Alert Lifecycle and Best Practices for Predictive Maintenance Alert Lifecycle Management

Justin Sindewald[1], Ryan Latini[2], and Joseph Rice[3]

[1,2,3]*The Boeing Company, Seattle, WA, 98108, USA*

justin.a.sindewald@boeing.com
ryan.j.latini@boeing.com
joseph.rice2@boeing.com

ABSTRACT

The predictive maintenance alert lifecycle is a critical topic in the aviation industry. Stakeholders, including operators, suppliers, and Original Equipment Manufacturers (OEMs), require effective frameworks to support the value proposition of predictive maintenance products and services. However, defining alert effectiveness is challenging due to the lack of industry standards for the end-to-end lifecycle of predictive maintenance alerts. Adding to the challenge, different stakeholders may want to optimize on different objectives. Often, alert performance is measured prematurely or not at all. To ensure high-quality alerts, all alerts should be managed through their entire lifecycle until obsolescence. This whitepaper outlines a clear conceptual framework for the predictive maintenance alert lifecycle and best practices alert lifecycle management. Future work will expand upon the qualitative benefits discussed herein with quantitative results achieved from applying this framework to assess alert effectiveness at scale within the Boeing alert catalog.

1. ALERT LIFECYCLE DEFINITION

The alert lifecycle consists of three phases illustrated in Figure 1.



Figure 1. Alert Lifecycle Major Phases

The maturation level of an aircraft program through its own product lifecycle dictates the type of data available for predictive maintenance alert development at the initialization phase. Aircraft Condition Monitoring System (ACMS) and Aircraft Condition Monitoring Function (ACMF) data have been a cornerstone of prognostics and health management in commercial aviation for many years. ACMS and ACMF data are recorded onboard and transmitted in-flight via satellite communications enabling near "real-time" alerting. This is often referred to as "snapshot" data and is typically the most limited dataset when compared to other data sources. Newer aircraft like the 737 MAX have systems such as Aircraft Health Management Onboard (AHMO), which is similarly captured onboard but is capable of selective high-resolution data capture. AHMO has an equivalent parameter set to full-flight data and has the benefit of being updated wirelessly in a matter of days or weeks. This capability enables rapid discovery and updates which ACMS or ACMF are not capable of as they rely on a more labor-intensive software update process. Full-flight data sources such as Quick Access Recorder (QAR) and Continuous Parameter Logging (CPL) data capture a subset of all available parameters at a typical sampling rate of 1 Hertz (Hz). The entire flight recording is transmitted on ground where it can then be translated from raw binary into engineering units. QAR and CPL data are the central focus of prognostics research in commercial aviation but are not capable of "real-time" alerting due to the transmission lag induced by the large file size. The last standard dataset is flight test port (FT-P) data. FT-P data is ultra-high resolution with some parameters being recorded at a rate of 100 Hz. All possible parameters are recorded, and the result is a high-fidelity view of normal operational behavior of the aircraft. FT-P data helps predictive maintenance engineers to create better hypotheses by understanding what normal behavior is without impacts of under sampling.

Alert lifecycle awareness and management helps predictive maintenance teams be intentional when using these various data sources to develop alerts. The initialization phase should use the highest fidelity data available such as flight test port data and QAR/CPL data to enable alert research. These high-fidelity datasets are transmitted for offboard parsing, offboard feature engineering, and offboard alerting.

Justin Sindewald et al. This is an open-access article distributed under the terms of the Creative Commons Attribution 3.0 United States License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Transmitting, processing, and analyzing these large datasets can lead to several hours, or days, of "lead time" prior to triggering an alert. Alerts implemented through lower fidelity datasets such as ACMF or AHMO benefit from onboard parsing and onboard feature engineering prior to transmission to enable real-time alerting. Throughout the Service Life and Optimization of an alert, the data capture of alerting features should migrate to onboard parsing and feature engineering systems, such as ACMF or AHMO. This harmonization should occur well before an alert reaches its obsolescence and understanding the end-to-end alert lifecycle will help predictive maintenance teams determine when this harmonization should occur.

