates and rollbacks. This is Domain 6's first dimension and the core of its original definition.

But Domain 6 has a second dimension that represents a deeper industry transformation: AI as the foundational platform for service delivery itself — using AI to process service requests, replace manual look-up, assist engineers in diagnosis, and auto-generate work orders and maintenance prescriptions. This dimension is becoming reality far faster than expected. A 2025 paper introduces an "AI Agent-driven predictive maintenance" framework in which four agent classes — fault-detection agent, RAG-based fault-classification agent, LLM-based fault-diagnosis agent, and digital-twin disposition-simulation agent — compose an autonomous maintenance workflow that requires the engineer to turn no page of any manual.

This chapter integrates both dimensions: §8.1 defines the AI dual-role framework; §8.2 builds the robot AI/ML component classification; §8.3 establishes AI-layer fault classification and preliminary diagnosis; §8.4 develops the model-performance-evaluation methodology; §8.5 defines full-lifecycle MLOps; §8.6 — the core new dimension — covers AI as the service-delivery platform, from automated service-request processing to intelligent diagnostic assistance as a complete system; §8.7 establishes the AI data-governance and compliance framework; §8.8 defines Domain 6 KPIs; §8.9 closes with the chapter summary and inter-domain interfaces.

8.1 AI Dual-Role Framework in Robot Service

Understanding the full value of Domain 6 requires a clear "AI dual-role" conceptual framework: AI is both the object of service (the robots being serviced carry AI) and the tool of service (AI is used to improve service quality and efficiency itself). The capability-build paths for the two roles differ but reinforce each other.

AI Role	Concrete Meaning	Service-Engineer Tasks	Core Capability Requirements
Role A — AI as Object of Service	The robot internally carries AI/ML components (vision perception, motion planning, reinforcement-learning policies) that themselves must be diagnosed, evaluated, updated, and maintained	Diagnose AI-layer faults (distinguishing AI from hardware); evaluate model performance (accuracy / drift / generalisation); manage model updates (OTA / version management / rollback); manage data quality	Understand the basic principles of AI/ML (not at the developer level, but able to read model metrics); master AI-layer diagnostic tools; be familiar with the service-engineer-applicable parts of MLOps
Role B — AI as Tool of Service	AI technologies (LLM, RAG, Agentic AI) are used in the service-delivery process itself: automated service-request processing, intelligent diagnostic assistance, accelerated knowledge retrieval, automated prescription generation	Use AI-assistance tools to raise their own service efficiency: obtain diagnostic recommendations from an AI knowledge base; submit fault descriptions to AI for an initial diagnostic path; use AI to auto-generate work orders and service reports	Capability to use and evaluate AI service tools; understanding of the limitations of AI recommendations (no blind trust — maintain professional judgement); provide AI systems with high-quality work-data input

Key Distinction Between the Two Roles Role A requires "AI technical knowledge" (ability to evaluate vision-model accuracy, recognise policy drift); Role B requires "AI-tool-use capability" (effective use of LLM knowledge bases; the ability to query the AI system and critically evaluate its recommendations). Neither requires the service engineer to possess AI-model development capability, but both require the engineer to possess "AI literacy" — understanding the capability boundaries of AI tools, knowing when to trust an AI recommendation and when to override it with professional judgement.

8.2 Robot AI/ML Component Classification System

Different types of AI/ML components have entirely different failure modes and service needs. The first task when facing a "robot AI fault" is to identify which class of AI component is involved, and then choose the corresponding diagnostic method.

AI-Component Category	Technical Principle	Typical Application Scenarios	Principal Failure Modes	Service-Intervention Frequency
Perception Layer	Computer-vision CNN / ViT; 3D point-cloud processing; multi-modal perception (vision + force + depth)	Object recognition and localisation; vision-guided picking; defect detection; person detection (cobot safety)	Distribution shift (new workpieces / lighting changes); accuracy degradation; sensor-to-model mismatch	High frequency (environmental changes trigger frequently)
Planning Layer	Reinforcement-learning (RL) policies; motion-planning networks; trajectory-generation neural networks	Complex path planning; assembly-strategy generation; grasp-pose planning	Policy drift (reward hacking); out-of-distribution failure; generalisation failure	Medium frequency (policy updates on a months-scale cycle)
Decision-Making Layer	Multi-task learning; large language models (VLA: Vision-Language-Action); reinforcement-learning policy networks	Natural-language instruction parsing; task-sequence decisions; multi-robot coordination	Semantic-understanding errors; ambiguity-induced wrong actions; safety constraints bypassed	High frequency (each task change can trigger)
Control Layer	Neural-network controllers; adaptive control; physics-informed neural networks (PINN)	Compliant control (force control); dynamic balance (humanoid gait); precision motion control	Model drift (actuator aging causes physical-model deviation); insufficient robustness	Low frequency (control models are relatively stable)
Predictive Layer	LSTM / Transformer time-series models; RUL-prediction models; anomaly-detection autoencoders	RUL prediction; health-state assessment; fault early warning	Data drift (sensor aging); model degradation (training distribution mismatched to current state)	Medium frequency (quarterly model-accuracy verification)

8.3 AI Layer Fault Classification and Initial Diagnosis

The greatest challenge of AI-layer fault diagnosis is "symptom disguise": AI-layer failures often manifest externally as mechanical or electrical-layer symptoms, leading engineers to spend significant

time investigating the wrong layer. Chapter 4 §4.6 provides the preliminary AI-layer classification method; this section deepens it by establishing the complete AI-layer fault classification system and a dedicated diagnostic decision tree.

8.3.1 AI-Layer Fault — Four-Quadrant Classification Framework

Quadrant	Fault Type	Typical Symptoms	Quick-Distinguishing Feature (vs Hardware Fault)
Q1 Deterministic AI Failure	**Distribution Shift**	Consistent failure under specific conditions (new workpieces / new lighting / new environment); full recovery when conditions revert	Key fingerprint: condition-dependence. The fault correlates strongly with specific input conditions, rather than occurring randomly; hardware faults usually correlate with time / temperature / vibration and are independent of input content
Q2 Probabilistic AI Degradation	**Model Drift; Accuracy Degradation**	Success rate declines slowly from the baseline (weeks to months); no obvious trigger event	Key fingerprint: time-trend dependence. Degradation correlates with time rather than operating hours; compare the long-term trend of model-confidence logs (scores) — hardware degradation usually has a corresponding physical-quantity change in sensor data
Q3 Strategic AI Failure	**Reinforcement-learning policy drift; reward-function misalignment**	The robot completes the task but via an "unintended path"; the success rate is acceptable but the manner is anomalous (e.g., taking a longer route, over-using torque)	Key fingerprint: anomalous behaviour with acceptable outcome. Hardware faults usually cause task failure; they do not produce the "anomalous-manner-but-acceptable-outcome" pattern
Q4 Data-Quality Failure	**Data contamination; labelling errors; sensor-to-model mismatch**	Model performance degrades suddenly after an OTA update, or model behaviour becomes anomalous after a sensor replacement	Key fingerprint: event-triggered. There is an explicit trigger event (update / replacement); unlike Q2's slow drift, Q4 is abrupt degradation

8.3.2 AI-Layer Fault Diagnostic Decision Tree (Five Steps)

- Step 1 Hardware Exclusion (per the three-step method in §4.6.2): sensor hardware check + standard-workpiece test + offline-inference verification. If hardware is normal, proceed to Step 2
- Step 2 Trigger-Condition Analysis: record the input conditions of every failure case (workpiece model, lighting, environment state, operating time). Determine whether there is a common trigger condition (→ Q1 distribution shift); whether failures are random with no pattern (→ Q2 model drift); or whether there is an explicit trigger event (→ Q4 data quality)
- Step 3 Confidence-Distribution Check: examine the confidence-score distribution in the model-inference log. Confidence persistently below threshold (e.g., < 0.7) → Q1 or Q2; confidence normal but behaviour anomalous → Q3 policy drift
- Step 4 Time-Series Comparison: compare current model-performance indicators against historical baselines (the commissioning baseline and the most recent verification). Slow decline (weeks / months scale) → Q2; abrupt jump → Q4; no temporal trend, condition-dependent only → Q1

- Step 5 Escalation and Remediation-Path Selection: Q1 (distribution shift) → data augmentation and model retraining (see §8.4); Q2 (model drift) → model-performance evaluation + retraining-trigger judgement (see §8.5); Q3 (policy drift) → policy reset or retraining; Q4 (data quality) → data audit + rollback if caused by OTA (see §8.5)

8.4 AI Model Performance Evaluation Methodology

Model-performance evaluation is the core dedicated capability of AI-layer service: periodically and systematically evaluate the actual operating performance of the robot's AI components, determine whether performance remains within acceptable range, and make service decisions of "continue use / strengthen monitoring / trigger retraining" based on the evaluation conclusion. Evaluation must rest on quantitative data, not on the engineer's subjective sense.

