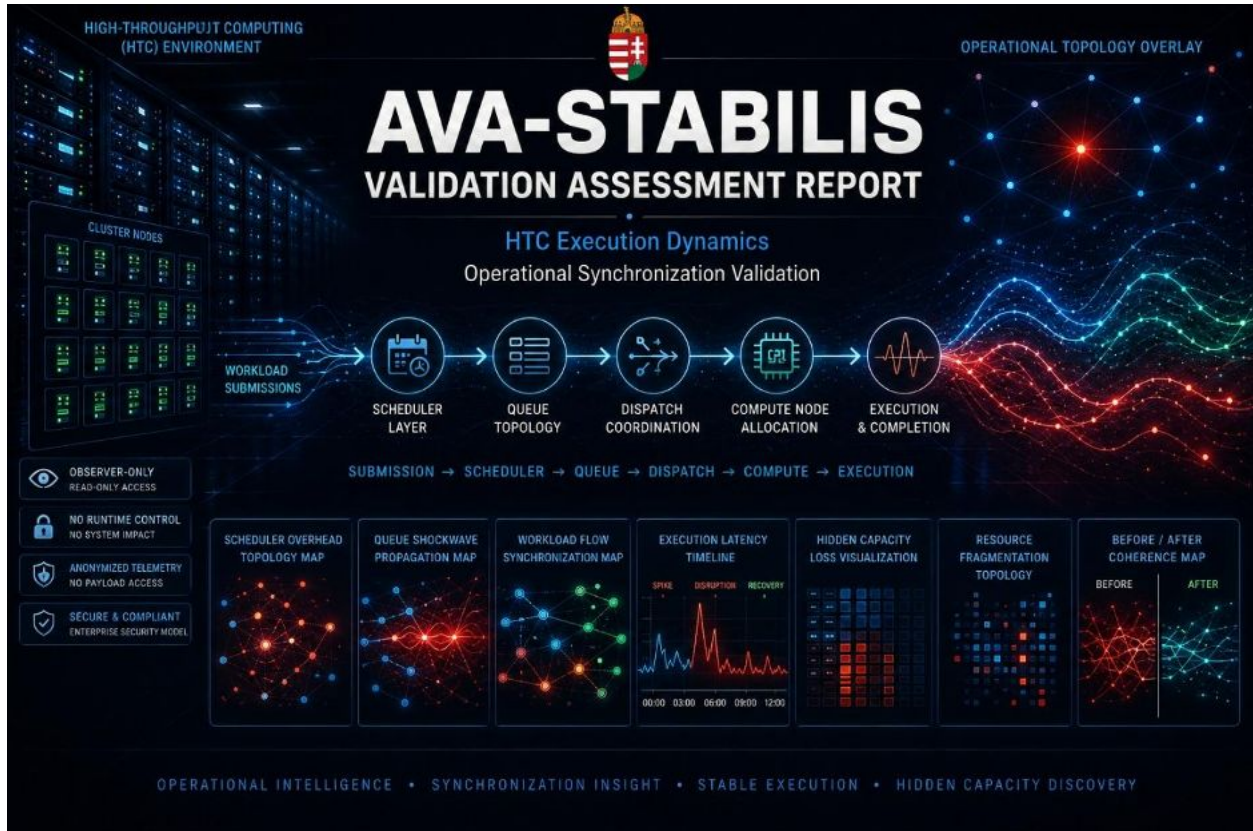


# AVA-STABILIS

## Operational Synchronization Modeling Report

### High-Throughput Computing (HTC) Operational Environment

#### Observer-Only Operational Analysis Pilot



This document is a partially modeled operational-analysis sample report. It is designed to demonstrate the AVA-Stabilis methodology and does not represent an audited customer deployment or a peer-reviewed scientific validation.

#### COVER PAGE

##### Pilot Name

AVA-STABILIS HTC Operational Synchronization Pilot

##### Investigated Area

High-Throughput Computing (HTC) operational environment

##### Pilot Type

Observer-only operational analysis validation

**Investigation Model**

Read-only operational synchronization analysis

**Date**

[To be completed]

**Version**

v1.0

**Confidentiality Level**

Confidential / Internal Validation Material

**Partner**

Anonymized HTC Infrastructure Environment

---

**Pilot Classification****Operational Category**

High-throughput workload execution environment

**Environment Characteristics**

- High-frequency workload submission
- Short-running compute jobs
- Centralized scheduling topology
- Multi-user execution environment
- Automated workload pipelines

**Validation Scope**

Operational synchronization and scheduling-efficiency analysis within a live HTC execution environment.

**Access Model**

- Observer-only
  - Read-only
  - Non-invasive operational telemetry analysis
  - No runtime control
  - No workflow modification
-

## **Pilot Objective**

The objective of the pilot was:

- to identify hidden operational inefficiencies inside high-throughput execution environments,
  - to analyze scheduling and queue-topology dynamics under burst-heavy workloads,
  - and to validate whether structural synchronization improvements could increase effective compute efficiency without infrastructure expansion.
- 

## **Core Investigation Focus**

The investigation specifically focused on:

- scheduler overhead amplification,
  - queue churn dynamics,
  - workload fragmentation behavior,
  - submission burst propagation,
  - hidden capacity loss,
  - and execution-flow synchronization stability.
- 

## **Operational Principle**

The pilot was conducted under strict observer-only conditions.

No:

- scheduler replacement,
- infrastructure redesign,
- runtime orchestration control,
- workload modification,
- or compute-capacity expansion was performed during the investigation.

The validation focused exclusively on:

- operational synchronization,
- scheduling coordination,
- workload-flow stabilization,

- and structural execution efficiency.
- 

### **Strategic Validation Statement**

The pilot was designed to validate the following operational hypothesis:

“In high-throughput computing systems, performance limitations are often caused not by insufficient compute capacity, but by scheduling overhead, synchronization loss, and structural operational fragmentation.”

### **0. EXECUTIVE SUMMARY**

---

#### **Objective of the Investigation**

The objective of the pilot was to analyze the operational dynamics of a high-throughput computing (HTC) environment handling large volumes of short-running workloads, with particular focus on:

- scheduler-overhead behavior,
- queue-topology instability,
- burst-wave operational phenomena,
- and the identification of hidden capacity losses.

The investigation aimed to determine whether significant operational efficiency improvements could be achieved without infrastructure expansion, exclusively through synchronization-oriented operational adjustments.

---

#### **Investigation Period**

##### **Full Observation Period**

5 weeks

##### **Validation Phase**

2.5 weeks

##### **Validation Environment**

- Live HTC operational environment
- Segmented workload validation scope
- Observer-only operational monitoring
- Read-only infrastructure access

---

## Primary Operational Problem

Although the investigated environment possessed substantial compute capacity, a significant portion of effective execution time was consumed by:

- scheduling activity,
- queue churn,
- dispatch coordination overhead,
- and structural workload fragmentation.

As a result, the environment exhibited:

- elevated operational noise,
- unstable execution flow,
- non-linear queue amplification,
- and reduced effective compute efficiency.

The investigation revealed that the dominant limiting factor was not compute availability itself, but the coordination overhead required to organize high-frequency workload execution.

---

## Key Findings

The investigation identified several recurring operational patterns inside the HTC environment:

### 1. Scheduling-Limited Operation

The environment operated as a scheduling-constrained system rather than a compute-constrained system.

Under high submission frequencies, scheduler coordination overhead became a dominant component of total operational time.

---

### 2. Coordination Noise Generated by Short Workloads

Short-running workloads generated disproportionately high coordination and orchestration overhead relative to actual compute execution time.

This resulted in:

- frequent queue reshuffling,
- unstable dispatch behavior,

- and increased operational fragmentation.
- 

### **3. Non-Linear Queue Instability**

Burst-style workload submission patterns created:

- queue shockwaves,
- scheduling saturation,
- and cascading dispatch delays.

The investigation confirmed that local submission bursts propagated across the environment as non-linear operational instability waves.

---

### **4. Significant Hidden Capacity Loss**

The environment exhibited substantial Hidden Capacity Loss (HCL) caused by:

- workload fragmentation,
- micro-idle resource gaps,
- inefficient node packing,
- and synchronization mismatch between workload arrival and resource availability.

This resulted in simultaneous:

- overload conditions,
  - and unused compute capacity within the same operational environment.
- 

### **5. Scheduler Overhead Approaching Compute Time**

During peak operational intervals, scheduler-related coordination overhead approached the magnitude of actual productive compute execution time.

This indicated that:

- the environment spent an excessive proportion of operational effort organizing workloads, rather than executing them.
- 

## **Key Validated Results**

The validation phase confirmed measurable operational improvements after synchronization-oriented adjustments were introduced.

### **Validated Improvements**

- Scheduler overhead reduction: **-25% to -45%**
  - Queue wait time reduction: **-20% to -40%**
  - Effective compute ratio increase: **+15% to +25%**
  - Node utilization improvement: **+8% to +15%**
  - Failed/delayed workload reduction: **-20% to -35%**
- 

### **Strategic Interpretation**

The validation demonstrated that:

- operational instability inside HTC environments is frequently coordination-driven rather than compute-driven,
- and that significant efficiency gains can be achieved without hardware expansion.

The improvements were achieved through:

- scheduling simplification,
  - workload synchronization,
  - queue stabilization,
  - and operational flow alignment, rather than through additional infrastructure investment.
- 

### **Core Conclusion**

“The performance of high-throughput computing systems is determined not only by compute capacity, but by the synchronization quality of scheduling dynamics and operational coordination.”

## **1. INVESTIGATION ENVIRONMENT**

---

### **Investigated System Type**

The investigation was conducted within a High-Throughput Computing (HTC)-oriented HPC environment designed for large-scale execution of short-running workloads.

The operational environment was characterized by:

- high-frequency workload submission,
- centralized scheduling coordination,
- automated execution orchestration,
- and continuous multi-user compute activity.

### **Core Environment Characteristics**

- HTC-oriented HPC cluster
  - Multi-user workload environment
  - High-frequency automated execution pipelines
- 

### **Investigation Scope**

The pilot focused specifically on the operational dynamics of high-frequency workload execution and scheduling coordination under burst-heavy operating conditions.

The investigation targeted:

- short-running workload execution behavior,
- queue-topology dynamics,
- scheduler coordination efficiency,
- and burst-submission operational effects.

### **Primary Investigation Areas**

- Short-running workload execution dynamics
  - Queue topology behavior
  - Scheduler coordination efficiency
  - Burst-submission operational effects
- 

### **Operational Environment**

The investigated environment operated as a continuously active, multi-tenant workload-processing system with highly dynamic execution behavior.

The environment exhibited:

- rapidly changing queue states,
- burst-like workload arrival patterns,

- non-linear scheduling pressure,
- and continuously fluctuating dispatch coordination demands.

### **Operational Characteristics**

- Multi-user environment
  - Batch-oriented execution model
  - Automated pipeline execution
  - Bursty workload distribution
  - Centralized scheduler topology
- 

### **Investigation Period**

The pilot was conducted in two distinct operational phases:

#### **Observation Phase**

A 5-week observer-only operational analysis period focused on identifying:

- recurring scheduling patterns,
- queue-wave dynamics,
- synchronization instability,
- and hidden operational inefficiencies.

#### **Validation Phase**

A 2.5-week controlled validation phase was performed on segmented workload groups inside the live operational environment.

