



FAA
Commercial Space
Transportation

Guide to Probability of Failure Analysis for New Expendable Launch Vehicles

Version 1.0

November 2005

HQ-032105

Guide to Probability of Failure Analysis for New Expendable Launch Vehicles

Version 1.0

November 2005

Federal Aviation Administration
Commercial Space Transportation
800 Independence Avenue, SW, Room 331
Washington, DC 20591

NOTICE

Use of trade names or names of manufacturers in this document does not constitute an official endorsement of such products or manufacturers, either expressed or implied, by the Federal Aviation Administration.

TABLE OF CONTENTS

1.0 INTRODUCTION	1
2.0 BACKGROUND	1
3.0 CANCELLATIONS.....	2
4.0 DEFINITIONS	2
5.0 PERFORMANCE STANDARDS	3
5.1 Account for launch vehicle failure probability in a consistent manner.....	3
5.2 Incorporate accurate data, scientific principles, and valid methodologies.....	4
5.3 Account for the outcomes of all previous flights.....	4
5.4 Account for changes to the vehicle configuration and other factors.....	5
6.0 PROBABILITY OF FAILURE ANALYSIS METHODOLOGY	5
6.1 Vehicle Design with Fewer Than Two Flights Completed.....	6
6.2 Vehicle Design with at Least Two Flights Completed.....	7
7.0 FAILURE PROBABILITY ADJUSTMENTS.....	8
7.1 Flight history.....	8
7.2 Corrective actions.....	8
7.3 Engineering approach.....	9
8.0 SUMMARY	9
APPENDIX: INDEPENDENT ASSESSMENT	10
NOTES.....	11

1.0 INTRODUCTION

This document provides definitions of key terms and performance standards for new expendable launch vehicle (ELV) failure probability analyses. A performance standard permits a launch operator to continue to employ alternative, potentially innovative methodologies if the results satisfy the performance standard. Current practice at the Federal launch ranges includes multiple methodologies for determining the probability of failure for new ELVs. The Federal Aviation Administration (FAA) here presents an acceptable method, but not necessarily the only method, to demonstrate compliance with the performance standard. The method suggested here is also intended to illustrate an acceptable level of fidelity for new ELV probability of failure analyses.

2.0 BACKGROUND

Recognizing the central importance of probability of failure estimates to launch risk analyses, the FAA, Office of Commercial Space Transportation (AST), and U.S. Air Force (USAF), through the Common Standards Working Group (CSWG), developed a guide for conducting valid probability of failure analyses for new ELVs¹. These analyses are crucial to ensuring public safety for launches that employ risk management or a combination of hazard isolation and risk management.

A probability of failure analysis for an ELV produces an estimate of the likelihood of occurrence of a hazardous event. A probability of failure analysis is an essential element of any launch risk analysis, such as the debris risk analysis required by 14 CFR § 415.35(a). Mathematically, risk is the product of the probability of occurrence and the consequences of a hazard to a population or installation. For launch risk analysis, the occurrence is a failure of a launch vehicle. In general, a launch risk analysis allocates the probability of launch vehicle failure to flight times and failure modes; however, allocation of the probability of launch vehicle failure exceeds the scope of this guide.

This guide is consistent with practice at the Federal launch ranges and represents a major improvement over the rigid method proposed in the FAA's Licensing and Safety Requirements for Launch, Notice of Proposed Rulemaking, the October 2000 NPRM [Federal Register, vol. 65, no. 207, October 25, 2000, pp. 63922 to 64123]. The October 2000 NPRM proposed to assign a fixed failure probability of 0.31 for the first 15 flights of a launch vehicle and a failure probability of 0.10 for the next 15 flights. For a launch vehicle with 30 or more flights, the October 2000 NPRM proposed that a launch operator "use the empirical failure probability determined from the actual flight history."

Given the set of practical constraints underlying an assessment of ELV failure probabilities, this guide provides a flexible approach that is technically valid and responds to the needs of the ELV industry. The focus is more on providing guiding principles than "recipe lists." This guide is intended to provide a commonly accepted framework for probability of failure analyses and should be useful to anyone who performs or evaluates launch risk analyses for new ELVs, including Federal range

safety personnel and launch operators. The FAA charted two independent assessments of this guide; the details of which are provided in the appendix.