2. RELATED WORK

The alert lifecycle framework presented herein is specifically designed to manage an alert catalog containing thousands of alerts across various aircraft types and components. The framework is intended to supplement accepted end-to-end condition monitoring processes and data mining processes such as those described in ISO (2003) and Chapman et. al (2000), respectively. The referenced processes are sufficient for developing predictive maintenance alerts but do not adequately address practices to manage alerts throughout their service life. One of the primary recommendations of this work is to limit the scope of alert effectiveness frameworks to focus specifically on the steady-state period of the alert service life. Industry standards published by SAE (2020) and Saxena et al. (2008) provide greater detail on accepted alert performance metrics that can be calculated to measure alert effectiveness. The definitions of such metrics are straightforward, but such standards do not fully consider the practicality of such calculations using in-service data. This paper compliments such standards by overlaying key performance metrics onto the applicable lifecycle phases to ease the management of an alert effectiveness framework.

3. ALERT LIFECYCLE PHASES

The typical lifecycle progression of an alert as measured by alert rate and alert utilization rate through the major lifecycle phases is shown in Figure 2. The alert rate is defined as the number of alerts per fleet flight cycles flown and the utilization rate is the number of alerts actioned per fleet flight cycles flown.

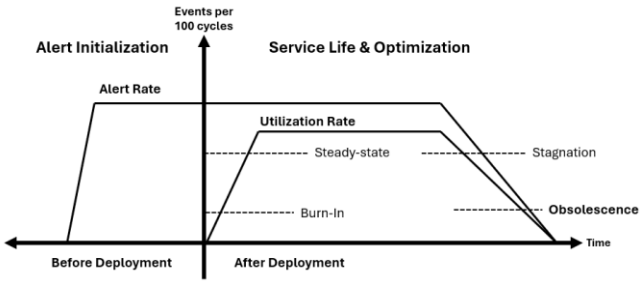


Figure 2. Lifecycle Progression of Alert Rate and Utilization

3.1. Alert Initialization

Alert Initialization involves developing a new predictive alert from the initial request to production deployment. A structured process is essential, including phases for ingestion, engineering understanding, research and development, and solution delivery. Boeing's Predictive Maintenance Content Team (PMCT) process exemplifies this structured approach and is outlined in Figure 3.

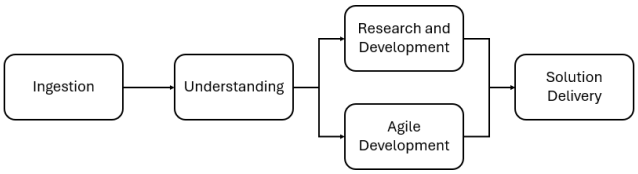


Figure 3. Predictive Maintenance Content Team Process

The Ingestion phase should assess existing solutions, deployment strategies, and project prioritization to put the operator experience at the forefront of any alert design to follow. Stakeholders must evaluate the suitability of current prognostic environments and the completeness of maintenance procedures. Experienced teams can forecast the complexity of an alert algorithm based on project proposals, ensuring that necessary changes to production environments or supporting maintenance procedures are initiated promptly.

The Understanding phase aims to build a knowledge base before extensive data science work begins. A structured engineering package summarizing Boeing's knowledge of the target failure mode is crucial for guiding the project.

Key outcomes should include:

1. A visual representation of operational norms, both at the system and component levels.

2. Identification methods for relevant target failure events.

3. Cost-based minimum alert performance requirements.

4. Recommendations for the most viable alert techniques to meet performance requirements.

The Research and Development process follows a typical data science process framework such as the cross-industry standard process for data mining (CRISP-DM). A critical milestone to be highlighted is the initial "beta" testing which marks the transition from historical data analysis to real-time detection. This testing is vital for identifying potential confirmation bias and verifying its usability to the personnel who are going to receive it.

Alert Solution Delivery is the final quality assurance phase that facilitates a smooth transition from early-adopter alert testing to fleet-wide operations. End-user training, fleet communications, and proper documentation updates are the

most important deliverables to support a healthy service life for a new alert.

3.2. Alert Service Life and Optimization

The alert service life spans from initial solution delivery to obsolescence. The progression of an alert through its service life and optimization cycle is illustrated in Figure 4 with recommended performance measures important to each phase.

