

MK Scoring Framework: Quantitative System Reliability Assessment

A Practical Methodology for SRE Implementation

Author: Manoj Kuppan

Document Type: Technical White Paper

1. Purpose

Modern distributed systems generate unprecedented volumes of telemetry data while operating in increasingly complex architectures. Organizations struggle to translate this data into actionable reliability improvements due to:

- **Fragmented insights:** Technical metrics isolated from business impact
- **Reactive posture:** Static threshold alerting generating false positives while missing subtle degradations
- **Inconsistent prioritization:** Lack of objective frameworks for prioritizing reliability work across diverse applications
- **Business misalignment:** Technical metrics failing to correlate with user experience and business outcomes

The MK Scoring Framework addresses these challenges by providing a quantitative, business-aligned methodology for assessing system reliability and prioritizing remediation efforts. The framework enables organizations to:

1. Systematically translate business goals into measurable reliability parameters
2. Calculate objective reliability scores enabling cross-application comparison
3. Automatically prioritize remediation based on weighted business impact
4. Track reliability improvements quantitatively over time

2. System Reliability Dimensions and Adoption Challenges

2.1 Six Dimensions of System Reliability

The framework evaluates systems across six interconnected dimensions that collectively define reliability:

Dimension 1: Observability - The ability to understand system internal state through external outputs including metrics, logs, traces, and events. This dimension encompasses monitoring infrastructure, Service Level Objectives (SLOs) and Indicators (SLIs), structured logging, and business metric instrumentation.


Dimension 2: Optimal Alerting - The effectiveness of incident notification systems in providing timely, relevant, and actionable alerts while minimizing noise and alert fatigue.

Dimension 3: Release Strategy - The robustness of deployment processes that minimize risk through comprehensive testing, gradual rollouts, and automated recovery mechanisms.

Dimension 4: Incident Management - The effectiveness of processes for detecting, responding to, and recovering from system incidents, including post-incident learning and improvement.

Dimension 5: Resiliency - The system's capacity to maintain acceptable performance despite failures through architectural patterns including circuit breakers, timeouts, retries, bulkheads, and graceful degradation.

Dimension 6: Risk Assessment - Systematic identification and mitigation of failure modes using Failure Mode and Effects Analysis (FMEA) methodology to quantify risks across critical customer journeys.



2.2 Common SRE Adoption Challenges

Organizations attempting to implement SRE practices face systematic barriers that prevent effective reliability improvement. Without quantitative assessment, teams operate on subjective

evaluations like "our monitoring is pretty good" rather than measurable baselines. This lack of objective measurement prevents tracking improvement over time and makes data-driven decisions impossible.

The prioritization challenge compounds this problem. With hundreds of potential reliability improvements across monitoring gaps, alerting enhancements, deployment safety, and resiliency patterns, teams lack frameworks for determining what to fix first. Technical preferences and whoever advocates most loudly often drive prioritization rather than actual business impact.

Inconsistent implementation creates another barrier. Different teams interpret SRE principles differently - one team's "comprehensive observability" means basic metrics while another implements distributed tracing and custom business instrumentation. This leads to incomparable reliability postures across application portfolios, preventing central teams from assessing organizational reliability systematically.

The missing business alignment represents the most critical gap. Technical reliability metrics like error rates and response times don't translate to business outcomes, preventing executive buy-in and appropriate resource allocation for reliability work. When incidents occur, translating technical impact to business terms requires manual analysis taking hours or days.

2.3 Hierarchical Observability Architecture

The MK Scoring Framework addresses these challenges through three-level hierarchical assessment, with each level serving a specific purpose in the overall reliability picture:

Level 1 - Individual Metrics: Statistical anomaly detection using z-scores and percentile analysis identifies deviations in specific metrics such as response time, error rate, and resource saturation. This level answers "what is abnormal right now?" through automated threshold calculations that adapt to actual system behavior rather than static limits.

Level 2 - Customer Journey Risk Assessment: Failure Mode and Effects Analysis (FMEA) methodology calculates Risk Priority Numbers ($RPN = \text{Severity} \times \text{Occurrence} \times \text{Detection}$) for critical user flows, identifying high-risk failure modes that warrant dedicated monitoring and mitigation. This level answers, "what failures would most impact customers?" and directly informs which parameters require monitoring in Level 3.

