International Journal of Engineering and Information Systems (IJEAIS)
ISSN: 2643-640X
Vol. 10 Issue 10, October - 2026, Pages: 202-213

Trust-Aware and Uncertainty-Calibrated Resource Selection in Federated Cloud Computing: A Conceptual Framework

Dr Neha Paliwal

HOD (Computer Science Department)
Shri Mahaveer College , Jaipur

Abstract: *Federated cloud computing enables organizations to coordinate resources across multiple cloud providers while retaining independent administrative control. However, selecting an appropriate provider becomes difficult when performance measurements are incomplete, service conditions fluctuate, and trustworthiness cannot be established through a single indicator. Existing resource-selection decisions may be affected by outdated monitoring information, conflicting service objectives, and differences between provider-reported performance and observed behaviour. This paper proposes a Trust-Aware and Uncertainty-Calibrated Resource Selection (TUCRS) framework to address these challenges. The framework introduces a conceptual distinction between resource suitability, evidence confidence, and provider trust. Rather than treating estimated performance values as completely reliable, TUCRS represents the uncertainty associated with resource observations and incorporates it into the selection process. Mandatory security and policy requirements are evaluated independently of preference-based optimization. Eligible providers are subsequently assessed using performance, cost, availability, trust evidence, and uncertainty-adjusted utility. The framework also introduces an evidence-refresh mechanism, a decision-stability principle, and a structured explanation of resource-selection outcomes. Five research propositions are developed to guide future theoretical and empirical investigations. The paper contributes a conceptual model and formal decision structure rather than an experimentally validated algorithm. Its proposed relationships require evaluation through mathematical analysis, simulation, and real-world multi-cloud deployments.*

Keywords: Federated cloud computing, resource selection, uncertainty modelling, provider trust, cloud monitoring, multi-criteria decision-making, service reliability, cloud orchestration.

1. Introduction

Cloud computing has become an important foundation for delivering computing infrastructure, software platforms, and digital services. Organizations may distribute their workloads across several cloud providers to improve flexibility, support different application requirements, and reduce dependence on a single infrastructure. Cloud federation extends this approach by allowing independently managed cloud environments to coordinate resource provisioning and service delivery.

However, coordination across providers introduces a difficult decision problem. A resource-selection mechanism must evaluate not only the technical suitability of a resource but also the reliability of the information used to assess it. Provider A may report low latency, while Provider B may provide more recent and complete monitoring measurements. Provider C may offer a lower price but have limited evidence of consistent service performance.

A decision system that considers only reported values may select a resource that appears optimal but performs poorly after deployment. Conversely, rejecting every resource associated with incomplete information may unnecessarily restrict the available choices. Effective selection therefore requires a distinction between a resource's estimated performance and the confidence that can reasonably be placed in that estimate.

Trust presents a related challenge. Provider trust may depend on historical service behaviour, the consistency of monitoring records, compliance with agreed policies, and the integrity of reported information. Trust should not be confused with a single availability percentage or a general reputation score. A provider may deliver reliable performance while failing to satisfy a particular security requirement. Consequently, mandatory security conditions must remain independent of any aggregated trust score.

Previous research has examined resource management in federated clouds, including pricing, discovery, selection, monitoring, and allocation. Trust-based cloud service selection has also investigated the combination of objective monitoring and subjective user feedback [1, 2]. These research directions provide a foundation for further investigation, but they do not remove the need to formulate selection decisions around the uncertainty and limitations of available evidence.

This paper proposes the Trust-Aware and Uncertainty-Calibrated Resource Selection (TUCRS) framework. Its central premise is that a resource-selection decision should depend on three distinct questions:

1. Is the resource eligible under the application's mandatory requirements?
2. How suitable is the resource according to the application's service objectives?
3. How reliable is the evidence supporting the estimated suitability and provider trust?

The purpose of this paper is to establish a theoretical foundation for addressing these questions. It does not claim that the proposed framework has been implemented or that its effectiveness has been experimentally demonstrated.

1.1 Research problem

The research problem can be expressed as follows:

How can a federated cloud orchestration system select an eligible resource when performance estimates, monitoring confidence, and provider trust differ across candidates, without allowing uncertainty or high performance scores to override mandatory security requirements?

