Data Driven Life Cycle Analysis to Optimize Cost, Risk, and Sustainability

George O. Bayer, Jr.

Brian Carroll

Cobec Consulting, Inc.

February 20, 2024

Contents

Abstract

Many government infrastructure investments adhere to a standard life cycle to estimate program cost, plan replacement timing, and compare business cases to one another in a cost-benefit analysis. What if those life cycle replacement timelines are inconsistent with system sustainability and are not cost-effective? Some infrastructure systems which are replaced according to an end-of-life schedule can be sustained more cost-effectively for longer periods of time via preventative maintenance. Our team examined multiple infrastructure program replacement timelines and analyzed operational effectiveness, cost/risk trade-offs, system redundancy, and sustainability, and we recommended life cycle adjustments based on those considerations. We reduced overall program cost by extending replacement timelines, eliminating system redundancy without compromising sustainability, and reprioritizing maintenance portfolios on critical backlogs. We document a comprehensive process to customize program life cycles to optimize cost, risk, and sustainability.

1 Introduction

When companies and government agencies implement infrastructure and operational systems, they plan to maintain those systems for an estimated useful life, the duration for which they can sustain operations without a significant increase in system failures and without operational impairment. During project acquisitions, government agencies estimate system life cycles according to manufacturers' estimates and business case timelines. They adhere to a standard life cycle to estimate program cost, plan replacement timing, and compare business cases to one another in a cost-benefit analysis. What if those life cycle replacement timelines are inconsistent with system sustainability and are not cost-effective? Are these timelines customized to the useful life of each individual system, or are standard life cycle replacement timelines not aligned with actual need?

Our team examined multiple infrastructure program replacement timelines and analyzed operational effectiveness, cost/risk trade-offs, system redundancy, and sustainability, and we recommended life cycle adjustments based on those considerations. We conducted sustainment analyses, using historical failure and inventory data to estimate system end-of-life. Utilizing operational runtime standards and observed runtimes of active systems, we estimated average utilization against manufacturing standards and the life cycle mean. To understand system life cycles and realize operational efficiencies, our team examined secondary power systems and conducted system trade-off analyses between legacy systems and alternative systems to evaluate the right fit.

In this analysis, we define the challenges in conducting both (1) alternative trade-off analyses and (2) life cycle estimation as it is applied to corporate and government operational systems. We examine the utility of sustainment analyses and system end-of-life estimation and the shortfalls associated with relying too much on a single data source for those analyses. We examine challenging data that is either incomplete or which in isolation may draw the wrong conclusions and walk through a step-by-step process for analyzing trade-offs and system life cycles with a detailed use case, and we present the recommendations of our analyses.

2 Sustainment of Aging Legacy Government Systems and Operations

2.1 Government Keeps Systems for a Long Time

Many government agencies when developing business cases estimate a standard life cycle for the analysis and acquisition, estimating program support for a set number of years, planning a certain number of parts spares to sustain the existing infrastructure and system operations, and writing contracts for program support and maintenance. In reality, at many of these agencies, the government actually retains these legacy infrastructure systems much longer than the life cycle estimate, and sustainment of those systems becomes increasingly difficult over time. Manufacturers of consumer-off-the-shelf (COTS) products estimate the total life cycle of the system, anticipating the end-of-life at a specific date. Similarly, third party maintenance contracts extend as long as the system parts are readily available or manufactured, and these third parties limit how long they will continue to maintain legacy systems, either with incentives to sell the government a newer system, or the margins on maintaining a very old system erode the business case and disincentivize industry.

Government agencies still retain some of these older systems due to their unique criticality, mutually exclusive project prioritization, and budget shortfalls. As a result, to continue sustainment of these systems without operational interruption, the agency needs analysts to proactively estimate system end-of-life (EOL) and end-of-service dates (EOS) by part and for the system as a whole. The government's cost/benefit analysis about project life cycle is a balance between system sustainability, cost of a Tech Refresh or system replacement, and continued operational viability. If the public or agency is relying on a systems uninterrupted operation, an accurate estimation of system life cycle and operational risk is critical.

2.2 Government Supports Projects Until the End of Life Cycle

A life cycle is the duration by which a company or government agency will deploy, operate, and maintain a system that provides a capability, function, or service, and it is standardized based on system type, estimated longevity, and a cost trade-off between continued maintenance and replacement cost. Government agencies, in order to avoid loss of service, wasteful F&E spending, or disruption, estimate a life cycle for acquisitions and infrastructure systems based on type (hardware-based systems or software-based systems), precedent (previous deployments and sustainment of similar systems historically), and manufacturer recommendations. Agencies also set life cycles at standard durations which allow them to compare between prospective projects or business case solutions. By establishing standard life cycle durations, agencies can compare and prioritize projects, especially when the decision is mutually exclusive, and they have to decide between one business case and another. During tighter budget cycles, mutually exclusive decisions are more prominent.

The FAA uses a 20-year life cycle for most hardware or mechanical system life cycles, and this allows enough time for the system to "pay off" and provide a return on investment (ROI) to the users and agency, usually without significant sustainability issues. This is sometimes complicated when projects include a mix of hardware and software solutions, and the standard replacement cycle might be shorter for the software than the hardware, or when part of the system is customized for the FAA, meaning limited parts production, extended sourcing challenges, and limited data about system sustainability.

While the 20-year standard is useful in comparing and choosing between business cases for agency solutions, the standard does not universally apply, and agencies should consider customized life cycles or exceptions to the rule.

2.3 Government Sometimes Keeps Projects Operating Well Beyond Intended Life cycle, Sometimes 40-50 Years

As referenced previously, government agencies sometimes keep systems well beyond their intended life cycles, retaining some operational systems 40-50 years. While maintenance on some of these older systems is actually less complex (analog systems are easier to repair, while digital systems sometimes require large component swap outs), aging system sustainment can be difficult to predict.

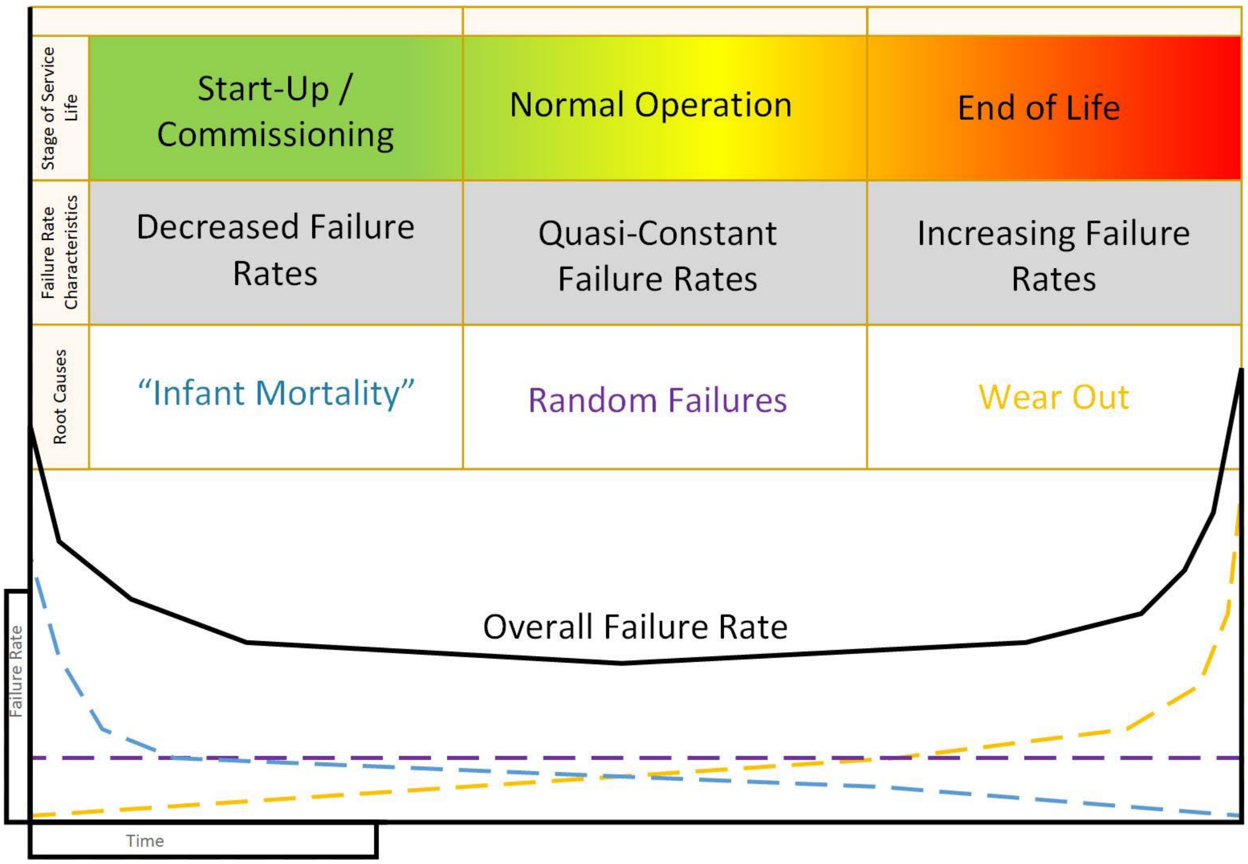

Figure 1: Supply Chain Bathtub Curve

Part failures usually follow a bathtub curve with disproportionately large quantities of failures occurring in a short time at the end of the system's life, it is difficult to predict the end of that bathtub curve. As a result, sustainment risk volatility increases the longer a system operates beyond its intended life cycle.

