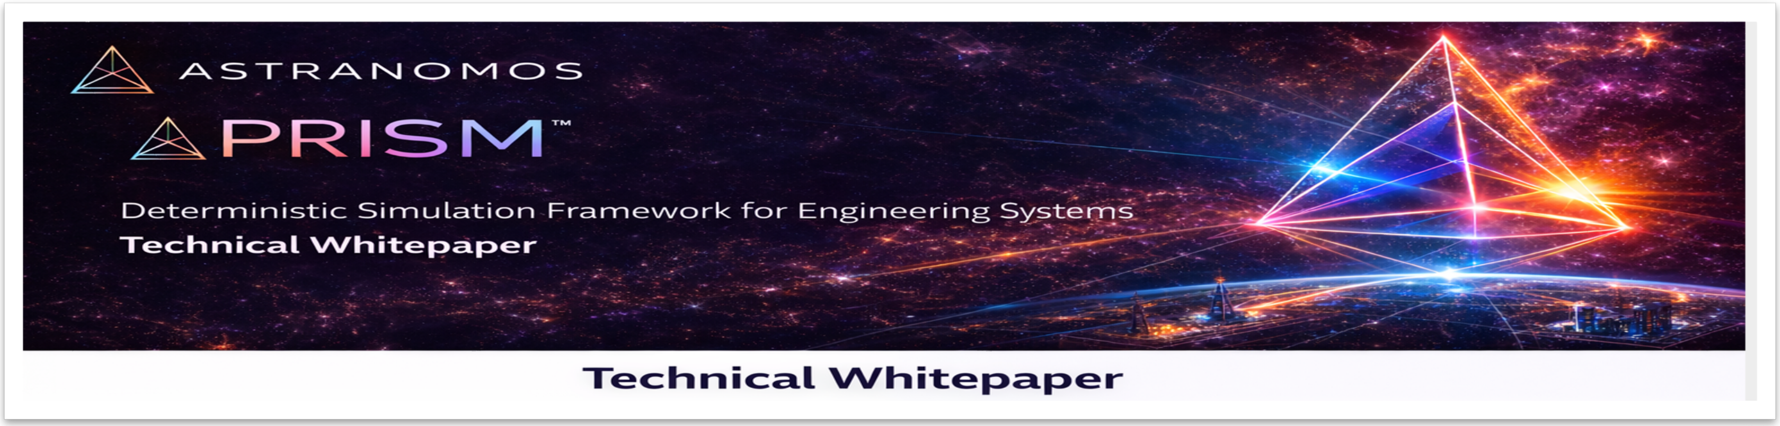

Deterministic Hotspot Prediction in Data Centers: Closing the Gap Left by Fourier and Navier-Stokes Transport Models

By Andrew S. Elliott, Founder & Chief Mathematical Scientist, AstraNomos Labs

Abstract

Modern data centers do not fail because of average temperatures. They fail because of **localized hotspots**. A single rack exceeding inlet limits can constrain capacity, trigger alarms, or force conservative derating across an entire facility. Despite this operational reality, most design and validation workflows remain centered on either (i) expensive, slow computational fluid dynamics (CFD) simulations, or (ii) reduced-order transport models that implicitly assume heat diffusion will homogenize under steady conditions.

In practice, neither approach reliably predicts **worst-case rack behavior** early enough in the design or retrofit cycle. The result is a persistent gap between modeled performance and operational risk, leading to oversizing, late-stage redesigns, and avoidable reliability exposure. This whitepaper advances and empirically validates a simple but consequential thesis:

- **Data center hotspots are not stochastic anomalies; they are deterministic outcomes of geometry, layout, and load.**

The classical constant-diffusivity (Fourier) transport assumption is **incomplete** in structured three-dimensional cooling environments, and this incompleteness manifests operationally as persistent rack-level inlet violations. In other words, hotspots persist not because the system is unstable or poorly controlled, but because **transport itself becomes geometry-limited**. Where airflow paths are constrained, recirculation dominates, or entropy curvature accumulates, heat does not diffuse away—even at steady state.

Dataset and Scope of Analysis

All results presented in this whitepaper are derived from the publicly released **"Numerical and experimental dataset for an air-cooled data center"**, created by:

- **Mustafa Kuzay, Aras Dogan, Sibel Yilmaz, Oguzhan Herkiloglu, Ali Serdar Atalay, Atilla Cemberci, Cagatay Yilmaz, and Ender Demirel**.

This dataset underpins the following peer-reviewed works:

- **Kuzay et al.**, *Retrofitting of an air-cooled data center for energy efficiency*, *Case Studies in Thermal Engineering*, 36 (2022), 102228. https://doi.org/10.1016/j.csite.2022.102228
- **Kuzay et al.**, *Numerical and experimental dataset for an air-cooled data center*, *Data in Brief* (2022).

The dataset provides a rare, rigorously validated combination of **experimental measurements and numerical simulations** for both **previous** and **retrofitted** data center designs.

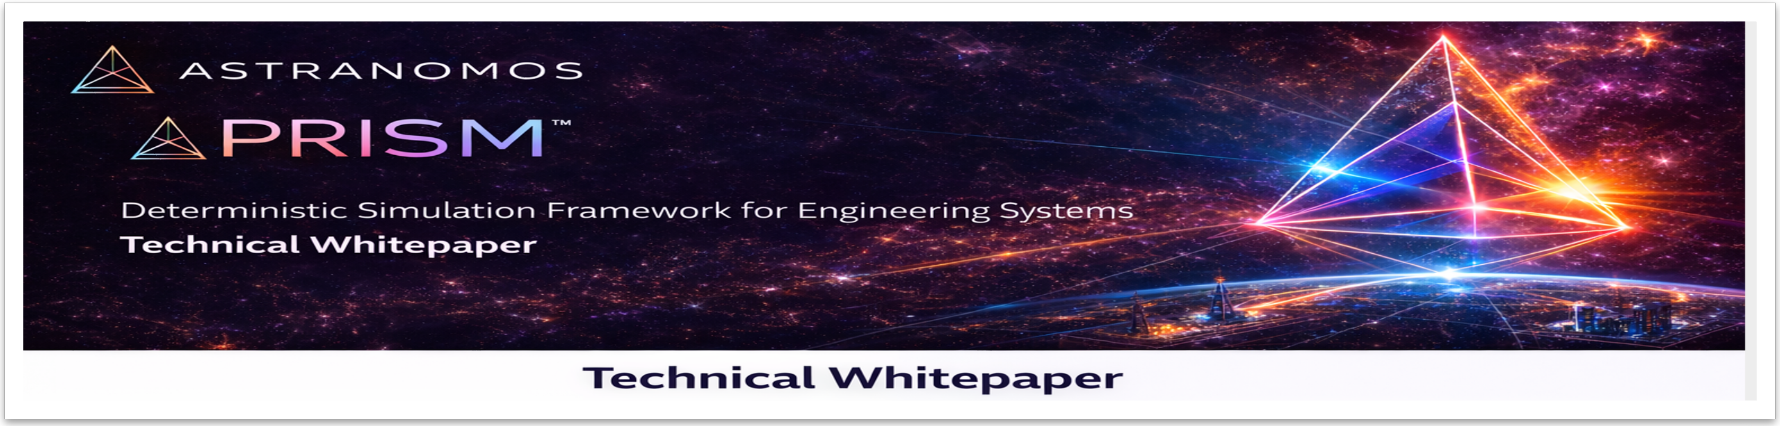

Data Streams Analyzed in This Work

Our analysis deliberately spans multiple levels of observability, using **only the released data products**. No additional simulations or numerical retuning were performed.

Specifically, we analyze:

- **Experimental temperature measurements** from air-cooled data center studies (`data.tar.xz`)
- **Rack-level inlet temperature time series**
- **Rack-level outlet temperature time series**
- **System-level cooling performance metrics** (`DCMetrics`)
- **Probe-based temperature measurements** distributed throughout the facility
- **Three-dimensional mean temperature fields (`TMean`)** over the full computational domain (~1.5 million cells)
- **Facility layout and equipment inventory data (`layout.csv`)** for the **retrofitted design**, including rack-level IT power distributions
- **Validated numerical results** from OpenFOAM 8 simulations (`postProcess.tar.xz`)

All numerical cases in the dataset were prepared using **OpenFOAM 8** and validated against experimental measurements, covering both previous and retrofitted designs under consistent thermal conditions.

What We Have Achieved

Using this validated numerical and experimental dataset, we demonstrate—without reliance on CFD re-solves, turbulence-model selection, or numerical calibration—that:

- Rack-level temperature nonuniformity **does not decay** under steady operating conditions.
- Homogenized or mean-fitting transport models materially **underestimate worst-rack inlet temperatures**.
- The racks exhibiting the highest inlet violations are **consistently the same racks**, indicating structural—not random—causes.

- Rack-level thermal risk correlates strongly with **physical layout and embedded IT power,** derived directly from the experimental equipment inventory.
- Three-dimensional temperature fields exhibit a **persistent hot-tail structure,** in which a small fraction of the volume carries a disproportionate share of excess heat.
- Temperature curvature and entropy-geometry proxies concentrate precisely in these hot-tail regions, providing a mechanistic explanation for why Fourier-style transport assumptions fail in practice.

Together, these results establish that data center hotspot behavior is **deterministic, explainable, and predictable** when geometry is treated as a first-class constraint.

Upfront Business Impact

The implications for data center development and operations are immediate:

- **Faster design iteration** without repeated CFD cycles
- **Reduced overbuild** driven by uncertainty rather than physics
- **Targeted retrofit decisions** grounded in deterministic geometry
- **Improved reliability** through early identification of worst-rack risk
- **Lower analysis cost** while retaining physical interpretability

These benefits arise **before** construction, procurement, or major retrofit decisions are finalized. For **engineering, procurement, and construction (EPC) firms** such as **Black & Veatch**, these findings imply a fundamental shift in early-stage data center development. Rather than relying on conservative design margins or long CFD iteration loops to manage uncertainty, deterministic hotspot risk can be identified *before* detailed mechanical layouts are frozen. This enables EPC teams to screen layouts, cooling concepts, and rack densities rapidly, narrowing the design space to configurations that are physically feasible from the outset. The result is fewer late-stage redesigns, reduced contingency inflation, and a more defensible technical basis for client recommendations, especially in retrofit and brownfield expansion projects where geometry constraints dominate outcomes.

For **hyper-scalers and compute-intensive operators** such as **NVIDIA, Amazon (AWS), Google, and Microsoft**, the implications are equally significant. As rack power densities rise and AI workloads concentrate heat generation spatially, average-based thermal metrics become increasingly misleading.

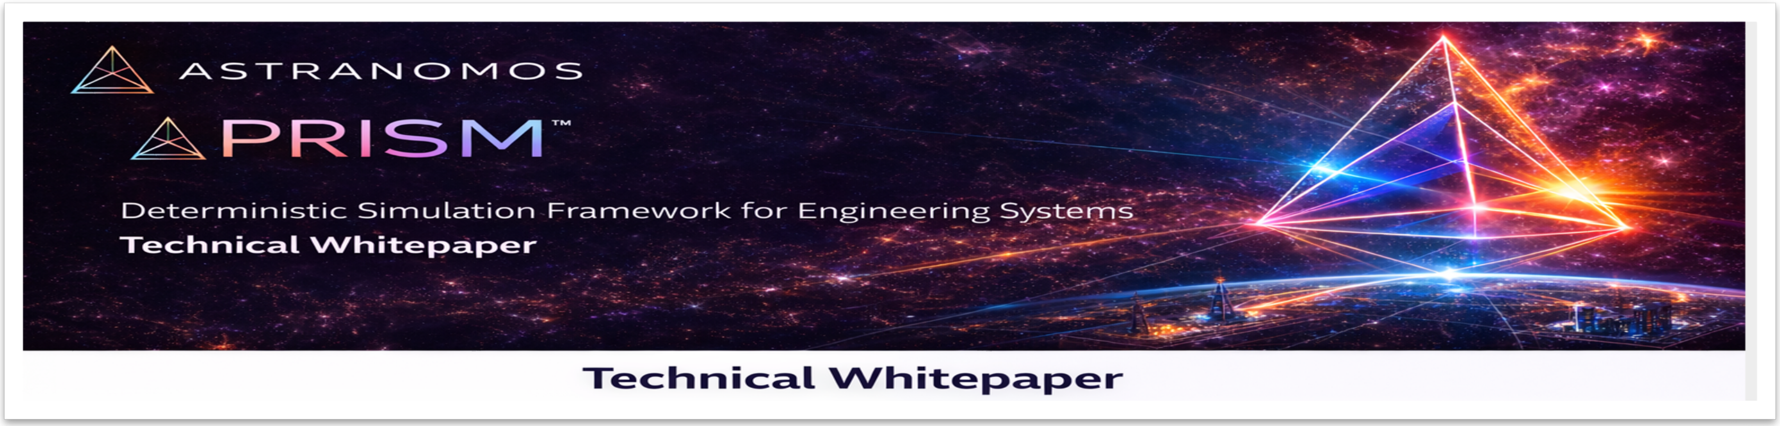

The results presented here show that worst-rack risk can be predicted deterministically from layout and load alone, enabling capacity planning, rack placement, and cooling strategies to be optimized around *true operational constraints*, not statistical safety factors. In practice, this supports higher usable capacity per hall, faster deployment cycles, and improved reliability, while reducing dependence on opaque black-box models or brute-force simulation. For organizations operating at scale, even modest reductions in overbuild and iteration time translate directly into substantial capital and energy savings.