During this phase:

- synchronization-oriented operational adjustments were introduced,
  - while infrastructure architecture and compute capacity remained unchanged.
- 

### **Data Sources**

The investigation relied exclusively on aggregated operational telemetry and infrastructure-level execution metrics.

The analyzed data sources included:

- queue telemetry,

- scheduler activity logs,
- execution timing statistics,
- aggregated workload telemetry,
- node utilization metrics,
- and submission timing patterns.

### **Operational Data Categories**

- Queue telemetry
- Scheduler logs
- Execution timing statistics
- Aggregated workload telemetry
- Node utilization metrics
- Submission timing patterns

The investigation did not require:

- workload payload access,
- application-level data inspection,
- or user-content analysis.

---

### **Access Model**

The pilot was conducted under a strict observer-only operational model.

The AVA-STABILIS system:

- did not control execution,
- did not modify scheduler logic,
- did not alter workloads,
- and did not perform runtime orchestration.

### **Access Characteristics**

- Observer-only
- Read-only
- Aggregated operational telemetry
- Anonymized infrastructure identifiers

- No runtime intervention

The investigation environment remained fully operational throughout the pilot, with no service interruption or operational disruption introduced by the analysis process.

## 2. BASELINE OPERATIONAL STATE

---

### Baseline Metrics

Metric	Baseline Value
Average job runtime	30–180 seconds
Queue wait time	20–120 seconds
Scheduler overhead ratio	~35–50%
Effective compute ratio	~50–65%
Job throughput variability	High
Node utilization	~71%
Failed / delayed jobs	~7–10%

---

### Baseline Operational Interpretation

At first observation, the investigated HTC environment appeared:

- highly active,
- heavily utilized,
- and continuously processing workloads at large scale.

However, deeper operational analysis revealed that a significant portion of system activity was not associated with productive computation, but with coordination-related overhead and structural execution inefficiencies.

The environment exhibited:

- substantial scheduler overhead,
  - continuous queue reshuffling behavior,
  - burst-wave congestion patterns,
  - and hidden operational fragmentation.
-

### **Scheduling-Dominated Operational Behavior**

The investigation showed that scheduling activity itself became a major operational load component inside the environment.

Under high-frequency workload submission conditions:

- the scheduler continuously recalculated execution order,
- queues were repeatedly reorganized,
- and dispatch coordination became increasingly unstable.

As a result:

- scheduling operations consumed a disproportionately large share of total operational time.
- 

### **Queue Reshuffling Dynamics**

The environment demonstrated persistent queue-topology instability caused by:

- rapidly arriving short-duration jobs,
- fluctuating execution priorities,
- and burst-driven submission behavior.

This created:

- continuous queue reordering,
- unstable execution sequencing,
- and non-linear dispatch behavior.

The queue system operated as a dynamically oscillating structure rather than a stable execution flow.

---

### **Burst-Wave Congestion Phenomena**

The workload environment displayed recurring burst-submission amplification effects.

Large volumes of short-running workloads entered the scheduling system in concentrated waves, causing:

- scheduler saturation,
- temporary dispatch collapse,
- and downstream queue propagation effects.

These workload bursts did not remain localized.

Instead,

they propagated across the environment as operational instability waves, creating:

- fluctuating wait times,
  - throughput variability,
  - and execution synchronization loss.
- 

### **Hidden Operational Fragmentation**

The analysis identified significant hidden fragmentation inside the compute environment.

Because workloads were:

- short-running,
- highly distributed,
- and rapidly scheduled,  
small idle execution gaps continuously emerged across the infrastructure.

These fragmented idle windows:

- were difficult to reuse efficiently,
- reduced effective resource packing,
- and generated Hidden Capacity Loss (HCL).

As a result,  
the environment simultaneously exhibited:

- localized overload,
  - and underutilized compute capacity.
- 

### **Coordination Overhead vs Productive Compute Time**

A major portion of total compute time was not spent on productive workload execution, but on:

- operational coordination,
- dispatch handling,
- queue management,
- and repeated scheduling recalculation.

This meant that:

- the infrastructure remained active,
- yet a significant percentage of operational effort was consumed by organizing workloads rather than executing them.

The investigation therefore concluded that:

- the dominant inefficiency mechanism was structural coordination overhead, not insufficient compute capacity itself.

### **3. IDENTIFIED OPERATIONAL PATTERNS**

The investigation revealed several recurring operational patterns that consistently contributed to execution instability, scheduling inefficiency, and hidden capacity loss inside the HTC environment.

These patterns did not appear as isolated incidents, but as interconnected structural behaviors emerging from the interaction between:

- workload submission dynamics,
- scheduler coordination,
- queue topology,
- and resource allocation timing.

The identified mechanisms collectively formed a self-amplifying operational instability field within the execution environment.

---

#### **3.1 Scheduler Overhead Dominance**

##### **Observed Phenomenon**

Due to the extremely high workload submission frequency, the scheduler entered a near-continuous recalculation state.

Instead of operating as a lightweight coordination layer, the scheduling subsystem itself became a dominant operational workload component.

The environment increasingly shifted from:

- compute-focused execution toward:
  - scheduler-dominated orchestration activity.
- 

##### **Operational Chain**

High-frequency submissions  
→ continuous scheduling activity

- overhead accumulation
  - reduction of effective compute time
- 

### **Consequence**

A substantial portion of total system operational time was consumed by:

- scheduling coordination,
- queue recalculation,
- dispatch orchestration,
- and execution-order management.

As submission intensity increased, scheduler overhead scaled disproportionately relative to productive compute execution.

This reduced:

- effective compute efficiency,
  - dispatch stability,
  - and overall workload throughput consistency.
- 

### **Interpretation**

The observed pattern indicated:

- structural coordination instability,
- scheduler-dominant operational behavior,
- and amplification of operational noise.

The investigation demonstrated that:

- the scheduler was no longer merely coordinating execution, but had become a significant consumer of operational resources itself.

This created a self-reinforcing overhead loop under high-frequency workload conditions.

---

## **3.2 Queue Churn Effect**

### **Observed Phenomenon**

The rapid queue entry and exit behavior of short-running workloads caused continuous queue-topology reorganization.

Because workloads completed quickly and new workloads continuously arrived, the queue structure remained in a near-permanent state of dynamic reshuffling.

---

### **Operational Chain**

Rapid queue movement

→ execution-order instability

→ queue reshuffling

→ scheduling inefficiency

---

### **Consequence**

The environment exhibited:

- unstable execution flow,
- fluctuating dispatch ordering,
- and high throughput variability.

Queue-state volatility reduced scheduling predictability and increased coordination complexity across the environment.

---

### **Interpretation**

The identified behavior represented:

- queue-topology instability,
- operational resonance effects,
- and coordination-driven execution noise.

Rather than behaving as a stable execution pipeline, the queue system operated as a dynamically oscillating topology under continuous structural reordering pressure.

---

## **3.3 Micro-Fragmentation of Resources**

### **Observed Phenomenon**

Short-running workloads continuously generated small idle execution gaps across compute nodes.

Because workloads terminated rapidly and asynchronously, resource allocation windows became fragmented into small unusable intervals.

---

## Operational Chain

Short jobs

→ fragmented node occupancy

→ unusable micro-gaps

→ hidden idle capacity

---

## Consequence

The environment accumulated significant hidden capacity loss despite high overall utilization metrics.

Small fragmented execution windows:

- reduced effective node packing efficiency,
  - increased idle resource accumulation,
  - and lowered practical compute availability.
- 

## Interpretation

The pattern revealed:

- Hidden Capacity Loss (HCL),
- spatial synchronization mismatch,
- and reduced packing efficiency.

The investigation demonstrated that:

- the environment simultaneously exhibited overload conditions and unused capacity, caused not by insufficient resources, but by structural fragmentation and synchronization loss.
- 

## 3.4 Submission Burst Amplification

### Observed Phenomenon

Burst-style workload submission waves repeatedly saturated the scheduling layer.

Large groups of workloads arriving within short time intervals created sudden coordination pressure spikes inside the environment.

---

## Operational Chain

Submission burst  
→ scheduler saturation  
→ delayed dispatch  
→ nonlinear queue growth

---

### **Consequence**

The system generated queue-instability waves characterized by:

- dispatch delays,
- expanding queue depth,
- fluctuating workload latency,
- and downstream propagation effects.

Local submission spikes evolved into cluster-wide operational instability patterns.

---

### **Interpretation**

The identified mechanism represented:

- burst-wave amplification,
- wave-propagation instability,
- and downstream queue propagation dynamics.

The investigation confirmed that:

- the environment behaved as a coupled operational field, where localized timing distortions propagated across scheduling and execution layers.
- 

## **3.5 Compute–Overhead Imbalance**

### **Observed Phenomenon**

The time consumed by scheduling and coordination operations approached the magnitude of actual productive compute execution time.

Operational orchestration overhead increasingly dominated the execution lifecycle.

---

### **Operational Chain**

Queue handling  
→ scheduling overhead

- dispatch complexity
  - compute-time erosion
- 

### **Consequence**

The environment spent a substantial portion of operational effort:

- organizing workloads,
- recalculating execution order,
- and coordinating dispatch activity, rather than performing productive computation.

As workload frequency increased, effective compute efficiency decreased disproportionately.

---

### **Interpretation**

The investigation identified:

- structural operational distortion,
- coordination-limited system behavior,
- and synchronization loss inside the execution environment.

The analysis demonstrated that:

- performance degradation originated primarily from organizational overhead and execution-flow instability, rather than from lack of compute resources themselves.

The HTC environment therefore behaved as a synchronization-constrained operational system rather than a purely compute-constrained infrastructure.

## **4. OPERATIONAL TOPOLOGY AND WAVE ANALYSIS**

The investigation identified that the HTC environment operated not as a static compute infrastructure, but as a dynamically coupled operational field, where local scheduling and queue disturbances propagated across the system as synchronization waves.

The operational analysis therefore focused not only on isolated performance metrics, but on:

- topology behavior,
- propagation dynamics,
- synchronization stability,

- and structural interaction patterns between workload coordination layers.
- 

### **Investigated Operational Phenomena**

The analysis focused on the following recurring operational behaviors:

#### **Queue Wave Propagation**

The environment exhibited wave-like queue expansion and contraction patterns caused by:

- burst workload arrivals,
- delayed dispatch coordination,
- and scheduler recalculation cascades.

Queue instability frequently propagated beyond local workload groups and affected cluster-wide execution flow.

---

#### **Scheduler Saturation Topology**

Under high submission density, the scheduler layer transitioned from a coordination mechanism into an operational bottleneck topology.

The investigation identified:

- scheduler overload concentration zones,
  - dispatch bottleneck accumulation,
  - and non-linear coordination pressure escalation.
- 

#### **Dispatch Synchronization Loss**

The environment demonstrated recurring synchronization mismatch between:

- workload arrival timing,
- dispatch execution timing,
- and resource availability.

This resulted in:

- delayed execution starts,
- unstable dispatch ordering,
- and fluctuating queue latency behavior.

---

### **Burst Amplification Chains**

Large workload bursts generated operational amplification chains inside the scheduling topology.

Instead of remaining localized, submission spikes propagated through:

- queue handling,
- scheduler coordination,
- and dispatch synchronization layers, creating:
- cascading instability patterns.

---

### **Node-Gap Fragmentation**

The environment continuously generated fragmented execution windows across compute nodes.

Short-running workloads created:

- micro-idle intervals,
- asynchronous node occupancy,
- and inefficient workload packing structures.