2.4 How Do We Help Government PM's and Organizations Estimate End-of-Life?

2.4.1 Sustainment Analysis

Sustainment analysts try to estimate parts failures, parts failure growth rates, as depicted in Figure 3, and inventory depletion, but the accuracy of these forecasts decline the longer a system continues in operation, risking an unexpected loss of service outside of government agency risk tolerances. In our 2022 ICEAA paper, *Sustainment Analysis Methodology for Cost Models and Business Cases*, we outlined, defined, and provided examples of how cost estimators, data analysts, and engineers develop comprehensive approaches and complex models to estimate a system's end-of-life (EOL) and EOS, and program managers utilize this data to plan the system replacement with a Tech Refresh or new business case which replaces the legacy system function.

To evaluate the needs of existing infrastructure and determine the system life cycle and best timing of replacing infrastructure capital investments, the agency must measure sustainment needs and establish a balance between the five factors of infrastructure life cycle decision-making:

1) **Cost to Sustain** – What is the cost of sustaining operations with existing operational expenses versus replacing aging infrastructure in the NAS? This is a trade-off between continued and increasing sustainment costs or capital costs of a replacement system. When the cost to sustain the legacy system exceeds that of a replacement over a set number of years of continued operation, the agency can justify replacement and determine its data-driven life cycle.
2) **Ability to Sustain** – At what point will continuing existing operations risk loss of service, or at what point will sustainment without significant investment no longer be feasible? With data-driven end-of-life forecasting and parts shortfalls are conducted, Second Level Engineering and program offices can make parts lifetime buys or replace critical parts before expending all available inventory. At some point, these options cannot be entertained, and the system requires replacement. When the loss of service risk exceeds the tolerance level of the agency or stakeholders (for the FAA – airlines, flying public), the agency can justify replacement and determine its data-driven life cycle.
3) **Timing of Replacement** – When is the best time to invest new capital to replace existing infrastructure? How long can the system be maintained before replacement? This is our life cycle estimate, and sustainment data originates with manufacturer estimates, similar agency historically baselined projects, and data analyses which estimate ability to sustain and useful life.
4) **Sustainment Methods to Extend System Useful Life** – Since agencies tend to keep legacy systems beyond their useful lives, understanding how to sustain these systems to avoid operational failure and provide operational solutions is critical to agency planning. New capital investments can take two to three years to get through investment analysis, so the government must determine sustainment options for life cycle extension to prevent shutting down a system prior to deployment of a replacement system. This includes analyzing parts failures for cause and replacing problem components, making lifetime buys of high-risk parts, and cannibalizing parts from other systems to extend the useful life of a system and delay system replacement.
5) **Cost/Benefit Analysis** – When do the costs of continued sustainment with an increase in parts failure or loss of service risk outweigh the cost of replacement? How do we justify capital investment in infrastructure? Life cycle analysis and estimation determines the breakeven point between sustainment cost and replacement cost, and the justification of the investment can be determined by estimating the operational impact of system failure by the probability of that occurrence. For the FAA, which might mean reduced efficiency and airport capacity due to increased separation of aircraft and inability to utilize equipment.

2.4.2 Failure Analysis & End-of-Life Forecasting

To determine how long an infrastructure system can remain in operation before going end-of-life, analysts estimate the system life cycle:

- Conducting failure analysis of legacy system parts,
- Forecasting failure rates and estimating growth rates,
- Estimating parts scrap (which parts are unrepairable),
- Calculating inventory depletion (until inventory is fully exhausted).

Most of an end-of-life forecast focuses on the sustainment and supply of system parts, the need to make repairs, and the ability to make repairs. To estimate end-of-life, government agencies need to collect and analyze data from supply and demand factors. Not all information critical for end-of-life analyses has a data source or might be collected. The primary factors considered in EOL analyses are the following:

- **Failure rates** – How often a part fails and the system as a result requires repair.
- **Failure growth rates** – How much more frequently are parts failing than they were previously and how much worse parts failures will continue to get.
- **Scrap rate** – Some parts are repairable and are recycled back into the supply. Some percentage of these repairable parts reach "beyond economic repair" (BER), and that percentage represents parts which are scrapped.
- **Inventory** – There is an existing inventory for a starting point of supply and an ongoing inventory which is maintained and represents which parts remain after parts fail and are scrapped.
- **Procurement** – This represents the availability of parts to procure. As systems age, the spare parts to operate those systems may no longer be manufactured and may not be available for procurement to rebuild inventory levels. This accelerates system end-of-life.
- **Substitution** – Some systems can utilize a variety of parts, so when one part can no longer be sourced, another one can be substituted for it instead. However, government agencies often have unique parts or have a high threshold of testing for compatibility and fit before a substitute can be chosen and enter a system architecture.

2.4.3 Failure Analysis and End-of-Life Forecasting Methodology

The risk of a system going End-of-Life (EOL) or reaching End-of-Service (EOS), where the system is no longer operational, are the primary determinates of life cycle estimation. These estimates forecast the limits to system sustainability, especially when parts inventory is finite. The methodology we use to estimate EOS depends on an initial risk assessment of parts demand and supply and the agency risk threshold.

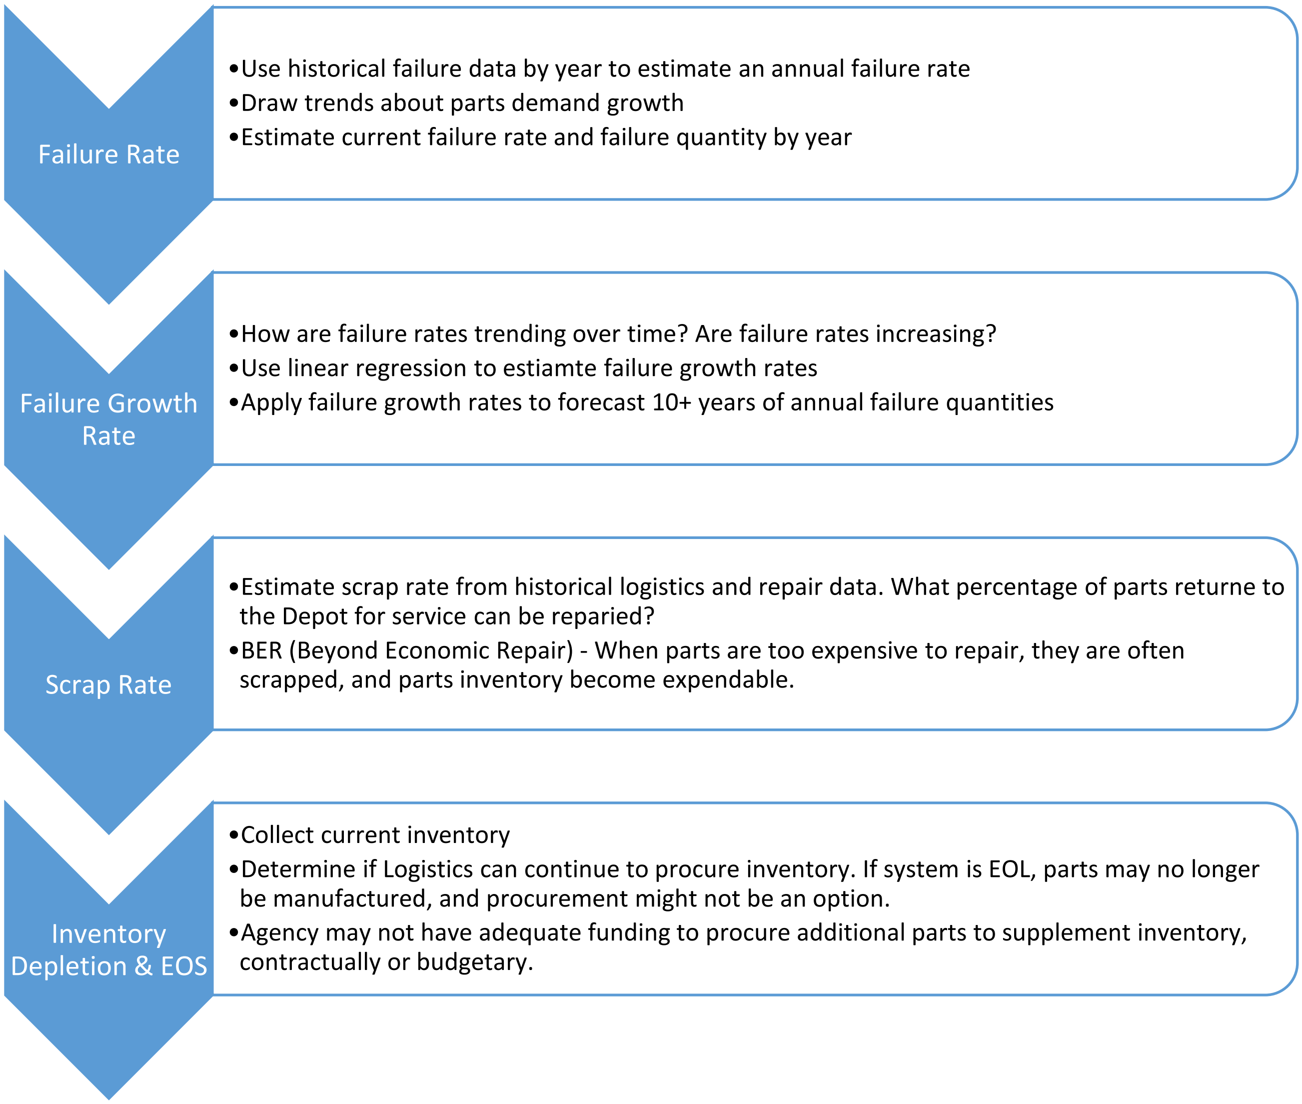

Figure 2: Failure Analysis & End-of-Life Forecasting

2.4.3.1 Failure Rate – Current Observed

To forecast annual failure rates for replaceable units in a system during sustainment analysis, the cost estimator will utilize historical failure data to forecast future parts demand. By observed failure growth trends in historical demand data, analysts can forecast failure growth rates using regression analysis to try to forecast how future parts demand will change over time and, in some cases, replicate the "bathtub curve."