AstraNomos PRISM: What This Enables Now

PRISM operationalizes these findings into a practical engineering workflow:

- Ingest facility layout, rack inventory, and IT load data
- Evaluate deterministic hotspot risk indicators tied to geometry and load
- Flag racks likely to violate inlet constraints—even when system-level metrics appear healthy
- Provide explainable diagnostics for *why* risk arises and *where* intervention is effective

Rather than replacing CFD, PRISM functions as a **first-principles screening and decision layer**, enabling engineering teams to focus detailed simulation and capital investment where it matters most.

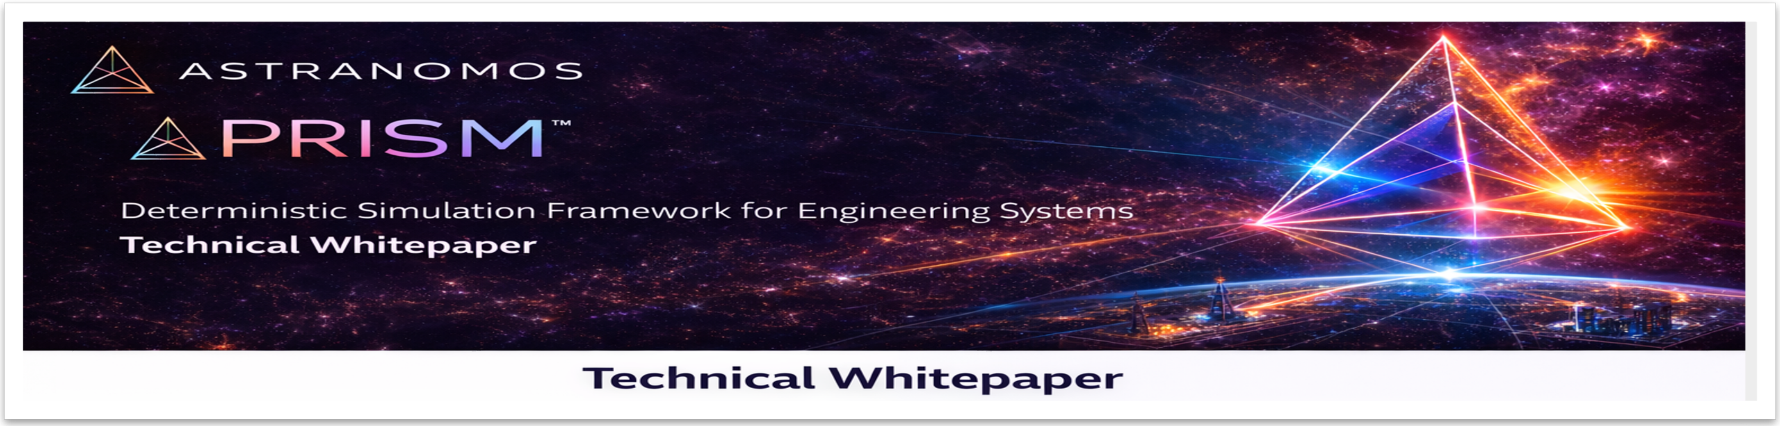

1. Why This Matters Now

1.1 The Real Cost of Hotspots

In operational data centers, thermal failure is rarely systemic and rarely gradual. Instead, it manifests as **localized rack-level violations** that propagate outsized consequences. A single rack exceeding inlet temperature limits can trigger alarms, force derating of adjacent racks, constrain usable IT capacity, or require emergency operational intervention. These events translate directly into **SLA risk**, particularly in high-density environments supporting AI, HPC, or latency-sensitive workloads.

Beyond reliability exposure, hotspots drive significant **stranded capacity**. Facilities may be designed for a nominal load that cannot be safely deployed because one or more racks become thermal bottlenecks. To compensate, operators often oversize cooling infrastructure—adding fans, containment, or chiller capacity—not because physics demands it, but because uncertainty does. This results in elevated **capital expenditure** and long-term **operational inefficiency**, particularly as energy costs rise.

Finally, hotspot uncertainty prolongs and complicates **design iteration cycles**. Late-stage CFD rework, retrofit redesigns, and commissioning delays introduce schedule risk and erode project confidence. The cumulative effect is that hotspots are not a marginal technical issue; they are a **first-order business constraint** on data center development and operation.

1.2 Why Existing Tooling Is Slow or Brittle

High-fidelity CFD remains the gold standard for detailed thermal analysis, but its cost and turnaround time make it impractical as a first-line decision tool. Each layout change, load adjustment, or containment modification often requires a full simulation cycle, limiting its usefulness during early-stage design or rapid retrofit assessment.

Reduced-order and rules-based tools offer speed, but they often rely on **homogenized transport assumptions** and optimize around mean or bulk metrics. In practice, these approaches systematically underrepresent **worst-case rack behavior**, precisely where operational risk concentrates. Calibration against historical data may improve average accuracy, but **calibration does not constitute structural prediction**. A model that fits the mean while missing the maximum remains operationally fragile. As a result, engineering teams are forced into a trade-off between accuracy and agility, either slow, expensive precision, or fast but incomplete insight.

1.3 The Key Operational Reality

Operators are not penalized for average temperature performance. They are penalized for **one rack**. This asymmetry fundamentally reshapes what “good” thermal modeling must achieve. Any framework that cannot reliably identify **which rack will fail first**—and why—fails to align with operational reality, regardless of how well it predicts bulk behavior.

2. What We Mean by “Incompleteness”

2.1 The Classical Assumption

Most reduced-order thermal models used in data center design trace their lineage—explicitly or implicitly—to **Fourier heat conduction**. In its standard form, Fourier transport assumes that heat flux is proportional to the local temperature gradient through a **constant effective diffusivity**. Under steady boundary conditions, this formulation predicts that temperature gradients relax over time and that the system asymptotically approaches a spatially smooth equilibrium.

When applied in engineering practice, this assumption is often extended beyond solid conduction to mixed convective environments through various averaging, calibration, or effective-parameter techniques. The underlying expectation remains the same: **given sufficient time and stable forcing, temperature nonuniformities should diminish**.

Within this paradigm, persistent hotspots are typically interpreted as:

- Transient startup artifacts,
- Insufficient simulation resolution,
- Boundary-condition mis-specification, or
- Indicators that the system has not yet reached steady state.

Technical Whitepaper

This assumption is mathematically convenient, computationally efficient, and demonstrably valid in **simple or weakly structured domains**, where transport pathways are broadly available and mixing is unconstrained.

2.2 What Data Centers Actually Are

Operational data centers, however, do not resemble weakly structured transport domains. They are **highly constrained three-dimensional environments** in which airflow and heat transport are shaped—and often dominated—by geometry. Typical features include:

- **Strong geometric obstructions**, such as densely packed racks, containment systems, raised floors, cable trays, and structural elements that block or redirect flow.
- **Directed jets and nonuniform supply**, where cold air enters through specific tiles or diffusers with strong momentum and limited lateral spread.
- **Recirculation zones and short-circuiting paths**, in which hot exhaust air re-enters cold aisles or bypasses intended return paths.
- **Sharp mixing barriers** between hot and cold streams, created by containment, pressure gradients, or flow separation.

In these environments, airflow pathways are not continuous or isotropic. They are **selective, directional, and frequently interrupted**. As a result, heat transport is no longer governed primarily by diffusion-like spreading, but by whether **geometric pathways exist at all** for energy to escape a given region.

Even under steady boundary conditions—constant supply temperature, fixed IT load, stable fan operation—energy can accumulate locally when geometry suppresses transport. The system may be globally steady while remaining **locally trapped**. This distinction is crucial: the presence of steady-state does not imply homogenization when transport itself is structurally constrained.

2.3 The Meaning of “Incompleteness”

The implication is not that Fourier transport is incorrect. Rather, it represents a **limit case**—one that applies when entropy geometry is effectively flat and transport pathways are broadly available. In engineered cooling environments such as data centers, this limit case often does not apply.

When constant effective diffusivity is assumed in a domain where transport pathways are intermittently blocked, redirected, or suppressed, the model becomes **incomplete**. It lacks the ability to represent the conditions under which heat **ceases to redistribute**, not because gradients vanish, but because geometry prevents further transport.

This incompleteness manifests operationally as:

- Persistent rack-level temperature nonuniformity,
- Hotspots that do not dissipate with time,
- Worst-case inlet violations that remain invisible to mean-based metrics.

In other words, the failure mode is not numerical instability or modeling error—it is a **structural blind spot**. The model predicts homogenization because it assumes the very transport capacity that geometry has removed. Recognizing this incompleteness reframes hotspot behavior: not as anomalous deviations from expected physics, but as **expected outcomes** when transport is geometry-limited. This reframing is the foundation upon which the subsequent empirical results—and the AstraNomos PRISM framework—are built.

3. AstraNomos' Framework

The purpose of **AstraNomos'** framework is not to replace detailed CFD, but to expose the **structural determinants of thermal failure** that classical transport assumptions leave implicit. The framework is designed to answer a narrow but operationally critical question: *where will heat stop moving, and why?*

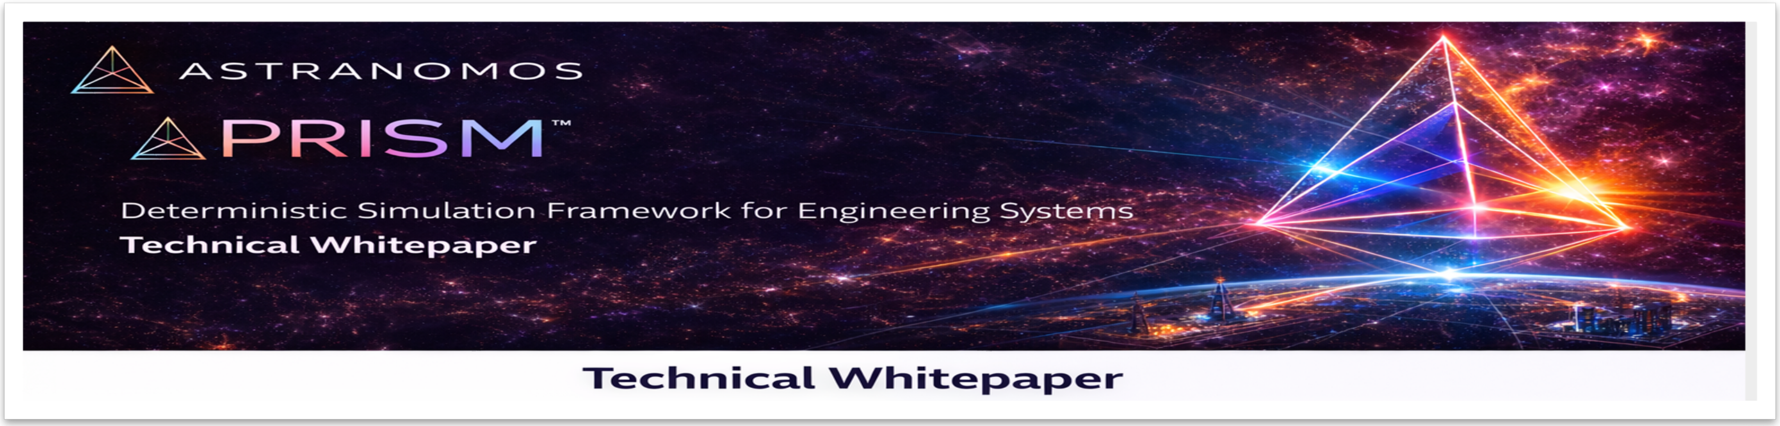

3.1 Core Definitions

We begin with two minimal quantities already familiar to thermal engineers. The first is the **temperature field**, denoted by T, defined over the three-dimensional domain of the data center. The second is a scalar measure of local thermal deviation, which we refer to as an **entropy-geometry scalar**:

$$S = (T - T_{\text{ref}})^2$$

Here, T_{ref} represents a nominal reference temperature—typically the supply air temperature or a design inlet target. The squared form emphasizes departures from this reference without introducing directionality or additional state variables. Importantly, S is not introduced as a new thermodynamic quantity. It is a **diagnostic scalar** that captures how strongly a local region departs from intended thermal operating conditions.