3.0 CANCELLATIONS

This document replaces the “FAA/AST Draft FAA Guidelines on Probability of Failure Analysis for New Expendable Launch Vehicles,” dated September 2004, in its entirety.

4.0 DEFINITIONS

The FAA uses the following definitions to evaluate any probability of failure analysis for an ELV.

Flight. For probability of failure analysis purposes, flight begins when a launch vehicle normally or inadvertently lifts off from a launch platform.

The FAA has an existing regulation (14 CFR §401.5) that defines the end of flight as follows: “For purposes of an ELV launch, flight ends after the licensee’s last exercise of control over its launch vehicle.” Therefore, when using this guide, the “flight” history of a subject vehicle should include all the in-flight failures and successes that occur from liftoff until after the licensee’s last exercise of control.

Liftoff. Liftoff occurs when there is any motion of the launch vehicle with respect to the launch platform.

The term “liftoff” is often used in the context of motion with respect to a fixed asset, such as a launch pad or sea platform, but here liftoff also includes separation from a carrier aircraft. For other types of launch platforms, the determination of liftoff will be on a case-by-case basis and may need to consider the threat to the general public before separation of the launch vehicle, such as when a balloon-launching craft is airborne.

In-Flight Failure. An in-flight failure occurs when a launch vehicle does not complete any phase of normal flight or when any anomalous condition exhibits the potential for a stage or its debris to impact the Earth or reenter the atmosphere during the mission or any future mission.

A launch accident* constitutes a failure. A launch incident[†] should be evaluated to determine if the anomalous condition exhibits the potential for a stage or its

*In 14 CFR §401.5, the FAA defines a launch accident as (1) a fatality or serious injury (as defined by 49 CFR 830.2) to any person who is not associated with the flight; (2) any damage to exceed \$25,000 to property not associated with the flight that is not located at the launch site or designated recovery area; (3) an unplanned event occurring during the flight of a launch vehicle resulting in the known impact of a launch vehicle, its payload or any component thereof: (a) for an expendable launch vehicle (ELV) outside designated impact limit lines; and (b) for a reusable launch vehicle (RLV), outside designated landing sites.

debris to impact the Earth or reenter the atmosphere during the mission or any future mission. An in-flight failure includes those cases where the failure occurs after the launch vehicle achieves an orbit, but it occurs in a stage that also operates or could operate in a future mission before achieving orbit. Hence, if the failure had occurred earlier in flight, it could have been a hazard to the public.

Initially, the FAA, in consultation with the CSWG, considered defining flight from the beginning of engine ignition to account for failures that resulted in liftoff or toppled the vehicle. However, there are times where a preplanned engine shutdown can occur that precludes liftoff but remains within the confines of planned, or normal, mission behavior. These types of occurrences would obviously not be considered an in-flight failure. As a result, although instances where anomalies in the final moments of a countdown have resulted in destruction of a vehicle, liftoff better serves to define the beginning of flight. Preflight anomalies exist that should be accounted for by launch risk analyses even though liftoff did not occur. If, for example, an anomaly occurring without liftoff had the potential to affect public safety, then it should be accounted for by a risk analyses as an on-pad failure. Note, however, such on-pad failures without liftoff should not be included in the “flight” history of a subject vehicle. The definition of flight used for the probability of failure analysis may be different from that used in other aspects of a license, such as in establishing terms and conditions of a license.

5.0 PERFORMANCE STANDARDS

The FAA uses the following performance standards to evaluate any probability of failure analysis for an ELV. A probability of failure analysis must meet all of these standards. Each of these standards will be described below. An ELV probability of failure analysis must:

- Account for launch vehicle failure probability in a consistent manner
- Incorporate accurate data, scientific principles, and valid methodologies
- Account for the outcomes of all previous flights
- Account for changes to the vehicle configuration and other factors

5.1 Account for launch vehicle failure probability in a consistent manner

Current practice promotes risk management as a means of protecting the public from a wide range of potential hazards during launch. Specifically, 14 CFR §417.107(b) of the draft final rule on Licensing and Safety Requirements for Launch published March 1, 2005 [Federal Register, vol. 70, no. 39, pp. 9885], would define acceptable risk levels for impacting inert and explosive debris, for toxic release (exposure to rocket propellant effluent), and for far field blast overpressure. The FAA’s performance standard

[†] In 14 CFR §401.5, the FAA defines a launch incident as an unplanned event occurring during flight of a launch vehicle, other than a launch accident, involving a malfunction of a flight safety system or safety critical system or failure of the licensee’s safety organization, design, or operations.

specifies that all flight safety analyses for a launch, regardless of hazard or phase of flight, should account for launch vehicle failure probability in a consistent manner.