8.4.1 Perception-Layer Model Performance Evaluation (Visual Perception Focus)

Evaluation Indicator	Definition and Calculation	Evaluation Frequency and Method	Criteria and Decision Triggers
Detection Accuracy	Object-detection accuracy on the standard test set (mAP@0.5)	Quarterly evaluation: run the model on the standard test-image set established at commissioning and record the mAP value	Below baseline by < 5 %: normal; 5–15 % drop: strengthen monitoring, evaluate monthly; > 15 % drop: trigger retraining
Confidence Distribution	Distribution shape of inference-confidence scores (ideal: a high / low bi-modal distribution without a heap of mid-range values)	Continuous monitoring: auto-compute daily confidence distribution from operating logs; mid-range share (0.4–0.6) exceeding 20 % is a warning	Degradation from bi-modal to uni-modal (mid-range pile-up): increased model uncertainty — an early signal of distribution shift
Miss Rate / False-Detection Rate (FN / FP)	Share of misses (should-detect but did not) and false detections (should-not-detect but did)	Auto-record during production (with a manual-review mechanism); aggregate monthly	For safety-relevant scenarios (person detection), miss rate must be ≤ 0.1 %; high false-detection rate causes invalid stoppages — keep within 5 %
Cross-Workpiece Generalisation Rate	First-time detection success rate on new workpieces (models not present in the training set)	Before introducing each new workpiece: run 100 standard tests with that workpiece	First-success-rate < 85 % requires data augmentation before production; < 70 % requires a dedicated new-workpiece training session

8.4.2 Planning and Control Layer Model Performance Evaluation

Evaluation of planning- and control-layer AI models is more complex than perception-layer evaluation because "success" is harder to quantify — an action sequence may complete the task but take longer, use more torque, and follow a more roundabout path; these regressions can be early signals of policy degradation:

- Task-Completion Rate: in a given time window, the share of tasks completed successfully — compared against the commissioning baseline. A drop of more than 5 % triggers evaluation
- Task-Efficiency Indicators: for similar tasks, the average completion time, average torque consumption, and average path length — compared with the commissioning baseline. Sustained efficiency deterioration (> 15 % slower than baseline) is an important signal of policy drift, even when the task-completion rate remains normal
- Behavioural-Consistency Assessment: for the same input state, tabulate the consistency (variance) of the robot's action decisions over time; rising variance indicates declining policy determinism — an early indicator of policy regression
- Safety-Boundary Compliance: record how often robot actions approach or touch safety-constraint boundaries (speed / torque / spatial limits); a rising frequency indicates declining robustness of the control policy near boundaries

8.4.3 Standardised Benchmark Test Suite

To ensure repeatability and comparability of AI-model evaluation, RSF recommends each AI-bearing robot establish a standardised Benchmark Suite at commissioning, for use in all subsequent periodic evaluations. Design principles for the Benchmark Suite:

- Standardised Test Conditions: standardised lighting (record intensity and colour temperature); standardised workpieces (record spec and condition of the standard workpieces at first delivery); standardised running program (a dedicated test program — production programs are not used as evaluation baselines)
- Test Coverage: must include "normal scenarios" (80 % of cases), "boundary scenarios" (15 %, near parameter limits), and "known hard cases" (5 %, typical inputs that historically failed) — a normal-scenario-only test set cannot detect performance regression
- Benchmark Version Management: the suite has a version number, and baseline results are bound to the suite version; if workpiece models or process parameters change materially, the suite must be updated synchronously and a new baseline established
- Automated Execution: ideally, the benchmark can be executed by the robot autonomously (no human intervention); results are auto-recorded to the digital-twin platform to support historical-trend tracking

8.5 MLOps — AI Model Full Lifecycle Management

MLOps (Machine-Learning Operations) is the full-lifecycle management methodology in the industrial-robot service context covering AI models from training to deployment, monitoring to update, and retirement to replacement. For the service engineer, MLOps does not require mastery of model-training technology — it requires understanding the lifecycle logic of an AI model so the right service action is taken at the right time.

8.5.1 Five Phases of the AI Model Lifecycle

Phase	Name	Principal Activities	Service-Engineer Role
M1	**Data Collection and**	Production-site data collection; data-	Ensure normal operation of production-

	Quality Management	labelling quality control; dataset version management; train / validation / test split	data collection equipment; collect and report raw data of failure cases (images / sensor data); feed back labelling-quality issues
M2	**Model Training and Validation**	Model-architecture selection; hyper-parameter tuning; training-convergence validation; over-/under-fit detection; benchmark-suite validation	Usually owned by the AI-development team; the service engineer is responsible for providing "acceptable-performance criteria" (based on real application needs, not academic metrics)
M3	**Deployment and Version Management**	OTA-package preparation (with code signing); staged rollout (canary → small-scope → full); post-deployment function verification; rollback-capability verification	Execute the standard benchmark after deployment; confirm deployment success; record the version update in the equipment passport
M4	**Continuous Monitoring and Drift Detection**	Real-time tracking of model-performance indicators; statistical tests for data drift; concept-drift detection; anomaly-pattern recognition	Daily monitoring-dashboard review; alarm response (§8.3 four-quadrant classification diagnosis); collect failure cases as data for model retraining
M5	**Model Retraining and Retirement Decisions**	Retraining-trigger judgement; new-dataset preparation; post-retraining evaluation; model-replacement decision; archiving of the retired model	Decide whether retraining-trigger conditions (see §8.5.3) are met; provide field-feedback data; verify the new model's performance after actual deployment

8.5.2 Model Drift Detection Methodology

Model Drift is the phenomenon of AI-component performance degrading over time after deployment, in three classes: input-data-distribution shift (Data Drift), model-output distribution change (Concept Drift), and overall performance regression (Performance Drift). Recent research (Adaptive Multi-Dimensional Monitoring, AMDM) lowers anomaly-detection latency from 12.3 s to 5.6 s and the false-alarm rate from 4.5 % to 0.9 %, indicating that real-time drift detection is now industrially deployable.

Drift Type	Detection Method	Robot-Service-Context Example	Response Strategy
Data Drift	Statistical-distribution tests (KS test, PSI index); monitoring of input-feature distribution changes	Factory changes workpiece supplier → workpiece-surface-gloss distribution changes → vision-model input distribution shifts	Alert + raise evaluation frequency; if the shift is significant, trigger data-augmented retraining
Concept Drift	Monitoring of model-output distribution (predicted-class distribution / confidence distribution); comparison of output differences between production data and training data	Process change causes the definition of the "correct answer" to change → the model's correct-classification criteria are out of date	Re-label data; update the training set; retrain the model
Performance Drift	Periodic execution of the benchmark suite; trend analysis of key performance indicators (KPIs)	Over time, vision-model accuracy declines slowly from 92 % to 83 %; no obvious trigger event	Analyse the root cause of the decline (data quality? sensor aging?); take targeted action once the cause is identified

8.5.3 Model Retraining Triggers and Handling Procedures

Model-retraining decisions must rest on objective trigger conditions, not on the subjective sense that "the model isn't as good as before." Retraining evaluation must be initiated if any of the following is satisfied:

- Decline in key performance indicators: core evaluation indicators (mAP, task-completion rate, etc.) decline beyond a preset threshold from the commissioning baseline (typically 15 %)
- Sustained drift-detection alarms: the PSI (Population Stability Index) exceeds 0.25 for two consecutive evaluation cycles (PSI > 0.25 is the industry-accepted major-drift threshold)
- Accumulation of failure cases: the system has accumulated more than 50 human-confirmed real failure cases — sufficient to form meaningful retraining supplementary data
- Major change in business conditions: new workpiece types exceed the current training-set coverage by more than 30 %; the production environment (lighting / background / workpiece-size distribution) changes fundamentally
- Performance regression introduced by OTA: after a new model-version deployment, benchmark-suite evaluation shows performance below the previous version — execute rollback and trigger a new-version quality review

Model-Version Rollback Procedure

- Rollback Trigger: within 24 hours of new-version deployment, the benchmark-suite result is below 95 % of the previous-version baseline (i.e., performance drop > 5 %)
- Rollback Execution: via the OTA channel, revert the model version to the most recent validated version (the previous two versions must be retained in device storage); after rollback, re-execute the benchmark for verification
- Rollback Record: record rollback reason, description of the new version's issue, and benchmark results before and after rollback; submit to the AI-development team for new-version issue analysis
- Rollback Notification: notify the customer (with disposition and expected impact); record in the equipment passport

8.6 AI as Service Delivery Platform — Disruptive Restructuring of Service Delivery

Major Expansion Perspective of Domain 6 AI is not only a component being serviced on the robot — AI itself is becoming the foundational platform of robot-service delivery. From intelligent service-request processing to AI-assisted diagnosis, to real-time knowledge-base retrieval, to auto-generated maintenance prescriptions — AI is upgrading robot service from "experience-driven" to "data + intelligence-driven." This transformation is not the distant future; it is the reality that entered industrial deployment in 2024–2025.