This formulation places evidence quality and trust at the centre of the selection process.

1.2 Research objectives

The objectives are to:

- Define a conceptual distinction between resource suitability, observation uncertainty, and provider trust.
- Develop a formal model for uncertainty-aware resource selection.
- Establish a feasibility-first mechanism for enforcing mandatory requirements.
- Introduce principles for refreshing evidence and avoiding unstable placement decisions.
- Formulate research propositions that can be tested in future studies.

1.3 Research contributions

The paper proposes a conceptual decision framework, a set of formal definitions, an uncertainty-adjusted utility model, and a collection of research propositions. Its intended contribution is the integration of these elements into one explicit decision structure. The originality of this integration must ultimately be assessed through a systematic review of existing research and subsequent validation.

2. Conceptual Foundations

2.1 Federated cloud environments

A federated cloud environment consists of independently managed computing resources that cooperate to support application requirements. Participating providers may differ in infrastructure, service interfaces, pricing, security policies, geographical location, and monitoring capabilities.

These differences mean that resource selection cannot always rely on direct comparisons of provider-reported values. Measurements may have different observation periods, definitions, levels of completeness, or degrees of uncertainty.

A common decision framework must therefore normalize the relevant information while preserving its origin, observation time, and limitations.

2.2 Resource suitability

Resource suitability describes how well a candidate resource satisfies the requirements of a particular workload. It may depend on latency, availability, cost, capacity, data location, and application-specific requirements.

Suitability is contextual rather than universal. A resource appropriate for a background analytics workload may be unsuitable for a latency-sensitive application. Therefore, a resource-selection mechanism should evaluate candidates against the requirements of the current workload instead of assigning every resource a permanent global ranking.

2.3 Observation uncertainty

Observation uncertainty refers to limitations in the available evidence about a resource. It may arise from:

- Missing monitoring measurements.
- Outdated performance observations.
- Short observation periods.
- Variations in workload conditions.

- Differences between measurement methods.
- Inconsistent reports from multiple monitoring sources.

Uncertainty is not identical to poor performance. A resource may have a promising estimated response time but insufficient evidence to support a confident decision. Similarly, a resource may have a high-confidence estimate indicating that its response time is consistently poor.

The model must represent these situations separately.

2.4 Provider trust

Provider trust represents the degree to which available evidence supports confidence in a provider's service behaviour and compliance with agreed expectations.

Trust may incorporate historical service-level compliance, monitoring consistency, incident handling, audit evidence, and verified security practices. The exact components depend on the application and the federation's governance arrangements.

A trust score should not be interpreted as proof of security. Mandatory security requirements must be verified independently, and uncertainty about a critical security control should lead to additional verification or rejection rather than being compensated for by a high score elsewhere.

2.5 Decision stability

Resource selection is a repeated decision in dynamic environments. Small fluctuations in monitoring values may cause the preferred provider to change repeatedly, creating unnecessary migrations and operational overhead.

Decision stability concerns the ability to avoid frequent placement changes when the evidence does not demonstrate a meaningful improvement. It requires the selection process to account for switching costs, minimum improvement thresholds, and the consequences of migration.

3. Research Gap and Conceptual Positioning

Research on cloud federation has already examined resource discovery, monitoring, pricing, resource selection, and allocation [1]. Trust-oriented service-selection research has also explored combining monitoring evidence with user feedback [2]. Therefore, neither cloud federation nor trust-based selection is, by itself, a new research contribution.

The conceptual question addressed here is narrower: how should a resource-selection model distinguish the estimated quality of a candidate from the uncertainty of the evidence supporting that estimate, while also enforcing trust-related eligibility requirements and maintaining decision stability?