5.2 Incorporate accurate data, scientific principles, and valid methodologies

This standard is a key element of an acceptable failure probability analysis. Accurate data means exactness and fidelity to the maximum extent possible. In this context, the FAA uses “scientific principles” to refer to knowledge, based on the scientific method, such as that established in the fields of physics, chemistry, and engineering. A failure probability analysis based on non-scientific principles, such as astrology, would not be consistent with this guide. A probability of failure estimate that is statistically and probabilistically valid should at least be the result of a sound application of mathematics. A sound application of mathematics uses correct premises and makes only conclusions that are properly derived from those premises.

The principles of probability are a mathematical theory concerned with the analysis of random events. Probability is a mathematical basis for prediction of the ratio of outcomes that would produce a given event to the total number of outcomes. Statistics refers to a branch of mathematics dealing with the collection, analysis, interpretation, and presentation of numerical data. A valid statistical analysis should account for the uncertainty in a statistical inference caused by sample size limits, the degree of applicability of data to a particular system, and the degree of homogeneity of the data.

5.3 Account for the outcomes of all previous flights

For a launch vehicle with fewer than two flights, a failure probability estimate should account for the outcomes of all previous launches of vehicles developed and launched in similar circumstances. The following five factors may be considered as part of the determination of what constitutes all previous flights of vehicles developed and launched under similar circumstances:

- Design characteristics of the vehicle.
- Development and integration processes of the vehicle, including especially the extent of integrated system testing.
- Related work experience of the launch and development team members.
- Outcomes of all previous flights of similar vehicles developed and launched by the launch operator.
- Country where the vehicle was developed and launched.

Because of the small data set available on launches of new ELVs, parsing the flight history database using the five factors described above may prove impractical. A CSWG investigation of historical failure probabilities revealed that the probability of failure on the first and second launches of a new launch vehicle depend greatly upon the launch experience of the developer. Specifically, the worldwide flight history of

ELVs from 1980 to 2002 reveals that launch operators who have never launched vehicles successfully before had 8 failures in 11 launch attempts. Worldwide flight history for “experienced launch vehicle developers” over the same period indicates 5 failures in 18 launch attempts. Many factors influence the level of experience of a launch vehicle developer. However, in the results of the recent CSWG investigation, the term “experienced launch vehicle developer” corresponded to developers who had produced at least one launch vehicle with a demonstrated probability of failure less than or equal to 33 percent. The probability of failure was based on the reference values in table A.

For a launch vehicle with two or more flights, a launch vehicle failure probability estimate should account for the outcomes of all previous flights of the subject vehicle in a valid manner. These outcomes should include all of the in-flight failures and successes that occur from liftoff until after the last exercise of control over the launch vehicle.

5.4 Account for changes to the vehicle configuration and other factors

The estimate should also account for changes in the vehicle configuration, integration and processing of the vehicle, and other factors that affect the launch vehicle development or production.

The family of Evolved Expendable Launch Vehicles (EELVs), the Delta IV and Atlas V, are examples of launch vehicles designed to fly in various configurations. For the medium class of EELVs, “changes in the vehicle configuration” include flights with various combinations of payload fairings and solid rocket motors. A valid probability of failure analysis might consider some configurations sufficiently similar to treat as Bernoulli trials of a subject vehicle, such as the EELVs that use a single common core booster. A valid analysis might consider other configurations, such as a heavy class EELV, as distinct because of important differences that may influence the probability of failure, such as flight loads, flight environment, vehicle design characteristics, and vehicle processing. To permit the development of different approaches in this area, this guide does not specify how to account for changes in the vehicle, merely that such changes should be accounted for in a valid probability of failure analysis.

6.0 PROBABILITY OF FAILURE ANALYSIS METHODOLOGY

This sample methodology satisfies the FAA’s performance standard for ELVs for the ascent phase of flight². Vehicle designs with fewer than two flights and those that have completed at least two flights are addressed.