8.6.1 AI-Powered Service Request Management

The traditional service-request flow: customer phones in fault → operator records → manual classification → dispatch. This flow has three systemic bottlenecks: incomplete information extraction (the operator may not understand technical detail, missing key diagnostic information); classification accuracy depends on individual experience; dispatch decisions lack real-time optimisation against engineer state and geographic information. AI service-request systems solve all three:

AI Capability	Implementation	Improvement over Traditional Flow
Multi-Modal Fault-Information Extraction	NLP extracts key information from the fault description (alarm codes, symptom description, occurrence time, severity); vision AI recognises fault phenomena from customer-uploaded site photos / video	Information-extraction completeness rises from about 60 % (manual recording) to 90 %+; reduces pre-arrival information-confirmation round-trips; first-visit diagnostic efficiency rises significantly
Intelligent Fault Classification and Prioritisation	A classification model trained on historical work-order data auto-maps the fault description to priority (P1–P4) and preliminary layer localisation (L1–L4); matches the Known Fault Library (RAG retrieval, see §8.6.2)	Classification accuracy above 85 % (Specialist level); P1-fault recognition speed drops from 3–5 minutes (manual) to under 30 seconds; reduces response delays caused by priority misjudgement
Intelligent Dispatch Optimisation	Jointly considers engineer certification tier (must cover the platform), current location (nearest available engineer), current workload (avoid overload), and parts availability (whether critical spares are on the engineer's vehicle)	Theoretical dispatch efficiency rises 30–40 % over manual optimisation; "nearest engineer" is upgraded to "best-matched engineer"
Automated Customer Communication	After work-order opening, auto-send the customer a confirmation (work-order number, expected response time, engineer info); auto-push progress updates; auto-generate reports	Higher customer satisfaction (transparent information); the service manager is freed from heavy routine communication to focus on exception cases

8.6.2 RAG-Powered Intelligent Knowledge Assistant (Beyond Manual Lookup)

"Look up the manual when a fault occurs" has long been the base working model of the robot-service industry. The model has a fundamental efficiency bottleneck: engineers must manually search hundreds of pages of technical manuals; after manual updates, they must re-learn; knowledge across multiple brands and platforms must be accumulated separately and cannot be reused. RAG (Retrieval-Augmented Generation) fundamentally restructures this model:

RAG Knowledge-System Architecture

Knowledge Layer	Knowledge Source	Role in Diagnostic Assistance
Structured Knowledge	OEM technical manuals across brands (vector-stored); international-standards text (ISO 10218 / ISO/TS 15066, etc.); alarm-code database; RSF framework documents	Engineer asks "What is FANUC SRVO-023 alarm?" → RAG retrieves the relevant manual sections from the vector database → LLM synthesises concise diagnostic suggestions and operating steps, rather than asking the engineer to leaf through the manual

Experiential Knowledge	Internal work-order history (structured fault records); Known Fault Library (KFL); engineer repair notes (unstructured)	Engineer describes the symptom → RAG retrieves similar historical work-order cases → recommends "the known root cause and most effective remediation for this fault type" → the engineer saves independent-reasoning time and directly verifies the historical conclusion
Real-Time Operational Knowledge	Digital-twin real-time state data; current equipment-health index; recent alarm history; real-time spare-parts inventory state	"What is this unit's current health state? What anomalies appeared in the past 30 days? Is part X in stock?" — the RAG system fuses real-time operational data with structured knowledge to give scenario-specific diagnostic recommendations rather than generic answers

Norms for Using RAG Knowledge Systems (Critical: engineers must not blindly trust AI recommendations)

> Trust Boundaries of AI Recommendations A RAG knowledge system is a diagnostic-assistance tool, not a diagnostic-decision tool. The distinction: tool output must be verified by the engineer's professional judgement, not executed directly.
>
> In particular: (1) AI-retrieved historical cases may not apply to the current concrete situation (equipment-model differences, application-scenario differences); (2) AI recommendations for "novel fault types" (no matching record in the knowledge base) are significantly less reliable and require special caution; (3) safety-related operations (diagnostics and repairs involving the safety circuit) must be independently judged by the engineer — sole reliance on AI recommendations is forbidden. AI raises diagnostic efficiency, but the responsibility for professional judgement remains with the engineer.

8.6.3 Agentic AI-Driven Predictive Maintenance Workflow

Recent research published in 2025 ("AI Agent-Enabled Predictive Maintenance," MDPI Computers, August 2025) proposes an autonomous-maintenance workflow architecture composed of several specialised AI agents — the most frontier application form of AI in robotic service. RSF integrates this architecture into a practical guide:

Fault-Detection Agent

Core Responsibilities	Continuously analyse multi-sensor data streams (vibration / current / temperature / torque) and recognise anomalous patterns
Data Input	Real-time sensor data (obtained from the digital-twin data infrastructure)
Typical Output	"J3 axis current rose 18 % over the past 72 hours, exceeding the normal fluctuation range (±10 %); judged as an early-degradation alarm (P3)."
Service-Engineer Notes	Replaces the traditional "engineer periodic inspection"; 24/7 uninterrupted monitoring; false-alarm rate must be controlled within 5 % through continuous training

Fault-Classification Agent

Core Responsibilities	Use RAG to retrieve the Known Fault Library, match the detected anomaly to historical fault modes, and output the most-likely root-cause hypothesis
Data Input	Fault-Detection Agent alarm + historical-alarm sequence + the Known Fault Library (KFL)
Typical Output	"Best-matched historical pattern: KFL-001 (J3 reducer grease depletion, historical hit rate

	82 %). Recommended diagnostic path: inspect grease condition + measure backlash."
Service-Engineer Notes	RAG retrieval is efficient (millisecond-class); classification accuracy depends on KFL entry quality; novel fault types require human intervention for classification

Diagnosis-Reasoning Agent

Core Responsibilities	Integrate multi-source information (detection data + classification conclusion + equipment passport + real-time state), perform combined reasoning, and output a diagnostic report and remediation recommendation
Data Input	Classification-Agent output + equipment passport (commissioning baselines) + digital-twin current state + relevant ISO standards clauses
Typical Output	A complete AI Diagnostic Report: symptom description, root-cause hypothesis (with confidence), recommended remediation steps, required spares, estimated labour hours — for engineer review and confirmation
Service-Engineer Notes	The key agent in the evolution from "AI-assisted" to "AI-led preliminary diagnosis"; output must be submitted to the engineer for review and is not auto-executed

Simulation Agent

Core Responsibilities	In the digital-twin environment, simulate the proposed maintenance operation and predict operational risk and expected outcomes
Data Input	Diagnosis-Reasoning-Agent remediation recommendation + digital-twin model
Typical Output	"Simulation result: after the J3-axis grease replacement, current is expected to return to within ±5 % of baseline and the health index to rise from 62 to 85; operational risk: low (standard PM operation, no change to safety-critical steps)."
Service-Engineer Notes	Verify the correctness of the remediation plan in the digital twin to reduce real-operation risk; simulations of complex operations still require Expert-tier engineer review

8.6.4 AI-Driven Automatic Service Report Generation

Writing service reports occupies about 20–30 % of an engineer's non-technical time in the traditional model. An AI LLM can auto-generate RSF-format-compliant service reports from work-order records, operation logs, and diagnostic data. Two key design points for automated reports:

- Source Traceability of Content: each data point in the auto-generated report (MTTR, part numbers, operating steps) must trace to a corresponding work-order record or sensor data; AI must not "fill in" unsupported content (the hallucination risk of LLMs must be strictly guarded against here)
- Engineer-Review-and-Sign-Off Mechanism: the AI-generated draft must be reviewed, confirmed, and signed by the executing engineer before being sent to the customer; the engineer's signature represents professional confirmation of report content, not merely format confirmation — a legal and professional requirement for completeness of the responsibility chain

8.7 AI Data Governance and Compliance Framework

Data governance for robot AI systems must establish norms at two layers: the service-object layer (AI-model data of the serviced robot) and the service-platform layer (data used by AI service tools). Both touch the core issues of data privacy, data ownership, and AI-system reliability.