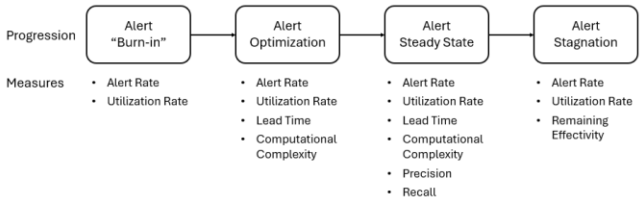


Figure 4. Recommended Alert Performance Measures

Alert effectiveness metrics measure how well an alert is preventing the target events while alert utilization metrics measure how often an alert is actioned by a customer. A measurement framework should ensure that deployed alerts meet minimum performance requirements for alert effectiveness and alert utilization, with an allowed "burn-in" period before formal performance reporting. During the "burn-in" period, it is important to identify whether the alert is under-alerting or over-alerting compared to expectations which can be measured by alert rate (e.g. # of alerts per fleet flight cycles). Operators will need time to digest alerts during the "burn-in" phase and may not start utilizing alerts right away. The allowed "burn-in" period should be defined by the expected frequency of both target events and alerts. An uptick in utilization of an alert signals the start of the "optimization" period. Feedback on a particular alert will start to come in from operators and opportunities for optimization will start to be identified. Measures of "lead time" and "computational complexity" can be used to determine when onboard harmonization may be appropriate, if not already accomplished. The optimization period ends when onboard data collection is implemented and alert utilization rate by operators has stabilized. In addition to metrics captured during "burn-in" and "optimization," the "steady state" period is where alert effectiveness metrics can be measured most efficiently and where alert performance should be most predictable. The "stagnation" period begins once a significant decrease in alert rate and/or utilization rate compared to the steady-state period is identified. If an obsolescence trigger such as a reliability modification is not yet known, then the alert should be reviewed against the obsolescence criteria in this paper to determine whether the alert is becoming obsolete.

Alert performance measurements throughout its service life can be difficult to do efficiently and effectively depending on the availability of component removal and logbook data. The alert performance measurement metrics must accommodate large variations in the intended use of the alert, the target users, customer operational norms, and the overlap of alerts with similar failure targets that some users may have access to.

1. Intended Use: the intended use of an alert that aids in diagnosing a complex set of airplane faults is different from an alert that predicts a specific component failure, and, if established incompletely, the performance metrics designed around predictive alerts may look poor for a very successful diagnostic alert.

2. Target users: alerts that are exposed to users in a variety of roles may provide high value for a user in a predictive maintenance engineering role, and low value for a user in a maintenance controller role. Some target alert users may be choosing between alerts that target similar failure modes depending on the tools and data they have access to.

3. Alert target overlap: for a target user that has access to multiple alerts targeting the same failure mode, an individual alert may have a lower value than for an operator that only has access to one alert option for that failure mode. That access can vary depending on the tools and data available to that user.

An effective alert performance measurement system must recognize this complexity and enable relevant evaluation of that alert throughout its lifecycle. A scalable engagement-based alert performance framework can be created based on alert utilization. An engagement-based framework assumes that alerts with a high expiration and rejection rate are low-performing and that alerts with a high completion and acceptance rate are high-performing (Figures 5a and 5b). However, there must also be an evaluation for potential extenuating circumstances that may unfairly be penalizing alert effectiveness, such as data dropouts, inventory shortages, etc. The lowest performing alerts are then evaluated in more detail where other metrics which require more effort to accurately assess are calculated.