Table 1. Conceptual challenges and proposed responses

Challenge	Potential decision error	TUCRS design response
Incomplete monitoring	A candidate is ranked using insufficient evidence	Explicit observation-confidence assessment
Stale measurements	Historical performance is treated as current	Evidence-age tracking and refresh requirements
Conflicting objectives	One attractive metric dominates the decision	Multi-criteria utility model
Unverified trust	Reputation is treated as proof of compliance	Independent policy and security checks
Rapid metric fluctuations	Placement changes repeatedly	Stability threshold and switching-cost consideration
Heterogeneous providers	Incomparable metrics produce misleading rankings	Common measurement definitions and normalization
Opaque selection	Operators cannot understand the chosen provider	Recorded evidence and decision explanations

Source: Proposed conceptual analysis.

The table identifies design requirements for the proposed model. It does not claim that every existing resource-selection system exhibits these weaknesses.

The principal distinction of TUCRS is its treatment of evidence quality as a first-class decision variable rather than as an implicit assumption.

4. Proposed TUCRS Framework

4.1 Framework overview

The framework contains six functional components:

1. Workload and policy specification.
2. Provider discovery and eligibility verification.
3. Evidence collection and normalization.
4. Uncertainty and trust assessment.
5. Resource ranking and decision-stability control.
6. Placement execution and feedback.

These components define logical responsibilities and do not require a particular programming language, cloud platform, or deployment topology.

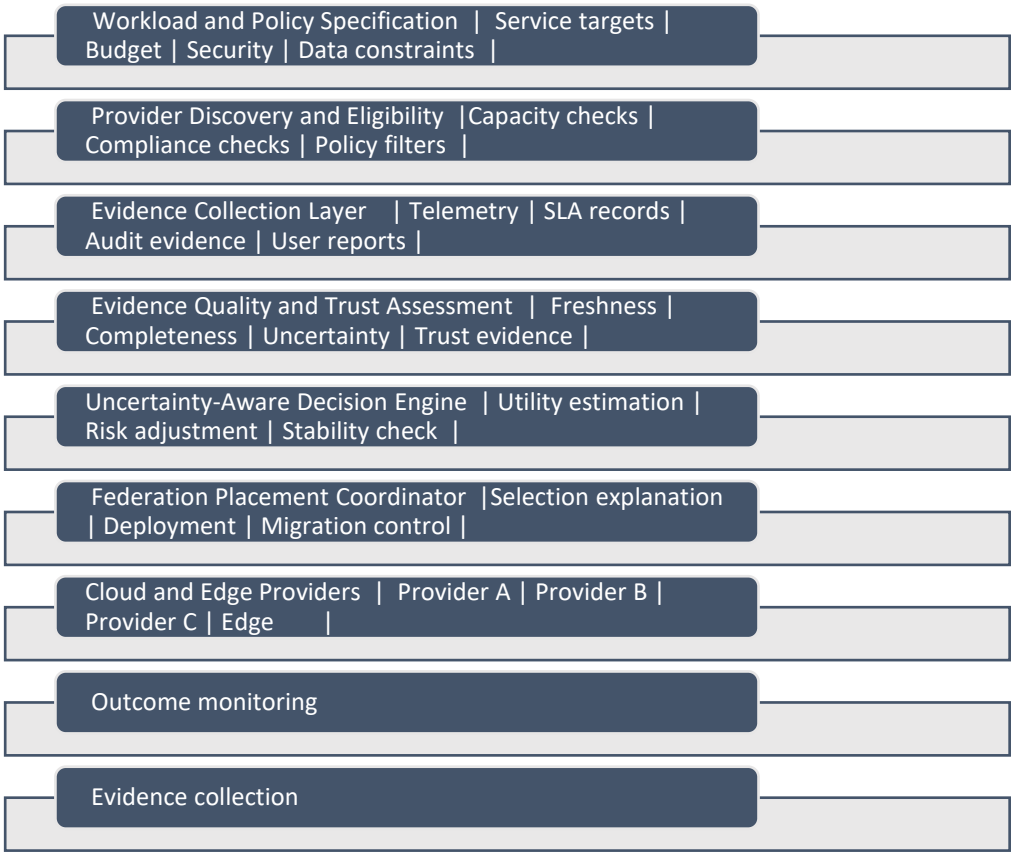


Figure 1. Conceptual architecture of TUCRS

Source: Original conceptual diagram proposed in this paper.

4.2 Workload and policy specification

The first component converts an application request into a structured requirement profile.