Level 3 - System-Wide Reliability Scoring: The MK Scoring Framework aggregates weighted parameters across six dimensions, providing a normalized percentage score that enables objective comparison, tracks improvement over time, and automatically prioritizes remediation based on business impact. This level answers "how reliable is this system overall, and what should we improve first?"




Figure 2.1. Observability Hierarchy leading into MK Scoring for system wide reliability

This hierarchical approach ensures that lower-level technical findings connect to higher-level business context, enabling teams to trace from a business impact (Level 3 low score) through specific customer journey risks (Level 2 high RPN) down to individual technical metrics requiring attention (Level 1 anomalies).

3. MK Scoring Framework Overview

Having established the challenges organizations face in implementing SRE practices and the hierarchical observability architecture required to address them, we now examine how the MK Scoring Framework transforms these concepts into a quantitative assessment methodology. The framework provides a systematic approach to quantifying system reliability by translating qualitative reliability assessments into measurable, comparable scores.

Traditional approaches to reliability assessment rely on subjective evaluations and inconsistent measurement practices. Teams might describe their systems as "pretty reliable" or "mostly monitored," but these qualitative assessments prevent objective comparison across applications and don't enable tracking improvement over time.

The framework addresses this challenge through three interconnected methodologies operating at different levels of system abstraction. At the lowest level, statistical anomaly detection identifies deviations in individual metrics. At the mid level, Failure Mode and Effects Analysis quantifies risks across customer journeys. At the highest level, weighted scoring across six reliability dimensions provides an overall system reliability assessment. This hierarchical approach ensures that technical findings connect to business context, enabling teams to prioritize work based on actual impact rather than technical preference.

3.1 Core Methodology

The MK Scoring Framework provides quantitative reliability assessment through weighted, normalized scoring:

$$\text{MK Score} = (\sum(w_i \times A_i) / \sum(w_i \times E_i)) \times 100\%$$

Where:

- **w_i** = weight assigned to parameter i (reflects business criticality, typically 1-5)
- **A_i** = actual measured value of parameter i (1-5 scale)
- **E_i** = expected optimal value (always 5)

Score Interpretation:

- **Above 85%:** Highly reliable system meeting optimal standards
- **70-85%:** Acceptable system requiring targeted improvements
- **Below 70%:** System requiring critical remediation attention

3.2 Score-Based Prioritization

The framework automatically identifies highest-impact improvement opportunities:

$$\text{Gap Score} = (w_i \times E_i) - (w_i \times A_i)$$

Parameters with highest gap scores receive priority attention, ensuring engineering efforts focus on business-critical improvements.

Example:

- Payment MTTD: weight=5, actual=2 → Gap = (5×5) - (5×2) = 15 points
- Logging Coverage: weight=2, actual=3 → Gap = (2×5) - (2×3) = 4 points

Payment MTTD receives priority despite both having room for improvement, because its business weight (5 vs 2) reflects higher impact.

4. Parameter Derivation Using GQM Method

While the MK Score formula provides the mathematical foundation for reliability assessment, the critical question remains: how do organizations determine which parameters to measure and what weights to assign? The Goal-Question-Metric (GQM) methodology answers this question by providing a systematic approach that maintains mathematical rigor while enabling context-specific customization.

4.1 Goal-Question-Metric Methodology

The framework uses GQM theory to derive application-specific parameters while maintaining mathematical rigor:

Goal Level: Define strategic reliability objectives

Question Level: Decompose goals into answerable questions

Metric Level: Map questions to measurable parameters with appropriate weights

4.2 Parameter Derivation Matrix

Organizations derive parameters systematically from established SRE principles by mapping each principle to specific questions for development teams, which then translate into measurable parameters with appropriate business weights. The GQM methodology ensures that every parameter traces back to a strategic goal, preventing metric collection for its own sake while maintaining mathematical rigor in the scoring framework.

4.3 GQM Application Example: E-commerce Checkout

Goal: "Maintain 99.9% checkout completion rate during peak shopping periods"

Questions:

- What is current payment gateway response time?

Metrics derived: latency with high weightage (5)

- How frequently do cart sessions timeout?

Metrics derived: Error rate with high weightage (5)

- What percentage of checkouts complete successfully?

Metrics derived: Volume and error rate with high weightage (5)

- Can we detect payment failures within 5 mins?