8.7.1 Key Data Governance Rules for AI/ML Systems

- Ownership of AI Training Data: image and sensor data collected from the customer's production site for AI-model training must have the customer's explicit authorisation written into the contract (explicit, not default); how the data are stored after training and whether they may be used to train models for other customers must be explicitly agreed
- Explainability Requirements for AI Systems: AI components involved in safety decisions (cobot force-protection, AMR obstacle avoidance) must have decision logic that can be audited and explained; "a black-box AI made a safety decision but cannot explain why" carries compliance risk under ISO 10218 and ISO/TS 15066
- Change Control for AI-System Changes: AI-model updates (OTA) must be incorporated into the controlled-change process of Chapter 3 (Domain 1); an unauthorised model update violates change-management procedures in the same way as an unauthorised hardware-parameter modification
- Data-Input Restrictions on AI-Assist Tools: when using a RAG knowledge system or an LLM diagnostic assistant, engineers must not feed customer-confidential production data (process parameters, output volume, product specifications) into third-party cloud AI systems; sensitive data must be handled with privately deployed AI tools
- Applicable Clauses of ISO/IEC 42001: ISO/IEC 42001:2023 (AI Management System) provides a compliance-management framework for robotic AI systems; its core requirements include AI risk assessment, AI-system change management, AI-system monitoring, and AI-incident recording — these align closely with RSF service procedures and can be satisfied in parallel

8.7.2 AI System Security Considerations

AI Safety-Risk Type	Manifestation in Robotic Service	Defensive Measures in the Service Procedure
Adversarial Attack	Malicious inputs (specific-pattern stickers) deceive the vision model, causing wrong object recognition or person-detection failure	Run adversarial-robustness assessment on safety-relevant vision components (person detection); vision sensors in the production environment require physical protection
Model Poisoning	Supply-chain attack: malicious samples are injected into the training data of a pre-trained model, causing anomalous behaviour under specific conditions	Model-source verification (digital signatures); full benchmark-suite verification before new-model deployment; OTA transport-layer encryption
AI Decision Opacity	An RL policy executes a dangerous action, but the engineer cannot understand from the log "why this decision was made"	Critical safety decisions require explainable logs; high-risk actions require an explainable "decision-reason" output beforehand; a Safety Monitor system (independent of the AI) provides hard-real-time safety assurance
Over-Automation Risk	An AI agent autonomously triggers a maintenance operation, but the conclusion is based on an incorrect diagnosis, leading to wrong parameter modifications or unnecessary downtime	All AI-agent outputs must be reviewed by an engineer before execution; AI agents are forbidden to modify safety parameters without human confirmation

8.8 Domain 6 Key Performance Indicator Framework

Domain 6's KPI system covers two dimensions: AI-system service quality (Role A — performance of AI components being serviced) and AI-service-platform effectiveness (Role B — the contribution of AI as a service tool).

Table 8.8-1 Domain 6 KPI System (Dual-Dimensional)

Dimension	Indicator	Definition	Industry Benchmark	Excellent Target	Management Significance
Role A — Servicing AI	**AI-Model Benchmark Pass Rate**	Share of AI components meeting commissioning-baseline performance targets in periodic evaluations	**≥ 85%**	**≥ 95%**	Core quality indicator; a low pass rate indicates undermaintained AI components or product-reliability issues
Role A — Servicing AI	**Drift-Detection Timeliness Rate**	Share of cases where retraining is initiated within the prescribed time (30 days) after model performance hits the retraining-trigger threshold	**≥ 90%**	**≥ 98%**	Measures MLOps response efficiency; failure to handle drift in time leads to continuing performance deterioration
Role A — Servicing AI	**Model-Update Success Rate**	Share of OTA model updates that succeed on first push (no rollback required)	**≥ 95%**	**≥ 99%**	Reflects MLOps-process maturity; a low success rate impacts service continuity
Role A — Servicing AI	**AI-Layer Fault-Classification Accuracy**	Share of AI-layer faults correctly classified into the four quadrants (Q1–Q4)	**≥ 80%**	**≥ 90%**	Measures the engineer's AI-layer diagnostic capability; a low rate wastes remediation resources
Role B — Service Platform	**AI-Processed Service-Request Rate**	Share of work orders for which AI autonomously completes classification and initial dispatch	**≥ 60%**	**≥ 85%**	Measures the maturity of the AI service-request platform; a high rate frees service-manager resources
Role B — Service Platform	**RAG Knowledge Hit Rate**	Share of engineer queries to the RAG system whose returned results are judged "of reference value"	**≥ 70%**	**≥ 85%**	Measures knowledge-base quality; a low hit rate indicates the knowledge base needs expansion or optimisation
Role B — Service Platform	**AI-Assisted-Diagnosis Acceptance Rate**	Share of AI diagnostic recommendations accepted (referenced or adopted) by engineers	**≥ 50%**	**≥ 70%**	A rate too low indicates poor AI-recommendation quality; too high may indicate engineer over-reliance (AI-recommendation quality must remain a focus)
Role B — Service Platform	**Report-Automation Share**	Share of work-order reports where the AI generates a draft and the engineer revises and confirms	**≥ 50%**	**≥ 80%**	Measures the AI penetration of report work; the engineer's revision amount must be monitored (a large revision implies poor draft quality)

8.9 Chapter Summary and Inter-Domain Interfaces

Chapter Summary Domain 6 (AI/ML System Service Methods) carries the most technologically frontier content of the RSF framework and is also the capability domain that will see the fastest change over the next decade. This chapter establishes the "dual-role" framework for AI in robotic service: as AI components being serviced (Role A), it builds the four-quadrant fault classification,

the model-performance-evaluation methodology, and full-lifecycle MLOps procedures; as the platform of service delivery (Role B), it builds a complete system covering AI service-request processing, RAG knowledge assistance, four classes of Agentic-AI maintenance workflow, and AI report automation.

The core contribution of AI is to upgrade robotic service from the efficiency-bottleneck mode of "manual lookup + experience judgement" into the high-efficiency collaboration mode of "data-driven + intelligent assistance + engineer professional judgement." However the technology evolves, the engineer's responsibility for professional judgement cannot be replaced by AI.

Interfaces with Other Capability Domains

Interface Direction	Interfacing Domain	Content
Domain 6 ← Domain 2	Fault Diagnosis & Technical Support	Domain 2's AI-layer preliminary diagnosis (§4.6) is Domain 6's upstream trigger: once Domain 2 completes hardware exclusion and confirms an AI-layer fault, the fault enters Domain 6's dedicated handling flow; Domain 6's four-quadrant classification deepens Domain 2's preliminary layer-localisation conclusion
Domain 6 ← Domain 3	Remote Service & Digital Operations	Domain 3's multi-sensor data infrastructure is the data source for Domain 6's model monitoring; Domain 3's OTA framework is the execution channel for Domain 6's model-update deployment; Domain 3's digital twin is the simulation environment for Domain 6's disposition-simulation agent
Domain 6 → Domain 1	Service Lifecycle Management	Domain 6's AI-layer health-assessment conclusions influence Domain 1's PM plan (AI-layer degradation is one of the triggers for entering Phase 4); Domain 6's AI model-version history must be recorded in the equipment passport (Domain 1's configuration-baseline system)
Domain 6 → Domain 4	SLA Design & Service Commitment	AI-component performance degradation affects OEE (quality-rate loss); the complexity of AI-layer maintenance must be handled separately in SLA-responsibility allocation; the model-performance KPIs established by Domain 6 form the technical basis for AI-related SLA clauses
Domain 6 ↔ Domain 5	Cross-Platform Service Procedures	AI-component architectures differ substantially across robot families (industrial-arm vision modules vs AMR SLAM algorithms vs humanoid VLA models); Domain 5 supplies platform-specific procedural interfaces for each platform's AI layer; Domain 6 supplies the generic methodology framework


RSF
Robotics Service Framework

Appendices

Appendix A	**Standards-to-Six Capability Domains Mapping Matrix**	*Standards Mapping Matrix*
Appendix B	**Automotive 4S System — Adoption Logic and Transformation Boundaries**	*Automotive 4S System Reference*
Appendix C	**RSF Innovation Contribution and Competitive Positioning**	*RSF Innovation Contribution*
Appendix D	**Glossary of Terms**	*Glossary*


RSF
Robotics Service Framework

Appendix A

Standards-to-Six Capability Domains Mapping Matrix

This matrix systematically presents the correspondence between the existing international standards system and the six RSF capability domains: what each standard covers, which RSF capability domain it primarily relates to, and what specialized content RSF supplements in that domain beyond the standard's scope. Reading guide: the domain listed under "Primary Associated Domain" is where the standard contributes most directly to RSF; the "RSF Supplementary Content" column shows content that existing standards cannot cover and that has been filled in originally by RSF.

