C3.ai

Enterprise AI for Aircraft Predictive Maintenance

The Rapid Sustainment Office (RSO) is the USAF's sustainment technology and digital transformation accelerator, charged with optimizing warfighter readiness by exploiting technologies to revolutionize sustainment operations. RSO's mission supports maintenance, repair, overhaul, and supply chain management for the USAF's 5,400+ military aircraft.

The Predictive Analytics and Decision Assistant (PANDA) is a suite of predictive logistics applications developed within the joint C3 AI and RSO Center of Excellence (CoE). Built on the C3 AI Platform & C3 AI Readiness and deployed on CloudOne powered by AWS GovCloud, PANDA is now comprised of eight mission applications that support over 775+ active users. The applications focus on the enhanced Reliability Centered Maintenance approach, that forecasts component failure based on historical maintenance data.

The latest PANDA application, the Sensor Based Algorithm (SBA) application, was developed to detect and address failure when it occurs outside the historical forecasted cadence. The SBA application within PANDA utilizes machine learning models to analyze aircraft sensor telemetry data that detects degradation, enabling the proactive detection of potential failure to aid maintenance recommendations

All C3 AI SBA models employ a virtual sensor, a machine learning technique that predicts sensor values using surrounding sensors, allowing for the detection of deviations from the expected behavior and offering advanced warning and detection of degradation. Unlike threshold-based solutions, this approach provides the context of the system at every moment in time to detect anomalous behavior.

The B-1B, a long-range bomber, was the first platform onboarded into PANDA for SBA beginning in September 2022. Within 16 weeks (by December 2022), the SBA application and first B-1B SBA models were deployed in PANDA, including the ability to upload new data directly into the application for analysis, review alert evidence packages, and enable maintenance actions. Additional capabilities to efficiently manage and address alerts on a large scale have since been implemented and fielded.

After successfully deploying SBAs across the B-1B platform, RSO expedited the roadmap to onboard more aircraft platforms, starting with the C-5 Galaxy in January 2023. This and all future SBA expansions will reside in PANDA, powered by the C3 AI Platform.

Project Objectives

- Proactively detect signs of degradation prior to failure
- Reduce time between mission end and alert review
- Rapidly scale SBA models across any system or aircraft platform

Results

Up to 85%
time reduction from aircraft data extract to alert analysis

100s
of failures successfully detected

1000+
components actively monitored

92%
accuracy of alerts scored

Challenges

The U.S. Air Force faced challenges with its legacy predictive maintenance systems, as there was no government owned, consolidated, and authorized application that could scale to meet all use cases. Legacy predictive maintenance applications had critical gaps: (1) Lack of USAF ownership of their predictive maintenance applications, (2) Lack of Authorization to Operate (ATO) prevented establishment of automated data ingestion pipelines , (3) lack of data rights which hindered trust and adoption, (4) Traditional rules-based SBA models held less value compared to AI/ML approaches, and (5) SBA model development was not scalable without a unified data image or integrated data science environment.

The RSO already had access to data for multiple aircraft platforms, including the B-1B. However, each aircraft platform presented unique challenges with its telemetry sensor data, including requiring decoding, overcoming data transfer limitations, and dealing with different formats. Specifically for the B-1B, telemetry data was encoded and lacked a scalable, automated data pipeline for stakeholders positioned globally to access data quickly.

To address these issues, RSO sought a scalable solution to enhance the USAF's predictive maintenance workflows. The desired solution needed to:

- Utilize AI/ML techniques to extract deep insights from aircraft telemetry sensor data and predict system and component failures.
- Handle a large volume of data from disparate sources, including co-mingling telemetry sensor data with maintenance, supply, and flight logs, in a unified model.
- Allow easy deployment, management, and maintenance of SBA models using advanced ML techniques to enhance monitoring performance.
- Enable streamlined investigative and case management workflows for stakeholders to coordinate predictive maintenance efforts.
- Include robust data ingestion capabilities to facilitate rapid analysis by allowing direct uploads into the application or ingestion from cloud sources.

