Right Service, Right Cost: An Outcome-Aligned Governance Framework for Managed Services Contracts

A Cross-Sector Framework for Right-Sizing SOWs, SLAs and KPIs

Anonymised Practitioner Manuscript by **Yogesh Gaur**

Abstract

A significant proportion of organisational expenditure is committed through contracts, making it essential to understand the underlying cost drivers and ensure that contractual obligations remain aligned with business requirements. Managed services contracts can become costly and difficult to govern when Statements of Work (SOWs), Service Level Agreements (SLAs) and Key Performance Indicators (KPIs) are designed primarily for maximum contractual control rather than sustainable business and customer outcomes.

Although outsourcing governance, performance-based contracting, service-level management, reliability and performance measurement are well established, these elements are often applied separately in practice. This practitioner article addresses that integration gap by proposing a six-layer Outcome-Aligned Managed Services Governance Framework. The framework segments services before contract design, right-sizes scope, classifies SLAs by criticality and operational feasibility, combines Mean Time to Restore (MTTR) with Mean Time Between Failures (MTBF), aligns KPIs with customer impact, and validates premium commitments through cost-to-serve.

An anonymised telecom managed-services application illustrates how contractual ambition can be aligned with operational reality while preserving supplier accountability. Although the framework is designed for broad application across multiple sectors, the author also recognises that no single governance model can fit every organisational context. Minor adjustments may therefore be required to reflect sector-specific operating conditions, regulatory requirements, service criticality, risk exposure and customer expectations. With appropriate adaptation, the framework helps organisations replace uniform, penalty-focused contracting with differentiated service expectations, transparent accountability, improved reliability and commercially justified performance. It also supports more informed contract design by linking service commitments with

operational feasibility, customer impact and cost-to-serve, enabling organisations to pay for outcomes that create measurable value rather than for contractual provisions that merely appear stronger on paper.

Keywords

Managed Services (MS); Statement of Work (SOW); Service Level Agreement (SLA); Key Performance Indicator (KPI); Supplier Governance; Service Reliability; Performance Governance; Cost-to-Serve; Outcome-based Contracting

1. Introduction and Management Problem

Managed services is an operating model in which an organisation transfers defined operational activities to a specialised external partner while retaining ownership of business outcomes, customer experience and strategic control. The partner manages day-to-day delivery in accordance with an agreed SOW, SLAs and KPIs. However, the organisation remains accountable for ensuring that the arrangement delivers reliable, sustainable and commercially justified service value.

The model is widely used across telecommunications, information technology, enterprise infrastructure, utilities, banking operations and other service environments of varying scale and complexity. Its effectiveness depends less on the volume of contractual clauses than on whether scope, service levels, performance measures and commercial mechanisms are aligned with business requirements within one coherent operating model.

The modern telecom managed-services model began to take shape in the early 2000s, with one of the earliest documented network outsourcing arrangements, established between Telfort and Ericsson (Ericsson, 2011). The landmark Bharti Airtel–Ericsson MS agreement, signed in 2004, subsequently demonstrated the model at a significantly larger and more transformational scale (Bharti Enterprises, 2004). The original purpose of managed services was not merely to transfer workforce to a lower-cost supplier or create a penalty-driven contractual relationship. It was designed to capitalise on the service provider's specialist technical capability, specialised tools, economies of scale, operational expertise, and standardised delivery practices, while enabling the

operator to achieve rapid scalability, reliable network performance, a more scalable and predictable cost model, and greater strategic focus on customers, markets, and growth. The underlying principle was to assign operational execution to the party best equipped to perform it, while the operator retained control over strategy, investment priorities, service expectations and business outcomes.

As managed services arrangements have evolved, organisations have placed greater emphasis on cost discipline, standardised service expectations, and formal accountability mechanisms. The resulting management challenge is to ensure that these mechanisms do not become ends in themselves but remain aligned with the model's strategic purpose: accessing specialised capability, strengthening operational reliability, enabling scalability and creating long-term business value.

Many managed-services contracts are designed to maximise contractual protection and may therefore appear robust on paper. Over time, the SOW may accumulate activities introduced after incidents or escalations; uniform SLAs may be applied across services and locations with different operating realities, and ambitious KPIs may be agreed without sufficient analysis of baseline performance, root causes, restoration feasibility, alignment with business requirements, or impact on contract cost. The result can be a contract that appears strong but becomes more expensive, creates avoidable disputes and still fails to improve customer outcomes.

A uniform SLA illustrates the problem. A remotely managed digital platform, a distributed field asset, a restricted facility and a service dependent on multiple third parties may require different response and restoration logic. Likewise, a low MTTR may indicate rapid recovery, yet a low MTBF or equivalent recurrence measure may reveal repeated instability. High-level availability can also conceal capacity, resilience, transaction or service-quality degradation that customers continue to experience.

The management challenge is therefore not simply to demand more from the managed-services partner. It is to determine the appropriate scope, service levels, reliability expectations, and KPIs for each operating context, while ensuring that the associated costs are justified by customer value, business risk, and long-term sustainability (Cullen et al., 2005; Lacity et al., 2009).

2. Contribution of the Article

This article contributes to managed-services practice by treating service segmentation, SOW design, SLA classification, MTTR, MTBF, KPIs, and cost-to-serve as an integrated governance system rather than as separate contractual annexes. The proposed framework helps managers identify over-specification, distinguish controllable from dependency-based performance, link measures to customer impact, and test whether premium commitments are commercially justified, aligned with business requirements, and in line with long-term sustainability.

The article does not argue for weak contracts. It argues for intelligent contracts. A strong managed-services contract is not the one that includes the maximum scope, the shortest SLA and the most ambitious KPI targets. It is the one that defines the right operating expectation for the right service, in the right context, at the right cost and with clear accountability.

The framework is sector-neutral in design and application-specific in calibration. It synthesises anonymised professional experience in telecom operations, financial planning (Gaur, 2026), and managed-services governance, with established research on contracting, outsourcing, service levels, reliability, and performance measurement. The telecom application is used to illustrate the framework rather than to define its boundary. The article offers an actionable managerial model, not a statistically validated causal claim.

3. Conceptual Foundations and Practice Gap

Four closely related perspectives inform managed-services governance. Transaction-cost economics examines whether outsourcing and contractual controls reduce or increase the total cost of managing a service. Relational contracting explains how formal obligations and structured working relationships operate together throughout the contract period. Outsourcing governance clarifies how responsibilities, risks, dependencies, performance and decision rights are allocated between the organisation and the service provider. Performance-based contracting links payment and accountability to agreed outcomes rather than only to activities performed. Together, these perspectives connect cost discipline, contractual clarity, structured governance and outcome-based accountability.

Transaction-cost economics highlights that contractual protection is not cost-free. Expanding the SOW, shortening restoration targets, increasing the number of KPIs and strengthening penalties may reduce risk. However, they can also increase monitoring effort, resource requirements and supplier risk premiums. Where performance depends on access, data availability, utilities, client decisions, regulation, third parties or wider operating conditions, the managed services partner may price this uncertainty into the contract without delivering a proportional improvement in outcomes (Williamson, 1979).