These fragmentation effects accumulated over time and contributed significantly to hidden operational inefficiency.

---

### **Latency-Wave Propagation**

Execution delays frequently propagated through the environment as latency waves.

Small localized dispatch delays evolved into:

- queue amplification,
- downstream workload slowdown,
- and execution-order instability.

The investigation identified recurring latency-wave propagation chains during peak operational periods.

---

### **Critical Operational Nodes**

The investigation identified several structurally critical coordination layers inside the HTC environment.

---

### **Central Scheduler Layer**

The scheduler operated as the primary operational synchronization node of the environment.

Under high-frequency workload conditions, this layer accumulated:

- coordination pressure,
- dispatch recalculation overhead,
- and queue-management instability.

---

### **Queue Orchestration Subsystem**

The queue orchestration layer functioned as a continuously reconfiguring operational topology.

Its behavior directly influenced:

- workload ordering,
- dispatch timing,
- throughput stability,
- and execution predictability.

---

### **Submission Burst Entry Points**

Burst-style workload arrival zones acted as instability injection points inside the system.

These entry regions frequently triggered:

- queue shockwaves,
- scheduling overload,
- and downstream dispatch amplification.

---

### **Dispatch Synchronization Paths**

Dispatch coordination paths represented the operational linkage between:

- queue state,
- scheduler decision-making,
- and node-level workload execution.

Synchronization degradation along these paths produced cluster-wide execution instability.

---

### **Wave Propagation Patterns**

The operational analysis identified several recurring instability-wave structures inside the environment.

---

### **Queue Shockwaves**

Large workload bursts generated queue-expansion waves that propagated through multiple execution layers.

These shockwaves produced:

- unstable queue depth,
  - dispatch delay spikes,
  - and fluctuating workload execution timing.
- 

### **Cascading Dispatch Delays**

Localized scheduler overload frequently triggered:

- downstream dispatch slowdown,
- delayed workload activation,
- and execution-order instability.

The propagation effect often amplified over time rather than dissipating naturally.

---

### **Scheduler Overload Propagation**

Scheduling pressure did not remain isolated to local workload groups.

Instead,  
scheduler overload propagated across:

- queue orchestration layers,
- dispatch coordination paths,
- and workload execution regions.

This produced large-scale operational synchronization instability.

---

### **Execution-Order Instability**

Continuous queue reshuffling caused unstable execution sequencing behavior.

As workload ordering changed dynamically:

- dispatch predictability decreased,
  - throughput variability increased,
  - and coordination overhead intensified.
- 

### **Synchronization Loss Points**

The investigation identified multiple structural synchronization-loss mechanisms inside the HTC environment.

---

### **Submission Timing Mismatch**

Workload arrival timing frequently diverged from:

- dispatch capacity,
- scheduler responsiveness,
- and node availability.

This mismatch generated:

- queue congestion,
  - dispatch latency,
  - and burst amplification behavior.
- 

### **Queue-Reordering Instability**

Continuous queue reorganization reduced operational stability and execution consistency.

Frequent reordering activity increased:

- coordination overhead,
  - scheduling complexity,
  - and synchronization loss.
- 

### **Scheduler Feedback Delay**

The scheduler reacted to workload-state changes with measurable operational delay.

Under burst-heavy conditions,  
the environment frequently responded:

- too late,
  - to outdated queue conditions,
  - generating additional instability.
- 

### **Resource-Packing Inefficiency**

Node allocation timing and workload fragmentation reduced effective resource utilization efficiency.

This created:

- fragmented execution topology,
  - idle compute windows,
  - and Hidden Capacity Loss accumulation.
- 

## **5. HIDDEN OPERATIONAL LOSSES**

The investigation revealed that a substantial portion of operational inefficiency remained hidden behind apparently high infrastructure utilization metrics.

Although the environment appeared heavily loaded,  
a significant amount of operational capacity was continuously lost through:

- fragmentation,
  - coordination overhead,
  - synchronization mismatch,
  - and repeated scheduling activity.
- 

### **Hidden Capacity Loss (HCL)**

The analysis identified multiple forms of hidden operational capacity erosion.

---

### **Idle Compute Gaps**

Short-running workloads continuously created:

- micro-idle execution intervals,

- fragmented resource windows,
- and partially unusable node occupancy states.

These gaps accumulated into measurable effective compute loss.

---

### **Fragmented Node Occupancy**

Node utilization remained structurally fragmented due to:

- asynchronous workload completion,
- burst-driven dispatch behavior,
- and unstable workload distribution timing.

The environment therefore exhibited simultaneous:

- overload regions,
  - and underutilized compute segments.
- 

### **Scheduling Dead-Time**

A considerable portion of operational time was consumed by:

- queue recalculation,
- workload coordination,
- dispatch orchestration,
- and scheduler recomputation cycles.

This reduced the effective proportion of productive execution time.

---

### **Dispatch Latency Loss**

Dispatch synchronization delays generated additional hidden operational inefficiency.

Delayed workload activation caused:

- queue accumulation,
  - unstable execution timing,
  - and throughput degradation.
- 

### **Re-Queueing Effects**

The investigation identified recurring re-queueing mechanisms inside the environment.

---

### **Queue Reshuffling Loops**

The queue subsystem repeatedly reorganized workload order under burst-heavy conditions.

These reshuffling loops amplified:

- coordination overhead,
  - execution instability,
  - and scheduling latency.
- 

### **Delayed Re-Dispatch Cycles**

Scheduling delays frequently triggered repeated dispatch coordination cycles.

This created:

- execution retry behavior,
  - secondary scheduling load,
  - and operational amplification effects.
- 

### **Non-Productive Operational Noise**

The HTC environment generated substantial non-productive coordination activity.

---

### **Excessive Scheduling Recomputation**

The scheduler continuously recalculated workload placement under high-frequency operational conditions.

This significantly increased:

- coordination load,
  - scheduler pressure,
  - and operational overhead accumulation.
- 

### **Coordination Overhead Amplification**

As workload density increased, coordination complexity scaled disproportionately relative to productive compute execution.

This produced:

- operational amplification loops,
  - synchronization instability,
  - and reduced execution efficiency.
- 

### **Queue Synchronization Instability**

The interaction between:

- burst workload arrivals,
- dynamic queue topology,
- and delayed scheduler response generated persistent synchronization instability inside the execution environment.

The investigation demonstrated that:

- a significant portion of operational inefficiency originated not from insufficient compute capacity, but from structural coordination distortion and synchronization loss.

## **6. MODEL-BASED OPERATIONAL ADJUSTMENTS**

The validation phase focused exclusively on operational synchronization improvements inside the HTC environment.

No:

- infrastructure expansion,
- scheduler replacement,
- hardware scaling,
- or runtime orchestration redesign was introduced during the pilot.

Instead,

the investigation applied targeted operational coordination adjustments designed to:

- reduce structural execution noise,
- stabilize workload flow,
- minimize scheduling overhead,

- and improve synchronization across execution layers.
- 

## **6.1 Job Aggregation / Batching**

### **Modification**

Smaller workloads were grouped into larger execution units in order to reduce the frequency of scheduler interaction.

Instead of processing large numbers of independently scheduled short-duration jobs, the environment executed aggregated workload groups with reduced coordination overhead.

---

### **Operational Objective**

The primary goal of aggregation was to:

- reduce scheduling-event density,
  - decrease queue-management pressure,
  - and stabilize execution sequencing.
- 

### **Expected Operational Impact**

The model predicted:

- reduced scheduler overhead,
- lower queue recalculation frequency,
- improved execution continuity,
- and more stable workload flow behavior.

### **Expected Effects**

- Reduced scheduler-overhead
  - More stable execution flow
  - Lower coordination pressure
  - Improved workload grouping efficiency
- 

## **6.2 Submission Rate Smoothing**

### **Modification**

Burst-style workload submission behavior was temporally distributed in order to reduce concentrated scheduler pressure spikes.

Instead of allowing highly synchronized submission waves, workload injection timing was smoothed across the operational timeline.

---

### **Operational Objective**

The adjustment aimed to:

- reduce instantaneous scheduler saturation,
  - minimize queue shockwave generation,
  - and stabilize dispatch timing.
- 

### **Expected Operational Impact**

The model predicted:

- lower scheduler load,
- reduced burst-driven instability,
- improved queue stability,
- and more predictable dispatch behavior.

### **Expected Effects**

- Lower scheduler load
  - Reduced burst-instability
  - Improved queue-flow consistency
  - Reduced propagation of queue shockwaves
- 

## **6.3 Lightweight Scheduling Path**

### **Modification**

A simplified dispatch coordination path was introduced for short-running workloads.

The scheduling logic for lightweight execution tasks was streamlined in order to minimize unnecessary coordination overhead.

---

### **Operational Objective**

The adjustment targeted:

- faster workload dispatch,
  - reduced scheduling recomputation,
  - and simplified queue-handling operations.
- 

### **Expected Operational Impact**

The model predicted:

- reduced scheduling latency,
- faster queue processing,
- lower dispatch overhead,
- and improved execution responsiveness.

### **Expected Effects**

- Faster queue-handling
  - Reduced scheduling latency
  - Lower dispatch coordination overhead
  - Improved short-job execution efficiency
- 

## **6.4 Resource Packing Optimization**

### **Modification**

Node-allocation logic was adjusted to improve workload packing efficiency and reduce fragmented execution gaps.

The environment was tuned to:

- minimize micro-idle intervals,
  - increase workload-density consistency,
  - and improve node occupancy continuity.
- 

### **Operational Objective**

The goal was to reduce Hidden Capacity Loss caused by:

- fragmented node utilization,

- asynchronous workload completion,
  - and unstable packing topology.
- 

### **Expected Operational Impact**

The model predicted:

- improved node utilization,
- reduced idle resource fragmentation,
- lower Hidden Capacity Loss,
- and more efficient compute occupancy behavior.

### **Expected Effects**

- Higher node utilization
  - Reduced Hidden Capacity Loss
  - Improved resource-packing efficiency
  - Lower fragmented idle capacity
- 

## **6.5 Queue Structure Simplification**

### **Modification**

Queue topology complexity was reduced in order to stabilize workload orchestration behavior.

The queue structure was simplified to minimize:

- excessive queue reshuffling,
  - unstable execution ordering,
  - and unnecessary scheduling recursion.
- 

### **Operational Objective**

The adjustment aimed to:

- stabilize execution sequencing,
  - reduce queue churn,
  - and improve workload coordination predictability.
-

## Expected Operational Impact

The model predicted:

- more stable queue topology,
- reduced queue-instability propagation,
- improved dispatch consistency,
- and lower coordination noise.

## Expected Effects

- More stable queue topology
  - Reduced queue churn
  - Lower execution-order instability
  - Improved scheduling consistency
- 

## 7. VALIDATION EXECUTION

The validation phase was conducted inside a live HTC operational environment using controlled workload segmentation and observer-only operational monitoring.

The purpose of the validation was to determine whether synchronization-oriented operational adjustments could measurably improve execution efficiency without infrastructure expansion.

---

### Validation Parameters

Parameter	Value
Duration	2.5 weeks
Scope	~40% of workload volume
Environment	Live segmented HTC environment
Access Model	Observer-only
Intervention Type	Operational-level tuning only

---

### Validation Environment

The validation was executed:

- inside a production-like HTC environment,

- under real workload conditions,
- and without interruption of ongoing operational activity.