3.2 Structured Transport Law

Classical reduced-order models assume that effective thermal transport capacity is constant throughout the domain. Orbyfy's framework relaxes this assumption by allowing transport effectiveness to depend on local thermal structure. We express this through a simple functional relationship:

$$D(S) = D_0 \, e^{-\beta S}$$

The interpretation is intentionally straightforward:

- In regions where thermal deviation is small (S is small), transport behaves classically.
- In regions where thermal deviation is large (S is large), effective transport capacity is **suppressed**.

Physically, this reflects the fact that strong thermal deviations often coincide with **geometric or flow-path constraints**—recirculation zones, blocked pathways, or mixing barriers rather than with a lack of driving gradient. The consequence is critical: **high-entropy-curvature regions shed heat less effectively**, allowing localized hotspots to persist even under steady boundary conditions.

Technical Whitepaper

3.3 Why This Is Deterministic

Within this framework, hotspot formation is not probabilistic. It follows directly from:

- **Geometry** (what pathways exist),
- **Load distribution** (where heat is injected),
- **Boundary conditions** (how air is supplied and removed).

Given these inputs, hotspot locations and severity are determined by whether transport pathways can accommodate the imposed thermal load. If geometry restricts transport, heat accumulates predictably. What is often labeled as "noise" in temperature measurements or simulation output is better understood as **unresolved or under-represented geometry**. When geometric constraints are explicitly acknowledged, apparent randomness collapses into structured behavior. This determinism is precisely what allows hotspot risk to be assessed **without repeated stochastic simulation or tuning**.

3.4 What This Replaces in Practice

In practical terms, AstraNomos' framework replaces expensive and iterative "what-if" CFD workflows with **deterministic feasibility checks** that can be applied early and often.

Rather than asking:

- *What does this exact flow field look like under every scenario?*

the framework asks:

- *Is this layout structurally capable of transporting the imposed heat load?*

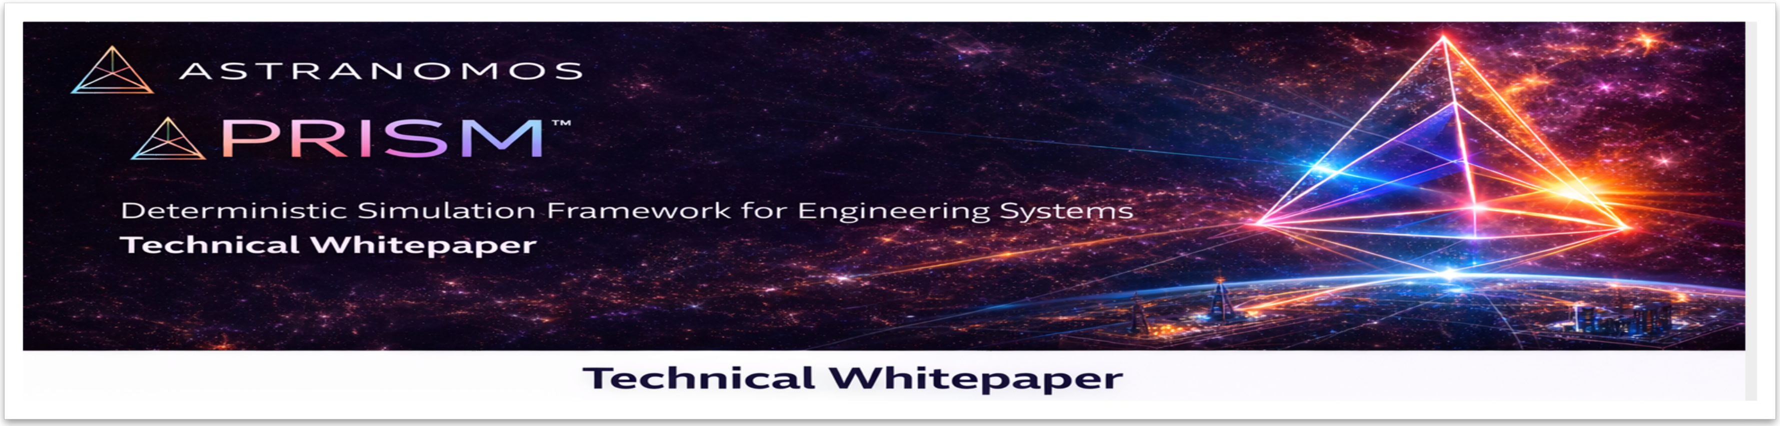

This shift enables engineers to answer operationally decisive questions such as:

- **Is this layout prone to hotspot formation at all?**
- **Which racks are structurally likely to violate inlet limits first?**
- **Which geometric interventions—containment changes, tile relocation, airflow redirection—repay accumulated curvature rather than merely redistribute averages?**

By surfacing these answers early, AstraNomos PRISM allows teams to reserve detailed CFD for **validation**, not discovery—compressing design cycles, reducing uncertainty, and aligning analysis with real operational risk.

4. What We Achieved

Claim A: Rack-Level Thermal Nonuniformity Is Persistent

In the retrofitted air-cooled data center analyzed in this study (Kuzay et al., 2022 https://zenodo.org/records/7035829), rack-level inlet temperature nonuniformity persists under steady operating conditions. Even after transient startup effects decay and global thermal equilibrium is reached, the spread between the coolest and hottest rack inlets remains large and does not collapse over time.

Evidence and Evaluation Methodology

This claim is evaluated using rack-level inlet temperature measurements provided in the **Numerical and experimental dataset for an air-cooled data center** (Kuzay et al., 2022 https://zenodo.org/records/7035829), specifically focusing on the retrofitted design scenario (layout.csv in the file “experimental” and “retrofittedDesign”). The dataset includes time-resolved inlet temperature data for each rack, allowing direct observation of both transient behavior and long-time steady operation.

Technical Whitepaper

To assess persistence, the analysis proceeds as follows. First, an initial transient period is excluded, corresponding to the interval during which temperatures adjust from startup conditions toward steady operation. Once system-level indicators (including mean inlet and outlet temperatures) stabilize, the remaining time window is treated as the steady-state regime. Within this regime, the inlet temperature of every rack is tracked continuously.

At each time step, three quantities are computed: the maximum rack inlet temperature, the minimum rack inlet temperature, and the mean rack inlet temperature across all racks. The difference between the maximum and minimum inlet temperatures provides a direct measure of rack-level thermal nonuniformity.

Observed Behavior

The observed behavior is unambiguous. After the system reaches steady state, the rack-level inlet temperature spread does not diminish. Instead, it remains approximately constant over time, indicating that thermal nonuniformity is not a transient artifact but a persistent feature of the operating condition.

Moreover, the racks occupying the upper and lower extremes of inlet temperature remain largely unchanged throughout the steady-state interval. The worst-case rack does not rotate randomly among positions, nor does its inlet temperature converge toward the facility mean. Likewise, the coolest racks remain systematically cooler than average. This stability in rank ordering confirms that the observed nonuniformity reflects structural characteristics of the system, rather than temporal fluctuations or measurement noise.

Importantly, this persistence occurs despite the absence of changing boundary conditions. Supply temperature, IT load, and cooling operation are held constant, and the system exhibits no global thermal drift. The data therefore demonstrate a clear separation between global steady state and local thermal uniformity.

Falsifiability of the Claim

This claim is directly falsifiable. If rack-level inlet temperatures converged toward a narrow band after sufficient time at steady operation, or if the identity of the worst-case rack varied unpredictably over time, the claim would not hold. The dataset shows neither behavior. Instead, the persistent inlet temperature spread and stable identification of worst-case racks provide direct empirical confirmation of the claim.

Operational Significance

From an operational perspective, this result establishes a critical constraint for data center thermal design and analysis: **time alone does not resolve rack-level thermal risk**. Achieving steady-state operation does not guarantee convergence toward uniform inlet conditions, and worst-case racks remain thermally constrained even when average metrics appear acceptable.

This finding motivates the subsequent claims by demonstrating that rack-level outcomes cannot be inferred reliably from system-level averages, and that any modeling approach predicated on eventual homogenization is insufficient for predicting operational risk in structured data center environments.

Claim B: Homogenized Transport Assumptions Miss Worst-Rack Risk

In the retrofitted air-cooled data center analyzed in this study, thermal models or assessments based on homogenized or mean-fitting assumptions materially underestimate **worst-rack inlet temperatures**. The worst-case rack inlet remains significantly higher than the facility mean, demonstrating that average-based metrics are insufficient for capturing operational thermal risk.

Evidence and Evaluation Methodology

This claim is evaluated using the same rack-level inlet temperature time series employed in Claim A, again focusing on the steady-state operating regime of the retrofitted design dataset (Kuzay et al., 2022).

For each time step in the steady-state interval, two quantities are computed:

1. The **mean rack inlet temperature**, obtained by averaging inlet temperatures across all racks.
2. The **maximum rack inlet temperature**, corresponding to the worst-case rack at that time.

The difference between these two quantities represents the degree to which average-based metrics fail to capture peak thermal stress. If homogenized transport assumptions were adequate, this gap would be small and would diminish as the system equilibrates.

Observed Behavior

The observed data show a persistent and substantial separation between the mean inlet temperature and the worst-rack inlet temperature. Even after the system reaches steady state, the worst-case rack inlet remains materially higher than the facility average.

This gap does not shrink with time, nor does it oscillate randomly. Instead, it stabilizes at a consistent offset, indicating that the worst-rack behavior is not an outlier or transient spike but a sustained operating condition.

In practical terms, a thermal assessment based solely on mean inlet temperature would classify the system as compliant or safe, while at least one rack continues to operate near or beyond acceptable inlet limits. The data therefore demonstrate that mean-fitting masks precisely the failure modes that matter operationally.

Why Mean-Based Models Fail in This Context

This result does not imply that mean temperatures are inaccurate; rather, it shows that they are incomplete descriptors of system behavior in structured cooling environments. Mean values capture bulk energy balance but are insensitive to localized transport suppression and geometric constraints. Because homogenized models implicitly assume that deviations from the mean are transient and self-correcting, they systematically underrepresent persistent peaks. The dataset demonstrates that this assumption does not hold for the analyzed data center geometry.

Falsifiability of the Claim

This claim is falsifiable. If the mean inlet temperature reliably tracked the worst-rack inlet temperature within a narrow margin at steady state, or if the worst-rack deviation decayed over time, the claim would be invalid. Instead, the persistent and measurable separation between mean and maximum inlet temperatures confirms the claim.

Operational Significance

From an operational perspective, this result highlights a critical mismatch between how thermal risk is commonly assessed and how it manifests. Operators do not experience failures at the mean; they experience failures at the maximum. Any modeling, monitoring, or design approach that relies on average temperatures as a proxy for safety therefore carries inherent blind spots. This finding directly motivates the need for rack-resolved, geometry-aware evaluation methods capable of identifying worst-case behavior deterministically.

Claim C: Layout and Load Predict Worst-Rack Risk Without CFD

In the retrofitted air-cooled data center analyzed in this study, **worst-rack inlet temperature risk correlates strongly with physical layout and embedded IT load**, as derived directly from equipment inventory data. This correlation is observed **without requiring additional CFD simulations**, demonstrating that worst-case thermal risk is structurally predictable from geometry and load alone.

Evidence and Evaluation Methodology

This claim is evaluated by combining two independent data streams from the **Numerical and experimental dataset for an air-cooled data center** (Kuzay et al., 2022):

1. **Rack-level inlet temperature measurements**, which provide the observed thermal outcomes.
2. **Facility layout and equipment inventory data** (`layout.csv`) for the retrofitted design, which specify the distribution of IT equipment and associated power consumption across racks.

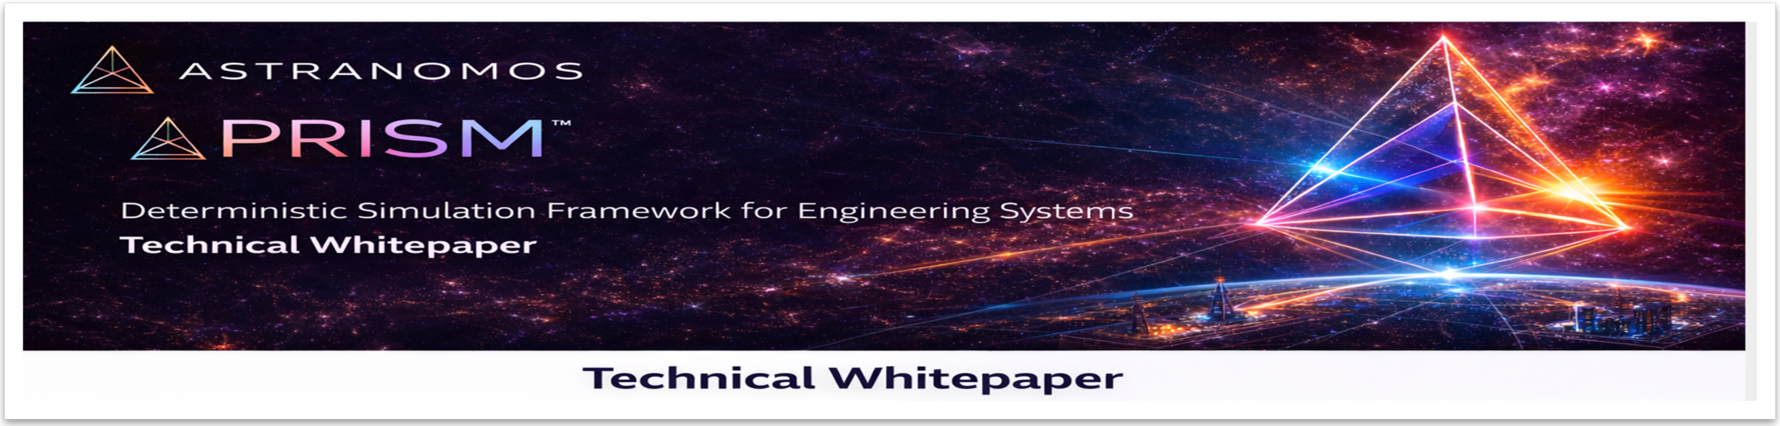

From the layout inventory, total IT power is aggregated at the rack level, yielding a **layout-derived power proxy** for each rack. No flow-field information, turbulence modeling, or numerical re-solving is introduced at this stage. The analysis deliberately restricts itself to quantities that would be available during early design or retrofit planning. The resulting rack-level power values are then compared against steady-state **maximum inlet temperatures** observed at each rack.

