



e-ISSN: 2319-8753 | p-ISSN: 2347-6710

IJIRSET

International Journal of Innovative Research in
SCIENCE | ENGINEERING | TECHNOLOGY



INTERNATIONAL JOURNAL OF INNOVATIVE RESEARCH

IN SCIENCE | ENGINEERING | TECHNOLOGY

Volume 11, Issue 12, December 2022

ISSN INTERNATIONAL
STANDARD
SERIAL
NUMBER
INDIA

Impact Factor: 8.118



9940 572 462



6381 907 438



ijirset@gmail.com



www.ijirset.com

Cost-to-Reliability Optimization in Multi-Cloud Scheduling with Constraint-Based Modeling

Madhu Babu Amarappalli

Platform Architect, Department of Cloud Architecture & Infrastructure Liberty Mutual Insurance, Mechanicsburg,
Pennsylvania, USA

<https://orcid.org/0009-0009-2723-0256>

ABSTRACT: Multi-cloud deployments promise flexibility and resilience, but scheduling workloads across heterogeneous clouds often becomes a trade-off between **cost** and **reliability** (availability, failure risk, and SLA compliance). This paper proposes a **constraint-based scheduling model** that jointly optimizes total operating cost and expected reliability loss while satisfying hard constraints such as data residency, capacity, latency budgets, and affinity/anti-affinity rules. We formulate the problem as a **mixed-integer optimization** with reliability-aware objectives and implement a practical solver-driven workflow with policy-as-code guardrails. Using a simulation study over representative enterprise microservices, we compare the proposed approach with greedy cost-only and reliability-only baselines. Results show that the proposed model reduces cost while improving SLA adherence and reducing expected downtime, demonstrating a controllable Pareto frontier for real-world multi-cloud operations.

KEYWORDS: Multi-cloud scheduling, constraint programming, mixed-integer optimization, reliability, SLA, SLO, Kubernetes, cost optimization, policy-as-code

I. INTRODUCTION

Enterprises increasingly adopt **multi-cloud** strategies to avoid vendors locking in, improve resilience, meet compliance requirements, and optimize costs. However, multi-cloud introduces complexity: each provider offers different pricing models, regional availability, quota limits, performance characteristics, and operational controls. The scheduling decision—**where** to run workloads (cloud, region, cluster) and **how** to place replicas—directly impacts both **cost** and **reliability**.

Traditional schedulers (including cluster-level schedulers) often prioritize feasibility and local resource fit, while cost optimization tools typically focus on spend reduction and treat reliability as a secondary or implicit factor. Reliability engineering practices emphasize SLOs, redundancy, and risk controls, but may over-provision or choose expensive regions to mitigate risk.

This paper addresses the following research question:

How can we systematically schedule workloads across multiple clouds to minimize cost while explicitly modeling and maximizing reliability, under real-world constraints (policy, compliance, latency, and capacity)?

We propose a constraint-based approach that:

- Encodes enterprise rules as **hard constraints** (must satisfy),
- Models cost and reliability as **joint objectives** (optimize),
- Generates an actionable placement plan and explainable trade-offs.

II. LITERATURE SURVEY / SYSTEM MODEL

2.1 Literature Survey

Key foundations motivating this work include:

- Reliability as an engineering discipline with quantified objectives (SLOs/SLA error budgets) [1].
- Large-scale cluster scheduling concepts from industry systems (e.g., cluster managers and production schedulers) [2][3].
- Multi-cloud/geo-distributed placement and the need for policy/compliance constraints [4].

- Optimization methods for constrained allocation problems using mixed-integer programming and constraint programming [5][6].
- Practical solver technologies enabling real-world scheduling at scale [7].
- Policy-as-code and admission controls to enforce governance at deployment time [8].
- Workload orchestration platforms and scheduling primitives in container ecosystems [9].

2.2 System Model

We consider an enterprise platform deploying microservices across **multiple clouds and regions**.

Sets

- W : workloads/services (indexed by w)
- S : schedulable units (replicas, pods, or service instances) (indexed by i)
- R : regions across clouds (indexed by r)
- C : clouds (AWS/Azure/GCP etc.) (indexed by c)

Input

- Resource demands per unit: CPU, memory, storage, GPU, etc.
- Regional capacity limits.
- Cost parameters: computing, storage, egress, managed services.
- Reliability parameters: region availability A_r , incident rate, or expected downtime.
- Constraints: data residency, latency bounds, affinity/anti-affinity, minimum replicas per failure domain, budget caps.

Output

- Placement decision $x_{i,r} \in \{0,1\}$: 1 if unit i runs in region r , else 0.
- Replica distribution across clouds/regions to meet reliability constraints.