Metrics derived: alert exists with less than or equal to 5 min interval with high weightage (5)

5. Framework Methodology

With the GQM method establishing how to derive appropriate parameters and weights, we now examine the six reliability dimensions that form the framework's assessment structure. Each dimension contributes specific measurable parameters that, when weighted and aggregated, produce the overall MK Score while enabling granular identification of improvement opportunities.

5.1 Six Dimensions - Detailed Parameters

The MK Scoring Framework evaluates system reliability across six interconnected dimensions, each contributing specific aspects to overall reliability assessment. Understanding these dimensions and their measurable parameters enables teams to systematically identify gaps and prioritize improvements based on business impact.

Dimension 1: Observability

Observability provides the foundation for all reliability work by enabling teams to understand system behavior through external signals. Without comprehensive observability, teams operate blind to degradations until customers report issues. The dimension encompasses not just metrics collection, but the entire infrastructure enabling rapid problem diagnosis and proactive issue detection.

Core Parameters:

- **Golden Signals Monitoring:** Latency (response time tracking with percentile analysis), Traffic (request volume and pattern analysis), Errors (error rate monitoring by type and severity), Saturation (resource utilization relative to limits)
- **SLO/SLI Framework:** Service Level Objectives defined for critical flows, Service Level Indicators monitored continuously, SLO compliance dashboards, Error budget tracking and consumption alerts
- **Logging Infrastructure:** Structured logging with correlation IDs enabling end-to-end transaction tracing, Log aggregation and centralized searchability, Log retention policies aligned with compliance requirements
- **Distributed Tracing:** Trace coverage across microservice boundaries, Trace sampling rates balancing cost and visibility, Trace-to-log correlation for deep debugging
- **Infrastructure Monitoring:** Database performance metrics (query latency, connection pool saturation), Message queue health (lag, throughput, dead letter queues), Custom business metrics instrumentation (checkout completion rate, payment success rate)

SLO Integration Framework

Service Level Objectives provide the critical link between technical observability and business expectations. Well-defined SLOs enable teams to make data-driven decisions about when to prioritize reliability work versus feature development through error budget management.

Table: SLO Integration Assessment

SLO Example	SLI (Measurement)	Derived Observability Parameters	Weight	Scoring Criteria
99.9% checkout success rate	Successful checkouts / Total attempts	Checkout Success Rate Monitoring,	5	5: SLO met with >20% error budget remaining 3: SLO

		Error Budget Tracking Dashboard		met with <10% error budget 1: SLO violated
95% of checkouts complete <2sec	P95 checkout latency	Checkout Latency P95 Monitoring, Latency SLO Alert Configuration	5	5: P95 <1.8sec (10% headroom) 3: P95 1.8-2.0sec 1: P95 >2.0sec
Payment gateway 99.95% available	Payment gateway uptime	Gateway Availability Monitoring, SLO Compliance Dashboard, Downtime Alert	5	5: >99.95% with multi-region failover 3: >99.95% single region 1: <99.95%

Organizations implementing comprehensive SLO frameworks achieve several benefits: clear definition of acceptable service quality levels enables informed trade-off decisions between velocity and reliability; error budgets provide mathematical justification for when to focus on stability versus new features; standardized SLI measurement across services enables portfolio-wide reliability comparison; and automated SLO violation alerts ensure teams detect degradations before error budgets exhaust.

Effective observability combined with SLO management enables sub-15-minute MTTR through automated anomaly detection and provides the data foundation for all other reliability dimensions. Teams should aim for SLO coverage on all customer-facing services and critical internal dependencies.

Dimension 2: Optimal Alerting

Alerting translates observability data into actionable notifications, but excessive alerts create fatigue while insufficient alerts miss critical issues. This dimension focuses on alert quality over quantity.

Core Parameters:

- **Alert Quality Metrics:** Precision assessment (percentage of alerts requiring action vs. false positives), Recall validation (percentage of real issues generating alerts vs. missed detections), Alert-to-incident ratio tracking
- **Alert Management:** Mean time to acknowledgment measurement, Alert noise evaluation and reduction initiatives, Runbook availability for common alerts, Alert escalation policy completeness
- **On-Call Health:** On-call rotation coverage and fairness, Engineer satisfaction with alert quality, Incident handoff documentation quality

Organizations achieving optimal alerting maintain precision above 85% and recall above 90%, ensuring engineers respond quickly to genuine issues without experiencing alert fatigue from false positives.