Relational contracting and outsourcing governance emphasise that formal contracts should be supported by structured governance. The contract should clearly define what must be delivered and the consequences of non-performance, while governance determines how performance, dependencies, escalations, and changes are reviewed and managed. This does not reduce the partner's accountability for obligations within its control. Rather, contractual clarity and structured governance work together to support effective and sustainable service delivery (Poppo & Zenger, 2002; Cullen et al., 2005; Lacity et al., 2009).

Performance-based contracting and SLA governance shift attention from completed activities to achieved outcomes. Outcome-based contracting research further recognises that service results often depend on interconnected actors, resources and decisions, requiring a system-wide view rather than isolated supplier measures (Ng et al., 2009). Effective arrangements therefore require clearly defined outcomes, measurable indicators and a transparent distinction between controllable and dependency-based risks. An aggressive SLA may appear strong, but accountability becomes unclear when acknowledgement, diagnosis, remote action, dispatch, access, and restoration are combined into a single target, even though they involve different owners and constraints. SLAs should therefore be assessed by their relevance, measurability and operational feasibility, not by strictness alone (Goo et al., 2009; Selviaridis & Wynstra, 2015).

Reliability and performance-measurement research add two further requirements. MTTR indicates how quickly service is restored, while MTBF or an equivalent recurrence or stability measure shows how long performance remains reliable before the next failure. Considering both prevents

repeated incident closure from being mistaken for sustained reliability improvement (O'Connor & Kleyner, 2012). Similarly, KPIs should reflect customer and business impact rather than the convenience of reporting. High-level availability may remain within target while lower-level capacity, resilience, accessibility, transaction performance or service quality continues to deteriorate (Kaplan & Norton, 1992; Neely et al., 2002).

The practical gap lies in the way these concepts are applied. SOWs, SLAs, reliability measures, KPIs and cost-to-serve are often designed as separate contractual components, although each directly influences the others. This article addresses that integration gap through the Outcome-Aligned Managed Services Governance Framework, which combines service segmentation, SOW right-sizing, SLA classification, reliability governance, customer-impacting KPIs, and cost-to-serve validation within a single, connected operating model.

4. The Outcome-Aligned Managed Services Governance Framework

The Outcome-Aligned Managed Services Governance Framework converts contract design into six connected managerial decisions that can be applied across digital, infrastructure, field, regulated and hybrid service models. It begins with the business outcome and operating context rather than with a standard contractual template. Each layer uses defined evidence, produces a governance output and establishes the basis for the next decision.

The sequence matters. Scope cannot be right-sized until services and operating contexts are segmented. A fair SLA cannot be set until partner ownership and dependencies are clear. Restoration or recovery performance cannot be judged without separating incident stages from recurring reliability. KPIs cannot be meaningful unless they measure where customer impact occurs. Premium commitments should then be approved only after their cost-to-serve and expected value are tested.