Table A. Failure Probability Reference Values and Confidence Limits for Launch Vehicle That Have Completed at Least Two Flights

Next Launch	<----- Success Failure ----->															
3						0.55	0.89	1.00								
						0.28	0.50	0.72								
						0.00	0.11	0.45								
4						0.42	0.71	0.93	1.00							
						0.21	0.39	0.61	0.79							
						0.00	0.07	0.29	0.58							
5						0.33	0.58	0.79	0.95	1.00						
						0.17	0.32	0.50	0.68	0.83						
						0.00	0.05	0.21	0.42	0.67						
6						0.28	0.49	0.67	0.83	0.96	1.00					
						0.14	0.27	0.42	0.58	0.73	0.86					
						0.00	0.04	0.17	0.33	0.51	0.72					
7						0.24	0.42	0.59	0.73	0.86	0.96	1.00				
						0.12	0.23	0.36	0.50	0.64	0.77	0.88				
						0.00	0.04	0.14	0.27	0.41	0.58	0.76				
8						0.21	0.37	0.52	0.65	0.77	0.88	0.97	1.00			
						0.10	0.20	0.32	0.44	0.56	0.68	0.80	0.90			
						0.00	0.03	0.12	0.23	0.35	0.48	0.63	0.79			
9						0.18	0.33	0.46	0.58	0.70	0.80	0.90	0.97	1.00		
						0.09	0.18	0.28	0.39	0.50	0.61	0.72	0.82	0.91		
						0.00	0.03	0.10	0.20	0.30	0.42	0.54	0.67	0.82		
10						0.16	0.30	0.42	0.53	0.63	0.73	0.82	0.91	0.98	1.00	
						0.08	0.16	0.26	0.35	0.45	0.55	0.65	0.74	0.84	0.92	
						0.00	0.02	0.09	0.18	0.27	0.37	0.47	0.58	0.70	0.84	
11						0.15	0.27	0.38	0.48	0.58	0.67	0.76	0.84	0.92	0.98	1.00
						0.07	0.15	0.23	0.32	0.41	0.50	0.59	0.68	0.77	0.85	0.93
						0.00	0.02	0.08	0.16	0.24	0.33	0.42	0.52	0.62	0.73	0.85

6.1 Vehicle Design with Fewer Than Two Flights Completed

For a launch vehicle with fewer than two flights completed, the analysis should use a baseline value[‡] for the launch vehicle failure probability estimate equal to the upper limit of the 60-percent, two-sided confidence limits of the binomial distribution for the outcomes of all previous flights of vehicles developed and launched in similar circumstances. The FAA may adjust the failure probability estimate away from the

[‡] A baseline value is the estimated launch vehicle failure probability for the first two flights unless adjustments away from the baseline value are justified to account for particular circumstances.

baseline value to account for the level of experience demonstrated by the launch operator and other factors that affect the probability of failure.

The FAA may also consider other approaches. Under certain circumstances, a failure probability analysis for a launch vehicle with fewer than two flights can satisfy the FAA's performance standard using expert opinion. The FAA approves such adjustments on a case-by-case basis.

Certain applications of Bayesian statistics, with input data from the generic flight history of vehicles developed and launched under similar circumstances and qualitative measures associated with the launch developer or operator, constitute at least one potentially valid statistical method to make failure probability estimates for a launch vehicle with fewer than two flights³. Because the validity of a statistical analysis depends greatly on the specific data under consideration, the FAA evaluates the statistical validity of a failure probability estimate on the basis of the circumstances and data available.

6.2 Vehicle Design with at Least Two Flights Completed

For a vehicle with two or more flights, the failure probability estimate should be based on table A and the flight history of the vehicle. Table A shows the failure probability reference values and confidence limits for launch vehicles that have completed at least two flights. Reference values are shown in bold. The reference values are the midpoints between 60-percent, two-sided confidence limits⁴ of the binomial distribution. For the special cases of zero failures or all failures, the reference values are equal to the midpoints between the 80-percent, one-sided confidence limit of the binomial distribution and zero or one, respectively. Values listed on the far left of table A apply when no launch failures were experienced. Values on the far right apply when only launch failures are experienced. Values in between apply to flight histories that include both failures and successes⁵. Upper and lower confidence bounds in table A are shown directly above and below each reference value. These confidence bounds are based on 60-percent, two-sided confidence limits of the binomial distribution. For the special cases of zero failures or all failures, the upper and lower confidence bounds are equal to the 80-percent, one-sided confidence limit and zero or one, respectively. The midpoint between the 60-percent, two-sided confidence limits and, for zero failures, the midpoint between the 80-percent, one-sided confidence limit provide answers that are reasonable and consistent with current practice.