At the FAA, the primary source of failure data originates from the FAA Depot where failures are recorded based on orders in the logistics system or based on recording of parts that are returned after they fail in the field. This failure data approximates parts failures (based on orders for replenishment), if parts were repaired, and if they were repaired and sent back out to the field as replacement spares. With this information, historical annual failures are constructed by using the time of arrival as a proxy for the time the part failed. This assumption is necessary as the actual time of failure is not recorded by the FAA currently. After collecting historical failures from the depot data, the analyst can try to forecast how part failures behave over time.

Metrics such as failure rate and failure growth rate can be estimated using regression techniques on historical annual failure data. These metrics can be utilized to predict how failures will evolve over time and what impact these failures will have on future inventory levels.

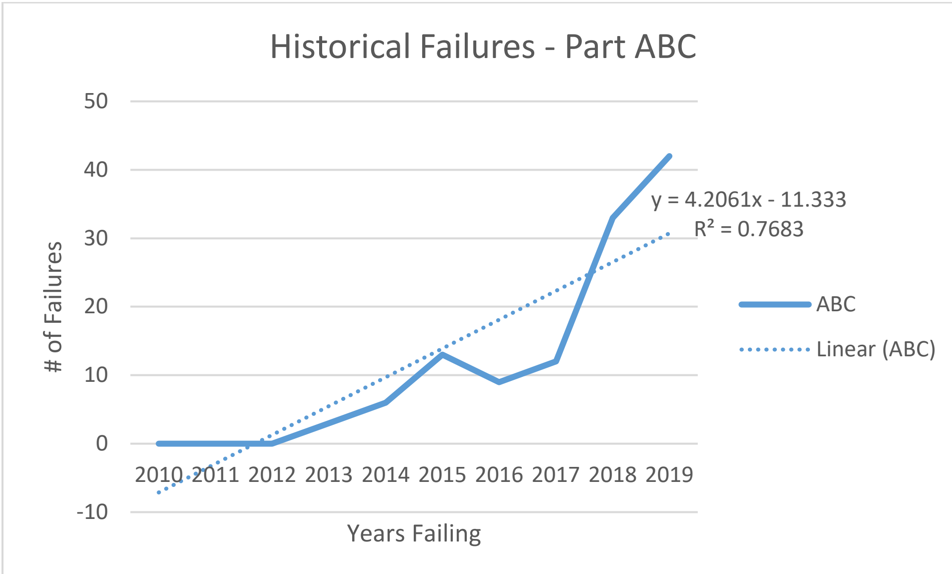

Figure 3: Example of Historical Parts Failures and Regression Analysis Used to Forecast Failure Growth and Future Failure Quantities

The failure rate is measured on a yearly basis and is defined as the number of annual failures over the number operating in service. The number of parts currently in-service can change on an annual basis which affects the number of failures that may occur. Thus, it is imperative to look at the number of failures with the context of the failure rate to differentiate between an increase in failures due to an increase in operational quantity versus innate properties of the part itself. Parts that experience an overall increase in failure rates are said to have failure growth, and this growth rate is estimated from regression.

2.4.3.2 Failure Growth Rate

Using historical failures as a basis to forecast rate of failure, regression analysis generates future year-over-year failures which can be used to estimate failure growth rate. We define the failure growth rate as the slope of the regression line fitted to the historical failure data. This number tells the rate at which failures increase year-to-year.

Each part has unique observable trends in its historical failure data from which the analyst must interpret repeatable trends. This historical failure data on a part-level basis is like a part “personality” related to its quality, durability, and operability within a system, estimating how long it will last in operation – a Mean Time Between Failure (MTBF). Understanding how a part behaves allows a model to predict how it will behave in the future. However, some data may convolute predictability based on misrepresentations of demand (parts ordered not due to failure or delayed failure data collection). This misrepresentation comes in the form of outliers that hinder quality analyses.

To overcome this obstacle, the analyst uses statistical methods and data validation to identify noise in the data.

Once the data selection is complete and the noise is removed, the analyst begins regression analysis to forecast future demand which then can be used to estimate the failure growth rate. The regression line is fitted on the conditional data sets, and the slope of this regression line estimates failure growth.

2.4.3.3 Scrap Rate – Beyond Economic Repair (BER) for Exchange & Repair Parts (E&R)

There exist two types of parts in the FAA cataloguing system:

1) Expendable Parts
    a. These parts cannot be repaired and are thrown away upon failing.
    b. Each failure depletes one unit of inventory.
2) Exchange & Repair (E&R) Parts
    a. **Repair** – These parts are intended to be repaired to original state.
    b. **Exchange –** Broken parts are returned to the Depot for repair and exchanged for repaired parts (or new parts) to restore field inventory.
    c. An Exchange & Repair part (E&R) failure only depletes inventory if it cannot be repaired any further.
        i. A part may be repaired too many times to point where it becomes expendable.
        ii. Sourcing of component parts to make repairs may become difficult over time. This will adversely impact the repairability of E&R parts as systems age.
    d. The proportion of parts that cannot be repaired is known as the "Beyond Economic Repair" (BER) or "scrap" rate.
        i. This rate represents what percentage of failed parts will, on average, not be able to be repaired.

Identifying a part classification as E&R or Expendable is critical for simulating inventory depletions and end-of-life. If a part is E&R, the next step is to estimate this proportion of parts that cannot be repaired. This estimation process is challenging and warrants additional analyses to achieve accurate estimates.

The BER rate for an E&R part is critical to determine the percentage of failures that will deplete inventory. Therefore, it is imperative to acquire the most accurate estimates for the BER rate, as different BER rates can result in wildly different parts depletion forecasts.

2.4.3.4 Estimate Annual Scrap Rate

For parts which are repaired at the depot, government agencies, like the FAA, do not immediately repair a failed part as soon as it arrives. In some cases, they do not repair until all serviceable inventory is exhausted and more parts need to be repaired to replenish inventory. This is one data misinterpretation we must avoid. Since we do not know how many of the parts returned to the depot for repair can be fixed, if we analyze the number of scrapped parts against the number of total failures (the total number scrapped divided by the total number failed), we may overestimate the scrap rate.

To illustrate this notion, suppose 100 failed units of part ABC arrived at the depot in 2020. The depot only attempted to repair 50 of these failed units and had to throw 25 parts away. As common misconception, analysts might estimate scrap rate as 25 scrapped parts divided by 100 total failures and a 25% scrap rate (25 / 100 = 25%). However, only half of these parts were tested, leaving the other 50 parts as potential additions to the scrap count. After testing these additional failures, the scrap rate could increase to as much as 50%, which could drastically change the outcome of the sustainment analyses.

To avoid this underestimate, the total number of units tested rather than units failed is used as the baseline to compare scrap. The outcome of testing a part is: (1) It is repaired, or (2) it is scrapped. Thus, the new baseline to compare scrap to becomes the number of units repaired plus the number of units scrapped. In the above example, rather than a scrap rate of 25/100 = 25%, the new estimated scrap rate becomes 25 / (25 + 25) = 25 / 50 = 50%.

Although not perfect, this new estimation provides a more accurate depiction of how many units of a part should be expected to be scrapped. Still, estimates can vary greatly depending on how many units are scrapped from the untested batch. To capture these scenarios, the analyst performs a sensitivity analysis on the scrap rate/BER rate to observe changes and its impact on sustainment.

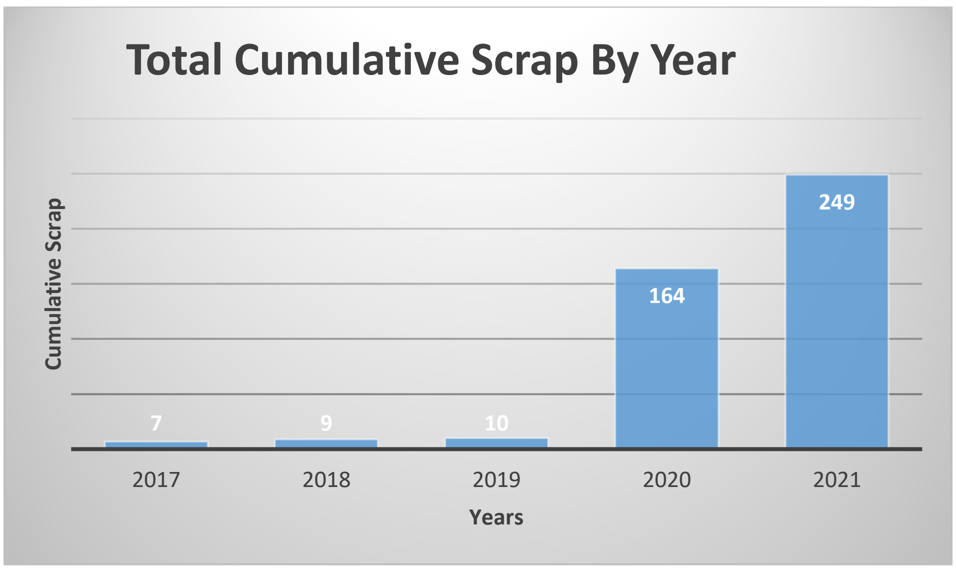

Figure 4: Example of Cumulative Scrap Rates and Impact of Data Collection on Inventory Forecasts

2.4.3.5 Estimating Inventory Depletion

Forecasts of future failures are developed from the historical annual failures of E&R parts in service. First, the historical failures are compared with the in-service numbers to deduce an overall increase in failure rates over time. If failure growth is detected, the analyst conducts a regression analysis of the historical failures.