A.1 Safety and Compliance Standards

Table A-1 Mapping of Safety and Compliance Standards

Standard No.	Scope	Primary Associated Domain	Standard Coverage	RSF Supplementary Content (beyond the standard)
ISO 10218-1/-2 :2025	Industrial robot-arm safety requirements (design layer + system-integration layer)	Domain 1 (Lifecycle), Domain 2 (Diagnosis), Domain 5 (Platform Procedures)	Robot design safety requirements; system-integration safety specifications; safety-guard design; emergency-stop requirements; safety-function classification (STO / SS1 / SS2)	Tiered competency assessment of service engineers; on-site maintenance SOPs; step-by-step safety-state verification procedures during the maintenance phase; service-operation safety procedures for cobot scenarios
ISO/TS 15066 :2016	Collaborative-robot (cobot) safety requirements	Domain 2 (Diagnosis), Domain 5 (cobot-specific)	Power and Force Limiting (PFL) human-contact limits; Speed and Separation Monitoring (SSM); Hand Guiding (HG); Safety-Rated Monitored Stop (SRMS)	Torque-sensor temperature-compensation management procedure; quarterly torque-compliance verification method; standardized procedure for establishing torque baselines; early-warning mechanism for cobot safety-parameter degradation
IEC 61508 (SIL assessment)	Functional safety of electrical / electronic / programmable safety systems	Domain 1 (Installation Verification), Domain 2 (Safety-Circuit Diagnosis)	Quantitative SIL evaluation method; functional-safety lifecycle; failure-rate calculation for safety functions	Safety-function test procedure for robot service engineers; step-by-step re-verification checklist for safety functions after repair; SIL applicability guide for service scenarios
ISO 45001 :2018	Occupational health and safety management system	Domain 1 (Full Lifecycle), Domain 5 (Universal Safety Procedures)	Occupational health-and-safety risk assessment; LOTO (energy-isolation) management requirements; safety-training record management	Dedicated LOTO procedure for robot field service; pre-condition checklist for safe service operations; SVP-1 Universal Safety Verification Procedure (a core component of UCSP)
IEC 62443 :2024	Industrial control-system cybersecurity	Domain 3 (Remote-Service Security Architecture)	OT-network Zone / Conduit segmentation model; Security Level (SL 1–4) definitions; organizational cybersecurity-program requirements (added in the 2024 edition)	DMZ architecture specification for robot remote service; zero-trust remote-access implementation for robot service scenarios; six-step remote-session management procedure; IEC 62443 compliance path for OTA updates

ISO/IEC 42001 :2023	AI management system	Domain 6 (AI / ML Service)	AI-system lifecycle management requirements; AI risk-assessment framework; organizational governance requirements for AI systems	Model performance-evaluation procedure for robot AI components; four-quadrant classification framework for AI-layer faults; engineer-applicable MLOps clauses for robot-service scenarios; compliance boundaries of Agentic AI maintenance workflows

A.2 Asset and Quality Management Standards

Table A-2 Mapping of Asset and Quality Management Standards

Standard No.	Scope	Primary Associated Domain	Standard Coverage	RSF Supplementary Content
ISO 55001 :2014	Full-lifecycle asset management	Domain 1 (Lifecycle Management — main axis)	Asset-management objectives and planning; asset-performance monitoring; asset-retirement decision framework; maintenance-strategy requirements	Five-phase robot service-lifecycle model; phase-transition Stage Gate criteria; four-dimensional retirement-assessment matrix; Equipment Configuration Baseline system
ISO 9001 :2015	Quality management system	Domain 1 (Service-Documentation System), Domain 4 (SLA Compliance)	Process control and documentation requirements (Clause 8.5); customer-satisfaction management; non-conformance management; continual improvement (PDCA)	Five-field standard for robot work-order records; work-order closure standard (dual-signature regime); service-improvement closed loop (fault response → root cause → controlled change); format standard for monthly SLA reports
ISO 31101 :2023	Safety management for service-robot application scenarios	Domain 1 (Lifecycle), Domain 4 (SLA Design)	Operating-environment risk assessment for service robots; personnel-protection requirements for operating sites; incident-reporting and analysis framework	Risk-adjusted SLA design for service robots; on-site service safety-assessment steps for service robots; adaptation of service procedures to non-manufacturing environments (hospitals / retail / home)
ISO 3691-4 :2023	Mobile-robot (AMR) safety requirements	Domain 5 (AMR-specific procedures)	AMR navigation-safety requirements; emergency-stop specification; geofencing requirements; speed limits when co-existing with personnel	AMR map-management procedure (mapping requirements / validity period / version control); AMR battery-health management procedure; VDA 5050 interface-compatibility verification; service-engineer operating standards for AMR fleet coordination
ISO 13482 :2014	Personal-care robot (service-robot) safety	Domain 5 (Humanoid-Robot Reference)	Safety requirements for personal-care robots; contact-force limits; personnel-detection requirements	Starting-point framework for humanoid-robot service (extending the spirit of ISO 13482); SVP-1 specialized safety-verification steps for

				humanoid robots; fall-risk assessment procedure for humanoid robots

A.3 Information Security and Digital-Operations Standards

Standard No.	Scope	Primary Associated Domain	Standard Coverage	RSF Supplementary Content
ISO 27001 :2022	Information-security management system	Domain 3 (Remote-Service Data Security), Domain 6 (AI Data Governance)	Information-asset classification and access control; remote-access security requirements (Clause A.6.7, added in the 2022 edition); audit-log requirements; supplier-relationship management	Six-step management procedure for robot remote-diagnostic sessions; remote-access data-classification framework (L1–L4); compliance path for cross-border data transfer in robot service; authorization standard for use of AI training data
ISO 23247 (Digital Twin)	Manufacturing digital-twin reference architecture	Domain 3 (Digital Twin)	Reference architecture for manufacturing digital twins; data-acquisition interface specification; model-update mechanism	Service-function definitions for the four-layer Robot Digital Twin (RDT) system architecture; four mechanisms of physical-digital synchronization; digital-twin-based AI maintenance-agent workflows; RDT Health Index (HI) calculation and alerting rules
VDA 5050 (AMR Communication)	AGV / AMR fleet-communication protocol (industry standard)	Domain 5 (AMR Fleet Management), Domain 3 (Remote Service)	Communication-interface definition between AGV / AMR and fleet-management systems; order management; status-report format; error classification	Multi-brand AMR service-data aggregation method based on VDA 5050; VDA 5050 version-compatibility verification procedure; migration-path guide between ISO 21423 (in development) and VDA 5050

A.4 Summary of Standards Coverage Density Across the Six Capability Domains

Table A-4 Standards Coverage Density and RSF Contribution Weight by Capability Domain

Capability Domain	Existing Standards Coverage Density	Primary Covering Standards	RSF Contribution Weight
Domain 1 Service Lifecycle Management	★★★★ (High)	ISO 55001 (Asset Management), ISO 9001 (Quality Management), ISO 45001 (Safety Management), ISO 10218 (On-Site Operational Safety)	Medium: standards provide the principle framework; RSF provides the five-phase model, Stage-Gate criteria, and operational procedures
Domain 2 Fault Diagnosis & Technical Support	★★☆ (Low–Medium)	ISO 10218 (Safety-Function Verification), IEC 61508 (Functional-Safety Assessment)	High: four-layer stratified diagnostic framework, diagnostic-funnel model, seven-category tool-capability framework, case library — all RSF originals