Dimension 3: Release Strategy

Release strategy determines how safely organizations can deploy changes, directly impacting both innovation velocity and system stability. Poor release strategies force trade-offs between speed and safety.

Core Parameters:

- **Deployment Safety:** Change failure rate (percentage of deployments causing incidents), Automated rollback capability and testing, Deployment frequency as measure of process maturity
- **Progressive Delivery:** Blue/green deployment implementation for zero-downtime releases, Canary deployment with statistical validation, Feature flag usage for progressive rollouts and rapid rollback
- **Testing Coverage:** Unit test coverage (target >80% of business logic), Integration test coverage (target >70% of API contracts), End-to-end test coverage (target >60% of critical user journeys), Pre-production environment parity with production

Organizations with mature release strategies achieve change failure rates below 15% while maintaining high deployment frequency, enabling rapid innovation without compromising stability.

Dimension 4: Incident Management

Incident management effectiveness determines how quickly teams detect and resolve issues, directly impacting customer experience during reliability events. Mature incident management turns potential disasters into minor disruptions.

Core Parameters:

- **Detection Speed:** Time to Detect (TTD) from incident start to first alert or discovery, Automated detection coverage (percentage of incidents detected by systems vs. customers)
- **Response Effectiveness:** Time to Mitigate (TTM) from detection to service restoration, Time to Resolve (TTR) from detection to root cause remediation, Incident severity classification accuracy
- **Learning and Improvement:** Post-incident review completion rate (target 100% for high-severity incidents), Incident runbook coverage, Action item completion tracking from retrospectives

High-performing teams achieve TTD under 15 minutes, TTM under 1 hour, and TTR under 4 hours for critical incidents through combination of automated detection, clear runbooks, and practiced response procedures.

Dimension 5: Resiliency

Resiliency determines whether systems gracefully handle failures or cascade into complete outages. This dimension focuses on architectural patterns and practices that maintain service quality despite component failures. Unlike simple availability measurement (which tells you the system was down), resiliency assessment evaluates the system's ability to prevent, detect, isolate, and recover from failures automatically.

Core Parameters:

- **Availability Measurement:** System uptime tracking against targets (typically 99.9%+ for critical services), Availability by customer segment for prioritized support
- **Resiliency Patterns Implementation:** Circuit breaker implementation with state monitoring and automatic recovery, Retry policies with exponential backoff and jitter to prevent thundering herds, Timeout configuration and enforcement across all external dependencies, Bulkhead isolation separating critical from non-critical resources, Graceful degradation modes with automated mode switching
- **Fault Tolerance Validation:** Chaos engineering adoption with regular game days, Disaster recovery procedures documented and tested quarterly, Multi-region or multi-availability-zone deployment for geographic redundancy

Resiliency Patterns Assessment Framework

Resiliency patterns represent proven architectural solutions to common failure scenarios. The table below provides a systematic assessment approach for each pattern.

Table: Resiliency Patterns Evaluation

Pattern	Question for Dev Team	Optimal Implementation (Score: 5)	Partial Implementation (Score: 3)	Missing/Inadequate (Score: 1)
Circuit Breaker	Is circuit breaker implemented for external dependencies? How is state monitored?	All external dependencies protected, state visible in dashboards, auto-recovery tested monthly, thresholds tuned to dependency SLOs	Critical dependencies only, manual state checking, quarterly recovery tests	Not implemented or no monitoring
Timeout	Are timeouts configured for all network calls? How are violations tracked?	All network calls have timeouts, values documented, violations monitored and alerted, timeout	Most calls have timeouts, some documentation, occasional monitoring	Inconsistent timeouts or no monitoring

		analysis in retrospectives		
Retry with Backoff	Are retry policies implemented? How is exhaustion handled?	Exponential backoff with jitter, exhaustion tracked and alerted, policies tested in chaos experiments	Fixed retry intervals, basic exhaustion tracking	No retry policy or naive implementation
Bulkhead Isolation	Are resources isolated by criticality? How is saturation monitored?	Thread/connection pools isolated by function, limits enforced, per-pool saturation monitoring and alerts	Some resource isolation, partial monitoring	Shared pools without isolation
Graceful Degradation	Are fallback behaviors defined? How is mode switching triggered?	Degraded modes documented, automatic switching based on health checks, customer notifications, regular testing	Modes defined but manual activation, limited testing	No degradation strategy

This condensed format enables rapid assessment during implementation reviews or architecture decision sessions while maintaining clear scoring criteria.