Project Highlights

- Ingested 5000 B-1B sorties with 75 billion rows of data for use in SBA model training.
- Built an extensible data model to support USAF aircraft platforms.
- Created a machine learning model pipeline that includes autoencoders, hierarchical models, transformers, and post-processors to generate SBA models.
- Deployed C3 AI Readiness for Aircraft Predictive Maintenance supporting two USAF personas: System Program Office (SPO) Engineers and Major Command (MAJCOM) maintenance managers.
- Implemented flight line data uploader enabling ingestion of new sensor data as it’s offloaded from the aircraft and blending sensor data with other data sources already integrated into PANDA’s unified data image.
- Developed SBAs for 11 failure modes, spanning 29 models, to detect system and component degradation across the B-1B.

Approach

The B-1B was chosen as the first platform to onboard into the SBA application due to its robust onboard sensor suite, extensive historical data spanning five years, and access to a decoder. The project began by ingesting and transforming historical telemetry sensor data from 5000 sorties into a unified data image. This data model formed the foundation for future onboarding of additional aircraft platforms and addressing other AI use cases.

All SBA models developed utilized the virtual sensor approach, which predicts sensor values using surrounding sensors for a healthy system. By comparing actual sensor measurements to the predictions, the model can determine if it resembles a healthy or degraded state. The virtual sensor model is trained solely on healthy system data, ensuring it only predicts healthy values.

SBA models have been developed at various levels: the system-level monitors whether the system is generating the desired symptomatic output, such as "Temperature Too Hot," while the component-level identifies issues specific to components and assists in root cause analysis for fault isolation, for example, "Faulty Modulating Valve."

The initial stages of SBA model development focus on data quality, analyzing each relevant sensor for significance, and addressing any missing or invalid measurements. To train the virtual sensor model exclusively on healthy data, a training mask is applied to remove suspected degraded system data. Thousands of models, each varying slightly, are then trained and evaluated to find the optimal one. Additional ML techniques are used to aggregate the virtual sensor model output for precise alert triggering.

The Virtual Sensor Toolkit was built on the C3 AI Platform's integrated data science environment, serving as the foundation for all future SBAs, including models developed for other platforms. This toolkit fully leverages the extensibility and scalability of the C3 AI Platform for rapid model development.

For the user interface, the CoE collaboratively designed and developed an application to manage and analyze alerts. It includes key data visualizations to provide System Program Office (SPO) engineers with confidence in the virtual sensor model's output and ensure that alerts accurately detect failures.