Domain 3 Remote Service & Digital Operations	★☆ (Low)	IEC 62443 (OT Cybersecurity), ISO 27001 (Information Security), ISO 23247 (Digital-Twin Reference Architecture)	Very high: S1–S4 evolution model, four-layer RDT architecture, predictive-maintenance system, OTA management procedures — all RSF originals
Domain 4 SLA Design & Service Commitment	☆ (Very Low)	No ISO / IEC standard directly applicable	Very high: robot SLA-metrics framework, four-tier SLA product system, SLA strategies for emerging companies, compensation mechanisms — all RSF originals
Domain 5 Cross-Platform Service Procedures	★★ (Low–Medium)	ISO 10218 (industrial arms), ISO/TS 15066 (cobots), ISO 3691-4 (AMR), ISO 13482 (service robots)	High: standards are platform-isolated; RSF provides a unified cross-platform framework (UCSP), the PSC platform-card system, and a heterogeneous-fleet management architecture
Domain 6 AI/ML System Service Methods	★☆ (Very Low)	ISO/IEC 42001 (AI Management System, governance level only)	Very high: AI four-quadrant fault classification, model-performance evaluation, MLOps service-engineer applicability framework, Agentic AI service platform — all RSF originals

Matrix Pattern Existing standards cover Domain 1 (Service Lifecycle Management) most densely, and RSF's relative contribution there is lowest (mainly operational procedures). Domain 4 (SLA Design) and Domain 6 (AI / ML Service) have almost no existing-standard coverage, and RSF's contribution in those two domains is almost entirely original. This distribution maps directly onto the five capability-layer needs described in §1.1 — technical evidence of where RSF's contribution is most needed.


RSF
Robotics Service Framework

Appendix B

Automotive 4S System — Adoption Logic and Transformation Boundaries

The inner layer of the RSF three-tier nested model — the Operational-Pattern Reference layer — takes its primary reference from the automotive 4S service system. This appendix systematically explains what RSF references from the 4S system, why, and where the boundaries lie. Both questions matter equally: the first establishes the legitimacy of the reference; the second prevents distortion from mechanical transplantation.

B.1 Premise of Adoption: Structural Isomorphism of Business Models

The fundamental reason the automotive 4S system serves as a core reference for RSF is the high structural isomorphism of the business model. In service economics, an industrial robot and an automobile are very similar: both are high-unit-cost hardware products (complex purchase decision, service life of 10+ years), both have maintenance costs that represent a significant share of total lifecycle cost, and both customers lack self-repair capability. These shared properties give rise to the same commercial demands: a specialized service network, authorized-technician certification, spare-parts supply-chain management, and quantifiable service-quality assessment. The automotive industry spent decades validating and iterating a service-operations system that meets those demands across millions of service touchpoints worldwide — and that is exactly where the 4S system's core value to RSF lies.

B.2 Four Core Items Adopted

Item 1 — Tiered Technician Certification Structure (ASE Certification Logic)

The core design of ASE (Automotive Service Excellence) certification is: classification by specialty, tiering by capability level, and a mandatory link between certification status and authorized scope of work. Holding a specific ASE certificate means the holder is authorized to independently perform the corresponding type of repair — certification is permission, not merely a credential. RSF's four-tier certification system (Professional / Specialist / Expert / Master) adopts this structural logic directly: the four tiers correspond to a progression in cognitive dimension; the six capability domains replace the specialty-direction classification used in automotive repair; each certification tier maps to a clearly defined authorized scope of work. The design principle of "linking certification to authorization" is transferred intact from the ASE system.

Item 2 — Service Closed-Loop Process Architecture

The 4S service blueprint standardizes a single repair as a reproducible end-to-end workflow: appointment registration → vehicle reception → fault diagnosis → repair authorization → repair execution → quality inspection → customer delivery → service follow-up. RSF converts this closed-loop logic into the basic architecture of robot field service (see §3.5): on-site assessment → fault diagnosis → work-authorization confirmation → maintenance execution → functional verification → service-record archiving → customer confirmation → remote follow-up. The conversion preserves the closed-loop logic and the data-accumulation principle while replacing the technical content of each node.

Item 3 — Quantified Customer-Satisfaction Tracking Mechanism

The automotive industry's CSI (Customer Satisfaction Index) and NPS (Net Promoter Score) systems provide quantifiable, geographically comparable, and decision-actionable service-quality measures. RSF brings the CSI / NPS mechanism into two layers: the annual review of ATC certification training centers (student-satisfaction NPS ≥ 8.0 / 10 is one of the pass criteria), and the customer-satisfaction metric (CSAT) within the SLA performance-tracking system (see §6.8).

Item 4 — Parts-Inventory Management Model (ABC Classification)

The automotive 4S system's ABC classification for parts inventory: Class A — high-frequency, high-value parts (high safety stock, highest priority); Class B — medium-frequency, medium-value parts (standard safety stock); Class C — low-frequency, low-value parts (order-triggered procurement). The parts-strategy design in RSF Domain 1 (Service Lifecycle Management) adopts this classification framework, applying differentiated safety-stock policies to common robot consumables (reducer grease, batteries, encoders, driver boards) under the ABC principle.

B.3 Transformation Boundaries: Three Non-Transferable Areas

There are clear technical boundaries on how far the 4S system can be adopted. Crossing these boundaries distorts the framework, because the technical assumptions of automotive service differ fundamentally from those of robot service:

Non-Transferable Area	4S System Assumption	Reality of Robot Service and RSF Handling
Fault-Diagnosis Logic	Standard OBD-II interface; fault modes discrete, finite, and deterministic; clear mapping from fault codes to troubleshooting trees	Robot fault modes (especially in AI-integrated systems) are probabilistic — vision-perception distribution shift and RL policy drift cannot be mapped to deterministic fault codes. RSF Domain 2 establishes the four-layer stratified diagnostic framework (original); Domain 6 establishes the AI four-quadrant fault classification (original).
Software and Firmware Service Procedures	ECU flashing is relatively simple; vehicle-model update cycles are measured in years; a clear rollback procedure exists for failed software updates	Robot software service is far more complex than automotive: ROS-node dependency management, AI-model version control, multi-system firmware-compatibility verification, safe rollback after a failed OTA push — none of which has any counterpart in the 4S system. The OTA-management section of RSF Domain 3 is entirely original.
Multi-Brand Mixed-Fleet Management	A 4S dealership serves a single brand; technician training, tooling, and spare-parts supply chains are all optimized for that brand	Robot service providers typically manage mixed fleets of industrial arms, AMRs, and cobots — cross-brand service coordination has no precedent in the 4S system. The Unified Cross-Platform Service Procedure (UCSP) framework in RSF Domain 5 draws its reference from service-integration methodology (the spirit of the SIAM framework) rather than from the 4S system.

B.4 Summary Table of Adoption Items

Table B-1 Summary of 4S Adoption Items and Transformation Boundaries

Adoption Dimension	Source in the 4S /	RSF-Transformed Content	Transferability

	Automotive System		
Tiered Technician Certification	ASE certification tier linked to authorized scope of work	Four-tier Professional / Specialist / Expert / Master certification linked to authorized domain scope	Structure transferred intact, content replaced
Service Closed-Loop Process	Seven-step 4S service-blueprint closed loop	Eight-step robot field-service procedure	Logic transferred, node content rewritten
Lean Site Management	Toyota TPS · 5S · SOW standardized work	Service-SOP documentation · standardized tool carts	Methodology applied directly
Quantified Customer Satisfaction	CSI / NPS linked to authorization status	Training NPS · ATC review system · SLA CSAT metric	Mechanism transferred, metrics localized
Parts-Inventory Management	ABC classification · safety-stock model	Tiered high-frequency robot-parts inventory and stock baseline	Model transferred, categories redefined
Fault-Diagnosis Logic (✗)	OBD-II · deterministic fault tree	Not transferable; Domain 2 and Domain 6 are entirely original	Not transferable
Software / Firmware Service Procedures (✗)	Standard ECU-flashing workflow	Not transferable; Domain 3 OTA management is entirely original	Not transferable
Multi-Brand Mixed-Fleet Management (✗)	Single-brand 4S exclusive-dealership model	Adopts the spirit of the SIAM framework, not the 4S system	Not transferable; reference frame replaced


RSF
Robotics Service Framework

Appendix C

RSF Innovation Contribution and Competitive Positioning

Formal Innovation Statement

Using the ISO / IEC compliance system as its outer boundary and the operational maturity of the automotive 4S service system as a reference, RSF is the first framework to construct a professional service-capability certification covering the full range of modern robotic systems (including AI / ML integration), filling the systemic industry gaps in service-personnel assessment, AI-layer service procedures, and robot SLA design.