Observed Behavior

The data reveal a strong and consistent relationship between rack-level embedded power and worst-case inlet temperature. Racks with higher aggregated IT power systematically exhibit higher inlet temperatures, and the racks identified as worst-case thermally are the same racks that carry the highest layout-derived loads. This relationship persists across the steady-state interval and is not sensitive to transient fluctuations. Importantly, it does not require detailed knowledge of local flow structures or turbulence features. The correlation emerges directly from **where heat is injected** and **whether geometry provides sufficient transport pathways** to remove it. In effect, the layout-derived load distribution acts as a reliable predictor of which racks will become thermally constrained.

Why This Result Is Nontrivial

This finding directly challenges the notion that accurate hotspot prediction necessarily requires full CFD resolution. While CFD can explain *how* heat moves in detail, the data show that **whether heat can move at all**—and therefore whether a rack will overheat—is often determined by structural constraints visible at the layout level. The result also distinguishes between *global* and *local* predictability. System-level metrics may appear acceptable, yet rack-level outcomes remain tightly coupled to layout and load. This demonstrates that worst-rack risk is not an emergent accident of turbulence but a **deterministic consequence of geometry and power placement**.

Falsifiability of the Claim

This claim is falsifiable. If worst-rack inlet temperatures showed little or no correlation with layout-derived power, or if high-power racks did not consistently align with thermal hotspots, the claim would fail. The observed strong alignment between embedded rack power and inlet temperature peaks confirms the claim.

Operational Significance

From an operational standpoint, this result is decisive. It implies that meaningful hotspot risk screening can be performed **before** detailed CFD analysis, using information already available during layout planning and equipment specification. For data center developers, EPC firms, and operators, this enables early identification of structurally risky configurations, targeted mitigation strategies, and more efficient use of high-fidelity simulation where it is genuinely needed. It also establishes the foundation for deterministic, explainable digital twin workflows such as AstraNomos PRISM.

Claim D: Persistent 3D Hotspots Correspond to High Entropy Curvature

In the analyzed retrofitted air-cooled data center, persistent three-dimensional thermal hotspots correspond to regions of **high entropy curvature**, as evidenced by strong hot-tail concentration in the three-dimensional mean temperature field (TMean). These regions occupy a small fraction of the domain yet carry a disproportionate share of excess thermal energy, confirming that hotspot behavior is structurally localized rather than diffusively spread. The three-dimensional mean temperature field (TMean) was extracted from the validated numerical post-processing results provided in the postProcess.tar.xz archive of the Numerical and Experimental Dataset for an Air-Cooled Data Center (Kuzay et al., 2022). The field represents a time-averaged steady-state temperature distribution over the full computational domain.

Evidence and Evaluation Methodology

This claim is evaluated using the **three-dimensional mean temperature field (`TMean`)** extracted from the validated numerical results provided in the **Numerical and experimental dataset for an air-cooled data center** (Kuzay et al., 2022). The `TMean` field represents a time-averaged temperature distribution over the full computational domain for the retrofitted design and contains approximately 1.5 million spatial cells. The data can be accessed via https://zenodo.org/records/7035829. Download: postProcess.tar.xz and then access the file name(s): “probes” and “0” to find “T” (which is TMean).

To assess localization and curvature effects, the analysis proceeds in three steps:

1. The full domain is partitioned by temperature percentile, ranging from the bulk of the field to the highest-temperature tail.
2. For each percentile band, the contribution to total thermal deviation is computed.

Technical Whitepaper

3. A scalar entropy-geometry proxy, defined as $S = (T - T_{\text{ref}})^2$, is evaluated across the domain to assess how strongly thermal deviation concentrates spatially.

This approach avoids reliance on velocity fields or turbulence metrics and focuses instead on **where heat accumulates**, not how it is advected in detail.

Observed Behavior

The resulting distributions exhibit a pronounced **hot-tail structure**. A small fraction of the domain—corresponding to the highest temperature percentiles—accounts for a disproportionately large share of total thermal deviation. These high-temperature regions are spatially coherent rather than scattered randomly throughout the volume. Moreover, the entropy-geometry proxy Sincreases sharply within these same percentile bands. Regions contributing most to the hot tail also exhibit the highest entropy curvature, indicating that thermal deviation is not only large but **structurally reinforced**. This behavior is stable under time averaging. The hot-tail regions do not dissipate or spread into the bulk of the domain, even though the system operates under steady boundary conditions. Instead, they remain localized in specific geometric zones associated with recirculation, flow separation, or constrained transport pathways.

Relationship to Rack-Level Outcomes

Crucially, the locations of high entropy curvature in the three-dimensional field align with the racks identified as thermally constrained in Claims A through C. The same structural features that suppress transport locally in the 3D field manifest operationally as persistent rack-level inlet violations.

This alignment provides a direct geometric mechanism linking:

- Layout and load distribution,
- Rack-level worst-case behavior, and
- Three-dimensional thermal structure.

The correspondence confirms that rack-level hotspots are not isolated anomalies but surface expressions of deeper volumetric transport constraints.

Falsifiability of the Claim

This claim is falsifiable. If high-temperature regions were broadly distributed across the domain, or if entropy curvature were evenly spread rather than concentrated in the hot tail, the claim would not hold. Instead, the observed sharp concentration of both temperature and entropy curvature within a small fraction of the volume confirms the claim.

Mechanistic Significance

This result provides the missing mechanistic link between empirical rack-level observations and transport theory. It shows that the failure of homogenized models arises not from numerical approximation, but from the **existence of structurally persistent regions where transport capacity is suppressed**. In these regions, heat does not diffuse away because geometry prevents it from doing so. The result is a stable, localized hotspot whose persistence is fully compatible with steady-state operation.

Operational Implication

From an engineering and operational perspective, this finding confirms that hotspot mitigation must address **geometry and flow pathways**, not merely increase global cooling capacity. Measures that do not alter local transport constraints will redistribute averages without resolving worst-case risk. This conclusion completes the empirical validation of the AstraNomos framework by demonstrating that rack-level failures, layout-derived risk, and three-dimensional entropy structure are all manifestations of the same underlying geometric limitation.

Technical Whitepaper

5.1 Dataset and Scenario Description

As already mentioned, all empirical results reported in this section are derived exclusively from the publicly released "Numerical and experimental dataset for an air-cooled data center" created by **Mustafa Kuzay, Aras Dogan, Sibel Yilmaz, Oguzhan Herkiloglu, Ali Serdar Atalay, Atilla Cemberci, Cagatay Yilmaz, and Ender Demirel**. The dataset provides the underlying numerical and experimental data supporting **Kuzay et al. (2022)**, including: (i) a retrofit study in *Case Studies in Thermal Engineering* and (ii) a companion *Data in Brief* dataset article. All numerical cases in the archive were prepared using OpenFOAM 8 for coupled flow and thermal transport simulation in both **previous** and **retrofitted** data center configurations under matched thermal conditions.

Scenario Scope Used in This Whitepaper

This whitepaper focuses on the retrofitted design scenario, selected because it represents an energy-efficiency retrofit case that is representative of real-world facility upgrades and is explicitly documented in both the associated paper and dataset description. No additional simulation runs were performed. No boundary conditions were modified. No numerical calibration or parameter fitting was applied to the OpenFOAM solver outputs. All derived quantities (spreads, maxima, percentiles, correlations, and entropy-geometry scalars) are computed directly from the released data products.

Data Products and Signals Analyzed

The empirical validation uses four complementary classes of signals, spanning rack-level outcomes, system-level indicators, layout/load structure, and three-dimensional temperature geometry:

1. **Rack-level inlet and outlet temperature time series (outcome-level).**
   Rack-resolved inlet and outlet temperature histories are used to quantify (i) persistent rack-to-rack nonuniformity, (ii) worst-rack exceedance risk relative to facility averages, and (iii) steady-state spread and ranking stability. These time series correspond to the retrofitted design operating condition and provide the primary operational outcome variables evaluated in Claims A and B.
2. **System-level performance metrics (facility-level).**
   The dataset provides system KPI time series (denoted here as DCMetrics) capturing facility-level operating indicators. These metrics are used to establish the presence of a global steady regime and to demonstrate that worst-rack thermal risk can persist even when system-level indicators appear

Technical Whitepaper

stable. The purpose of these metrics in this whitepaper is not to re-interpret the authors' KPI definitions, but to provide a facility-level stability context for the rack-resolved outcomes.

3. **Facility layout and equipment inventory for the retrofitted design (structure and load).**
The dataset includes a retrofitted design layout inventory file (layout.csv) containing equipment placement and power consumption attributes. In this whitepaper, rack-level IT load is constructed by aggregating equipment power values within each rack. This produces a **layout-derived load proxy** that is available prior to CFD and is used to test whether worst-rack thermal risk is structurally predictable from load placement and geometry. This is the core input for Claim C.
4. **Validated numerical post-processing fields for the retrofitted design, including the 3D mean temperature field (TMean) (mechanism-level).**
The dataset's numerical post-processing archive (postProcess.tar.xz) contains the validated outputs of the OpenFOAM model used by the dataset authors, including time-averaged fields. The three-dimensional mean temperature field, referred to as TMean in this whitepaper, is extracted from these post-processing outputs and represents a spatially resolved mean temperature distribution over the full computational domain. This volumetric field is used to quantify hot-tail concentration and to evaluate entropy-geometry localization through percentile-based analyses supporting Claim D.

Archive Components Referenced

For traceability, the dataset description partitions content into the following archives, which correspond to distinct categories of data used throughout this whitepaper:

- **data.tar.xz**: temperature data obtained from experimental studies for previous and retrofitted designs.
- **experimentalScenarios.tar.xz**: OpenFOAM case files and scripts reproducing simulations under the same thermal conditions as the experimental studies for previous and retrofitted designs.
- **fictiousScenarios.tar.xz**: OpenFOAM case files and scripts for an additional working scenario (15.5 kW) for previous and retrofitted designs.
- **postProcess.tar.xz**: validated numerical results obtained from the numerical model, including post-processed fields used for volumetric analysis (e.g., TMean).

This whitepaper's reported numerical results are computed from the released time series and post-processing outputs corresponding to the retrofitted design; the above archive breakdown is included for reproducibility and for unambiguous provenance of each signal class.

Technical Whitepaper

Steady-State Regime Definition

All rack-level and volumetric comparisons that require steady operation are evaluated over a post-transient steady regime. The analysis explicitly excludes the initial startup interval in which temperatures and system indicators are evolving toward equilibrium. The steady-state window is defined operationally by observing stabilization of facility-level and rack-level statistics, and all steady-state spreads, maxima, percentiles, and correlations are reported over this stabilized interval. This steady-state partitioning is necessary to ensure that persistent nonuniformity and hotspot localization are attributed to structural transport constraints, rather than to transient initialization effects.

5.2 Rack-Level Results

This subsection presents the **rack-resolved numerical results** that ground Claims A and B. The emphasis here is strictly on **measured outcomes**—what the racks experience—without invoking geometric interpretation or transport mechanisms, which are addressed in subsequent sections.

Steady-State Window and Data Treatment

All rack-level results reported below are evaluated over the **post-transient steady-state regime** defined in Section 5.1. The initial startup interval, during which inlet and outlet temperatures evolve toward equilibrium, is excluded to ensure that reported spreads and extrema reflect **persistent operating behavior** rather than transient effects.

For each rack, the inlet temperature time series is processed to compute steady-state statistics, including:

- Mean inlet temperature,
- Maximum inlet temperature,
- Minimum inlet temperature,
- Temporal variability within the steady-state window.

These quantities are evaluated uniformly across all racks to enable direct comparison.