The regression analysis tells us how many failures will occur on an annual basis up to the projected end year. Alongside the forecasted failures is the estimated scrap rate of the part. Over time as parts get older, the ability to repair the parts may decrease, causing an increase in the scrap rate. Therefore, an estimated scrap rate is forecasted from the historical scrap growth rate point estimate.

Analysts forecast inventory needs and required parts procurements by estimates of forecasted demand and scrap. Parts that cannot be repaired are the proportion of forecasted failures that are beyond economic repair. This number is simply the estimated scrap rate multiplied by the forecasted failures.

Unrepairable parts = Future parts demand X scrap rate

This calculation is conducted each forecasted year, and the total amount of part needs is actualized. This analysis allows for the identification of how many parts will need to be acquired to sustain an operation until the specified year. Rather than procuring an excess number of parts and hoping things go well, the forecasts tell how many parts will be needed and by what date.

Starting from the current year, the number of failures from the regression analysis is actualized and the number of units scrapped is calculated from the scrap rate. This amount scrapped is then used to deplete the current inventory for that year. Going into the next year, the newly depleted inventory becomes the starting inventory and the process repeats. This process allows for the identification of the exact year where inventory levels will fall below a critical threshold. This is the time when a decision must be made to address the inventory issue. Figure 5 depicts the date at which all inventory is exhausted and when a part can be a single source of failure and system loss of service.

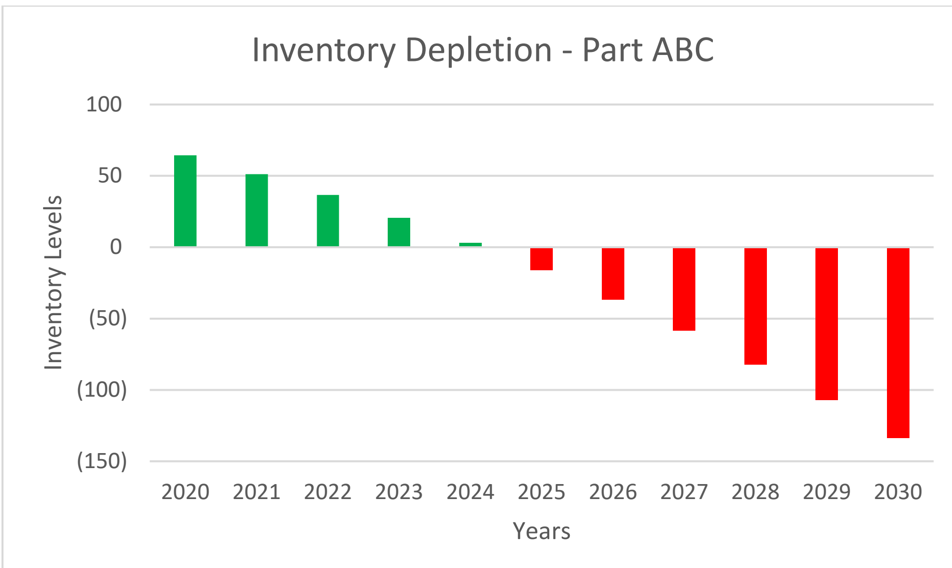

Figure 5: Example of Forecasted Inventory Depletion/Exhaustion

2.5 How Does Predictive Analysis and Sustainment Help Measure "System Life cycle"?

2.5.1 Data-Driven Probabilistic Analyses

Using historical data on annual maintenance data and system failures, root cause analysis of those failures, statistical analyses about system failures or interruptions, and data trend forecasts to forecast operational interruptions, agencies measure and evaluate risk in terms of probabilities. What is the likelihood a 10-year-old secondary power system will fail? What is the probability that the failure of a secondary power system will impact operations?

These probabilities of occurrence have to be measured via use cases to quantify potential consequences of system failures. If an operational system were multiple times per year as the associated infrastructure around it ages, but its secondary power and operational redundancy prevents an operational impact, the agency might elect to continue its operation beyond its intended or useful life. Life cycles are impacted by risk and the agency's tolerance of risk.

2.5.2 Risk Management and Trade-offs

Once we evaluate the sustainability risk of a government infrastructure system and the probability of its failure and maintenance requirements over time, to determine the optimum system life cycle, analysts must assess the agency's risk tolerance, set risk thresholds according to each critical factor, and assess operational impacts to users.

2.5.2.1 Level of Criticality

If the uninterrupted operation of a government infrastructure system is ever compromised, analysts must identify, quantify, and monetize the impact of that potential outage to each user and stakeholder. Breaking down risk events by user group allows analysts to measure each use case separately and find the means by which to measure system outages and those outage impacts on users. Often historical outage events can be used by comparison to estimate future outages.

By evaluating use cases and using specific examples, analysts can measure the level of criticality of a system and the impact of its service operation to the government, users, and stakeholders. To the FAA, the primary stakeholders would be the flying public and airlines.

2.5.2.2 Level of Redundancy & Operational Resilience

The FAA like many other government agencies has multiple layers of redundancy to maintain the highest levels of resiliency and continued operations. By understanding how redundant systems operate and are activated, we can better understand the magnitude and triggers of those risks and set thresholds by which we can measure risk tolerance. By measuring historical system failures of primary and backup systems, we can better measure operational resilience and account for system redundancy in life cycle standards.

When systems have enough layers of redundancy, and when the probability of a catastrophic failure is very low, system failures may have little to no operational impact at all.

2.5.2.3 Risk Tolerance Formula

The agency's tolerated standard of operational impact and the frequency of that impact occurrence is its measure of risk tolerance. The tolerance level is applied to multiple decision variables and use cases to comprehensively evaluate risk.

$$Risk\ Tolerance = Probability\ of\ Failure\ X\ Frequency\ of\ Failure\ X\ Impact$$

2.5.2.4 Cost Estimation in Risk Evaluation

To evaluate life cycle project risk against tolerance levels, one objective basis by which the agency can determine the optimum life cycle duration is using a cost/benefit analysis. This includes measuring project end-of-life and the associated cost/impact of operational failure, the operational cost for sustaining the system and extending the life cycle, and by comparing the cost of replacement to the cost of continuing to sustain the system for several more years.

1) **Estimating End-of-Life** – By estimating the system end-of-life, cost estimators can determine how long an existing system can be sustained until it runs out of spare parts or when spare parts can no longer be procured. End-of-life should correlate with a potential loss of service where at

least some stakeholders or customers can be impacted. In the case of the FAA, EOL can indicate a potential loss of service for the flying public. By quantifying and monetizing this impact, analysts can objectively compare the risk of operational system failure against the cost of replacement to optimize a life cycle standard.

2) **Sustainment** – Cost estimators can also measure and estimate the cost to continue to sustain the legacy system over a period of time. Usually this is measured over a project life cycle, a predetermined time for which the program office compares the legacy case sustainment efforts to the anticipated life of a new procurement replacement system. The cost to sustain the legacy system over the life of the project is used as a baseline of comparison for the new procurement business case. Sometimes, the legacy system cannot be sustained over the project life cycle without significant capital investment. In addition, operational costs to continue corrective and preventative maintenance can become prohibitively expensive.
3) **Cost Avoidance** – Estimating the cost of the legacy case allows cost estimators to have a baseline of comparison to the replacement system procurement. When building a business case for this legacy replacement project, the cost estimator can claim that the cost to sustain the legacy system over the same time period would be the "avoided cost" of the new system. This "avoided cost" would be a quantifiable benefit for the capital investment business case. Used to justify the new procurement, the sustainment cost of the new system over the project life cycle is usually significantly less than the sustainment cost of the legacy system.

3 System Trade-Offs – Analyzing Backup Systems to Optimize Cost, Minimize Risk, and Retain Sustainability

3.1 Overview of System Trade-offs

When organizations estimate system end-of-life and evaluate life cycles, instead of just replacing legacy systems with like-for-like replacements, they consider similar alternatives that might be more efficient and reduce systems costs. Cost estimators measure these improvements as cost avoidance in business case analyses. Those more efficient alternatives might not have the same features, capabilities, resiliency, or operations as the original systems they are replacing. As a result, the program office must consider **program trade-offs**, where a legacy system can only be replaced by a more efficient system if the circumstances are right.

For secondary power systems, these trade-offs can be different durations of power generation in the event that primary power is compromised. It could also include different reliability factors, layers of redundancy, duration of single event runtime, and a variety of other distinct features between the alternatives. Like a traditional alternatives analysis, the analyst most weigh the advantages and disadvantages of each system to determine if the system being evaluated, its location, circumstances, history, and risk tolerance makes it a feasible candidate for an alternative system replacement.

Staying within risk tolerances, the more efficient alternative system might offer major cost savings and other more modern capabilities compared to the legacy system it is replacing, while systems in other locations and circumstances might not be candidates for alternative system replacement at all.

3.2 Secondary Power Systems for Government Systems – Legacy Secondary Power Versus More Efficient Alternative Secondary Power Systems

In our analysis, we considered trade-offs and multiple factors and data-driven thresholds to identify the life cycles of existing secondary power systems and alternative environmentally friendly secondary power solutions. In our evaluation of existing secondary power systems and alternative solutions, we established multiple factors by which we could analyze the trade-offs between systems, and which would be primary candidates for alternative solutions, which would provide multiple benefits to the agency and a significant reduction in operational cost over its life cycle.

3.2.1 Challenges

In the analysis of system trade-offs, we encountered challenges in data collection, data availability, and interpretation. Some data on which we relied in previous failure analyses and EOL forecasting did not exist in the same form, or the failure data without context or interpretation might lead analysts to a misinterpretation of the results. Where in previous failure analyses, we would utilize one database for historical parts demand or system outage records, during this secondary power analysis, we were not afforded the same frequency of data collection, especially due to the nature of secondary power, which is not always running, but which activates when needed. System outage data in its raw form could vary greatly from location to location, depending on the reliability of commercial power, instead of the reliability of the secondary power system we were analyzing.

