signal65

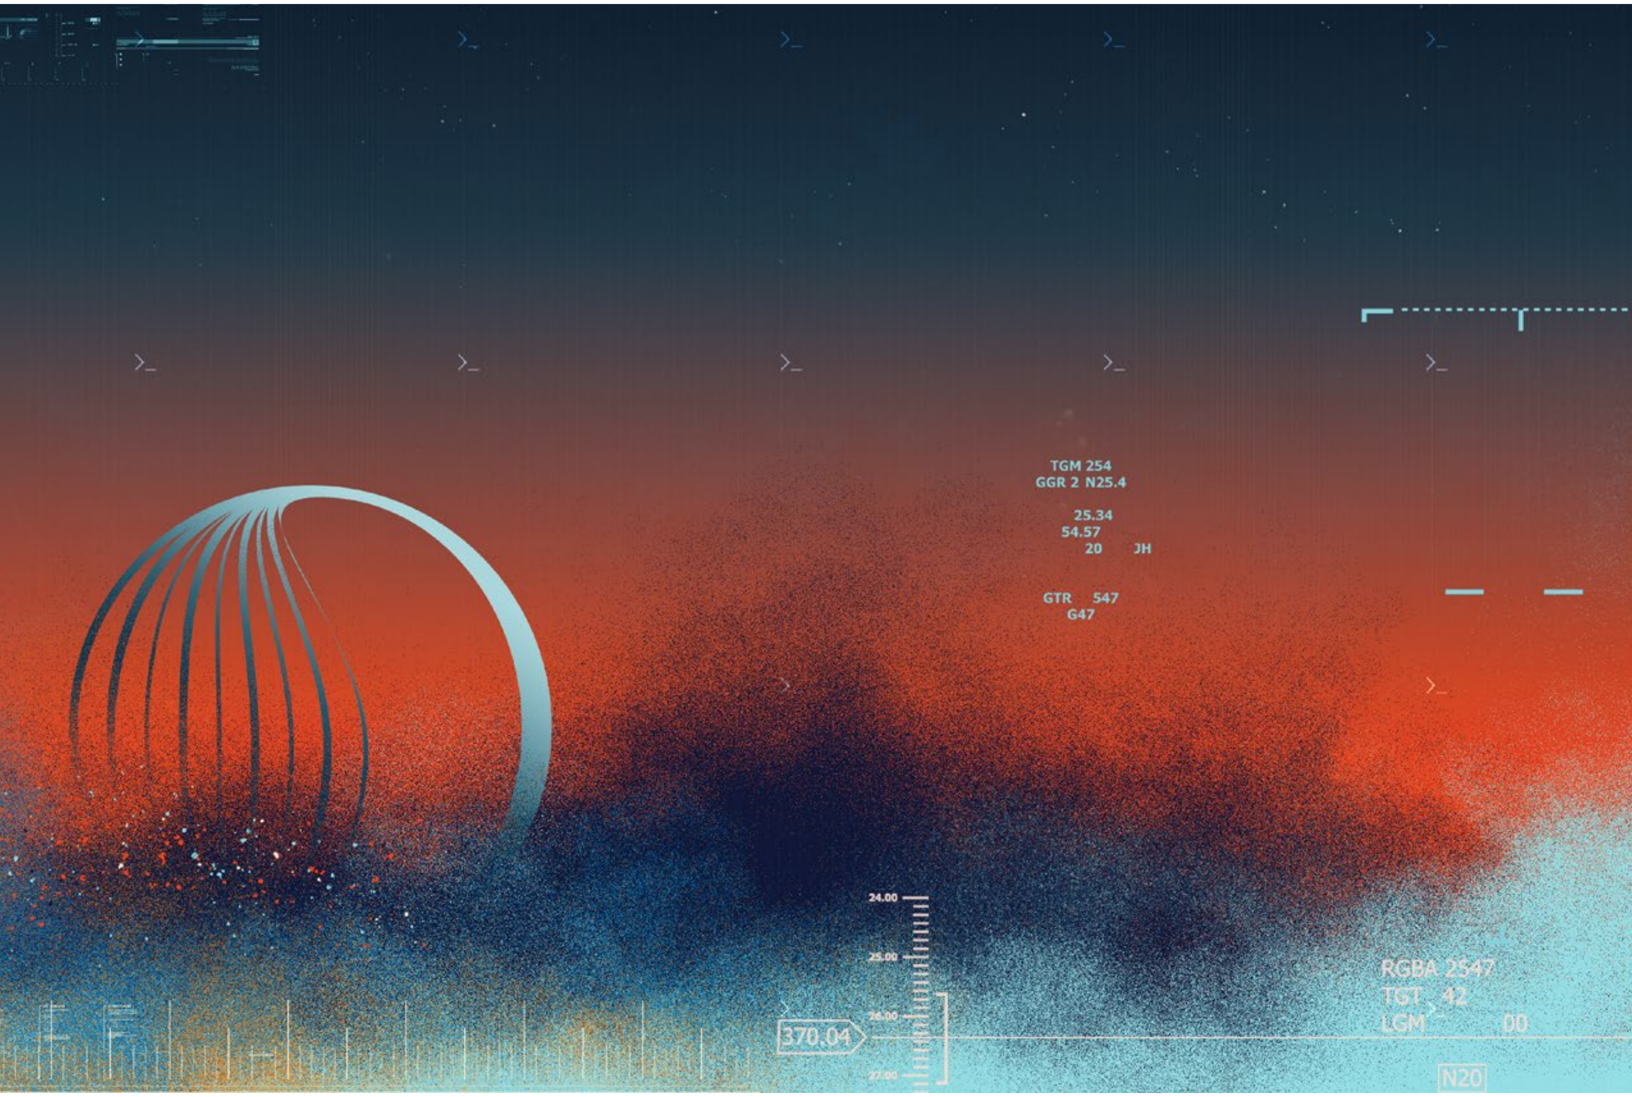


EXECUTIVE SUMMARY

# Digital Twins as the Control Plane for Physical AI

## Fleet-Scale Deployment with Dell AI Infrastructure

AUTHOR

**Brian Martin**

AI Data Center Performance | Signal65

**MARCH 2026**

IN PARTNERSHIP WITH

Dell Technologies

550 AERO

# Operational Intelligence for Physical Systems

Organizations operating safety-critical physical systems such as aviation fleets, energy infrastructure, industrial equipment, or autonomous platforms face a critical data imbalance. The events that matter most are rare by design. Enterprises therefore accumulate vast telemetry reflecting normal operation and very little representing abnormal or failure states.

Traditional approaches struggle in this environment. Threshold alerting reacts after limits are crossed. Fleet-average models dilute asset-specific behavior. The result is a persistent gap between available telemetry and actionable, predictive intelligence. Generative digital twins close that gap by learning the unique behavioral signature of each asset and continuously evaluating live performance enabling the twin to be the control plane and the asset sensor.

## Key Highlights

**Reduce Safety Risk**
Surface emerging performance deviations before they become safety incidents.

**Increase Asset Uptime**
Transition from reactive maintenance to predictive intervention.

**Lower Maintenance Cost**
Align maintenance activity with actual asset behavior.

# Generative Digital Twins

The generative digital twin functions as a continuously learning control plane for physical systems. Per-asset diffusion-based models are trained on real telemetry to reproduce normal operational behavior across the asset's full envelope. Once validated, the twin generates synthetic operating scenarios without introducing physical risk. Live telemetry is continuously compared against model predictions, and periodic retraining incorporates maintenance events, operational drift, and environmental variability, closing the loop between physical performance and digital intelligence.