The profile may contain:

- Maximum acceptable latency.
- Minimum availability.
- Required processing capacity.
- Maximum operational cost.
- Permitted data locations.

- Mandatory security controls.
- Application-specific performance priorities.

The profile distinguishes **hard constraints** from **soft preferences**. Hard constraints determine whether a candidate is eligible. Soft preferences determine how eligible candidates are ranked.

4.3 Provider discovery and eligibility verification

The discovery component identifies resources that could potentially execute the workload. It then evaluates their capacity, access permissions, security controls, and policy compatibility.

A candidate that violates a mandatory requirement is excluded from subsequent preference-based ranking. This prevents a low price or attractive latency estimate from compensating for a critical compliance failure.

4.4 Evidence collection and normalization

The evidence component gathers available measurements and supporting records from participating providers.

Each observation should ideally include:

- The metric and its unit.
- The observation timestamp.
- The measurement interval.
- The measurement source.
- The applicable workload conditions.
- The method used to derive the value.

A common representation is necessary because provider measurements may otherwise be incomparable. Missing values should be marked explicitly rather than silently replaced with assumed values.

4.5 Uncertainty and trust assessment

The framework evaluates two distinct aspects of the evidence.

Observation confidence describes the quality and reliability of the measurements available for a particular decision.

Provider trust describes the degree of support for confidence in the provider's behaviour and adherence to agreed expectations.

A provider may have a strong historical trust record but currently lack fresh performance measurements. Conversely, a provider may supply recent telemetry but have insufficient evidence of long-term reliability.

Keeping these dimensions separate prevents a strong historical record from automatically eliminating the need for current monitoring.

4.6 Decision and placement coordination

The decision engine first filters ineligible resources. It then calculates an uncertainty-adjusted score for the remaining candidates.

Before changing an existing placement, the coordinator considers whether the proposed alternative offers a sufficiently meaningful improvement after accounting for switching costs and operational risks.

The system records the criteria, evidence, and constraints responsible for the decision. This information supports auditing and future model evaluation.

5. Mathematical Formulation

This section presents a formal conceptual model. Its equations define a proposed decision structure; they are not a claim that a validated algorithm has been implemented.

5.1 Resource set and observations

Let the federation contain a finite set of candidate resources:

[R={r_1,r_2,\ldots,r_n}]

For each resource (r_i), define an observation vector:

[X_i=(L_i,C_i,A_i,K_i)]

where:

- (L_i) denotes estimated latency.
- (C_i) denotes estimated cost.
- (A_i) denotes estimated availability.
- (K_i) denotes a normalized capacity or workload-fit measure.

Let (q_i) represent the confidence associated with the available performance evidence, and let (t_i) represent the provider-trust estimate.

Both (q_i) and (t_i) are defined on the interval ([0,1]), where larger values indicate stronger confidence or trust evidence under the chosen assessment procedure.

These scores must be derived from documented rules or calibrated statistical methods. They should not be assigned arbitrarily in an empirical implementation.

5.2 Feasibility conditions

Let (F_i) indicate whether candidate (r_i) satisfies all mandatory requirements:

[F_i =\begin{cases}1, & \text{if all required conditions are satisfied},\0, & \text{otherwise}.\end{cases}]

The eligible set is:

[R_F={r_i\in R=1}]

If (R_F) is empty, the model reports that no feasible candidate is available under the current requirements.

This is an important property of the framework: the model does not select a non-compliant resource merely because all available alternatives perform poorly.

5.3 Normalized suitability

Let each candidate have normalized benefit scores:

[b_{ij}\in[0,1]]

where (b_{ij}) represents the suitability of candidate (i) on criterion (j). Higher values are defined to indicate greater suitability.

For latency and cost, lower raw values would normally produce higher normalized benefit scores. For availability and workload fit, higher raw values would normally produce higher scores.

The baseline suitability score is:

[S_i=\sum_{j=1}^{m}w_jb_{ij}]

subject to:

[w_j\geq0,\qquad \sum_{j=1}^{m}w_j=1]

The weights represent the priorities of the current workload. Different applications may use different weights.