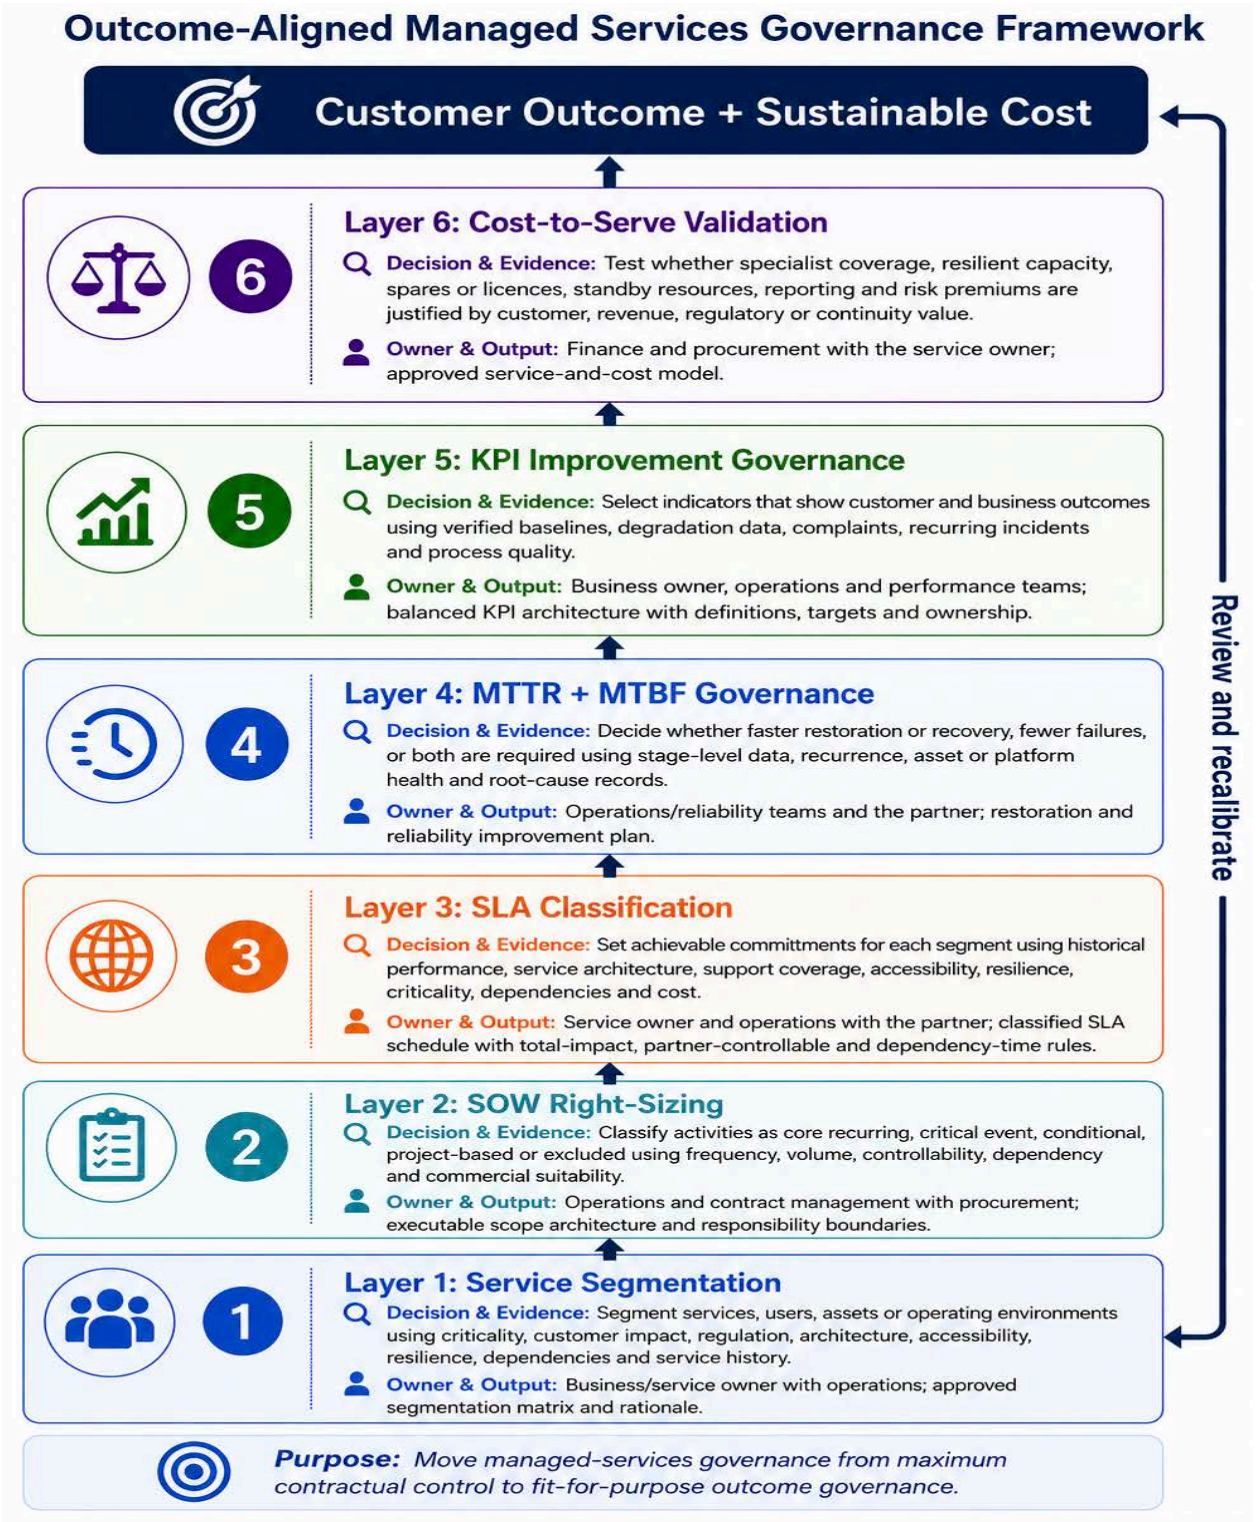

Figure 1. Outcome-Aligned Managed Services Governance Framework

Note. Ownership is cross-functional, but each decision requires a clearly accountable approver.

The framework is general in sequence but contextual in evidence. In cloud and IT services, segmentation may emphasise data criticality, cybersecurity, capacity, and resilience; in utilities and industrial maintenance, safety, physical access, and asset condition; and in banking or public services, transaction criticality, regulation, and continuity. The six decisions remain consistent, while the segmentation variables, SLA stages and KPIs are adapted to the service environment.

The framework is cyclical rather than static: segment, right-size, classify, restore and stabilise, measure and improve, validate value, and recalibrate. Performance results, customer impact, incident patterns and cost data should feed back into segmentation and contract review.

No layer should be optimised in isolation. Tightening SLAs without testing feasibility can raise cost and disputes; reducing scope without understanding risk can weaken service; improving MTTR without MTBF can reward repeated recovery; and increasing the number of KPIs can create reporting effort without better decisions. The framework balances customer value, operational feasibility, partner accountability and sustainable cost.

5. From Over-Specification to Periodic Recalibration

One hidden reason managed-services contracts become expensive is over-specification. Organisations often assume that stronger protection means broader scope, shorter SLAs, more KPIs and heavier penalties. This response is understandable after service failures, but it may increase cost without improving customer outcomes (Cullen et al., 2005; Williamson, 1979).

Over-specification typically appears in three forms. SOW over-coverage occurs when low-frequency, duplicated, conditional, or dependency-based activities are absorbed into the recurring fixed fee rather than assigned a commercial treatment appropriate to their nature. SLA over-commitment occurs when aggressive targets are applied across services and operating contexts with different delivery realities. KPI over-demand occurs when ambitious targets are set without a verified baseline, root-cause analysis, feasibility assessment or defined improvement path.

The commercial paradox is that the organisation believes it has negotiated stronger protection, while the managed services partner prices the additional complexity, standby resources and penalty

exposure into the contract. Selective precision is more effective: define clearly what matters, where premium service is justified, what the partner can control, and what costs are reasonable for the outcome.

Figure 2. From Contract Control to Outcome Governance

Over-specification can also emerge over time. Multi-year contracts remain fixed while technology, volumes, demand, operating models, process maturity, and organisational priorities change. The organisation may continue paying for activities, capacity or performance standards that no longer reflect current needs.

A formal annual recalibration, supplemented by review whenever a material change occurs, should examine activity volumes, utilisation, responsibilities, dependencies, service criticality, operating conditions, historical performance, KPI relevance and cost-to-serve. Recalibration is not a mechanism for weakening service. It keeps the operating model relevant and allows commitments to be strengthened when evidence and business value justify them.

Because penalties address non-performance rather than its underlying causes, periodic recalibration should remain focused on scope clarity, achievable commitments, repeat-fault reduction and verified improvement.

6. Service Segmentation Before Contract Design

A common mistake is to begin contract design directly with the SOW, SLA and KPI schedule. A more mature approach first segments the managed-services environment. Without segmentation, the same contractual expectations may be applied to services, users, assets, or operating conditions that differ in customer impact, delivery feasibility, and business value.

Segmentation should consider service criticality, customer impact, revenue exposure, regulatory sensitivity, demand pattern, architecture, resilience, accessibility, dependency, service maturity, incident history, MTTR, MTBF or equivalent recurrence measures, and cost-to-serve. Geography may be important in distributed field services; data classification, cyber risk and capacity may be more important in cloud services; and safety or transaction criticality may dominate in utilities, industrial or financial operations.

The objective is not to create excessive categories. Differentiation is justified only where variations materially affect service risk, feasibility or cost. The output should be a concise segmentation matrix approved by the business or service owner and validated by operations. The principle is simple: do not buy premium service everywhere; buy it where the business case justifies it.

7. SOW: From Over-Coverage to Outcome Clarity

The SOW defines responsibilities, deliverables, exclusions, dependencies and operating boundaries. A well-designed SOW provides clarity of execution; a poorly designed one creates hidden costs, duplicated or unclear responsibilities, and operational rigidity. Managed services may include monitoring, incident management, problem management, service desk, remote support, field support, preventive maintenance, capacity management, access administration, compliance reporting, optimisation, continuity and emergency response. The issue is not whether these activities are valid, but whether they are appropriately classified and priced.

A right-sized SOW separates core recurring work from critical event scope, conditional activity, project work and excluded or dependency-based responsibilities. This prevents the organisation from paying a permanent recurring cost for occasional work while ensuring that critical services cannot be rejected merely because a trigger or responsibility is unclear.

Before placing an activity in the fixed scope, managers should ask whether it directly supports customer experience, continuity, risk reduction or efficiency; whether it is recurring or conditional; whether historical data supports its frequency and volume; whether the partner controls the required inputs; and whether fixed-fee, unit-rate, event-based, project or change-control treatment is most appropriate.