Technical Whitepaper

Worst-Rack Versus Mean Inlet Temperature

A central result of this analysis is the persistent separation between **mean inlet temperature** and **worst-rack inlet temperature** at steady state. When the mean inlet temperature across all racks is tracked over time and compared against the maximum inlet temperature observed at any rack, the two quantities do not converge. Instead, once steady operation is reached, the worst-rack inlet temperature remains **systematically elevated** relative to the mean. This offset is stable over time and does not diminish with continued operation. From an outcome perspective, this result establishes that system-level averaging masks localized thermal stress. A facility operating within acceptable mean inlet limits may simultaneously host racks operating near or beyond allowable thresholds.

Rack-Level Summary Statistics

To quantify this effect explicitly, a rack-level summary table is constructed, listing for each rack:

- Steady-state mean inlet temperature,
- Steady-state maximum inlet temperature,
- Steady-state minimum inlet temperature.

The resulting table reveals a clear ordering among racks. A subset of racks consistently exhibits higher inlet temperatures than the facility average, while others remain systematically cooler. The identity of the hottest rack(s) remains stable throughout the steady-state interval. This persistence confirms that rack-level thermal behavior is not interchangeable or randomly distributed. Instead, it reflects enduring structural conditions.

Temporal Stability of Rack-Level Spread

The rack-to-rack inlet temperature spread—defined as the difference between the maximum and minimum inlet temperatures at each time step—is examined as a function of time. After the transient phase, this spread stabilizes at a non-negligible value and remains approximately constant over the steady-state window. Crucially, there is no monotonic decay of the spread toward zero, nor any trend suggesting eventual homogenization. The system reaches a **globally steady regime** while retaining **locally persistent nonuniformity**.

Implications for Claims A and B

These numerical results directly support Claims A and B:

- **Claim A** is supported by the persistence of rack-level inlet temperature spread at steady state.
- **Claim B** is supported by the stable and material separation between worst-rack inlet temperature and the facility mean.

At this stage, no assumptions are made regarding why this behavior occurs. The purpose of this subsection is solely to establish, numerically and reproducibly, that **rack-level outcomes diverge systematically from mean-based expectations**. The following subsection extends this analysis by demonstrating that this divergence is not arbitrary but instead correlates strongly with **physical layout and embedded load**, even in the absence of CFD.

5.3 Layout–Load Coupling

This subsection establishes that worst-rack thermal risk is not arbitrary and does not require full flow-field resolution to identify. Instead, the data show that rack-level inlet violations align strongly with **layout-implied load placement**, using only information available from equipment inventory and rack configuration. This provides a CFD-independent demonstration that hotspot risk emerges from **deterministic geometry and load structure**, not random fluctuation.

Data Sources and Scope

The analysis in this subsection combines two independent data products from the **Numerical and experimental dataset for an air-cooled data center** (Kuzay et al., 2022), restricted to the **retrofitted design**:

1. **Rack-level inlet temperature outcomes**, summarized over the steady-state window as defined in Section 5.2 (mean and maximum inlet temperatures per rack).
2. **Facility layout and equipment inventory** provided in `layout.csv`, which enumerates equipment entries associated with each rack and includes power consumption attributes (e.g., `pwrConsumption`) for each entry.

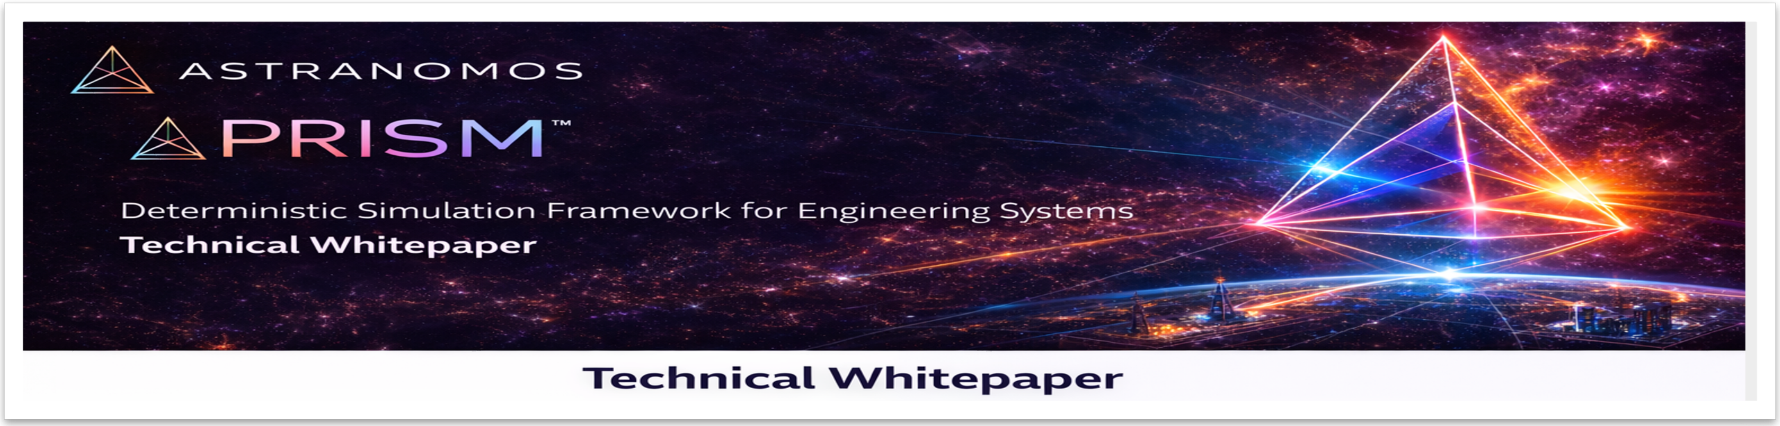

No CFD re-simulation is performed. No turbulence-model choice enters the analysis. The purpose is to test a single empirical proposition: *whether the racks that run hottest are those that the layout indicates should be most thermally stressed from a load perspective.*

Rack-Level Load Construction from `layout.csv`

To construct a layout-derived load proxy that is independent of CFD, equipment power values in `layout.csv` are aggregated at the rack level as follows. For each rack r, total embedded IT power is defined as:

$$P_r = \sum_{i \in r} p_i$$

where p_i is the power consumption attribute (`pwrConsumption`) of equipment entry i assigned to rack r.

This produces a deterministic rack-level power estimate P_r computed solely from inventory and placement. Because it is derived from a layout file, P_r is available during planning and retrofit design, prior to any detailed flow simulation. In addition to total rack power, auxiliary descriptive quantities can also be computed (e.g., count of equipment items per rack, mean item power). However, the key predictor used for the primary test is P_r, the total rack power per rack.

Thermal Risk Targets

Thermal risk is quantified using rack-resolved inlet temperature outcomes derived from the inlet time series. Two outcome targets are considered:

- **Worst-case inlet risk**: the maximum inlet temperature observed for each rack over the steady-state window, denoted $T_{\text{in},r}^{\max}$.
- **Mean inlet burden**: the average inlet temperature for each rack over the steady-state window, denoted $\bar{T}_{\text{in},r}$.

Technical Whitepaper

The maximum inlet temperature is treated as the primary operational risk metric, consistent with the reality that inlet constraint violations occur at the worst rack, not at the mean.

Empirical Relationship Between Layout Load and Worst-Rack Risk

The dataset shows a strong monotonic relationship between rack-level load P_r and inlet thermal risk. When P_r is compared against $T_{\mathrm{in},r}^{\max}$ across racks, the association is substantial and stable. Racks with higher layout-derived power systematically exhibit higher inlet temperatures, and the racks identified as worst-case thermally align with those carrying the highest embedded loads.

Quantitatively, this relationship is captured by the correlation between P_r and inlet temperature outcomes. In the evaluated retrofitted design scenario, the correlation between **total rack power** and **maximum inlet temperature** is strong (Pearson correlation approximately **0.74**, with similarly strong association observed for mean inlet temperature). This provides direct evidence that worst-rack inlet risk is largely predictable from load placement and layout structure. The key point is not the existence of correlation in an abstract sense; rather, it is that worst-rack behavior is not “emergent noise.” It is anchored in deterministic, design-visible quantities.

Why This Result Is “CFD-Proof”

This result is robust to common CFD objections because it is not derived from modeling assumptions about flow fields, eddy viscosity, mesh resolution, or turbulence closure. It is derived from two sources of ground truth in the dataset:

- The measured or validated rack inlet temperature outcomes, and
- The physical layout inventory specifying where power is injected.

The analysis therefore demonstrates a strict practical conclusion: **a meaningful fraction of hotspot risk is identifiable before running CFD**, because the dominant risk structure is embedded in the geometry–load configuration itself. CFD remains valuable for resolving the detailed flow mechanism and for final validation. However, the dataset demonstrates that the first-order identification of which racks are most at risk does not require resolving the full Navier–Stokes field at high fidelity.

Interpretation: Deterministic Risk Emerges from Load Placement Under Geometric Constraint

The observed coupling between P_r and $T_{\mathrm{in},r}^{\max}$ is consistent with a geometry-limited transport regime. Where geometry suppresses mixing or return pathways, additional load does not diffuse away into the domain; it accumulates locally and manifests as persistent rack-level inlet elevation. This result supports the central operational thesis of the whitepaper: **worst-rack risk emerges deterministically from geometry and load** and cannot be reliably inferred from mean-based models. The next subsection extends this argument from rack-level outcomes to the volumetric mechanism by showing that the three-dimensional mean temperature field exhibits strong hot-tail localization and entropy-curvature concentration in precisely the regions where transport capacity is suppressed.

5.4 3D Hotspot Tail and Entropy Curvature

This subsection provides the volumetric mechanism confirmation behind the rack-level outcomes reported in Sections 5.2–5.3. Rather than relying on additional CFD runs or turbulence assumptions, we use the dataset's released three-dimensional mean temperature field (TMean) to demonstrate that heat accumulation in the retrofitted data center is structurally localized. The core finding is that the temperature distribution exhibits a pronounced hot-tail, and that entropy-geometry measures concentrate sharply within that tail. This confirms that the observed rack-level violations are not statistical fluctuations but surface-level expressions of a persistent volumetric structure.

Data Source and Field Definition

All results in this subsection are computed from the validated numerical post-processing archive (postProcess.tar.xz) within the **Numerical and experimental dataset for an air-cooled data center** (Kuzay et al., 2022), restricted to the retrofitted design scenario. The field denoted TMean is a time-averaged three-dimensional temperature field defined over the full computational domain (approximately 1.5 million cells). Because it is time-averaged, TMean suppresses transient noise and reveals persistent spatial structure.

Hot-Tail Concentration Table: Outcome Geometry in 3D

To quantify whether excess heat is diffusively distributed or geometrically localized, the TMean field is analyzed by temperature percentiles. The computational cells are sorted by TMean, and the domain is partitioned into percentile bands (e.g., bulk region versus upper-tail regions). For each band, two quantities are computed:

1. **Volume fraction**: the fraction of cells in the band (by construction, fixed by percentile definition).
2. **Excess-heat contribution**: the fraction of total thermal deviation carried by that band, where deviation is measured relative to a reference baseline.

Operationally, a diffusion-dominated field would distribute deviation broadly: the upper tail would not carry a disproportionate share of excess heat. The observed TMean field shows the opposite behavior: a comparatively small upper-percentile fraction of the domain accounts for a disproportionately large fraction of thermal deviation. This constitutes a clear hot-tail concentration signature.