Resiliency Pattern Implementation Example

Consider an e-commerce checkout service implementing comprehensive resiliency: circuit breakers prevent cascading payment gateway failures, 3-second timeouts with fallback to secondary processors handle slow responses, exponential backoff retries (100ms, 200ms, 400ms, 800ms with jitter) manage transient errors, dedicated payment thread pools (50 threads) isolate from browse operations (150 threads), and graceful degradation queues payments asynchronously when gateways fail while displaying "Payment processing experiencing delays" to customers. This multi-layered approach achieves blast radius containment (failures isolated to <10% of capacity) and 95%+ automatic recovery for transient issues.

Organizations with strong resiliency maintain customer experience even during partial outages through intelligent fallback behaviors and automated recovery mechanisms. The goal is to achieve "blast radius" containment where failures isolate to less than 10% of system capacity, and automatic recovery success rates exceeding 95% for transient failures.

Dimension 6: Risk Assessment

Risk assessment uses Failure Mode and Effects Analysis (FMEA) methodology to systematically identify, quantify, and mitigate potential failures across critical customer journeys before they impact users.

Core Parameters:

- **FMEA Coverage:** Critical customer journeys documented and analyzed (typically 5-10 journeys covering 80% of business value), RPN calculation for each failure mode ($RPN = \text{Severity} \times \text{Occurrence} \times \text{Detection}$)
- **High-Risk Mitigation:** Monitoring parameters derived for all high-RPN failure modes ($RPN > 200$), Mitigation strategies implemented and tested for critical risks, Regular FMEA review and updates as system architecture evolves
- **Risk-Based Prioritization:** Engineering backlog prioritization incorporating RPN scores, Architecture decision records documenting risk trade-offs

FMEA methodology ensures that monitoring and resiliency investments focus on failure modes with highest potential customer impact, preventing wasteful effort on low-risk scenarios while addressing critical vulnerabilities.

6. Sample Scoring and Statistical Scenarios

Having established the six dimensions and their constituent parameters, we now demonstrate how the framework operates in practice through concrete examples. These scenarios illustrate both the scoring calculation mechanics and the statistical methods underlying anomaly detection at Level 1 of the hierarchical observability architecture.

6.1 Complete CUJ Scoring Example

Table 6.1: E-commerce Checkout Flow - MK Score Calculation

Pillar	Parameter	Weight	Actual	Expected	Weighted Actual	Weighted Expected	Gap
Observability	Payment Gateway Latency	5	3	5	15	25	10
Observability	Checkout Flow Tracing	4	4	5	16	20	4
Observability	Error Rate Monitoring	5	5	5	25	25	0
Optimal Alerting	Payment Failure Alerts	5	5	5	25	25	0
Optimal Alerting	Cart Abandonment Detection	3	2	5	6	15	9

Release Strategy	Deployment Success Rate	5	4	5	20	25	5
Release Strategy	Automated Rollback	4	5	5	20	20	0
Incident Metrics	Payment MTTD	5	2	5	10	25	15
Incident Metrics	Checkout MTTR	5	3	5	15	25	10
Resiliency	Payment Circuit Breaker	4	5	5	20	20	0
Resiliency	Database Failover	4	3	5	12	20	8
TOTALS		49			184	245	

MK Score = $(184 / 245) \times 100\% = 75.1\%$

Interpretation: Acceptable system requiring targeted improvements

Priority Remediation (by Gap Score):

1. Payment MTTD (Gap: 15) - Improve detection from current 15min to <5min target
2. Payment Gateway Latency (Gap: 10) - Optimize response time
3. Checkout MTTR (Gap: 10) - Faster incident resolution
4. Cart Abandonment Detection (Gap: 9) - Implement monitoring

6.2 Statistical Anomaly Detection Scenarios

Table 6.2: Response Time Distribution Analysis

Percentile	Response Time	Interpretation
P50 (Median)	210ms	Typical transaction performance
P75	270ms	Acceptable performance range
P95	400ms	Threshold for performance degradation alerts
P99	485ms	Threshold for critical investigation
Current: 520ms	Exceeds P99	Action Required: Investigate transaction