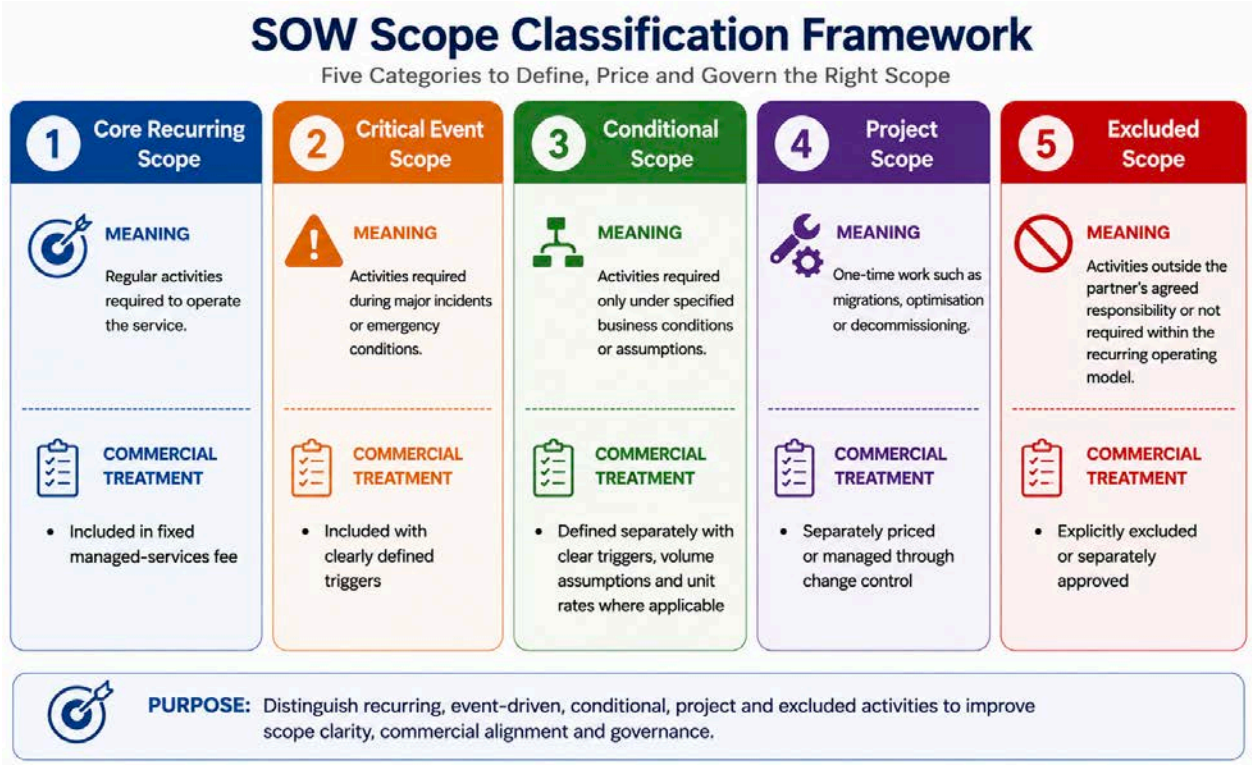

Figure 3. SOW Scope Classification Framework

The SOW should also identify client-owned and third-party dependencies. Clear dependencies do not reduce partner accountability; they prevent the partner from being held fully responsible for activities it cannot control and prevent internal blockers from remaining invisible. The purpose of

the SOW is not to store every operational concern. It is to define an executable, measurable and commercially aligned operating model.

8. SLA: From Uniform Service Levels to Classified Commitments

SLAs define service expectations, response and restoration targets, escalation rules, measurement logic and commercial consequences. They are also major cost drivers. A stricter SLA creates value only when it is linked to customer impact, operational feasibility, service criticality and business priority. Otherwise, it becomes an expensive promise (Goo et al., 2009).

A service incident may pass through acknowledgement, diagnosis, remote recovery, dispatch, access, temporary restoration, permanent repair and closure. Some stages can be completed quickly and remotely; others may depend on physical access, specialist resources, approvals or third parties. Combining the entire journey into a single short target may look strong in the contract, but it can blur accountability and materially increase the cost of service delivery.

A classified SLA should separate, where relevant, acknowledgement, diagnosis, remote action, dispatch, temporary restoration, permanent restoration and closure. The total customer-impact clock should continue until service is restored, because the customer and organisation remain affected regardless of the source of delay. For accountability purposes, partner-controllable time and dependency-based time should be measured separately, with clear ownership, evidence and escalation requirements. This preserves visibility into the complete outage while ensuring that the commercial consequences are fairly applied to the party responsible for the delay.

Uniform SLAs are attractive because they simplify administration and comparison. However, a single restoration commitment across remotely recoverable services, distributed field assets, restricted facilities and services dependent on external parties can create operational and commercial distortion. The partner may price the most difficult delivery conditions into every service, even when only a limited segment requires premium support.