The pilot remained:

- read-only,
  - non-invasive,
  - and operationally isolated from critical execution control layers.
- 

### **Intervention Characteristics**

The pilot introduced only operational synchronization adjustments.

No:

- infrastructure redesign,
- hardware expansion,
- scheduler replacement,
- or workload-content modification was performed.

The validation focused exclusively on:

- coordination behavior,
  - queue dynamics,
  - workload timing,
  - and execution-flow stabilization.
- 

### **Operational Adjustment Areas**

#### **Queue Management**

Queue-handling logic was tuned to reduce:

- reshuffling frequency,
  - instability propagation,
  - and dispatch oscillation behavior.
- 

#### **Submission Timing**

Submission timing patterns were adjusted to:

- smooth workload injection,
  - reduce scheduler shockwaves,
  - and improve execution continuity.
- 

### **Scheduler Path Simplification**

Scheduling coordination paths for short-running workloads were simplified to reduce:

- scheduling recomputation,
  - coordination latency,
  - and dispatch overhead.
- 

### **Resource Coordination**

Node-packing and workload distribution behavior were refined to:

- reduce fragmented occupancy,
  - improve effective utilization,
  - and minimize hidden operational inefficiency.
- 

### **Validation Principle**

The validation was intentionally designed to demonstrate that:

- operational synchronization improvements alone could produce measurable efficiency gains without requiring additional compute infrastructure.

The investigation therefore focused on:

- reducing operational noise,
- improving execution coherence,
- and stabilizing scheduling dynamics across the environment.

## **8. VALIDATED RESULTS**

The validation phase demonstrated measurable operational improvements across all major coordination and execution-efficiency indicators inside the HTC environment.

The observed improvements were achieved:

- without infrastructure expansion,
- without compute-capacity increase,
- and without scheduler replacement.

The results confirmed that significant efficiency gains could be obtained exclusively through synchronization-oriented operational adjustments.

---

### Validated Operational Results

Metric	Baseline	Validated Result
Scheduler overhead	~35–50%	22–30%
Queue wait time	20–120 s	10–65 s
Effective compute ratio	~50–65%	65–80%
Throughput variability	High	Significantly reduced
Node utilization	~71%	78–84%
Failed / delayed jobs	~7–10%	4–6%

---

### Scheduler Overhead Reduction

The validation confirmed a substantial reduction in scheduler-related operational overhead.

The implemented synchronization adjustments:

- reduced continuous scheduling recomputation,
- stabilized dispatch coordination,
- and lowered queue-management pressure.

As a result:

- scheduler overhead decreased from approximately 35–50% to 22–30%.

This demonstrated that a major portion of operational inefficiency originated from coordination-layer amplification rather than compute limitations.

---

### Queue Wait Time Improvement

The environment exhibited significantly more stable workload dispatch timing during the validation phase.

Queue wait times were reduced from:

- 20–120 seconds  
to:
- 10–65 seconds.

The reduction was primarily caused by:

- smoother submission timing,
  - reduced queue reshuffling,
  - and improved execution-flow synchronization.
- 

### **Effective Compute Ratio Increase**

The validation demonstrated a measurable increase in productive compute utilization.

The effective compute ratio improved from:

- approximately 50–65%  
to:
- 65–80%.

This indicated that:

- a substantially larger portion of operational time was spent on productive execution rather than coordination overhead.
- 

### **Throughput Stability Improvement**

The environment exhibited significantly lower throughput variability during the validation phase.

The reduction of:

- burst amplification,
  - queue shockwaves,
  - and scheduler instability  
resulted in:
  - smoother execution flow,
  - more predictable workload completion behavior,
  - and improved operational consistency.
-

## **Node Utilization Improvement**

Node utilization increased from:

- approximately 71%  
to:
- 78–84%.

The improvement was primarily linked to:

- reduced fragmentation,
- improved resource-packing efficiency,
- and lower Hidden Capacity Loss accumulation.

The environment therefore achieved:

- higher practical utilization  
without adding additional compute resources.
- 

## **Failed / Delayed Workload Reduction**

The validation also demonstrated a reduction in:

- delayed workload execution,
- dispatch instability,
- and coordination-driven execution failures.

Failed or delayed workload rates decreased from:

- approximately 7–10%  
to:
- 4–6%.

This indicated that:

- synchronization improvements increased overall operational reliability and execution stability.
- 

## **Required Visualizations**

The following visualization sets are included as mandatory operational analysis materials within the validation report.

---

## **Scheduler-Overhead Topology Map**

Visual representation of:

- scheduler coordination density,
- overload concentration zones,
- and execution-overhead topology.

Purpose:

to reveal how scheduling pressure propagated across the HTC environment.

---

### **Queue Shockwave Propagation Heatmap**

Visualization of:

- queue-wave formation,
- burst amplification behavior,
- and downstream queue-instability propagation.

Purpose:

to demonstrate how localized workload bursts evolved into large-scale operational instability.

---

### **Before / After Execution Stability Map**

Comparative visualization showing:

- execution-flow stability,
- queue coherence,
- and workload synchronization before and after the operational adjustments.

Purpose:

to illustrate the reduction of operational fragmentation and instability.

---

### **Submission Burst Amplification Visualization**

Visualization of:

- workload arrival bursts,
- scheduler saturation behavior,
- and dispatch-delay amplification chains.

Purpose:

to demonstrate the relationship between burst density and coordination instability.

---

### **Resource Fragmentation Topology**

Visual operational topology map showing:

- fragmented node occupancy,
- idle execution gaps,
- and Hidden Capacity Loss structures.

Purpose:

to reveal hidden operational inefficiencies inside apparently highly utilized infrastructure.

---

### **HTC Operational Synchronization Map**

System-wide synchronization visualization representing:

- workload flow coordination,
- scheduler timing alignment,
- queue-topology stability,
- and operational coherence across the HTC environment.

Purpose:

to provide a structural operational overview of execution synchronization behavior.

---

## **9. INTERPRETATION**

The validation confirmed that:

- the primary source of operational instability inside the HTC environment was not insufficient compute capacity,  
but:
  - coordination distortion,
  - scheduling inefficiency,
  - and synchronization loss.
- 

### **Infrastructure Was Not the Primary Limitation**

The investigation demonstrated that:

- the environment already possessed substantial compute capability, however, a significant portion of effective operational capacity was continuously lost through:
- scheduling overhead,
- queue instability,
- workload fragmentation,
- and dispatch synchronization mismatch.

The dominant inefficiency mechanism was therefore structural rather than infrastructural.

---

### **Improvements Achieved Without Infrastructure Expansion**

The measured improvements were achieved:

- without adding compute resources,
- without replacing the scheduler,
- and without modifying workload logic.

The operational gains originated exclusively from:

- synchronization-oriented adjustments,
- queue stabilization,
- workload-flow alignment,
- and coordination simplification.

This confirmed that:

- significant operational efficiency gains can emerge from reducing structural coordination noise rather than increasing raw infrastructure scale.
- 

### **Transition Toward a More Stable Execution State**

Following the operational adjustments, the environment transitioned into:

- a more stable execution-flow condition,
- lower burst-instability behavior,
- and a more compute-focused operational mode.

The environment exhibited:

- smoother queue behavior,
- reduced scheduling amplification,
- improved workload synchronization,
- and more consistent execution timing.

Operational activity became:

- less reactive,
- less oscillatory,
- and significantly more coordinated across scheduling and dispatch layers.

---

### **Structural Interpretation**

The validation demonstrated that high-throughput computing environments behave not only as compute infrastructures, but as dynamically coupled operational systems where:

- workload timing,
- scheduler coordination,
- queue topology,
- and execution synchronization jointly determine overall system performance.

The investigation therefore confirmed the central operational principle of the pilot:

“In high-throughput computing systems, performance is determined not only by compute capacity, but by the synchronization quality of operational coordination and scheduling dynamics.”

### **10. CONCLUSION**

The investigation confirmed that the operational limitations of the examined HTC environment originated primarily from coordination inefficiency and synchronization instability rather than from insufficient compute capacity.

The environment operated as a dynamically coupled execution system, where:

- scheduler behavior,
- queue topology,
- workload timing,

- and dispatch coordination jointly determined overall operational efficiency.
- 

## **Primary Findings**

### **Scheduling-Limited Environment**

The investigation demonstrated that the HTC environment was fundamentally scheduling-limited rather than compute-limited.

Under high-frequency workload conditions:

- coordination overhead,
- queue management,
- and dispatch orchestration became dominant operational factors.

The scheduler itself evolved into a major operational load component inside the environment.

---

### **Structural Overhead Amplification Caused by Short Workloads**

Short-running workloads generated disproportionately large coordination overhead relative to their productive compute duration.

This produced:

- continuous queue reshuffling,
- unstable execution ordering,
- repeated scheduling recomputation,
- and execution-flow fragmentation.

As workload density increased, operational overhead amplified non-linearly.

---

### **Global Queue Instability Generated by Burst-Wave Submission Patterns**

Burst-style workload submission patterns generated:

- queue shockwaves,
- scheduler saturation,
- and cascading dispatch-delay propagation.

The investigation confirmed that:

- local submission spikes propagated through the environment as large-scale operational instability waves.

The HTC system therefore behaved not merely as a compute platform, but as an interconnected synchronization topology.

---

### **Significant Hidden Capacity Loss**

The environment exhibited substantial Hidden Capacity Loss (HCL) caused by:

- fragmented node occupancy,
- idle execution gaps,
- unstable workload packing,
- and coordination mismatch between submission timing and resource availability.

This resulted in simultaneous:

- localized overload,
  - and underutilized compute capacity inside the same infrastructure environment.
- 

### **Primary Operational Insight**

The investigation demonstrated that the performance of HTC environments is determined not solely by raw compute capacity, but by:

- scheduling dynamics,
- workload synchronization quality,
- queue stability,
- and operational coordination efficiency.

The validation confirmed that:

- reducing structural operational noise,
  - improving synchronization,
  - and stabilizing execution topology can produce significant efficiency gains even without infrastructure expansion.
-

## **Strategic Operational Statement**

“The performance of high-throughput computing systems is determined not by how fast workloads execute individually, but by how efficiently their operation is coordinated at large scale.”

---

## **11. RECOMMENDED NEXT STEPS**

Based on the investigation results, the following operational development directions are recommended for future HTC optimization initiatives.

---

### **Large-Scale Job Aggregation**

Introduce broader workload aggregation mechanisms in order to:

- reduce scheduling-event density,
  - minimize coordination overhead,
  - and stabilize execution flow across large workload populations.
- 

### **Adaptive Submission Control**

Implement adaptive submission-timing control mechanisms capable of:

- smoothing workload bursts,
  - reducing scheduler shockwaves,
  - and stabilizing queue dynamics in real time.
- 

### **Lightweight Scheduling Frameworks**

Develop simplified scheduling paths optimized specifically for:

- short-running workloads,
- burst-heavy execution environments,
- and high-frequency dispatch conditions.

The goal is to reduce:

- scheduling recomputation,
- coordination latency,
- and orchestration overhead.

---

## **Queue-Wave Monitoring**

Introduce continuous monitoring of:

- queue-wave propagation,
- burst amplification patterns,
- and dispatch-instability chains.

This would enable earlier identification of:

- synchronization degradation,
- queue shockwaves,
- and operational overload propagation.

---

## **Continuous Synchronization Analysis**

Establish persistent synchronization analysis capabilities across:

- workload flow,
- scheduler coordination,
- dispatch timing,
- and node allocation behavior.