For these types of trade-off analyses where data is not readily available or complex, analysts must determine how to interpret data from each data source and set statistical thresholds by which to measure and evaluate substitutes.

3.2.1.1 Data is Scarce and System Failure History is Infrequent or Mutually Dependent

Sometimes the data we utilize to collect system failure history or track outages does not collect all of the data that we need to develop trend analyses and interpret data to forecast future behavior. When this data is incomplete, the analyst must find another means by which to supplement that information, whether through meeting with subject matter experts (SMEs), correlating data from other sources to fill in the gaps, or examining system and user behavior to better predict outcomes.

3.2.1.2 Hard to Collect Historical Data on System Failures to Estimate Sustainability of Existing System and Conduct Failure Analysis

Secondary power systems, unlike primary systems, are not constantly running. They turn on when the primary power is interrupted, and most data while idle does not indicate system operability. For instance, if commercial power to a government or commercial system does not fail over a six-month period, the secondary power system would remain mostly idle during that time, and little to no data about its operability would be available. In contrast, whenever a primary system in operation turns off or operates outside of its threshold, alarms, operations, and users are informed and impacted. Those events are easily recorded in databases to track not only the outage but the reason for its occurrence, its root cause.

3.2.1.3 Limited Data Points – Periodic Maintenance as Means of Data Collection

For a secondary power system, both the primary power and the secondary power have to fail for a main recording of secondary power system failure. This is an exercise in probabilities, and some secondary power systems support equipment on very stable and reliable commercial power grids, meaning that the probability of a secondary power outage during a commercial power outage is extraordinarily remote. The secondary power system may still have the same reliability as others, but if it is never required to turn on, we cannot collect operations data to predict its reliability or useful life.

The other way to measure secondary power reliability is to visit the site for periodic regularly scheduled maintenance and test the system for operability. This is done periodically and as periodic maintenance is planned, but it does not offer comprehensive data points for system reliability and failure data. If in between visits, the secondary system becomes inoperable, and the primary system does not fail, the organization may be unaware that there is a system failure, and the data will not reflect it.

3.2.1.4 Complicated Risk Analysis

Failure analyses with these types of complexities and layers are not as straight forward as primary systems, so to get a full picture and to forecast system reliability, risk, and life cycle sustainability, the analyst has to combine:

1) Reliability of the primary system, in this case commercial power,
2) Reliability of the secondary system being analyzed, in this case secondary power,
3) Levels of redundancy, where the primary system has another primary system as redundancy.

3.2.1.5 Risk Analysis and Sustainability Can Be Analyzed as a Combination of Many Factors, Not Single Point of Failure

Without one data source to conduct a full risk and reliability analysis for forecasting system trade-offs and life cycle estimation, the analyst must consider multiple factors in aggregate to develop risk recommendations and defendable conclusions. For our example of secondary power where performance and failure data are not as robust to forecast end-of-life and life cycle extension estimations alone, we use a combination of factors and steps in our analyses.

For system trade-offs, each factor we analyze has a threshold of risk tolerance which determine whether or not an alternative secondary power system qualifies. We considered several factors, including:

- Primary system that secondary power supports
- Collocation of secondary power
- Types of secondary power
- Existing power redundancies
- System redundancies
- Primary power outage frequency
- Secondary power failure analysis
- Maintenance travel time
- Impacts of system outages
- Operational runtime
- Combined load on power systems
- Secondary power limitations

For life cycle analysis, we consider many of the same factors, but instead of trade-offs between the secondary power capabilities, we analyze factors that help us estimate end-of-life and increases in system failures, and we measure these against organizational risk tolerances and impacts. In Figure 6, we define the process flow in a decision tree for our trade-off analysis.

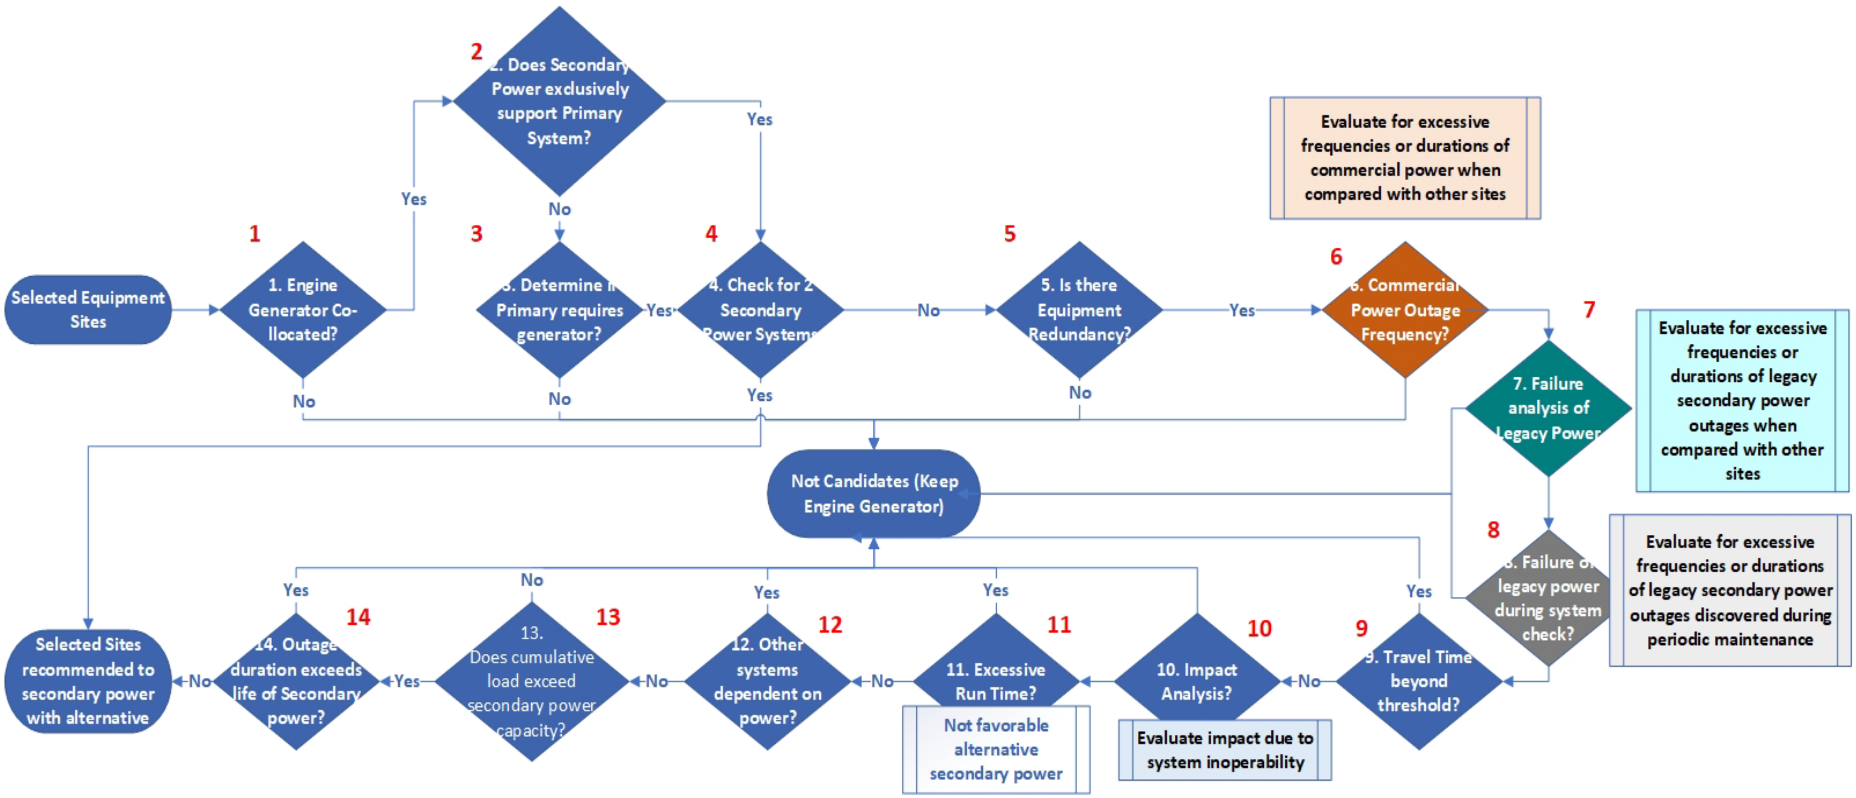

Figure 6: Trade-off Analysis Decision Tree & Steps

3.2.2 How to Measure Trade-Offs – System Architecture & Factors

3.2.2.1 Examine/Document Pros and Cons of Each Secondary System

In our example in Figure 6, we analyze the trade-offs between different secondary power systems by measuring the pros, cons, capabilities, and constraints of each system. What are the limitations for each system? In steps 1 and 2, we determine if the secondary power system is collocated with the selected equipment we are analyzing. If it is not collocated, the analysis of the secondary power is irrelevant for our selected systems. If the secondary power supports some other equipment, we will include it in a different analysis for secondary equipment.

To make sure the system is a candidate for alternative secondary power, we analyze if the current secondary power system is absolutely required. In step 3, if the current secondary power cannot have a substitute, we do not go further in our analysis.

3.2.2.2 Redundancy

As we continue our trade-off analysis for system resiliency and life cycle, we determine what level of system redundancy and backup system redundancy we have. In step 4, we check for a redundant secondary backup system. It is unlikely that a system would have two backup power systems, but under certain requirements and considerations, commercial and government systems may have layers of redundant secondary power. There could be circumstances that would result in duplicate secondary power. Perhaps a secondary power system was replaced in the past with a more efficient alternative secondary power source, but instead of removing the original secondary power system, it was left in place due to budgetary constraints or project prioritization. This analysis would help organizations to identify opportunities to remove extra layers of redundancy and right-size system infrastructure. For cost estimation, this presents an opportunity for cost avoidance analysis.