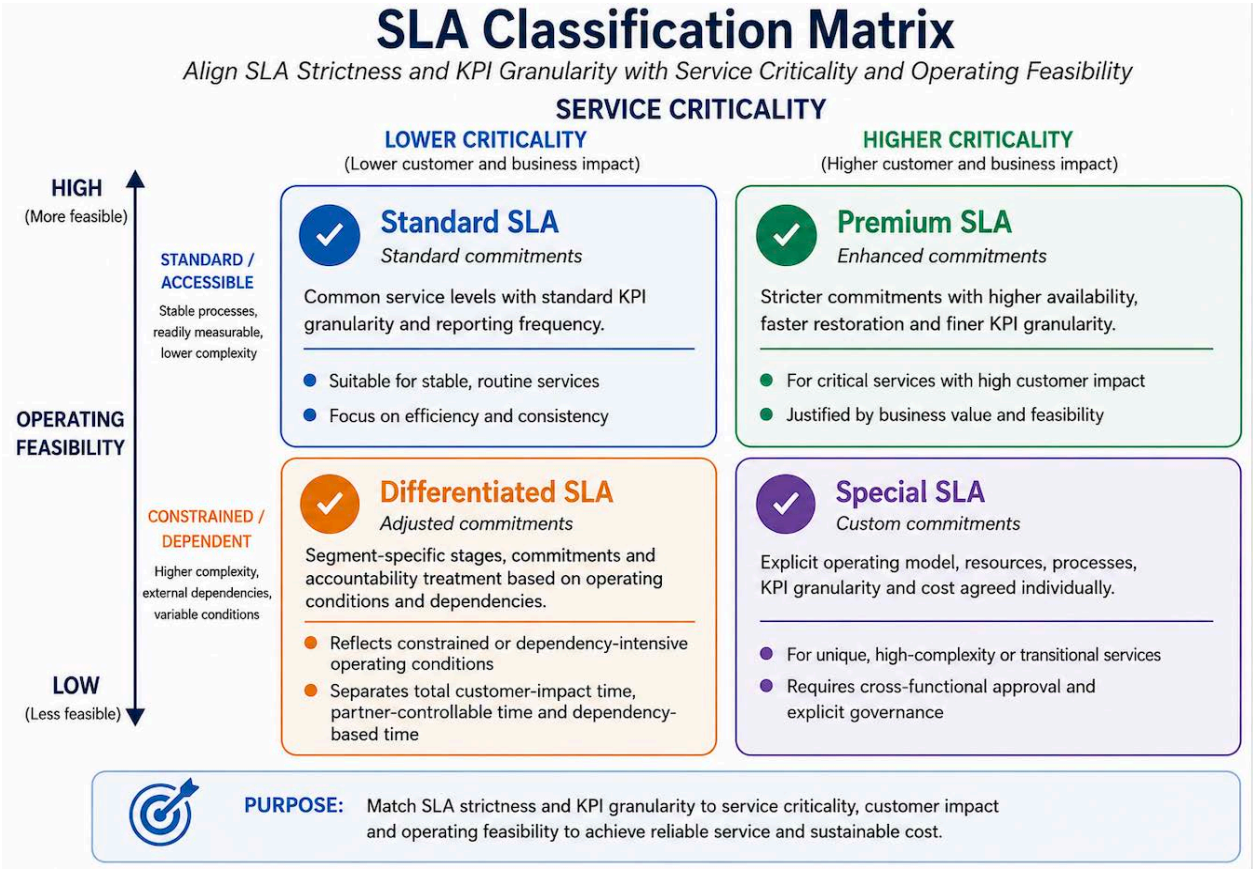

Figure 4. SLA Classification by Service Criticality and Operating Feasibility

Classification does not mean automatically relaxing service in constrained or dependency-intensive environments. A critical service may justify specialist coverage, resilient capacity, local resources, priority access or standby support. The requirement should, however, be explicitly approved based on customer impact, revenue exposure, regulatory importance, safety or continuity risk, and the delivery model should make the commitment achievable.

The resulting SLA is stronger because it is measurable and enforceable. It preserves high expectations where they matter, distinguishes partner-controlled performance from dependencies and makes the commercial premium visible.

9. MTTR and MTBF: Balancing Restoration and Reliability

Managed-services governance should examine not only how quickly services are restored but also how often they fail. MTTR answers, “How quickly did service recover?” MTBF answers, “How long did the service remain stable before the next failure?” Both are necessary.

A partner may repeatedly meet restoration SLAs, yet the same site, platform, or component fails again within a short period. Contract reporting may appear compliant, yet the customer continues to experience an unreliable service. In that case, the governance problem is not only about restoration speed; it is also about weak reliability (O’Connor & Kleyner, 2012).

Historical MTTR should be decomposed, where relevant, into detection, ticket creation, acknowledgement, diagnosis, remote recovery, dispatch, access, repair or replacement, external dependency and validation. Not every stage applies to every service, but the breakdown shows where time is consumed, which delays are controllable and where the organisation must resolve its own or third-party dependencies.

MTBF and recurrence analysis should then identify chronic services or assets, weak preventive maintenance, capacity or configuration instability, poor repair quality, environmental conditions and design weaknesses. A strong reliability model asks three connected questions: how quickly the incident was acknowledged, how quickly service was restored, and how long performance remained stable afterwards. This shifts governance from repeated closure to sustained reliability improvement.

10. KPI Hierarchy and Outcome-Based Performance

KPIs should answer a simple question: is the organisation receiving the outcome and value for which it pays? In many contracts, KPIs become symbols of ambition rather than management tools. The better question is not the highest target that can be demanded, but the performance problem being solved, the verified baseline and the realistic improvement path.

A balanced hierarchy separates strategic, operational, process and financial measures. This prevents high-level outcomes from being disconnected from the operational causes and commercial consequences that management must address (Kaplan & Norton, 1992; Neely et al., 2002).

KPI level	Purpose	Illustrative measures
Strategic	Measures business and customer outcomes	Availability, customer experience, revenue-impacting outage, business continuity
Operational	Measures service performance and reliability	MTTR, MTBF, repeat faults, SLA compliance, degradation duration
Process	Measures execution discipline and cause removal	Ticket ageing, preventive-maintenance quality, root-cause analysis and closure, action ageing
Financial	Measures cost, efficiency and commercial alignment	Cost per service, asset, transaction or incident; service credits; incentives; cost-to-serve

Table 1. KPI hierarchy and illustrative measures

KPI granularity should match the level at which customers or business operations experience impact. A service may appear available while transaction latency, processing capacity, resilience, queue time, response quality or a critical user journey is degraded. In field-based services, an asset may remain technically operational while delivering reduced capacity or quality. The answer is not to measure every component, but to identify the lowest meaningful level at which material degradation can be detected and acted upon. The telecom application in Section 13 illustrates this principle through transmission and radio-capacity indicators.

Targets should be improvement-oriented. A partner should not be rewarded merely for accepting an unrealistic KPI, nor should it benefit from a weak baseline. Rewards should be linked to verified, sustained and business-relevant improvement, including root-cause closure and repeat-fault reduction.

KPIs should also be progressive rather than static. Targets that are appropriate at the beginning of a managed-services arrangement may become too easy as processes mature, tools improve,

automation increases, recurring causes are eliminated, and the partner gains greater operational knowledge. Conversely, material changes in technology, demand, scope, or service risk may require reconsideration of KPI definitions and targets. Performance expectations should therefore follow an evidence-based improvement path, moving from an agreed baseline towards progressively stronger outcomes as capability matures. This progression should be established through transparent review and mutual validation rather than through arbitrary target increases unsupported by operational feasibility or business value.

11. Cost-to-Serve, Accountability and Incentives

Every SOW, SLA and KPI has a cost. A broad scope may require additional people, tools, licences, reporting effort, or specialist capability. A strict SLA may require resilient capacity, local spares, standby resources, 24/7 coverage, priority support, faster escalation or greater field presence. A complex KPI structure creates data, monitoring and governance effort. Premium commitments should therefore be approved only where customer impact, revenue exposure, regulation, safety, continuity or strategic importance justifies the additional cost (Lacity et al., 2009; Selviaridis & Wynstra, 2015).

Cost-to-serve analysis should test whether the same commitment is required across all segments; whether root-cause removal, automation, redundancy, or redesign could achieve the outcome more economically; and whether the service credit is meaningful relative to the premium paid. If an aggressive commitment attracts a high recurring premium but the consequence of missing it is small or capped, the economics may favour the partner even when service levels are missed.

Governance should distinguish partner-controllable failures from dependency-based failures, with clear, verifiable evidence. Controllable failures may include delayed acknowledgement, weak preventive maintenance, poor repair quality, inadequate field readiness and recurring incidents caused by unresolved root causes. Dependency-based failures may include denied access, utility interruptions, third-party outages, customer-side constraints and force majeure. Failures within the partner's control should remain subject to SLA obligations and commercial consequences. At the same time, dependency-related delays should be recorded separately with defined ownership,

evidence, escalation paths, resolution timelines, and corrective actions, rather than being treated as invisible exemptions.