C.1 Technical Evidence for "First-Constructed"

"First-constructed" is a factual statement, not a value claim. Across the following three dimensions, no public framework with equivalent coverage existed prior to RSF:

Dimension	Industry State before RSF	RSF First Contribution
Service-Engineer Competence Certification	ISO 10218 calls for "competent persons" but does not define capability levels; no ISO / IEC standard provides a tiered assessment framework for robot service engineers; no comparable international certification system exists	Established the Professional / Specialist / Expert / Master four-tier certification framework; defined knowledge requirements, practical-skill standards, and authorized work scope for each tier; established certification-linked continuing professional development (CPD)
AI / ML Component Service Procedures	ISO 10218:2025 covers deterministic control systems; ISO/IEC 42001 stops at the governance layer; no standard or framework defines field-diagnosis methods for AI-layer faults, model-performance evaluation procedures, or MLOps procedures applicable to service engineers	Established the four-quadrant AI-layer fault classification framework (Q1–Q4); defined the standard benchmark-suite specification for model-performance evaluation; established the four-agent architecture for Agentic AI maintenance workflows; first to systematize the dual role of AI as service platform (Role B) and AI as service object (Role A)
Robotics Service SLA Framework	No ISO / IEC standard defines availability targets, response-time tiers, SLA performance metrics, or breach-compensation mechanisms for robot service; service contracts rely mainly on non-quantified "reasonable efforts" clauses	Established an OEE-centered robot SLA-metrics system (separating outcome, process, and data metrics); designed a four-tier SLA product system (L1–L4); proposed eight SLA strategies for emerging robotics companies; defined the structural restructuring of the SLA system under the RaaS model

C.2 Three Historical Precedents and Their Relationship to RSF

RSF was not established in a vacuum; it follows the historical path of three industries that produced service-management frameworks after technology reached scale. Understanding these three precedents clarifies RSF's positioning and boundaries. They are design references for RSF, not substitutes — the particularities of robot service (statutory physical-safety requirements, AI uncertainty, cross-platform complexity) make none of the existing frameworks directly transferable.

Table C-1 Three Historical Precedents: Core Insights and Reference Value to RSF

Precedent	Historical Context	Core Contribution and Reference Value to RSF	Key Differences between RSF and the Precedent

ASE Certification + Automotive 4S	1970s–2000s: as automotive ECUs became widespread, traditional mechanics lacked electronic-diagnostic capability; customers could not judge the skill level of dealership technicians	◦ ASE established tiered certification across 8+ specialties, turning capability into verifiable qualification ◦ The closed-loop structure "certification → authorized scope of work → customer trust" is a convention RSF also follows ◦ The closed-loop logic of the 4S service blueprint is paralleled by RSF's eight-step field-service procedure	RSF must additionally handle probabilistic AI-layer diagnosis, heterogeneous cross-platform fleets, and AI-model version management — none of which has any parallel in automotive service
CMMI (Software Capability Maturity)	1990s: software projects were widely over schedule or failing, but process capability could not be assessed in advance; the industry lacked a common language for capability maturity	◦ Defined organizational process-capability levels 1–5, letting organizations know "where they stand" ◦ RSF follows the same tiering convention and will establish the Robot Service Maturity Model (RSF-MM Level 0–5, planned for release in v2.0) ◦ The logic "capability can be tiered → assessed → improved" holds equally for robot service	CMMI targets pure software processes with no physical-safety requirements; RSF must coexist with statutory baselines such as IEC 61508; robot-service "processes" include physical operations, so CMMI's metric system must be rebuilt from scratch
ITIL® / ITSM (External Reference Only)	1980s: IT service was chaotic; the UK government commissioned a service-management framework that later became the de facto global standard for IT service management	◦ The ITIL® development history from industry initiative → framework document → certification system → international influence is used only as an external development comparison for RSF ◦ RSF's three-tier closed loop — fault response → root-cause analysis → service improvement — answers a comparable operational need to IT service management's incident / problem / change practices; the RSF loop is defined independently, for physical systems ◦ RSF's lifecycle-management chapter adopts a full-lifecycle, value-centred view consistent with modern	ITIL® is primarily oriented to IT service environments (logical failures, often remotely recoverable); robot service targets physical systems (physical intervention may be required); common ITSM SLA examples often use RTO / RPO, while robot SLAs must be based on MTTR / OEE — the metric system must be rebuilt from scratch

		service-management thinking, while remaining robotics-specific and independent of official ITIL content	

C.3 RSF's Contemporary Value: Avoiding the IT Industry's Twenty-Year Wait

It took the IT industry roughly twenty years to establish ITIL® as a widely adopted service-management framework after technology had already been deployed at scale. During those twenty years, the IT service industry endured large amounts of avoidable disorder: no capability standard, no quality commitment, no incident procedure. In 2025 the robotics industry is going through the technology-scaling inflection point that the IT industry went through around 1985.

RSF's strategic significance is to move the establishment of a service-capability framework forward — from "10–20 years after technology deployment" to "contemporaneous with deployment." This is not only an efficiency issue but a safety issue: deploying robots at scale without a standardized service-capability framework would mean tens of thousands of engineers maintaining industrial equipment capable of real physical harm, without standard training, without capability assessment, and without safety procedures. RSF's emergence is the industry's systemic response to that risk.

Appendix D

Glossary of Terms

This glossary collects the core professional terms used in the RSF framework, grouped by capability domain and topic. Each entry includes: abbreviation (if any), English full name, and a concise definition. This glossary is the standard terminology reference for RSF certification examinations and ATC training courses.

D.1 Framework and Certification System

Abbreviation	English Term	Definition
RSF	Robotics Service Framework	The professional capability framework and certification system for industrial-robot service defined in this document
6D	Six Capability Domains	The six core capability modules (Domain 1 – Domain 6) of the RSF framework, covering both the service-execution layer and the management-architecture layer
RSF-MM	RSF Maturity Model	A Level 0–5 maturity assessment model for service organizations, similar in logic to CMMI; planned for release in v2.0
ATC	Authorized Training Center	An institution authorized by RSF to deliver certification training courses and administer examinations
CPD	Continuing Specialist Development	Credit-based continuing-learning system required to maintain RSF certification (30 credits every 3 years)
T3	Train-the-Trainer	Authorized training program for ATC trainers, qualifying them to deliver RSF certification courses
PSC	Platform Service Card	A standardized one-page reference document for platform-specific service requirements, used alongside the UCSP universal procedure
UCSP	Unified Cross-Platform Service Procedures	RSF's two-layer framework for cross-platform service management (common layer + platform-specific layer)

D.2 Domain 1 — Service Lifecycle Management

Abbreviation	English Term	Definition
—	Service Lifecycle	The full process of a robot system from delivery and installation to controlled decommissioning; RSF divides it into five phases
PM	Preventive Maintenance	Scheduled maintenance performed before failures occur, aimed at maintaining high equipment availability
LOTO	Lockout/Tagout	Safety procedure that physically isolates and labels all energy sources before maintenance work begins
TCP	Tool Center Point	Reference point of the robot end-effector and the core measurement baseline for accuracy verification
OEE	Overall Equipment Effectiveness	Availability × Performance × Quality — composite performance metric for manufacturing equipment (world-class benchmark ≥ 85%)
MTTR	Mean Time to Repair	Average time from work-order opening to equipment return to normal operation
MTBF	Mean Time Between Failures	Average time during which equipment operates normally between two failures
EOL	End of Life	OEM announces that a part is no longer produced; spare-parts strategy and equipment-retirement schedule must be reassessed
DAC	Delivery Acceptance Certificate	Gate document for Phase 2 → Phase 3 transition; must be co-signed by the service provider and the customer
KFL	Known Fault Library	Organizational knowledge asset that stores past successful diagnoses in structured form

SVP	Safety Verification Procedure	Four-step standardized safety pre-check that must be performed before any service operation on any platform (a component of UCSP)