Statistical Summary:

- Mean (μ): 200ms
- Standard Deviation (σ): 100ms
- Distribution: Right-skewed (occasional slow outliers)

Detection Method Selection:

Table 6.3: Anomaly Detection Approach

Metric Characteristic	Recommended Method	Threshold Example	Rationale
Bell-curve distribution	3-sigma ($\mu \pm 3\sigma$)	200ms $\pm$ 300ms = 500ms upper limit	Statistical validity for normal distributions
Right-skewed (long tail)	P99 percentile	485ms	Accounts for expected slow transactions
Bimodal (distinct peaks)	Separate percentiles per mode	P99 per traffic pattern	Different thresholds for high/low traffic
Seasonal patterns	Time-windowed percentiles	P99 per hour-of-day	Accounts for expected variation

Practical Recommendation: For most system metrics showing non-normal distributions (response times, latency), use percentile-based thresholds (P95 or P99) rather than 3-sigma calculations.

7. Framework Differentiation

The MK Scoring Framework addresses three fundamental limitations of traditional reliability assessment:

Table 7.1: Traditional Assumptions vs. MK Framework Approach

Conventional Assumption	Traditional Approach	MK Framework Approach		Measurable Impact
Static Thresholds Suffice	Fixed thresholds (e.g., "alert when CPU >80%") applied regardless of system context or behavior patterns	Statistical boundaries (percentiles, 3-sigma) that adapt to actual system behavior and account for normal variation		67% reduction in false positives (evidenced by MTTD improvement: 30min→10min)
Technical and Business Metrics Are Separate	Technical teams measure error rates independently; business teams separately assess revenue impact with manual correlation	Mathematical correlation: $\text{Business_Impact_Score} = \sum(\text{Anomaly} \times \text{Volume} \times \text{Revenue})$ provides automated real-time business impact assessment		Business impact visibility: hours/days → seconds, enabling immediate executive decision-making
Output Format	Maturity level (CMMI Level 3)	Percentage score with component breakdown	Intuitive interpretation, trend tracking	

8. Demonstrated Impact

Validation across multiple enterprise environments demonstrates measurable impact: 67% MTTD reduction, 57-67% MTTR reduction, 40-60% cost optimization, and quantified business outcomes including multi-million-dollar cost avoidance. **Automotive Finance Sector (2019):**

- MTTD: 67% reduction (30 minutes → 10 minutes)
- MTTR: 57% reduction (3.5 hours → 1.5 hours)
- Business impact assessment: hours/days → real-time
- Industry recognition: Methodology showcased to 100+ Fortune 500 organizations via observability platform provider

Healthcare Manufacturing (2021-2022):

- MK Score improvement: 25% increase within 6 months
- MTTR: 67% reduction (3 hours → 1 hour)
- MTTD: <15 minutes through automated detection
- Cost optimization: 40% overall logging cost reduction, 60% for specific services
- Business outcomes: 99% reduction in surgical pack shortages, price assurance 75%→98%
- Operational impact: 30 human hours saved per major incident × 100 incidents/year

Financial Services (2024):

- Gamified CUJ evaluation being assessed by Global Technology SRE panel.

9. Conclusion

The MK Scoring Framework provides a systematic, quantitative approach to reliability engineering that integrates statistical analysis, risk assessment, and business-aligned scoring into a unified methodology. Through hierarchical observability architecture spanning individual metrics, customer journeys, and system-wide assessment, organizations gain unprecedented capability for objective reliability measurement and data-driven prioritization.

The framework's core strength lies in resolving the traditional tension between universal applicability and organizational customization. Through Goal-Question-Metric theory integration, the same mathematical foundation adapts to diverse contexts while maintaining scoring comparability. Organizations in automotive finance, healthcare manufacturing, and financial services have deployed identical methodology with context-specific weights, producing comparable reliability scores that enable portfolio-wide assessment.

For organizations seeking to mature reliability engineering practices, the MK Scoring Framework offers a proven methodology combining established SRE principles (Golden Signals, SLO/SLI, FMEA, resiliency patterns) with quantitative rigor and business alignment. The framework transforms reliability from qualitative aspiration to measurable, improvable organizational capability.