Penalties and service credits reinforce minimum accountability, but they do not replace improvement. They are intended to ensure that the managed services partner treats agreed commitments with the required seriousness and should not be regarded as an opportunity for the organisation to recover money after service failure. The contract is signed to secure reliable delivery, not to generate penalty income. Although financial deductions may provide limited compensation, they cannot reverse customer impact, operational disruption or loss of confidence. Governance should therefore remain focused on service restoration, root-cause elimination and sustained improvement. A stronger commercial model combines consequences for controllable failure with incentives for verified performance beyond the agreed baseline. Incentives should reward outcomes such as sustained MTBF improvement, higher availability, consistent performance beyond agreed service levels, sustained KPI achievement, reduced customer impact and verified efficiency gains—not merely additional reporting or acceptance of more ambitious targets.

The aim is commercial alignment: the organisation should understand what it is paying for, the partner should understand what it controls, and both parties should be able to trace premium cost to a defined customer or business outcome.

12. Governance, Ownership and Decision Rights

The framework requires cross-functional input, but cross-functional participation should not blur accountability. Each decision needs an owner, supporting evidence, approval authority and review mechanism. Contract governance manages dependencies and change; it does not transfer the partner’s contractual obligations back to the client.

Role	Primary decision rights and responsibilities
Business or service owner	Approves service criticality, customer outcomes, risk tolerance and the rationale for premium commitments.
Operations and reliability	Validates feasibility, incident baselines, restoration or recovery stages, MTBF or recurrence, operating conditions, dependencies and required resources.
Finance	Tests affordability, cost-to-serve, expected value, penalty economics and the financial impact of differentiated service levels.
Procurement	Converts the approved operating model into fixed-fee, unit-rate, event-based, project, incentive or hybrid commercial treatment.
Legal and contract management	Defines responsibilities, exclusions, evidence, SLA clock rules, change control, dispute resolution and recalibration rights.
Managed services partner	Delivers obligations within its control, maintains readiness, provides transparent data and closes root-cause and improvement actions.
Joint governance committee	Approves material segmentation, scope, SLA, KPI, incentive and recalibration decisions requiring cross-functional authority.

Table 2. Governance roles and decision rights

This allocation prevents a common failure: one function demands a premium commitment, another funds it, a third manages delivery and the partner is held accountable for dependencies controlled elsewhere. Structured governance aligns the decision without diluting responsibility.

13. Illustrative Application: Telecom Managed Services

13.1. Context

The framework is intended for cross-sector use. Telecom is presented as an illustrative application because it combines distributed assets, service criticality, field access, external dependencies and measurable customer impact. The application is drawn from the author's professional experience in telecom operations, financial planning and managed-services governance. Organisational

names, contract values and commercially sensitive details are anonymised, and no confidential organisational data is used.

A large telecom operator in an African market had a managed-services arrangement covering network operations and field support across dense urban areas, semi-urban clusters, rural areas, remote locations, hilly terrain and service-critical sites. The distributed assets had materially different levels of accessibility, redundancy, customer impact and field-support availability.

Over successive contract cycles, the arrangement had expanded through renewals, operational escalations and additional requirements. Core recurring work, occasional activities, project work, client dependencies, and third-party responsibilities were not always clearly distinguished from one another. Many activities had effectively become part of the fixed fee even when their frequency, ownership or value differed.

13.2. Contractual Problem

The SLA model had also become increasingly ambitious. Stakeholders discussed restoration commitments as short as 30 minutes. The objective was understandable, but the target did not clearly distinguish between acknowledgement, remote diagnosis, dispatch, arrival at the site, and physical restoration.

In an urban location, acknowledgement, ticket creation, remote diagnosis, or traffic rerouting might be completed quickly. Physical restoration could still require field dispatch, site access, spares and repair. In rural, remote or hilly locations, travel time, road conditions, weather, restricted access and limited local resources created additional constraints. A single target across all stages and geographies, therefore, combined activities involving different owners and varying levels of control.

The partner may accept an aggressive commitment by incorporating additional workforce, standby coverage, local resources, logistics and risk contingency into the contract price. Even then, some incidents will still depend on access, utilities, travel time, permissions, or third parties. Where penalty exposure is capped, the organisation may pay a recurring premium for an SLA that is

physically difficult to deliver, while receiving only limited compensation when the target is missed. In such cases, the commercial structure may create an asymmetry in which the partner prices the risk into the recurring contract value. At the same time, its financial exposure for non-compliance remains limited or capped. Moreover, severe incidents may occur relatively infrequently because telecom networks are often designed with multiple layers of resilience and protection. Managers should therefore distinguish genuine service requirements from aspirational commitments and demand the level of service that is operationally necessary, commercially justified and proportionate to the actual risk.

13.3. Analysis and Framework Application

The practitioner experience illustrates how the arrangement could be reviewed as a single integrated operating model rather than as separate SOW, SLA, and KPI annexes. A sustainable contract review should begin by segmenting services and locations by geography, criticality, customer impact, accessibility, incident history, and restoration feasibility. This establishes whether materially different operating environments require differentiated scope, service levels or support models.

The next step is to classify SOW activities by nature and frequency. Core recurring and critical-event activities may require permanent delivery readiness. In contrast, conditional, project-based or dependency-based work may be better managed through defined triggers, unit pricing or change control. This distinction helps prevent occasional or externally dependent activities from being absorbed unnecessarily into the recurring fixed fee.

SLA commitments should then be separated into the operational stages that the partner can meaningfully control. Short targets may be appropriate for acknowledgement, ticket creation, diagnosis and remote action, while field arrival and physical restoration should reflect access conditions, geography, service criticality and relevant dependencies. Difficult operating conditions should not automatically result in weaker service commitments. Critical services may still justify premium support when the business rationale, resources, and costs are explicitly approved.

Historical MTTR should be decomposed into detection, acknowledgement, diagnosis, dispatch, travel, access, repair and validation. This analysis makes visible where restoration time is consumed and distinguishes partner-controllable delay from access, spares, client or third-party dependencies. MTBF and repeat-fault analysis should be considered alongside MTTR to avoid misinterpreting repeated restorations as sustained reliability. The governance question should therefore extend beyond whether an outage was restored within the target to whether the underlying cause was permanently resolved.

KPI design should also move below high-level availability, where customer impact requires greater granularity. In telecom networks, for example, a sector may remain technically available while one or more transceivers, carriers or transmission paths are unavailable. Although the service is not fully down, reduced capacity may cause congestion, call blocking, reduced accessibility, or degraded quality during busy periods. The objective is not to multiply KPIs, but to identify the lowest meaningful level at which material customer impact can be detected, assigned and addressed.

The final step should consolidate the segmentation, scope, SLA, reliability and KPI decisions into an approved service-and-cost model. Each premium requirement should be linked to a defined business rationale, an accountable owner, a delivery assumption, and an expected outcome. This ensures that the commercial proposal reflects the service required rather than accumulating additional resources and obligations without clear justification.