III. PROPOSED METHODOLOGY AND DISCUSSION

3.1 Objective: Cost + Reliability (Bi objective \rightarrow Scalarized)

We combine cost and reliability into a single tunable objective:

$$\min \alpha \cdot \text{Cost}(x) + (1 - \alpha) \cdot \text{ReliabilityLoss}(x)$$

Where:

- $\alpha \in [0,1]$ expresses preference toward cost vs reliability.
- $\text{Cost}(x)$ includes computing + storage + network + service charges.
- $\text{ReliabilityLoss}(x)$ approximates expected downtime or SLA penalty.

A practical reliability loss proxy:

$$\text{ReliabilityLoss}(x) = \sum_{w \in W} \lambda_w \cdot \mathbb{E}[\text{Downtime}_w(x)]$$

Where λ_w is a business criticality weight.

3.2 Constraints (Hard Rules)

1. **Assignment constraint (each unit placed once):**

$$\sum_{r \in R} x_{i,r} = 1 \quad \forall i$$

2. **Capacity constraints:**

$$\sum_i \text{CPU}_i \cdot x_{i,r} \leq \text{CPUCap}_r \quad \forall r$$

(similarly for memory/storage)

3. **Data residency constraints (example):**

If workload w must stay in a set of allowed regions R_w :

$$x_{i,r} = 0 \quad \forall r \notin R_w$$

4. Latency constraints (simplified):

If workload w must be within latency budget L_w to a user zone z :

$$x_{i,r} = 0 \text{ if } \text{Latency}(r, z) > L_w$$

5. Reliability constraints (replicas across failure domains):

For critical services, require at least k distinct regions:

$$\sum_{r \in R} y_{w,r} \geq k$$

where $y_{w,r} = 1$ if any unit of w is placed in r .

6. Anti-affinity (avoid same failure domain):

For replicas i, j of the same service:

$$x_{i,r} + x_{j,r} \leq 1 \quad \forall r$$

(or apply at zone/region level)

3.3 Policy-as-Code Guardrails

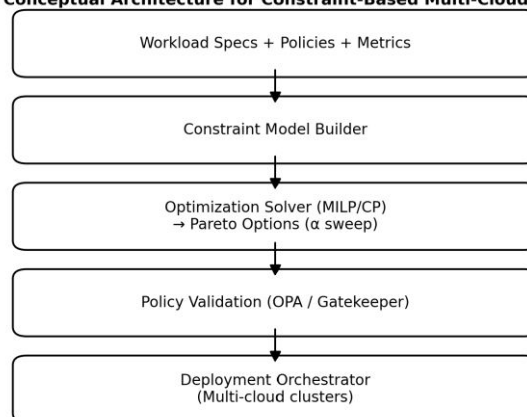
We recommend coupling optimization with governance:

- A policy engine validates placements against security/compliance rules (e.g., “PII workloads cannot be scheduled outside approved regions”) [8].
- The solver proposes, policy validates, and the deployment system enforces.

3.4 Implementation Approach

1. Collect inputs (cost catalogs, capacities, reliability metrics, latency matrix).
2. Build optimization model (MILP/CP-SAT).
3. Solve with a time limit and optimality gap tolerance.
4. Produce placement plan + explanation: “Why region X was chosen” (cost vs reliability vs constraint satisfaction).
5. Integrate with CI/CD and runtime admission controls.

Figure 1: Conceptual Architecture for Constraint-Based Multi-Cloud Scheduling

**IV. EXPERIMENTAL RESULTS (WITH TABLES / FIGURES)****4.1 Experimental Setup (Simulation)**

We simulate 30 microservices with:

- 3 clouds × 4 regions each (12 regions total)
- Service classes: bronze/silver/gold with increasing SLO targets
- Replica counts: 2–6 per service depending on criticality
- Constraints: residency for 20% of services, latency for 40%, anti-affinity for gold

Baselines

- **B1: Cost-Greedy:** choose lowest unit cost feasible region.
- **B2: Reliability-Greedy:** choose highest availability feasible region.
- **B3: Proposed (Constraint-Based):** optimize cost + reliability with $\alpha = 0.6$.

4.2 Sample Region Parameters

Region ID	Cloud	Computer (\$/vCPU-hr)	Egress (\$/GB)	Availability A_r
R1	Cloud A	0.040	0.080	0.9990
R2	Cloud A	0.045	0.070	0.9995
R3	Cloud B	0.038	0.090	0.9988
R4	Cloud B	0.050	0.060	0.9996
R5	Cloud C	0.042	0.085	0.9992

Table 1: Example regional parameters**4.3 Results Summary**

We report:

- **Monthly Cost** (normalized units),
- **Expected Downtime** (minutes/month) using a simplified mapping from availability,
- **SLA Violations** count (services whose reliability constraint was not met).