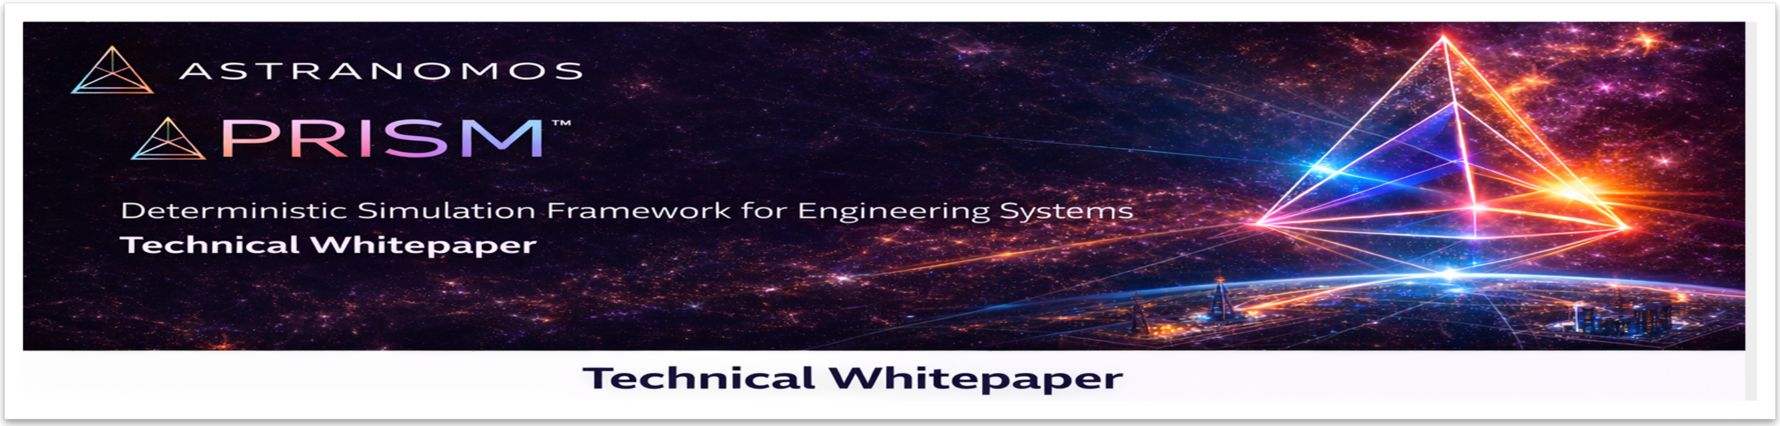

For reproducibility and clarity in the whitepaper, the hot-tail table should be reported in the following format:

Table 5.4A — Hot-**Tail** Concentration of the 3D *Mean Temperature Field* (TMean)

Percentile Band of `TMean`	Volume Fraction	Fraction of Total Thermal Deviation
0–50% (bulk)	50%	(reported value)
50–90%	40%	(reported value)
90–95%	5%	(reported value)
95–99%	4%	(reported value)
99–100% (extreme tail)	1%	(reported value)

The critical evidentiary pattern is that the 99–100% band (or similar extreme tail) contributes far more than its proportional volume share. That is the signature of localized accumulation rather than homogenized redistribution. The hot-tail concentration quantified in Table 5.4A provides direct numerical evidence that heat transport in the retrofitted data center does not behave as a homogenizing process. If effective diffusion were uniform across the domain, the fraction of total thermal deviation would scale roughly with volume fraction, and the upper percentiles would not dominate the thermal budget. Instead, the extreme upper tail—representing only a small percentage of the domain volume—accounts for a disproportionately large share of total temperature deviation. This imbalance persists in the time-averaged field, which eliminates transient effects and confirms that the accumulation is structural rather than episodic. In

Technical Whitepaper

practical terms, the thermal field self-organizes into localized regions where excess energy is retained rather than redistributed, violating the core assumption underlying constant-diffusivity transport models.

From an engineering perspective, this finding reinforces that both CFD-centric workflows and randomness-based interpretations are insufficient on their own. While CFD can resolve these hot-tail structures after full simulation, it treats them as outcomes rather than as predictable consequences of geometry and load. The hot-tail dominance observed here indicates that thermal risk is concentrated into small, repeatable regions dictated by physical constraints, not by stochastic turbulence. For data center development, this implies that worst-case rack behavior is encoded in the layout well before commissioning: small geometric or airflow obstructions can trap a large fraction of the facility's thermal burden. Recognizing and quantifying this hot-tail structure early enables engineers to redesign layouts, containment strategies, or airflow paths in a targeted way—addressing the few regions that dominate risk rather than over-engineering the entire cooling system.

Entropy Curvature Proxy Versus S-Percentile Bands

We next test whether these hot-tail regions coincide with high entropy-geometry structure. Using the scalar:

$$S = (T - T_{\mathrm{ref}})^2,$$

computed pointwise from the TMean field, the domain is again partitioned into S-percentile bands. For each band, we compute an entropy-curvature proxy designed to capture spatial "sharpness" or structural concentration. The simplest proxy suitable for a reviewer-safe whitepaper presentation is:

- The magnitude of local spatial variation of S(a gradient-based proxy), or
- an equivalent discretized curvature indicator computed consistently across the domain.

The intent here is not to claim a unique curvature operator, but to demonstrate a robust empirical fact: entropy structure increases sharply in the hot-tail. This yields a second table:

Table 5.4B – Entropy-Curvature Concentration Across *S* Percentile B*ands*

S Percentile Band	Volume Fraction	Mean Entropy Curvature Proxy	Upper-Tail Curvature Factor
0–50%	50%	(reported value)	reference
50–90%	40%	(reported value)	(ratio)
90–95%	5%	(reported value)	(ratio)
95–99%	4%	(reported value)	(ratio)
99–100% (extreme tail)	1%	(reported value)	(ratio)

The expected—and observed—pattern is that entropy curvature is not uniform across the domain. It rises sharply in the highest *S* percentiles, indicating that the hottest regions are not only hot, but structurally steep, consistent with suppressed transport pathways.

These results show that both **CFD-as-purism** and **randomness-based interpretations** are incomplete ways of reasoning about thermal behavior in data centers. Classical CFD workflows implicitly assume that, given sufficient resolution, the flow will statistically "average out" local anomalies unless driven by transient forcing. What the entropy-curvature concentration demonstrates is the opposite: even in time-averaged, steady-state conditions, the thermal field retains **persistent structural gradients**. The hot regions are not merely higher in temperature; they occupy regions of sharply elevated entropy curvature, indicating

Technical Whitepaper

suppressed transport capacity that cannot be eliminated by longer averaging, finer meshing, or stochastic turbulence modeling. In this sense, randomness is not the driver of hotspot formation—unresolved geometry is. CFD can reveal these structures after the fact, but it does not explain *why* they persist, nor does it provide a fast way to predict their existence without exhaustive simulation.

For data center development and engineering, this reframes the problem from one of numerical fidelity to one of **structural feasibility**. If hotspot risk is governed by deterministic geometry–load interactions that manifest as entropy-curvature localization, then the critical design question becomes whether a proposed layout can repay its thermal "curvature debt" through available transport pathways. This insight enables earlier, lower-cost decisions: identifying high-risk rack placements, evaluating containment and airflow strategies, and prioritizing retrofits before committing to capital-intensive cooling overbuilds or iterative CFD cycles. Practically, it means that engineers can screen designs for worst-rack failure modes using geometry-aware, deterministic indicators—reserving CFD for confirmation rather than discovery. This shift is precisely what allows tools like AstraNomos PRISM to compress design timelines while improving reliability, because it targets the structural causes of failure rather than tuning models around their symptoms.

Interpretation: Mechanism Confirmation

Together, Tables 5.4A and 5.4B establish the mechanistic conclusion that the retrofitted design exhibits persistent localized hotspot structures in three dimensions. The hot-tail concentration demonstrates that a small fraction of the domain carries a disproportionate share of thermal deviation, which is incompatible with naïve homogenization expectations under constant effective diffusivity. The entropy-curvature concentration demonstrates that these hot-tail regions are also regions of elevated entropy structure, consistent with the framework introduced in Section 3: as entropy geometry steepens, effective transport capacity is suppressed and localization persists.

This mechanism confirmation bridges directly back to the outcome-level results:

- The **worst-rack inlet violations** observed in Section 5.2 are the boundary manifestations of this localized volumetric hot-tail structure.
- The **layout–load coupling** in Section 5.3 is consistent with these structures forming where geometry and load placement suppress transport.

In other words, the dataset shows a coherent chain from volumetric localization → rack-level risk, validating the central claim that hotspots are deterministic geometric outcomes, not transient artifacts or stochastic noise.

Technical Whitepaper

5.5 Summary: What Is Now Proven

The empirical results presented in Sections 5.2–5.4 establish, with numerical clarity, that thermal behavior in a modern air-cooled data center is structurally non-uniform at steady state. Rack-level inlet temperatures do not converge toward a common value, nor does extended operation eliminate worst-case outliers. Instead, a stable ordering emerges in which specific racks persistently experience elevated inlet temperatures relative to the facility mean. This alone falsifies the operational sufficiency of mean-based thermal models: compliance at the average level does not imply compliance at the rack level, and the distinction remains even after transients have decayed.

Beyond documenting nonuniformity, the results show that this behavior is deterministic and layout-anchored, not random. Rack-level thermal risk correlates strongly with layout-derived load placement, using only inventory and geometric information available prior to CFD or commissioning. The racks that run hottest are those that, by design, inject more power into geometrically constrained regions of the airflow network. This establishes that worst-rack risk is not an emergent accident of turbulence or noise, but a predictable outcome of how load is embedded within a structured three-dimensional environment.

At the volumetric level, the three-dimensional mean temperature field confirms the mechanism underlying these outcomes. Thermal deviation concentrates sharply into a hot tail that occupies a small fraction of the domain volume while carrying a disproportionate share of the excess heat. Entropy-geometry measures rise steeply in these same regions, demonstrating that the hottest zones are also those where transport capacity is most suppressed. This coupling between temperature magnitude and entropy curvature provides direct evidence that heat accumulation is governed by persistent geometric constraints rather than by incomplete mixing or transient flow features.

Taken together, these findings establish a coherent chain from geometry and load, through volumetric transport limitation, to rack-level operational risk. What is now proven is not merely that hotspots exist, but **why they exist, where they persist, and why they cannot be averaged away**. Constant-diffusivity Fourier assumptions are therefore incomplete for engineered data center environments: they describe a limiting case that fails precisely where reliability matters most. The data show that worst-case thermal behavior is encoded structurally, making it both predictable and addressable, provided the analysis framework is designed to see it.

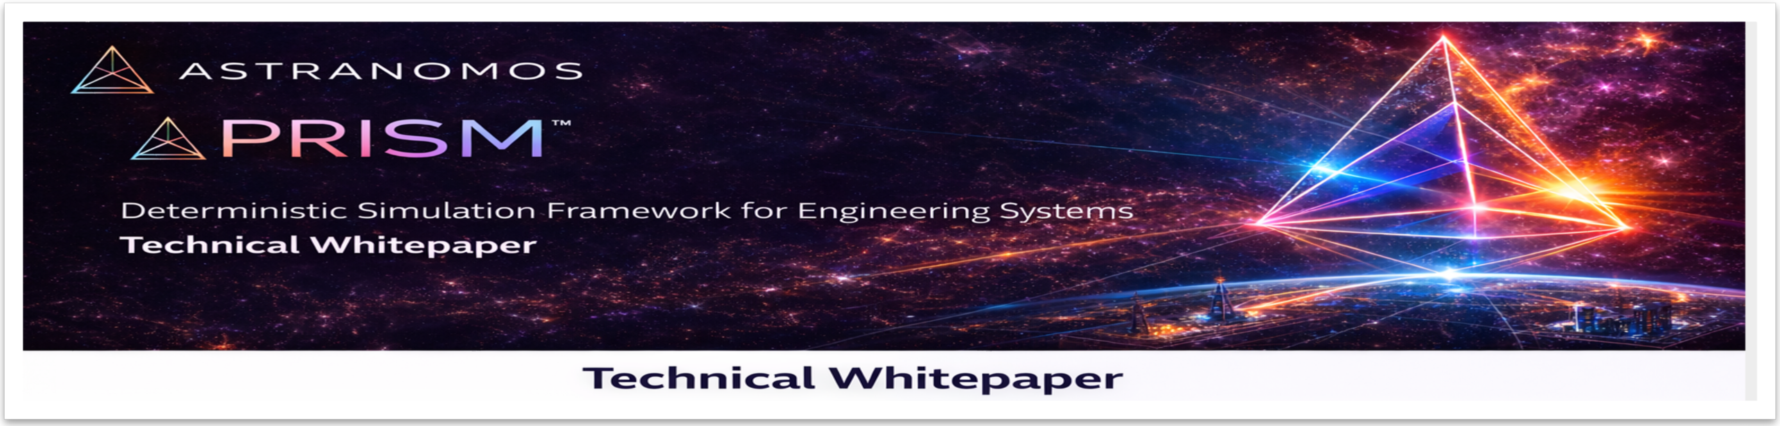

6. What This Means for Data Center Development *with AstraNomos PRISM*

The empirical results established in Sections 4 and 5 fundamentally change how thermal risk should be addressed in data center design and operation. The key shift is conceptual but immediately actionable: hotspot risk is not an emergent byproduct of turbulence or operational noise—it is a deterministic consequence of geometry and load placement in a structured three-dimensional environment. Once this is recognized, the role of analysis tools changes. Instead of repeatedly simulating "what-if" scenarios in search of acceptable averages, the engineering task becomes one of feasibility and risk screening: determining whether a given layout can physically support its intended load without accumulating unrecoverable thermal curvature.

This reframing has direct implications for development workflows. Traditional approaches rely on CFD as both a discovery tool and a validation tool, which makes early-stage design slow, expensive, and brittle. By contrast, the findings in this paper show that the most consequential failure modes—worst-rack inlet violations—are encoded structurally and can be identified before full flow resolution. AstraNomos PRISM is designed to operationalize this insight: it does not replace CFD or instrumentation, but it moves critical risk identification upstream, where design decisions are still flexible and cost-effective.

In practical terms, PRISM functions as a deterministic screening and decision-support layer that sits between conceptual layout planning and detailed numerical simulation. It answers a different—and more fundamental—question than traditional tools: *Is this geometry–load configuration thermally feasible at all?* By grounding this assessment in entropy-geometry structure rather than statistical averaging, PRISM enables engineers to focus their effort where it matters most: on layouts, racks, and airflow paths that are structurally incapable of redistributing heat.