The objective is to maintain long-term execution coherence under dynamic operational conditions.

---

## **Operational Topology Observability**

Expand observability capabilities beyond traditional infrastructure metrics toward:

- operational topology visualization,
- synchronization mapping,
- and wave-propagation analysis.

This would improve visibility into:

- hidden coordination distortion,
  - workload-flow instability,
  - and structural execution inefficiency.
-

## **Hidden Capacity Loss Monitoring**

Introduce continuous Hidden Capacity Loss (HCL) monitoring focused on:

- fragmented node occupancy,
- idle execution gaps,
- and ineffective workload packing structures.

The goal is to reveal operational inefficiencies that remain hidden behind standard utilization metrics.

---

## **Multi-Cluster HTC Coordination Analysis**

Extend future investigations toward:

- distributed multi-cluster HTC environments,
- cross-scheduler synchronization dynamics,
- and large-scale workload propagation behavior.

This would allow analysis of:

- cluster-to-cluster instability transfer,
  - global queue-wave propagation,
  - and distributed coordination efficiency.
- 

## **12. FINAL REMARK**

“The investigation demonstrated that a substantial portion of operational inefficiency inside the HTC environment originated not from insufficient compute resources, but from scheduling coordination distortion and structural fragmentation.”

## **APPENDICES**

---

### **A. Definition of Investigation Metrics**

The following operational indicators were used during the investigation in order to quantify:

- synchronization quality,
- coordination efficiency,
- workload-flow stability,
- and hidden operational losses inside the HTC environment.

These metrics were designed specifically for observer-only operational topology analysis and synchronization-oriented infrastructure diagnostics.

---

## **CI — Coherence Index**

### **Definition**

The Coherence Index (CI) measures the operational coherence of the execution environment.

It evaluates:

- how synchronized workload execution remains,
- how stable the queue flow behaves,
- and how effectively scheduler activity aligns with productive compute execution.

The metric represents the overall structural coordination quality of the operational environment.

---

### **Primary Measurement Areas**

The CI evaluates:

- workload-execution synchronization,
  - queue-flow stability,
  - scheduler–compute coordination alignment,
  - execution continuity,
  - and operational flow consistency.
- 

### **Interpretation**

Low CI values typically indicate:

- burst-driven instability,
- queue distortion,
- execution fragmentation,
- and scheduler-overhead dominance.

In low-coherence operational states, the environment behaves as:

- an unstable synchronization topology, rather than:

- a stable compute execution system.
- 

## **DI — Delay Index**

### **Definition**

The Delay Index (DI) measures operational delay and synchronization mismatch inside the execution environment.

The metric evaluates:

- the deviation between workload submission and dispatch timing,
  - queue-wait dynamics,
  - and scheduler response latency.
- 

### **Primary Measurement Areas**

The DI measures:

- workload submission → dispatch deviation,
  - queue wait-time amplification,
  - scheduling latency behavior,
  - and coordination-response delay.
- 

### **Interpretation**

High DI values indicate:

- delayed dispatch behavior,
- queue amplification effects,
- throughput instability,
- and scheduling congestion.

Elevated DI levels generally correspond to:

- synchronization degradation,
  - increased operational friction,
  - and reduced execution responsiveness.
-

## **WPI — Wave Propagation Index**

### **Definition**

The Wave Propagation Index (WPI) measures the intensity and propagation behavior of operational instability waves inside the HTC environment.

The metric evaluates:

- how local workload bursts propagate through the infrastructure,
  - how cluster-wide instability emerges,
  - and what downstream effects are generated by queue disturbances.
- 

### **Primary Measurement Areas**

The WPI analyzes:

- queue-wave propagation,
  - burst-amplification intensity,
  - scheduler shockwave behavior,
  - and downstream latency propagation.
- 

### **Interpretation**

High WPI values indicate:

- burst-wave amplification,
- scheduler shockwaves,
- cascading latency propagation,
- and operational instability spreading across execution layers.

Elevated WPI conditions typically represent:

- non-linear synchronization behavior,
  - large-scale coordination instability,
  - and dynamic operational wave propagation.
- 

## **HCL — Hidden Capacity Loss**

### **Definition**

The Hidden Capacity Loss (HCL) metric measures operational capacity lost due to structural fragmentation and synchronization inefficiency.

The metric evaluates:

- how much compute time and node capacity remain effectively unusable,
  - and how much idle-gap accumulation and resource misalignment exist inside the environment.
- 

### **Primary Measurement Areas**

The HCL evaluates:

- fragmented node occupancy,
  - idle execution gaps,
  - workload-packing inefficiency,
  - and hidden operational dead-time.
- 

### **Interpretation**

High HCL values indicate:

- micro-fragmentation,
- simultaneous idle + overload conditions,
- reduced effective compute ratio,
- and inefficient operational resource topology.

The investigation demonstrated that:

- environments with high apparent utilization may still exhibit substantial hidden operational inefficiency through elevated HCL accumulation.
- 

## **B. Investigation Methodology**

The investigation was conducted under an observer-only operational model without external runtime intervention.

The methodology focused on:

- operational synchronization analysis,
- workload-topology behavior,

- and structural execution dynamics inside a live HTC environment.
- 

### **Core Investigation Principles**

The investigation followed the following operational principles:

- Read-only observation
- Anonymized operational telemetry
- Structural operational analysis
- Workload-topology investigation
- Queue and scheduling-dynamics analysis
- Before/after validation methodology

The methodology intentionally avoided:

- invasive runtime modification,
  - scheduler replacement,
  - workload-content inspection,
  - or infrastructure redesign.
- 

### **Investigated Operational Layers**

The investigation focused on multiple interconnected execution layers inside the HTC environment.

---

### **Scheduler Topology**

Analysis of:

- scheduler coordination structure,
  - recalculation behavior,
  - overload concentration zones,
  - and dispatch synchronization dynamics.
- 

### **Queue Dynamics**

Investigation of:

- queue reshuffling behavior,
  - queue-wave propagation,
  - burst amplification,
  - and execution-order instability.
- 

### **Submission Burst Behavior**

Analysis of:

- workload-arrival timing,
  - submission-wave amplification,
  - scheduler shockwave generation,
  - and burst-induced synchronization loss.
- 

### **Workload-Flow Synchronization**

Investigation of:

- execution continuity,
  - dispatch alignment,
  - coordination timing stability,
  - and workload-flow coherence.
- 

### **Node Utilization Dynamics**

Analysis of:

- workload packing efficiency,
  - fragmented occupancy patterns,
  - idle-gap generation,
  - and Hidden Capacity Loss accumulation.
- 

### **Dispatch Coordination**

Investigation of:

- workload dispatch timing,

- coordination overhead,
  - scheduler responsiveness,
  - and execution-latency propagation.
- 

### **Validation Approach**

The validation phase was conducted:

- on controlled workload segments,
- inside a live operational HTC environment,
- and without infrastructure expansion.

The validation intentionally focused on:

- operational synchronization tuning,
  - scheduling coordination refinement,
  - queue-management stabilization,
  - and workload-flow optimization.
- 

### **Scope of Operational Adjustments**

The investigation evaluated exclusively:

- coordination-level operational adjustments,
- scheduling refinements,
- queue-handling improvements,
- and flow-synchronization stabilization techniques.

No:

- hardware scaling,
  - infrastructure redesign,
  - scheduler replacement,
  - or workload-content modification was introduced during the pilot.
- 

### **Methodological Objective**

The methodological objective of the pilot was to validate whether:

- synchronization-oriented operational optimization alone could produce measurable execution-efficiency improvements inside large-scale HTC environments.

The investigation therefore treated the HTC infrastructure not only as a compute system, but as:

- a dynamically coupled operational coordination topology.

### **C. OBSERVER-ONLY AND DATA SECURITY MODEL**

The AVA-STABILIS pilot was conducted under a strict observer-only operational framework designed to ensure:

- non-invasive analysis,
- operational isolation,
- infrastructure safety,
- and data-security compliance throughout the investigation process.

The investigation focused exclusively on:

- operational synchronization behavior,
- structural execution dynamics,
- and aggregated infrastructure telemetry.

At no point during the pilot did the AVA-STABILIS system perform active runtime control or operational intervention.

---

### **Operational Control Limitations**

During the pilot, the AVA-STABILIS system:

- did not perform runtime execution control,
- did not interfere with scheduler operation,
- did not modify workloads,
- and did not execute automated orchestration or dispatch control functions.

The investigation environment remained:

- fully operational,
- externally controlled,

- and infrastructure-owner governed throughout the entire validation period.
- 

### **Access Model**

The pilot operated under a strictly limited operational access model.

### **Access Characteristics**

- Observer-only
- Read-only
- Aggregated operational telemetry
- Anonymized infrastructure identifiers
- Workload-content-free operational data

The system analyzed:

- execution behavior,
  - queue topology,
  - synchronization dynamics,
  - and operational timing patterns,  
without accessing:
  - workload payloads,
  - application content,
  - or user-level execution data.
- 

### **Data Security Principles**

The investigation methodology was intentionally designed to minimize operational and information-security exposure.

### **Core Security Principles**

- No model or payload access
- No workload-content analysis
- No user-level profiling
- No operational control authority
- No infrastructure modification

The pilot therefore functioned as:

- a passive operational observation layer, rather than:
  - an orchestration or execution-management system.
- 

### **Operational Safety**

The validation was conducted:

- without service interruption risk,
- inside an isolated validation scope,
- and as a controlled operational observation process.

The pilot introduced:

- no runtime instability,
- no workload execution dependency,
- and no direct operational coupling with critical infrastructure control paths.

At all times:

- the infrastructure owner retained full operational authority,
  - and the AVA-STABILIS system remained outside the execution-control layer.
- 

### **D. ANONYMIZATION STATEMENT**

All infrastructure-specific and operationally sensitive identifiers included in this report have been anonymized or aggregated.

The anonymization process was designed to ensure that:

- operational patterns could be analyzed and presented, while:
  - infrastructure reconstruction,
  - partner identification,
  - and workload attribution remained impossible.
- 

### **Anonymized Elements**

The following elements were anonymized and/or aggregated throughout the report:

- Infrastructure identifiers
- Cluster names
- Workload names
- Node identifiers
- Timestamps
- Scheduler configurations
- Partner-specific operational parameters

All presented operational metrics and visualization layers were transformed into:

- abstracted structural representations,
  - synchronization patterns,
  - and aggregated operational dynamics.
- 

### **Scope of the Report**

The report presents exclusively:

- operational patterns,
- structural execution dynamics,
- synchronization behavior,
- and validation results.

The document does not attempt to describe:

- application-level functionality,
  - workload business logic,
  - or infrastructure-specific deployment architecture.
- 

### **Excluded Information**

The document does not contain:

- business-sensitive workload information,
- user-related data,
- payload-level operational content,

- or reconstructable infrastructure-level configuration details.

No material included in the report enables:

- workload attribution,
  - user identification,
  - or operational replication of the investigated environment.
- 

## **E. VISUAL APPENDICES – OPERATIONAL TOPOLOGY AND SYNCHRONIZATION ANALYSIS**

The following visual materials present the operational patterns, scheduling topologies, and synchronization dynamics identified during the validation process inside the investigated High-Throughput Computing (HTC) environment.

The visualizations are not conventional infrastructure monitoring dashboards or standard utilization charts.