For the primary equipment for which we are providing commercial and secondary power, we next check to see if that equipment has a backup as well. This equipment redundancy will add resilience to the primary system, so even if the commercial power fails and the secondary power either fails or stops providing backup service, a redundant system that does not draw off of the same power grid will be unaffected. Understanding how the systems work will help the analyst evaluate secondary power system trade-offs and will facilitate life cycle estimation.

Multiple layers of redundancy would all have to be impaired before any operational impact, including:

- Primary commercial power outage
- Secondary power system outage
- Any duplicate secondary power system outage
- Redundant operational system outage
- Redundant system commercial power outage and any secondary power outage

All these layers of redundancy provide an extra layer of resiliency and reduced system risk, which are considered when estimating life cycles and operational risk tolerances.

3.2.2.3 Failure Frequency Analysis of Primary/Commercial Power

As we detailed in Section 2.4 for failure analysis, we can utilize failure analysis to predict future failures. When evaluating a secondary system that operates only when the primary system fails, the utility of that failure analysis is limited. To calculate the probability of an operational impact, in our analysis, we have to understand the reliability of primary power, secondary power, and any redundant systems to estimate an impact to operations. Comparing these aggregate probabilistic analyses to measure risk is challenging and requires scenario and sensitivity analyses.

In step 6, we measure commercial power outage frequency. This allows us to understand the frequency upon which secondary power is required to operate. Using historical data, we analyze data trends and conduct a regression analysis to estimate the frequency of commercial power outages and if these power outages are increasing in frequency over time. If we observe trends we can model, we apply linear regression or other types of regression (log normal, exponential) to model the future changes.

Commercial power does not act uniformly at different geographic locations. Remote locations may have more power interruptions than those close to large cities or on emergency power grids (like those shared with an airport or hospital). Power supplied to a remote mountain might be less reliable. Commercial power on the grid is not uniform over time either. Some areas change their mix of power sources, and population increases can impact grid stability as can weather phenomenon. All these factors must be considered when modeling power reliability.

3.2.2.4 Failure Duration Analysis of Primary/Commercial Power

In additional to frequency of power outages, we estimate and analyze data regarding the duration of historical outages and project whether or not there is a likelihood of these outage durations continuing. We conduct the same probabilistic analysis and forecast outage durations as average and maximum observed outages. Maximum outages are important because if agencies do not tolerate any operational interruption, the secondary power system being considered must supply power as long as any previously observed power outage. These types of factors and deciding between average outages and maximum outages impact our evaluation threshold.

3.2.2.5 Failure Analysis/ Failure Frequency of Secondary Power System

In step 7, we analyze the frequency of failures of secondary power. The challenge to using historical data on secondary power outages is that these outage occurrences are largely dependent on the reliability of the commercial or primary power, which, in part, is dependent geographically on the area of the grid it operates. The company or organization monitoring performance of the secondary power may only collect data on secondary power reliability and failures when commercial power goes out. If the commercial power never fails, we may not have data on the reliability or probability of inoperability of the secondary power system.

Organizations have other means of measuring the reliability of secondary systems, like power, utilizing remote monitoring equipment. On power systems which do not turn on unless required, periodic maintenance is a good indicator of operational reliability as maintenance teams test equipment when they perform periodic maintenance checks at regular intervals. As indicated in step 8, this data should supplement the pure failure data which is recorded when both primary and secondary systems fail.

3.2.2.6 Travel Time to Repair

In large companies or government organizations which have large geographic maintenance coverage areas, the distance from the equipment to the maintenance center is another critical factor of evaluation for the life cycle analyst. If the travel distance to a site is very long (as compared to average (mean) travel times), the potential risk for longer duration outages and system operational impact is much larger than those systems which are collocated with the maintenance centers.

In our analysis, time is a factor for how long secondary power will operate and provide necessary power for the system. If there is a power capacity distinction between the current secondary power system and an alternative secondary power system being considered, travel time to a site plays a disproportionately large role in our power system trade-off analysis. In our analysis, we developed thresholds of travel time beyond which we used as a cutoff between the current secondary power and the alternative secondary power being considered.

3.2.2.7 Impact Analysis

In step 10, we conduct an impact analysis in the event that multiple layers of system and power redundancy all simultaneously or sequentially fail to operate or stop working, and we estimate that impact of inoperability on stakeholders who depend on that system's operation.

Criticality, shortfall analysis, and impact to operations due to system inoperability are measurable and quantifiable. Analysis teams use simulation, process analysis, and shortfall analysis techniques to measure differences between operations with and without systems to multiple users and external stakeholders. Some systems have much larger potential impacts than others. Therefore, by measuring the potential impact of outages, we can estimate the dollar impact of an outage and calculate its cost to the government, the public, and stakeholders. These are risk trade-off calculations, and a certain threshold of cost risk is established to evaluate the tolerance of outages by equipment type. A risk intolerance might dictate which secondary power system is acceptable or viable.

In our analysis, we separated impacts into quantified ranges. For those estimated to be medium or high impact we would not consider secondary power alternatives. Any risk from alternative secondary power that might increase the risk of failure would not be tolerated. Where outage impacts were lower, we developed a sensitivity analysis to consider which impacts would be within the agency's risk tolerance.

3.2.2.8 Runtime

Another critical measure for trade-off and life cycle analyses for mechanical systems is runtime, which we used for our secondary power trade-off analysis in step 11 of Figure 6. Runtime can be used in life cycle analysis as a finite measure of total system life, where mechanical systems often have an estimated total number of hours of runtime during their useful life before major repairs have to be made. It can also be used for trade-off analysis for secondary power systems by correlating total runtime to the number of primary power outages to estimate average outage duration. If the runtime significantly exceeds the overall average, this may be an indication that outages have a long duration, and any alternative secondary power system we are evaluating in our trade-off analysis must have the capacity to operate as long as these estimated historical outages.

4 Life Cycle Estimation – How Do We Use Incomplete Data to Estimate System Life Cycles?

4.1 Standard Life cycle Durations as Means of Investment Comparison

The government has provided standard life cycle durations for which the agency could compare capital investments and acquisitions, estimating the cost and benefits of each business case over the same 20-year duration for a business case and utilizing specific finance metric thresholds: Net Present Value (NPV), Internal Rate of Return (IRR), Payback, and Benefit/Cost Ratio (B/C Ratio).

By planning to standard 20-year life cycles, the government can compare investments equally. However, using an average or standard life cycle might not be an accurate measure of the infrastructure system's life cycle and sustainability. By not aligning the standard life cycle used for business case evaluation with the actual system useful life, the agency risks a wasteful premature reinvestment, replacing the system with another acquisition before actually necessary. Or, the agency might risk replacing the system too late, resulting in possible operational interruption.

4.1.1 Risking System Loss of Service

If a company or government agency uses a standard but not-data-driven life cycle for an operational system, that arbitrary life cycle might not align with the actual useful life of the system. If the system reaches EOL and incurs failures at a frequency or impact that exceeds the agency's risk tolerance, it could risk loss of service. The organization would have waited too late for system replacement, and the operational consequence could be very expensive and impact many users.

4.1.2 Premature System Replacement and Waste of Agency F&E Dollars

If a company or agency follows a life cycle standard replacement schedule for an operational system that is well before the actual useful life of the system, they might replace the system too soon. The consequences of premature replacement would be:

1) Given a constrained capital budget where the capital projects required exceed that of funding available, the agency must prioritize some projects at the expense or delay of other projects. If the first project is not absolutely necessary and is replaced too soon, its funding comes at the expense of another project which might have more immediate value or urgency. Delaying the second project delays real value to stakeholders, and this could be avoided using data-driven life cycle analysis.
2) If a system has an actual data-driven life cycle and useful life greater than the arbitrary standard life cycle imposed on capital projects, premature system replacement wastes capital dollars unnecessarily and might compromise a more strategic objective or enterprise solution.

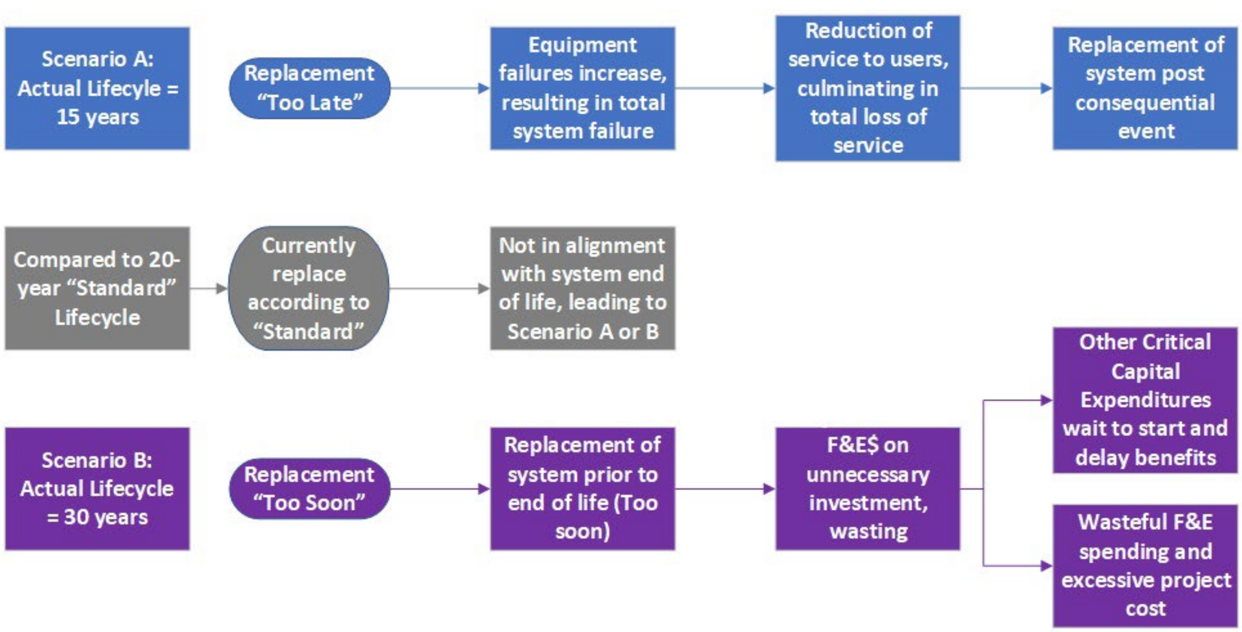

Figure 7: Consequences of Misalignment of System Life cycle Standard with System Actual Useful Life