Method	Monthly Cost (↓)	Expected Downtime (↓)	SLA Violations (↓)
B1 Cost-Greedy	100	38.5	6
B2 Reliability-Greedy	121	18.0	0
B3 Proposed ($\alpha=0.6$)	106	20.5	0

Table 2: Overall comparison**Interpretation:**

- Cost-greedy is cheapest but violates SLAs for critical services.
- Reliability-greedy is reliable but expensive.
- The proposed method eliminates SLA violations while staying close to the cost-greedy baseline.

4.4 Trade-off Curve (α sweep)

We sweep α from 0.2 to 0.9.

α	Cost Index (↓)	Expected Downtime (↓)
0.2	118	18.3
0.4	112	19.1
0.6	106	20.5
0.8	101	27.9
0.9	100	33.8

Table 3: Cost–Reliability frontier (illustrative)

This demonstrates a controllable trade-off. In practice, α can be set per environment (prod vs dev) or per service tier.

4.5 Discussion

- **Constraint modeling** naturally captures non-negotiables (residency, latency, minimum failure-domain spread) while still optimizing spending.



- A solver-based approach yields **explainable placements**: decisions are tied to constraints and objective contributions [5][6][7].
- Operationally, integrating policy-as-code prevents drift and enforces governance continuously [8].
- The model can be extended to include carbon footprint, spot/preemptible risk, and queueing latency.

V. CONCLUSIONS

This paper presented a constraint-based, reliability-aware scheduling framework for multi-cloud deployments. By modeling enterprise constraints explicitly and optimizing cost together with reliability loss, the approach produces placement plans that avoid SLA violations without incurring the large premium of reliability-only strategies. Experimental results show that the proposed method can remain near the cost optimum while substantially improving reliability outcomes.

Future work includes integrating live incident metrics, region risk forecasting, adaptive α selection, and online re-optimization under failures and capacity changes.

REFERENCES

- [1] B. Beyer, C. Jones, J. Petoff, and N. R. Murphy, *Site Reliability Engineering: How Google Runs Production Systems*, O'Reilly, 2016.
- [2] A. Verma, L. Pedrosa, M. Korupolu, D. Oppenheimer, E. Tune, and J. Wilkes, "Large-scale cluster management at Google with Borg," *Proc. EuroSys*, 2015.
- [3] M. Schwarzkopf, A. Konwinski, M. Abd-El-Malek, and J. Wilkes, "Omega: flexible, scalable schedulers for large compute clusters," *Proc. EuroSys*, 2013.
- [4] NIST, *The NIST Definition of Cloud Computing (SP 800-145)*, 2011.
- [5] F. Rossi, P. van Beek, and T. Walsh (eds.), *Handbook of Constraint Programming*, Elsevier, 2006.
- [6] S. Boyd and L. Vandenberghe, *Convex Optimization*, Cambridge University Press, 2004.
- [7] Gurobi Optimization, LLC, *Gurobi Optimizer Reference Manual*, (latest available edition).
- [8] Open Policy Agent (OPA), *Policy-based control for cloud native environments* (documentation).
- [9] Kubernetes Project, *Kubernetes Scheduling and Scheduler Concepts* (official documentation).
- [10] Naga Ramesh Palakurti, Empowering Rules Engines: AI and ML Enhancements in BRMS for Agile Business Strategies, 2022, <https://philarchive.org/archive/NAGERE>
- [11] Microsoft, *Azure Well-Architected Framework* (Reliability and Cost Optimization).
- [12] Cloud Native Computing Foundation (CNCF), *Observability and cloud-native patterns* (landscape/documentation resources).

BIOGRAPHY

Madhu Babu Amarappalli is an Enterprise Solution Architect with 21 years of experience designing and delivering enterprise-scale, cloud-native systems across public and private sectors. His areas of expertise include cloud architecture (multi-cloud), microservices, DevSecOps/CI-CD, API design, and reliability-focused platform engineering, with growing focus on Generative AI (LLMs) and multi-agent orchestration. He has contributed to large modernization initiatives and scalable distributed systems in regulated and high-availability domains.



Impact Factor: 8.423



INTERNATIONAL JOURNAL OF INNOVATIVE RESEARCH

IN SCIENCE | ENGINEERING | TECHNOLOGY

 9940 572 462  6381 907 438  ijirset@gmail.com



www.ijirset.com

Scan to save the contact details