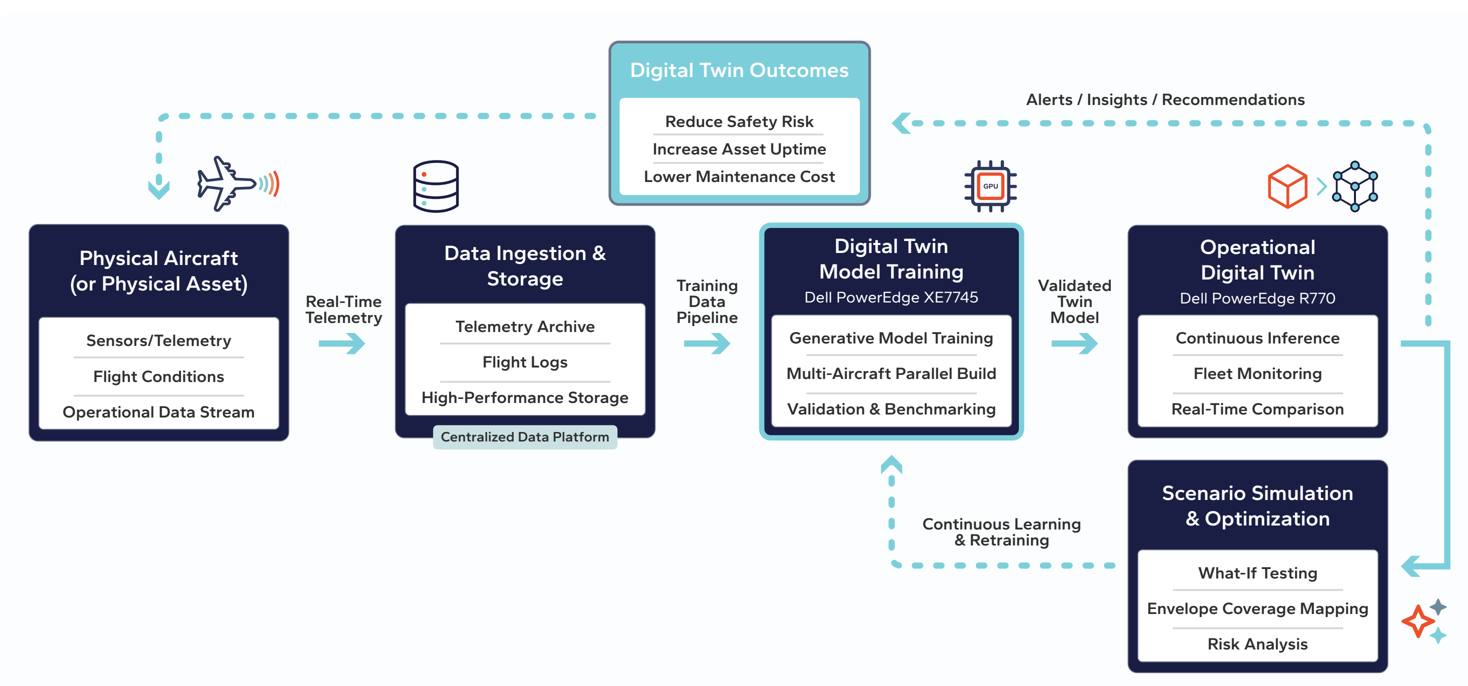


***Figure 1:*** *Architectural Overview*

# Continuous Learning Cycle

Physical Asset > Data Ingestion > Digital Twin Training (PowerEdge XE7745) > Operational Digital Twin (PowerEdge R770) > Scenario Simulation & Optimization > Continuous Learning

The architecture follows a closed-loop lifecycle:

1. Real telemetry trains a per-asset digital twin.
2. The twin generates operational scenarios across the asset's full envelope.
3. Live telemetry is continuously compared against twin predictions.
4. Residual deviations surface emerging anomalies.
5. Periodic retraining incorporates operational drift and maintenance events.

This lifecycle transforms monitoring from static threshold enforcement to dynamic behavioral comparison.

# Train Big, Deploy Smart

Generative digital twin training is compute-intensive; continuous inference is not. Dell's infrastructure portfolio aligns directly to these lifecycle phases. The Dell PowerEdge XE7745 provides the GPU density, memory bandwidth, and NVMe throughput required for parallel, multi-asset generative training at fleet scale. Once validated, operational twins transition to the Dell PowerEdge R770, a storage-rich, cost-efficient platform optimized for continuous telemetry ingestion, real-time inference, and anomaly detection.

Measured scaling demonstrates predictable economics and clean horizontal expansion: a single R770 node supports approximately 50 aircraft in real time; four nodes scale to roughly 200 assets; twenty nodes approach 1,000. Because inference workloads parallelize independently by asset, GPU-dense infrastructure is concentrated where it delivers the most value while inference tiers remain efficient and scalable. This dual-platform strategy accelerates time-to-intelligence while maintaining disciplined infrastructure utilization.

# Conclusion

Generative digital twins establish a scalable control plane for physical AI. By modeling assets individually, leveraging synthetic data to overcome anomaly scarcity, and aligning infrastructure to distinct training and inference phases, enterprises can move from reactive monitoring to predictive governance of high-value physical systems. Demonstrated on Dell AI infrastructure, this architecture demonstrates fleet-scale, per-asset intelligence is technically feasible, operationally practical, and economically scalable, delivering measurable reductions in risk, improved uptime, and lower maintenance cost.

While validated in aviation, the closed-loop generative twin model applies wherever safety-critical systems generate time-series telemetry and anomaly data is limited, including energy generation, industrial equipment, maritime propulsion, automotive powertrains, autonomous platforms, and industrial IoT environments. In each case real telemetry trains the twin; the twin generates realistic operating scenarios; deviations surface emerging risk; retraining incorporates operational evolution. The transition from observation to operational authority is achievable today across industries.

# Important Information About this Report

## CONTRIBUTORS

**Brian Martin**
AI Data Center Performance | Signal65

## PUBLISHER

**Ryan Shrout**
President and GM | Signal65

## INQUIRIES

Contact us if you would like to discuss this report and Signal65 will respond promptly.

## CITATIONS

This paper can be cited by accredited press and analysts, but must be cited in-context, displaying author's name, author's title, and "Signal65." Non-press and non-analysts must receive prior written permission by Signal65 for any citations.

## LICENSING



View the full report here »

## DISCLOSURES

Signal65 provides research, analysis, advising, and consulting to many high-tech companies, including those mentioned in this paper. No employees at the firm hold any equity positions with any companies cited in this document.

## ABOUT SIGNAL65

Signal65 is a leading research organization specializing in enterprise AI infrastructure optimization and deployment strategies. Our lab focuses on evaluating and optimizing AI hardware and software solutions for real-world enterprise applications, with particular expertise in large language models, retrieval-augmented generation systems, and distributed AI architectures.

For more information, visit signal65.com
or contact research@signal65.com

## IN PARTNERSHIP WITH

## CONTACT INFORMATION

Signal65 I signal65.com