This weighted-sum model is a simplifying assumption. Its suitability should be examined where criteria are correlated or where a very poor result on one dimension must not be compensated for by strong results on another.

5.4 Uncertainty-adjusted suitability

A central design principle of TUCRS is that uncertain estimates should not be treated as equally dependable as well-supported estimates.

Let (u_i\in[0,1]) denote the estimated uncertainty associated with the suitability assessment, where larger values indicate greater uncertainty.

Define the uncertainty-adjusted suitability as:

[S_i^{*}=S_i-\lambda u_i]

where (\lambda\geq0) controls the degree to which uncertainty affects the decision.

The parameter (\lambda) should reflect the application's tolerance for uncertain evidence. A safety-sensitive or interruption-sensitive application may require a more conservative policy than a low-risk background workload.

The expression above is a conceptual penalty model. Its linear form is not universally optimal. Future studies should compare it with confidence intervals, robust optimization, Bayesian decision models, and distribution-based risk measures.

5.5 Trust-aware selection

Let (t_i) represent the provider-trust estimate for an eligible candidate. A combined decision score may be written as:

[J_i=S_i^{*}+\mu t_i]

where (\mu\geq0) represents the influence of trust evidence.

The candidate is selected according to:

[r^{*}=\arg\max_{r_i\in R_F}J_i]

This formulation is intended for candidates that have already passed all mandatory policy checks. Trust scoring does not replace security verification.

The parameter values and score calibration require careful investigation. A high trust score should not be allowed to conceal uncertainty about a critical control or override a mandatory application requirement.

5.6 Evidence freshness

Let (\Delta_i) represent the elapsed time since the latest relevant observation for resource (r_i).

One possible freshness function is:

[f_i=e^{-\kappa\Delta_i}]

where (\kappa>0) controls the rate at which the value decreases as evidence becomes older.

The function expresses the assumption that older observations may become less informative about current conditions. The appropriate decay rate depends on the metric: network latency may change quickly, while some contractual attributes change infrequently.

Freshness should therefore be metric-specific where necessary. It should not be interpreted as a universal measure of accuracy.

5.7 Decision stability

Let (r_{\mathrm{current}}) be the resource currently hosting the workload and (r_{\mathrm{new}}) be a proposed alternative.

A migration should be considered only when the expected benefit exceeds the combined switching cost and a stability threshold:

[J_{\mathrm{new}}-J_{\mathrm{current}}

C_{\mathrm{switch}}+\tau]

where:

- (C_{\mathrm{switch}}) represents the normalized cost or penalty of migration.
- (\tau) represents a minimum improvement threshold.

This rule is intended to reduce unnecessary migrations caused by small score fluctuations.

A more complete implementation would distinguish monetary costs from service disruption, state-transfer overhead, security risks, and recovery requirements. Those components would need appropriate normalization or a constrained optimization treatment.

6. Decision Procedure

The following procedure summarizes the proposed selection logic.