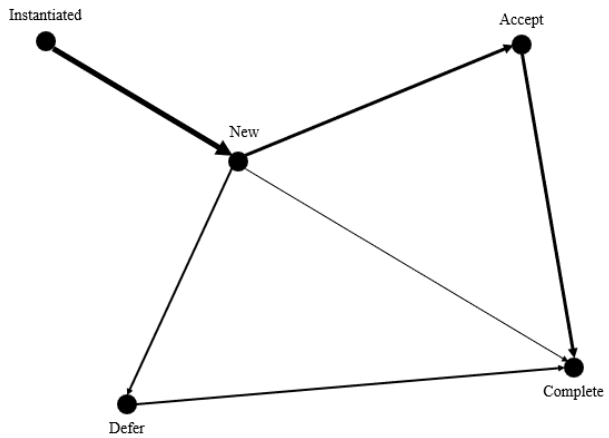


Figure 5a. Desirable Workflow Progression

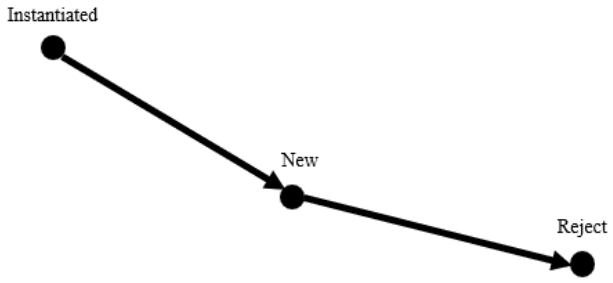


Figure 5b. Undesirable Workflow Progression

Minor adjustments to the understanding of the target failure mode or alert methodology can occur throughout the service life without initiating obsolescence. However, significant negative performance should prompt a review for potential obsolescence if no minor adjustments can recover performance. Additionally, especially with machine learning based alerts, model drift would be expected over time and allowable thresholds should be established early and monitored. Changes to component design, system configurations, or integrated vehicle health management (IVHM) aspects should trigger a review of the alert's relevance.

3.3. Alert Obsolescence

Obsolescence is the final phase of the alert lifecycle and is crucial for maintaining a healthy predictive maintenance alert catalog. It can be challenging to determine when to retire an alert but avoiding obsolescence altogether can lead to operator fatigue and distrust in alert content. Alerts may become obsolete after reaching significant milestones, such as:

1. Minimum performance requirements did not meet the following defined maturation period.

2. Incorporation of a Service Bulletin or other design change that addresses the target failure mode.

3. Updates to onboard IVHM software expanding monitoring capabilities.

4. Obsolescence of features in the selected prognostics environment.

Predictive maintenance engineers should proactively drive alert obsolescence for offboard alerts by making updates to onboard alerting capabilities. Delaying or avoiding the harmonization of onboard alerting will ultimately diminish the total value capture of an alert. A technical review process should guide obsolescence decisions, involving a board representing predictive maintenance strategy, alert execution, and maintenance execution. Documenting obsolescence reasons is vital for knowledge transfer, ensuring future teams do not repeat fruitless ingestion and understanding activities, or recreate obsolete alerts.

3.4. Alert Lifecycle Example

In practice, the alert lifecycle when framed from the operational issue it is trying to address is not as linear as presented here. The lifecycle of one alert iteration is likely to branch into related alert content lifecycles for either the same failure mode target or other adjacent problems. The 737 MAX Fan Air Modulating Valve (FAMV) prognostic alert lifecycle is exemplary of the natural lifecycle of predictive maintenance alert content. The FAMV controls cooling air to the cold side of the precooler heat exchanger and regulates the temperature of the engine bleed air supply to be below 390 degrees Fahrenheit. The lifecycle of alert content for the 737 MAX FAMV was seeded by the design of the IVHM system and initial in-service data collection. The 737 MAX FAMV IVHM System elements are:

1. Sensors installed for bleed pressures, bleed temperatures, and FAMV position

2. Integrated Air System Controller (IASC) captures sensor data and sends it to the Digital Flight Data Acquisition Unit (DFDAU) via an ARINC 429 connection protocol

3. 737 MAX equipped with multiple data recording formats for health management data including Quick Access Recorder (QAR), AHMO, and Aircraft Condition Monitoring System (ACMS)

4. 737 MAX equipped with multiple data transmission system options including via Aircraft Communications Addressing and Reporting System (ACARS) satellite, WiFi, or Cellular connections

The 737 MAX Fan Air Modulating Valve In-Service Data Collection findings which led to the alert research were:

1. MAINT LIGHT related to FAMV Maintenance Messages became a top driver for schedule interruptions for the worldwide 737 MAX fleet