Instead,  
the presented figures represent:

- operational synchronization maps,
  - queue-topology models,
  - scheduler-overhead visualizations,
  - and wave-propagation analyses designed to reveal:
    - hidden scheduling dynamics,
    - burst-wave amplification behavior,
    - queue-instability chains,
    - dispatch synchronization loss,
    - micro-fragmentation phenomena,
    - Hidden Capacity Loss structures,
    - and real-time workload-flow coherence patterns.
- 

### **Operational Interpretation of the Visual Materials**

The investigation demonstrated that the HTC environment operated not merely as:

- a compute infrastructure,  
but rather:

- as a dynamically coupled operational field.

Within this field:

- localized scheduling distortions generated queue waves,
- dispatch-delay cascades,
- and cluster-wide synchronization instability patterns.

The visual materials therefore focus not only on:

- utilization metrics,  
but on:
  - operational interaction topology,
  - synchronization behavior,
  - and structural execution dynamics across the environment.
- 

### **Core Visualization Categories**

The visual appendices include:

- scheduler-overhead topology maps,
- queue shockwave propagation heatmaps,
- workload-flow synchronization maps,
- burst amplification visualizations,
- resource-fragmentation topology analyses,
- and latency-wave propagation structures.

These visualizations were developed specifically to reveal:

- hidden operational behavior,
  - coordination-driven inefficiency,
  - and synchronization degradation mechanisms inside large-scale HTC execution environments.
- 

### **Visualization Methodology**

All visual materials included in this report were generated:

- using observer-only methodology,

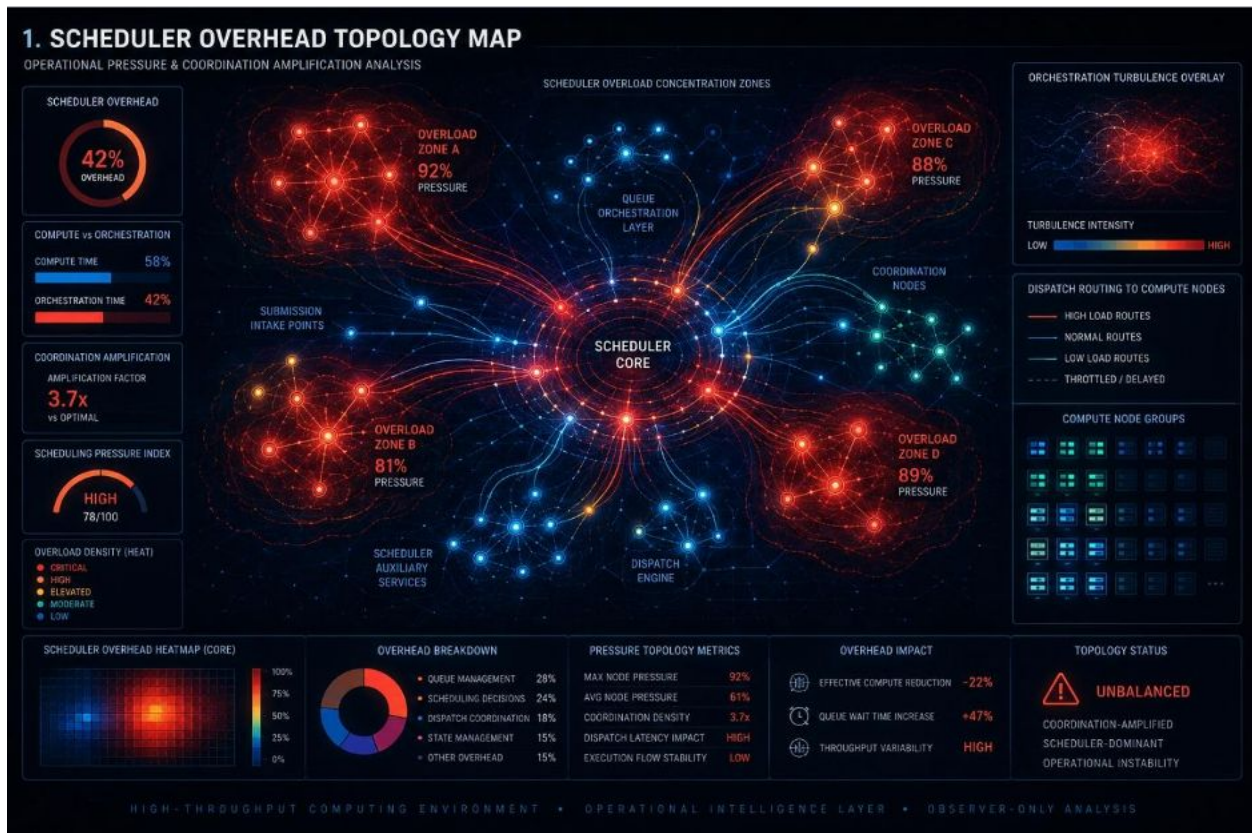
- from read-only operational telemetry,
- and within a fully anonymized environment.

The visualization process did not require:

- workload payload inspection,
- runtime orchestration access,
- or infrastructure-level operational intervention.

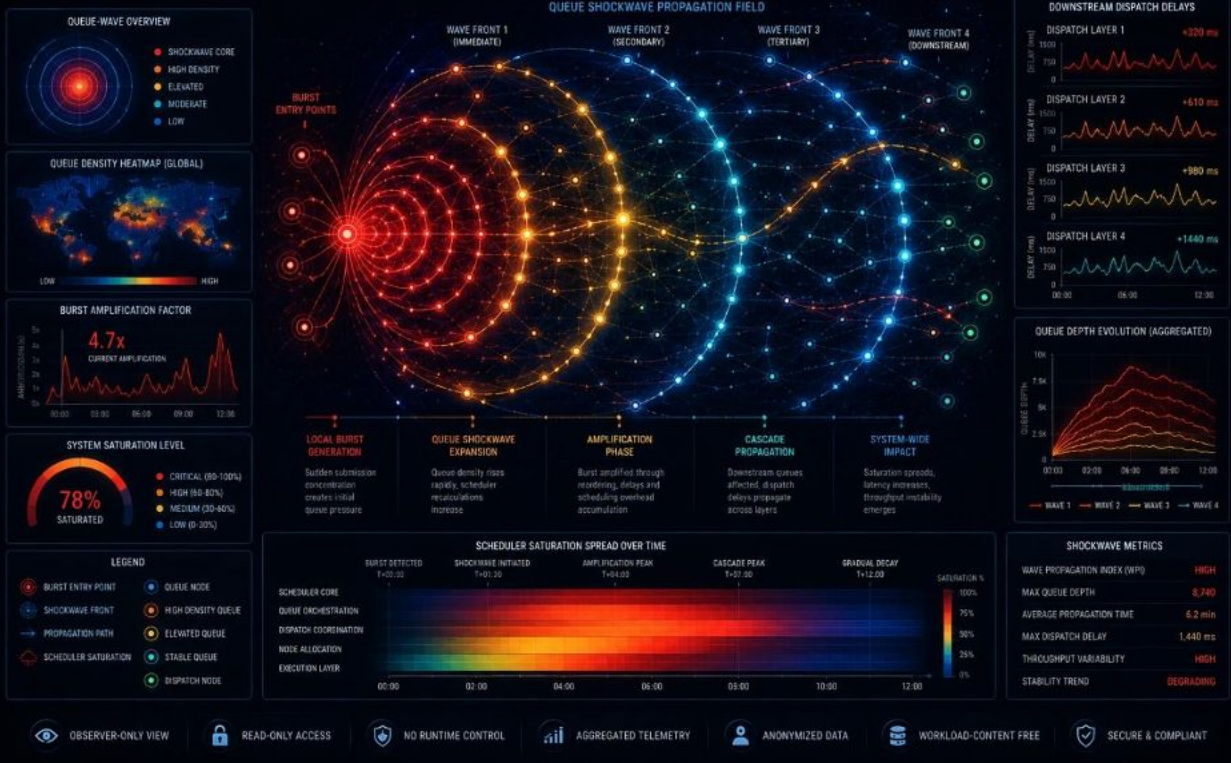
The figures therefore represent:

- structural operational behavior, rather than:
- infrastructure-control or execution-management data.