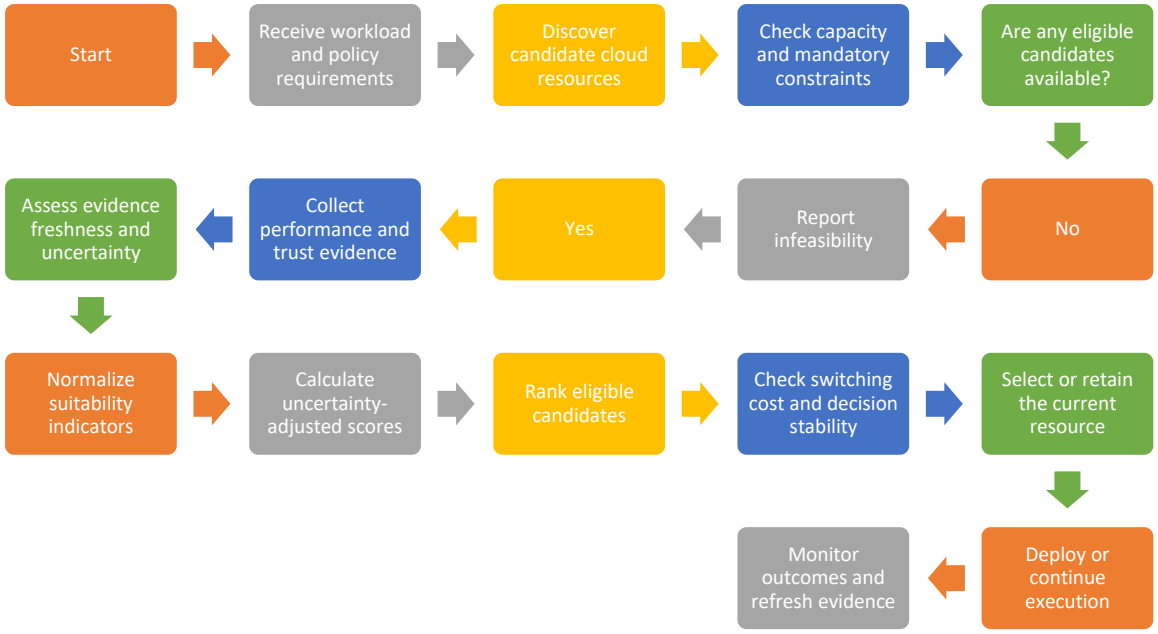


Figure 2. TUCRS decision workflow

Source: Proposed conceptual workflow.

Table 2. Decision stages and expected outputs

Stage	Input	Expected output
Requirement analysis	Workload profile and policies	Structured constraints and preferences
Candidate discovery	Provider registry	Candidate resource set
Feasibility evaluation	Capacity and compliance evidence	Eligible resource set
Evidence assessment	Telemetry and provider records	Confidence, uncertainty, and trust estimates
Suitability evaluation	Normalized indicators	Baseline and adjusted scores
Stability evaluation	Current placement and candidate scores	Migration or retention decision
Outcome monitoring	Runtime measurements	Updated evidence for subsequent decisions

The decision procedure is intentionally modular. The same conceptual structure could accommodate different monitoring systems, trust models, and optimization algorithms.

7. Research Propositions

Because the paper is conceptual, its contributions are expressed as research propositions rather than empirical findings.

Proposition 1: Evidence quality and selection reliability

When candidate resources have materially different levels of observation uncertainty, explicitly accounting for uncertainty should reduce the likelihood of selecting a candidate whose apparent advantage depends on unreliable evidence, provided that the uncertainty estimates are reasonably calibrated.

Rationale: A ranking based only on point estimates ignores differences in the strength of the underlying evidence. An uncertainty-aware model makes those differences visible.

Possible test: Compare point-estimate ranking with uncertainty-aware ranking using known resource-performance distributions and measure decision error.

Proposition 2: Separation of feasibility and utility

Separating mandatory eligibility checks from preference-based scoring should make policy compliance easier to verify than a single weighted score that combines mandatory and optional criteria.

Rationale: A weighted score can allow a high performance or cost score to compensate numerically for a low security score. A feasibility-first model prevents such compensation for requirements designated as mandatory.

Possible test: Construct cases in which an otherwise attractive resource violates one required policy and verify that the decision mechanism consistently rejects it.

Proposition 3: Trust and performance are distinct decision dimensions

Provider trust and observed performance should be represented separately when they describe different aspects of service suitability.

Rationale: Strong historical trust evidence does not guarantee current performance, while strong recent performance does not establish long-term policy compliance.

Possible test: Compare a joint trust-performance model against a model that uses performance alone across scenarios with changing performance and trust evidence.

Proposition 4: Decision stability and switching overhead

In environments where resource measurements fluctuate around similar values, incorporating switching costs and minimum improvement thresholds should reduce avoidable migrations.

Rationale: A decision system that reacts to every small score change may cause additional operational cost without producing a meaningful service benefit.

Possible test: Measure migration frequency, total switching cost, and service performance under different threshold settings.

Proposition 5: Evidence freshness and decision confidence

A metric-specific evidence-refresh policy should be more appropriate than a universal freshness threshold when monitoring indicators have different rates of change.

Rationale: The usefulness of an observation depends partly on how quickly the measured property changes.