D.3 Domain 2 — Fault Diagnosis and Technical Support

Abbreviation	English Term	Definition
L1	Mechanical Layer	First layer of the four-layer diagnostic architecture: joint mechanisms, reducers, bearings, end-effectors, and other physical mechanical components
L2	Electrical Control Layer	Second layer of the four-layer diagnostic architecture: servo drives, safety controllers, sensors, communication interfaces
L3	Software Middleware Layer	Third layer of the four-layer diagnostic architecture: controller firmware, ROS, communication protocols, coordinate-frame management
L4	AI Decision Layer	Fourth layer of the four-layer diagnostic architecture: perception models, planning networks, reinforcement-learning policies
FTFR	First-Time Fix Rate	Share of work orders whose root cause is correctly identified on the first visit (target ≥ 75 %)
MTDA	Mean Time to Diagnose (Accurately)	Average time from diagnosis start to root-cause confirmation
RCA	Root Cause Analysis	Systematic analytical method for identifying the root cause of a fault (rather than its surface symptoms), including 5-Why, Ishikawa, and FTA
MCSA	Motor Current Signature Analysis	Diagnostic method that analyzes the motor-current spectrum to detect winding and bearing faults
OSCE	Objective Structured Clinical Examination	Multi-station standardized scoring framework used in the RSF practical-skills assessment

D.4 Domain 3 — Remote Service and Digital Operations

Abbreviation	Chinese Term	English Term	Definition
RDT	**Robot Digital Twin**	Robot Digital Twin	Virtual digital model synchronized in real time with the physical robot system, supporting state assessment, simulation, and prediction
RUL	**Remaining Useful Life**	Remaining Useful Life	Predicted remaining normal-operation time of a critical component before its maintenance threshold is reached
HI	**Health Index**	Health Index	Composite equipment-state metric that fuses multi-dimensional sensor data into a single value in the 0–100 range
OTA	**Over-The-Air Update**	Over-The-Air	Technology for remotely deploying equipment firmware, control software, or AI models via network
CBM	**Condition-Based Maintenance**	Condition-Based Maintenance	Proactive maintenance strategy triggered when sensor data or diagnostic indicators exceed defined thresholds
PdM	**Predictive Maintenance**	Predictive Maintenance	Maintenance approach that uses machine-learning models to predict failure timing and intervene in advance
PsM	**Prescriptive Maintenance**	Prescriptive Maintenance	Extension of predictive maintenance that automatically prescribes the optimal maintenance action (when, where, and what to do)
LSTM	**Long Short-Term Memory network**	Long Short-Term Memory	Deep-learning model that excels at time-series data and is well suited to equipment degradation-trend prediction

PINN	**Physics-Informed Neural Network**	Physics-Informed Neural Network	AI model that embeds physical-equation constraints into a neural network — currently the state-of-the-art approach for RUL prediction
ZTA	**Zero-Trust Architecture**	Zero Trust Architecture	Cybersecurity architectural principle of "never trust, always verify," applicable to robot remote-access scenarios
DMZ	**Industrial DMZ**	Industrial Demilitarized Zone	Isolated buffer zone between the OT network and the IT / Internet network through which all remote access must be relayed
SOH	**State of Health**	State of Health	Battery's current full-charge capacity as a percentage of its rated capacity — a measure of battery aging
TSN	**Time-Sensitive Networking**	Time-Sensitive Networking	Industrial-Ethernet protocol providing deterministic low latency; the network infrastructure for real-time robot control
MFA	**Multi-Factor Authentication**	Multi-Factor Authentication	Authentication mechanism requiring two or more verification factors — a baseline requirement for RSF remote access
PSI	**Population Stability Index**	Population Stability Index	Statistical measure of data-distribution shift (> 0.25 marks the major-drift warning threshold)

D.5 Domain 4 — SLA Design and Service Commitment

Abbreviation	English Term	Definition
SLA	Service Level Agreement	Contract document that quantifies service-quality commitments, including metric definitions, target values, and breach-compensation clauses
KPI	Key Performance Indicator	Quantifiable indicators used to measure service performance (e.g., availability, MTTR, PM compliance rate)
RaaS	Robots as a Service	Subscription deployment model in which the service provider retains equipment ownership — the strongest incentive structure for SLA fulfillment
CSAT	Customer Satisfaction Score	Survey-based metric of customers' subjective evaluation of service quality (RSF target ≥ 7.5 / 10)
NPS	Net Promoter Score	Metric of customers' likelihood to recommend; used in ATC training-quality assessment (target ≥ 8.0 / 10)
P&L	Profit and Loss	Revenue and cost management of service operations; the core financial responsibility of Specialist-tier Service Managers
TCO	Total Cost of Ownership	Complete cost over the equipment's full lifecycle, including acquisition, maintenance, energy, and end-of-life disposal

D.6 Domain 5 — Cross-Platform Service Procedures

Abbreviation	English Term	Definition
AMR	Autonomous Mobile Robot	Mobile robot that uses SLAM or other navigation technology to autonomously plan paths
Cobot	Collaborative Robot	Robot designed for safe collaboration with humans in a shared workspace; governed by ISO/TS 15066
SLAM	Simultaneous Localization and Mapping	Algorithm by which a mobile robot simultaneously builds an environment map and localizes itself within it
SEA	Series Elastic Actuator	Joint design with an elastic element placed in series between motor and load, used for force control and collision compliance

VLA	Vision-Language-Action Model	Core AI architecture for humanoid robots fusing visual perception, natural-language understanding, and action generation
FMS	Fleet Management System	Software platform that uniformly dispatches and monitors multiple AMRs, typically using the VDA 5050 protocol for communication
STO	Safe Torque Off	Robot safety function that immediately removes all motor torque without waiting for deceleration
SS1	Safe Stop 1	Controlled stop: decelerate to halt, then activate STO (Category 1 stop)
SS2	Safe Stop 2	Controlled stop: decelerate to halt and maintain the stopped position (motor continues to apply holding torque)

D.7 Domain 6 — AI / ML System Service Methods

Abbreviation	English Term	Definition
MLOps	Machine Learning Operations	Methodology for full-lifecycle management of AI models, from training to deployment, monitoring, and updating
RAG	Retrieval-Augmented Generation	AI technique combining knowledge-base semantic search with large-language-model generation, used for intelligent diagnostic assistants
LLM	Large Language Model	General-purpose language AI model trained on large-scale text (e.g., GPT-4, Claude), applicable to service-knowledge assistance
mAP	Mean Average Precision	Composite performance-evaluation metric for object-detection models (typically expressed as mAP@0.5)
Q1–Q4	Four-Quadrant AI Fault Classification	RSF-defined AI-layer fault classification framework: Q1 distribution shift / Q2 model drift / Q3 policy drift / Q4 data-quality failure
AMDM	Adaptive Multi-Dimensional Monitoring	Agentic AI monitoring framework that reduces anomaly-detection latency to 5.6 seconds
KFL	Known Fault Library	Organizational asset that stores diagnostic experience in structured form; the core data source for RAG knowledge systems

D.8 General Technical Terms

Abbreviation	English Term	Definition
EOAT	End of Arm Tooling	Tool mounted at the wrist of a robot (gripper, suction cup, welding torch, vision system, etc.)
TCP	Tool Center Point	Reference point of the robot end-effector and the core measurement baseline for accuracy verification
EtherCAT	EtherCAT	Real-time industrial bus protocol based on Ethernet, widely used for servo-drive communication
PROFINET	PROFINET	Industrial-Ethernet communication protocol promoted by Siemens and other vendors
OPC UA	OPC Unified Architecture	Cross-platform data-exchange standard for the Industrial IoT, supporting robot-to-cloud integration
MQTT	Message Queuing Telemetry Transport	Lightweight IoT messaging protocol, suitable for low-bandwidth transmission of high-volume sensor data
ROS	Robot Operating System	Mainstream middleware framework for robot software development (ROS / ROS2), widely used in AMRs and research robots
FTA	Fault Tree Analysis	Top-down logic analysis method that traces all possible root causes of a top-level failure event
FMEA	Failure Mode and Effects Analysis	Preventive analysis tool that systematically identifies potential failure modes and their effects
PPE	Personal Protective Equipment	Safety protective equipment required during robot field service (safety shoes, safety glasses, gloves, etc.)
BOM	Bill of Materials	Complete list of all parts and materials required for a product or

		maintenance task
CMMS	Computerized Maintenance Management System	Software platform for managing work orders, PM schedules, spare-parts inventory, and service performance

ITIL® is a registered trademark of PeopleCert. References to ITIL are for identification, descriptive, and comparative purposes only. RSF is independent and is not affiliated with, endorsed by, sponsored by, certified by, or licensed by PeopleCert. RSF does not reproduce official ITIL materials. All third-party framework, certification, and standard names (including ASE, PMI, ISO/IEC) are used for identification only; no endorsement is implied.