2. Root cause investigation at the supplier determined the primary failure mode to be an actuator seal tear

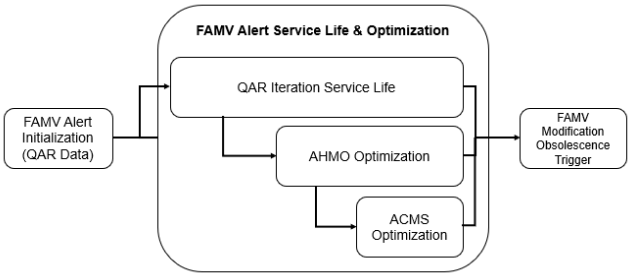


Figure 6. 737 MAX FAMV Alert Lifecycle

The end-to-end alert lifecycle for the FAMV is shown in Figure 6. The first prognostic alert developed for the FAMV used QAR data and offboard logic to identify when the actuator seal was starting to tear and likely to drive a MAINT LIGHT. Due to the typical time delay of 1-2 days for QAR data transmission & processing, the success of the offboard QAR alert branched off into an AHMO alert which was capable of the same prognostic detection with the advantage of real-time alerting. Additionally, an ACMS report is in the process of being created that mimics the AHMO report to enable a parallel alert for operators who do not have the AHMO option. The FAMV supplier has introduced a modification to the FAMV which mitigates the actuator seal tear and significantly reduces FAMV failures. The QAR, AHMO, and ACMS-based alerts are rendered obsolete on modified FAMVs, and Boeing is in the process of suppressing any FAMV alerts for operators who have reported that they have incorporated the FAMV modification. The number of FAMVs requiring modification was significant and the obsolescence period has taken many years and is still ongoing.

Operators are now requesting prognostic alerts for the modified FAMV, but failure modes are not as clear as with the actuator seal tear on the original FAMV. The supplier is modifying the IASC to transmit bleed valve torque motor currents to the DFDAU which will make these values available for prognostic alerting. The availability of the torque motor currents in the future is expected to be the key to prognostic alerting on the modified FAMV and will lead to a new lifecycle of yet another 737 MAX FAMV alert variant.

4. CONCLUSION

Establishing a standardized lifecycle for predictive maintenance alerts is essential for enhancing the effectiveness and reliability of these critical systems within Boeing's operations. The conceptual structure of the alert lifecycle established by the paper takes a step towards a scalable alert effectiveness framework that can be repeated by predictive maintenance providers throughout the industry. The 737 MAX FAMV prognostic example discussed illustrates the complexity of alert lifecycles in practice and underscores the need for responsible lifecycle management. Future work will focus on quantitative measures that further prove the alert lifecycle model and demonstrate the success of lifecycle management best practices recommended herein. As we move forward, continuous refinement of this lifecycle will be vital to content curation in Boeing's predictive maintenance products, ultimately leading to improved operational efficiency for all stakeholders involved.

REFERENCES

Chapman, P., Clinton, J., Kerber, R., Khabaza, T., Reinartz, T., Shearer, C., & Wirth, R. (2000). CRISP-DM 1.0: Step-by-Step Data Mining Guide. CRISP-DM Consortium. Burlington, MA: SPSS Inc.

International Organization for Standardization (ISO) (2003). Condition monitoring and diagnostics of machines — Data processing, communication and presentation — Part 1: General guidelines. Geneva, Switzerland: International Organization for Standardization (ISO).

Society of Automotive Engineers (SAE) (2020). Diagnostic and Prognostic Metrics for Aerospace Propulsion Health Management Systems. In SAE, AIR 7999. Warrendale, PA: Society of Automotive Engineers.

Saxena, A., Celaya, J., Balaban, E., Goebel, K., Saha, B., Saha, S., & Schwabacher, M. (2008). Metrics for Evaluating Performance of Prognostic Techniques. In 2008 International Conference on Prognostics and Health Management Conferences Proceedings.

BIOGRAPHIES

Justin A. Sindewald is a lead predictive maintenance engineer at Boeing. He completed his bachelor's degree in mechanical engineering at the University of Illinois Urbana-Champaign. Prior to Boeing, Justin worked at United Airlines for 6 years as a predictive maintenance engineer where he focused on environmental control systems across all the Boeing fleet types operated by United.

Ryan J. Latini is a senior predictive maintenance engineer at Boeing focused on the Boeing 787. He holds a bachelor's and a master's degree in mechanical engineering from the University of Colorado Boulder and the University of Washington Seattle respectively. Prior to predictive maintenance engineering, he worked at Boeing as a design and analysis engineer in environmental control systems on the 737 in pneumatics and air conditioning systems.

Joseph Rice is a senior predictive maintenance data scientist at Boeing. He holds a bachelor's degree in mathematics from the University of Connecticut and a master's degree in data analytics from Southern New Hampshire University. Prior to Boeing, he worked at Pratt & Whitney for 5 years as a simulation architect and data scientist focused on engine durability and maintenance modeling for GTF powerplant family power-by-the-hour contracts.