4.2 Life cycle Extension Factors

4.2.1 Failure Analysis with Scarce or Incomplete Data Is a Challenge to Life cycle Estimation

When considering life cycle extension for government infrastructure systems, we often utilize failure analysis of the system's components to estimate how long we can sustain the systems, especially for systems which are near or at end-of-life (EOL), where parts cannot be procured and are no longer manufactured. As we defined in section 2.4, we measure and analyze the following factors during parts failure and inventory analyses to measure system sustainability and life cycle:

- Failure rates
- Failure growth rates
- Scrap rate
- Inventory
- Procurement
- Inventory depletion

Some systems on which we conduct failure analysis and try to estimate end-of-life and the full system's data-driven life cycle have incomplete data or limited data which can cloud our conclusions. Secondary power systems are an example of a type of system that no matter how comprehensive the data collection, failure data is limited. To analyze system reliability, we need to record when the secondary power system is operating and when it fails to operate for any reason.

The challenge with secondary power failure data as with other systems which are utilized in the event of primary power interruption is that the secondary power only turns on when the primary power fails to operate or has a service interruption. For most operational infrastructure systems at agencies and commercial entities, the systems are constantly running, and any interruption of their operation would be recorded, analyzed, and a cause attributed. All this data would facilitate our life cycle analysis and provide critical details to system resiliency, how to improve resiliency, and how resiliency might erode over time.

For secondary power systems which only turn on when primary power is interrupted, if primary power is very reliable and interruption exceedingly infrequent, there would be no data on the reliability or failure of the secondary power. Some systems at agencies have status monitors that alert maintenance and oversight operations immediately if there is a failure of that system. This type of monitoring works for many redundant systems as well, but for secondary power, unless the company or agency can remotely turn on the secondary power at regular intervals, its reliability is unknown and data not generated until regular periodic maintenance and field testing is conducted, which may not be frequent.

This absence of data for secondary power systems makes system failure analysis, modeling the trends of resiliency over time very difficult to calculate. As a result, an agency or company with these types of secondary power systems must either increase system testing to generate data they can analyze regularly or use other factors about sustained system operability to estimate data-driven life cycles.

4.2.2 Historical System Failure Analysis

When we conduct failure analysis for a hardware-based system to estimate its end-of-service (EOS), we estimate the point at which we would run out of spare parts and critical inventory and compromise continued operation. This is the end of its natural system life cycle (See Figure 8). If the system's risk of EOS exceeds the agency's risk tolerance, accurate lifecycle estimation can help avoid that outcome.

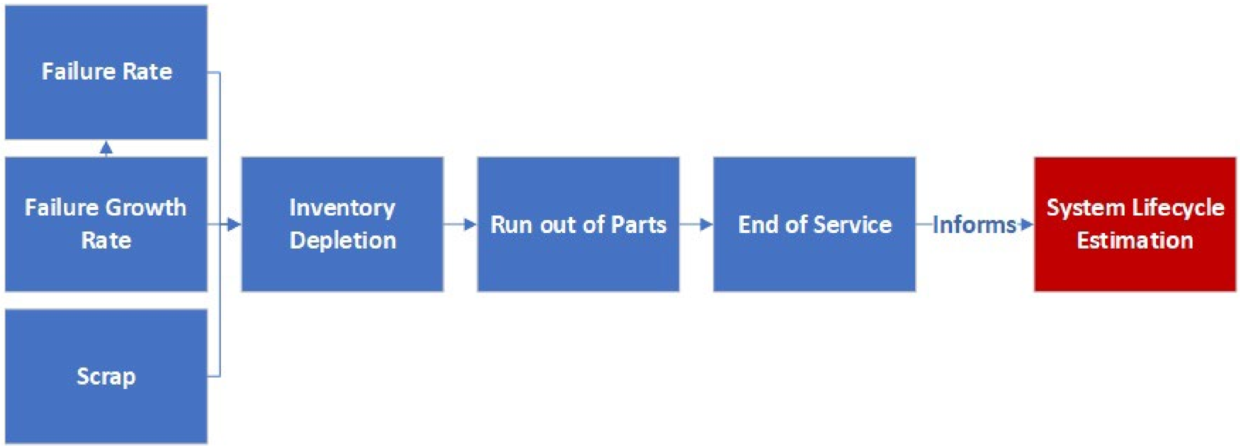

Figure 8: Life cycle Estimation – Consequences of Misalignment of System Life cycle Standard with System Actual Useful Life

If the risk tolerance for failure of any kind is low and the impact of operational failure high, estimating overall system failure frequency using regression analysis of historical failures is critical. In this case, life cycle would be estimated not by the full depletion of parts' inventory but, instead, by when failure frequency and duration exceed the failure thresholds determined by the agency's risk tolerance, highlighted in Figure 9.

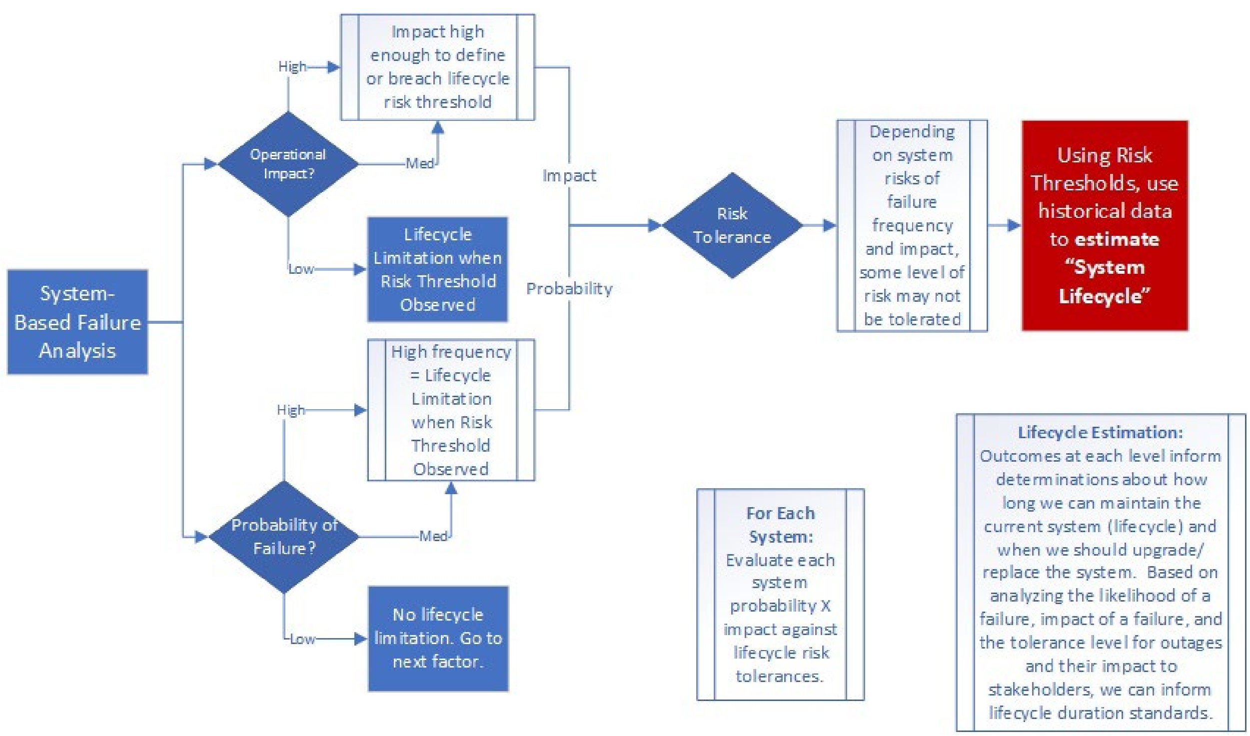

Figure 9: Life cycle Estimation – System-Based Failure and Impact Analysis vs. Risk Tolerances

4.2.3 Analyze Run Times or Other Measures of Operation

Manufacturers of infrastructure and operations equipment usually provide recommended life cycles that estimate the intended life of systems that they sell. Even systems created specifically for a customer or government agency will have some context to a portfolio of like equipment useful to estimate its life cycle. Manufacturing companies estimate these life cycles based on historical operations of similar systems they sell. These standards of reliability and useful life are measured in years of life, hours of run time, or a similar measure.

4.2.3.1 Age of System as Life cycle Indicator

For our analysis of secondary power, we utilized both historical years in context of that first standard life cycle from the government and the relative performance and reliability of those systems at different age benchmarks. That way, if age of the system gradually impacted system reliability and failure frequency, we could set statistical standards at each age benchmark to measure risk.

