

# Structural Diagnostic Report

---

ai.01 — Interconnect Stability Control

## Scenario S3: Heterogeneous Accelerator Fabric

**Document Type:** Structural Diagnostic Scenario Report  
**Application ID:** ai.01  
**Scenario ID:** ai.01.S3  
**Schema Version:** 0.5.1

---

Source: SORT AI Structural Diagnostics Demo  
Application: ai.01 Interconnect Stability Control  
Scenario: S3  
Version: 1.0.1  
Generated: 2025-01-18  
Web: <https://independent-research-systems-modeling.com>

**Scope and Limits:** This report presents a structural diagnostic scenario analysis based on pre-computed, normalized projection runs. It is not a complete Architecture Risk Assessment and does not contain implementation guidance.

## 1. Scenario Overview

---

### System Class

Mixed accelerator fleet combining GPU, TPU, and NPU devices with heterogeneous execution characteristics and varying progress rates.

### Scale Abstraction

Straggler-dominated regime with reactive overprovisioning, high device heterogeneity, and significant progress rate variance across the fleet.

### Operational Context

Mixed training and inference workloads distributed across heterogeneous pipeline configurations. Variable synchronization intensity with capability-aware scheduling policies attempting to manage device diversity.

## 2. Observed Structural Pattern

---

The following structural effects emerge from the interaction of correctly functioning devices with different execution characteristics:

- Device capability differences create systematic rather than random straggler patterns, with predictable slowdowns emerging from scheduling-device mismatches.
- Homogeneous scheduling policies applied to heterogeneous fleets amplify coupling effects by treating devices with different progress rates as interchangeable.
- Synchronization points transform device heterogeneity into global execution coupling where the slowest device path determines system throughput regardless of fleet capacity.
- Overprovisioning increases device diversity within the fleet and may worsen coupling effects rather than resolve them.
- Straggler identification based on device identity misses the structural root cause, leading to ineffective mitigation strategies that address symptoms rather than coupling dynamics.

## 3. Stability Assessment

---

### Baseline Structural Condition

System operates in reactive mode with overprovisioning failing to resolve instability. Cost trajectory accelerating despite continued investment. Stability reserve is negative.

### Observed Instability Class

Unstable — characterized by declining efficiency despite investment, politicized performance diagnosis, and capacity additions that fail to yield proportional benefit.

### Post-Projection Stability Class

Marginal — straggler cascades contained through coupling-aware heterogeneous scheduling. Cost trajectory stabilized with limited but positive stability reserve.

## Transition Type

Partial stabilization from unstable to marginal operating conditions.

## 4. Aggregated Indicators

---

All values are normalized ratios. No absolute values or reconstructable parameters are provided.

Indicator	Baseline	Comparison	Direction
Fleet Utilization Efficiency	0.52	0.74	Improvement
Device Idle Time Ratio	0.41	0.19	Improvement
Straggler Cascade Frequency	0.38	0.12	Improvement
Progress Rate Variance	0.47	0.21	Improvement
Overprovisioning Effectiveness	0.34	0.67	Improvement
Scheduling Coherence Index	0.48	0.78	Improvement

## 5. Interpretation

---

### Systemic Relevance

The observed instability pattern is systemically relevant because it represents a fundamental mismatch between scheduling assumptions and physical device characteristics. The fleet contains devices that perform correctly to specification, yet system-level efficiency declines because scheduling decisions designed for homogeneous populations create systematic coordination failures when applied to heterogeneous ones. This is not a device problem or a capacity problem — it is a structural coupling problem.

### Detection Challenge

This instability class remains undetected in practice because device-level metrics show individual units performing to specification. The problem exists in the interaction between scheduling decisions and device population heterogeneity, a domain that standard monitoring does not observe. Performance diagnosis becomes political because different teams observe their components working correctly while system-level outcomes degrade.

Fleet efficiency declines despite investment, straggler discussions intensify without resolution, and overprovisioning provides only temporary relief — yet no single component or team appears responsible.

## 6. Decision Relevance

---

If a heterogeneous accelerator fleet shows declining efficiency despite investment, with straggler discussions becoming political and overprovisioning providing only temporary relief, the underlying cause is likely structural coupling rooted in scheduling-device mismatch.

Additional capacity increases device diversity and may worsen coupling effects. The problem cannot be provisioned away through investment in more hardware.

Structural visibility into heterogeneous coupling dynamics enables scheduling policies that work with rather than against device diversity, converting population heterogeneity from a liability into a manageable characteristic.

**Related Document:** [SORT AI Interconnect Application Context Brief](#)

---

*Such structural findings are typically contextualized through a scoped architecture risk assessment.*