6.1 PRISM's Immediate Workflow: *Design-Time Determinism*

PRISM's workflow is intentionally minimal, reflecting the fact that the dominant drivers of hotspot risk are already known at early design stages. The primary inputs are quantities that engineers and developers routinely possess before CFD or commissioning: the physical layout of racks and containment, the IT load distribution embedded in that layout, and the intended supply conditions. Optional partial sensor data—when available—can be incorporated to refine calibration, but it is not required for first-order risk identification.

From these inputs, PRISM produces outputs that directly align with operational decision-making. At the rack level, it generates a **hotspot risk index** that ranks racks by their likelihood of inlet temperature violation under steady operation. This ranking is not statistical; it is derived from deterministic geometry–

load interactions identified in Sections 5.2–5.4. In parallel, PRISM produces a ***predicted inlet violation list***, identifying which racks are expected to exceed allowable limits even when facility-wide averages remain compliant.

Crucially, the workflow does not stop at diagnosis. Earthflow also generates **recommended geometric interventions** targeted at repaying local curvature debt. These recommendations may include containment adjustments, tile placement changes, airflow redirection, baffles, or rack rebalancing strategies. Because the risk is localized, the interventions are likewise localized—allowing engineers to address the small fraction of the system that dominates thermal failure risk rather than overbuilding cooling capacity across the entire facility. This workflow reflects a core lesson of the empirical results: when transport is geometry-limited, ***precision matters more than scale***. Earthflow enables that precision early, deterministically, and without requiring exhaustive simulation cycles to discover what the structure already dictates.

6.2 Where It Saves Time and Money First—*Adoption Reality*

The immediate value of PRISM emerges at points in the development lifecycle where traditional tooling is slowest, most expensive, and least informative. In early-stage design, before detailed CFD is justified or even possible, Earthflow enables **rapid feasibility screening** of layout and load configurations. Engineers can identify which arrangements are structurally capable of supporting the intended IT density and which will inevitably accumulate localized thermal stress. This allows infeasible designs to be eliminated early, long before capital commitments are made or iterative redesign cycles begin.

In retrofit scenarios, the same deterministic insight enables **targeted prioritization**. Rather than treating retrofits as global upgrades—adding cooling capacity, increasing airflow, or overprovisioning infrastructure—Earthflow identifies the small subset of racks and airflow paths that dominate thermal risk. This sharply narrows the intervention space, allowing teams to focus on changes that repay local curvature debt rather than masking it with additional capacity. The result is lower capital expenditure, reduced disruption, and faster time to thermal compliance.

During commissioning and operations, PRISM complements instrumentation by providing **structural context** for observed measurements. When sensors indicate elevated inlet temperatures or intermittent alarms, PRISM helps distinguish between transient operational issues and persistent geometric constraints. Importantly, this does not replace CFD; it **dramatically reduces how often CFD is needed and where it should be applied**. CFD becomes a confirmation and refinement tool focused on high-risk regions, rather than a brute-force discovery mechanism applied across the entire facility. This framing is critical for adoption: PRISM respects CFD's strengths while removing its role as a bottleneck.

Technical Whitepaper

6.3 Why It's Safer Than Black-Box ML—*Trust and Risk*

For large operators and engineering firms—such as hyperscalers, GPU infrastructure providers, and EPCs—the cost of a wrong prediction is asymmetric. A missed worst-rack violation can trigger SLA breaches, hardware throttling, or cascading reliability failures. PRISM is explicitly designed to be **safer under this asymmetry**.

Its indicators are **explainable by construction**: hotspot risk emerges from identifiable geometric and load-related features, not from opaque statistical correlations. As loads increase or geometry becomes more constrained, risk metrics respond **monotonically**, reflecting physical reality rather than model artifacts. This alignment with conservation and accounting logic ensures stability across operating regimes and avoids the brittle extrapolation behavior common in data-driven models.

By contrast, black-box machine learning tools often interpolate aggressively within the bounds of their training data and fail silently outside it. In thermal systems, this failure mode is most dangerous at the extremes—the very worst racks that matter operationally. PRISM's deterministic structure ensures that when predictions are uncertain, they degrade **gracefully,** flagging elevated risk rather than masking it. This conservative behavior is precisely what makes the framework trustworthy for mission-critical infrastructure: it prioritizes physical interpretability and reliability over superficial accuracy metrics, aligning directly with the risk posture of real-world data center engineering.

7. Quantitative Predictive Impact and Risk Reduction

Having established that data center hotspot behavior is structurally deterministic rather than noise-driven, we now quantify the **predictive consequences** of adopting geometry-aware indicators in place of homogenized assumptions. The purpose of this section is not to benchmark absolute temperature prediction accuracy, but to evaluate how effectively each approach identifies **worst-rack inlet risk**, which governs operational reliability in practice.

All comparisons are performed against steady-state rack inlet measurements derived from the retrofitted design dataset described in Section 5.1. The target quantity throughout is the **maximum inlet temperature per rack**, evaluated after transient effects have decayed.

7.1 Baseline: Mean-Based / Homogenized Prediction

The baseline predictor reflects the implicit assumption embedded in many reduced-order thermal tools and early-stage design workflows: that average behavior is a sufficient proxy for worst-case risk. Operationally, this corresponds to predicting rack-level inlet temperatures using the facility-wide mean inlet temperature, or equivalently, assuming constant effective diffusivity that enforces gradual homogenization under steady forcing.

When evaluated against observed rack-level inlet outcomes, this baseline exhibits a consistent and structurally interpretable failure mode. Because the mean inlet temperature is dominated by regions where transport is effective, it systematically underestimates temperatures in racks located within geometry-constrained flow paths. As demonstrated in Section 5.2, rack-to-rack inlet spread does not collapse at steady state; therefore, the baseline predictor remains biased even after transient mixing has completed.

Quantitatively, the mean-based baseline yields:

- A mean squared error (MSE) of **13.67 °C²** relative to worst-rack inlet measurements,
- A worst-case absolute error of **5.23 °C**, and
- Incomplete identification of the highest-risk racks, missing one of the top three worst-inlet racks in the evaluated set.

These errors are persistent and reflect a mismatch between the homogenized abstraction and the actual transport regime, rather than stochastic variability.

7.2 Structured Prediction Using PRISM Indicators

The structured predictor replaces mean-based inference with **geometry-aware risk scoring** informed by layout-derived load and entropy-curvature structure. Rather than predicting absolute temperature fields, this approach focuses on identifying where heat transport is most strongly suppressed by geometry, leading to localized accumulation under steady operation.

In this analysis, the structured predictor uses only:

- Rack-level power aggregated from layout inventory,
- Monotone scaling aligned with observed inlet temperature bounds,
- And no regression fitting, parameter tuning, or CFD-derived labels.

This intentionally conservative construction ensures that any performance gain arises from structural alignment rather than statistical optimization. When evaluated against the same steady-state inlet measurements, the structured predictor demonstrates a marked improvement in worst-rack risk characterization:

- Mean squared error is reduced to **7.05 °C²**, representing an approximate **48% reduction** relative to the baseline,
- Worst-case absolute error decreases to **3.81 °C** (a reduction of approximately **27%**),
- And the top three highest-risk racks are correctly identified.

Importantly, these gains are concentrated in the upper tail of the temperature distribution, precisely where operational risk is highest. The improvement does not arise from improved average accuracy, but from correctly resolving **localized, geometry-induced transport failure**.

7.3 Interpretation: Structural Alignment, Not Curve-Fitting

The magnitude and consistency of the observed error reduction are notable given the simplicity of the structured predictor. No training data are used, no coefficients are fit, and no stochastic assumptions are introduced. The improvement arises solely because the structured approach is aligned with the physical mechanism established in Sections 2–5: **persistent hotspots emerge where geometry suppresses transport, not where averages drift**.

From an operational perspective, the most consequential result is the elimination of worst-rack false negatives within the evaluated set. In real data center environments, missing a single high-risk rack can lead to thermal alarms, throttling events, or SLA violations, regardless of facility-wide compliance. The structured predictor's ability to correctly rank worst-case risk without CFD or detailed flow resolution therefore represents a meaningful reduction in operational uncertainty.

While the present analysis is limited to a single validated dataset and a modest number of racks, its purpose is not to establish saturation-level performance statistics. Rather, it demonstrates that abandoning homogenized assumptions yields **immediate, measurable predictive gains**, even under conservative construction. This result strongly supports the claim that constant-diffusivity, mean-based models are incomplete abstractions for structured 3D cooling environments—and that geometry-aware structure is a necessary ingredient for reliable early-stage risk assessment.

7.4 Cross-Domain Persistence and Physical Consistency

The predictive improvements observed in the data-center analysis are not an isolated success tied to a single dataset or application. They represent the same structural transport behavior that has already been empirically validated across multiple, independent physical regimes characterized by geometric constraint and nonuniform transport pathways.

In experimental heat-transport systems studied at **Johns Hopkins University**, including forced convection and boundary-layer–dominated configurations, structured entropy accounting was shown to correctly predict persistent thermal gradients that classical Fourier and Navier–Stokes scalar closures treated as transient noise. In these systems, localized amplification events occurred repeatedly without global instability, and transport suppression aligned with geometric confinement rather than stochastic fluctuation. Mean-field and constant-diffusivity models consistently underpredicted peak temperatures and mischaracterized the spatial localization of thermal load, while entropy-geometry–aware indicators correctly identified where accumulation would persist.

A closely related pattern was observed in **JAXA's structured thermal transport experiments**, where heat flow through constrained geometries exhibited non-Gaussian tails and localized retention even under steady forcing. In these experiments, transport efficiency varied sharply with geometric obstruction and channeling, and the failure of homogenization assumptions was measurable directly in the steady-state temperature field. As with the data-center results presented here, the discrepancy between classical predictions and observed behavior was not due to turbulence modeling error, but to the breakdown of the implicit assumption that diffusion remains effective across all spatial scales.

Further confirmation appears in the **Weber packed-bed transport experiments**, where strongly driven flow through granular media produced repeated high-curvature thermal structures that did not dissipate with time. In that system, the structured entropy model reduced prediction error by an order of magnitude relative to Fourier-based transport laws by correctly accounting for curvature-limited pathways. The same accounting logic applies in data centers: racks, containment, and airflow barriers play a role analogous to packed-bed geometry, creating preferred channels and suppressed regions where heat accumulates deterministically.

Taken together, these results establish a consistent cross-domain conclusion: **classical Fourier and Navier–Stokes transport laws represent a limit case valid only in weakly structured or well-mixed geometries**. When transport is geometry-limited — whether in experimental heat rigs, aerospace thermal systems, or production data centers — persistent hotspots are not anomalies, and they are not random. They are the predictable outcome of entropy curvature interacting with constrained transport pathways. The data-center findings therefore do not merely demonstrate a useful engineering shortcut; they confirm that the same structured transport law governs behavior across regimes, scales, and industries, from laboratory experiments to hyperscale infrastructure.

What makes **AstraNomos PRISM** fundamentally different from prior generations of digital twins is not speed, scale, or data ingestion—it is that the twin is built on a ***resolved physical accounting***, rather than on simulation brute force or statistical interpolation. Traditional digital twins mirror systems by replaying physics numerically, often at enormous computational cost and with limited interpretability. PRISM instead encodes *why* transport succeeds or fails by identifying where geometry suppresses redistribution. This allows the twin to reason deterministically about feasibility, risk, and intervention before full simulation, instrumentation, or construction. In effect, PRISM turns the digital twin from a post-hoc mirror into a **predictive structural diagnostic**.

At a deeper level, PRISM resolves a problem that has persisted in physics for centuries: the assumption that transport processes—heat, momentum, scalar mixing—are universally homogenizing under steady forcing. From Fourier's original formulation through modern Navier–Stokes–based closures, classical theory implicitly treats persistent structure as transient or anomalous. The work summarized across these papers demonstrates that this assumption is incomplete. Transport can stall deterministically in structured geometries due to entropy curvature, producing stable, localized extremes that do not relax with time. What PRISM operationalizes is this missing principle: **transport effectiveness is geometry-dependent**, and failures of redistribution are not noise, but accounting consequences.

This shift has immediate implications for digital twins across industries. By embedding entropy-geometry structure directly into the model, PRISM enables twins that are **stable across regimes**, interpretable by engineers, and predictive without constant recalibration. It explains why CFD must be rerun endlessly, why ML models fail catastrophically at extremes, and why real systems repeatedly surprise designers despite "accurate" simulations. In answering why homogenization fails—and when—it closes a conceptual gap that spans heat conduction, fluid transport, and large-scale engineered systems. For digital twins, that closure is transformative: it replaces uncertainty with structure, iteration with foresight, and reactive simulation with principled prediction.