Possible test: Compare metric-specific refresh intervals with a fixed interval under workloads with different temporal variability.

These propositions are theoretically motivated hypotheses. Their validity and practical importance remain open questions until they are tested.

8. Conceptual Application Scenarios

8.1 Healthcare information services

A healthcare information service may require strict access controls, approved data locations, and high service availability. TUCRS would first reject resources that fail mandatory requirements. It would then compare eligible resources using performance evidence, cost, monitoring confidence, and trust-related records.

The framework would not automatically infer that a provider is safe merely because it has a high trust score. Compliance evidence would remain a separate requirement.

8.2 Online learning platforms

An online learning platform may prioritize interactive responsiveness during live classes while placing greater emphasis on cost for background content processing.

The workload profile could assign different preferences to different application components. The selection mechanism would evaluate those preferences while retaining mandatory data-protection requirements.

8.3 Scientific and batch computing

A batch-processing workload may tolerate longer execution times but require predictable capacity and controlled expenditure. In this setting, uncertainty in completion-time estimates and the reliability of capacity information may matter more than very low network latency.

These scenarios illustrate how the conceptual model could be configured for different application requirements. They do not demonstrate that TUCRS has achieved any specific performance improvement.

9. Security, Governance, and Ethical Considerations

9.1 Independent verification of mandatory controls

Security and compliance requirements should be assessed through appropriate evidence, such as configuration checks, verified policies, and relevant audit records. A composite trust score is not sufficient evidence that a critical control is functioning correctly.

9.2 Integrity of monitoring information

A provider or monitoring component may supply incorrect, incomplete, or misleading observations. The framework should preserve the source and timestamp of important measurements and apply validation checks where possible.

Where monitoring information cannot be sufficiently verified, the system should communicate that limitation rather than treating the estimate as certain.

9.3 Accountability and explainability

Each decision should retain a record of:

- The workload requirements used.
- The candidates considered.
- The constraints that excluded candidates.
- The evidence used to calculate the ranking.
- The uncertainty and trust assessments applied.
- The reason for selecting, retaining, or migrating a workload.

Such records can help operators review incorrect decisions and improve the model over time.

9.4 Privacy and data minimization

Trust assessment may involve provider records, user feedback, and operational logs. The framework should collect only information necessary for the decision, restrict access to sensitive records, and apply suitable retention policies.

9.5 Avoiding unfair provider exclusion

A provider with limited historical data may receive a lower confidence estimate even when its actual performance is acceptable. A poorly designed uncertainty penalty could systematically disadvantage newer or smaller providers.

Future implementations should investigate how to distinguish genuine evidence of poor reliability from a simple lack of observations. Where appropriate, the system may request further evidence or conduct a limited, low-risk assessment rather than automatically excluding a provider solely because it is less established.

10. Proposed Validation Strategy

A conceptual framework requires a validation plan even when no experiments are reported in the present paper.

Table 3. Suggested validation programme

Research question	Evaluation approach	Potential measures
Does uncertainty-aware ranking improve selection quality?	Compare with point-estimate ranking	Decision regret, constraint satisfaction
Does trust add useful information?	Compare trust-aware and trust-free variants	Failure frequency, SLA violations
Does evidence freshness matter?	Vary monitoring age and update intervals	Prediction error, selection stability
Does stability control reduce unnecessary migrations?	Vary switching thresholds	Migration count, switching overhead
Are decisions explainable?	Review recorded decision traces	Trace completeness, operator assessment
Does the model generalize across workloads?	Evaluate distinct application profiles	Performance by workload category

A future study should define its baseline methods, workload distributions, provider characteristics, uncertainty-generation process, and evaluation metrics before collecting results.

The study should also report the limitations of its test environment. Results from a simulated environment would not automatically establish performance in production systems.

10.1 Analytical validation

The mathematical formulation should first be examined for consistency. This includes verifying that normalized scores remain within their defined ranges, mandatory constraints are always enforced, and the decision rule behaves sensibly when monitoring information is missing or stale.

Sensitivity analysis should investigate how the uncertainty penalty, trust weight, and migration threshold influence resource selection.

