

Failures are Rare – Operational Risk is Not

Safety-critical systems generate vast telemetry but almost no failure data. Traditional AI cannot train on what rarely happens. Generative digital twins close the gap.

THE ENTERPRISE CHALLENGE

The Data You Have vs. The Intelligence You Need

Organizations operating fleets, infrastructure, or industrial equipment face a structural data imbalance: abundant normal operations data, limited anomaly data, and high consequences when deviations go undetected.

- ✗ Thresholds react only after limits are exceeded
- ✗ Fleet-average models obscure asset-specific behavior
- ✗ Physics-only simulations miss operational variability
- ✗ Waiting for failures to accumulate is impractical

THE DIGITAL TWIN SHIFT

A Continuously Learning Operational Replica

A generative digital twin ingests real telemetry, learns each asset's behavioral signature, generates safe synthetic edge scenarios, and serves as the reference against which live performance is continuously evaluated.



Reduce Safety Risk

Identify emerging performance deviations before they escalate into safety incidents or in-service failures.

Increase Asset Uptime

Shift from reactive repair to predictive, condition-based intervention that maintains operational continuity.

Lower Maintenance Cost

Align maintenance with real asset behavior, reducing unnecessary inspections and preventing costly breakdowns.

ARCHITECTURE & SCALE

Closed-Loop Digital Twin Lifecycle

Real telemetry trains a per-asset digital twin. The twin generates operational scenarios. Live data is compared against predictions. Deviations surface anomalies. Retraining incorporates drift and maintenance events.



Applicable Across Telemetry-Rich, Failure-Scarce Industries

AVIATION & eVTOL

ENERGY INFRASTRUCTURE

INDUSTRIAL MACHINERY

MARITIME PROPULSION

AUTOMOTIVE POWERTRAINS

AUTONOMOUS PLATFORMS

VALIDATED IN A SAFETY-CRITICAL ENVIRONMENT

High-Fidelity Behavioral Learning from Limited Real-World Data

Using real flight telemetry across multiple aircraft platforms, generative diffusion models reconstructed engine behavior within single-digit error rates during stabilized operations.



10–20
FLIGHTS TO BUILD A TWIN

~ 144K
DATA POINTS MINIMUM

>98%
IN-FLIGHT MODEL ACCURACY

TRAINING PHASE

Dell PowerEdge XE7745

GPU-dense infrastructure for parallel per-asset model training across fleets.

Form Factor	4U Air-Cooled
CPU	Dual EPYC 9555
GPU	8x RTX Pro 6000
Memory	2.3 TB DDR5
Storage	8 Gen5 NVMe SSD

DEPLOYMENT PHASE

Dell PowerEdge R770

Storage-rich, cost-efficient inference for continuous fleet-scale anomaly detection.

Form Factor	2U Rack
CPU	Dual Xeon 6760P
GPU	2x NVIDIA L40S
Memory	2 TB DDR5
Storage	32 Gen5 NVMe SSD

Predictable Horizontal Scale

1 Node

OPERATOR SCALE
~ 50 aircraft fleet

4 Nodes

REGIONAL SCALE
~ 200 aircraft fleet

20 Nodes

ENTERPRISE SCALE
~ 1,000 aircraft fleet

Inference scales linearly across fleets – Dell PowerEdge R770 nodes

From Reactive Monitoring to Predictive Operational Authority

Generative digital twins enable organizations to model assets individually, reduce dependence on rare failure data, and align infrastructure to lifecycle phases. The result is measurable improvement in safety, uptime, and cost efficiency across fleets.

Explore the full Signal65 Insights Paper »

IN PARTNERSHIP WITH
Dell Technologies

DEMONSTRATED BY
550 Aero