4.2.3.2 Runtime as Life cycle Indicator

We also utilized runtime for some secondary systems in hours of operation as that is recorded via data and sensors. Runtime measures how much wear is imposed on the system. This level of wear can vary dramatically between systems. For instance, if two people were to buy the same manufacturer and model car from a car dealer the same day and drive the same number of miles, but driver A would drive mostly highway miles and run the car for a much fewer number of hours, arriving to work quickly and turning off the car, and driver B would drive in the city, taking twice as long to reach his or her destination, operating the car in idle in traffic every day, the car from Driver B would exhibit greater wear and would likely incur more frequent operability issues or failures than Driver A's car. In fact, the life cycle for Driver B's car would very likely be shorter than that of Driver A. This same concept of runtime is measured to determine the wear and tear on corporate and government equipment and can be a very valuable measure from which we can estimate risk for systems which far exceed the mean runtime compared to the rest of the pool of systems in operation. Setting a specific threshold of runtime can be used as preventative maintenance and allows program offices and second level engineering to estimate how soon a system needs to be replaced. In our analysis, we used this measure to determine system life cycle extension and a replacement date.

4.2.3.3 Age of Systems that Qualify for Consideration

One of the considerations for life cycle analysis is to bound the pool of candidates for life cycle extension or analysis by specific factors and limit scope. This is especially important if an organization is analyzing an especially large pool of candidate systems, and the program office or engineering team wants to consider a subset as an initial prototype or to prioritize systems which are most important.

The purpose of life cycle extension analysis is two-parts:

1) **Repair Prioritization** – Organizations (maintenance, technical operations, second level engineering) are trying to develop data-driven life cycle to forecast and prioritize work needs, save significant dollars by not prematurely replacing systems, and identifying the greatest replacement/repair needs based on their useful life.

2) **Timing of Investment and System Replacement** – Other organizations, program offices, Investment Planning & Analysis (review organizations which evaluate mutually exclusive business cases), and decision councils, like the FAA's Joint Resources Council (JRC), are interested in accurately estimating the timing of system replacement as a Tech Refresh or Sustainment (like-for-like system acquisition) investment and to avoid loss of service if the system fails.

What systems we analyze by age matters. If the intended or standard life cycle for comparative purposes is 20 years, analyzing deployed systems that are 3-5 years old will have little utility to serve either purpose. They are not priorities for replacement or repair to Technical Operations or Second Level Engineering business units unless they are not needed at their current locations and could be relocated to another location with greater need.

Age of the candidate pool for life cycle extension also matters when evaluating systems according to sustainment and end-of-life risk. Replacement analysis for systems less than 10 years old results in acquisitions or Tech Refreshes in 10-20 years in the future. These decisions are not a priority for program managers and acquisition teams now; there are many more nearer term decisions that need to be evaluated, so choosing the age pool of candidate systems is critical for life cycle analysis and portfolio prioritization.

4.2.4 Risk Analysis Factors that Adjust Risk Tolerance Levels

Outside of failure analyses and estimation of wear, other risk tolerances which might impact the potential life cycle measure of infrastructure systems include factors that impact operations or prevent operational impact.

4.2.4.1 Redundancy – Back Up Systems

Many government and FAA systems have multiple layers of redundancy. If the equipment covered by continuous power has very little operational impact if an outage occurs, the risk tolerance may be higher and allow a higher number or duration of failures. If in the event that primary or backup power fails, if the system has a backup or redundancy that will take over and prevent operational interruption, this adds to the resiliency and reduces the dependency on failure risk and primary system resilience. For life cycle estimation, system redundancy and backup systems reduce overall operational interruption risk, and agencies could be more willing to extend life cycles of systems with redundancies than those which have a single source of failure. The risk tolerance in the case of system redundancy is higher.

4.2.4.2 Multiple Layers of Redundancy Adds Operational Complexity

Although system redundancy avoids a single point of a failure and adds system resiliency, multiple layers of redundancy and at different failure points can be complicated. Some backup systems are only activated under certain circumstances, perhaps triggered by the duration of the primary system failure.

Other backup systems provide similar functionality of the primary system, but with a shorter range or a reduced functionality. Even though those backup systems will allow uninterrupted operations, those operations may have to be constrained, and in the case of an organization like the FAA, this could impact flight operations and air traffic control. If the backup systems or level of redundancy does not provide the equivalent level of operation as the primary system, the analyst needs to measure that risk to operations and monetize that impact in terms of cost to the organization, the public, and stakeholders.

Similarly, some backup systems will not operate indefinitely. When considering secondary power systems, some backup power will only last for hours or days. If an outage is anticipated to last several days, the agency must consider the impact of loss of service or a contingency work around to restore power.

Finally, not all layers of redundancy are created equal, and they may perform in practice differently than intended. For instance, one system may have full functionality for mechanical operation and data transfer operating on primary power; however, when operating on secondary power, the full functionality might be reduced. Analysts and organizations must estimate this impact to determine if the risk of power failure is within operational risk tolerances or adjust secondary power accordingly. This is important in life cycle analysis and risk analyses in right-sizing systems. If the risk tolerances require fully functional systems and all capabilities without interruption, the reliability of the primary system must remain at the same level as when it was installed, and an increase in system failures associated with age cannot be tolerated. Accordingly, we cannot extend the life cycle of those systems without compromising risk standards.

5 Conclusions and Use Case Results

5.1 Secondary Power Trade-off Analysis

In our trade-off analysis between legacy secondary power systems and an alternative more efficient replacement where we analyzed the pros and cons, cost and benefits, design constraints, and system-specific requirements, we recommended the replacement of 34% of legacy systems with a more cost-efficient modern replacement. These systems provide the same reliability and are rightsized to the equipment they support.

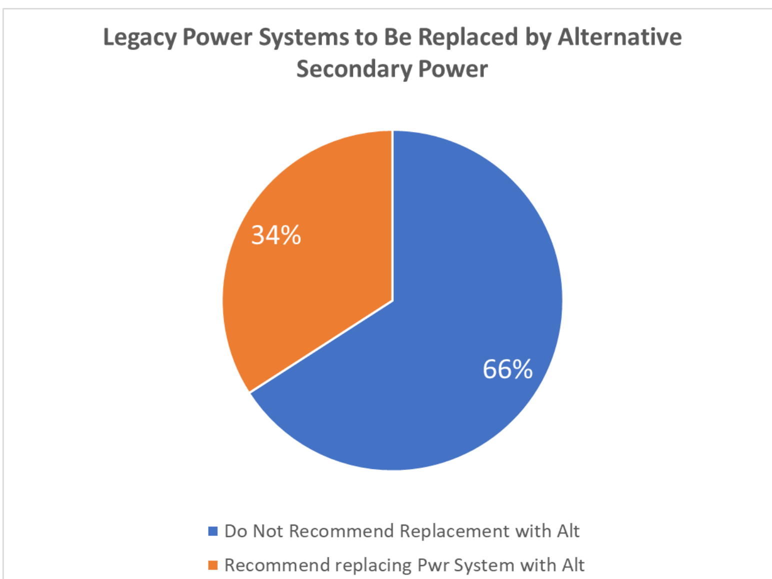

5.2 Secondary Power Life Cycle Extension Analysis

In the life cycle extension analysis of secondary power systems, the analysis team considered a sample size of secondary power systems supporting a specific equipment type and those systems which were within five years of the standard life cycle which systems currently follow. Analyzing multiple factors for resiliency, failure risk, and continued sustainability, the team analyzed system data against risk thresholds of secondary power failure frequency and frequency forecast, cumulative runtime hours, and impact analysis and concluded that 59% of analyzed systems could be extended another five years without major investment, increased sustainment cost, or excessive frequency of failure. More than 50% could also be extended a total of 10 years.

5.3 Conclusions

When companies and government agencies implement infrastructure and operational systems, they plan to maintain those systems for an estimated useful life, the duration for which they can sustain operations without a significant increase in system failures and without operational impairment. During project acquisitions, government agencies utilize a standard life cycle to estimate program cost, plan replacement timing, and compare business cases to one another in a cost-benefit analysis. Based on analyses of system trade-offs and life cycles, using historical data and risk thresholds, life cycle replacement timelines are not always consistent with system sustainability or cost-effective. In our use case, we concluded that standard life cycle replacement schedules are convenient for comparative business case analyses but are not aligned with systems' useful life or optimized to reduce operational risk and costs.

Our team examined multiple infrastructure program replacement timelines and analyzed operational effectiveness, cost/risk trade-offs, system redundancy, and sustainability, and we recommended life cycle adjustments based on those considerations. We conducted sustainment analyses, using historical failure and inventory data to estimate system end-of-life. Utilizing operational runtime standards and observed runtimes of active systems, we estimated average utilization against manufacturing standards and the life cycle mean. To understand system life cycles and realize operational efficiencies, our team examined secondary power systems and conducted system trade-off analyses between legacy systems and alternative systems to evaluate the right fit.

In this analysis, we define the challenges in conducting both (1) alternative trade-off analyses and (2) life cycle estimation as it is applied to corporate and government operational systems. We examine the utility of sustainment analyses and system end-of-life estimation and the shortfalls associated with relying too much on a single data source for those analyses. We examine challenging data that is either incomplete or which in isolation may draw the wrong conclusions and walk through a step-by-step process for analyzing trade-offs and system life cycles with a detailed use case, and we present the recommendations of our analyses.

Standardized life cycles are useful for comparing mutually exclusive business cases by standardizing life cycle timelines by which we measure project costs and benefits. However, operational life cycles are most efficient when they are estimated against the historical useful life of individual systems, where agencies and corporations can maximize the useful life of systems, minimize costs, and optimize replacement timelines. In our use case, we recommended more than a third of candidate systems be replaced by a more efficient alternative system and that more than half of operational systems we analyzed could apply a life cycle 5 or 10 years longer than the current standard. By taking a methodical and comprehensive approach to system trade-offs and life cycle analyses, we can plan and adjust system life cycles to optimize budgets, prioritize maintenance backlogs, and preserve funding for other capital projects.