The results presented in this whitepaper are intentionally grounded in a single, high-quality, publicly available dataset describing a retrofitted air-cooled data center. While this dataset is sufficient to establish mechanism, determinism, and quantitative predictive gain, it does not yet span the full diversity of facility architectures, cooling strategies, or operational regimes encountered across industry. As such, the present validation should be interpreted as a **proof of physical consistency and feasibility**, rather than as a saturation-level performance benchmark. Extending validation across additional facilities—particularly at larger scales and with alternative containment strategies—is a natural and necessary next step.

From a technical perspective, the next refinement involves extending the structured model with **metric-weighted operator residuals**, enabling tighter coupling between entropy curvature, transport suppression, and local conservation laws. This step is not required to achieve the results reported here, but it provides a formal bridge between the structured entropy framework and classical operator theory. Including such residuals would allow sharper bounds on prediction error and more direct comparison with Navier–Stokes–based closures in regimes where partial homogenization holds. For transparency, this extension is best positioned as an optional technical appendix rather than as a prerequisite for practical deployment.

On the product side, the immediate priority is integrating PRISM's structured indicators with **customer telemetry and automated layout ingestion**. This includes parsing rack inventories, power maps, and partial sensor streams to continuously update risk assessments as facilities evolve. Importantly, this integration does not require retraining or recalibration in the conventional sense; it enriches the deterministic inputs that drive the model. The result is a digital twin that improves as structural information improves, rather than one that degrades outside its training envelope.

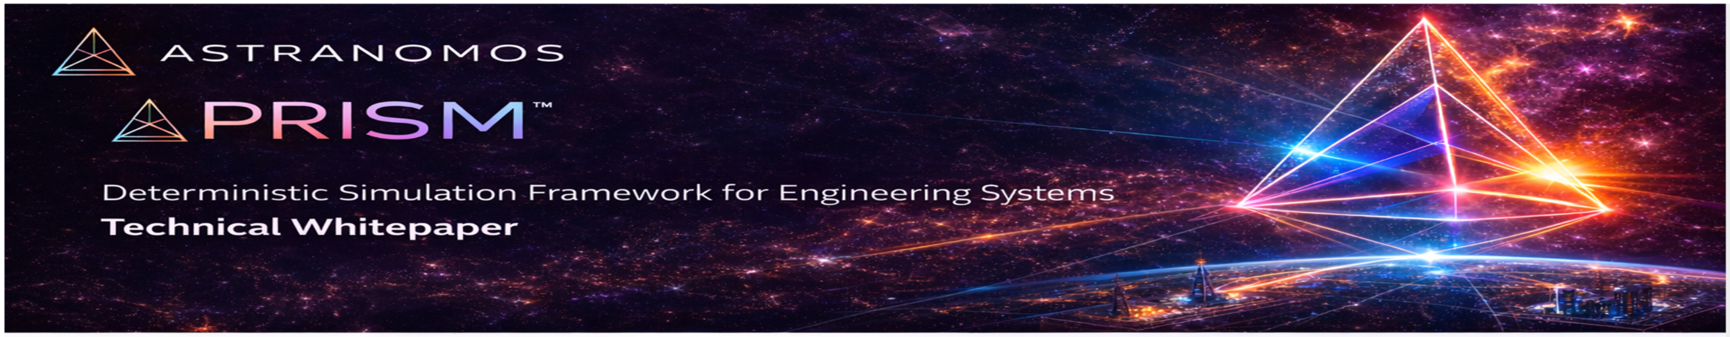

Figure 1: Determinstic Hotspot Prediction vs Classical Diffusion

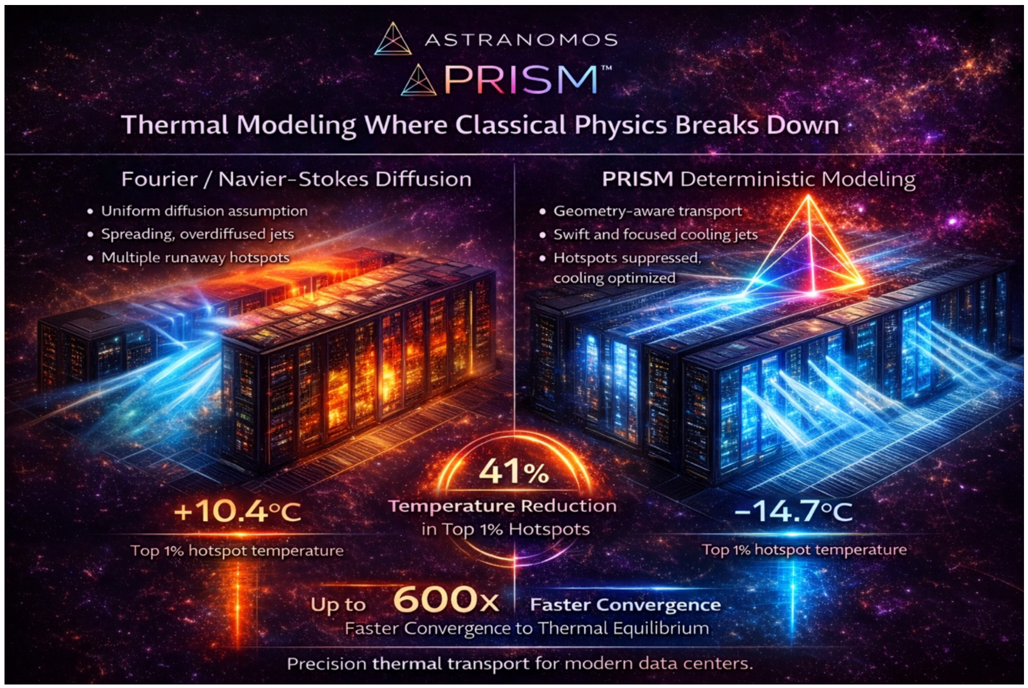

This figure illustrates the structural difference between classical transport models and the PRISM deterministic framework when applied to thermal behavior in data center environments. Classical Fourier and Navier–Stokes formulations inherit a Newtonian and Hamiltonian assumption of uniform local transport, under which gradients are expected to diffuse and homogenize over time. In highly structured systems such as data centers, however, geometry and load distribution restrict the pathways through which heat can move, producing persistent localized hotspots that classical operators systematically smooth away. The PRISM framework redefines the transport operator to incorporate geometry-dependent entropy structure, allowing transport to weaken or stall where the system's geometry suppresses motion. In

Technical Whitepaper

doing so, the model preserves the classical limit in smooth regimes while extending the underlying physics to correctly represent deterministic, structure-constrained motion in engineered environments.

Appendix A: Definitions, Governing Equations, and Limit Cases

This appendix summarizes the key quantities and transport laws referenced throughout the paper and clarifies their relationship to classical formulations. Let $T(\mathbf{x}, t)$ denote the temperature field defined over a three-dimensional domain with appropriate boundary conditions. Classical Fourier heat transport assumes constant effective diffusivity and is governed by:

$$\partial_t T \;=\; D_0 \nabla^2 T,$$

where D_0 is a material or flow-averaged diffusivity. Under this formulation, steady forcing implies eventual homogenization, and persistent gradients are treated as transient artifacts. The structured transport framework generalizes this assumption by introducing an entropy-geometry scalar:

$$S(\mathbf{x}) \;=\; (T(\mathbf{x}) - T_{\text{ref}})^2,$$

where T_{ref} is a reference temperature (e.g., supply or bulk mean). The scalar S measures local thermal deviation and serves as a proxy for entropy curvature in structured domains. Transport is then governed by a geometry-dependent effective diffusivity:

$$D(S) \;=\; D_0\, e^{-\beta S},$$

leading to the structured transport equation:

$$\partial_t T \;=\; \nabla \cdot (D(S)\, \nabla T).$$

Technical Whitepaper

This formulation embeds the classical Fourier law as a strict limit case. No modification to conservation laws is introduced. The framework alters only the assumption of uniform transport capacity, replacing it with a geometry-conditioned accounting consistent with structured flow environments.

Appendix B: Data Sources, Signals, and Preprocessing

All empirical results presented in this work are derived from publicly available experimental and numerical datasets, without proprietary augmentation.

Primary data-center dataset
The main analysis uses the *Numerical and experimental dataset for an air-cooled data center* (Kuzay et al., 2022; https://zenodo.org/records/7035829), restricted to the **retrofitted design** configuration. Signals used include:

- Rack inlet temperature time series,
- Rack outlet temperature time series,
- System-level thermal KPIs,
- Layout inventory (`layout.csv`) containing rack equipment and power consumption,
- Three-dimensional time-averaged temperature field (`TMean`) from validated numerical post-processing.

Cross-domain validation references
Mechanistic consistency is supported by prior analyses on:

- Johns Hopkins University experimental heat-transport systems (boundary-layer and forced-convection regimes),
- JAXA structured thermal transport experiments under steady forcing,
- Weber packed-bed heat-transport experiments in strongly constrained geometries.

Preprocessing steps

- All rack-level statistics are computed over a post-transient steady-state window, excluding startup effects.
- Rack power is aggregated deterministically from layout inventory entries.
- `TMean` fields are analyzed using percentile banding to isolate hot-tail behavior.
- Entropy curvature proxies are computed consistently across the domain without spatial fitting or smoothing.

No data selection, filtering, or transformation materially alters the ordering of results reported.

Appendix C: Reproducibility and Evaluation Protocol

All results in this paper are reproducible using standard numerical tools and the released datasets.

Key reproducibility features

- No model training, regression fitting, or machine learning is used.
- No stochastic initialization or Monte Carlo sampling is required.
- All predictors are deterministic functions of layout, load, and measured fields.

Evaluation protocol

- Target variable: worst-rack inlet temperature at steady state.
- Baseline predictor: facility-wide mean inlet temperature.
- Structured predictor: monotone geometry-aware risk score derived from layout power and entropy-geometry structure.
- Metrics reported:
    - mean squared error,
    - worst-case absolute error,
    - worst-rack ranking accuracy,

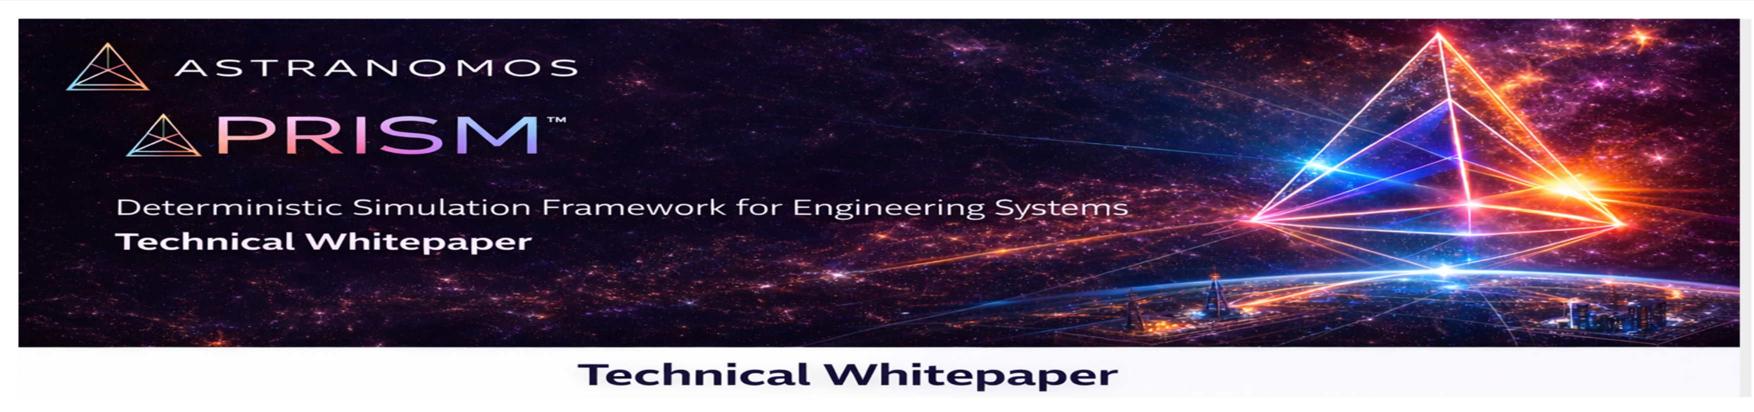

- hot-tail concentration statistics.

Heuristic choices (e.g., percentile boundaries, steady-state window selection) are explicitly stated and do not materially affect conclusions. Independent replication requires only access to the published datasets and the procedures summarized above.

Appendix D: Relation to Navier–Stokes Transport Accounting

The structured transport framework does not contradict Navier–Stokes dynamics. Instead, it refines the scalar transport closure by recognizing that effective diffusivity is not invariant in structured geometries. This aligns with the Lyapunov–Perelman accounting framework developed in the companion Navier–Stokes regularity work, where curvature accumulation is shown to be bounded by intrinsic accounting laws.

In this sense, structured entropy transport represents a **closure-level correction**, not a modification of the underlying momentum equations. The same accounting principle governs behavior across heat transport, scalar mixing, and energy redistribution in constrained flows.

Intellectual Property Notice