Applied in this sequence, the framework supports contract design that is evidence-based, operationally realistic and commercially justified. It helps organisations preserve clear partner accountability while avoiding commitments that appear strong contractually but add limited value to service reliability or customer outcomes.

13.4. Governance Decisions and Learning

The illustrative application shows how the basis for approving managed-services commitments can be strengthened without simply weakening contractual requirements. Fixed recurring scope

should be distinguished from conditional and project-based work; SLA stages and dependency treatment should be made explicit; geography and service criticality should inform restoration commitments; MTBF and repeat faults should complement MTTR; customer-impacting degradation should remain visible; and a clear cost-and-value rationale should support every premium commitment.

The central learning is that contractual strictness and operational effectiveness are not the same. A short SLA creates value when the partner controls the activity, and faster restoration materially protects the customer or business. The same target may increase cost and disputes when physical constraints or external dependencies make consistent delivery unrealistic.

A further learning is that managed-services governance should focus on the quality of contractual decisions rather than the apparent intensity of contractual clauses. Commitments should be approved only when supported by evidence of business criticality, partner controllability, operational feasibility and commercial value. These assumptions should also be reviewed as technology, service maturity, demand and risk evolve, ensuring that the contract remains demanding where value justifies it and adaptable where operating conditions change.

Taken together, the illustrative telecom application shows that effective managed-services governance depends on a connected sequence of decisions: segment the operating environment, classify scope, differentiate SLA commitments, govern restoration and recurring reliability together, align KPIs with customer impact, and validate premium requirements against cost and value. Although the operating variables will differ across sectors, the underlying governance logic remains transferable. The next section translates these lessons into a practical implementation roadmap that organisations can apply during new procurement, contract renewal, service transformation or periodic recalibration.

14. Implementation Roadmap

The framework can be applied during a new procurement, a contract renewal, a service transformation, or an annual recalibration. Implementation should be evidence-led and cross-functional, with clear outputs at each stage.

Step	Action	Required output
1	Establish the current scope, service volume, performance, incident, penalty, incentive and cost baseline. Reconcile inconsistent definitions and data sources.	Agreed current-state baseline and problem statement.
2	Segment services, users, assets or operating environments by criticality, customer impact, regulation, architecture, accessibility, resilience, dependency and service history.	Approved segmentation matrix and rationale.
3	Classify SOW activities by frequency, value, controllability, dependency and commercial treatment.	Revised scope architecture and responsibility model.
4	Redesign SLA stages; define the start and end of the total customer-impact clock; and separately measure partner-controllable and dependency-based time, supported by evidence and ownership.	Classified SLA schedule and measurement rules.
5	Decompose MTTR and analyse MTBF, repeat faults, preventive maintenance and asset or platform health.	Restoration and reliability improvement plans.
6	Redesign KPIs around customer and business impact; confirm definitions, baselines, data sources, owners and targets.	Balanced KPI architecture.
7	Validate specialist coverage, resilient capacity, spares or licences, standby resources, tools, logistics, reporting and supplier risk premium against expected value.	Service-and-cost business case.

Step	Action	Required output
8	Secure approval from the service owner, operations, finance, procurement, legal and relevant dependency owners.	Approved operating and commercial model.
9	Implement through contract amendment, renewal or new award, supported by transition, data-readiness and mobilisation plans.	Executed contract and transition plan.
10	Review outcomes annually and whenever a material change occurs in technology, demand, operating context, responsibility, risk or service maturity.	Recalibration decision and updated framework inputs.

Table 3. Ten-step implementation roadmap

Where baselines or measurement systems are immature, the contract may use a defined stabilisation period before full commercial consequences apply. That period should contain dated actions and evidence requirements rather than becoming an open-ended exemption.

Recalibration should be treated as part of good governance, not as proof that the original contract failed. Multi-year arrangements operate in changing environments; the governance model must retain the ability to adapt while preserving accountability.

15. Contract Review and Practitioner Checklist

The following checklist provides a concise test for contract design, renewal and annual recalibration. It should be used after the detailed analysis, not as a substitute for evidence.

Area	Key question	Governance check
Segmentation	Are materially different services or operating contexts being treated alike?	Differentiate by criticality, customer impact, operating feasibility, dependency and cost.
SOW	Is each fixed-fee activity recurring, valuable and controllable?	Classify scope and define ownership, dependency and commercial treatment.

Area	Key question	Governance check
SLA	Does the commitment reflect the operational stage and service segment?	Separate acknowledgement, diagnosis, dispatch, access and restoration where needed.
Feasibility	Can the SLA be delivered under actual operating and dependency conditions?	Test history, accessibility, support coverage, resilience, dependencies and external constraints.
MTTR	Where is restoration time consumed?	Decompose the restoration journey and assign controllable actions.
MTBF	Is service stability improving?	Track repeat faults, chronic services, platforms or assets, and root-cause closure.
KPI	Does the measure reflect customer or business impact?	Use the lowest meaningful level; remove measures that do not support decisions.
Cost	Is the premium worth paying?	Compare incremental cost with customer, revenue, regulatory and continuity value.
Incentives	Are consequences and rewards economically meaningful?	Penalise controllable failure and reward verified, sustained improvement.
Recalibration	Is the model still relevant?	Conduct annual review and additional review after material change.

Table 4. Practitioner contract-review checklist

16. Managerial and Broader Service Implications

For senior management, the central implication is to view managed-services contracts as operating models rather than procurement documents. Leaders should ask whether the organisation is buying

the highest service level or the right service level, and whether the additional commitment protects a defined customer, business or risk outcome.

For operations, incident history, restoration stages, MTBF, repeat faults, access, redundancy and customer-impact data should shape contractual requirements. For finance and procurement, broader scope and stricter targets are not automatic negotiation wins; they may increase the base fee and supplier risk premium. For legal and contract-management teams, accountability, evidence requirements, dependency treatment, SLA clock rules and recalibration mechanisms should be defined with sufficient precision to support effective execution.

For managed services partners, value does not come from accepting every clause. It comes from delivering controllable obligations, maintaining transparency, eliminating recurring causes and proposing improvements that can be verified. Clear dependency treatment also provides fairer governance without weakening partner accountability.

A well-structured contract framework does more than define obligations and enforce accountability. It creates a shared understanding of expected outcomes, decision rights, performance standards, dependencies and commercial consequences. Such clarity strengthens trust among business owners, operations, finance, procurement, legal teams and managed services partners by ensuring that each party understands what must be delivered, what evidence will be used, what remains within its control and how changes will be governed. This trust is built through transparent information and fair accountability, but it is ultimately validated through the partner's consistent delivery and sustained performance.

By linking scope, service levels, reliability expectations, KPIs and cost-to-serve to actual business requirements, the framework also supports business sustainability. Organisations can avoid paying for excessive commitments that add limited value while ensuring that critical services receive the protection they genuinely require. Periodic recalibration keeps the contract relevant as technology, demand, risk and operating conditions evolve, positioning contract governance as an operating foundation for reliable delivery, stakeholder confidence, responsible cost management and long-term organisational value.