10.2 Simulation-based validation

A subsequent study could use a cloud or edge simulation environment to represent multiple providers with varying performance, prices, monitoring delays, and reliability characteristics.

The simulation should compare TUCRS with clearly defined baselines under identical workloads. All generated data and assumptions should be documented.

10.3 Real-world validation

A later deployment could evaluate the framework using actual telemetry and operational records. Such a study would need to consider provider API limitations, measurement overhead, privacy obligations, and the costs associated with migrating active workloads.

The present paper reports no simulation or deployment results.

11. Limitations and Future Research

The proposed framework has several limitations.

First, the uncertainty-adjusted score is a conceptual formulation. Its linear penalty may not represent all forms of uncertainty or risk. Alternative probabilistic and robust decision models should be investigated.

Second, trust is difficult to define universally. The evidence that supports trust in a commercial web service may differ from the evidence required for a regulated healthcare workload. Trust components must therefore be justified in the relevant domain.

Third, uncertainty and trust estimates may themselves be unreliable. A poorly calibrated uncertainty model could penalize suitable resources, while a weak trust model could create unjustified confidence.

Fourth, the model assumes that relevant indicators can be normalized and combined. The choice of weights may materially influence the ranking, and no universal weighting scheme is likely to suit every workload.

Fifth, provider participation may be constrained by business agreements, technical interoperability, data residency, and administrative policies. These conditions may limit the candidates that can actually be selected.

Future work should investigate probabilistic uncertainty estimates, domain-specific trust evidence, explainable resource selection, fairness for providers with limited historical data, and the relationship between decision stability and service-level compliance. Empirical studies should establish whether the proposed conceptual distinctions improve selection quality under realistic conditions.

12. Conclusion

This paper introduced TUCRS, a conceptual framework for resource selection in federated cloud environments where performance evidence, monitoring confidence, and provider trust vary across candidates.

The framework distinguishes mandatory eligibility from preference-based optimization and separates estimated resource suitability from uncertainty in the supporting evidence. It further introduces a conceptual trust-aware ranking mechanism, evidence-freshness considerations, and a stability rule intended to reduce unnecessary placement changes.

The principal contribution is a formalized research direction in which resource-selection decisions are evaluated not only by their predicted outcomes but also by the reliability of the evidence supporting them. Five research propositions and a validation programme were developed to guide future investigations.

The framework remains theoretical. Its equations, decision principles, and anticipated benefits must be tested against alternative methods before claims about improved reliability, reduced migration overhead, or better resource selection can be substantiated. A systematic literature review is also necessary to determine the novelty of the integrated framework relative to existing uncertainty-aware and trust-based cloud-selection methods.

References

1. Liaqat, M., Chang, V., Gani, A., Hamid, S. H. A., Toseef, M., Shoaib, U., & Ali, R. L. (2017). Federated cloud resource management: Review and discussion. *Journal of Network and Computer Applications, 77*, 87–105. https://doi.org/10.1016/j.jnca.2016.10.008
2. *Towards a trust evaluation middleware for cloud service selection*. (2017). *Future Generation Computer Systems, 74*, 302–312. https://doi.org/10.1016/j.future.2016.01.009
3. Mell, P., & Grance, T. (2011). *The NIST Definition of Cloud Computing* (NIST Special Publication 800-145). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-145
4. Armbrust, M., Fox, A., Griffith, R., Joseph, A. D., Katz, R., Konwinski, A., Lee, G., Patterson, D., Rabkin, A., Stoica, I., & Zaharia, M. (2010). A view of cloud computing. *Communications of the ACM, 53*(4), 50–58. https://doi.org/10.1145/1721654.1721672
5. Zadeh, L. A. (1965). Fuzzy sets. *Information and Control, 8*(3), 338–353. https://doi.org/10.1016/S0019-9958(65)90241-X

Reference note: References 1–4 provide background on federated cloud management, trust-oriented selection, and cloud computing. Reference 5 is relevant only if fuzzy modelling is developed in subsequent work. Before submission, verify the full author information for Reference 2 and conduct a broader literature search for research directly addressing uncertainty-aware cloud selection, risk-sensitive optimization, and evidence-calibrated trust.