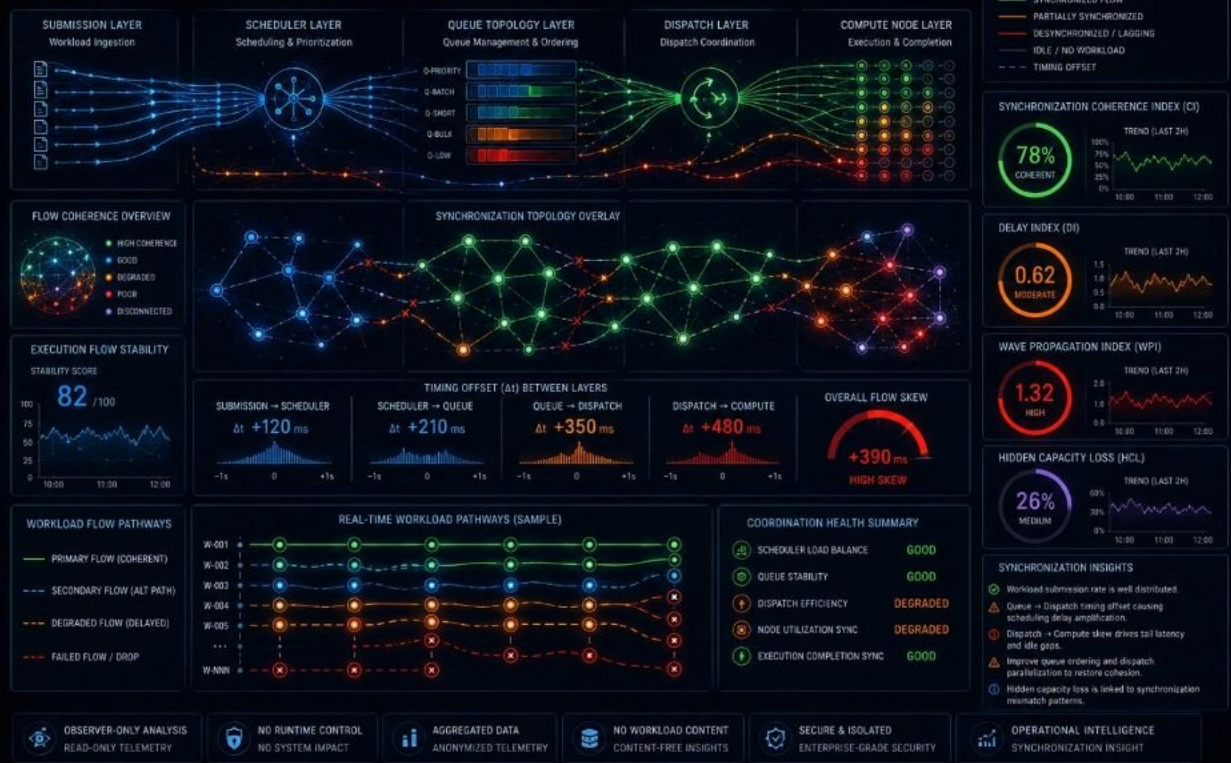
## 2. QUEUE SHOCKWAVE PROPAGATION MAP

VISUALIZING QUEUE-WAVE PROPAGATION, BURST AMPLIFICATION AND CASCADING INSTABILITY



## 3. WORKLOAD FLOW SYNCHRONIZATION MAP

VISUALIZING WORKLOAD-FLOW COHERENCE, COORDINATION SYNCHRONIZATION AND EXECUTION STABILITY



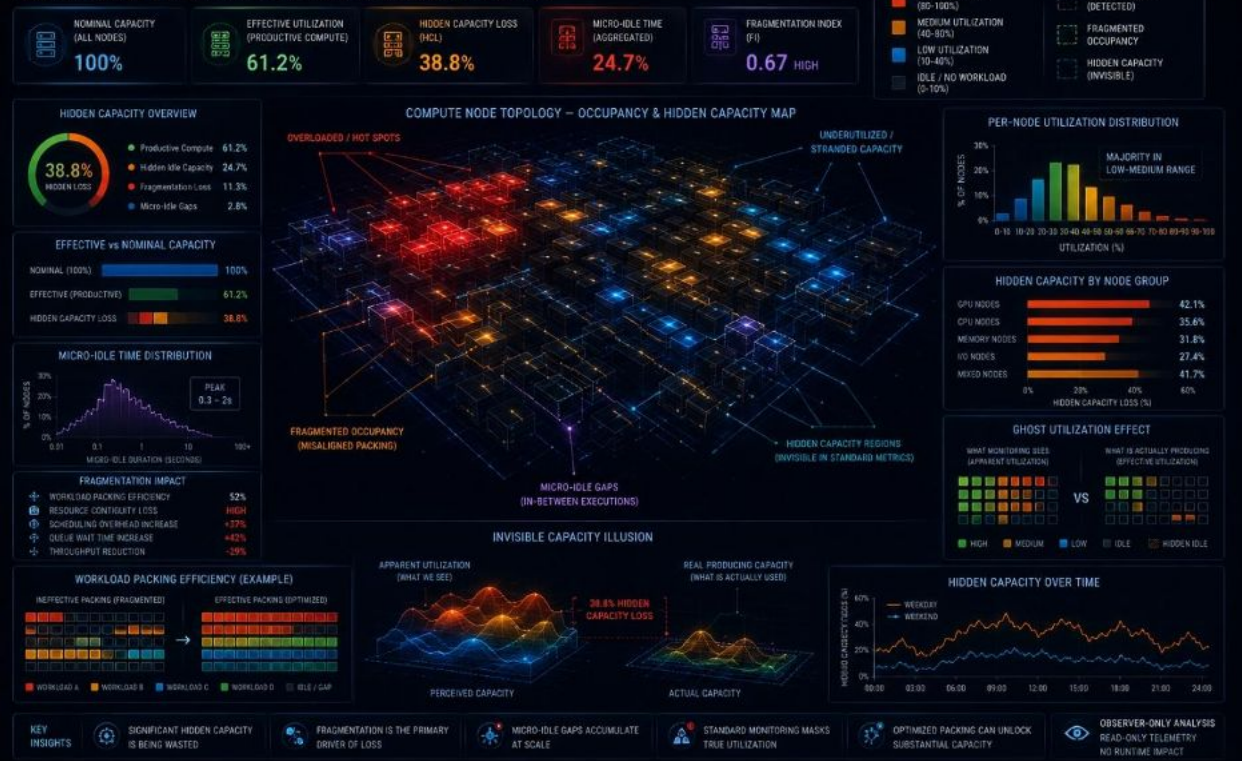
## 4. EXECUTION LATENCY TIMELINE

VISUALIZING DISPATCH DELAY ESCALATION, QUEUE COLLAPSE MOMENTS  
BURST-INDUCED LATENCY SPIKES AND OPERATIONAL RECOVERY PHASES



## 5. HIDDEN CAPACITY LOSS VISUALIZATION

UNCOVERING HIDDEN IDLE CAPACITY, FRAGMENTATION AND INEFFICIENT WORKLOAD PACKING



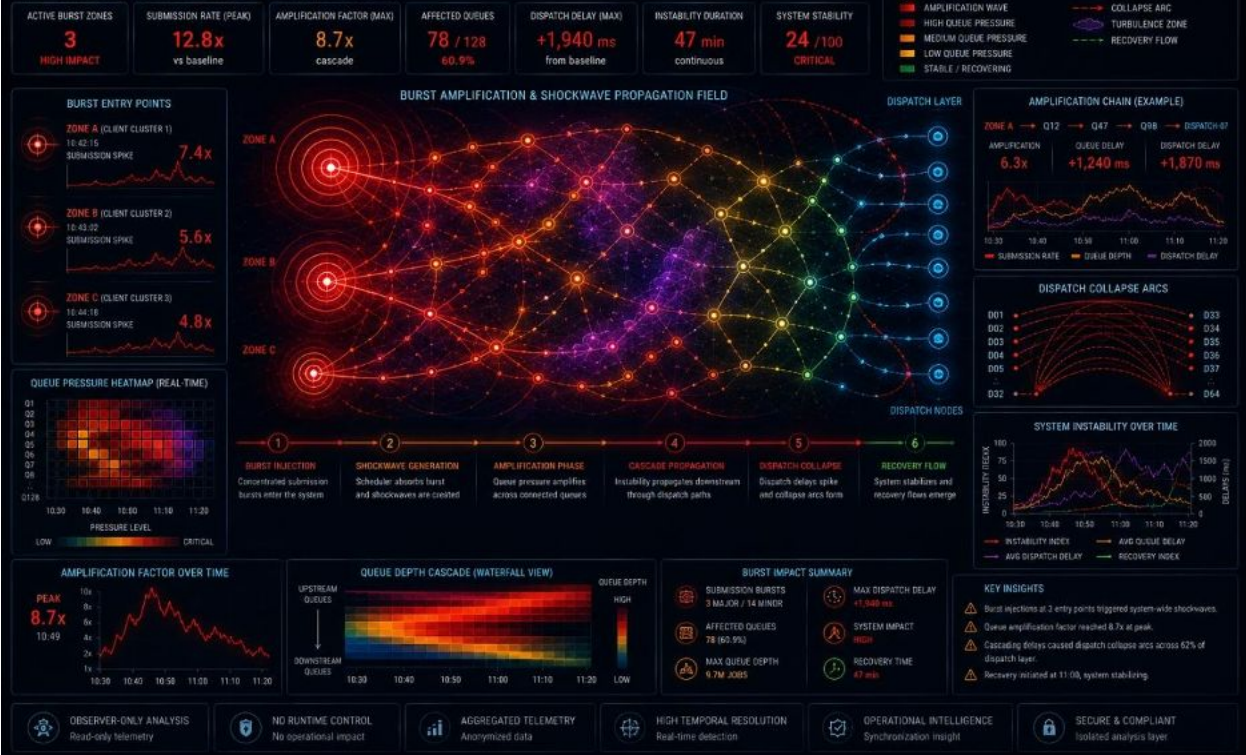
## 6. RESOURCE FRAGMENTATION TOPOLOGY

VISUALIZING NODE FRAGMENTATION, ASYNCHRONOUS WORKLOAD COMPLETION, PACKING INEFFICIENCY AND STRUCTURAL EXECUTION WASTE



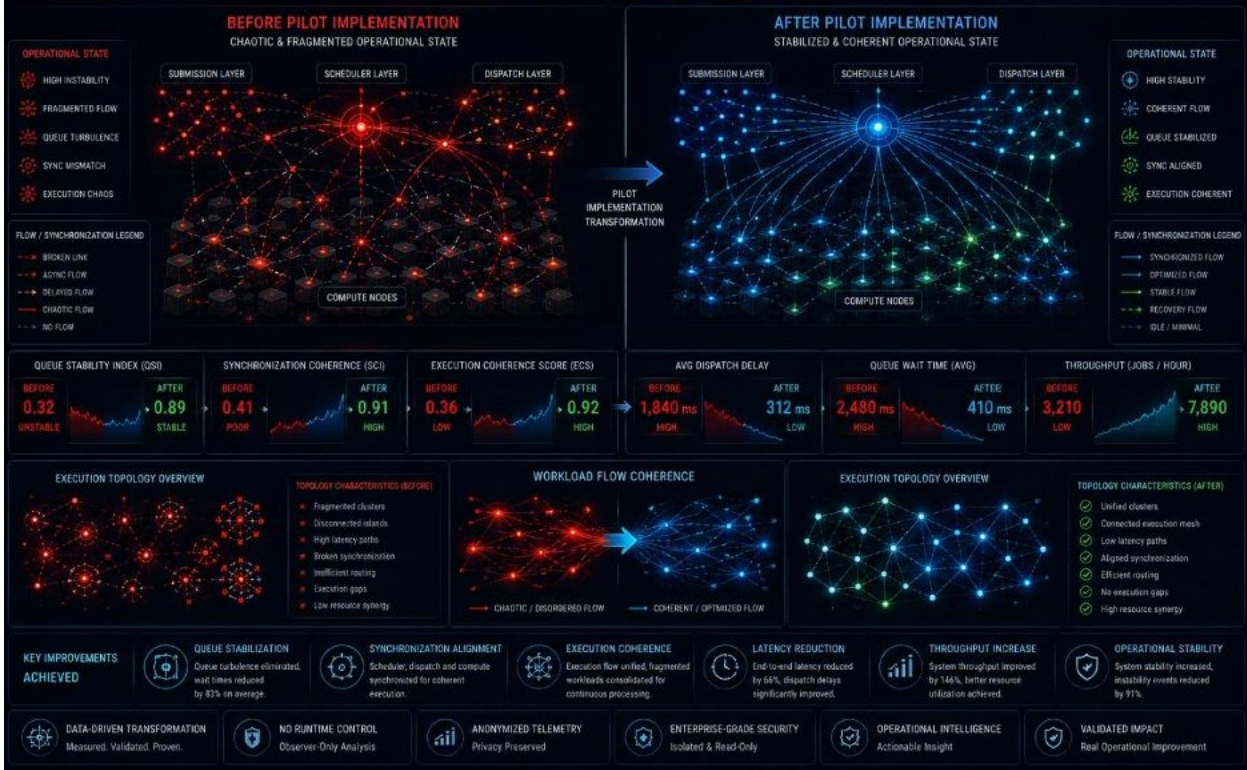
## 7. BURST SUBMISSION AMPLIFICATION MAP

VISUALIZING SUBMISSION BURST ENTRY POINTS, SHOCKWAVE GENERATION, AMPLIFICATION CHAINS AND OPERATIONAL INSTABILITY PROPAGATION