The broader service implication is important. Managed services support digital platforms, enterprise operations, telecom networks, utilities, banking systems, industrial assets and public services that affect daily economic and social activity. Poorly designed contractual measures can encourage reporting compliance while customers or users continue to experience unreliable or degraded service. Outcome-aligned governance connects contractual commitments to actual service impact, supporting reliability, continuity, access and public trust.

17. Generalisability, Limitations and Research Agenda

The Outcome-Aligned Managed Services Governance Framework is sector-neutral in design but practice-derived. It is supported by established literature and structured practitioner reasoning, yet it has not been tested through a large-scale survey, a controlled comparison, or a multi-organisation longitudinal study. It should therefore be interpreted as an integrative managerial framework rather than as a proven causal model.

Its generality lies in the sequence of decisions: segment, right-size, classify, govern reliability, measure outcomes, and validate value, not in identical metrics or contract clauses. Each sector must calibrate the relevant variables, thresholds, measures and evidence requirements. Cloud services may emphasise cybersecurity, data residency and capacity; utilities may prioritise safety, access and asset condition; banking may focus on transaction continuity and regulatory exposure; and industrial services may require maintenance, spares and operational-risk measures.

The telecom application is illustrative rather than restrictive. It makes the framework visible because geography, access, redundancy, customer impact and third-party dependency vary materially across the service environment. Other sectors may reveal the same governance tension through different variables, such as platform resilience, transaction integrity, safety, demand volatility or vendor concentration.

The framework does not support automatic SLA relaxation in difficult or complex environments. A demanding commitment may be fully justified where failure would create serious consequences for customers, safety, regulatory requirements, revenue, or continuity. The framework requires that

such commitments be explicitly justified, operationally enabled and transparently priced. It also cautions against excessive segmentation, which can create its own administrative and reporting costs.

In strictly regulated environments, cost-to-serve validation may require a different decision logic. Certain expenditures—such as mandated resilience, cybersecurity controls, data residency, safety measures, auditability, continuity capability or minimum service obligations—may be non-discretionary even when their immediate financial return is difficult to quantify. In such contexts, validation should first establish the minimum cost of regulatory compliance and risk protection and then assess whether any additional premium commitment creates incremental customer, operational or commercial value. The framework should therefore distinguish mandatory compliance costs from discretionary service enhancements, while continuing to test whether each requirement is delivered through the most efficient and operationally sustainable model.

Future research could test the framework across sectors, compare uniform and segmented SLA models, evaluate the relationship between MTTR–MTBF or recurrence governance and customer outcomes, and assess whether cost-to-serve validation reduces supplier risk premiums, disputes or service-cost escalation. Comparative applications would also help determine which segmentation variables and KPI architectures are most effective in different service environments. Comparative research in highly regulated sectors could further examine how mandatory compliance costs, discretionary service premiums and risk-reduction value should be distinguished within cost-to-serve validation.

18. Conclusion

Managed-services contracts are often judged by the apparent strength of their SOW, SLA and KPI clauses. Stronger wording, however, does not automatically create stronger performance. Excessive scope, uniformly aggressive SLAs, and unrealistic KPIs can increase costs, obscure dependencies, and create disputes without improving customer outcomes.

This article proposes a shift from contract control to outcome-aligned governance. Organisations should segment services before contracting, right-size the SOW, classify SLA stages and service commitments, govern MTTR alongside MTBF, select customer-impacting KPIs, and validate premium obligations through cost-to-serve. These decisions should be supported by clear ownership, transparent evidence and annual recalibration. **A well-structured contract framework, therefore, does more than enforce accountability; it strengthens stakeholder trust and business sustainability by aligning service commitments, operational realities, and commercial value within a single transparent governance model.**

The framework changes the management question from, "How much can we demand from the partner?" to, "What outcome do we need, what can the partner control, what is the outcome worth, and how should it be governed?" Managed services should not become a high-cost penalty game. They should operate as a disciplined model for reliable service, transparent accountability, continuous improvement and long-term organisational value.

19. Research Integrity and Data Statement

This is a conceptual practitioner article supported by an anonymised application. No confidential organisational data, personal data, patient information or identifiable participant information is used. The article does not report research involving human participants; ethical approval and participant consent are therefore not applicable.

20. References

Bharti Enterprises. (2004). *Forging global partnerships*. https://www.bharti.com/PDFs/English-Bharti-Today/2004/Bharti-Today_Vol_8_Issue_2-2004_Forging-Global-Partnerships.pdf

Cullen, S., Seddon, P. B., & Willcocks, L. P. (2005). IT outsourcing configuration: Research into defining and designing outsourcing arrangements. *The Journal of Strategic Information Systems*, *14*(4), 357–387. https://doi.org/10.1016/j.jsis.2005.07.001

Ericsson. (2011). *Ericsson global services: Strengthening operator competitiveness*. https://cdn-netpr.pl/file/attachment/282012/ec/bugs_press_backgrounder_q4_2010.pdf

Gaur, Y. (2026). Strategic budgeting in the digital era: A governance and agility framework. *International Journal of Business and Management*, *21*(4), 174. https://doi.org/10.5539/ijbm.v21n4p174

Goo, J., Kishore, R., Rao, H. R., & Nam, K. (2009). The role of service level agreements in relational management of information technology outsourcing: An empirical study. *MIS Quarterly*, *33*(1), 119–145. https://doi.org/10.2307/20650281

Kaplan, R. S., & Norton, D. P. (1992). The balanced scorecard—Measures that drive performance. *Harvard Business Review*, *70*(1), 71–79. https://hbr.org/1992/01/the-balanced-scorecard-measures-that-drive-performance-2

Lacity, M. C., Khan, S. A., & Willcocks, L. P. (2009). A review of the IT outsourcing literature: Insights for practice. *The Journal of Strategic Information Systems*, *18*(3), 130–146. https://doi.org/10.1016/j.jsis.2009.06.002

Neely, A., Adams, C., & Kennerley, M. (2002). *The performance prism: The scorecard for measuring and managing business success*. Financial Times Prentice Hall.

Ng, I. C. L., Maull, R., & Yip, N. (2009). Outcome-based contracts as a driver for systems thinking and service-dominant logic in service science: Evidence from the defence industry. *European Management Journal*, *27*(6), 377–387. https://doi.org/10.1016/j.emj.2009.05.002

O'Connor, P. D. T., & Kleyner, A. (2012). *Practical reliability engineering* (5th ed.). John Wiley & Sons. https://doi.org/10.1002/9781119961260

Poppo, L., & Zenger, T. R. (2002). Do formal contracts and relational governance function as substitutes or complements? *Strategic Management Journal*, *23*(8), 707–725. https://doi.org/10.1002/smj.249

Selviaridis, K., & Wynstra, F. (2015). Performance-based contracting: A literature review and future research directions. *International Journal of Production Research*, *53*(12), 3505–3540. https://doi.org/10.1080/00207543.2014.978031

Williamson, O. E. (1979). Transaction-cost economics: The governance of contractual relations. *The Journal of Law and Economics*, *22*(2), 233–261. https://doi.org/10.1086/466942