Virtual Sensor Toolkit to Generate SBA Models

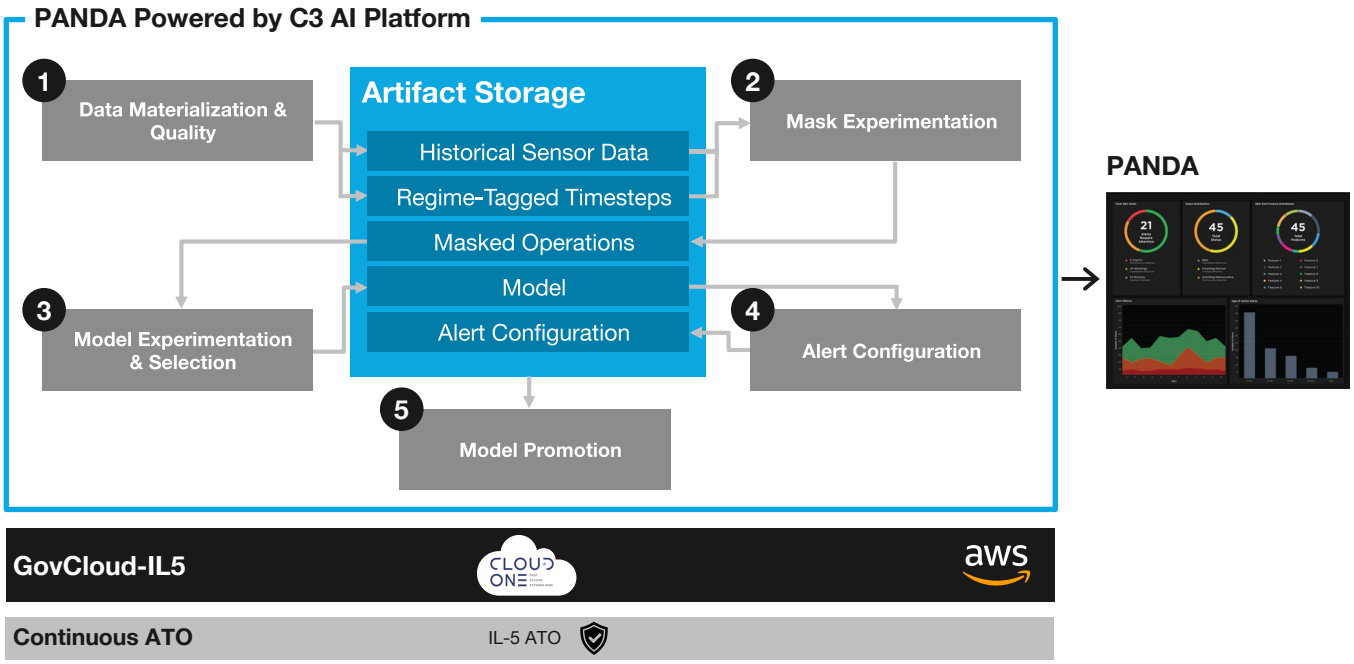

Solution Architecture

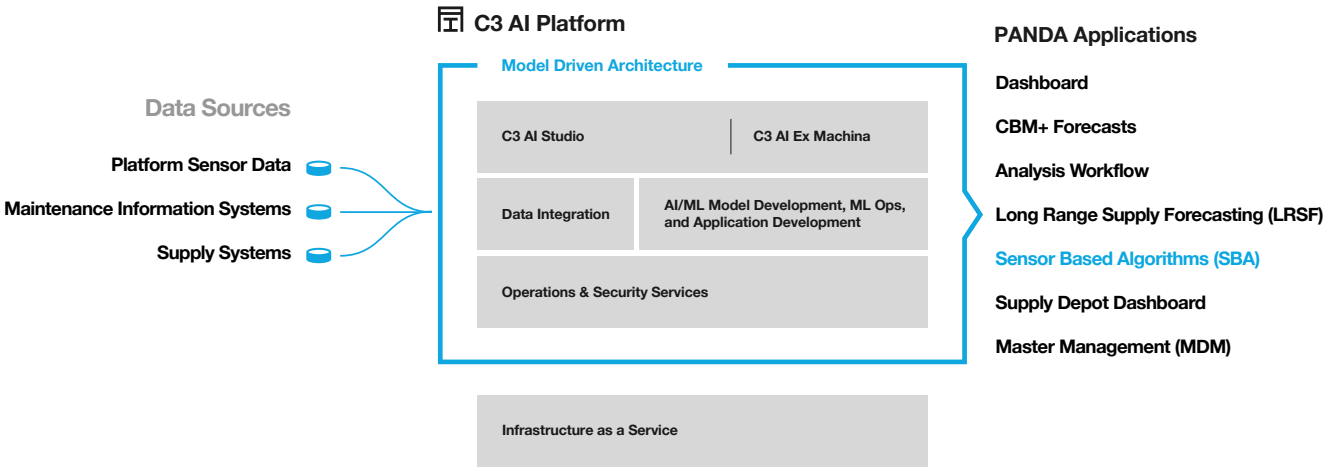

Benefits

Built on the C3 AI Platform and extending the C3 AI Readiness application, PANDA delivers the following benefits to RSO and the USAF:

- Develop SBA models rapidly using the Virtual Sensor Toolkit in days, rather than weeks or months with traditional data science tools.
- Access the application from any internet-connected device with a CAC, allowing sensor data upload, alert analysis, and maintenance actioning directly from the flight line.
- Upload encoded sensor data immediately following a sortie, automatically decode and trigger alerts within PANDA, greatly reducing time to review alerts.
- Analyze the evidence package of each alert to confirm its validity as a true detection of degradation.
- Efficiently manage and address alerts at scale using case management capabilities.

About United States Air Force Rapid Sustainment Office (RSO)

Established in 2020, the RSO accelerates the delivery of critical operational solutions to the Department of the Air Force sustainment enterprise. https://www.aflcmc.af.mil/RSO

Proven Results in 8-12 Weeks | **Visit c3.ai/get-started**