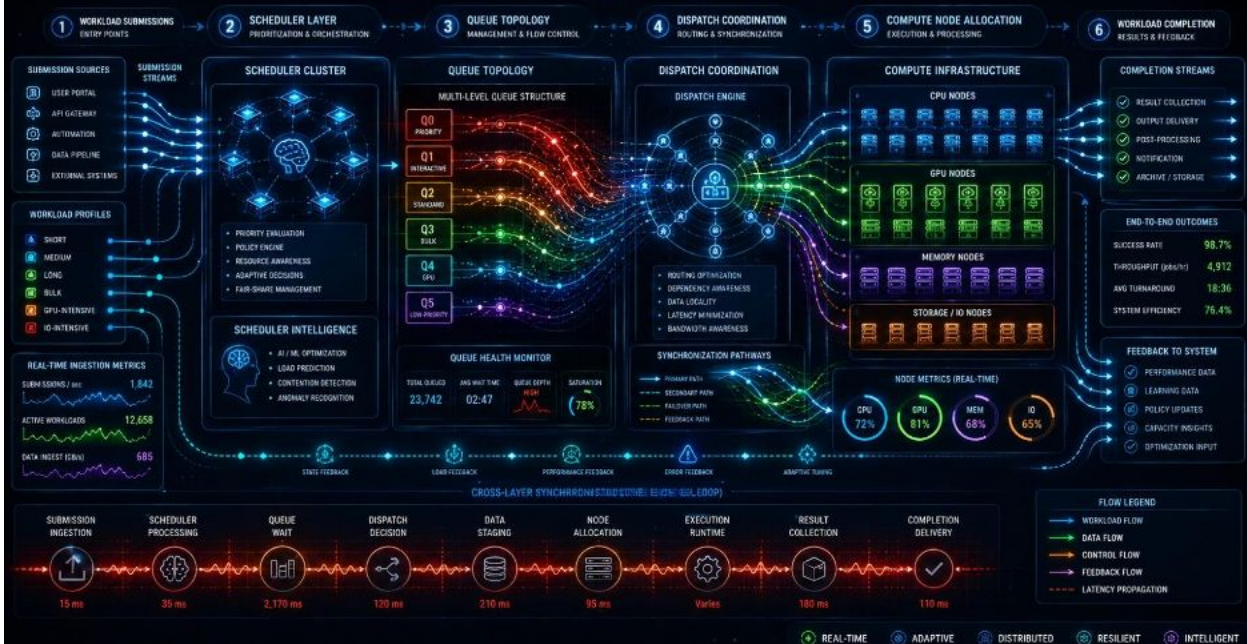
## 8. BEFORE / AFTER OPERATIONAL COHERENCE MAP

VISUALIZING OPERATIONAL TRANSFORMATION, QUEUE STABILIZATION, SYNCHRONIZATION IMPROVEMENT AND EXECUTION COHERENCE RECOVERY



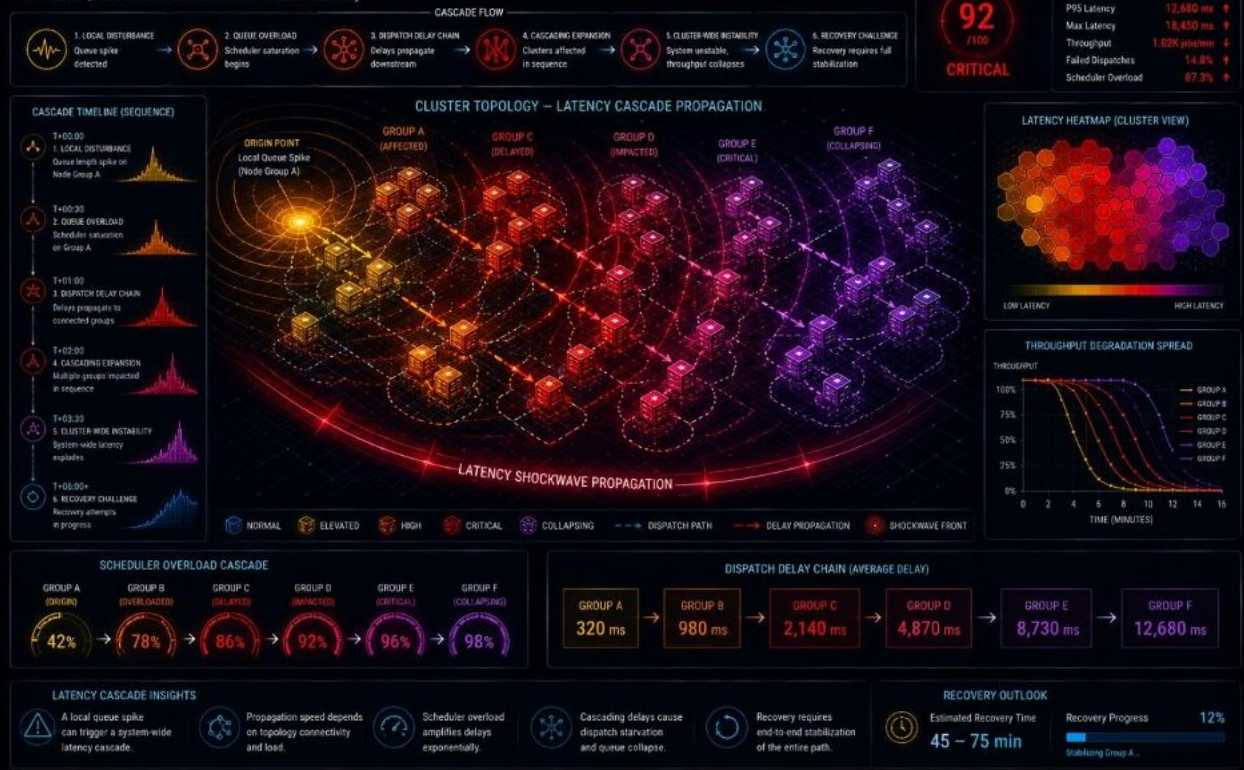
## 9. HTC OPERATIONAL TOPOLOGY DIAGRAM

END-TO-END HIGH-THROUGHPUT COMPUTING (HTC) OPERATIONAL ECOSYSTEM

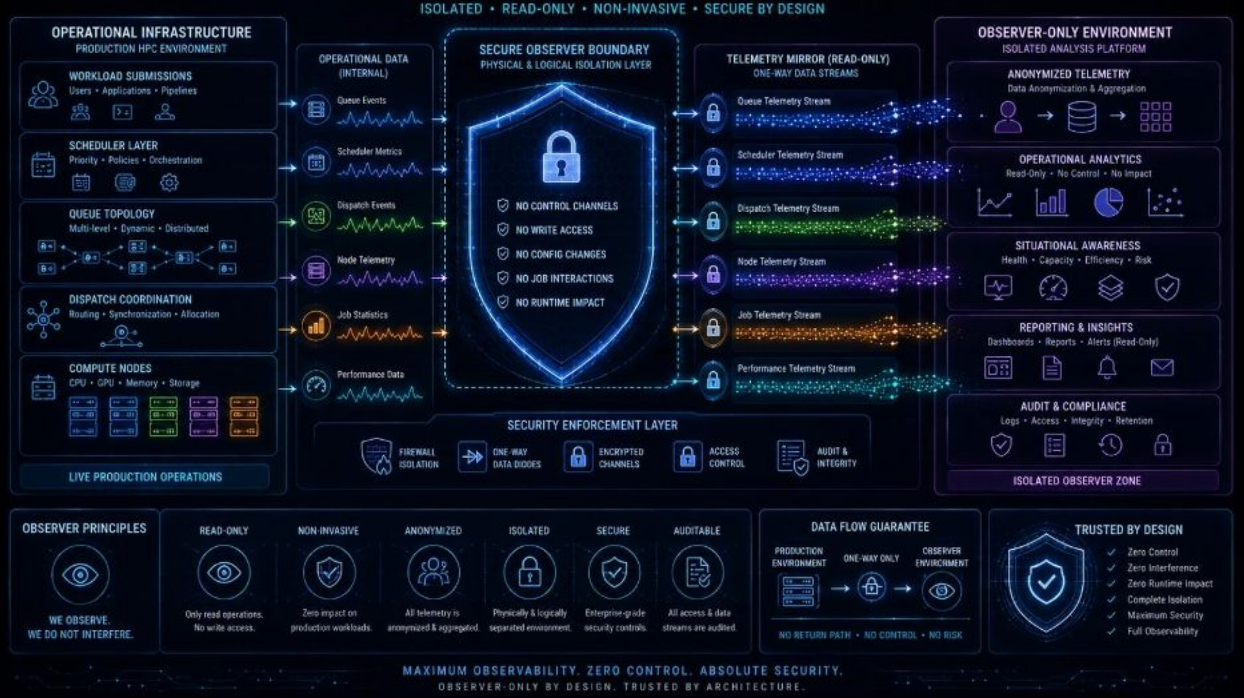


# 10. OPTIONAL EXTRA IMAGE — LATENCY CASCADE MAP

From Local Queue Disturbance to Cluster-Wide Instability



# 11. OBSERVER-ONLY SECURITY MODEL

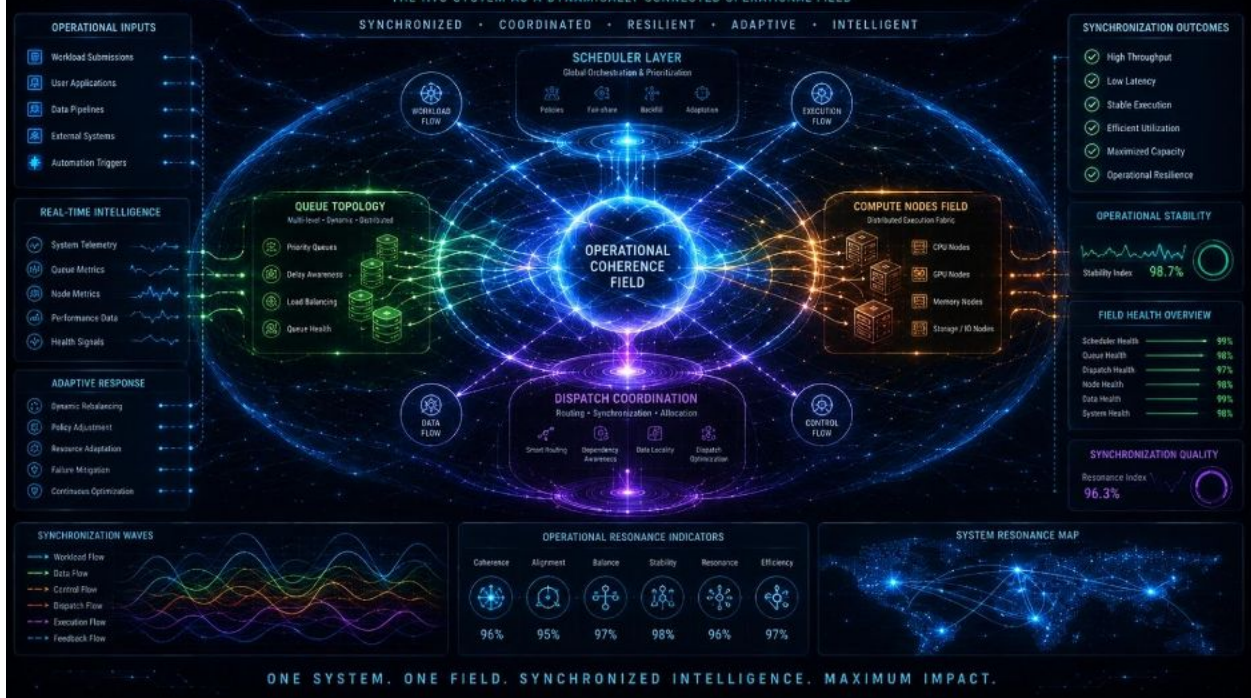


MAXIMUM OBSERVABILITY. ZERO CONTROL. ABSOLUTE SECURITY.  
OBSERVER-ONLY BY DESIGN. TRUSTED BY ARCHITECTURE.

# 12. OPTIONAL EXTRA IMAGE — HTC SYNCHRONIZATION FIELD

THE HTC SYSTEM AS A DYNAMICALLY CONNECTED OPERATIONAL FIELD

SYNCHRONIZED · COORDINATED · RESILIENT · ADAPTIVE · INTELLIGENT



ONE SYSTEM. ONE FIELD. SYNCHRONIZED INTELLIGENCE. MAXIMUM IMPACT.