An analysis for a vehicle with at least two flights completed should use the reference value⁶ for the launch vehicle failure probability of table A based on the outcomes of all previous flights of the subject vehicle unless an adjustment is warranted (see paragraph 7.0). For example, the values in the row labeled launch number three of this table can be used to estimate the failure probability for the third launch.

Consider a vehicle that has experienced two failures in eight launches. The reference value for the probability of failure on the ninth launch would be 0.28, with a

lower bound of 0.10 and an upper bound of 0.46. Had that vehicle completed those eight launches without failure, the reference value for the ninth launch would be 0.09 with a lower bound of 0.00 and an upper bound of 0.18.

The FAA may also consider other approaches. Once a launch vehicle completes at least two flights, the FAA will accept a Bayesian estimate based on a uniform prior distribution of one hypothetical failure in two hypothetical flights updated with the outcomes of all previous flights of the subject vehicle. The reference probability estimate will be the final estimate input to any launch risk analysis unless the FAA has a reason to make an adjustment away from the reference value.

7.0 FAILURE PROBABILITY ADJUSTMENTS

Adjustments away from the reference value of the failure probability shown in table A may prove necessary for various reasons. For example, the FAA may adjust the failure probability estimate to account for the following:

- Evidence obtained from the flight history of the vehicle
- Corrective actions taken in response to a failure of the vehicle
- Vehicle configuration changes
- Other modifications that affect vehicle reliability
- Demonstrated quality of the engineering approach to launch vehicle processing and associated hazard mitigation.

In all cases, the launch risk analysis should use a final failure probability estimate that falls within the confidence limits given in table A.

7.1 Flight history

Failure probability adjustments away from the reference value may account for the nature of launch outcomes in the flight history of the subject vehicle. For example, a failure might be weighted heavily if the failure mode demonstrated a lack of quality control on the part of the launch vehicle developer, while failures for other reasons might not justify a significant departure from the reference value. In addition, a subject launch vehicle or launch vehicle subsystem may demonstrate a high degree of reliability when operated within a limited and well-defined parameter range, but the same item may demonstrate less reliability when operated outside that parameter range. The observed correlation between ambient temperature and incidence of O-ring blow-by observed for the Space Transportation System (STS) solid rocket boosters before STS 51-L (Challenger) provides such an example. In such cases, making adjustments away from the reference value to account for the nature of launch outcomes under distinct environmental conditions may be reasonable.

7.2 Corrective actions

Corrective actions taken in response to a failure of the subject launch vehicle or other subject launch vehicle modifications that affect reliability may warrant failure

probability adjustments away from the reference value. For example, a subject launch vehicle or launch vehicle subsystem may demonstrate a relatively high degree of reliability after a corrective action, such as the addition of improved guidance and navigation systems, is taken. The demonstrated reliability based on actual flight history could serve as a useful guide for the adjusted failure probability estimate. In such cases, making adjustments away from the reference value may prove reasonable.

7.3 Engineering approach

Demonstrated quality of the engineering approach to launch vehicle processing may guide failure probability adjustments away from the reference value. If, for example, a launch operator demonstrates a substandard level of quality control in its engineering approach to launch vehicle processing, then making adjustments away from the reference value may be reasonable. Also, if a launch operator demonstrates a change in quality control, then adjustments away from the reference value may be reasonable.

8.0 SUMMARY

Conducting valid probability of failure analyses of new expendable launch vehicles (ELVs) is essential to ensuring public safety for launches that employ risk management or a combination of hazard isolation and risk management. More than one way to establish an acceptable estimate of probability of failure for launch vehicles exists. Current practice at the Federal launch ranges includes multiple methodologies. Paragraph 5.0 provides performance standards that will permit launch operators to continue to employ alternative, potentially innovative methodologies that comply with the performance standard. This flexible approach is technically valid, responds to the needs of the ELV industry, and focuses on providing guiding principles, not “recipe lists.”

Paragraph 6.0 provides a sample methodology that satisfies the performance standard. For a launch vehicle with fewer than two flights completed, the sample method uses a baseline value for the launch vehicle failure probability estimate equal to the upper limit of the 60-percent, two-sided confidence limits of the binomial distribution for the outcomes of all previous flights of vehicles developed and launched in similar circumstances. For a vehicle with two or more flights, the sample method produces a failure probability estimate equal to the reference value listed in table A based on the outcomes of all previous flights of the subject vehicle.

The FAA may approve adjustments away from the baseline or reference values. Typically, such adjustments are made to account for various factors, including the quality of the engineering approach, flight history, corrective actions in response to launch vehicle failures, configuration modifications, and level of experience demonstrated by the launch operator. In all cases the launch risk analysis should use a final failure probability estimate that falls within the confidence limits given in table A.

APPENDIX: INDEPENDENT ASSESSMENT

Recognizing the central importance of probability of failure estimates to any risk analysis, the FAA and USAF obtained independent assessments of a draft version of this guide. The Common Standards Working Group (CSWG) obtained reviews from two outside teams:

1. An Independent Assessment Team (IAT), composed of Futron Corporation (Bethesda, Maryland) and Professor Ali Mosleh of the University of Maryland, was tasked to provide a thorough, well-documented, and objective assessment of the proposed requirements and supporting documentation.
2. Professor Valen Johnson of the University of Michigan reviewed the draft guidelines and developed a hierarchical Bayesian statistical approach.

Both review teams were asked to:

- Evaluate the validity of the FAA's performance standard and, if necessary, make recommendations for improvement or suggest an alternative;
- Evaluate the validity of the methodology for producing launch vehicle failure probability estimates; and
- Evaluate the clarity, usefulness, and validity of the proposed supporting documentation.

The IAT also performed a validation of the launch event outcome database developed for the CSWG by ACTA Incorporated (Torrance, California). Professor Johnson and the IAT examined the validity of the observations made on historical data, statistical interpretation, approach and criteria used to develop reference and baseline launch failure probabilities, and proposed use of the resulting numbers in this guide.

Key conclusions of both of the evaluations were that the proposed performance standards, methodological framework, and supporting documentation are, as a whole, technically valid given the practical constraints underlying an assessment of failure probabilities for new ELVs. Positive features include simplicity, justifiable conservative tendencies, and allowance for alternative methods of estimating launch failure probability, subject to reasonable quality requirements.

The IAT found that some of the methodological proposals by FAA do not adhere to the mathematical rigor of the classical statistical theory. However, the IAT also found that mathematical rigor based on classical statistical theory cannot be fully adhered to because the scarcity of the directly relevant statistical data limits practical use of such methods.

NOTES

¹ A formal definition of “new launch vehicles” does not exist. An integrated vehicle design with no or limited flight experience is generally considered new. The FAA will determine the applicability of these guidelines on a case-by-case basis. These guidelines do not necessarily apply to an integrated vehicle design with extensive flight experience.

² In this context the ascent phase of flight is from liftoff through orbital insertion, including each planned impact for an orbital launch and through final impact for a suborbital launch.

³ For example, see S.D. Guikema and M.E. Pate-Cornell, *Journal of Spacecraft and Rockets*, vol. 41, no. 1, pp. 93-102 (Jan-Feb 2004). Also, see the “degree of belief interpretation” of probability in NASA’s “Probabilistic Risk Assessment Procedures Guide for NASA Managers and Practitioners” available at <http://www.hq.nasa.gov/office/codeq/doctree/praguide.pdf>

⁴ If there are only two possible outcomes (success or failure) for a repeatable process that remains unchanged (i.e. Bernoulli trials), then there is a 40-percent chance that the actual probability of failure is outside the range specified by the 60-percent, two-sided confidence limits in table A. Specifically, a 20-percent chance exists that the actual probability of failure is above the range specified by the 60-percent, two-sided confidence limits, and a 20-percent chance that the actual probability of failure is below the range specified by the 60-percent, two-sided confidence limits.

⁵ For example, a vehicle that experienced one failure in seven launches would have a reference value of 0.20 (the bold value in the second column of the row for launch eight).

⁶ A reference value is the estimated launch vehicle failure probability for greater than two flights unless adjustments away from the reference value are justified to account for particular circumstances.